

Informal Notes on Mathematics

reorganized English version

Runnel Zhang

INM draft 2026-07-08



Read it, absorb it, and forget it.

— P. R. Halmos

Preface

This book is a reorganized edition of my informal mathematical notes, written across a long stretch of time from middle school to the present. In that sense, *INM* is not only a mathematical document, but also a personal record. It preserves the way a subject first appeared to me, the way a computation was once hard, the way a theorem gradually stopped being a name and became a tool, and the way one topic kept opening a door into another.

The original notes were not planned as a textbook. They were accumulated from exercises, self-study, lectures, fragments of reading, exam preparation, conversations, and attempts to explain ideas to myself. Some entries begin with very elementary problems: set identities, inequalities, Euclidean geometry, congruences, derivatives, determinants, or multivariable integrals. Other entries move toward more advanced objects: quotient spaces, exact sequences, cyclotomic fields, exterior algebra, Hilbert spaces, differential forms, de Rham cohomology, Fourier expansions, operator methods, topos-flavored foundations, mathematical physics, and the mathematics around machine learning. The purpose of this edition is to let these layers coexist rather than to hide one behind the other.

A recurring intention in the reorganization is to connect elementary techniques with the higher mathematics they quietly anticipate. A Venn diagram is also Boolean algebra and measurable-event logic. A gcd computation is also ideal arithmetic in a principal ideal domain. A determinant is also the action on the top exterior power. A line integral is also the pullback of a differential form. Fourier coefficients are also Hilbert-space coordinates. These connections are not meant to make the elementary material look more sophisticated than it is; they are meant to show why the elementary material is worth keeping. The first calculation is often the most honest place where a later abstraction can be seen.

The present volume arrangement follows the main themes of the original notes. The book is divided into Foundations, Arithmetic and Algebra, Combinatorics, Linear Algebra, Calculus and Inequalities, Multivariable Analysis, Fourier and ODEs, Geometry, Topology, and Research Vignettes. The last part is deliberately kept for material that is exploratory, cross-disciplinary, or difficult to classify cleanly: small mathematical-physics notes, AI-related mathematical explanations, and research-flavored digressions. I prefer to keep such pieces visible, because they record moments when mathematics was not a syllabus but a landscape.

There is also a memorial reason for making this collection. A notebook from one's earlier years is full of unevenness: some pages are naive, some are overambitious, some are clumsy, and some still contain the excitement of first contact. To reorganize these notes is not to erase that unevenness, but to give it form. The book marks a path from contest-style manipulation and classroom calculus toward algebra, analysis, topology, geometry, and the modern languages that now feel unavoidable. It is a record of mathematical growth, but also of the persistence of certain questions: what is the right structure, what is invariant, what is being measured, and why should one care?

This edition is therefore best read as a working companion rather than a polished authority. Definitions and theorems are included where they clarify the original notes, but the spirit remains informal. I have tried to preserve computations, examples, and small motivations, while adding enough surrounding structure to make the notes reusable. Some chapters are closer to textbook exposition; others are closer to annotated memory.

That mixture is intentional. It is the shape of the archive.

Runnel Zhang
July 8, 2026

Contents

Preface	v
I Sets and Foundations	1
N-20210204	5
1.1 Why classes enter before functions	5
1.2 Russell’s class as the first warning	5
1.3 Power sets and indexed families	5
1.4 Set identities	6
1.5 Images, inverse images, and functions	7
1.6 One-sided inverses	8
1.7 Set-level bookkeeping behind functions	8
N-20210206	9
2.1 Relations and equivalence classes	9
2.2 Partitions	10
2.3 Indexed products and projections	10
2.4 The first universal property	11
2.5 Integers and induction	11
2.6 Division, gcd, and congruence	12
N-20210207	15
3.1 Choice enters the story	15
3.2 Partially ordered sets	15
3.3 Well-ordering	16
3.4 Equipollence and cardinal numbers	16
3.5 Cantor’s theorem	17
3.6 Comparing cardinals	17
3.7 Infinite cardinals and finite disturbance	17
3.8 Concrete encodings behind the arithmetic	19
N-20220131	21
4.1 Why this translation/note matters	21
4.2 The discomfort with ZFC	21
4.3 Three misunderstandings about ETCS	22
4.4 Elements as special functions	22
4.5 The informal ETCS-style principles	23
4.6 Subsets via characteristic functions	23
4.7 Products, exponentials, and inverse images	24
4.8 Choice as right inverses	24
4.9 What I take from the article	24
4.10 Translation note	25
4.11 ETCS intuition: membership is not primitive practice	25
4.12 Element reasoning inside categorical set theory	25
4.13 Why this matters philosophically	25
N-20220205	27
5.1 Why there are several axiomatizations of set theory	27
5.2 The main systems	27
5.3 ZFC	27
5.4 NBG	27
5.5 Kelley–Morse set theory	28

5.6	New Foundations	28
5.7	How to compare the systems	29
5.8	What remains unclear	29
5.9	Where the idea leads	29
5.10	Why classes are introduced	29
5.11	New Foundations and stratification	30
5.12	Comparing foundational choices	30
N-20220710		31
6.1	Plane vectors	31
6.2	Field axioms	31
6.3	Ordered fields	32
6.4	Elementary consequences of the order axioms	32
6.5	Embedding $\mathbb{Q}$ into $\mathbb{R}$	33
6.6	Nested intervals	33
6.7	Ordered fields and the real line	33
6.8	Vectors as algebraicized geometry	34
6.9	The conceptual payoff	34
6.10	Ordered fields and what they cannot prove	34
6.11	Archimedean property and density	34
6.12	Completeness as the real-number axiom	34
6.13	From ordered algebra to the real line	34
N-20220713		35
7.1	Large and small positive real numbers	35
7.2	The geometric view of the real line	35
7.3	Supremum property	35
7.4	Nested intervals imply completeness	36
7.5	Dedekind cuts	36
7.6	Metric spaces	36
7.7	Subspaces and density	37
7.8	Supremum as least upper bound	37
7.9	Dedekind cuts and completeness	38
7.10	Looking ahead	38
7.11	Supremum as order-completeness	38
7.12	Cauchy completion and metric completion	38
7.13	Dedekind cuts versus Cauchy sequences	38
7.14	Completeness as the no-gap principle	38
N-20230225		39
8.1	Why this topic appears	39
8.2	The real-number proof	39
8.3	Hyperreal numbers	39
8.4	Infinitesimals and finite hyperreals	40
8.5	Hyperfinite strings of nines	40
8.6	Transfer principle	41
8.7	Derivative through infinitesimals	41
8.8	Integral through hyperfinite sums	42
8.9	What nonstandard analysis is not saying	42
8.10	Source repair: internal and external sets	42
8.11	Source repair: standard part and Loeb measure	42
8.12	Source repair: transfer has limits	43
8.13	A final synthesis	43
8.14	Real equality and limiting process	43
8.15	Hyperreals and infinitesimal remainders	43
8.16	Transfer and caution	43
8.17	Decimals, limits, and infinitesimal models	43

N-20230710	45
9.1 The question of foundations	45
9.2 Set theory	45
9.3 Category theory	46
9.4 Type theory	48
9.5 Homotopy type theory and univalence	48
9.6 Comparative analysis	49
9.7 Modern applications	49
9.8 Toward a unified perspective	49
9.9 Conclusion	50
9.10 Equality in the three languages	50
9.11 Universal properties versus encodings	50
9.12 Computation and formal proof	50
 II Categories and Logic	 51
N-20210721	55
10.1 Starting point	55
10.2 The table is not a list	55
10.3 The dualizing object	56
10.4 Stone duality as the first complete example	56
10.5 What Stone duality says in one sentence	58
10.6 Gelfand duality as a continuous analogue	58
10.7 Affine schemes and the affine line	59
10.8 Locales and the Sierpiński space	59
10.9 Topoi, logoi, and higher versions	60
10.10 The formal meaning of duality	60
10.11 What I should have understood in 2021	61
 N-20210804	 63
11.1 Boolean algebra beyond computing	63
11.2 A first definition via Heyting algebras	63
11.3 A small amount of algebra	63
11.4 Boolean rings	64
11.5 Examples of Boolean rings	65
11.6 Boolean algebras directly	66
11.7 Characteristic functions and the meaning of 2^X	66
11.8 Moving between the two structures	67
11.9 The logical shadow	67
11.10 Filters, ultrafilters, and Stone spaces	67
11.11 Boolean rings from Boolean algebras	68
11.12 Stone duality as a first hint	68
 N-20221005	 69
12.1 First orientation	69
12.2 From truth to resources	69
12.3 Narrative as a resource pipeline	70
12.4 Tensor, additives, and branching stories	72
12.5 Dialogue as a game	72
12.6 Schopenhauer's stratagem as a ludic interaction	72
12.7 Ludics: loci instead of signifieds	73
12.8 Orthogonality and negation	74
12.9 Geometry of interaction	74
12.10 From narrative to social relations	75
12.11 The general pattern	75

N-20221227	77
13.1 More examples of categories	77
13.2 Subcategories and functors	77
13.3 Forgetful functors	77
13.4 A topology generated by a distance-like question	77
13.5 Posets as categories, with an example	78
13.6 Functorial thinking	78
13.7 Generated topology	78
13.8 Why category theory appears beside topology	78
13.9 The larger pattern	79
13.10 Posets as categories and logic	79
13.11 Forgetful functors and free constructions	79
13.12 Initial topology generated by questions	79
13.13 Product topology as an initial topology	80
13.14 Quotient topology as a final topology	80
13.15 Functors from topology to algebra	80
13.16 Initial and final structures in topology	80
N-20230313	81
14.1 What this note continues	81
14.2 Posets are categories	81
14.3 A small poset example	81
14.4 Isomorphisms in a poset category	82
14.5 Groups as one-object categories	82
14.6 Why this example matters	82
14.7 Products of categories	82
14.8 The opposite category	83
14.9 Monomorphisms and epimorphisms	83
14.10 The lesson of the note	83
N-20230314	85
15.1 Free and forgetful functors	85
15.2 Faithful and full functors	85
15.3 Concrete categories	86
15.4 Functors out of a one-object group category	86
15.5 Functors between one-object group categories	86
15.6 Why the note says “learn algebra first”	87
15.7 Looking ahead: natural transformations	87
N-20240211	89
16.1 Why topoi enter the story	89
16.2 Categories	89
16.3 Universal properties and limits	90
16.4 Exponentials and cartesian closure	90
16.5 Subobjects and the subobject classifier	91
16.6 Elementary topoi	91
16.7 A toy presheaf topos	92
16.8 Grothendieck topos and sheaves	93
16.9 Internal logic	93
16.10 Power objects	94
16.11 Topos as generalized universe of sets	94
16.12 A high-level comparison	94
16.13 Further directions	94
16.14 Subobjects and characteristic maps	95
16.15 Why logic becomes intuitionistic	95
16.16 Sheaves as variable sets	95

N-20251109	97
17.1 The setting: maps with zero maps	97
17.2 Kernel: the universal object killed by a map	97
17.3 Cokernel: the universal quotient killing a map	98
17.4 Image: the part of the codomain actually reached	98
17.5 Coimage: the effective part of the domain	99
17.6 The standard factorization	99
17.7 A vector-space sanity check	100
17.8 A module example: multiplication by an integer	100
17.9 Exactness language	101
17.10 Warnings: what depends on being abelian	101
17.11 Synthesis: from element chasing to universal properties	101
 III Number Theory Basics	 103
N-20220717	109
18.1 A beginning point: recursive definitions	109
18.2 Arithmetic progressions	109
18.3 Determining an arithmetic progression	110
18.4 Geometric progressions	110
18.5 Converting recurrences	111
18.6 A factorial recurrence	111
18.7 Groups	111
18.8 Symmetric and dihedral groups	112
18.9 Products and isomorphisms	112
18.10 Linear recurrences	112
18.11 Groups from repeated operations	113
18.12 Why the pattern matters	113
18.13 Recurrences as iteration of maps	113
18.14 First-order linear recurrences	113
18.15 Groups behind repeated operations	113
18.16 Sequences as orbits of operations	113
 N-20220723	 115
19.1 The ant on a decagon	115
19.2 Roots of unity	115
19.3 The polynomial $x^{2n} + x^n + 1$	115
19.4 A broader pattern	116
19.5 Cyclotomic viewpoint	116
19.6 Why this is more than a trick	116
19.7 Roots of unity viewpoint	117
19.8 Cyclotomic divisibility	117
 N-20220803	 119
20.1 A trigonometric opening example	119
20.2 Logarithmic sequences	119
20.3 Recovering terms from partial sums	120
20.4 Sums and recurrence relations	120
20.5 Geometric progression products	120
20.6 A harmonic-looking sequence	121
20.7 A final synthesis	121
20.8 Sequences as discrete dynamical systems	121
20.9 Partial sums and finite differences	121
20.10 Generating functions for sequence exercises	122
20.11 Telescoping as a coboundary	122
20.12 Second-order recurrences and characteristic roots	122

20.13	Discrete calculus of elementary sequences	123
N-20220913		125
21.1	From guessing to proving	125
21.2	Divisibility	125
21.3	Example: reducing the dividend	126
21.4	Example: polynomial divisibility in one variable	126
21.5	Congruences	126
21.6	A large power sum	127
21.7	A wider interpretation	127
21.8	Divisibility as ideal containment	127
21.9	GCD and LCM as ideal operations	127
21.10	Congruences as quotient rings	128
21.11	Units and Bezout inverses	128
21.12	Valuations as prime coordinates	128
21.13	Chinese remainder theorem as local-to-global	129
21.14	Polynomial divisibility as the same story	129
21.15	Divisibility as quotient algebra	129
N-20220918		131
22.1	A general odd-power divisibility principle	131
22.2	A divisibility problem with 512	131
22.3	Why finite differences work	132
22.4	An alternating harmonic numerator	132
22.5	The next abstraction	132
22.6	Finite differences as discrete derivatives	133
22.7	The binomial basis and integer-valued polynomials	133
22.8	Pairing as group symmetry	133
22.9	P-adic valuation and lifting intuition	134
22.10	Alternating harmonic sums and denominator control	134
22.11	Discrete summation by parts	134
22.12	Discrete calculus and local divisibility	134
N-20221013		135
23.1	Shoelace formula	135
23.2	Why the shoelace formula works	135
23.3	A circle-area maximum problem	136
23.4	Remainder theorem	136
23.5	A polynomial remainder example	136
23.6	Congruence arithmetic	137
23.7	Fermat and Euler	137
23.8	Shoelace formula as oriented area	138
23.9	Evaluation as quotient by $x-a$	138
23.10	Where the calculation points	138
23.11	Shoelace as a discrete line integral	138
23.12	Remainder theorem and quotient rings	139
23.13	Fermat–Euler and unit groups	139
23.14	Quotients in area, polynomials, and congruences	139
N-20221030		141
24.1	The well-ordering principle	141
24.2	Divisibility, gcd, and lcm	141
24.3	Division algorithm	142
24.4	Euclidean algorithm	142
24.5	Example	142
24.6	Fundamental theorem of arithmetic	142
24.7	The Euler phi function	143

24.8	The broader lesson	143
24.9	Well-ordering and Euclidean domains	143
24.10	Bezout and principal ideals	143
24.11	Euler function and unit groups	144
24.12	Euclidean arithmetic as ideals and unit groups	144
24.13	Euclidean domains	144
24.14	PID and Bezout domains	144
24.15	Euler phi as unit counting	144
24.16	Multiplicative functions	145
24.17	Euclidean algorithm and ideals	145
24.18	Euler's theorem as group theory	145
N-20221031		147
25.1	Divisibility	147
25.2	Division algorithm	147
25.3	Greatest common divisor and least common multiple	148
25.4	Euclidean algorithm	148
25.5	Bezout theorem	148
25.6	Gauss lemma and divisibility consequences	149
25.7	Prime numbers	149
25.8	Congruences	149
25.9	Exercises and common methods	149
25.10	From example to method	150
25.11	Division algorithm as normal form	150
25.12	Euclidean algorithm and invariant gcd	150
25.13	Gauss lemma and prime ideals	150
25.14	Normal forms, gcd invariants, and prime divisibility	150
25.15	$\mathbb{Z}$ as a PID	151
25.16	Modular inverses as units	151
25.17	Smith normal form over integers	151
25.18	CRT as product decomposition	151
25.19	Solving congruences by image and kernel	152
25.20	Smith normal form and congruences	152
N-20221119		153
26.1	Prime and composite numbers	153
26.2	Euclid's theorem	153
26.3	Prime divisibility	153
26.4	Fundamental theorem of arithmetic	154
26.5	Perfect numbers and square numbers	154
26.6	Exponent of a prime in n factorial	155
26.7	Congruence exercises	155
26.8	Broader perspective	155
26.9	Prime exponents and valuations	155
26.10	Divisor functions as multiplicative functions	155
26.11	Congruences and finite groups	156
26.12	Prime coordinates and finite unit groups	156
26.13	Arithmetic functions and convolution	156
26.14	Möbius inversion in number theory	156
26.15	Dirichlet series and Euler products	157
26.16	p -adic valuations and multiplicativity	157
26.17	Euler products and prime independence	157
26.18	Möbius inversion over primes	158
N-20221126		159
27.1	Congruence	159
27.2	Basic properties	159

27.3	Residue classes	159
27.4	Complete and reduced residue systems	160
27.5	Example: last digit	160
27.6	Example: modulo 17	160
27.7	Example: reducing a long expression	161
27.8	A polynomial congruence by induction	161
27.9	Quadratic residues modulo 8	161
27.10	Where the idea leads	161
27.11	Residue systems as representatives	161
27.12	Euler function through prime powers	162
27.13	Orders and exponent reduction	162
27.14	Residues, units, and exponent cycles	162
27.15	Residue systems as groups of units	162
27.16	CRT decomposition of unit groups	162
27.17	Primitive roots and characters	163
27.18	Returning to residue-system exercises	163
27.19	Primitive roots as logarithms in finite groups	163
27.20	Characters as Fourier analysis on unit groups	163
N-20230127		165
28.1	The original problem and the first successful method	165
28.2	The direct factorization attempt	165
28.3	The attempted generalization	166
28.4	The computational experiment	166
28.5	Explicit root formula	167
28.6	Common real factors	167
28.7	The organizing idea	168
28.8	Cyclotomic summary of the experiment	168
28.9	Cyclotomic polynomials	168
28.10	Galois action on roots of unity	168
28.11	Resultants and common roots	169
28.12	Finite fields and character sums	169
28.13	Returning to the original experiments	169
28.14	Primitive roots and cyclotomic factor tests	169
28.15	Cyclotomic identities through Möbius inversion	169
IV	Arithmetic Geometry	171
N-20211007		175
29.1	The question is an inequality, not a slogan	175
29.2	A first chain of rigidity results	175
29.3	From Mordell finiteness to effective inequalities	176
29.4	The Szpiro inequality as a concrete target	176
29.5	Vojta's inequality and the hyperbolic curve viewpoint	177
29.6	Arithmetic fundamental groups	178
29.7	The local elliptic curve model: Tate curves and q -parameters	178
29.8	Normalised degree as log-volume	179
29.9	Why ordinary scheme-theoretic geometry is not enough	179
29.10	The route from deformation to inequality	180
29.11	What this first note establishes	181
N-20220123		183
30.1	Why this note starts with algebraic number theory	183
30.2	Number fields and algebraic integers	183
30.3	Ideals and factorization	183
30.4	Quadratic fields and class number	184

30.5	Localization, completion, and p-adic numbers	184
30.6	Modular forms and L-functions	186
30.7	First understanding of Langlands	186
30.8	Local and global correspondences	187
30.9	Cross-field intuitions	187
30.10	Numerical experiments	187
30.11	Topics to learn next	189
30.12	Synthesis	189
30.13	Local-global intuition	189
30.14	Why L-functions enter	189
30.15	A first Langlands slogan	190
30.16	Structural viewpoint	190
30.17	Ideals as repaired unique factorization	190
30.18	Local fields and arithmetic microscopes	190
30.19	Modular forms and Galois representations	190
30.20	Arithmetic symmetry as the bridge to Langlands	191
N-20221002		193
31.1	Why finite fields matter	193
31.2	Fields	193
31.3	Field extensions	193
31.4	Constructing finite fields	194
31.5	Important properties of finite fields	194
31.6	Polynomial factorization over finite fields	195
31.7	Example over F_3	195
31.8	Algorithms	195
31.9	Applications	196
31.10	Irreducibility over finite fields	196
31.11	Constructing extensions	196
31.12	Frobenius map	196
31.13	What survives later	197
31.14	How the pieces fit	197
31.15	Frobenius and the first Galois group	197
31.16	Irreducible polynomials and orbit size	198
31.17	Counting irreducible polynomials	198
31.18	Square-free and distinct-degree factorization	198
31.19	Berlekamp's algorithm as linear algebra	199
31.20	Coding theory connection	199
31.21	Algebraic geometry shadow	199
31.22	Frobenius structure behind finite fields	199
N-20221120		201
32.1	Definition and nonsingularity	201
32.2	Why the discriminant appears	201
32.3	Cryptographic motivation	202
32.4	Geometric addition law	202
32.5	Algebraic formula	202
32.6	Special cases and the point at infinity	203
32.7	Why associativity is difficult	203
32.8	Elliptic curves over finite fields	203
32.9	Why nonsingularity matters for the group law	204
32.10	Explicit addition formulas	204
32.11	Cryptographic intuition	204
32.12	Structural viewpoint	205
32.13	Nonsingularity and the discriminant	205
32.14	The group law as divisor arithmetic	205
32.15	Associativity and geometry	205

32.16	Cryptographic meaning	205
32.17	Elliptic curves as geometry with a group law	205
N-20251113		207
33.1	What has to be explained	207
33.2	The nonarchimedean theta-function	207
33.3	The local theta-link as a non-ring-theoretic map	208
33.4	The bridge supplied by Kummer theory	209
33.5	Two symmetries and the bookkeeping of labels	209
33.6	Conjugate synchronization and the basepoint problem	210
33.7	The log-link: recovering additive structure from multiplicative structure	210
33.8	The log-theta-lattice	211
33.9	The log-shell: a common container for incompatible links	212
33.10	Mono-anabelian reconstruction and multiradiality	212
33.11	The three indeterminacies	213
33.12	What the second note adds	214
N-20251114		215
34.1	Why the word Teichmueller matters	215
34.2	Classical Teichmueller theory: fixed topology, variable complex structure	215
34.3	Distortion is measured in one universe	216
34.4	Why the arithmetic analogy cannot be literal	217
34.5	Hodge theaters as arithmetic stages	217
34.6	Etale-like and Frobenius-like directions	218
34.7	The theta-link is the point where one universe breaks	218
34.8	The log-link does not restore naive equality	219
34.9	Classical marking versus multiradial reconstruction	219
34.10	The mathematical source of the controversy	220
34.11	A compact model of the logic	221
34.12	What this third note establishes	221
V	Functions and Inequalities	223
N-20220417		229
35.1	Proof strategies	229
35.2	Cyclic and symmetric sums	229
35.3	AM–GM	230
35.4	Muirhead’s inequality	230
35.5	Power means	230
35.6	Hölder and Cauchy	231
35.7	Jensen and Karamata	231
35.8	Cyclic sums	231
35.9	Proof strategy before computation	231
35.10	Next viewpoint	232
35.11	Inequalities as convexity statements	232
35.12	Majorization and Karamata	232
35.13	Holder as norm geometry	232
35.14	Convexity, duality, and equality cases	232
35.15	Inequality methods as choice of language	233
35.16	Convexity behind Jensen and Karamata	233
35.17	Schur convexity	233
35.18	Entropy and the log-sum inequality	234
N-20220718		235
36.1	Fixed-point iteration	235
36.2	Möbius recurrences	235

36.3	Example	236
36.4	Second-order linear recurrences	236
36.5	Dedekind cuts and metric exercises	236
36.6	Fixed points of recurrence maps	237
36.7	Monotone bounded sequences	237
36.8	After the examples	237
36.9	Fixed points and stability	237
36.10	Möbius recurrences and matrices	237
36.11	Second-order recurrences	238
36.12	Recurrences as one-dimensional dynamics	238
N-20220721		239
37.1	Random experiments and sample spaces	239
37.2	Classical model and Venn diagrams	239
37.3	Conditional probability	239
37.4	A transition to substitution tricks	240
37.5	A classical symmetric equation	240
37.6	Inequality examples	241
37.7	Jensen, trigonometry, and contest style	241
37.8	The hidden structure	241
37.9	Conditional probability as changing measure	241
37.10	Bayes formula and inverse problems	241
37.11	Inequality substitutions as coordinate choices	242
37.12	Changing measures and changing coordinates	242
N-20220801		245
38.1	Subsets and characteristic functions	245
38.2	Set laws	245
38.3	Venn diagram exercises	246
38.4	Inclusion–exclusion with literature examples	246
38.5	Set constructions with squares	246
38.6	Floor functions	247
38.7	Functions and monotonicity	247
38.8	Characteristic functions and subsets	247
38.9	Functions as rules with domains	247
38.10	The method behind the trick	248
38.11	Set operations as Boolean algebra	248
38.12	Characteristic functions and algebra of predicates	248
38.13	Preimages as Boolean homomorphisms	248
38.14	Sigma-algebras and measurable events	249
38.15	Stone duality intuition	249
38.16	Difference of squares as algebraic structure	249
38.17	Sets as algebra, logic, and measurement	249
38.18	Boolean algebra beyond Venn diagrams	249
38.19	Characteristic functions	250
38.20	Sigma-algebras as observable event systems	250
38.21	Stone duality as topology of logic	250
38.22	Returning to the original set exercises	251
38.23	From Venn diagrams to algebra of information	251
38.24	Why sigma-algebras restrict what can be asked	251
N-20220806		253
39.1	Images, preimages, and finite set examples	253
39.2	Closure under addition	253
39.3	A combinatorial maximal-set problem	253
39.4	Functional and algebraic exercises	254
39.5	Matrix and coordinate side computations	254

39.6	Images and preimages behave differently	254
39.7	Closure operations	255
39.8	The reusable pattern	255
39.9	Closure operators	255
39.10	Generated objects as intersections	255
39.11	Fixed points of closure	255
39.12	Images, preimages, and logic	256
39.13	Posets and extremal chains	256
39.14	Galois connections	256
39.15	Tarski fixed-point theorem	256
39.16	Closure, order, and fixed-point thinking	257
39.17	Closure operators across mathematics	257
39.18	Closed sets as fixed points	257
39.19	Galois connections and closure	257
39.20	Tarski fixed points and monotone iteration	258
39.21	Returning to the original exercises	258
39.22	Generation as an intersection construction	258
39.23	Closure and logic of consequences	258
N-20220808		261
40.1	A set-counting problem	261
40.2	A modular arithmetic construction	261
40.3	Why not every construction is finished	261
40.4	Sequence exercises	262
40.5	Mathematical induction	262
40.6	Examples of induction	262
40.7	Induction structure	263
40.8	Modular patterns	263
40.9	A note on scope	263
40.10	Construction versus upper bound	263
40.11	Residue classes as finite quotient spaces	263
40.12	Induction as well-founded recursion	264
40.13	Finite differences in divisibility induction	264
40.14	Floor functions and lattice points	264
40.15	Pigeonhole and averaging	265
40.16	Finite structure from quotients and induction	265
N-20220810		267
41.1	Basic inequalities	267
41.2	AM–GM	267
41.3	Cauchy’s inequality	268
41.4	Triangle inequality and its variants	268
41.5	Worked AM–GM examples	268
41.6	A Cauchy example	269
41.7	Coordinate geometry at the end of the note	269
41.8	What the examples reveal: three structures	269
41.9	Cauchy–Schwarz and Gram determinants	270
41.10	Triangle inequality and normed spaces	270
41.11	AM–GM, convexity, and cones	270
N-20220811		271
42.1	How to read this note	271
42.2	A fraction problem with a constraint	271
42.3	A weighted square-root comparison	271
42.4	Maximizing a radical expression	272
42.5	Equality in Cauchy–Schwarz	273
42.6	Triangle-angle inequalities	273

42.7	Two spheres and a linear equation	273
42.8	A product inequality	274
42.9	Vector-style radical inequality	274
42.10	The deeper habit	274
42.11	Equality cases as structural information	275
42.12	Optimization on constrained domains	275
42.13	Vectors and geometric inequalities	275
42.14	Equality cases as geometric diagnostics	275
42.15	Equality cases as geometry	275
42.16	Lagrange multipliers and inequalities	276
42.17	Semidefinite viewpoint	276
42.18	Returning to the original problems	276
N-20221012		277
43.1	Bernoulli's inequality	277
43.2	Power means	277
43.3	Cauchy–Schwarz	277
43.4	Minkowski inequality	278
43.5	A three-variable inequality	278
43.6	Chebyshev's inequality	278
43.7	Jensen's inequality	279
43.8	Young, Hölder, and Minkowski	279
43.9	Cauchy–Schwarz as a master inequality	279
43.10	Equality conditions	280
43.11	The lesson behind it	280
43.12	The hierarchy of means	280
43.13	Convex order and Jensen	280
43.14	Minkowski and norm geometry	280
43.15	Classical inequalities in one convex landscape	280
N-20221105		281
44.1	A short note about elementary methods	281
44.2	Periodicity from a transformed argument	281
44.3	A quadratic minimum	281
44.4	A logarithmic comparison	282
44.5	A parameterized polynomial	282
44.6	An inequality under $abc=1$	282
44.7	The conceptual turn	283
44.8	Periodicity from transformed arguments	283
44.9	Quadratic and logarithmic estimates	283
44.10	Functional equations and orbit reduction	283
44.11	Functions through symmetry and normal form	283
N-20221217		285
45.1	Structure of the exercise sheet	285
45.2	Domain problems	285
45.3	Floor functions	286
45.4	Range problems	286
45.5	Dirichlet function	286
45.6	Graph transformations and zero counting	287
45.7	Logarithmic parameter equation	287
45.8	Functional equations with period-like recurrence	287
45.9	A reciprocal functional equation	288
45.10	Quadratic iteration problem	288
45.11	Max and absolute value	288
45.12	Parameter problems and critical values	288
45.13	Functional equations as reduction rules	289

45.14	What the pattern suggests	289
45.15	Domains as constraint intersections	289
45.16	Dirichlet function through density and measure	289
45.17	Thomae's function as a nearby comparison	290
45.18	Parameter problems as bifurcation diagrams	290
45.19	Functional equations as group actions on the domain	290
45.20	Max, min, and nonsmooth points	290
45.21	Ranges as images and compactness	291
45.22	Function exercises through topology and dynamics	291
N-20230107		293
46.1	A note on graph problems	293
46.2	Centers of symmetry	293
46.3	Symmetry about a line	293
46.4	Absolute value transformations	293
46.5	Zero counting by graphs	294
46.6	How to find a center of symmetry in practice	294
46.7	Absolute value as folding	294
46.8	Parameter motion and intersection counting	294
46.9	Graph transformations as group actions	295
46.10	Absolute value as folding, not algebraic noise	296
46.11	Zero counting as a stability problem	296
N-20230116		297
47.1	What a function is	297
47.2	Image and preimage	297
47.3	Injective and surjective	297
47.4	Composition and inverse functions	298
47.5	Functional equations on finite sets	298
47.6	Functional equations over real numbers	298
47.7	Looking ahead	298
47.8	Functions as morphisms	299
47.9	Images, preimages, and logic	299
47.10	Functional equations as constraints on orbits	299
47.11	Functions as structure-preserving maps	299
47.12	Morphisms and structural information loss	299
47.13	Image and preimage as an adjunction	299
47.14	Finite functional graphs	300
47.15	Returning to the original examples	300
47.16	Injections, surjections, and factors	300
47.17	Fibers as local data of a function	300
N-20230117		301
48.1	Continuing the function chapter	301
48.2	Piecewise functions	301
48.3	Periodicity from functional equations	301
48.4	Graph transformations	302
48.5	Image of an interval	302
48.6	Functional equations and substitution	302
48.7	Counting functions between finite sets	302
48.8	A compact bridge	303
48.9	Piecewise functions as glued maps	303
48.10	Periodicity from generated transformations	303
48.11	Images of intervals	303
48.12	Local definitions and global function behavior	303
48.13	Functional equations as dynamics	303
48.14	Cauchy's equation and regularity	303

48.15	Symbolic dynamics	304
48.16	Rigidity from regularity	304
48.17	Piecewise functions and partitions	304
48.18	Functional equations and invariants	305
N-20230316		307
49.1	Fixed points	307
49.2	Iteration	307
49.3	Metric spaces and contractions	307
49.4	Banach fixed point theorem	308
49.5	Error estimate	308
49.6	Applications	309
49.7	Beyond Banach	309
49.8	What the example reveals	309
49.9	Why completeness is the hidden hypothesis	309
49.10	Picard iteration as fixed-point theory	310
49.11	Source repair: non-contractive fixed-point theorems	310
49.12	Source repair: Brouwer and Schauder	310
VI	Combinatorics and Models	313
N-20220720		319
50.1	Counting principles	319
50.2	A divisibility example	319
50.3	Permutations and combinations	320
50.4	Binomial theorem	320
50.5	Polynomial coefficients	320
50.6	Further contest examples	321
50.7	Binomial identities by stories	321
50.8	When to use multiplication and when to use addition	321
50.9	A broader reading	321
50.10	Counting principles as operations on combinatorial classes	321
50.11	Inclusion–exclusion as Möbius inversion	322
50.12	Indicator functions and probability language	322
50.13	Generating functions for binomial identities	322
50.14	Counting up to symmetry: Burnside’s lemma	323
50.15	Counting through structure and symmetry	323
N-20221205		325
51.1	Summation and product notation	325
51.2	De Morgan’s laws	325
51.3	Inclusion–exclusion	325
51.4	Pigeonhole principle	326
51.5	Addition and multiplication principles	326
51.6	Permutations	326
51.7	Combinations	327
51.8	Pascal identity	327
51.9	Binomial theorem	327
51.10	Counting subsets	327
51.11	Examples and exercises	328
51.12	Indicator functions as a unifying proof method	328
51.13	Combinatorial identities by double counting	329
51.14	Conceptual synthesis	329
51.15	Indicator functions and inclusion–exclusion	329
51.16	Inclusion–exclusion as Möbius inversion	329
51.17	Double counting as proof by projection	330

51.18	Generating functions	330
51.19	Pigeonhole principle and the probabilistic method	330
51.20	Asymptotics of binomial coefficients	330
51.21	Combinatorial species	331
51.22	Counting as inversion, projection, and generation	331
N-20221207		333
52.1	A set construction problem	333
52.2	What this problem teaches	334
52.3	A Venn-diagram necessary-and-sufficient problem	334
52.4	Difference identities	335
52.5	Coordinate and algebraic exercises	335
52.6	A note on proof writing	335
52.7	Extremal set problems: why constructions matter	335
52.8	Venn diagrams versus proofs	336
52.9	A wider frame	336
52.10	Extremal problems: upper bound plus construction	336
52.11	Additive combinatorics viewpoint	336
52.12	Cauchy–Davenport shadow	337
52.13	Boolean atoms and Venn regions	337
52.14	Symmetric difference as addition mod two	337
52.15	Sperner and intersecting-family context	337
52.16	Shifting and stability	338
52.17	Extremal sets, Boolean algebra, and stability	338
N-20221209		339
53.1	What this page is doing	339
53.2	A floor-function example	339
53.3	A binomial counting identity	339
53.4	Solving a system by discriminants	340
53.5	A prime-number exercise	340
53.6	A nested-set exercise	340
53.7	A recurrence of sets	340
53.8	An intersection-counting problem	341
53.9	Set-builder notation as a translation problem	341
53.10	Recursive set construction	341
53.11	Further direction	341
53.12	Floor functions as lattice-point counts	341
53.13	Finite differences and binomial basis	342
53.14	Recursive set construction and fixed points	342
53.15	Incidence counting	343
53.16	Polynomial method	343
53.17	Group actions and corrected division	343
53.18	Counting by lattices, incidences, and symmetries	343
N-20221210		345
54.1	Binomial theorem	345
54.2	Absolute value of binomial sums	345
54.3	A summation exercise	345
54.4	Set exercises	346
54.5	Counting subsets with properties	346
54.6	Set-builder notation	347
54.7	Closing perspective	347
54.8	Binomial identities as counting in two languages	347
54.9	Binomial transform and Möbius inversion	347
54.10	Finite differences and the binomial basis	348
54.11	Probability distribution behind binomial coefficients	348

54.12	Asymptotic shape: from binomial coefficients to Gaussians	349
54.13	Generating functions and labelled structures	349
54.14	Combinatorial species and relabelling	349
54.15	Binomial algebra from sets to species	349
N-20230118		351
55.1	Pigeonhole principle	351
55.2	Subsets and binomial coefficients	351
55.3	Binomial identities by double counting	351
55.4	A subset-sum style example	352
55.5	Generating functions	352
55.6	Helly-type remark	352
55.7	The next layer	352
55.8	Pigeonhole as averaging	352
55.9	Generating functions and subset sums	352
55.10	Helly-type thinking	353
55.11	Finite compactness principles in combinatorics	353
55.12	Pigeonhole as compactness principle	353
55.13	Ramsey theory	353
55.14	Helly theorem	353
55.15	Fractional and topological Helly	354
55.16	Local-to-global intersection principles	354
55.17	Ramsey as structured inevitability	354
N-20240323		355
56.1	Binomial identities	355
56.2	Good subsets	355
56.3	Generating functions	356
56.4	A probability example with repeated trials	356
56.5	Catalan numbers	357
56.6	Convex sets and convex hulls	357
56.7	Separating convex sets	357
56.8	Planar graphs and Euler-type counting	358
56.9	Graphs and adjacency matrices	358
56.10	The deeper habit	358
56.11	Generating functions as algebraic counting	358
56.12	Catalan structures	359
56.13	Convex hulls and separation	359
56.14	Generating functions, Catalan recursion, and spectra	359
56.15	Catalan numbers as recursive geometry	359
56.16	Lattice paths and reflection principle	360
56.17	Graph walks and adjacency matrices	360
56.18	Random walks and Markov chains	360
56.19	Catalan recursion and parenthesization	360
56.20	Adjacency spectra and expansion	361
N-20240330		363
57.1	Good triangles	363
57.2	Why vertices must lie on the convex hull	363
57.3	Cyclic arc encoding	363
57.4	The geometric idea behind the proof	364
57.5	Euler formula for convex polyhedra	364
57.6	Regular convex polyhedra	364
57.7	Euler characteristic	365
57.8	Degree of a map	365
57.9	Homology group and quotient group reminders	365
57.10	A bridge outward	365

57.11	Euler characteristic as invariant	366
57.12	Degree counting on graphs	366
57.13	Quotients and surface construction	366
57.14	Euler bookkeeping from graphs to surfaces	366
57.15	Euler characteristic as topological bookkeeping	366
57.16	Euler characteristic and homology	367
57.17	Degree theory	367
57.18	Quotient spaces and identifications	367
57.19	Euler characteristic as alternating rank	367
57.20	Degree and fixed-point consequences	368
N-20251126		369
58.1	What the model is trying to measure	369
58.2	Submultiplicativity and the connective constant	369
58.3	Bridges and why they help	370
58.4	Polygons	372
58.5	Critical exponents	373
58.6	How dimension changes the story	374
58.7	The broader lesson	374
N-20251127		377
59.1	Why the hexagonal lattice is special	377
59.2	A domain version of the counting problem	377
59.3	The parafermionic observable	378
59.4	From local cancellation to a global identity	378
59.5	The upper bound on the connective constant	380
59.6	The lower bound on the connective constant	380
59.7	Relation with SLE	380
59.8	Loop models and the same mechanism	381
59.9	From example to method	381
N-20251128		383
60.1	The high-dimensional philosophy	383
60.2	The two-point function	383
60.3	The lace expansion as corrected random walk	384
60.4	Bubble diagrams and diagrammatic estimates	384
60.5	Differential inequalities	385
60.6	Functional integral representations	385
60.7	Dimension four and logarithmic corrections	386
60.8	How the three methods fit together	387
60.9	The conceptual turn	387
VII	Linear Algebra	389
N-20210621		397
61.1	First orientation	397
61.2	Props as typed wiring diagrams	397
61.3	The alphabet: copy, discard, add, zero, and scalars	398
61.4	Two colours: a comonoid and a monoid	398
61.5	The bialgebra law	399
61.6	Scalars as ring operations	400
61.7	How a matrix becomes a diagram	400
61.8	Completeness: a proof system for matrices	401
61.9	Why the calculus hides indices and exposes structure	401
61.10	From matrices to linear relations	402
61.11	From linear relations to Lagrangian relations	402

61.12	Affine shifts and stabilizers	403
61.13	Relation with ZX-calculus	403
61.14	Broader perspective	403
N-20220407		405
62.1	Complex arithmetic	405
62.2	Field axioms in $\mathbb{C}$	405
62.3	Vector space identities	406
62.4	A non-example with infinity	406
62.5	Subspaces	406
62.6	Bridge to later ideas	407
62.7	How to read basic vector-space exercises	407
62.8	Beyond real and complex scalar structures	407
62.9	Quotient spaces	408
62.10	Dual spaces and coordinates	408
62.11	Returning to the original elementary exercises	408
62.12	Why fields make linear algebra work	408
62.13	The universal property of quotient spaces	409
N-20220411		411
63.1	Invertibility	411
63.2	The 2×2 formula	411
63.3	Diagonal matrices	411
63.4	Products of inverses	412
63.5	Adjugate formula	412
63.6	Row reduction for inverses	412
63.7	Rotation matrices	412
63.8	Beyond the first calculation	413
63.9	Invertibility as isomorphism	413
63.10	Adjugate and exterior powers	413
63.11	Rotations as orthogonal matrices	413
63.12	Invertibility, volume, and orthogonal motion	414
63.13	Inverse as undoing a linear process	414
63.14	Invertible matrices as a Lie group	414
63.15	Determinant as a volume form	414
63.16	Polar decomposition	415
63.17	Numerical stability of inverses	415
63.18	Row operations and the general linear group	415
63.19	Exponential map and infinitesimal motion	416
N-20220412		417
64.1	Unit vectors and Kronecker delta	417
64.2	Einstein summation convention	417
64.3	Rank	417
64.4	Transpose and Hermitian conjugate	418
64.5	Trace	418
64.6	A more structural view	418
64.7	Transpose and dual maps	419
64.8	Hermitian adjoints in Hilbert spaces	419
64.9	Trace as invariant	419
64.10	Schatten norms	419
64.11	Indices as tensor bookkeeping	420
64.12	Transpose as pullback on dual spaces	420
64.13	Trace as contraction	420
64.14	Rank and factorization	420
64.15	Coordinate indices as intrinsic linear algebra	420

N-20220725	421
65.1 Inversions and signs	421
65.2 Definition of determinant	421
65.3 Basic determinant properties	422
65.4 Sarrus and triangular matrices	422
65.5 Cofactor expansion	422
65.6 A structured determinant	422
65.7 The determinant line and exterior powers	423
65.8 Why permutation signs are forced	423
65.9 Minors as exterior coordinates	424
65.10 Cramer's rule from volume ratios	424
N-20230209	425
66.1 From vector identities to linear representations	425
66.2 Coordinates as isomorphisms	425
66.3 The basis-change formula	426
66.4 A worked basis-change example	426
66.5 Which quantities are invariant?	427
66.6 Covectors and the inverse-transpose rule	427
66.7 Bilinear forms are not linear operators	428
66.8 Symmetric and skew parts	429
66.9 Tensor slots and transformation laws	429
66.10 Tensor products and universal property	430
66.11 Coordinate-free reading of the original matrix facts	430
N-20230217	431
67.1 Tensor products and bases	431
67.2 Bilinearity	431
67.3 Metric, dot product, and components	432
67.4 Reciprocal basis	432
67.5 Covariant and contravariant components	432
67.6 Vector product and Levi-Civita symbol	433
67.7 Scalar triple product	433
67.8 What changes next	433
67.9 Tensor product by its universal property	433
67.10 Components under basis change	433
67.11 Metric as an index-lowering map	434
67.12 Tensor notation as transformation law	434
67.13 Dual spaces and reciprocal bases	434
67.14 Musical isomorphisms	434
67.15 Frames and moving coordinates	435
67.16 Continuum mechanics notation	435
67.17 Metric tensors and Gram matrices	435
67.18 Covariant and contravariant components from Gram matrices	435
67.19 Why tensor products linearize bilinear data	436
N-20251108	437
68.1 Vectors beyond arrows	437
68.2 The minus-one trick	437
68.3 Example of the trick	437
68.4 Linear mappings	438
68.5 Complex numbers as a real vector space	438
68.6 Finite-dimensional isomorphism theorem	438
68.7 Basis change	439
68.8 Equivalence and similarity	439
68.9 Kernel and image	439
68.10 Categorical side remarks	440

68.11	The general pattern	440
68.12	Vectors as elements of a linear structure	440
68.13	The minus-one trick as a kernel graph	440
68.14	Kernel-image factorization	441
68.15	Coordinates, kernels, and affine linearization	441
68.16	Exact sequences behind rank-nullity	442
68.17	Quotient description of the image	442
68.18	Cokernel and obstructions	442
68.19	Returning to the MML basis-change exercises	442
68.20	Row reduction as choosing a splitting	443
68.21	Cokernel as the missing equation space	443
N-20251110		445
69.1	Cofactor orthogonality	445
69.2	Vandermonde determinant	445
69.3	Block triangular determinants	446
69.4	General Laplace expansion	446
69.5	Determinant splitting	446
69.6	Adding rows or columns	447
69.7	Augmenting to a Vandermonde determinant	447
69.8	Row and column transformations	447
69.9	Triangularization by row operations	448
69.10	Recursive determinants	448
69.11	Cramer's rule	448
69.12	Rank and minors	449
69.13	Solvability of linear systems	449
69.14	Handwritten comments preserved	450
69.15	The lesson behind it	450
69.16	Determinant as alternating multilinear invariant	450
69.17	Vandermonde and evaluation maps	450
69.18	Rank and linear systems	450
69.19	Determinants and rank as invariants	450
69.20	Determinant through exterior powers	451
69.21	Schur complement	451
69.22	Matrix determinant lemma	451
69.23	Rank varieties	451
69.24	Returning to Cramer and rank	452
69.25	Determinant identities as low-rank geometry	452
69.26	Schur complements and eliminating variables	452
69.27	Determinantal varieties and singular loci	452
N-20251202		453
70.1	Diagonalizable matrices	453
70.2	Diagonalizability is independent from other properties	453
70.3	Algebraic and geometric multiplicity	454
70.4	Procedure for diagonalization	454
70.5	Inner product and vector length	454
70.6	Orthogonal and orthonormal systems	455
70.7	Schmidt orthogonalization	455
70.8	Example of Schmidt orthogonalization	456
70.9	QR decomposition	456
70.10	Orthogonal matrices	456
70.11	Real symmetric matrices	457
70.12	Quadratic forms	457
70.13	Positive definite quadratic forms	457
70.14	Conic and quadric interpretations	458
70.15	Common traps recorded in the notes	458

70.16	A closing lens	458
70.17	Diagonalization as decomposition into invariant lines	459
70.18	Algebraic versus geometric multiplicity	459
70.19	Inner products and orthogonal projection	459
70.20	Quadratic forms and spectral theorem	459
70.21	Spectral decomposition and quadratic geometry	459
70.22	Spectral theorem and orthogonal bases	459
70.23	QR as algorithmic Gram-Schmidt	460
70.24	Sylvester's law of inertia	460
70.25	Hilbert-space projection	460
70.26	Orthogonal diagonalization and axes	460
70.27	Gram-Schmidt as factorization of coordinates	461
70.28	Spectral theorem beyond finite dimension	461
N-20260103		463
71.1	Determinant expansion and cofactors	463
71.2	Recursive determinants	463
71.3	Circulant-style determinant	464
71.4	Cancellation through orthogonal matrices	464
71.5	Determinant product and block simplification	464
71.6	Schur-complement style block transformation	465
71.7	Rank-two decomposition	465
71.8	Adjugate identities	466
71.9	Skew-symmetric matrices via quadratic forms	466
71.10	Absorption and cancellation	467
71.11	Rank identities	467
71.12	Commutation from factorization	468
71.13	Projection-like matrices from bases	468
71.14	Nilpotent polynomial inverse	468
71.15	Equivalent vector groups	469
71.16	Computing matrix powers	469
71.17	Homogeneous systems and adjugate	470
71.18	Null space basis and augmented rank	470
71.19	Rank of products and equality of solution spaces	471
71.20	General solution from known solutions	471
71.21	Equivalence of row vector groups	472
71.22	Diagonal dominance and invertibility	472
71.23	Where the idea leads	472
71.24	Determinant tricks as invariant-preserving transformations	472
71.25	Block determinants and Schur complements	473
71.26	Matrix powers and minimal polynomials	473
71.27	Matrix tricks from invariants and polynomials	473
71.28	Rank factorization and normal forms	473
71.29	Idempotents as projections	473
71.30	Spectral asymptotics of matrix powers	474
71.31	Null spaces and orthogonal complements	474
71.32	Matrix powers and stable limits	474
71.33	Row space, column space, and rank factorization	475
71.34	Positivity through Gram forms	475
N-20260105		477
72.1	Eigenvalues of matrix functions	477
72.2	No real eigenvector from a quadratic relation	477
72.3	Independence of eigenvectors	478
72.4	Real symmetric reconstruction from eigenvalues	478
72.5	Similarity invariants: trace and determinant	479
72.6	Trace shortcuts	479

72.7	Row-sum eigenvectors	479
72.8	Diagonalization problems	480
72.9	Real symmetric eigenvalue problems	480
72.10	Odd powers of real symmetric matrices	480
72.11	Rayleigh quotient extrema	481
72.12	Orthogonal determinant trick	481
72.13	Quadratic forms through diagonalization	481
72.14	Positive definite equations	482
72.15	Commuting positive definite products	482
72.16	Principal minors and congruence	483
72.17	Ordering eigenvalues and positive definiteness	483
72.18	Transition matrices	483
72.19	Same coordinates under two bases	484
72.20	Conceptual synthesis	484
72.21	Eigenvalues through functional calculus	484
72.22	Similarity invariants	484
72.23	Positive definiteness and quadratic geometry	485
72.24	Row-sum eigenvectors and invariant aggregate directions	485
72.25	Spectral geometry of positive matrices	485
72.26	Functional calculus for matrices	485
72.27	Loewner order and positive definiteness	485
72.28	Lyapunov equations	486
72.29	Convex and semidefinite optimization	486
72.30	Matrix monotone warning	486
72.31	Positive definite matrices and ellipsoids	486
72.32	Lyapunov equations and stability	487
72.33	Functional calculus and matrix inequalities	487
N-20260303		489
73.1	Concept map of the chapter	489
73.2	Five model transformations	489
73.3	A spectral example: a biological neural network	490
73.4	SVD as a sequence of geometric operations	491
73.5	SVD and latent structure in data	492
73.6	Rank-one blocks and image compression	492
73.7	A taxonomy of matrix properties	494
73.8	Closing perspective	494
N-20260305		497
74.1	Projection onto a subspace	497
74.2	Projection matrix and least squares	497
74.3	Pseudoinverse	497
74.4	Projection geometry in the source diagram	498
74.5	Bridge to later ideas	498
74.6	Projection theorem	498
74.7	Normal equations	498
74.8	Pseudoinverse and singular values	499
74.9	Least squares as projection geometry	499
N-20260311		501
75.1	Why size matters	501
75.2	Metrics and completeness	501
75.3	Normed vector spaces	501
75.4	Standard vector norms	502
75.5	Matrix norms	502
75.6	Unitary invariance	502
75.7	Induced norms	503

75.8	Normal matrices	503
75.9	Compatibility	503
75.10	Worked constrained maxima	503
75.11	ML-facing interpretation	504
75.12	What each norm is measuring	504
N-20260312		505
76.1	The long-term question	505
76.2	Spectral radius	505
76.3	Spectral radius is not a norm	505
76.4	Approximating spectral radius	506
76.5	Convergent matrices	506
76.6	Jordan boundary behavior	506
76.7	Neumann series	506
76.8	Matrix power series	506
76.9	Condition number	507
76.10	Perturbation bound	507
76.11	Ill-conditioned example	507
76.12	Numerical and ML interpretation	507
76.13	Power growth, inversion, and sensitivity	508
N-20260318		509
77.1	Why canonical forms matter	509
77.2	Polynomial matrices	509
77.3	Smith normal form	509
77.4	Jordan data from elementary divisors	509
77.5	Jordan chains	510
77.6	Computing powers	510
77.7	Gershgorin discs	510
77.8	Gershgorin figure	511
77.9	Separated components and scaling	511
77.10	Power iteration	511
77.11	Inverse and shifted inverse iteration	511
77.12	Numerical traps	512
77.13	ML-facing interpretation	512
77.14	Choosing an eigenvalue tool	512
N-20260602		513
78.1	Why decompositions are algorithms	513
78.2	Doolittle LU	513
78.3	Existence and uniqueness	513
78.4	Elimination matrices	513
78.5	Doolittle formulas	513
78.6	Pivoting	514
78.7	Storage of LU	514
78.8	Cholesky decomposition	514
78.9	Cholesky example	514
78.10	Full-rank decomposition	514
78.11	Hermite standard form	514
78.12	Singular values	515
78.13	Singular value decomposition	515
78.14	SVD geometry	515
78.15	Which decomposition matches the matrix	515
N-20260603		517
79.1	Why ordinary inverse is not enough	517
79.2	Four regimes	517

79.3	Penrose equations	517
79.4	Fifteen generalized inverses	517
79.5	Projection transformations	517
79.6	Idempotence and orthogonality	517
79.7	Column and row-space projections	518
79.8	SVD computation	518
79.9	Full-rank computation	518
79.10	Greville recursion	518
79.11	Iterative approximations	518
79.12	Solvability	519
79.13	Least squares and minimum norm	519
79.14	Identities	519
79.15	Reading the Moore–Penrose inverse geometrically	519
N-20260604		521
80.1	The computational problem	521
80.2	Triangular solves	521
80.3	Gauss elimination	521
80.4	Pivoting	521
80.5	LU solution	521
80.6	Cholesky and LDL	522
80.7	Tridiagonal chasing	522
80.8	Stationary iteration	522
80.9	Jacobi iteration	522
80.10	Gauss-Seidel iteration	522
80.11	Diagonal dominance	523
80.12	SOR iteration	523
80.13	Optimization viewpoint	523
80.14	Steepest descent	523
80.15	Method map	523
80.16	How the linear-system method is selected	524
VIII	Calculus	525
N-20220315		531
81.1	A scattered algebraic note	531
81.2	A quadratic with parameters	531
81.3	A parabola and a hyperbola	532
81.4	A constrained extremum	532
81.5	Tangency as a double-root condition	533
81.6	Completing the square	533
81.7	A wider frame	533
81.8	Tangency as multiplicity	534
81.9	Constrained extrema and normal directions	534
81.10	Degeneracy and second-order data	534
81.11	Multiplicity, normality, and local geometry	534
N-20220319		535
82.1	Sequences and convergence	535
82.2	Limits of functions	535
82.3	Derivatives	536
82.4	Small-o and big-O	536
82.5	Integration	537
82.6	Why the epsilon definition is shaped this way	537
82.7	Continuity by sequences	538
82.8	Derivative as best linear approximation	538

82.9	Integral as accumulation	538
82.10	Convergence as topology	538
82.11	Derivative as a linear map	539
82.12	Returning to the original examples	539
N-20220320		541
83.1	From the mean value theorem to Cauchy's form	541
83.2	Continuity at a point	541
83.3	How this note fits the previous one	542
83.4	Rolle's theorem as the engine	542
83.5	Cauchy's theorem and L'Hopital's rule	543
83.6	Uniform continuity	543
83.7	Further direction	543
83.8	Mean value theorems as compactness plus local linearity	543
83.9	Cauchy's theorem and ratios of differentials	544
83.10	Uniform continuity and why closed intervals matter	544
83.11	From local slopes to global control	544
83.12	Mean value theorem via compactness	544
83.13	Lipschitz estimates and uniqueness	545
83.14	Taylor formula with remainder	545
N-20220409		547
84.1	Motivation	547
84.2	Faà di Bruno formula	547
84.3	Leibniz's differential notation	547
84.4	Parametric derivatives	548
84.5	Differentiating determinants	548
84.6	Matrix multiplication	548
84.7	Schur complement	549
84.8	What the pattern suggests	549
84.9	Faà di Bruno as chain rule for higher jets	549
84.10	Leibniz rules and derivations	549
84.11	Matrix multiplication as composition	549
84.12	Schur complements	550
84.13	Composition, derivations, and elimination	550
N-20220628		551
85.1	Linearity	551
85.2	Partial fractions	551
85.3	Integration by parts	552
85.4	Substitution	552
85.5	Partial fractions as algebra before calculus	553
85.6	Integration by parts as product rule reversed	553
85.7	A second lens	554
85.8	Partial fractions as algebra of poles	554
85.9	Integration by parts as adjointness	554
85.10	Substitution as pullback of one-forms	554
85.11	Integration techniques as structural transformations	554
85.12	Integration as reverse differentiation	554
85.13	Substitution as pullback	555
85.14	Integration by parts and adjoints	555
85.15	Partial fractions and residues	555
N-20221107		557
86.1	Trigonometric limits	557
86.2	Derivative of sine and cosine	557
86.3	Tangent, cotangent, secant, and cosecant	557

86.4	Inverse trigonometric derivatives	558
86.5	Examples from the handwritten page	558
86.6	A closing lens	558
86.7	The circle as a Lie group	558
86.8	Trigonometric derivatives from the exponential	559
86.9	Inverse trigonometric branches	559
86.10	Trigonometry from circle symmetry and branches	559
N-20221211		561
87.1	Average velocity and instantaneous velocity	561
87.2	Derivative notation	561
87.3	Power functions	561
87.4	Exponential and logarithmic derivatives	562
87.5	Trigonometric derivatives	562
87.6	Inverse trigonometric derivatives	562
87.7	Examples	563
87.8	Context and outlook	563
87.9	Derivative as first-order linearization	563
87.10	Inverse trigonometric derivatives and the inverse function theorem	563
87.11	Endpoint blow-up and failure of local invertibility	564
87.12	Logarithm as inverse and as group homomorphism	564
87.13	Automatic differentiation and computation graphs	564
87.14	Jets and higher-order local data	565
87.15	Differentials and one-forms	565
87.16	Derivatives as linearization and pullback	565
N-20230310		567
88.1	The derivative from secants to tangents	567
88.2	Derivative notation and local definition	568
88.3	Rules of differentiation	568
88.4	Elementary derivative formulas	568
88.5	Extrema, Fermat's lemma, and critical points	569
88.6	A derivative proof of a basic inequality	569
88.7	Young's inequality	569
88.8	Holder's inequality	570
88.9	Minkowski's inequality	570
88.10	Convex functions and Jensen's inequality	570
88.11	Typical exercise themes	571
88.12	The reusable pattern	572
88.13	Derivative as local linear approximation	572
88.14	Extrema and first variation	572
88.15	Young, Holder, and convex duality	572
88.16	Local linearity and convex control	572
N-20251112		573
89.1	Cauchy–Schwarz and Minkowski	573
89.2	Neighborhood notation	573
89.3	Trigonometric identities	573
89.4	Hyperbolic functions	574
89.5	Sequence limits: epsilon–N language	574
89.6	Infinite limits and one-sided limits	575
89.7	Heine criterion for function limits	575
89.8	Infinitesimals and infinitesimal orders	575
89.9	Squeeze theorem and monotone convergence	576
89.10	Cauchy criteria	576
89.11	Continuity	576
89.12	Derivatives and differentials	577

89.13	Taylor-type expansions	577
89.14	Where the calculation points	577
89.15	Limits through filters and sequences	578
89.16	Hyperbolic functions and exponential coordinates	578
89.17	Cauchy–Schwarz, Minkowski, and normed spaces	578
89.18	Calculus review through limits and norm geometry	578
N-20251210		579
90.1	Strict monotonicity from the derivative	579
90.2	First derivative test for extrema	579
90.3	Second derivative test	579
90.4	Concavity and inflection points	580
90.5	Asymptotes	580
90.6	General procedure for sketching a curve	581
90.7	Basic antiderivative table	581
90.8	Substitution	581
90.9	Integration by parts	582
90.10	Rational functions and partial fractions	582
90.11	Quadratic denominator reduction	582
90.12	Trigonometric rational functions	583
90.13	Definite integrals and Riemann sums	583
90.14	Basic properties of definite integrals	584
90.15	Integral mean value theorem	584
90.16	Applications: area, volume, and arc length	584
90.17	Improper integrals	585
90.18	Broader perspective	585
90.19	Monotonicity and derivative sign	585
90.20	Convexity and supporting lines	586
90.21	Integration methods as inverse rules	586
90.22	Error terms and numerical integration	586
90.23	Local derivative data and global shape	586
N-20260106		587
91.1	Convexity and the Hermite–Hadamard inequality	587
91.2	Kantorovich-type integral inequality	587
91.3	Integral mean value with odd symmetry	588
91.4	A lower bound involving f double-prime over f	589
91.5	Symmetry of integrals about the midpoint	589
91.6	Midpoint and trapezoidal error terms	590
91.7	First special integral	590
91.8	Positivity of a Fresnel-type integral	591
91.9	A symmetry integral with arcsine	591
91.10	A logarithmic sine integral	592
91.11	Integrals of $1/(a + b \cos x)$	592
91.12	Wallis integral and Wallis product	593
91.13	A wider frame	593
91.14	Convexity as second-order positivity	593
91.15	Integral inequalities as weighted averages	593
91.16	Symmetry and cancellation	594
91.17	Wallis formula and asymptotic products	594
91.18	Integral inequalities through convexity and asymptotics	594
IX	Multivariable Calculus	595
N-20220405		599
92.1	Fréchet differentiability and complex differentiability	599

92.2	Fréchet derivative	599
92.3	Visualizing complex maps	600
92.4	Hessian matrices and Taylor expansion	600
92.5	Bordered Hessian	600
92.6	Example	601
92.7	Finite fields and Galois theory	602
92.8	Why complex analysis meets Hessians	602
92.9	Frechet derivative as linear approximation	602
92.10	Hessian as second derivative	602
N-20220501		603
93.1	Why this note was too compressed before	603
93.2	Lagrange multipliers: the basic idea	603
93.3	Example: minimizing distance to a line	603
93.4	Why the method is valid	604
93.5	Hessian matrix and the second derivative test	604
93.6	Sylvester criterion and eigenvalues	605
93.7	Bordered Hessian	605
93.8	Example revisited with a bordered Hessian	605
93.9	Degenerate cases and higher-order tests	606
93.10	Saddle points and geometric visualization	606
93.11	Optimization beyond equality constraints	606
93.12	Global topology and extrema	607
93.13	Computation and MATLAB	607
93.14	Second variation on the tangent space	607
93.15	A useful afterword	608
93.16	Lagrange multipliers as cotangent annihilators	608
93.17	Bordered Hessian and restricted second variation	608
93.18	Degenerate extrema and higher order	608
93.19	First and second variations under constraints	608
N-20220728		609
94.1	Spatial vectors	609
94.2	Coordinate formulas	609
94.3	Cross product	610
94.4	From products to universal properties	610
94.5	Categories and groupoids	610
94.6	Cross product geometry	611
94.7	Categorical product versus vector product	611
94.8	What changes next	611
94.9	Cross product from the Hodge star	611
94.10	Scalar triple product and volume	611
94.11	Universal products versus vector products	611
94.12	Metric, orientation, and universal products	612
94.13	Cross product and Hodge duality	612
94.14	Scalar triple product as volume form	612
94.15	Lie algebra of rotations	612
94.16	Categorical versus vector products	613
94.17	Why the cross product is exceptional	613
94.18	Angular velocity and infinitesimal rotations	613
94.19	Universal property versus geometric product	613
N-20230205		615
95.1	The vector triple product	615
95.2	Scalar triple product	615
95.3	From vector identities to exterior algebra	615
95.4	Levi-Civita notation	616

95.5	Lagrange identity	616
95.6	Why tensors enter naturally	617
95.7	Coordinate-free moral	617
95.8	The hidden structure	617
95.9	Levi-Civita symbol as a tensor	618
95.10	Contraction identities	618
95.11	Coordinate-free meaning of vector identities	618
95.12	Returning to the original vector identities	618
95.13	From vector identities to index-free proofs	619
95.14	Why tensor notation matters	619
N-20240413		621
96.1	The mapmaker's problem	621
96.2	Derivative first, partial derivatives second	621
96.3	The chain rule is not a formula trick	622
96.4	A small laboratory: polar coordinates	622
96.5	Tangent planes from equations	623
96.6	A manifold is a space with local Euclidean windows	623
96.7	Two pictures of tangent vectors	623
96.8	Regular values: one clean machine for making manifolds	624
96.9	One chart in detail: stereographic projection	624
96.10	Vector fields: tangent spaces glued along a space	625
96.11	Pushforward in a concrete example	625
96.12	What curvature will later add	625
96.13	How the local-linear picture is used	626
N-20260304		627
97.1	Why point-set language appears	627
97.2	Balls and neighborhoods	627
97.3	Interior, exterior, and boundary	627
97.4	Open and closed sets	628
97.5	Accumulation points and closure	629
97.6	Boundedness and compactness	629
97.7	Connectedness and path-connectedness	630
97.8	What this vocabulary controls	630
N-20260306		631
98.1	What changes from one variable to many	631
98.2	How to prove and disprove limits	631
98.3	Continuity	632
98.4	Partial and directional derivatives	632
98.5	Differentiability as linear approximation	633
98.6	Tangent planes and normal vectors	633
98.7	Mixed partial derivatives	634
98.8	The main lesson	635
N-20260313		637
99.1	The chain rule is composition of linear maps	637
99.2	One input variable, two intermediate variables	637
99.3	Two input variables	638
99.4	Total differentials	638
99.5	Implicit differentiation in one equation	639
99.6	Implicit differentiation for systems	639
99.7	A shape table	640
99.8	What to carry forward	640
N-20260314		641
100.1	Purpose of this guide	641

100.2	Classifying points relative to a set	641
100.3	Limit templates	641
100.4	Degree comparison for rational expressions	642
100.5	Differentiability checklist	642
100.6	Tangent and normal templates	643
100.7	Local extrema template	643
100.8	Common failure modes	644
N-20260325		645
101.1	Local extrema	645
101.2	First-order necessary condition	645
101.3	Quadratic model and the Hessian	645
101.4	Second derivative test	646
101.5	A picture of the Hessian test	647
101.6	Constrained extrema	647
101.7	Endpoint and boundary warnings	648
X	Vector Calculus	649
N-20230123		655
102.1	Indefinite and definite integrals	655
102.2	Area by slicing	655
102.3	Work as an integral	655
102.4	Line integrals	656
102.5	Green-type intuition	656
102.6	Polar coordinates	656
102.7	Spherical coordinates	657
102.8	Dot product and cross product	657
102.9	After the examples: what object is being integrated?	657
102.10	Line integrals as pullbacks	658
102.11	Exact forms and potentials	658
102.12	Conservation laws	658
102.13	De Rham preview	659
102.14	Work integrals and gauge freedom	659
102.15	Polar and spherical adapted charts	659
N-20230201		661
103.1	First orientation	661
103.2	Statement of Newton's shell theorem	661
103.3	Why symmetry alone is not enough	661
103.4	Solid-angle proof in full detail	662
103.5	Potential proof: doing the integral	663
103.6	The direct force integral	664
103.7	Gauss-law viewpoint	665
103.8	Harmonic potential	665
103.9	Why the inverse-square law is special	666
103.10	Application: a uniform solid sphere	666
103.11	Multipole viewpoint	666
103.12	A broader reading	667
N-20230311		669
104.1	Why two theories of integration exist	669
104.2	Riemann sums	669
104.3	Darboux upper and lower sums	669
104.4	A simple integrable discontinuity	670
104.5	Dirichlet's function	670

104.6	The Lebesgue idea	670
104.7	Characteristic functions and sets	671
104.8	Convergence theorems	671
104.9	Improper Riemann integrals versus Lebesgue integrals	672
104.10	Comparison table	672
104.11	Source repair: the measure-space starting point	672
104.12	Source repair: simple functions as building blocks	672
104.13	Source repair: probability viewpoint	673
104.14	A note on scope	673
104.15	Riemann integration partitions the domain	673
104.16	Lebesgue integration partitions the values	673
104.17	Dirichlet function as separator	673
104.18	Convergence theorems as exchange principles	673
104.19	Integration theory through measurable structure	675
104.20	Riemann sums and measure spaces	675
104.21	Dominated and monotone convergence in practice	676
104.22	Lp spaces and duality	676
104.23	Returning to the original Riemann examples	676
N-20240427		677
105.1	A theorem looking for its natural language	677
105.2	One-forms: measuring motion	677
105.3	The pullback experiment	677
105.4	Two-forms: measuring oriented area	678
105.5	Exterior derivative as one operator	678
105.6	A familiar theorem reappears	678
105.7	Orientation is not decoration	679
105.8	Why arc length is a different kind of geometry	679
105.9	Closed forms can detect holes	679
105.10	A table of classical shadows	680
105.11	Flux as a form	680
105.12	The moral of pullback	680
105.13	Exit line	680
N-20260403		681
106.1	From area sums to double integrals	681
106.2	Geometric meanings	681
106.3	Basic properties	682
106.4	Existence of the integral	682
106.5	Symmetry before computation	683
106.6	Mean value theorem for double integrals	683
106.7	A worked average-value example	684
106.8	Relation with product measure	684
106.9	A real Fubini warning: conditional two-dimensional accumulation	684
106.10	Product measure behind rectangular sums	685
106.11	Symmetry as projection onto invariant functions	685
106.12	Average value as a pushforward distribution	686
106.13	How the double-integral setup is checked	686
N-20260405		687
107.1	The real task: describe the region	687
107.2	Changing the order of integration	687
107.3	Polar coordinates	688
107.4	Jacobian as local area scaling	688
107.5	Why the absolute value appears	689
107.6	A linear change example	689
107.7	Choosing coordinates	689

107.8	Slicing as pushforward of area	689
107.9	Exterior algebra proof of the Jacobian factor	690
107.10	Polar coordinates as a singular chart	690
107.11	Coarea formula in the simplest radial case	690
107.12	When a clever substitution is actually a quotient	691
107.13	How a change of variables earns its bounds	691
N-20260406		693
108.1	Triple integrals as weighted volume	693
108.2	Rectangular-coordinate example	693
108.3	Cylindrical coordinates	693
108.4	Spherical coordinates	694
108.5	Coordinate-choice table	694
108.6	Mass, center of mass, and moments	694
108.7	Symmetry	695
108.8	Three ways to understand volume factors	695
108.9	Toward divergence theorem	695
108.10	Metric tensor derivation of spherical volume	696
108.11	Divergence as infinitesimal volume expansion	696
108.12	Pushforward of densities and probability	696
108.13	Volume forms and orientation	697
108.14	How the solid determines the coordinates	697
N-20260408		699
109.1	Volume as an integral	699
109.2	Describing spatial regions	699
109.3	A reconstruction workflow	700
109.4	Changing order of integration	700
109.5	Cross-sections	701
109.6	Top minus bottom	701
109.7	Cylindrical and spherical descriptions	701
109.8	Absolute extrema over solids	702
109.9	Fibre integration: top-minus-bottom as a pushforward	702
109.10	Cavalieri principle and layer-cake formula	702
109.11	Why regions must be split: projections can change topology	703
109.12	Order of integration as choosing a foliation	703
109.13	Reconstructing the solid before integrating	703
N-20260416		705
110.1	Two meanings of integration on curves	705
110.2	Line integrals with respect to arc length	705
110.3	Line integrals of vector fields	705
110.4	Surface integrals of scalar fields	706
110.5	Flux integrals	706
110.6	First kind versus second kind	706
110.7	Parametrization invariance	707
110.8	Differential forms viewpoint	707
110.9	Orientability warning	707
110.10	Densities versus differential forms	707
110.11	Flux and the Hodge star	707
110.12	Boundary orientation is not a sign convention	708
110.13	Non-orientability as a real obstruction	708
110.14	What the parametrization is responsible for	708
N-20260417		709
111.1	Green's theorem	709
111.2	Proof idea: rectangles and cancellation	709

111.3	Conservative fields and path independence	710
111.4	The punctured-plane warning	711
111.5	Area by Green's theorem	711
111.6	Differential forms notation	711
111.7	Closed forms, exact forms, and topology	712
111.8	Connection with de Rham cohomology	712
111.9	Winding number hidden in the punctured-plane example	712
111.10	From Green's theorem to residues	713
111.11	Cohomology as bookkeeping for failed potentials	713
111.12	Boundary cancellation and the algebra $\partial^2 = 0$	713
111.13	What Green's theorem checks in practice	714
N-20260615		715
112.1	Why this note is split this way	715
112.2	Four similar-looking integrals	715
112.3	First-kind surface integrals	715
112.4	Template surfaces	716
112.5	Example pattern: spherical cap	717
112.6	Second-kind surface integrals as flux	718
112.7	Model flux computations: closed, capped, singular, and sheeted	718
112.8	Gauss theorem	720
112.9	Singular fields and missing points	720
112.10	Stokes theorem	720
112.11	Stokes model computations: area vector and spanning-surface freedom	721
112.12	Decision table	722
112.13	Conceptual endpoint	722
112.14	The de Rham complex behind the vector formulas	722
112.15	General Stokes theorem	722
112.16	Singular sources and distributions	723
112.17	Conservation laws	723
112.18	Currents	724
112.19	Gauss and Stokes as dual local-to-global laws	724
112.20	Weak forms and PDE	724
XI	Complex Analysis	725
N-20220321		729
113.1	The complex exponential as motion	729
113.2	Roots and angles	730
113.3	A divisibility result	730
113.4	Closing perspective	732
113.5	Why multiplication by i is rotation	732
113.6	Roots of unity as symmetry	732
113.7	Context and outlook	732
113.8	Roots of unity as a finite symmetry group	733
113.9	Cyclotomic factorization	733
113.10	Fourier viewpoint	733
113.11	Cyclic symmetry on the unit circle	733
113.12	Roots of unity as characters	733
113.13	Discrete Fourier transform	734
113.14	Line-to-circle covering map	734
113.15	Cyclotomic fields	734
113.16	Finite Fourier analysis from roots of unity	734
113.17	Cyclotomic fields and constructibility	735

N-20220323	737
114.1 Why complex analysis feels different	737
114.2 Complex derivative	737
114.3 Basic examples	738
114.4 Closure properties	739
114.5 The main intuition	739
114.6 Complex versus real differentiability	739
114.7 Local power series intuition	739
N-20220813	741
115.1 Why complex integration is different	741
115.2 Real integrals as a warm-up	741
115.3 Contours and parametrization	741
115.4 Riemann sums and midpoint sums	742
115.5 Complex Riemann sums	743
115.6 Integrals of powers on the unit circle	743
115.7 Cauchy–Goursat theorem	743
115.8 Contour integrals depend on paths	744
115.9 Primitive functions and path independence	744
115.10 Why singularities matter	744
115.11 A bridge outward	744
115.12 Contour integrals as integrals of one-forms	745
115.13 Cauchy theorem and exactness	745
115.14 Residues as obstruction	745
115.15 Residues, topology, and exactness	745
N-20260616	747
116.1 Why these topics belong together	747
116.2 Numerical series	747
116.3 Positive-term series	747
116.4 Ratio and root tests	748
116.5 Alternating and arbitrary-sign series	748
116.6 Function series and uniform convergence	749
116.7 Power series	749
116.8 Taylor series	749
116.9 Improper integrals	750
116.10 Unified viewpoint	750
116.11 Uniform convergence as norm convergence	750
116.12 Analytic functions and convergence radius	751
116.13 Dominated convergence and improper integrals	751
116.14 Tauberian flavor	751
N-20260617	753
117.1 Why Fourier series is a separate theme	753
117.2 Orthogonality of trigonometric functions	753
117.3 Fourier coefficients	753
117.4 Dirichlet convergence theorem	753
117.5 Even and odd functions	754
117.6 Half-range expansions	754
117.7 Concrete coefficient models	754
117.8 Dirichlet kernel: what partial sums really do	756
117.9 Fejer means: averaging partial sums improves convergence	756
117.10 Hilbert-space viewpoint	756
117.11 Fourier series in Hilbert space	757
117.12 Parseval identity	757
117.13 Gibbs phenomenon	757
117.14 Sturm–Liouville and Fourier bases	757

117.15	Heat and wave equations	758
117.16	Convergence modes of Fourier series	758
117.17	Smoothness is decay of coefficients	758
117.18	Sobolev norm from Fourier coefficients	759
117.19	Fourier transform as the non-periodic limit	759
117.20	Distributions and weak Fourier series	759
N-20260618		761
118.1	Why the differential-equation part is one note	761
118.2	Basic concepts	761
118.3	Separable equations	761
118.4	Homogeneous first-order equations	761
118.5	First-order linear equations	762
118.6	Bernoulli equations	762
118.7	Exact equations	762
118.8	Reducible higher-order equations	763
118.9	Second-order linear equations	763
118.10	Constant-coefficient homogeneous equations	763
118.11	Constant-coefficient nonhomogeneous equations	764
118.12	Initial-value problems	764
118.13	Functional-analytic viewpoint	764
118.14	Differential equations as operators	764
118.15	Green functions	765
118.16	Semigroups and evolution	765
118.17	Sturm–Liouville and Fourier expansion	765
N-20260619		767
119.1	What a first-order PDE asks for	767
119.2	The unforced transport equation	767
119.3	Example: translating a parabola	768
119.4	Forced transport	769
119.5	What can go wrong	770
119.6	General linear first-order equations	770
119.7	Quasilinear equations	770
119.8	Crossing characteristics and shocks	771
119.9	Initial curves and well-posedness	772
119.10	Conservation-law viewpoint	772
N-20260620		773
120.1	The equation and its physical meaning	773
120.2	Factorization into transport operators	773
120.3	Initial-value problem on the whole line	774
120.4	Finite strings and Fourier series	775
120.5	Energy conservation	775
120.6	Inhomogeneous wave equation and Duhamel’s principle	776
120.7	Boundary conditions on a finite interval	776
120.8	Deriving the Fourier coefficients	776
120.9	Standing waves and travelling waves	777
120.10	The wave equation with damping	777
N-20260621		779
121.1	Diffusion versus wave motion	779
121.2	Heat kernel on the whole line	779
121.3	Similarity solution on a semi-infinite rod	779
121.4	Fourier series for a finite rod	780
121.5	Maximum principle	781
121.6	Energy decay on an interval	781

121.7	The maximum principle with proof idea	781
121.8	Smoothing and the heat semigroup	782
121.9	Fourier transform solution on the real line	782
121.10	Boundary conditions and conservation	783
121.11	Steady states and long-time behavior	783
121.12	Comparison with the wave equation	784
XII	Classical Geometry	785
N-20220615		791
122.1	What the note contains	791
122.2	Problem 1: two distances under a linear constraint	791
122.3	Problem 2: parabola and distance	791
122.4	Problem 3: packing unit circles	792
122.5	Problem 4: a symmetric Diophantine equation	792
122.6	The larger pattern	793
122.7	Coordinate choice and invariants	793
122.8	Packing and density	793
122.9	Symmetric Diophantine equations	793
122.10	Olympiad structure behind coordinates and symmetry	793
122.11	What an olympiad sheet trains	793
122.12	Coordinate geometry as optimization	794
122.13	Packing and convex bodies	794
122.14	Group actions and symmetry reduction	794
122.15	Convexity and supporting hyperplanes	795
122.16	Returning to the original coordinate methods	795
122.17	Coordinate choice as quotienting symmetry	795
122.18	Packing bounds as volume plus local obstruction	795
122.19	Equality cases as rigid configurations	795
N-20220708		797
123.1	Reflecting the orthocenter	797
123.2	The incenter–excenter lemma	797
123.3	Directed angles modulo 180°	798
123.4	Miquel point of a triangle	799
123.5	Directed angles reduce case distinctions	799
123.6	Miquel point proof pattern	799
123.7	Further context	799
123.8	Directed angles as projective bookkeeping	800
123.9	Reflections and isogonal structure	800
123.10	Miquel point as a compatibility condition	800
123.11	Angle invariants and cyclic compatibility	800
123.12	Why directed angles are worth learning	800
123.13	Directed angles and invariants	800
123.14	Circle pencils	801
123.15	Cross ratio	801
123.16	Moduli of configurations	801
123.17	Miquel configurations and spiral similarity	801
123.18	Inversion as normalization of circles	802
N-20220715		803
124.1	Conic sections	803
124.2	Latus rectum and tangent geometry	803
124.3	Rational and irrational density	804
124.4	Vector spaces	804
124.5	Dedekind cuts and order	804

124.6	Addition of Dedekind cuts	805
124.7	Additive inverses	805
124.8	Conics through focus and directrix	805
124.9	Dedekind cuts and addition	806
124.10	A compact bridge	806
124.11	Conics through affine and projective lenses	806
124.12	Density and completion	806
124.13	Dedekind addition as forced operation	806
124.14	Coordinates, conics, and the completed line	806
N-20220716		807
125.1	Line and ellipse	807
125.2	Tangents from a point to an ellipse	807
125.3	Line and hyperbola	808
125.4	Parabola tangent through a point	808
125.5	A parabola and a circle	809
125.6	An ellipse with prescribed quadrilateral area	809
125.7	The next layer	810
125.8	Tangency via discriminants	810
125.9	Polarity with respect to a conic	810
125.10	Area constraints and affine invariance	811
125.11	Tangency, polarity, and affine normalization	811
N-20220719		813
126.1	Why solid geometry feels different	813
126.2	Basic axioms of incidence	813
126.3	Relations between two lines	813
126.4	A line and a plane	814
126.5	Parallel planes and parallel lines	814
126.6	Perpendicularity	814
126.7	Projection theorem	814
126.8	A coordinate viewpoint	815
126.9	Skew lines	815
126.10	Planes by normal vectors	815
126.11	The organizing idea	815
126.12	Incidence geometry	816
126.13	Orthogonality as inner-product structure	816
126.14	Projection as least-distance decomposition	816
126.15	Affine incidence plus metric projection	816
126.16	Affine geometry of points, lines, and planes	816
126.17	Projective completion	817
126.18	Grassmannians	817
126.19	Exterior algebra and incidence	817
126.20	Returning to solid geometry	817
126.21	Incidence via ranks	817
126.22	Projective duality of points and planes	818
N-20220802		819
127.1	Angles and radians	819
127.2	Sine, cosine, and tangent	819
127.3	Special angles	820
127.4	Addition formulas	820
127.5	Double-angle and half-angle formulas	820
127.6	Graphs of trigonometric functions	820
127.7	Triangle formulas	821
127.8	Radians are not a convention only	821
127.9	Addition formulas as rotation composition	821

127.10	From technique to structure	821
127.11	Radians and the exponential map	822
127.12	Addition formulas from rotation matrices	822
127.13	Triangle formulas and invariant quantities	822
127.14	Trigonometry as rotation geometry	822
127.15	Trigonometry and rotations	822
127.16	Complex exponentials	823
127.17	Spherical trigonometry	823
127.18	Harmonic analysis	823
127.19	Trigonometric identities as algebra	823
127.20	Curvature and spherical trigonometry	824
N-20220911		825
128.1	Line formulas	825
128.2	Parallel and perpendicular lines	825
128.3	Direction and normal vectors	826
128.4	Angles, distances, and parallel lines	826
128.5	Circle equations	827
128.6	A fractional-linear range problem	827
128.7	A graph-slope counting problem	828
128.8	A moving-line maximum	828
128.9	Triangle coordinate problem	829
128.10	Coordinate tricks as changes of viewpoint	829
128.11	Power of a point	829
128.12	The structural moral	829
128.13	Lines as affine equations	830
128.14	Circles and quadratic forms	830
128.15	Fractional-linear transformations	830
128.16	Analytic geometry through algebraic invariants	830
128.17	Conics as quadratic forms	830
128.18	Affine transformations	831
128.19	Projective viewpoint on conics	831
128.20	Duality	831
128.21	Returning to analytic geometry	831
128.22	Distance formulas as dual norms	831
128.23	Conics and degenerations	832
N-20230119		833
129.1	Why this geometry chapter matters	833
129.2	Ceva's theorem	833
129.3	Menelaus' theorem	834
129.4	Ptolemy's theorem	834
129.5	Generalized Ptolemy inequality	834
129.6	Simson line	835
129.7	Euler line and Euler formula	835
129.8	Why the pattern matters	836
129.9	An invariant map for the theorem list	836
129.10	Barycentric coordinates	837
129.11	Projective transformations	837
129.12	Ptolemy and inversive geometry	837
129.13	Elliptic and hyperbolic analogues	837
129.14	Returning to the theorem list	837
129.15	Ceva and Menelaus as dual affine statements	838
129.16	Euler line and affine covariance	838
129.17	Simson line through projection	838

N-20230220	839
130.1 Motivation	839
130.2 Sine difference formula by geometry	839
130.3 Rotation-matrix proof	839
130.4 Euler formula proof	840
130.5 Tangent formula	840
130.6 The method behind the trick	840
130.7 Rotation proof as the structural proof	840
130.8 Euler formula and characters	840
130.9 Tangent formula as projective coordinate	841
130.10 Addition formulas as rotation composition	841
N-20230223	843
131.1 Correcting a trigonometric mistake	843
131.2 A trigonometric system	843
131.3 A rational-parametrization idea	843
131.4 A direct algebraic route	844
131.5 Multiple-angle formulas	844
131.6 The determinant recurrence	845
131.7 From technique to structure	845
131.8 Trigonometric systems as algebraic systems	845
131.9 Rational parametrization of the circle	845
131.10 Multiple-angle formulas and Chebyshev polynomials	845
131.11 Determinant recurrences	846
131.12 Trigonometry through algebraic parametrization	846
N-20240518	847
132.1 The annoyance of parallel lines	847
132.2 Homogeneous coordinates are ratios	847
132.3 A coordinate scene: where parallel lines meet	847
132.4 Duality: points and lines speak the same language	848
132.5 Transformations and what survives	848
132.6 A first invariant: the cross-ratio	848
132.7 Conics as one projective family	848
132.8 The projective plane as all lines through the origin	849
132.9 Desargues as a projective-flavored theorem	849
132.10 Projective transformations in coordinates	849
132.11 The conic as a disguised projective line	849
132.12 Final checkpoint	850
N-20250211	851
133.1 The organizing question	851
133.2 Affine transformations	851
133.3 Conics in matrix form	851
133.4 Projective geometry	853
133.5 Cross-ratio	853
133.6 Projective duality and conics	854
133.7 Inversion	854
133.8 Möbius transformations	854
133.9 Conic sections	855
133.10 The Erlangen program	855
133.11 A hierarchy of invariants	855
133.12 Quadratic forms and matrix language	855
133.13 Affine classification versus projective classification	856
133.14 Polar lines	856
133.15 The structural moral	856
133.16 How to choose the right geometry	856

133.17	Using the transformation group in a conic problem	857
XIII	Curves and Structures	859
N-20220212		863
134.1	Structure beyond topology	863
134.2	From smooth structure to complex structure	864
134.3	Definition of a Riemann surface	864
134.4	Examples	864
134.5	Holomorphic maps between Riemann surfaces	866
134.6	Meromorphic functions as maps to the sphere	867
134.7	Degree and multiplicity	868
134.8	Local power form is not global	869
134.9	Holomorphic transition maps	869
134.10	Examples beyond the plane	870
134.11	Meromorphic functions	870
134.12	Conceptual synthesis	870
134.13	Complex structure is stronger than topology	870
134.14	Meromorphic functions and branched coverings	871
134.15	Riemann–Hurwitz	871
134.16	Complex curves through topology and ramification	871
N-20221214		873
135.1	The route of the note	873
135.2	Equidecomposability and dissection congruence	873
135.3	Hilbert’s third problem and the Dehn invariant	875
135.4	Measure-theoretic background	875
135.5	Tarski’s circle-squaring problem	875
135.6	Scissors congruence and why nice pieces fail	876
135.7	Laczkovich’s positive theorem	876
135.8	First idea: work in the torus	876
135.9	Random translations and a free action	877
135.10	The graph and bounded-distance bijection	877
135.11	Orbit matching	878
135.12	Discrepancy estimates	879
135.13	From Laczkovich to Borel circle squaring	879
135.14	Bridge to later ideas	880
135.15	Scissors congruence and invariants	880
135.16	Measurable versus nonmeasurable decompositions	880
135.17	Circle squaring and group actions	880
135.18	Equidecomposition under measure and group action	881
N-20230801		883
136.1	Starting point	883
136.2	Connections and curvature	883
136.3	Higgs bundles	883
136.4	The moduli space	884
136.5	The Hitchin fibration	884
136.6	Spectral curves	884
136.7	Geometric Langlands perspective	886
136.8	Synthesis	886
136.9	Moment-map reading of the equations	886
136.10	Spectral curves as global eigenvalues	886
136.11	Four descriptions to keep together	886
136.12	Source repair: Yang–Mills background	887
136.13	Source repair: dimensional reduction	887

136.14	Source repair: gauge invariance	887
136.15	Source repair: Higgs bundles and stability	888
136.16	Source repair: spectral data and Hitchin fibration	888
136.17	Source repair: geometric Langlands and S-duality	888
136.18	Source repair: local work still missing	889
N-20240608		891
137.1	The curve is not only the drawing	891
137.2	The function ring of a curve	891
137.3	When factorization becomes geometry	892
137.4	Smooth points and singular points	892
137.5	Tangent cones: the first nonzero shadow	892
137.6	The cusp as a ring	893
137.7	Intersections: one visible point may count twice	893
137.8	Projective closure repairs missing intersections	893
137.9	Bezout as a promise of clean accounting	893
137.10	Normalization: repairing a singular curve	894
137.11	Resultants: eliminating a variable	894
137.12	A projective calculation with a line and a conic	894
137.13	Real pictures and algebraically closed accounting	895
137.14	How to read a plane curve from its equation	895
N-20240914		897
138.1	A ledger for meromorphic functions	897
138.2	Principal divisors: the ledger of one function	897
138.3	Reading a divisor as permission	897
138.4	Line bundles: permission becomes geometry	898
138.5	The canonical divisor	898
138.6	Riemann–Roch without intimidation	899
138.7	Hurwitz as topology with branch points	899
138.8	A divisor game on the sphere	899
138.9	Degree as a first estimate	899
138.10	Canonical divisors through examples	900
138.11	Why Riemann–Roch has a correction term	900
138.12	Final caution	900
N-20241026		901
139.1	Starting from functions, not pictures	901
139.2	Coordinate rings: a space through its functions	901
139.3	Nullstellensatz: the dictionary becomes reliable	901
139.4	The Zariski topology is coarse on purpose	902
139.5	Distinguished opens and localization	902
139.6	Sheaves: local data with a gluing rule	902
139.7	Stalks: looking through a microscope	902
139.8	A local example: two axes crossing	903
139.9	Why the strange topology helps	903
139.10	Products and unions through rings	903
139.11	Radical ideals and forgotten nilpotents	903
139.12	Regular functions versus continuous functions	904
139.13	Morphisms as pullback of functions	904
139.14	A compact dictionary	904
N-20250118		905
140.1	A new kind of point	905
140.2	The affine line spectrum	905
140.3	Closed sets and distinguished opens	905
140.4	Localization is zooming in	906

140.5	The structure sheaf	906
140.6	Nilpotents: invisible but geometric	906
140.7	The arithmetic curve	906
140.8	Maximal ideals are not enough	907
140.9	Specialization as movement from generic to concrete	907
140.10	Residue fields	907
140.11	A tiny nonreduced calculation	907
140.12	Reading $\text{Spec } R$ in examples	908
N-20250301		909
141.1	Projective geometry returns with functions attached	909
141.2	The projective line as the model	909
141.3	Graded rings and homogeneous data	909
141.4	Projective space as Proj	909
141.5	The projective plane through three windows	910
141.6	A conic across charts	910
141.7	Projective closure with scheme memory	910
141.8	What Proj adds	911
141.9	Why the irrelevant ideal is excluded	911
141.10	Standard opens as spectra	911
141.11	A second conic chart calculation	911
141.12	Line bundles are already nearby	912
141.13	Closing the geometry arc	912
N-20250705		913
142.1	First orientation	913
142.2	The topological object and its planar shadow	913
142.3	Some basic knot-theory language	914
142.4	The pieces of a diagram	916
142.5	Nikonov's bridge: diagram elements as probes	917
142.6	Invariants, coinvariants, and labels	918
142.7	What this note is not trying to do	919
N-20250707		921
143.1	Reading only one small part of a long paper	921
143.2	Where this sits inside Nikonov's paper	921
143.3	A region is not just a blank patch	922
143.4	At a crossing, four regions want to talk to each other	923
143.5	Why the operation is partial	924
143.6	Probe clearing and color rules	925
143.7	What about crossings themselves?	926
143.8	The paper in one playful sentence	927
XIV	Topology and Homotopy	929
N-20210806		933
144.1	Purpose of the note	933
144.2	Group actions	933
144.3	Homotopy	934
144.4	Cofibrations, fibrations, and deformation retracts	935
144.5	Spheres and homotopy groups	938
144.6	The Hopf fibration	939
144.7	References to return to	940
144.8	Group actions as symmetry bookkeeping	941
144.9	Hopf fibration intuition	941

N-20221218	943
145.1 Why the note starts with metrics	943
145.2 Boundedness: diameter and radius	943
145.3 Metric spaces	943
145.4 Sequences and continuity	944
145.5 Standard metrics	944
145.6 Norms and inner products	945
145.7 Open and closed balls	945
145.8 Open sets, closed sets, and limit points	945
145.9 Cauchy sequences and completeness	946
145.10 Total boundedness and compactness	946
145.11 Beyond the first calculation	946
145.12 Equivalent metrics and what topology forgets	946
145.13 Completion as adding ideal limit points	947
145.14 Uniform structure hidden inside metric topology	947
145.15 Compactness through total boundedness	947
145.16 Baire category as a completeness phenomenon	948
145.17 Metric spaces through completion and category	948
N-20221221	949
146.1 Completion	949
146.2 Let the buyer beware	949
146.3 Homeomorphisms and embeddings	949
146.4 Quotient intuition	950
146.5 A more structural view	950
146.6 Completion and its universal property	950
146.7 Homeomorphism versus uniform equivalence	950
146.8 Continuous bijections that are not homeomorphisms	950
146.9 Quotient topology as final topology	951
146.10 Endpoint gluing and local neighborhoods	951
146.11 Completion, gluing, and quotient control	951
N-20221222	953
147.1 Re-defining continuity	953
147.2 Closed sets and closure	953
147.3 Subspaces	953
147.4 Connectedness	953
147.5 Path-connectedness	954
147.6 Continuity: why preimages appear	954
147.7 Closure through neighborhoods	954
147.8 Connectedness by continuous images	954
147.9 Path-connectedness and components	955
147.10 Next viewpoint	955
147.11 Continuity as a contravariant test	955
147.12 Closure as an operator	955
147.13 Why sequences are not enough	955
147.14 Connectedness as maps to a two-point space	956
147.15 Path components versus connected components	956
147.16 Local connectedness as the missing hypothesis	956
147.17 Continuity, closure, and connectedness	956
N-20221223	957
148.1 Homotopy intuition	957
148.2 Simply connected spaces	957
148.3 Semigroups, monoids, and groups	957
148.4 Congruence, quotients, and homomorphisms	957
148.5 Loops and the beginning of the fundamental group	958

148.6	Why the punctured plane has nontrivial loops	958
148.7	Homomorphisms as structure-preserving maps	958
148.8	Quotients in topology and algebra	959
148.9	Higher viewpoint: from holes to algebra	959
148.10	A useful afterword	959
148.11	Concatenation of loops and why homotopy classes form a group	959
148.12	The exponential covering and the integer invariant	959
148.13	Punctured plane deformation retracts to the circle	960
148.14	Functoriality: maps of spaces give homomorphisms	960
148.15	Van Kampen as a computation machine	960
148.16	Group completion and why inverses matter	961
148.17	Loops, coverings, and algebraic invariants	961
N-20221228		963
149.1	Interval modulo endpoints	963
149.2	More quotient spaces	963
149.3	Definition of quotient topology	963
149.4	Product topology	964
149.5	Universal property of quotient topology	964
149.6	Endpoint identification in detail	964
149.7	Product topology via projections	964
149.8	Quotient and product as opposite directions	965
149.9	A second lens	965
149.10	Saturated sets and quotient openness	965
149.11	Quotients can destroy Hausdorffness	965
149.12	Product topology and the mapping property	965
149.13	From squares to surfaces through edge words	966
149.14	Products and quotients need not commute blindly	966
149.15	Products, quotients, and gluing data	966
N-20221229		967
150.1	What this note is really doing	967
150.2	The torus as a parametrized product	967
150.3	Disjoint union	968
150.4	Wedge sum	968
150.5	CW complexes: building spaces by cells	968
150.6	The torus as a quotient of a square	969
150.7	Cosets as algebraic quotients	970
150.8	The Klein bottle	971
150.9	Polygon schemata for surfaces	971
150.10	Projective spaces as quotient spaces	972
150.11	Cellular boundary words and fundamental groups	972
150.12	Why polygon schemata are efficient	973
150.13	Further context	973
150.14	Cellular chain complex of the torus	973
150.15	Klein bottle: same cells, different boundary	974
150.16	Euler characteristic as cell bookkeeping	974
150.17	Orientability from polygon words	974
150.18	Van Kampen and presentations from gluing	974
150.19	Projective spaces through cells	975
150.20	Cellular surfaces through homology and presentations	975
N-20240727		977
151.1	What this chapter is for	977
151.2	The first analytic reflex	977
151.3	Why the Hopf and K-theory hints changed the search	978
151.4	The Eilenberg swindle as the central guess	979

151.5	How this looked like a K-theory statement	980
151.6	The role I imagined for the Hopf fibration	980
151.7	The meme as a learning ladder	981
151.8	What the original interpretation claimed	982
151.9	A slightly uneasy stopping point	982
N-20251217		983
152.1	Why this second note exists	983
152.2	The lecture starts from ordinary equations	984
152.3	The Hopf fibration as the first nontrivial map	984
152.4	Homotopy groups and the first stable class	985
152.5	From natural numbers to finite sets	986
152.6	The space of finite sets	986
152.7	The K-space of finite signed sets	987
152.8	The loop h	987
152.9	Why the first guess was close but reversed	988
152.10	Finite sets and stable spheres	989
152.11	Relation with the classical Hopf map	990
152.12	Cobordism: the movie becomes a circle	990
152.13	Algebraic K-theory of a field	991
152.14	Intrinsic skein-lasagna	992
152.15	The V-k cobordism group	992
152.16	Generalized Pontryagin–Thom	993
152.17	The K2 slide and Steinberg relations	994
152.18	The final lasagna diagrams	995
152.19	What the equation finally says	996
152.20	Structural viewpoint	997
XV	Machine Learning	999
N-20250926		1003
153.1	The setting	1003
153.2	The theoretical tensor	1003
153.3	Computing the sparse derivative	1004
153.4	Contraction by the chain rule	1004
153.5	Why the tensor disappears	1005
153.6	Contravariant reading of the same calculation	1005
153.7	Outer-product expression	1006
153.8	Numerical example	1006
153.9	Einstein notation proof	1006
153.10	Kronecker product viewpoint	1007
153.11	Softmax–cross-entropy simplification	1007
153.12	A wider interpretation	1007
153.13	Backpropagation as cotangent pullback	1007
153.14	Why the three-dimensional tensor collapses	1007
153.15	Batch form	1008
153.16	Backpropagation as sparse tensor contraction	1008
N-20251008		1009
154.1	Why probability enters generative modeling	1009
154.2	Random variables, distributions, and densities	1009
154.3	Joint, marginal, and conditional distributions	1010
154.4	Bayes’ rule as inversion of a generative story	1010
154.5	Gaussian distributions	1010
154.6	Expectation and Monte Carlo approximation	1011
154.7	KL divergence	1011

154.8	Jensen's inequality	1012
154.9	The next abstraction	1012
N-20251009		1013
155.1	From autoencoders to latent variable models	1013
155.2	The generative story	1013
155.3	Graphical model notation	1014
155.4	Marginal likelihood	1014
155.5	The posterior problem	1015
155.6	Why naive Monte Carlo is not enough	1015
155.7	The recognition model	1015
155.8	Amortized inference	1016
155.9	The VAE model diagram	1016
155.10	Remaining derivation	1016
N-20251010		1017
156.1	The approximate posterior as an optimization problem	1017
156.2	The ELBO identity	1017
156.3	Jensen's inequality derivation	1018
156.4	Reconstruction term and regularization term	1018
156.5	The ELBO gap	1019
156.6	Closed-form KL for diagonal Gaussians	1019
156.7	Negative ELBO as a loss	1020
156.8	Empirical lower bounds	1020
156.9	Information-theoretic reading	1021
156.10	Limits of the ELBO	1021
156.11	What survives later	1021
N-20251011		1023
157.1	The gradient problem	1023
157.2	Score-function estimator and pathwise estimator	1023
157.3	Gaussian reparameterization	1024
157.4	The stochastic ELBO estimator	1024
157.5	Why implementations output log variance	1024
157.6	Training as ELBO maximization	1025
157.7	Decoder likelihoods and reconstruction losses	1025
157.8	Generation after training	1025
157.9	Empirical marginal likelihood curves	1026
157.10	Limitations and extensions	1026
157.11	How the pieces fit	1027
N-20260302		1029
158.1	Backpropagation as cotangent pullback	1029
158.2	Geometric data analysis	1030
158.3	Manifold hypothesis and embeddings	1030
158.4	Optimal transport and stochastic dynamics	1031
158.5	Symmetry, equivariance, and category-like thinking	1032
158.6	Further direction	1032
158.7	Backpropagation and differential geometry	1033
158.8	Manifold hypothesis and representation learning	1033
158.9	Symmetry and equivariance	1033
158.10	Optimal transport and distributions	1033
158.11	Mathematical structures for learning systems	1033
158.12	Rebuilt reading path: what pure mathematics contributes	1033
158.13	Topological data analysis in usable form	1034
158.14	Wasserstein geometry and stochastic dynamics	1034
158.15	Category language without exaggeration	1034

I

Sets and Foundations

Part I: Contents

N-20210204	5
1.1 Why classes enter before functions	5
1.2 Russell's class as the first warning	5
1.3 Power sets and indexed families	5
1.4 Set identities	6
1.5 Images, inverse images, and functions	7
1.6 One-sided inverses	8
1.7 Set-level bookkeeping behind functions	8
N-20210206	9
2.1 Relations and equivalence classes	9
2.2 Partitions	10
2.3 Indexed products and projections	10
2.4 The first universal property	11
2.5 Integers and induction	11
2.6 Division, gcd, and congruence	12
N-20210207	15
3.1 Choice enters the story	15
3.2 Partially ordered sets	15
3.3 Well-ordering	16
3.4 Equipollence and cardinal numbers	16
3.5 Cantor's theorem	17
3.6 Comparing cardinals	17
3.7 Infinite cardinals and finite disturbance	17
3.8 Concrete encodings behind the arithmetic	19
N-20220131	21
4.1 Why this translation/note matters	21
4.2 The discomfort with ZFC	21
4.3 Three misunderstandings about ETCS	22
4.4 Elements as special functions	22
4.5 The informal ETCS-style principles	23
4.6 Subsets via characteristic functions	23
4.7 Products, exponentials, and inverse images	24
4.8 Choice as right inverses	24
4.9 What I take from the article	24
4.10 Translation note	25
4.11 ETCS intuition: membership is not primitive practice	25
4.12 Element reasoning inside categorical set theory	25
4.13 Why this matters philosophically	25
N-20220205	27
5.1 Why there are several axiomatizations of set theory	27
5.2 The main systems	27
5.3 ZFC	27
5.4 NBG	27
5.5 Kelley–Morse set theory	28
5.6 New Foundations	28
5.7 How to compare the systems	29
5.8 What remains unclear	29
5.9 Where the idea leads	29
5.10 Why classes are introduced	29
5.11 New Foundations and stratification	30
5.12 Comparing foundational choices	30

N-20220710	31
6.1 Plane vectors	31
6.2 Field axioms	31
6.3 Ordered fields	32
6.4 Elementary consequences of the order axioms	32
6.5 Embedding $\mathbb{Q}$ into $\mathbb{R}$	33
6.6 Nested intervals	33
6.7 Ordered fields and the real line	33
6.8 Vectors as algebraicized geometry	34
6.9 The conceptual payoff	34
6.10 Ordered fields and what they cannot prove	34
6.11 Archimedean property and density	34
6.12 Completeness as the real-number axiom	34
6.13 From ordered algebra to the real line	34
N-20220713	35
7.1 Large and small positive real numbers	35
7.2 The geometric view of the real line	35
7.3 Supremum property	35
7.4 Nested intervals imply completeness	36
7.5 Dedekind cuts	36
7.6 Metric spaces	36
7.7 Subspaces and density	37
7.8 Supremum as least upper bound	37
7.9 Dedekind cuts and completeness	38
7.10 Looking ahead	38
7.11 Supremum as order-completeness	38
7.12 Cauchy completion and metric completion	38
7.13 Dedekind cuts versus Cauchy sequences	38
7.14 Completeness as the no-gap principle	38
N-20230225	39
8.1 Why this topic appears	39
8.2 The real-number proof	39
8.3 Hyperreal numbers	39
8.4 Infinitesimals and finite hyperreals	40
8.5 Hyperfinite strings of nines	40
8.6 Transfer principle	41
8.7 Derivative through infinitesimals	41
8.8 Integral through hyperfinite sums	42
8.9 What nonstandard analysis is not saying	42
8.10 Source repair: internal and external sets	42
8.11 Source repair: standard part and Loeb measure	42
8.12 Source repair: transfer has limits	43
8.13 A final synthesis	43
8.14 Real equality and limiting process	43
8.15 Hyperreals and infinitesimal remainders	43
8.16 Transfer and caution	43
8.17 Decimals, limits, and infinitesimal models	43
N-20230710	45
9.1 The question of foundations	45
9.2 Set theory	45
9.3 Category theory	46
9.4 Type theory	48
9.5 Homotopy type theory and univalence	48
9.6 Comparative analysis	49
9.7 Modern applications	49
9.8 Toward a unified perspective	49
9.9 Conclusion	50
9.10 Equality in the three languages	50

9.11	Universal properties versus encodings	50
9.12	Computation and formal proof	50

Sets, Classes, and the First Warnings About Size

N-20210204

§1.1 Why classes enter before functions

The point of this note is to set up the basic language of sets, classes, indexed families, and functions. The original PDF begins very much in the style of a reading note: it is following the opening sections of a book, but it also pauses at the places where the distinction between a set and a proper class starts to feel less innocent.

Definition 1.1.1. A **set** is a class which can be an element of another class. A class which is not a set is called a **proper class**.

This is not meant to be the whole foundation of set theory, but it explains the working convention of the note: some collections are small enough to be members of other collections, while others are too large and must remain proper classes.

Caution

The original note says that the distinction between sets and proper classes is “not too clear”. Informally this is understandable, but formally this distinction is exactly what prevents the usual paradoxes. In a class theory such as NBG, sets are precisely the classes that may occur as elements of classes; proper classes may not be elements.

§1.2 Russell’s class as the first warning

The first serious example is Russell’s construction. Consider

$$M = \{X \mid X \text{ is a set and } X \notin X\}.$$

At first sight this is a natural class: many sets are not members of themselves. For instance, the set of all books is not itself a book.

Proposition 1.2.1

The class M is not a set.

Proof. Assume, for contradiction, that M is a set. Then it makes sense to ask whether $M \in M$. By the defining condition of M ,

$$M \in M \iff M \notin M,$$

which is impossible. Hence M must be a proper class. $\square$

This is the first place where the note is not merely listing definitions. The example shows why unrestricted comprehension cannot simply produce a set from every condition.

§1.3 Power sets and indexed families

Definition 1.3.1. The **power set axiom** asserts that for every set A , the class

$$\mathcal{P}(A) = \{X \mid X \subseteq A\}$$

of all subsets of A is itself a set. It is also common to write 2^A for this power set.

The next construction is an indexed family. Given a family of classes or sets

$$\{A_i \mid i \in I\},$$

its union and intersection are written

$$\bigcup_{i \in I} A_i = \{x \mid x \in A_i \text{ for some } i \in I\},$$

and

$$\bigcap_{i \in I} A_i = \{x \mid x \in A_i \text{ for all } i \in I\}.$$

Remark 1.3.2 — The original note explicitly points out that the indexing collection I need not be finite or countable. This is a useful mental shift: the notation $\{A_i\}_{i \in I}$ is not just a long list. It may be a genuinely infinite family.

Caution

If I is allowed to be a proper class, then extra care is needed. In ordinary ZFC, one usually speaks of set-indexed families. In a class-theoretic setting, class-indexed families can be discussed, but one must check separately when their unions or intersections are sets.

If A and B are classes, the relative complement of A in B is

$$B \setminus A = \{x \mid x \in B \text{ and } x \notin A\}.$$

If all classes under discussion are subsets of a fixed universe U , then $U \setminus A$ is often denoted A^c .

§1.4 Set identities

The original note lists the standard identities and leaves the proof as easy. A cleaned-up version is:

Proposition 1.4.1

Let A , B , and B_i be subsets of a fixed universe U . Then

$$A \cap \left(\bigcup_{i \in I} B_i \right) = \bigcup_{i \in I} (A \cap B_i),$$

$$A \cup \left(\bigcap_{i \in I} B_i \right) = \bigcap_{i \in I} (A \cup B_i),$$

$$\left(\bigcup_{i \in I} A_i \right)^c = \bigcap_{i \in I} A_i^c, \quad \left(\bigcap_{i \in I} A_i \right)^c = \bigcup_{i \in I} A_i^c.$$

Moreover, $A \cup B = B \cup A$, $A \subseteq B$ means every element of A lies in B , and $A \cap B = A$ when $A \subseteq B$.

Proof. All of these follow by chasing membership. For instance, an element x lies in $A \cap (\bigcup_i B_i)$ exactly when $x \in A$ and $x \in B_i$ for at least one i , which is exactly the condition that $x \in \bigcup_i (A \cap B_i)$. The De Morgan laws are proved similarly after replacing “there exists” by “for all” under negation. $\square$

§1.5 Images, inverse images, and functions

Let $f : A \rightarrow B$ be a function. If $S \subseteq A$, the image of S under f is

$$f(S) = \{b \in B \mid b = f(a) \text{ for some } a \in S\}.$$

If $T \subseteq B$, the inverse image of T is

$$f^{-1}(T) = \{a \in A \mid f(a) \in T\}.$$

Example 1.5.1

For $f : \mathbb{R} \rightarrow \mathbb{R}$, $f(x) = x^2$, we have

$$f(\mathbb{R}) = [0, \infty), \quad f^{-1}(\{9\}) = \{-3, 3\}.$$

A map $f : A \rightarrow B$ is **injective** if $f(a) = f(a')$ implies $a = a'$. It is **surjective** if $f(A) = B$. It is **bijective** if it is both injective and surjective.

Proposition 1.5.2

For composable maps $A \xrightarrow{f} B \xrightarrow{g} C$:

- (i) if f and g are injective, then $g \circ f$ is injective;
- (ii) if f and g are surjective, then $g \circ f$ is surjective;
- (iii) if $g \circ f$ is injective, then f is injective;
- (iv) if $g \circ f$ is surjective, then g is surjective.

§1.6 One-sided inverses

The last theorem in the original note is about recognizing injective and surjective maps through one-sided inverses.

Theorem 1.6.1

Let $f : A \rightarrow B$ be a function and assume $A \neq \emptyset$.

- (i) f is injective if and only if there is a map $g : B \rightarrow A$ such that $g \circ f = 1_A$.
- (ii) f is surjective if and only if there is a map $h : B \rightarrow A$ such that $f \circ h = 1_B$, provided the required choices of elements from the fibres of f can be made.

Proof. For (i), suppose f is injective. For $b \in f(A)$, there is a unique $a \in A$ with $f(a) = b$. Choose once and for all an element $a_0 \in A$. Define

$$g(b) = \begin{cases} a, & b \in f(A) \text{ and } f(a) = b, \\ a_0, & b \notin f(A). \end{cases}$$

Then $g(f(a)) = a$ for all $a \in A$. Conversely, if $g \circ f = 1_A$, then $f(a) = f(a')$ implies

$$a = g(f(a)) = g(f(a')) = a',$$

so f is injective.

For (ii), if h exists and $f \circ h = 1_B$, then every $b \in B$ equals $f(h(b))$, so f is surjective. Conversely, if f is surjective, each fibre $f^{-1}(b)$ is nonempty. Choosing $h(b) \in f^{-1}(b)$ gives a right inverse. $\square$

Caution

The last choice of one element from each fibre is exactly where the Axiom of Choice may enter. The original note itself flags this point in a footnote, and it will reappear in the later note on choice and cardinal numbers.

§1.7 Set-level bookkeeping behind functions

The calculations in this note are elementary only after the ambient set has been fixed. Russell's class shows why the notation

$$\{x \mid P(x)\}$$

cannot be used without asking whether the collection is actually a set. Once a set A is available, the usual constructions are safe: subsets live in $\mathcal{P}(A)$, indexed unions and intersections are formed from a family $(A_i)_{i \in I}$, and membership proofs reduce identities such as De Morgan's laws to quantifier manipulations.

For functions the same bookkeeping reappears in a more operational form. The inverse image $f^{-1}(T)$ is always defined for $T \subseteq B$, while the image $f(S)$ depends on the part of the domain actually hit. Injectivity is detected by a left inverse $g \circ f = 1_A$; surjectivity is detected by a right inverse $f \circ h = 1_B$. The construction of g for an injection uses only a single default element of A , whereas the construction of h for a surjection asks for one chosen element from each fibre $f^{-1}(b)$. That fibre-by-fibre choice is the exact point to remember when the later notes discuss the Axiom of Choice and cardinal comparison.

Relations, Products, and the First Appearance of Universal Properties

N-20210206

§2.1 Relations and equivalence classes

This note continues the foundational setup, moving from sets and functions to relations, products, and the integers. The original title says “Integers”; I will silently correct this to *integers*.

Definition 2.1.1. For sets or classes A and B , their Cartesian product is

$$A \times B = \{(a, b) \mid a \in A, b \in B\}.$$

A **relation from A to B** is a subclass of $A \times B$. A **relation on A** is a subclass of $A \times A$.

Caution

The original note has the phrase “a subset of R is a relation on $A \times B$ ”, which is clearly a slip. The intended statement is that a relation is a subset of $A \times B$.

Definition 2.1.2. A relation $R \subseteq A \times A$ is an **equivalence relation** if it is:

- (i) reflexive: $(a, a) \in R$ for all $a \in A$;
- (ii) symmetric: $(a, b) \in R$ implies $(b, a) \in R$;
- (iii) transitive: $(a, b), (b, c) \in R$ implies $(a, c) \in R$.

When R is an equivalence relation, we write $a \sim b$ for $(a, b) \in R$, and define the equivalence class of a by

$$[a] = \{x \in A \mid x \sim a\}.$$

The quotient class is

$$A/R = \{[a] \mid a \in A\}.$$

Proposition 2.1.3

If R is an equivalence relation on A , then:

- (i) $[a] \neq \emptyset$ for every $a \in A$;
- (ii) $A = \bigcup_{a \in A} [a] = \bigcup_{C \in A/R} C$;
- (iii) $[a] = [b]$ if and only if $a \sim b$.

Lemma 2.1.4

For any $a, b \in A$, either $[a] \cap [b] = \emptyset$, or $[a] = [b]$.

Proof. If $x \in [a] \cap [b]$, then $x \sim a$ and $x \sim b$. By symmetry and transitivity, $a \sim b$, and hence any element equivalent to a is equivalent to b , and conversely. Thus $[a] = [b]$. $\square$

§2.2 Partitions

Definition 2.2.1. Let A be a nonempty class. A family $\{A_i \mid i \in I\}$ of subclasses of A is a **partition** of A if

- (i) $A_i \neq \emptyset$ for every $i \in I$;
- (ii) $\bigcup_{i \in I} A_i = A$;
- (iii) $A_i \cap A_j = \emptyset$ whenever $i \neq j$.

Remark 2.2.2 — This is one of the most useful ways to remember equivalence relations: an equivalence relation cuts A into disjoint blocks, and those blocks form a partition. Conversely, a partition defines an equivalence relation by declaring two elements equivalent exactly when they lie in the same block.

§2.3 Indexed products and projections

The note then generalizes finite Cartesian products. For a set-indexed family $\{A_i \mid i \in I\}$, define

$$\prod_{i \in I} A_i$$

to be the set of functions

$$f : I \rightarrow \bigcup_{i \in I} A_i$$

such that $f(i) \in A_i$ for every $i \in I$.

When $I = \{1, 2, \dots, n\}$, this agrees with the usual product

$$A_1 \times A_2 \times \cdots \times A_n,$$

because such a function is the same information as the ordered n -tuple

$$(f(1), f(2), \dots, f(n)).$$

Remark 2.3.1 — The original note pauses over the case $I = \{1, 2\}$. A point of $A_1 \times A_2$ may be written as a pair (a_1, a_2) , but it may also be regarded as a function f with $f(1) = a_1$ and $f(2) = a_2$. This is a small but important shift: products are not just lists; they are choice-like functions over an index set.

For each $k \in I$, the **canonical projection**

$$\pi_k : \prod_{i \in I} A_i \rightarrow A_k$$

is defined by

$$\pi_k(f) = f(k).$$

Example 2.3.2

If $I = \mathbb{R}$ and $A_i = \mathbb{C}$ for all i , then an element of $\prod_{i \in I} A_i$ is a function $f : \mathbb{R} \rightarrow \mathbb{C}$. For example, $f(x) = ix$. The projection at $k \in \mathbb{R}$ returns

$$\pi_k(f) = f(k) = ik.$$

§2.4 The first universal property

Here the original note becomes excited: “Appears! Universal property!” This is worth preserving. Products are not only sets of tuples; they are characterized by how maps into them behave.

Theorem 2.4.1 (Universal property of products)

Let $\{A_i \mid i \in I\}$ be a family of sets. There is a set D , equipped with maps

$$\pi_i : D \rightarrow A_i \quad (i \in I),$$

such that for every set C and every family of maps $f_i : C \rightarrow A_i$, there is a unique map

$$f : C \rightarrow D$$

with

$$\pi_i \circ f = f_i$$

for every $i \in I$. Moreover, D is unique up to a unique bijection compatible with the projections.

Proof. Take $D = \prod_{i \in I} A_i$. Given maps $f_i : C \rightarrow A_i$, define

$$f(c)(i) = f_i(c).$$

Then $f(c) \in \prod_i A_i$, and $\pi_i(f(c)) = f_i(c)$. Uniqueness follows because a point of the product is determined by all of its coordinates. $\square$

Remark 2.4.2 — The original note says that the proof is a little long and should maybe be postponed until systematic study of category theory. The cleaned-up proof above is short because we already know the concrete product. The deeper point is that the same pattern will later define products in arbitrary categories, where elements and coordinates may no longer be available.

§2.5 Integers and induction

The note then turns to the natural numbers and the integers through Peano’s axioms. In modern notation the Peano axioms say, roughly:

- (i) 0 is a natural number;
- (ii) every natural number has a successor;
- (iii) 0 is not the successor of any natural number;

- (iv) distinct natural numbers have distinct successors;
- (v) any property true of 0 and inherited by successors is true of all natural numbers.

The fifth axiom is the principle of mathematical induction.

Theorem 2.5.1 (Induction and strong induction)

Let $S \subseteq \mathbb{N}$ with $0 \in S$. If either

- (i) $n \in S$ implies $n + 1 \in S$ for all $n \in \mathbb{N}$, or
- (ii) $\{0, 1, \dots, n - 1\} \subseteq S$ implies $n \in S$ for all $n \in \mathbb{N}$,

then $S = \mathbb{N}$.

Caution

The original note proves this by taking the least element of $\mathbb{N} \setminus S$. That proof is clean if the well-ordering principle for $\mathbb{N}$ is already available. If one is deriving the theory strictly from Peano's axioms, one should first justify the well-ordering principle or prove the statement directly by induction.

§2.6 Division, gcd, and congruence

Theorem 2.6.1 (Division algorithm)

If $a, b \in \mathbb{Z}$ and $a \neq 0$, then there exist unique integers q, r such that

$$b = aq + r, \quad 0 \leq r < |a|.$$

Definition 2.6.2. For $a \neq 0$, we write $a \mid b$ if there exists $k \in \mathbb{Z}$ such that $ak = b$. Otherwise we write $a \nmid b$.

Definition 2.6.3. The positive integer c is the greatest common divisor of $a_1, a_2, \dots, a_n$ if

- (i) $c \mid a_i$ for every $1 \leq i \leq n$;
- (ii) whenever $d \in \mathbb{Z}$ and $d \mid a_i$ for every i , then $d \mid c$.

It is denoted $(a_1, a_2, \dots, a_n)$.

Definition 2.6.4. Let $m > 0$. If $a, b \in \mathbb{Z}$ and $m \mid (a - b)$, then a is congruent to b modulo m , written

$$a \equiv b \pmod{m}.$$

Theorem 2.6.5

Let $m > 0$ and $a, b, c, d \in \mathbb{Z}$.

- (i) Congruence modulo m is an equivalence relation on $\mathbb{Z}$, with precisely m equivalence classes.
- (ii) If $a \equiv b \pmod{m}$ and $c \equiv d \pmod{m}$, then

$$a + c \equiv b + d \pmod{m}, \quad ac \equiv bd \pmod{m}.$$

- (iii) If $ab \equiv ac \pmod{m}$ and $(a, m) = 1$, then $b \equiv c \pmod{m}$.

Proof. The original note leaves this as an exercise. The key point is divisibility. For example,

$$m \mid (a - b), \quad m \mid (c - d)$$

imply

$$m \mid (a + c) - (b + d).$$

For multiplication, write

$$ac - bd = c(a - b) + b(c - d).$$

For cancellation, $ab \equiv ac \pmod{m}$ means $m \mid a(b - c)$. If $(a, m) = 1$, Euclid's lemma gives $m \mid (b - c)$, hence $b \equiv c \pmod{m}$. $\square$

Choice, Well-Ordering, and the Arithmetic of Cardinal Size

N-20210207

§3.1 Choice enters the story

This note begins exactly where the previous chapter warned us: sometimes a construction requires choosing one element from each set in a family. The original note states the Axiom of Choice in product form.

Axiom 3.1.1 (Axiom of Choice)

The product of a family of nonempty sets indexed by a nonempty set is nonempty.

Remark 3.1.2 — Equivalently, if S is a set, then there is a choice function defined on the collection of all nonempty subsets of S :

$$f(A) \in A$$

for every nonempty $A \subseteq S$. This is the alternate version recorded in the original note.

§3.2 Partially ordered sets

Definition 3.2.1. A **partially ordered set** is a nonempty set A together with a relation $\leq$ which is reflexive, transitive, and antisymmetric. Antisymmetry means that

$$a \leq b \text{ and } b \leq a \implies a = b.$$

Two elements $a, b \in A$ are **comparable** if either $a \leq b$ or $b \leq a$. If every pair of elements is comparable, then the order is a **total order** or **linear order**.

Example 3.2.2

Let $A = \mathcal{P}(\{1, 2, 3, 4, 5\})$, ordered by inclusion. Then $C \leq D$ means $C \subseteq D$. This is a partial order, but not a total order: for example, $\{1, 2\}$ and $\{2, 3\}$ are not comparable.

Let $(A, \leq)$ be a partially ordered set. An element $a \in A$ is **maximal** if whenever $c \in A$ is comparable with a and $a \leq c$, then $c = a$. An **upper bound** of a nonempty subset $B \subseteq A$ is an element $d \in A$ such that $b \leq d$ for every $b \in B$. A nonempty subset $B \subseteq A$ is a **chain** if it is totally ordered by the inherited order.

Theorem 3.2.3 (Zorn's lemma)

Let A be a nonempty partially ordered set. If every chain in A has an upper bound in A , then A has a maximal element.

Caution

The original note says “local maximum”. The standard term is *maximal element*. It need not be a greatest element: there may be no element above all elements of A .

§3.3 Well-ordering

Definition 3.3.1. A partially ordered set A is **well ordered** if every nonempty subset of A has a least element.

Proposition 3.3.2 (Well-ordering principle)

Every nonempty set A admits a total ordering under which it is well ordered.

Remark 3.3.3 — The original note notes that this extends mathematical induction from $\mathbb{N}$ to arbitrary well-ordered sets. Conceptually this is the beginning of transfinite induction, even if the note does not yet develop ordinals.

§3.4 Equipollence and cardinal numbers

Definition 3.4.1. Two sets A and B are **equipollent**, written $A \sim B$, if there exists a bijection $A \rightarrow B$.

Definition 3.4.2. If a set A is equipollent to

$$I_n = \{1, 2, \dots, n\}$$

for some $n \in \mathbb{N}$, or to $I_0 = \emptyset$, then A is finite. Otherwise A is infinite.

Definition 3.4.3. The **cardinal number** of a set A , denoted $|A|$, is its equivalence class under equipollence. It is a finite or infinite cardinal according as A is finite or infinite.

Caution

Strictly speaking, the “equivalence class of all sets bijective with A ” is usually too large to be a set. A formal treatment either uses representatives such as initial ordinals, or works in a class theory where such collections are proper classes. The original note is using the usual informal convention.

If A and B are disjoint sets with $|A| = \alpha$ and $|B| = \beta$, define

$$\alpha + \beta = |A \cup B|.$$

Similarly, cardinal multiplication is defined by

$$\alpha\beta = |A \times B|.$$

§3.5 Cantor's theorem

Theorem 3.5.1 (Cantor)

For every set A ,

$$|A| < |\mathcal{P}(A)|.$$

Proof. The map $a \mapsto \{a\}$ is injective from A into $\mathcal{P}(A)$, so $|A| \leq |\mathcal{P}(A)|$.

Suppose for contradiction that $f : A \rightarrow \mathcal{P}(A)$ is surjective. Define

$$B = \{a \in A \mid a \notin f(a)\}.$$

Since $B \subseteq A$, surjectivity gives $a_0 \in A$ with $f(a_0) = B$. Then

$$a_0 \in B \iff a_0 \notin f(a_0) = B,$$

which is impossible. Hence no surjection $A \rightarrow \mathcal{P}(A)$ exists, and so $|A| < |\mathcal{P}(A)|$. $\square$

This proof deliberately echoes Russell's paradox from the first chapter. In one case the self-reference shows that a class is too large to be a set; in the other, diagonalization shows that the power set is strictly larger than the original set.

§3.6 Comparing cardinals

Theorem 3.6.1 (Cantor–Bernstein)

If A and B are sets such that $|A| \leq |B|$ and $|B| \leq |A|$, then

$$|A| = |B|.$$

Theorem 3.6.2 (Trichotomy for cardinals)

The class of cardinal numbers is linearly ordered by $\leq$. If α and β are cardinal numbers, then exactly one of the following holds:

$$\alpha < \beta, \quad \alpha = \beta, \quad \beta < \alpha.$$

Remark 3.6.3 — The original note calls this the Law of Trichotomy and mentions that, in the proof, one orders a family of functions by extension. This is one of the standard places where Zorn's lemma enters the comparison of cardinalities.

§3.7 Infinite cardinals and finite disturbance

The final part of the original note deliberately becomes a theorem list after the writing process was interrupted. I keep the statements, but the better way to read them is through explicit encodings: pairing functions and binary strings let us put our hands on the size of these infinite sets.

Theorem 3.7.1

Every infinite set has a denumerable subset. In particular,

$$\aleph_0 \leq \alpha$$

for every infinite cardinal α , assuming the usual choice principles being used in this chapter.

Lemma 3.7.2

If A is an infinite set and F is a finite set, then

$$|A \cup F| = |A|.$$

In particular,

$$\alpha + n = \alpha$$

for every infinite cardinal α and finite cardinal n .

Theorem 3.7.3

If $\alpha \leq \beta$ and β is infinite, then

$$\alpha + \beta = \beta.$$

Theorem 3.7.4

If $0 < \alpha \leq \beta$ and β is infinite, then

$$\alpha\beta = \beta.$$

In particular,

$$\aleph_0 \cdot \aleph_0 = \aleph_0,$$

and if B is finite then

$$\aleph_0 \cdot |B| = \aleph_0$$

provided $B \neq \emptyset$.

Theorem 3.7.5

Let A be a set and, for each integer $n \geq 1$, let

$$A^n = \underbrace{A \times A \times \cdots \times A}_{n \text{ factors}}.$$

If A is finite, then

$$|A^n| = |A|^n,$$

and if A is infinite, then

$$|A^n| = |A|.$$

Moreover,

$$\left| \bigcup_{n \in \mathbb{N}} A^n \right| = \aleph_0 |A|.$$

Corollary 3.7.6

If A is an infinite set and $\mathcal{F}(A)$ denotes the set of all finite subsets of A , then

$$|\mathcal{F}(A)| = |A|.$$

§3.8 Concrete encodings behind the arithmetic

The equality

$$\aleph_0 \cdot \aleph_0 = \aleph_0$$

is not just a theorem name. A direct bijection is the Cantor pairing function

$$\pi : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}, \quad \pi(m, n) = \frac{(m+n)(m+n+1)}{2} + m.$$

It enumerates lattice points by diagonals. For instance,

$$\pi(0, 0) = 0, \quad \pi(1, 0) = 2, \quad \pi(0, 1) = 1.$$

To invert it, start with $k \in \mathbb{N}$. Find the unique integer $s \geq 0$ such that

$$\frac{s(s+1)}{2} \leq k < \frac{(s+1)(s+2)}{2}.$$

Then

$$m = k - \frac{s(s+1)}{2}, \quad n = s - m.$$

This recovers a unique pair (m, n) , so π is a bijection.

There is an equally concrete encoding of finite subsets of $\mathbb{N}$. If

$$S = \{s_1 < s_2 < \cdots < s_k\} \subseteq \mathbb{N}$$

is finite, define

$$f(S) = \sum_{i=1}^k 2^{s_i},$$

with $f(\emptyset) = 0$. Binary expansion is unique, so two different finite subsets have different codes. Hence

$$|\mathcal{F}(\mathbb{N})| \leq \aleph_0.$$

The reverse injection is the singleton map

$$\mathbb{N} \hookrightarrow \mathcal{F}(\mathbb{N}), \quad n \mapsto \{n\}.$$

By Cantor–Bernstein,

$$|\mathcal{F}(\mathbb{N})| = \aleph_0.$$

This is the finite-subset theorem in a form one can actually compute: a finite subset is just the support of the 1's in a binary expansion.

Gap

The general statement $|\mathcal{F}(A)| = |A|$ for every infinite set A still uses the comparison machinery above, and with arbitrary A one must pay attention to the choice principles being used. But for the countable model case, the pairing function and the binary encoding already show the mechanism rather than merely quoting the cardinal-arithmetic slogan.

Rethinking Set Theory Through Functions Rather Than Membership

N-20220131

§4.1 Why this translation/note matters

This note began as a translation and reading note on Tom Leinster’s article *Rethinking Set Theory*. The article won the Chauvenet Prize, and the reason is easy to feel while reading it: it is not trying to make set theory more technical, but to ask why the set theory used by working mathematicians feels so different from the official ZFC story.

The guiding question is simple:

Question

If most mathematicians manipulate sets safely without remembering the ZFC axioms, what principles are they actually using?

The answer explored here is Lawvere’s Elementary Theory of the Category of Sets, usually abbreviated ETCS. The important point is that ETCS is not an attempt to replace set theory by category theory. It is another axiomatization of set theory, written in a language closer to everyday mathematical practice.

§4.2 The discomfort with ZFC

ZFC treats every element of every set as itself a set. Formally, this is convenient. But in ordinary mathematics it is strange. If I ask what the elements of π are, or what the elements of a real number are, most mathematicians would say that the question is meaningless.

In ZFC, however, every object has been coded as a set, so asking for the elements of a real number is formally allowed. This does not mean ZFC is wrong. It means that ZFC uses the word “set” differently from how mathematicians often use it in practice.

Remark 4.2.1 — The point is not merely that real numbers are encoded arbitrarily. The deeper point is linguistic: ZFC makes membership primitive, while ordinary mathematics often treats functions, products, subsets, and inverse images as the more natural operations.

A typical ZFC axiom such as foundation says that every nonempty set has an element disjoint from it. This makes formal sense only because all elements are also sets. But when the set under discussion is something like $\mathbb{R}$, the ordinary mathematical intuition does not ask whether $\pi \cap \mathbb{R} = \emptyset$.

§4.3 Three misunderstandings about ETCS

The original article emphasizes three misunderstandings.

- (i) ETCS is not a competitor to set theory; it is a form of set theory.
- (ii) ETCS does not require more sophisticated mathematics than ZFC. It can be stated elementarily, even though its origin is categorical.
- (iii) ETCS is not circular merely because it is inspired by categories. ZFC says, roughly, “there are things called sets and a relation called membership.” ETCS says, roughly, “there are things called sets and things called functions, and functions compose.” Both are first-order theories.

Caution

It is easy to overstate the category-theoretic flavor. The point here is not to presuppose the whole theory of categories. The point is to take functions and composition as basic, because that is often closer to actual mathematical use.

§4.4 Elements as special functions




Figure 4.1: The motivating picture: elements can be represented as maps from a singleton.

The key idea is to define elements using functions. Suppose there is a singleton set

$$1 = \{\bullet\}.$$

Then an element of a set X can be identified with a function

$$1 \rightarrow X.$$

Indeed, such a function is determined by the image of $\bullet$, and that image is an element of X .

At first this looks like a formal trick, but it is a very general mathematical pattern:

- a loop in a topological space X is a continuous map $S^1 \rightarrow X$;
- a sequence in X is a function $\mathbb{N} \rightarrow X$;
- a solution of equations in a ring can be represented by a suitable homomorphism out of a quotient ring;
- a point of a space is often best treated as a map from a test object.

Remark 4.4.1 — This is the first philosophical shift: instead of saying that functions are built out of elements, we may say that elements are special cases of functions.

§4.5 The informal ETCS-style principles

The note lists the principles in an informal table. Rewritten cleanly, they say roughly:

- (i) functions compose associatively and have identity functions;
- (ii) there is a singleton set;
- (iii) there is an empty set;
- (iv) a function is determined by its effect on elements;
- (v) products $X \times Y$ exist;
- (vi) function sets Y^X exist;
- (vii) inverse images of elements or subsets exist;
- (viii) subsets of X correspond to functions $X \rightarrow \{0, 1\}$;
- (ix) the natural numbers form a set;
- (x) every surjection has a right inverse, i.e. a choice principle.

These principles are striking because they sound like the operations used constantly in ordinary mathematics. They do not begin by asking what the elements of elements are.

§4.6 Subsets via characteristic functions

A subset $A \subseteq X$ may be represented by its characteristic function

$$\chi_A : X \rightarrow \{0, 1\}, \quad \chi_A(x) = \begin{cases} 1, & x \in A, \\ 0, & x \notin A. \end{cases}$$

Thus the power set of X is expressed as a function set

$$2^X.$$

This fits the ETCS point of view perfectly: subsets are not introduced by a membership relation alone, but through maps into a classifier of truth values.

§4.7 Products, exponentials, and inverse images

The ordinary constructions of mathematics appear directly:

$$X \times Y$$

represents pairs; the exponential

$$Y^X$$

represents functions from X to Y ; and inverse images let us pull subsets back along functions. These are exactly the operations that mathematicians use constantly when building structures.

Remark 4.7.1 — In this sense, ETCS feels less like an encoding scheme and more like a formalization of actual mathematical behavior.

§4.8 Choice as right inverses

The choice principle in this language says that every surjection has a right inverse. If

$$f : X \twoheadrightarrow Y$$

is surjective, then there exists

$$g : Y \rightarrow X$$

with

$$f \circ g = \text{id}_Y.$$

This is exactly the idea of choosing one element from each fibre of f .

Caution

This formulation is often more transparent than the usual family-of-nonempty-sets version, but it is the same choice principle in a categorical dress.

§4.9 What I take from the article

The main lesson is not that ZFC is useless or wrong. ZFC is powerful and foundationally successful. The lesson is that there is a mismatch between the official membership-first universe of ZFC and the function-centered practice of most mathematics.

ETCS suggests that one can axiomatize set theory using:

sets, functions, composition,
products, function spaces, subsets,
natural numbers, choice.

That language feels closer to how I actually use sets.

Question

Is ETCS only a more convenient presentation of ordinary mathematics, or does it also change how one should teach foundations? The article strongly suggests the latter, and this is the part I find most provocative.

§4.10 Translation note

This chapter is a reconstruction of a translation/reading note rather than a verbatim translation. I keep the mathematical thread: why ZFC feels alien in ordinary practice, how ETCS starts from functions, and why elements can be treated as maps from a singleton.

§4.11 ETCS intuition: membership is not primitive practice

One of the central insights of ETCS is that working mathematicians rarely use the membership relation as their main tool. They use functions. A set is understood by the maps into it, the maps out of it, its products, its subsets, and its quotients.

For example, instead of defining an ordered pair through nested braces, one can characterize the product $A \times B$ by its universal property. Instead of treating a subset as a special kind of element relation, one can treat it through its characteristic map to a subobject classifier. This does not deny membership; it changes which notions are primitive.

§4.12 Element reasoning inside categorical set theory

A common worry is that categorical set theory loses elements. It does not. In a category of sets, an element of A can be represented as a map

$$1 \rightarrow A$$

from a terminal object. Thus element reasoning is recovered as morphism reasoning from the one-point set.

This is a typical categorical move: replace internal membership statements by external mapping properties. The result often looks more abstract, but it is closer to how structures are transported by functions.

§4.13 Why this matters philosophically

ZFC gives a reductionist foundation: everything is encoded as a set. ETCS gives a structural foundation: sets are objects in a category satisfying axioms that support ordinary set operations. The philosophical difference is that ETCS does not ask what a set is made of at the lowest membership level; it asks what role sets play in mathematics.

This is why Leinster's article is compelling. It shows that the official foundation and the working foundation are not identical. ZFC is powerful as a background theory, but ETCS more closely describes the everyday language of functions, products, quotients, and subsets.

A First Comparison of ZFC, NBG, KM, and New Foundations

N-20220205

§5.1 Why there are several axiomatizations of set theory

This note compares several axiomatic systems for set theory: ZFC, NBG, Kelley–Morse set theory, and Quine’s New Foundations. The underlying question is not merely historical. Different systems express different compromises between intuition, rigor, large collections, and self-reference.

§5.2 The main systems

The main systems in this note are:

ZFC. Zermelo–Fraenkel set theory with Choice, the standard foundation for most modern mathematics.

NBG. Von Neumann–Bernays–Goedel set theory, which distinguishes sets from classes.

KM. Kelley–Morse set theory, a stronger class theory than NBG.

NF. Quine’s New Foundations, which uses stratification to control comprehension.

§5.3 ZFC

ZFC is built from several familiar axioms: extensionality, empty set, pairing, union, power set, foundation, replacement, infinity, and choice.

Remark 5.3.1 — ZFC’s strength is that it is compact, standard, and sufficient for almost all ordinary mathematics. Replacement and Infinity allow one to build complicated structures such as ordinals, cardinals, and the continuum.

Caution

ZFC does not directly treat proper classes as objects of the theory. This is often fine, but it becomes awkward when discussing objects such as the class of all sets, large categories, or global constructions in category theory.

§5.4 NBG

NBG introduces classes as first-class objects while still distinguishing sets from proper classes.

Definition 5.4.1. In NBG, every set is a class, but not every class is a set. A proper class is a class that is too large to be a set.

This makes it possible to speak about the universe of all sets as a proper class. It also makes certain category-theoretic discussions feel more natural, since large categories can be treated using classes.

Remark 5.4.2 — NBG is conservative over ZFC for statements purely about sets. In ordinary set-level mathematics, it usually proves the same theorems as ZFC. Its advantage is expressive convenience when classes are involved.

§5.5 Kelley–Morse set theory

KM extends the class-theoretic viewpoint further. It allows stronger comprehension principles for classes and therefore gives a richer theory of proper classes.

Remark 5.5.1 — KM can be thought of as a stronger language for discussing global class-sized structures. It is useful when NBG’s class comprehension is too weak for the intended argument.

The tradeoff is that KM is logically stronger and less minimal. It may be more expressive, but it also introduces more foundational weight.

Question

Is there a clean presentation of KM that remains beginner-friendly while preserving its extra expressive power? This is still unclear in the note, and it is worth returning to after learning more model theory and class theory.

§5.6 New Foundations

Quine’s New Foundations takes a more radical route. It uses a stratification condition on formulas to control paradoxes while allowing constructions that look very different from ZFC.

Caution

The original wording describes NF as allowing self-reference in a loose way. More precisely, NF restricts comprehension to stratified formulas. The point is not that arbitrary self-reference is allowed, but that the theory draws the safety boundary differently from ZFC.

NF differs from ZFC in several ways:

- (i) ZFC uses Foundation to rule out membership cycles; NF uses stratification.
- (ii) ZFC builds sets through iterative hierarchy intuition; NF modifies the logic of allowable comprehension.
- (iii) NF has subtle consistency questions, and variants such as NFU are often easier to handle.

§5.7 How to compare the systems

The comparison can be organized along three axes.

§5.7.i Expressive power

ZFC is clean for ordinary set-level mathematics. NBG and KM are better suited for talking about proper classes and large structures. NF explores a different response to self-reference.

§5.7.ii Logical strength

NBG is conservative over ZFC for set statements. KM is stronger than NBG. NF and its variants sit in a different part of the foundational landscape and should not be casually compared by intuition alone.

§5.7.iii Use in practice

For most everyday mathematics, ZFC is enough. For category theory or discussions involving universes and proper classes, NBG-style language is often more comfortable. KM becomes relevant when one truly needs stronger class principles.

§5.8 What remains unclear

Gap

The original note contains several broad claims about conservativity and consistency that need more precise model-theoretic formulation. In particular, KM should not be casually described as conservative over ZFC in the same way NBG is. A later revision should add exact statements about relative consistency and conservativity.

Question

Which foundation is best for actual mathematical practice: a minimal set theory such as ZFC, a class theory such as NBG, a stronger class theory such as KM, or a category-inspired system such as ETCS? The answer probably depends on what one wants to do.

§5.9 Where the idea leads

ZFC is the mainstream baseline. NBG keeps the same set-level mathematics while adding proper classes. KM strengthens the class theory. NF changes the handling of comprehension and self-reference. The deeper lesson is that foundations are not only about avoiding paradoxes; they also shape what kinds of mathematical objects feel natural to discuss.

§5.10 Why classes are introduced

The main pressure on ZFC is that some collections are too large to be sets. The collection of all sets, the collection of all groups, or the collection of all ordinals cannot be sets

without recreating paradoxes. Class theories introduce a formal distinction:

sets are classes that can be elements of other classes,
proper classes are too large to be elements.

In NBG, classes are first-class objects of the language, but quantification over classes is restricted in a way that keeps the theory conservative over ZFC for set-level statements. KM allows stronger class comprehension and is therefore stronger.

This explains the practical difference. ZFC can simulate classes by formulas, but NBG and KM allow one to talk about classes directly. For category theory this is convenient: the category of all sets is not small, but it can be treated as a class-sized category.

§5.11 New Foundations and stratification

Quine's NF tries a different repair of naive comprehension. Instead of restricting comprehension to subsets of already existing sets, it allows comprehension only for stratified formulas. A formula is stratified if one can assign type levels to variables so that membership $x \in y$ raises type by one, while equality $x = y$ keeps type the same.

For example, $x \notin x$ is not stratified, because it would require the type of x to be one higher than itself. Thus Russell's paradox is blocked. But a formula like $x = x$ is stratified, so NF has a universal set.

This is philosophically striking: ZFC avoids a universal set; NF permits one by controlling the syntax of comprehension.

§5.12 Comparing foundational choices

The systems differ not only in strength but in taste. ZFC is minimal enough to serve as a common background for ordinary mathematics. NBG is convenient when large classes are used explicitly. KM is stronger and supports more class-level reasoning. NF preserves a universal-set intuition but pays for it with syntactic restrictions.

The useful conclusion is not that one system is simply correct and the others wrong. A foundation is a formal environment. Different environments make different constructions easy or hard.

Vectors, Ordered Fields, and the First Axioms of Real Numbers

N-20220710

§6.1 Plane vectors

The first page contains exercises on plane vectors. A vector in the plane may be written

$$\vec{a} = (x, y) = x\mathbf{i} + y\mathbf{j}.$$

Addition and scalar multiplication are coordinatewise:

$$(x, y) + (z, w) = (x + z, y + w), \quad \lambda(x, y) = (\lambda x, \lambda y).$$

The length and dot product are

$$|(x, y)| = \sqrt{x^2 + y^2}, \quad (x, y) \cdot (z, w) = xz + yw.$$

Thus the angle between two nonzero vectors satisfies

$$\cos \theta = \frac{\vec{a} \cdot \vec{b}}{|\vec{a}| |\vec{b}|} = \frac{xz + yw}{\sqrt{x^2 + y^2} \sqrt{z^2 + w^2}}.$$

If

$$\vec{x} = (\cos \alpha, \sin \alpha), \quad \vec{y} = (\cos \beta, \sin \beta),$$

then

$$\cos \theta = \cos \alpha \cos \beta + \sin \alpha \sin \beta = \cos(\alpha - \beta).$$

This gives a vector proof of the cosine subtraction formula.

§6.2 Field axioms

Axiom 6.2.1 (Field axioms)

A field is a set F with operations $+$ and $\cdot$ such that $(F, +)$ is an abelian group, $(F \setminus \{0\}, \cdot)$ is an abelian group, and multiplication distributes over addition.

The second page moves to the algebraic structure of the real numbers. The field axioms are:

(F1) $x + (y + z) = (x + y) + z;$

(F2) $x + y = y + x;$

(F3) there exists 0 such that $0 + x = x;$

(F4) for each x , there is $-x$ such that $x + (-x) = 0;$

(F5) $x(yz) = (xy)z;$

- (F6) $xy = yx$;
- (F7) there exists $1 \neq 0$ such that $1x = x$;
- (F8) for $x \neq 0$, there is x^{-1} such that $xx^{-1} = 1$;
- (F9) $x(y + z) = xy + xz$.

The note emphasizes that a “space” in mathematics is usually a set equipped with extra structure. A field is not merely a set; it is a set equipped with addition, multiplication, identities, inverses, and compatibility laws.

§6.3 Ordered fields

The order axioms include:

- (O1) Transitivity: if $x \leq y$ and $y \leq z$, then $x \leq z$.
- (O2) Antisymmetry: if $x \leq y$ and $y \leq x$, then $x = y$.
- (O3) Totality: for all x, y , either $x \leq y$ or $y \leq x$.
- (O4) Compatibility with addition: $x \leq y \Rightarrow x + z \leq y + z$.
- (O5) Compatibility with multiplication: $0 \leq x$ and $0 \leq y \Rightarrow 0 \leq xy$.

The Archimedean axiom says that for every $x > 0$ and every y , there exists a positive integer n such that

$$nx \geq y.$$

A typical consequence is that every interval $(0, a)$ with $a > 0$ is nonempty: choose n so large that $n > a^{-1}$, then $1/n < a$.

§6.4 Elementary consequences of the order axioms

Proposition 6.4.1 (Consequences of an ordered field)

In an ordered field, addition preserves order and multiplication by a positive element preserves order. In particular, squares are nonnegative:

$$x^2 \geq 0.$$

The page records several exercises:

- (i) $x \geq 0$ iff $-x \leq 0$, and $y > 1$ implies $0 < 1/y < 1$.
- (ii) $1 > 0$ and $-1 \neq 1$.
- (iii) If $x \leq y$ and $a \leq 0$, then $ax \geq ay$.
- (iv) If $a \leq b$ and $x \leq y$, then $a + x \leq b + y$, with equality iff $a = b$ and $x = y$.
- (v) If $a < x$ for every $a < x$, then $x \leq y$; this is a typical contradiction argument using an interval between two distinct real numbers.
- (vi) $x^2 \geq 0$ for every real x .

(vii) If $a^2 < a$, then $0 < a < 1$.

(viii) If nonzero x, y have the same sign, then

$$(x + y)^2 > (x - y)^2.$$

§6.5 Embedding $\mathbb{Q}$ into $\mathbb{R}$

The note then proves that $\mathbb{R}$ contains a copy of $\mathbb{Q}$. Define for integers n

$$\iota(n) = \underbrace{1 + \cdots + 1}_{n \text{ times}}$$

for $n > 0$, $\iota(0) = 0$, and $\iota(n) = -\iota(-n)$ for $n < 0$. Then define

$$\iota\left(\frac{p}{q}\right) = \frac{\iota(p)}{\iota(q)}.$$

One must check that this is well-defined: if $p/q = s/t$, then $pt = qs$, so

$$\iota(p)\iota(t) = \iota(q)\iota(s),$$

and therefore

$$\frac{\iota(p)}{\iota(q)} = \frac{\iota(s)}{\iota(t)}.$$

This embedding preserves addition, multiplication, and order.

§6.6 Nested intervals

The final part introduces the nested interval axiom. If

$$I_n = [a_n, b_n]$$

is a descending sequence of closed intervals with real endpoints, then

$$I_1 \supseteq I_2 \supseteq I_3 \supseteq \cdots$$

implies

$$\bigcap_{n \geq 1} I_n \neq \emptyset.$$

Together with the field, order, and Archimedean axioms, this is one way to characterize the real numbers.

§6.7 Ordered fields and the real line

An ordered field is a field equipped with an order compatible with addition and multiplication. Compatibility means

$$a < b \Rightarrow a + c < b + c, \quad 0 < a, 0 < b \Rightarrow 0 < ab.$$

The rational numbers form an ordered field, but they are not complete. The real numbers are obtained by adding the missing limits, for example through Dedekind cuts or Cauchy sequences.

This is why early vector and coordinate exercises naturally lead to axioms of the real numbers. Geometry needs coordinates, coordinates need arithmetic, and analysis needs completeness.

§6.8 Vectors as algebraicized geometry

A plane vector (x, y) packages a displacement into two real numbers. Vector addition corresponds to composing displacements; scalar multiplication corresponds to stretching. The algebraic laws of vectors are therefore not arbitrary axioms. They express basic geometric operations in coordinate form.

§6.9 The conceptual payoff

The note moves from computational vector exercises to the axiomatic construction of the real line. The important conceptual shift is that familiar formulas are only the visible part of an axiomatic structure. A rigorous real-number theory must explain why rational numbers embed, why intervals contain points, and why irrational numbers such as $\sqrt{2}$ can exist.

§6.10 Ordered fields and what they cannot prove

An ordered field supports algebra and order, but not necessarily completeness. The rationals satisfy the ordered-field axioms but have gaps. Thus many analytic theorems need more than order and algebra.

§6.11 Archimedean property and density

In the real numbers, the Archimedean property says no positive infinitesimal exists inside $\mathbb{R}$: for every x , some integer n satisfies $n > x$. This leads to density of rationals and helps connect arithmetic approximations with the geometry of the line.

§6.12 Completeness as the real-number axiom

The supremum property, nested intervals, Cauchy completeness, and Dedekind completeness are equivalent ways of saying that the real line has no gaps. This is the feature that turns ordered algebra into analysis.

§6.13 From ordered algebra to the real line

Vectors need linear structure; ordered fields need compatibility between addition, multiplication, and order; the real numbers need completeness. The early axioms are not formal decoration: each one supports a different layer of geometry and analysis.

Completeness, Supremum, Dedekind Cuts, and Metric Spaces

N-20220713

§7.1 Large and small positive real numbers

The page begins with a simple but important fact: for every positive real A , there is a positive real $M > A$; for every positive real a , there is a positive real $\varepsilon < a$. For instance,

$$M = A + 1, \quad \varepsilon = \frac{a}{2}.$$

This is the first appearance of the ε -language of analysis.

§7.2 The geometric view of the real line

We think of $\mathbb{R}$ as a line, where order is represented by left and right. This picture is helpful, but the note emphasizes that all conclusions must ultimately come from axioms and rigorous reasoning, not from the drawing itself.

A set $X \subseteq \mathbb{R}$ is bounded above if there exists $A \in \mathbb{R}$ such that

$$x \leq A \quad \text{for all } x \in X.$$

Such an A is an upper bound. Similarly, X is bounded below if there exists a such that

$$a \leq x \quad \text{for all } x \in X.$$

§7.3 Supremum property

The least upper bound property says: if $X \subseteq \mathbb{R}$ is nonempty and bounded above, then the set of upper bounds has a least element. This least upper bound is denoted

$$\sup X.$$

Equivalently, $M = \sup X$ iff:

- (i) $x \leq M$ for all $x \in X$;
- (ii) for every $\varepsilon > 0$, there exists $x \in X$ such that

$$x > M - \varepsilon.$$

The note sketches a proof of the supremum property from the nested interval principle by bisecting intervals and selecting the half that still contains upper bounds or points of the set.

The dual notion is the infimum:

$$\inf X.$$

One always has

$$\inf X \leq \sup X$$

for a nonempty bounded set.

§7.4 Nested intervals imply completeness

Let $X \subseteq \mathbb{R}$ be nonempty and bounded above. One can construct nested intervals

$$I_n = [a_n, b_n]$$

so that a_n is not an upper bound of X , while b_n is an upper bound, and

$$b_n - a_n \leq 2^{-n}.$$

The nested interval axiom gives a point

$$M \in \bigcap_{n=1}^{\infty} I_n.$$

The shrinking length forces M to be the least upper bound of X . If M were not an upper bound, some $x \in X$ would lie above it; for large n , the right endpoint b_n would be too close to M , contradiction. If M were not least, a smaller upper bound would contradict the construction of the a_n 's.

§7.5 Dedekind cuts

A Dedekind cut is a subset $X \subset \mathbb{Q}$ such that:

- (i) $X \neq \emptyset$ and $\mathbb{Q} \setminus X \neq \emptyset$;
- (ii) if $x \in X$ and $y < x$, then $y \in X$;
- (iii) X has no greatest element.

A rational number p/q defines the cut

$$X_{p/q} = \{x \in \mathbb{Q} : x < p/q\}.$$

An irrational such as $\sqrt{2}$ is represented by

$$X_{\sqrt{2}} = \{x \in \mathbb{Q}_{>0} : x^2 < 2\} \cup \mathbb{Q}_{\leq 0}.$$

Thus real numbers can be constructed from rational cuts.

§7.6 Metric spaces

A metric space is a set X equipped with a function

$$d : X \times X \rightarrow \mathbb{R}_{\geq 0}$$

satisfying:

- (i) $d(x, y) = 0$ iff $x = y$;
- (ii) $d(x, y) = d(y, x)$;
- (iii) $d(x, z) \leq d(x, y) + d(y, z)$.

The usual metric on $\mathbb{R}$ is

$$d(x, y) = |x - y|.$$

On $\mathbb{R}^q$, common metrics are:

$$d_2(x, y) = \left(\sum_{n=1}^q |x_n - y_n|^2 \right)^{1/2},$$

$$d_1(x, y) = \sum_{n=1}^q |x_n - y_n|,$$

and

$$d_\infty(x, y) = \max_{1 \leq n \leq q} |x_n - y_n|.$$

The note remarks that d_2 requires Cauchy–Schwarz for the triangle inequality, d_1 is the Manhattan metric, and d_∞ is often called the Chebyshev metric.

A final example is the discrete metric:

$$d(x, y) = \begin{cases} 0, & x = y, \\ 1, & x \neq y. \end{cases}$$

In the discrete metric, open balls are especially simple:

$$B_r(x) = \begin{cases} \{x\}, & 0 < r \leq 1, \\ X, & r > 1. \end{cases}$$

The note points out that boundary points are not included in an open ball, explaining why the case $r = 1$ gives $\{x\}$, not all of X .

§7.7 Subspaces and density

If $Y \subseteq X$ and X has metric d , then Y inherits the subspace metric

$$d_Y(y_1, y_2) = d(y_1, y_2).$$

A subset $Y \subseteq X$ is dense if for every $x \in X$ and every $\varepsilon > 0$, there exists $y \in Y$ such that

$$d(x, y) < \varepsilon.$$

This formalizes the idea that Y comes arbitrarily close to every point of X .

§7.8 Supremum as least upper bound

The supremum of a set $A \subseteq \mathbb{R}$ is not merely a maximum. It is the least number s such that every element of A is at most s . A maximum must belong to A ; a supremum need not.

For example,

$$\sup(0, 1) = 1,$$

but $(0, 1)$ has no maximum. This distinction is exactly what makes completeness subtle.

§7.9 Dedekind cuts and completeness

A Dedekind cut separates rationals into a lower part and an upper part with no gap. Irrational numbers appear as cuts not produced by a rational boundary. Completeness says that every such cut corresponds to a real number.

This is the philosophical point of constructing $\mathbb{R}$: the real line is obtained by filling all rational gaps.

§7.10 Looking ahead

The note moves from the order-completeness of $\mathbb{R}$ to the metric-space language of distance and density. Supremum, Dedekind cuts, nested intervals, and metrics are different ways of making precise the intuition that the real line has no holes.

§7.11 Supremum as order-completeness

The supremum property says every nonempty bounded-above set has a least upper bound. This is the order-theoretic heart of the real line. It is what makes monotone convergence and intermediate-value reasoning possible.

§7.12 Cauchy completion and metric completion

A metric space can be completed by adding equivalence classes of Cauchy sequences. Two Cauchy sequences are equivalent if their distance tends to zero. The real numbers can be constructed as the completion of $\mathbb{Q}$ under the usual metric.

§7.13 Dedekind cuts versus Cauchy sequences

Dedekind cuts construct reals by filling order gaps; Cauchy completion constructs reals by adding missing limits of Cauchy sequences. The two constructions produce isomorphic complete ordered fields, but they emphasize different structures: order versus metric.

§7.14 Completeness as the no-gap principle

Completeness is the bridge from algebraic number systems to analysis. Supremum, nested intervals, Dedekind cuts, density, and metric completion are not separate facts; they are equivalent ways of ruling out gaps.

Nonstandard Analysis and the Meaning of $0.999 \dots = 1$

N-20230225

§8.1 Why this topic appears

The note begins from a familiar confusion: why should

$$0.999 \dots = 1$$

be exactly true? The standard real-number answer is short: the decimal is the infinite series

$$\sum_{k=1}^{\infty} 9 \cdot 10^{-k} = 1.$$

But the note wants to explore a different intuition: if one stops after a very large number of nines, the number is slightly less than one. Nonstandard analysis gives a formal language in which “very large finite” and “infinitesimally small” can be discussed without contradiction.

§8.2 The real-number proof

Let

$$x = 0.999 \dots$$

Then

$$10x = 9.999 \dots$$

Subtracting gives

$$9x = 9, \quad x = 1.$$

Equivalently,

$$0.999 \dots = \sum_{k=1}^{\infty} 9 \cdot 10^{-k} = 9 \cdot \frac{1/10}{1 - 1/10} = 1.$$

In the real numbers there is no positive real number smaller than all 10^{-n} . Therefore the “gap” between $0.999 \dots$ and 1 is not a real number.

§8.3 Hyperreal numbers

Nonstandard analysis extends $\mathbb{R}$ to a larger field ${}^*\mathbb{R}$, containing infinitesimal and infinite numbers. One construction uses an ultrapower:

$${}^*\mathbb{R} = \mathbb{R}^{\mathbb{N}}/\mathcal{U},$$

where $\mathcal{U}$ is a nonprincipal ultrafilter on $\mathbb{N}$. A hyperreal number is represented by a sequence of real numbers; two sequences are considered equal if they agree on a set belonging to the ultrafilter.

For example,

$$H = [1, 2, 3, 4, \dots]$$

is an infinite hyperinteger, because it is larger than every standard integer. Its reciprocal

$$\varepsilon = 1/H$$

is a positive infinitesimal: it is greater than zero but smaller than every positive standard real number.

§8.4 Infinitesimals and finite hyperreals

A hyperreal x is infinitesimal if

$$|x| < r$$

for every positive standard real r . A hyperreal x is finite if

$$|x| < N$$

for some standard integer N . Every finite hyperreal x is infinitely close to a unique real number, called its standard part:

$$\text{st}(x).$$

We write $x \approx y$ if $x - y$ is infinitesimal.

The standard-part map is the bridge back to ordinary real analysis. Nonstandard analysis often computes with infinitesimal quantities and then takes standard part to return to a real value.

§8.5 Hyperfinite strings of nines

Let H be an infinite hyperinteger and define the hyperfinite decimal

$$x_H = \sum_{k=1}^H 9 \cdot 10^{-k}.$$

The finite geometric-sum formula still applies by transfer:

$$x_H = 1 - 10^{-H}.$$

Here 10^{-H} is a positive infinitesimal. Therefore

$$x_H < 1, \quad x_H \approx 1, \quad \text{st}(x_H) = 1.$$

This distinguishes two statements:

- (i) the real infinite decimal $0.999\dots$ equals 1;
- (ii) a hyperfinite string of H nines equals $1 - 10^{-H}$, which is infinitesimally below 1.

The confusion comes from mixing these two interpretations.

§8.6 Transfer principle

The transfer principle says that every first-order statement true in $\mathbb{R}$ is true in ${}^*\mathbb{R}$ after replacing objects by their hyperreal extensions. This is why ordinary algebraic identities, order laws, and finite-sum formulas can be used with hyperintegers.

For example, the formula

$$\sum_{k=1}^n 9 \cdot 10^{-k} = 1 - 10^{-n}$$

for every natural number n transfers to every hypernatural H :

$$\sum_{k=1}^H 9 \cdot 10^{-k} = 1 - 10^{-H}.$$

This is the formal justification behind the hyperfinite decimal calculation.

§8.7 Derivative through infinitesimals

If f is differentiable at a , then for every nonzero infinitesimal ε ,

$$f'(a) = \text{st} \left(\frac{f(a + \varepsilon) - f(a)}{\varepsilon} \right).$$

For example, if $f(x) = x^2$, then

$$\frac{(a + \varepsilon)^2 - a^2}{\varepsilon} = 2a + \varepsilon.$$

Taking standard part gives

$$f'(a) = 2a.$$

This recovers the intuitive infinitesimal calculation that historically motivated calculus, but now in a rigorous model.

For a trigonometric example, let $f(x) = \sin x$ and let $\varepsilon \neq 0$ be infinitesimal. Then

$$\begin{aligned} \frac{\sin(x + \varepsilon) - \sin x}{\varepsilon} &= \frac{\sin x \cos \varepsilon + \cos x \sin \varepsilon - \sin x}{\varepsilon} \\ &= \sin x \frac{\cos \varepsilon - 1}{\varepsilon} + \cos x \frac{\sin \varepsilon}{\varepsilon}. \end{aligned}$$

The rigorous nonstandard meaning of the familiar approximations is:

$$\frac{\sin \varepsilon}{\varepsilon} \approx 1, \quad \frac{\cos \varepsilon - 1}{\varepsilon} \approx 0,$$

because the standard limits

$$\lim_{h \rightarrow 0} \frac{\sin h}{h} = 1, \quad \lim_{h \rightarrow 0} \frac{\cos h - 1}{h} = 0$$

transfer to infinitesimal h . Taking standard part therefore gives

$$\text{st} \left(\frac{\sin(x + \varepsilon) - \sin x}{\varepsilon} \right) = \cos x.$$

This is the same derivative as in standard calculus, but the computation is genuinely infinitesimal rather than merely symbolic.

§8.8 Integral through hyperfinite sums

For a continuous function f on $[a, b]$, take an infinite hyperinteger H and set

$$\Delta x = \frac{b-a}{H}, \quad x_i = a + i\Delta x.$$

Then

$$\int_a^b f(x) dx = \text{st} \left(\sum_{i=0}^{H-1} f(x_i) \Delta x \right).$$

This is a hyperfinite Riemann sum. It makes the intuitive phrase “add infinitely many infinitesimal rectangles” precise.

§8.9 What nonstandard analysis is not saying

Nonstandard analysis does not say that the real number $0.999\dots$ is less than 1. In $\mathbb{R}$, it equals 1. What nonstandard analysis adds is a separate object: a hyperfinite decimal with an infinite but not completed number of nines. That object is infinitesimally below 1, and its standard part is 1.

This distinction is exactly why nonstandard analysis is useful. It can formalize intuitive infinitesimal reasoning without changing the standard real results.

§8.10 Source repair: internal and external sets

The uploaded note includes the important distinction between internal and external sets. Internal sets arise from the ultrapower construction, roughly by taking sequences of standard sets and passing to the quotient. External sets are subsets of ${}^*\mathbb{R}$ that do not arise in this way.

A standard example is the set of infinitesimals

$$\mu(0) = \{x \in {}^*\mathbb{R} : x \text{ is infinitesimal}\}.$$

This set is external. The distinction matters because the transfer principle applies to internal objects, not arbitrary external subsets. Many false proofs in nonstandard analysis come from accidentally applying transfer to an external set.

§8.11 Source repair: standard part and Loeb measure

For every finite hyperreal x , there is a unique real number infinitely close to it. This number is the standard part:

$$\text{st}(x) \in \mathbb{R}, \quad x \approx \text{st}(x).$$

Infinitesimal calculations produce finite hyperreal approximations; taking standard part extracts the real answer.

The source also points toward Loeb measure. The rough pattern is

$$\text{hyperfinite counting measure} \longrightarrow \text{standard part} \longrightarrow \text{Loeb measure}.$$

This turns an internally finite-looking probability space into a standard measure space rich enough for analysis.

§8.12 Source repair: transfer has limits

The uploaded note correctly worries that transfer applies to first-order statements, not arbitrary higher-order assertions. Algebraic identities, order properties, and finite-sum formulas transfer cleanly. Statements quantifying over all subsets or all functions require more care.

Thus the safe habit is: use transfer for first-order statements about standard structures, use internal-set reasoning inside the hyperreal universe, and use standard part only when returning to ordinary real analysis.

§8.13 A final synthesis

Nonstandard analysis is a model-theoretic enlargement of the real line. It does not replace epsilon-delta analysis; it gives an equivalent language for many arguments. The standard real continuum has no infinitesimals, but a sufficiently rich extension can contain them while preserving first-order real algebra and order facts by transfer. The decimal example is a good gateway because it separates completed infinite limits from hyperfinite approximations.

§8.14 Real equality and limiting process

The identity $0.999\ldots = 1$ in the real numbers follows because the difference after n digits is 10^{-n} , and this tends to zero. Equality is a statement in the completed real number system.

§8.15 Hyperreals and infinitesimal remainders

In a hyperreal extension, a hyperfinite decimal with H nines differs from 1 by 10^{-H} , a positive infinitesimal when H is infinite. Its standard part is still 1.

§8.16 Transfer and caution

The transfer principle lets first-order statements move between reals and hyperreals. But one must distinguish finite hyperreal approximations, standard real limits, and standard-part projection.

§8.17 Decimals, limits, and infinitesimal models

The decimal question is really about completion and interpretation. In $\mathbb{R}$, the infinite decimal denotes a limit; in hyperreals, hyperfinite truncations may differ infinitesimally while having the same standard part.

Foundations of Mathematics: Set Theory, Category Theory, and Type Theory

N-20230710

§9.1 The question of foundations

This chapter is a reconstructed version of an expository article comparing three major foundational languages: set theory, category theory, and type theory. The point is not to declare a single winner. Rather, I want to understand how each framework answers the same basic questions:

- (i) What are mathematical objects?
- (ii) What counts as equality?
- (iii) What kind of reasoning is allowed?
- (iv) How does computation enter mathematics?
- (v) How does one organize structures and transformations between them?

Set theory emphasizes membership; category theory emphasizes morphisms and universal properties; type theory emphasizes construction, judgment, and computation. The most interesting modern developments, especially univalent foundations and homotopy type theory, no longer keep these worlds separate.

§9.2 Set theory

§9.2.i From naive comprehension to axioms

Modern set theory begins with the intuition that mathematical objects can be collected into sets. In naive set theory, one might try to form

$$\{x \mid P(x)\}$$

for any property P . Russell's paradox shows that unrestricted comprehension is impossible: the class

$$R = \{x \mid x \notin x\}$$

leads to contradiction if it is treated as a set.

The response is axiomatization. Zermelo–Fraenkel set theory with Choice, ZFC, restricts set formation and supplies a controlled universe in which ordinary mathematics can be developed.

§9.2.ii ZFC as the standard background

The familiar axioms include extensionality, pairing, union, power set, infinity, replacement, foundation, separation, and choice. Their roles may be summarized as follows:

Extensionality. A set is determined by its elements.

Pairing, union, power set. These build new sets from old ones.

Infinity. This guarantees an infinite set, usually giving the natural numbers.

Replacement. Definable images of sets are sets.

Foundation. Membership is well-founded and circular membership is excluded.

Choice. Products of nonempty families are nonempty; equivalently, every surjection has a right inverse.

Remark 9.2.1 — The strength of ZFC is not that it matches everyday mathematical language perfectly. Its strength is that it gives a stable formal universe large enough for most mathematics and disciplined enough to avoid naive paradoxes.

§9.2.iii Applications and limitations

Set theory supports analysis, algebra, topology, measure theory, model theory, and essentially all ordinary mathematical constructions. It also enables transfinite methods: ordinals, cardinals, large cardinals, forcing, independence, and descriptive set theory.

But ZFC has philosophical and practical limitations. It represents everything as sets, even when this representation feels artificial. A group is not experienced as a nested set; it is experienced as a structure with operations. A topological space is not merely a set of ordered pairs defining open sets; it is a space with continuous maps. This motivates category theory.

§9.3 Category theory

§9.3.i From objects to morphisms

Category theory changes the focus from internal membership to external relations. A category consists of objects and morphisms, with associative composition and identities. In this language, mathematical structures are understood by the maps they admit.

Definition 9.3.1. A category $\mathcal{C}$ consists of objects, morphism sets $\mathcal{C}(X, Y)$, identity morphisms id_X , and composition

$$\mathcal{C}(Y, Z) \times \mathcal{C}(X, Y) \rightarrow \mathcal{C}(X, Z),$$

satisfying associativity and identity laws.


Example 9.3.2

Set, **Grp**, **Top**, **Ring**, and **Vect_k** are standard categories. Each has a different notion of structure-preserving map.

§9.3.ii Universal properties

The central categorical idea is the universal property. Products, coproducts, pullbacks, pushouts, quotients, tensor products, free objects, and completions can all be characterized by mapping properties.

For example, a product $X \times Y$ is characterized by



This diagram says that giving a map into a product is the same as giving compatible maps into its two factors.

Remark 9.3.3 — Universal properties explain why the same construction appears in many different subjects. They reveal that the construction is not about the elements of a particular model, but about a structural role.

§9.3.iii Functors and natural transformations

Definition 9.3.4. A functor $F : \mathcal{C} \rightarrow \mathcal{D}$ sends objects to objects and morphisms to morphisms, preserving identities and composition.

Definition 9.3.5. A natural transformation $\eta : F \Rightarrow G$ between functors $F, G : \mathcal{C} \rightarrow \mathcal{D}$ assigns to each object X a morphism $\eta_X : F(X) \rightarrow G(X)$ such that for every $f : X \rightarrow Y$,

$$\begin{array}{ccc} F(X) & \xrightarrow{F(f)} & F(Y) \\ \eta_X \downarrow & & \downarrow \eta_Y \\ G(X) & \xrightarrow{G(f)} & G(Y) \end{array}$$

commutes.

This square is one of the simplest and most important categorical diagrams. It encodes the idea that a construction is canonical, not dependent on arbitrary choices.

§9.3.iv Categorical foundations

Category theory can function as a language for organizing mathematics, but it can also become foundational. ETCS axiomatizes sets categorically. Topos theory generalizes set-like universes. Higher category theory and infinity-categories extend the categorical language to homotopical and derived contexts.

Remark 9.3.6 — Category theory does not merely replace set theory. It clarifies invariance. It asks not only what an object is made of, but how it transforms and how it participates in universal constructions.

§9.4 Type theory

§9.4.i Judgments and types

Type theory begins from judgments such as

$$a : A,$$

meaning that a is a term of type A . Instead of putting every object into one universe of sets, type theory stratifies objects by the types they inhabit.

Remark 9.4.1 — The phrase “ $a \in A$ ” in set theory and the judgment “ $a : A$ ” in type theory look similar but behave differently. Membership is a proposition; typing is a judgment of well-formedness.

§9.4.ii Dependent types

Dependent type theory allows types to depend on terms. If A is a type and $B(x)$ is a type depending on $x : A$, then one can form:

$$\prod_{x:A} B(x), \quad \sum_{x:A} B(x).$$

The product type corresponds to universal quantification, while the sum type corresponds to existential quantification.

§9.4.iii Curry–Howard

The Curry–Howard correspondence identifies propositions with types and proofs with terms:

$$\text{proposition} \leftrightarrow \text{type}, \quad \text{proof} \leftrightarrow \text{term}.$$

For example,

$$A \times B$$

corresponds to conjunction, while

$$A \rightarrow B$$

corresponds to implication.

This makes type theory naturally computational. A proof is not merely a static certificate; it may carry an algorithm.

§9.4.iv Constructivity and computation

Classical set theory freely uses excluded middle and choice. Type theory is often constructive: to prove existence, one should construct a witness. This is not only a philosophical choice. It is what makes type theory suitable for proof assistants such as Coq, Lean, Agda, and Isabelle/HOL.

§9.5 Homotopy type theory and univalence

Homotopy type theory interprets types as spaces, terms as points, and identity proofs as paths. Higher identity proofs become higher homotopies. In this interpretation, equality is no longer merely a binary relation; it has structure.

The univalence axiom says, roughly, that equivalent structures may be identified:

$$(A = B) \simeq (A \simeq B).$$

This expresses a structuralist principle: isomorphic or equivalent objects should be interchangeable in mathematics.

Remark 9.5.1 — Univalence is one of the most striking attempts to synthesize set-theoretic rigor, categorical invariance, and type-theoretic computation. It turns the informal mathematical habit of identifying isomorphic objects into a formal principle.

§9.6 Comparative analysis

Framework	Primitive emphasis	Equality	Strength
Set theory	membership	extensional equality	universal coding power
Category theory	morphisms	isomorphism/ equivalence	structural invariance
Type theory	judgments/terms	identity types	computation and construction

Set theory is excellent as a common foundational background. Category theory is excellent for organizing structures and transformations. Type theory is excellent when proofs, programs, and constructions are meant to interact.

§9.7 Modern applications

§9.7.i Proof assistants

Proof assistants make foundations operational. Lean’s mathlib, Coq’s constructive type theory, Agda’s dependent types, and Isabelle’s higher-order logic show that foundational choices affect how mathematics is formalized.

§9.7.ii Computer science

Type theory underlies programming language theory, dependent types, formal verification, compiler correctness, and proof-carrying code. Category theory appears in semantics, monads, functional programming, and compositional systems. Set theory remains important in logic, model theory, semantics, and databases.

§9.7.iii Geometry and higher structures

Modern geometry often uses categorical and type-theoretic ideas: stacks, sheaves, derived categories, infinity-categories, and higher topos theory. Foundations are no longer merely a safety net beneath mathematics; they become active tools inside research.

§9.8 Toward a unified perspective

A mature view should not treat the three foundations as enemies. They answer different needs:

- (i) ZFC gives a robust global universe;
- (ii) category theory reveals structural patterns and invariance;
- (iii) type theory connects proof, construction, and computation;
- (iv) HoTT and univalent foundations attempt to synthesize equivalence-based mathematics with formal computation.

§9.9 Conclusion

The foundational landscape is plural. Set theory gives a powerful ontology of membership. Category theory gives a language of structure and transformation. Type theory gives a constructive and computational account of mathematical reasoning. The deepest modern lesson is not that one of them eliminates the others, but that mathematics becomes clearer when one understands which foundational lens is being used and why.

§9.10 Equality in the three languages

The three foundations treat equality differently. In set theory, equality is extensional:

$$A = B \iff \forall x (x \in A \leftrightarrow x \in B).$$

In category theory, sameness is usually replaced by isomorphism or equivalence. Two groups may not be literally equal as sets, but if they are isomorphic, group theory treats them as the same structure.

Type theory makes equality internal. A statement $a = b$ is itself a type, and a proof of equality is a term of that type. Homotopy type theory then interprets such proofs as paths. This explains why univalence is so conceptually strong: it aligns equality of types with equivalence of structures.

§9.11 Universal properties versus encodings

A set-theoretic construction often gives a concrete encoding. A categorical construction gives a role. For example, a product is not fundamentally a set of ordered pairs; it is an object receiving projections and satisfying a universal mapping property.

This matters foundationally because encodings are arbitrary but universal properties are invariant. Kuratowski's ordered pair and other encodings are useful inside ZFC, but no working mathematician cares which coding was used once the universal property is available.

§9.12 Computation and formal proof

Type theory is not only a philosophical alternative. It is the foundation behind many formal proof systems. A proof is a term, and checking a proof becomes type checking. Dependent types allow mathematical statements to be expressed with high precision: a proof of existence contains a witness, and a proof of a universal statement behaves like a function.

This computational character is one reason type theory feels different from classical set theory. It is not only about what objects exist, but also about how they are constructed and verified.

II

Categories and Logic

Part II: Contents

N-20210721	55
10.1 Starting point	55
10.2 The table is not a list	55
10.3 The dualizing object	56
10.4 Stone duality as the first complete example	56
10.5 What Stone duality says in one sentence	58
10.6 Gelfand duality as a continuous analogue	58
10.7 Affine schemes and the affine line	59
10.8 Locales and the Sierpiński space	59
10.9 Topoi, logoi, and higher versions	60
10.10 The formal meaning of duality	60
10.11 What I should have understood in 2021	61
N-20210804	63
11.1 Boolean algebra beyond computing	63
11.2 A first definition via Heyting algebras	63
11.3 A small amount of algebra	63
11.4 Boolean rings	64
11.5 Examples of Boolean rings	65
11.6 Boolean algebras directly	66
11.7 Characteristic functions and the meaning of 2^X	66
11.8 Moving between the two structures	67
11.9 The logical shadow	67
11.10 Filters, ultrafilters, and Stone spaces	67
11.11 Boolean rings from Boolean algebras	68
11.12 Stone duality as a first hint	68
N-20221005	69
12.1 First orientation	69
12.2 From truth to resources	69
12.3 Narrative as a resource pipeline	70
12.4 Tensor, additives, and branching stories	72
12.5 Dialogue as a game	72
12.6 Schopenhauer's stratagem as a ludic interaction	72
12.7 Ludics: loci instead of signifieds	73
12.8 Orthogonality and negation	74
12.9 Geometry of interaction	74
12.10 From narrative to social relations	75
12.11 The general pattern	75
N-20221227	77
13.1 More examples of categories	77
13.2 Subcategories and functors	77
13.3 Forgetful functors	77
13.4 A topology generated by a distance-like question	77
13.5 Posets as categories, with an example	78
13.6 Functorial thinking	78
13.7 Generated topology	78
13.8 Why category theory appears beside topology	78
13.9 The larger pattern	79
13.10 Posets as categories and logic	79
13.11 Forgetful functors and free constructions	79
13.12 Initial topology generated by questions	79
13.13 Product topology as an initial topology	80
13.14 Quotient topology as a final topology	80
13.15 Functors from topology to algebra	80
13.16 Initial and final structures in topology	80

N-20230313	81
14.1 What this note continues	81
14.2 Posets are categories	81
14.3 A small poset example	81
14.4 Isomorphisms in a poset category	82
14.5 Groups as one-object categories	82
14.6 Why this example matters	82
14.7 Products of categories	82
14.8 The opposite category	83
14.9 Monomorphisms and epimorphisms	83
14.10 The lesson of the note	83
N-20230314	85
15.1 Free and forgetful functors	85
15.2 Faithful and full functors	85
15.3 Concrete categories	86
15.4 Functors out of a one-object group category	86
15.5 Functors between one-object group categories	86
15.6 Why the note says “learn algebra first”	87
15.7 Looking ahead: natural transformations	87
N-20240211	89
16.1 Why topoi enter the story	89
16.2 Categories	89
16.3 Universal properties and limits	90
16.4 Exponentials and cartesian closure	90
16.5 Subobjects and the subobject classifier	91
16.6 Elementary topoi	91
16.7 A toy presheaf topos	92
16.8 Grothendieck topoi and sheaves	93
16.9 Internal logic	93
16.10 Power objects	94
16.11 Topos as generalized universe of sets	94
16.12 A high-level comparison	94
16.13 Further directions	94
16.14 Subobjects and characteristic maps	95
16.15 Why logic becomes intuitionistic	95
16.16 Sheaves as variable sets	95
N-20251109	97
17.1 The setting: maps with zero maps	97
17.2 Kernel: the universal object killed by a map	97
17.3 Cokernel: the universal quotient killing a map	98
17.4 Image: the part of the codomain actually reached	98
17.5 Coimage: the effective part of the domain	99
17.6 The standard factorization	99
17.7 A vector-space sanity check	100
17.8 A module example: multiplication by an integer	100
17.9 Exactness language	101
17.10 Warnings: what depends on being abelian	101
17.11 Synthesis: from element chasing to universal properties	101

A Gateway to Dualities: Space, Algebra, and the Dualizing Object

N-20210721

§10.1 Starting point

In July 2021, a table titled “Some dualities” appeared in our Geek College chat. It was shared by Dr. Keyao Peng, and at first glance it looked like a compact list of examples. The table had three columns: a geometric object, an algebraic object, and a mysterious third entry called the dualizing object or gauge object.

The table became one of the first places where I felt that modern mathematics was not merely a collection of topics, but a repeated pattern. Geometry and algebra were not standing separately. They seemed to be two languages for the same structure. The first row, Stone spaces versus Boolean algebras, was the entry point that most immediately caught my attention. It led to my later attempt to write about Boolean algebra, but that earlier writing mostly explained the algebraic operations. It did not yet explain the missing spatial half: why logic should have a topology.

This note rewrites that summer question in a more systematic form. The aim is not merely to list several dualities. The aim is to understand the mechanism behind the table:

a space is studied through its observations,
an algebra is studied through its models or points.

The third column of the table is the key. The dualizing object is the interpreter through which the two sides talk to each other.

§10.2 The table is not a list

Here is the table in a cleaned-up form.

Geometry	Algebra	Dualizing object / gauge object
Stone spaces	Boolean algebras	the Boolean values $\mathbf{2} = \{0, 1\}$
compact Hausdorff spaces	commutative C^* -algebras	the complex numbers $\mathbb{C}$
affine schemes	commutative rings	the affine line $\mathbb{A}^1$
locales	frames	the Sierpiński space $\mathbb{S}$
topoi	logoi	the topos Set of sets
∞ -topoi	∞ -logoi	the ∞ -topos of ∞ -groupoids

The table should not be read as six unrelated examples. It is an architectural diagram. The left column contains objects that behave like spaces. The middle column contains algebraic or logical presentations of those spaces. The right column contains the object used to measure, test, or observe them.

The shortest slogan is:

geometry is recovered from algebra, and algebra is recovered from geometry.

But this slogan is still too vague. The real mechanism is contravariant homming into a dualizing object.

§10.3 The dualizing object

Let Ω be a dualizing object. A geometric object X gives an algebra of Ω -valued observations:

$$X \mapsto \text{Hom}(X, \Omega).$$

A point of $\text{Hom}(X, \Omega)$ is a way of asking a question about X , with answer living in Ω .

Conversely, an algebraic object A gives a geometric space of Ω -valued models:

$$A \mapsto \text{Hom}(A, \Omega).$$

A point of this new space is a way to interpret the algebra A inside the measuring object Ω .

The reason a duality reverses arrows is already visible here. If

$$f : X \rightarrow Y$$

is a map of spaces, then precomposition gives

$$\text{Hom}(Y, \Omega) \longrightarrow \text{Hom}(X, \Omega), \quad \varphi \longmapsto \varphi \circ f.$$

So a map $X \rightarrow Y$ produces a map in the opposite direction on the algebras of observations. This elementary reversal is the seed of the whole table.

Remark 10.3.1 — The third column is therefore not an accessory. It tells us what kind of observation is allowed. Changing the dualizing object changes the kind of geometry that can be seen.

§10.4 Stone duality as the first complete example

The first row is the best place to begin because it is the most logical and the most concrete:

$$\text{Stone spaces} \quad \longleftrightarrow \quad \text{Boolean algebras}, \quad \Omega = \mathbf{2} = \{0, 1\}.$$

Here $\mathbf{2}$ is the space of truth values. A map into $\mathbf{2}$ is a yes-or-no observation.

§10.4.i Boolean algebras as propositional logic

A Boolean algebra is an algebraic model of classical propositional logic.

Definition 10.4.1 (Boolean algebra). A Boolean algebra is a structure

$$B = (B, \vee, \wedge, ', 0, 1)$$

where $\vee$ and $\wedge$ are join and meet, 0 and 1 are bottom and top, and each element $a \in B$ has a complement a' . It satisfies the usual associativity, commutativity, absorption, distributivity, and complementation laws.

The example to keep in mind is the power set $\mathcal{P}(X)$. Then

$$A \vee B = A \cup B, \quad A \wedge B = A \cap B, \quad A' = X \setminus A.$$

So Boolean algebra is the algebra of regions that can be combined by union, intersection, and complement.

§10.4.ii Stone spaces as spaces with Boolean topology

A Stone space is a topological space whose topology is controlled by clopen sets.

Definition 10.4.2 (Stone space). A Stone space is a compact Hausdorff totally disconnected space. Equivalently, it is a compact Hausdorff space with a basis of clopen subsets.

A subset is clopen if it is both closed and open. In ordinary connected spaces, clopen sets are rare. In a Stone space, clopen sets are abundant enough to form the logical basis of the space.

For any Stone space X , the clopen subsets form a Boolean algebra:

$$\text{Clop}(X).$$

The Boolean operations are exactly the set operations:

$$U \vee V = U \cup V, \quad U \wedge V = U \cap V, \quad U' = X \setminus U.$$

This gives the direction

$$X \longmapsto \text{Clop}(X).$$

§10.4.iii Why $\mathbf{2}$ is the dualizing object

The set $\mathbf{2} = \{0, 1\}$ is given the discrete topology. A continuous map

$$\chi : X \rightarrow \mathbf{2}$$

is the characteristic function of a clopen subset of X . Indeed,

$$\chi^{-1}(\{1\})$$

is clopen because $\{1\} \subset \mathbf{2}$ is clopen.

Thus

$$\text{Hom}_{\mathbf{Top}}(X, \mathbf{2}) \cong \text{Clop}(X).$$

This is the first half of the row: the algebra of logic on a Stone space is the algebra of truth-valued continuous observations.

§10.4.iv Ultrafilters as points

Now start from a Boolean algebra B . A point of the Stone space of B is not given in advance. It must be reconstructed from the algebra.

Definition 10.4.3 (Filter and ultrafilter). A filter $F \subseteq B$ is a subset such that $1 \in F$, $0 \notin F$, it is closed under finite meets, and it is upward closed. An ultrafilter is a maximal proper filter.

An ultrafilter is a complete truth assignment. For every proposition $a \in B$, exactly one of a and a' belongs to the ultrafilter. So an ultrafilter decides every proposition in a way compatible with Boolean logic.

Let

$$S(B) = \{\mathfrak{u} : \mathfrak{u} \text{ is an ultrafilter on } B\}.$$

For each $a \in B$, define

$$U_a = \{\mathfrak{u} \in S(B) : a \in \mathfrak{u}\}.$$

The sets U_a form a basis of clopen subsets. This turns $S(B)$ into a Stone space.

The map

$$a \longmapsto U_a$$

preserves Boolean operations:

$$U_{a \wedge b} = U_a \cap U_b, \quad U_{a \vee b} = U_a \cup U_b, \quad U_{a'} = S(B) \setminus U_a.$$

Theorem 10.4.4 (Stone representation theorem)

For every Boolean algebra B , the map

$$B \longrightarrow \text{Clop}(S(B)), \quad a \longmapsto U_a$$

is an isomorphism of Boolean algebras.

This is the missing geometric half of Boolean algebra. A Boolean algebra is not only a symbolic calculus of propositions. It is the algebra of clopen regions of a canonically reconstructed space.

§10.5 What Stone duality says in one sentence

Stone duality says:

logic has a topology.

The propositions are clopen subsets. The truth assignments are points. The Boolean algebra is the algebra of observable yes-or-no regions of the Stone space.

The dualizing object **2** explains both directions:

$$X \longmapsto \text{Hom}(X, \mathbf{2}) \cong \text{Clop}(X),$$

$$B \longmapsto \text{Hom}(B, \mathbf{2}) \cong S(B).$$

On the geometric side, maps to **2** are clopen predicates. On the algebraic side, homomorphisms to **2** are truth assignments, equivalently ultrafilters.

This is what I had not fully seen in 2021. I understood the algebraic rules of Boolean algebra, but I had not yet seen that the rules reconstruct a space.

§10.6 Gelfand duality as a continuous analogue

The second row replaces truth-valued observations by complex-valued observations:

$$\text{compact Hausdorff spaces} \quad \longleftrightarrow \quad \text{commutative } C^*\text{-algebras}, \quad \Omega = \mathbb{C}.$$

If X is compact Hausdorff, then $C(X)$, the algebra of continuous complex-valued functions on X , is a commutative unital C^* -algebra. This is the direction

$$X \longmapsto C(X) = \text{Hom}(X, \mathbb{C})$$

in a suitably interpreted sense.

Conversely, if A is a commutative unital C^* -algebra, its spectrum is the set of characters

$$\text{Spec}(A) = \text{Hom}_{C^*\text{-alg}}(A, \mathbb{C}).$$

These characters are the points of the reconstructed compact Hausdorff space.

Theorem 10.6.1 (Gelfand duality)

The category of compact Hausdorff spaces is contravariantly equivalent to the category of commutative unital C^* -algebras. In particular,

$$A \cong C(\operatorname{Spec}(A))$$

for commutative unital C^* -algebras A .

The analogy with Stone duality is direct:

Stone duality	Gelfand duality
$\mathbf{2}$ -valued observations	$\mathbb{C}$ -valued observations
clopen predicates	continuous functions
ultrafilters	characters
Boolean logic	commutative functional analysis

§10.7 Affine schemes and the affine line

The third row is algebraic geometry:

$$\text{affine schemes} \longleftrightarrow \text{commutative rings}, \quad \Omega = \mathbb{A}^1.$$

The affine line $\mathbb{A}^1$ is the object that represents regular functions. For an affine scheme X , maps

$$X \rightarrow \mathbb{A}^1$$

are regular functions on X . Thus the ring of functions on X is recovered as

$$\operatorname{Hom}(X, \mathbb{A}^1).$$

Conversely, a commutative ring R gives the affine scheme

$$\operatorname{Spec} R,$$

whose points are prime ideals of R . The slogan is the same: an algebraic space is reconstructed from its algebra of functions.

This row is close in spirit to Gelfand duality, but the geometry is algebraic rather than analytic. The measuring object is no longer $\mathbb{C}$ as a topological field of values, but $\mathbb{A}^1$ as the universal carrier of regular functions.

§10.8 Locales and the Sierpiński space

The fourth row is

$$\text{locales} \longleftrightarrow \text{frames}, \quad \Omega = \mathbb{S}.$$

A locale is a point-free space. Instead of starting with points and then defining open sets, one starts with the lattice of opens.

Definition 10.8.1 (Frame). A frame is a complete lattice L in which finite meets distribute over arbitrary joins:

$$a \wedge \bigvee_i b_i = \bigvee_i (a \wedge b_i).$$

For a topological space X , the lattice $\mathcal{O}(X)$ of open subsets is a frame. Locale theory studies this frame as the primary object.

The dualizing object here is the Sierpiński space $\mathbb{S}$, with two points $\{0, 1\}$, where $\{1\}$ is open but $\{0\}$ is not. The key fact is:

$$\mathrm{Hom}(X, \mathbb{S}) \cong \mathcal{O}(X).$$

A continuous map $X \rightarrow \mathbb{S}$ is exactly the characteristic function of an open subset of X .

This makes the third-column idea completely visible. The Sierpiński space is not merely a two-point curiosity. It is the classifier of open sets.

§10.9 Topoi, logoi, and higher versions

The last rows are much larger in scope:

$$\text{topoi} \quad \longleftrightarrow \quad \text{logoi}, \quad \Omega = \mathbf{Set},$$

and

$$\infty\text{-topoi} \quad \longleftrightarrow \quad \infty\text{-logoi}.$$

A topos can be read as a generalized space, often presented as a category of sheaves. It is a place where mathematics can be carried out internally. A logos is the corresponding logical or algebraic presentation.

The row is harder than Stone duality, but the analogy is still useful:

$$\text{space of models} \quad \longleftrightarrow \quad \text{logic/theory of that space}.$$

The topos of sets $\mathbf{Set}$ acts as a reference universe. For ∞ -topoi, ordinary sets are replaced by spaces or ∞ -groupoids, so that higher homotopical information is retained rather than collapsed.

For this note, the details of logoi are less important than the direction of generalization. The table begins with Boolean truth values and ends with entire universes of mathematics as measuring objects.

§10.10 The formal meaning of duality

Categorically, a duality between categories $\mathcal{C}$ and $\mathcal{D}$ is an equivalence

$$\mathcal{C} \simeq \mathcal{D}^{op}.$$

Equivalently, there are contravariant functors

$$F : \mathcal{C} \rightarrow \mathcal{D}, \quad G : \mathcal{D} \rightarrow \mathcal{C}$$

which recover the original objects up to natural isomorphism.

In many rows of the table, the functors are built by homming into a dualizing object. This is why the third column matters. It is the object that converts a space into its algebra of observations and an algebra into its space of models.

So the statement is not simply:

algebra is the dual of geometry.

A more precise statement is:

for suitable categories, a chosen measuring object
produces a contravariant equivalence.

The geometry/algebra distinction is then a matter of viewpoint. A space can be treated through its functions. An algebra can be treated through its spectrum of points.

§10.11 What I should have understood in 2021

The table's first row was already enough to contain the lesson. A Boolean algebra is not only a formal calculus of $\vee$, $\wedge$, and complement. It has a Stone space of complete truth assignments. Conversely, a Stone space is not only a topological object. It has a Boolean algebra of clopen predicates.

The right summary is:

points are models or truth assignments,
propositions are regions,
the dualizing object tells us what kind of question is being asked.

In Stone duality, the question is yes-or-no, so the dualizing object is **2**. In Gelfand duality, the question is complex-valued measurement, so the dualizing object is $\mathbb{C}$. In locale theory, the question is openness, so the dualizing object is the Sierpiński space. In topos theory, the question becomes internal mathematics itself.

This is the missing half of the 2021 Boolean algebra study. I had understood the syntax of Boolean operations, but not yet the topology of truth assignments. The real lesson of the table is:

space is recovered from its observations,
logic is recovered as the topology of its truth assignments.

Boolean Algebras, Boolean Rings, and the First Hint of Stone Duality

N-20210804

§11.1 Boolean algebra beyond computing

This note began from a small but common misunderstanding: Boolean algebra is often first seen through computer science, so it is easy to think that it is mainly a computational convention rather than a mathematical structure. I want to push against that impression.

Boolean algebra belongs simultaneously to logic, set theory, topology, measure theory, ring theory, and later to categorical duality. The original motivation was a comparison between Stone duality for Boolean algebras and broader dualities involving topological or logical structures. I also had Heyting algebras in mind, although at this stage it is better not to start there.

Remark 11.1.1 — The planned route is deliberately elementary: first collect the algebraic vocabulary, then introduce Boolean rings and Boolean algebras, and only later move toward Stone spaces, Stone duality, and perhaps Heyting algebras. The point is not to pretend that Boolean algebra is difficult at the entry level, but to show that it has more mathematical depth than its first appearance suggests.

§11.2 A first definition via Heyting algebras

One compact definition is the following.

Definition 11.2.1. A Boolean algebra is a Heyting algebra satisfying the law of excluded middle. Equivalently, for every element x ,

$$\neg\neg x = x,$$

or equivalently

$$x \vee \neg x = \top.$$

Caution

This definition is elegant but not beginner-friendly. If one has not learned Heyting algebras, it hides the concrete content. So in the rest of this chapter I reconstruct Boolean algebras through groups, rings, lattices, subsets, and characteristic functions.

§11.3 A small amount of algebra

Let S be a nonempty set. A binary operation on S is a map

$$S \times S \rightarrow S.$$

Depending on context we may denote the value on (a, b) by ab , $a + b$, $a - b$, $a * b$, or $a \circ b$.

Definition 11.3.1. A semigroup is a nonempty set G with an associative binary operation:

$$a(bc) = (ab)c.$$

If there is an element $e \in G$ such that $ae = ea = a$, and every element $a \in G$ has an inverse a^{-1} with

$$aa^{-1} = a^{-1}a = e,$$

then G is a group. If $ab = ba$ for all $a, b \in G$, then G is an abelian group.

I will not need the full theory of group homomorphisms, kernels, images, cosets, normal subgroups, or quotient groups here, but the rough memory is useful: a quotient group G/N consists of cosets of a normal subgroup $N \triangleleft G$, with multiplication

$$(aN)(bN) = abN.$$

Definition 11.3.2. A ring R is a set with addition and multiplication such that $(R, +)$ is an abelian group, multiplication is associative, and multiplication distributes over addition:

$$a(b + c) = ab + ac, \quad (b + c)a = ba + ca.$$

If multiplication is commutative, the ring is commutative. If there is an element $1 \in R$ with $1a = a1 = a$, the ring has a unit.

Definition 11.3.3. A lattice is a partially ordered set in which finite meets and finite joins exist. We write

$$p \wedge q, \quad p \vee q$$

for meet and join.

Remark 11.3.4 — This algebra review is intentionally minimal. The important point is that Boolean algebra can be approached from two sides: as a special ring and as a special lattice.

§11.4 Boolean rings

The smallest nontrivial ring with unit is

$$\mathbb{F}_2 = \{0, 1\},$$

with addition and multiplication modulo 2:

$+$	0	1	$\cdot$	0	1
0	0	1	0	0	0
1	1	0	1	0	1

Every element satisfies $p^2 = p$. This condition is the key.

Definition 11.4.1. A Boolean ring is a ring R in which every element is idempotent:

$$x^2 = x \quad (x \in R).$$

Some definitions also require a unit, but the idempotence condition is the essential feature.

Proposition 11.4.2

Every Boolean ring satisfies $x + x = 0$ for every $x \in R$, and every Boolean ring is commutative.

Proof. Since $x + x$ is idempotent,

$$x + x = (x + x)^2 = x^2 + xx + xx + x^2 = x + x + 2x^2.$$

Thus $2x^2 = 0$, and since $x^2 = x$, we get $2x = 0$.

Next,

$$x + y = (x + y)^2 = x^2 + xy + yx + y^2 = x + y + xy + yx.$$

So $xy + yx = 0$. But every element is its own additive inverse, hence $xy = yx$. $\square$

§11.5 Examples of Boolean rings

The two-dimensional example is

$$\mathbb{F}_2^2 = \{(0, 0), (0, 1), (1, 0), (1, 1)\},$$

with coordinatewise operations:

$$(p_0, p_1) + (q_0, q_1) = (p_0 + q_0, p_1 + q_1),$$

$$(p_0, p_1)(q_0, q_1) = (p_0q_0, p_1q_1).$$

The zero and unit are

$$(0, 0), \quad (1, 1).$$

The same construction gives $\mathbb{F}_2^n$.

Another standard example is a ring of sets.

Example 11.5.1

Let $\mathcal{R}$ be a class of sets such that if $E, F \in \mathcal{R}$, then

$$E \cup F \in \mathcal{R}, \quad E - F \in \mathcal{R}.$$

Define

$$E + F = (E \setminus F) \cup (F \setminus E),$$

that is, symmetric difference, and define

$$EF = E \cap F.$$

Then the set-theoretic laws match the Boolean ring laws.

Remark 11.5.2 — This is also why Boolean algebra naturally appears in measure theory. Collections of measurable sets, especially modulo null sets, behave like Boolean algebras or Boolean rings.

§11.6 Boolean algebras directly

Let X be a set. The power set $\mathcal{P}(X)$ has union, intersection, and complement:

$$P \cup Q, \quad P \cap Q, \quad P^c = X \setminus P.$$

This is the concrete model for Boolean algebra.

Definition 11.6.1. A Boolean algebra is a nonempty set A equipped with two binary operations $\vee, \wedge$, a unary operation $(\prime)$, and elements $0, 1$, satisfying the identities modeled on union, intersection, and complement. In particular,

$$\begin{aligned} p \wedge 0 &= 0, & p \wedge 1 &= p, & p \wedge p' &= 0, & p \wedge p &= p, \\ p \vee 1 &= 1, & p \vee 0 &= p, & p \vee p' &= 1, & p \vee p &= p, \\ (p')' &= p, & (p \wedge q)' &= p' \vee q', & (p \vee q)' &= p' \wedge q', \end{aligned}$$

together with commutativity, associativity, and distributivity.

Caution

One sentence in the original note has an obvious slip: it says, in effect, that every Boolean algebra can be turned into a Boolean ring and every Boolean algebra can be turned into a Boolean algebra. The intended statement is that Boolean algebras and Boolean rings with unit are two equivalent presentations of the same structure.

§11.7 Characteristic functions and the meaning of 2^X

Every subset $P \subseteq X$ corresponds to a characteristic function

$$\chi_P : X \rightarrow \mathbb{F}_2, \quad \chi_P(x) = \begin{cases} 1, & x \in P, \\ 0, & x \notin P. \end{cases}$$

This gives a natural bijection

$$\mathcal{P}(X) \cong 2^X.$$

Thus writing the power set as 2^X is not merely a notational habit: a subset is the same thing as a 0/1-valued function.

Under this identification, pointwise operations on functions become set operations:

$$P + Q = (P \cap Q^c) \cup (P^c \cap Q),$$

which is symmetric difference, and

$$PQ = P \cap Q.$$

The zero and unit are

$$0 = \emptyset, \quad 1 = X.$$

§11.8 Moving between the two structures

From a Boolean ring R with unit, define

$$p \wedge q = pq, \quad p \vee q = p + q + pq, \quad p' = p + 1.$$

These operations make R a Boolean algebra.

Conversely, from a Boolean algebra A , define

$$p + q = (p \wedge q') \vee (p' \wedge q), \quad pq = p \wedge q.$$

This makes A a Boolean ring.

Remark 11.8.1 — The formulas are not magic. In the subset model, addition is exactly symmetric difference, multiplication is intersection, and complement is symmetric difference with the whole universe.

§11.9 The logical shadow

The operations $\wedge, \vee, '$ correspond to “and”, “or”, and “not”. The elements 0 and 1 correspond to false and true. This is why Boolean algebra is the algebra of classical propositional logic.

Remark 11.9.1 — This also explains why Heyting algebras are the natural next object: Boolean algebras model classical logic, while Heyting algebras model intuitionistic logic. Stone duality then brings topology into the same picture.

§11.10 Filters, ultrafilters, and Stone spaces

The next natural step is to introduce filters, ultrafilters, Stone spaces, and the Stone representation theorem. This does not replace the dictionary between Boolean rings, Boolean algebras, sets, and logic; it explains why that dictionary has a topological shadow.

Let B be a Boolean algebra. A filter $F \subseteq B$ is a collection of propositions regarded as “large” or “accepted”. It satisfies three basic rules:

$$1 \in F, \quad 0 \notin F,$$

if $a, b \in F$, then $a \wedge b \in F$, and if $a \in F$ and $a \leq b$, then $b \in F$. Thus a filter is closed under strengthening by consequence and under finite conjunction.

An ultrafilter is a maximal proper filter. Equivalently, for every $a \in B$, exactly one of a and a' belongs to the ultrafilter. This makes an ultrafilter look like a complete truth assignment: every Boolean proposition is judged either true or false, compatibly with the Boolean operations.

The Stone space of B , denoted $S(B)$, is the set of all ultrafilters on B . For each $a \in B$, define

$$U_a = \{u \in S(B) : a \in u\}.$$

These subsets form a basis of clopen sets. The Boolean operations are reflected topologically by

$$U_{a \wedge b} = U_a \cap U_b, \quad U_{a \vee b} = U_a \cup U_b, \quad U_{a'} = S(B) \setminus U_a.$$

So the algebra B is represented inside the topology of $S(B)$ as the Boolean algebra of clopen sets.

Stone's representation theorem says, in one common form, that every Boolean algebra is isomorphic to the Boolean algebra of clopen subsets of its Stone space. The finite example is easy to remember: if $B = \mathcal{P}(X)$, then ultrafilters are just choices of points of X , and the Stone space recovers X with the discrete topology. For infinite Boolean algebras, non-principal ultrafilters appear, and the resulting Stone space is compact, Hausdorff, and totally disconnected.

§11.11 Boolean rings from Boolean algebras

Given a Boolean algebra, define addition by symmetric difference

$$A + B = (A \setminus B) \cup (B \setminus A)$$

and multiplication by intersection

$$AB = A \cap B.$$

Then every element satisfies $A^2 = A$, so the resulting ring is Boolean. Conversely, a Boolean ring gives a Boolean algebra by defining meet as multiplication and join by

$$a \vee b = a + b + ab.$$

This equivalence explains why logic, sets, and algebra keep producing the same identities.

§11.12 Stone duality as a first hint

Stone duality says that Boolean algebras correspond to certain topological spaces, namely Stone spaces. The rough idea is that algebraic operations encode clopen subsets of a space. Thus Boolean logic can be represented geometrically. This is one of the earliest examples of algebra and topology reflecting each other.

Linear Logic: Resources, Dialogue, Narrative, and Geometry

N-20221005

§12.1 First orientation

This note is reconstructed from a ten-page PDF titled *The Geometry of Dialogue and the Logic of Narrative: A Taste of Linear Logic*, dated 2022-10-05. The original text was not a conventional mathematical note. It was closer to a reflective synthesis: a recent immersion into linear logic, several remarks attributed to Dr. Keyao Peng, an interest in game design and emergent narrative, and a set of visual fragments from Ludics, sequent calculus, and Geometry of Interaction.

The background matters. The original memo was written at the moment when logic stopped looking like a static taxonomy of truth values and began to look like a theory of interaction. The guiding words were:

interaction, resource, game, narrative.

The reconstruction below keeps this origin, but changes the form. Instead of an essay moving from metaphor to metaphor, this note is organized as a usable mathematical notebook: definitions first, then small calculi, then code-like models, and finally the more philosophical interpretation.

Remark 12.1.1 — The original PDF repeatedly tried to connect three worlds: a child’s fable, a philosophical debate, and high-level topology or geometry of interaction. The stable mathematical bridge between them is linear logic’s treatment of resources and duality.

§12.2 From truth to resources

Classical logic treats propositions as reusable truths. If A is available, then one may use A as often as needed. This is encoded by the structural rules:

weakening: ignore a hypothesis, contraction: duplicate a hypothesis.

Linear logic removes these structural rules by default. A hypothesis is a resource. If it is consumed by an inference, it is no longer freely available.

Definition 12.2.1 (Linear implication). The linear implication

$$A \multimap B$$

means: given one resource of type A , one may consume it to produce one resource of type B . It is not the same as the classical implication $A \rightarrow B$, because A is not automatically reusable.

The exponential modality $!$ marks reusable resources. A useful slogan is:

$$A \rightarrow B \text{ behaves like } !A \multimap B.$$

The ordinary implication first makes A reusable, then spends it linearly.

classical reading	linear reading	informal meaning
$A \rightarrow B$	$!A \multimap B$	reusable assumption produces B
$A \otimes B$	tensor	both resources are present together
$A \& B$	additive conjunction	one must be ready for either choice
$A \oplus B$	additive disjunction	one side is chosen by the prover

The essential change is not symbolic. It is operational. Linear logic asks what a proof *does* to its premises.

§12.3 Narrative as a resource pipeline

The PDF's clearest concrete example is the story of the Three Little Pigs. The story is treated as a multiset-rewriting system. Characters and materials are resources. Events are linear implications.

Let the initial state be

$$\Delta_0 = \{\text{pig}, \text{pig}, \text{pig}, \text{straw}, \text{sticks}, \text{bricks}, \text{wolf}\}.$$

The story rules are:

$$\begin{aligned} r_1 &: \text{pig} \otimes \text{straw} \multimap \text{straw\_house}, \\ r_2 &: \text{pig} \otimes \text{sticks} \multimap \text{stick\_house}, \\ r_3 &: \text{pig} \otimes \text{bricks} \multimap \text{brick\_house}, \\ r_4 &: \text{wolf} \otimes \text{straw\_house} \multimap \text{wolf}, \\ r_5 &: \text{wolf} \otimes \text{stick\_house} \multimap \text{wolf}, \\ r_6 &: \text{wolf} \otimes \text{brick\_house} \multimap \text{brick\_house}. \end{aligned}$$

The canonical ending of the fable is the sequent

$$\Delta_0 \vdash \text{brick\_house}.$$

This says: starting from the initial resources, there is a linear proof whose final surviving resource is the brick house.

Example 12.3.1 (A code-level model of the fable)

The following pseudo-Python is often clearer than a large proof tree.

```
from collections import Counter

state = Counter({
    "pig": 3,
    "straw": 1,
    "sticks": 1,
    "bricks": 1,
    "wolf": 1,
})

rules = [
    (("pig", "straw"),      "straw_house"),
    (("pig", "sticks"),     "stick_house"),
    (("pig", "bricks"),     "brick_house"),
    (("wolf", "straw_house"), "wolf"),
    (("wolf", "stick_house"), "wolf"),
    (("wolf", "brick_house"), "brick_house"),
]

def apply(rule, state):
    inputs, output = rule
    new_state = state.copy()
    for x in inputs:
        if new_state[x] <= 0:
            return None
        new_state[x] -= 1
    new_state[output] += 1
    return +new_state

story = [rules[0], rules[1], rules[2], rules[3], rules[4], rules[5]]
for rule in story:
    state = apply(rule, state)

assert state == Counter({"brick_house": 1})
```

This is exactly the resource-sensitive reading: a pig used to build the straw house is consumed; it cannot simultaneously remain an unused pig. The wolf survives the destruction of the straw and stick houses, but not the encounter with the brick house.

The proof tree in the original PDF should therefore be read as a storyboard. Each inference step is not only a logical step but a plot event. The proof search is the narrative search.

§12.4 Tensor, additives, and branching stories

The tensor $A \otimes B$ means that both resources are available together. In the fable,

pig $\otimes$ **straw**

means that a pig and straw must both be present for the event r_1 to fire.

Additives model choice. The connective $A \& B$ means that the environment may choose which side to test; the proof must be ready for both. This is useful for games and branching narratives.

A compact narrative grammar is:

linear-logic object	narrative reading
Δ	current multiset of actors and props
$A \otimes B$	co-presence of resources
$A \multimap B$	event consuming A and producing B
$A \& B$	branch where the environment chooses
$A \oplus B$	branch where the system chooses
$!A$	persistent rule or reusable condition

This explains why linear logic is attractive for game design. It tracks what is actually available after each action.

§12.5 Dialogue as a game

The second major theme of the PDF is Dr. Peng’s comment that linear logic opened a “new world” because logic ascends to the level of interaction, or even to the level of a game. In game semantics:

propositions are games, proofs are strategies.

A proof of A is not merely a certificate that A is true. It is a strategy for the Proponent to win the game associated with A , no matter how the Opponent plays.

Definition 12.5.1 (Game-semantical slogan). A formula specifies a protocol of interaction. A proof specifies a winning strategy in that protocol.

The cut rule has a dynamic meaning. If one proof produces A and another proof consumes A , cutting them together makes the two strategies interact. Cut elimination is the execution of that interaction until no internal communication remains.

proof of A
 $\Downarrow$
 strategy producing moves of type A
 $\Downarrow$
 cut with a counter-strategy
 $\Downarrow$
 interaction / normalization / play

§12.6 Schopenhauer’s stratagem as a ludic interaction

The PDF’s second example uses a Schopenhauer-style debate. The Opponent says:

“He is a child, so you must make allowance for him.”

The Proponent answers:

“Precisely because he is a child, I must correct him; otherwise he will persist in bad habits.”

The content of the debate is less important than its structure. The Proponent accepts one premise and flips the other.

A code-like representation is:

```
O = {
  "P1": "he is a child",
  "P2": "one must make allowance for children",
  "claim": "do not correct him",
}

P = {
  "accept": "P1",
  "negate": "P2",
  "counterclaim": "because he is a child, correct him",
  "support": "otherwise he will persist in bad habits",
}

def retorsio_argumenti(opponent, proponent):
    assert proponent["accept"] == "P1"
    return {
        "kept": opponent["P1"],
        "reversed": opponent["P2"],
        "new_claim": proponent["counterclaim"],
    }
```

This is close to the Ludics interpretation. A debate is a tree of possible moves. A winning strategy is a path through that tree that answers every possible continuation by the opponent.

§12.7 Ludics: loci instead of signifieds

The PDF’s phrase “the signifier chain without the signifier” should not be read merely as literary theory. In Girard’s Ludics, formulas are replaced by loci: addresses at which interaction takes place.

Definition 12.7.1 (Locus, informal). A locus is an address in a design. Instead of saying that a proof manipulates formulas with fixed semantic content, Ludics says that designs interact at addresses.

The meaning of a proposition is therefore not hidden behind the symbol as a static signified. Meaning is generated by interaction: by how a design responds to its dual.

This gives a cleaner version of the PDF’s Lacanian metaphor:

meaning is not stored inside a sign;
meaning appears through interaction with its negation.

§12.8 Orthogonality and negation

A central operation in Ludics and linear logic is orthogonality. Very roughly, two designs are orthogonal if their interaction terminates properly.

Definition 12.8.1 (Orthogonality, notebook version). Write $D \perp E$ when the interaction between designs D and E converges successfully. Then the orthogonal of a collection A is

$$A^\perp = \{E : \forall D \in A, D \perp E\}.$$

A behavior is often recovered by biorthogonal closure:

$$A = (A^\perp)^\perp.$$

This is the mathematically controlled form of the PDF's claim that a proposition is defined by its negation. One does not understand A by staring at A alone. One understands A through the class of all counter-designs with which it can interact.

§12.9 Geometry of interaction

Geometry of Interaction is the diagrammatic and dynamical side of the same idea. A proof is not a static derivation tree but a flow of information. The PDF's single most useful visual artifact is a colored morphism diagram from Melliès-style topologies.

$$La \longrightarrow L(*m) \otimes L((Rm) \otimes a)$$

may be depicted in the following way:




Figure 12.1: Geometry-of-Interaction morphism.

The visible morphism has the form

$$La \longrightarrow L(*m) \otimes L((Rm) \otimes a).$$

For a first reading, the notation is less important than the picture. The colored ribbons behave like wires. They route input types into output types, bend around each other, and meet at nodes. This is why the original PDF compares human interaction to Feynman diagrams: logical equivalence becomes a kind of topological deformation of processes.

Remark 12.9.1 — The diagram should not be treated as a decorative metaphor. In *Geometry of Interaction*, cut elimination is modeled by the circulation of information through a network. The graphical form is an attempt to make proof normalization visible.

§12.10 From narrative to social relations

The original PDF ends by extending the same logic to narrative generation and social relations. This is speculative, but it has a precise core. A social or narrative process can be modeled as a resource transformation system:

$$\text{inputs} \otimes \text{labor/process} \multimap \text{outputs}.$$

The important constraint is that resources do not duplicate for free. If a theory of narrative or society assumes arbitrary reuse of agents, time, attention, or material goods, it is silently using weakening and contraction. Linear logic forces these assumptions into the open.

A minimal production rule looks like:

```
production_rule = {
  "inputs":  ["labor", "wood", "tools"],
  "output":  "table",
  "logic":   "labor tensor wood tensor tools -o table",
}
```

The same pattern covers the fable rule

$$\text{pig} \otimes \text{bricks} \multimap \text{brick\_house}.$$

This is why the PDF's jump from the Three Little Pigs to social production is not arbitrary. Both are resource-sensitive narratives.

§12.11 The general pattern

The reconstructed message is simple:

$$\begin{array}{c}
 \text{logic is not only truth classification} \\
 \Downarrow \\
 \text{a proof is a resource-sensitive process} \\
 \Downarrow \\
 \text{a process can be read as a game} \\
 \Downarrow \\
 \text{a game can be read as dialogue} \\
 \Downarrow \\
 \text{dialogue can be visualized as Geometry of Interaction.}
 \end{array}$$

The first version of the PDF used large philosophical language: signifier chains, narrative, game, topology, dialectics, social relations. The cleaned-up mathematical version keeps the ambition but anchors it in four concrete devices:

$$A \multimap B, \quad A \otimes B, \quad A^\perp, \quad A = (A^\perp)^\perp.$$

Linear logic is the logic of finite resources; Ludics is the logic of interaction at addresses; game semantics is the logic of strategies; Geometry of Interaction is the logic of information flow. Together they explain why a proof can look like a story, a debate, or a diagram of motion.

Category Theory Continued and the Question Topology

N-20221227

§13.1 More examples of categories

This note continues the category-theory examples. Abelian groups with group homomorphisms form **Ab**. Modules over a ring A form **Mod** $_A$. Rings with unital ring homomorphisms form **Rings**. Topological spaces with continuous maps form **Top**, whose isomorphisms are homeomorphisms.

A partially ordered set $(S, \geq)$ can be regarded as a category: the objects are elements of S , and there is exactly one morphism $x \rightarrow y$ iff $x \geq y$. Reflexivity gives identities, transitivity gives composition, and antisymmetry ensures isomorphisms are trivial.

§13.2 Subcategories and functors

A subcategory $\mathcal{A} \subseteq \mathcal{C}$ has some objects and some morphisms of $\mathcal{C}$, closed under identities and composition. A functor

$$F : \mathcal{C} \rightarrow \mathcal{D}$$

sends objects to objects and morphisms to morphisms, preserving identities and composition:

$$F(1_A) = 1_{F(A)}, \quad F(g \circ f) = F(g) \circ F(f).$$

The diagram in the original page is a standard functoriality square: the image of a composable path under F is still composable, and the composite maps to the composite.

§13.3 Forgetful functors

A forgetful functor forgets structure. For example,

$$\mathbf{Grp} \rightarrow \mathbf{Set}$$

sends a group to its underlying set and a homomorphism to its underlying function. Similarly,

$$\mathbf{Top} \rightarrow \mathbf{Set}$$

forgets open sets, and $\mathbf{Vect}_k \rightarrow \mathbf{Set}$ forgets vector-space operations.

The note's warning is that a forgetful functor is not innocent: once structure is forgotten, many questions cannot be answered anymore. A bijection of underlying sets need not be an isomorphism of groups, vector spaces, or topological spaces.

§13.4 A topology generated by a distance-like question

The second half moves into topology. The page considers a collection of subsets and asks what topology it generates. Given a basis-like family $\mathcal{B}$, the topology generated by it consists of arbitrary unions of finite intersections of elements of $\mathcal{B}$.

The displayed rational-looking function and its graph on the page are used as a reminder: topology is not only about familiar Euclidean open intervals. One can impose a topology by declaring certain sets to be basic neighborhoods.

§13.5 Posets as categories, with an example

A partially ordered set becomes a category in a very economical way. There is a unique morphism $x \rightarrow y$ exactly when $x \leq y$. Identity morphisms come from reflexivity $x \leq x$, and composition comes from transitivity: if $x \leq y$ and $y \leq z$, then $x \leq z$.

This example is important because it shows that a category need not look like a collection of algebraic objects and homomorphisms. A category can encode logic or order. Meets and joins in a poset become limits and colimits in the corresponding category.

§13.6 Functorial thinking

A functor is not merely a map between categories. It is a structure-preserving translation. If $F : \mathcal{C} \rightarrow \mathcal{D}$, then for every composable pair

$$A \xrightarrow{f} B \xrightarrow{g} C$$

we require

$$F(g \circ f) = F(g) \circ F(f).$$

This condition says that doing a construction after composing maps gives the same result as first applying the construction to each map and then composing.

Examples clarify the point. The fundamental group construction is a functor from pointed topological spaces to groups. Homology is a functor from spaces to graded abelian groups. The forgetful functor from groups to sets forgets multiplication but still sends homomorphisms to functions.

§13.7 Generated topology

The topology part of the note concerns how open sets can be generated from a chosen collection. If $\mathcal{B}$ is intended as a basis, two conditions are expected: every point lies in some basis element, and intersections of basis elements are locally refined by basis elements. Then every open set is a union of basis elements.

If one starts with a subbasis $\mathcal{S}$, the generated topology consists of arbitrary unions of finite intersections of elements of $\mathcal{S}$. This explains the phrase “generated by”: we first force the chosen sets to be open, then close under the axioms of topology.

§13.8 Why category theory appears beside topology

Topological constructions are full of universal properties. Product topology is the coarsest topology making the projection maps continuous. Quotient topology is the finest topology making the quotient map continuous. These statements are categorical: they characterize objects by how maps into or out of them behave.

Thus the note’s shift from categories to topology is natural. Topology is not only a study of open sets; it is a study of spaces through continuous maps, and that is exactly categorical thinking.

§13.9 The larger pattern

Category theory and topology meet through structure-preserving maps. A functor preserves categorical structure; a continuous map preserves open-set structure. A forgetful functor deliberately discards structure. This is why the note naturally moves from categories to topology: both ask what data must be preserved for two objects to be considered the same.

§13.10 Posets as categories and logic

A poset viewed as a category has at most one morphism between any two objects. In this setting, categorical statements become order-theoretic statements. A product is a meet, a coproduct is a join, an initial object is a least element, and a terminal object is a greatest element.

For example, in the poset of open subsets of a space ordered by inclusion, intersections behave like meets and unions behave like joins. This is why topology and order theory are so close: open sets form a lattice, and categorical language packages its universal properties.

§13.11 Forgetful functors and free constructions

A forgetful functor loses structure, but it often has a left adjoint that freely adds structure. For instance,

$$U : \mathbf{Grp} \rightarrow \mathbf{Set}$$

forgets the group operation. Its left adjoint sends a set S to the free group $F(S)$. The adjunction says that a set map

$$S \rightarrow U(G)$$

corresponds uniquely to a group homomorphism

$$F(S) \rightarrow G.$$

This is the categorical form of "define a homomorphism out of a free object by choosing images of generators."

§13.12 Initial topology generated by questions

The topology generated by a family of maps can be understood as an initial topology. Suppose we have functions

$$f_i : X \rightarrow Y_i$$

where each Y_i is already topological. The initial topology on X is the coarsest topology making all f_i continuous. Its subbasis is

$$f_i^{-1}(U), \quad U \subseteq Y_i \text{ open.}$$

This matches the note's "topology generated by a question" intuition: each map f_i is a way of observing X , and the generated topology is exactly the open-set structure forced by those observations.

§13.13 Product topology as an initial topology

The product topology is the initial topology for the projection maps

$$\pi_X : X \times Y \rightarrow X, \quad \pi_Y : X \times Y \rightarrow Y.$$

A map $h : Z \rightarrow X \times Y$ is continuous iff the coordinate maps $\pi_X \circ h$ and $\pi_Y \circ h$ are continuous. This is a categorical universal property, not a visual fact about rectangles.

The rectangle basis $U \times V$ appears because

$$U \times V = \pi_X^{-1}(U) \cap \pi_Y^{-1}(V).$$

Thus the usual construction of product open sets is forced by the initial-topology viewpoint.

§13.14 Quotient topology as a final topology

Dually, quotient topology is a final topology. For a quotient map

$$q : X \rightarrow Y,$$

the topology on Y is the finest topology making q continuous, and a map $f : Y \rightarrow Z$ is continuous iff $f \circ q : X \rightarrow Z$ is continuous.

Thus product and quotient are opposite-looking constructions:

initial topology	maps into known spaces	coarsest topology
final topology	maps from known spaces	finest topology

This is the precise categorical version of the topology examples in the note.

§13.15 Functors from topology to algebra

The note mentions that the fundamental group is functorial. More generally, algebraic topology studies functors such as

$$\pi_1 : \mathbf{Top}_* \rightarrow \mathbf{Grp}, \quad H_n : \mathbf{Top} \rightarrow \mathbf{Ab}.$$

A continuous map of spaces induces a homomorphism of groups or abelian groups. This turns geometric questions into algebraic ones.

The power of this move is that algebraic invariants can obstruct homeomorphisms. If two spaces have non-isomorphic fundamental groups or homology groups, they cannot be the same topological space.

§13.16 Initial and final structures in topology

The category-theory material is not separate from topology. Posets show that categories can encode order and logic. Functors express structure-preserving translation. Generated topologies, products, and quotients are governed by universal properties. Algebraic topology then uses functors to extract algebra from spaces. The common theme is: understand objects through the maps they admit.

Category Theory Continued: Posets, Groups as One-Object Categories, Products, and Monomorphisms

N-20230313

§14.1 What this note continues

This note continues the elementary category-theory reading. The source pages are not just a list of definitions. They contain several examples whose purpose is to change one's intuition about what an arrow may be. In ordinary algebra courses, a morphism is often a function preserving structure. Category theory is more flexible: arrows can be order relations, group elements, inclusions, or constructions whose nature depends on the category.

The main examples are:

posets as categories, groups as one-object categories,
products of categories, monomorphisms and epimorphisms.

§14.2 Posets are categories

Let $(P, \leq)$ be a partially ordered set. We can build a category, still denoted P , as follows. The objects are the elements of P . For two objects $p, q \in P$, define

$$\text{Hom}(p, q) = \begin{cases} \{*\}, & p \leq q, \\ \emptyset, & p \not\leq q. \end{cases}$$

Thus there is at most one arrow from p to q . The identity arrow exists because $p \leq p$. Composition exists because order is transitive: if $p \leq q$ and $q \leq r$, then $p \leq r$.

This example is surprisingly important. It shows that arrows in a category need not be functions. In the poset category, the arrow $p \rightarrow q$ is the fact that $p \leq q$.

§14.3 A small poset example

Suppose P has four objects a, b, c, d with relations

$$a \leq b, \quad a \leq c \leq d, \quad a \leq d,$$

and identities at every object. The corresponding category has arrows exactly for these order relations. In a Hasse diagram, upward paths encode arrows, but the category contains the composite arrow as well. For example, from $a \leq c$ and $c \leq d$, the composite gives the arrow $a \rightarrow d$.

The handwritten note emphasizes the point: the same symbol $\leq$ can be treated as an arrow. This is not a metaphor; it satisfies the category axioms.

§14.4 Isomorphisms in a poset category

An isomorphism in a category is an arrow

$$f : A \rightarrow B$$

for which there exists

$$g : B \rightarrow A$$

such that

$$g \circ f = \text{id}_A, \quad f \circ g = \text{id}_B.$$

In a poset category, an arrow $p \rightarrow q$ means $p \leq q$. An isomorphism between p and q would require arrows both ways:

$$p \leq q, \quad q \leq p.$$

By antisymmetry of a partial order, this forces $p = q$. Therefore no two distinct objects of a poset category are isomorphic.

This explains the source's short handwritten answer, "Obviously." It is obvious only after translating categorical isomorphism back into the order axioms.

§14.5 Groups as one-object categories

A group G can be regarded as a category $\mathcal{G}$ with one object $*$. The arrows from $*$ to itself are the elements of G :

$$\text{Hom}_{\mathcal{G}}(*, *) = G.$$

Composition of arrows is group multiplication. The identity arrow is the identity element $e \in G$. Inverses of arrows are group inverses.

Thus a group is exactly a category with one object in which every arrow is an isomorphism. This is not a trick. It says that a group can be understood as the collection of all symmetries of one object, with composition as the operation.

§14.6 Why this example matters

If one first learns category theory, a group may look like a special category: one object, only reversible arrows. This reverses the usual learning order. The note quotes the idea that a group is what one gets by collecting all symmetries of an object. From the categorical viewpoint, a group is the extreme case where there is only one object and every transformation is invertible.

The slogan is:

$$\text{monoid} = \text{one-object category}, \quad \text{group} = \text{one-object groupoid}.$$

A groupoid is a category in which every arrow is invertible.

§14.7 Products of categories

Given categories $\mathcal{C}$ and $\mathcal{D}$, their product category $\mathcal{C} \times \mathcal{D}$ has objects

$$(C, D), \quad C \in \mathcal{C}, \quad D \in \mathcal{D},$$

and morphisms

$$(f, g) : (C, D) \rightarrow (C', D')$$

where $f : C \rightarrow C'$ is a morphism in $\mathcal{C}$ and $g : D \rightarrow D'$ is a morphism in $\mathcal{D}$. Composition is componentwise:

$$(f', g') \circ (f, g) = (f' \circ f, g' \circ g).$$

The identity on (C, D) is $(\text{id}_C, \text{id}_D)$.

This is the categorical version of taking a product structure coordinate by coordinate.

§14.8 The opposite category

For any category $\mathcal{C}$, the opposite category $\mathcal{C}^{op}$ has the same objects as $\mathcal{C}$, but every arrow is reversed:

$$f : A \rightarrow B \text{ in } \mathcal{C} \quad \longleftrightarrow \quad f^{op} : B \rightarrow A \text{ in } \mathcal{C}^{op}.$$

Composition is reversed accordingly. The opposite category is a formal way to turn statements about maps out of an object into statements about maps into an object.

For posets, passing to the opposite category corresponds to reversing the order. If P is viewed as a category, then P^{op} is the poset with

$$p \leq_{op} q \quad \Longleftrightarrow \quad q \leq p.$$

§14.9 Monomorphisms and epimorphisms

In ordinary set theory, an injective function is one that does not identify distinct elements. Category theory cannot always talk about elements, so it defines a left-cancellation property instead.

A morphism $f : X \rightarrow Y$ is a monomorphism if for every pair of arrows $g, h : Z \rightarrow X$,

$$f \circ g = f \circ h \quad \implies \quad g = h.$$

It is an epimorphism if for every pair of arrows $g, h : Y \rightarrow Z$,

$$g \circ f = h \circ f \quad \implies \quad g = h.$$

In **Set**, monomorphisms are exactly injective functions, and epimorphisms are exactly surjective functions. In other categories, this may fail, which is why the categorical definitions are more fundamental.

§14.10 The lesson of the note

The note's guiding idea is that category theory asks us to stop identifying arrows with functions. In a poset, arrows are inequalities. In a group-as-category, arrows are group elements. In a product category, arrows are pairs of arrows. In an opposite category, arrows are formally reversed. Monomorphism and epimorphism then express injective-like and surjective-like behavior without mentioning elements.

The transition is from element-based thinking to arrow-based thinking. This is the first real shift in category theory.

Category Theory Continued: Free and Forgetful Functors, Concrete Categories, and Group Actions

N-20230314

§15.1 Free and forgetful functors

The first page discusses free and forgetful functors. These words are often used informally before one learns the general adjunction language. The basic pattern is simple: a forgetful functor removes structure, while a free functor freely adds structure.

For example, the forgetful functor

$$U : \mathbf{Vect}_k \rightarrow \mathbf{Set}$$

sends a vector space to its underlying set and a linear map to its underlying function. Similarly,

$$\mathbf{Grp} \rightarrow \mathbf{Set}, \quad \mathbf{CRing} \rightarrow \mathbf{Set}$$

forget group or ring structure. There is also a forgetful functor

$$\mathbf{CRing} \rightarrow \mathbf{Grp}$$

that sends a commutative ring R to its additive abelian group $(R, +)$.

A free functor goes in the opposite direction. The free vector space functor

$$F : \mathbf{Set} \rightarrow \mathbf{Vect}_k$$

sends a set S to the vector space with basis S . A function $f : S \rightarrow T$ induces a linear map $F(S) \rightarrow F(T)$ by sending the basis vector s to the basis vector $f(s)$ and extending linearly.

The slogan is:

free = add the least structure needed,
forgetful = discard structure but keep underlying data.

§15.2 Faithful and full functors

For a functor $F : \mathcal{A} \rightarrow \mathcal{B}$, each pair of objects A_1, A_2 gives a map on hom-sets:

$$F : \mathrm{Hom}_{\mathcal{A}}(A_1, A_2) \rightarrow \mathrm{Hom}_{\mathcal{B}}(F(A_1), F(A_2)).$$

The functor is faithful if this map is injective for every pair of objects. It is full if this map is surjective for every pair.

Faithful means that distinct arrows remain distinct after applying F . Full means that every arrow between the images comes from an arrow upstairs. A forgetful functor is often faithful but not full. For instance, not every function between underlying sets of groups is a group homomorphism.

§15.3 Concrete categories

A concrete category is a category equipped with a faithful functor to **Set**:

$$U : \mathcal{A} \rightarrow \mathbf{Set}.$$

This means every object has an underlying set and every arrow has an underlying function, in a way that does not collapse distinct arrows.

The handwritten note is careful here: many familiar categories are concrete, such as groups, rings, vector spaces, and topological spaces. But the word “concrete” is extra structure, not just a vibe. It depends on choosing a faithful functor to sets.

§15.4 Functors out of a one-object group category

Let G be a group and let $\mathcal{G}$ be its associated one-object category. There is one object $*$, and

$$\mathrm{Hom}_{\mathcal{G}}(*, *) = G.$$

Consider a functor

$$F : \mathcal{G} \rightarrow \mathbf{Set}.$$

The object $*$ is sent to a set

$$S = F(*).$$

Each group element $g \in G$, regarded as an arrow $* \rightarrow *$, is sent to a function

$$F(g) : S \rightarrow S.$$

Since functors preserve identity and composition,

$$F(e) = \mathrm{id}_S, \quad F(gh) = F(g) \circ F(h)$$

up to the convention of left or right actions. This is exactly a group action of G on S .

Thus:

$$\text{functor } \mathcal{G} \rightarrow \mathbf{Set} \quad \Longleftrightarrow \quad \text{a } G\text{-set}.$$

This is an early example of category theory turning algebra into actions that preserve structure.

§15.5 Functors between one-object group categories

If G and H are groups and $\mathcal{G}, \mathcal{H}$ their one-object categories, then a functor

$$F : \mathcal{G} \rightarrow \mathcal{H}$$

must send the unique object of $\mathcal{G}$ to the unique object of $\mathcal{H}$. On arrows, it gives a function

$$G \rightarrow H.$$

Functoriality requires

$$F(gh) = F(g)F(h), \quad F(e_G) = e_H.$$

So functors $\mathcal{G} \rightarrow \mathcal{H}$ are exactly group homomorphisms $G \rightarrow H$.

This shows again that functoriality is not an extra decoration. It is exactly the condition that algebraic composition is preserved.

§15.6 Why the note says “learn algebra first”

The handwritten page remarks that one should learn enough algebra before pushing further. That is reasonable: the example of a group as a one-object category is conceptually simple only if one already understands group operations, homomorphisms, and group actions. Without those, the category-theoretic formulation can feel like empty abstraction.

The actual learning path is circular in a productive way. Algebra supplies examples for category theory; category theory then reorganizes algebra.

§15.7 Looking ahead: natural transformations

The final handwritten comment says that natural transformations are coming next. This is the right next step. If functors are structure-preserving maps between categories, then natural transformations are structure-preserving maps between functors.

In the group-action example, a natural transformation between two functors $\mathcal{G} \rightarrow \mathbf{Set}$ is a function between the underlying G -sets that commutes with the G -action. In other words, it is an equivariant map. This is the categorical reason natural transformations are not optional: they are the morphisms between representations of the same structure.

Topos Theory: Categorical Preliminaries, Internal Logic, and the Geometry of Sets

N-20240211

§16.1 Why topoi enter the story

A topos is one of the places where several lines of mathematics meet: category theory, logic, geometry, sheaf theory, and foundations. I want to regard it not merely as an abstract generalization of set theory, but as a way of making precise the idea that “sets may vary over a space” and that logic may depend on the ambient universe in which it is interpreted.

The route of this chapter is therefore:

categories $\rightarrow$ limits and universal properties
 $\rightarrow$ cartesian closure $\rightarrow$ subobject classifiers
 $\rightarrow$ elementary topoi $\rightarrow$ sheaf topoi and internal logic.

§16.2 Categories

Definition 16.2.1. A category $\mathcal{C}$ consists of:

- (i) a class of objects $\text{Ob}(\mathcal{C})$;
- (ii) for each pair of objects X, Y , a collection $\text{Hom}_{\mathcal{C}}(X, Y)$ of morphisms from X to Y ;
- (iii) a composition operation

$$\text{Hom}_{\mathcal{C}}(Y, Z) \times \text{Hom}_{\mathcal{C}}(X, Y) \longrightarrow \text{Hom}_{\mathcal{C}}(X, Z),$$

written $(g, f) \mapsto g \circ f$;

- (iv) for each object X , an identity morphism $\text{id}_X : X \rightarrow X$.

These data satisfy associativity and identity laws:

$$h \circ (g \circ f) = (h \circ g) \circ f, \quad f \circ \text{id}_X = f, \quad \text{id}_Y \circ f = f.$$

Example 16.2.2

The category **Set** has sets as objects and functions as morphisms. The category **Grp** has groups and group homomorphisms. The category **Top** has topological spaces and continuous maps.

Remark 16.2.3 — The categorical viewpoint is not that objects become unimportant, but that objects are understood through their maps. This is exactly the shift that later makes generalized set theories possible.

§16.3 Universal properties and limits

The most important categorical habit is to define objects not by their internal elements, but by how maps into or out of them behave.

Definition 16.3.1. A terminal object in $\mathcal{C}$ is an object 1 such that for every object X , there is a unique morphism

$$X \rightarrow 1.$$

Dually, an initial object is an object 0 such that for every X , there is a unique morphism

$$0 \rightarrow X.$$

In **Set**, a singleton is terminal and the empty set is initial.

Definition 16.3.2. A product of X and Y is an object $X \times Y$, equipped with projections

$$\pi_X : X \times Y \rightarrow X, \quad \pi_Y : X \times Y \rightarrow Y,$$

such that for every object Z and every pair of maps $f : Z \rightarrow X$, $g : Z \rightarrow Y$, there exists a unique map $\langle f, g \rangle : Z \rightarrow X \times Y$ making the diagram commute:

$$\begin{array}{ccc} & Z & \\ f \swarrow & \downarrow \langle f, g \rangle & \searrow g \\ & X \times Y & \\ \pi_X \swarrow & & \searrow \pi_Y \\ X & & Y. \end{array}$$

Definition 16.3.3. An equalizer of parallel morphisms $f, g : X \rightarrow Y$ is a morphism $e : E \rightarrow X$ such that

$$f \circ e = g \circ e,$$

and such that any other map $h : Z \rightarrow X$ with $f \circ h = g \circ h$ factors uniquely through e :

$$\begin{array}{ccccc} Z & \xrightarrow{h} & E & \xrightarrow{e} & X & \xrightarrow[f]{g} & Y \\ & \dashrightarrow & & & & & \end{array}$$

Definition 16.3.4. A pullback of $f : X \rightarrow Z$ and $g : Y \rightarrow Z$ is an object $X \times_Z Y$ fitting into a cartesian square

$$\begin{array}{ccc} X \times_Z Y & \longrightarrow & Y \\ \downarrow & & \downarrow g \\ X & \xrightarrow{f} & Z \end{array}$$

with the universal property that any object mapping compatibly to X and Y factors uniquely through the pullback.

Remark 16.3.5 — Products, equalizers, and pullbacks are not isolated gadgets. They are instances of limits. A limit is a universal cone over a diagram; it is the categorical way to say that an object satisfies all the compatibility conditions encoded by the diagram.

§16.4 Exponentials and cartesian closure

In set theory, the set of functions from X to Y is denoted Y^X . Categorically, this is characterized by an adjunction.

Definition 16.4.1. A category with finite products is cartesian closed if for every pair of objects X, Y , there exists an exponential object Y^X and a natural bijection

$$\mathrm{Hom}_{\mathcal{C}}(Z, Y^X) \cong \mathrm{Hom}_{\mathcal{C}}(Z \times X, Y).$$

The corresponding evaluation map is

$$\mathrm{ev} : Y^X \times X \rightarrow Y.$$

A morphism $Z \rightarrow Y^X$ is the same thing as a morphism $Z \times X \rightarrow Y$. This is currying in categorical form:

$$\begin{array}{ccc} Z \times X & \xrightarrow{\tilde{f}} & Y \\ & \searrow \hat{f} \times \mathrm{id}_X & \nearrow \mathrm{ev} \\ & Y^X \times X & \end{array}$$

Example 16.4.2

In **Set**, Y^X is the ordinary set of functions $X \rightarrow Y$. The bijection

$$\mathrm{Hom}_{\mathbf{Set}}(Z, Y^X) \cong \mathrm{Hom}_{\mathbf{Set}}(Z \times X, Y)$$

says that a Z -indexed family of functions $X \rightarrow Y$ is the same thing as a function of two variables $Z \times X \rightarrow Y$.

§16.5 Subobjects and the subobject classifier

A subset $A \subseteq X$ may be represented by its characteristic function

$$\chi_A : X \rightarrow \{0, 1\}.$$

The categorical version replaces $\{0, 1\}$ by an object Ω , called the subobject classifier.

Definition 16.5.1. A subobject classifier in a category with a terminal object is an object Ω , together with a map

$$\mathrm{true} : 1 \rightarrow \Omega,$$

such that for every monomorphism $m : A \hookrightarrow X$, there is a unique characteristic morphism $\chi_A : X \rightarrow \Omega$ making the square a pullback:

$$\begin{array}{ccc} A & \longrightarrow & 1 \\ m \downarrow & & \downarrow \mathrm{true} \\ X & \xrightarrow{\chi_A} & \Omega. \end{array}$$

In **Set**, $\Omega = \{0, 1\}$, true picks 1, and χ_A is the ordinary characteristic function.

Remark 16.5.2 — This is one of the decisive steps toward topos theory. A topos has an internal object of truth values, so logic becomes an internal structure of the category rather than an external metalanguage imposed from outside.

§16.6 Elementary topoi

Definition 16.6.1. An elementary topos is a category $\mathcal{E}$ satisfying:

- (i) it has finite limits;
- (ii) it is cartesian closed;
- (iii) it has a subobject classifier.

This definition is strikingly short. It says that a category behaves enough like **Set** to support finite constructions, function spaces, and a logic of predicates.

Example 16.6.2

The category **Set** is an elementary topos. More generally, for any small category $\mathcal{C}$, the presheaf category

$$\widehat{\mathcal{C}} = \mathbf{Set}^{\mathcal{C}^{op}}$$

is an elementary topos.

Example 16.6.3

If X is a topological space, the category $\mathbf{Sh}(X)$ of sheaves on X is a topos. It behaves like a category of sets varying continuously over X .

§16.7 A toy presheaf topos

The smallest example that is not simply **Set** is already useful. Let

$$\mathbf{2} = (0 \rightarrow 1)$$

be the category associated to the two-element poset. Consider the presheaf category

$$\widehat{\mathbf{2}} = \mathbf{Set}^{\mathbf{2}^{op}}.$$

Every presheaf category is a topos, so this is a genuine elementary topos.

An object $P \in \widehat{\mathbf{2}}$ is exactly the data

$$P(1) \xrightarrow{r} P(0),$$

because the arrow $0 \rightarrow 1$ in $\mathbf{2}$ becomes a restriction map $P(1) \rightarrow P(0)$ in $\mathbf{2}^{op}$. It is safest to think of this as a two-stage set: elements visible at stage 1 restrict to elements visible at stage 0. It is not automatically an equivalence relation or a quotient; it is simply a function carrying information from the higher stage to the lower stage.

The subobject classifier is made of sieves. With this convention, the sieves on 0 are

$$\Omega(0) = \{\emptyset, \{id_0\}\},$$

while the sieves on 1 are

$$\Omega(1) = \{\emptyset, \{u : 0 \rightarrow 1\}, \{id_1, u\}\}.$$

Thus the three-valued part occurs at stage 1, and the restriction map

$$\Omega(1) \longrightarrow \Omega(0)$$

is

$$\emptyset \mapsto \emptyset, \quad \{u\} \mapsto \{id_0\}, \quad \{id_1, u\} \mapsto \{id_0\}.$$

The element $\{u\}$ is the interesting third truth value: not true at stage 1 itself, but true after restricting along $0 \rightarrow 1$.

Here is a characteristic map in this toy universe. Let

$$P(1) = \{a, b\}, \quad P(0) = \{*\},$$

with both a and b restricting to $*$. Define a subpresheaf $Q \hookrightarrow P$ by

$$Q(1) = \{a\}, \quad Q(0) = \{*\}.$$

The characteristic morphism $\chi_Q : P \rightarrow \Omega$ sends

$$\chi_Q(1)(a) = \{id_1, u\}, \quad \chi_Q(1)(b) = \{u\},$$

and

$$\chi_Q(0)(*) = \{id_0\}.$$

So a is fully true, while b is not true at stage 1 but becomes true after restriction to stage 0. This is the point at which internal logic stops being two-valued. In $\mathbf{Set}^{2^{op}}$, a proposition may have a stage-dependent truth value; truth is not merely “yes” or “no”, but is tracked through the restriction structure.

§16.8 Grothendieck topoi and sheaves

A Grothendieck topos is, roughly, a category equivalent to sheaves on a site. A site is a category equipped with a notion of covering. Thus Grothendieck topoi generalize sheaves on topological spaces.

Remark 16.8.1 — This is the geometric origin of topos theory. Instead of studying a space only through its points, one studies its sheaves. In many situations, the sheaf category remembers more robust information than the point set itself.

In algebraic geometry, this becomes essential. The étale topos of a scheme, for example, packages arithmetic and geometric information that is invisible from the underlying topological space alone.

§16.9 Internal logic

Every elementary topos carries an internal higher-order intuitionistic logic. Subobjects of an object X behave like predicates on X , and the subobject classifier Ω behaves like the object of truth values.

Remark 16.9.1 — The internal logic of a general topos is usually intuitionistic, not necessarily classical. The law of excluded middle

$$P \vee \neg P$$

need not hold internally. Thus moving from $\mathbf{Set}$ to a general topos also changes the logic that is naturally available.

For a sheaf topos, this has a geometric meaning: truth may be local. A statement may hold on an open cover without having a single global witness. This is exactly the kind of phenomenon that sheaf theory is designed to express.

§16.10 Power objects

In set theory, $\mathcal{P}(X)$ is the set of subsets of X . In a topos, the analogue is the power object

$$\Omega^X.$$

Indeed, by cartesian closure and the subobject classifier,

$$\mathrm{Hom}_{\mathcal{E}}(B, \Omega^X) \cong \mathrm{Hom}_{\mathcal{E}}(B \times X, \Omega),$$

which classifies subobjects of $B \times X$. Thus Ω^X parametrizes subobjects of X , just as 2^X parametrizes subsets of X in ordinary set theory.

§16.11 Topos as generalized universe of sets

A topos can be read in two complementary ways:

- (i) as a generalized universe of sets, where one can do mathematics internally;
- (ii) as a generalized space, where sheaves and local data become fundamental.

This double reading is the source of its power. In one direction, topos theory generalizes set-theoretic foundations. In the other direction, it turns geometry into logic.

§16.12 A high-level comparison

Setting	Truth values	Subsets/predicates
Set	$\{0, 1\}$	characteristic functions $X \rightarrow 2$
Boolean topos	Boolean algebra object	classical internal logic
General topos	Ω	subobjects classified by $X \rightarrow \Omega$
Sheaf topos	local truth values	predicates varying over a space

Remark 16.12.1 — From a higher viewpoint, a topos is not merely a category satisfying three axioms. It is a world in which mathematics can be interpreted. The shift from sets to topoi is analogous to the shift from Euclidean geometry to geometry relative to a transformation group: the ambient universe becomes variable.

§16.13 Further directions

Several directions naturally follow:

- (i) Lawvere–Tierney topologies, which internalize sheafification;
- (ii) geometric morphisms, the correct maps between topoi;
- (iii) classifying topoi, which classify models of geometric theories;

- (iv) Grothendieck topoi in algebraic geometry;
- (v) elementary topoi as foundations for constructive mathematics.

Question

The most important question for me is not only what a topos is, but what it means to do mathematics inside a topos. Once truth, subsets, and function spaces become internal categorical objects, foundations stop being a fixed background and become part of the geometry.

§16.14 Subobjects and characteristic maps

In **Set**, a subset $A \subseteq X$ is classified by its characteristic function

$$\chi_A : X \rightarrow \{0, 1\}.$$

A subobject classifier generalizes this. In a topos, there is an object Ω and a map

$$\text{true} : 1 \rightarrow \Omega$$

such that every subobject $A \hookrightarrow X$ is obtained as a pullback of **true** along a unique characteristic map $\chi_A : X \rightarrow \Omega$.

This is the categorical replacement for membership predicates. A subobject is not merely a subset; it is something classified by a truth-value object internal to the topos.

§16.15 Why logic becomes intuitionistic

In a general topos, the object Ω of truth values need not behave like the two-element Boolean algebra. As a result, the internal logic is generally intuitionistic. In particular, the law of excluded middle

$$P \vee \neg P$$

may fail internally.

This is not a defect. It means that truth can vary with geometric or contextual information. In a sheaf topos, a proposition may be true locally on some open sets without being globally decidable.

§16.16 Sheaves as variable sets

A sheaf on a topological space assigns data to each open set, with restriction maps and gluing conditions. One can think of a sheaf as a set varying over the space. The topos of sheaves on X therefore behaves like a universe of variable sets over X .

This is the geometric intuition behind topoi. They are not arbitrary abstract categories; they generalize the category of sets in a way that remembers variation, locality, and gluing.

Kernels, Cokernels, Images, and Coimages in Abelian Categories

N-20251109

§17.1 The setting: maps with zero maps

The uploaded note studies a morphism

$$f : A \rightarrow B$$

from the categorical viewpoint. The words domain and codomain mean exactly what they mean for an ordinary function:

$$A = \text{domain}, \quad B = \text{codomain}.$$

For linear maps one may think of A as the input space and B as the declared output space. The actual values reached by f may form a smaller object inside B ; that smaller object is the image.

The categorical definitions below need a zero object and zero morphisms. A zero object is simultaneously initial and terminal. Once a category has a zero object 0 , every pair of objects X, Y has a distinguished zero morphism

$$0_{X,Y} : X \rightarrow 0 \rightarrow Y.$$

This lets equations such as $f \circ k = 0$ make sense without mentioning elements.

The most familiar examples are:

$$\mathbf{Vect}_k, \quad \mathbf{Ab}, \quad R\text{-}\mathbf{Mod}.$$

These are abelian categories. In them, kernels, cokernels, images, and coimages behave like the objects from linear algebra and module theory.

§17.2 Kernel: the universal object killed by a map

For a morphism $f : A \rightarrow B$, a kernel of f is a morphism

$$k : K \rightarrow A$$

such that

$$f \circ k = 0,$$

and universal with this property: whenever $k' : K' \rightarrow A$ also satisfies $f \circ k' = 0$, there is a unique morphism $u : K' \rightarrow K$ with

$$k' = k \circ u.$$

In diagram form:

$$\begin{array}{ccccc} K' & & & & \\ \downarrow u & \searrow k' & & & \\ K & \xrightarrow{k} & A & \xrightarrow{f} & B. \end{array}$$

The condition is that both routes from K' to B are zero after applying f .

In vector spaces this recovers

$$\text{Ker}(f) = \{a \in A : f(a) = 0\}.$$

The categorical definition says something stronger than just listing elements: the kernel is the best, largest, and universal way to factor all inputs that become zero.

In an abelian category, the kernel morphism $k : K \rightarrow A$ is a monomorphism. It is therefore legitimate to think of K as a subobject of the domain A , though the invariant datum is the arrow k , not merely the object K .

§17.3 Cokernel: the universal quotient killing a map

The cokernel is the dual construction. A cokernel of $f : A \rightarrow B$ is a morphism

$$c : B \rightarrow C$$

such that

$$c \circ f = 0,$$

and universal with this property: whenever $c' : B \rightarrow C'$ satisfies $c' \circ f = 0$, there is a unique morphism $u : C \rightarrow C'$ with

$$c' = u \circ c.$$

The diagram is:

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & \xrightarrow{c} & C \\ & & \searrow c' & & \downarrow u \\ & & & & C'. \end{array}$$

In vector spaces or modules,

$$\text{Coker}(f) = B / \text{Im}(f).$$

Thus the cokernel measures what remains in the codomain after the actual output of f has been collapsed to zero.

In an abelian category, the cokernel morphism $c : B \rightarrow C$ is an epimorphism. It is therefore the categorical version of a quotient object of the codomain.

§17.4 Image: the part of the codomain actually reached

In linear algebra the image is

$$\text{Im}(f) = \{f(a) : a \in A\} \subseteq B.$$

In an abelian category, one defines the image without element language by first taking the cokernel

$$c : B \rightarrow \text{Coker}(f)$$

and then taking the kernel of that cokernel:

$$\text{Im}(f) = \text{Ker}(B \xrightarrow{c} \text{Coker}(f)).$$

So the image is represented by a monomorphism

$$i : \text{Im}(f) \rightarrow B.$$

The diagram is:

$$\text{Im}(f) \xhookrightarrow{i} B \xrightarrow{c} \text{Coker}(f).$$

The equation $c \circ i = 0$ says that everything in the image is killed by the quotient map c . The universal property says that it is the largest subobject of B killed by c .

This gives a useful slogan:

$$\text{Im}(f) = \text{the part of the codomain not visible to the cokernel.}$$

§17.5 Coimage: the effective part of the domain

The coimage is dual to the image. Begin with the kernel

$$k : \text{Ker}(f) \rightarrow A.$$

Then define the coimage as the cokernel of the kernel:

$$\text{Coim}(f) = \text{Coker}(\text{Ker}(f) \xrightarrow{k} A).$$

So it is represented by an epimorphism

$$q : A \rightarrow \text{Coim}(f).$$

The diagram is:

$$\text{Ker}(f) \xhookrightarrow{k} A \xrightarrow{q} \text{Coim}(f).$$

In vector spaces this is the quotient

$$\text{Coim}(f) = A / \text{Ker}(f).$$

It is the domain after removing precisely the part that f sends to zero. In that sense, it is the effective input object of f .

§17.6 The standard factorization

In an abelian category, every morphism $f : A \rightarrow B$ factors through its coimage and image:

$$A \xrightarrow{q} \text{Coim}(f) \xrightarrow{\bar{f}} \text{Im}(f) \xrightarrow{i} B.$$

Here:

- $q : A \rightarrow \text{Coim}(f)$ is the cokernel of $\text{Ker}(f) \rightarrow A$, hence an epimorphism;
- $i : \text{Im}(f) \rightarrow B$ is the kernel of $B \rightarrow \text{Coker}(f)$, hence a monomorphism;
- the middle map $\bar{f} : \text{Coim}(f) \rightarrow \text{Im}(f)$ is an isomorphism in an abelian category;
- the original morphism is recovered by

$$f = i \circ \bar{f} \circ q.$$

A compact diagram is:

$$\text{Ker}(f) \hookrightarrow A \xrightarrow{q} \text{Coim}(f) \xrightarrow[\cong]{\bar{f}} \text{Im}(f) \hookrightarrow B \xrightarrow{c} \text{Coker}(f).$$

f

This is the categorical form of the first isomorphism theorem:

$$A / \text{Ker}(f) \cong \text{Im}(f).$$

The uploaded note's phrase that coimage is effective input and image is actual output is exactly this factorization: q discards useless input, i embeds the actual output into the declared codomain, and $\bar{f}$ identifies these two descriptions.

§17.7 A vector-space sanity check

Let $T : V \rightarrow W$ be a linear map. Then

$$\text{Ker}(T) = \{v \in V : T(v) = 0\}, \quad \text{Coker}(T) = W / \text{Im}(T).$$

The coimage is

$$\text{Coim}(T) = V / \text{Ker}(T),$$

and the induced map

$$\bar{T} : V / \text{Ker}(T) \rightarrow \text{Im}(T), \quad [v] \mapsto T(v),$$

is an isomorphism. It is well-defined because if $v - v' \in \text{Ker}(T)$, then $T(v) = T(v')$. It is injective because $T(v) = 0$ means $[v] = 0$, and it is surjective by definition of image.

For a matrix $A : k^n \rightarrow k^m$, this says that the rank is the dimension of both sides of the middle isomorphism:

$$\dim(k^n / \text{Ker } A) = \dim \text{Im } A = \text{rank } A.$$

This is rank-nullity in categorical clothing.

§17.8 A module example: multiplication by an integer

Consider the group homomorphism

$$f : \mathbb{Z} \rightarrow \mathbb{Z}, \quad f(x) = nx.$$

If $n \neq 0$, then

$$\text{Ker}(f) = 0, \quad \text{Im}(f) = n\mathbb{Z}, \quad \text{Coker}(f) = \mathbb{Z}/n\mathbb{Z}.$$

The coimage is

$$\text{Coim}(f) = \mathbb{Z}/0 \cong \mathbb{Z}.$$

The middle isomorphism is

$$\mathbb{Z} \xrightarrow{\sim} n\mathbb{Z}, \quad x \mapsto nx.$$

This example is useful because the image is literally a subgroup of the codomain, while the coimage is a quotient of the domain. They look different as subobject and quotient, but in an abelian category they are canonically isomorphic through f .

§17.9 Exactness language

Kernel and image are also the grammar of exact sequences. A sequence

$$X \xrightarrow{g} Y \xrightarrow{h} Z$$

is exact at Y when

$$\text{Im}(g) = \text{Ker}(h)$$

as subobjects of Y . The equality means that the elements, vectors, or generalized inputs that h kills are exactly those that came from g .

In categorical language, exactness replaces element chasing by universal properties. For example, saying that

$$\text{Ker}(f) \rightarrow A \rightarrow B$$

is exact at A is precisely saying that the first arrow is the kernel of f . Saying that

$$A \rightarrow B \rightarrow \text{Coker}(f)$$

is exact at B is precisely saying that the second arrow is the cokernel of f , after passing through the image.

§17.10 Warnings: what depends on being abelian

The statements above should not be read as valid in every category. Several layers are being used.

- A category must have a zero object and zero morphisms before kernels and cokernels can even be formulated in this form.
- A category may have kernels and cokernels but still fail to make every monomorphism a kernel or every epimorphism a cokernel.
- The canonical map $\text{Coim}(f) \rightarrow \text{Im}(f)$ need not be an isomorphism outside abelian categories.

Thus the beautiful factorization

$$\text{Coim}(f) \cong \text{Im}(f)$$

is not a formal tautology. It is one of the defining strengths of abelian categories.

For intuition, **Set** is not the right environment for this theory: it has no zero object in the needed sense. Pointed sets have a zero object and kernels in a pointed sense, but they are not abelian. Vector spaces, abelian groups, and modules are the safe models.

§17.11 Synthesis: from element chasing to universal properties

The original note's definitions can be compressed into the following table:

object	categorical construction	linear algebra model	where it lives
$\text{Ker}(f)$	universal map into A killed by f	$\{a : f(a) = 0\}$	subobject of A
$\text{Coker}(f)$	universal quotient of B killing f	$B / \text{Im}(f)$	quotient of B
$\text{Im}(f)$	$\text{Ker}(B \rightarrow \text{Coker}(f))$	$\{f(a) : a \in A\}$	subobject of B
$\text{Coim}(f)$	$\text{Coker}(\text{Ker}(f) \rightarrow A)$	$A / \text{Ker}(f)$	quotient of A

The standard factorization

$$A \twoheadrightarrow \operatorname{Coim}(f) \xrightarrow{\sim} \operatorname{Im}(f) \hookrightarrow B$$

turns a morphism into three conceptual pieces: quotient out the part of the input killed by f , identify the remaining effective input with the subobject actually reached in B , and then include that subobject in the codomain.

For vector spaces this says exactly the familiar rank–nullity picture: $A/\ker f$ is isomorphic to $\operatorname{im} f$. For modules and abelian groups the same statement is the first isomorphism theorem. The categorical formulation is useful because it records which part of the argument is universal and which part depends on being in an abelian category: the comparison map $\operatorname{Coim}(f) \rightarrow \operatorname{Im}(f)$ is an isomorphism in the safe models listed above, but it should not be assumed in an arbitrary category with kernels and cokernels.

III

Number Theory Basics

Part III: Contents

N-20220717	109
18.1 A beginning point: recursive definitions	109
18.2 Arithmetic progressions	109
18.3 Determining an arithmetic progression	110
18.4 Geometric progressions	110
18.5 Converting recurrences	111
18.6 A factorial recurrence	111
18.7 Groups	111
18.8 Symmetric and dihedral groups	112
18.9 Products and isomorphisms	112
18.10 Linear recurrences	112
18.11 Groups from repeated operations	113
18.12 Why the pattern matters	113
18.13 Recurrences as iteration of maps	113
18.14 First-order linear recurrences	113
18.15 Groups behind repeated operations	113
18.16 Sequences as orbits of operations	113
N-20220723	115
19.1 The ant on a decagon	115
19.2 Roots of unity	115
19.3 The polynomial $x^{2n} + x^n + 1$	115
19.4 A broader pattern	116
19.5 Cyclotomic viewpoint	116
19.6 Why this is more than a trick	116
19.7 Roots of unity viewpoint	117
19.8 Cyclotomic divisibility	117
N-20220803	119
20.1 A trigonometric opening example	119
20.2 Logarithmic sequences	119
20.3 Recovering terms from partial sums	120
20.4 Sums and recurrence relations	120
20.5 Geometric progression products	120
20.6 A harmonic-looking sequence	121
20.7 A final synthesis	121
20.8 Sequences as discrete dynamical systems	121
20.9 Partial sums and finite differences	121
20.10 Generating functions for sequence exercises	122
20.11 Telescoping as a coboundary	122
20.12 Second-order recurrences and characteristic roots	122
20.13 Discrete calculus of elementary sequences	123
N-20220913	125
21.1 From guessing to proving	125
21.2 Divisibility	125
21.3 Example: reducing the dividend	126
21.4 Example: polynomial divisibility in one variable	126
21.5 Congruences	126
21.6 A large power sum	127
21.7 A wider interpretation	127
21.8 Divisibility as ideal containment	127
21.9 GCD and LCM as ideal operations	127
21.10 Congruences as quotient rings	128
21.11 Units and Bezout inverses	128
21.12 Valuations as prime coordinates	128
21.13 Chinese remainder theorem as local-to-global	129

21.14	Polynomial divisibility as the same story	129
21.15	Divisibility as quotient algebra	129
N-20220918		131
22.1	A general odd-power divisibility principle	131
22.2	A divisibility problem with 512	131
22.3	Why finite differences work	132
22.4	An alternating harmonic numerator	132
22.5	The next abstraction	132
22.6	Finite differences as discrete derivatives	133
22.7	The binomial basis and integer-valued polynomials	133
22.8	Pairing as group symmetry	133
22.9	P-adic valuation and lifting intuition	134
22.10	Alternating harmonic sums and denominator control	134
22.11	Discrete summation by parts	134
22.12	Discrete calculus and local divisibility	134
N-20221013		135
23.1	Shoelace formula	135
23.2	Why the shoelace formula works	135
23.3	A circle-area maximum problem	136
23.4	Remainder theorem	136
23.5	A polynomial remainder example	136
23.6	Congruence arithmetic	137
23.7	Fermat and Euler	137
23.8	Shoelace formula as oriented area	138
23.9	Evaluation as quotient by $x-a$	138
23.10	Where the calculation points	138
23.11	Shoelace as a discrete line integral	138
23.12	Remainder theorem and quotient rings	139
23.13	Fermat–Euler and unit groups	139
23.14	Quotients in area, polynomials, and congruences	139
N-20221030		141
24.1	The well-ordering principle	141
24.2	Divisibility, gcd, and lcm	141
24.3	Division algorithm	142
24.4	Euclidean algorithm	142
24.5	Example	142
24.6	Fundamental theorem of arithmetic	142
24.7	The Euler phi function	143
24.8	The broader lesson	143
24.9	Well-ordering and Euclidean domains	143
24.10	Bezout and principal ideals	143
24.11	Euler function and unit groups	144
24.12	Euclidean arithmetic as ideals and unit groups	144
24.13	Euclidean domains	144
24.14	PID and Bezout domains	144
24.15	Euler phi as unit counting	144
24.16	Multiplicative functions	145
24.17	Euclidean algorithm and ideals	145
24.18	Euler’s theorem as group theory	145
N-20221031		147
25.1	Divisibility	147
25.2	Division algorithm	147
25.3	Greatest common divisor and least common multiple	148
25.4	Euclidean algorithm	148
25.5	Bezout theorem	148
25.6	Gauss lemma and divisibility consequences	149
25.7	Prime numbers	149

25.8	Congruences	149
25.9	Exercises and common methods	149
25.10	From example to method	150
25.11	Division algorithm as normal form	150
25.12	Euclidean algorithm and invariant gcd	150
25.13	Gauss lemma and prime ideals	150
25.14	Normal forms, gcd invariants, and prime divisibility	150
25.15	$\mathbb{Z}$ as a PID	151
25.16	Modular inverses as units	151
25.17	Smith normal form over integers	151
25.18	CRT as product decomposition	151
25.19	Solving congruences by image and kernel	152
25.20	Smith normal form and congruences	152
N-20221119		153
26.1	Prime and composite numbers	153
26.2	Euclid's theorem	153
26.3	Prime divisibility	153
26.4	Fundamental theorem of arithmetic	154
26.5	Perfect numbers and square numbers	154
26.6	Exponent of a prime in n factorial	155
26.7	Congruence exercises	155
26.8	Broader perspective	155
26.9	Prime exponents and valuations	155
26.10	Divisor functions as multiplicative functions	155
26.11	Congruences and finite groups	156
26.12	Prime coordinates and finite unit groups	156
26.13	Arithmetic functions and convolution	156
26.14	Möbius inversion in number theory	156
26.15	Dirichlet series and Euler products	157
26.16	p -adic valuations and multiplicativity	157
26.17	Euler products and prime independence	157
26.18	Möbius inversion over primes	158
N-20221126		159
27.1	Congruence	159
27.2	Basic properties	159
27.3	Residue classes	159
27.4	Complete and reduced residue systems	160
27.5	Example: last digit	160
27.6	Example: modulo 17	160
27.7	Example: reducing a long expression	161
27.8	A polynomial congruence by induction	161
27.9	Quadratic residues modulo 8	161
27.10	Where the idea leads	161
27.11	Residue systems as representatives	161
27.12	Euler function through prime powers	162
27.13	Orders and exponent reduction	162
27.14	Residues, units, and exponent cycles	162
27.15	Residue systems as groups of units	162
27.16	CRT decomposition of unit groups	162
27.17	Primitive roots and characters	163
27.18	Returning to residue-system exercises	163
27.19	Primitive roots as logarithms in finite groups	163
27.20	Characters as Fourier analysis on unit groups	163
N-20230127		165
28.1	The original problem and the first successful method	165
28.2	The direct factorization attempt	165
28.3	The attempted generalization	166
28.4	The computational experiment	166

28.5	Explicit root formula	167
28.6	Common real factors	167
28.7	The organizing idea	168
28.8	Cyclotomic summary of the experiment	168
28.9	Cyclotomic polynomials	168
28.10	Galois action on roots of unity	168
28.11	Resultants and common roots	169
28.12	Finite fields and character sums	169
28.13	Returning to the original experiments	169
28.14	Primitive roots and cyclotomic factor tests	169
28.15	Cyclotomic identities through Möbius inversion	169

Recurrences, Arithmetic and Geometric Progressions, and the First Examples of Groups

N-20220717

§18.1 A beginning point: recursive definitions

This note begins with the feeling that recurrence relations look like guessing games. In fact, most elementary recurrences are solved by trying to convert them into one of a few familiar sequences.

The first example is a recurrence of the form

$$a_{n+2} = 5a_{n+1} - 6a_n.$$

If one expects powers, the natural test is to try $a_n = r^n$. Substitution gives

$$r^{n+2} = 5r^{n+1} - 6r^n,$$

so

$$r^2 - 5r + 6 = 0, \quad (r - 2)(r - 3) = 0.$$

Thus the general shape is

$$a_n = A2^n + B3^n.$$

This method is not magic: the recurrence is linear, and exponentials are eigenvectors for the shift operator.

§18.2 Arithmetic progressions

An arithmetic progression is a sequence whose consecutive differences are constant:

$$a_{n+1} - a_n = d.$$

Equivalently,

$$a_n = a_1 + (n - 1)d.$$

The sum of the first n terms is

$$S_n = \frac{n}{2}(a_1 + a_n) = na_1 + \frac{n(n-1)}{2}d.$$

The note emphasizes an important point: if several indices have the same sum, then the corresponding terms of an arithmetic progression also have the same sum. For example, if

$$x_1 + x_2 = y_1 + y_2 = m,$$

then

$$a_{x_1} + a_{x_2} = 2a_1 + (x_1 + x_2 - 2)d = 2a_1 + (m - 2)d = a_{y_1} + a_{y_2}.$$

More generally, if

$$x_1 + \cdots + x_k = y_1 + \cdots + y_k,$$

then

$$a_{x_1} + \cdots + a_{x_k} = a_{y_1} + \cdots + a_{y_k}.$$

This is a small but useful structural observation: arithmetic progressions turn index addition into term addition, up to a constant shift.

§18.3 Determining an arithmetic progression

A typical exercise is: given a_2 and a_5 , find a_1 and d . For instance, if

$$a_2 = 5, \quad a_5 = 11,$$

then

$$a_1 + d = 5, \quad a_1 + 4d = 11.$$

Subtracting gives $3d = 6$, so $d = 2$, and then $a_1 = 3$.

Another exercise gives several sums, for example

$$a_1 + a_2 + a_3 = 24, \quad a_n + a_{n+1} + a_{n+2} = 146,$$

and asks for the number of terms or a particular partial sum. The thought process is always the same: express every term as $a_1 + (k-1)d$, reduce to equations in a_1, d, n , then use the sum formula. The original page highlights the formula

$$S_n = \frac{n}{2}(a_1 + a_n)$$

because it often avoids writing every term.

§18.4 Geometric progressions

A geometric progression is a sequence whose ratios are constant:

$$\frac{a_{n+1}}{a_n} = q.$$

Hence

$$a_n = a_1 q^{n-1}.$$

If $q \neq 1$, then

$$S_n = a_1 \frac{q^n - 1}{q - 1}.$$

If $q = 1$, then

$$S_n = na_1.$$

The note gives the derivation by subtracting:

$$S_n = a_1 + a_1 q + \cdots + a_1 q^{n-1},$$

$$qS_n = a_1 q + a_1 q^2 + \cdots + a_1 q^n,$$

so

$$(q-1)S_n = a_1(q^n - 1).$$

This derivation is better than memorization because it also explains why the case $q = 1$ must be separated.

§18.5 Converting recurrences

The page contains the example

$$a_{n+1} = 2a_n + 1.$$

There are two standard ways to solve it.

First, shift the sequence by a constant. Try $b_n = a_n + c$, and choose c so that the recurrence becomes homogeneous. Since

$$a_{n+1} + 1 = 2(a_n + 1),$$

we let $b_n = a_n + 1$. Then

$$b_{n+1} = 2b_n,$$

so $b_n = 2^{n-1}b_1$. If $a_1 = 1$, then $b_1 = 2$, and

$$a_n = 2^n - 1.$$

Second, one may divide by a power of 2. This often works for recurrences of the form

$$a_{n+1} = qa_n + r_n.$$

The goal is to make the left side telescope.

§18.6 A factorial recurrence

Another example is

$$a_n = na_{n-1} + 1, \quad a_1 = 1.$$

Dividing by $n!$ gives

$$\frac{a_n}{n!} = \frac{a_{n-1}}{(n-1)!} + \frac{1}{n!}.$$

Thus if

$$b_n = \frac{a_n}{n!},$$

then

$$b_n = b_{n-1} + \frac{1}{n!}.$$

Therefore

$$a_n = n! \left(1 + \frac{1}{2!} + \frac{1}{3!} + \cdots + \frac{1}{n!} \right).$$

The point is not the final formula alone. The point is that multiplication by n suggests normalizing by $n!$.

§18.7 Groups

The second half of the note switches to group theory. A group is a set G with a binary operation $*$ satisfying:

- (i) associativity: $(a * b) * c = a * (b * c)$;
- (ii) identity: there is 1_G such that $1_G * g = g * 1_G = g$;
- (iii) inverses: for each $g \in G$, there is g^{-1} with $gg^{-1} = g^{-1}g = 1_G$.

Commutativity is not required. If it holds, the group is abelian.

Standard examples include:

$$(\mathbb{Z}, +), \quad \mathbb{Z}/n\mathbb{Z}, \quad S^1 = \{z \in \mathbb{C} : |z| = 1\}.$$

The note explicitly warns that closure is already included in the phrase “binary operation on G ”: if $*$: $G \times G \rightarrow G$, then $g * h \in G$ automatically.

§18.8 Symmetric and dihedral groups

The symmetric group S_n consists of all permutations of n objects. Composition is the group operation. The note writes permutations in cycle notation and emphasizes that the order of composition matters.

The dihedral group D_{2n} is the symmetry group of a regular n -gon. It contains rotations and reflections. It has presentation

$$D_{2n} = \langle r, s \mid r^n = s^2 = 1, \, srs = r^{-1} \rangle.$$

For $n = 5$, for example, r is a rotation by 72° , and s is a reflection. The relation $srs = r^{-1}$ says that reflecting reverses the direction of rotation.

§18.9 Products and isomorphisms

If G, H are groups, their product is

$$G \times H = \{(g, h) : g \in G, \, h \in H\},$$

with operation

$$(g_1, h_1)(g_2, h_2) = (g_1g_2, h_1h_2).$$

The inverse is

$$(g, h)^{-1} = (g^{-1}, h^{-1}).$$

Two groups are isomorphic if there is a bijection $\varphi : G \rightarrow H$ such that

$$\varphi(g_1g_2) = \varphi(g_1)\varphi(g_2).$$

An isomorphism means that the two groups have the same structure with different labels.

§18.10 Linear recurrences

A first-order linear recurrence has the form

$$a_{n+1} = ra_n + b.$$

If $re1$, look for a fixed point L satisfying $L = rL + b$. Then subtract it:

$$a_{n+1} - L = r(a_n - L).$$

So

$$a_n - L = r^n(a_0 - L).$$

This method explains why geometric progressions appear inside many non-homogeneous recurrences.

§18.11 Groups from repeated operations

The note's first group examples should be read as abstraction from repeated operations. Integers under addition, nonzero rationals under multiplication, rotations of a polygon, and permutations are all systems where operations can be composed and undone. The group axioms isolate exactly that pattern: associativity, identity, and inverses.

§18.12 Why the pattern matters

This note is valuable because it places recurrences and groups next to each other. A recurrence is about repeatedly applying an operation; a group is about operations that can be composed and undone. Both topics are early appearances of the same general idea: mathematics often studies not isolated objects, but transformations and their iterates.

§18.13 Recurrences as iteration of maps

A recurrence $a_{n+1} = F(a_n)$ is iteration of a map. Arithmetic progressions iterate a translation, and geometric progressions iterate a scaling. Their closed forms are orbit formulas:

$$a_n = a_1 + (n-1)d, \quad a_n = a_1 q^{n-1}.$$

§18.14 First-order linear recurrences

A recurrence

$$a_{n+1} = ra_n + b$$

has a fixed point $a_* = b/(1-r)$ when $r \neq 1$. Subtracting the fixed point gives

$$a_{n+1} - a_* = r(a_n - a_*).$$

This turns an affine recurrence into a geometric progression.

§18.15 Groups behind repeated operations

Groups formalize reversible repeated operations. Integers under addition record repeated translation; nonzero real numbers under multiplication record repeated scaling. The recurrence examples are therefore early examples of group actions on a set.

§18.16 Sequences as orbits of operations

Sequences become clearer as dynamics. Arithmetic and geometric progressions are translation and scaling orbits, affine recurrences become geometric after shifting to a fixed point, and groups record the reversible operations behind iteration.

Roots of Unity, Decagons, and Cyclotomic Divisibility

N-20220723

§19.1 The ant on a decagon

The first problem features an ant standing on vertex 1 of a regular decagon. It hops 1 space at a time for 10 hops, then 2 spaces at a time for 10 hops, and so on until it hops 10 spaces at a time for 10 hops.

Modulo 10, doing 10 hops of length k moves the ant by

$$10k \equiv 0 \pmod{10}.$$

Thus after each block of ten hops, the ant returns to the same vertex. Since it starts at vertex 1, it ends at vertex 1. This is the most elementary roots-of-unity idea: motion around a regular polygon is arithmetic modulo the number of vertices.

§19.2 Roots of unity

For a positive integer n , the n -th roots of unity are the complex solutions of

$$z^n = 1.$$

They are

$$U_n = \left\{ e^{2\pi i k/n} : k = 0, 1, \dots, n-1 \right\}.$$

Geometrically, these are the vertices of a regular n -gon on the unit circle.

If n is even, $z^n = 1$ has two real roots, 1 and -1 . If n is odd, the only real root is 1.

The example

$$\left(\frac{-1 + i\sqrt{3}}{2} \right)^3$$

is easiest in polar form. Since

$$\frac{-1 + i\sqrt{3}}{2} = \cos \frac{2\pi}{3} + i \sin \frac{2\pi}{3},$$

its cube is

$$\cos 2\pi + i \sin 2\pi = 1.$$

§19.3 The polynomial $x^{2n} + x^n + 1$

The note includes a Math StackExchange question about when

$$x^2 + x + 1$$

divides

$$x^{2n} + x^n + 1.$$

Let ω be a primitive third root of unity. Then

$$\omega^2 + \omega + 1 = 0.$$

The polynomial $x^2 + x + 1$ divides $x^{2n} + x^n + 1$ iff both ω and ω^2 are roots of the latter. Now

$$\omega^{2n} + \omega^n + 1 = 0$$

iff ω^n is a nontrivial third root of unity. This happens iff

$$n \not\equiv 0 \pmod{3}.$$

§19.4 A broader pattern

The natural generalization is

$$1 + x^n + x^{2n} + \cdots + x^{kn}$$

and the divisor

$$1 + x + x^2 + \cdots + x^k = \frac{x^{k+1} - 1}{x - 1}.$$

The roots of the divisor are the nontrivial $(k+1)$ -st roots of unity. If $\zeta^{k+1} = 1$ and $\zeta \neq 1$, then

$$1 + \zeta^n + \cdots + \zeta^{kn} = 0$$

provided $\zeta^n \neq 1$. If the map

$$\zeta \mapsto \zeta^n$$

permutes the nontrivial $(k+1)$ -st roots, which occurs when

$$\gcd(n, k+1) = 1,$$

then divisibility follows.

§19.5 Cyclotomic viewpoint

The polynomial $x^m - 1$ factors as

$$x^m - 1 = \prod_{d|m} \Phi_d(x),$$

where Φ_d is the d -th cyclotomic polynomial. The quotient

$$\frac{x^{k+1} - 1}{x - 1}$$

is the product of cyclotomic factors $\Phi_d(x)$ for $d \mid k+1$, $d > 1$. Divisibility questions of this type are therefore really questions about how exponentiation acts on roots of unity.

§19.6 Why this is more than a trick

The note's handwritten comment says that this is not merely about solving a problem. That is exactly right. Roots of unity turn algebra into geometry: powers become rotations, divisibility becomes vanishing on polygons, and congruences become angle arithmetic. This is the same philosophy behind Fourier analysis, cyclic groups, and cyclotomic fields.

§19.7 Roots of unity viewpoint

A regular decagon is naturally encoded by the tenth roots of unity:

$$\zeta^k = e^{2\pi i k/10}, \quad k = 0, 1, \dots, 9.$$

Moving by m vertices is multiplication by ζ^m . After ten such moves, the total displacement is multiplication by $(\zeta^m)^{10} = 1$. This explains why the ant returns after each block of ten equal jumps.

The advantage of this model is that a geometric hopping problem becomes modular arithmetic in $\mathbb{Z}/10\mathbb{Z}$.

§19.8 Cyclotomic divisibility

The same roots of unity explain polynomial divisibility. To show that a polynomial $Q(x)$ divides $P(x)$, one may show that every root of Q is also a root of P , provided Q has no repeated roots and the field is appropriate.

For example, divisibility by $x^2 + x + 1$ can be checked at primitive third roots of unity. This is often cleaner than expanding and factoring by force.

Arithmetic and Geometric Sequence Exercises

N-20220803

§20.1 A trigonometric opening example

The first exercise concerns a triangle ABC such that

$$\lg(\sin A), \quad \lg(\sin B), \quad \lg(\sin C)$$

form an arithmetic progression, and A, B, C themselves also form an arithmetic progression. Since

$$A + B + C = \pi,$$

if A, B, C are in arithmetic progression, then

$$B = \frac{\pi}{3}.$$

The logarithmic condition means

$$2 \lg(\sin B) = \lg(\sin A) + \lg(\sin C),$$

so

$$\sin^2 B = \sin A \sin C.$$

Because $\sin B = \sqrt{3}/2$, we get

$$\sin A \sin C = \frac{3}{4}.$$

Using $A + C = 2\pi/3$, one can analyze

$$\cos(A - C) = \cos A \cos C + \sin A \sin C.$$

The computation in the page leads to the conclusion that the triangle is equilateral. The underlying idea is that two independent arithmetic-progressions conditions force symmetry.

§20.2 Logarithmic sequences

If

$$a_n = 10 + \lg 2^n,$$

then

$$a_n = 10 + n \lg 2.$$

Hence

$$a_{n+1} - a_n = \lg 2,$$

so $\{a_n\}$ is an arithmetic progression.

This is a typical use of logarithms: they turn powers into multiples. A geometric pattern in 2^n becomes an arithmetic pattern after applying $\lg$.

§20.3 Recovering terms from partial sums

If S_n is the sum of the first n terms of a sequence, then

$$a_n = S_n - S_{n-1}.$$

The note uses this in problems involving two arithmetic progressions $\{a_n\}$, $\{b_n\}$ whose sums S_n, T_n satisfy

$$\frac{S_n}{T_n} = \frac{2n}{3n+1}.$$

For arithmetic progressions,

$$S_n = \frac{n}{2}(a_1 + a_n),$$

but the cleaner trick is to use

$$a_n = S_n - S_{n-1}, \quad b_n = T_n - T_{n-1}.$$

The page obtains

$$\frac{a_n}{b_n} = \frac{4n-2}{6n-2} = \frac{2n-1}{3n-1}.$$

§20.4 Sums and recurrence relations

A common type of problem gives a relation between S_n and a_n . For example,

$$S_n = 2a_n - 2.$$

Then

$$a_{n+1} = S_{n+1} - S_n = 2a_{n+1} - 2 - (2a_n - 2),$$

so

$$a_{n+1} = 2a_n.$$

Thus the sequence is geometric. This is a useful principle: whenever a problem relates a term to a partial sum, subtract consecutive equations.

§20.5 Geometric progression products

If $\{a_n\}$ is a geometric progression, then

$$a_m^2 = a_{m-r}a_{m+r}.$$

This appears in several exercises. It follows from

$$a_n = a_1 q^{n-1}.$$

For example,

$$a_6^2 = a_5 a_7.$$

The note uses such identities to compute missing terms and ratios from partial information.

§20.6 A harmonic-looking sequence

One page studies

$$a_n = \frac{1}{1+2} + \frac{1}{2+3} + \cdots + \frac{1}{n+n+1}$$

or similar telescoping sums. The key transformation is

$$\frac{1}{k(k+1)} = \frac{1}{k} - \frac{1}{k+1}.$$

Telescoping is the additive analogue of cancellation in products. When the denominator factors into consecutive terms, always look for a difference of reciprocals.

§20.7 A final synthesis

This note is mostly exercise-driven, but the underlying methods are stable: logarithms convert multiplication into addition, partial sums are inverted by differences, arithmetic progressions have linear terms, geometric progressions have exponential terms, and telescoping makes hidden cancellation explicit. These are exactly the same patterns that later become finite-difference calculus.

§20.8 Sequences as discrete dynamical systems

An arithmetic progression and a geometric progression are both generated by repeatedly applying a simple transformation. For an arithmetic progression,

$$a_{n+1} = a_n + d,$$

so the update is translation. For a geometric progression,

$$a_{n+1} = qa_n,$$

so the update is scaling. This explains why logarithms turn geometric progressions into arithmetic ones:

$$\log a_{n+1} = \log q + \log a_n.$$

Multiplicative iteration becomes additive iteration after applying $\log$.

This is the same structural idea behind linear recurrences. A sequence is not only a list of numbers; it is often an orbit under an update rule.

§20.9 Partial sums and finite differences

If $S_n = a_1 + \cdots + a_n$, then

$$a_n = S_n - S_{n-1} = \Delta S_{n-1}.$$

So recovering terms from partial sums is finite differentiation. Conversely,

$$S_n = S_0 + \sum_{k=1}^n a_k$$

is finite integration.

For an arithmetic progression $a_n = a_1 + (n - 1)d$, the first difference is constant:

$$\Delta a_n = d.$$

For a quadratic sequence, the second difference is constant. This is why finite-difference tables detect polynomial sequences, just as derivatives detect polynomial degree in calculus.

§20.10 Generating functions for sequence exercises

A sequence (a_n) can be packaged into a generating function

$$A(x) = \sum_{n \geq 0} a_n x^n.$$

Arithmetic and geometric progressions have especially simple forms. For example,

$$\sum_{n \geq 0} q^n x^n = \frac{1}{1 - qx}.$$

A linear sequence produces a rational generating function with denominator $(1 - x)^2$. Recurrences translate into algebraic equations for generating functions.

This turns many sequence problems into algebra: instead of manipulating terms one by one, encode the whole sequence and operate on the generating function.

§20.11 Telescoping as a coboundary

A telescoping sum has the form

$$a_k = F(k + 1) - F(k).$$

Then

$$\sum_{k=m}^n a_k = F(n + 1) - F(m).$$

Only the boundary terms survive. In this language, a telescoping term is a discrete derivative or coboundary.

The identity

$$\frac{1}{k(k+1)} = \frac{1}{k} - \frac{1}{k+1}$$

is exactly this pattern with $F(k) = 1/k$. The cancellation is not accidental; it is the discrete fundamental theorem of calculus.

§20.12 Second-order recurrences and characteristic roots

Some sequence problems go beyond arithmetic and geometric progressions. A linear recurrence such as

$$a_{n+2} = pa_{n+1} + qa_n$$

is solved by trying $a_n = \lambda^n$. This gives the characteristic equation

$$\lambda^2 = p\lambda + q.$$

When the roots are distinct, the solution has the form

$$a_n = A\lambda_1^n + B\lambda_2^n.$$

Thus geometric progressions are the eigenvectors of the shift recurrence. This is the bridge from elementary sequences to linear algebra and dynamical systems.

§20.13 Discrete calculus of elementary sequences

The exercises are unified by discrete calculus. Logarithms linearize multiplication, finite differences invert partial sums, telescoping records boundary cancellation, generating functions encode whole sequences, and linear recurrences are controlled by characteristic roots. The elementary tricks are the visible form of a general principle: choose the representation in which the sequence's update rule becomes simple.

Divisibility, Congruences, and the Habit of Proving Necessity

N-20220913

§21.1 From guessing to proving

The note opens with a useful warning: in number theory, guessing examples is not enough. One must separate necessity and sufficiency. If a condition is necessary, it must follow from the assumption. If it is sufficient, it must imply the desired conclusion. Many false solutions prove only one direction.

The first example asks for a prime p of the form

$$p = x(x + 1) - 4, \quad x \in \mathbb{N}_+.$$

Since $x(x + 1)$ is always even,

$$p = x(x + 1) - 4$$

is even. If p is prime, then the only even prime is 2. Hence

$$p = 2, \quad x(x + 1) = 6,$$

so

$$x = 2.$$

This is a clean example of using parity as a necessary condition before solving.

§21.2 Divisibility

For integers a, b , with $a \neq 0$, we write

$$a \mid b$$

if there exists $q \in \mathbb{Z}$ such that

$$b = aq.$$

The page lists standard properties:

- (i) $a \mid b \iff -a \mid b \iff a \mid -b \iff |a| \mid |b|$;
- (ii) if $a \mid b$ and $b \mid c$, then $a \mid c$;
- (iii) if $a \mid b$ and $a \mid c$, then $a \mid bx + cy$ for all integers x, y ;
- (iv) if $m \neq 0$, then $a \mid b \iff ma \mid mb$;
- (v) if $a \mid b$ and $b \mid a$, then $b = \pm a$;
- (vi) if $b \neq 0$ and $a \mid b$, then $|a| \leq |b|$.

The handwritten note points out that property (iii) is often the most useful. It lets us combine divisibility relations linearly.

§21.3 Example: reducing the dividend

Suppose

$$x^2 + 1 \mid x^3 - x, \quad x \in \mathbb{Z}_+.$$

Modulo $x^2 + 1$, we have

$$x^2 \equiv -1,$$

so

$$x^3 - x = x(x^2) - x \equiv -x - x = -2x.$$

Thus

$$x^2 + 1 \mid 2x.$$

For $x > 1$,

$$x^2 + 1 > 2x,$$

so divisibility is impossible unless $2x = 0$, which cannot happen for positive x . Checking $x = 1$, we get

$$1^3 - 1 = 0,$$

so $x = 1$ works.

This example illustrates the handwritten strategy: reduce the large expression by using the divisor itself, then use size estimates to force a small list of cases.

§21.4 Example: polynomial divisibility in one variable

For divisibility such as

$$x^2 + x + 1 \mid x^4 - 2,$$

one uses

$$x^2 + x + 1 = 0 \implies x^3 - 1 = (x - 1)(x^2 + x + 1) \equiv 0.$$

Thus $x^3 \equiv 1$, and

$$x^4 - 2 \equiv x - 2.$$

So the divisor must divide $x - 2$. This leaves only very small candidates, because for large $|x|$, the quadratic divisor is too large in absolute value to divide the linear remainder.

The general method is Euclidean division in disguise: replace the dividend by its remainder modulo the divisor.

§21.5 Congruences

Congruence notation compresses divisibility:

$$a \equiv b \pmod{m}$$

means

$$m \mid a - b.$$

This turns many divisibility problems into arithmetic with remainders. The note emphasizes that congruences can be added, subtracted, and multiplied:

$$a \equiv b \pmod{m}, \quad c \equiv d \pmod{m} \implies a + c \equiv b + d, \quad ac \equiv bd \pmod{m}.$$

§21.6 A large power sum

The last problem asks to prove a divisibility statement involving

$$1^{1983} + 2^{1983} + \cdots + 1983^{1983}.$$

The intended thought is pairing and congruence. When the exponent is odd,

$$(m - k)^{1983} \equiv -k^{1983} \pmod{m}.$$

So terms can be paired to cancel modulo a suitable modulus. To prove divisibility by

$$1 + 2 + \cdots + 1983 = \frac{1983 \cdot 1984}{2},$$

one decomposes the modulus into factors and proves divisibility factor by factor, using pairing and residue classes.

The handwritten solution is valuable because it records a strategy rather than just a theorem: for large power sums, first look for symmetry $k \leftrightarrow m - k$, then reduce to congruences.

§21.7 A wider interpretation

Divisibility is the first place where algebraic thinking becomes economical. Instead of expanding huge expressions, we work modulo the divisor and keep only remainders. This is exactly the same principle behind quotient rings such as

$$\mathbb{Z}/m\mathbb{Z}$$

and polynomial quotients

$$F[x]/(f).$$

The note's small examples are therefore early versions of modular algebra.

§21.8 Divisibility as ideal containment

In $\mathbb{Z}$, the statement $a \mid b$ is equivalent to

$$b \in (a),$$

where $(a) = a\mathbb{Z}$ is the ideal generated by a . Divisibility reverses inclusion of principal ideals:

$$a \mid b \iff (b) \subseteq (a).$$

This explains why linear combinations are central. If $a \mid b$ and $a \mid c$, then every integer linear combination $bx + cy$ lies in (a) .

§21.9 GCD and LCM as ideal operations

For integers a, b , Bezout's theorem says

$$(a) + (b) = (\gcd(a, b)).$$

The ideal generated by all linear combinations of a and b is generated by their greatest common divisor. Similarly,

$$(a) \cap (b) = (\text{lcm}(a, b)).$$

Thus gcd and lcm are not just computational devices; they are sum and intersection of ideals in $\mathbb{Z}$.

§21.10 Congruences as quotient rings

The congruence

$$a \equiv b \pmod{n}$$

means that a and b define the same element of the quotient ring

$$\mathbb{Z}/n\mathbb{Z}.$$

The projection

$$\pi : \mathbb{Z} \rightarrow \mathbb{Z}/n\mathbb{Z}$$

forgets the exact integer and remembers only its residue class. Its kernel is $n\mathbb{Z}$, so this is the first isomorphism theorem in its simplest form:

$$\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/\ker \pi.$$

Modular arithmetic is therefore quotient algebra, not merely a remainder shortcut.

§21.11 Units and Bezout inverses

An integer a has a multiplicative inverse modulo n exactly when

$$\gcd(a, n) = 1.$$

By Bezout, there exist integers x, y such that

$$ax + ny = 1.$$

Reducing modulo n gives

$$ax \equiv 1 \pmod{n},$$

so x is the inverse of a modulo n . This is the concrete connection between gcd computations and units in $\mathbb{Z}/n\mathbb{Z}$.

§21.12 Valuations as prime coordinates

For nonzero integers, $v_p(n)$ denotes the exponent of the prime p in n . Divisibility becomes coordinatewise comparison:

$$a \mid b \iff v_p(a) \leq v_p(b) \text{ for all primes } p.$$

Gcd and lcm become coordinatewise minimum and maximum:

$$v_p(\gcd(a, b)) = \min(v_p(a), v_p(b)), \quad v_p(\text{lcm}(a, b)) = \max(v_p(a), v_p(b)).$$

This turns divisibility into geometry on the exponent lattice of prime factorizations.

§21.13 Chinese remainder theorem as local-to-global

Theorem 21.13.1 (Chinese remainder theorem)

If $m_1, \dots, m_r$ are pairwise coprime positive integers, then

$$\mathbb{Z}/(m_1 \cdots m_r)\mathbb{Z} \cong \prod_{i=1}^r \mathbb{Z}/m_i\mathbb{Z}.$$

If

$$n = \prod_i p_i^{e_i},$$

then the Chinese remainder theorem gives

$$\mathbb{Z}/n\mathbb{Z} \cong \prod_i \mathbb{Z}/p_i^{e_i}\mathbb{Z}.$$

So a congruence problem modulo n can often be solved prime-power by prime-power and then recombined.

This is a first local-to-global principle: understand arithmetic at each prime separately, then assemble the global answer.

§21.14 Polynomial divisibility as the same story

Over a field k , the ring $k[x]$ has Euclidean division:

$$f = qg + r, \quad \deg r < \deg g.$$

Therefore the methods of integer divisibility have polynomial analogues. In particular,

$$x - a \mid f(x) \iff f(a) = 0.$$

Working modulo a polynomial $g(x)$ means working in the quotient ring $k[x]/(g)$. The earlier examples that reduce $x^4 - 2$ modulo $x^2 + x + 1$ are exactly quotient-ring computations in elementary disguise.

§21.15 Divisibility as quotient algebra

Divisibility problems become clearer when the divisor is treated as a structure rather than an obstacle. Ideals organize divisibility, quotient rings organize congruence, valuations localize divisibility at primes, and Euclidean division transfers the same ideas to polynomials. The elementary habit "reduce by the divisor" is the first step toward quotient algebra.

Divisibility by Finite Differences and Pairing

N-20220918

§22.1 A general odd-power divisibility principle

The page begins with a useful theorem: for odd r , the sum

$$1^r + 2^r + \cdots + k^r$$

has a strong divisibility property by the triangular number

$$1 + 2 + \cdots + k = \frac{k(k+1)}{2}.$$

The natural proof idea is pairing. Pair j with $k+1-j$. Since r is odd,

$$j^r + (k+1-j)^r$$

is divisible by $k+1$, because

$$k+1-j \equiv -j \pmod{k+1}.$$

A second argument handles the factor k or the remaining factor according to parity. The handwritten note records the right high-level method in blue: pairing, oddness, and treating the whole expression rather than each term separately.

§22.2 A divisibility problem with 512

The main worked example is:

$$\text{prove that } 512 \mid 3^{2n} - 32n^2 + 24n - 1, \quad n \in \mathbb{N}_+.$$

Set

$$f(n) = 3^{2n} - 32n^2 + 24n - 1.$$

The note uses finite differences. Since $512 \mid f(1)$, it is enough to show

$$512 \mid f(n+1) - f(n)$$

for all n . Compute

$$\begin{aligned} f(n+1) - f(n) &= 3^{2n+2} - 32(n+1)^2 + 24(n+1) - 1 \\ &\quad - 3^{2n} + 32n^2 - 24n + 1 \\ &= 8 \cdot 3^{2n} - 64n - 8 \\ &= 8(3^{2n} - 8n - 1). \end{aligned}$$

Thus it remains to prove

$$64 \mid 3^{2n} - 8n - 1.$$

Define

$$h(n) = 3^{2n} - 8n - 1.$$

Then

$$\begin{aligned} h(n+1) - h(n) &= 3^{2n+2} - 8(n+1) - 1 - 3^{2n} + 8n + 1 \\ &= 8 \cdot 3^{2n} - 8 \\ &= 8(3^{2n} - 1). \end{aligned}$$

So it is enough to know

$$8 \mid 3^{2n} - 1.$$

But $3^2 = 9 \equiv 1 \pmod{8}$, so

$$3^{2n} \equiv 1 \pmod{8}.$$

Therefore $64 \mid h(n)$, and hence $512 \mid f(n)$.

§22.3 Why finite differences work

The solution is essentially induction, but written in a more structural way. If $m \mid f(1)$ and $m \mid f(n+1) - f(n)$ for every n , then

$$f(n) = f(1) + \sum_{j=1}^{n-1} (f(j+1) - f(j))$$

is divisible by m . This is the discrete analogue of recovering a function from its derivative.

The page is valuable because it shows a common trick for expressions involving n both in an exponent and a polynomial. Direct induction may be messy, but finite differences reduce the polynomial part and often reveal a smaller congruence.

§22.4 An alternating harmonic numerator

The second problem is:

$$\frac{p}{q} = 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \cdots - \frac{1}{1318} + \frac{1}{1319},$$

where p, q are coprime positive integers, and one must prove

$$1979 \mid p.$$

The note rewrites the alternating sum by separating odd and even denominators:

$$1 + \frac{1}{3} + \cdots + \frac{1}{1319} - 2 \left(\frac{1}{2} + \frac{1}{4} + \cdots + \frac{1}{1318} \right).$$

After simplification, this becomes a sum of paired fractions in which a factor 1979 appears in the numerator. The important observation is that 1979 is prime and is coprime to the denominators appearing after cancellation, so if the rational number is written in lowest terms, the factor 1979 cannot disappear into the denominator. Hence it must divide p .

§22.5 The next abstraction

Both problems use the same philosophy: do not expand everything. For power sums, pair symmetric terms. For congruence sequences, take finite differences. For rational sums, reorganize the denominator structure so that a prime factor becomes visible. These are all examples of choosing the right invariant before calculating.

§22.6 Finite differences as discrete derivatives

The finite difference operator

$$\Delta f(n) = f(n+1) - f(n)$$

is the discrete analogue of differentiation. If $m \mid f(1)$ and $m \mid \Delta f(n)$ for every n , then

$$f(n) = f(1) + \sum_{j=1}^{n-1} \Delta f(j)$$

is divisible by m . This is induction written as a discrete fundamental theorem of calculus.

For a polynomial of degree d ,

$$\Delta^{d+1} f = 0.$$

So finite differences lower polynomial complexity in exactly the way derivatives do.

§22.7 The binomial basis and integer-valued polynomials

The binomial polynomials

$$\binom{x}{0}, \quad \binom{x}{1}, \quad \binom{x}{2}, \quad \dots$$

are adapted to finite differences:

$$\Delta \binom{x}{k} = \binom{x}{k-1}.$$

A classical theorem says that every integer-valued polynomial has a unique expansion

$$f(x) = \sum_k c_k \binom{x}{k}, \quad c_k \in \mathbb{Z}.$$

This basis is often better than the monomial basis $1, x, x^2, \dots$ when proving divisibility for all integers.

The reason is simple: $\binom{n}{k}$ is automatically integral for integer n , and finite differences act on it cleanly.

§22.8 Pairing as group symmetry

In odd-power sums, the pairing

$$j \longleftrightarrow m - j$$

is an action of a two-element symmetry group on residue classes modulo m . If r is odd, then

$$(m - j)^r \equiv -j^r \pmod{m},$$

so each orbit contributes zero modulo m , except possibly fixed points.

Thus the cancellation is not a lucky trick. It is orbit cancellation under the involution $j \mapsto m - j$. Similar symmetry arguments appear in character sums and finite group actions.

§22.9 P-adic valuation and lifting intuition

For a prime p , $v_p(N)$ records the exponent of p dividing N . Divisibility by p^k becomes the inequality

$$v_p(N) \geq k.$$

In problems such as proving divisibility by $512 = 2^9$, the real task is to track how many factors of 2 are gained at each finite-difference step.

The deeper principle is lifting. A congruence modulo p may sometimes be lifted to a congruence modulo higher powers p^k . Hensel lifting is the systematic version of this idea. Even when the theorem is not invoked, the intuition is valuable: prime-power divisibility is controlled locally and incrementally.

§22.10 Alternating harmonic sums and denominator control

For rational sums, the numerator after reduction depends on which prime factors can survive cancellation with the denominator. If a prime p appears in the numerator of a common-denominator expression and is coprime to the denominator after simplification, then p must remain in the reduced numerator.

This is a valuation statement. Writing a rational number as A/B , one compares $v_p(A)$ and $v_p(B)$. Reduction cancels the smaller valuation from both sides. If $v_p(B) = 0$ and $v_p(A) > 0$, then the reduced numerator is still divisible by p .

§22.11 Discrete summation by parts

There is also a discrete analogue of integration by parts. For sequences a_n, b_n , one form is

$$\sum_{n=m}^{N-1} a_n \Delta b_n = a_N b_N - a_m b_m - \sum_{n=m}^{N-1} b_{n+1} \Delta a_n.$$

This technique turns products inside sums into boundary terms plus a simpler sum. It is the natural extension of telescoping and finite differences.

§22.12 Discrete calculus and local divisibility

The note's divisibility tricks are forms of discrete calculus and local arithmetic. Finite differences lower complexity, telescoping keeps boundary terms, pairing exploits symmetry, valuations measure prime-power divisibility, and rational-sum arguments track which primes can cancel. The correct invariant is usually not the whole expression, but its difference, orbit pairing, or valuation.

Shoelace Formula, Remainder Theorem, and Modular Arithmetic

N-20221013

§23.1 Shoelace formula

For a polygon with vertices

$$A_1(x_1, y_1), A_2(x_2, y_2), \dots, A_n(x_n, y_n),$$

listed cyclically, its area is

$$S = \frac{1}{2} \left| \sum_{i=1}^n (x_i y_{i+1} - x_{i+1} y_i) \right|,$$

where

$$x_{n+1} = x_1, \quad y_{n+1} = y_1.$$

This is the shoelace formula.

For the example

$$A_1 = (3, 4), \quad A_2 = (5, 11), \quad A_3 = (12, 8), \quad A_4 = (9, 5), \quad A_5 = (5, 6),$$

the note computes

$$\begin{aligned} 2S &= |3 \cdot 11 - 5 \cdot 4 + 5 \cdot 8 - 12 \cdot 11 + 12 \cdot 5 - 8 \cdot 9 \\ &\quad + 9 \cdot 6 - 5 \cdot 5 + 5 \cdot 4 - 3 \cdot 6| \\ &= 60. \end{aligned}$$

Thus

$$S = 30.$$

§23.2 Why the shoelace formula works

Each term

$$x_i y_{i+1} - x_{i+1} y_i$$

is a 2×2 determinant. Geometrically, it is the signed area of the parallelogram spanned by the two position vectors. The shoelace formula adds signed triangle areas around the origin. Interior cancellations leave the polygon area.

This determinant viewpoint explains why orientation matters. Reversing the order of vertices changes the sign, and the absolute value recovers ordinary area.

§23.3 A circle-area maximum problem

One exercise asks: if three points lie on

$$x^2 + y^2 = 2x - 4y,$$

what is the maximum of

$$x_1y_2 - x_2y_1 + x_2y_3 - x_3y_2?$$

First complete the square:

$$(x - 1)^2 + (y + 2)^2 = 5.$$

So the points lie on a circle of radius $\sqrt{5}$. Expressions such as

$$x_iy_j - x_jy_i$$

are twice signed triangle areas with the origin. The handwritten solution interprets the expression as twice the area of a triangle determined by the points and then maximizes the possible area.

The largest triangle area in a circle occurs for a suitably positioned chord/diameter configuration; the note obtains the maximum

$$20.$$

The main idea is not the number alone: determinants convert coordinate expressions into signed areas.

§23.4 Remainder theorem

If a polynomial $p(x)$ is divided by $x - c$, then the remainder is

$$p(c).$$

Indeed, polynomial division gives

$$p(x) = (x - c)q(x) + r,$$

where r is constant. Substituting $x = c$ yields

$$r = p(c).$$

The note includes this proof because it is often more useful than the statement. It explains why evaluating at roots kills divisible parts.

§23.5 A polynomial remainder example

Suppose $P(x)$ leaves remainder 99 when divided by $x - 19$, and remainder 19 when divided by $x - 99$. If $P(x)$ is divided by

$$(x - 19)(x - 99),$$

the remainder must be linear:

$$R(x) = ax + b.$$

The conditions are

$$R(19) = 99, \quad R(99) = 19.$$

Thus

$$19a + b = 99, \quad 99a + b = 19.$$

Solving gives

$$a = -1, \quad b = 118,$$

so

$$R(x) = 118 - x.$$

This example is the remainder theorem plus interpolation.

§23.6 Congruence arithmetic

The note then switches to congruences. We write

$$a \equiv b \pmod{m}$$

if

$$m \mid a - b.$$

Useful properties include:

- (i) reflexivity: $a \equiv a \pmod{m}$;
- (ii) symmetry and transitivity;
- (iii) compatibility with addition and multiplication;
- (iv) if $a \equiv b \pmod{m}$, then $a^k \equiv b^k \pmod{m}$.

The page also records elementary digit facts:

- (i) $n^2 \pmod{3}$ is 0 or 1;
- (ii) $n^2 \pmod{8}$ is 0, 1, 4;
- (iii) $n \equiv \text{sum of digits of } n \pmod{9}$.

These are quick residue filters for contest problems.

§23.7 Fermat and Euler

Fermat's little theorem says: if p is prime and $p \nmid a$, then

$$a^{p-1} \equiv 1 \pmod{p}.$$

Euler's theorem generalizes it: if

$$\gcd(a, m) = 1,$$

then

$$a^{\varphi(m)} \equiv 1 \pmod{m}.$$

Here $\varphi(m)$ is Euler's totient function: the number of positive integers up to m that are coprime to m .

The note's example asks for the remainder of

$$2^{1000}$$

modulo 13. Since

$$\varphi(13) = 12,$$

we reduce

$$1000 \equiv 4 \pmod{12}.$$

Therefore

$$2^{1000} \equiv 2^4 = 16 \equiv 3 \pmod{13}.$$

§23.8 Shoelace formula as oriented area

For polygon vertices (x_i, y_i) , the shoelace formula is

$$A = \frac{1}{2} \left| \sum_i (x_i y_{i+1} - x_{i+1} y_i) \right|.$$

The summand

$$x_i y_{i+1} - x_{i+1} y_i$$

is the two-dimensional determinant of the consecutive vertex vectors. Thus the formula computes oriented area by summing triangular pieces based at the origin. The absolute value at the end forgets orientation and returns ordinary area.

§23.9 Evaluation as quotient by $x-a$

The polynomial remainder theorem says that the remainder of $f(x)$ on division by $x-a$ is $f(a)$. Thus $x-a$ divides $f(x)$ iff $f(a) = 0$. This is the simplest root-factor principle and is the foundation for many later factorization arguments.

§23.10 Where the calculation points

This note connects three determinant-like ideas. The shoelace formula uses 2×2 determinants to measure area. The remainder theorem uses evaluation to collapse polynomial division. Congruences use remainders to collapse integer arithmetic. In all three cases, the strategy is to replace a complicated object by an invariant that preserves exactly the information needed.

§23.11 Shoelace as a discrete line integral

The shoelace formula

$$\text{Area} = \frac{1}{2} \left| \sum_i x_i y_{i+1} - x_{i+1} y_i \right|$$

can be read as a discrete version of Green's theorem:

$$\text{Area}(D) = \frac{1}{2} \int_{\partial D} x \, dy - y \, dx.$$

The polygon formula samples this boundary integral along straight edges.

§23.12 Remainder theorem and quotient rings

The remainder theorem says

$$f(x) \equiv f(a) \pmod{x - a}.$$

This is evaluation as a quotient map:

$$k[x] \rightarrow k[x]/(x - a) \cong k.$$

So polynomial remainders are quotient algebra in elementary form.

§23.13 Fermat–Euler and unit groups

Euler’s theorem

$$a^{\varphi(n)} \equiv 1 \pmod{n}$$

for $\gcd(a, n) = 1$ says that the class of a lies in the finite unit group $(\mathbb{Z}/n\mathbb{Z})^\times$. Fermat’s little theorem is the prime case.

§23.14 Quotients in area, polynomials, and congruences

The note links three quotient ideas: polygon area as boundary integral, polynomial remainders as quotient rings, and modular exponentiation as finite unit-group arithmetic.

Integer Arithmetic: Well-Ordering, Euclidean Algorithm, and the Euler Function

N-20221030

§24.1 The well-ordering principle

The note begins with a concise list of basic facts about integers. The first is the well-ordering principle:

Theorem 24.1.1 (Well-ordering of $\mathbb{Z}_+$)

Every nonempty subset $A \subseteq \mathbb{Z}_+$ has a least element.

This statement looks elementary, but it is the engine behind many proofs in number theory. It is the reason the Euclidean algorithm terminates, the reason minimal-counterexample arguments work, and one of the standard formulations equivalent to induction.

§24.2 Divisibility, gcd, and lcm

For integers a, b , with $a \neq 0$, we say

$$a \mid b$$

if there exists $c \in \mathbb{Z}$ such that

$$b = ac.$$

The greatest common divisor of nonzero integers a, b is the positive integer d such that

$$d \mid a, \quad d \mid b,$$

and if $e \mid a$ and $e \mid b$, then $e \mid d$. We write

$$d = (a, b).$$

If $(a, b) = 1$, then a and b are relatively prime.

Similarly, the least common multiple l is the positive integer such that

$$a \mid l, \quad b \mid l,$$

and every common multiple of a, b is a multiple of l . For positive a, b ,

$$(a, b)[a, b] = ab.$$

The relation is best understood prime-by-prime: gcd takes the smaller exponent, lcm takes the larger exponent.

§24.3 Division algorithm

The division algorithm says that for $a, b \in \mathbb{Z}$, $b \neq 0$, there exist unique integers q, r such that

$$a = qb + r, \quad 0 \leq r < |b|.$$

Here q is the quotient and r is the remainder. This is not merely long division. It is the formal foundation of modular arithmetic and the Euclidean algorithm.

§24.4 Euclidean algorithm

Starting with $a, b \neq 0$, repeatedly divide:

$$\begin{aligned} a &= q_0b + r_0, \\ b &= q_1r_0 + r_1, \\ r_0 &= q_2r_1 + r_2, \\ &\vdots \\ r_{n-2} &= q_nr_{n-1} + r_n, \\ r_{n-1} &= q_{n+1}r_n. \end{aligned}$$

The last nonzero remainder r_n is the gcd of a, b .

The reason is the invariant

$$(a, b) = (b, r_0) = (r_0, r_1) = \cdots = (r_{n-1}, r_n) = r_n.$$

At each step, replacing (a, b) by $(b, a - qb)$ does not change the set of common divisors. This is the real idea; the numerical algorithm is just repeated use of this invariant.

§24.5 Example

The note gives a Euclidean-algorithm example such as

$$57970 \quad \text{and} \quad 10353.$$

The computation proceeds by repeated division until the remainder becomes zero. The last nonzero remainder is the gcd. The same backward substitutions produce integers x, y such that

$$ax + by = (a, b).$$

This is Bézout's identity.

§24.6 Fundamental theorem of arithmetic

Theorem 24.6.1 (Fundamental theorem of arithmetic)

Every integer $n > 1$ can be written as a product of primes, and this factorization is unique up to the order of the factors.

Every integer $n > 1$ factors uniquely into primes:

$$n = p_1^{\alpha_1} p_2^{\alpha_2} \cdots p_k^{\alpha_k}.$$

The note presents this as an essential property of integers. Uniqueness is what makes gcd, lcm, and the Euler function computable by exponents.

For example, if

$$a = \prod p^{\alpha_p}, \quad b = \prod p^{\beta_p},$$

then

$$(a, b) = \prod p^{\min(\alpha_p, \beta_p)}, \quad [a, b] = \prod p^{\max(\alpha_p, \beta_p)}.$$

§24.7 The Euler phi function

The Euler function $\varphi(n)$ counts the positive integers not exceeding n that are coprime to n . If

$$n = p_1^{\alpha_1} \cdots p_k^{\alpha_k},$$

then

$$\varphi(n) = n \left(1 - \frac{1}{p_1}\right) \cdots \left(1 - \frac{1}{p_k}\right).$$

This follows by inclusion–exclusion: remove multiples of each prime divisor of n .

For a prime power,

$$\varphi(p^a) = p^a - p^{a-1}.$$

For a prime p ,

$$\varphi(p) = p - 1.$$

§24.8 The broader lesson

This note is a compact bridge from elementary arithmetic to algebra. The Euclidean algorithm is not just a computation; it is an algorithmic proof that $\mathbb{Z}$ is a Euclidean domain. Bézout identity says ideals in $\mathbb{Z}$ are principal. Unique factorization says primes control divisibility. The Euler function prepares the way for the multiplicative group $(\mathbb{Z}/n\mathbb{Z})^\times$, where number theory becomes group theory.

§24.9 Well-ordering and Euclidean domains

The well-ordering principle makes the division algorithm possible in $\mathbb{Z}$. The Euclidean algorithm works because each step strictly decreases a nonnegative remainder. This pattern generalizes to Euclidean domains.

§24.10 Bezout and principal ideals

The gcd of a and b is the positive generator of the ideal

$$(a) + (b) = \{ax + by : x, y \in \mathbb{Z}\}.$$

Bezout's identity says the gcd is an integer linear combination of a and b .

§24.11 Euler function and unit groups

Euler's function counts units modulo n :

$$\varphi(n) = |(\mathbb{Z}/n\mathbb{Z})^\times|.$$

Euler's theorem is Lagrange's theorem in this finite group:

$$a^{\varphi(n)} \equiv 1 \pmod{n}$$

whenever $\gcd(a, n) = 1$.

§24.12 Euclidean arithmetic as ideals and unit groups

Integer arithmetic is organized by order, ideals, and groups. Well-ordering drives division, Euclidean algorithms compute ideal generators, unique factorization gives valuations, and Euler's theorem is group theory modulo n .

§24.13 Euclidean domains

The Euclidean algorithm works in $\mathbb{Z}$ because there is a division process with decreasing remainder. A Euclidean domain is an integral domain R with a size function N such that for $a, b \in R$, $b \neq 0$, one can write

$$a = qb + r, \quad r = 0 \text{ or } N(r) < N(b).$$

The integers and polynomial rings $k[x]$ over a field are Euclidean domains.

In every Euclidean domain, the Euclidean algorithm produces greatest common divisors and Bezout identities. Thus the elementary gcd algorithm is one example of a general algebraic mechanism.

§24.14 PID and Bezout domains

A principal ideal domain is an integral domain in which every ideal is generated by one element. In a PID, gcds exist because

$$(a) + (b) = (d)$$

for some d , and that d is a greatest common divisor of a and b .

Bezout's identity

$$ax + by = d$$

means precisely that $d \in (a) + (b)$ and generates the sum ideal. This reframes gcd as an ideal-theoretic object.

§24.15 Euler phi as unit counting

Euler's function counts units in a quotient ring:

$$\varphi(n) = |(\mathbb{Z}/n\mathbb{Z})^\times|.$$

If

$$n = \prod_i p_i^{e_i},$$

then the Chinese remainder theorem gives

$$(\mathbb{Z}/n\mathbb{Z})^\times \cong \prod_i (\mathbb{Z}/p_i^{e_i}\mathbb{Z})^\times.$$

Therefore

$$\varphi(n) = n \prod_{p|n} \left(1 - \frac{1}{p}\right).$$

The familiar formula is a group-size computation.

§24.16 Multiplicative functions

A function $f(n)$ is multiplicative if $f(mn) = f(m)f(n)$ whenever $\gcd(m, n) = 1$. Euler's phi, the divisor-counting function $\tau(n)$, and the sum-of-divisors function $\sigma(n)$ are multiplicative.

The reason is again Chinese remaindering or prime factorization: arithmetic data often decomposes independently over primes. This is the elementary beginning of Euler products and analytic number theory.

§24.17 Euclidean algorithm and ideals

The Euclidean algorithm does more than compute a gcd. It constructs coefficients x, y such that

$$ax + by = \gcd(a, b).$$

The back-substitution process is a constructive proof that the ideal (a, b) is principal.

This construction also gives modular inverses. If $\gcd(a, n) = 1$, then

$$ax + ny = 1,$$

so x is the inverse of a modulo n . The original algorithmic procedure is therefore the computational heart of quotient-ring arithmetic.

§24.18 Euler's theorem as group theory

Euler's theorem

$$a^{\varphi(n)} \equiv 1 \pmod{n}$$

for $\gcd(a, n) = 1$ follows from Lagrange's theorem in the finite group $(\mathbb{Z}/n\mathbb{Z})^\times$. The elementary proof by reduced residue systems is already a group action proof: multiplication by a permutes the reduced residues.

This viewpoint clarifies why the exponent $\varphi(n)$ appears: it is the order of the group of units.

Divisibility, Euclidean Algorithm, Greatest Common Divisor, Bezout, and Congruences

N-20221031

§25.1 Divisibility

The first page begins with divisibility. For integers a, b , with $a \neq 0$, we write

$$a \mid b$$

if there exists $k \in \mathbb{Z}$ such that

$$b = ak.$$

If no such integer exists, we write $a \nmid b$.

The basic properties recorded in the note are:

- (i) $a \mid a$ and $a \mid 0$;
- (ii) if $a \mid b$ and $b \mid a$, then $a = \pm b$;
- (iii) if $a \mid b$ and $b \mid c$, then $a \mid c$;
- (iv) if $a \mid b$ and $a \mid c$, then $a \mid bx + cy$ for all integers x, y ;
- (v) if $a \mid b$, then $ac \mid bc$ for every integer c .

The handwritten proofs repeatedly use the definition. For instance, if $b = am$ and $c = an$, then

$$bx + cy = a(mx + ny),$$

so $a \mid bx + cy$.

§25.2 Division algorithm

The division algorithm says: if $a, b \in \mathbb{Z}$ and $b > 0$, then there exist unique integers q, r such that

$$a = bq + r, \quad 0 \leq r < b.$$

The note calls attention to uniqueness. If

$$a = bq + r = bq' + r', \quad 0 \leq r, r' < b,$$

then

$$b(q - q') = r' - r.$$

The right side has absolute value less than b , while the left side is a multiple of b . Therefore $r = r'$ and then $q = q'$.

§25.3 Greatest common divisor and least common multiple

The greatest common divisor is written

$$(a, b),$$

and the least common multiple is written

$$[a, b].$$

The note records the identity

$$(a, b)[a, b] = |ab|.$$

It also records useful gcd properties:

$$(a, 1) = 1, \quad (a, 0) = |a|, \quad (a, a) = |a|,$$

$$(a, b) = (b, a) = (ta, tb)/|t|,$$

and most importantly the Euclidean step

$$(a, b) = (a - b, b),$$

or more generally

$$(a, b) = (a - qb, b).$$

The proof of the Euclidean step is inclusion both ways. Any common divisor of a and b divides $a - qb$. Conversely, any common divisor of $a - qb$ and b divides a .

§25.4 Euclidean algorithm

Repeatedly applying

$$(a, b) = (b, r), \quad a = bq + r,$$

produces the Euclidean algorithm. The last nonzero remainder is the gcd.

The note's red proof makes the descent explicit: the remainders strictly decrease and remain nonnegative, so the process must terminate. This is not just an algorithm; it is a proof that the gcd can be computed by repeated division.

§25.5 Bezout theorem

Theorem 25.5.1 (Bezout identity)

For integers a, b , there exist integers x, y such that

$$ax + by = \gcd(a, b).$$

Bezout's theorem says that if $a, b \in \mathbb{Z}$ are not both zero, then there exist $s, t \in \mathbb{Z}$ such that

$$as + bt = (a, b).$$

The note proves this by induction on $a + b$ and also implicitly through the Euclidean algorithm. When the gcd is obtained as the last nonzero remainder, one substitutes backward to express that remainder as an integer linear combination of a and b .

A key special case is:

$$(a, b) = 1 \iff \exists s, t \in \mathbb{Z}, as + bt = 1.$$

This is the form most often used in proofs.

§25.6 Gauss lemma and divisibility consequences

Lemma 25.6.1 (Gauss lemma)

If p is prime and $p \mid ab$, then $p \mid a$ or $p \mid b$.

The next page proves Euclid's lemma or Gauss's lemma:

$$(a, b) = 1, \quad a \mid bc \implies a \mid c.$$

Using Bezout, choose s, t such that

$$as + bt = 1.$$

Multiplying by c gives

$$asc + btc = c.$$

Since $a \mid asc$ and $a \mid bc$ implies $a \mid btc$, we get $a \mid c$.

This is the standard tool behind unique factorization and many modular arguments.

§25.7 Prime numbers

A prime number p satisfies: if

$$p \mid ab,$$

then

$$p \mid a \quad \text{or} \quad p \mid b.$$

This follows from Gauss's lemma. If $p \nmid a$, then $(p, a) = 1$, so from $p \mid ab$ we get $p \mid b$.

The note also uses the property repeatedly in exercises involving prime powers and factorization.

§25.8 Congruences

For an integer $m > 0$, write

$$a \equiv b \pmod{m}$$

if $m \mid (a - b)$. Congruence is compatible with addition and multiplication:

$$a \equiv b, \quad c \equiv d \pmod{m} \implies a + c \equiv b + d \pmod{m},$$

$$ac \equiv bd \pmod{m}.$$

The handwritten examples include reductions of powers modulo m and using congruences to rule out equations.

§25.9 Exercises and common methods

The printed exercise pages ask one to prove identities such as

$$(a, b) = (a - b, b), \quad (a^m - 1, a^n - 1) = a^{(m, n)} - 1,$$

and divisibility statements involving expressions like

$$2^{2n+1} + 1, \quad 3^{2n} - 1.$$

The repeated methods are:

- (i) apply the Euclidean algorithm to gcd expressions;
- (ii) use Bezout when coprimality appears;
- (iii) factor expressions like $x^m - y^m$;
- (iv) reduce powers modulo a suitable modulus;
- (v) use prime factorization when divisor functions appear.

§25.10 From example to method

The note develops number theory from one simple relation: divisibility. The division algorithm gives remainders; remainders give the Euclidean algorithm; the Euclidean algorithm gives Bezout; Bezout gives Gauss's lemma; Gauss's lemma supports prime factorization and congruence arguments. This chain is the backbone of elementary number theory.

§25.11 Division algorithm as normal form

For integers a, b with $b > 0$, division gives

$$a = qb + r, \quad 0 \leq r < b.$$

The remainder r is a normal form for the congruence class of a modulo b .

§25.12 Euclidean algorithm and invariant gcd

The key invariant is

$$\gcd(a, b) = \gcd(b, a - qb).$$

Each Euclidean step preserves the gcd while reducing size. Back-substitution then expresses the gcd as a Bezout combination.

§25.13 Gauss lemma and prime ideals

Lemma 25.13.1 (Gauss lemma, ideal form)

For a prime p , the ideal $(p) \subset \mathbb{Z}$ is prime; equivalently, $p \mid ab$ implies $p \mid a$ or $p \mid b$.

Gauss's lemma says if p is prime and $p \mid ab$, then $p \mid a$ or $p \mid b$. In ideal language, (p) is a prime ideal of $\mathbb{Z}$. This is the algebraic property that makes unique factorization work.

§25.14 Normal forms, gcd invariants, and prime divisibility

Divisibility, gcd, Bezout, congruence, and primes form one system. Remainders are quotient normal forms, the Euclidean algorithm preserves gcd, and prime divisibility is the ideal-theoretic foundation of unique factorization.

§25.15 $\mathbb{Z}$ as a PID

The full chain of gcd, lcm, Bezout identity, congruence, and modular inverse is unified by the fact that $\mathbb{Z}$ is a principal ideal domain. Every ideal is of the form

$$(n) = n\mathbb{Z}.$$

For two integers,

$$(a) + (b) = (\gcd(a, b)), \quad (a) \cap (b) = (\text{lcm}(a, b)).$$

This gives a structural proof of many gcd-lcm identities.

For example,

$$\gcd(a, b) \text{lcm}(a, b) = |ab|$$

is the numerical reflection of how prime valuations take minima and maxima.

§25.16 Modular inverses as units

The congruence equation

$$ax \equiv 1 \pmod{n}$$

asks whether the residue class of a is a unit in the ring $\mathbb{Z}/n\mathbb{Z}$. It has a solution iff $\gcd(a, n) = 1$.

More generally,

$$ax \equiv b \pmod{n}$$

has a solution iff $\gcd(a, n) \mid b$. This is because the image of multiplication by a on $\mathbb{Z}/n\mathbb{Z}$ has size $n/\gcd(a, n)$ and consists exactly of multiples of that gcd class.

§25.17 Smith normal form over integers

Linear equations over $\mathbb{Z}$ require more than rank. For an integer matrix A , Smith normal form says there exist unimodular integer matrices P, Q such that

$$PAQ = \text{diag}(d_1, \dots, d_r, 0, \dots, 0), \quad d_i \mid d_{i+1}.$$

This diagonal form solves integer linear systems and classifies finitely generated abelian groups.

Thus gcd and Bezout operations are the 1×1 and row-column-operation shadows of Smith normal form.

§25.18 CRT as product decomposition

When m, n are coprime,

$$\mathbb{Z}/mn\mathbb{Z} \cong \mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}.$$

The Chinese remainder theorem is therefore not simply a trick for solving congruences. It says the quotient ring modulo a product of coprime moduli decomposes as a product of quotient rings.

This decomposition explains why arithmetic functions and congruence problems split prime by prime.

§25.19 Solving congruences by image and kernel

Multiplication by a defines a group homomorphism

$$m_a : \mathbb{Z}/n\mathbb{Z} \rightarrow \mathbb{Z}/n\mathbb{Z}, \quad x \mapsto ax.$$

The congruence $ax \equiv b \pmod{n}$ asks whether b lies in the image of this map. The kernel has size $\gcd(a, n)$, so the image has size $n/\gcd(a, n)$.

The solvability condition $\gcd(a, n) \mid b$ is therefore rank-nullity for a finite abelian group. Once solvable, there are exactly $\gcd(a, n)$ solutions modulo n .

§25.20 Smith normal form and congruences

A system of integer linear equations

$$Ax = b$$

can be simplified by invertible integer row and column operations. Smith normal form diagonalizes A into invariant factors d_i . The system reduces to independent divisibility conditions

$$d_i y_i = c_i.$$

Thus the one-variable congruence theorem is the first case of a full normal-form theory for homomorphisms of finitely generated abelian groups.

Prime Numbers, Composite Numbers, Divisor Functions, Congruences, and Number-Theoretic Exercises

N-20221119

§26.1 Prime and composite numbers

The first page defines prime and composite numbers. An integer $a > 1$ is prime if its only positive divisors are 1 and a . If $a > 1$ is not prime, it is composite.

A basic observation is that if a is composite, then it has a prime divisor not exceeding $\sqrt{a}$. Indeed, if

$$a = pq, \quad 1 < p \leq q,$$

then

$$p^2 \leq pq = a,$$

so $p \leq \sqrt{a}$. This justifies trial division up to $\sqrt{a}$.

The note then mentions the sieve method. To find primes up to 50, begin with 2, 3, 5, 7 as sieving primes because

$$\sqrt{50} < 8.$$

Cross out multiples of these primes. The remaining numbers are prime.

§26.2 Euclid's theorem

Theorem 26.2.1 (Euclid's theorem)

There are infinitely many prime numbers.

Euclid's theorem says that there are infinitely many primes. Suppose the primes were

$$p_1, p_2, \dots, p_k.$$

Let

$$N = p_1 p_2 \cdots p_k + 1.$$

No p_i divides N , since division by p_i leaves remainder 1. But every integer $N > 1$ has a prime divisor. This prime divisor is not among the listed primes, contradiction.

The handwritten version uses exactly this remainder idea: the product of all known primes plus one creates a new prime factor.

§26.3 Prime divisibility

The page records several prime-divisibility facts. If p is prime and

$$p \mid ab,$$

then

$$p \mid a \quad \text{or} \quad p \mid b.$$

More generally, if

$$p \mid a_1 a_2 \cdots a_n,$$

then p divides at least one factor. This is proved by induction from the two-factor statement.

The note also records the contrapositive: if p divides none of the factors, then it does not divide their product.

§26.4 Fundamental theorem of arithmetic

Theorem 26.4.1 (Fundamental theorem of arithmetic)

Every integer greater than 1 has a unique prime factorization, up to the order of the prime factors.

Every integer $n > 1$ can be written uniquely as

$$n = p_1^{\alpha_1} p_2^{\alpha_2} \cdots p_k^{\alpha_k},$$

where the p_i are distinct primes and $\alpha_i > 0$. The order of factors is the only ambiguity.

The note uses this factorization to compute divisor functions. If

$$d \mid n,$$

then

$$d = p_1^{\beta_1} p_2^{\beta_2} \cdots p_k^{\beta_k}, \quad 0 \leq \beta_i \leq \alpha_i.$$

Therefore the number of positive divisors is

$$\tau(n) = \prod_{i=1}^k (\alpha_i + 1).$$

The sum of positive divisors is

$$\sigma(n) = \prod_{i=1}^k (1 + p_i + p_i^2 + \cdots + p_i^{\alpha_i}) = \prod_{i=1}^k \frac{p_i^{\alpha_i+1} - 1}{p_i - 1}.$$

§26.5 Perfect numbers and square numbers

A number n is perfect if the sum of its positive divisors is $2n$:

$$\sigma(n) = 2n.$$

The note lists the first examples:

$$6, \quad 28, \quad 496, \quad 8128, \dots$$

Another important fact is:

$$n \text{ is a square} \iff \tau(n) \text{ is odd.}$$

If $n = p_1^{\alpha_1} \cdots p_k^{\alpha_k}$, then n is a square iff every α_i is even. This is equivalent to every $\alpha_i + 1$ being odd, so their product $\tau(n)$ is odd.

§26.6 Exponent of a prime in $n!$

The note records Legendre's formula. If p is prime, then the exponent of p in $n!$ is

$$\sum_{j \geq 1} \left\lfloor \frac{n}{p^j} \right\rfloor.$$

This counts multiples of p , then multiples of p^2 , then multiples of p^3 , and so on. For example, numbers divisible by p^2 contribute an extra factor of p beyond the first count.

§26.7 Congruence exercises

The printed pages contain exercises on congruences. The basic definition is

$$a \equiv b \pmod{m} \iff m \mid (a - b).$$

The operations of addition and multiplication are compatible with congruence:

$$a \equiv b, c \equiv d \pmod{m} \implies a + c \equiv b + d, \quad ac \equiv bd \pmod{m}.$$

A typical problem asks for a digit or a remainder of a large power. The method is to find the period of powers modulo m . For example, powers of 2 modulo 10 cycle through

$$2, 4, 8, 6.$$

Once the period is known, the exponent can be reduced modulo the period.

§26.8 Broader perspective

This note is a transition from divisibility to arithmetic structure. Prime factorization turns every positive integer into exponent data. Divisor counting and divisor sums become products over primes. Congruences turn equality into equality up to remainders. Together, these tools make large integer problems finite and computable.

§26.9 Prime exponents and valuations

For a prime p , the exponent $v_p(n)$ turns divisibility into arithmetic:

$$a \mid b \iff v_p(a) \leq v_p(b) \text{ for all } p.$$

Legendre's formula gives

$$v_p(n!) = \sum_{k \geq 1} \left\lfloor \frac{n}{p^k} \right\rfloor.$$

It counts multiples of p , then multiples of p^2 , and so on.

§26.10 Divisor functions as multiplicative functions

If $n = \prod p_i^{e_i}$, then

$$d(n) = \prod_i (e_i + 1), \quad \sigma(n) = \prod_i \frac{p_i^{e_i+1} - 1}{p_i - 1}.$$

These formulas hold because choosing a divisor is choosing each prime exponent independently. This is multiplicativity coming from unique factorization.

§26.11 Congruences and finite groups

Reduced residue classes modulo n form the unit group $(\mathbb{Z}/n\mathbb{Z})^\times$. Euler's theorem is the statement that every element of this finite group has order dividing the group size.

§26.12 Prime coordinates and finite unit groups

Prime factorization supplies coordinates, divisor functions multiply across primes, factorial valuations count prime powers, and congruence problems become group theory in finite unit groups.

§26.13 Arithmetic functions and convolution

The divisor-counting functions in the original note are examples of arithmetic functions, functions

$$f : \mathbb{N} \rightarrow \mathbb{C}.$$

They form an algebra under Dirichlet convolution:

$$(f * g)(n) = \sum_{d|n} f(d)g(n/d).$$

The identity element is

$$\varepsilon(n) = \begin{cases} 1, & n = 1, \\ 0, & n > 1. \end{cases}$$

For example, if $1(n) = 1$ for all n , then the divisor-counting function is

$$\tau = 1 * 1.$$

The sum-of-divisors function is

$$\sigma = \text{id} * 1, \quad \text{id}(n) = n.$$

§26.14 Möbius inversion in number theory

The number-theoretic Möbius function is

$$\mu(n) = \begin{cases} 1, & n = 1, \\ (-1)^k, & n \text{ is a product of } k \text{ distinct primes,} \\ 0, & p^2 \mid n \text{ for some } p. \end{cases}$$

It is the inverse of the constant function 1 under Dirichlet convolution:

$$\mu * 1 = \varepsilon.$$

Therefore if

$$F(n) = \sum_{d|n} f(d),$$

then

$$f(n) = \sum_{d|n} \mu(d)F(n/d).$$

This is Möbius inversion on the divisibility poset.

§26.15 Dirichlet series and Euler products

To an arithmetic function f one may attach a Dirichlet series

$$D_f(s) = \sum_{n=1}^{\infty} \frac{f(n)}{n^s}.$$

Dirichlet convolution corresponds to multiplication of Dirichlet series when convergence is justified:

$$D_{f*g}(s) = D_f(s)D_g(s).$$

For the constant function 1,

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s} = \prod_p \frac{1}{1 - p^{-s}}.$$

The Euler product is unique factorization written analytically.

The elementary prime and divisor computations in the note therefore open the door to analytic number theory: primes are encoded in the analytic behavior of $\zeta(s)$ and related series.

§26.16 P-adic valuations and multiplicativity

If

$$n = \prod_p p^{v_p(n)},$$

then multiplicative functions are determined by their values on prime powers. For example,

$$\tau(n) = \prod_p (v_p(n) + 1), \quad \sigma(n) = \prod_p \frac{p^{v_p(n)+1} - 1}{p - 1}.$$

This is why divisor-function formulas are products over primes. Elementary prime factorization is already a coordinate system for arithmetic functions.

§26.17 Euler products and prime independence

The Euler product

$$\zeta(s) = \prod_p (1 - p^{-s})^{-1}$$

comes from multiplying the geometric series

$$1 + p^{-s} + p^{-2s} + \dots$$

for every prime p . Expanding the product chooses an exponent for each prime, hence chooses a positive integer.

This is unique factorization written as analysis. Elementary divisor formulas are multiplicative because the exponent choices at different primes are independent.

§26.18 Möbius inversion over primes

The Möbius function $\mu(n)$ performs inclusion–exclusion over the set of primes dividing n . If n is squarefree with k prime factors, then $\mu(n) = (-1)^k$. If a square divides n , it contributes zero because repeated prime factors break the simple subset structure.

Thus number-theoretic Möbius inversion and set-theoretic inclusion–exclusion are the same operation on different posets: divisors ordered by divisibility versus subsets ordered by inclusion.

Congruence Classes, Complete Residue Systems, and Euler's Function

N-20221126

§27.1 Congruence

Let m be a positive integer and $a, b \in \mathbb{Z}$. We say that a and b are congruent modulo m , written

$$a \equiv b \pmod{m},$$

if

$$m \mid (a - b).$$

This definition is just divisibility language with a new notation. It allows us to work with remainders rather than whole integers.

§27.2 Basic properties

The note lists the standard properties:

- (i) Reflexivity: $a \equiv a \pmod{m}$.
- (ii) Symmetry: if $a \equiv b \pmod{m}$, then $b \equiv a \pmod{m}$.
- (iii) Transitivity: if $a \equiv b \pmod{m}$ and $b \equiv c \pmod{m}$, then $a \equiv c \pmod{m}$.
- (iv) Compatibility with addition and multiplication:

$$a \equiv b, c \equiv d \pmod{m} \Rightarrow a + c \equiv b + d, \quad ac \equiv bd \pmod{m}.$$

- (v) If $a \equiv b \pmod{m}$, then

$$a^n \equiv b^n \pmod{m}$$

for every positive integer n .

The last two rules are what make congruence useful: one can replace a large number by a small remainder and keep computing.

§27.3 Residue classes

A residue class modulo n is a set

$$M_i = \{x \in \mathbb{Z} : x \equiv i \pmod{n}\}.$$

The integers split into n disjoint classes:

$$M_0, M_1, \dots, M_{n-1}.$$

This is a partition of $\mathbb{Z}$. The note records two important consequences:

- (i) among any $n + 1$ integers, two are congruent modulo n ;
- (ii) if $(i, n) = 1$, then every integer in M_i is coprime to n .

The first is the pigeonhole principle; the second follows because congruent numbers have the same gcd with n .

§27.4 Complete and reduced residue systems

A complete residue system modulo n is any set of n representatives, one from each residue class. The standard example is

$$0, 1, 2, \dots, n - 1.$$

A reduced residue system modulo n consists of one representative from each residue class coprime to n . The number of such classes is Euler's function

$$\varphi(n).$$

If

$$n = p_1^{\alpha_1} \cdots p_k^{\alpha_k},$$

then

$$\varphi(n) = n \left(1 - \frac{1}{p_1}\right) \cdots \left(1 - \frac{1}{p_k}\right).$$

In particular,

$$\varphi(p) = p - 1, \quad \varphi(p^a) = p^a - p^{a-1}.$$

§27.5 Example: last digit

To compute the last digit of 15^{200} , work modulo 10. Since

$$15 \equiv 5 \pmod{10},$$

we get

$$15^{200} \equiv 5^{200} \equiv 5 \pmod{10}.$$

Thus the last digit is 5.

The note's method is to reduce early. Do not expand powers; replace the base by its residue.

§27.6 Example: modulo 17

To compute

$$15^{100} \pmod{17},$$

notice

$$15 \equiv -2 \pmod{17}.$$

Then

$$15^{100} \equiv (-2)^{100} = 2^{100} \pmod{17}.$$

Since

$$2^4 = 16 \equiv -1 \pmod{17},$$

we have

$$2^8 \equiv 1 \pmod{17}.$$

Therefore

$$2^{100} = 2^{96} 2^4 \equiv 1^{12} (-1) \equiv 16 \pmod{17}.$$

§27.7 Example: reducing a long expression

For an expression such as

$$2001 \cdot 2002 + 2003 \cdot 2004 + 2005 \cdot 2006$$

modulo 45, reduce every factor modulo 45 first. The arithmetic becomes small. The main habit is:

$$\text{large expression} \longrightarrow \text{small residues.}$$

This is the same principle behind all modular computation.

§27.8 A polynomial congruence by induction

The note also includes a divisibility proof for an expression depending on k , of the form

$$12^{2k+1} + 36^k + 5^k + 3.$$

The standard method is induction or modular periodicity. One chooses the modulus, checks the base case, and then shows the expression changes by a multiple of the modulus when k increases. Equivalently, reduce each base modulo the modulus and use periodicity of powers.

§27.9 Quadratic residues modulo 8

A consolidation exercise proves that a perfect square modulo 8 is congruent to

$$0, 1, \text{ or } 4.$$

Indeed, every integer is congruent modulo 8 to one of

$$0, 1, 2, 3, 4, 5, 6, 7.$$

Squaring gives residues

$$0, 1, 4, 1, 0, 1, 4, 1.$$

Thus only 0, 1, 4 occur. This is a powerful small filter in number theory problems.

§27.10 Where the idea leads

Congruence classes are quotient thinking. Instead of remembering the whole integer, we remember only its class in

$$\mathbb{Z}/n\mathbb{Z}.$$

Complete residue systems are representatives of all classes; reduced residue systems are the invertible classes. Euler's function counts the units. This note is therefore the elementary entrance to the ring $\mathbb{Z}/n\mathbb{Z}$ and its unit group.

§27.11 Residue systems as representatives

A complete residue system modulo n is a choice of one representative from each class in $\mathbb{Z}/n\mathbb{Z}$. A reduced residue system is a choice of representatives from the unit group

$$(\mathbb{Z}/n\mathbb{Z})^\times.$$

Thus the difference between complete and reduced systems is the difference between the whole quotient ring and its invertible elements.

§27.12 Euler function through prime powers

For a prime power,

$$\varphi(p^k) = p^k - p^{k-1} = p^k \left(1 - \frac{1}{p}\right).$$

Using the Chinese remainder theorem gives

$$\varphi(n) = n \prod_{p|n} \left(1 - \frac{1}{p}\right).$$

§27.13 Orders and exponent reduction

If a is a unit modulo n , its order is the smallest d with $a^d \equiv 1 \pmod{n}$. Exponent reduction is really reduction modulo the order of a , not blindly modulo $\varphi(n)$.

§27.14 Residues, units, and exponent cycles

Residue classes form quotient rings, reduced residue systems form unit groups, Euler's function counts units, and exponent tricks are statements about orders in finite groups.

§27.15 Residue systems as groups of units

A reduced residue system modulo n is the unit group

$$(\mathbb{Z}/n\mathbb{Z})^\times.$$

Its order is Euler's function $\varphi(n)$. Euler's theorem says that if $\gcd(a, n) = 1$, then

$$a^{\varphi(n)} \equiv 1 \pmod{n}.$$

This is simply Lagrange's theorem applied to the finite group of units.

Thus the elementary reduced-residue computations are group theory in disguise.

§27.16 CRT decomposition of unit groups

If m, n are coprime, then

$$\mathbb{Z}/mn\mathbb{Z} \cong \mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}.$$

Taking units gives

$$(\mathbb{Z}/mn\mathbb{Z})^\times \cong (\mathbb{Z}/m\mathbb{Z})^\times \times (\mathbb{Z}/n\mathbb{Z})^\times.$$

This explains the multiplicativity of φ .

For prime powers, the structure of the unit group is more subtle. For odd primes p , the group $(\mathbb{Z}/p^k\mathbb{Z})^\times$ is cyclic of order $p^{k-1}(p-1)$. Powers of 2 require separate treatment.

§27.17 Primitive roots and characters

A primitive root modulo n is a generator of the unit group. When the unit group is cyclic, every unit is a power of one element. This converts multiplicative congruence problems into exponent arithmetic.

Dirichlet characters are homomorphisms

$$\chi : (\mathbb{Z}/n\mathbb{Z})^\times \rightarrow \mathbb{C}^\times$$

extended to all integers by setting $\chi(a) = 0$ when $\gcd(a, n) > 1$. They are Fourier characters on the finite abelian group of units and are central in Dirichlet's theorem on primes in arithmetic progressions.

§27.18 Returning to residue-system exercises

The note's residue tables, reduced systems, and Euler-function formulas are not just modular arithmetic. They are computations inside finite groups. The next levels are group decomposition, primitive roots, and characters.

§27.19 Primitive roots as logarithms in finite groups

When $(\mathbb{Z}/n\mathbb{Z})^\times$ is cyclic with generator g , every unit has the form g^k . This turns multiplication into addition of exponents modulo the group order. In other words, a primitive root supplies a discrete logarithm coordinate.

This is powerful but also limited: not every modulus has primitive roots. The classification says primitive roots exist for

$$n = 1, 2, 4, p^k, 2p^k$$

where p is an odd prime. This explains why some modular arithmetic problems become simple after choosing a primitive root, while others require decomposing the unit group into several cyclic factors.

§27.20 Characters as Fourier analysis on unit groups

Dirichlet characters are not arbitrary functions; they are group homomorphisms from the unit group to complex roots of unity. Orthogonality of characters gives filters for congruence classes. This is the multiplicative analogue of using roots of unity to filter residues modulo n .

The elementary reduced-residue system is therefore the domain on which finite harmonic analysis is performed.

Factorization of $\sum_{i=0}^k x^{in}$: Roots of Unity, False Generalization, and Cyclotomic Structure

N-20230127

§28.1 The original problem and the first successful method

The paper begins with the factorization of

$$x^{2n} + x^n + 1$$

by

$$x^2 + x + 1.$$

Let ω be a primitive cubic root of unity. Then

$$\omega^2 + \omega + 1 = 0, \quad \omega^3 = 1, \quad \omega \neq 1.$$

The polynomial $x^2 + x + 1$ divides $x^{2n} + x^n + 1$ iff both ω and ω^2 are roots. Substituting $x = \omega$ gives

$$\omega^{2n} + \omega^n + 1.$$

This is zero exactly when ω^n is a nontrivial cubic root of unity. Therefore the divisibility holds iff

$$n \not\equiv 0 \pmod{3}.$$

This proof is short because roots of unity convert a divisibility question into an exponent congruence question.

§28.2 The direct factorization attempt

The note then tries to prove the same statement by direct algebra. Since

$$x^{2n} + x^n + 1 = \frac{x^{3n} - 1}{x^n - 1},$$

one can try to factor numerator and denominator using difference of squares or cubes. For example, when $n = 2$,

$$x^4 + x^2 + 1 = (x^2 + x + 1)(x^2 - x + 1).$$

When $n = 4$, one obtains several quadratic-looking factors, including again $x^2 + x + 1$.

But when n changes parity or becomes divisible by 3, the algebraic route splits into cases. The original article carefully lists cases such as $n = 3m \pm 1$ and tracks which geometric-series quotients appear. The lesson is that direct factorization can prove the theorem, but it obscures the simple reason: the exponent map $z \mapsto z^n$ acts on the three roots of unity.

§28.3 The attempted generalization

The natural next question is whether

$$1 + x^n + x^{2n} + \cdots + x^{kn}$$

is divisible by

$$1 + x + x^2 + \cdots + x^k$$

whenever n is not a multiple of $k+1$. This first guess is too weak. The correct condition is coprimality.

Theorem 28.3.1 (Root-of-unity test for the generalization)

Over a field of characteristic zero,

$$1 + x + x^2 + \cdots + x^k \quad \text{divides} \quad 1 + x^n + x^{2n} + \cdots + x^{kn}$$

if and only if

$$\gcd(n, k+1) = 1.$$

Proof. The roots of

$$1 + x + x^2 + \cdots + x^k = \frac{x^{k+1} - 1}{x - 1}$$

are exactly the nontrivial $(k+1)$ -st roots of unity. Let $\xi \neq 1$ satisfy $\xi^{k+1} = 1$. We need

$$1 + \xi^n + \xi^{2n} + \cdots + \xi^{kn} = 0$$

for every such ξ .

This is a geometric sum. It vanishes exactly when ξ^n is a nontrivial $(k+1)$ -st root of unity. Therefore the map

$$\xi \mapsto \xi^n$$

must send every nontrivial $(k+1)$ -st root to another nontrivial $(k+1)$ -st root. Equivalently, the power map on the cyclic group μ_{k+1} must have trivial kernel. This happens exactly when n is invertible modulo $k+1$, i.e.

$$\gcd(n, k+1) = 1.$$

□

The original false guess $n \not\equiv 0 \pmod{k+1}$ fails because a nonzero residue may still not be invertible. For instance, modulo 6, the residue 2 is nonzero but the map $z \mapsto z^2$ has a nontrivial kernel on sixth roots of unity.

§28.4 The computational experiment

The note then turns to MATLAB. It computes roots of

$$x^{kn} + x^{(k-1)n} + \cdots + x^n + 1 = 0$$

for different values of n and plots them in the complex plane, sometimes using a third coordinate to record n . This is not just decoration. The picture reveals periodicity in the arguments of the roots.

The observed pattern is: the roots lie on the unit circle, their arguments are fractions of 2π , and as n varies the same angular patterns repeat. For $k = 2$, the pattern has two angles per block; for $k = 3$, three angles per block; in general, the $(k + 1)$ -st-root structure controls the blocks.

This experimental section is important because it records a genuine mathematical process: a false conjecture is tested, numerical pictures suggest the right invariant, and then the invariant is expressed exactly by roots of unity.

§28.5 Explicit root formula

The equation

$$1 + x^n + x^{2n} + \cdots + x^{kn} = 0$$

can be rewritten as

$$\frac{x^{n(k+1)} - 1}{x^n - 1} = 0.$$

Thus the roots are those $n(k+1)$ -st roots of unity whose n th power is not 1. Equivalently,

$$x = \exp\left(\frac{2\pi i}{n(k+1)}(r(k+1) + m)\right),$$

where $r = 0, \dots, n-1$ and $m = 1, \dots, k$. This gives exactly kn roots, as expected from the degree.

The original note writes these roots as products of exponentials:

$$e^{2\pi i r/n} e^{2\pi i m/(n(k+1))}.$$

Substituting into the polynomial gives a finite geometric sum whose ratio is a nontrivial $(k+1)$ -st root of unity, hence the sum is zero.

§28.6 Common real factors

The final theorem concerns common real-coefficient factors among these polynomials as n varies. Since complex roots come in conjugate pairs, a set of roots stable under conjugation gives a real polynomial factor. When $d = \gcd(n, k+1) > 1$, some of the angular positions line up across different n , producing common roots and hence common real factors.

This is the cyclotomic viewpoint in disguise. The polynomial

$$x^m - 1$$

factors as

$$x^m - 1 = \prod_{d|m} \Phi_d(x),$$

and the geometric sum

$$1 + x + \cdots + x^k = \frac{x^{k+1} - 1}{x - 1}$$

contains the cyclotomic factors $\Phi_d(x)$ for divisors $d > 1$ of $k+1$. Replacing x by x^n pulls these cyclotomic conditions back along the power map.

§28.7 The organizing idea

The value of this note is the exploration, not only the final condition. It begins with a correct theorem, tries a direct factorization proof, guesses a generalization, finds that the guess is false, uses computation to see the roots, and finally explains the pattern by complex exponentials. That is a miniature version of mathematical research: algebraic manipulation, conjecture, counterexample, visualization, and structural explanation.

§28.8 Cyclotomic summary of the experiment

The computational pictures, the failed generalization, and the final divisibility criterion all express one group-theoretic fact: roots of unity form finite cyclic groups, and power maps on cyclic groups are controlled by greatest common divisors.

Remark 28.8.1 (From roots to factors) — A rational polynomial divisibility statement can be checked on all Galois conjugate roots. For roots of unity, the conjugates of a primitive m -th root are

$$\zeta_m \mapsto \zeta_m^a, \quad a \in (\mathbb{Z}/m\mathbb{Z})^\times.$$

Thus a common factor is not merely one common complex root; it is a conjugation-stable set of common roots.

The cyclotomic factorization

$$x^m - 1 = \prod_{d|m} \Phi_d(x)$$

separates roots by exact order. Geometric sums select all nontrivial roots of a given order divisor, while substitutions $x \mapsto x^n$ pull those root conditions back along the power map. This is the structural explanation of the root plots in the original experiment.

§28.9 Cyclotomic polynomials

The roots of $x^n - 1$ are the n -th roots of unity. The primitive n -th roots are those with exact order n . The cyclotomic polynomial is

$$\Phi_n(x) = \prod_{\substack{1 \leq k \leq n \\ \gcd(k,n)=1}} (x - e^{2\pi i k/n}).$$

It has degree $\varphi(n)$ and integer coefficients.

The factorization is

$$x^n - 1 = \prod_{d|n} \Phi_d(x).$$

This identity is the structural version of many elementary root-of-unity factorization exercises.

§28.10 Galois action on roots of unity

The field $\mathbb{Q}(\zeta_n)$ generated by a primitive root of unity ζ_n is the n -th cyclotomic field. Its Galois group is

$$\text{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \cong (\mathbb{Z}/n\mathbb{Z})^\times.$$

An element $a \in (\mathbb{Z}/n\mathbb{Z})^\times$ acts by

$$\zeta_n \mapsto \zeta_n^a.$$

Thus the same reduced residue system that appears in elementary number theory controls the symmetries of roots of unity.

§28.11 Resultants and common roots

Two polynomials have a common root exactly when their resultant is zero. For example, studying when a cyclotomic polynomial divides another expression can be reframed as asking which roots are common.

The elementary method says: substitute a primitive root ω and check whether the expression vanishes. The algebraic version says: divisibility is detected by vanishing on all Galois conjugate roots, or by a resultant condition.

§28.12 Finite fields and character sums

Roots of unity also appear in finite fields. The multiplicative group $\mathbb{F}_q^\times$ is cyclic of order $q - 1$. Therefore roots of $x^n = 1$ in $\mathbb{F}_q$ are controlled by $\gcd(n, q - 1)$.

Character sums such as

$$\sum_{x \in \mathbb{F}_q} \chi(x) \psi(f(x))$$

are research-level descendants of elementary root-of-unity sums. They connect finite fields, Fourier analysis on finite groups, and algebraic geometry.

§28.13 Returning to the original experiments

The note's factorization experiments with roots on the complex plane are the concrete beginning of cyclotomic theory. The advanced bridge is precise: roots of unity form finite cyclic groups, cyclotomic polynomials package primitive roots, and Galois groups act by exponentiation modulo n .

§28.14 Primitive roots and cyclotomic factor tests

To decide whether a polynomial expression is divisible by $\Phi_n(x)$, it is enough to substitute a primitive n -th root of unity ζ and check whether the expression vanishes. Because Φ_n is the minimal polynomial of ζ over $\mathbb{Q}$, vanishing at one primitive root forces vanishing at all Galois conjugates.

This explains why root-of-unity substitution is not just a clever complex-number trick. It is a minimal-polynomial test.

§28.15 Cyclotomic identities through Möbius inversion

From

$$x^n - 1 = \prod_{d|n} \Phi_d(x)$$

one obtains

$$\Phi_n(x) = \prod_{d|n} (x^d - 1)^{\mu(n/d)}.$$

This is Möbius inversion on the divisor lattice. The alternating exponents in this formula are the same inclusion–exclusion pattern seen in counting and arithmetic functions.

Thus the note’s factorization experiments sit at an intersection of complex roots, integer polynomials, Galois symmetry, and Möbius inversion.

IV

Arithmetic Geometry

Part IV: Contents

N-20211007	175
29.1 The question is an inequality, not a slogan	175
29.2 A first chain of rigidity results	175
29.3 From Mordell finiteness to effective inequalities	176
29.4 The Szpiro inequality as a concrete target	176
29.5 Vojta's inequality and the hyperbolic curve viewpoint	177
29.6 Arithmetic fundamental groups	178
29.7 The local elliptic curve model: Tate curves and q -parameters	178
29.8 Normalised degree as log-volume	179
29.9 Why ordinary scheme-theoretic geometry is not enough	179
29.10 The route from deformation to inequality	180
29.11 What this first note establishes	181
 N-20220123	 183
30.1 Why this note starts with algebraic number theory	183
30.2 Number fields and algebraic integers	183
30.3 Ideals and factorization	183
30.4 Quadratic fields and class number	184
30.5 Localization, completion, and p -adic numbers	184
30.6 Modular forms and L -functions	186
30.7 First understanding of Langlands	186
30.8 Local and global correspondences	187
30.9 Cross-field intuitions	187
30.10 Numerical experiments	187
30.11 Topics to learn next	189
30.12 Synthesis	189
30.13 Local-global intuition	189
30.14 Why L -functions enter	189
30.15 A first Langlands slogan	190
30.16 Structural viewpoint	190
30.17 Ideals as repaired unique factorization	190
30.18 Local fields and arithmetic microscopes	190
30.19 Modular forms and Galois representations	190
30.20 Arithmetic symmetry as the bridge to Langlands	191
 N-20221002	 193
31.1 Why finite fields matter	193
31.2 Fields	193
31.3 Field extensions	193
31.4 Constructing finite fields	194
31.5 Important properties of finite fields	194
31.6 Polynomial factorization over finite fields	195
31.7 Example over F_3	195
31.8 Algorithms	195
31.9 Applications	196
31.10 Irreducibility over finite fields	196
31.11 Constructing extensions	196
31.12 Frobenius map	196
31.13 What survives later	197
31.14 How the pieces fit	197
31.15 Frobenius and the first Galois group	197
31.16 Irreducible polynomials and orbit size	198
31.17 Counting irreducible polynomials	198
31.18 Square-free and distinct-degree factorization	198
31.19 Berlekamp's algorithm as linear algebra	199
31.20 Coding theory connection	199
31.21 Algebraic geometry shadow	199

31.22	Frobenius structure behind finite fields	199
N-20221120		201
32.1	Definition and nonsingularity	201
32.2	Why the discriminant appears	201
32.3	Cryptographic motivation	202
32.4	Geometric addition law	202
32.5	Algebraic formula	202
32.6	Special cases and the point at infinity	203
32.7	Why associativity is difficult	203
32.8	Elliptic curves over finite fields	203
32.9	Why nonsingularity matters for the group law	204
32.10	Explicit addition formulas	204
32.11	Cryptographic intuition	204
32.12	Structural viewpoint	205
32.13	Nonsingularity and the discriminant	205
32.14	The group law as divisor arithmetic	205
32.15	Associativity and geometry	205
32.16	Cryptographic meaning	205
32.17	Elliptic curves as geometry with a group law	205
N-20251113		207
33.1	What has to be explained	207
33.2	The nonarchimedean theta-function	207
33.3	The local theta-link as a non-ring-theoretic map	208
33.4	The bridge supplied by Kummer theory	209
33.5	Two symmetries and the bookkeeping of labels	209
33.6	Conjugate synchronization and the basepoint problem	210
33.7	The log-link: recovering additive structure from multiplicative structure	210
33.8	The log-theta-lattice	211
33.9	The log-shell: a common container for incompatible links	212
33.10	Mono-anabelian reconstruction and multiradiality	212
33.11	The three indeterminacies	213
33.12	What the second note adds	214
N-20251114		215
34.1	Why the word Teichmueller matters	215
34.2	Classical Teichmueller theory: fixed topology, variable complex structure	215
34.3	Distortion is measured in one universe	216
34.4	Why the arithmetic analogy cannot be literal	217
34.5	Hodge theaters as arithmetic stages	217
34.6	Etale-like and Frobenius-like directions	218
34.7	The theta-link is the point where one universe breaks	218
34.8	The log-link does not restore naive equality	219
34.9	Classical marking versus multiradial reconstruction	219
34.10	The mathematical source of the controversy	220
34.11	A compact model of the logic	221
34.12	What this third note establishes	221

From Diophantine Inequalities to Arithmetic Deformation

N-20211007

§29.1 The question is an inequality, not a slogan

Fesenko’s survey presents inter-universal Teichmueller theory under another name: arithmetic deformation theory. This name is useful because it puts the reader’s attention on the mathematical target before the new language appears. The target is not first of all a mystical “inter-universal” philosophy. It is a family of Diophantine inequalities in which one wants to control the arithmetic size of a curve, an elliptic curve, or a rational point.

The basic shape is always the same. One side measures size: height, discriminant, or a quantity built out of bad reduction. The other side measures the places where the object is allowed to be complicated: conductor, ramification, different, or a divisor of excluded points. The problem is to prove that the first kind of size cannot grow too fast compared with the second kind of size.

Definition 29.1.1 (Number field and absolute Galois group). A number field is a finite extension $K/\mathbb{Q}$. If K^{alg} is an algebraic closure of K , its absolute Galois group is

$$G_K = \text{Gal}(K^{\text{alg}}/K),$$

regarded as a profinite topological group.

The appearance of G_K is not decoration. A large part of the story leading to IUT asks how much arithmetic geometry can be recovered from groups of symmetries. In ordinary algebraic geometry one starts with equations and then studies automorphisms or fundamental groups. Anabelian geometry tries to reverse this direction: from a suitable fundamental group, recover the geometric or arithmetic object.

§29.2 A first chain of rigidity results

The first rigidity theorem in this note is about number fields alone.

Theorem 29.2.1 (Neukirch–Ikedda–Uchida, schematic form)

Let K_1 and K_2 be number fields. An isomorphism of topological groups

$$G_{K_1} \cong G_{K_2}$$

comes from a unique field-theoretic identification of the corresponding algebraic closures, in a way compatible with the two number fields. Thus, for number fields, the absolute Galois group remembers the field.

This theorem should be read as a warning against the naive view that a Galois group is only an auxiliary invariant. In the number-field case, the group is so rigid that the field can be reconstructed from it.

The next theorem connects curves with a surprisingly small amount of ramification.

Theorem 29.2.2 (Belyi)

Let C be an irreducible smooth projective algebraic curve over a field of characteristic zero. Then C can be defined over $\overline{\mathbb{Q}}$ if and only if there exists a finite morphism

$$C \longrightarrow \mathbb{P}^1$$

which is ramified over at most three points of $\mathbb{P}^1$.

Such maps are called Belyi maps. For this note, the important point is not the proof but the direction of thought. Curves over number fields, coverings of $\mathbb{P}^1 \setminus \{0, 1, \infty\}$, and Galois/fundamental-group data are not separate topics. They are different faces of the same arithmetic geometry.

Remark 29.2.3 (Why this belongs in the IUT entrance) — Fesenko emphasizes that versions of Belyi maps are used later in arithmetic deformation theory and its application to Diophantine geometry. This is why the story does not begin directly with theta-links. Before one reaches theta-links, one has already entered a world where curves, number fields, and fundamental groups are designed to remember one another.

§29.3 From Mordell finiteness to effective inequalities

Faltings' theorem gives a finiteness result.

Theorem 29.3.1 (Faltings–Mordell)

If C is a smooth projective curve of genus > 1 over a number field K , then $C(K)$ is finite.

Finiteness is a qualitative statement. The effective Mordell problem asks for more: not only that rational points are finite, but that their heights can be bounded in terms of computable arithmetic data. In the background of Fesenko's survey, this effective direction sits beside several closely related conjectures:

effective Mordell, Szpiro, *abc*, Frey, Vojta for hyperbolic curves.

These are not identical statements, but Fesenko presents them as a closely connected cluster. The important organizing idea is that they all seek an inequality strong enough to control Diophantine complexity.

§29.4 The Szpiro inequality as a concrete target

The Szpiro conjecture is a good first model because it is phrased directly in terms of elliptic curves and bad reduction. Let E_K be an elliptic curve over a number field K . In the split multiplicative situation, the bad fibers of the minimal regular proper model have type I_{n_v} . Here v runs over nonarchimedean valuations of K , $k(v)$ is the residue field, and n_v is the number of components in the bad fiber.

Definition 29.4.1 (Conductor-size and discriminant-size, split multiplicative model). In the split multiplicative case, the logarithmic discriminant-size and conductor-size are, schematically,

$$\sum_{v \text{ bad}} n_v \log |k(v)|, \quad \sum_{v \text{ bad}} \log |k(v)|.$$

The first expression counts bad fibers with multiplicity n_v . The second remembers only the bad places themselves.

The conjectural Szpiro inequality then has the form

$$\sum_{v \text{ bad}} n_v \log |k(v)| \leq c + (6 + \epsilon) \sum_{v \text{ bad}} \log |k(v)|,$$

for every $\epsilon > 0$, with a constant c depending on K and ϵ . Equivalently, using the absolute norms of the minimal discriminant and the conductor, one writes

$$N(\text{Disc } E_K) \leq c' N(\text{Cond } E_K)^{6+\epsilon}.$$

Remark 29.4.2 (The meaning of the exponent) — The number 6 is not a random constant inserted into the conjecture. In Fesenko’s discussion it is connected with the geometric Szpiro-type inequality for elliptic surfaces and with the comparison between the arithmetic and geometric sides of the problem. For the present note, the safe takeaway is that Szpiro is an inequality bounding discriminant-size by conductor-size.

A very simple reading is this: the conductor records where the elliptic curve behaves badly; the discriminant records how badly it behaves there. Szpiro predicts that “how badly” cannot be much larger than a fixed power of “where badly”.

§29.5 Vojta’s inequality and the hyperbolic curve viewpoint

The Vojta form is closer to the geometry of punctured curves.

Definition 29.5.1 (Hyperbolic curve). Over a field of characteristic zero, a hyperbolic curve is a smooth projective geometrically connected curve of genus g with r points removed, such that

$$2 - 2g - r < 0.$$

Examples include $\mathbb{P}^1 \setminus \{0, 1, \infty\}$ and an elliptic curve with the origin removed.

The condition is not merely topological decoration. Hyperbolic curves have nonabelian algebraic fundamental groups, and these groups are the kind of objects anabelian geometry can hope to use for reconstruction.

Let C be a smooth proper geometrically connected curve over a number field K , and let D be a reduced effective divisor on C such that $C \setminus D$ is hyperbolic. Vojta-type inequalities compare a height attached to $\omega_C(D)$ with terms measuring ramification and contact with D :

$$h_{\omega_C(D)}(x) \leq c + (1 + \epsilon)(\log\text{-diff}_C(x) + \log\text{-cond}_D(x)).$$

Here $\log\text{-diff}$ measures the arithmetic complexity of the field of definition of x , while $\log\text{-cond}$ measures intersection with the divisor D .

Remark 29.5.2 (Why Belyi maps reappear) — Fesenko notes that Belyi maps reduce the Vojta conjecture for such (C, D) to the special case

$$C = \mathbb{P}^1, \quad D = [0] + [1] + [\infty].$$

Thus the same three-punctured projective line that appears in Belyi's theorem also acts as a universal testing ground for the Diophantine inequality.

§29.6 Arithmetic fundamental groups

For a geometrically integral scheme X over a perfect field K , there is a fundamental exact sequence

$$1 \longrightarrow \pi_1^{\text{geom}}(X) \longrightarrow \pi_1(X) \longrightarrow G_K \longrightarrow 1,$$

where

$$\pi_1^{\text{geom}}(X) = \pi_1(X \times_K K^{\text{alg}}).$$

For schemes over number fields or related fields, $\pi_1(X)$ is often called an arithmetic fundamental group.

Definition 29.6.1 (Anabelian reconstruction, guiding form). Anabelian reconstruction is the idea that, for suitable schemes X , one can recover the arithmetic or geometric object from its algebraic fundamental group together with the projection

$$\pi_1(X) \longrightarrow G_K.$$

The slogan is not that every space is determined by its fundamental group, but that certain highly nonabelian arithmetic schemes are rigid enough for this to happen.

In the IUT entrance, this matters because the later theory needs structures that survive when ordinary ring-theoretic comparison is no longer available. Fesenko summarizes the transition as follows: arithmetic deformation theory is related to anabelian geometry, but uses a different package of tools, including mono-anabelian geometry, nonarchimedean theta-functions, monoid-theoretic categories, and deconstruction/reconstruction of ring structures.

§29.7 The local elliptic curve model: Tate curves and q -parameters

The first truly arithmetic-deformation object in the survey is not the abc equation itself. It is an elliptic curve with split multiplicative reduction and its local q -parameters.

Let E_F be an elliptic curve over a number field F . Suppose that, at a bad nonarchimedean valuation v , the curve has split multiplicative reduction. Let F_v be the completion of F at v . Then locally one has a Tate uniformization

$$E_F(F_v) \cong F_v^\times / \langle q_v \rangle,$$

where $q_v \in F_v^\times$ is the local q -parameter and $\langle q_v \rangle$ is the cyclic subgroup generated by q_v .

Definition 29.7.1 (The idele q_{E_F}). Define an idele $q_{E_F} \in \mathbb{A}_F^\times$ by setting its component to be

$$(q_{E_F})_v = \begin{cases} q_v, & v \text{ a bad split multiplicative valuation,} \\ 1, & v \text{ archimedean or good reduction.} \end{cases}$$

after choosing the local q_v at each bad valuation.

The valuation of q_v records the number n_v of components of the corresponding bad fiber. Therefore the degree of this idele is not an artificial invariant: it packages the left-hand side of the Szpiro inequality.

§29.8 Normalised degree as log-volume

Fesenko uses ideles to define a normalised degree. Let k be a number field and $\mathbb{A}_k^\times$ its idele group. For $\alpha \in \mathbb{A}_k^\times$, the non-normalised degree is

$$\deg_k(\alpha) = -\log |\alpha|,$$

where $|\alpha|$ is the module by which multiplication by α scales Haar measure on the adèle ring. The normalised degree is

$$\deg(\alpha) = \frac{1}{[k : \mathbb{Q}]} \deg_k(\alpha).$$

This normalisation is stable under finite field extension, which is why the same symbol $\deg$ can be used after passing to larger number fields.

Remark 29.8.1 (Why log-volume enters) — Because of the minus sign, $-\deg(\alpha)$ is naturally interpreted as a log-volume. In the later IUT inequality, the left side is written in terms of $-\deg q_E$. This is why the problem of bounding Szpiro-type data becomes a problem about bounding log-volumes.

For the elliptic curve above, Fesenko identifies the goal as giving a suitable upper bound on

$$\deg q_{E_F}.$$

This is the cleanest first target of arithmetic deformation theory in the survey.

§29.9 Why ordinary scheme-theoretic geometry is not enough

The next move is the point where the theory stops looking like ordinary algebraic geometry. Fix a prime $l > 3$, relatively prime to the bad valuations and to the relevant local valuation numbers. For a bad valuation v , Fesenko records a monoid-theoretic transformation on the submonoid of $F_v^\times$ generated by units and q_v :

$$q_v \mapsto q_v^{m^2}, \quad u \mapsto u \quad (u \in \mathcal{O}_{F_v}^\times),$$

where

$$1 \leq m \leq \frac{l-1}{2}.$$

The element $q_v^{m^2}$ is later viewed as a special value of a nonarchimedean theta-function.

Proposition 29.9.1 (The first deformation warning)

The transformation

$$q_v \mapsto q_v^{m^2}, \quad u \mapsto u$$


should be read as monoid-theoretic rather than ring-theoretic. It is not a morphism of schemes and does not preserve the full ring structure.

Proof. The point is conceptual rather than computational. A ring morphism must respect both multiplication and addition. The displayed transformation is defined on multiplicative data: the local units and the generator q_v of the Tate period subgroup. It is designed to manipulate multiplicative/theta data, not to preserve addition in F_v . Thus it belongs to the monoid-theoretic side of the theory. $\square$

This is the first reason the word “deformation” is appropriate. The ring structure is not transported intact. It is partly dismantled so that certain multiplicative and theta-theoretic information can be moved, after which one needs reconstruction tools.

§29.10 The route from deformation to inequality

The full theta-link belongs to later notes. But the route can already be drawn.



The diagram is not a proof. It is a navigation map. Starting from Diophantine inequalities, one passes to elliptic curves and their bad reduction. The bad reduction is encoded by q_E . IUT then introduces a deformation machine involving arithmetic fundamental groups, mono-anabelian reconstruction, theta-links, log-links, and log-theta-lattices. The intended output is a log-volume inequality which returns to the original Diophantine problem.

Theorem 29.10.1 (Main IUT inequality, schematic form)

In Fesenko’s account, the main theorem of IUT gives a bound on log-volumes of the form

$$-\deg q_E \leq -\deg \Theta_E,$$

where the right-hand side is the log-volume of a union of possible images of theta-data after applying the theta-link, subject to specified indeterminacies.

The statement above is deliberately schematic. It records the mathematical role of the theorem, not the construction of the objects. The important point for this first note is that the theorem does not directly say “abc is true”. It gives a bound on a deformation volume. A further computation of the right-hand side is then used to obtain a bound of the desired Diophantine kind.

Remark 29.10.2 (The bridge to later notes) — The words theta-link, log-link, log-theta-lattice, indeterminacy, and multiradiality all belong to the mechanism behind Θ_E . They should not be treated as decorative terminology. They are the devices by which the theory tries to compare arithmetic data after the ordinary ring structure has been deconstructed.

§29.11 What this first note establishes

The first entrance to IUT should therefore be remembered through four questions.

1. What is being bounded? The initial concrete target is $\deg q_E$, which packages the bad-reduction side of a Szpiro-type inequality.
2. What is rigid? Absolute Galois groups and arithmetic fundamental groups can, in suitable settings, remember number fields or curves.
3. What is allowed to vary? The ordinary ring structure is no longer transported as a single indivisible object; multiplicative and additive aspects are separated.
4. Why is advanced machinery needed? Once ordinary scheme morphisms are not enough, one needs mono-anabelian reconstruction, Kummer-type methods, nonarchimedean theta-functions, and log-volume control.

This is already a mathematical picture, but not yet the full theory. The next note should begin where this one stops: with the moment when the map involving $q_v^{m^2}$ becomes part of the theta-link, and when the lost additive structure has to be handled through log-links and log-shells.

From Algebraic Number Theory to the First Shadow of Langlands

N-20220123

§30.1 Why this note starts with algebraic number theory

This note was written as preparation for a seminar on the Langlands program. Since the seminar was expected to involve algebraic number theory, I first collect the basic vocabulary: number fields, algebraic integers, ideals, local fields, modular forms, L -functions, and finally the first shadow of the Langlands correspondence.

Remark 30.1.1 — This is not a proof of the Langlands program. It is a personal attempt to see why many different objects—number fields, ideals, local completions, modular forms, Galois representations, and automorphic representations—might belong to one story.

§30.2 Number fields and algebraic integers

Definition 30.2.1. A number field is a finite extension $K/\mathbb{Q}$. Equivalently, K is a field that is a finite-dimensional vector space over $\mathbb{Q}$. A basic example is a quadratic field

$$K = \mathbb{Q}(\sqrt{d}),$$

where d is a nonsquare integer.

Definition 30.2.2. Let K be a number field. An element $\alpha \in K$ is an algebraic integer if it satisfies a monic polynomial with integer coefficients:

$$x^n + a_{n-1}x^{n-1} + \cdots + a_0 = 0, \quad a_i \in \mathbb{Z}.$$

The algebraic integers in K form a ring, denoted $\mathcal{O}_K$.

The reason for introducing algebraic integers is that they play in K the role played by ordinary integers in $\mathbb{Q}$. Unique factorization of elements may fail, but that failure leads to a deeper replacement: factorization of ideals.

§30.3 Ideals and factorization

Definition 30.3.1. An ideal $\mathfrak{a} \subseteq \mathcal{O}_K$ is an additive subgroup such that for every $\alpha \in \mathfrak{a}$ and every $\gamma \in \mathcal{O}_K$,

$$\gamma\alpha \in \mathfrak{a}.$$

Every nonzero ideal of $\mathcal{O}_K$ factors uniquely into prime ideals. Thus even when element factorization fails, ideal factorization remains well behaved.

Remark 30.3.2 — This is the first conceptual jump: instead of forcing unique factorization of elements, one changes the objects being factored. Ideals see the arithmetic more cleanly.

§30.4 Quadratic fields and class number

For $K = \mathbb{Q}(\sqrt{d})$, the ring of integers is

$$\mathcal{O}_K = \begin{cases} \mathbb{Z}[\sqrt{d}], & d \equiv 2, 3 \pmod{4}, \\ \mathbb{Z}\left[\frac{1+\sqrt{d}}{2}\right], & d \equiv 1 \pmod{4}. \end{cases}$$

Definition 30.4.1. The class group of K is the group of nonzero fractional ideals modulo principal fractional ideals. Its order is the class number h_K .

If $h_K = 1$, then every ideal class is principal. If $h_K > 1$, nonprincipal ideal classes exist, and element factorization can fail.

Example 30.4.2

For

$$K = \mathbb{Q}(\sqrt{-5}), \quad \mathcal{O}_K = \mathbb{Z}[\sqrt{-5}],$$

one has $h_K = 2$. The familiar failure of unique factorization is

$$6 = 2 \cdot 3 = (1 + \sqrt{-5})(1 - \sqrt{-5}).$$

Caution

The original calculation of the ideal factorization of (6) is too compressed and drops one conjugate prime. A corrected form is

$$(2) = \mathfrak{p}_2^2, \quad (3) = \mathfrak{p}_3\mathfrak{p}_3',$$

where

$$\mathfrak{p}_2 = (2, 1 + \sqrt{-5}), \quad \mathfrak{p}_3 = (3, 1 + \sqrt{-5}), \quad \mathfrak{p}_3' = (3, 1 - \sqrt{-5}).$$

Thus

$$(6) = \mathfrak{p}_2^2\mathfrak{p}_3\mathfrak{p}_3'.$$

§30.5 Localization, completion, and p -adic numbers

A recurring idea in number theory is the comparison between local and global information. Fix a prime p . Completing $\mathbb{Q}$ under the p -adic metric gives

$$\mathbb{Q}_p.$$

For a number field K , one can complete at a prime ideal $\mathfrak{p}$, producing a local field $K_{\mathfrak{p}}$.

Remark 30.5.1 — The local–global theme is one of the first hints of Langlands. Global arithmetic objects often decompose into local pieces, and the Langlands program repeatedly asks how local and global data determine each other.

Definition 30.5.2 (Hensel’s lemma, informal version). Hensel’s lemma says that, under suitable nondegeneracy hypotheses, a root of a polynomial equation modulo p can be lifted to a root over the p -adic integers or p -adic field.

This resembles Newton iteration: an approximate solution is improved step by step in a complete space.

Example 30.5.3 (A small Hensel lifting computation)

Consider

$$f(x) = x^2 - 2$$

over the 7-adic integers. Modulo 7, one has

$$3^2 - 2 = 7 \equiv 0 \pmod{7},$$

so $x_0 = 3$ is a root mod 7. Also

$$f'(x) = 2x, \quad f'(3) = 6 \not\equiv 0 \pmod{7}.$$

Hensel’s lemma therefore says that this root lifts uniquely to a 7-adic root.

To lift from mod 7 to mod 49, write

$$x = 3 + 7t.$$

Then

$$f(3 + 7t) = (3 + 7t)^2 - 2 = 7 + 42t + 49t^2 \equiv 7(1 + 6t) \pmod{49}.$$

We need $1 + 6t \equiv 0 \pmod{7}$. Since $6^{-1} \equiv 6 \pmod{7}$, this gives

$$t \equiv -6 \equiv 1 \pmod{7}.$$

Thus the lifted root modulo 49 is

$$x_1 = 3 + 7 = 10, \quad 10^2 - 2 = 98 \equiv 0 \pmod{49}.$$

The computation shows the Newton-like nature of Hensel lifting: a root modulo p is corrected by one p -adic digit at a time.

In Python, the first lifting step can be written as:

```
p = 7
x = 3
# lift x from mod p to mod p^2 for f(x)=x^2-2
t = -(x*x - 2)//p * pow(2*x, -1, p) % p
x = x + p*t
print(x, (x*x - 2) % (p*p))    # 10, 0
```

§30.6 Modular forms and L -functions

Modular forms are analytic functions with strong symmetry. A typical transformation law is

$$f\left(\frac{az+b}{cz+d}\right) = (cz+d)^k f(z),$$

for matrices

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

in a suitable discrete subgroup such as $\mathrm{SL}_2(\mathbb{Z})$. The integer k is the weight.

If a modular form has Fourier expansion

$$f(z) = \sum_{n \geq 0} a(n) e^{2\pi i n z},$$

then one can attach an L -function

$$L(f, s) = \sum_{n \geq 1} \frac{a(n)}{n^s}.$$

Analytic continuation and functional equations of such functions encode arithmetic information.

Example 30.6.1 (Ramanujan's Delta function)

The modular discriminant is

$$\Delta(z) = q \prod_{n=1}^{\infty} (1 - q^n)^{24} = \sum_{n \geq 1} \tau(n) q^n, \quad q = e^{2\pi i z}.$$

It is a modular form of weight 12. The coefficients $\tau(n)$ form the Ramanujan tau function, whose growth properties were predicted by Ramanujan and later proved by Deligne.

§30.7 First understanding of Langlands

The Langlands program is often described as a unifying theory linking number theory, representation theory, geometry, and analysis. A safer first slogan is

$$\text{Galois-side symmetry} \quad \leftrightarrow \quad \text{automorphic-side symmetry},$$

with L -functions as a major bridge.

On one side are Galois representations, which encode actions of Galois groups on vector spaces. On the other side are automorphic representations, which organize modular forms and their higher-dimensional analogues.

Caution

It is too imprecise to say simply that an automorphic representation is a homomorphism from a group to a general linear group. That phrase is a useful beginner's memory of representation theory, but automorphic representations are representations of adelic groups on spaces of functions. For this note, the essential idea is that representation theory becomes a language for arithmetic.

§30.8 Local and global correspondences

Local Langlands studies local fields such as $\mathbb{Q}_p$ and finite extensions. Global Langlands studies number fields and their adèle rings. The local theory is more complete in many settings, while the global theory contains many deep conjectures.

Caution

A naive statement such as

$$\{\rho : G_F \rightarrow \mathrm{GL}_n(\mathbb{C})\} \longleftrightarrow \{\text{automorphic representations}\}$$

is only a shadow of the real statement. For GL_n , local Langlands relates suitable Langlands parameters to irreducible admissible representations of $\mathrm{GL}_n(F)$.

§30.9 Cross-field intuitions

The original train of thought then branches into several speculative but useful directions.

§30.9.i Category theory

Category theory offers a language for comparing structures through functors and natural transformations. It is tempting to imagine Langlands correspondences as functorial relationships or equivalences between categories.

Gap

This is currently only a guiding intuition. A rigorous categorical formulation must specify the categories, functors, and compatibility conditions. Without those, it remains a helpful but informal picture.

§30.9.ii Algebraic topology

The idea of turning arithmetic into topology is not absurd: cohomology, sheaves, stacks, and derived geometry are everywhere in modern number theory. But the naive idea of treating ideals as cells of a CW complex needs an actual construction before it becomes mathematics rather than metaphor.

§30.9.iii Linear algebra and representation theory

Representation theory can be seen as the extension of linear algebra into group theory and symmetry. Eigenvalues, traces, matrices, and vector spaces occur constantly in Galois representations and automorphic forms.

§30.10 Numerical experiments

Computational experiments cannot prove Langlands statements, but they can build intuition. Natural experiments include:

- (i) computing initial values of the Ramanujan tau function;
- (ii) plotting the oscillation of modular-form coefficients;

- (iii) implementing a Hensel-lifting example;
- (iv) testing simple finite-dimensional representations in toy cases.

The first object to compute is Ramanujan's discriminant modular form

$$\Delta(z) = q \prod_{n=1}^{\infty} (1 - q^n)^{24} = \sum_{n=1}^{\infty} \tau(n) q^n, \quad q = e^{2\pi iz}.$$

Its first coefficients are

$$\tau(1) = 1, \quad \tau(2) = -24, \quad \tau(3) = 252, \quad \tau(4) = -1472, \quad \tau(5) = 4830, \quad \tau(6) = -6048.$$

These numbers are already arithmetic, but one should not match them with an arbitrary elliptic curve.

Now take the elliptic curve

$$E : y^2 + y = x^3 - x^2,$$

Cremona label 11a3, of conductor 11. For a prime $p \neq 11$, define

$$a_p = p + 1 - \#E(\mathbb{F}_p).$$

A direct enumeration gives

p	$\#E(\mathbb{F}_p)$	a_p
2	5	-2
3	5	-1
5	5	1
7	10	-2

For example, over $\mathbb{F}_5$ the affine solutions are only

$$(0, 0), (0, 4), (1, 0), (1, 4),$$

together with the point at infinity, so $\#E(\mathbb{F}_5) = 5$ and $a_5 = 1$.

The important correction is that

$$\tau(2), \tau(3), \tau(5), \tau(7) = -24, 252, 4830, -16744$$

do *not* equal

$$a_2, a_3, a_5, a_7 = -2, -1, 1, -2.$$

This mismatch is not a failure of modularity; it is a failure to choose the matching modular form. Δ is a weight-12, level-1 modular form. The elliptic curve of conductor 11 corresponds instead to a weight-2, level-11 newform whose p -th coefficient is a_p . The concrete Langlands-style lesson is therefore sharper than a slogan: once the level, weight, and normalization are correct, Fourier coefficients on the automorphic side match Frobenius traces or point counts on the Galois side.

A short SageMath check is enough to reproduce the finite-field side:

```
E = EllipticCurve([0,-1,1,0,0])    # y^2 + y = x^3 - x^2
for p in [2,3,5,7,13,17,19]:
    Np = E.change_ring(GF(p)).cardinality()
    print(p, p + 1 - Np, E.ap(p))
```

For the Δ -function coefficients one can use a truncated product:

```
R.<q> = PowerSeriesRing(ZZ, default_prec=12)
Delta = q * prod((1 - q^n)^24 for n in range(1,12))
print(Delta)
```

§30.11 Topics to learn next

Several missing foundations are clear: complex analysis for analytic continuation and functional equations of L -functions; measure theory and integration for automorphic forms and harmonic analysis; Galois theory as the classical source of arithmetic symmetry; Sato–Tate phenomena and equidistribution; and random matrix heuristics for zeros of L -functions.

Question

Why should modular forms and Galois representations, which look completely different at first, speak a common language? This question is not answered here, but it is the right surprise to keep.

§30.12 Synthesis

The route of the note is

number fields $\rightarrow$ ideals $\rightarrow$ local fields $\rightarrow$ modular forms and L -functions $\rightarrow$ Langlands.

Some formulations in the original note were too optimistic or imprecise, but the underlying instinct is good: Langlands is not one isolated theorem, but a web of correspondences linking arithmetic, analysis, geometry, and representation theory.

§30.13 Local-global intuition

A recurring theme in algebraic number theory is to understand a global object by all of its local shadows. For a number field K , the local fields K_v arise by completing K at its archimedean and nonarchimedean places. The rational field gives familiar examples:

$$\mathbb{Q} \hookrightarrow \mathbb{R}, \quad \mathbb{Q} \hookrightarrow \mathbb{Q}_p.$$

The real completion measures ordinary size; the p -adic completion measures divisibility by p . The collection of all completions records many different ways a number can be close to zero.

This local-global idea is one reason adeles and ideles appear. They package all completions at once, allowing arithmetic problems to be studied through harmonic analysis on locally compact groups.

§30.14 Why L -functions enter

An L -function is a generating function that encodes arithmetic data. For example, the Riemann zeta function

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}$$

has the Euler product

$$\zeta(s) = \prod_p \frac{1}{1 - p^{-s}}.$$

The sum sees positive integers; the product sees primes. This equality is the prototype for the idea that analytic behavior of a function can reveal arithmetic structure.

In Langlands-style thinking, Galois representations and automorphic forms have associated L -functions. A correspondence is often tested by matching these L -functions.

§30.15 A first Langlands slogan

A very rough slogan is:

$$\text{Galois representations} \quad \leftrightarrow \quad \text{automorphic representations.}$$

The left side is arithmetic and studies symmetries of field extensions. The right side is analytic and representation-theoretic, involving functions on groups over adeles. The Langlands program predicts a deep dictionary between them.

The modularity theorem for elliptic curves over $\mathbb{Q}$ is a famous instance: arithmetic information from an elliptic curve matches a modular form. This is why modular forms enter a note that begins with number fields and ideals.

§30.16 Structural viewpoint

The note is a map, not a proof. Number fields generalize $\mathbb{Q}$; algebraic integers generalize $\mathbb{Z}$; ideals repair failure of unique factorization; local fields isolate one prime at a time; adeles recombine all primes; modular forms and automorphic representations provide analytic symmetry; L -functions act as fingerprints. Langlands is the claim that these fingerprints match across arithmetic and analysis.

§30.17 Ideals as repaired unique factorization

The central obstruction in algebraic number theory is that elements may fail to factor uniquely. Ideals repair this: in the ring of integers $\mathcal{O}_K$ of a number field, nonzero ideals factor uniquely into prime ideals,

$$\mathfrak{a} = \prod_i \mathfrak{p}_i^{e_i}.$$

The class group measures the gap between principal ideals and all ideals. Class number one means ideal factorization descends all the way back to element factorization.

§30.18 Local fields and arithmetic microscopes

Completion at a prime ideal gives a local field $K_{\mathfrak{p}}$. This is an arithmetic microscope: it sees divisibility and congruence near one prime at a time. Global statements are often proved by comparing all local completions.

The product formula and local-global principles explain why p -adic numbers naturally appear next to number fields.

§30.19 Modular forms and Galois representations

A modular form has Fourier expansion

$$f(q) = \sum_{n \geq 0} a_n q^n.$$

The coefficients a_p often encode arithmetic information. In the Langlands viewpoint, modular forms correspond to Galois representations, so analytic objects and symmetry actions on algebraic numbers describe the same data.

§30.20 Arithmetic symmetry as the bridge to Langlands

The path from algebraic number theory to Langlands is a passage from factorization to symmetry. Ideals repair arithmetic, local fields isolate primes, L -functions package global data, and Langlands predicts that arithmetic Galois symmetries are mirrored by automorphic analytic objects.

Finite Fields and Polynomial Factorization

N-20221002

§31.1 Why finite fields matter

This note is a more systematic essay on finite fields and factorization. The central question is: how do polynomials factor when the coefficients live in a finite field rather than in $\mathbb{Q}$, $\mathbb{R}$, or $\mathbb{C}$?

This question is not only theoretical. It appears in coding theory, cryptography, computational algebra, and the construction of algebraic extensions. The note correctly starts from the definition of a field, because factorization depends heavily on the coefficient domain.

§31.2 Fields

A field F is a set with addition and multiplication such that:

- (i) $(F, +)$ is an abelian group with identity 0;
- (ii) $F \setminus \{0\}$ is an abelian group under multiplication with identity 1;
- (iii) multiplication distributes over addition.

The exclusion of 0 from the multiplicative group is essential. If 0 had a multiplicative inverse, then from $0 \cdot a = 0$ one would get contradictions such as $1 = 0$. So the phrase “every nonzero element has an inverse” is not a technical nuisance; it is what makes division possible.

Examples are

$$\mathbb{Q}, \quad \mathbb{R}, \quad \mathbb{C},$$

and, for a prime p ,

$$\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}.$$

The requirement that p be prime is crucial. If p is composite, zero divisors appear. For example in $\mathbb{Z}/6\mathbb{Z}$,

$$2 \cdot 3 \equiv 0 \pmod{6},$$

even though neither factor is zero.

§31.3 Field extensions

A field extension K/F means F is embedded in a larger field K . The point of passing to K is often to make more polynomial equations solvable. For example,

$$x^2 + 1$$

has no root in $\mathbb{R}$, but it splits in $\mathbb{C}$:

$$x^2 + 1 = (x - i)(x + i).$$

If K is finite-dimensional as a vector space over F , its dimension is the degree of the extension:

$$[K : F] = \dim_F K.$$

The note's intuition is good: finite extensions are tractable because linear algebra can be used. Infinite extensions require much heavier tools.

§31.4 Constructing finite fields

For a prime p , $\mathbb{F}_p$ is the simplest finite field. To construct fields with p^n elements, use irreducible polynomials. If

$$f(x) \in \mathbb{F}_p[x]$$

is irreducible of degree n , then

$$\mathbb{F}_p[x]/(f(x))$$

is a field with p^n elements.

For example, over $\mathbb{F}_2$, the polynomial

$$f(x) = x^3 + x + 1$$

has no root in $\mathbb{F}_2$, hence is irreducible of degree 3. Therefore

$$\mathbb{F}_2[x]/(x^3 + x + 1)$$

is a field with $2^3 = 8$ elements.

In this quotient field, if α denotes the class of x , then

$$\alpha^3 + \alpha + 1 = 0,$$

so

$$\alpha^3 = \alpha + 1$$

in characteristic 2. Every element can be uniquely written as

$$a + b\alpha + c\alpha^2, \quad a, b, c \in \mathbb{F}_2.$$

§31.5 Important properties of finite fields

The standard facts are:

- (i) finite fields have order p^n ;
- (ii) for every prime power p^n , there is a finite field of that order, unique up to isomorphism;
- (iii) the nonzero elements $\mathbb{F}_{p^n}^\times$ form a cyclic group of order $p^n - 1$;
- (iv) the Frobenius map

$$\text{Frob}(a) = a^p$$

is an automorphism of $\mathbb{F}_{p^n}$.

The cyclicity of the multiplicative group is surprisingly powerful. It means every nonzero element is a power of one primitive element. This is why finite field computations often become exponent arithmetic modulo $p^n - 1$.

§31.6 Polynomial factorization over finite fields

In $\mathbb{F}_q[x]$, every nonzero polynomial factors uniquely into irreducible polynomials, up to order and multiplication by nonzero constants. This is analogous to unique prime factorization of integers.

However, a polynomial need not split into linear factors over its base finite field. For example, a polynomial may be irreducible over $\mathbb{F}_p$ but split over an extension field. This is exactly like $x^2 + 1$ over $\mathbb{R}$, except that the extension fields are finite and highly structured.

§31.7 Example over $\mathbb{F}_3$

Consider

$$f(x) = x^4 + x^2 + 2 \in \mathbb{F}_3[x].$$

First test roots:

$$f(0) = 2 \neq 0,$$

$$f(1) = 1 + 1 + 2 = 4 \equiv 1 \pmod{3},$$

$$f(2) = 2^4 + 2^2 + 2 = 16 + 4 + 2 = 22 \equiv 1 \pmod{3}.$$

So f has no linear factor. It might still factor as a product of two quadratics. A systematic check over $\mathbb{F}_3$ shows that no such factorization exists, so f is irreducible over $\mathbb{F}_3$.

This illustrates the basic workflow: root testing only detects linear factors. For degree 4, one must also check quadratic factors.

§31.8 Algorithms

The note lists several factorization algorithms:

- (i) trial division, useful for very small degrees;
- (ii) Berlekamp's algorithm, which converts factorization into linear algebra over finite fields;
- (iii) Cantor–Zassenhaus, a randomized algorithm widely used in computer algebra systems.

The high-level idea behind these algorithms is that finite fields make the Frobenius map computable. Since every element of $\mathbb{F}_q$ satisfies

$$a^q = a,$$

the polynomial

$$x^q - x$$

contains all elements of $\mathbb{F}_q$ as roots. This identity becomes a machine for detecting factors.

§31.9 Applications

Finite-field factorization appears in:

- (i) error-correcting codes;
- (ii) cryptographic protocols;
- (iii) construction of extension fields;
- (iv) computational number theory;
- (v) symbolic algebra systems.

The note's final reflection is that finite fields are small enough for computation but structured enough to retain deep algebraic behavior. That is why they are so central.

§31.10 Irreducibility over finite fields

A polynomial can factor differently depending on the coefficient field. For example, $x^2 + 1$ is irreducible over $\mathbb{R}$ only if one insists on real linear factors, but over $\mathbb{C}$ it splits. Over finite fields the same phenomenon becomes arithmetic.

For $\mathbb{F}_p$, a quadratic polynomial

$$ax^2 + bx + c$$

factors iff its discriminant

$$\Delta = b^2 - 4ac$$

is a square in $\mathbb{F}_p$. Thus irreducibility is controlled by quadratic residues. This links polynomial factorization directly to modular arithmetic.

§31.11 Constructing extensions

If $f(x) \in F[x]$ is irreducible of degree n , then

$$F[x]/(f(x))$$

is a field extension of F of degree n . The class of x becomes an element α satisfying $f(\alpha) = 0$. This is the algebraic meaning of “adjoining a root.”

For finite fields, this construction produces fields with p^n elements:

$$\mathbb{F}_{p^n} \cong \mathbb{F}_p[x]/(f(x))$$

for any irreducible f of degree n . Every finite field has prime-power size, and for each prime power there is a unique finite field up to isomorphism.

§31.12 Frobenius map

In characteristic p , the Frobenius map is

$$\varphi(a) = a^p.$$

It is a field homomorphism because

$$(a + b)^p = a^p + b^p$$

in characteristic p . In a finite field, Frobenius is an automorphism. Its fixed field inside $\mathbb{F}_{p^n}$ is $\mathbb{F}_p$.

The polynomial

$$x^{p^n} - x$$

has all elements of $\mathbb{F}_{p^n}$ as roots, and it factors over $\mathbb{F}_p$ as the product of all monic irreducible polynomials whose degrees divide n . This is a central organizing fact for finite-field factorization.

§31.13 What survives later

Finite fields turn algebra into arithmetic. Roots of polynomials become elements of finite extensions, irreducibility becomes a modular phenomenon, and Frobenius gives a canonical symmetry. This is why finite fields appear in coding theory and cryptography: their algebra is rich but completely finite and computable.

§31.14 How the pieces fit

A finite field is where group theory, ring theory, linear algebra, and number theory meet. The additive group of $\mathbb{F}_{p^n}$ is an n -dimensional vector space over $\mathbb{F}_p$. The nonzero elements form a cyclic multiplicative group of order $p^n - 1$. Polynomial quotients construct the field, and Frobenius supplies a canonical symmetry.

Thus a finite field computation often has three simultaneous readings:

operation	structure being used
addition	$\mathbb{F}_p$ -linear algebra
multiplication	cyclic unit group and polynomial reduction
Frobenius $x \mapsto x^p$	Galois symmetry

§31.15 Frobenius and the first Galois group

In characteristic p , the Frobenius map

$$F : x \mapsto x^p$$

is a field homomorphism because

$$(x + y)^p = x^p + y^p.$$

In $\mathbb{F}_{p^n}$, it is an automorphism, and

$$F^n = 1.$$

The Galois group is cyclic:

$$\text{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p) = \langle F \rangle \cong \mathbb{Z}/n\mathbb{Z}.$$

This is the cleanest first example of a Galois group: roots are moved by raising to powers of p , and the fixed field of Frobenius is $\mathbb{F}_p$.

§31.16 Irreducible polynomials and orbit size

An element $\alpha \in \mathbb{F}_{q^n}$ has conjugates

$$\alpha, \quad \alpha^q, \quad \alpha^{q^2}, \quad \dots$$

under Frobenius. Its minimal polynomial over $\mathbb{F}_q$ is

$$\prod_{i=0}^{d-1} (x - \alpha^{q^i}),$$

where d is the size of the Frobenius orbit. Thus irreducible polynomials over finite fields are Frobenius orbits in disguise.

This explains why the degrees of irreducible factors of a polynomial are controlled by Frobenius cycles.

§31.17 Counting irreducible polynomials

Every element of $\mathbb{F}_{q^n}$ is a root of

$$x^{q^n} - x.$$

Over $\mathbb{F}_q$, this polynomial is the product of all monic irreducible polynomials whose degrees divide n . If N_d is the number of monic irreducible polynomials of degree d , then

$$q^n = \sum_{d|n} dN_d.$$

By Möbius inversion,

$$N_n = \frac{1}{n} \sum_{d|n} \mu(d) q^{n/d}.$$

This formula turns factorization into enumeration. The factor d appears because each degree- d irreducible polynomial contributes d roots.

§31.18 Square-free and distinct-degree factorization

Finite-field factorization algorithms exploit Frobenius. The polynomial

$$x^{q^m} - x$$

has as roots exactly the elements of $\mathbb{F}_{q^m}$. Therefore

$$\gcd(f(x), x^{q^m} - x)$$

extracts from f the product of irreducible factors whose degrees divide m .

This is the idea behind distinct-degree factorization. First remove repeated factors by checking $\gcd(f, f')$. Then use Frobenius powers to separate factors by degree. Finally use algorithms such as Cantor–Zassenhaus to split equal-degree parts.

§31.19 Berlekamp's algorithm as linear algebra

Berlekamp's algorithm turns factorization into a nullspace computation. Over $\mathbb{F}_q$, consider the map on the quotient algebra $A = \mathbb{F}_q[x]/(f)$

$$B : g \mapsto g^q - g.$$

Elements in the kernel satisfy $g^q = g$, and this kernel contains information about decomposing f into factors.

Thus the algorithm uses the finite-dimensional $\mathbb{F}_q$ -vector space structure of the quotient ring. This is a concrete instance of a general computational theme: algebraic structure becomes linear algebra after passing to a finite quotient.

§31.20 Coding theory connection

Finite fields are essential in coding theory because code symbols can be added, multiplied, and evaluated polynomially. Reed–Solomon codes evaluate polynomials over $\mathbb{F}_q$ at distinct field points:

$$f \mapsto (f(a_1), \dots, f(a_m)).$$

A nonzero polynomial of degree $< k$ has at most $k - 1$ roots, so two different low-degree polynomials cannot agree at too many evaluation points. This gives error detection and correction.

Here the same polynomial-root principle used in factorization becomes a robustness principle for information transmission.

§31.21 Algebraic geometry shadow

The identity

$$\#\mathbb{A}^1(\mathbb{F}_q) = q$$

and the polynomial $x^q - x$ already hint at a larger story: equations over finite fields define finite sets of points, and Frobenius acts on their solutions. In higher-dimensional algebraic geometry, counting points over $\mathbb{F}_{q^n}$ reveals cohomological information.

Thus elementary finite-field polynomial factorization is a first step toward modern arithmetic geometry: finite fields provide a world where algebraic equations, symmetries, and point counts interact exactly.

§31.22 Frobenius structure behind finite fields

Finite fields are not just small fields. They are finite vector spaces, cyclic multiplicative groups, quotient rings, Frobenius dynamical systems, and computational domains all at once. Irreducible polynomials build extensions; Frobenius organizes roots; gcds with $x^{q^m} - x$ detect factor degrees; and the same algebra supports coding theory and cryptography.

Elliptic Curves: Nonsingularity, Group Law, and Cryptographic Motivation

N-20221120

§32.1 Definition and nonsingularity

Definition 32.1.1 (Elliptic curve in Weierstrass form). An elliptic curve over a field is a nonsingular cubic curve with a chosen base point. In Weierstrass form it is written as

$$y^2 = x^3 + ax + b,$$

with discriminant nonzero.

An elliptic curve in short Weierstrass form is the set of solutions

$$E : y^2 = x^3 + ax + b$$

together with a point at infinity. The printed page emphasizes a crucial restriction: elliptic curves are not allowed to have double or triple roots in the cubic polynomial

$$x^3 + ax + b.$$

Geometrically, a triple root gives a cusp, while a double root gives a self-intersection or node. These singularities destroy the clean group law.

The nonsingularity condition is

$$4a^3 + 27b^2 \neq 0.$$

Equivalently, the discriminant

$$\Delta = -16(4a^3 + 27b^2)$$

is nonzero. The note's highlighted line is important: this condition is not cosmetic. It ensures that the tangent line is well-behaved and that the addition law can be defined without pathological exceptions.

§32.2 Why the discriminant appears

If the cubic has a double root r_1 , then

$$x^3 + ax + b = (x - r_1)^2(x - r_2).$$

Expanding and comparing coefficients gives a relation between a, b, r_1, r_2 . Eliminating the roots yields

$$4a^3 + 27b^2 = 0.$$

Conversely, if this expression vanishes, the cubic has a repeated root. This is the cubic discriminant calculation.

A useful conceptual memory: repeated roots occur exactly when a polynomial and its derivative have a common root. Here

$$f(x) = x^3 + ax + b, \quad f'(x) = 3x^2 + a.$$

So singularity is the failure of f and f' to be coprime.

§32.3 Cryptographic motivation

The printed note then moves to cryptography. Earlier cryptosystems such as Diffie–Hellman use discrete exponentiation in a group. To use elliptic curves cryptographically, we need the points on an elliptic curve to form a group. The advantage is that elliptic-curve groups can give strong security with smaller parameters than classical finite-field multiplicative groups.

This is the key shift:

$$\text{number-theoretic cryptography} \rightsquigarrow \text{hard problems in groups.}$$

For elliptic curves, the hard problem is the elliptic-curve discrete logarithm problem.

§32.4 Geometric addition law

Let $P_1, P_2 \in E$. If $P_1 \neq P_2$, draw the line through them. It intersects the cubic in a third point R . Reflect R across the x -axis; the reflected point is defined to be

$$P_1 + P_2.$$

If $P_1 = P_2$, use the tangent line at P_1 instead of the secant line. This is point doubling.




Figure 32.1: Schematic chord-and-tangent addition on an elliptic curve.

§32.5 Algebraic formula

For

$$P = (x_1, y_1), \quad Q = (x_2, y_2),$$

with $P \neq Q$, the slope is

$$m = \frac{y_2 - y_1}{x_2 - x_1}.$$

Then

$$x_3 = m^2 - x_1 - x_2,$$

and

$$y_3 = m(x_1 - x_3) - y_1.$$

Thus

$$P + Q = (x_3, y_3).$$

For point doubling,

$$m = \frac{3x_1^2 + a}{2y_1}.$$

These formulas come from substituting the line equation into the cubic and using the fact that a line meets a cubic in three points counted with multiplicity.

§32.6 Special cases and the point at infinity

If $Q = -P$, then the line through P and Q is vertical, so it does not meet the affine curve in a third finite point. We define

$$P + (-P) = \mathcal{O},$$

where $\mathcal{O}$ is the point at infinity. This point acts as the identity:

$$P + \mathcal{O} = P.$$

The printed pages point out that all vertical lines on the elliptic curve meet at the point at infinity in the projective closure. This is not just a convenient fiction; projective geometry is the natural setting in which a line always meets a cubic in three points.

§32.7 Why associativity is difficult

The group law is visually simple, but associativity

$$(P + Q) + R = P + (Q + R)$$

is hard to prove by direct geometry. Algebraically it becomes a long calculation. Conceptually, elliptic curves are genus-one algebraic curves, and the chord-and-tangent law is related to divisor classes. That higher viewpoint explains why the operation is genuinely associative rather than merely plausible from the picture.

§32.8 Elliptic curves over finite fields

For cryptography one usually works over a finite field:

$$E(\mathbb{F}_p) : y^2 = x^3 + ax + b \pmod{p}.$$

The same addition formulas work, with division interpreted as multiplication by inverses modulo p . The group $E(\mathbb{F}_p)$ is finite, and scalar multiplication

$$nP = P + P + \cdots + P$$

is computationally easy, while recovering n from P and nP is believed hard for well-chosen curves.

§32.9 Why nonsingularity matters for the group law

The chord-and-tangent construction needs every line to meet the cubic in three points counted with multiplicity. If the curve has a cusp or a node, the tangent behavior at the singular point is not well-defined in the smooth sense, and the group law becomes degenerate.

Algebraically, singularity occurs when the cubic $x^3 + ax + b$ has a repeated root. This happens exactly when the polynomial and its derivative

$$3x^2 + a$$

have a common root. The discriminant condition

$$4a^3 + 27b^2 \neq 0$$

excludes that. Thus the printed discriminant is not a decorative formula; it is what guarantees that the geometric addition law behaves smoothly.

§32.10 Explicit addition formulas

If $P = (x_1, y_1)$ and $Q = (x_2, y_2)$ with $P \neq Q$, the slope of the line through them is

$$m = \frac{y_2 - y_1}{x_2 - x_1}.$$

The third intersection point has x -coordinate

$$x_3 = m^2 - x_1 - x_2.$$

After reflecting across the x -axis, the sum is

$$P + Q = (x_3, m(x_1 - x_3) - y_1).$$

For doubling $P = Q$, use the tangent slope

$$m = \frac{3x_1^2 + a}{2y_1}.$$

These formulas are what make elliptic curves computationally usable.

§32.11 Cryptographic intuition

Over a finite field, the set $E(\mathbb{F}_p)$ is finite but still carries the elliptic-curve group law. Repeated addition gives scalar multiplication:

$$nP = P + P + \cdots + P.$$

It is easy to compute nP by double-and-add, but hard to recover n from P and nP for suitable curves and parameters. This is the elliptic curve discrete logarithm problem.

The note's cryptographic motivation should therefore be understood as a computational asymmetry: the group law is efficient in one direction and hard to invert in the other.

§32.12 Structural viewpoint

This note is a good meeting point of algebraic geometry and cryptography. A cubic equation becomes a smooth curve; smoothness is a discriminant condition; a geometric chord construction becomes a group operation; projective geometry supplies the identity element; finite fields turn the group into a computational object. The remarkable fact is that a picture of a cubic curve encodes a serious algebraic group.

§32.13 Nonsingularity and the discriminant

An elliptic curve in Weierstrass form must be nonsingular. The discriminant condition prevents cusps and self-intersections. Geometrically, nonsingularity ensures that every point has a well-defined tangent line, which is needed for the group law.

§32.14 The group law as divisor arithmetic

The chord-and-tangent rule says that three collinear intersection points add to zero. Behind the picture is divisor theory: rational functions have divisors of zeros and poles whose total degree is zero. The group law on the curve packages this linear equivalence relation.

§32.15 Associativity and geometry

Associativity is visually nontrivial because the chord construction depends on intermediate points. Algebraically it follows from the structure of the Picard group of degree-zero divisor classes:

$$E \cong \text{Pic}^0(E).$$

§32.16 Cryptographic meaning

Elliptic-curve cryptography uses the difficulty of recovering n from $Q = nP$. The geometry supplies a finite abelian group over $\mathbb{F}_q$, and arithmetic security depends on the discrete logarithm problem in that group.

§32.17 Elliptic curves as geometry with a group law

Elliptic curves are not merely cubic equations. Nonsingularity gives tangent geometry, chord-and-tangent gives a group, divisor theory explains associativity, and finite-field points make the group computationally useful.

When Addition Starts Floating: Theta-Links, Log-Links, and Arithmetic Deformation

N-20251113

§33.1 What has to be explained

The previous note ended with a deliberately strange operation. At a bad split multiplicative valuation v , one has a Tate parameter q_v , and the arithmetic-deformation viewpoint asks us to take seriously transformations of the shape

$$q_v \mapsto q_v^{m^2}, \quad u \mapsto u \quad (u \in \mathcal{O}_{F_v}^\times).$$

This is not a ring homomorphism. It is a transformation of multiplicative data. The central question of this note is therefore precise:

How can one deform multiplicative theta-data, lose ordinary additive compatibility, and still recover enough structure to prove an inequality?

The answer in Fesenko's survey passes through nonarchimedean theta-functions, generalized Kummer theory, theta-links, log-links, log-shells, multiradial reconstruction, and indeterminacies. The goal here is not to pretend to define the full IUT apparatus. The goal is to keep the mathematical objects visible and to explain why each object is forced by the previous one.

§33.2 The nonarchimedean theta-function

Let L be a local field of characteristic zero with finite residue field, and let C_L be the completion of an algebraic closure of L . Let $q \in L$ be a nonzero element of the maximal ideal of $\mathcal{O}_L$. In the application to elliptic curves, q is the Tate parameter q_v .

Definition 33.2.1 (Theta-function of type au^m). A holomorphic function $f : C_L^\times \rightarrow C_L$, defined over L , is a theta-function of type au^m , where $a \in C_L^\times$ and $m \in \mathbb{Z}$, if

$$f(u) = au^m f(qu).$$

A meromorphic function invariant under $u \mapsto qu$ descends to a function on the Tate quotient

$$C_L^\times / \langle q \rangle.$$

The basic nonarchimedean theta-function used in the survey is

$$\theta(u) = \sum_{n \in \mathbb{Z}} (-1)^n q^{n(n-1)/2} u^n.$$

It also has the product expansion

$$\theta(u) = (1-u) \prod_{n \geq 1} (1-q^n)(1-q^n u)(1-q^n u^{-1}),$$

which is the nonarchimedean analogue of the Jacobi triple product expression. The function is chosen because its zeros and poles interact well with the Tate curve and with the theta-values used in IUT.

Remark 33.2.2 (Why a theta-function appears) — The theta-function is not added to the story for analogy with complex analysis alone. On a Tate curve, elliptic functions with period q can be expressed in terms of theta-functions. Thus theta-functions are a natural language for functions on the local analytic uniformization

$$E(L) \cong L^\times / \langle q \rangle.$$

Fesenko also records a modified theta-function, written here schematically as $\ddot{\Theta}(u)$, whose special values at points separated by periods from $\pm i$ satisfy a relation of the form

$$q^{m^2/2} \ddot{\Theta}(i\sqrt{q^m}) = \ddot{\Theta}(i).$$

After choosing a $2l$ -th root of q , this leads to theta-values of the shape

$$\Theta(\sqrt{-q^m}) = q^{m^2}, \quad 1 \leq m \leq \frac{l-1}{2},$$

up to multiplication by roots of unity of order dividing $2l$.

Remark 33.2.3 (A controlled simplification) — The formula above is the correct structural feature needed here: special theta-values encode powers q^{m^2} . The complete IUT construction keeps track of coverings, labels, roots of unity, Galois actions, and compatible local-global data. Those details are not being suppressed because they are unimportant; they are being postponed because the first conceptual step is to see why theta-values can serve as deformation data.

§33.3 The local theta-link as a non-ring-theoretic map

At a bad reduction valuation of odd residue characteristic, the main version of the theta-link described by Fesenko revolves around a local monoid-theoretic transformation. In simplified notation it sends

$$q \longmapsto \{\Theta(\sqrt{-q^m}) = q^{m^2}\}_{1 \leq m \leq (l-1)/2},$$

and it sends units to units by the identity map, up to the appropriate torsion ambiguity.

Definition 33.3.1 (Local theta-link, structural form). A local theta-link is a monoid-theoretic comparison which, at a bad split multiplicative valuation, keeps the unit part essentially fixed while replacing the Tate-period generator q by a family of theta-values behaving like q^{m^2} . It is a comparison of multiplicative and theta-theoretic data, not a homomorphism of local rings.

Thus the local theta-link acts on two kinds of local data:

$$\mathcal{O}_L^\times / \text{Tor}(\mathcal{O}_L^\times), \quad \{q^{m^2} : 1 \leq m \leq (l-1)/2\},$$

and it is glued to global monoid-theoretic data satisfying the product formula.

Proposition 33.3.2 (Why the theta-link is not scheme-theoretic)

The theta-link is not compatible with the full ring structure of L . In particular, it should not be interpreted as a morphism of schemes or as a ring isomorphism between two local fields.

Proof. A ring homomorphism must preserve both multiplication and addition. The theta-link is built from the multiplicative monoid of units, powers of the Tate parameter, and special values of a theta-function. It transports multiplicative/theta data, but the additive operation of the local field is not transported by the displayed map. Fesenko therefore describes this as an arithmetic deformation of the local structure rather than as an ordinary scheme-theoretic morphism. $\square$

This is the first honest meaning of the phrase “addition starts floating”. It does not mean that the equation $1 + 1 = 2$ becomes false inside a fixed field. It means that after theta-deformation one is no longer allowed to identify the additive coordinates on the two sides as if a ring homomorphism had carried them across.

§33.4 The bridge supplied by Kummer theory

If the theta-link only destroyed additive compatibility, it would not be useful. The next ingredient explains how monoid-theoretic data can still communicate with Galois and fundamental-group data.

Definition 33.4.1 (Kummer map, schematic form). Let H be an open subgroup of a Galois or arithmetic fundamental group acting on an abelian group M . Kummer theory studies maps of the form

$$M^H \longrightarrow H^1(H, \operatorname{Hom}(\mathbb{Q}/\mathbb{Z}, M)),$$

obtained by considering divisibility of elements of M .

For ordinary intuition, Kummer theory turns the question of extracting roots into a cohomological Galois datum. In IUT, the relevant versions are generalized and truncated. They are applied not merely to elements of fields, but also to monoid-theoretic structures associated with theta-functions, theta-values, number fields, and their completions.

Remark 33.4.2 (Kummer theory as a translator) — The theta-link is monoid-theoretic. Mono-anabelian reconstruction works on the side of arithmetic fundamental groups. Kummer theory is one of the translators between these languages. It relates theta-values and monoid data to Galois or fundamental-group structures that can be reconstructed and compared.

A key issue here is cyclotomic rigidity. Various copies of $\widehat{\mathbb{Z}}$, or quotients of such copies, arise from roots of unity, from monoid-theoretic data, and from subquotients of fundamental groups. IUT needs algorithms that identify the relevant cyclotomic structures in a canonical way. Without such rigidity, the roots of unity introduced by extracting l -th or $2l$ -th roots would create uncontrolled ambiguity.

§33.5 Two symmetries and the bookkeeping of labels

The theta-link is built after choosing a prime $l > 3$. The l -torsion points of the elliptic curve provide labels modeled on $\mathbb{F}_l$. Fesenko emphasizes two different symmetries.

Definition 33.5.1 (The two label symmetries). The geometric-additive symmetry is

$$\mathbb{F}_l^{\times\pm} = \mathbb{F}_l \rtimes \{\pm 1\},$$

acting on labels by

$$z \mapsto \pm z + a.$$

The arithmetic-multiplicative symmetry is

$$\mathbb{F}_l^{*\pm} = \mathbb{F}_l^\times / \{\pm 1\}.$$

The first symmetry is essentially geometric. It is related to the geometric part of the arithmetic fundamental groups and to Kummer theory around theta-values. It is also used to synchronize conjugacy ambiguities of local Galois groups attached to different labels.

The second symmetry is essentially arithmetic. It is related to the Galois action of number fields and is multiplicative by definition. It helps separate the zero label from the nonzero labels and is more naturally multiradial.

Symmetry	Nature	Role
$\mathbb{F}_l^{\times\pm}$	geometric/additive	conjugate synchronization
$\mathbb{F}_l^{*\pm}$	arithmetic/multiplicative	label bookkeeping and descent

Remark 33.5.2 (Why labels are not cosmetic) — The labels are not merely indices in a long construction. They keep track of different copies of local Galois data and theta-values. If these labels are identified too naively, one loses the distinction between the geometric-additive and arithmetic-multiplicative parts of the construction.

§33.6 Conjugate synchronization and the basepoint problem

Fundamental groups depend on basepoints. In ordinary expositions this dependence is often hidden because different basepoints lead to groups identified up to inner automorphism. In IUT this ambiguity cannot simply be ignored, because the theta-link and log-link compare structures that live in different theaters.

Definition 33.6.1 (Conjugate synchronization, informal). Conjugate synchronization is a system of identifications between local absolute Galois groups attached to different labels, arranged so that one does not introduce uncontrolled conjugacy indeterminacy.

The need for such synchronization is one reason the theory becomes much more rigid than a naive “copy-and-paste” of local fields. The comparison is not allowed to choose arbitrary basepoint identifications. It must respect the label symmetries and the Kummer-theoretic data surrounding the theta-values.

§33.7 The log-link: recovering additive structure from multiplicative structure

The theta-link uses multiplicative theta-values. To regain access to additive structure, Fesenko turns to the nonarchimedean logarithm.

Definition 33.7.1 (Nonarchimedean logarithm). For a local field L , the nonarchimedean logarithm

$$\log : \mathcal{O}_L^\times \longrightarrow L$$

is defined on elements $1 - \alpha$, with α in the maximal ideal of $\mathcal{O}_L$, by

$$\log(1 - \alpha) = - \sum_{n \geq 1} \frac{\alpha^n}{n},$$

and sends multiplicative representatives of the finite residue field to 0.

The key point is that $\log$ transforms multiplication into addition where it is defined. Thus the logarithm lets multiplicative unit data speak an additive language.

Definition 33.7.2 (Log-link, structural form). A log-link is the comparison which shifts multiplicative unit data through the nonarchimedean logarithm, so that additive structures can be reconstructed from logarithmic images of multiplicative structures.

Fesenko explains the relation between the two links as follows. The multiplicative structures on the two sides of the theta-link are related through value-group data. The additive structures are related through unit-group data after applying the log-link.


Remark 33.7.3 (Why the logarithm is not optional) — Theta-values naturally act on multiplicative data, not directly on the additive group of the local field. The logarithm creates a setting where multiplicative unit data has an additive image. This is why the log-link is not an auxiliary simplification; it is part of the reconstruction mechanism.

§33.8 The log-theta-lattice

The theta-link and log-link do not appear as isolated arrows. They are organized into a two-dimensional pattern called the log-theta-lattice.

Definition 33.8.1 (Log-theta-lattice, schematic). The log-theta-lattice is a two-dimensional array of theaters. Vertical arrows are log-links, and horizontal arrows are theta-links. A horizontal theta-link at one level depends on the log-link structure immediately below it.

The picture is:



This diagram is only a guide. Fesenko notes that the main results use just two neighboring infinite vertical lines and the horizontal arrow between the lattice points $(0,0)$ and $(1,0)$. The important idea is that the theory does not compare a single source and target once. It studies a controlled pattern of comparisons.

§33.9 The log-shell: a common container for incompatible links

A technical difficulty now appears. The log-link does not commute with the theta-link. If one expected a clean commutative square, one would be disappointed. IUT instead introduces a common structure that can hold the relevant images even when elementwise compatibility fails.

Definition 33.9.1 (Log-shell). For a nonarchimedean local field L , the nonarchimedean part of the log-shell is a compact subgroup of the form

$$(p^*)^{-1} \log(\mathcal{O}_L^\times),$$

where $p^* = p$ if p is odd and $2^* = 4$. At a complex archimedean place, the corresponding log-shell is the closed ball of radius π .

The log-shell should not be imagined as an error bar in a numerical approximation. It is a structural container. The relevant Kummer isomorphisms are not compatible with the log-link at the level of individual elements, but their images are contained in the log-shell. This is exactly the kind of weakened compatibility that arithmetic deformation needs.

Remark 33.9.2 (Compatibility without equality) — In ordinary algebra one often tries to prove that a diagram commutes element by element. The log-shell records a weaker but controlled form of compatibility: one compares inclusions of certain subsets or regions rather than exact equality of individual transported elements.

§33.10 Mono-anabelian reconstruction and multiradiality


The words “reconstruction” and “multiradiality” are easy to mystify, but their roles are concrete.

Definition 33.10.1 (Mono-anabelian reconstruction). Mono-anabelian reconstruction is the construction of ring-theoretic data from a single topological group that is known to be isomorphic to an arithmetic fundamental group of a suitable hyperbolic curve or orbicurve. The reconstruction is algorithmic and functorial in the sense needed for localization and change of base field.

The difference from ordinary anabelian comparison is important. Classical anabelian geometry often compares two geometric objects through an isomorphism of their fundamental groups. Mono-anabelian geometry asks for algorithms that start from one suitable group and recover the ring structure associated to it.

Definition 33.10.2 (Multiradial algorithm). A reconstruction algorithm is multiradial if an object constructed from one spoke of a wheel-like system is described in terms that still make sense from the viewpoint of other spokes. Equivalently, it is compatible with simultaneous execution at multiple spokes.

Here the “wheel and spokes” language is only an analogy, but a useful one. A spoke corresponds to one theater or one radial perspective. The center corresponds to a core object reconstructed from data that can be described from more than one spoke.



multiradial descriptions must remain meaningful from several spokes at once

Remark 33.10.3 (Why indeterminacy appears) — A multiradial description is stronger than a one-spoke description. To make such descriptions possible, IUT allows certain controlled indeterminacies. These are not defects added after the fact; they are part of how the theory avoids making illegal identifications between different theaters.

§33.11 The three indeterminacies

Fesenko singles out three indeterminacies that enter the computation of volume deformation.

Definition 33.11.1 (The three indeterminacies, structural form). The three main indeterminacies are:

1. Ind1: an indeterminacy related to permutation symmetries of Galois and arithmetic fundamental groups along vertical lines of the log-theta-lattice;
2. Ind2: an indeterminacy related to the horizontal theta-link and to the action of a compact group on $\log(\mathcal{O}_L^\times)$;
3. Ind3: an indeterminacy arising from upper semi-compatibility of Kummer isomorphisms with log-links in a single vertical line.

The third phrase, “upper semi-compatibility”, is deliberately weaker than commutativity of a diagram. Instead of asking for equality of transported elements, one asks for controlled inclusion relations among the relevant subsets.

Proposition 33.11.2 (Why indeterminacies are mathematically productive)

The indeterminacies do not merely obscure the comparison. They make it possible to compare data across distinct theaters without falsely identifying the full ring structures on the two sides.

Proof. The theta-link is not ring-theoretic, and the log-link does not commute with it. Therefore an elementwise comparison of ordinary rings would impose more structure than the construction provides. The indeterminacies specify exactly which ambiguity is allowed: permutation ambiguity, horizontal theta-link ambiguity, and weakened Kummer/log compatibility. A volume computation can then be performed after enlarging the target by these controlled ambiguities. $\square$

This is where the intuitive phrase “floating addition” becomes mathematically safer. The additive structure is not discarded into chaos. It is reconstructed through logarithmic images and fundamental-group algorithms, while the unavoidable ambiguity is measured by the three indeterminacies.

§33.12 What the second note adds

We can now summarize the actual mechanism more accurately than by saying “IUT uses many universes”.

1. The Tate parameter q gives local multiplicative data at bad reduction valuations.
2. Nonarchimedean theta-functions produce distinguished values behaving like q^{m^2} .
3. The theta-link transports this theta/multiplicative data but is not a ring morphism.
4. Generalized Kummer theory relates monoid-theoretic data to Galois and arithmetic fundamental-group data.
5. The two label symmetries separate the geometric-additive side from the arithmetic-multiplicative side.
6. The log-link uses $\log : \mathcal{O}_L^\times \rightarrow L$ to recover additive information from multiplicative units.
7. The log-shell supplies a common container when the log-link and theta-link fail to commute elementwise.
8. Mono-anabelian reconstruction and multiradial algorithms make it possible to describe reconstructed objects from several theaters at once.
9. The three indeterminacies specify the allowed ambiguity and enter the final volume estimate.

The next note should therefore not treat the controversy around IUT as a piece of mathematical news. The foundational issue is already visible here: once ordinary ring-theoretic identification is forbidden, the theory must replace it by theta-links, log-links, Kummer correspondences, and multiradial reconstruction. The controversy is ultimately about whether this replacement is strong enough, not about whether the vocabulary is unusual.

Teichmueller Theory Without One Universe: Classical Deformation, Hodge Theaters, and Identification

N-20251114

§34.1 Why the word Teichmueller matters

The word “Teichmueller” in inter-universal Teichmueller theory should not be treated as a poetic label. It points to a deformation problem. In classical Teichmueller theory, one fixes a topological surface and studies how complex structures on that surface vary. In IUT, the analogy is not that arithmetic curves literally form a classical Teichmueller space. The analogy is that some structures are treated as rigid reference data, while other structures are deformed and measured.

This note compares the two situations. The comparison is useful only if it is kept precise. Classical Teichmueller theory varies complex structures inside a common topological background. IUT tries to compare arithmetic data across distinct theaters where ordinary ring-theoretic identification is not allowed. The controversial strength of the theory is already visible at this level: it replaces the ordinary demand for one common universe by a system of reconstructed, multiradial, and indeterminacy-controlled comparisons.

§34.2 Classical Teichmueller theory: fixed topology, variable complex structure

Let S be an oriented topological surface of finite type. To remember which complex surface is supposed to be a deformation of S , one uses a marking.

Definition 34.2.1 (Marked Riemann surface). A marked Riemann surface modeled on S is a pair

$$(X, f),$$

where X is a Riemann surface and

$$f : S \longrightarrow X$$

is an orientation-preserving homeomorphism, or more generally a quasiconformal map, regarded up to homotopy.

The marking is the key device. Without it, one has only a collection of complex curves. With it, different complex structures can still be compared as structures on the same underlying topological surface.

Definition 34.2.2 (Teichmueller space). The Teichmueller space $T(S)$ is the set of equivalence classes of marked Riemann surfaces (X, f) , where

$$(X, f) \sim (Y, g)$$

if there is a biholomorphism $h : X \rightarrow Y$ such that $h \circ f$ is homotopic to g .

Thus classical Teichmueller theory separates two layers:

topological surface S and complex structure on S .

The topology is the reference frame. The complex structure is what moves.

Remark 34.2.3 (The role of marking) — A marking is an authorized identification. It tells us when two points of Teichmueller space represent two complex structures on the same underlying surface. This will be the main contrast with IUT: IUT cannot simply use a single marking to identify all ring structures across its theaters.

§34.3 Distortion is measured in one universe

Teichmueller theory does not only parametrize complex structures. It measures deformation by quasiconformal distortion.

Definition 34.3.1 (Quasiconformal distortion). Let $f : X \rightarrow Y$ be a quasiconformal map between Riemann surfaces. In a local coordinate it has distributional derivatives f_z and $f_{\bar{z}}$. Its Beltrami coefficient is

$$\mu_f = \frac{f_{\bar{z}}}{f_z},$$

and the maximal dilatation is

$$K(f) = \frac{1 + \|\mu_f\|_\infty}{1 - \|\mu_f\|_\infty}.$$

The Teichmueller distance is then defined by minimizing distortion among maps in the correct homotopy class:

$$d_T((X, f), (Y, g)) = \frac{1}{2} \inf_h \log K(h),$$

where $h : X \rightarrow Y$ ranges over quasiconformal maps such that $h \circ f$ is homotopic to g .

The important point is structural: this distance makes sense because the comparison is made inside one clear framework. Both surfaces are marked by the same S . The phrase “same underlying topological surface” has a precise meaning.

Proposition 34.3.2 (Classical deformation has a common reference object)

In classical Teichmueller theory, the comparison of two complex structures is mediated by the fixed surface S . The marking supplies a legitimate way to say that the two complex structures are structures on the same object.

Proof. A point of $T(S)$ is not merely a Riemann surface X ; it is a pair (X, f) . Given two points (X, f) and (Y, g) , the maps f and g identify both surfaces with the same topological source S , up to homotopy. Therefore quasiconformal maps $h : X \rightarrow Y$ can be tested against the homotopy condition $h \circ f \simeq g$. The fixed S is the common reference object. $\square$

§34.4 Why the arithmetic analogy cannot be literal

In the arithmetic setting, the objects are not Riemann surfaces equipped with varying complex structures. The basic geometric objects are hyperbolic curves, elliptic curves with punctures or orbifold quotients, number fields, local fields, arithmetic fundamental groups, and Galois groups. The rigid structure is no longer a topological surface. It is closer to a package of fundamental-group-like data from which ring structures can be reconstructed.

Definition 34.4.1 (Arithmetic fundamental-group reference data, schematic). For a hyperbolic curve X over a number field K , the arithmetic fundamental group fits into an exact sequence

$$1 \longrightarrow \pi_1^{\text{geom}}(X) \longrightarrow \pi_1(X) \longrightarrow G_K \longrightarrow 1.$$

In the IUT entrance, such group-theoretic data acts as a rigid reference structure from which arithmetic and ring-theoretic data may be reconstructed by mono-anabelian methods.

The analogy with classical Teichmueller theory is therefore shifted:

classical theory : fixed topological type, variable complex structure,

IUT direction : rigid group-theoretic data, deformed ring-theoretic presentations.

This is only an analogy, but it is a mathematically meaningful one.

Remark 34.4.2 (What should not be inferred) — The word “Teichmueller” does not mean that IUT constructs a Teichmueller space of number fields in the classical sense. It means that the theory is organized around deformation, rigidity, and comparison of structures, but the structures being compared are arithmetic and group-theoretic rather than complex-analytic.

§34.5 Hodge theaters as arithmetic stages

The language of theaters is one of the points where IUT becomes unfamiliar. A theater should not be imagined as a parallel universe in a science-fiction sense. It is a package of data within which certain arithmetic, geometric, and group-theoretic structures can be discussed internally.

Definition 34.5.1 (Hodge theater, working schematic). In this note, a Hodge theater means a package of arithmetic data containing several mutually related layers: local fields, number-field data, arithmetic fundamental groups, monoid-theoretic data, and theta-data. The full definition in IUT is much more elaborate. The working point is that a theater provides an internal stage on which ring structures, Galois data, theta-values, and reconstruction algorithms coexist.

The word “theater” is useful because IUT does not simply compare two rings by a ring isomorphism. It compares data internal to one theater with data internal to another theater through special links.

Definition 34.5.2 (Inter-universal comparison, schematic). An inter-universal comparison is a comparison between data belonging to distinct Hodge theaters, where one does not assume a global ring-theoretic identification between the theaters. The comparison is mediated by structures such as theta-links, log-links, Kummer correspondences, and mono-anabelian reconstruction.

In ordinary algebraic geometry, the most comfortable comparison has the form

$$R \cong R'$$

or at least a morphism of schemes. In IUT, the central horizontal comparison is not of that kind. The theta-link is precisely a link that deforms ring-theoretic structure rather than preserving it.

§34.6 Etale-like and Frobenius-like directions

Fesenko’s survey uses the distinction between structures that behave like rigid etale or fundamental-group data and structures that behave more like monoid-theoretic or Frobenius-type data. This is one of the cleanest ways to understand why ordinary scheme morphisms are not enough.

Definition 34.6.1 (Etale-like and Frobenius-like structures, schematic). An etale-like structure is a structure on the fundamental-group side: it is functorial, rigid, and invariant enough to be used for reconstruction. A Frobenius-like structure is a monoid-theoretic structure used to build deformation links, especially when no ordinary Frobenius morphism of number fields exists.

The terminology is suggestive. Over a field of positive characteristic, the Frobenius map $x \mapsto x^p$ is a genuine morphism. For number fields, there is no such global Frobenius endomorphism. IUT instead builds categorical and monoid-theoretic analogues that can play a deformation role.

Remark 34.6.2 (The arithmetic reason this is radical) — The theory tries to get a Frobenius-like deformation effect in a setting where the usual Frobenius morphism does not exist. This is part of the reason the comparison cannot remain inside ordinary scheme-theoretic geometry.

§34.7 The theta-link is the point where one universe breaks

The local model already appeared in the previous two notes:

$$q_v \mapsto q_v^{m^2}, \quad u \mapsto u.$$

This operation is meaningful on the multiplicative/monoid side, but it is not a ring homomorphism. It therefore cannot serve as an identification of local fields as rings.

Proposition 34.7.1 (No common ring is carried by the theta-link)

If a theta-link is interpreted as the assignment that carries a Tate parameter q to theta-values behaving like q^{m^2} while keeping unit data on the monoid side, then it cannot be regarded as a ring-theoretic isomorphism between the source and target local structures.

Proof. A ring-theoretic isomorphism must preserve addition and multiplication simultaneously. The theta-link is constructed from multiplicative monoids, unit groups, Tate parameters, and theta-values. Its displayed local behavior is not an additive operation. Therefore treating it as a ring isomorphism would add a compatibility that the construction is not designed to provide. $\square$

This is the exact mathematical sense in which the “one universe” picture breaks. If one tries to force the two sides of the theta-link into the same ring-theoretic universe, one loses the deformation that the theta-link was built to express.

§34.8 The log-link does not restore naive equality

The log-link exists because the theta-link moves multiplicative data while the final Diophantine inequality needs additive and volume information. The nonarchimedean logarithm

$$\log : \mathcal{O}_L^\times \longrightarrow L$$

translates unit data into additive data. This does not mean that the original ring structure has been transported unchanged.

Definition 34.8.1 (Recovered additive data, schematic). Recovered additive data means additive structure obtained after applying logarithmic and mono-anabelian reconstruction procedures to multiplicative or group-theoretic data. It is not the same thing as an additive structure transported by a ring isomorphism.

The distinction is small in words but large in mathematics. Classical Teichmueller theory deforms complex structures while markings keep track of the same underlying topological surface. IUT deforms arithmetic structures in a way that prevents one from identifying the two sides as the same ring. Additive information must therefore be reconstructed rather than simply carried over.

Remark 34.8.2 (Why the log-shell is needed) — The log-link and the theta-link do not give a clean commutative square of elementwise identifications. The log-shell provides a controlled container for the resulting ambiguity. It is a way of replacing exact equality by a bounded or contained region in which the relevant images live.

§34.9 Classical marking versus multiradial reconstruction

The most useful comparison between classical Teichmueller theory and IUT is not a slogan about deformation. It is a comparison between two ways of controlling identity.

The comparison can be kept in five pairs.

1. Classical theory chooses a fixed topological surface S ; IUT treats arithmetic fundamental-group-like data as rigid and reconstructive.
2. Classical theory uses a marking $f : S \rightarrow X$; IUT uses mono-anabelian algorithms to recover ring-theoretic data from group-theoretic input.
3. Classical theory moves the complex structure on S ; IUT moves a ring-theoretic or monoid-theoretic presentation of arithmetic data.
4. Classical theory measures distortion by quasiconformal maps; IUT creates and measures arithmetic deformation by theta-links and log-links.

5. Classical comparison remains inside one marked topological setting; IUT comparison moves across distinct Hodge theaters.

This list shows why IUT is not merely difficult notation. The role played by marking in classical Teichmueller theory is not replaced by one simple arithmetic marking. Instead, it is replaced by a much more elaborate system: reconstruction algorithms, label symmetries, Kummer theory, log-shells, and indeterminacies.

Definition 34.9.1 (Multiradial identification, schematic). A multiradial identification is not a direct equality of objects in one common ring universe. It is an identification of reconstructed structures that remains meaningful when viewed from several theaters or spokes at once.

This is why multiradiality is conceptually central. If a construction only makes sense from one spoke, it may depend on a hidden choice of coordinates. A multiradial construction is designed to survive the change of theater.

§34.10 The mathematical source of the controversy

This note should not become a news summary. The mathematical tension can be stated without naming any episode. It is a tension about legitimate identification.

Proposition 34.10.1 (The identification burden)

Once ordinary ring-theoretic identification is forbidden across theaters, every later comparison must justify itself by some replacement mechanism: mono-anabelian reconstruction, Kummer correspondence, theta-link compatibility, log-link compatibility, multiradiality, or a controlled indeterminacy.


Proof. Without a ring isomorphism, there is no default rule saying that an element, sum, unit, valuation, or theta-value in one theater is the same object as an element, sum, unit, valuation, or theta-value in another. Therefore each permitted comparison must be constructed and checked. The structures listed in the statement are precisely the devices introduced to make such comparisons possible without collapsing all theaters into one ordinary ring-theoretic setting. $\square$

The controversial foundational move is therefore not merely “there are many universes”. A vague multiverse would have no mathematical force. The real move is this: do not identify rings directly; instead recover and compare selected structures by controlled algorithms. The difficulty is that the replacement mechanisms must be strong enough to support the final log-volume inequality. If they are too weak, the deformation cannot be measured. If they are interpreted as ordinary equalities, the deformation disappears.

Remark 34.10.2 (Why a one-universe reading is dangerous) — A one-universe reading tries to place all theaters into a single background in which ordinary ring-theoretic equality is available. That may make the theory look more familiar, but it changes the object under study. The theta-link is designed precisely to avoid being a ring-theoretic identification.

§34.11 A compact model of the logic

The logical structure of the comparison can be summarized as follows:



The dashed arrow is the temptation that the theory resists. The comparison does not proceed by declaring the two reconstructed rings equal. It proceeds by controlling how their relevant images sit inside the logarithmic and theta-theoretic structures.

§34.12 What this third note establishes

The relationship with classical Teichmueller theory can now be stated cleanly.

1. Classical Teichmueller theory fixes a topological surface and deforms complex structures on it.
2. The marking is the mechanism that makes comparison legitimate in the classical theory.
3. IUT does not have a single arithmetic marking that identifies all ring structures across theaters.
4. Instead, IUT uses Hodge theaters, theta-links, log-links, mono-anabelian reconstruction, Kummer theory, multiradiality, and indeterminacies.
5. The phrase “inter-universal” means that comparisons are made without collapsing the theaters into one ordinary ring-theoretic universe.
6. The deepest mathematical pressure point is therefore the legitimacy and strength of these replacement identifications.

The theory is difficult not only because it contains many objects, but because it changes the usual rule of comparison. Classical Teichmueller theory asks how complex structures vary when the underlying topological object is held fixed. IUT asks how arithmetic structures can be deformed when the ordinary ring structure is not allowed to serve as a common background. That is the mathematical content behind the word “inter-universal”.



Functions and Inequalities

Part V: Contents

N-20220417	229
35.1 Proof strategies	229
35.2 Cyclic and symmetric sums	229
35.3 AM–GM	230
35.4 Muirhead’s inequality	230
35.5 Power means	230
35.6 Hölder and Cauchy	231
35.7 Jensen and Karamata	231
35.8 Cyclic sums	231
35.9 Proof strategy before computation	231
35.10 Next viewpoint	232
35.11 Inequalities as convexity statements	232
35.12 Majorization and Karamata	232
35.13 Holder as norm geometry	232
35.14 Convexity, duality, and equality cases	232
35.15 Inequality methods as choice of language	233
35.16 Convexity behind Jensen and Karamata	233
35.17 Schur convexity	233
35.18 Entropy and the log-sum inequality	234
N-20220718	235
36.1 Fixed-point iteration	235
36.2 Möbius recurrences	235
36.3 Example	236
36.4 Second-order linear recurrences	236
36.5 Dedekind cuts and metric exercises	236
36.6 Fixed points of recurrence maps	237
36.7 Monotone bounded sequences	237
36.8 After the examples	237
36.9 Fixed points and stability	237
36.10 Möbius recurrences and matrices	237
36.11 Second-order recurrences	238
36.12 Recurrences as one-dimensional dynamics	238
N-20220721	239
37.1 Random experiments and sample spaces	239
37.2 Classical model and Venn diagrams	239
37.3 Conditional probability	239
37.4 A transition to substitution tricks	240
37.5 A classical symmetric equation	240
37.6 Inequality examples	241
37.7 Jensen, trigonometry, and contest style	241
37.8 The hidden structure	241
37.9 Conditional probability as changing measure	241
37.10 Bayes formula and inverse problems	241
37.11 Inequality substitutions as coordinate choices	242
37.12 Changing measures and changing coordinates	242
N-20220801	245
38.1 Subsets and characteristic functions	245
38.2 Set laws	245
38.3 Venn diagram exercises	246
38.4 Inclusion–exclusion with literature examples	246
38.5 Set constructions with squares	246
38.6 Floor functions	247
38.7 Functions and monotonicity	247
38.8 Characteristic functions and subsets	247

38.9	Functions as rules with domains	247
38.10	The method behind the trick	248
38.11	Set operations as Boolean algebra	248
38.12	Characteristic functions and algebra of predicates	248
38.13	Preimages as Boolean homomorphisms	248
38.14	Sigma-algebras and measurable events	249
38.15	Stone duality intuition	249
38.16	Difference of squares as algebraic structure	249
38.17	Sets as algebra, logic, and measurement	249
38.18	Boolean algebra beyond Venn diagrams	249
38.19	Characteristic functions	250
38.20	Sigma-algebras as observable event systems	250
38.21	Stone duality as topology of logic	250
38.22	Returning to the original set exercises	251
38.23	From Venn diagrams to algebra of information	251
38.24	Why sigma-algebras restrict what can be asked	251
N-20220806		253
39.1	Images, preimages, and finite set examples	253
39.2	Closure under addition	253
39.3	A combinatorial maximal-set problem	253
39.4	Functional and algebraic exercises	254
39.5	Matrix and coordinate side computations	254
39.6	Images and preimages behave differently	254
39.7	Closure operations	255
39.8	The reusable pattern	255
39.9	Closure operators	255
39.10	Generated objects as intersections	255
39.11	Fixed points of closure	255
39.12	Images, preimages, and logic	256
39.13	Posets and extremal chains	256
39.14	Galois connections	256
39.15	Tarski fixed-point theorem	256
39.16	Closure, order, and fixed-point thinking	257
39.17	Closure operators across mathematics	257
39.18	Closed sets as fixed points	257
39.19	Galois connections and closure	257
39.20	Tarski fixed points and monotone iteration	258
39.21	Returning to the original exercises	258
39.22	Generation as an intersection construction	258
39.23	Closure and logic of consequences	258
N-20220808		261
40.1	A set-counting problem	261
40.2	A modular arithmetic construction	261
40.3	Why not every construction is finished	261
40.4	Sequence exercises	262
40.5	Mathematical induction	262
40.6	Examples of induction	262
40.7	Induction structure	263
40.8	Modular patterns	263
40.9	A note on scope	263
40.10	Construction versus upper bound	263
40.11	Residue classes as finite quotient spaces	263
40.12	Induction as well-founded recursion	264
40.13	Finite differences in divisibility induction	264
40.14	Floor functions and lattice points	264
40.15	Pigeonhole and averaging	265
40.16	Finite structure from quotients and induction	265

N-20220810	267
41.1 Basic inequalities	267
41.2 AM–GM	267
41.3 Cauchy’s inequality	268
41.4 Triangle inequality and its variants	268
41.5 Worked AM–GM examples	268
41.6 A Cauchy example	269
41.7 Coordinate geometry at the end of the note	269
41.8 What the examples reveal: three structures	269
41.9 Cauchy–Schwarz and Gram determinants	270
41.10 Triangle inequality and normed spaces	270
41.11 AM–GM, convexity, and cones	270
N-20220811	271
42.1 How to read this note	271
42.2 A fraction problem with a constraint	271
42.3 A weighted square-root comparison	271
42.4 Maximizing a radical expression	272
42.5 Equality in Cauchy–Schwarz	273
42.6 Triangle-angle inequalities	273
42.7 Two spheres and a linear equation	273
42.8 A product inequality	274
42.9 Vector-style radical inequality	274
42.10 The deeper habit	274
42.11 Equality cases as structural information	275
42.12 Optimization on constrained domains	275
42.13 Vectors and geometric inequalities	275
42.14 Equality cases as geometric diagnostics	275
42.15 Equality cases as geometry	275
42.16 Lagrange multipliers and inequalities	276
42.17 Semidefinite viewpoint	276
42.18 Returning to the original problems	276
N-20221012	277
43.1 Bernoulli’s inequality	277
43.2 Power means	277
43.3 Cauchy–Schwarz	277
43.4 Minkowski inequality	278
43.5 A three-variable inequality	278
43.6 Chebyshev’s inequality	278
43.7 Jensen’s inequality	279
43.8 Young, Hölder, and Minkowski	279
43.9 Cauchy–Schwarz as a master inequality	279
43.10 Equality conditions	280
43.11 The lesson behind it	280
43.12 The hierarchy of means	280
43.13 Convex order and Jensen	280
43.14 Minkowski and norm geometry	280
43.15 Classical inequalities in one convex landscape	280
N-20221105	281
44.1 A short note about elementary methods	281
44.2 Periodicity from a transformed argument	281
44.3 A quadratic minimum	281
44.4 A logarithmic comparison	282
44.5 A parameterized polynomial	282
44.6 An inequality under $abc=1$	282
44.7 The conceptual turn	283
44.8 Periodicity from transformed arguments	283
44.9 Quadratic and logarithmic estimates	283
44.10 Functional equations and orbit reduction	283

44.11	Functions through symmetry and normal form	283
N-20221217		285
45.1	Structure of the exercise sheet	285
45.2	Domain problems	285
45.3	Floor functions	286
45.4	Range problems	286
45.5	Dirichlet function	286
45.6	Graph transformations and zero counting	287
45.7	Logarithmic parameter equation	287
45.8	Functional equations with period-like recurrence	287
45.9	A reciprocal functional equation	288
45.10	Quadratic iteration problem	288
45.11	Max and absolute value	288
45.12	Parameter problems and critical values	288
45.13	Functional equations as reduction rules	289
45.14	What the pattern suggests	289
45.15	Domains as constraint intersections	289
45.16	Dirichlet function through density and measure	289
45.17	Thomae's function as a nearby comparison	290
45.18	Parameter problems as bifurcation diagrams	290
45.19	Functional equations as group actions on the domain	290
45.20	Max, min, and nonsmooth points	290
45.21	Ranges as images and compactness	291
45.22	Function exercises through topology and dynamics	291
N-20230107		293
46.1	A note on graph problems	293
46.2	Centers of symmetry	293
46.3	Symmetry about a line	293
46.4	Absolute value transformations	293
46.5	Zero counting by graphs	294
46.6	How to find a center of symmetry in practice	294
46.7	Absolute value as folding	294
46.8	Parameter motion and intersection counting	294
46.9	Graph transformations as group actions	295
46.10	Absolute value as folding, not algebraic noise	296
46.11	Zero counting as a stability problem	296
N-20230116		297
47.1	What a function is	297
47.2	Image and preimage	297
47.3	Injective and surjective	297
47.4	Composition and inverse functions	298
47.5	Functional equations on finite sets	298
47.6	Functional equations over real numbers	298
47.7	Looking ahead	298
47.8	Functions as morphisms	299
47.9	Images, preimages, and logic	299
47.10	Functional equations as constraints on orbits	299
47.11	Functions as structure-preserving maps	299
47.12	Morphisms and structural information loss	299
47.13	Image and preimage as an adjunction	299
47.14	Finite functional graphs	300
47.15	Returning to the original examples	300
47.16	Injections, surjections, and factors	300
47.17	Fibers as local data of a function	300
N-20230117		301
48.1	Continuing the function chapter	301
48.2	Piecewise functions	301

48.3	Periodicity from functional equations	301
48.4	Graph transformations	302
48.5	Image of an interval	302
48.6	Functional equations and substitution	302
48.7	Counting functions between finite sets	302
48.8	A compact bridge	303
48.9	Piecewise functions as glued maps	303
48.10	Periodicity from generated transformations	303
48.11	Images of intervals	303
48.12	Local definitions and global function behavior	303
48.13	Functional equations as dynamics	303
48.14	Cauchy's equation and regularity	303
48.15	Symbolic dynamics	304
48.16	Rigidity from regularity	304
48.17	Piecewise functions and partitions	304
48.18	Functional equations and invariants	305
N-20230316		307
49.1	Fixed points	307
49.2	Iteration	307
49.3	Metric spaces and contractions	307
49.4	Banach fixed point theorem	308
49.5	Error estimate	308
49.6	Applications	309
49.7	Beyond Banach	309
49.8	What the example reveals	309
49.9	Why completeness is the hidden hypothesis	309
49.10	Picard iteration as fixed-point theory	310
49.11	Source repair: non-contractive fixed-point theorems	310
49.12	Source repair: Brouwer and Schauder	310

Proof Strategies, Cyclic Sums, and Inequality Methods

N-20220417

§35.1 Proof strategies

The note begins with a meta-level reflection: when solving a problem, one should distinguish between “find all” and “prove all.” A complete solution often has two halves:

- (i) find the candidates;
- (ii) prove that these and only these candidates work.

For example, to solve

$$3x + 2 = 17,$$

one first finds $x = 5$, then verifies that $x = 5$ indeed works and no other value can.

A more olympiad-style example is the problem of finding all positive integers n such that

$$2^n + 12^n + 2011^n$$

is a perfect square. The note uses modular arithmetic to narrow the possibilities: if $n > 1$ is odd, one obtains a contradiction modulo 4; if $n > 1$ is even, one obtains a contradiction modulo 3. Thus only $n = 1$ remains.

The page lists several general proof moves:

- (i) prove both directions when the statement is iff;
- (ii) for optimization, prove a bound and prove that equality can occur;
- (iii) for existence, show an object can be constructed;
- (iv) for minimality, exhibit a candidate and show anything smaller fails.

§35.2 Cyclic and symmetric sums

For three variables, cyclic sums and symmetric sums differ. The cyclic sum runs over cyclic permutations:

$$\begin{aligned}\sum_{\text{cyc}} a^2 &= a^2 + b^2 + c^2, \\ \sum_{\text{cyc}} a^2 b &= a^2 b + b^2 c + c^2 a.\end{aligned}$$

The symmetric sum runs over all distinct permutations:

$$\begin{aligned}\sum_{\text{sym}} a^2 &= a^2 + b^2 + c^2, \\ \sum_{\text{sym}} a^2 b &= a^2 b + a^2 c + b^2 a + b^2 c + c^2 a + c^2 b.\end{aligned}$$

This notation greatly shortens inequality statements.

§35.3 AM–GM

For nonnegative numbers $a_1, \dots, a_n$, the arithmetic-geometric mean inequality is

$$\frac{a_1 + \dots + a_n}{n} \geq \sqrt[n]{a_1 \cdots a_n}.$$

A Canadian Olympiad example in the note is

$$\frac{a^3}{bc} + \frac{b^3}{ca} + \frac{c^3}{ab} \geq a + b + c$$

for $a, b, c > 0$. By AM–GM,

$$\frac{1}{4} \left(2\frac{a^3}{bc} + \frac{b^3}{ca} + \frac{c^3}{ab} \right) \geq a.$$

Repeating cyclically and adding gives the desired inequality.

§35.4 Muirhead's inequality

For nonnegative variables, Muirhead's inequality compares symmetric sums using majorization. If

$$(x_1, \dots, x_n) \succ (y_1, \dots, y_n),$$

then

$$\sum_{\text{sym}} a_1^{x_1} \cdots a_n^{x_n} \geq \sum_{\text{sym}} a_1^{y_1} \cdots a_n^{y_n}.$$

For three variables,

$$(3, 0, 0) \succ (2, 1, 0),$$

so

$$a^3 + b^3 + c^3 \geq \sum_{\text{cyc}} a^2b$$

is a typical Muirhead-type comparison after normalization of symmetric sums.

§35.5 Power means

The power mean of order r is

$$P(r) = \begin{cases} \left(\frac{a_1^r + \dots + a_n^r}{n} \right)^{1/r}, & r \neq 0, \\ \sqrt[n]{a_1 \cdots a_n}, & r = 0. \end{cases}$$

The power mean inequality says that $P(r)$ is increasing in r . The familiar chain is

$$\sqrt{\frac{a_1^2 + \dots + a_n^2}{n}} \geq \frac{a_1 + \dots + a_n}{n} \geq \sqrt[n]{a_1 \cdots a_n} \geq \frac{n}{\frac{1}{a_1} + \dots + \frac{1}{a_n}}.$$

The note records several practical tricks:

- (i) if a condition says $abc = 1$, set $a = x/y$, $b = y/z$, $c = z/x$;
- (ii) if a, b, c are sides of a triangle, replace them by $x + y$, $y + z$, $z + x$;
- (iii) Schur's inequality often converts a cyclic sum into a more tractable expression.

§35.6 Hölder and Cauchy

Hölder's inequality in the form appearing in the page is

$$\left(\sum_{i=1}^n a_i\right)^p \left(\sum_{i=1}^n b_i\right)^q \geq \left(\sum_{i=1}^n \sqrt[p+q]{a_i^p b_i^q}\right)^{p+q}.$$

When $p = q = 1$, it reduces to Cauchy's inequality:

$$\left(\sum_i a_i\right) \left(\sum_i b_i\right) \geq \left(\sum_i \sqrt{a_i b_i}\right)^2.$$

The page also includes the example

$$\left(\sum_{\text{cyc}} \frac{a}{\sqrt{b+c}}\right)^2 \left(\sum_{\text{cyc}} a(b+c)\right) \geq (a+b+c)^3,$$

which is a typical Hölder move: choose weights so that the radical disappears.

§35.7 Jensen and Karamata

For a convex function f , Jensen's inequality says

$$\frac{f(a_1) + \cdots + f(a_n)}{n} \geq f\left(\frac{a_1 + \cdots + a_n}{n}\right).$$

For a concave function, the inequality reverses. Taking $f(x) = \ln x$ on $(0, \infty)$ recovers AM–GM:

$$\frac{\ln a_1 + \cdots + \ln a_n}{n} \leq \ln\left(\frac{a_1 + \cdots + a_n}{n}\right).$$

Karamata's inequality states that if f is convex and

$$(x_1, \dots, x_n) \succ (y_1, \dots, y_n),$$

then

$$f(x_1) + \cdots + f(x_n) \geq f(y_1) + \cdots + f(y_n).$$

Muirhead is, in spirit, a multiplicative/symmetric-sum analogue of majorization; Jensen and Karamata are the functional version.

§35.8 Cyclic sums

Cyclic notation compresses repeated symmetric expressions. For variables a, b, c , a cyclic sum means

$$\sum_{\text{cyc}} f(a, b, c) = f(a, b, c) + f(b, c, a) + f(c, a, b).$$

It is weaker than a full symmetric sum but often exactly matches triangle or inequality problems where the variables play rotating roles.

§35.9 Proof strategy before computation

Many inequality problems become easier once one identifies the likely equality case. If equality occurs at $a = b = c$, then substitutions that separate average and deviations are natural. If equality occurs at a boundary, monotonicity or uvw-style methods may be better. Guessing the equality case is not cheating; it guides the proof search.

§35.10 Next viewpoint

The note is really about choosing the right inequality language. AM–GM is local and elementary; Muirhead is symmetric and algebraic; Hölder is functional and norm-like; Jensen and Karamata are convex-geometric. A good inequality proof often consists of recognizing which language matches the structure of the expression.

§35.11 Inequalities as convexity statements

Many classical inequalities are consequences of convexity. Jensen's inequality says that for convex f ,

$$f\left(\sum_i \lambda_i x_i\right) \leq \sum_i \lambda_i f(x_i), \quad \lambda_i \geq 0, \quad \sum_i \lambda_i = 1.$$

AM–GM follows by applying concavity of $\log x$, or convexity of $-\log x$. Power mean inequalities follow from comparing convex powers.

§35.12 Majorization and Karamata

Muirhead and Karamata are organized by majorization. A vector x majorizes y when partial sums of decreasing rearrangements of x dominate those of y , with equal total sum. Karamata says convex functions respect this order:

$$x \succ y \quad \Rightarrow \quad \sum_i f(x_i) \geq \sum_i f(y_i).$$

This turns many symmetric inequalities into statements about how spread out a vector is.

§35.13 Holder as norm geometry

Holder's inequality

$$\sum_i |a_i b_i| \leq \left(\sum_i |a_i|^p\right)^{1/p} \left(\sum_i |b_i|^q\right)^{1/q}, \quad \frac{1}{p} + \frac{1}{q} = 1,$$

is a statement about dual norms. Cauchy–Schwarz is the case $p = q = 2$, where inner-product geometry gives orthogonal projection.

Thus inequalities are not isolated tricks; they describe geometry of convex bodies and normed spaces.

§35.14 Convexity, duality, and equality cases

AM–GM, Cauchy, Holder, Jensen, Muirhead, and Karamata form a hierarchy. The elementary technique is choosing the right inequality; the advanced structure is convexity, duality, majorization, and equality-case geometry.

§35.15 Inequality methods as choice of language

The inequality section should not be read as a list of weapons. Each inequality method corresponds to a different kind of structure. AM–GM sees products inside sums. Cauchy and Hölder see norms and inner-product-like expressions. Jensen sees convexity. Muirhead sees symmetric homogeneous polynomials through majorization.

A good contest solution usually begins before the first line of algebra: one asks which structure is already present. If the expression is symmetric and homogeneous, Muirhead or uvw may be natural. If there are denominators, Cauchy or Titu Andreescu’s lemma may fit. If a logarithm or power function appears, Jensen may be the real explanation.

§35.16 Convexity behind Jensen and Karamata

Many inequalities in the original note are different faces of convexity. A function f is convex on an interval if

$$f(tx + (1-t)y) \leq tf(x) + (1-t)f(y), \quad 0 \leq t \leq 1.$$

Jensen’s inequality says that for weights $w_i \geq 0$ with $\sum w_i = 1$,

$$f\left(\sum_i w_i x_i\right) \leq \sum_i w_i f(x_i).$$

AM–GM is Jensen applied to $f(x) = -\log x$, or equivalently to the concavity of $\log x$.

Karamata’s inequality refines Jensen through majorization. If x majorizes y , written $x \succ y$, then for convex f ,

$$\sum_i f(x_i) \geq \sum_i f(y_i).$$

This explains why many olympiad inequalities ask one to compare how “spread out” two lists are.

§35.17 Schur convexity

A symmetric differentiable function $F(x_1, \dots, x_n)$ is Schur-convex if

$$x \succ y \Rightarrow F(x) \geq F(y).$$

A useful criterion is

$$(x_i - x_j) \left(\frac{\partial F}{\partial x_i} - \frac{\partial F}{\partial x_j} \right) \geq 0$$

for all i, j . This turns some inequality problems into checking monotonicity of partial derivatives.

For example, power sums

$$F_p(x) = \sum_i x_i^p$$

are Schur-convex for $p \geq 1$ on the positive orthant. This explains why more uneven distributions increase higher powers.

§35.18 Entropy and the log-sum inequality

The log-sum inequality states

$$\sum_i a_i \log \frac{a_i}{b_i} \geq \left(\sum_i a_i \right) \log \frac{\sum_i a_i}{\sum_i b_i}, \quad a_i, b_i > 0.$$

It is Jensen's inequality in disguise. In information theory, it implies nonnegativity of relative entropy:

$$D(p||q) = \sum_i p_i \log \frac{p_i}{q_i} \geq 0.$$

This shows how elementary convexity inequalities become the mathematical foundation of entropy, statistics, and variational inference.

Nonlinear Recurrences, Fixed Points, and Metric-Space Exercises

N-20220718

§36.1 Fixed-point iteration

The note begins with iterative sequences. If a function f has a fixed point x_0 , meaning

$$f(x_0) = x_0,$$

then the recurrence

$$a_{n+1} = f(a_n)$$

may converge to x_0 , depending on the behavior of f . The graphical picture is the intersection of $y = f(x)$ and $y = x$.

The elementary idea is this: if $a_n \rightarrow L$ and f is continuous, then

$$L = \lim a_{n+1} = \lim f(a_n) = f(L),$$

so any limit must be a fixed point. But finding a fixed point is not the same as proving convergence. A recurrence can have fixed points and still oscillate or diverge.

§36.2 Möbius recurrences

The page considers recurrences of the form

$$a_{n+1} = \frac{pa_n + q}{ra_n + t}.$$

These are fractional linear or Möbius transformations. Their fixed points solve

$$x = \frac{px + q}{rx + t},$$

so

$$rx^2 + (t - p)x - q = 0.$$

Suppose x_1, x_2 are two distinct fixed points. Then a standard calculation gives

$$\frac{a_{n+1} - x_1}{a_{n+1} - x_2} = \lambda \frac{a_n - x_1}{a_n - x_2}$$

for a constant λ . Thus the new sequence

$$b_n = \frac{a_n - x_1}{a_n - x_2}$$

is geometric.

This is the key idea in the note: convert a nonlinear recurrence into a geometric progression by using the fixed points as coordinates. This is not accidental. Möbius transformations act naturally on the projective line, and the cross-ratio-like expression above is the coordinate that linearizes the action.

§36.3 Example

For

$$a_1 = \frac{1}{2}, \quad a_{n+1} = \frac{5a_n - 1}{a_n + 3},$$

the fixed-point equation is

$$x = \frac{5x - 1}{x + 3},$$

so

$$x^2 - 2x + 1 = 0, \quad x = 1.$$

Here the two fixed points coincide, so the previous method degenerates. One then studies

$$a_{n+1} - 1 = \frac{4a_n - 4}{a_n + 3} = \frac{4(a_n - 1)}{a_n + 3}.$$

Taking reciprocals gives

$$\frac{1}{a_{n+1} - 1} = \frac{a_n + 3}{4(a_n - 1)} = \frac{1}{4} + \frac{1}{a_n - 1}.$$

Thus

$$b_n = \frac{1}{a_n - 1}$$

is arithmetic, not geometric. This is the repeated-root analogue of the fixed-point method.

§36.4 Second-order linear recurrences

The next topic is

$$a_{n+2} = pa_{n+1} + qa_n.$$

Try $a_n = \lambda^n$. The characteristic equation is

$$\lambda^2 - p\lambda - q = 0.$$

If it has two distinct roots λ_1, λ_2 , then

$$a_n = A\lambda_1^n + B\lambda_2^n.$$

If the root is repeated, then the second independent solution is $n\lambda^n$, so

$$a_n = (A + Bn)\lambda^n.$$

This is the recurrence analogue of solving a linear differential equation with constant coefficients.

§36.5 Dedekind cuts and metric exercises

The uploaded pages also continue exercises on Dedekind cuts and metric spaces. For Dedekind cuts X, Y , the sum is

$$X + Y = \{x + y : x \in X, y \in Y\}.$$

The notes check that this is again a cut: it is nonempty, not all of $\mathbb{Q}$, downward closed, and has no largest rational.

For a metric space (X, d) , if $Y \subseteq X$, the subspace metric is

$$d_Y(y_1, y_2) = d(y_1, y_2).$$

The note also checks that Euclidean distance on $\mathbb{R}^n$

$$d(x, y) = \sqrt{(x_1 - y_1)^2 + \cdots + (x_n - y_n)^2}$$

is a metric. The only nontrivial axiom is the triangle inequality, which ultimately follows from Cauchy–Schwarz.

§36.6 Fixed points of recurrence maps

A recurrence $x_{n+1} = f(x_n)$ is an iteration problem. A possible limit L must satisfy

$$L = f(L).$$

Thus the first step is to solve the fixed-point equation. The second step is to decide whether the iteration actually converges to that fixed point.

A common local test is $|f'(L)| < 1$. Then nearby points are pulled toward L . If $|f'(L)| > 1$, nearby points are pushed away. This explains why fixed points are not automatically limits.

§36.7 Monotone bounded sequences

Many elementary recurrence proofs use the theorem: a monotone bounded real sequence converges. The method is to prove by induction that the sequence stays in an interval, then prove monotonicity. Completeness of $\mathbb{R}$ is the hidden reason the limit exists.

§36.8 After the examples

The important common thread is change of coordinates. A Möbius recurrence becomes geometric after measuring position relative to fixed points. A repeated fixed point becomes arithmetic after taking a reciprocal. A Dedekind cut becomes a real number after using order-theoretic structure. A subset becomes a metric space after inheriting distance. The mathematical move is the same: choose the representation in which the structure becomes visible.

§36.9 Fixed points and stability

For an iteration $x_{n+1} = F(x_n)$, a fixed point satisfies $F(x_*) = x_*$. The behavior near x_* is controlled by $F'(x_*)$ in one dimension: if $|F'(x_*)| < 1$, the fixed point is locally attracting.

§36.10 Möbius recurrences and matrices

A recurrence of the form

$$x_{n+1} = \frac{ax_n + b}{cx_n + d}$$

is a fractional-linear transformation. It corresponds to the matrix

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

acting on the projective line. Iterating the recurrence corresponds to taking powers of this matrix.

§36.11 Second-order recurrences

A linear recurrence

$$a_{n+2} = pa_{n+1} + qa_n$$

is governed by the characteristic equation

$$\lambda^2 = p\lambda + q.$$

Distinct roots give combinations of geometric sequences. This is diagonalization of the shift dynamics.

§36.12 Recurrences as one-dimensional dynamics

Nonlinear and linear recurrences are dynamical systems. Fixed points organize stability, Möbius recurrences are projective matrix dynamics, and second-order linear recurrences are solved by eigenvalue methods.

Probability, Conditional Probability, and Substitution Tricks in Inequalities

N-20220721

§37.1 Random experiments and sample spaces

This note begins with the language of probability. A random experiment is a process that can be repeated under the same conditions, whose possible outcomes are known, but whose actual outcome is uncertain. Examples include tossing a coin, drawing a ball from a box, or choosing a card.

The set of all possible basic outcomes is the sample space, denoted

$$\Omega.$$

A subset $A \subseteq \Omega$ is an event. If the sample space is finite and all outcomes are equally likely, then

$$P(A) = \frac{|A|}{|\Omega|}.$$

The note stresses that events are sets. This is why union, intersection, complement, and inclusion–exclusion immediately become probability tools:

$$P(A \cup B) = P(A) + P(B) - P(A \cap B).$$

§37.2 Classical model and Venn diagrams

One example has seven labeled balls, with three properties represented by sets A_1, A_2, A_3 . The note uses a Venn diagram and factorial counts to compute a probability such as

$$P = \frac{7! - 3 \cdot 6! + 3 \cdot 5! - 4!}{7!}.$$

The exact numbers matter less than the method. If conditions overlap, count by inclusion–exclusion. A Venn diagram is not a proof by itself, but it helps prevent double-counting.

§37.3 Conditional probability

The conditional probability of B given A is

$$P(B \mid A) = \frac{P(A \cap B)}{P(A)}$$

when $P(A) > 0$. The note includes examples where one must be careful not to confuse $P(B \mid A)$ with $P(B)$.

Independence is the condition

$$P(A \cap B) = P(A)P(B),$$

or equivalently

$$P(B \mid A) = P(B).$$

This means that knowing A occurred does not change the probability of B .

A common trap appears in nested events. If $A \subseteq B$, then

$$P(B \mid A) = 1,$$

but

$$P(A \mid B) = \frac{P(A)}{P(B)}$$

may be much smaller. Conditional probability is not symmetric.

§37.4 A transition to substitution tricks

The note then changes direction and discusses substitutions used in inequalities. A recurring constraint is

$$abc = 1.$$

A standard substitution is

$$a = \frac{x}{y}, \quad b = \frac{y}{z}, \quad c = \frac{z}{x}.$$

Another useful substitution is

$$x = a + \frac{1}{a}, \quad y = b + \frac{1}{b}, \quad z = c + \frac{1}{c}.$$

Such substitutions are not just algebraic decoration. They encode hidden positivity and symmetry. The note points out that many complicated-looking inequalities become simpler after this change of variables.

§37.5 A classical symmetric equation

The pages discuss equations such as

$$xyz = x + y + z + 2.$$

This can be rewritten as

$$xyz + 1 = (x + 1)(y + 1)(z + 1) - xy - yz - zx - x - y - z.$$

A better-known substitution is

$$x = a + \frac{1}{a}, \quad y = b + \frac{1}{b}, \quad z = c + \frac{1}{c},$$

or trigonometric substitutions when the variables resemble

$$2 \cos A, \quad 2 \cos B, \quad 2 \cos C.$$

The note is circling around the idea that symmetric algebraic equations often become geometry or trigonometry after the right substitution.

§37.6 Inequality examples

One example asks to prove

$$\frac{a}{b+c+a} + \frac{b}{c+a+b} + \frac{c}{a+b+c} \geq \frac{3}{2}$$

in a normalized form. The exact expression on the page is a contest inequality, and the solution uses standard tools such as Cauchy, AM–GM, or substitution.

The point is that conditions like $abc = 1$, $x+y+z = \text{constant}$, or $x^2+y^2+z^2+2xyz = 1$ should not be treated as random constraints. They are invitations to choose a special parametrization.

§37.7 Jensen, trigonometry, and contest style

Later examples involve trigonometric inequalities in a triangle. A triangle condition $A + B + C = \pi$ naturally suggests identities such as

$$\cos^2 A + \cos^2 B + \cos^2 C + 2 \cos A \cos B \cos C = 1.$$

This identity is one of the reasons that substitutions by cosines appear in olympiad inequalities.

§37.8 The hidden structure

This note looks like probability followed by inequalities, but both topics are about choosing the right sample space. In probability, the sample space is literal: events are subsets of Ω . In inequalities, the “sample space” is the domain allowed by the constraint. A good substitution changes that domain into something more natural, just as choosing a good probability model turns a counting problem into a set problem.

§37.9 Conditional probability as changing measure

Conditional probability is not just a formula. It changes the universe from Ω to the event B :

$$\mathbb{P}(A \mid B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}.$$

The denominator renormalizes the measure so that B has probability 1.

This is why Venn diagrams help: conditioning restricts attention to one region and rescales its area.

§37.10 Bayes formula and inverse problems

Bayes’ formula

$$\mathbb{P}(A_i \mid B) = \frac{\mathbb{P}(B \mid A_i)\mathbb{P}(A_i)}{\sum_j \mathbb{P}(B \mid A_j)\mathbb{P}(A_j)}$$

turns likelihoods into posterior probabilities. It is a finite version of inference: update prior weights after observing evidence.

The formula is especially useful when B is easier to compute conditional on each case A_i than directly.

§37.11 Inequality substitutions as coordinate choices

The substitution tricks in the inequality part should be read as coordinate changes. Symmetric expressions often simplify under

$$u = x + y, \quad v = xy,$$

or under trigonometric substitutions that encode constraints. The good substitution is the one that turns the constraint into a simple domain.

§37.12 Changing measures and changing coordinates

Probability and inequality substitutions share a habit: change the viewpoint before computing. Conditioning changes the measure; Bayes reverses conditional information; substitutions choose coordinates adapted to the constraint. This last sentence is worth making more precise, because it is the higher-level bridge hidden inside the elementary examples.

Definition 37.12.1 (Conditioning as a new probability measure). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let $B \in \mathcal{F}$ satisfy $\mathbb{P}(B) > 0$. The conditional probability given B is the probability measure

$$\mathbb{P}_B(A) = \mathbb{P}(A \mid B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}.$$

Equivalently,

$$\mathbb{P}_B(A) = \int_A \frac{\mathbf{1}_B}{\mathbb{P}(B)} d\mathbb{P}.$$

Thus $\mathbb{P}_B$ is absolutely continuous with respect to $\mathbb{P}$, with Radon–Nikodym derivative

$$\frac{d\mathbb{P}_B}{d\mathbb{P}} = \frac{\mathbf{1}_B}{\mathbb{P}(B)}.$$

This is the clean meaning of the phrase “conditioning changes the measure”. The set Ω may stay formally the same, but the weights assigned to its parts have changed. Everything outside B receives weight zero, and the mass inside B is rescaled so that the total mass is again one.

Proposition 37.12.2 (Conditional expectation as integration in the new measure)

For every integrable random variable X ,

$$\mathbb{E}[X \mid B] = \int_{\Omega} X d\mathbb{P}_B = \frac{1}{\mathbb{P}(B)} \int_B X d\mathbb{P}.$$

Proof. By the definition of $\mathbb{P}_B$ and its density with respect to $\mathbb{P}$,

$$\int_{\Omega} X d\mathbb{P}_B = \int_{\Omega} X \frac{\mathbf{1}_B}{\mathbb{P}(B)} d\mathbb{P} = \frac{1}{\mathbb{P}(B)} \int_B X d\mathbb{P}.$$

This is precisely the usual formula for conditional expectation on the event B . $\square$

Bayes' formula is also clearer from this viewpoint. Suppose $A_1, \dots, A_n$ form a partition of Ω . After observing B , the old weights $\mathbb{P}(A_i)$ are replaced by

$$\mathbb{P}_B(A_i) = \frac{\mathbb{P}(B \mid A_i)\mathbb{P}(A_i)}{\sum_j \mathbb{P}(B \mid A_j)\mathbb{P}(A_j)}.$$

The numerator is the old weight multiplied by the likelihood of the observation. The denominator is just the normalization which makes the new weights add to one. So Bayes' formula is not a trick for reversing arrows; it is a reweighting rule.

The same pattern appears in substitutions for inequalities. A substitution changes how we integrate, compare, or optimize over a domain. If $T : U \rightarrow V$ is a change of variables, then a function on V can be pulled back to U by

$$f \mapsto f \circ T.$$

For integrals, the measure also changes through the Jacobian:

$$\int_V f(y) dy = \int_U f(T(x)) |\det DT(x)| dx.$$

For inequalities, even when no explicit integral is present, the same philosophy remains: choose coordinates in which the constraint and the symmetry become visible.

Remark 37.12.3 (The shared template) — Conditioning replaces $\mathbb{P}$ by the weighted measure $\mathbb{P}_B$. Bayes' formula replaces prior weights by posterior weights. A substitution replaces the original coordinates by coordinates adapted to the constraint. In all three cases, the computation becomes easier because the background viewpoint has changed before the algebra begins.

Set Exercises, Functions, and Elementary Problem-Solving

N-20220801

§38.1 Subsets and characteristic functions

The note begins with the idea of counting subsets. If a set A has n elements, then it has

$$2^n$$

subsets. One way to see this is to use the characteristic function

$$\chi_S(x) = \begin{cases} 1, & x \in S, \\ 0, & x \notin S. \end{cases}$$

Choosing a subset $S \subseteq A$ is the same as choosing a function

$$\chi_S : A \rightarrow \{0, 1\}.$$

There are 2 choices for each element of A , hence 2^n choices in total.

This is a useful conceptual improvement over simply memorizing 2^n . A subset is a yes/no decision for every element.

§38.2 Set laws

The page records De Morgan's laws:

$$C_U(A \cap B) = C_U A \cup C_U B,$$

$$C_U(A \cup B) = C_U A \cap C_U B.$$

It also recalls inclusion–exclusion:

$$|A \cup B| = |A| + |B| - |A \cap B|.$$

The exercises ask for set descriptions such as:

$$\{2n : 1 \leq n \leq 5\},$$

$$\{x : x \leq 5\},$$

and graphs such as

$$\{(x, y) : y = x^2 - x - 1\}.$$

This is a good reminder that sets can be finite lists, solution sets of inequalities, or geometric loci.

§38.3 Venn diagram exercises

One exercise gives

$$U = A \cup B = \{1, 2, 3, 4\}, \quad A \cap C_U B = \{1, 2\},$$

and asks for B . Since $A \setminus B = \{1, 2\}$, the remaining elements 3, 4 must lie in B , so

$$B = \{3, 4\}.$$

Another exercise has

$$A = \{x : 1 < x < 7\}, \quad B = \{x : x^2 - 4x - 5 \leq 0\}.$$

Since

$$x^2 - 4x - 5 = (x - 5)(x + 1),$$

we get

$$B = [-1, 5].$$

Therefore

$$A \cap C_{\mathbb{R}} B = (5, 7).$$

The note's drawings are not decorative. They are there to force the habit of separating “inside A , inside B , inside both, inside neither.”

§38.4 Inclusion–exclusion with literature examples

One problem describes a survey of 100 students reading classical novels. If 80 read either X or Y , 60 read X , and 40 read both, then the number reading Y is found by

$$|X \cup Y| = |X| + |Y| - |X \cap Y|.$$

Thus

$$80 = 60 + |Y| - 40,$$

so

$$|Y| = 60.$$

Again, the important point is not the story, but the relation among union, intersection, and individual counts.

§38.5 Set constructions with squares

The note includes exercises with sets such as

$$A = \{a : a = x^2 - y^2, \ x, y \in \mathbb{Z}\},$$

and

$$B = \{a : a = x^2 - y^2, \ x, y \in \mathbb{Q}\}.$$

Since

$$x^2 - y^2 = (x + y)(x - y),$$

every difference of squares is a product of two numbers of the same parity in the integer case. This explains why expressions such as

$$2k - 1$$

are easy to represent, while some even numbers require more care.

For rationals, the factorization is more flexible, and quotients of such expressions can often remain in B . The page is training the habit of factoring before trying random examples.

§38.6 Floor functions

A floor-function example studies values of

$$y = \lfloor x \rfloor + \lfloor 2x \rfloor + \lfloor 3x \rfloor, \quad 1 \leq y \leq 100.$$

The natural move is to write

$$x = k + t, \quad k \in \mathbb{Z}, \quad 0 \leq t < 1.$$

Then the three floor terms change only when t crosses multiples of $1/2$ or $1/3$. Hence the possible values in one period come from the intervals determined by

$$0, \quad \frac{1}{3}, \quad \frac{1}{2}, \quad \frac{2}{3}, \quad 1.$$

This is much cleaner than trying to graph the whole function at once.

§38.7 Functions and monotonicity

The later pages switch to functions. Several problems ask for domains, ranges, monotonicity, and maximum or minimum values. Typical tools include:

- (i) rewriting a square root or logarithm condition as a domain restriction;
- (ii) using monotonicity to locate maxima and minima;
- (iii) drawing a rough graph to understand intervals;
- (iv) using substitution when a composite expression repeats.

One example studies a piecewise or iterated function and asks how many cycles appear. The drawing in the note shows the value bouncing among intervals, which is an early form of dynamical thinking.

§38.8 Characteristic functions and subsets

A subset $S \subseteq A$ is the same data as a function

$$\chi_S : A \rightarrow \{0, 1\}.$$

Each element chooses either 0 or 1, so a set with n elements has 2^n subsets. This is more than a counting trick: it turns subsets into functions, which is the beginning of a very useful viewpoint.

§38.9 Functions as rules with domains

A function is not merely a formula. It consists of a domain, codomain, and assignment rule. The same formula can define different functions if the domain or codomain changes. This is why image, preimage, injectivity, and surjectivity must always be stated relative to a specified domain and codomain.

§38.10 The method behind the trick

This note is a collection of elementary set and function exercises, but the hidden agenda is representational flexibility. A set can be represented by a list, a predicate, a Venn diagram, or a characteristic function. A function can be represented by a formula, a graph, a domain/range description, or an iteration rule. Much of mathematics consists of switching to the representation in which the problem becomes visible.

§38.11 Set operations as Boolean algebra

For a fixed universe U , the power set $\mathcal{P}(U)$ is a Boolean algebra:

$$A \vee B = A \cup B, \quad A \wedge B = A \cap B, \quad \neg A = U \setminus A.$$

De Morgan's laws become algebraic identities:

$$\neg(A \vee B) = \neg A \wedge \neg B, \quad \neg(A \wedge B) = \neg A \vee \neg B.$$

Thus Venn-diagram manipulation is the visual form of Boolean algebra.

§38.12 Characteristic functions and algebra of predicates

A subset $A \subseteq U$ is equivalently a predicate or a characteristic function

$$\mathbf{1}_A : U \rightarrow \{0, 1\}.$$

Set operations become operations on 0-1 functions:

$$\mathbf{1}_{A \cap B} = \mathbf{1}_A \mathbf{1}_B, \quad \mathbf{1}_{A^c} = 1 - \mathbf{1}_A,$$

and

$$\mathbf{1}_{A \cup B} = \mathbf{1}_A + \mathbf{1}_B - \mathbf{1}_A \mathbf{1}_B.$$

For finite U , cardinality is summation:

$$|A| = \sum_{x \in U} \mathbf{1}_A(x).$$

This converts set identities into ordinary algebraic identities.

§38.13 Preimages as Boolean homomorphisms

For a function $f : X \rightarrow Y$, the preimage operation

$$f^{-1} : \mathcal{P}(Y) \rightarrow \mathcal{P}(X)$$

preserves Boolean operations:

$$f^{-1}(B \cup C) = f^{-1}(B) \cup f^{-1}(C),$$

$$f^{-1}(B \cap C) = f^{-1}(B) \cap f^{-1}(C), \quad f^{-1}(Y \setminus B) = X \setminus f^{-1}(B).$$

Images do not preserve intersections or complements in this exact way. This asymmetry explains why topology defines continuity using preimages of open sets.

§38.14 Sigma-algebras and measurable events

Probability does not always use the full power set. A sigma-algebra $\mathcal{A}$ is a family of subsets closed under complement and countable union. Its elements are measurable events. A probability measure assigns numbers to these events so that disjoint unions add:

$$P\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} P(A_n).$$

Thus the elementary set operations $\cup, \cap, ^c$ become the grammar of probability and measure theory.

§38.15 Stone duality intuition

Stone duality says that Boolean algebras can be represented as algebras of clopen subsets of topological spaces. Given a Boolean algebra B , its Stone space consists of ultrafilters on B . Each element $b \in B$ determines a clopen set

$$\widehat{b} = \{u : b \in u\}.$$

This shows that elementary set algebra, propositional logic, and certain zero-dimensional topological spaces are different faces of one structure.

§38.16 Difference of squares as algebraic structure

The exercise involving

$$x^2 - y^2 = (x + y)(x - y)$$

is a reminder that set-builder notation often hides algebraic structure. Over the integers, the two factors $x + y$ and $x - y$ have the same parity. Over the rationals, the parity obstruction disappears. The same formula therefore defines different sets depending on the ambient number system.

So even elementary set-builder problems force a structural question: what universe are the variables living in?

§38.17 Sets as algebra, logic, and measurement

The note's set and function exercises are unified by representation. A subset may be a list, predicate, characteristic function, measurable event, or Boolean-algebra element. A function is not just a formula but a map with domain and codomain, and its preimage operation is the structurally stable one. Most mistakes in these exercises come from forgetting which representation is being used.

§38.18 Boolean algebra beyond Venn diagrams

The original exercises with unions, intersections, complements, and Venn diagrams are the concrete beginning of Boolean algebra. A Boolean algebra is a structure with operations

$$a \vee b, \quad a \wedge b, \quad \neg a, \quad 0, \quad 1$$

satisfying the familiar laws of union, intersection, complement, empty set, and universal set. For a fixed universe U , the power set $\mathcal{P}(U)$ is the model example:

$$A \vee B = A \cup B, \quad A \wedge B = A \cap B, \quad \neg A = U \setminus A.$$

The elementary identities

$$A \cap (B \cup C) = (A \cap B) \cup (A \cap C), \quad (A \cup B)^c = A^c \cap B^c$$

are therefore not isolated set tricks. They are algebraic laws in a Boolean algebra.

§38.19 Characteristic functions

Every subset $A \subseteq U$ has a characteristic function

$$\mathbf{1}_A : U \rightarrow \{0, 1\}, \quad \mathbf{1}_A(x) = \begin{cases} 1, & x \in A, \\ 0, & x \notin A. \end{cases}$$

Set operations become arithmetic or logical operations:

$$\mathbf{1}_{A \cap B} = \mathbf{1}_A \mathbf{1}_B, \quad \mathbf{1}_{A \cup B} = \max(\mathbf{1}_A, \mathbf{1}_B), \quad \mathbf{1}_{A^c} = 1 - \mathbf{1}_A.$$

For finite sets,

$$|A| = \sum_{x \in U} \mathbf{1}_A(x).$$

This converts Venn-diagram counting into algebra. Inclusion–exclusion is just the expansion of a product of factors $1 - \mathbf{1}_{A_i}$.

§38.20 Sigma-algebras as observable event systems

In probability and measure theory, one does not always allow every subset of a universe. A sigma-algebra $\mathcal{A}$ on U is a collection of subsets satisfying

$$U \in \mathcal{A}, \quad A \in \mathcal{A} \Rightarrow A^c \in \mathcal{A}, \quad A_1, A_2, \dots \in \mathcal{A} \Rightarrow \bigcup_{n=1}^{\infty} A_n \in \mathcal{A}.$$

This is a Boolean algebra closed under countable unions. The elements of $\mathcal{A}$ are the measurable sets or events.

The elementary set laws still hold, but the important new question is: which subsets are allowed to be measured? Probability begins by assigning a number $P(A)$ to each $A \in \mathcal{A}$ such that

$$P\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} P(A_n)$$

for disjoint events. Thus basic set operations become the grammar of probability.

§38.21 Stone duality as topology of logic

Stone duality says, roughly, that every Boolean algebra can be represented as an algebra of clopen subsets of a topological space. For a Boolean algebra B , its Stone space consists of ultrafilters on B . Each element $b \in B$ determines a set

$$\hat{b} = \{u : b \in u\}.$$

These sets behave like the subsets of a space.

This tells us that elementary set algebra is not merely a notation for collections. It can encode topology. Logical operations, set operations, and certain topological spaces are three faces of the same structure.

§38.22 Returning to the original set exercises

The original exercises train manipulation of $\cup, \cap, ^c, \setminus$. At the advanced level, the same manipulations become Boolean algebra, characteristic-function algebra, measurable-event algebra, and even topological representation theory through Stone duality. Nothing in the elementary calculations is wasted: every identity is a structural law that survives into probability, logic, and topology.

§38.23 From Venn diagrams to algebra of information

A Venn diagram is a visual partition of a universe into atoms determined by a finite family of sets. For two sets, the atoms are

$$A \cap B, \quad A \cap B^c, \quad A^c \cap B, \quad A^c \cap B^c.$$

For n sets, there are at most 2^n atoms. Every Boolean expression in the sets is a union of some of these atoms.

This gives a systematic method for elementary set identities: reduce both sides to the same atoms. In logic, the same method is truth tables. In probability, the atoms are elementary events. In database theory, they are the records satisfying combinations of predicates.

Thus the elementary operation of shading regions is already the algebra of information partitions. A set is a yes/no question about elements; a finite family of sets is a finite collection of yes/no questions; the atoms are the complete answer patterns.

§38.24 Why sigma-algebras restrict what can be asked

In finite examples one usually allows every subset. In analysis, allowing every subset of $\mathbb{R}$ is incompatible with countably additive length if one accepts the usual choice principles. Sigma-algebras solve this by specifying which questions are measurable.

The elementary complement and union laws are retained, but countability enters:

$$A_1, A_2, \dots \in \mathcal{A} \implies \bigcup_{n=1}^{\infty} A_n \in \mathcal{A}.$$

This is exactly what is needed to discuss limiting events such as “infinitely many of these events occur.” The original finite set exercises therefore become the finite visible model of measurable-event logic.

Set Theory Exercises, Closures, Functions, and Algebraic Side Notes

N-20220806

§39.1 Images, preimages, and finite set examples

The first page contains set-theoretic exercises. For example, let

$$A = \{2, 0, 1, 3\}$$

and define

$$B = \{x : x \in A, 2 - x^2 \notin A\}.$$

Checking each element:

$$x = -2 \Rightarrow 2 - x^2 = -2 \notin A,$$

$$x = 0 \Rightarrow 2 - x^2 = 2 \in A,$$

$$x = -1 \Rightarrow 2 - x^2 = 1 \in A,$$

$$x = -3 \Rightarrow 2 - x^2 = -7 \notin A.$$

The page concludes with a finite set B after direct inspection. The point is that a set defined by a condition is not mysterious: test the condition element by element when the universe is small.

§39.2 Closure under addition

The next exercise asks for a subset $S \subseteq \mathbb{R}$ closed under addition:

$$x, y \in S \quad \Rightarrow \quad x + y \in S.$$

Examples include:

$$\mathbb{Z}, \quad \mathbb{Q}, \quad \mathbb{R}, \quad [0, \infty).$$

A finite example such as $\mathbb{Z}/6\mathbb{Z}$ also works if addition is understood modulo 6. The note correctly distinguishes ordinary addition from addition modulo an equivalence relation.

A related problem says: if $S_1, S_2 \subseteq \mathbb{R}$ are closed under addition and are proper in a suitable sense, prove there is an element c outside their union. This resembles the Dedekind-cut intuition: two proper additive-closed pieces cannot cover the entire real line unless one of them already absorbs too much structure.

§39.3 A combinatorial maximal-set problem

Another problem considers

$$M \subseteq \{1, 2, \dots, 2011\}$$

with the property that for any two elements $a, b \in M$, at least one of

$$a \mid b, \quad b \mid a, \quad |a - b| = \max M$$

holds. The note first explores powers of 2:

$$1, 2, 2^2, \dots, 2^{10},$$

since divisibility is automatic along a chain. It then observes that if one tries to add a number not on the same divisibility chain, the condition quickly fails.

The key structural idea is that divisibility partially orders the integers. A set in which every pair is comparable under divisibility is a chain. Powers of a fixed prime give the cleanest example of a long chain. This is an early encounter with posets and extremal combinatorics.

§39.4 Functional and algebraic exercises

The later pages mix several algebraic problems. One recurring theme is solving equations by rewriting them into factored or symmetric forms. For instance, equations involving

$$x^2 + y^2, \quad xy, \quad x + y$$

are usually better handled through substitutions like

$$u = x + y, \quad v = xy.$$

This is the elementary symmetric-polynomial viewpoint.

Another page discusses functions and constraints of the form

$$f(A \cap B), \quad f(A \cup B), \quad f^{-1}(C).$$

The essential warning is that images and preimages behave differently:

$$f^{-1}(C \cap D) = f^{-1}(C) \cap f^{-1}(D)$$

always, while

$$f(A \cap B) = f(A) \cap f(B)$$

need not hold unless f is injective on the relevant sets.

§39.5 Matrix and coordinate side computations

Several pages contain coordinate computations and small matrix reductions. These are not central enough to deserve separate chapters, but they reinforce a common habit: reduce a system to normal form. Whether one is checking membership in a set, solving a function equation, or reducing a matrix, the practical goal is to isolate independent parameters.

§39.6 Images and preimages behave differently

For a function $f : X \rightarrow Y$, the image of a set $A \subseteq X$ is

$$f(A) = \{f(a) : a \in A\},$$

while the preimage of $B \subseteq Y$ is

$$f^{-1}(B) = \{x \in X : f(x) \in B\}.$$

Preimages preserve unions, intersections, and complements exactly. Images preserve unions but not intersections in general. This asymmetry is why topology defines continuity using preimages of open sets.

§39.7 Closure operations

Many set exercises secretly ask whether an operation is a closure operator. A closure operator should be extensive, monotone, and idempotent:

$$A \subseteq \overline{A}, \quad A \subseteq B \Rightarrow \overline{A} \subseteq \overline{B}, \quad \overline{\overline{A}} = \overline{A}.$$

Recognizing these three properties helps organize otherwise scattered set-manipulation problems.

§39.8 The reusable pattern

This note is scattered, but its center is closure. A set may be closed under addition, a family of elements may be closed under a divisibility relation, a function preimage is closed under Boolean operations, and a solution set may be closed under linear combinations. Recognizing closure is one of the first steps toward algebraic structure.

§39.9 Closure operators

Many exercises in this note ask whether a set is stable under an operation. The abstract pattern is a closure operator. On a partially ordered set P , a closure operator is a map

$$c : P \rightarrow P$$

satisfying

$$x \leq c(x), \quad x \leq y \Rightarrow c(x) \leq c(y), \quad c(c(x)) = c(x).$$

These are extensivity, monotonicity, and idempotence.

Examples include topological closure $A \mapsto \overline{A}$, linear span $S \mapsto \text{span}(S)$, subgroup generation $S \mapsto \langle S \rangle$, and logical consequence from assumptions.

§39.10 Generated objects as intersections

When a problem asks for the smallest object containing S and closed under some rule, the general construction is

$$\langle S \rangle = \bigcap \{C : S \subseteq C, C \text{ is closed under the rule}\}.$$

This works because intersections of closed objects are closed in many common settings: intersections of subgroups are subgroups, intersections of subspaces are subspaces, and arbitrary intersections of closed sets are closed.

Thus "generated by" is not a vague phrase. It means minimality under inclusion plus closure under specified operations.

§39.11 Fixed points of closure

An object is closed when it is a fixed point of the closure operator:

$$c(x) = x.$$

The closed objects form a structured family. For example, topologically closed subsets are exactly fixed points of the topological closure operator; subspaces are fixed points of span; subgroups are fixed points of subgroup generation.

This viewpoint also explains iterative procedures. Start with a set, repeatedly add everything forced by the rule, and stop when no new elements appear. In finite settings this must stabilize; in infinite settings one may need transfinite or lattice-theoretic methods.

§39.12 Images, preimages, and logic

For $f : X \rightarrow Y$, preimage preserves all Boolean operations, while image only preserves unions in general. The failure

$$f(A \cap B) \subseteq f(A) \cap f(B)$$

can be strict: two different points may have the same image. Equality for all subsets is closely related to injectivity.

This is the set-theoretic reason that preimages are central in topology, measure theory, and logic. A predicate on Y pulls back to a predicate on X , and logical connectives are preserved exactly.

§39.13 Posets and extremal chains

The maximal-set problem based on divisibility is naturally a problem in a partially ordered set. Write

$$a \preceq b \iff a \mid b.$$

A family in which every pair is comparable is a chain. Powers of a fixed prime form a chain:

$$1, p, p^2, p^3, \dots$$

Extremal questions about such families lead toward Dilworth-type and Sperner-type theorems, where one bounds the size of chains or antichains in a poset.

§39.14 Galois connections

A Galois connection translates order conditions between two posets. One common form is

$$f(p) \leq q \iff p \leq g(q).$$

Composites of adjoint maps often produce closure operators. In linear algebra, taking annihilators reverses inclusion and double annihilators produce closure-like operations.

This explains why constraints and solution sets are dual. A set of equations determines a solution set; a solution set determines equations vanishing on it. Applying both directions gives a closure operation.

§39.15 Tarski fixed-point theorem

On a complete lattice, every monotone map has fixed points. Tarski's theorem says the set of fixed points is itself a complete lattice. This is the abstract version of repeatedly applying a monotone rule until it stabilizes.

The elementary finite-set exercises show the finite shadow of this theorem: if a process only adds elements inside a finite universe, it eventually reaches a fixed point.

§39.16 Closure, order, and fixed-point thinking

The scattered exercises are unified by closure and order. Sets closed under addition, divisibility chains, preimage stability, generated substructures, and repeated operations that stabilize are all examples of order-theoretic thinking. The advanced structure is not decoration; it explains why the same proof patterns appear across algebra, topology, logic, and combinatorics.

§39.17 Closure operators across mathematics

The closure-style exercises in the original note can be organized by the notion of a closure operator. On a partially ordered set P , a closure operator is a map

$$c : P \rightarrow P$$

satisfying

$$x \leq c(x), \quad x \leq y \Rightarrow c(x) \leq c(y), \quad c(c(x)) = c(x).$$

These are called extensivity, monotonicity, and idempotence.

Many familiar constructions have this form:

$$\begin{aligned} A &\mapsto \overline{A} && \text{in topology,} \\ S &\mapsto \text{span}(S) && \text{in linear algebra,} \\ S &\mapsto \langle S \rangle && \text{in group theory.} \end{aligned}$$

The elementary operation “generate the smallest object containing a given set” is exactly closure.

§39.18 Closed sets as fixed points

An element x is closed if

$$c(x) = x.$$

Thus closed sets are fixed points of the closure operator. The family of closed elements is closed under arbitrary meets. In set language, topologically closed sets are stable under arbitrary intersections; subspaces generated by vector sets are stable under intersections; subgroups generated by sets are stable under intersections.

This explains a standard construction:

$$c(S) = \bigcap \{C : S \subseteq C, C \text{ closed}\}.$$

The closure of S is the intersection of all closed objects containing it.

§39.19 Galois connections and closure

A Galois connection between posets P and Q is a pair of order-reversing or order-preserving maps that translate inequalities in one side into inequalities on the other. A common form is

$$f(p) \leq q \iff p \leq g(q).$$

The composite $g \circ f$ is then a closure operator.

In linear algebra, for a subset S of vectors, one can define the annihilator

$$S^\perp = \{\varphi \in V^* : \varphi(s) = 0 \text{ for all } s \in S\}.$$

Taking annihilators reverses inclusion and creates closure-like double-perp operations.

Thus the elementary habit of taking complements, generated sets, or common solution sets points toward a deep pattern: objects can be studied through the constraints they impose, and closure is often a double-constraint operation.

§39.20 Tarski fixed points and monotone iteration

On a complete lattice, every monotone map has a fixed point. This is Tarski's fixed-point theorem. A complete lattice is a partially ordered set in which every subset has a supremum and an infimum.

For a monotone map $F : L \rightarrow L$, the set of fixed points

$$\{x \in L : F(x) = x\}$$

is itself a complete lattice. This theorem appears in logic, semantics of programming languages, and inductive definitions.

The elementary idea of repeatedly applying a closure process until nothing changes is therefore a finite visible version of a fixed-point theorem for ordered structures.

§39.21 Returning to the original exercises

Whenever the original note asks for a smallest set satisfying some property, an invariant collection, or a repeated operation that stabilizes, the advanced structure is closure. The same pattern appears in topology, algebra, logic, and theoretical computer science. The elementary set manipulations are the computational entrance to closure operators and fixed-point mathematics.

§39.22 Generation as an intersection construction

Many elementary problems ask for the smallest set satisfying a property. The general construction is:

$$\langle S \rangle = \bigcap \{C : S \subseteq C, C \text{ has the desired closure property}\}.$$

This works because the intersection of closed objects is closed. For subspaces, intersections of subspaces are subspaces. For subgroups, intersections of subgroups are subgroups. For topological closed sets, arbitrary intersections are closed.

This explains why “generated by” is such a stable phrase across mathematics. The generated object is not guessed; it is forced by closure under operations and minimality under inclusion.

§39.23 Closure and logic of consequences

A closure operator can also be interpreted as logical consequence. If S is a set of assumptions, $c(S)$ is the set of all consequences of S . Extensivity says assumptions are consequences of themselves. Monotonicity says more assumptions give more consequences. Idempotence says taking consequences twice adds nothing new.

This logic viewpoint connects the original set exercises to algebraic logic and model theory. In algebra, closure under operations generates structures; in logic, closure under inference generates theories.

Set Counting, Modular Patterns, and Mathematical Induction

N-20220808

§40.1 A set-counting problem

The first page studies finite sets $A_1, A_2, \dots, A_{32}$ and asks how many subsets can satisfy certain intersection conditions. The handwritten solution builds a toy example in which each set contains a fixed number of elements and one carefully controls pairwise intersections.

The general strategy is inclusion–exclusion. If every set has m elements and pairwise overlaps have controlled size, then the size of the union is not simply the sum of sizes. One must subtract overlaps and add back higher overlaps. The note’s construction eventually counts a large number of possible sets, obtaining a count of the form

$$31 + 28 \cdot 30 = 871.$$

The exact model is less important than the idea: when many sets overlap, one should count by layers, not by naive addition.

§40.2 A modular arithmetic construction

Another problem constructs two sets:

$$A = \{1, 4, 6, 7, 9, \dots, 1998\},$$

and a reversed or complementary set near 2000. The note observes that the pattern is governed by residues modulo 11. Since

$$2000 = 11 \cdot 181 + 9,$$

one can write elements as

$$11t + r.$$

The allowed residues are something like

$$r \in \{1, 4, 6, 7, 9\}.$$

The real insight is that forbidden differences correspond to forbidden residue differences. Instead of checking hundreds of numbers, one checks a small table modulo 11. This is one of the most common olympiad uses of modular arithmetic.

§40.3 Why not every construction is finished

The page explicitly says “problem reserved” in a few places. That is an honest part of mathematical notes. The important thing is to record what was tried and why it seemed plausible. Here the attempt is to remove residues that conflict pairwise and leave a maximal collection. Whether the construction is optimal requires a separate upper-bound argument.

This distinction is crucial: constructing an example proves a lower bound, but proving maximality requires an upper bound.

§40.4 Sequence exercises

Later pages return to sequence problems. Examples include arithmetic and geometric sequences, partial sums, and recurrence-like identities. One problem has

$$a_n = a_1 + (n - 1)d$$

and conditions such as

$$a_1 + a_2 + \cdots + a_n = S_n.$$

The correct habit is to translate everything into a_1, d, n , solve, and then check the requested value.

Another problem uses a geometric progression and asks for a common ratio or a sum. The formula

$$S_n = a_1 \frac{q^n - 1}{q - 1}$$

again appears, with the warning that $q = 1$ is a separate case.

§40.5 Mathematical induction

The most substantial later pages discuss induction. The note distinguishes several forms:

- (i) ordinary induction: prove $P(1)$, then prove $P(n) \Rightarrow P(n + 1)$;
- (ii) strong induction: prove $P(1), \dots, P(n) \Rightarrow P(n + 1)$;
- (iii) induction with a shifted start: begin at $n = m$ instead of 1.

The proof skeleton is:

- (i) base case;
- (ii) induction hypothesis;
- (iii) induction step;
- (iv) conclusion.

The note emphasizes that the induction hypothesis must be used precisely. One cannot simply assume the desired statement for $n + 1$.

§40.6 Examples of induction

One exercise proves a formula for a sum, such as

$$1 + 2 + \cdots + n = \frac{n(n + 1)}{2}.$$

The induction step is

$$1 + \cdots + n + (n + 1) = \frac{n(n + 1)}{2} + (n + 1) = \frac{(n + 1)(n + 2)}{2}.$$

Another exercise proves a divisibility statement. For such problems, the usual technique is to write the $n + 1$ case in terms of the n case plus a visibly divisible remainder.

The page also contains a problem of the form

$$4^{2k+1} + 3^{k+1}$$

or similar, where the induction step uses modular or factorization identities. The guiding rule is: when proving divisibility by induction, preserve a factor rather than expand everything.

§40.7 Induction structure

A proof by induction has two parts: the base case and the inductive step. The inductive step is not simply checking the next case; it is proving

$$P(n) \Rightarrow P(n+1)$$

for arbitrary n . This is why writing the induction hypothesis explicitly matters.

Strong induction allows the step to use all previous cases:

$$P(1), P(2), \dots, P(n) \Rightarrow P(n+1).$$

It is useful when the object of size $n+1$ breaks into smaller pieces of varying sizes.

§40.8 Modular patterns

When a problem depends on remainders, the natural universe is $\mathbb{Z}/m\mathbb{Z}$. Instead of listing many cases, one should identify the period. If a sequence satisfies

$$a_{n+m} = a_n,$$

then all behavior is determined by the first m terms. This is the organizing idea behind many modular counting exercises.

§40.9 A note on scope

This note is about finite structure. Set problems use inclusion–exclusion and modular residues; sequence problems use closed forms; induction proves that a pattern persists for all n . The conceptual chain is:

$$\text{example} \Rightarrow \text{pattern} \Rightarrow \text{proof by induction or upper bound}.$$

A mature solution must know which stage it is in. Guessing a pattern is not proving it; constructing an example is not proving optimality; verifying a few cases is not induction.

§40.10 Construction versus upper bound

A combinatorial extremal problem has two halves. A construction proves a lower bound: it shows that at least N objects are possible. An upper bound proves that no construction can exceed N . Only when the two match is the extremal value determined.

The note’s “problem reserved” comments are therefore mathematically honest. Producing a residue-class construction is not the same as proving optimality. The missing step is usually an invariant, counting argument, or compression method that forbids larger examples.

§40.11 Residue classes as finite quotient spaces

When a construction is governed by residues modulo m , the natural universe is

$$\mathbb{Z}/m\mathbb{Z}.$$

A set of integers with forbidden differences becomes a subset of this finite cyclic group with forbidden difference set D . Equivalently, it is an independent set in the Cayley graph whose edges connect residues differing by elements of D .

This transforms a long integer-checking problem into a finite graph problem on residues. The allowed residues are not a random pattern; they form a configuration avoiding forbidden differences in a quotient group.

§40.12 Induction as well-founded recursion

Ordinary induction works because $\mathbb{N}$ is well-ordered. To prove $P(n)$ for all n , it is enough to prove a base case and a step $P(n) \Rightarrow P(n+1)$. Strong induction uses the stronger hypothesis that all smaller cases are known.

The common generalization is well-founded induction. If every nonempty set of objects has a minimal element under a relation $<$, then one may prove $P(x)$ by assuming $P(y)$ for all $y < x$. This is why induction works on integers, trees, expressions, and recursively defined sets.

§40.13 Finite differences in divisibility induction

Many divisibility inductions become cleaner with finite differences. If

$$f(n+1) - f(n)$$

is divisible by m for all n , and $f(n_0)$ is divisible by m , then every later $f(n)$ is divisible by m . This is the discrete fundamental theorem of calculus:

$$f(n) = f(n_0) + \sum_{k=n_0}^{n-1} (f(k+1) - f(k)).$$

So induction is often just summing a difference identity.

§40.14 Floor functions and lattice points

A floor expression records how many integer thresholds have been crossed. Sums such as

$$\sum_{k=0}^{m-1} \left\lfloor \frac{ak+b}{m} \right\rfloor$$

count lattice points under a rational-slope line. This connects floor-function exercises to lattice-point enumeration.

In higher dimensions, Ehrhart theory studies

$$L_P(n) = |nP \cap \mathbb{Z}^d|,$$

the number of lattice points in dilates of a polytope. For integral polytopes, $L_P(n)$ is a polynomial in n . Thus one-variable floor sums are small shadows of a general theory.

§40.15 Pigeonhole and averaging

Many finite counting arguments are averaging arguments. If N objects are distributed among r boxes, then some box has at least

$$\left\lceil \frac{N}{r} \right\rceil$$

objects. This is the quantitative pigeonhole principle.

In extremal problems, one often defines boxes by residues, intervals, or local configurations, then proves that too many objects force a forbidden collision inside one box.

§40.16 Finite structure from quotients and induction

The note's finite examples are unified by quotienting and persistence. Modular constructions live in finite quotient groups, induction proves that a pattern persists, finite differences turn induction into summation, floor functions count lattice points, and extremal constructions require matching lower and upper bounds.

Basic Inequalities, Cauchy, Triangle Inequality, and Coordinate Geometry

N-20220810

§41.1 Basic inequalities

The first page is a compact inequality note. It begins with the elementary observation that squares are nonnegative. For real numbers,

$$(a - b)^2 \geq 0,$$

so

$$a^2 + b^2 \geq 2ab,$$

with equality iff $a = b$. For three variables,

$$a^2 + b^2 + c^2 \geq ab + bc + ca,$$

because

$$2(a^2 + b^2 + c^2 - ab - bc - ca) = (a - b)^2 + (b - c)^2 + (c - a)^2 \geq 0.$$

The handwritten page marks equality conditions carefully: equality occurs when all variables involved in the squared differences coincide.

§41.2 AM–GM

For $a, b \geq 0$,

$$\frac{a + b}{2} \geq \sqrt{ab},$$

with equality iff $a = b$. Equivalently,

$$a + b \geq 2\sqrt{ab}.$$

The note also records the three-variable form

$$\frac{a + b + c}{3} \geq \sqrt[3]{abc}, \quad a, b, c \geq 0,$$

and uses it repeatedly in examples.

A typical application is to expressions such as

$$\frac{a + b + c}{3} \geq \sqrt[3]{abc} \implies a + b + c \geq 3\sqrt[3]{abc}.$$

The useful habit is to look for terms whose product is fixed. AM–GM then turns a fixed product into a lower bound for a sum.

§41.3 Cauchy's inequality

The note states Cauchy's inequality in two forms. The vector form is

$$(a_1b_1 + \cdots + a_nb_n)^2 \leq (a_1^2 + \cdots + a_n^2)(b_1^2 + \cdots + b_n^2).$$

The Engel or Titu Andreescu form is

$$\frac{x_1^2}{a_1} + \cdots + \frac{x_n^2}{a_n} \geq \frac{(x_1 + \cdots + x_n)^2}{a_1 + \cdots + a_n}, \quad a_i > 0.$$

This second form is often what the handwritten examples use. It is obtained by applying Cauchy to

$$\left(\sum_i \frac{x_i^2}{a_i} \right) \left(\sum_i a_i \right) \geq \left(\sum_i x_i \right)^2.$$

§41.4 Triangle inequality and its variants

The page then discusses absolute values:

$$|a + b| \leq |a| + |b|, \quad |a - b| \leq |a - c| + |c - b|.$$

The second formula is the triangle inequality with an intermediate point c . The note also uses the reverse triangle inequality

$$||a| - |b|| \leq |a - b|.$$

A useful proof is

$$|a| = |a - b + b| \leq |a - b| + |b|, \quad |b| = |b - a + a| \leq |a - b| + |a|.$$

Together these imply the reverse inequality.

The handwritten warning is that absolute value inequalities are not merely algebraic signs. They encode distances on the real line.

§41.5 Worked AM–GM examples

One example asks for a lower bound of expressions of the form

$$x + \frac{1}{x}, \quad x > 0.$$

By AM–GM,

$$x + \frac{1}{x} \geq 2.$$

Similarly,

$$x + y + z \geq 3\sqrt[3]{xyz}$$

when $x, y, z > 0$. If a problem fixes xyz , this gives a direct lower bound for the sum.

Another common pattern is

$$\frac{x}{y} + \frac{y}{z} + \frac{z}{x} \geq 3, \quad x, y, z > 0,$$

because the product of the three terms is 1.

§41.6 A Cauchy example

For positive a, b, c , a typical inequality is

$$\frac{a^2}{b+c} + \frac{b^2}{c+a} + \frac{c^2}{a+b} \geq \frac{(a+b+c)^2}{2(a+b+c)} = \frac{a+b+c}{2}.$$

This is exactly Engel form of Cauchy. The handwritten page uses this kind of manipulation to turn a sum of fractions into one clean square divided by one clean denominator.

§41.7 Coordinate geometry at the end of the note

The last page shifts to coordinate geometry. The diagrams show points, vectors, and lines, with formulas for distances and slopes. The essential formulas are:

$$AB = \sqrt{(x_A - x_B)^2 + (y_A - y_B)^2},$$

$$m = \frac{y_2 - y_1}{x_2 - x_1},$$

and the point-slope form of a line

$$y - y_0 = m(x - x_0).$$

A recurring idea is that geometric claims can be turned into algebraic equalities. For example, perpendicularity of two nonvertical lines is

$$m_1 m_2 = -1,$$

while parallelism is

$$m_1 = m_2.$$

Distance and slope are the coordinate versions of length and direction.

§41.8 What the examples reveal: three structures

This note is about choosing the right inequality language. Squares prove elementary inequalities, AM–GM handles fixed products, Cauchy handles dot products and sums of fractions, and triangle inequalities handle distances. The advanced version of the same habit is to identify the structure behind the inequality before calculating.

Elementary inequality	Higher structure	Equality meaning
Cauchy–Schwarz	inner product / projection / Gram matrix	vectors are dependent
Triangle inequality	normed vector space	same direction in strictly convex cases
AM–GM	convexity or concavity of log	all variables equal
Positive square	positive semidefinite cone	quadratic form degenerates

§41.9 Cauchy–Schwarz and Gram determinants

Cauchy–Schwarz says

$$|\langle x, y \rangle| \leq \|x\| \|y\|.$$

It follows from the nonnegativity of

$$\|x - ty\|^2 \geq 0$$

or, equivalently, from positivity of the Gram matrix

$$G(x, y) = \begin{pmatrix} \langle x, x \rangle & \langle x, y \rangle \\ \langle y, x \rangle & \langle y, y \rangle \end{pmatrix}.$$

The inequality $\det G(x, y) \geq 0$ is exactly Cauchy–Schwarz. Equality means G has rank one, i.e. the two vectors are linearly dependent.

In Hilbert spaces, the same inequality controls Fourier coefficients:

$$|\langle f, e_n \rangle| \leq \|f\|_2 \|e_n\|_2.$$

This is why an inequality first met in finite sums becomes indispensable in functional analysis.

§41.10 Triangle inequality and normed spaces

The triangle inequality

$$\|x + y\| \leq \|x\| + \|y\|$$

defines the geometry of a normed vector space. The usual absolute value, Euclidean length, sup norm, and L^p norms are all instances. Minkowski’s inequality is precisely the triangle inequality in L^p :

$$\|f + g\|_p \leq \|f\|_p + \|g\|_p, \quad \|f\|_p = \left(\int |f|^p \right)^{1/p}.$$

Thus a contest-level inequality about square roots and sums is structurally the same statement as the triangle inequality for functions.

§41.11 AM–GM, convexity, and cones

For positive variables, AM–GM follows from concavity of $\log$:

$$\log \left(\frac{1}{n} \sum_i x_i \right) \geq \frac{1}{n} \sum_i \log x_i.$$

The equality condition says all variables are equal, meaning there is no spread.

Many inequalities also describe membership in a convex cone. For instance,

$$A \succeq 0 \iff x^T A x \geq 0 \text{ for all } x$$

defines the cone of positive semidefinite matrices. Basic “square is nonnegative” arguments are the one-dimensional shadow of this cone viewpoint.

Inequality Exercises: Cauchy, AM–GM, Vectors, and Optimization

N-20220811

§42.1 How to read this note

This note is a compact inequality practice sheet. The solutions are handwritten, but the recurring pattern is clear: almost every problem is asking us to recognize which standard inequality is hiding behind the expression. The useful question is not “which formula should I memorize?”, but rather “what structure does the expression already have?”

If there are fractions with variable denominators, Cauchy or Engel form is natural. If there are products such as abc , AM–GM is natural. If there are square roots of quadratic expressions, think of vector lengths and the triangle inequality. If the problem asks for a maximum or minimum under a quadratic constraint, expect equality cases in Cauchy–Schwarz.

§42.2 A fraction problem with a constraint

One problem asks to handle expressions of the form

$$\frac{4}{4-x^2} + \frac{9}{9-y^2}$$

under constraints such as $xy = -1$ and bounded intervals for x, y . The handwritten solution uses a Cauchy-style lower bound:

$$\frac{4}{4-x^2} + \frac{9}{9-y^2} \geq \frac{(2+3)^2}{(4-x^2) + (9-y^2)}.$$

Then the constraint controls the denominator. The important point is that the numerator 4, 9 are squares, so the expression is inviting the Cauchy inequality

$$\frac{a^2}{u} + \frac{b^2}{v} \geq \frac{(a+b)^2}{u+v}.$$

This is one of the most common contest moves: when a sum of fractions has square numerators, try to combine them by Cauchy rather than simplifying separately.

§42.3 A weighted square-root comparison

Another exercise defines

$$P = \sqrt{ab} + \sqrt{cd},$$

and

$$Q = \sqrt{ma+nc} \sqrt{\frac{b}{m} + \frac{d}{n}},$$

where the variables are positive. Expanding Q^2 gives

$$\begin{aligned} Q^2 &= (ma + nc) \left(\frac{b}{m} + \frac{d}{n} \right) \\ &= ab + cd + \frac{mad}{n} + \frac{ncb}{m}. \end{aligned}$$

By AM–GM,

$$\frac{mad}{n} + \frac{ncb}{m} \geq 2\sqrt{abcd}.$$

Hence

$$Q^2 \geq ab + cd + 2\sqrt{abcd} = (\sqrt{ab} + \sqrt{cd})^2 = P^2.$$

So $P \leq Q$, with equality exactly when

$$\frac{mad}{n} = \frac{ncb}{m}.$$

The original note asks for the condition for strict inequality; it is precisely the negation of this equality condition.

§42.4 Maximizing a radical expression

The note then asks for the maximum of

$$f(x) = 2\sqrt{x-3} + \sqrt{5-x}, \quad 3 \leq x \leq 5.$$

One way is calculus. We have

$$f'(x) = \frac{1}{\sqrt{x-3}} - \frac{1}{2\sqrt{5-x}}.$$

Setting $f'(x) = 0$ gives

$$2\sqrt{5-x} = \sqrt{x-3},$$

so

$$4(5-x) = x-3, \quad x = \frac{23}{5}.$$

At this point,

$$f\left(\frac{23}{5}\right) = 2\sqrt{\frac{8}{5}} + \sqrt{\frac{2}{5}} = \sqrt{10}.$$

But the handwritten note also asks whether this can be done without derivatives. The answer is yes, by Cauchy–Schwarz:

$$\left(2\sqrt{x-3} + \sqrt{5-x}\right)^2 \leq (4+1)((x-3) + (5-x)) = 10.$$

Equality holds when

$$\frac{2}{1} = \frac{\sqrt{x-3}}{\sqrt{5-x}},$$

which again gives $x = 23/5$. This non-calculus proof is more elegant because it explains the form of the answer.

§42.5 Equality in Cauchy–Schwarz

A vector problem gives

$$a^2 + b^2 + c^2 = 25, \quad x^2 + y^2 + z^2 = 36, \quad ax + by + cz = 30.$$

By Cauchy–Schwarz,

$$(ax + by + cz)^2 \leq (a^2 + b^2 + c^2)(x^2 + y^2 + z^2) = 25 \cdot 36 = 900.$$

Since $ax + by + cz = 30$, equality holds. Therefore

$$(a, b, c) = k(x, y, z).$$

Comparing norms gives

$$k = \frac{5}{6}.$$

Thus

$$\frac{a + b + c}{x + y + z} = \frac{5}{6}.$$

This is an excellent example of why equality cases matter. The inequality is not just bounding a number; it tells us the exact geometry of the vectors.

§42.6 Triangle-angle inequalities

Let A, B, C be the angles of a triangle. Then

$$A + B + C = \pi.$$

By Cauchy–Schwarz,

$$(A + B + C) \left(\frac{1}{A} + \frac{1}{B} + \frac{1}{C} \right) \geq (1 + 1 + 1)^2 = 9.$$

Hence

$$\frac{1}{A} + \frac{1}{B} + \frac{1}{C} \geq \frac{9}{\pi}.$$

Similarly,

$$A^2 + B^2 + C^2 \geq \frac{(A + B + C)^2}{3} = \frac{\pi^2}{3}.$$

Both equalities occur only for

$$A = B = C = \frac{\pi}{3}.$$

The note's implicit lesson is that symmetric triangle inequalities often have equality at the equilateral triangle.

§42.7 Two spheres and a linear equation

Another problem asks to solve a system of two sphere equations. The trick is to subtract them. If two equations are

$$(x - a_1)^2 + (y - b_1)^2 + (z - c_1)^2 = R_1^2,$$

$$(x - a_2)^2 + (y - b_2)^2 + (z - c_2)^2 = R_2^2,$$

subtracting cancels $x^2 + y^2 + z^2$ and leaves a plane. The intersection of two spheres is therefore a circle, a point, or empty; algebraically it becomes a sphere-plane problem.

The handwritten solution then uses Cauchy–Schwarz to bound a linear expression such as

$$-8x + 6y - 24z.$$

The equality condition determines the direction of the vector (x, y, z) . Again the method is geometric: a linear function is maximized on a sphere in the direction of its coefficient vector.

§42.8 A product inequality

If

$$(1 + a)(1 + b)(1 + c) = 8, \quad a, b, c > 0,$$

then the note proves

$$abc \leq 1.$$

Indeed,

$$\begin{aligned} 8 &= 1 + a + b + c + ab + bc + ca + abc \\ &\geq 1 + 3\sqrt[3]{abc} + 3\sqrt{(ab)(bc)(ca)} + abc \\ &= 1 + 3t + 3t^2 + t^3 = (1 + t)^3, \end{aligned}$$

where $t = \sqrt[3]{abc}$. Thus $1 + t \leq 2$, so $t \leq 1$, and therefore $abc \leq 1$.

§42.9 Vector-style radical inequality

The pages also contain an inequality of the form

$$\sqrt{x^2 + xy + y^2} + \sqrt{x^2 + xz + z^2} \geq \sqrt{y^2 + yz + z^2}.$$

This is best understood as a disguised triangle inequality. One rewrites each quadratic as a squared Euclidean norm of a two-dimensional vector. For example,

$$x^2 + xy + y^2 = \left(x + \frac{y}{2}\right)^2 + \frac{3}{4}y^2.$$

Once each radical is interpreted as a vector length, the inequality becomes

$$\|u\| + \|v\| \geq \|u + v\|.$$

This is a typical high-level move: turn algebraic square roots into geometry.

§42.10 The deeper habit

The note is not merely a list of inequality tricks. It repeatedly uses the same three ideas: Cauchy detects equality of vectors, AM–GM detects equality of positive quantities, and radical inequalities often hide triangle inequalities. The more mature way to remember the sheet is therefore:

$$\text{fractions} \rightarrow \text{Cauchy}, \quad \text{products} \rightarrow \text{AM–GM}, \quad \text{radicals} \rightarrow \text{norms}.$$

§42.11 Equality cases as structural information

In inequality problems, equality conditions are not afterthoughts. They often reveal the correct substitution or normalization. For Cauchy–Schwarz,

$$\left(\sum_i a_i b_i\right)^2 = \left(\sum_i a_i^2\right) \left(\sum_i b_i^2\right)$$

holds exactly when a and b are proportional. If a proposed solution cannot recover the equality case, it is probably not using the natural structure.

§42.12 Optimization on constrained domains

Many radical or fraction inequalities are optimization problems on a constrained set. A constraint such as $x + y + z = 1$ suggests compactness if variables are nonnegative. Then a maximum exists, and one can use convexity, Cauchy, or Lagrange multipliers to locate it.

This separates two tasks: first prove existence or reduce the domain; then identify the inequality that gives the optimum.

§42.13 Vectors and geometric inequalities

Vector methods convert algebraic inequalities into length comparisons. If an expression can be written as a dot product, Cauchy applies. If it is a sum of lengths, triangle inequality or Minkowski applies.

The power of vectors is that equality becomes a geometric condition: parallel vectors, collinearity, or degenerate triangles.

§42.14 Equality cases as geometric diagnostics

The exercise sheet is about choosing the right geometry. Cauchy handles dot products, AM–GM handles balanced positive products, convexity handles averages, and Lagrange multipliers handle smooth constrained extrema. Equality cases are the map showing which method is natural.

§42.15 Equality cases as geometry

The original inequality exercises repeatedly ask when equality holds. This is not a side detail; equality describes the geometry of the inequality. In Cauchy–Schwarz,

$$\left(\sum a_i b_i\right)^2 \leq \left(\sum a_i^2\right) \left(\sum b_i^2\right),$$

equality holds exactly when the vectors (a_i) and (b_i) are linearly dependent. In AM–GM,

$$\frac{x_1 + \cdots + x_n}{n} \geq \sqrt[n]{x_1 \cdots x_n},$$

equality holds when all variables are equal. These equality sets often reveal the hidden symmetry of a problem.

§42.16 Lagrange multipliers and inequalities

Many inequality exercises are optimization problems with constraints. If one wants to extremize $f(x_1, \dots, x_n)$ under $g(x_1, \dots, x_n) = 0$, the smooth interior candidates satisfy

$$\nabla f = \lambda \nabla g.$$

For example, optimizing a quadratic form under $\|x\| = 1$ leads to

$$Ax = \lambda x,$$

so eigenvalues are extrema of quadratic forms on the unit sphere. This is the Rayleigh quotient principle:

$$\lambda_{\min} \leq \frac{x^T Ax}{x^T x} \leq \lambda_{\max}$$

for symmetric A .

§42.17 Semidefinite viewpoint

Some inequalities are equivalent to positive semidefiniteness. To prove

$$ax^2 + 2bxy + cy^2 \geq 0$$

for all x, y , one checks

$$\begin{pmatrix} a & b \\ b & c \end{pmatrix} \succeq 0.$$

That means

$$a \geq 0, \quad ac - b^2 \geq 0.$$

This connects elementary quadratic inequalities to semidefinite programming, where one optimizes linear functions over the cone of positive semidefinite matrices.

§42.18 Returning to the original problems

The contest techniques in the source—Cauchy, AM–GM, substitutions, and equality checking—should be read as finite-dimensional optimization. The advanced language does not replace the computations; it explains why the same few patterns recur: projection, convexity, eigenvalues, and positive semidefinite forms.

§43.1 Bernoulli's inequality

The note begins in the middle of a list of standard inequalities. Bernoulli's inequality says that if $x > -1$ and $n \in \mathbb{N}$, then

$$(1 + x)^n \geq 1 + nx.$$

A slightly more general form is: if $x_i > -1$ and the x_i 's have the same sign, then

$$(1 + x_1)(1 + x_2) \cdots (1 + x_n) \geq 1 + x_1 + x_2 + \cdots + x_n.$$

When $x_1 = \cdots = x_n = x$, this reduces to Bernoulli's inequality.

The handwritten note adds the useful memory: for $x > 0$, the graph of $y = (1 + x)^n$ lies above its tangent line at 0. This is the convexity explanation.

§43.2 Power means

For a positive vector

$$a = (a_1, \dots, a_n),$$

and $r \neq 0$, define the power mean

$$M_r(a) = \left(\frac{a_1^r + a_2^r + \cdots + a_n^r}{n} \right)^{1/r}.$$

The power mean inequality says that if $r \leq s$, then

$$M_r(a) \leq M_s(a).$$

Thus the familiar chain is

$$\text{harmonic mean} \leq \text{geometric mean} \leq \text{arithmetic mean} \leq \text{quadratic mean}.$$

This is a good example of a high-level unification. Instead of remembering many separate inequalities, one remembers that changing the exponent changes the mean monotonically.

§43.3 Cauchy–Schwarz

The Cauchy inequality is

$$(a_1^2 + \cdots + a_n^2)(b_1^2 + \cdots + b_n^2) \geq (a_1 b_1 + \cdots + a_n b_n)^2.$$

It is equality exactly when the two vectors are proportional:

$$(a_1, \dots, a_n) = \lambda(b_1, \dots, b_n).$$

The Engel form, also called Titu Andreescu's lemma, is

$$\frac{a_1^2}{b_1} + \frac{a_2^2}{b_2} + \cdots + \frac{a_n^2}{b_n} \geq \frac{(a_1 + a_2 + \cdots + a_n)^2}{b_1 + b_2 + \cdots + b_n},$$

for $b_i > 0$. The note marks this formula because it is one of the most frequently used forms in contest inequalities.

The equality condition is

$$\frac{a_1}{b_1} = \frac{a_2}{b_2} = \cdots = \frac{a_n}{b_n}.$$

§43.4 Minkowski inequality

The page also records the inequality

$$\sqrt{a_1^2 + b_1^2} + \cdots + \sqrt{a_n^2 + b_n^2} \geq \sqrt{(a_1 + \cdots + a_n)^2 + (b_1 + \cdots + b_n)^2}.$$

This is the triangle inequality in $\mathbb{R}^2$, applied repeatedly. Each pair (a_i, b_i) is a vector, and the left side is the sum of lengths, while the right side is the length of the sum.

This viewpoint is much better than memorizing the expression. Whenever square roots of sums of squares appear, try to see vectors.

§43.5 A three-variable inequality

One of the listed consequences has the form

$$\frac{x}{y+z}(a+b) + \frac{y}{z+x}(c+a) + \frac{z}{x+y}(a+b) \geq \sqrt{3(ab+bc+ca)},$$

under positivity assumptions. The exact printed expression is somewhat compressed, but the intended family is clear: combine Cauchy/Engel form with AM–GM to convert fractions into symmetric expressions.

The practical rule is: denominators such as $x+y$, $y+z$, $z+x$ usually suggest either Nesbitt-type transformations or Engel form.

§43.6 Chebyshev's inequality

If

$$a_1 \leq a_2 \leq \cdots \leq a_n, \quad b_1 \leq b_2 \leq \cdots \leq b_n,$$

then Chebyshev's sum inequality says

$$\left(\frac{a_1 + \cdots + a_n}{n} \right) \left(\frac{b_1 + \cdots + b_n}{n} \right) \leq \frac{a_1 b_1 + \cdots + a_n b_n}{n}.$$

Equivalently,

$$(a_1 + \cdots + a_n)(b_1 + \cdots + b_n) \leq n(a_1 b_1 + \cdots + a_n b_n).$$

The note writes this as “same monotonic order gives larger average product.” This is the right intuition: sorted sequences positively correlate.

§43.7 Jensen's inequality

For a convex function f , Jensen's inequality says

$$f(\alpha x + (1 - \alpha)y) \leq \alpha f(x) + (1 - \alpha)f(y),$$

where $0 \leq \alpha \leq 1$. More generally, if $\alpha_i \geq 0$ and $\sum \alpha_i = 1$, then

$$f(\alpha_1 x_1 + \cdots + \alpha_n x_n) \leq \alpha_1 f(x_1) + \cdots + \alpha_n f(x_n).$$

The note also records the second-derivative criterion: if

$$f''(x) \geq 0,$$

then f is convex.

Jensen is not just another inequality; it is a language. AM–GM follows by applying Jensen to $\ln x$, and many trigonometric inequalities follow by checking convexity or concavity on the relevant interval.

§43.8 Young, Hölder, and Minkowski

Young's inequality says that if $a, b > 0$, $p, q > 1$, and

$$\frac{1}{p} + \frac{1}{q} = 1,$$

then

$$ab \leq \frac{a^p}{p} + \frac{b^q}{q}.$$

Equality holds when

$$a^p = b^q.$$

Hölder's inequality says

$$a_1 b_1 + \cdots + a_n b_n \leq (a_1^p + \cdots + a_n^p)^{1/p} (b_1^q + \cdots + b_n^q)^{1/q}.$$

Minkowski's inequality says, for $p > 1$,

$$\left(\sum_{i=1}^n (a_i + b_i)^p \right)^{1/p} \leq \left(\sum_{i=1}^n a_i^p \right)^{1/p} + \left(\sum_{i=1}^n b_i^p \right)^{1/p}.$$

The note lists these at the end because they are higher-level inequalities. They belong naturally to normed vector spaces and L^p spaces.

§43.9 Cauchy–Schwarz as a master inequality

The Cauchy–Schwarz inequality

$$\left(\sum_i a_i b_i \right)^2 \leq \left(\sum_i a_i^2 \right) \left(\sum_i b_i^2 \right)$$

can be read geometrically as

$$|a \cdot b| \leq \|a\| \|b\|.$$

Many inequalities are disguised Cauchy–Schwarz. The skill is to choose the two vectors correctly.

§43.10 Equality conditions

For every inequality, the equality condition is part of the theorem. In Cauchy–Schwarz, equality holds iff the two vectors are linearly dependent. In AM–GM, equality holds iff all the compared positive numbers are equal. Tracking equality conditions is essential in optimization problems.

§43.11 The lesson behind it

This page is best read as a map rather than a list. Bernoulli comes from convexity of powers. Power means organize all classical means. Cauchy is inner-product geometry. Minkowski is the triangle inequality. Chebyshev is monotone correlation. Jensen is convexity. Young and Hölder are the algebra of dual exponents. Once one sees these meanings, choosing the right inequality becomes much less arbitrary.

§43.12 The hierarchy of means

Power means form a scale:

$$M_r(x) = \left(\frac{1}{n} \sum_i x_i^r \right)^{1/r}.$$

For positive variables, $M_r \leq M_s$ when $r < s$. The arithmetic, geometric, and harmonic means are special points on this scale.

This hierarchy organizes many inequalities: instead of proving each comparison separately, one proves monotonicity of the mean parameter.

§43.13 Convex order and Jensen

Jensen’s inequality is the central convexity principle:

$$f(\mathbb{E}X) \leq \mathbb{E}f(X)$$

for convex f . In finite form this is the weighted inequality. It tells us that spreading a variable increases the expectation of a convex function.

This idea connects Jensen, Karamata, and variance-type estimates.

§43.14 Minkowski and norm geometry

Minkowski’s inequality is the triangle inequality in L^p :

$$\left(\sum_i |x_i + y_i|^p \right)^{1/p} \leq \left(\sum_i |x_i|^p \right)^{1/p} + \left(\sum_i |y_i|^p \right)^{1/p}.$$

It is not merely an algebraic estimate; it says ℓ^p is a normed vector space.

§43.15 Classical inequalities in one convex landscape

Classical inequalities are best organized by structure: Bernoulli by convexity of powers, Cauchy and Holder by dual norms, Minkowski by norm geometry, Jensen by convex expectation, and Karamata by majorization. The named inequalities are landmarks in one landscape.

Function Exercises and Inequality Estimates

N-20221105

§44.1 A short note about elementary methods

The handwritten heading includes the sentiment that a person who understands mathematics sees the solution more clearly than someone who only imitates formulas. The actual page is a small exercise sheet: periodicity, quadratic minima, logarithmic comparison, parameter ranges, and an inequality with $abc = 1$. I keep the exercise style, but expand the reasoning so that the methods are visible.

§44.2 Periodicity from a transformed argument

The first problem has a function f satisfying a relation of the form

$$f(x+4) = f(x), \quad g(x+4) = g(x) + 2,$$

and asks about the periodicity of a composite or transformed function. The handwritten solution notices that the increment in x produces the same value of f and a controlled increment of g . The general principle is:

$$f(x+T) = f(x) \implies f \text{ has period } T.$$

If a formula contains $x+8$, use the period twice. For example,

$$f(x+8) = f((x+4)+4) = f(x+4) = f(x).$$

The important point is that a period is not necessarily the least positive period. A function with period 4 automatically has period 8, 12, ... Many errors come from confusing “a period” with “the fundamental period.”

§44.3 A quadratic minimum

The second problem studies

$$y = x^2 + x + \frac{5}{3}.$$

Complete the square:

$$y = \left(x + \frac{1}{2}\right)^2 + \frac{17}{12}.$$

Hence

$$y \geq \frac{17}{12},$$

with equality at

$$x = -\frac{1}{2}.$$

This is the cleanest form of a quadratic problem. Differentiation gives the same result, but completing the square better reveals the geometry: a parabola shifted horizontally and vertically.

§44.4 A logarithmic comparison

Another problem compares functions involving logarithms, for example

$$f(x) = \log_a(x+1)$$

and a related expression $g(x) - f(x)$. The handwritten solution rewrites the difference using logarithm rules:

$$\log_a u - \log_a v = \log_a \frac{u}{v}.$$

Then the sign of the difference depends on whether the base a is greater than 1 or between 0 and 1. This is the crucial trap:

$$a > 1 \quad \Rightarrow \quad \log_a x \text{ is increasing,}$$

whereas

$$0 < a < 1 \quad \Rightarrow \quad \log_a x \text{ is decreasing.}$$

So a logarithmic inequality must always track the base.

§44.5 A parameterized polynomial

The fourth problem considers a parameter p and a function with a specified zero or extremum. The handwritten computation substitutes the relevant value and reduces to an equation in p . The broad template is:

- (i) if x_0 is a zero, impose $f(x_0) = 0$;
- (ii) if x_0 is an extremum of a differentiable function, impose $f'(x_0) = 0$;
- (iii) if a graph is tangent to a line, impose both value equality and slope equality.

This page is mainly training the habit of translating a verbal condition into equations.

§44.6 An inequality under $abc = 1$

The last problem states: for $a, b, c > 0$, $abc = 1$, prove an inequality of the form

$$\frac{1}{a^3(cb+c)} + \frac{1}{b^3(ca+c)} + \frac{1}{c^3(a+b)} \geq \frac{3}{2}.$$

The handwritten solution uses the condition

$$abc = 1$$

to replace reciprocal expressions by symmetric sums such as

$$\frac{1}{a} + \frac{1}{b} + \frac{1}{c} = ab + bc + ca.$$

Then Cauchy–Schwarz or AM–GM gives a lower bound.

The useful lesson is that the constraint $abc = 1$ should not sit unused. It usually allows conversion between variables and reciprocals. A standard move is also

$$a = \frac{x}{y}, \quad b = \frac{y}{z}, \quad c = \frac{z}{x},$$

which automatically enforces $abc = 1$. Whether one uses this substitution or direct reciprocal rewriting depends on the shape of the expression.

§44.7 The conceptual turn

The exercises look miscellaneous, but the same strategy recurs: normalize the expression before attacking it. Periodicity reduces arguments modulo a period; completing the square normalizes a quadratic; logarithm laws turn products into sums; parameter conditions become equations; $abc = 1$ turns reciprocal expressions into symmetric ones. Elementary problem solving is often the art of finding the right normalization.

§44.8 Periodicity from transformed arguments

A relation such as $f(x+a) = f(x)$ is invariance under translation by a . More complicated relations, for example involving $f(c-x)$, combine translation and reflection. The group generated by these transformations controls the function's repeated values.

Thus many functional-equation exercises are really symmetry problems on the domain.

§44.9 Quadratic and logarithmic estimates

A quadratic minimum is controlled by completing the square:

$$ax^2 + bx + c = a \left(x + \frac{b}{2a} \right)^2 + c - \frac{b^2}{4a}.$$

A logarithmic comparison often uses monotonicity or concavity of log. In both cases, the point is to expose the structural feature: convexity for quadratics, monotone concavity for logarithms.

§44.10 Functional equations and orbit reduction

If a functional equation lets one transform x into simpler inputs, then the domain splits into orbits under the generated transformations. A fundamental domain is a set of representatives for these orbits.

Solving the equation means assigning values on the fundamental domain and transporting them consistently along the orbit.

§44.11 Functions through symmetry and normal form

Function exercises are unified by structure: periodicity is group invariance, quadratics are convex normal forms, logarithmic inequalities come from monotonicity and concavity, and functional equations reduce values along orbits.

Function Exercises: Domains, Ranges, Dirichlet Function, Parameters, and Functional Equations

N-20221217

§45.1 Structure of the exercise sheet

The note is a function exercise sheet, not a single theorem. It contains domain problems, range problems, a Dirichlet-function example, parameter problems, zero-counting by graphs, and functional equations. The repaired version keeps the exercise logic: each problem is a training example for a specific method.

§45.2 Domain problems

For a function domain, the first rule is to list all legality conditions. Square roots require nonnegative radicands, logarithms require positive arguments and valid bases, denominators cannot be zero, and iterated functions require every intermediate input to lie in the next domain.

For example,

$$y = \frac{\sqrt{x(3-x)}}{\lg(x-3)^2}$$

requires

$$x(3-x) \geq 0, \quad (x-3)^2 > 0, \quad \lg(x-3)^2 \neq 0.$$

The first gives $0 \leq x \leq 3$. The second excludes $x = 3$. The third excludes $(x-3)^2 = 1$, so $x = 2$ or 4 , but only 2 lies in the current interval. Hence

$$D = [0, 3) \setminus \{2\}.$$

The important step is not the answer; it is the intersection of all restrictions.

For an expression such as

$$\sqrt{a^x - kb^x}$$

in a denominator, the condition is strict:

$$a^x - kb^x > 0.$$

Since $b^x > 0$, divide by b^x :

$$\left(\frac{a}{b}\right)^x > k.$$

Now the inequality depends on whether $a/b > 1$ or $0 < a/b < 1$. This is a recurring trap in exponential and logarithmic domain problems.

§45.3 Floor functions

The exercise sheet also uses the Gauss or floor function $[x]$. The correct method is interval splitting:

$$[x] = m \iff m \leq x < m + 1.$$

A floor-function domain problem is therefore a finite or countable case split. One should not treat $[x]$ as a smooth algebraic expression.

§45.4 Range problems

For

$$y = \sqrt{x-4} + \sqrt{15-3x},$$

the domain is $4 \leq x \leq 5$. Put $u = x - 4$, so $0 \leq u \leq 1$ and

$$y = \sqrt{u} + \sqrt{3(1-u)}.$$

By Cauchy–Schwarz,

$$y^2 \leq (1+3)(u+1-u) = 4,$$

so $y \leq 2$. At the endpoints the smaller value is 1, and the maximum is achieved when the Cauchy equality condition holds. Thus the range is $[1, 2]$.

For

$$y = x + \sqrt{1-2x},$$

let $t = \sqrt{1-2x} \geq 0$. Then

$$x = \frac{1-t^2}{2}, \quad y = -\frac{1}{2}(t-1)^2 + 1.$$

Hence the range is

$$(-\infty, 1].$$

This substitution method is the natural way to remove a radical while preserving the domain.

§45.5 Dirichlet function

The Dirichlet function is

$$D(x) = \begin{cases} 1, & x \in \mathbb{Q}, \\ 0, & x \in \mathbb{R} \setminus \mathbb{Q}. \end{cases}$$

It is bounded, because $0 \leq D(x) \leq 1$. It is not monotone on any interval, because every interval contains both rational and irrational numbers. It has every nonzero rational period T , since

$$x \in \mathbb{Q} \iff x + T \in \mathbb{Q}$$

for $T \in \mathbb{Q}$. It has no irrational period, because adding an irrational changes rationality.

This exercise is included because it teaches that functions can be extremely discontinuous while still having simple algebraic properties.

§45.6 Graph transformations and zero counting

The problem involving f and

$$g(x) = b - f(2 - x)$$

asks for the values of b for which $f(x) - g(x)$ has four zeros. The transformation $x \mapsto 2 - x$ reflects the graph across the vertical line $x = 1$; then $b -$ reflects vertically and shifts. The number of zeros of $f - g$ is the number of intersections of two transformed graphs.

The method is graphical but not vague. The number of intersections changes only when a moving graph becomes tangent, passes through a vertex/corner, or crosses an endpoint. Thus parameter intervals are found by checking these critical positions.

§45.7 Logarithmic parameter equation

For an equation such as

$$\ln(ax) = 2\ln(x + 1),$$

first impose the domain:

$$ax > 0, \quad x + 1 > 0.$$

Then combine logarithms:

$$\ln(ax) = \ln((x + 1)^2)$$

so

$$ax = (x + 1)^2.$$

This becomes a quadratic in x . Having exactly one real solution means a discriminant or tangency condition, but the domain must still be checked. The sheet is training this two-step habit: algebraic transformation plus domain verification.

§45.8 Functional equations with period-like recurrence

Suppose f is defined on $[0, 1)$ by

$$f(x) = 2x - x^2$$

or another explicit expression, and for all real x satisfies

$$f(x) + f(x + 1) = 1.$$

Then values outside $[0, 1)$ are reduced by repeatedly using

$$f(x + 1) = 1 - f(x).$$

If $a = \log_2 3$, to compute $f(a) + f(2a) + f(3a)$, first determine the integer intervals containing $a, 2a, 3a$, reduce their fractional parts into $[0, 1)$, and then apply the recurrence each time.

This is the same idea as reducing modulo a period, except that shifting by 1 may flip the value rather than preserve it.

§45.9 A reciprocal functional equation

The sheet also has a functional equation of the form

$$f(x) = f(1)x + f(2)\frac{1}{x} - 1$$

for nonzero x . On $(0, +\infty)$, this becomes a problem about minimizing

$$Ax + \frac{B}{x} - 1.$$

The AM–GM inequality gives

$$Ax + \frac{B}{x} \geq 2\sqrt{AB}$$

when $A, B > 0$. The equality condition determines the point where the minimum occurs. This is a typical example where a functional equation reduces to a standard inequality.

§45.10 Quadratic iteration problem

For a quadratic

$$f(x) = x^2 + 2ax + 8,$$

conditions involving both $f(x) \leq 0$ and $f(f(x)) \leq 8$ should be treated by introducing the image variable $u = f(x)$. Then the second condition becomes $f(u) \leq 8$. The problem is about how the range of the first inequality is mapped by the quadratic again. Completing the square and tracking intervals is safer than expanding $f(f(x))$ blindly.

§45.11 Max and absolute value

The exercise about

$$M(x) = \max\{f(x), g(x)\}$$

uses the identity

$$\max\{u, v\} = \frac{u + v + |u - v|}{2}.$$

Thus maximum and minimum operations can be encoded using absolute value. This is useful when deciding whether a function built by max/min remains elementary.

§45.12 Parameter problems and critical values

When a parameter changes the number of solutions, the change usually happens at a critical value: a tangent contact, an endpoint contact, or a collision of two roots. This is why discriminants and graph tangencies appear so often.

For a quadratic equation depending on a parameter, the condition for exactly one real root is usually

$$\Delta = 0.$$

For a graphical equation, the same condition appears geometrically as tangency. The algebraic and geometric languages are two versions of the same idea.

§45.13 Functional equations as reduction rules

A functional equation such as

$$f(x+1) = 1 - f(x)$$

should be used as a reduction rule. It tells us how to move an input back to a fundamental interval. After two shifts,

$$f(x+2) = 1 - f(x+1) = f(x),$$

so the function becomes periodic with period 2. This kind of hidden periodicity is easy to miss if one treats the equation only algebraically.

§45.14 What the pattern suggests

The page is not just a set of answers. It teaches a toolbox: domains are legality intersections; ranges use substitution and inequalities; floor functions require interval splitting; Dirichlet's function uses density of rationals and irrationals; parameter problems are tangency/discriminant problems; functional equations reduce values to a fundamental interval; max and min can be represented by absolute values. The common theme is that a function is not merely a formula but a domain, graph, transformation rule, and set of constraints.

§45.15 Domains as constraint intersections

A domain problem is a constraint-satisfaction problem. Each operation contributes a legality condition: square roots impose inequalities, logarithms impose positivity, denominators impose nonvanishing, and compositions impose intermediate-domain conditions. The domain is the intersection of all these constraints.

This viewpoint prevents a common error: simplifying an expression before recording its original restrictions. For example, an expression with a denominator may algebraically simplify, but the excluded zeros of the original denominator still matter unless the problem explicitly defines the simplified expression as a new function.

§45.16 Dirichlet function through density and measure

The Dirichlet function

$$D(x) = \mathbf{1}_{\mathbb{Q}}(x)$$

is discontinuous at every real number. At any point a , every neighborhood contains rationals and irrationals, so along rational sequences D tends to 1, while along irrational sequences it tends to 0.

It is not Riemann integrable on any interval, because every subinterval has supremum 1 and infimum 0. But it is Lebesgue integrable, and its integral is 0, since $\mathbb{Q}$ has measure zero:

$$\int_0^1 \mathbf{1}_{\mathbb{Q}}(x) dx = 0.$$

Thus the same function separates two integration theories: Riemann integration sees oscillation on every interval; Lebesgue integration sees that the exceptional set is small in measure.

§45.17 Thomae's function as a nearby comparison

A useful comparison is Thomae's function

$$T(x) = \begin{cases} \frac{1}{q}, & x = \frac{p}{q} \text{ in lowest terms,} \\ 0, & x \notin \mathbb{Q}. \end{cases}$$

It is continuous at every irrational point and discontinuous at every rational point. Unlike the Dirichlet function, it is Riemann integrable with integral 0.

This comparison shows that dense discontinuities alone do not determine integrability. The size and height distribution of the discontinuities matter.

§45.18 Parameter problems as bifurcation diagrams

When a parameter changes the number of solutions, the change occurs at critical values. Algebraically, these are discriminant-zero cases, endpoint hits, or degeneracies. Geometrically, they are tangencies or changes in intersection pattern.

For a family $F(x, b) = 0$, a typical critical point satisfies

$$F(x, b) = 0, \quad \frac{\partial F}{\partial x}(x, b) = 0.$$

The first equation says the graph intersects the axis; the second says the intersection is tangent, so two roots can merge or disappear. This is the elementary version of a bifurcation diagram.

§45.19 Functional equations as group actions on the domain

A relation such as

$$f(x+1) = 1 - f(x)$$

should be read as an action of the translation group on the input line, together with an action on values. Applying the shift twice gives

$$f(x+2) = f(x),$$

so the function descends to a period-two structure after accounting for the value flip.

The interval $[0, 1)$ is a fundamental domain: every real input can be moved into it by integer shifts. The functional equation tells how the value changes during that reduction. This is the same logic behind periodic functions, modular arithmetic, and covering spaces.

§45.20 Max, min, and nonsmooth points

The identity

$$\max\{u, v\} = \frac{u + v + |u - v|}{2}$$

turns a maximum into an expression involving absolute value. It also explains where nonsmoothness can occur: at points where $u = v$, the absolute value has a corner unless the crossing is tangential in a compatible way.

In optimization language, maximum functions are piecewise-smooth. Their derivatives are governed by the active branch. At a switching point, one may need one-sided derivatives or subgradients rather than an ordinary derivative.

§45.21 Ranges as images and compactness

A range problem asks for the image $f(D)$. If D is compact and f is continuous, then $f(D)$ is compact. In $\mathbb{R}$, that means closed and bounded. This is why many range problems on closed intervals reduce to finding extrema.

If the domain is not compact, endpoints may be missing or infinity may be approached. Then the range must be described with open or half-open endpoints. This is not a cosmetic detail: it records whether a value is achieved or only approached.

§45.22 Function exercises through topology and dynamics

These function exercises are training a structural habit. A function is not just a formula; it is a formula together with a domain, a range, transformations, constraints, and possible exceptional sets. Domain questions are intersections of constraints, range questions are image problems, parameter questions track critical values, and functional equations reduce inputs by a symmetry or group action.

Function Graphs, Symmetry Centers, and Transformations

N-20230107

§46.1 A note on graph problems

This handwritten note studies function graphs and transformations: centers of symmetry, symmetry about vertical lines, absolute values, and graph-based zero counting. The important habit is to translate a visual symmetry into an algebraic identity.

§46.2 Centers of symmetry

A graph $y = f(x)$ has center of symmetry (a, b) if

$$f(a + t) + f(a - t) = 2b$$

whenever both sides are defined. Thus for a rational function like

$$f(x) = \frac{1}{x+1} + \frac{1}{x+2},$$

one shifts $x = a + t$ and looks for cancellation between t and $-t$. The vertical asymptotes also help: the center of symmetry often lies halfway between paired asymptotes.

§46.3 Symmetry about a line

A graph is symmetric about the vertical line $x = c$ iff

$$f(c + t) = f(c - t).$$

The note uses this for functions such as

$$f(x) = (1 - x^2)(x^2 + ax + b).$$

If the graph is symmetric about $x = -2$, substitute $x = -2 + t$ and $x = -2 - t$, then equate the two expressions. This is safer than trying to see the answer from the expanded graph.

§46.4 Absolute value transformations

Expressions involving

$$|x - a|, \quad ||x| - 2|, \quad |f(x)|$$

should be read as folding operations. The graph of $y = |f(x)|$ reflects the part of f below the x -axis upward. The graph of $y = f(|x|)$ copies the right half of the graph to the left symmetrically.

The problem involving

$$|||x| - 2| - 1| + |||y| - 2| - 1| = 1$$

is best attacked by symmetry. First solve in the first quadrant, then reflect across the coordinate axes. The curve is made of translated line segments, so the required string length is a sum of segment lengths rather than an area computation.

§46.5 Zero counting by graphs

A later exercise combines a periodic/symmetric function f with

$$g(x) = |x \cos(\pi x)|.$$

The number of zeros of $g(x) - f(x)$ is the number of intersections of two graphs. The correct method is to use symmetry and monotonicity on subintervals instead of solving a transcendental equation directly.

§46.6 How to find a center of symmetry in practice

For a rational function, a center of symmetry is often hidden in the denominator structure. Suppose two vertical asymptotes are paired around a point $x = a$. Then a is a natural candidate for the horizontal coordinate of the center. After shifting $x = a + t$, test whether

$$f(a + t) + f(a - t)$$

is constant. If it equals $2b$, then the center is (a, b) .

This method is better than expanding everything at once. The shift recenters the graph so that symmetry becomes ordinary oddness or evenness around the origin. In many problems, the algebra simplifies only after this recentering.

§46.7 Absolute value as folding

The note's absolute-value examples should be read visually. The transformation

$$y = |f(x)|$$

keeps the part of the graph above the x -axis and reflects the part below upward. The transformation

$$y = f(|x|)$$

copies the right half of the original graph to the left. Nested absolute values such as

$$||x| - 2|$$

are successive folds of the number line. Each fold creates new symmetry and new corner points.

For a curve such as

$$|||x| - 2| - 1| + |||y| - 2| - 1| = 1,$$

one should not immediately expand into many cases. First locate the breakpoints $x = 0, \pm 1, \pm 2, \pm 3$, solve the shape in one symmetric region, and reflect it. The algebraic expression is complicated, but the geometry is a repeated folding of straight lines.

§46.8 Parameter motion and intersection counting

Many graph problems ask for the number of roots of

$$f(x) - g(x) = 0.$$

This is the number of intersections of two graphs. When a parameter moves one graph, the number of intersections is stable except at critical moments: tangency, passing through a corner, or passing through an endpoint of a piece.

Thus parameter problems should be solved by first finding the exceptional parameter values. Between two exceptional values, the number of intersections cannot change. This converts an apparently algebraic zero-counting problem into a geometric classification problem.

§46.9 Graph transformations as group actions

The higher viewpoint is that many graph tricks are actions on the plane. Let G be a group of transformations of $\mathbb{R}^2$. If $g \in G$ and $\Gamma_f = \{(x, f(x)) : x \in D\}$ is the graph of f , then g sends Γ_f to another curve. When the image is again a graph, it gives a transformed function.

Definition 46.9.1 (Symmetry of a graph). A transformation $g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a symmetry of the graph Γ_f if

$$g(\Gamma_f) = \Gamma_f.$$

In this language, a graph property is an invariance property.

For example, reflection across the vertical line $x = c$ is

$$R_c(x, y) = (2c - x, y).$$

The graph is invariant under R_c iff

$$f(c + t) = f(c - t).$$

Thus vertical symmetry is evenness after translating the origin to c .

Similarly, the half-turn about (a, b) is

$$H_{a,b}(x, y) = (2a - x, 2b - y).$$

The graph is invariant under this half-turn iff

$$f(a + t) + f(a - t) = 2b.$$

Proposition 46.9.2 (Symmetry identities)

Let f be defined on a domain symmetric about a . Then:

- (i) Γ_f is symmetric about $x = a$ iff $f(a + t) = f(a - t)$;
- (ii) Γ_f has center (a, b) iff $f(a + t) + f(a - t) = 2b$.

Proof. For (i), reflection sends $(a + t, f(a + t))$ to $(a - t, f(a + t))$. This lies on the graph exactly when $f(a - t) = f(a + t)$. For (ii), the half-turn sends $(a + t, f(a + t))$ to $(a - t, 2b - f(a + t))$. This lies on the graph exactly when $f(a - t) = 2b - f(a + t)$. $\square$

§46.10 Absolute value as folding, not algebraic noise

The transformation

$$y = |f(x)|$$

keeps the part of the graph above the x -axis and reflects the lower part upward. The transformation

$$y = f(|x|)$$

forgets the left half of the original input and copies the right half across the y -axis.

Remark 46.10.1 (A quotient-like intuition) — An absolute value identifies two points t and $-t$. Thus nested absolute values are repeated foldings of the line. This is why expressions such as

$$||x| - 2| - 1|$$

are better studied by locating breakpoints and symmetries than by expanding all cases blindly.

§46.11 Zero counting as a stability problem

Many problems ask for the number of roots of

$$f(x) - g_t(x) = 0$$

as a parameter t changes. Geometrically, this is the number of intersections of two curves. The number is stable under small changes of t unless a special event occurs.

The common events are:

tangency, endpoint contact, crossing a corner or discontinuity.

For a smooth one-parameter equation

$$F(x, t) = 0,$$

a tangency or repeated-root event is detected by

$$F(x, t) = 0, \quad F_x(x, t) = 0.$$

This is the calculus version of the geometric statement that two curves touch without crossing transversely.

Remark 46.11.1 (The conceptual payoff) — A center of symmetry is a half-turn invariance. A vertical axis of symmetry is reflection invariance. An absolute value is a folding operation. A zero-counting problem is an intersection-stability problem. The advantage of this viewpoint is that it tells us what can change and what cannot change before any heavy algebra begins.

Functions, Mappings, Injections, Surjections, and Functional Equations

N-20230116

§47.1 What a function is

The handwritten pages continue the chapter on functions. A function is not just a formula. It consists of a domain, a codomain, and a rule assigning to each input exactly one output:

$$f : X \rightarrow Y.$$

For a finite set such as

$$X = \{1, 2, \dots, n\},$$

a function $f : X \rightarrow X$ can be recorded by the list

$$(f(1), f(2), \dots, f(n)).$$

This is why the page writes small arrays of values. The list representation is useful when studying injections, surjections, permutations, and iterates of f .

§47.2 Image and preimage

For $A \subseteq X$, the image is

$$f(A) = \{f(a) : a \in A\}.$$

For $B \subseteq Y$, the preimage is

$$f^{-1}(B) = \{x \in X : f(x) \in B\}.$$

The note repeatedly uses the fact that preimages behave better than images:

$$f^{-1}(B \cup C) = f^{-1}(B) \cup f^{-1}(C),$$

$$f^{-1}(B \cap C) = f^{-1}(B) \cap f^{-1}(C),$$

$$f^{-1}(Y \setminus B) = X \setminus f^{-1}(B).$$

Images preserve unions, but not intersections in general.

§47.3 Injective and surjective

A function is injective if

$$f(x_1) = f(x_2) \implies x_1 = x_2.$$

It is surjective if every element of the codomain is hit:

$$\forall y \in Y, \quad \exists x \in X, \quad f(x) = y.$$

For finite sets of the same size, injective and surjective are equivalent. The handwritten notes use this fact in examples where a function from a finite set to itself is shown to be a permutation.

§47.4 Composition and inverse functions

The composition of $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ is

$$g \circ f : X \rightarrow Z, \quad (g \circ f)(x) = g(f(x)).$$

If there is a function $g : Y \rightarrow X$ such that

$$g \circ f = \text{id}_X, \quad f \circ g = \text{id}_Y,$$

then f is bijective and $g = f^{-1}$.

The note includes finite examples where $f^2 = f \circ f$ or f^3 is computed. The orbit of an element under repeated application of f is

$$x, f(x), f^2(x), \dots$$

For a finite set, every orbit eventually repeats. If f is a permutation, every orbit is a cycle.

§47.5 Functional equations on finite sets

Several exercises ask for functions satisfying equations such as

$$f(f(x)) = x$$

or

$$f(f(x)) = f(x).$$

The first says that f is an involution. Its cycles have length 1 or 2. The second says that every value of f is fixed by f ; after one application, the function stabilizes.

These examples explain why drawing arrows is useful. A finite function is a directed graph where each vertex has exactly one outgoing arrow. Equations involving iterates impose restrictions on the possible directed graph.

§47.6 Functional equations over real numbers

The printed exercise page includes functional equations over $\mathbb{R}$. A common form is Jensen-like or Cauchy-like:

$$f(x + y) = f(x) + f(y),$$

possibly with extra regularity assumptions. Without assumptions, such equations can be wild; with continuity or monotonicity, they often force linearity.

For example, if f is continuous and additive, then

$$f(x) = cx.$$

The proof first shows $f(n) = nf(1)$ for integers, then $f(q) = qf(1)$ for rationals, then extends to reals by continuity.

§47.7 Looking ahead

The note moves from set-level functions to structural behavior. A function can be studied by its image/preimage behavior, its injectivity/surjectivity, its iterates, or the equations it satisfies. In finite settings, iteration gives cycles and trees. In real settings, functional equations often force rigid algebraic forms once continuity or monotonicity is added.

§47.8 Functions as morphisms

A function is data $f : X \rightarrow Y$, not only a formula. Injectivity says no information is identified; surjectivity says every target value is reached. Bijectivity says the function is an isomorphism of sets.

§47.9 Images, preimages, and logic

Preimages preserve all Boolean operations:

$$f^{-1}(A \cap B) = f^{-1}(A) \cap f^{-1}(B), \quad f^{-1}(Y \setminus A) = X \setminus f^{-1}(A).$$

Images preserve unions but not intersections in general. This asymmetry is why continuity and measurability are defined using preimages.

§47.10 Functional equations as constraints on orbits

A functional equation involving transformations of the input generates an action on the domain. Values are transported along orbits, and consistency must be checked around cycles.

§47.11 Functions as structure-preserving maps

The chapter's function exercises are about structure: maps, information loss, inverse images, composition, finite-set counting, and orbit constraints from functional equations.

§47.12 Morphisms and structural information loss

The original note studies injections, surjections, and functions. In category theory, functions are morphisms in the category **Set**. A function $f : A \rightarrow B$ is injective iff it is left-cancellable:

$$f \circ g_1 = f \circ g_2 \Rightarrow g_1 = g_2.$$

This is the categorical notion of monomorphism. It is surjective iff it is right-cancellable in **Set**:

$$h_1 \circ f = h_2 \circ f \Rightarrow h_1 = h_2.$$

This is epimorphism in **Set**.

The warning is that in other categories epimorphism need not mean surjectivity. For example, the inclusion $\mathbb{Z} \hookrightarrow \mathbb{Q}$ is an epimorphism in the category of rings, but not a surjective map. Thus the elementary set case is friendly but not universal.

§47.13 Image and preimage as an adjunction

For a function $f : X \rightarrow Y$, the direct image and inverse image satisfy

$$f(A) \subseteq B \iff A \subseteq f^{-1}(B).$$

This is a Galois connection between subsets of X and subsets of Y .

The inverse image preserves all Boolean operations:

$$f^{-1}(B \cup C) = f^{-1}(B) \cup f^{-1}(C), \quad f^{-1}(Y \setminus B) = X \setminus f^{-1}(B).$$

Direct image preserves unions but not generally intersections or complements. This asymmetry explains many elementary mistakes in set-function exercises.

§47.14 Finite functional graphs

A function $f : X \rightarrow X$ on a finite set defines a directed graph where each vertex has one outgoing edge $x \rightarrow f(x)$. Every connected component consists of a directed cycle with rooted trees feeding into it.

This structure explains iteration. For every x , the sequence

$$x, f(x), f^2(x), \dots$$

eventually becomes periodic. Injectivity means every component is a pure cycle; surjectivity on a finite set is equivalent to injectivity.

Thus elementary finite-function exercises are also the first examples of discrete dynamical systems.

§47.15 Returning to the original examples

Counting functions, checking injectivity, and computing images/preimages are not isolated set-theory tasks. They are the concrete form of morphisms, adjunctions between image and preimage, and finite dynamics under iteration. The original examples should be kept because they make these abstract ideas visible.

§47.16 Injections, surjections, and factors

Every function $f : X \rightarrow Y$ factors as

$$X \twoheadrightarrow X/\sim_f \xrightarrow{\cong} f(X) \hookrightarrow Y,$$

where $x \sim_f x'$ iff $f(x) = f(x')$. The first map collapses fibers, the middle map identifies equivalence classes with image points, and the last map includes the image in Y .

In elementary language, injectivity says the first collapse is trivial. Surjectivity says the last inclusion reaches all of Y . This factorization is the set-theoretic prototype of epi-mono factorization in category theory.

§47.17 Fibers as local data of a function

The fiber over $y \in Y$ is

$$f^{-1}(y) = \{x \in X : f(x) = y\}.$$

Counting a finite domain by fibers gives

$$|X| = \sum_{y \in f(X)} |f^{-1}(y)|.$$

This is the elementary version of disintegration: study a map by decomposing the domain into fibers. In topology and geometry, fiber bundles and covering spaces refine this idea when all fibers vary continuously over the base.

More Function Exercises: Periodicity, Piecewise Definitions, Images, and Functional Equations

N-20230117

§48.1 Continuing the function chapter

This note continues the function exercises from the previous day. The pages mix hand-written solutions with printed exercises. The recurring theme is to translate each problem into one of several languages:

formula, graph, image/preimage, iteration, functional equation.

Many errors in such problems come from using a formula while ignoring its domain or from treating a functional equation as if it held only for special values.

§48.2 Piecewise functions

The first page includes piecewise-defined functions. For example, a function may be defined differently on intervals such as

$$x < 0, \quad x = 0, \quad x > 0.$$

When studying continuity or solving equations involving such a function, one must first decide which case the input belongs to. If the input itself is an expression such as $f(x)$ or $x + 1$, the relevant case may change during the calculation.

The safe method is:

- (i) identify the possible regions of the input;
- (ii) solve the equation in each region;
- (iii) check that the solution actually belongs to the assumed region.

§48.3 Periodicity from functional equations

A relation such as

$$f(x + 1) = 1 - f(x)$$

implies

$$f(x + 2) = 1 - f(x + 1) = f(x).$$

Thus the function is periodic with period 2. The note uses this style of argument in several places: apply the same equation repeatedly until the input returns to a familiar form.

If instead

$$f(x + T) = f(x)$$

for all x , then T is a period. The smallest positive period, if it exists, is called the fundamental period.

§48.4 Graph transformations

The pages include small sketches of graphs. The standard transformations are:

$$\begin{aligned} y &= f(x - a) && \text{shift right by } a, \\ y &= f(x) + b && \text{shift up by } b, \\ y &= -f(x) && \text{reflect across the } x\text{-axis,} \\ y &= f(-x) && \text{reflect across the } y\text{-axis.} \end{aligned}$$

Absolute values fold graphs:

$$y = |f(x)|$$

reflects the negative part upward, while

$$y = f(|x|)$$

copies the right half to the left.

§48.5 Image of an interval

For a continuous function on an interval, the image is also an interval. This is the intermediate value principle in set language. If the function is monotone on the interval, then the endpoints determine the image directly.

The note's exercises often ask for the range of an expression. The method is to rewrite it as a function of one variable, find monotonicity or extrema, and then translate the result back into an interval.

§48.6 Functional equations and substitution

The handwritten solutions repeatedly substitute special values such as $x = 0$, $y = 0$, or $y = x$. This is a standard strategy. For a functional equation in two variables,

$$F(x, y, f(x), f(y), f(x + y)) = 0,$$

choosing simple values isolates pieces of the function.

For example, if

$$f(x + y) = f(x) + f(y) + c,$$

then setting $x = y = 0$ determines $f(0)$, and subtracting a constant can reduce the equation to the additive Cauchy equation.

§48.7 Counting functions between finite sets

The printed exercise page also touches finite functions. If $|A| = m$ and $|B| = n$, then the number of functions $A \rightarrow B$ is

$$n^m,$$

because each element of A independently chooses one image in B .

If one requires injectivity and $m \leq n$, the number is

$$n(n - 1) \cdots (n - m + 1).$$

If one requires bijectivity and $m = n$, the number is

$$n!.$$

§48.8 A compact bridge

The note is a training page for reading functions flexibly. A function may be a formula, a graph, a rule on a finite set, or an unknown object constrained by an equation. Each representation suggests a different method: casework for piecewise definitions, transformations for graphs, counting for finite domains, and substitution for functional equations.

§48.9 Piecewise functions as glued maps

A piecewise function is built by gluing formulas on regions. The important checks occur at boundaries: continuity, differentiability, or jump behavior depends on whether the pieces match there.

§48.10 Periodicity from generated transformations

Relations such as $f(x+a) = f(x)$ or $f(c-x) = f(x)$ generate translations and reflections of the line. The resulting group action partitions the domain into orbits.

§48.11 Images of intervals

For continuous functions, the image of an interval is an interval. If the interval is compact, the image is compact, so maxima and minima are attained. This is the topological reason range problems often reduce to endpoint and critical-point checks.

§48.12 Local definitions and global function behavior

Piecewise definitions, periodicity, graph transformations, images, and functional equations all ask how local rules glue into global behavior over the domain.

§48.13 Functional equations as dynamics

A functional equation often specifies a transformation rule rather than a direct formula. For example, equations involving $f(x+1)$, $f(2x)$, or $f(f(x))$ are naturally studied by iteration. The orbit of x under a transformation T is

$$x, T(x), T^2(x), \dots$$

A functional equation may prescribe how f behaves along these orbits.

For instance, if

$$f(x+1) = f(x),$$

then f is constant along orbits of the translation $T(x) = x+1$. Periodicity is invariance under a group action of $\mathbb{Z}$.

§48.14 Cauchy's equation and regularity

Cauchy's equation is

$$f(x+y) = f(x) + f(y).$$

If f is continuous, measurable, or bounded on an interval, then

$$f(x) = cx.$$

But without regularity assumptions, there are wild solutions constructed using a Hamel basis of $\mathbb{R}$ over $\mathbb{Q}$.

This shows why elementary functional equations often include hidden conditions such as continuity or monotonicity. These conditions are not technical decorations; they exclude pathological solutions arising from the axiom of choice.

§48.15 Symbolic dynamics

Piecewise-defined functions and iteration can lead to symbolic dynamics. If an interval map sends subintervals across one another, one can encode an orbit by a sequence of symbols recording which subinterval contains each iterate.

The shift map

$$\sigma((a_0, a_1, a_2, \dots)) = (a_1, a_2, a_3, \dots)$$

is a central model. Complicated interval dynamics can often be compared to shifts on symbol spaces.

Thus a piecewise function graph in an elementary note can be the beginning of chaos, coding of orbits, and topological entropy.

§48.16 Rigidity from regularity

Many functional equations have many algebraic solutions but few regular solutions. The pattern is:

Equation	Algebraic freedom	Regularity conclusion
$f(x + y) = f(x) + f(y)$	Hamel-linear maps	$f(x) = cx$ under continuity
$f(xy) = f(x) + f(y)$	Logarithmic freedom over the multiplicative group	$f(x) = c \log x$ under continuity
$f(x + T) = f(x)$	Arbitrary values on a fundamental domain	Periodic structure under smoothness

The original piecewise and functional-equation exercises should therefore be read as local training for a general principle: algebraic equations become meaningful only after the allowed regularity class is specified.

§48.17 Piecewise functions and partitions

A piecewise function is a function defined after partitioning the domain into regions. The mathematical issue is not only which formula applies, but how the pieces meet. Discontinuities occur at boundaries of the partition; differentiability requires matching first-order behavior across boundaries.

For example, if

$$f(x) = \begin{cases} f_1(x), & x < a, \\ f_2(x), & x \geq a, \end{cases}$$

then continuity at a requires

$$\lim_{x \rightarrow a^-} f_1(x) = f_2(a).$$

Differentiability additionally requires matching slopes. This is a local gluing condition.

§48.18 Functional equations and invariants

A functional equation often says that f is invariant or equivariant under transformations. The equation

$$f(x + T) = f(x)$$

says f is invariant under the translation action of $T\mathbb{Z}$. The equation

$$f(ax) = bf(x)$$

says f is homogeneous with respect to scaling.

Thus solving a functional equation often means finding the invariants of a group or semigroup action. This perspective makes the original examples more natural: periodicity, parity, and scaling are all symmetry constraints.

Fixed Points, Iteration, and the Banach Contraction Principle

N-20230316

§49.1 Fixed points

A fixed point of a map $T : X \rightarrow X$ is a point $x \in X$ such that

$$T(x) = x.$$

Intuitively, applying the transformation does not move the point.

A basic example is

$$T(x) = \cos x.$$

The equation

$$x = \cos x$$

has a unique real solution, approximately

$$x \approx 0.739085 \dots,$$

often called the Dottie number. Iteration

$$x_{n+1} = \cos x_n$$

converges to this fixed point for many starting values.

§49.2 Iteration

Given a map $T : X \rightarrow X$, an iterative process has the form

$$x_{n+1} = T(x_n).$$

If $x_n \rightarrow x$ and T is continuous, then

$$T(x) = T\left(\lim_{n \rightarrow \infty} x_n\right) = \lim_{n \rightarrow \infty} T(x_n) = \lim_{n \rightarrow \infty} x_{n+1} = x.$$

Thus the limit of a convergent iteration is automatically a fixed point.

The hard question is not this final step, but why the iteration converges at all.

§49.3 Metric spaces and contractions

Definition 49.3.1. A metric space is a set X equipped with a distance function $d : X \times X \rightarrow \mathbb{R}_{\geq 0}$ such that:

- (i) $d(x, y) = 0$ iff $x = y$;
- (ii) $d(x, y) = d(y, x)$;
- (iii) $d(x, z) \leq d(x, y) + d(y, z)$.

Definition 49.3.2. A map $T : X \rightarrow X$ is a contraction if there exists a constant $0 < c < 1$ such that

$$d(Tx, Ty) \leq c d(x, y)$$

for all $x, y \in X$.

§49.4 Banach fixed point theorem

Theorem 49.4.1 (Banach contraction principle)

Let X be a nonempty complete metric space, and let $T : X \rightarrow X$ be a contraction. Then T has a unique fixed point x^* . Moreover, for any starting point $x_0 \in X$, the sequence

$$x_{n+1} = T(x_n)$$

converges to x^* .

Proof. Let $x_{n+1} = T(x_n)$. Since T is a contraction,

$$d(x_{n+1}, x_n) = d(Tx_n, Tx_{n-1}) \leq c d(x_n, x_{n-1}).$$

Inductively,

$$d(x_{n+1}, x_n) \leq c^n d(x_1, x_0).$$

For $m > n$, the triangle inequality gives

$$d(x_m, x_n) \leq \sum_{k=n}^{m-1} d(x_{k+1}, x_k) \leq \sum_{k=n}^{m-1} c^k d(x_1, x_0) \leq \frac{c^n}{1-c} d(x_1, x_0).$$

Thus (x_n) is Cauchy. Completeness gives a limit x^* . By continuity of T ,

$$T(x^*) = x^*.$$

For uniqueness, if $Tx = x$ and $Ty = y$, then

$$d(x, y) = d(Tx, Ty) \leq c d(x, y).$$

Since $c < 1$, this forces $d(x, y) = 0$, so $x = y$. □

§49.5 Error estimate

The proof gives an explicit estimate:

$$d(x_n, x^*) \leq \frac{c^n}{1-c} d(x_1, x_0).$$

This is useful computationally because it tells us how quickly the iteration converges.

§49.6 Applications

Example 49.6.1 (Solving equations)

To solve $f(x) = 0$, one may rewrite it as $x = T(x)$. If T is a contraction on a complete interval or metric space, iteration converges to the solution.

Example 49.6.2 (Differential equations)

Picard iteration for ordinary differential equations is an application of the contraction principle in a function space. The local existence and uniqueness theorem for ODEs can be proved this way.

Example 49.6.3 (Functional equations)

Many recursive definitions can be interpreted as fixed point problems. In computer science, fixed point theorems appear in domain theory, semantics of recursion, and program analysis.

§49.7 Beyond Banach

Banach's theorem is powerful because it gives existence, uniqueness, and an algorithm. Other fixed point theorems weaken some of these conclusions:

- (i) Brouwer's fixed point theorem gives existence for continuous maps from a closed ball to itself, but not uniqueness or iteration convergence;
- (ii) Schauder's fixed point theorem extends Brouwer-type ideas to infinite-dimensional settings;
- (iii) Tarski's fixed point theorem concerns monotone maps on complete lattices.

§49.8 What the example reveals

The fixed point idea is a bridge between analysis, topology, computation, and dynamics. The Banach contraction principle is the cleanest analytic version: contraction gives a unique fixed point and a concrete iterative algorithm for finding it.

§49.9 Why completeness is the hidden hypothesis

The contraction estimate shows that the iterative sequence is Cauchy. It does not by itself produce a point inside the space. Completeness is exactly the condition that Cauchy sequences have limits in the space. This is why the theorem is false on an incomplete space: the iteration may approach a missing point.

For example, in $(0, 1)$ one may construct a contraction whose natural limit is 0, but $0 \notin (0, 1)$. The shrinking process is not enough unless the space contains its limiting points.

§49.10 Picard iteration as fixed-point theory

The same idea solves ordinary differential equations. The initial value problem

$$y'(t) = f(t, y(t)), \quad y(t_0) = y_0$$

is equivalent to

$$y(t) = y_0 + \int_{t_0}^t f(s, y(s)) ds.$$

Define an operator T on functions by the right-hand side. Under a Lipschitz condition and on a short interval, T is a contraction. Its fixed point is the solution. Thus the theorem is not merely about numerical iteration; it is the engine behind existence and uniqueness.

§49.11 Source repair: non-contractive fixed-point theorems

The uploaded note does not stop at Banach's theorem. It asks what happens when the map is not a strict contraction. Krasnoselskii-type theorems apply when an operator can be decomposed as

$$T = A + B,$$

where A is contractive and B is compact, usually on a closed convex subset of a Banach space. The guiding idea is that the contractive part controls distance, while the compact part supplies finite-dimensional approximation behavior.

Meir-Keeler's theorem weakens the fixed contraction constant. Instead of requiring one number $q < 1$ for all distances, it asks that distances in a small band above ε are pushed below ε . Geraghty-type theorems further replace a fixed constant by a control function. The common theme remains: some form of distance decrease forces Cauchy behavior.

§49.12 Source repair: Brouwer and Schauder

Brouwer's theorem says that a continuous map from a closed ball in finite-dimensional Euclidean space to itself has a fixed point. It gives existence but not uniqueness, and it does not give the Banach iteration estimate. For example, a rotation of a disk or a map that folds a ball into itself may have fixed points for topological reasons even though no distance is uniformly decreased.

Schauder's theorem is the infinite-dimensional version most relevant to analysis: a compact continuous map on a closed convex bounded subset of a Banach space has a fixed point under the usual hypotheses. In applications to integral and differential equations, the compactness often comes from smoothing or compact embedding: one first rewrites the equation as $u = T(u)$, then proves that T maps a convex set into itself and that the image is relatively compact.

Thus the final comparison is not just a slogan about "existence". In Banach's theorem one proves

$$d(x_{n+1}, x_n) \leq q^n d(x_1, x_0),$$

so the iterates themselves locate the unique fixed point and give an error estimate. In Brouwer-Schauder arguments the proof usually supplies no canonical iteration and no uniqueness: two different solutions may exist, and an approximating sequence needs an independent compactness or degree argument. When reading a fixed-point theorem

later, the first diagnostic question should therefore be whether the hypothesis is metric contraction, compactness/convexity, or a mixture such as Krasnoselskii's decomposition $T = A + B$.

VI

Combinatorics and Models

Part VI: Contents

N-20220720	319
50.1 Counting principles	319
50.2 A divisibility example	319
50.3 Permutations and combinations	320
50.4 Binomial theorem	320
50.5 Polynomial coefficients	320
50.6 Further contest examples	321
50.7 Binomial identities by stories	321
50.8 When to use multiplication and when to use addition	321
50.9 A broader reading	321
50.10 Counting principles as operations on combinatorial classes	321
50.11 Inclusion–exclusion as Möbius inversion	322
50.12 Indicator functions and probability language	322
50.13 Generating functions for binomial identities	322
50.14 Counting up to symmetry: Burnside’s lemma	323
50.15 Counting through structure and symmetry	323
 N-20221205	 325
51.1 Summation and product notation	325
51.2 De Morgan’s laws	325
51.3 Inclusion–exclusion	325
51.4 Pigeonhole principle	326
51.5 Addition and multiplication principles	326
51.6 Permutations	326
51.7 Combinations	327
51.8 Pascal identity	327
51.9 Binomial theorem	327
51.10 Counting subsets	327
51.11 Examples and exercises	328
51.12 Indicator functions as a unifying proof method	328
51.13 Combinatorial identities by double counting	329
51.14 Conceptual synthesis	329
51.15 Indicator functions and inclusion–exclusion	329
51.16 Inclusion–exclusion as Möbius inversion	329
51.17 Double counting as proof by projection	330
51.18 Generating functions	330
51.19 Pigeonhole principle and the probabilistic method	330
51.20 Asymptotics of binomial coefficients	330
51.21 Combinatorial species	331
51.22 Counting as inversion, projection, and generation	331
 N-20221207	 333
52.1 A set construction problem	333
52.2 What this problem teaches	334
52.3 A Venn-diagram necessary-and-sufficient problem	334
52.4 Difference identities	335
52.5 Coordinate and algebraic exercises	335
52.6 A note on proof writing	335
52.7 Extremal set problems: why constructions matter	335
52.8 Venn diagrams versus proofs	336
52.9 A wider frame	336
52.10 Extremal problems: upper bound plus construction	336
52.11 Additive combinatorics viewpoint	336
52.12 Cauchy–Davenport shadow	337
52.13 Boolean atoms and Venn regions	337
52.14 Symmetric difference as addition mod two	337
52.15 Sperner and intersecting-family context	337

52.16	Shifting and stability	338
52.17	Extremal sets, Boolean algebra, and stability	338
N-20221209		339
53.1	What this page is doing	339
53.2	A floor-function example	339
53.3	A binomial counting identity	339
53.4	Solving a system by discriminants	340
53.5	A prime-number exercise	340
53.6	A nested-set exercise	340
53.7	A recurrence of sets	340
53.8	An intersection-counting problem	341
53.9	Set-builder notation as a translation problem	341
53.10	Recursive set construction	341
53.11	Further direction	341
53.12	Floor functions as lattice-point counts	341
53.13	Finite differences and binomial basis	342
53.14	Recursive set construction and fixed points	342
53.15	Incidence counting	343
53.16	Polynomial method	343
53.17	Group actions and corrected division	343
53.18	Counting by lattices, incidences, and symmetries	343
N-20221210		345
54.1	Binomial theorem	345
54.2	Absolute value of binomial sums	345
54.3	A summation exercise	345
54.4	Set exercises	346
54.5	Counting subsets with properties	346
54.6	Set-builder notation	347
54.7	Closing perspective	347
54.8	Binomial identities as counting in two languages	347
54.9	Binomial transform and Möbius inversion	347
54.10	Finite differences and the binomial basis	348
54.11	Probability distribution behind binomial coefficients	348
54.12	Asymptotic shape: from binomial coefficients to Gaussians	349
54.13	Generating functions and labelled structures	349
54.14	Combinatorial species and relabelling	349
54.15	Binomial algebra from sets to species	349
N-20230118		351
55.1	Pigeonhole principle	351
55.2	Subsets and binomial coefficients	351
55.3	Binomial identities by double counting	351
55.4	A subset-sum style example	352
55.5	Generating functions	352
55.6	Helly-type remark	352
55.7	The next layer	352
55.8	Pigeonhole as averaging	352
55.9	Generating functions and subset sums	352
55.10	Helly-type thinking	353
55.11	Finite compactness principles in combinatorics	353
55.12	Pigeonhole as compactness principle	353
55.13	Ramsey theory	353
55.14	Helly theorem	353
55.15	Fractional and topological Helly	354
55.16	Local-to-global intersection principles	354
55.17	Ramsey as structured inevitability	354
N-20240323		355
56.1	Binomial identities	355

56.2	Good subsets	355
56.3	Generating functions	356
56.4	A probability example with repeated trials	356
56.5	Catalan numbers	357
56.6	Convex sets and convex hulls	357
56.7	Separating convex sets	357
56.8	Planar graphs and Euler-type counting	358
56.9	Graphs and adjacency matrices	358
56.10	The deeper habit	358
56.11	Generating functions as algebraic counting	358
56.12	Catalan structures	359
56.13	Convex hulls and separation	359
56.14	Generating functions, Catalan recursion, and spectra	359
56.15	Catalan numbers as recursive geometry	359
56.16	Lattice paths and reflection principle	360
56.17	Graph walks and adjacency matrices	360
56.18	Random walks and Markov chains	360
56.19	Catalan recursion and parenthesization	360
56.20	Adjacency spectra and expansion	361
N-20240330		363
57.1	Good triangles	363
57.2	Why vertices must lie on the convex hull	363
57.3	Cyclic arc encoding	363
57.4	The geometric idea behind the proof	364
57.5	Euler formula for convex polyhedra	364
57.6	Regular convex polyhedra	364
57.7	Euler characteristic	365
57.8	Degree of a map	365
57.9	Homology group and quotient group reminders	365
57.10	A bridge outward	365
57.11	Euler characteristic as invariant	366
57.12	Degree counting on graphs	366
57.13	Quotients and surface construction	366
57.14	Euler bookkeeping from graphs to surfaces	366
57.15	Euler characteristic as topological bookkeeping	366
57.16	Euler characteristic and homology	367
57.17	Degree theory	367
57.18	Quotient spaces and identifications	367
57.19	Euler characteristic as alternating rank	367
57.20	Degree and fixed-point consequences	368
N-20251126		369
58.1	What the model is trying to measure	369
58.2	Submultiplicativity and the connective constant	369
58.3	Bridges and why they help	370
58.4	Polygons	372
58.5	Critical exponents	373
58.6	How dimension changes the story	374
58.7	The broader lesson	374
N-20251127		377
59.1	Why the hexagonal lattice is special	377
59.2	A domain version of the counting problem	377
59.3	The parafermionic observable	378
59.4	From local cancellation to a global identity	378
59.5	The upper bound on the connective constant	380
59.6	The lower bound on the connective constant	380
59.7	Relation with SLE	380
59.8	Loop models and the same mechanism	381
59.9	From example to method	381

N-20251128	383
60.1 The high-dimensional philosophy	383
60.2 The two-point function	383
60.3 The lace expansion as corrected random walk	384
60.4 Bubble diagrams and diagrammatic estimates	384
60.5 Differential inequalities	385
60.6 Functional integral representations	385
60.7 Dimension four and logarithmic corrections	386
60.8 How the three methods fit together	387
60.9 The conceptual turn	387

Counting Principles, Binomial Identities, and Mixed Problem Solving

N-20220720

§50.1 Counting principles

The note begins with the standard counting principles:

Addition principle. If a task can be done in disjoint cases with $m_1, m_2, \dots, m_n$ possibilities, then the total number is $m_1 + \dots + m_n$.

Multiplication principle. If a task is done in stages with $m_1, m_2, \dots, m_n$ choices at each stage, then the total number is $m_1 m_2 \dots m_n$.

Inclusion–exclusion. For sets $A_1, \dots, A_n$, one adds single sizes, subtracts pairwise intersections, adds triple intersections, and so on.

A Venn diagram in the note reviews the three-set formula:

$$\begin{aligned} |A \cup B \cup C| &= |A| + |B| + |C| \\ &\quad - |A \cap B| - |A \cap C| - |B \cap C| + |A \cap B \cap C|. \end{aligned}$$

§50.2 A divisibility example

One problem asks how many positive integers not exceeding 2021 are neither multiples of 6 nor multiples of 8. The excluded set is

$$A \cup B,$$

where A is the set of multiples of 6, and B is the set of multiples of 8. Since

$$\text{lcm}(6, 8) = 24,$$

we have

$$|A| = 336, \quad |B| = 252, \quad |A \cap B| = 84.$$

Thus the number excluded is

$$336 + 252 - 84 = 504,$$

and the answer is

$$2021 - 504 = 1517.$$

This is exactly what inclusion–exclusion is for: avoid double-counting the common multiples.

§50.3 Permutations and combinations

The page distinguishes arrangements and selections:

$$A_n^m = \frac{n!}{(n-m)!}, \quad C_n^m = \frac{n!}{m!(n-m)!}.$$

A sequence problem asks for the number of 10-term sequences with $x + y = 10$ and $x - y = 4$, giving $x = 7, y = 3$. Therefore the number of sequences is

$$\binom{10}{3} = 120.$$

The thought process is: once the number of one kind of symbol is fixed, the sequence is determined by choosing its positions.

§50.4 Binomial theorem

The binomial theorem is

$$(a + b)^n = \sum_{k=0}^n \binom{n}{k} a^{n-k} b^k.$$

Several standard identities follow immediately:

$$\sum_{k=0}^n \binom{n}{k} = 2^n,$$

$$\sum_{k=0}^n (-1)^k \binom{n}{k} = 0,$$

and by differentiating

$$(1 + x)^n$$

and setting $x = 1$,

$$\sum_{k=0}^n k \binom{n}{k} = n2^{n-1}.$$

The page writes these identities as examples of how a coefficient formula becomes a counting tool.

§50.5 Polynomial coefficients

A worked example asks for the constant term or a particular coefficient in an expansion such as

$$(2x - 2y)^5.$$

The term involving $x^{5-k}y^k$ is

$$\binom{5}{k} (2x)^{5-k} (-2y)^k.$$

Thus the coefficient of x^2y^3 comes from $k = 3$:

$$\binom{5}{3} 2^2 (-2)^3 = -320.$$

The note's point is that coefficient problems become simple once the target exponent determines k .

§50.6 Further contest examples

The later pages mix counting with algebra, inequalities, and geometry. There are examples involving choosing books from different subjects, solving inequalities by sign charts, and geometry problems involving similar triangles. The unifying skill is case management: one must decide whether the problem is a sum of cases, a product of stages, or a coefficient extraction.

§50.7 Binomial identities by stories

The identity

$$\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1}$$

comes from choosing a committee of size k from n people. Either a distinguished person is not chosen, or that person is chosen. This kind of proof is often clearer than algebraic manipulation because it explains what the two terms mean.

§50.8 When to use multiplication and when to use addition

The addition principle applies to disjoint alternatives. The multiplication principle applies to successive independent choices. Many counting mistakes come from using multiplication when cases overlap, or using addition when a task has stages. A good habit is to write a sentence: “first choose ..., then choose ...” for multiplication, and “either ..., or ...” for addition.

§50.9 A broader reading

Elementary counting is a first meeting with combinatorial species and generating functions. Addition corresponds to disjoint union, multiplication corresponds to Cartesian product, and binomial coefficients are the coefficients of a generating function. This is why the binomial theorem is not just algebra: it is also a counting machine.

§50.10 Counting principles as operations on combinatorial classes

The addition and multiplication principles are not merely classroom rules. They are the first operations on combinatorial classes. If $\mathcal{A}$ and $\mathcal{B}$ are disjoint classes, then

$$|\mathcal{A} \sqcup \mathcal{B}| = |\mathcal{A}| + |\mathcal{B}|.$$

If an object is built by choosing one object from $\mathcal{A}$ and one from $\mathcal{B}$, then

$$|\mathcal{A} \times \mathcal{B}| = |\mathcal{A}||\mathcal{B}|.$$

Thus addition corresponds to disjoint union, and multiplication corresponds to Cartesian product.

This viewpoint is useful because it explains why formulas should follow from constructions. Before computing a number, identify whether the object is a disjoint alternative, a staged product, a quotient by symmetry, or a coefficient in a generating function.

§50.11 Inclusion–exclusion as Möbius inversion

The formula

$$|A \cup B| = |A| + |B| - |A \cap B|$$

is a small instance of Möbius inversion on the Boolean lattice. For a finite family $A_1, \dots, A_n$, the number of elements lying in at least one A_i is

$$\left| \bigcup_{i=1}^n A_i \right| = \sum_{\emptyset \neq S \subseteq [n]} (-1)^{|S|+1} \left| \bigcap_{i \in S} A_i \right|.$$

The alternating signs are the Möbius function of the subset poset:

$$\mu(S, T) = (-1)^{|T|-|S|}.$$

So inclusion–exclusion is not a trick for Venn diagrams; it is inversion of the operation that replaces exact membership data by cumulative membership data.

§50.12 Indicator functions and probability language

For a finite universe U , a set $A \subseteq U$ has indicator function

$$\mathbf{1}_A(x) = \begin{cases} 1, & x \in A, \\ 0, & x \notin A. \end{cases}$$

Then

$$|A| = \sum_{x \in U} \mathbf{1}_A(x).$$

The two-set inclusion–exclusion identity is the pointwise identity

$$\mathbf{1}_{A \cup B} = \mathbf{1}_A + \mathbf{1}_B - \mathbf{1}_A \mathbf{1}_B.$$

Summing over U gives the cardinality formula.

This is also the beginning of probability. If X is uniformly chosen from U , then

$$\mathbb{P}(X \in A) = \frac{|A|}{|U|} = \mathbb{E}[\mathbf{1}_A(X)].$$

Counting and expectation become the same operation after normalization.

§50.13 Generating functions for binomial identities

The binomial theorem says

$$(1+x)^n = \sum_{k=0}^n \binom{n}{k} x^k.$$

The coefficient of x^k counts subsets of size k . Differentiating and setting $x = 1$ gives

$$\sum_{k=0}^n k \binom{n}{k} = n2^{n-1},$$

which counts pairs (S, s) where $S \subseteq [n]$ and $s \in S$: choose s first, then choose the rest of S .

This illustrates the right use of generating functions: algebraic operations on series encode operations on counted objects.

§50.14 Counting up to symmetry: Burnside's lemma

If a finite group G acts on a finite set X , the number of orbits is

$$|X/G| = \frac{1}{|G|} \sum_{g \in G} |\text{Fix}(g)|.$$

This is Burnside's lemma. It explains why one cannot always divide by the number of symmetries: objects with nontrivial stabilizers have smaller orbits.

Elementary examples include necklaces, colorings of polygons, and arrangements considered up to rotation or reflection. The advanced lesson is that symmetry changes the unit of counting from elements to orbits.

§50.15 Counting through structure and symmetry

The counting principles become powerful when read structurally. Addition is disjoint union, multiplication is product, inclusion–exclusion is Möbius inversion, binomial coefficients are generating-function coefficients, and symmetry counting is orbit counting. The elementary skill is to decide which structure the problem is really using before doing arithmetic.

Lecture Notes on Sets, Counting, and Binomial Coefficients

N-20221205

§51.1 Summation and product notation

This note is a lecture-style introduction to sets and elementary combinatorics. It begins with notation. If $m \leq n$, then

$$\sum_{k=m}^n a_k = a_m + a_{m+1} + \cdots + a_n,$$

and

$$\prod_{k=m}^n a_k = a_m a_{m+1} \cdots a_n.$$

Nested sums are allowed. For example,

$$\sum_{j=1}^n \sum_{k=1}^j a_k = a_1 + (a_1 + a_2) + \cdots + (a_1 + a_2 + \cdots + a_n).$$

The notation is not merely shorthand. It teaches us to control indices, which is essential in combinatorics and analysis.

§51.2 De Morgan's laws

For subsets A, B of a universe U , De Morgan's laws say

$$\overline{A \cup B} = \overline{A} \cap \overline{B},$$

and

$$\overline{A \cap B} = \overline{A} \cup \overline{B}.$$

The proof is an element-chasing proof. For the first identity:

$$\begin{aligned} x \in \overline{A \cup B} &\iff x \notin A \cup B \\ &\iff x \notin A \text{ and } x \notin B \\ &\iff x \in \overline{A} \cap \overline{B}. \end{aligned}$$

The handwritten margin proves the second identity in the same style. The important habit is to translate set operations into logical words: union means “or”, intersection means “and”, complement means “not.”

§51.3 Inclusion–exclusion

For finite sets,

$$|A \cup B| = |A| + |B| - |A \cap B|.$$

For three sets,

$$\begin{aligned} |A \cup B \cup C| &= |A| + |B| + |C| \\ &\quad - |A \cap B| - |A \cap C| - |B \cap C| \\ &\quad + |A \cap B \cap C|. \end{aligned}$$

The note derives the three-set formula by applying the two-set formula to

$$A \cup (B \cup C).$$

The overcounting pattern is worth remembering: count singles, subtract pairwise overlaps, add triple overlaps back. Higher versions continue with alternating signs.

§51.4 Pigeonhole principle

The pigeonhole principle states: if $n+1$ objects are placed into n boxes, then at least one box contains at least two objects. This is almost obvious, but it is extremely powerful.

The useful mental form is contrapositive: if every box contains at most one object, then n boxes contain at most n objects. Whenever a problem asks us to prove two objects share a property, try to define boxes by that property.

§51.5 Addition and multiplication principles

The addition principle says: if a task can be completed in disjoint ways, with m_k possibilities in case k , then the total number is

$$\sum_k m_k.$$

The note's example is choosing a route from home to school: if the alternatives are parallel choices, add them.

The multiplication principle says: if a task is completed in stages, with m_k choices at stage k , then the total number is

$$\prod_k m_k.$$

This is sequential choice. The key distinction is:

$$\text{cases} \Rightarrow \text{add}, \quad \text{stages} \Rightarrow \text{multiply}.$$

Many counting mistakes come from confusing these two.

§51.6 Permutations

To arrange n distinct objects without repetition, there are

$$n! = n(n-1)(n-2) \cdots 2 \cdot 1$$

ways. If we choose and arrange m objects from n , the number is

$$A_n^m = n(n-1) \cdots (n-m+1) = \frac{n!}{(n-m)!}.$$

The multiplication principle explains the formula: first choose the first object, then the second, and so on, with one fewer option each time.

§51.7 Combinations

If order does not matter, we divide by the number of internal arrangements. Thus the number of m -element subsets of an n -element set is

$$\binom{n}{m} = \frac{n!}{m!(n-m)!}.$$

The note remarks that different exams may write this as C_n^m or $\binom{n}{m}$. The mathematical object is the same.

The symmetry

$$\binom{n}{m} = \binom{n}{n-m}$$

has a simple interpretation: choosing the m elements included is the same as choosing the $n - m$ elements excluded.

§51.8 Pascal identity

The note asks the reader to verify

$$\binom{n}{m} = \binom{n-1}{m} + \binom{n-1}{m-1}.$$

A combinatorial proof: count m -subsets of an n -element set according to whether they contain a distinguished element x . If they do not contain x , choose all m elements from the remaining $n - 1$. If they do contain x , choose the remaining $m - 1$ elements from the remaining $n - 1$.

This proof is better than algebraic manipulation because it explains why the identity exists.

§51.9 Binomial theorem

The binomial theorem is

$$(a + b)^n = \sum_{j=0}^n \binom{n}{j} a^{n-j} b^j.$$

The note suggests two proofs. The combinatorial proof expands

$$(a + b)(a + b) \cdots (a + b).$$

To get $a^{n-j} b^j$, choose exactly j of the n brackets to contribute b . This can be done in $\binom{n}{j}$ ways.

The induction proof is longer but mechanical. Both are useful, but the combinatorial proof reveals the meaning of the coefficients.

§51.10 Counting subsets

If a set S has n elements, then it has

$$2^n$$

subsets. One proof is to count by size:

$$\sum_{m=0}^n \binom{n}{m} = (1+1)^n = 2^n.$$

Another proof is binary choice: each element is either included or excluded.

The second proof is usually the best mental model. A subset is a yes/no decision for each element.

§51.11 Examples and exercises

The later pages contain examples involving set-builder notation, floor functions, integer lattice points, and set operations. The method is always to translate definitions into simpler conditions.

For instance, if

$$A = \{x : x^2 - 6 < 0\},$$

then

$$-\sqrt{6} < x < \sqrt{6}.$$

If a set uses the floor function $[x]$, split the real line into intervals

$$k \leq x < k+1.$$

If a set contains integer pairs satisfying

$$x^2 + y^2 \leq 1,$$

then draw the tiny lattice disk and count directly.

The note also has exercises asking whether certain sets intersect a circle or line family. The principle is again translation: set notation becomes equations, equations become geometry.

§51.12 Indicator functions as a unifying proof method

Many counting identities become transparent with indicator functions. For a set A , define

$$1_A(x) = \begin{cases} 1, & x \in A, \\ 0, & x \notin A. \end{cases}$$

Then

$$|A| = \sum_x 1_A(x)$$

when the universe is finite. De Morgan's laws, inclusion-exclusion, and complement counting can all be proved by checking the identity pointwise for each element x .

For two sets,

$$1_{A \cup B} = 1_A + 1_B - 1_{A \cap B}.$$

Summing over the universe gives

$$|A \cup B| = |A| + |B| - |A \cap B|.$$

For many sets, the same pointwise idea gives the full inclusion-exclusion formula.

§51.13 Combinatorial identities by double counting

The binomial identity

$$\binom{n}{r} = \binom{n-1}{r} + \binom{n-1}{r-1}$$

can be proved by deciding whether a distinguished element is chosen. If it is not chosen, choose all r elements from the remaining $n-1$. If it is chosen, choose $r-1$ more.

This is the basic style of double counting: count the same collection in two ways. It is more robust than algebraic manipulation because it explains why the identity is true.

§51.14 Conceptual synthesis

This lecture is not just “sets and counting.” It introduces three ways of thinking that remain important everywhere: logical translation, structural counting, and representation switching. De Morgan turns set algebra into propositional logic; inclusion–exclusion counts by correcting overcounting; binomial coefficients count subsets and coefficients at once. These are the first signs that combinatorics is both discrete and algebraic.

§51.15 Indicator functions and inclusion–exclusion

For a finite universe U , define

$$\mathbf{1}_A(x) = \begin{cases} 1, & x \in A, \\ 0, & x \notin A. \end{cases}$$

Then

$$|A| = \sum_{x \in U} \mathbf{1}_A(x).$$

The identity

$$\mathbf{1}_{A \cup B} = \mathbf{1}_A + \mathbf{1}_B - \mathbf{1}_A \mathbf{1}_B$$

proves inclusion–exclusion after summing. For many sets, the product

$$\prod_{i=1}^n (1 - \mathbf{1}_{A_i})$$

encodes the complement of the union, and expanding it gives the full alternating formula.

§51.16 Inclusion–exclusion as Möbius inversion

On the Boolean lattice $\mathcal{P}([n])$, the zeta transform sends exact data to cumulative data:

$$g(S) = \sum_{T \subseteq S} f(T).$$

Möbius inversion recovers f :

$$f(S) = \sum_{T \subseteq S} (-1)^{|S|-|T|} g(T).$$

Inclusion–exclusion is precisely this inversion applied to membership in intersections of events.

This explains the alternating signs structurally. They are the inverse of summing over subsets.

§51.17 Double counting as proof by projection

A double-counting proof counts the same finite set of pairs in two projections. For example, the identity

$$\sum_{k=0}^n k \binom{n}{k} = n2^{n-1}$$

counts pairs (S, s) where $S \subseteq [n]$ and $s \in S$. Counting by S gives the left side; counting by s gives the right side.

This viewpoint is powerful because it forces one to name the object being counted. Once the object is named, the identity is no longer mysterious.

§51.18 Generating functions

Ordinary generating functions package a sequence into

$$A(x) = \sum_{n \geq 0} a_n x^n.$$

The coefficient extraction notation

$$[x^n]A(x) = a_n$$

turns algebra into counting. For example,

$$(1+x)^n = \sum_{k=0}^n \binom{n}{k} x^k$$

records subset counts by size.

Products of generating functions correspond to splitting an object into parts. This is the algebraic form of the multiplication principle.

§51.19 Pigeonhole principle and the probabilistic method

The pigeonhole principle is an averaging statement. If N objects are placed into r boxes, some box contains at least $\lceil N/r \rceil$ objects. The probabilistic method generalizes this: to prove that an object with a desired property exists, choose a random object and show that the probability of success is positive.

For instance, if the expected number of bad events is less than 1, then there must be a choice with no bad events. This is a research-level continuation of the same counting intuition: average behavior forces existence.

§51.20 Asymptotics of binomial coefficients

Pascal's triangle has a limiting shape. Stirling's formula gives

$$\binom{n}{\lfloor n/2 \rfloor} \sim \frac{2^n}{\sqrt{\pi n/2}}.$$

Probabilistically, the binomial coefficients are the distribution of a sum of independent Bernoulli variables. After centering and scaling, the central limit theorem gives a Gaussian approximation.

So the finite identity $\sum_k \binom{n}{k} = 2^n$ is the exact count, while asymptotic analysis explains the shape of the whole row for large n .

§51.21 Combinatorial species

A combinatorial species assigns to each finite label set U a set $F[U]$ of structures, compatibly with relabelling bijections. The species of subsets sends U to $\mathcal{P}(U)$. A bijection of label sets transports subsets.

Species theory records not just how many objects exist, but how constructions behave under relabelling. Addition, product, and composition of species categorify the elementary counting principles.

§51.22 Counting as inversion, projection, and generation

The lecture's elementary tools have precise advanced continuations. Set logic becomes Boolean algebra, inclusion–exclusion becomes Möbius inversion, double counting becomes projection of incidence sets, generating functions encode sequences, pigeonhole becomes averaging and probabilistic existence, and binomial coefficients connect exact counting with asymptotic probability.

Set Problems, Extremal Counting, and Venn-Diagram Reasoning

N-20221207

§52.1 A set construction problem

The first problem defines a finite set

$$A = \{a_1 < a_2 < \cdots < a_k\}$$

of positive integers and builds two derived sets, one from sums and one from differences. The handwritten solution records the important idea: use the union $S \cup T$, and do not try to identify every element exactly. Bound its size from below and above.

The sum set contains at least the ordered chain

$$2a_1, a_1 + a_2, 2a_2, \dots, a_{k-1} + a_k, 2a_k,$$

so

$$|S| \geq 2k - 1.$$

The difference set contains

$$0, a_2 - a_1, a_3 - a_1, \dots, a_k - a_1,$$

so

$$|T| \geq k.$$

Under the condition of the problem, $S \cap T = \emptyset$, hence

$$|S \cup T| = |S| + |T| \geq 3k - 1.$$

On the other hand, all these numbers lie between 0 and $2a_k$, so

$$|S \cup T| \leq 2a_k + 1.$$

Thus

$$3k - 1 \leq 2a_k + 1.$$

If $a_k \leq 2021$, then

$$k \leq \frac{2 \cdot 2021 + 2}{3} = 1348.$$

The construction

$$A = \{674, 675, \dots, 2021\}$$

attains this bound, so the maximum is

$$|A| = 1348.$$

§52.2 What this problem teaches

The note's blue comment says that using $S \cup T$ is a good way to avoid proving a complicated extremal statement directly. This is the key lesson. When a problem asks for the largest possible size of a set, it often helps to build auxiliary sets whose sizes are forced by the structure.

A lower bound comes from a construction. An upper bound comes from counting constraints. A complete solution needs both.

§52.3 A Venn-diagram necessary-and-sufficient problem

Another problem defines set difference

$$A - B = \{x : x \in A, x \notin B\}$$

and asks for the relation among A, B, C when

$$(A - B) \cup (B - A) \subseteq C.$$

The question asks whether this is equivalent to a condition involving

$$A \subseteq (C - B) \cup (B - C)$$

or


$$A \cap B \cap C = \emptyset.$$

The handwritten solution first draws a Venn diagram, then rewrites the pieces algebraically. The symmetric difference

$$(A - B) \cup (B - A)$$

is the part where membership in A and B differs. If this is contained in C , then outside C the sets A and B must agree.

The correct way to solve such problems is to analyze each of the eight Venn regions. The note records a false start and then corrects it, which is valuable: set identities are easy to get wrong if one reasons only verbally.



A Venn diagram is a region bookkeeping device.

Figure 52.1: For three sets one should reason region by region, not by visual guesswork alone.

§52.4 Difference identities

One useful identity appearing in the handwritten reasoning is

$$B - C = B \cap C^c.$$

Therefore

$$B - C - (A \cap B) = B \cap C^c \cap (A \cap B)^c.$$

Using De Morgan,

$$(A \cap B)^c = A^c \cup B^c,$$

so

$$B \cap C^c \cap (A^c \cup B^c) = (B \cap C^c \cap A^c) \cup (B \cap C^c \cap B^c).$$

The second term is empty, hence

$$B - C - (A \cap B) = B \cap A^c \cap C^c.$$

This is exactly the kind of algebraic cleanup that a Venn diagram suggests but does not replace.

§52.5 Coordinate and algebraic exercises

The later pages contain a mixture of algebraic and coordinate exercises. They include problems about quadratic expressions, tangent or chord diagrams, and conditions that must be translated into equations. The handwritten computations repeatedly use the same pattern:

- (i) introduce the unknowns explicitly;
- (ii) translate each condition into an equation or inequality;
- (iii) reduce by substitution or by completing the square;
- (iv) check equality or boundary cases.

Several sketches show lines touching curves or intersecting regions. Rather than reproducing them as raw images, the important content is the method: a picture tells us which algebraic condition to impose, such as tangency, parallelism, or containment.

§52.6 A note on proof writing

The red handwritten reflection says that a purely diagrammatic method is intuitive but can be unreliable. The mature solution uses the diagram to discover the statement, then uses set algebra to prove it. This is an excellent general lesson. Geometry and diagrams are for finding the path; formal manipulations are for making the path safe.

§52.7 Extremal set problems: why constructions matter

In an extremal counting problem, an upper bound is only half the solution. One must also give a construction showing the bound is sharp. The note's sum-set and difference-set arguments follow this pattern: first prove that the derived sets must contain enough elements; then exhibit a set A for which the lower bound is actually attained.

This is a common combinatorial habit. Inequalities show impossibility beyond a threshold. Constructions show that the threshold is real.

§52.8 Venn diagrams versus proofs

Venn diagrams are excellent for discovering set identities, but the proof should eventually be written elementwise. For example,

$$A \setminus (B \cup C) = (A \setminus B) \cap (A \setminus C)$$

is proved by taking an arbitrary x and translating membership:

$$x \in A \setminus (B \cup C) \iff x \in A, x \notin B, x \notin C.$$

This is exactly the condition for membership in the right-hand side. The diagram gives intuition; the element chase gives proof.

§52.9 A wider frame

This note is about extremal and logical reasoning. The extremal problem uses auxiliary sets and cardinality bounds. The Venn problem uses Boolean algebra and region decomposition. Both are early forms of a common method: replace a vague configuration by a finite list of regions or quantities that can be counted exactly. This is the same spirit behind combinatorial proofs, measure decompositions, and even sigma-algebra calculations later in analysis.

§52.10 Extremal problems: upper bound plus construction

The sum-set/difference-set problem has the standard shape of extremal combinatorics. One proves an upper bound by showing that certain auxiliary objects must fit inside a finite container. Then one gives a construction meeting the bound.

This two-part structure is essential:

$$\text{upper bound} + \text{matching construction} = \text{exact extremal value}.$$

A construction alone proves possibility; an upper bound alone proves limitation. A complete extremal solution needs both.

§52.11 Additive combinatorics viewpoint

For a finite set $A \subseteq \mathbb{Z}$, the sum set and difference set are

$$A + A = \{a + a' : a, a' \in A\}, \quad A - A = \{a - a' : a, a' \in A\}.$$

The size of these sets measures additive structure. Arithmetic progressions have unusually small sumsets. Random sets usually have much larger sumsets.

A central theorem in additive combinatorics is Freiman's philosophy: if $|A + A|$ is small compared with $|A|$, then A must resemble a generalized arithmetic progression. The elementary problem is a small concrete case of studying what set structure is forced by sum and difference constraints.

§52.12 Cauchy–Davenport shadow

In a prime cyclic group $\mathbb{Z}/p\mathbb{Z}$, the Cauchy–Davenport theorem says

$$|A + B| \geq \min(p, |A| + |B| - 1).$$

This is the modular analogue of the simple lower bound that a sumset of ordered integer sets contains a chain of at least $|A| + |B| - 1$ elements.

The theorem shows that the lower-bound logic in the note is not accidental. Sumsets are forced to grow unless the ambient group wraps around or the set has special structure.

§52.13 Boolean atoms and Venn regions

For three sets A, B, C , the universe is partitioned into at most eight atoms:

$$A^{\epsilon_1} \cap B^{\epsilon_2} \cap C^{\epsilon_3}, \quad \epsilon_i \in \{0, 1\},$$

where $A^1 = A$ and $A^0 = A^c$. Every Boolean expression in A, B, C is a union of some of these atoms.

This gives a systematic proof method: reduce both sides of a set identity to the same atoms. Venn diagrams are pictures of this atomic decomposition; elementwise proofs are the formal version.

§52.14 Symmetric difference as addition mod two

The symmetric difference

$$A \triangle B = (A \setminus B) \cup (B \setminus A)$$

behaves like addition modulo 2 on characteristic functions:

$$\mathbf{1}_{A \triangle B} = \mathbf{1}_A + \mathbf{1}_B \pmod{2}.$$

Thus the power set $\mathcal{P}(U)$ becomes a vector space over $\mathbb{F}_2$ under symmetric difference. Intersection acts like multiplication.

This explains why expressions such as $(A - B) \cup (B - A)$ have algebraic structure: they are not just set differences, but mod-two sums of predicates.

§52.15 Sperner and intersecting-family context

Many extremal set problems ask for the largest family avoiding a containment or disjointness pattern. Sperner's theorem says that an antichain $\mathcal{F} \subseteq \mathcal{P}([n])$ satisfies

$$|\mathcal{F}| \leq \binom{n}{\lfloor n/2 \rfloor}.$$

The largest example is the middle layer.

For intersecting families, the Erdős–Ko–Rado theorem says that if $\mathcal{F} \subseteq \binom{[n]}{k}$ is pairwise intersecting and $n \geq 2k$, then

$$|\mathcal{F}| \leq \binom{n-1}{k-1}.$$

The extremal example is a star: all k -sets containing a fixed element.

These theorems are advanced versions of the note's method: identify the forbidden configuration, prove an upper bound, and classify sharp examples.

§52.16 Shifting and stability

Shifting transforms a set family into a more ordered one without changing its size or destroying the relevant forbidden configuration. Repeated shifting often reduces an extremal problem to a canonical family.

A stability theorem asks more than the maximum: if a family is close to maximum size, must it be structurally close to an extremal example? This is the research-level continuation of checking equality cases in an elementary inequality.

§52.17 Extremal sets, Boolean algebra, and stability

The note combines two important habits. Extremal counting uses auxiliary sets, upper bounds, and sharp constructions. Venn reasoning uses Boolean atoms and set algebra. At a higher level these become additive combinatorics, extremal set theory, vector-space structure over $\mathbb{F}_2$, and stability methods.

Set-Theoretic and Combinatorial Exercises after the First Lecture

N-20221209

§53.1 What this page is doing

This note is a handwritten exercise continuation of the set-and-counting lecture. The problems involve floor functions, set-builder notation, binomial identities, lattice or coordinate sets, and counting arguments. The common goal is to learn how to translate a compact set expression into something countable.

§53.2 A floor-function example

One example asks for a set defined by inequalities and floor functions. The handwritten solution separates the problem into intervals because

$$\lfloor x \rfloor = k \iff k \leq x < k + 1.$$

This is the essential method. Any problem involving $\lfloor x \rfloor$ should first be split into integer strips.

For instance, if a condition contains

$$3\lfloor x \rfloor - 5 < 0,$$

then

$$\lfloor x \rfloor < \frac{5}{3},$$

so

$$\lfloor x \rfloor \leq 1.$$

One then intersects this with the remaining condition, such as a quadratic inequality. The final answer is an interval or a finite union of intervals.

§53.3 A binomial counting identity

Another exercise invokes the binomial theorem. The handwritten note uses the idea that terms in a sum can be counted by choosing which positions have a particular property. A typical identity is

$$\sum_{j=0}^n \binom{n}{j} = 2^n,$$

or, with signs,

$$\sum_{j=0}^n (-1)^j \binom{n}{j} = 0.$$

The important thought is not the formula alone. The coefficient $\binom{n}{j}$ counts ways of choosing j positions among n , and the binomial theorem packages all choices simultaneously.

§53.4 Solving a system by discriminants

One problem gives a two-variable system whose equations can be transformed into a quadratic with a parameter p . The handwritten solution examines the discriminant

$$\Delta \geq 0$$

to decide when real solutions exist.

This is a recurring analytic-algebraic pattern: if a system can be reduced to a quadratic equation, existence of real solutions is controlled by the discriminant. The solution set changes when the discriminant becomes zero; this is the boundary between two real roots and no real roots.

§53.5 A prime-number exercise

Another exercise involves primes p, q . The handwritten reasoning uses parity and basic divisibility. The standard approach is:

- (i) reduce the equation modulo a small integer such as 2, 3, 4;
- (ii) use the fact that the only even prime is 2;
- (iii) separate small exceptional cases before treating odd primes.

This is the same habit developed in earlier divisibility notes: do not solve a Diophantine equation blindly. First use congruences to force the shape of the variables.

§53.6 A nested-set exercise

The note includes a Russian sentence asking to “describe the next set” or “write the following set in ordinary language.” The examples involve finite sets such as

$$\emptyset, \quad \{a_1\}, \quad \{a_1, a_2\}, \quad \dots$$

and collections of subsets. The intended lesson is the difference between an element and a subset:

$$a \in A \quad \text{versus} \quad \{a\} \subseteq A.$$

Confusing these two is one of the most common early set-theory mistakes.

§53.7 A recurrence of sets

One exercise defines sets recursively, for example

$$A_{k+1} = A_k \cup \{x_{k+1}\}$$

or by adding all new subsets involving a new element. The handwritten solution proves the pattern by induction: assume the formula for A_k , then show the construction gives A_{k+1} .

This is the set-theoretic analogue of a recurrence relation. Instead of numbers changing with k , collections of objects change with k .

§53.8 An intersection-counting problem

The final exercise on the second page asks whether a selected element belongs to an intersection of many sets. The handwritten solution builds sets $B(A_i)$ and counts their sizes, arriving at a large product-like number. The exact numerical expression is less important than the logic:

- (i) translate membership into a condition;
- (ii) count how many sets contain a fixed element;
- (iii) use complements or intersections to isolate the desired elements.

This is a typical finite-set counting method: count incidence pairs in two ways.

§53.9 Set-builder notation as a translation problem

A complicated set-builder expression should be read in stages. First identify the universe of discourse. Then translate each condition into an inequality, divisibility condition, or membership condition. Finally decide whether the problem is asking for the set itself, its size, or a proof of equality.

For example, a condition involving $\lfloor x \rfloor$ is not smooth; it divides the real line into half-open intervals. A condition involving binomial coefficients often asks for a counting interpretation rather than algebraic simplification.

§53.10 Recursive set construction

When a set is defined recursively, the first few stages should be computed explicitly. This reveals whether the construction is increasing, periodic, stabilizing, or producing genuinely new elements. After that, induction is the natural proof method.

The general pattern is:

$$S_0 \subseteq S_1 \subseteq S_2 \subseteq \cdots$$

if the rule is monotone. If the universe is finite, an increasing chain must eventually stabilize. This is a useful idea behind many finite combinatorial processes.

§53.11 Further direction

The exercises are elementary, but they are training the language of discrete mathematics. Floor functions partition the real line into strips; binomial coefficients count choices; discriminants detect existence; recursive set definitions invite induction; incidence counting turns membership questions into arithmetic. The high-level skill is always the same: replace notation by a structure that can actually be counted.

§53.12 Floor functions as lattice-point counts

A floor expression such as

$$\left\lfloor \frac{an + b}{m} \right\rfloor$$

counts how many integer thresholds lie below a rational expression. Sums of floor functions often count lattice points below a line. For instance,

$$\sum_{k=0}^{m-1} \left\lfloor \frac{ak}{m} \right\rfloor$$

counts integer points under a rational-slope diagonal in a rectangle.

This links elementary floor exercises to Ehrhart theory, where one studies

$$L_P(n) = |nP \cap \mathbb{Z}^d|$$

for a polytope P . For integral polytopes, $L_P(n)$ is a polynomial in n . Floor-sum exercises are one-dimensional shadows of lattice-point enumeration.

§53.13 Finite differences and binomial basis

For a sequence $a(n)$, define

$$\Delta a(n) = a(n+1) - a(n).$$

If $a(n)$ is a polynomial of degree d , then Δa has degree $d-1$, and

$$\Delta^{d+1} a(n) = 0.$$

The binomial basis is adapted to this:

$$\Delta \binom{n}{k} = \binom{n}{k-1}.$$

Thus integer-valued polynomials are naturally expressed in the basis

$$\binom{n}{0}, \quad \binom{n}{1}, \quad \binom{n}{2}, \quad \dots$$

rather than ordinary powers. This explains why binomial coefficients repeatedly appear in discrete counting formulas.

§53.14 Recursive set construction and fixed points

A recursive set construction such as

$$S_{k+1} = F(S_k)$$

should first be checked for monotonicity. If $S_k \subseteq S_{k+1}$ inside a finite universe, then the chain

$$S_0 \subseteq S_1 \subseteq S_2 \subseteq \dots$$

must stabilize. The final set is a fixed point:

$$F(S) = S.$$

This is the finite version of closure-operator and fixed-point thinking.

§53.15 Incidence counting

If a problem asks how many sets contain a fixed element, it often helps to count incidence pairs:

$$I = \{(x, A) : x \in A\}.$$

Counting by x gives

$$|I| = \sum_x \#\{A : x \in A\},$$

while counting by A gives

$$|I| = \sum_A |A|.$$

Equating the two is double counting. Many intersection and membership problems are incidence problems in disguise.

§53.16 Polynomial method

The polynomial method proves combinatorial facts by encoding conditions into polynomial vanishing. The elementary fact is that a nonzero polynomial of degree d over a field has at most d roots. If a low-degree polynomial is forced to vanish at too many points, it must be the zero polynomial.

Lagrange interpolation gives an explicit reconstruction formula:

$$P(x) = \sum_{i=1}^n P(a_i) \prod_{j \neq i} \frac{x - a_j}{a_i - a_j}.$$

This is the advanced extension of discriminant and coefficient arguments: existence questions can be turned into degree and vanishing constraints.

§53.17 Group actions and corrected division

When counting objects up to symmetry, one must account for stabilizers. If a finite group G acts on X , orbit-stabilizer says

$$|Gx| = \frac{|G|}{|\text{Stab}(x)|}.$$

Burnside's lemma says

$$|X/G| = \frac{1}{|G|} \sum_{g \in G} |\text{Fix}(g)|.$$

This is why symmetry counting cannot always be done by simply dividing by $|G|$.

§53.18 Counting by lattices, incidences, and symmetries

The exercises train translation. Floor functions translate to lattice-point counts, binomial identities to finite differences and coefficient extraction, recursive set definitions to fixed points, membership counts to incidence pairs, and symmetry counts to group actions. The mature solution first finds the structure being counted, then chooses the appropriate algebraic language.

Binomial Theorem, Absolute-Value Estimates, and Set Exercises

N-20221210

§54.1 Binomial theorem

The first page begins with the binomial theorem:

$$(a + b)^n = \sum_{i=0}^n \binom{n}{i} a^i b^{n-i}.$$

The handwritten derivation emphasizes selection. When expanding

$$(a + b)(a + b) \cdots (a + b),$$

one chooses either a or b from each factor. A term $a^i b^{n-i}$ occurs when a is chosen from exactly i of the n factors, and this can be done in $\binom{n}{i}$ ways.

The same idea proves the recurrence

$$\binom{n+1}{i} = \binom{n}{i} + \binom{n}{i-1},$$

because in choosing i objects from $n + 1$, either the new object is not chosen or it is chosen.

§54.2 Absolute value of binomial sums

The note then estimates expressions such as

$$|a + b|^n.$$

Using the triangle inequality and the binomial theorem,

$$|a + b|^n \leq (|a| + |b|)^n = \sum_{i=0}^n \binom{n}{i} |a|^i |b|^{n-i}.$$

The important habit is to move absolute values inside after expanding only when one has a legitimate inequality. The handwritten page marks the triangle inequality as the justification.

§54.3 A summation exercise

One exercise asks for a sum of the form

$$\sum_{i=0}^n \binom{n}{i}.$$

Putting $a = b = 1$ in the binomial theorem gives

$$2^n = \sum_{i=0}^n \binom{n}{i}.$$

Similarly, putting $a = 1, b = -1$ gives

$$0 = (1 - 1)^n = \sum_{i=0}^n (-1)^{n-i} \binom{n}{i},$$

which shows that the sum of even-indexed binomial coefficients equals the sum of odd-indexed ones when $n > 0$.

§54.4 Set exercises

The printed pages list elementary set exercises. The notes use standard notation:

$$A \cup B, \quad A \cap B, \quad A \setminus B, \quad A^c.$$

One recurring task is to translate a verbal condition into set operations. For example,

$$A \setminus (B \cup C) = (A \setminus B) \cap (A \setminus C).$$

This is proved elementwise:

$$x \in A \setminus (B \cup C) \iff x \in A, x \notin B, x \notin C.$$

The exercises also ask for necessary and sufficient conditions for equalities such as

$$A \cup B = A, \quad A \cap B = B.$$

Both are equivalent to

$$B \subseteq A.$$

§54.5 Counting subsets with properties

For a finite set X with $|X| = n$, the number of subsets is

$$2^n.$$

If one imposes parity conditions, the count is often half of the total. For example, the number of subsets of an n -element set with even cardinality is

$$2^{n-1},$$

when $n \geq 1$. This follows from

$$\sum_{k \text{ even}} \binom{n}{k} = \sum_{k \text{ odd}} \binom{n}{k} = 2^{n-1}.$$

The printed exercises include conditions involving positive and negative numbers, odd and even elements, and sums of elements. The general method is to split the set into independent categories and multiply the numbers of choices.

§54.6 Set-builder notation

The note also works with sets such as

$$U = \{x : x = 2n, n \in \mathbb{Z}\}, \quad V = \{x : x = 3n, n \in \mathbb{Z}\}.$$

Here U is the set of even integers and V is the set of multiples of 3. Then

$$U \cap V = \{x : x = 6n, n \in \mathbb{Z}\}$$

and

$$U \cup V$$

is the set of integers divisible by 2 or by 3.

The point is to read set-builder notation as a logical predicate.

§54.7 Closing perspective

The note links binomial coefficients with set counting. The binomial theorem counts choices from factors; subset counting counts choices from elements. The same coefficients $\binom{n}{k}$ appear in both languages. Set identities should be proved elementwise, while counting identities should be proved by telling a story about what is being chosen.

§54.8 Binomial identities as counting in two languages

The binomial theorem has two simultaneous meanings. Algebraically, it is coefficient extraction:

$$(1+x)^n = \sum_{k=0}^n \binom{n}{k} x^k.$$

Combinatorially, it counts subsets by size. The coefficient of x^k is the number of ways to choose k marked factors, equivalently a k -element subset of an n -element set.

This double meaning is why identities should be checked in two ways whenever possible. For example, Vandermonde's identity

$$\sum_k \binom{r}{k} \binom{s}{n-k} = \binom{r+s}{n}$$

is coefficient extraction from

$$(1+x)^r(1+x)^s = (1+x)^{r+s},$$

but it also counts choosing n elements from a disjoint union of an r -set and an s -set: choose k from the first part and $n-k$ from the second.

§54.9 Binomial transform and Möbius inversion

The binomial transform of a sequence (a_k) is

$$b_n = \sum_{k=0}^n \binom{n}{k} a_k.$$

The inverse is

$$a_n = \sum_{k=0}^n (-1)^{n-k} \binom{n}{k} b_k.$$

This is not a coincidence; it is Möbius inversion on the Boolean lattice of subsets.

Indeed, if $B(S) = \sum_{T \subseteq S} A(T)$, then

$$A(S) = \sum_{T \subseteq S} (-1)^{|S|-|T|} B(T).$$

When $A(T)$ only depends on $|T|$, this collapses exactly to the binomial inversion formula above. Thus inclusion–exclusion, binomial inversion, and subset counting are the same mechanism at different resolutions.

§54.10 Finite differences and the binomial basis

The polynomials

$$\binom{x}{0}, \quad \binom{x}{1}, \quad \binom{x}{2}, \quad \dots$$

form a natural basis for polynomial functions on integers. The finite-difference operator

$$\Delta f(x) = f(x+1) - f(x)$$

acts especially simply:

$$\Delta \binom{x}{k} = \binom{x}{k-1}.$$

This mirrors ordinary calculus, where differentiation lowers powers:

$$\frac{d}{dx} x^k = kx^{k-1}.$$

The binomial basis is therefore the right coordinate system for discrete calculus. Identities involving $\binom{n}{k}$ often become transparent after applying Δ or summing finite differences.

§54.11 Probability distribution behind binomial coefficients

If $X \sim \text{Bin}(n, p)$, then

$$\mathbb{P}(X = k) = \binom{n}{k} p^k (1-p)^{n-k}.$$

The binomial theorem says the probabilities sum to one:

$$\sum_{k=0}^n \binom{n}{k} p^k (1-p)^{n-k} = (p + 1 - p)^n = 1.$$

The probability generating function is

$$G_X(t) = \mathbb{E}[t^X] = (1 - p + pt)^n.$$

Differentiating gives

$$\mathbb{E}[X] = G'_X(1) = np,$$

and a second derivative computation gives

$$\text{Var}(X) = np(1-p).$$

So the same binomial expansion that counts subsets also describes the sum of n independent Bernoulli trials.

§54.12 Asymptotic shape: from binomial coefficients to Gaussians

For large n , the row $\binom{n}{0}, \binom{n}{1}, \dots, \binom{n}{n}$ has a bell shape. Probabilistically, this is the central limit theorem applied to

$$X = X_1 + \dots + X_n, \quad X_i \sim \text{Bernoulli}(p).$$

After centering and scaling,

$$\frac{X - np}{\sqrt{np(1-p)}}$$

approaches the standard normal distribution.

For $p = 1/2$, the largest coefficient is near $k = n/2$. Stirling's formula gives the classical estimate

$$\binom{n}{\lfloor n/2 \rfloor} \sim \frac{2^n}{\sqrt{\pi n/2}}.$$

Thus a finite algebraic row of Pascal's triangle already contains the first shadow of probability limit theory.

§54.13 Generating functions and labelled structures

For labelled combinatorial structures, the exponential generating function is

$$A(x) = \sum_{n \geq 0} a_n \frac{x^n}{n!}.$$

The factor $n!$ compensates for labels. Products of exponential generating functions correspond to splitting a labelled set into parts. The exponential formula says that if a structure is a set of connected components, then

$$A(x) = \exp(C(x)),$$

where $C(x)$ counts connected components.

This is the advanced version of the elementary rule "split the set into independent choices and multiply." The multiplication principle becomes a structural statement about labelled decomposition.

§54.14 Combinatorial species and relabelling

A combinatorial species is a functor from finite sets and bijections to sets. For each finite label set U , a species F assigns a set $F[U]$ of structures on U . A bijection $U \rightarrow V$ transports structures from $F[U]$ to $F[V]$.

The species of subsets sends U to $\mathcal{P}(U)$. Choosing a subset is functorial under relabelling: if labels are renamed, chosen elements are renamed with them. The binomial theorem is the decategorified counting shadow of this functorial behavior.

§54.15 Binomial algebra from sets to species

The binomial theorem is not merely an expansion formula. It is a meeting point of algebra, subset counting, probability, finite differences, generating functions, and functorial relabelling. The elementary exercises train the same habit in many forms: translate a formula into choices, translate a set identity into element logic, and translate a sum into a structural decomposition.

Combinatorics: Pigeonhole Principle, Subsets, Binomial Sums, and Helly-Type Ideas

N-20230118

§55.1 Pigeonhole principle

The note begins with the pigeonhole principle. If $n + 1$ objects are put into n boxes, then some box contains at least two objects. More generally, if N objects are placed into k boxes, then some box contains at least

$$\left\lceil \frac{N}{k} \right\rceil$$

objects.

The handwritten examples use this principle to force coincidences: same remainder, same subset size, same color, or same interval. The key is to decide what the boxes are. Often the boxes are congruence classes modulo m or possible values of a sum.

§55.2 Subsets and binomial coefficients

For $[n] = \{1, 2, \dots, n\}$, the number of k -element subsets is

$$\binom{n}{k}.$$

The note uses the identity

$$\sum_{k=0}^n \binom{n}{k} = 2^n$$

to count all subsets. It also uses

$$\sum_{k=0}^n k \binom{n}{k} = n2^{n-1}.$$

This identity counts pairs (A, a) where $A \subseteq [n]$ and $a \in A$. Choose a first in n ways, then choose the rest of A freely in 2^{n-1} ways.

§55.3 Binomial identities by double counting

The handwritten page contains identities of the type

$$\sum_j \binom{a}{j} \binom{b}{n-j} = \binom{a+b}{n}.$$

This is Vandermonde's identity. The combinatorial proof is to choose n objects from two groups of sizes a and b . If j objects are chosen from the first group, then $n - j$ are chosen from the second.

The note marks this as a better proof than manipulating factorials, because it explains why the identity is true.

§55.4 A subset-sum style example

One example lists several subsets A_i of a finite set and asks for a conclusion about intersections or unions. The general method is to encode each element by the list of sets containing it. If there are too many subsets or too many incidences, two elements must share a pattern, or some intersection must be nonempty.

This is a typical pigeonhole principle application: the objects are elements or subsets, and the boxes are possible membership patterns.

§55.5 Generating functions

The page also uses generating functions for binomial sums. The coefficient of x^r in

$$(1+x)^n$$

is $\binom{n}{r}$. The coefficient of x^r in

$$(1+x)^a(1+x)^b = (1+x)^{a+b}$$

gives Vandermonde's identity.

This is the algebraic version of double counting. The combinatorial story and the generating-function coefficient extraction are two descriptions of the same structure.

§55.6 Helly-type remark

The last page mentions Helly's theorem. In the plane, Helly's theorem says that for a finite family of convex sets, if every three of them have nonempty intersection, then the whole family has nonempty intersection. The note does not prove the full theorem, but it points toward the idea that local intersection data can force global intersection.

This belongs naturally after the pigeonhole and subset examples: it is another local-to-global principle, now in convex geometry rather than finite counting.

§55.7 The next layer

The note is about forcing structure from too many objects and too few possibilities. The pigeonhole principle forces repetition. Binomial coefficients count choices. Double counting proves identities by counting the same set twice. Generating functions package these identities algebraically. Helly-type theorems show that similar local-to-global principles also appear in geometry.

§55.8 Pigeonhole as averaging

The pigeonhole principle is the statement that an average forces at least one large value. If N objects occupy r boxes, some box contains at least $\lceil N/r \rceil$ objects.

§55.9 Generating functions and subset sums

Subset-counting problems are naturally encoded by

$$\prod_{i=1}^n (1+x^{a_i}).$$

The coefficient of x^m counts subsets with sum m . This turns subset-sum counting into coefficient extraction.

§55.10 Helly-type thinking

Helly's theorem says that for convex sets in $\mathbb{R}^d$, global intersection can be tested by subfamilies of size $d + 1$. This is the geometric analogue of reducing a global condition to small local tests.

§55.11 Finite compactness principles in combinatorics

Pigeonhole, binomial sums, subset sums, generating functions, and Helly-type ideas are all compactness principles for finite configurations: too much local data forces a global structure.

§55.12 Pigeonhole as compactness principle

The finite pigeonhole principle says that too many objects placed into too few boxes force repetition. A compactness version says that if infinitely many points lie in a compact metric space, then some subsequence converges. In $[0, 1]$, this is Bolzano–Weierstrass.

The proof is a repeated pigeonhole argument: divide the interval into finitely many pieces, choose one containing infinitely many points, subdivide again, and continue. The nested intervals shrink to a point.

Thus compactness is an infinite pigeonhole principle. This is why elementary pigeonhole arguments often foreshadow topology and analysis.

§55.13 Ramsey theory

Ramsey theory says that sufficiently large structures contain ordered substructures. The simplest theorem is: for every r, s , there exists $R(r, s)$ such that every red-blue coloring of the edges of $K_{R(r, s)}$ contains either a red K_r or a blue K_s .

The pigeonhole principle is the smallest Ramsey theorem. It guarantees equality of colors or boxes; Ramsey theory guarantees large monochromatic patterns.

This connects elementary “two objects must share a property” problems with a research field studying unavoidable order inside large disorder.

§55.14 Helly theorem

Theorem 55.14.1 (Helly's theorem)

For a finite family of convex subsets of $\mathbb{R}^d$, if every $d + 1$ of them have nonempty intersection, then the whole family has nonempty intersection.

Helly's theorem in $\mathbb{R}^d$ says that for a finite family of convex sets, if every $d + 1$ of them have nonempty intersection, then the whole family has nonempty intersection.

For $d = 1$, this says that if intervals pairwise intersect, then all intervals intersect. That case is a strong one-dimensional pigeonhole-like fact. In higher dimensions, convexity replaces linear order.

The original Helly-type ideas are therefore not just clever geometry. They are the finite-dimensional form of a general intersection principle.

§55.15 Fractional and topological Helly

Fractional Helly theorems weaken the condition: if a positive fraction of all small subfamilies intersect, then a large subfamily has common intersection. Topological Helly theorems replace convexity by conditions on homology or contractibility of intersections.

This is a research-level continuation of the elementary principle: local intersection data can force global intersection. The advanced theory asks exactly how much local data is enough and what kind of geometry can replace convexity.

§55.16 Local-to-global intersection principles

Helly's theorem is powerful because it turns local intersection data into global intersection. In $\mathbb{R}^d$, checking every $d + 1$ members of a finite convex family suffices. The number $d + 1$ is not arbitrary: it is the size at which affine dependence begins, by Radon's theorem.

Radon's theorem says that any $d + 2$ points in $\mathbb{R}^d$ can be partitioned into two sets whose convex hulls intersect. Helly's theorem can be proved from Radon's theorem. This links elementary convex geometry to the combinatorics of affine dependence.

§55.17 Ramsey as structured inevitability

Pigeonhole says repetition is inevitable. Ramsey says organized substructure is inevitable. The shift from repetition to substructure is the main conceptual jump. In graph terms, edge colorings of large complete graphs must contain large monochromatic complete subgraphs.

This point should be connected back to the original exercises: whenever an argument says “among many objects, two must share a property,” one can ask the Ramsey-level question, “among sufficiently many objects, must a whole structured subset share the property?”

Combinatorics: Identities, Generating Functions, Catalan Numbers, and More

N-20240323

§56.1 Binomial identities

The note opens with several binomial identities and, more importantly, their combinatorial meanings. The basic identity

$$\sum_{i=0}^n \binom{n}{i} = 2^n$$

counts all subsets of an n -element set: choosing a subset means choosing its size i and then choosing i elements.

Pascal's identity

$$\binom{n}{i} + \binom{n}{i+1} = \binom{n+1}{i+1}$$

can be proved by fixing a distinguished element in an $(n+1)$ -element set. An $(i+1)$ -subset either contains the distinguished element or it does not. The two cases give the two terms on the left.

Vandermonde's identity

$$\sum_i \binom{m}{i} \binom{n}{r-i} = \binom{m+n}{r}$$

counts choosing r balls from two boxes with m and n balls. If i balls are chosen from the first box, then $r-i$ are chosen from the second.

Another useful identity is

$$\binom{n}{r} \binom{n-r}{k-r} = \binom{n}{k} \binom{k}{r}.$$

Both sides count a pair (R, K) with $R \subseteq K \subseteq [n]$, $|R| = r$, and $|K| = k$.

§56.2 Good subsets

A nonempty subset $A \subseteq \{1, 2, \dots, n\}$ is called good if

$$|A| \leq \min_{x \in A} x.$$

Let a_n be the number of good subsets. The problem asks one to prove

$$a_{n+2} = a_{n+1} + a_n + 1.$$

A clean proof splits good subsets of $[n+2]$ into cases.

First, if $n+2 \notin A$, then A is just a good subset of $[n+1]$, giving a_{n+1} possibilities.

Second, if $A = \{n + 2\}$, this gives one extra subset.

Third, suppose $n + 2 \in A$ and $|A| \geq 2$. Remove $n + 2$ and subtract 1 from every remaining element. If

$$A' = A \setminus \{n + 2\},$$

then the condition becomes

$$|A'| + 1 \leq \min A'.$$

After shifting $B = \{x - 1 : x \in A'\}$, this is exactly

$$|B| \leq \min B,$$

so B is a good subset of $[n]$. This gives a_n possibilities. Hence

$$a_{n+2} = a_{n+1} + a_n + 1.$$

The handwritten page also gives a summation proof by fixing $|A| = k$. If $|A| = k$, then every element must be at least k , so one chooses the remaining data from a tail interval. The recurrence then follows by repeated Pascal identities.

§56.3 Generating functions

For a sequence $(a_n)_{n \geq 0}$, its ordinary generating function is

$$f(x) = \sum_{n \geq 0} a_n x^n.$$

The note emphasizes a key principle: coefficients can encode counting problems.

For example, the number of nonnegative integer solutions of

$$x_1 + x_2 + \cdots + x_k = n$$

is the coefficient of x^n in

$$(1 + x + x^2 + \cdots)^k = \frac{1}{(1 - x)^k}.$$

Using the binomial expansion,

$$\frac{1}{(1 - x)^k} = \sum_{n \geq 0} \binom{n + k - 1}{k - 1} x^n,$$

so the number of solutions is

$$\binom{n + k - 1}{k - 1}.$$

The notes also use generating functions to separate parity or color constraints. A typical trick is to multiply several series, each recording the allowed choices for one box or color, and then read off the coefficient of the target power.

§56.4 A probability example with repeated trials

One example considers repeated trials and asks for a probability after n operations. The handwritten solution introduces a generating function and then rewrites nested binomial sums using exponentials. The important method is not the final closed form alone but the transformation:

$$\text{complicated convolution} \longrightarrow \text{product/composition of generating functions}.$$

This is a standard way to turn repeated independent operations into algebra. Binomial coefficients record choices; exponential series record independent counts; coefficient extraction recovers the desired probability.

§56.5 Catalan numbers

The note then turns to Catalan numbers. One model is the number of ways to triangulate a convex $(n + 2)$ -gon:

$$C_n = \frac{1}{n+1} \binom{2n}{n}.$$

For example, a hexagon has 14 triangulations.

Let h_n denote the number of triangulations of an $(n + 2)$ -gon. Choose the triangle incident with a fixed side. If its third vertex cuts the polygon into two smaller polygons, the numbers of triangulations multiply. Therefore

$$h_n = \sum_{i=0}^{n-1} h_i h_{n-1-i}, \quad h_0 = 1.$$

If

$$H(x) = \sum_{n \geq 0} h_n x^n,$$

then the recurrence gives

$$H(x) = 1 + xH(x)^2.$$

Solving the quadratic and choosing the branch with $H(0) = 1$ gives

$$H(x) = \frac{1 - \sqrt{1 - 4x}}{2x}, \quad h_n = C_n.$$

This derivation explains why Catalan numbers appear in triangulations, parentheses, Dyck paths, and many noncrossing structures: all of them are recursively split into a left and a right part.

§56.6 Convex sets and convex hulls

The geometry part defines a convex set S by the condition that for any $x, y \in S$ and any $t \in [0, 1]$,

$$tx + (1 - t)y \in S.$$

Equivalently, S contains the line segment between any two of its points. The note contrasts convex and non-convex examples.

For finitely many points $x_1, \dots, x_k$, their convex hull is

$$\left\{ \sum_{i=1}^k t_i x_i : t_i \geq 0, \sum_{i=1}^k t_i = 1 \right\}.$$

The handwritten proof uses induction on k : the case $k = 2$ is the definition of a segment; the induction step writes the convex combination as a combination of one point and a convex combination of the remaining points.

§56.7 Separating convex sets

The note sketches a separation idea. For two disjoint convex sets, one tries to find a point pair realizing minimal distance and then uses the perpendicular bisector of the segment between them as a separating line. In symbols, if $u \in S_1$ and $v \in S_2$ are closest

points, then the hyperplane perpendicular to $v - u$ near the midpoint separates the two sets.

In finite-dimensional Euclidean geometry, this is the intuitive core of the separating hyperplane theorem. The handwritten diagrams show projections onto a line and inequalities involving squared distances. The algebra expands

$$\|u + t(a - u) - v\|^2$$

for small t and extracts a sign condition by letting $t \rightarrow 0$.

§56.8 Planar graphs and Euler-type counting

The later pages move toward graph theory and planar configurations. The recurring idea is that a planar drawing imposes counting restrictions. Vertices, edges, and faces cannot vary independently; they are constrained by Euler-type formulas.

This connects naturally with the previous convex-geometry material: triangulations, planar graphs, polygon decompositions, and convex hulls are different languages for controlling incidence data.

§56.9 Graphs and adjacency matrices

A graph G consists of a vertex set $V = \{v_1, \dots, v_n\}$ and an edge relation. For a simple graph, the adjacency matrix $A(G) = (a_{ij})$ is defined by

$$a_{ij} = \begin{cases} 1, & \text{if there is an edge from } v_i \text{ to } v_j, \\ 0, & \text{otherwise.} \end{cases}$$

For a multigraph, a_{ij} records the number of edges from v_i to v_j .

The source quotes the standard theorem: the (i, j) -entry of $A(G)^\ell$ equals the number of walks of length ℓ from v_i to v_j . The proof is matrix multiplication. For $\ell = 2$,

$$(A^2)_{ij} = \sum_k a_{ik}a_{kj},$$

which counts all intermediate vertices v_k . The general case follows by induction.

§56.10 The deeper habit

This note collects several ways of counting: direct binomial identities, recurrence, generating functions, Catalan decomposition, convex geometry, and graph adjacency matrices. The common habit is to replace a raw counting problem by a structure-preserving encoding: subsets become binomial coefficients, recursions become generating functions, triangulations become Catalan recurrences, convex hulls become linear combinations, and walks become matrix powers.

§56.11 Generating functions as algebraic counting

A generating function packages a sequence into a formal series. Products encode decomposition of objects, and coefficient extraction reads off counts:

$$[x^n]A(x)B(x) = \sum_{k=0}^n a_k b_{n-k}.$$

§56.12 Catalan structures

Catalan numbers count many recursively decomposable objects:

$$C_n = \frac{1}{n+1} \binom{2n}{n}.$$

Their generating function satisfies

$$C(x) = 1 + xC(x)^2.$$

The quadratic equation records the root-subtree decomposition.

§56.13 Convex hulls and separation

The separating hyperplane theorem says disjoint convex sets can often be separated by a linear functional. This is the geometric dual of optimization: a linear functional witnesses non-intersection.

§56.14 Generating functions, Catalan recursion, and spectra

Binomial identities, generating functions, Catalan numbers, probability trials, and convex separation all share one habit: encode the structure so that algebraic or linear operations reveal the count or obstruction.

§56.15 Catalan numbers as recursive geometry

Catalan numbers count many objects: correctly matched parentheses, triangulations of polygons, rooted binary trees, Dyck paths, and noncrossing partitions. A standard formula is

$$C_n = \frac{1}{n+1} \binom{2n}{n}.$$

The recursion

$$C_{n+1} = \sum_{i=0}^n C_i C_{n-i}$$

comes from choosing the first return of a Dyck path or the root split of a binary tree.

The generating function $C(x) = \sum_{n \geq 0} C_n x^n$ satisfies

$$C(x) = 1 + xC(x)^2.$$

Solving gives

$$C(x) = \frac{1 - \sqrt{1 - 4x}}{2x}.$$

Thus an elementary recursion becomes an algebraic generating function.

§56.16 Lattice paths and reflection principle

Many Catalan formulas are proved by the reflection principle. The number of paths from $(0, 0)$ to (n, n) that never go above the diagonal is

$$\binom{2n}{n} - \binom{2n}{n+1} = \frac{1}{n+1} \binom{2n}{n}.$$

The subtracted term counts bad paths by reflecting the part up to the first diagonal crossing.

This is a model example of an involutive or bijective proof. Rather than evaluating a sum, one constructs a symmetry pairing bad objects with easier objects.

§56.17 Graph walks and adjacency matrices

If G is a finite graph with adjacency matrix A , then

$$(A^k)_{ij}$$

counts walks of length k from vertex i to vertex j . This turns walk counting into matrix powers.

If A is diagonalizable with eigenvalues $\lambda_1, \dots, \lambda_n$, the growth of walks is controlled by the largest eigenvalues. In a regular connected graph, the largest eigenvalue is the degree, and the second-largest eigenvalue controls expansion and mixing.

Thus an elementary graph-walk count leads directly to spectral graph theory.

§56.18 Random walks and Markov chains

Normalize the adjacency matrix of a d -regular graph by

$$P = \frac{1}{d} A.$$

Then P is the transition matrix of a random walk. The distribution after k steps is

$$\pi_k = \pi_0 P^k.$$

Eigenvalues of P determine convergence to stationarity. If the graph is connected and non-bipartite, the walk converges to the uniform distribution.

The same object therefore has two readings: A^k counts deterministic walks, while P^k describes probabilities. The note's graph-walk calculations are the finite beginning of Markov-chain theory.

§56.19 Catalan recursion and parenthesization

Polygon triangulations and parenthesizations are the same combinatorial structure. A triangulation of an $(n+2)$ -gon corresponds to a way of multiplying $n+1$ factors with binary parentheses. Choosing the triangle incident to a fixed edge splits the polygon into two smaller polygons, producing the Catalan recursion.

This is the combinatorics behind the associahedron, a polytope whose vertices are parenthesizations and whose edges are single reassociations:

$$(ab)c \leftrightarrow a(bc).$$

The elementary Catalan count therefore touches the geometry of moduli spaces of brackets.

§56.20 Adjacency spectra and expansion

The theorem that $(A^k)_{ij}$ counts walks becomes much stronger when combined with eigenvalues. If a graph is d -regular with adjacency eigenvalues

$$d = \lambda_1 \geq |\lambda_2| \geq \cdots,$$

then $|\lambda_2|$ controls how quickly random walks mix and how well the graph expands.

Expander graphs are sparse graphs that behave like highly connected graphs. They are central in theoretical computer science, coding theory, and combinatorics. The elementary walk-counting matrix is the first tool for studying them.

Combinatorial Geometry, Euler Formula, Platonic Solids, Degree, and Quotients

N-20240330

§57.1 Good triangles

The first problem concerns n points in the plane with no three collinear. A triangle formed by three of these points is called good if its area is maximal among all triangles formed by the point set. The goal is to show that there cannot be more than n good triangles.

The note begins with small cases. For $n = 4$, a convex quadrilateral gives four good-looking candidate triangles by deleting one vertex. For $n = 5$, a pentagon suggests five candidates. These pictures motivate the expected upper bound.

§57.2 Why vertices must lie on the convex hull

If a triangle uses a point strictly inside the convex hull, it cannot be maximal. Moving that interior vertex outward toward a hull vertex while keeping the opposite side fixed increases area. Therefore every good triangle has all three vertices on the boundary of the convex hull.

This is the first reduction: the problem becomes cyclic. We may list the hull vertices in clockwise order and encode a triangle by the three arc lengths between its consecutive vertices.

§57.3 Cyclic arc encoding

Let a good triangle have vertices A, B, C on the convex hull, ordered cyclically. Let

$$a, b, c$$

be the numbers of boundary steps along the three arcs from A to B , from B to C , and from C back to A . Then

$$a + b + c = m,$$

where m is the number of hull vertices.

The handwritten argument tries to order all good triangles by their arc data. If one arc is too large compared with another, then shifting a vertex along the circle or convex boundary can produce a triangle of larger area, contradicting maximality. The intended conclusion is that good triangles are controlled by a cyclic balancing condition, so there can be at most one good triangle associated with each starting vertex. Hence the number of good triangles is at most $m \leq n$.

§57.4 The geometric idea behind the proof

The core geometric principle is this: with one side fixed, area is proportional to the distance from the third vertex to the line containing the side. Therefore, if a vertex is not as far as possible from the opposite side, the triangle cannot be maximal.

The diagrams in the note use this repeatedly. If a candidate triangle sits inside a larger triangle or if one can replace a vertex by a farther hull vertex on the same side of a line, the area increases. These local replacements are the geometric engine of the proof.

§57.5 Euler formula for convex polyhedra

The next part reviews Euler's formula for a convex polyhedron:

$$V - E + F = 2.$$

The source gives an induction proof by contracting an edge. Suppose a polyhedron has $v' = k + 1$ vertices. Contract one edge, merging its two endpoints into one vertex. Then

$$v' - 1 = k, \quad e' - 1 = e, \quad f' = f.$$

If the contracted polyhedron satisfies $v - e + f = 2$, then the original one does too:

$$(v' - 1) - (e' - 1) + f' = 2 \implies v' - e' + f' = 2.$$

The proof is schematic rather than fully formal, but it captures the invariant idea: a safe local contraction changes V and E together and keeps $V - E + F$ unchanged.

§57.6 Regular convex polyhedra

The note then considers regular convex polyhedra. Suppose every face is a regular a -gon and exactly b faces meet at each vertex. Counting incidences gives

$$aF = 2E, \quad bV = 2E.$$

Substitute into Euler's formula:

$$\frac{2E}{b} - E + \frac{2E}{a} = 2.$$

Dividing by $2E$ gives

$$\frac{1}{a} + \frac{1}{b} > \frac{1}{2}.$$

Since $a, b \geq 3$, the only possibilities are

$$(a, b) = (3, 3), (3, 4), (4, 3), (3, 5), (5, 3).$$

These correspond to the tetrahedron, octahedron, cube, icosahedron, and dodecahedron.

The handwritten page lists sample triples (F, E, V) such as

$$(4, 6, 4), \quad (8, 12, 6), \quad (6, 12, 8), \quad (20, 30, 12), \quad (12, 30, 20).$$

§57.7 Euler characteristic

The note marks the extension to Euler characteristic:

$$\chi = V - E + F.$$

For a convex polyhedron, $\chi = 2$. More generally, χ is a topological invariant: under suitable deformations or subdivisions it remains unchanged. This explains why Euler's formula is not merely a counting trick but a topological statement.

§57.8 Degree of a map

The last page quotes the degree of a map between spheres. For any $n > 0$ and a map

$$f : S^n \rightarrow S^n,$$

the induced map on top homology is multiplication by some integer d :

$$f_* : H_n(S^n) \rightarrow H_n(S^n), \quad f_*(1) = d.$$

This integer is the degree of f .

The key property is multiplicativity:

$$\deg(g \circ f) = \deg(g) \deg(f).$$

The note's margin comment asks why this should be true. The answer is that induced maps on homology compose functorially:

$$(g \circ f)_* = g_* \circ f_*.$$

If f_* multiplies by d_f and g_* multiplies by d_g , their composition multiplies by $d_g d_f$.

§57.9 Homology group and quotient group reminders

The page also records the quotient definition

$$H_n(X) = Z_n(X)/B_n(X),$$

meaning that homology is cycles modulo boundaries. It then recalls quotient groups. If $N \trianglelefteq G$, the quotient group G/N consists of cosets, and multiplication is defined by

$$(g_1 N)(g_2 N) = (g_1 g_2) N.$$

Normality is exactly the condition needed for this product to be well-defined.

The handwritten note draws a circle and a cylinder-like picture, suggesting the intuition that quotienting identifies points and changes the topology. Algebraic quotient and topological quotient are different constructions, but both create new objects by identifying equivalent elements.

§57.10 A bridge outward

This note moves from extremal planar geometry to polyhedral topology and then to homological algebra. The bridge is invariant thinking. A maximal-area triangle is controlled by convex hull geometry. A polyhedron is controlled by $V - E + F$. A map of spheres is controlled by degree. A quotient group is controlled by which elements are declared equivalent. In each case, the problem becomes manageable after finding the invariant that survives the allowed transformations.

§57.11 Euler characteristic as invariant

For a convex polyhedron,

$$V - E + F = 2.$$

This number is the Euler characteristic of the sphere. More generally, triangulated surfaces have Euler characteristic independent of the triangulation.

This is the topological invariant behind the combinatorial formula.

§57.12 Degree counting on graphs

The identity

$$\sum_v \deg(v) = 2E$$

counts incidences between vertices and edges. In polyhedron problems, this turns local face/vertex constraints into global inequalities.

The classification of Platonic solids follows by combining Euler's formula with degree and face-size bounds.

§57.13 Quotients and surface construction

Identifying edges of polygons constructs surfaces. The torus, sphere, projective plane, and Klein bottle can all be represented by edge-gluing diagrams. Euler characteristic tracks how cells survive the quotient.

§57.14 Euler bookkeeping from graphs to surfaces

Combinatorial geometry becomes topology when one asks which quantities survive deformation or subdivision. Euler characteristic, incidence counting, convex hull structure, and quotient surfaces are the advanced frame behind the finite diagrams.

§57.15 Euler characteristic as topological bookkeeping

Euler's formula for a connected planar graph is

$$V - E + F = 2.$$

The quantity

$$\chi = V - E + F$$

is the Euler characteristic. For a triangulated surface, the same alternating count generalizes to

$$\chi = \sum_{k \geq 0} (-1)^k f_k,$$

where f_k is the number of k -simplices.

The important point is invariance: if the triangulation is refined without changing the underlying surface, χ does not change. Thus Euler's formula is not only a planar counting trick; it is a topological invariant.

§57.16 Euler characteristic and homology

For a finite simplicial complex, homology groups H_k measure k -dimensional holes. Their dimensions

$$b_k = \dim H_k$$

are Betti numbers. The Euler characteristic also equals

$$\chi = \sum_{k \geq 0} (-1)^k b_k.$$

For a sphere, $b_0 = 1$, $b_2 = 1$, and all other Betti numbers vanish, so $\chi = 2$. For a torus, $b_0 = 1$, $b_1 = 2$, $b_2 = 1$, so $\chi = 0$.

This explains why the same number appears in graph embeddings, triangulations, and surface classification.

§57.17 Degree theory

The degree of a map is an integer measuring how many times the domain wraps around the target, counted with orientation. For a map $f : S^1 \rightarrow S^1$, the degree is the winding number. If

$$f(e^{it}) = e^{int},$$

then $\deg f = n$.

In higher dimensions, degree theory is used to prove existence theorems. A nonzero degree often forces a solution because the map cannot be deformed away from a target point without crossing it.

Thus the elementary “degree” and quotient-space discussions in the note connect to a robust topological invariant.

§57.18 Quotient spaces and identifications

A quotient space is obtained by declaring some points equivalent and treating each equivalence class as one point. For example, gluing the two ends of an interval gives a circle:

$$[0, 1]/(0 \sim 1) \cong S^1.$$

Gluing edges of a square can produce a cylinder, Möbius strip, torus, or projective plane depending on the identifications.

The elementary quotient examples are therefore not mere pictures. They are the language by which topology builds new spaces from old ones.

§57.19 Euler characteristic as alternating rank

For a chain complex

$$0 \rightarrow C_n \xrightarrow{\partial_n} C_{n-1} \rightarrow \cdots \rightarrow C_0 \rightarrow 0,$$

with finite-dimensional chain groups, the Euler characteristic satisfies

$$\sum_k (-1)^k \dim C_k = \sum_k (-1)^k \dim H_k.$$

This identity is the algebraic reason a count of vertices, edges, and faces equals a count of homological holes.

For a polyhedral sphere, the chain groups count cells and the homology counts one connected component and one enclosed void. Hence $\chi = 2$.

§57.20 Degree and fixed-point consequences

Degree theory is not just a way to label maps. It proves existence. If a continuous map $f : S^n \rightarrow S^n$ has nonzero degree, it is surjective. Otherwise a missing point would allow the map to contract through $S^n \setminus \{p\}$, whose top homology vanishes, contradicting nonzero degree.

This is the topological form of a familiar elementary idea: a nonzero invariant prevents deformation to a trivial configuration. Degree is one such invariant.

Self-Avoiding Walks I: Connective Constants, Bridges, Polygons, and Critical Exponents

N-20251126

§58.1 What the model is trying to measure

A self-avoiding walk is the simplest discrete model of a polymer chain with excluded volume. The walk is allowed to move from one lattice point to a neighbouring lattice point, but it is not allowed to visit the same vertex twice. Thus the model keeps the most basic feature of a real chain: two monomers cannot occupy the same place.

Let $\mathbb{Z}^d$ be the hypercubic lattice. An n -step walk is a sequence

$$\omega = (\omega_0, \omega_1, \dots, \omega_n),$$

where $\omega_i \in \mathbb{Z}^d$, $\omega_0 = 0$, and

$$|\omega_{i+1} - \omega_i| = 1.$$

It is self-avoiding if

$$\omega_i \neq \omega_j \quad (i \neq j).$$

Write c_n for the number of n -step self-avoiding walks starting at the origin.

The ordinary simple random walk has exactly $(2d)^n$ possible n -step paths. For self-avoiding walks, there is no such closed formula, because each new step depends on the entire past of the path. This is the first reason the subject is difficult: the model is local in its rule but non-local in its memory.

Remark 58.1.1 — The phrase “self-avoiding” sounds like a small constraint, but it changes the problem qualitatively. The number of walks is no longer obtained by multiplying independent choices, and the endpoint distribution is no longer governed by the usual central limit theorem in low dimensions.

§58.2 Submultiplicativity and the connective constant

The first invariant is the exponential growth rate of c_n . Concatenation gives the basic inequality

$$c_{m+n} \leq c_m c_n.$$

Indeed, if an $(m+n)$ -step self-avoiding walk is cut after m steps, then the first part is an m -step self-avoiding walk, and the translated second part is an n -step self-avoiding walk. Not every pair can be concatenated without intersection, so this gives an inequality rather than an equality.

By Fekete’s lemma, the limit

$$\mu = \lim_{n \rightarrow \infty} c_n^{1/n}$$

exists. This number is called the connective constant. Equivalently,

$$\log \mu = \lim_{n \rightarrow \infty} \frac{1}{n} \log c_n.$$

The rough meaning is

$$c_n \approx \mu^n \times \text{a subexponential correction.}$$

Definition 58.2.1. The connective constant μ is the exponential growth rate of self-avoiding walks on the lattice. It depends on the lattice, not only on the dimension.

The reciprocal

$$z_c = \mu^{-1}$$

is the critical point of the generating function

$$\chi(z) = \sum_{n \geq 0} c_n z^n.$$

For $z < z_c$, the series converges. For $z > z_c$, it diverges. Thus the connective constant is both a combinatorial growth constant and the location of a statistical-mechanical phase transition.

§58.3 Bridges and why they help

Counting all self-avoiding walks directly is hard because a walk can wander in all directions. A bridge is a self-avoiding walk with a preferred direction. For example, in the first coordinate direction, a bridge is a walk whose first coordinate starts at its minimum and ends at its maximum:

$$\omega_0^{(1)} < \omega_i^{(1)} \leq \omega_n^{(1)} \quad (1 \leq i \leq n).$$

The precise convention varies slightly, but the guiding idea is stable: a bridge progresses from left to right without going below its starting level in that coordinate.

The following picture is the right mental picture for the Hammersley–Welsh argument. A half-space walk is not yet a bridge, but it can be chopped into alternating bridge-like pieces. The dotted reflected part indicates the trick: reflect suitable tails of the walk so that the pieces line up into one bridge.

Let b_n denote the number of n -step bridges. Bridges are useful because they concatenate better than arbitrary walks. If one bridge is followed by another translated bridge, the result is still a bridge. Hence

$$b_{m+n} \geq b_m b_n.$$

This supermultiplicativity implies that the bridge growth rate exists:

$$\lim_{n \rightarrow \infty} b_n^{1/n} = \sup_n b_n^{1/n}.$$

A central fact is that bridges have the same exponential growth rate as all self-avoiding walks:

$$\lim_{n \rightarrow \infty} b_n^{1/n} = \mu.$$

The Hammersley–Welsh method gives a much sharper comparison. Its philosophy is to decompose a general self-avoiding walk into bridge-like pieces by reflecting parts of




Figure 58.1: a half-space walk decomposed into bridges and reflected into a single bridge.

The polygon viewpoint is important because it connects self-avoiding walks with loop models and planar statistical mechanics. In two dimensions, loops are not just a counting device; they interact with conformal geometry.

The bridge picture has an analogue for polygons. The following construction turns two bridges into a longer self-avoiding walk closely related to a polygon-counting lower bound. The useful point is again not exact enumeration. It is a controlled way to build large objects out of smaller ones while keeping self-avoidance visible.




Figure 58.3: bridge data used in a polygon lower-bound construction.

§58.5 Critical exponents

The connective constant describes only the exponential part of c_n . The more delicate question is the subexponential correction. The expected form is

$$c_n \sim A\mu^n n^{\gamma-1},$$

where γ is a critical exponent. Another exponent ν describes the typical size of an n -step walk:

$$\mathbb{E}(|\omega_n|^2) \approx Dn^{2\nu}.$$

The exponent ν measures how swollen the polymer is. For simple random walk one has $\nu = 1/2$. Self-avoidance pushes the walk away from itself, so in low dimensions ν is expected to be larger.

There is also a two-point function

$$G_z(x) = \sum_{\omega: 0 \rightarrow x} z^{|\omega|},$$

where the sum is over self-avoiding walks from 0 to x . Summing over endpoints gives the susceptibility

$$\chi(z) = \sum_{x \in \mathbb{Z}^d} G_z(x) = \sum_{n \geq 0} c_n z^n.$$

Near z_c , one expects

$$\chi(z) \approx (z_c - z)^{-\gamma}.$$

So the same exponent γ can be read from either the coefficient asymptotics of c_n or the singularity of the susceptibility.

§58.6 How dimension changes the story

The dimension d controls how strongly a walk intersects itself. This is why the self-avoiding walk behaves differently in different dimensions.

In dimension 1, the model is trivial: after the first step, the walk cannot turn around without revisiting a vertex. Thus there are only two n -step self-avoiding walks for $n \geq 1$.

In dimension 2, self-avoidance is strong. The expected scaling limit is related to Schramm–Loewner evolution with parameter $8/3$. The predicted exponents are

$$\nu = \frac{3}{4}, \quad \gamma = \frac{43}{32}.$$

These values are part of the conformal-invariance picture, but the full square-lattice scaling limit remains a major open problem.

In dimension 3, the model is physically central but mathematically difficult. The exponents are not expected to be mean-field, and no exact solution is known.

Dimension 4 is critical. The leading power laws are mean-field-like, but logarithmic corrections appear. This is the dimension where renormalisation group methods become essential.

In dimensions $d \geq 5$, the model has mean-field behaviour. The lace expansion proves, in suitable settings, that self-avoiding walks behave diffusively:

$$\nu = \frac{1}{2}, \quad \gamma = 1.$$

The walk is still self-avoiding, but high-dimensional space gives it enough room that self-intersections become perturbative rather than dominant.

§58.7 The broader lesson

The first layer of the theory is the following dictionary:

combinatorial object	analytic quantity
c_n	$\chi(z) = \sum c_n z^n$
μ	$z_c = 1/\mu$
bridges	controlled concatenation
polygons	closed-loop version of the model
$G_z(x)$	endpoint-resolved two-point function

The connective constant is the exponential growth rate. Critical exponents describe what remains after this exponential growth has been factored out. Bridges and polygons are not side topics: they are the combinatorial tools that make the basic growth theory work.

Self-Avoiding Walks II: The Hexagonal Lattice, Holomorphic Observables, and the Exact Connective Constant

N-20251127

§59.1 Why the hexagonal lattice is special

For most lattices, the connective constant of self-avoiding walks is not known exactly. The hexagonal lattice is exceptional. Duminil-Copin and Smirnov proved that its connective constant is

$$\mu_{\text{hex}} = \sqrt{2 + \sqrt{2}}.$$

Equivalently, the critical step weight is

$$x_c = \frac{1}{\mu_{\text{hex}}} = \frac{1}{\sqrt{2 + \sqrt{2}}}.$$

This result is not an enumeration trick in the elementary sense. It comes from a discrete complex-analytic observable. The proof is one of the cleanest examples of the modern two-dimensional philosophy: at the critical point, a lattice model may possess a discrete holomorphic structure that forces exact identities.

Remark 59.1.1 — The number $\sqrt{2 + \sqrt{2}}$ should not be memorised as a miracle constant. It is the value for which the parafermionic observable satisfies a local cancellation identity on the hexagonal lattice.

§59.2 A domain version of the counting problem

Instead of counting walks in the whole plane immediately, one works in a finite domain cut from the hexagonal lattice. Fix a boundary point a . For another boundary point or mid-edge z , consider self-avoiding walks γ starting at a and ending at z , staying inside the domain. Their weighted sum has the form

$$\sum_{\gamma: a \rightarrow z} x^{|\gamma|} \times \text{a complex phase depending on the winding of } \gamma.$$

The length weight $x^{|\gamma|}$ remembers the usual self-avoiding-walk generating function. The complex phase is the new ingredient.

The winding of a planar path records how much the direction of the path turns. On the hexagonal lattice the allowed turns are highly constrained, and this makes it possible to tune the phase so that local contributions cancel.

The first picture for this part is especially helpful. On the left, the domain is made of hexagonal-lattice cells and the boundary is described by mid-edges. On the right, the same picture explains winding: the path is not only counted by its length, but also by the total amount it turns.




Figure 59.1: a hexagonal-lattice domain and the winding of a self-avoiding walk.

§59.3 The parafermionic observable

The observable can be remembered schematically as

$$F(z) = \sum_{\gamma: a \rightarrow z} x^{|\gamma|} e^{-i\sigma W(\gamma)},$$

where $W(\gamma)$ is the total winding and σ is a spin parameter. The exact theorem uses the special choice

$$\sigma = \frac{5}{8}, \quad x = x_c = \frac{1}{\sqrt{2 + \sqrt{2}}}.$$

At this value, the observable satisfies a discrete Cauchy–Riemann type relation around each vertex.

The word “holomorphic” here should be read carefully. The observable is defined on a lattice, so it is not holomorphic in the classical analytic sense. Rather, it satisfies a local linear relation that becomes the Cauchy–Riemann equation in a scaling limit. This is why the observable is often called discrete holomorphic or preholomorphic.

Remark 59.3.1 — The observable is not merely a generating function. The phase remembers geometry. Without the winding phase, the local cancellation identity would not hold.

§59.4 From local cancellation to a global identity

The local identity is summed over all vertices in a finite domain. Interior terms cancel, leaving only boundary contributions. This is the discrete analogue of using Green’s theorem or Cauchy’s theorem: local holomorphicity turns an interior statement into a boundary statement.

The cancellation is proved by grouping local contributions. The picture below shows the two types of groupings: pairs of walks that visit all three adjacent mid-edges, and triplets obtained by extending or erasing the last step. At the special values of x and σ , each such group has total contribution zero.




Figure 59.2: pair and triplet cancellations in the proof of the local observable identity.

For a suitable trapezoidal domain on the hexagonal lattice, the resulting identity relates generating functions of walks ending on different parts of the boundary. Very schematically it has the form

Here is the actual domain picture used in the lecture notes. The boundary pieces are labelled separately because the final identity distinguishes walks according to which part of the boundary they hit.




Figure 59.3: the trapezoidal domain $S_{T,L}$ and its boundary pieces.

$$1 = A_T(x_c) + B_T(x_c) + E_T(x_c),$$

where the terms count walks from the marked boundary edge to various boundary arcs, with explicit positive coefficients after taking the correct real part.

The important structural fact is positivity. Once the boundary identity has positive

terms, it gives inequalities for ordinary generating functions. Letting the size of the domain tend to infinity then produces information about convergence and divergence at the critical value.

§59.5 The upper bound on the connective constant

To prove

$$\mu_{\text{hex}} \leq \sqrt{2 + \sqrt{2}},$$

one shows that the self-avoiding-walk generating function cannot converge beyond

$$x_c = \frac{1}{\sqrt{2 + \sqrt{2}}}.$$

Equivalently, if $x > x_c$, then the full generating function must diverge. The boundary identity supplies enough positive mass at x_c to force criticality at or before this value.

The proof is subtle because finite-domain identities do not immediately determine the infinite-volume connective constant. One has to pass from boundary generating functions to walks in the full lattice. The Hammersley–Welsh bridge technology from the previous note is part of the background that makes this passage possible.

§59.6 The lower bound on the connective constant

The other inequality

$$\mu_{\text{hex}} \geq \sqrt{2 + \sqrt{2}}$$

amounts to showing that the generating function still converges below x_c . In practice one proves that for $x < x_c$, the weighted sum of self-avoiding walks is finite. This uses the same domain identities together with limiting arguments and standard consequences of submultiplicativity.

Combining the two bounds gives the exact value

$$\mu_{\text{hex}} = \sqrt{2 + \sqrt{2}}.$$

Theorem 59.6.1 (Connective constant of the hexagonal lattice)

For self-avoiding walks on the hexagonal lattice,

$$\lim_{n \rightarrow \infty} c_n^{1/n} = \sqrt{2 + \sqrt{2}}.$$

§59.7 Relation with SLE

The same observable also points toward the expected scaling limit. The conjectural statement is that two-dimensional self-avoiding walk, after suitable rescaling, converges to $\text{SLE}_{8/3}$. This would imply many predicted critical exponents, such as

$$\nu = \frac{3}{4}, \quad \gamma = \frac{43}{32}.$$

The exact connective constant result is much more modest than a full scaling-limit theorem, but it belongs to the same conceptual world. The local observable is designed so that at criticality the lattice model begins to look conformally natural.

Remark 59.7.1 — The theorem gives the exact exponential growth constant on the hexagonal lattice. It does not by itself prove the $\text{SLE}_{8/3}$ scaling limit. The former is a rigorous theorem; the latter is the guiding two-dimensional conjectural picture.

§59.8 Loop models and the same mechanism

Self-avoiding walks can be viewed as the $n = 0$ member of the $O(n)$ loop model. In the loop model, configurations consist of nonintersecting loops, and each loop receives a weight n . Formally setting $n = 0$ kills configurations with closed loops and leaves a single open path, which is the self-avoiding walk.

This viewpoint explains why holomorphic observables appear. For special values of the loop weight and the step weight, the model has a parafermionic observable satisfying local relations. The hexagonal self-avoiding walk is the most famous case, but the mechanism belongs to a broader family of two-dimensional lattice models.

The phase diagram situates this in the $O(n)$ phase diagram. For this note, the picture should be read only as orientation: self-avoiding walk corresponds to the $n = 0$ edge of the loop-model story, while the exactly solvable point belongs to a broader two-dimensional critical curve.




Figure 59.4: phase diagram for $O(n)$ loop models.

§59.9 From example to method

The proof of the hexagonal connective constant has a clear conceptual shape:

choose a winding-weighted observable
 $\Downarrow$
 tune x and σ so local cancellations hold
 $\Downarrow$
 sum over a finite domain
 $\Downarrow$
 obtain positive boundary identities
 $\Downarrow$
 pass to the infinite lattice and identify μ .

The striking part is that an exact enumerative constant is obtained through a complex-analytic idea rather than through exact enumeration of c_n .

Self-Avoiding Walks III: Lace Expansion, Functional Integrals, and the Critical Dimension Four

N-20251128

§60.1 The high-dimensional philosophy

In low dimensions, self-avoidance dominates the geometry of the walk. In high dimensions, the walk has more room, and self-intersections become less frequent. The expected result is that the self-avoiding walk behaves like simple random walk at large scales:

$$|\omega_n| \text{ is typically of order } n^{1/2}.$$

This is called mean-field behaviour.

The rigorous tool that proves this in dimensions $d \geq 5$, for nearest-neighbour self-avoiding walk, is the lace expansion. It is a systematic inclusion–exclusion method that turns the global self-avoidance constraint into a perturbative correction to the simple random walk equation.

§60.2 The two-point function

The central analytic object is the two-point function

$$G_z(x) = \sum_{\omega: 0 \rightarrow x} z^{|\omega|},$$

where the sum is over self-avoiding walks from 0 to x . The susceptibility is

$$\chi(z) = \sum_{x \in \mathbb{Z}^d} G_z(x).$$

The critical point is $z_c = 1/\mu$.

For simple random walk, the two-point function satisfies a renewal equation. If D is the one-step distribution, then the generating function C_z satisfies

$$C_z = \delta_0 + zD * C_z.$$

Equivalently, in Fourier space,

$$\hat{C}_z(k) = \frac{1}{1 - z\hat{D}(k)}.$$

For self-avoiding walk, this exact renewal equation fails because the future path is constrained by the past. The lace expansion repairs the equation by adding correction terms.

§60.3 The lace expansion as corrected random walk

The output of the lace expansion can be remembered schematically as

$$G_z = \delta_0 + zD * G_z + \Pi_z * G_z,$$

where Π_z is the lace-expansion correction. This formula is not meant as a definition of Π_z ; it is the structural form to remember. The term $zD * G_z$ is what simple random walk would have. The term $\Pi_z * G_z$ records the accumulated effect of forbidding self-intersections.

The correction terms are represented by small diagrams. The following two examples should be read as diagrammatic shorthand rather than literal pictures of the walk. A line represents a two-point function, and extra chords represent constraints coming from self-intersection corrections.




Figure 60.1: typical lace-expansion diagrams.

In Fourier space the relation becomes morally

$$\hat{G}_z(k) = \frac{1}{1 - z\hat{D}(k) - \hat{\Pi}_z(k)}.$$

Thus the proof of mean-field behaviour becomes a problem of showing that Π_z is small and regular enough near criticality.

Remark 60.3.1 — The lace expansion is often intimidating because the combinatorics of laces is elaborate. Conceptually, however, its role is simple: it isolates the first random-walk term and packages all self-intersection corrections into a controlled perturbation.

§60.4 Bubble diagrams and diagrammatic estimates

The perturbative parameter is measured by diagrammatic sums. A basic example is the bubble diagram

The word “bubble” is literal in the diagrammatic notation. One sums over two independent-looking connections between the same endpoints. Finiteness of the corresponding sum is what makes the high-dimensional expansion behave like a small perturbation.




Figure 60.2: the bubble diagram.

$$B(z) = \sum_{x \in \mathbb{Z}^d} G_z(x)^2.$$

For simple random walk at criticality, this sum is finite exactly above the critical dimension 4. This is the first signal that dimension 4 is the borderline.

In dimensions $d > 4$, one can prove estimates of the following kind:

$$\sum_x G_{z_c}(x)^2 < \infty.$$

Such bounds imply that the lace-expansion correction behaves perturbatively. They lead to the mean-field critical exponents

$$\gamma = 1, \quad \nu = \frac{1}{2}.$$

More concretely, the susceptibility diverges like

$$\chi(z) \asymp \frac{1}{1 - z/z_c},$$

and the mean-square displacement grows linearly in n .

§60.5 Differential inequalities

One route to controlling the susceptibility is through differential inequalities. The generating function $\chi(z)$ is increasing in z , and differentiating it gives information about the average length of a walk under the z -weighted measure. The lace expansion supplies inequalities that compare $\chi'(z)$ to powers of $\chi(z)$.

The mean-field picture corresponds to the ordinary differential equation

$$\chi'(z) \approx C\chi(z)^2.$$

Solving this heuristic equation gives a simple pole at the critical point, hence $\gamma = 1$. The rigorous proof replaces the heuristic approximation by upper and lower differential inequalities with controlled error terms.

§60.6 Functional integral representations

The lecture notes also explain a different representation of the model. A weakly self-avoiding walk can be written using a self-intersection local time. In continuous time, one can weight a walk by

$$\exp\left(-g \sum_x L_x^2\right),$$

where L_x is the local time spent at x and $g > 0$ penalises self-intersections. The strict self-avoiding walk corresponds morally to the limit where self-intersections are forbidden completely.

There is a supersymmetric functional integral representation in which walk quantities are encoded by integrals involving bosonic and fermionic variables. The point is not that one wants to compute these integrals explicitly. The point is that the representation turns the probabilistic walk model into a field-theoretic object suitable for renormalisation group analysis.

The renormalisation-group part organises space into blocks and then into polymers made of blocks. This terminology is unfortunately overloaded: these “polymers” are not the original self-avoiding walks. They are bookkeeping regions used to localise the effective interaction at each scale.




Figure 60.3: a polymer at one scale and its closure at the next scale.

Remark 60.6.1 — The fermionic variables are not an extra physical particle in the polymer. They are an algebraic device that makes cancellations exact. In this representation, supersymmetry is a bookkeeping symmetry for the walk model.

§60.7 Dimension four and logarithmic corrections

Dimension 4 is the upper critical dimension for self-avoiding walk. Above four dimensions, the bubble diagram is finite and the lace expansion gives mean-field behaviour. At four dimensions, the relevant diagram is marginal: it fails only logarithmically. This is why the leading powers look mean-field but acquire logarithmic corrections.

For the continuous-time weakly self-avoiding walk in $d = 4$, the renormalisation group tracks how coupling constants change across length scales. One decomposes the covariance into scale pieces and integrates out fluctuations scale by scale. The effective interaction evolves under a map of the form

$$Z_0 \mapsto Z_1 \mapsto Z_2 \mapsto \cdots .$$

The essential phenomenon is that the self-interaction coupling is marginally irrelevant: it flows slowly to zero. The walk becomes asymptotically Gaussian, but only after logarithmic corrections are taken into account.

A typical conclusion has the form

$$\chi(g, \nu) \sim A \varepsilon^{-1} (\log \varepsilon^{-1})^{1/4}$$

as the killing parameter $\varepsilon \downarrow 0$, with model-dependent constants. The exact formulation depends on the version of the weakly self-avoiding walk, but the memory aid is stable: four dimensions means mean-field powers multiplied by logarithms.

§60.8 How the three methods fit together

The three major techniques in these lectures have different roles:

method	dimension or setting	main output
bridges and polygons	all dimensions	growth constants and subexponential bounds
holomorphic observable	$d = 2$, hexagonal lattice	exact connective constant
lace expansion	$d > 4$	mean-field critical behaviour
renormalisation group	$d = 4$	logarithmic corrections

This is a useful mental map because the methods are not interchangeable. The two-dimensional proof uses special planar structure. The high-dimensional proof uses small intersection probabilities. The four-dimensional proof needs a scale-by-scale analysis because the perturbation is exactly marginal at first order.

§60.9 The conceptual turn

The lace expansion should be remembered as the high-dimensional counterpart of the two-dimensional holomorphic observable. Both are ways of making the self-avoidance constraint tractable, but they exploit opposite regimes. In two dimensions, geometry is rigid and special identities appear. In high dimensions, geometry is loose and self-intersections become a small perturbation.

The critical dimension 4 is the boundary between these worlds. Above it, self-avoiding walk is close to simple random walk. At it, the same statement remains morally true, but logarithms record the slow decay of the self-interaction.

VII

Linear Algebra

Part VII: Contents

N-20210621	397
61.1 First orientation	397
61.2 Props as typed wiring diagrams	397
61.3 The alphabet: copy, discard, add, zero, and scalars	398
61.4 Two colours: a comonoid and a monoid	398
61.5 The bialgebra law	399
61.6 Scalars as ring operations	400
61.7 How a matrix becomes a diagram	400
61.8 Completeness: a proof system for matrices	401
61.9 Why the calculus hides indices and exposes structure	401
61.10 From matrices to linear relations	402
61.11 From linear relations to Lagrangian relations	402
61.12 Affine shifts and stabilizers	403
61.13 Relation with ZX-calculus	403
61.14 Broader perspective	403
 N-20220407	 405
62.1 Complex arithmetic	405
62.2 Field axioms in $\mathbb{C}$	405
62.3 Vector space identities	406
62.4 A non-example with infinity	406
62.5 Subspaces	406
62.6 Bridge to later ideas	407
62.7 How to read basic vector-space exercises	407
62.8 Beyond real and complex scalar structures	407
62.9 Quotient spaces	408
62.10 Dual spaces and coordinates	408
62.11 Returning to the original elementary exercises	408
62.12 Why fields make linear algebra work	408
62.13 The universal property of quotient spaces	409
 N-20220411	 411
63.1 Invertibility	411
63.2 The 2×2 formula	411
63.3 Diagonal matrices	411
63.4 Products of inverses	412
63.5 Adjugate formula	412
63.6 Row reduction for inverses	412
63.7 Rotation matrices	412
63.8 Beyond the first calculation	413
63.9 Invertibility as isomorphism	413
63.10 Adjugate and exterior powers	413
63.11 Rotations as orthogonal matrices	413
63.12 Invertibility, volume, and orthogonal motion	414
63.13 Inverse as undoing a linear process	414
63.14 Invertible matrices as a Lie group	414
63.15 Determinant as a volume form	414
63.16 Polar decomposition	415
63.17 Numerical stability of inverses	415
63.18 Row operations and the general linear group	415
63.19 Exponential map and infinitesimal motion	416
 N-20220412	 417
64.1 Unit vectors and Kronecker delta	417
64.2 Einstein summation convention	417
64.3 Rank	417
64.4 Transpose and Hermitian conjugate	418

64.5	Trace	418
64.6	A more structural view	418
64.7	Transpose and dual maps	419
64.8	Hermitian adjoints in Hilbert spaces	419
64.9	Trace as invariant	419
64.10	Schatten norms	419
64.11	Indices as tensor bookkeeping	420
64.12	Transpose as pullback on dual spaces	420
64.13	Trace as contraction	420
64.14	Rank and factorization	420
64.15	Coordinate indices as intrinsic linear algebra	420
N-20220725		421
65.1	Inversions and signs	421
65.2	Definition of determinant	421
65.3	Basic determinant properties	422
65.4	Sarrus and triangular matrices	422
65.5	Cofactor expansion	422
65.6	A structured determinant	422
65.7	The determinant line and exterior powers	423
65.8	Why permutation signs are forced	423
65.9	Minors as exterior coordinates	424
65.10	Cramer's rule from volume ratios	424
N-20230209		425
66.1	From vector identities to linear representations	425
66.2	Coordinates as isomorphisms	425
66.3	The basis-change formula	426
66.4	A worked basis-change example	426
66.5	Which quantities are invariant?	427
66.6	Covectors and the inverse-transpose rule	427
66.7	Bilinear forms are not linear operators	428
66.8	Symmetric and skew parts	429
66.9	Tensor slots and transformation laws	429
66.10	Tensor products and universal property	430
66.11	Coordinate-free reading of the original matrix facts	430
N-20230217		431
67.1	Tensor products and bases	431
67.2	Bilinearity	431
67.3	Metric, dot product, and components	432
67.4	Reciprocal basis	432
67.5	Covariant and contravariant components	432
67.6	Vector product and Levi-Civita symbol	433
67.7	Scalar triple product	433
67.8	What changes next	433
67.9	Tensor product by its universal property	433
67.10	Components under basis change	433
67.11	Metric as an index-lowering map	434
67.12	Tensor notation as transformation law	434
67.13	Dual spaces and reciprocal bases	434
67.14	Musical isomorphisms	434
67.15	Frames and moving coordinates	435
67.16	Continuum mechanics notation	435
67.17	Metric tensors and Gram matrices	435
67.18	Covariant and contravariant components from Gram matrices	435
67.19	Why tensor products linearize bilinear data	436
N-20251108		437
68.1	Vectors beyond arrows	437
68.2	The minus-one trick	437

68.3	Example of the trick	437
68.4	Linear mappings	438
68.5	Complex numbers as a real vector space	438
68.6	Finite-dimensional isomorphism theorem	438
68.7	Basis change	439
68.8	Equivalence and similarity	439
68.9	Kernel and image	439
68.10	Categorical side remarks	440
68.11	The general pattern	440
68.12	Vectors as elements of a linear structure	440
68.13	The minus-one trick as a kernel graph	440
68.14	Kernel-image factorization	441
68.15	Coordinates, kernels, and affine linearization	441
68.16	Exact sequences behind rank-nullity	442
68.17	Quotient description of the image	442
68.18	Cokernel and obstructions	442
68.19	Returning to the MML basis-change exercises	442
68.20	Row reduction as choosing a splitting	443
68.21	Cokernel as the missing equation space	443
N-20251110		445
69.1	Cofactor orthogonality	445
69.2	Vandermonde determinant	445
69.3	Block triangular determinants	446
69.4	General Laplace expansion	446
69.5	Determinant splitting	446
69.6	Adding rows or columns	447
69.7	Augmenting to a Vandermonde determinant	447
69.8	Row and column transformations	447
69.9	Triangularization by row operations	448
69.10	Recursive determinants	448
69.11	Cramer's rule	448
69.12	Rank and minors	449
69.13	Solvability of linear systems	449
69.14	Handwritten comments preserved	450
69.15	The lesson behind it	450
69.16	Determinant as alternating multilinear invariant	450
69.17	Vandermonde and evaluation maps	450
69.18	Rank and linear systems	450
69.19	Determinants and rank as invariants	450
69.20	Determinant through exterior powers	451
69.21	Schur complement	451
69.22	Matrix determinant lemma	451
69.23	Rank varieties	451
69.24	Returning to Cramer and rank	452
69.25	Determinant identities as low-rank geometry	452
69.26	Schur complements and eliminating variables	452
69.27	Determinantal varieties and singular loci	452
N-20251202		453
70.1	Diagonalizable matrices	453
70.2	Diagonalizability is independent from other properties	453
70.3	Algebraic and geometric multiplicity	454
70.4	Procedure for diagonalization	454
70.5	Inner product and vector length	454
70.6	Orthogonal and orthonormal systems	455
70.7	Schmidt orthogonalization	455
70.8	Example of Schmidt orthogonalization	456
70.9	QR decomposition	456
70.10	Orthogonal matrices	456
70.11	Real symmetric matrices	457

70.12	Quadratic forms	457
70.13	Positive definite quadratic forms	457
70.14	Conic and quadric interpretations	458
70.15	Common traps recorded in the notes	458
70.16	A closing lens	458
70.17	Diagonalization as decomposition into invariant lines	459
70.18	Algebraic versus geometric multiplicity	459
70.19	Inner products and orthogonal projection	459
70.20	Quadratic forms and spectral theorem	459
70.21	Spectral decomposition and quadratic geometry	459
70.22	Spectral theorem and orthogonal bases	459
70.23	QR as algorithmic Gram–Schmidt	460
70.24	Sylvester’s law of inertia	460
70.25	Hilbert-space projection	460
70.26	Orthogonal diagonalization and axes	460
70.27	Gram–Schmidt as factorization of coordinates	461
70.28	Spectral theorem beyond finite dimension	461
N-20260103		463
71.1	Determinant expansion and cofactors	463
71.2	Recursive determinants	463
71.3	Circulant-style determinant	464
71.4	Cancellation through orthogonal matrices	464
71.5	Determinant product and block simplification	464
71.6	Schur-complement style block transformation	465
71.7	Rank-two decomposition	465
71.8	Adjugate identities	466
71.9	Skew-symmetric matrices via quadratic forms	466
71.10	Absorption and cancellation	467
71.11	Rank identities	467
71.12	Commutation from factorization	468
71.13	Projection-like matrices from bases	468
71.14	Nilpotent polynomial inverse	468
71.15	Equivalent vector groups	469
71.16	Computing matrix powers	469
71.17	Homogeneous systems and adjugate	470
71.18	Null space basis and augmented rank	470
71.19	Rank of products and equality of solution spaces	471
71.20	General solution from known solutions	471
71.21	Equivalence of row vector groups	472
71.22	Diagonal dominance and invertibility	472
71.23	Where the idea leads	472
71.24	Determinant tricks as invariant-preserving transformations	472
71.25	Block determinants and Schur complements	473
71.26	Matrix powers and minimal polynomials	473
71.27	Matrix tricks from invariants and polynomials	473
71.28	Rank factorization and normal forms	473
71.29	Idempotents as projections	473
71.30	Spectral asymptotics of matrix powers	474
71.31	Null spaces and orthogonal complements	474
71.32	Matrix powers and stable limits	474
71.33	Row space, column space, and rank factorization	475
71.34	Positivity through Gram forms	475
N-20260105		477
72.1	Eigenvalues of matrix functions	477
72.2	No real eigenvector from a quadratic relation	477
72.3	Independence of eigenvectors	478
72.4	Real symmetric reconstruction from eigenvalues	478
72.5	Similarity invariants: trace and determinant	479
72.6	Trace shortcuts	479

72.7	Row-sum eigenvectors	479
72.8	Diagonalization problems	480
72.9	Real symmetric eigenvalue problems	480
72.10	Odd powers of real symmetric matrices	480
72.11	Rayleigh quotient extrema	481
72.12	Orthogonal determinant trick	481
72.13	Quadratic forms through diagonalization	481
72.14	Positive definite equations	482
72.15	Commuting positive definite products	482
72.16	Principal minors and congruence	483
72.17	Ordering eigenvalues and positive definiteness	483
72.18	Transition matrices	483
72.19	Same coordinates under two bases	484
72.20	Conceptual synthesis	484
72.21	Eigenvalues through functional calculus	484
72.22	Similarity invariants	484
72.23	Positive definiteness and quadratic geometry	485
72.24	Row-sum eigenvectors and invariant aggregate directions	485
72.25	Spectral geometry of positive matrices	485
72.26	Functional calculus for matrices	485
72.27	Loewner order and positive definiteness	485
72.28	Lyapunov equations	486
72.29	Convex and semidefinite optimization	486
72.30	Matrix monotone warning	486
72.31	Positive definite matrices and ellipsoids	486
72.32	Lyapunov equations and stability	487
72.33	Functional calculus and matrix inequalities	487
N-20260303		489
73.1	Concept map of the chapter	489
73.2	Five model transformations	489
73.3	A spectral example: a biological neural network	490
73.4	SVD as a sequence of geometric operations	491
73.5	SVD and latent structure in data	492
73.6	Rank-one blocks and image compression	492
73.7	A taxonomy of matrix properties	494
73.8	Closing perspective	494
N-20260305		497
74.1	Projection onto a subspace	497
74.2	Projection matrix and least squares	497
74.3	Pseudoinverse	497
74.4	Projection geometry in the source diagram	498
74.5	Bridge to later ideas	498
74.6	Projection theorem	498
74.7	Normal equations	498
74.8	Pseudoinverse and singular values	499
74.9	Least squares as projection geometry	499
N-20260311		501
75.1	Why size matters	501
75.2	Metrics and completeness	501
75.3	Normed vector spaces	501
75.4	Standard vector norms	502
75.5	Matrix norms	502
75.6	Unitary invariance	502
75.7	Induced norms	503
75.8	Normal matrices	503
75.9	Compatibility	503
75.10	Worked constrained maxima	503
75.11	ML-facing interpretation	504

75.12	What each norm is measuring	504
N-20260312		505
76.1	The long-term question	505
76.2	Spectral radius	505
76.3	Spectral radius is not a norm	505
76.4	Approximating spectral radius	506
76.5	Convergent matrices	506
76.6	Jordan boundary behavior	506
76.7	Neumann series	506
76.8	Matrix power series	506
76.9	Condition number	507
76.10	Perturbation bound	507
76.11	Ill-conditioned example	507
76.12	Numerical and ML interpretation	507
76.13	Power growth, inversion, and sensitivity	508
N-20260318		509
77.1	Why canonical forms matter	509
77.2	Polynomial matrices	509
77.3	Smith normal form	509
77.4	Jordan data from elementary divisors	509
77.5	Jordan chains	510
77.6	Computing powers	510
77.7	Gershgorin discs	510
77.8	Gershgorin figure	511
77.9	Separated components and scaling	511
77.10	Power iteration	511
77.11	Inverse and shifted inverse iteration	511
77.12	Numerical traps	512
77.13	ML-facing interpretation	512
77.14	Choosing an eigenvalue tool	512
N-20260602		513
78.1	Why decompositions are algorithms	513
78.2	Doolittle LU	513
78.3	Existence and uniqueness	513
78.4	Elimination matrices	513
78.5	Doolittle formulas	513
78.6	Pivoting	514
78.7	Storage of LU	514
78.8	Cholesky decomposition	514
78.9	Cholesky example	514
78.10	Full-rank decomposition	514
78.11	Hermite standard form	514
78.12	Singular values	515
78.13	Singular value decomposition	515
78.14	SVD geometry	515
78.15	Which decomposition matches the matrix	515
N-20260603		517
79.1	Why ordinary inverse is not enough	517
79.2	Four regimes	517
79.3	Penrose equations	517
79.4	Fifteen generalized inverses	517
79.5	Projection transformations	517
79.6	Idempotence and orthogonality	517
79.7	Column and row-space projections	518
79.8	SVD computation	518
79.9	Full-rank computation	518
79.10	Greville recursion	518

79.11	Iterative approximations	518
79.12	Solvability	519
79.13	Least squares and minimum norm	519
79.14	Identities	519
79.15	Reading the Moore–Penrose inverse geometrically	519
N-20260604		521
80.1	The computational problem	521
80.2	Triangular solves	521
80.3	Gauss elimination	521
80.4	Pivoting	521
80.5	LU solution	521
80.6	Cholesky and LDL	522
80.7	Tridiagonal chasing	522
80.8	Stationary iteration	522
80.9	Jacobi iteration	522
80.10	Gauss-Seidel iteration	522
80.11	Diagonal dominance	523
80.12	SOR iteration	523
80.13	Optimization viewpoint	523
80.14	Steepest descent	523
80.15	Method map	523
80.16	How the linear-system method is selected	524

Prop Presentations for Linear Algebra: Diagrams, Matrices, and Relations

N-20210621

§61.1 First orientation

The small diagram that first caught my attention looked almost too simple: a few white nodes, a few grey nodes, some scalar labels, and many wires. The surprising claim was that this was not merely a mnemonic for linear algebra. It was a presentation of linear algebra as a graphical calculus.

The original fascination was this:

ordinary linear algebra writes maps as grids of numbers;
graphical linear algebra writes them as processes made of wires and nodes.

This note decodes that special symbolic system. The point is not to admire the pictures, but to understand why the pictures form a proof system.

The guiding slogan is:

matrices are compiled diagrams, and diagrams are structured linear processes.

Once this is understood, the later movement toward linear relations, Lagrangian relations, circuits, and stabilizer calculi becomes much less mysterious.

§61.2 Props as typed wiring diagrams

The right ambient language is a prop.

Definition 61.2.1 (Prop). A prop is a strict symmetric monoidal category whose objects are natural numbers

$$0, 1, 2, \dots$$

and whose monoidal product on objects is addition.

The object n should be read as n parallel wires. A morphism

$$f : n \rightarrow m$$

is a process with n input wires and m output wires.

There are two basic operations on morphisms.

First, vertical composition plugs outputs into inputs:

$$f : n \rightarrow m, \quad g : m \rightarrow p, \quad g \circ f : n \rightarrow p.$$

Second, the monoidal product puts diagrams side by side:

$$f : n \rightarrow m, \quad g : n' \rightarrow m', \quad f \oplus g : n + n' \rightarrow m + m'.$$

This makes a prop feel like a typed programming language for wiring diagrams:

prop notation	programming / wiring intuition
n	n wires
$f : n \rightarrow m$	process with n inputs and m outputs
$g \circ f$	run f , then run g
$f \oplus g$	run f and g in parallel

The prop $\mathbf{Mat}_k$ has natural numbers as objects and matrices over k as morphisms. A morphism

$$n \rightarrow m$$

is an $m \times n$ matrix, regarded as a linear map

$$k^n \rightarrow k^m.$$

Composition is matrix multiplication, and the monoidal product is block direct sum.

§61.3 The alphabet: copy, discard, add, zero, and scalars

The graphical calculus $\mathbf{cb}_k$ is generated by a small alphabet. In the diagrams below, time flows upward: inputs are at the bottom and outputs are at the top.

symbol	type	linear meaning	intuition
 copy Δ	$1 \rightarrow 2$	$x \mapsto (x, x)$	duplicate one wire into two copies
 discard ϵ	$1 \rightarrow 0$	$x \mapsto *$	forget the value on a wire
 add μ	$2 \rightarrow 1$	$(x, y) \mapsto x + y$	combine two wires by addition
 zero η	$0 \rightarrow 1$	$* \mapsto 0$	create the additive zero
 scalar a	$1 \rightarrow 1$	$x \mapsto ax$	multiply a wire by $a \in k$

One must be careful about the word “copy.” In ordinary vector spaces, there is no natural linear map $V \rightarrow V \oplus V$ for every abstract vector space V unless one has fixed the diagonal map. In this prop, one wire represents the free rank-one module k . The copy generator is the linear map

$$k \rightarrow k \oplus k, \quad x \mapsto (x, x).$$

Arbitrary matrices are then built by composing and tensoring these elementary maps.

§61.4 Two colours: a comonoid and a monoid

The white structure is the copying structure:

$$\Delta : 1 \rightarrow 2, \quad \epsilon : 1 \rightarrow 0.$$

Algebraically, it is a commutative comonoid. The equations say that copying twice does not depend on which copy is copied first, that the two copies can be swapped, and that discarding a copied output returns the original wire.

The grey structure is the adding structure:

$$\mu : 2 \rightarrow 1, \quad \eta : 0 \rightarrow 1.$$

Algebraically, it is a commutative monoid. The equations say that adding three terms does not depend on bracketing, that order does not matter, and that adding zero changes nothing.

A good way to remember the two colours is:

colour	structure	operation
white	comonoid	copy / discard
grey	monoid	add / zero

The calculus becomes interesting because these two structures are not independent. They interact through the bialgebra law.

§61.5 The bialgebra law

The central equation says:

$$\text{copy after add} = \text{extadd after copying both inputs}.$$

In ordinary algebra, this is the statement

$$\Delta(x + y) = (x + y, x + y) = (x, x) + (y, y).$$

The two sides can be read as dataflow programs.

Left side:

```
input:  (x, y)
step 1: add          -> x + y
step 2: copy        -> (x + y, x + y)
output: (x + y, x + y)
```

Right side:

```
input:  (x, y)
step 1: copy x and copy y      -> (x, x, y, y)
step 2: swap the middle wires -> (x, y, x, y)
step 3: add pairwise          -> (x + y, x + y)
output: (x + y, x + y)
```

This is the first moment where the topology of wires matters. The swap of the middle wires is not decoration. It is the wiring needed to pair the first copy of x with the first copy of y , and the second copy of x with the second copy of y .

Remark 61.5.1 — The bialgebra law is the graphical form of distributivity. It says that copying respects addition. Without this law, the white and grey nodes would be two unrelated gadgets; with it, they become the algebraic core of linear algebra.

§61.6 Scalars as ring operations

A scalar bead labelled $a \in k$ is the one-wire map

$$[a] : 1 \rightarrow 1, \quad x \mapsto ax.$$

The scalar equations force these beads to behave like elements of the ring k .

Stacking scalar beads corresponds to multiplication:

$$[b] \circ [a] = [ba].$$

If k is commutative, this is just $[ab]$.

Parallel scalar branches combined by the add/copy structure express addition:

$$[a] + [b] = [a + b].$$

The identity scalar is a plain wire:

$$[1] = \text{id}_1.$$

The zero scalar factors through discard and zero:

$$[0] = \eta \circ \epsilon.$$

So the calculus has two layers:

structural layer		copy, discard, add, zero
coefficient layer		scalar beads labelled by elements of k

Together these layers generate matrices.

§61.7 How a matrix becomes a diagram

Consider the matrix

$$A = \begin{bmatrix} 2 & 3 \\ 5 & 7 \end{bmatrix} : k^2 \rightarrow k^2.$$

It sends

$$(x_1, x_2) \mapsto (y_1, y_2), \quad y_1 = 2x_1 + 3x_2, \quad y_2 = 5x_1 + 7x_2.$$

The diagrammatic recipe is:

columns	input wires
rows	output wires
A_{ij}	scalar bead on the branch from input j to output i
copy nodes	fan one input out to all rows
add nodes	sum all contributions landing at one output

In code-like form:

```
def compile_matrix(A):
    # A has m rows and n columns.
    # Input wire j carries x_j.
    # Output wire i should carry sum_j A[i][j] * x_j.
    for each input wire j:
        copy x_j into m branches
    for each pair (i, j):
        put scalar bead A[i][j] on branch j -> i
    for each output wire i:
        add all branches landing at i
```

For the example above, the diagram computes

$$\begin{aligned}y_1 &= 2x_1 + 3x_2, \\y_2 &= 5x_1 + 7x_2.\end{aligned}$$

This is why diagrams can represent arbitrary matrices. Copy creates enough branches; scalars weight the branches; adders collect contributions.

§61.8 Completeness: a proof system for matrices

The key theorem can be stated as follows.

Theorem 61.8.1 (Completeness of the matrix prop)

For a suitable ring k , the graphical theory $\mathbf{cb}_k$ presents the prop $\mathbf{Mat}_k$ of matrices over k under direct sum.

This theorem has two parts.

Soundness says that every equation derivable by diagram rewriting is true as a matrix equation. If two diagrams can be transformed into each other using the graphical rules, then they denote the same matrix.

Completeness says the converse: if two diagrams denote the same matrix, then the graphical rules are strong enough to prove them equal.

Thus the diagrams are not merely suggestive notation. They form a proof system for linear algebra:

$$\text{diagram equality} \iff \text{matrix equality}.$$

§61.9 Why the calculus hides indices and exposes structure

Matrix notation is extremely efficient, but it hides the process structure behind indices. The product of two matrices is written

$$(AB)_{ij} = \sum_k A_{ik} B_{kj}.$$

The summation index k is an internal wire: output from one process becomes input to another. Diagrammatic notation makes that internal connection visible.

The comparison is:

matrix notation	diagram notation
entries and indices	wires and nodes
matrix multiplication	vertical plugging
block direct sum	side-by-side placement
transpose-like operations	reflection or reversal, depending on convention
trace / feedback	connecting an output back to an input

It would be too strong to say that $\mathbf{Mat}_k$ is completely basis-free: wire counts still refer to finite free modules k^n . The better statement is:

the calculus hides indices and exposes compositional structure.

This is exactly what makes it useful in networks, signal flow graphs, circuits, tensor contractions, and other settings where the shape of composition matters more than individual coordinates.

§61.10 From matrices to linear relations

A matrix is a total linear function

$$k^n \rightarrow k^m.$$

A linear relation from k^n to k^m is more general. It is a linear subspace

$$R \subseteq k^n \oplus k^m.$$

Instead of saying each input has exactly one output, a relation says which input-output pairs are compatible.

For example, the equation

$$x + y = z$$

defines a linear relation among x, y, z . It is not best understood as a function until one chooses which variables are inputs and which are outputs.

This is why graphical calculi naturally move beyond matrices. Wires and nodes are good at expressing constraints. A relation can be composed by connecting shared variables and existentially hiding the internal wires. This is the relational analogue of matrix multiplication.

The broader graphical-linear-algebra story uses interacting Hopf-algebra structures to present such relations. The matrix calculus $\mathbf{cb}_k$ is the entry point; linear relations are the next layer.

§61.11 From linear relations to Lagrangian relations

The later part of the story enters symplectic linear algebra. A symplectic vector space is a vector space V equipped with a nondegenerate skew-symmetric form

$$\omega : V \times V \rightarrow k.$$

The standard example is k^{2n} with form represented by the block matrix

$$J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}.$$

For a subspace $W \subseteq V$, define its symplectic orthogonal by

$$W^\omega = \{v \in V : \omega(v, w) = 0 \text{ for all } w \in W\}.$$

A subspace is Lagrangian if

$$W = W^\omega.$$

Definition 61.11.1 (Lagrangian relation, informal). A Lagrangian relation is a linear relation that is Lagrangian inside the appropriate symplectic input-output space.

The conceptual shift is:

linear relation	linear constraint
Lagrangian relation	symplectic or phase-space-compatible constraint

This is why the calculus becomes relevant to physics. Phase space, circuits, and stabilizer-like systems are naturally symplectic.

§61.12 Affine shifts and stabilizers

The affine version allows translations. Instead of linear subspaces passing through the origin, one allows affine subspaces:

$$v + W.$$

Diagrammatically, this means adding generators that can shift a relation by a constant.

This affine Lagrangian-relational world connects to several process theories:

domain	role of the calculus
electrical circuits	constraints between currents and potentials
signal flow graphs	linear dynamical interconnection
Spekkens' toy theory	epistemic states as affine Lagrangian relations
qudit stabilizer circuits	symplectic/affine structure over finite fields

For this note, these applications should be read as the research-level endpoint rather than the entry point. The entry point is still the alphabet: copy, discard, add, zero, and scalars.

§61.13 Relation with ZX-calculus

The resemblance to ZX-calculus is real but should be stated carefully. Both calculi use coloured nodes, wires, and bialgebra-like interaction laws. Both are props or prop-like symmetric monoidal graphical languages. Both replace long index manipulations by diagrammatic rewriting.

The difference is the intended semantics:

graphical linear algebra	ZX-calculus
matrices, relations, linear constraints	quantum processes
copy/add structures	complementary observables
linear algebra and circuits	categorical quantum mechanics

Thus $\mathbf{cb}_k$ is not simply ZX-calculus. It is better seen as a training ground for understanding why ZX-like diagrammatic calculi work: algebraic equations are replaced by local topological rewrites.

§61.14 Broader perspective

The main lesson is not that matrices can be drawn prettily. The main lesson is that a small symbolic system can present all matrix algebra.

The core dictionary is:

diagrammatic object	linear-algebraic meaning
n wires	k^n
diagram $n \rightarrow m$	$m \times n$ matrix
vertical composition	matrix multiplication
side-by-side composition	block direct sum
copy	$x \mapsto (x, x)$
add	$(x, y) \mapsto x + y$
scalar bead a	$x \mapsto ax$

The bialgebra law is the heart of the system:

copying a sum is the same as summing the copies.

The completeness theorem says that this is not only notation but a genuine proof system:

diagrammatic equality is matrix equality.

After this is understood, the path toward linear relations, Lagrangian relations, affine shifts, circuits, and stabilizer calculi is no longer a jump. It is the natural extension of the same idea: algebra can be presented as topology of wires.

Exercises on $\mathbb{R}^n$, $\mathbb{C}^n$, Vector Spaces, and Subspaces

N-20220407

§62.1 Complex arithmetic

The first page is a worked set of exercises from the beginning of linear algebra, where $\mathbb{C}$ is treated as a field.

If $a, b \in \mathbb{R}$ are not both zero, then

$$\frac{1}{a + bi} = \frac{a - bi}{a^2 + b^2} = \frac{a}{a^2 + b^2} - \frac{b}{a^2 + b^2}i.$$

Thus the required real numbers c, d in

$$\frac{1}{a + bi} = c + di$$

are

$$c = \frac{a}{a^2 + b^2}, \quad d = -\frac{b}{a^2 + b^2}.$$

A cube root of 1 other than 1 is

$$\omega = \frac{-1 + \sqrt{3}i}{2},$$

since

$$\omega^2 + \omega + 1 = 0, \quad \omega^3 = 1.$$

The two nontrivial cube roots are

$$\frac{-1 + \sqrt{3}i}{2}, \quad \frac{-1 - \sqrt{3}i}{2}.$$

Two square roots of i are obtained from polar form:

$$i = e^{i\pi/2}, \quad \sqrt{i} = e^{i\pi/4} = \frac{1 + i}{\sqrt{2}}, \quad -\sqrt{i} = -\frac{1 + i}{\sqrt{2}}.$$

§62.2 Field axioms in $\mathbb{C}$

If

$$\alpha = a + bi, \quad \beta = c + di, \quad \lambda = x + yi,$$

then commutativity and associativity of addition follow from those of real numbers. For example,

$$\alpha + \beta = (a + c) + (b + d)i = (c + a) + (d + b)i = \beta + \alpha.$$

Similarly,

$$\lambda(\alpha + \beta) = \lambda\alpha + \lambda\beta$$

follows by expanding both sides and using distributivity in $\mathbb{R}$.

Every $\alpha \in \mathbb{C}$ has an additive inverse, namely $-\alpha$. If $\alpha \neq 0$, its multiplicative inverse is

$$\alpha^{-1} = \frac{\bar{\alpha}}{|\alpha|^2}.$$

§62.3 Vector space identities

In a vector space V over a field F , one proves

$$-(-v) = v$$

because $-v + v = 0$, so the additive inverse of $-v$ is v .

If $a \in F$, $v \in V$, and $av = 0$, then either $a = 0$ or $v = 0$. Indeed, if $a \neq 0$, multiply by a^{-1} :

$$v = a^{-1}(av) = a^{-1}0 = 0.$$

Given $v, w \in V$, the equation

$$v + 3x = w$$

has the unique solution

$$x = \frac{1}{3}(w - v),$$

provided $3 \neq 0$ in the field.

§62.4 A non-example with infinity

The page asks whether

$$\mathbb{R} \cup \{\infty\} \cup \{-\infty\}$$

can be made into a real vector space with the suggested arithmetic. It cannot. One obstruction is associativity. If one declares

$$\infty + (-\infty) = 0,$$

then

$$(\infty + \infty) + (-\infty) = \infty + (-\infty) = 0,$$

while

$$\infty + (\infty + (-\infty)) = \infty + 0 = \infty.$$

Thus associativity fails.

§62.5 Subspaces

A subset $U \subseteq V$ is a subspace iff it is nonempty and closed under addition and scalar multiplication. Equivalently,

$$x, y \in U, \ a, b \in F \quad \Rightarrow \quad ax + by \in U.$$

Examples from the page include:

- (i) $\{(x_1, x_2, x_3, x_4) \in F^4 : x_3 = 5x_4 + b\}$ is a subspace iff $b = 0$.
- (ii) The continuous real-valued functions on $[0, 1]$ form a subspace of $\mathbb{R}^{[0,1]}$.
- (iii) The differentiable real-valued functions on $\mathbb{R}$ form a subspace of $\mathbb{R}^{\mathbb{R}}$.
- (iv) The differentiable functions f on $(0, 3)$ satisfying $f'(2) = b$ form a subspace iff $b = 0$.
- (v) The complex sequences converging to 0 form a subspace of $\mathbb{C}^{\infty}$.

§62.6 Bridge to later ideas

These exercises are elementary, but they encode several structural distinctions that reappear throughout algebra and geometry.

Elementary exercise	Later structure
checking closure under linear combinations	subspaces and kernels
conditions like $x_3 = 5x_4 + b$	linear subspace versus affine translate
real versus complex scalar multiplication	scalar restriction and complex linearity
solving constraints with free variables	quotients and cosets of kernels
coordinates of vectors	dual basis and linear functionals

For example, multiplication by $a+bi$ on $\mathbb{C}$, viewed as a real vector space, is represented by

$$\begin{pmatrix} a & -b \\ b & a \end{pmatrix}.$$

Its determinant is $a^2 + b^2$, the squared modulus. Thus complex scalar multiplication already contains the geometry of rotation and scaling.

§62.7 How to read basic vector-space exercises

The purpose of these exercises is not to check axioms forever. The real skill is to recognize when an object is already known to be a vector space and when a proposed subset passes the subspace test. The fastest mental template is:

$$0 \in U, \quad u, v \in U \Rightarrow u + v \in U, \quad \lambda u \in U.$$

The examples $f'(2) = b$, $x_3 = 5x_4 + b$, and sequences converging to zero are all testing the same issue: affine conditions usually fail to be subspaces unless the constant term is zero. In later linear algebra this becomes the distinction between a linear subspace and an affine translate.

§62.8 Beyond real and complex scalar structures

The elementary note introduces vector spaces through closure under addition and scalar multiplication. The reason this structure is so rigid is that the scalars form a field. A vector space over a field k is a module over k , and field scalars make every linearly independent set extend to a basis. This is why every finite-dimensional vector space has a dimension and why all bases have the same number of vectors.

If the scalars are changed from a field to a ring, the same axioms define a module. Modules over rings need not have bases. For example, $\mathbb{Z}/2\mathbb{Z}$ is a module over $\mathbb{Z}$, but it is not a free module: no element can serve as an integer-basis vector because $2x = 0$ for every x .

This explains why elementary linear algebra is unusually clean. It is not merely because the definitions are simple; it is because fields remove torsion and many divisibility obstacles.

§62.9 Quotient spaces

If $U \subseteq V$ is a subspace, the quotient space V/U consists of cosets

$$v + U = \{v + u : u \in U\}.$$

Two vectors represent the same coset iff their difference lies in U . The quotient forgets directions inside U and keeps only the remaining information.

The dimension formula is

$$\dim(V/U) = \dim V - \dim U.$$

This is the structural version of “free variables after constraints.” Solving a linear system often means identifying a particular affine coset of the null space. The quotient-space language makes this precise.

§62.10 Dual spaces and coordinates

The dual space is

$$V^* = \text{Hom}(V, k),$$

the vector space of linear functionals $V \rightarrow k$. If $e_1, \dots, e_n$ is a basis of V , the dual basis $e^1, \dots, e^n$ is defined by

$$e^i(e_j) = \delta_j^i.$$

Coordinates are evaluations by dual basis vectors:

$$v = \sum_i x_i e_i, \quad x_i = e^i(v).$$

Thus a coordinate is not the vector itself; it is the value of a linear functional on the vector.

This distinction becomes essential for tensors, differential forms, gradients, and coordinate changes. The elementary row/column distinction is the first shadow of the distinction between vectors and covectors.

§62.11 Returning to the original elementary exercises

Subspaces, bases, linear independence, and dimension are not merely bookkeeping. They are the field-module case of a broader algebraic story. Quotients explain constraints; duals explain coordinates; modules explain why vector spaces are the friendly special case. The original exercises train the stable finite-dimensional core from which these advanced notions grow.

§62.12 Why fields make linear algebra work

A useful way to reread the original vector-space exercises is to ask which steps secretly use division. When we prove that a nonzero vector can be rescaled, that a pivot can be normalized to 1, or that a linearly independent list can be extended to a basis, we use the fact that every nonzero scalar has an inverse. This is exactly the field condition.

Over a general ring, the same-looking statements may fail. The equation $2x = 1$ has no solution over $\mathbb{Z}$, so one cannot normalize a pivot 2 by division. The module $\mathbb{Z}/2\mathbb{Z}$

has torsion because $2x = 0$ for every element. This is why row reduction over fields is clean, while integer linear algebra needs Smith normal form.

The elementary notion of a basis also becomes more delicate over rings. A free module has a basis, but not every module is free. Therefore the theorem

$$\dim V = \text{number of basis vectors}$$

is not just a counting convenience; it is a privilege of vector spaces over fields. This explains why linear algebra is often the first genuinely structural algebra course: it gives a setting where generation, independence, coordinates, and quotienting all behave perfectly.

§62.13 The universal property of quotient spaces

The quotient V/U is not merely a set of cosets. It is characterized by a universal property. Let $q : V \rightarrow V/U$ be the quotient map. If a linear map $T : V \rightarrow W$ kills U , meaning $U \subseteq \ker T$, then there is a unique linear map $\bar{T} : V/U \rightarrow W$ such that

$$T = \bar{T} \circ q.$$

This says that any linear map unable to distinguish vectors differing by elements of U must factor through the quotient.

This universal property is the conceptual reason quotient spaces appear in homology, representation theory, and functional analysis. In homology, boundaries are declared trivial, so cycles differing by boundaries become the same class. In linear algebra exercises, free variables often describe a coset of a kernel; quotient language explains what information survives after ignoring kernel directions.

Invertible Matrices, Adjoins, Row Reduction, and Rotations

N-20220411

§63.1 Invertibility

A square matrix A is invertible if there exists a matrix A^{-1} such that

$$A^{-1}A = I, \quad AA^{-1} = I.$$

Some immediate consequences recorded in the note are:

- (i) If A is invertible, then Gaussian elimination has a pivot in every column and every row.
- (ii) A may have many distinct powers, but it has only one inverse.
- (iii) If A is invertible, then the equation $Ax = b$ has the unique solution

$$x = A^{-1}b.$$

- (iv) If the homogeneous equation $Ax = 0$ has a nonzero solution, then A is not invertible.

The last point is fundamental: invertibility is equivalent to trivial kernel for a square matrix.

§63.2 The 2×2 formula

For a 2×2 matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

we have

$$A^{-1} = \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix},$$

provided

$$ad - bc \neq 0.$$

Thus the determinant is the obstruction to invertibility.

§63.3 Diagonal matrices

If

$$A = \text{diag}(d_1, \dots, d_n),$$

then A is invertible iff every $d_i \neq 0$, and

$$A^{-1} = \text{diag}(d_1^{-1}, \dots, d_n^{-1}).$$

This is the simplest case where the inverse is completely visible.

§63.4 Products of inverses

For invertible matrices,

$$(AB)^{-1} = B^{-1}A^{-1}.$$

More generally,

$$(ABC)^{-1} = C^{-1}B^{-1}A^{-1}.$$

The order reverses because composition reverses when undoing operations.

§63.5 Adjugate formula

Let $A = (a_{ij})$. Delete row i and column j to get the minor M_{ij} . The cofactor is

$$C_{ij} = (-1)^{i+j} \det M_{ij}.$$

The adjugate matrix is the transpose of the cofactor matrix:

$$\text{adj}(A) = (C_{ji}).$$

Then

$$A^{-1} = \frac{1}{\det A} \text{adj}(A)$$

when $\det A \neq 0$.

The note calls this formula computationally expensive but conceptually important. It proves that entries of the inverse are rational functions of entries of A .

§63.6 Row reduction for inverses

The practical method is row reduction:

$$[A \mid I] \longrightarrow [I \mid A^{-1}].$$

For example, for

$$M = \begin{pmatrix} 1 & 4 \\ 2 & 9 \end{pmatrix},$$

one has

$$\det M = 9 - 8 = 1,$$

and hence

$$M^{-1} = \begin{pmatrix} 9 & -4 \\ -2 & 1 \end{pmatrix}.$$

The note also illustrates obtaining this inverse by elementary row operations on the augmented matrix.

§63.7 Rotation matrices

The final inserted page discusses planar rotations. The rotation matrix is

$$R_\theta = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}.$$

It sends

$$\begin{pmatrix} x \\ y \end{pmatrix}$$

to

$$\begin{pmatrix} x \cos \theta - y \sin \theta \\ x \sin \theta + y \cos \theta \end{pmatrix}.$$

Its inverse is

$$R_\theta^{-1} = R_{-\theta} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}.$$




Figure 63.1: A planar rotation acts linearly on vectors and is represented by R_θ .

§63.8 Beyond the first calculation

Invertible matrices are automorphisms of finite-dimensional vector spaces. Gaussian elimination, determinants, adjugates, and block decompositions are different ways to detect or compute the same object: a reversible linear transformation.

§63.9 Invertibility as isomorphism

A matrix is invertible exactly when the associated linear map is an isomorphism. This means three equivalent things in finite dimensions:

$$\ker A = 0, \quad \operatorname{im} A = V, \quad \det A \neq 0.$$

Row reduction tests the same fact by constructing inverse operations.

§63.10 Adjugate and exterior powers

The adjugate formula

$$A^{-1} = \frac{1}{\det A} \operatorname{adj}(A)$$

comes from cofactors. Conceptually, determinants describe the action of A on top exterior power, while cofactors encode how $(n-1)$ -dimensional volume elements transform.

§63.11 Rotations as orthogonal matrices

A rotation matrix R satisfies

$$R^T R = I, \quad \det R = 1.$$

The first condition preserves lengths and angles; the second preserves orientation. Thus rotations are the orientation-preserving part of the orthogonal group.

§63.12 Invertibility, volume, and orthogonal motion

Inverse formulas, row reduction, adjugates, diagonal matrices, and rotations are all different faces of isomorphism: solve uniquely, preserve or rescale volume, and understand what the map does independently of coordinates.

§63.13 Inverse as undoing a linear process

The inverse matrix should be read as an operation that undoes a linear process. The formula

$$(AB)^{-1} = B^{-1}A^{-1}$$

then becomes obvious: if one first applies B and then A , one must undo A first and then undo B . Row reduction, adjugates, and determinants are three different languages for this same reversibility.

The rotation matrix is a particularly good example because the inverse is visible geometrically. Rotating by θ is undone by rotating by $-\theta$, so

$$R_\theta^{-1} = R_{-\theta} = R_\theta^T.$$

This is the first glimpse of orthogonal matrices, where algebraic inverse equals geometric transpose.

§63.14 Invertible matrices as a Lie group

The set of all invertible $n \times n$ matrices over $\mathbb{R}$ is

$$GL(n, \mathbb{R}) = \{A : \det A \neq 0\}.$$

It is a group under multiplication and also an open subset of $\mathbb{R}^{n^2}$. Therefore it is a Lie group: a group with a smooth manifold structure.

The tangent space at the identity is the Lie algebra

$$\mathfrak{gl}(n, \mathbb{R}) = M_n(\mathbb{R}).$$

The exponential map

$$\exp X = I + X + \frac{X^2}{2!} + \cdots$$

connects infinitesimal generators to invertible transformations. Elementary matrix powers and exponentials are thus the entry point to continuous symmetry.

§63.15 Determinant as a volume form

The determinant is not merely a formula for invertibility. It is the factor by which a linear map scales oriented volume:

$$\text{Vol}(Av_1, \dots, Av_n) = \det(A) \text{Vol}(v_1, \dots, v_n).$$

The multiplicative identity

$$\det(AB) = \det A \det B$$

says that volume scaling composes multiplicatively.

This is why row operations affect determinants in exactly the familiar ways. Swapping rows reverses orientation, scaling a row scales volume, and adding one row to another shears volume without changing it.

§63.16 Polar decomposition

Every invertible real matrix has a polar decomposition

$$A = QP,$$

where Q is orthogonal and P is symmetric positive definite. One can take

$$P = (A^T A)^{1/2}, \quad Q = AP^{-1}.$$

Geometrically, P stretches along orthogonal principal axes, and Q rotates or reflects. This separates deformation into “pure strain” and “rigid motion.”

The elementary inverse/rotation examples become clearer through this decomposition. Orthogonal matrices preserve lengths; positive definite matrices perform anisotropic stretching; a general invertible matrix is both.

§63.17 Numerical stability of inverses

Although A^{-1} exists when $\det A \neq 0$, computing it may be unstable. The condition number

$$\kappa_2(A) = \|A\|_2 \|A^{-1}\|_2 = \frac{\sigma_{\max}(A)}{\sigma_{\min}(A)}$$

measures how much errors may be amplified. A tiny determinant is not by itself the best diagnostic; singular values describe directional collapse.

Thus the elementary criterion $\det A \neq 0$ answers existence, while numerical linear algebra asks sensitivity.

§63.18 Row operations and the general linear group

Gaussian elimination writes an invertible matrix as a product of elementary matrices. Geometrically, elementary matrices are simple transformations: shears, scalings, and swaps. Thus an invertible linear transformation can be built from elementary geometric moves. This is the computational skeleton behind the group $GL(n)$.

The determinant map

$$\det : GL(n, \mathbb{R}) \rightarrow \mathbb{R}^\times$$

is a group homomorphism. Its kernel is the special linear group

$$SL(n, \mathbb{R}) = \{A : \det A = 1\}.$$

This group consists of volume-preserving linear transformations. The elementary determinant tests in the original note therefore separate invertible maps by how they scale volume.

For rotations, the relevant subgroup is

$$O(n) = \{Q : Q^T Q = I\}, \quad SO(n) = O(n) \cap SL(n, \mathbb{R}).$$

Orthogonal matrices preserve the inner product, hence all lengths and angles. The chain

$$SO(n) \subset O(n) \subset GL(n)$$

organizes the original examples: rotations are the most rigid, orthogonal maps allow reflections, and invertible maps allow arbitrary linear distortion without collapse.

§63.19 Exponential map and infinitesimal motion

A smooth path $A(t)$ in $GL(n)$ with $A(0) = I$ has derivative $X = A'(0)$, an arbitrary matrix. Conversely, a matrix X generates the path

$$A(t) = e^{tX}.$$

If X is skew-symmetric, then e^{tX} is orthogonal. Indeed,

$$\frac{d}{dt}(e^{tX})^T e^{tX} = 0$$

when $X^T = -X$. This gives a precise meaning to “skew-symmetric matrices generate rotations.”

The elementary rotation matrices in the note are therefore not isolated formulas; they are one-parameter subgroups of $GL(n)$.

Linear Algebra Continued: Indices, Rank, Transpose, Hermitian Conjugate, and Trace

N-20220412

§64.1 Unit vectors and Kronecker delta

This note follows the previous linear algebra material and collects notation that will be useful later.

Definition 64.1.1. A unit vector is a vector of length 1. I write such a vector as $\mathbf{v}$ when the direction matters more than its magnitude.

Definition 64.1.2 (Kronecker delta). The Kronecker delta is

$$\delta_{ij} = \begin{cases} 1, & i = j, \\ 0, & i \neq j. \end{cases}$$

Thus the matrix (δ_{ij}) is the identity matrix:

$$(\delta_{ij}) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

§64.2 Einstein summation convention

Definition 64.2.1. Einstein summation convention says that when an index is repeated exactly twice in a term, summation over that index is understood:

$$x_i y_i = \sum_i x_i y_i.$$

A repeated index is called a dummy index, and its name does not matter:

$$x_i y_i = x_j y_j = x_k y_k.$$

The basic rules are:

- (i) an index appearing once is a free index;
- (ii) an index appearing twice is summed over;
- (iii) an index appearing three or more times in one term is not allowed.

§64.3 Rank

Definition 64.3.1. The column rank of a matrix is the maximal number of linearly independent columns. The row rank is the maximal number of linearly independent rows. These two numbers are equal, and their common value is the rank of the matrix.

Gaussian elimination is the standard way to compute rank: reduce the matrix to row echelon form and count the nonzero rows.

§64.4 Transpose and Hermitian conjugate

Definition 64.4.1. If A is an $m \times n$ matrix, its transpose A^T is the $n \times m$ matrix defined by

$$(A^T)_{ij} = A_{ji}.$$

Definition 64.4.2. The Hermitian conjugate of A is

$$A^* = (A^T)^- = \overline{A}^T.$$

It satisfies

$$(AB)^* = B^* A^*.$$

Indeed,

$$(AB)^* = \overline{(AB)}^T = (\overline{A} \overline{B})^T = \overline{B}^T \overline{A}^T = B^* A^*.$$

Definition 64.4.3. A matrix is symmetric if

$$A^T = A.$$

A matrix is Hermitian if

$$A^* = A.$$

A matrix is skew-Hermitian if

$$A^* = -A.$$

In a skew-Hermitian matrix, the diagonal entries are purely imaginary.

§64.5 Trace

Definition 64.5.1. The trace of an $n \times n$ matrix A is the sum of its diagonal entries:

$$\text{tr}(A) = \sum_i A_{ii}.$$

In Einstein notation this is simply

$$\text{tr}(A) = A_{ii}.$$

§64.6 A more structural view

The common theme is index notation. The Kronecker delta represents the identity, repeated indices encode summation, and matrix operations such as transpose, adjoint, and trace become compact when written with indices.

§64.7 Transpose and dual maps

If $A : V \rightarrow W$ is a linear map, the dual map is

$$A^* : W^* \rightarrow V^*, \quad A^*(\varphi) = \varphi \circ A.$$

In bases, this dual map is represented by the transpose matrix. This explains why transposition reverses order:

$$(AB)^T = B^T A^T.$$

Composition of maps reverses when pulled back to functionals.

The elementary transpose is therefore not just flipping a matrix; it is the coordinate form of passing from vectors to covectors.

§64.8 Hermitian adjoints in Hilbert spaces

In an inner product space, the adjoint A^* is defined by

$$\langle Ax, y \rangle = \langle x, A^*y \rangle.$$

Over $\mathbb{C}$, its matrix is the conjugate transpose. A Hermitian matrix satisfies $A = A^*$ and has real eigenvalues. A unitary matrix satisfies $A^*A = I$ and preserves the inner product.

The elementary real symmetric matrix is thus the real finite-dimensional case of self-adjoint operators. This becomes central in quantum mechanics, Fourier analysis, and spectral theory.

§64.9 Trace as invariant

The trace is the sum of diagonal entries:

$$\text{tr } A = \sum_i a_{ii}.$$

It is invariant under similarity because

$$\text{tr}(P^{-1}AP) = \text{tr}(APP^{-1}) = \text{tr } A.$$

The cyclic rule

$$\text{tr}(AB) = \text{tr}(BA)$$

is the key identity.

In category-theoretic language, trace abstracts the idea of closing an input-output wire. In finite-dimensional vector spaces, that abstract trace reduces to the familiar matrix trace. This is why traces appear in representation theory, statistical mechanics, and index theory.

§64.10 Schatten norms

Rank counts nonzero singular values, while trace and determinant combine eigenvalues. A unified family of matrix norms uses singular values:

$$\|A\|_{S_p} = \left(\sum_i \sigma_i(A)^p \right)^{1/p}.$$

For $p = 2$ this is the Frobenius norm; for $p = 1$ it is the nuclear norm. These norms are finite-dimensional models of operator ideals.

Thus rank, trace, Hermitian adjoint, and transpose are all pieces of a larger operator-theoretic vocabulary.

§64.11 Indices as tensor bookkeeping

The Kronecker delta δ_i^j is the coordinate symbol for the identity map. Einstein summation records contractions, such as

$$y_i = A_i^j x_j.$$

The repeated index is summed because a linear map consumes one coordinate direction and produces another.

§64.12 Transpose as pullback on dual spaces

For $A : V \rightarrow W$, the dual map is

$$A^* : W^* \rightarrow V^*, \quad A^*(\varphi) = \varphi \circ A.$$

In bases, this is represented by the transpose. Hermitian conjugate adds complex conjugation because complex inner products are conjugate-linear in one argument.

§64.13 Trace as contraction

The trace

$$\mathrm{tr}(A) = A_i^i$$

is the contraction of one upper and one lower index. It is invariant under change of basis because it is not really a coordinate sum; it is the categorical trace of an endomorphism in finite-dimensional vector spaces.

§64.14 Rank and factorization

Rank is the dimension of the image. A rank- r map factors through an r -dimensional space:

$$V \twoheadrightarrow \mathrm{im} A \hookrightarrow W.$$

This explains why rank measures the true number of independent output directions.

§64.15 Coordinate indices as intrinsic linear algebra

Indices, transpose, Hermitian conjugate, rank, and trace are coordinate shadows of intrinsic operations: duality, contraction, image factorization, and inner-product-dependent adjoints.

Permutations, Inversions, and Determinants

N-20220725

§65.1 Inversions and signs

The note begins with inversions of a permutation. For a permutation

$$\sigma = (\sigma_1, \dots, \sigma_n),$$

an inversion is a pair $i < j$ such that

$$\sigma_i > \sigma_j.$$

The number of inversions is denoted $\tau(\sigma)$. The sign of the permutation is

$$\text{sgn}(\sigma) = (-1)^{\tau(\sigma)}.$$

For example, for 426315, the inversions are counted by checking how many later entries are smaller than each entry. The note obtains

$$\tau(426315) = 8.$$

In determinant expansions, the sign of each term is exactly the sign of a permutation. Thus a term like

$$a_{21}a_{53}a_{16}a_{42}a_{65}a_{34}$$

corresponds to a row/column permutation, and its sign is determined by its inversion number.

§65.2 Definition of determinant

Definition 65.2.1 (Determinant). The determinant is the alternating multilinear function of the columns of a square matrix normalized by

$$\det I = 1.$$

It measures oriented volume scaling.

The determinant of an $n \times n$ matrix $A = (a_{ij})$ is

$$\det A = \sum_{\sigma \in S_n} \text{sgn}(\sigma) a_{1\sigma(1)} a_{2\sigma(2)} \cdots a_{n\sigma(n)}.$$

This formula is not meant for computation in large dimensions. Its value is conceptual: it explains why row swaps change sign, why repeated rows give zero, and why the determinant is multilinear in rows and columns.

§65.3 Basic determinant properties

The note records the following standard properties:

- (i) Transposition does not change the determinant:

$$\det A^T = \det A.$$

- (ii) Swapping two rows changes the sign.
- (iii) Pulling a scalar out of one row pulls it out of the determinant.
- (iv) The determinant is additive in each row separately.
- (v) Adding a multiple of one row to another row does not change the determinant.

From these one immediately gets the zero properties:

- (i) if a row is zero, then the determinant is zero;
- (ii) if two rows are proportional, then the determinant is zero.

§65.4 Sarrus and triangular matrices

For a 3×3 matrix, Sarrus' rule gives a convenient visual computation. But the note also emphasizes that triangular matrices are easier: if A is upper or lower triangular, then

$$\det A = a_{11}a_{22} \cdots a_{nn}.$$

This is why row reduction is so powerful. We use row operations to simplify a determinant into triangular form while tracking sign changes and scalar factors.

§65.5 Cofactor expansion

The cofactor of a_{ij} is

$$C_{ij} = (-1)^{i+j} M_{ij},$$

where M_{ij} is the determinant obtained by deleting row i and column j . Expanding along row i ,

$$\det A = \sum_{j=1}^n a_{ij} C_{ij}.$$

A page in the note works through determinant computations by choosing rows or columns with many zeros. This is the main practical advice: do not expand randomly; expand where the determinant is sparse.

§65.6 A structured determinant

A common determinant in the note has the form

$$D_n = \begin{vmatrix} x & a & a & \cdots & a \\ a & x & a & \cdots & a \\ \vdots & \vdots & \vdots & & \vdots \\ a & a & a & \cdots & x \end{vmatrix}.$$

This matrix can be written

$$D_n = (x - a)I + aJ,$$

where J is the all-one matrix. The vector $(1, \dots, 1)$ is an eigenvector with eigenvalue

$$x + (n - 1)a,$$

and any vector whose coordinates sum to 0 has eigenvalue

$$x - a.$$

Therefore

$$D_n = (x + (n - 1)a)(x - a)^{n-1}.$$

The note obtains this by row and column operations; the eigenvalue argument is the higher-level explanation.

§65.7 The determinant line and exterior powers

The previous computational rules all point to one structure. Let V be an n -dimensional vector space. The top exterior power

$$\Lambda^n V$$

is one-dimensional and is called the determinant line. A nonzero element of this line is a volume form; a choice up to positive scalar is an orientation.

Definition 65.7.1 (Determinant via the top exterior power). For a linear map $A : V \rightarrow V$, the induced map

$$\Lambda^n A : \Lambda^n V \rightarrow \Lambda^n V$$

acts on a one-dimensional space, hence is multiplication by a scalar. That scalar is $\det A$.

In a basis $e_1, \dots, e_n$, if the columns of A are $v_1, \dots, v_n$, then

$$v_1 \wedge \cdots \wedge v_n = (\det A)e_1 \wedge \cdots \wedge e_n.$$

This single identity explains the main determinant rules:

- (i) multilinearity comes from multilinearity of the wedge product;
- (ii) repeated or dependent columns wedge to zero;
- (iii) swapping two columns changes the sign;
- (iv) multiplicativity follows from $\Lambda^n(AB) = (\Lambda^n A)(\Lambda^n B)$.

§65.8 Why permutation signs are forced

The sign of a permutation is not an arbitrary convention. It is the unique homomorphism

$$S_n \longrightarrow \{\pm 1\}$$

that sends every transposition to -1 . The inversion number gives a concrete formula:

$$\text{sgn}(\sigma) = (-1)^{\# \text{ inversions of } \sigma}.$$

Because the determinant is alternating, permuting the columns by σ must multiply the volume form by this sign. This is why the Leibniz formula has exactly the sign pattern it has.

§65.9 Minors as exterior coordinates

The same exterior-power viewpoint also explains minors. If $A : V \rightarrow W$, then

$$\Lambda^k A : \Lambda^k V \rightarrow \Lambda^k W$$

is represented, in the induced exterior bases, by the $k \times k$ minors of A . Therefore

$$\text{rank } A < k \iff \Lambda^k A = 0 \iff \text{all } k \times k \text{ minors vanish.}$$

Thus minors are not arbitrary subdeterminants; they are coordinates of an exterior-power map.

§65.10 Cramer's rule from volume ratios

Cramer's rule replaces one column of A by b . Geometrically, solving $Ax = b$ asks for the coordinates of b in the parallelepiped basis given by the columns of A . The coordinate x_i is the ratio of two oriented volumes:

$$x_i = \frac{\det A_i(b)}{\det A}.$$

This interpretation makes the formula less mysterious: a coordinate is measured by replacing the corresponding basis vector and comparing volumes.

Linear Maps, Matrix Representations, and Tensor Invariance

N-20230209

§66.1 From vector identities to linear representations

The previous vector-identity note explains the cross product, scalar triple product, exterior algebra, and Hodge star. This note uses that example as motivation, but its main subject is different: how coordinate arrays represent invariant linear or multilinear objects.

A linear map $T : V \rightarrow W$ is an object independent of coordinates. A matrix is what appears only after choosing a basis of V and a basis of W . Similarly, a tensor is not the displayed array of numbers; the array is a coordinate representation whose transformation law tells us what object is being represented.

Remark 66.1.1 (Division of labor between the two notes) — The note N-20230205 is about why three-dimensional vector identities are really statements about metric, orientation, wedge product, and Hodge star. The present note is about basis change, matrices, invariants, and tensor transformation laws. This separation prevents the same vector triple product discussion from being repeated twice.

§66.2 Coordinates as isomorphisms

A basis is not just a list of vectors; it is an isomorphism from an abstract vector space to a coordinate space. If $B = (b_1, \dots, b_n)$ is a basis of V , then

$$[\]_B : V \longrightarrow k^n$$

sends a vector to its coordinate column in that basis. The vector v is intrinsic; the column $[v]_B$ depends on B .

This distinction is the source of all basis-change formulas. A matrix is not the linear map itself. A matrix is the expression of a linear map after choosing coordinate isomorphisms.

Definition 66.2.1 (Matrix of a linear map). Let $T : V \rightarrow W$ be a linear map, let B be a basis of V , and let C be a basis of W . The matrix of T in these bases is the matrix $[T]_{C,B}$ satisfying

$$[T(v)]_C = [T]_{C,B}[v]_B$$

for all $v \in V$. Equivalently,

$$[T]_{C,B} = [\]_C \circ T \circ [\]_B^{-1}.$$

The last formula is often the cleanest one. It says that the bottom arrow in coordinate space is obtained by transporting the actual map T through coordinate isomorphisms.

§66.3 The basis-change formula

Suppose B, B' are two bases of V , and C, C' are two bases of W . Let

$$[v]_B = P[v]_{B'}, \quad [w]_C = Q[w]_{C'}.$$

Thus P converts new V -coordinates into old V -coordinates, and Q converts new W -coordinates into old W -coordinates.

Proposition 66.3.1 (Matrix transformation rule)

If

$$A = [T]_{C,B}, \quad A' = [T]_{C',B'},$$

then

$$A' = Q^{-1}AP.$$

Proof. Starting with new coordinates $[v]_{B'}$, first convert to old coordinates:

$$[v]_B = P[v]_{B'}.$$

Apply the old matrix A :

$$[T(v)]_C = A[v]_B = AP[v]_{B'}.$$

Finally convert the output from old C -coordinates to new C' -coordinates. Since $[w]_C = Q[w]_{C'}$, we have $[w]_{C'} = Q^{-1}[w]_C$. Hence

$$[T(v)]_{C'} = Q^{-1}AP[v]_{B'}.$$

Therefore $A' = Q^{-1}AP$. □

For an endomorphism $T : V \rightarrow V$, if the same basis change is used on domain and codomain, then $P = Q$ and

$$A' = P^{-1}AP.$$

This is similarity. Similar matrices are not two different operators; they are the same operator written in two coordinate systems.

§66.4 A worked basis-change example

Let $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ have standard-basis matrix

$$A = \begin{pmatrix} 2 & 1 \\ 0 & 3 \end{pmatrix}.$$

Use the new basis

$$b_1 = (1, 0), \quad b_2 = (1, 1).$$

The change-of-basis matrix whose columns are the new basis vectors in the old basis is

$$P = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

Then

$$A' = P^{-1}AP = \begin{pmatrix} 2 & 0 \\ 0 & 3 \end{pmatrix}.$$

The operator did not change. The new basis reveals that b_1 and b_2 are eigenvectors, so the matrix becomes diagonal.

This example also explains the meaning of diagonalization: it is not a magical simplification of numbers; it is the search for a coordinate system adapted to the invariant directions of the operator.

§66.5 Which quantities are invariant?

For an endomorphism, changing basis replaces A by $P^{-1}AP$. Any expression unchanged under this replacement belongs to the operator rather than to the coordinate matrix.

Proposition 66.5.1 (Similarity invariants)

If $A' = P^{-1}AP$, then

$$\operatorname{tr}(A') = \operatorname{tr}(A), \quad \det(A') = \det(A),$$

and A and A' have the same characteristic polynomial.

Proof. For trace,

$$\operatorname{tr}(P^{-1}AP) = \operatorname{tr}(APP^{-1}) = \operatorname{tr}(A),$$

using $\operatorname{tr}(XY) = \operatorname{tr}(YX)$. For determinant,

$$\det(P^{-1}AP) = \det(P^{-1}) \det(A) \det(P) = \det(A).$$

For the characteristic polynomial,

$$\det(\lambda I - P^{-1}AP) = \det(P^{-1}(\lambda I - A)P) = \det(\lambda I - A).$$

□

Rank is invariant for maps between different spaces as well, because left and right multiplication by invertible matrices do not change the dimension of the image. Trace and determinant, however, are naturally invariants of endomorphisms, not arbitrary maps $V \rightarrow W$.

§66.6 Covectors and the inverse-transpose rule

A covector is a linear functional $\alpha : V \rightarrow k$. It pairs with a vector to give a scalar:

$$\alpha(v) \in k.$$

The scalar must not depend on coordinates.

Assume again that

$$[v]_B = P[v]_{B'}.$$

Write a_B and $a_{B'}$ for the coordinate columns of α in the dual bases, so that

$$\alpha(v) = a_B^T [v]_B = a_{B'}^T [v]_{B'}.$$

Substituting $[v]_B = P[v]_{B'}$, we get

$$a_B^T P[v]_{B'} = a_{B'}^T [v]_{B'}$$

for all v , hence

$$a_{B'} = P^T a_B.$$

Equivalently, if one writes coordinate changes in the opposite direction, the familiar inverse-transpose appears. The important point is not the slogan “covectors transform oppositely”, but the invariant pairing

$$\alpha(v)$$

that forces the formula.

§66.7 Bilinear forms are not linear operators

A common source of confusion is that both a linear operator and a bilinear form may be represented by square matrices. Their basis-change laws are different.

Let $g : V \times V \rightarrow k$ be a bilinear form. In basis B , write

$$g(v, w) = [v]_B^T G [w]_B.$$

If $[v]_B = P[v]_{B'}$, then

$$g(v, w) = [v]_{B'}^T P^T G P [w]_{B'}.$$

Thus the matrix of the bilinear form in the new basis is

$$G' = P^T G P.$$

Compare this with the operator rule $A' = P^{-1} A P$. A matrix representing a bilinear form changes by congruence; a matrix representing an endomorphism changes by similarity.

Example 66.7.1 (The Euclidean metric in a non-orthonormal basis)

In the standard basis of $\mathbb{R}^2$, the Euclidean inner product has matrix I . For the basis

$$b_1 = (1, 0), \quad b_2 = (1, 1),$$

with

$$P = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix},$$

the metric matrix becomes

$$G' = P^T I P = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}.$$

The inner product is still the same geometric object, but because the new basis is not orthonormal, its coordinate matrix is no longer the identity.

This is why statements such as “ $A^T = A$ ” must be interpreted carefully. A symmetric matrix may represent a symmetric bilinear form in a chosen basis, or it may represent a self-adjoint operator relative to a chosen inner product. These are related but not identical ideas.

§66.8 Symmetric and skew parts

For an ordinary square matrix over a field of characteristic not 2, one can write

$$A = \frac{A + A^T}{2} + \frac{A - A^T}{2}.$$

The first part is symmetric, and the second part is skew-symmetric. But the coordinate-free meaning depends on what A represents.

If A represents a bilinear form, the symmetric part corresponds to

$$g_s(v, w) = \frac{g(v, w) + g(w, v)}{2},$$

and the skew part corresponds to

$$g_a(v, w) = \frac{g(v, w) - g(w, v)}{2}.$$

The skew part is an alternating 2-form.

If A represents a linear operator on an inner-product space, then self-adjointness is the condition

$$\langle Av, w \rangle = \langle v, Aw \rangle.$$

In an orthonormal basis this becomes $A^T = A$. In a non-orthonormal basis with metric matrix G , it becomes

$$A^T G = G A.$$

So the naive transpose test is valid only after the metric and basis have been fixed.

§66.9 Tensor slots and transformation laws

A tensor is best understood by its slots. The basic examples are:

object	tensor type
vector	V
covector	V^*
bilinear form	$V^* \otimes V^*$
linear operator	$V \otimes V^*$

The components transform according to which slots are vector slots and which slots are covector slots.

If $[v]_B = P[v]_{B'}$, then schematically:

object	components	basis-change law
vector	x	$x' = P^{-1}x$
covector	a	$a' = P^T a$
operator	A	$A' = P^{-1} A P$
bilinear form	G	$G' = P^T G P$

The formulas differ because the objects have different tensor types. Treating every square array as the same kind of object is the conceptual mistake tensor notation is designed to prevent.

§66.10 Tensor products and universal property

A bilinear map

$$B : V \times W \rightarrow U$$

is equivalently a linear map

$$\tilde{B} : V \otimes W \rightarrow U$$

satisfying

$$\tilde{B}(v \otimes w) = B(v, w).$$

This is not just notation. It says that the tensor product is the linear object that represents bilinear dependence.

For example, matrix multiplication can be viewed as contraction of repeated indices:

$$(AB)^i_j = A^i_k B^k_j.$$

The repeated index k is summed. The output has one upper and one lower index, so it is again an operator-type tensor.

Trace is another contraction:

$$\text{tr}(A) = A^i_i.$$

This explains why trace is coordinate-invariant for an endomorphism: it contracts one vector slot with one covector slot, producing a scalar.

§66.11 Coordinate-free reading of the original matrix facts

The scattered formulas involving A^{-1} , A^T , and A^{-T} should now be read as transformation laws:

- (i) $P^{-1}AP$ belongs to linear operators under basis change;
- (ii) P^TGP belongs to bilinear forms and quadratic forms;
- (iii) inverse-transpose behavior belongs to covectors and dual bases;
- (iv) trace is contraction of an operator-type tensor;
- (v) determinant is the induced scalar on the determinant line.

The mature lesson is not merely that “coordinates may change”. The real lesson is that each mathematical object has a type, and its type determines the correct transformation law. Tensor notation is useful because it keeps this type information visible.

A Brief Note on Tensors II: Tensor Products, Components, and More

N-20230217

§67.1 Tensor products and bases

The note continues the tensor discussion. If V and W are vector spaces with bases

$$e_1, \dots, e_m, \quad f_1, \dots, f_n,$$

then the tensor product $V \otimes W$ has basis

$$e_i \otimes f_j, \quad 1 \leq i \leq m, 1 \leq j \leq n.$$

Therefore

$$\dim(V \otimes W) = \dim V \cdot \dim W.$$

A general tensor in $V \otimes W$ can be written

$$T = \sum_{i,j} T^{ij} e_i \otimes f_j.$$

The page emphasizes that $e_i \otimes f_j$ should not be confused with ordered pairs. Tensor product is bilinear and universal for bilinear maps.

§67.2 Bilinearity

The tensor product satisfies

$$(u + v) \otimes w = u \otimes w + v \otimes w,$$

$$u \otimes (w + z) = u \otimes w + u \otimes z,$$

$$(cu) \otimes w = u \otimes (cw) = c(u \otimes w).$$

This is why a bilinear map

$$B : V \times W \rightarrow U$$

can be replaced by a linear map

$$\tilde{B} : V \otimes W \rightarrow U$$

with

$$\tilde{B}(v \otimes w) = B(v, w).$$

This universal property is the conceptual definition of tensor product.

§67.3 Metric, dot product, and components

The note then returns to geometry. Let g_{ij} be the metric components in a basis e_i :

$$g_{ij} = e_i \cdot e_j.$$

If

$$u = u^i e_i, \quad v = v^j e_j,$$

then

$$u \cdot v = g_{ij} u^i v^j$$

using Einstein summation convention.

In an orthonormal basis, $g_{ij} = \delta_{ij}$ and this reduces to the usual dot product. In a non-orthonormal basis, the metric matrix is needed.

§67.4 Reciprocal basis

Given a basis g_1, g_2, g_3 in Euclidean space, its reciprocal basis g^1, g^2, g^3 is defined by

$$g^i \cdot g_j = \delta_j^i.$$

The handwritten page computes reciprocal bases by solving linear systems or by using cross products. In three dimensions,

$$g^1 = \frac{g_2 \times g_3}{g_1 \cdot (g_2 \times g_3)},$$

with cyclic formulas for g^2 and g^3 .

The reciprocal basis lets one recover components:

$$v^i = g^i \cdot v, \quad v_i = g_i \cdot v.$$

This is the distinction between contravariant and covariant components.

§67.5 Covariant and contravariant components

The note stresses that vectors and components are not the same thing. A vector v is basis-independent, but its coordinate list changes when the basis changes.

Contravariant components v^i are the coefficients in

$$v = v^i g_i.$$

Covariant components are obtained by lowering an index with the metric:

$$v_i = g_{ij} v^j.$$

Conversely,

$$v^i = g^{ij} v_j,$$

where (g^{ij}) is the inverse metric matrix.

This is the concrete meaning of raising and lowering indices.

§67.6 Vector product and Levi-Civita symbol

The final page introduces the Levi-Civita symbol:

$$\varepsilon_{ijk} = \begin{cases} 1, & (i, j, k) \text{ is an even permutation of } (1, 2, 3), \\ -1, & (i, j, k) \text{ is an odd permutation of } (1, 2, 3), \\ 0, & \text{if two indices are equal.} \end{cases}$$

In an orthonormal right-handed basis,

$$(u \times v)_i = \varepsilon_{ijk} u_j v_k.$$

The note uses this notation to compress determinant and cross-product formulas.

§67.7 Scalar triple product

The scalar triple product is

$$u \cdot (v \times w).$$

It equals the determinant of the matrix with columns u, v, w , and it measures oriented volume. The reciprocal-basis formulas above are built from this quantity because the denominator

$$g_1 \cdot (g_2 \times g_3)$$

is the volume of the parallelepiped spanned by the basis vectors.

§67.8 What changes next

The note's main lesson is that tensor notation is coordinate bookkeeping with invariant meaning. Tensor products linearize bilinear operations. Metrics convert vectors and covectors. Reciprocal bases separate geometric vectors from their components. Levi-Civita symbols encode orientation. The formulas look index-heavy, but they are all ways of protecting geometric meaning under a change of basis.

§67.9 Tensor product by its universal property

The tensor product $V \otimes W$ is characterized by bilinear maps:

$$\text{Bil}(V \times W, U) \cong \text{Lin}(V \otimes W, U).$$

It is the linear space that turns bilinear dependence into linear dependence.

§67.10 Components under basis change

Vectors and covectors transform oppositely under basis change. This is why indices matter: upper indices follow the basis, lower indices follow the dual basis. Contractions pair one of each.

§67.11 Metric as an index-lowering map

An inner product gives an isomorphism

$$V \rightarrow V^*, \quad v \mapsto \langle v, - \rangle.$$

This allows one to lower indices. Without a metric, vectors and covectors should not be identified.

§67.12 Tensor notation as transformation law

Tensor notation records multilinearity and transformation behavior. Tensor products linearize bilinear maps, metrics identify vectors with covectors, and cross products are special metric/orientation constructions.

§67.13 Dual spaces and reciprocal bases

Given a basis g_1, g_2, g_3 , the reciprocal basis g^1, g^2, g^3 satisfies

$$g^i \cdot g_j = \delta_j^i.$$

This is a metric identification of vectors with covectors. Abstractly, the dual basis lives in V^* and is defined by

$$\gamma^i(g_j) = \delta_j^i.$$

The Euclidean dot product gives an isomorphism

$$V \rightarrow V^*, \quad v \mapsto (w \mapsto v \cdot w).$$

Under this isomorphism, the covectors γ^i correspond to the reciprocal vectors g^i .

Thus reciprocal bases are not a computational curiosity; they are dual bases transported through the metric.

§67.14 Musical isomorphisms

In differential geometry, a metric g identifies vector fields and covector fields. The maps are often denoted by musical symbols:

$$\flat : TM \rightarrow T^*M, \quad X^\flat(Y) = g(X, Y),$$

$$\sharp : T^*M \rightarrow TM.$$

Lowering an index is applying $\flat$; raising an index is applying $\sharp$.

The elementary formulas

$$v_i = g_{ij}v^j, \quad v^i = g^{ij}v_j$$

are exactly these operations in coordinates.

§67.15 Frames and moving coordinates

On a manifold, a basis of tangent vectors varying from point to point is a frame. Reciprocal frames live in the cotangent bundle. If e_i is a local frame and θ^i is the dual coframe, then

$$\theta^i(e_j) = \delta_j^i.$$

Differential forms are built from the coframe, while vector fields are built from the frame.

This generalizes the note's reciprocal basis from one vector space to geometry on curved spaces.

§67.16 Continuum mechanics notation

In continuum mechanics, covariant and contravariant bases appear when coordinates deform with a material body. The deformation gradient maps reference tangent vectors to current tangent vectors. Metrics and reciprocal bases are needed because curvilinear coordinate directions are not usually orthonormal.

The tensor notation in the original note is therefore not just abstract algebra. It is the language needed when coordinates themselves become part of the geometry.

§67.17 Metric tensors and Gram matrices

Given a non-orthonormal basis $g_1, \dots, g_n$, the metric tensor has components

$$g_{ij} = g_i \cdot g_j.$$

The inverse matrix (g^{ij}) satisfies

$$g^{ik} g_{kj} = \delta_j^i.$$

The reciprocal basis is

$$g^i = \sum_j g^{ij} g_j.$$

Then

$$g^i \cdot g_j = \delta_j^i.$$

This shows exactly how reciprocal bases are computed from the Gram matrix.

§67.18 Covariant and contravariant components from Gram matrices

A vector v may be written as

$$v = v^i g_i.$$

The coefficients v^i are contravariant components. The covariant components are

$$v_i = v \cdot g_i = g_{ij} v^j.$$

Raising and lowering indices uses the metric and its inverse:

$$v_i = g_{ij} v^j, \quad v^i = g^{ij} v_j.$$

The original reciprocal-basis computations are therefore a concrete finite-dimensional model for tensor index work in geometry and physics.

§67.19 Why tensor products linearize bilinear data

If $B : V \times W \rightarrow k$ is bilinear, it is inconvenient because it has two inputs. The tensor product turns it into a linear functional

$$\tilde{B} : V \otimes W \rightarrow k.$$

This is useful because linear maps are easier to compose, dualize, and represent by matrices. Many physical quantities, such as stress and inertia, are naturally bilinear before being encoded as tensors.

Notes on MML: Vectors, the Minus-One Trick, Linear Maps, Basis Change, Kernel, and Image

N-20251108

§68.1 Vectors beyond arrows

The first page stresses that vectors are not only geometric arrows. Polynomials can form a vector space: they can be added and scaled. Audio signals can also be treated as vectors: an audio signal is represented by a list of numbers, and adding or scaling those lists gives another signal.

The handwritten remark connects this to embeddings. In machine learning, word vectors or sentence vectors are not geometric arrows in the elementary sense, but they live in vector spaces and obey vector operations. Their meaning comes from the representation learned by the model, not from a literal physical arrow.

The lesson is that a vector is an element of a vector space. Geometry is only one model.

§68.2 The minus-one trick

The source discusses a practical trick for reading solutions of a homogeneous system

$$Ax = 0, \quad A \in \mathbb{R}^{k \times n}.$$

Assume A is in reduced row echelon form and has no zero rows. Pivot columns correspond to leading ones; free columns correspond to variables that can be chosen freely.

The minus-one trick augments the row-reduced matrix into an $n \times n$ matrix by inserting rows with a single -1 in the diagonal positions corresponding to free variables. Then the columns containing -1 on the diagonal directly give basis vectors for the null space.

The point is not magic. It is a compact way of doing the usual free-variable substitution. Setting one free variable to 1 and the others to 0 gives one kernel basis vector; the inserted -1 row records that choice.

§68.3 Example of the trick

The note's example starts with

$$A = \begin{bmatrix} 1 & 3 & 0 & 0 & 3 \\ 0 & 0 & 1 & 0 & 9 \\ 0 & 0 & 0 & 1 & -4 \end{bmatrix}.$$

The free variables are in columns 2 and 5. Augmenting gives

$$\tilde{A} = \begin{bmatrix} 1 & 3 & 0 & 0 & 3 \\ 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 9 \\ 0 & 0 & 0 & 1 & -4 \\ 0 & 0 & 0 & 0 & -1 \end{bmatrix}.$$

The columns with diagonal -1 give

$$x = \lambda_1 \begin{bmatrix} 3 \\ -1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \lambda_2 \begin{bmatrix} 3 \\ 0 \\ 9 \\ -4 \\ -1 \end{bmatrix}.$$

This is the same solution one obtains by solving by free variables, but the trick makes the basis easy to read.

§68.4 Linear mappings

A mapping $\Phi : V \rightarrow W$ between vector spaces is linear if

$$\Phi(\lambda x + \psi y) = \lambda \Phi(x) + \psi \Phi(y)$$

for all $x, y \in V$ and scalars λ, ψ .

The note lists special cases:

- an isomorphism is linear and bijective;
- an endomorphism is a linear map $V \rightarrow V$;
- an automorphism is a bijective endomorphism;
- the identity map $\text{id}_V : V \rightarrow V$ is the identity automorphism.

The handwritten annotations translate these as “same structure,” “same domain and codomain,” and “self-isomorphism.”

§68.5 Complex numbers as a real vector space

The example

$$\Phi : \mathbb{R}^2 \rightarrow \mathbb{C}, \quad \Phi(x_1, x_2) = x_1 + ix_2$$

is linear over $\mathbb{R}$:

$$\Phi(x + y) = \Phi(x) + \Phi(y), \quad \Phi(\lambda x) = \lambda \Phi(x).$$

It is also bijective, so it is an isomorphism of real vector spaces.

The note marks this as a homomorphism. The important distinction is that it preserves the vector-space addition and scalar multiplication. It is not claiming that every additional structure is automatically preserved.

§68.6 Finite-dimensional isomorphism theorem

The theorem quoted from Axler says that finite-dimensional vector spaces V and W are isomorphic iff

$$\dim V = \dim W.$$

This is one of the most important simplifications in finite-dimensional linear algebra: after choosing bases, all n -dimensional vector spaces over the same field look like $\mathbb{F}^n$.

But this isomorphism is not canonical unless a basis is chosen. The abstract vector space and its coordinate representation are related but not identical.

§68.7 Basis change

Let $\Phi : V \rightarrow W$ be linear. Suppose V has old basis

$$B = (b_1, \dots, b_n)$$

and new basis

$$\tilde{B} = (\tilde{b}_1, \dots, \tilde{b}_n).$$

Similarly, let W have old basis

$$C = (c_1, \dots, c_m)$$

and new basis

$$\tilde{C} = (\tilde{c}_1, \dots, \tilde{c}_m).$$

If A_Φ is the matrix of Φ with respect to B, C , and $\tilde{A}_\Phi$ is the matrix with respect to $\tilde{B}, \tilde{C}$, then

$$\tilde{A}_\Phi = T^{-1} A_\Phi S.$$

Here S converts coordinates from the new basis of V to the old basis of V , and T converts coordinates from the new basis of W to the old basis of W .

The diagram in the source page says exactly this: to use the old matrix, first translate the input coordinates into the old basis, then apply A_Φ , then translate the output coordinates into the new output basis.

§68.8 Equivalence and similarity

Two rectangular matrices $A, \tilde{A} \in \mathbb{R}^{m \times n}$ are equivalent if there exist invertible matrices $S \in \mathbb{R}^{n \times n}$ and $T \in \mathbb{R}^{m \times m}$ such that

$$\tilde{A} = T^{-1} A S.$$

This corresponds to changing the basis in both domain and codomain.

Two square matrices $A, \tilde{A} \in \mathbb{R}^{n \times n}$ are similar if

$$\tilde{A} = S^{-1} A S.$$

This corresponds to changing basis for a linear map $V \rightarrow V$ using the same coordinate change in domain and codomain.

The handwritten table emphasizes the difference:

Concept	Situation	Invariant meaning
Equivalence	linear map $V \rightarrow W$	change bases in domain and codomain
Similarity	linear map $V \rightarrow V$	same operator in a new basis

§68.9 Kernel and image

For a linear map $\Phi : V \rightarrow W$, the kernel is

$$\ker \Phi = \{v \in V : \Phi(v) = 0_W\},$$

and the image is

$$\operatorname{Im} \Phi = \{w \in W : \exists v \in V, \Phi(v) = w\}.$$

The kernel records what is collapsed to zero. The image records what outputs are reachable.

The source diagram shows the domain on the left, codomain on the right, a yellow region for the kernel, and a blue/yellow region for the image. The note also records the injectivity criterion:

$$\Phi \text{ is injective} \iff \ker \Phi = \{0_V\}.$$

§68.10 Categorical side remarks

The margin notes compare kernel, image, and cokernel with categorical language. In ordinary linear algebra,

$$\ker f \rightarrow A \rightarrow B$$

means the kernel is a subobject of the domain. The image is a subobject of the codomain. The cokernel, when defined, is a quotient-like object on the codomain side.

This is a useful preview: linear algebra already contains many category-theoretic patterns, even before they are named abstractly.

§68.11 The general pattern

This note links computational linear algebra with structural thinking. Row reduction gives the null space; the minus-one trick packages free-variable substitution. Linear maps become matrices only after choosing bases. Basis change explains why the same map can have different matrices. Kernel and image describe what a map destroys and what it reaches. These ideas are also the language behind machine-learning embeddings, where vectors need not be arrows but still live in vector spaces with meaningful linear structure.

§68.12 Vectors as elements of a linear structure

A vector is not essentially an arrow; it is an element of a vector space. Coordinates appear only after choosing a basis. The same vector has different coordinate columns in different bases.

§68.13 The minus-one trick as a kernel graph

The phrase “minus-one trick” should be understood carefully. In this note it is not primarily a homogeneous-coordinate trick and it is not an affine shift. It is a compact way of reading the kernel of a row-reduced homogeneous system.

Proposition 68.13.1 (Kernel graph form of an RREF system)

After possibly reordering variables, suppose a homogeneous system has row-reduced matrix

$$\begin{bmatrix} I_r & A \end{bmatrix}.$$

Write the unknown vector as

$$x = (x_p, x_f),$$

where $x_p \in k^r$ are pivot variables and $x_f \in k^{n-r}$ are free variables. Then the solution space is

$$\ker \begin{bmatrix} I_r & A \end{bmatrix} = \left\{ \begin{pmatrix} -Au \\ u \end{pmatrix} : u \in k^{n-r} \right\}.$$

Thus the kernel is the graph of the linear map

$$u \mapsto -Au.$$

Proof. The equation is

$$x_p + Ax_f = 0.$$

Hence $x_p = -Ax_f$. If we write $u = x_f$, every solution has the displayed form, and every displayed vector clearly satisfies the equation. $\square$

The inserted -1 rows in the practical trick are a bookkeeping device for the free variables. Choosing $u = e_j$, the j -th standard basis vector of the free-variable space, gives one kernel basis vector:

$$\begin{pmatrix} -Ae_j \\ e_j \end{pmatrix}.$$

So the trick is best read as a way of writing the graph basis of a kernel.

Remark 68.13.2 (Do not confuse two uses of extra coordinates) — Homogeneous coordinates do turn affine transformations into linear transformations one dimension higher. That is a real and useful idea. But the minus-one trick in this row-reduction context is different: it does not linearize an affine map; it packages the free-variable choices in a homogeneous linear system.

§68.14 Kernel-image factorization

For a linear map $T : V \rightarrow W$, the kernel records lost directions and the image records reachable directions. Rank-nullity says

$$\dim V = \dim \ker T + \dim \operatorname{im} T.$$

§68.15 Coordinates, kernels, and affine linearization

The review should distinguish vectors from coordinates, affine maps from linear maps, and a matrix from the linear transformation it represents. Kernel, image, and basis change are the structural core.

§68.16 Exact sequences behind rank-nullity

For a linear map $T : V \rightarrow W$, the kernel and image fit into the exact sequence

$$0 \rightarrow \ker T \rightarrow V \xrightarrow{T} \operatorname{im} T \rightarrow 0.$$

Exactness means that the image of each arrow equals the kernel of the next. Taking dimensions gives

$$\dim V = \dim \ker T + \dim \operatorname{im} T.$$

This is rank-nullity.

The elementary theorem is therefore the dimension shadow of an exact sequence. Exact sequences become one of the central languages of algebra, topology, and geometry.

§68.17 Quotient description of the image

The first isomorphism theorem says

$$V / \ker T \cong \operatorname{im} T.$$

The quotient collapses all vectors that differ by a null vector. What remains is precisely the information visible after applying T .

This gives a clean meaning to free variables and solution spaces. The kernel is what the map forgets; the image is what survives.

§68.18 Cokernel and obstructions

The cokernel of $T : V \rightarrow W$ is

$$\operatorname{coker} T = W / \operatorname{im} T.$$

It measures the failure of T to be onto. A system

$$Tx = b$$

is solvable exactly when b maps to zero in the cokernel.

Thus kernel measures non-uniqueness, while cokernel measures obstruction to existence. This is the conceptual form of the elementary distinction between infinitely many solutions and no solution.

§68.19 Returning to the MML basis-change exercises

The minus-one trick, basis conversion, and kernel/image computations in the note are all examples of one structure: a linear map forgets some directions, preserves some information, and may fail to reach some outputs. Exact sequences and quotients express this without depending on a chosen matrix.

§68.20 Row reduction as choosing a splitting

When a matrix is row-reduced, the pivot variables are expressed in terms of free variables. Abstractly, this chooses a splitting of the domain into a complement of the kernel plus the kernel:

$$V \cong U \oplus \ker T.$$

The map T is injective on the chosen complement U and maps U isomorphically onto $\operatorname{im} T$.

This splitting is not canonical: different row-reduction choices or different bases may choose different complements. The exact sequence

$$0 \rightarrow \ker T \rightarrow V \rightarrow \operatorname{im} T \rightarrow 0$$

is canonical; a particular parametric solution is a coordinate choice.

§68.21 Cokernel as the missing equation space

For $T : V \rightarrow W$, the cokernel $W / \operatorname{im} T$ measures which right-hand sides are not reachable. In matrix terms, left null vectors produce compatibility conditions. If $y \in \ker T^T$, then any solvable equation $Tx = b$ must satisfy

$$y^T b = 0.$$

Thus the left null space is the dual of the obstruction space.

This gives a precise interpretation of consistency conditions in linear systems: they are not extra tricks, but equations defining whether b vanishes in the cokernel.

Linear Algebra Review I: Determinants, Rank, Linear Systems, and More

N-20251110

§69.1 Cofactor orthogonality

For an $n \times n$ matrix

$$A = (a_{ij})_{n \times n},$$

let A_{ij} denote the cofactor of a_{ij} . The page starts with the standard cofactor identity:

$$\sum_{k=1}^n a_{ik} A_{jk} = \begin{cases} |A|, & i = j, \\ 0, & i \neq j, \end{cases}$$

and similarly

$$\sum_{k=1}^n a_{ki} A_{kj} = \begin{cases} |A|, & i = j, \\ 0, & i \neq j. \end{cases}$$

When $i = j$, this is exactly Laplace expansion along the i -th row or column. When $i \neq j$, the sum is the determinant of a matrix with two equal rows or two equal columns, hence it is zero. This is the conceptual reason the adjugate identity

$$A A^* = A^* A = |A| I$$

holds.

§69.2 Vandermonde determinant

The note then proves the Vandermonde determinant:

$$D_n(x_1, x_2, \dots, x_n) = \begin{vmatrix} 1 & 1 & \cdots & 1 \\ x_1 & x_2 & \cdots & x_n \\ x_1^2 & x_2^2 & \cdots & x_n^2 \\ \vdots & \vdots & & \vdots \\ x_1^{n-1} & x_2^{n-1} & \cdots & x_n^{n-1} \end{vmatrix} = \prod_{1 \leq i < j \leq n} (x_j - x_i).$$

The proof strategy shown on the page is the usual elimination proof. Starting from the last row, subtract suitable multiples of the previous row so that each column except the first acquires a factor $x_j - x_1$. After pulling these factors out, the remaining determinant is a Vandermonde determinant of order $n - 1$ in $x_2, \dots, x_n$. Thus

$$D_n = \prod_{j=2}^n (x_j - x_1) D_{n-1}(x_2, \dots, x_n),$$

and induction gives the formula.

The page emphasizes that this formula may be used directly later. The practical meaning is that a determinant built from powers of variables often hides a product of pairwise differences.

§69.3 Block triangular determinants

For a block matrix of the form

$$\begin{vmatrix} A & O \\ * & B \end{vmatrix},$$

where A is $n \times n$ and B is $m \times m$, the determinant is

$$\begin{vmatrix} A & O \\ * & B \end{vmatrix} = |A| |B|.$$

Similarly,

$$\begin{vmatrix} A & * \\ O & B \end{vmatrix} = |A| |B|.$$

But when the zero block is on the other diagonal position, one must account for row or column swaps:

$$\begin{vmatrix} O & A \\ B & * \end{vmatrix} = (-1)^{mn} |A| |B|, \quad \begin{vmatrix} * & A \\ B & O \end{vmatrix} = (-1)^{mn} |A| |B|.$$

The handwritten comment warns that one cannot simply write every block matrix determinant as $|A||B|$. The location of the zero block and the parity of the swap matter.

§69.4 General Laplace expansion

A more general Laplace expansion is recorded. Choose k rows with indices $i_1, \dots, i_k$. Then

$$|A| = \sum_{1 \leq j_1 < \dots < j_k \leq n} (-1)^{i_1 + \dots + i_k + j_1 + \dots + j_k} MN,$$

where M is the k -minor formed by the selected rows and columns $j_1, \dots, j_k$, and N is the complementary cofactor determinant obtained after deleting those rows and columns.

The page notes that the case $k = 1$ is the usual row or column expansion. The formula is useful when many entries in a group of rows or columns are zero.

§69.5 Determinant splitting

One green block is labelled as splitting a row or column. The principle is multilinearity. If one row is a sum of two rows, then the determinant is the sum of two determinants. For example,

$$R_i = u + v \implies D(\dots, R_i, \dots) = D(\dots, u, \dots) + D(\dots, v, \dots).$$

The same is true for columns.

The example on the page uses a tridiagonal-looking determinant and splits the first row into a simple row plus a row containing the parameter a . The resulting recurrence has the form

$$D_n = a^n + aD_{n-1},$$

and iteration gives a closed expression. The lesson is that splitting is useful when one part becomes triangular and the other part reduces the order.

§69.6 Adding rows or columns

Another determinant trick is to add multiples of one row or column to another. This does not change the determinant. The note uses this to transform matrices into Vandermonde-like form.

A typical pattern is

$$C_j \leftarrow C_j + C_1$$

or

$$R_i \leftarrow R_i + \lambda R_j.$$

If the operation produces a row or column of ones or zeros, the determinant becomes easier to expand.

The handwritten margin says that this method is especially useful for symmetric or nearly Vandermonde matrices: add all rows, add all columns, then expand along the newly simplified row or column.

§69.7 Augmenting to a Vandermonde determinant

One page expands a determinant by adding an extra column containing a new variable y . The determinant becomes a Vandermonde determinant in

$$x_1, x_2, \dots, x_n, y.$$

After evaluating the enlarged determinant as

$$\prod_{1 \leq i < j \leq n} (x_j - x_i) \prod_{k=1}^n (y - x_k),$$

one compares the coefficient of a suitable power of y . This coefficient comparison recovers the original determinant.

The important technique is coefficient comparison after embedding the target determinant into a known determinant. This is a common trick for determinants containing missing powers or sums of variables.

§69.8 Row and column transformations

The scanned pages collect several row/column patterns:

- (i) proportional differences, where adjacent columns are subtracted to create repeated columns or zeros;
- (ii) proportional additions, where columns are successively modified by adding multiples from one side;
- (iii) total row or column addition, where all columns are added to one column to reveal a common factor;
- (iv) paired addition, where the upper half of rows is added to the lower half and the right half of columns is subtracted from the left half.

For example, a determinant with rows

$$a^2, (a+1)^2, (a+2)^2, (a+3)^2$$

for several values of a can be reduced by subtracting adjacent columns. After repeated differences, low-degree polynomial dependence appears; eventually the determinant may vanish because the rows live in a space of too small dimension.

This is the general principle: if entries are values of a polynomial of degree less than the number of rows, finite differences can expose linear dependence.

§69.9 Triangularization by row operations

Several examples reduce a determinant to triangular form. One page shows a matrix whose first row is added to every other row; the determinant becomes upper triangular and the product of diagonal entries is obtained, for instance

$$D = n!.$$

Another example has a diagonal matrix plus a rank-one row of ones. If all diagonal entries a_i are nonzero, add suitable multiples of the other columns to the first column to get

$$D = \left(a_0 - \frac{1}{a_1} - \cdots - \frac{1}{a_n} \right) a_1 a_2 \cdots a_n.$$

If one of the a_i is zero, expansion along the corresponding row or column gives the exceptional formula. The source page explicitly separates the two cases.

§69.10 Recursive determinants

A determinant with ones on the diagonal or near the diagonal leads to a recurrence such as

$$D_n = D_{n-1} - D_{n-2} = -D_{n-3} = \cdots.$$

The page lists the resulting periodic cases depending on $n \bmod 3$. The key method is expansion along a sparse row or column after simplifying the matrix.

Another recursive example uses a tridiagonal determinant with diagonal $2a$ and off-diagonal a^2 and 1. Swapping rows or columns to make the matrix anti-diagonal creates a sign

$$(-1)^{n(n-1)/2},$$

which the handwritten note circles. This is a reminder that permutation signs are often the hidden danger in determinant shortcuts.

§69.11 Cramer's rule

The note then turns to linear systems. For the 2×2 system

$$\begin{cases} a_{11}x_1 + a_{12}x_2 = b_1, \\ a_{21}x_1 + a_{22}x_2 = b_2, \end{cases}$$

let

$$\Delta = \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix}, \quad \Delta_1 = \begin{vmatrix} b_1 & a_{12} \\ b_2 & a_{22} \end{vmatrix}, \quad \Delta_2 = \begin{vmatrix} a_{11} & b_1 \\ a_{21} & b_2 \end{vmatrix}.$$

If $\Delta \neq 0$, there is a unique solution

$$x_1 = \frac{\Delta_1}{\Delta}, \quad x_2 = \frac{\Delta_2}{\Delta}.$$

If $\Delta = 0$ but not all numerator determinants vanish, the system has no solution. If all determinants vanish, there may be infinitely many solutions.

For an $n \times n$ system $Ax = b$, if

$$D = |A| \neq 0,$$

then

$$x_j = \frac{D_j}{D},$$

where D_j is obtained by replacing the j -th column of A by the right-hand-side vector b .

§69.12 Rank and minors

The later pages review rank. The rank of a matrix is the largest order of a nonzero minor:

$$r(A) = \max\{k : \text{extsomekext-minor of } A \text{ is nonzero}\}.$$

Equivalently, it is the dimension of the column space or row space.

The note states the standard consequences:

$$r(A) = r(A^T),$$

row and column elementary transformations preserve rank, and for square A ,

$$|A| \neq 0 \iff r(A) = n.$$

If $r(A) < n$, the columns are linearly dependent and the determinant vanishes.

§69.13 Solvability of linear systems

For a linear system

$$Ax = b,$$

let $[A|b]$ be the augmented matrix. The solvability criterion is

$$Ax = b \text{ is solvable} \iff r(A) = r([A|b]).$$

If the common rank is r and the number of unknowns is n , then:

$$r = n \implies \text{extunique solution},$$

$$r < n \implies \text{extinfinitely many solutions}.$$

If

$$r(A) < r([A|b]),$$

then the system is inconsistent.

For the homogeneous system $Ax = 0$, the augmented rank is always the same as $r(A)$, so the system always has the zero solution. It has nonzero solutions iff

$$r(A) < n.$$

§69.14 Handwritten comments preserved

The note repeatedly warns against applying formulas mechanically. The important comments are:

- block determinant formulas require the zero block to be in the right place;
- determinant transformations are safe only if one tracks sign changes;
- row/column tricks are not independent recipes but attempts to create zeros, common factors, or known determinant patterns;
- Cramer's rule is clean but only useful when the coefficient determinant is nonzero;
- rank is the robust language for the singular cases where determinants vanish.

§69.15 The lesson behind it

This note is a determinant-method toolbox. Vandermonde handles power columns; block triangular form factors determinants; Laplace expansion exploits zeros; row and column transformations simplify structure; recursion evaluates patterned determinants; Cramer's rule solves nonsingular systems; rank handles singular systems. The conceptual progression is from explicit determinant computation to the structural question of linear dependence.

§69.16 Determinant as alternating multilinear invariant

The determinant is characterized by multilinearity, alternation, and normalization. This explains row operation rules, cofactor expansion, and why repeated rows give determinant zero.

§69.17 Vandermonde and evaluation maps

The Vandermonde determinant

$$\prod_{i < j} (x_j - x_i)$$

measures whether evaluation at distinct points gives independent conditions on polynomials. If two points coincide, two rows coincide and the determinant vanishes.

§69.18 Rank and linear systems

Rank measures the number of independent equations or independent output directions. A system is solvable exactly when the right-hand side lies in the column space.

§69.19 Determinants and rank as invariants

Determinant tricks, block forms, Vandermonde structures, and rank computations all track the same objects: volume scaling, independence, and solvability.

§69.20 Determinant through exterior powers

The determinant tricks in this note become more coherent through exterior algebra. A matrix $A : V \rightarrow V$ induces a map on the top exterior power:

$$\Lambda^n A : \Lambda^n V \rightarrow \Lambda^n V.$$

Since $\Lambda^n V$ is one-dimensional, this induced map is multiplication by a scalar. That scalar is $\det A$.

Laplace expansion, row operations, and block triangular formulas are all coordinate expressions of how A acts on oriented volume. For instance, if two columns become dependent, their wedge product is zero, so the determinant is zero.

§69.21 Schur complement

For a block matrix

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$$

with A invertible, block elimination gives

$$\det M = \det A \det(D - CA^{-1}B).$$

The matrix

$$D - CA^{-1}B$$

is the Schur complement of A in M .

This single formula explains many block determinant tricks. It also appears in Gaussian elimination, conditional covariance matrices, numerical linear algebra, and optimization.

§69.22 Matrix determinant lemma

A rank-one update has a particularly simple determinant formula:

$$\det(A + uv^T) = \det(A)(1 + v^T A^{-1}u)$$

when A is invertible. More generally,

$$\det(A + UV^T) = \det(A) \det(I + V^T A^{-1}U).$$

This is useful when a complicated determinant is a simple matrix plus a low-rank perturbation.

Many contest determinant problems secretly create low-rank updates by adding rows or columns. The advanced formula names the structure behind the trick.

§69.23 Rank varieties

The condition $\text{rank } A \leq r$ is algebraic: it is equivalent to the vanishing of all $(r+1) \times (r+1)$ minors. Therefore matrices of rank at most r form an algebraic variety in affine space.

For example, singular $n \times n$ matrices are the hypersurface

$$\det A = 0.$$

The elementary rank/minor criteria are thus equations defining geometric spaces. This is the doorway from linear algebra to algebraic geometry.

§69.24 Returning to Cramer and rank

Cramer's rule solves a nonsingular system by ratios of determinants. Rank criteria handle the singular cases. The advanced view says: the nonsingular matrices form the open set where the determinant volume form is nonzero; singular matrices lie on the determinantal hypersurface; rank strata describe how badly invertibility fails.

§69.25 Determinant identities as low-rank geometry

Many determinant problems are solved by adding all rows to one row, subtracting rows, or recognizing that a matrix has the form

$$D + uv^T.$$

The matrix determinant lemma says

$$\det(D + uv^T) = \det(D)(1 + v^T D^{-1}u)$$

when D is invertible. If D is diagonal, this turns a large determinant into one scalar correction term.

Geometrically, a rank-one update changes the image of the unit cube only in one coupled direction. Most independent volume scaling is still done by D ; the determinant lemma measures the one additional shear/stretch contribution.

§69.26 Schur complements and eliminating variables

Block Gaussian elimination is not merely a determinant shortcut. If

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} u \\ v \end{pmatrix}$$

and A is invertible, then the first block row gives $x = A^{-1}(u - By)$. Substituting into the second block row yields

$$(D - CA^{-1}B)y = v - CA^{-1}u.$$

The Schur complement is the effective system after eliminating x .

This interpretation connects determinant computation, solving linear systems, conditional covariance in probability, and constrained quadratic optimization.

§69.27 Determinantal varieties and singular loci

The rank condition $\text{rank } A \leq r$ defines a determinantal variety. Its equations are all $(r+1)$ -minors. The singular matrices form the hypersurface $\det A = 0$, and matrices of lower rank form deeper singular strata inside it.

Thus the elementary condition “determinant equals zero” is not merely a yes/no invertibility test. It defines a geometric boundary between invertible transformations and maps with collapsed directions. Cramer's rule works on the open complement of this hypersurface; rank methods describe what happens on the hypersurface.

Linear Algebra Review II:

Diagonalization, Inner Products, Quadratic Forms, and More

N-20251202

§70.1 Diagonalizable matrices

A square matrix A is diagonalizable if it is similar to a diagonal matrix:

$$P^{-1}AP = D.$$

Equivalently,

$$AP = PD.$$

If the columns of P are $\alpha_1, \dots, \alpha_n$ and

$$D = \text{diag}(\lambda_1, \dots, \lambda_n),$$

then

$$A\alpha_i = \lambda_i\alpha_i.$$

Thus the columns of the similarity matrix must be eigenvectors.

The highlighted theorem on the first page is:

$$A_n \text{ is diagonalizable} \iff A_n \text{ has } n \text{ linearly independent eigenvectors.}$$

The diagonal entries are the corresponding eigenvalues, in whatever order the eigenvectors are placed in P .

§70.2 Diagonalizability is independent from other properties

The first page compares diagonalizability with several related but independent properties.

- A matrix can be diagonalizable but not full-rank: the zero matrix is diagonal and hence diagonalizable, but its rank is zero.
- A matrix can be full-rank but not diagonalizable: a Jordan block such as

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

is invertible but has only one independent eigenvector.

- Repeated eigenvalues do not prevent diagonalization: the identity matrix has a repeated eigenvalue but is already diagonal.
- Nonzero eigenvalues do not guarantee diagonalization: the same Jordan block has nonzero eigenvalue 1 but is not diagonalizable.

The handwritten point is that determinant, rank, eigenvalue multiplicity, and diagonalizability are different kinds of information.

§70.3 Algebraic and geometric multiplicity

Let λ_0 be a k -fold root of the characteristic polynomial. The eigenspace is

$$E_{\lambda_0} = \ker(\lambda_0 I - A).$$

The number of linearly independent eigenvectors belonging to λ_0 is

$$\dim E_{\lambda_0} = n - r(\lambda_0 I - A).$$

The note states the inequality

$$\dim E_{\lambda_0} \leq k.$$

In words: geometric multiplicity is at most algebraic multiplicity.

A matrix is diagonalizable iff for every eigenvalue λ_i with algebraic multiplicity k_i ,

$$\dim E_{\lambda_i} = k_i,$$

or equivalently

$$r(\lambda_i I - A) = n - k_i.$$

§70.4 Procedure for diagonalization

The scanned page gives a concrete algorithm:

- (i) Solve the characteristic equation

$$|\lambda I - A| = 0$$

to find eigenvalues $\lambda_1, \dots, \lambda_m$ with multiplicities $s_1, \dots, s_m$.

- (ii) For each λ_i , solve

$$(\lambda_i I - A)x = 0$$

to obtain a basis of the eigenspace.

- (iii) If the number of independent eigenvectors for some λ_i is less than s_i , the matrix is not diagonalizable.

- (iv) If the total number is n , put all eigenvectors into P . Then

$$P^{-1}AP = \text{diag}(\lambda_1, \dots, \lambda_1, \lambda_2, \dots).$$

§70.5 Inner product and vector length

The next part reviews inner products. For real vectors

$$\alpha = (a_1, \dots, a_n)^T, \quad \beta = (b_1, \dots, b_n)^T,$$

the real inner product is

$$(\alpha, \beta) = a_1 b_1 + \dots + a_n b_n = \alpha^T \beta.$$

For complex vectors one uses the Hermitian inner product, with conjugation in one variable.

The length is

$$\|\alpha\| = \sqrt{(\alpha, \alpha)}.$$

The angle between nonzero real vectors is

$$\cos \theta = \frac{(\alpha, \beta)}{\|\alpha\| \|\beta\|}.$$

If

$$(\alpha, \beta) = 0,$$

then the vectors are orthogonal.

§70.6 Orthogonal and orthonormal systems

A vector group is orthogonal if every two distinct nonzero vectors are orthogonal. It is orthonormal if, in addition, every vector has norm 1.

The note states that a nonzero orthogonal vector group is linearly independent. The proof is short: if

$$c_1\alpha_1 + \cdots + c_k\alpha_k = 0,$$

take inner product with α_i . All cross terms vanish, so

$$c_i\|\alpha_i\|^2 = 0,$$

and hence $c_i = 0$.

§70.7 Schmidt orthogonalization

Given a linearly independent set

$$\alpha_1, \alpha_2, \dots, \alpha_n,$$

Schmidt orthogonalization constructs an equivalent orthogonal set

$$\xi_1, \xi_2, \dots, \xi_n.$$

The construction is

$$\begin{aligned}\xi_1 &= \alpha_1, \\ \xi_2 &= \alpha_2 - \frac{(\alpha_2, \xi_1)}{(\xi_1, \xi_1)}\xi_1,\end{aligned}$$

and in general

$$\xi_i = \alpha_i - \sum_{k=1}^{i-1} \frac{(\alpha_i, \xi_k)}{(\xi_k, \xi_k)}\xi_k.$$

Then normalize:

$$\beta_i = \frac{\xi_i}{\|\xi_i\|}.$$

The scanned proof checks two things: $\xi_i \neq 0$ and $(\xi_i, \xi_k) = 0$ for $k < i$. If $\xi_i = 0$, then α_i would be a linear combination of the previous vectors, contradicting independence.

§70.8 Example of Schmidt orthogonalization

The example orthogonalizes three vectors in $\mathbb{R}^4$:

$$\alpha_1 = (1, 0, 0, 1)^T, \quad \alpha_2 = (1, 1, 0, 1)^T, \quad \alpha_3 = (1, 1, 1, 0)^T.$$

The result after orthogonalization and normalization is displayed as an orthonormal group

$$\begin{aligned} \beta_1 &= \frac{1}{\sqrt{2}}(1, 0, 0, 1)^T, \\ \beta_2 &= (0, 1, 0, 0)^T, \\ \beta_3 &= \frac{1}{\sqrt{6}}(1, 0, 2, -1)^T. \end{aligned}$$

The point of the example is not the arithmetic alone; it shows how each new vector is stripped of its projections onto the earlier directions.

§70.9 QR decomposition

The note presents the matrix form of Schmidt orthogonalization. If

$$A = (\alpha_1, \dots, \alpha_n) \in \mathbb{R}^{m \times n}, \quad r(A) = n,$$

then

$$A = QR,$$

where

$$Q^T Q = I,$$

and R is an upper triangular invertible matrix.

The columns of Q are the orthonormal vectors produced by Schmidt orthogonalization. The matrix R records the coefficients expressing the original columns in the orthonormal basis.

§70.10 Orthogonal matrices

A real square matrix A is orthogonal if

$$A^T A = I.$$

Equivalent conditions are:

$$A^{-1} = A^T, \quad A A^T = I,$$

the columns of A form an orthonormal basis, and the rows of A form an orthonormal basis.

If A and B are orthogonal matrices of the same order, then AB is also orthogonal:

$$(AB)^T (AB) = B^T A^T A B = B^T I B = I.$$

If A is orthogonal, then

$$|A|^2 = 1,$$

so $|A| = \pm 1$. Orthogonal matrices preserve lengths:

$$\|A\alpha\| = \|\alpha\|.$$

Therefore their eigenvalues, when real, have absolute value 1.

§70.11 Real symmetric matrices

For a real symmetric matrix $A = A^T$, the note records:

- (i) all eigenvalues are real;
- (ii) eigenvectors corresponding to distinct eigenvalues are orthogonal;
- (iii) there exists an orthogonal matrix P such that

$$P^T A P = D$$

is diagonal.

The proof that eigenvalues are real uses

$$A\xi = \lambda\xi$$

and the scalar $\bar{\xi}^T A\xi$. Symmetry forces it to equal its conjugate, hence $\lambda = \bar{\lambda}$.

For orthogonality, if

$$A\xi_1 = \lambda_1\xi_1, \quad A\xi_2 = \lambda_2\xi_2,$$

then symmetry gives

$$\lambda_2\xi_1^T\xi_2 = \xi_1^T A\xi_2 = (A\xi_1)^T\xi_2 = \lambda_1\xi_1^T\xi_2.$$

If $\lambda_1 \neq \lambda_2$, then $\xi_1^T\xi_2 = 0$.

§70.12 Quadratic forms

A quadratic form in variables $x_1, \dots, x_n$ is

$$f(x) = x^T A x,$$

where A is symmetric. By an orthogonal change of variables

$$x = Py, \quad P^T P = I,$$

one obtains

$$f = x^T A x = y^T (P^T A P) y = \lambda_1 y_1^2 + \dots + \lambda_n y_n^2.$$

This is the principal-axis form of the quadratic form.

The note distinguishes ordinary diagonalization from orthogonal diagonalization. For real symmetric matrices, the diagonalizing matrix can be chosen orthogonal, which preserves lengths and angles. This is why quadratic forms can be geometrically classified by rotations.

§70.13 Positive definite quadratic forms

A quadratic form is positive definite if

$$x^T A x > 0 \quad (x \neq 0).$$

For a real symmetric matrix, this is equivalent to all eigenvalues being positive. Another criterion is Sylvester's criterion: all leading principal minors are positive:

$$\Delta_1 > 0, \quad \Delta_2 > 0, \quad \dots, \quad \Delta_n > 0.$$

Negative definiteness and semidefiniteness are treated similarly, with signs or non-strict inequalities.

The page also gives the standard rank and signature language: after a nonsingular linear substitution, a real quadratic form can be reduced to

$$z_1^2 + \cdots + z_p^2 - z_{p+1}^2 - \cdots - z_r^2.$$

The numbers of positive and negative squares are invariants.

§70.14 Conic and quadric interpretations

The later pages connect quadratic forms to second-degree equations. A conic can be written in matrix form as

$$ax^2 + 2bxy + cy^2 + dx + ey + f = 0.$$

The quadratic part

$$\begin{pmatrix} x & y \end{pmatrix} \begin{pmatrix} a & b \\ b & c \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}$$

can be diagonalized by an orthogonal transformation. Translation then removes linear terms when possible. This explains why conics are classified by eigenvalues of the quadratic part and by remaining constants.

The scanned diagrams show arrows from original coordinates to rotated and translated coordinate systems. The mathematical point is that cross terms disappear after choosing principal axes.

§70.15 Common traps recorded in the notes

The handwritten comments emphasize several traps:

- diagonalizable does not mean full-rank;
- repeated eigenvalues do not automatically prevent diagonalization;
- a matrix with real eigenvalues need not be diagonalizable unless enough eigenvectors exist;
- for real symmetric matrices, orthogonal diagonalization is stronger and more geometric;
- Schmidt orthogonalization must subtract projections onto all previous orthogonal vectors, not only the immediately previous one;
- positive definiteness is a statement about all nonzero vectors, not only about diagonal entries.

§70.16 A closing lens

This note links eigenvalue theory and geometry. Diagonalization decomposes a linear map into independent eigen-directions. Inner products measure lengths and angles. Schmidt orthogonalization turns independence into orthogonality. Orthogonal matrices preserve geometry. Real symmetric matrices are the ideal case: they have real eigenvalues, orthogonal eigenspaces, and orthogonal diagonalization. Quadratic forms use this structure to classify conics, extrema, and definiteness.

§70.17 Diagonalization as decomposition into invariant lines

A matrix is diagonalizable when the space decomposes as a direct sum of eigenspaces. In that basis, the operator acts by independent scalings.

§70.18 Algebraic versus geometric multiplicity

The algebraic multiplicity counts root multiplicity in the characteristic polynomial; geometric multiplicity counts eigenvector directions. Defect occurs when there are not enough independent eigenvectors.

§70.19 Inner products and orthogonal projection

An inner product adds metric geometry to linear algebra. Orthogonal projection is the best approximation in a subspace, and Gram–Schmidt constructs orthonormal coordinates adapted to that metric.

§70.20 Quadratic forms and spectral theorem

Real symmetric matrices define quadratic forms and are orthogonally diagonalizable. Positive definiteness is positivity of all eigenvalues.

§70.21 Spectral decomposition and quadratic geometry

Diagonalization, inner products, orthogonality, and quadratic forms are one story: decompose the space into independent directions, preferably orthogonal ones when a metric is present.

§70.22 Spectral theorem and orthogonal bases

Theorem 70.22.1 (Spectral theorem, finite-dimensional real form)

Every real symmetric matrix is orthogonally diagonalizable. Equivalently, it has an orthonormal basis of eigenvectors and real eigenvalues.

The note's real symmetric matrix theorem is the finite-dimensional spectral theorem. If $A = A^T$, then there exists an orthogonal matrix Q such that

$$Q^T A Q = \text{diag}(\lambda_1, \dots, \lambda_n).$$

Equivalently, $\mathbb{R}^n$ has an orthonormal basis of eigenvectors of A .

This theorem is deeper than ordinary diagonalization because the basis can be chosen orthonormal. Orthogonality preserves lengths and angles, so quadratic forms can be classified geometrically rather than only algebraically.

§70.23 QR as algorithmic Gram–Schmidt

Gram–Schmidt turns independent vectors into an orthonormal set. Written in matrix form, it gives QR decomposition:

$$A = QR, \quad Q^T Q = I, \quad R \text{ upper triangular.}$$

In numerical linear algebra, QR is used to solve least-squares problems and compute eigenvalues. Modified Gram–Schmidt and Householder reflections are introduced because the naive formula can be numerically unstable.

Thus the elementary orthogonalization exercise is not just a hand computation. It is the conceptual beginning of stable matrix algorithms.

§70.24 Sylvester’s law of inertia

A real quadratic form can be diagonalized by congruence:

$$x^T A x = y^T (P^T A P) y.$$

Sylvester’s law of inertia says that the numbers of positive, negative, and zero squares in a diagonal form are invariant under invertible linear changes of variables.

Therefore a quadratic form has an intrinsic signature

$$(n_+, n_-, n_0).$$

Positive definiteness means signature $(n, 0, 0)$. This explains why completing squares and diagonalizing symmetric matrices lead to the same classification.

§70.25 Hilbert-space projection

In an inner product space, projection onto a closed subspace is characterized by orthogonality of the error. In finite dimensions this is ordinary projection. In Hilbert spaces it becomes the projection theorem: if M is a closed subspace of a Hilbert space H , then every $x \in H$ decomposes uniquely as

$$x = m + m^\perp, \quad m \in M, \quad m^\perp \in M^\perp.$$

Least squares, Fourier series, and orthogonal diagonalization all rely on this same projection geometry.

§70.26 Orthogonal diagonalization and axes

For a real symmetric matrix A , the quadratic form $q(x) = x^T A x$ defines a family of level sets. Orthogonal diagonalization chooses axes along which the cross terms vanish:

$$q(x) = \lambda_1 y_1^2 + \cdots + \lambda_n y_n^2.$$

The eigenvectors are the principal axes, and the eigenvalues are the curvatures or scaling factors along them.

This interpretation unifies conic classification, second-derivative tests, inertia of quadratic forms, and PCA. In PCA, the covariance matrix is symmetric and positive semidefinite. Its eigenvectors give directions of maximal variance.

§70.27 Gram–Schmidt as factorization of coordinates

Suppose the columns of A are independent vectors $a_1, \dots, a_n$. Gram–Schmidt constructs orthonormal vectors $q_1, \dots, q_n$ and upper-triangular coefficients r_{ij} such that

$$a_j = \sum_{i \leq j} r_{ij} q_i.$$

This is exactly

$$A = QR.$$

The diagonal entries r_{jj} are the lengths of the components remaining after projecting away earlier directions.

This makes the connection to least squares immediate. Solving $Ax \approx b$ becomes solving

$$Rx = Q^T b,$$

which is stable because Q preserves lengths.

§70.28 Spectral theorem beyond finite dimension

Remark 70.28.1 — In infinite-dimensional Hilbert spaces, the spectral theorem is no longer only a statement about eigenvectors; it is formulated through projection-valued measures or functional calculus for self-adjoint operators.

The finite-dimensional spectral theorem has infinite-dimensional descendants. A bounded self-adjoint operator on a Hilbert space has a spectral measure. This lets one define functions of the operator by integration over its spectrum. Differential operators such as $-d^2/dx^2$ are unbounded. With suitable domains, however, they are self-adjoint and have eigenfunction expansions.

Thus diagonalizing a symmetric matrix is the first finite version of decomposing an operator into modes.

Linear Algebra Review III: Determinant Tricks, Rank, Matrix Powers, and More

N-20260103

§71.1 Determinant expansion and cofactors

The first group of problems uses determinant expansion. If A_{ij} denotes the cofactor of a_{ij} , then

$$|A| = \sum_{j=1}^n a_{ij} A_{ij}$$

for expansion along the i -th row. A useful extension is the cofactor orthogonality relation:

$$\sum_{j=1}^n b_j A_{ij}$$

is the determinant obtained by replacing the i -th row of A by the row $(b_1, \dots, b_n)$.

The first example asks for a sum such as

$$M_{41} + M_{42} + M_{43}$$

from a given 4×4 determinant. The note's hint says to regard it as replacing the fourth row by a row such as $(-1, 1, -1, 0)$. This is the key move: a linear combination of cofactors is usually not computed one cofactor at a time; it is reassembled into one determinant.

A second example gives entries in one row and cofactors in another row. The identity

$$\sum_j a_{ij} A_{kj} = 0, \quad i \neq k,$$

allows the unknown entry to be solved quickly. This is again not a calculation trick but a structural fact: the determinant with two equal rows is zero.

§71.2 Recursive determinants

One determinant has diagonal entries $a+b$, adjacent entries ab and 1, and zeros elsewhere. Expanding along the first column or splitting the first column gives a recurrence of the form

$$D_n = (a+b)D_{n-1} - abD_{n-2}.$$

The characteristic equation is

$$x^2 - (a+b)x + ab = 0,$$

with roots a and b . Therefore the closed form is of the type

$$D_n = a^n + b^n$$

up to the indexing convention used in the original determinant.

The page's point is that determinant recurrences are often solved exactly like linear recurrences. After expansion, the determinant order becomes the sequence index.

§71.3 Circulant-style determinant

The next determinant is a 4×4 matrix built from cyclic shifts of

$$a, b, c, d.$$

The hint uses a block form

$$\begin{pmatrix} C & D \\ D & C \end{pmatrix}$$

and the block transformation

$$\begin{vmatrix} C & D \\ D & C \end{vmatrix} = |C + D| |C - D|.$$

The resulting eigenvalue-like factors are

$$a + c \pm \sqrt{(a + b)(c + d)}, \quad a - c \pm \sqrt{(b - a)(d - b)}$$

with the exact signs depending on the row ordering.

The important idea is that cyclic or block-symmetric structure should be attacked by adding and subtracting matching blocks. This diagonalizes the symmetry.

§71.4 Cancellation through orthogonal matrices

One problem states: if A and B are same-order orthogonal matrices and

$$|A| + |B| = 0,$$

prove

$$|A + B| = 0.$$

Since A and B are orthogonal,

$$|A|, |B| \in \{1, -1\}.$$

The hint uses

$$A^T(A + B)B^T = A^T B^T + E.$$

A simpler determinant argument is

$$|A^T(A + B)B^T| = |A^T| |A + B| |B^T|.$$

The transformed matrix has determinant proportional to $|A + B|$, and the condition $|A| + |B| = 0$ forces singularity. The source proof is a typical example of using invertible multiplication to preserve the zero/nonzero status of a determinant.

§71.5 Determinant product and block simplification

A problem with column vectors

$$A = (\alpha_1, \alpha_2, \alpha_3), \quad B = (\alpha_3 - 2\alpha_1, 3\alpha_2, \alpha_1)$$

uses the factorization

$$B = AC$$

for a concrete matrix C . Then

$$|B| = |A| |C|.$$

If $|B| = 6$, one computes $|A|$ from $|C|$, and then

$$|A + B| = |A(I + C)| = |A| |I + C|.$$

The handwritten solution does exactly this: once the matrix relation is seen, the determinant becomes a product problem rather than an expansion problem.

The page also includes a block determinant simplification: for suitable rectangular blocks,

$$\begin{vmatrix} A & O \\ E_n & B^T \end{vmatrix} = 0$$

when the block triangular formula leaves a zero determinant factor. The intended memory is the block triangular determinant rule and the fact that rectangular dimensions can force a zero block determinant.

§71.6 Schur-complement style block transformation

For square blocks A, B, C, D with A invertible and $AC = CA$, the note proves

$$\begin{vmatrix} A & B \\ C & D \end{vmatrix} = |AD - CB|.$$

The hint multiplies by the block matrix

$$\begin{pmatrix} E & O \\ -CA^{-1} & E \end{pmatrix}$$

to eliminate the lower-left block:

$$\begin{pmatrix} E & O \\ -CA^{-1} & E \end{pmatrix} \begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} A & B \\ O & D - CA^{-1}B \end{pmatrix}.$$

Thus

$$\begin{vmatrix} A & B \\ C & D \end{vmatrix} = |A| |D - CA^{-1}B| = |AD - CB|,$$

where the commutation condition is needed to pass from $A(D - CA^{-1}B)$ to $AD - CB$.

§71.7 Rank-two decomposition

The vector-and-matrix section begins with a standard-form/rank decomposition problem. If

$$A \in \mathbb{R}^{m \times n}, \quad r(A) = 2,$$

then there exist column vectors

$$\alpha_1, \alpha_2 \in \mathbb{R}^m, \quad \beta_1, \beta_2 \in \mathbb{R}^n$$

such that

$$A = \alpha_1 \beta_1^T + \alpha_2 \beta_2^T.$$

The hint is the rank normal form

$$A = P \operatorname{diag}(E_2, O) Q.$$

Writing the first two columns of P as α_1, α_2 and the first two rows of Q as β_1^T, β_2^T gives the result.

This is the finite-dimensional version of saying that a rank- r matrix is a sum of r rank-one matrices.

§71.8 Adjugate identities

The pages list the key adjugate formulas:

$$\begin{aligned} AA^* &= A^*A = |A|E, \\ (kA)^{-1} &= k^{-1}A^{-1}, \quad (AB)^{-1} = B^{-1}A^{-1}, \\ |A^*| &= |A|^{n-1}. \end{aligned}$$

One problem gives A^* and asks for A . From

$$|A^*| = |A|^{n-1}$$

one first finds $|A|$, then uses

$$(A^*)^{-1} = \frac{1}{|A|}A.$$

Thus

$$A = |A|(A^*)^{-1}.$$

The handwritten solution tracks the possible sign of $|A|$ and then writes the final matrix.

§71.9 Skew-symmetric matrices via quadratic forms

A problem states: if for all $\xi \in \mathbb{R}^n$,

$$\xi^T A \xi = 0,$$

prove that A is skew-symmetric. The note's hint says to take $\xi = e_i$ and $\xi = e_i + e_j$.

Taking $\xi = e_i$ gives

$$a_{ii} = 0.$$

Taking $\xi = e_i + e_j$ gives

$$a_{ii} + a_{ij} + a_{ji} + a_{jj} = 0.$$

Since the diagonal entries are zero,

$$a_{ij} + a_{ji} = 0.$$

Hence

$$A^T = -A.$$

§71.10 Absorption and cancellation

If

$$A^2 = A, \quad B^2 = B, \quad (A + B)^2 = A + B,$$

then expanding gives

$$AB + BA = 0.$$

Multiplying on the left by A gives

$$AAB + ABA = 0 \implies AB + ABA = 0.$$

Multiplying on the right by A gives

$$ABA + BA = 0.$$

Combining these relations yields

$$AB = BA = 0.$$

The source labels this as using absorption and cancellation: idempotent matrices behave like projections, and the condition that their sum is again idempotent forces their interaction to vanish.

§71.11 Rank identities

The rank block on the page records:

$$\begin{aligned} 0 \leq r(A) &\leq \min\{m, n\}, \\ r(A^T) &= r(A), \quad r(kA) = r(A) \quad (k \neq 0), \\ r(A) + r(B) - n &\leq r(AB) \leq \min\{r(A), r(B)\}, \\ r(A + B) &\leq r(A) + r(B). \end{aligned}$$

It also records that elementary transformations preserve rank and that

$$|A| \neq 0 \iff A \text{ is full-rank}.$$

Two examples use these identities. If

$$A^2 = -A,$$

then

$$A(E + A) = 0$$

and the rank inequality gives

$$r(A) + r(E + A) = n.$$

Another example with $A = MN^T$ and $|N^T M| \neq 0$ proves

$$r(A^2) = r(A)$$

by bounding ranks from both sides and using the invertibility of $N^T M$.

§71.12 Commutation from factorization

If

$$BA = A + 2B,$$

then

$$(B - E)(A - 2E) = 2E.$$

Hence $B - E$ is invertible and

$$(B - E)^{-1} = \frac{1}{2}(A - 2E).$$

Since a matrix commutes with its inverse and with polynomials in itself, one obtains the commutation needed to show

$$AB = BA.$$

The trick is to factor the equation into an invertible product.

§71.13 Projection-like matrices from bases

A problem defines

$$A = (\alpha_1, \dots, \alpha_n), \quad B = A^{-1}, \quad B^T = (\beta_1, \dots, \beta_n),$$

and

$$C = \alpha_1 \beta_1^T + \dots + \alpha_k \beta_k^T.$$

The hint writes

$$C = (\alpha_1, \dots, \alpha_k)(\beta_1, \dots, \beta_k)^T = A \operatorname{diag}(E_k, O) A^{-1}.$$

Therefore

$$C^2 = C.$$

The equation $Cx = 0$ is then equivalent to killing the first k coordinates in the A -basis. The solution space is spanned by the remaining basis vectors.

§71.14 Nilpotent polynomial inverse

If

$$A^k = O,$$

then

$$(E - A)(E + A + A^2 + \dots + A^{k-1}) = E - A^k = E.$$

Therefore $E - A$ is invertible and

$$(E - A)^{-1} = E + A + \dots + A^{k-1}.$$

This page labels the method as using polynomial formulas. It is the finite geometric series identity for matrices.

§71.15 Equivalent vector groups

One problem gives two independent vector groups

$$\alpha_1, \alpha_2, \alpha_3, \quad \beta_1, \beta_2,$$

with all cross inner products zero:

$$(\alpha_i, \beta_j) = 0.$$

It asks to prove that the combined group is linearly independent. If

$$\alpha + \beta = 0,$$

where α is in the span of the α_i and β is in the span of the β_j , then substituting $\beta = -\alpha$ gives

$$(\alpha, \alpha) = 0.$$

By positive definiteness of the inner product, $\alpha = 0$, and then $\beta = 0$. Hence all coefficients vanish.

Another problem asks for a vector γ so that two vector groups become equivalent. The note's method is to express each known vector in terms of the other group and then add one suitable missing vector so both spans coincide.

§71.16 Computing matrix powers

The note gives several methods for A^n .

For

$$A = E + C, \quad C^3 = 0,$$

use the binomial formula:

$$A^n = (E + C)^n = E + nC + \frac{n(n-1)}{2}C^2.$$

For a rank-one matrix

$$B = \alpha\beta^T,$$

we have

$$B^n = (\beta^T \alpha)^{n-1} \alpha \beta^T.$$

This appears in the source hint as

$$(\alpha\beta^T)^n = \alpha(\beta^T \alpha)^{n-1} \beta^T.$$

A permutation/shear example computes

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix}^{100} A \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}^{99}.$$

The left factor adds 100 times the first row to the third row, while the right factor permutes columns. The source solution explicitly rewrites the effect as row and column operations.

For a stochastic-like matrix

$$A = \begin{pmatrix} 1-a & a \\ b & 1-b \end{pmatrix}, \quad 0 < a, b < 1,$$

find eigenvalues

$$1, \quad 1 - a - b.$$

Since $|1 - a - b| < 1$, the second eigenvalue tends to zero in powers. Thus A^n tends to the projection onto the eigenline for eigenvalue 1.

§71.17 Homogeneous systems and adjugate

If

$$|A| = 0, \quad A^* \neq O,$$

then

$$AA^* = 0.$$

Any nonzero column of A^* is a nonzero solution of

$$Ax = 0.$$

The note also observes that $A^* \neq 0$ and $|A| = 0$ imply

$$r(A) = n - 1,$$

so the null space is one-dimensional. Thus the general solution is

$$x = k\xi,$$

where ξ is any nonzero column of A^* .

§71.18 Null space basis and augmented rank

If $A \in \mathbb{R}^{m \times n}$ has rank $r < n$, and N is the matrix whose columns form a basis of the solution space of $Ax = 0$, then the note proves

$$r(A^T, N) = n.$$

The idea is to show that the only solution of

$$Ax = 0, \quad N^T x = 0$$

is $x = 0$. Since $x \in \ker A$, write $x = Nc$. Then

$$N^T Nc = 0.$$

Because the columns of N are independent, $N^T N$ is invertible, so $c = 0$.

§71.19 Rank of products and equality of solution spaces

The page uses the dimension formula

$$r(A) + \dim N(A) = n.$$

If

$$r(AB) = r(B),$$

then the systems

$$ABx = 0, \quad Bx = 0$$

have the same nullity and one null space is contained in the other; hence their solution spaces are equal. This is then used to show

$$r(ABC) = r(BC).$$

Similarly,

$$r(AA^T) = r(A)$$

is proved by comparing the equations

$$AA^T x = 0$$

and

$$A^T x = 0.$$

Indeed,

$$x^T AA^T x = (A^T x)^T (A^T x) = 0$$

forces $A^T x = 0$.

§71.20 General solution from known solutions

One problem gives a nonhomogeneous system $Ax = b$ with a known solution and additional vectors satisfying relations such as

$$\eta_1 + \eta_2 = (2, 3, 4, 5)^T, \quad \eta_2 + \eta_3 = (1, 2, 3, 4)^T.$$

The note uses the rule:

$$\text{extgeneral solution} = \text{extoneparticular solution} + \text{extgeneral solution of } Ax = 0.$$

If the rank is 3 in four unknowns, the homogeneous solution space is one-dimensional, so the final answer has one free parameter.

Another problem combines bases of $Ax = 0$ and $Bx = 0$ to prove that the block system

$$\begin{pmatrix} A \\ B \end{pmatrix} x = 0$$

has a nonzero solution. This is essentially an intersection-of-null-spaces question.

§71.21 Equivalence of row vector groups

The last page proves a criterion of the form:

$$\begin{aligned} Ax = 0, Bx = 0 \text{ have the same solutions} \\ \iff \text{the row groups of } A, B \text{ are equivalent.} \end{aligned}$$

The rank condition behind it is

$$r(A) = r(A, B) = r(B)$$

for the stacked row matrix. If the solution spaces are equal, the orthogonal complements of those solution spaces, namely the row spaces, are equal. Conversely, equal row spaces clearly define the same homogeneous equations.

§71.22 Diagonal dominance and invertibility

One problem states: if

$$|a_{ii}| > \sum_{j \neq i} |a_{ij}|, \quad i = 1, \dots, n,$$

then A is invertible. The proof is by contradiction. If $Ax = 0$ has a nonzero solution, choose k so that

$$|x_k| = \max_i |x_i|.$$

The k -th equation gives

$$a_{kk}x_k = - \sum_{j \neq k} a_{kj}x_j.$$

Taking absolute values,

$$|a_{kk}||x_k| \leq \sum_{j \neq k} |a_{kj}||x_j| \leq \left(\sum_{j \neq k} |a_{kj}| \right) |x_k|.$$

Since $x_k \neq 0$, this contradicts strict diagonal dominance. Hence $Ax = 0$ only has the zero solution and A is invertible.

§71.23 Where the idea leads

This self-review page is not a random list. Almost every problem asks: can the object be transformed into a standard form without changing the essential quantity? Determinant problems use row/column operations and block decompositions. Rank problems use normal forms and null spaces. Vector-group problems use span equivalence. Matrix-power problems use eigenvalues, nilpotence, or rank-one structure. Linear-system problems use the principle that solution sets are affine spaces built from one particular solution plus the null space.

§71.24 Determinant tricks as invariant-preserving transformations

Row and column operations are useful because their effect on determinant is controlled. The art is to simplify the matrix while tracking whether the determinant is preserved, multiplied, or sign-changed.

§71.25 Block determinants and Schur complements

For a block matrix with invertible A , elimination gives the Schur complement

$$D - CA^{-1}B.$$

This reduces determinant and inverse computations by eliminating variables blockwise.

§71.26 Matrix powers and minimal polynomials

High powers of a matrix are controlled by the minimal polynomial. Cayley–Hamilton lets one reduce A^n to a combination of lower powers.

§71.27 Matrix tricks from invariants and polynomials

Determinant tricks, recursive determinants, block simplification, rank decompositions, and matrix powers all exploit invariants under controlled transformations and polynomial relations of matrices.

§71.28 Rank factorization and normal forms

If A has rank r , then it factors as

$$A = BC, \quad B \in k^{m \times r}, \quad C \in k^{r \times n}, \quad \text{rank } B = \text{rank } C = r.$$

This is the abstract form behind the note's rank-two decomposition. It says a rank- r map first projects into an r -dimensional image coordinate space and then embeds that information into the output.

Over a field, elementary row and column operations reduce A to the normal form

$$\begin{pmatrix} I_r & 0 \\ 0 & 0 \end{pmatrix}.$$

Over $\mathbb{Z}$, the analogue is Smith normal form, where divisibility data remain. This shows why rank is enough over fields but not enough over rings.

§71.29 Idempotents as projections

A matrix P satisfying

$$P^2 = P$$

is idempotent. Its eigenvalues are only 0 and 1, and

$$V = \text{im } P \oplus \ker P.$$

If P is also symmetric, it is an orthogonal projection. Without symmetry, it is projection along one subspace onto another.

The note contains several absorption/idempotent problems. The advanced reading is that these are problems about decomposing a vector space into complementary subspaces.

§71.30 Spectral asymptotics of matrix powers

Many self-review problems ask for A^n . The long-term behavior is governed by eigenvalues and Jordan blocks. If

$$A = PJP^{-1}, \quad J = \lambda I + N, \quad N^m = 0,$$

then

$$J^n = \sum_{k=0}^{m-1} \binom{n}{k} \lambda^{n-k} N^k.$$

Thus powers are controlled by exponential factors λ^n and polynomial factors from nilpotent parts.

For stochastic-like matrices, the eigenvalue 1 often gives a limiting projection, while all eigenvalues of modulus less than 1 decay. This is the finite-dimensional model of convergence to equilibrium in Markov chains.

§71.31 Null spaces and orthogonal complements

The identity

$$\text{rank}(AA^T) = \text{rank } A$$

comes from equality of null spaces:

$$AA^T x = 0 \iff A^T x = 0.$$

Indeed,

$$x^T AA^T x = |A^T x|^2.$$

This is a Hilbert-space idea in finite-dimensional form: positivity of an inner product turns an algebraic equation into a norm-square equation.

§71.32 Matrix powers and stable limits

When computing A^n , diagonalization is only one case. More generally, the behavior of powers is governed by spectral projectors. If the eigenvalues split into dominant and decaying parts, then A^n often approaches a projection onto the dominant eigenspace after normalization.

For example, if P is a projection, then

$$P^n = P$$

for all $n \geq 1$. If $A = P + N$ with $PN = NP = 0$ and $N^n \rightarrow 0$, then

$$A^n \rightarrow P.$$

This pattern explains many self-review problems where a matrix has a stable part and a nilpotent or contracting part.

§71.33 Row space, column space, and rank factorization

A rank factorization $A = BC$ can be constructed by choosing a basis of the column space. The matrix B embeds the coordinate space of the image into the output; the matrix C records the coordinate functionals that produce those image coordinates from the input.

This is a concrete form of the factorization

$$V \rightarrow V/\ker A \cong \operatorname{im} A \rightarrow W.$$

The map first forgets kernel directions, then identifies the quotient with the image, then includes the image into W .

Thus rank factorization is not just a computational compression. It is the canonical quotient-image structure written in matrices.

§71.34 Positivity through Gram forms

Identities such as $\operatorname{rank}(AA^T) = \operatorname{rank} A$ rely on positivity:

$$x^T AA^T x = \|A^T x\|^2.$$

A Gram matrix has entries

$$G_{ij} = \langle v_i, v_j \rangle.$$

It is always positive semidefinite. Conversely, every positive semidefinite matrix is a Gram matrix of some vectors. This connects rank, inner products, and geometric realization of matrices.

Linear Algebra Review IV: Eigenvalue Tricks, Positive Definiteness, and More

N-20260105

§72.1 Eigenvalues of matrix functions

The first highlighted method says that if $A\xi = \lambda\xi$, then many matrices built from A have the same eigenvector. For a polynomial f ,

$$f(A)\xi = f(\lambda)\xi.$$

If A is invertible, then

$$A^{-1}\xi = \lambda^{-1}\xi.$$

If A is invertible, its adjugate satisfies

$$A^* = |A|A^{-1},$$

so

$$A^*\xi = \frac{|A|}{\lambda}\xi.$$

This is the principle behind the first example: if $A^k = O$, then all eigenvalues satisfy $\lambda^k = 0$, hence all eigenvalues are zero. Therefore

$$|E + 3A| = \prod_i (1 + 3\lambda_i) = 1.$$

The handwritten note identifies this as nilpotent behavior: all eigenvalues of a nilpotent matrix are zero.

§72.2 No real eigenvector from a quadratic relation

Another example assumes a real 2×2 matrix satisfies

$$A^2 + aA + bE = O, \quad a^2 - 4b < 0.$$

If there were a real eigenvector β with real eigenvalue λ , then

$$(\lambda^2 + a\lambda + b)\beta = 0.$$

Since $\beta \neq 0$, one would have

$$\lambda^2 + a\lambda + b = 0.$$

But the discriminant is negative, so there is no real root. Hence the matrix has no real eigenvector. The page notes this as a useful way to rule out real eigenvectors by reducing the problem to the scalar polynomial satisfied by A .

§72.3 Independence of eigenvectors

The second highlighted block uses standard eigenvector facts:

- eigenvectors belonging to distinct eigenvalues are linearly independent;
- an $n \times n$ matrix is diagonalizable iff it has n independent eigenvectors;
- for a real symmetric matrix, eigenvectors belonging to different eigenvalues are orthogonal;
- the geometric multiplicity of an eigenvalue is at most its algebraic multiplicity.

For example, if A has three distinct eigenvalues $\lambda_1, \lambda_2, \lambda_3$ with eigenvectors $\alpha_1, \alpha_2, \alpha_3$, and if

$$B = \alpha_1 + \alpha_2 + \alpha_3,$$

then expressions such as

$$A^3 B, \quad (R - A)B$$

can be evaluated by applying A to each eigendirection separately. The source uses the characteristic identity

$$(A^3 - A^2)\alpha_i = (\lambda_i^3 - \lambda_i^2)\alpha_i.$$

§72.4 Real symmetric reconstruction from eigenvalues

One example states that a real symmetric 3×3 matrix has eigenvalues

$$-1, 1, 1$$

and that an eigenvector belonging to -1 is

$$(0, 1, 1)^T.$$

To reconstruct A , choose an orthonormal basis with first vector in the -1 eigenspace and the other two spanning its orthogonal complement. If

$$P = (\xi_1, \xi_2, \xi_3)$$

is orthogonal and

$$D = \text{diag}(-1, 1, 1),$$

then

$$A = PDP^T.$$

The note's handwritten solution chooses simple orthogonal vectors, forms P , and computes A explicitly. The key is that real symmetric matrices are orthogonally diagonalizable, so reconstruction from eigenspaces is geometrically natural.

§72.5 Similarity invariants: trace and determinant

If A and B are similar, then they have the same characteristic polynomial and hence the same eigenvalues with multiplicity. Therefore

$$\operatorname{tr} A = \operatorname{tr} B, \quad |A| = |B|.$$

One problem uses this with matrices such as

$$3E + 2A, \quad 2E + B, \quad E - 2B$$

being noninvertible. Noninvertibility gives eigenvalue conditions. For example,

$$|3E + 2A| = 0$$

means that $-3/2$ is an eigenvalue of A . After collecting the shared eigenvalues, the trace of the adjugate is computed using

$$\lambda(A^*) = \frac{|A|}{\lambda(A)}.$$

The page writes the final trace as a sum of transformed eigenvalues.

§72.6 Trace shortcuts

The note includes a boxed reminder:

- (i) for a concrete matrix, the trace is the sum of diagonal entries;
- (ii) for a matrix with known eigenvalues, the trace is the sum of eigenvalues;
- (iii) a skew-symmetric matrix has diagonal entries zero, so its trace is zero;
- (iv) a nilpotent matrix has all eigenvalues zero, so its trace is zero.

The handwritten comment says that these are not different facts; they are the same invariant read in different languages.

§72.7 Row-sum eigenvectors

A highlighted method uses equations such as

$$A\xi = \lambda\xi, \quad |\lambda E - A| = 0.$$

If every row sum of A is 2, then

$$A(1, 1, \dots, 1)^T = 2(1, 1, \dots, 1)^T.$$

So 2 is an eigenvalue and the all-ones vector is an eigenvector. The note treats this as a standard quick observation: constant row sums produce the eigenvector $\mathbf{1}$.

§72.8 Diagonalization problems

One problem gives

$$A = \begin{pmatrix} 1 & -1 & 1 \\ a & 4 & b \\ -3 & -3 & 5 \end{pmatrix}$$

and says it has three independent eigenvectors, with 2 a double eigenvalue. Since A is diagonalizable, the eigenspace for $\lambda = 2$ must have dimension 2. Therefore

$$r(2E - A) = 1.$$

Row reduction of $2E - A$ gives conditions on a, b , and the page obtains

$$a = 2, \quad b = -2.$$

The method is important: when a repeated eigenvalue occurs in a diagonalizable matrix, its eigenspace dimension must equal its algebraic multiplicity.

Another example uses $A^3 = A$. The polynomial relation is

$$A(A - E)(A + E) = O.$$

The possible eigenvalues are $-1, 0, 1$. The note compares ranks of factors to count geometric multiplicities and concludes diagonalizability by showing that the total number of independent eigendirections is n .

§72.9 Real symmetric eigenvalue problems

For a real symmetric matrix with rank 2 and relation

$$A^3 + 2A^2 = O,$$

the eigenvalues satisfy

$$\lambda^3 + 2\lambda^2 = \lambda^2(\lambda + 2) = 0.$$

Thus possible eigenvalues are 0 and -2 . Since $r(A) = 2$, exactly two eigenvalues are nonzero, and for a real symmetric matrix the diagonal form is similar to

$$\text{diag}(-2, -2, 0).$$

The source warns not to infer algebraic multiplicities only from the scalar equation; rank supplies the missing count.

§72.10 Odd powers of real symmetric matrices

One problem states: if A is a real symmetric matrix, prove that for every positive integer k there exists a real symmetric matrix B such that

$$B^{2k+1} = A.$$

Since A is real symmetric, there is an orthogonal matrix Q with

$$Q^T A Q = D = \text{diag}(\lambda_1, \dots, \lambda_n).$$

Every real number has a real $(2k+1)$ -st root, so define

$$C = \text{diag}(\sqrt[2k+1]{\lambda_1}, \dots, \sqrt[2k+1]{\lambda_n}).$$

Then

$$B = QCQ^T$$

is real symmetric and satisfies $B^{2k+1} = A$. The source points out that the odd power matters; even roots of negative eigenvalues would fail over the reals.

§72.11 Rayleigh quotient extrema

For a real symmetric matrix A , the problem asks for

$$\max_{x^T x=1} x^T A x, \quad \min_{x^T x=1} x^T A x.$$

If

$$A = Q^T \text{diag}(\lambda_1, \dots, \lambda_n) Q$$

with Q orthogonal, set $y = Qx$. Then $y^T y = 1$ and

$$x^T A x = y^T \text{diag}(\lambda_i) y = \sum_i \lambda_i y_i^2.$$

Thus the maximum is the largest eigenvalue and the minimum is the smallest eigenvalue. In the source example, the eigenvalues are

$$1, 4, 10,$$

so the maximum is 10 and the minimum is 1.

§72.12 Orthogonal determinant trick

If A is an orthogonal matrix and $|A| = -1$, then the note proves that A has eigenvalue -1 . Since A is orthogonal,

$$|A| |E + A| = |A + A^T A| = |E + A|.$$

Because $|A| = -1$, this forces

$$|E + A| = 0.$$

Therefore -1 is an eigenvalue.

§72.13 Quadratic forms through diagonalization

The quadratic-form section uses diagonalization. If

$$f(x_1, x_2, x_3) = x^T A x$$

is brought by an orthogonal change $x = Py$ to

$$2y_1^2 - y_2^2 - y_3^2,$$

then A is orthogonally similar to

$$D = \text{diag}(2, -1, -1).$$

A problem asks for

$$|3A^* - 2A^{-1}|.$$

Using similarity and adjugate eigenvalues, reduce the determinant to the diagonal case:

$$3A^* - 2A^{-1} \sim (3|D| - 2)D^{-1},$$

so the determinant is computed from the diagonal entries. The source obtains a numerical value by multiplying the resulting diagonal factors.

§72.14 Positive definite equations

One problem states: if A, C are positive definite symmetric matrices and the equation

$$AX + XA = C$$

has a unique solution B , prove B is positive definite. The note observes first that transposing the equation gives another solution, so uniqueness implies

$$B^T = B.$$

If $B\xi = \lambda\xi$, then using the equation gives

$$\xi^T C \xi = \xi^T A B \xi + \xi^T B A \xi = 2\lambda \xi^T A \xi.$$

Since A and C are positive definite, both quadratic forms are positive for $\xi \neq 0$, so $\lambda > 0$. Hence all eigenvalues of the symmetric matrix B are positive, and B is positive definite.

§72.15 Commuting positive definite products

If A, B are symmetric positive definite and

$$AB = BA,$$

then AB is symmetric and positive definite. Symmetry follows from

$$(AB)^T = B^T A^T = BA = AB.$$

For positive definiteness, diagonalize $B = P^T D P$ with $D > 0$, or use the fact that commuting real symmetric matrices can be simultaneously orthogonally diagonalized. Then the eigenvalues of AB are products of positive eigenvalues, hence positive.

The source also proves an inequality: for positive definite A, B ,

$$(A + B)(A^{-1} + B^{-1})$$

has all eigenvalues at least 4. Writing

$$(A + B)(A^{-1} + B^{-1}) = 2E + C + C^{-1}, \quad C = AB^{-1},$$

and reducing to the positive eigenvalues of a similar symmetric positive definite matrix gives

$$\lambda + \lambda^{-1} \geq 2.$$

§72.16 Principal minors and congruence

The page uses positive definiteness through principal minors. One example has $Ax = 0$ with nonzero solution, so

$$|A| = 0.$$

Candidate values for a parameter are tested by the leading principal minors of a given positive definite matrix B . The source chooses the value making both the determinant condition and positive definiteness true.

Another problem uses congruence. If B is symmetric, prove that

$$A = \begin{pmatrix} E & B \\ B & E \end{pmatrix}$$

is positive definite iff all eigenvalues of B satisfy

$$|\lambda_i(B)| < 1.$$

The hint uses a congruent transformation with a matrix like

$$P = \begin{pmatrix} E & -B \\ O & E \end{pmatrix}$$

to reduce the form to

$$\text{diag}(E, E - B^2).$$

Thus positive definiteness is equivalent to

$$E - B^2 > 0,$$

which is equivalent to $1 - \lambda_i(B)^2 > 0$.

§72.17 Ordering eigenvalues and positive definiteness

If A, B are symmetric positive definite and the smallest eigenvalue of A is larger than the largest eigenvalue of B , then $A - B$ is positive definite. The note's hint takes a midpoint k between those two eigenvalue bounds. Then

$$A - kE > 0, \quad kE - B > 0,$$

so

$$A - B = (A - kE) + (kE - B) > 0.$$

§72.18 Transition matrices

The linear-space section studies change of basis. Suppose in $\mathbb{R}^3$ there are bases

$$\varepsilon_1, \varepsilon_2, \varepsilon_3$$

and

$$\alpha_1 = \varepsilon_1, \quad \alpha_2 = \varepsilon_1 + \varepsilon_2, \quad \alpha_3 = \varepsilon_1 + \varepsilon_2 + \varepsilon_3,$$

as well as another basis

$$\beta_1 = \varepsilon_1 + \varepsilon_2, \quad \beta_2 = 2\varepsilon_2 + \varepsilon_3, \quad \beta_3 = \varepsilon_1 + 3\varepsilon_3.$$

Let

$$(\alpha_1, \alpha_2, \alpha_3) = (\varepsilon_1, \varepsilon_2, \varepsilon_3)P_1,$$

and

$$(\beta_1, \beta_2, \beta_3) = (\varepsilon_1, \varepsilon_2, \varepsilon_3)P_2.$$

Then the transition matrix from the α -basis to the β -basis is

$$P = P_1^{-1}P_2,$$

because

$$(\beta_1, \beta_2, \beta_3) = (\alpha_1, \alpha_2, \alpha_3)P.$$

The handwritten solution computes these matrices explicitly.

§72.19 Same coordinates under two bases

The last page asks for vectors whose coordinates are the same under two bases. Let the standard basis be

$$\varepsilon_1 = (1, 0, 0)^T, \quad \varepsilon_2 = (0, 1, 0)^T, \quad \varepsilon_3 = (0, 0, 1)^T,$$

and another basis be

$$\alpha_1 = (1, 0, 0)^T, \quad \alpha_2 = (1, 1, 0)^T, \quad \alpha_3 = (1, 1, 1)^T.$$

A vector with coordinate column $t = (t_1, t_2, t_3)^T$ has the same coordinates in both bases iff

$$t_1\varepsilon_1 + t_2\varepsilon_2 + t_3\varepsilon_3 = t_1\alpha_1 + t_2\alpha_2 + t_3\alpha_3.$$

This gives a homogeneous linear system in t . The source row-reduces the system and obtains a one-dimensional family, hence all such vectors are scalar multiples of one vector.

§72.20 Conceptual synthesis

This note is a review of how eigenvalues and bases organize linear algebra. Polynomial identities for matrices become scalar identities on eigenvalues. Diagonalization translates operators into independent directions. Real symmetric matrices make the translation orthogonal and geometric. Positive definiteness is tested by eigenvalues, principal minors, or congruence. Change of basis problems are solved by writing both bases in a common reference basis and multiplying the corresponding coordinate matrices.

§72.21 Eigenvalues through functional calculus

If $Av = \lambda v$, then for a polynomial p ,

$$p(A)v = p(\lambda)v.$$

This is why eigenvalues of matrix functions can often be read from eigenvalues of the matrix.

§72.22 Similarity invariants

Trace, determinant, characteristic polynomial, and eigenvalues are invariant under similarity. They depend on the linear operator, not the chosen basis.

§72.23 Positive definiteness and quadratic geometry

A symmetric matrix A is positive definite exactly when

$$x^T A x > 0 \quad (x \neq 0).$$

Spectrally, this means all eigenvalues are positive. Geometrically, its level sets are ellipsoids.

§72.24 Row-sum eigenvectors and invariant aggregate directions

If every row sum equals c , then the all-ones vector is an eigenvector with eigenvalue c . This is a symmetry-reduction trick: a common aggregate direction is preserved.

§72.25 Spectral geometry of positive matrices

Eigenvalue tricks work because invariant subspaces, similarity invariants, spectral functions, and quadratic-form geometry constrain what a matrix can do independent of coordinates.

§72.26 Functional calculus for matrices

If A is diagonalizable, $A = PDP^{-1}$, then a polynomial or suitable function of A is defined by

$$f(A) = Pf(D)P^{-1}, \quad f(D) = \text{diag}(f(\lambda_1), \dots, f(\lambda_n)).$$

For real symmetric matrices, this is especially stable because P can be chosen orthogonal.

This explains many eigenvalue tricks in the note. If $A\xi = \lambda\xi$, then

$$f(A)\xi = f(\lambda)\xi.$$

Nilpotent, idempotent, inverse, adjugate, and odd-root problems are all examples of applying functions to eigenvalues.

§72.27 Loewner order and positive definiteness

For symmetric matrices, write

$$A \succeq B$$

if $A - B$ is positive semidefinite. This is the Loewner order. Positive definiteness is then $A \succ 0$.

The Rayleigh quotient characterizes eigenvalue bounds:

$$\lambda_{\min}(A) \leq \frac{x^T A x}{x^T x} \leq \lambda_{\max}(A).$$

Thus a statement such as “the smallest eigenvalue of A is larger than the largest eigenvalue of B ” immediately implies

$$A - B \succ 0.$$

This is exactly the order-theoretic content of one of the positive-definite examples.

§72.28 Lyapunov equations

The equation

$$AX + XA = C$$

in the note is a special Lyapunov/Sylvester equation. More generally,

$$AX + XB = C$$

has a unique solution when the spectra of A and $-B$ are disjoint.

If A is positive definite and symmetric, the equation

$$AX + XA = C$$

can be diagonalized by an orthogonal matrix. If $A = QDQ^T$ and $Y = Q^T X Q$, $K = Q^T C Q$, then

$$DY + YD = K, \quad (d_i + d_j)Y_{ij} = K_{ij}.$$

Because $d_i + d_j > 0$, the solution is unique. Positivity of C then forces positivity properties of X .

§72.29 Convex and semidefinite optimization

Positive semidefinite matrices form a convex cone:

$$S_+^n = \{A = A^T : x^T A x \geq 0 \text{ for all } x\}.$$

A semidefinite program optimizes a linear function over affine slices of this cone. Many quadratic-form inequalities can be rewritten as semidefinite constraints.

The elementary leading-principal-minor tests are therefore finite-dimensional certificates for membership in this cone. The research-level continuation is optimization over matrix cones.

§72.30 Matrix monotone warning

For positive numbers, $0 < a \leq b$ implies $a^2 \leq b^2$. For positive definite matrices, this implication may fail when matrices do not commute. Functions preserving Loewner order are called operator monotone.

The function $t \mapsto t^{-1}$ reverses order on positive definite matrices:

$$0 \prec A \preceq B \implies B^{-1} \preceq A^{-1}.$$

This kind of phenomenon explains why matrix inequalities are subtler than scalar inequalities, even when the notation looks similar.

§72.31 Positive definite matrices and ellipsoids

A symmetric positive definite matrix A defines an ellipsoid

$$E_A = \{x : x^T A x \leq 1\}.$$

If $A = QDQ^T$ with $D = \text{diag}(\lambda_i)$, then the principal semi-axis lengths are $1/\sqrt{\lambda_i}$. Larger eigenvalues mean tighter curvature in the corresponding eigenvector directions.

This turns inequalities between positive definite matrices into inclusions of ellipsoids. If

$$A \succeq B,$$

then $x^T A x \geq x^T B x$ for all x , so

$$E_A \subseteq E_B.$$

The Loewner order is therefore geometric containment reversed through quadratic forms.

§72.32 Lyapunov equations and stability

For a continuous-time linear system

$$x' = Mx,$$

a Lyapunov function may have the form

$$V(x) = x^T P x, \quad P \succ 0.$$

Then

$$\frac{d}{dt} V(x(t)) = x^T (M^T P + P M) x.$$

If one can solve

$$M^T P + P M = -Q, \quad Q \succ 0,$$

then V decreases along nonzero trajectories, proving stability.

This gives an applied meaning to Sylvester and Lyapunov equations. They are not only matrix exercises; they construct quadratic energies for dynamical systems.

§72.33 Functional calculus and matrix inequalities

For a positive definite matrix A , one defines

$$A^{1/2} = Q \operatorname{diag}(\sqrt{\lambda_i}) Q^T.$$

This square root is the unique positive definite matrix satisfying $(A^{1/2})^2 = A$. Similarly, one defines $\log A$, e^A , and A^t spectrally.

However, scalar intuition can fail for noncommuting matrices. Even if $A \preceq B$, it is not always true for arbitrary scalar functions f that $f(A) \preceq f(B)$. Operator monotone functions are precisely those for which this implication holds. This warning is essential when extending elementary inequalities to matrices.

Two-Dimensional Linear Maps, Spectral Geometry, and Singular Value Decomposition

N-20260303

§73.1 Concept map of the chapter

The source note begins with a map of the main linear-algebra concepts used later in the machine-learning-oriented discussion. Determinants test invertibility and also appear before eigenvalues; eigenvalues determine eigenvectors; eigenvectors construct orthogonal matrices in favorable cases; orthogonal matrices are used in diagonalization; diagonalization and singular value decomposition then become the computational language for dimensionality reduction.




Figure 4.1 A mind map of the concepts introduced in this chapter, along with where they are used in other parts of the book.

Figure 73.1: A concept map linking determinant, invertibility, eigenvalues, eigenvectors, orthogonal matrices, diagonalization, Cholesky factorization, SVD, and dimensionality reduction.

The diagram should not be read as a strict dependency graph. It is a guide to how these objects are reused: determinant and eigenvalues are algebraic invariants; eigenvectors and orthogonal matrices give coordinate systems; SVD is the robust replacement when diagonalization is unavailable or when the matrix is rectangular.

§73.2 Five model transformations

A 2×2 matrix is not only an array of four numbers; it is a rule that moves every point of the plane. The determinant records signed area scaling, while eigenvectors record

directions that survive the transformation without turning:

$$|A(e_1), A(e_2)| = \det A, \quad Av = \lambda v.$$

The source figure compares five concrete linear maps by applying each to a grid of colored points. The left column is the input, the right column is the transformed output, and the middle column marks eigenvector directions and eigenvalue data where they are real.




Figure 73.2: Five planar linear mappings and their eigenspace behavior. Each row shows colored input points, eigenvalue/eigenvector information, and the transformed output points.

The five rows illustrate the standard possibilities: axis-aligned stretching, shear, rotation with complex eigenvalues, collapse onto a line, and symmetric stretching. The determinant and the eigenvalues answer different questions. The determinant says how area changes; eigenvectors say which directions are preserved.

§73.3 A spectral example: a biological neural network

The note then moves from geometry to data. For the *C. elegans* neural network, one may form a connectivity matrix. Since the directed connection matrix is generally not symmetric, one often studies a symmetrized matrix, for example

$$A_{\text{sym}} = A + A^T.$$

The eigenvalues of this matrix describe dominant modes of connectivity.

Figure 4.5
Caenorhabditis
elegans neural
network (Kaiser and
Hilgetag,
2006). (a) Sym-
metrized
connectivity matrix;
(b) Eigenspectrum.




Figure 73.3: The symmetrized connectivity matrix of the *C. elegans* neural network and its sorted eigenspectrum.

This example is a useful warning: eigenvectors are not only arrows in the plane. In a network or a dataset, an eigenvector can indicate a coherent global mode, and the corresponding eigenvalue measures its strength.

§73.4 SVD as a sequence of geometric operations

Eigenvalue decomposition is powerful but fragile: it is best behaved for square matrices with enough eigenvectors, and especially for symmetric matrices. Singular value decomposition is more stable. Every real matrix $A \in \mathbb{R}^{m \times n}$ admits

$$A = U\Sigma V^T,$$

where U and V are orthogonal and Σ has nonnegative diagonal entries $\sigma_1 \geq \sigma_2 \geq \dots \geq 0$. Geometrically, V^T changes the input coordinates, Σ stretches by the singular values, and U changes the output coordinates.




Figure 73.4: Geometric interpretation of SVD: the unit disk is first expressed in the right-singular-vector coordinates, then stretched by the singular values, and finally rotated into the output coordinates.

Thus SVD does not ask for directions that remain invariant under A . Instead, it asks for orthogonal input directions that are sent to orthogonal output axes. The unit circle becomes an ellipse, and the semi-axis lengths are the singular values.

§73.5 SVD and latent structure in data

A rating table can be viewed as a matrix. Rows may represent movies and columns may represent users. The raw entries are ratings, but the singular vectors can reveal latent directions such as genre preference or taste profile.




Figure 4.10 Movie ratings of three people for four movies and its SVD decomposition.

Figure 73.5: A movie-rating matrix and its SVD decomposition. Dominant singular directions highlight latent structure in user and movie preferences.

The labels placed on singular directions are interpretations, not part of the theorem. The theorem says that the first singular directions give the most efficient orthogonal explanation of the matrix energy.

§73.6 Rank-one blocks and image compression

SVD decomposes a matrix into rank-one pieces:

$$A = \sum_{i=1}^r \sigma_i u_i v_i^T.$$

For an image matrix, each term $\sigma_i u_i v_i^T$ is a weighted rank-one visual pattern.




Figure 73.6: The first rank-one SVD components of an image. Each component captures one weighted visual pattern rather than a recognizable full image.

Truncated SVD keeps only the first k terms:

$$\hat{\mathbf{A}}(k) = \sum_{i=1}^k \sigma_i u_i v_i^T.$$

The Eckart–Young theorem says that this is the best rank- k approximation in the spectral norm and the Frobenius norm.




Figure 73.7: Low-rank image reconstruction from the first few singular components. As the rank increases, the approximation recovers more recognizable structure.

This example explains why SVD is useful beyond pure linear algebra: it gives a principled compression method by separating large-scale structure from smaller-scale details.

For the examples in this note, the practical reading rule is:

question	tool
area/orientation	$\det A$
invariant directions	eigenvectors
collapse or information loss	rank A
principal stretching	singular values
best low-rank compression	truncated SVD

This is why the images are not decorative. They show the same matrix simultaneously as an algebraic object, a geometric transformation, and a data-processing device.

Orthogonal Projection, Least Squares, and the Pseudoinverse

N-20260305

§74.1 Projection onto a subspace

Let $U \subseteq \mathbb{R}^n$ have ordered basis $b_1, \dots, b_m$ and put

$$B = (b_1, \dots, b_m) \in \mathbb{R}^{n \times m}.$$

The orthogonal projection of x onto U has the form $B\lambda$. The error must be orthogonal to U :

$$B^T(x - B\lambda) = 0.$$

Thus

$$B^T B \lambda = B^T x, \quad \lambda = (B^T B)^{-1} B^T x,$$

assuming the columns of B are independent. Therefore

$$\pi_U(x) = B(B^T B)^{-1} B^T x.$$

§74.2 Projection matrix and least squares

The matrix

$$P_U = B(B^T B)^{-1} B^T$$

is symmetric and idempotent:

$$P_U^T = P_U, \quad P_U^2 = P_U.$$

For an overdetermined system $B\lambda \approx x$, minimizing $\|B\lambda - x\|^2$ gives the same normal equations. Least squares is orthogonal projection written in coordinates.

§74.3 Pseudoinverse

When B has full column rank, the Moore–Penrose pseudoinverse is

$$B^+ = (B^T B)^{-1} B^T.$$

For a general matrix $A = U\Sigma V^T$, the pseudoinverse is $A^+ = V\Sigma^+ U^T$, where nonzero singular values are inverted. This makes SVD the stable language of projection and least-squares computation.

§74.4 Projection geometry in the source diagram

The diagram is kept because projection and least squares are most transparent when the orthogonal error is visible. The vector being approximated is decomposed into its projection onto the column space and an error vector perpendicular to that space.




Figure 3.11
Projection onto a two-dimensional subspace U with basis b_1, b_2 . The projection $\pi_U(x)$ of $x \in \mathbb{R}^3$ onto U can be expressed as a linear combination of b_1, b_2 and the displacement vector $x - \pi_U(x)$ is orthogonal to both b_1 and b_2 .

Figure 74.1: Projection onto a two-dimensional subspace: the residual $x - \pi_U(x)$ is orthogonal to the basis directions of U .

§74.5 Bridge to later ideas

Projection is the geometric heart of least squares: the best approximation is characterized not by making the error small in every coordinate, but by making the error orthogonal to every allowed correction direction.

§74.6 Projection theorem

Theorem 74.6.1 (Projection theorem)

For a closed subspace W of an inner-product space and a vector y , the best approximation $p \in W$ is characterized by

$$y - p \perp W.$$

For a closed subspace W of an inner-product space, the best approximation to y in W is the point $p \in W$ such that

$$y - p \perp W.$$

This orthogonality condition is the geometric core of least squares.

§74.7 Normal equations

For $Ax \approx b$, the residual $b - Ax$ must be orthogonal to the column space of A :

$$A^T(Ax - b) = 0.$$

Hence

$$A^T Ax = A^T b.$$

§74.8 Pseudoinverse and singular values

The Moore–Penrose pseudoinverse is most transparent through the singular value decomposition. If

$$A = U\Sigma V^T,$$

then

$$A^+ = V\Sigma^+U^T.$$

Small singular values reveal unstable directions.

§74.9 Least squares as projection geometry

Least squares is projection geometry. Normal equations express orthogonality of residuals, the pseudoinverse selects the minimum-norm solution, and SVD exposes the stable and unstable directions of the problem.

Matrix Computation I: Norms, Operators, and Numerical Size

N-20260311

§75.1 Why size matters

Matrix computation starts from a simple question: how large can an error become after a linear operation acts on it? A vector norm measures the size of data; a matrix norm measures the size of an operator. In machine learning language, these are the quantities behind amplification, Lipschitz constants, gradient growth, and numerical stability.

A matrix is an operator

$$x \mapsto Ax.$$

To control it we need both

vector size and operator size.

§75.2 Metrics and completeness

A metric on a set X is a function $d : X \times X \rightarrow [0, \infty)$ satisfying positivity, symmetry, and the triangle inequality. It defines convergence. A sequence is Cauchy if

$$\forall \varepsilon > 0, \exists N, \quad m, n > N \Rightarrow d(x_m, x_n) < \varepsilon.$$

A metric space is complete if every Cauchy sequence converges in it. Completeness means that limiting procedures stay inside the space.

The discrete metric is a useful warning: not every metric expresses the geometry we want. For numerical work, metrics induced by norms are the central examples.

§75.3 Normed vector spaces

A norm satisfies

$$\|x\| = 0 \Leftrightarrow x = 0, \quad \|\alpha x\| = |\alpha| \|x\|, \quad \|x + y\| \leq \|x\| + \|y\|.$$

It induces the metric

$$d(x, y) = \|x - y\|.$$

A complete normed vector space is a Banach space.

In finite dimension, all norms are equivalent: for two norms $\|\cdot\|_a, \|\cdot\|_b$, there exist constants $c, C > 0$ such that

$$c\|x\|_a \leq \|x\|_b \leq C\|x\|_a.$$

Thus finite-dimensional convergence does not depend on the chosen norm, although numerical constants do.

§75.4 Standard vector norms

For $x = (x_1, \dots, x_n)^T$,

$$\|x\|_1 = \sum_i |x_i|, \quad \|x\|_2 = \left(\sum_i |x_i|^2 \right)^{1/2}, \quad \|x\|_\infty = \max_i |x_i|.$$

More generally,

$$\|x\|_p = \left(\sum_i |x_i|^p \right)^{1/p}, \quad 1 \leq p < \infty.$$

As $p \rightarrow \infty$,

$$\|x\|_p \rightarrow \|x\|_\infty.$$

Holder's inequality gives

$$\sum_i |x_i y_i| \leq \|x\|_p \|y\|_q, \quad \frac{1}{p} + \frac{1}{q} = 1,$$

and Minkowski's inequality gives the triangle inequality for p -norms.

§75.5 Matrix norms

A matrix norm satisfies the norm axioms and, for computation, should also be submultiplicative:

$$\|AB\| \leq \|A\| \|B\|.$$

This is what allows repeated operations to be bounded:

$$\|A^k\| \leq \|A\|^k.$$

Entrywise examples include

$$m_1(A) = \sum_{i,j} |a_{ij}|, \quad \|A\|_F = \left(\sum_{i,j} |a_{ij}|^2 \right)^{1/2}.$$

The Frobenius norm also satisfies

$$\|A\|_F^2 = \text{tr}(A^H A).$$

§75.6 Unitary invariance

A matrix U is unitary if $U^H U = I$. It preserves Euclidean length:

$$\|Ux\|_2 = \|x\|_2.$$

The Frobenius norm is unitarily invariant:

$$\|UAV\|_F = \|A\|_F$$

for unitary U, V . This follows from cyclic invariance of trace:

$$\|UAV\|_F^2 = \text{tr}(V^H A^H AV) = \text{tr}(A^H A).$$

Unitary transformations are changes of orthonormal coordinates, not changes of size.

§75.7 Induced norms

Given vector norms on domain and codomain, the induced matrix norm is

$$\|A\| = \max_{x \neq 0} \frac{\|Ax\|}{\|x\|} = \max_{\|x\|=1} \|Ax\|.$$

It is the maximal amplification factor of A . The standard formulas are

$$\|A\|_1 = \max_j \sum_i |a_{ij}|, \quad \|A\|_\infty = \max_i \sum_j |a_{ij}|,$$

and

$$\|A\|_2 = \sigma_{\max}(A) = \sqrt{\lambda_{\max}(A^H A)}.$$

The 2-norm is the spectral norm.

§75.8 Normal matrices

A matrix is Hermitian if $A^H = A$, skew-Hermitian if $A^H = -A$, and normal if

$$A^H A = A A^H.$$

Normal matrices are unitarily diagonalizable. If

$$A = U \Lambda U^H,$$

then

$$\|A\|_2 = \max_i |\lambda_i|.$$

For nonnormal matrices, eigenvalues may seriously underestimate transient amplification.

§75.9 Compatibility

A vector norm and a matrix norm are compatible if

$$\|Ax\| \leq \|A\| \|x\|.$$

Every induced matrix norm is compatible with the vector norm that induces it. Compatibility is the bridge from operator size to error propagation:

$$\|A(x + \delta x) - Ax\| \leq \|A\| \|\delta x\|.$$

§75.10 Worked constrained maxima

Let

$$A = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 2 \end{bmatrix}.$$

Then A is Hermitian with eigenvalues 2, 2, 0. Hence

$$\|A\|_2 = 2.$$

The column sums are all 2, so $\|A\|_1 = 2$. Therefore

$$\max_{\|x\|_1=4} \|Ax\|_1 = 8.$$

If U is unitary and $\|Ux\|_2 = 3$, then $\|x\|_2 = 3$, hence

$$\max_{\|Ux\|_2=3} \|Ax\|_2 = 6.$$

The method is: translate the constraint into a norm ball and apply the induced norm.

§75.11 ML-facing interpretation

A linear layer has Lipschitz constant at most $\|W\|_2$ in Euclidean norm. A stack obeys

$$\|W_L \cdots W_1\|_2 \leq \prod_{\ell=1}^L \|W_\ell\|_2.$$

This inequality is crude but important: moderate layerwise amplification can accumulate over depth. Spectral normalization, orthogonal initialization, residual scaling, and weight decay are all connected to the same numerical question: how large can a perturbation become?

§75.12 What each norm is measuring

object	what it measures	typical use
metric	distance between two points	convergence and continuity
vector norm	size of one vector	error magnitude
induced matrix norm	$\max_{x \neq 0} \ Ax\ /\ x\ $	worst-case amplification
Frobenius norm	$\sqrt{\sum_{ij} a_{ij} ^2}$	entrywise energy
spectral norm	$\sigma_{\max}(A)$	Euclidean Lipschitz bound

The worked constrained-maximum examples explain the difference between Frobenius and spectral size. Frobenius norm counts all entries at once, while the spectral norm asks for the single input direction that is stretched the most. In the linear-layer interpretation, $\|W\|_2$ is therefore the relevant bound for adversarial or perturbation amplification, whereas $\|W\|_F$ is a global weight-size penalty. The inequality

$$\|W_L \cdots W_1\|_2 \leq \prod_{\ell=1}^L \|W_\ell\|_2$$

should be read as a warning about composition: even if each layer is controlled separately, the product controls how errors can accumulate through the whole map.

Matrix Computation II: Spectral Radius, Stability, and Conditioning

N-20260312

§76.1 The long-term question

A norm measures one-step amplification. Spectral radius measures long-term exponential behavior. When a matrix is applied repeatedly,

$$x, Ax, A^2x, \dots,$$

which directions survive is controlled by eigenvalues.

§76.2 Spectral radius

For $A \in \mathbb{C}^{n \times n}$ with eigenvalues $\lambda_1, \dots, \lambda_n$, define

$$\rho(A) = \max_j |\lambda_j|.$$

For any operator norm,

$$\rho(A) \leq \|A\|.$$

Indeed, if $Ax = \lambda x$, $x \neq 0$, then

$$|\lambda| \|x\| = \|Ax\| \leq \|A\| \|x\|.$$

For normal matrices,

$$\rho(A) = \|A\|_2.$$

§76.3 Spectral radius is not a norm

The nonzero nilpotent matrix

$$N = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

has $\rho(N) = 0$. Thus positive definiteness fails.

Let

$$N_1 = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad N_2 = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}.$$

Then $\rho(N_1) = \rho(N_2) = 0$, but

$$\rho(N_1 + N_2) = 1.$$

So the triangle inequality fails.

Finally, take

$$E = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad F = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}.$$

Then $\rho(E) = \rho(F) = 1$, while $\rho(EF) = (1 + \sqrt{5})/2 > 1$. So submultiplicativity fails.

§76.4 Approximating spectral radius

For every $\varepsilon > 0$, there exists a matrix norm such that

$$\|A\| \leq \rho(A) + \varepsilon.$$

The proof uses Jordan form. Write $P^{-1}AP = J$. On each Jordan block, conjugate by

$$D = \text{diag}(1, \delta, \delta^2, \dots),$$

which shrinks the superdiagonal entries from 1 to δ . Choosing δ small makes the norm of the block close to the modulus of its eigenvalue.

§76.5 Convergent matrices

A matrix is convergent if

$$A^k \rightarrow 0.$$

The criterion is

$$A^k \rightarrow 0 \iff \rho(A) < 1.$$

Necessity follows because eigenvalues of A^k are λ^k . Sufficiency follows by choosing a norm with $\|A\| < 1$, using the approximation theorem, and then

$$\|A^k\| \leq \|A\|^k \rightarrow 0.$$

§76.6 Jordan boundary behavior

If $J = \lambda I + N$, $N^m = 0$, then

$$J^k = \sum_{j=0}^{m-1} \binom{k}{j} \lambda^{k-j} N^j.$$

Polynomial factors cannot defeat exponential decay when $|\lambda| < 1$, but they matter when $|\lambda| = 1$. This is why the strict inequality $\rho(A) < 1$ is essential.

§76.7 Neumann series

If $\rho(A) < 1$, then

$$\sum_{k=0}^{\infty} A^k = (I - A)^{-1}.$$

This follows from

$$(I - A)(I + A + \dots + A^m) = I - A^{m+1}.$$

Neumann series turns a static inverse into repeated feedback accumulation.

§76.8 Matrix power series

If $\sum a_k z^k$ has scalar radius of convergence R , then

$$\sum a_k A^k$$

converges when $\rho(A) < R$ and diverges when $\rho(A) > R$. Eigenvalues decide whether the matrix lies inside the convergence disk.

§76.9 Condition number

For invertible A , define

$$\kappa(A) = \|A\| \|A^{-1}\|.$$

In the Euclidean norm,

$$\kappa_2(A) = \frac{\sigma_{\max}(A)}{\sigma_{\min}(A)}.$$

If A is normal and invertible, then

$$\kappa_2(A) = \frac{\max |\lambda_j|}{\min |\lambda_j|}.$$

A large condition number means solving $Ax = b$ is intrinsically sensitive.

§76.10 Perturbation bound

If $(A + \delta A)\hat{x} = b + \delta b$, then a standard estimate is

$$\frac{\|x - \hat{x}\|}{\|x\|} \leq \frac{\kappa(A)}{1 - \kappa(A) \frac{\|\delta A\|}{\|A\|}} \left(\frac{\|\delta A\|}{\|A\|} + \frac{\|\delta b\|}{\|b\|} \right),$$

provided

$$\kappa(A) \frac{\|\delta A\|}{\|A\|} < 1.$$

The denominator is part of the theorem, not a decoration.

§76.11 Ill-conditioned example

Let

$$A_\varepsilon = \begin{bmatrix} 1 & -1 \\ -1 & 1 + \varepsilon \end{bmatrix}.$$

The eigenvalues are

$$\lambda_{\pm} = \frac{2 + \varepsilon \pm \sqrt{4 + \varepsilon^2}}{2}.$$

For small ε ,

$$\lambda_+ \approx 2, \quad \lambda_- \approx \varepsilon/2,$$

so

$$\kappa_2(A_\varepsilon) \approx \frac{4}{\varepsilon}.$$

The matrix is invertible for $\varepsilon > 0$, but nearly singular when ε is small.

§76.12 Numerical and ML interpretation

Repeated multiplication appears in recurrent systems, graph propagation, linearized training dynamics, and iterative solvers. Spectral radius tells whether powers decay or grow asymptotically. Condition number tells whether solving or inversion amplifies perturbations.

Mixed precision increases the relative size of rounding perturbations. If the problem is ill-conditioned, those perturbations can become wrong directions. The safe statement is not that condition number explains all training failures, but that it reveals a basic numerical instability channel.

§76.13 Power growth, inversion, and sensitivity

quantity	meaning	where it appears
$\rho(A)$	asymptotic eigenvalue growth	A^k and stability
$\ A\ $	one-step amplification	single perturbation bound
$A^k \rightarrow 0$	$\rho(A) < 1$	decaying iteration
$(I - A)^{-1}$	$\sum_{k \geq 0} A^k$	Neumann series
$\kappa(A)$	$\ A\ \ A^{-1}\ $	sensitivity of solving

The distinction between $\rho(A)$ and $\|A\|$ is the main computational point here. A norm controls what can happen in one step, while the spectral radius controls the eventual behavior of powers. Jordan blocks explain why the boundary case $\rho(A) = 1$ is delicate: polynomial growth can remain even when all eigenvalues have modulus one.

For solving $Ax = b$, the condition number measures how much relative error can be magnified by inversion. The ill-conditioned example is therefore not a numerical accident; it shows that a small perturbation in b can point in a direction where A^{-1} is large. This is the same instability channel that appears in iterative dynamics and mixed-precision computations.

Matrix Computation III: Canonical Forms, Localization, and Iteration

N-20260318

§77.1 Why canonical forms matter

Canonical forms explain what a matrix cannot hide under a change of basis. For computation they answer: how do powers behave, where can eigenvalues be, and which direction dominates an iteration?

§77.2 Polynomial matrices

A λ -matrix is a matrix whose entries lie in $\mathbb{C}[\lambda]$. Its rank is the largest order of a minor that is not the zero polynomial. A square λ -matrix is nonsingular if its determinant is not the zero polynomial.

The elementary operations are row or column exchange, multiplication of a row or column by a nonzero constant, and adding $\varphi(\lambda)$ times one row or column to another.

§77.3 Smith normal form

Every rank r polynomial matrix is equivalent to

$$\text{diag}(d_1(\lambda), \dots, d_r(\lambda), 0, \dots, 0),$$

where each d_i is monic and

$$d_1 \mid d_2 \mid \dots \mid d_r.$$

This is the Smith normal form.

Let $D_k(\lambda)$ be the monic greatest common divisor of all nonzero k -minors. The invariant factors are

$$d_1 = D_1, \quad d_k = \frac{D_k}{D_{k-1}}.$$

Factoring the nonconstant invariant factors gives elementary divisors.

§77.4 Jordan data from elementary divisors

For $A \in \mathbb{C}^{n \times n}$, apply Smith form to

$$\lambda I - A.$$

An elementary divisor

$$(\lambda - \lambda_0)^s$$

corresponds to a Jordan block of size s for eigenvalue λ_0 .

Thus Smith form is an algebraic route to Jordan structure.

§77.5 Jordan chains

A Jordan block is

$$J_s(\lambda) = \begin{bmatrix} \lambda & 1 & & 0 \\ & \lambda & \ddots & \\ & & \ddots & 1 \\ 0 & & & \lambda \end{bmatrix}.$$

If $A = PJP^{-1}$, the columns in one Jordan chain satisfy

$$Ap_1 = \lambda p_1, \quad Ap_i = \lambda p_i + p_{i-1}.$$

The later vectors are generalized eigenvectors.

§77.6 Computing powers

If $J = \lambda I + N$ with $N^s = 0$, then

$$J^k = \sum_{j=0}^{s-1} \binom{k}{j} \lambda^{k-j} N^j.$$

For $A = PJP^{-1}$,

$$A^k = PJ^kP^{-1}.$$

Jordan form explains polynomial factors in matrix powers.

§77.7 Gershgorin discs

For $A = (a_{ij})$, define

$$R_i = \sum_{j \neq i} |a_{ij}|.$$

Every eigenvalue lies in

$$\bigcup_i \{z \in \mathbb{C} : |z - a_{ii}| \leq R_i\}.$$

Proof: choose i with $|x_i| = \max_j |x_j|$ for an eigenvector x . From $Ax = \lambda x$,

$$(\lambda - a_{ii})x_i = \sum_{j \neq i} a_{ij}x_j,$$

so $|\lambda - a_{ii}| \leq R_i$.

§77.8 Gershgorin figure




Figure 77.1: Gershgorin discs localize eigenvalues. A separated component containing one disc contains exactly one eigenvalue.

§77.9 Separated components and scaling

If a connected component of the union of Gershgorin discs consists of exactly m discs and is separated from the others, then it contains exactly m eigenvalues.

Column Gershgorin discs can be obtained by applying the theorem to A^T . Diagonal similarity

$$D^{-1}AD$$

does not change eigenvalues but changes radii, since

$$(D^{-1}AD)_{ij} = d_i^{-1}a_{ij}d_j.$$

This is a practical way to improve localization.

§77.10 Power iteration

If

$$|\lambda_1| > |\lambda_2| \geq \cdots,$$

and the initial vector has a nonzero component in the dominant eigenvector direction, then

$$u_{k+1} = \frac{Au_k}{\|Au_k\|}$$

converges toward the dominant eigenvector direction. The convergence rate is governed by

$$\left| \frac{\lambda_2}{\lambda_1} \right|.$$

§77.11 Inverse and shifted inverse iteration

Inverse iteration applies power iteration to A^{-1} , so it finds eigenvalues of small modulus. Shifted inverse iteration applies inverse iteration to

$$A - pI.$$

Eigenvalues of $(A - pI)^{-1}$ are

$$\frac{1}{\lambda_j - p}.$$

Thus the eigenvalue closest to p becomes dominant.

§77.12 Numerical traps

Shifted inverse iteration requires solving systems with $A - pI$, which may be nearly singular. This is why the method can be powerful and unstable at the same time. Pivoting and residual checks are essential.

Gershgorin discs are cheap localization, not exact computation. Overlap does not mean eigenvalues are evenly spread through the overlap.

§77.13 ML-facing interpretation

Power iteration appears in PageRank, spectral normalization, graph propagation, and monitoring dominant eigenvalue behavior. State-space and recurrent models depend on polynomial and spectral structure: the minimal polynomial, Jordan blocks, and spectral radius all describe memory and long-run dynamics.

§77.14 Choosing an eigenvalue tool

Tool	What it exposes	When to use it
Smith form	invariant factors	module or integer-structure questions
Jordan form	eigenvalue chains	powers, nilpotent parts, generalized eigenvectors
Gershgorin	discs containing eigenvalues	quick localization before computation
Power iteration	dominant eigenpair	one eigenvalue separated in modulus
Shifted inverse iteration	eigenvalue near a target	refining a chosen spectral region

The order matters in practice. Gershgorin gives a cheap first picture of where eigenvalues can be and can sometimes separate clusters. Power iteration is then useful only if the dominant eigenvalue is separated and the starting vector has a component in its eigendirection. Shifted inverse iteration reverses the question: after a target μ has been chosen, solving $(A - \mu I)y = x$ amplifies the eigendirection whose eigenvalue is closest to μ . The numerical traps in the note are exactly the failure modes of these assumptions: poor separation, defective eigenspaces, or a shift that makes the linear solve unstable.

Matrix Decompositions: LU, Cholesky, Full Rank, and SVD

N-20260602

§78.1 Why decompositions are algorithms

A decomposition rewrites a matrix into factors that are easier to compute with. LU reduces a system to triangular solves; Cholesky uses positive definiteness; full-rank decomposition exposes rank; SVD exposes singular directions.

§78.2 Doolittle LU

If $A \in \mathbb{C}^{n \times n}$ can be written as

$$A = LU,$$

where L is unit lower triangular and U is upper triangular, this is the Doolittle or LU decomposition.

§78.3 Existence and uniqueness

If the first $n - 1$ leading principal minors of A are nonzero, then $A = LU$ exists. If A is nonsingular, the decomposition is unique. The condition says that Gaussian elimination without row exchange never meets a zero pivot.

§78.4 Elimination matrices

At step k , an elimination matrix L_k zeros entries below the pivot. After all steps,

$$L_{n-1} \cdots L_1 A = U, \quad A = L_1^{-1} \cdots L_{n-1}^{-1} U.$$

The product of the inverses is the lower triangular factor.

§78.5 Doolittle formulas

Comparing $A = LU$ gives

$$u_{kj} = a_{kj} - \sum_{t=1}^{k-1} l_{kt} u_{tj}, \quad j \geq k,$$
$$l_{ik} = \frac{1}{u_{kk}} \left(a_{ik} - \sum_{t=1}^{k-1} l_{it} u_{tk} \right), \quad i > k.$$

The pivot u_{kk} must be nonzero.

§78.6 Pivoting

If a pivot is zero or small, row exchange is needed. With partial pivoting,

$$PA = LU,$$

where P is a permutation matrix. Pivoting controls large multipliers and improves numerical stability.

§78.7 Storage of LU

The upper triangle stores U ; the strictly lower triangle stores the multipliers l_{ij} . The diagonal of L is omitted because it is all ones.

§78.8 Cholesky decomposition

If A is Hermitian positive definite, there is a unique lower triangular L with positive diagonal such that

$$A = LL^H.$$

Solving $Ax = b$ becomes

$$Ly = b, \quad L^H x = y.$$

Cholesky costs about $\frac{1}{3}n^3 + O(n^2)$, roughly half of LU.

§78.9 Cholesky example

For

$$A = \begin{bmatrix} 4 & 2 \\ 2 & 5 \end{bmatrix},$$

we obtain

$$L = \begin{bmatrix} 2 & 0 \\ 1 & 2 \end{bmatrix}, \quad A = LL^T.$$

§78.10 Full-rank decomposition

If $\text{rank } A = r > 0$, a full-rank decomposition is

$$A = FG, \quad F \in \mathbb{C}^{m \times r}, \quad G \in \mathbb{C}^{r \times n},$$

with both factors rank r . It always exists and is not unique.

§78.11 Hermite standard form

Row operations reduce A to a Hermite-style form H . If pivot columns are $j_1, \dots, j_r$, choose F as the corresponding columns of A , and G as the nonzero rows of H . Then $A = FG$.

§78.12 Singular values

The singular values of A are

$$\sigma_i = \sqrt{\lambda_i(A^H A)}.$$

Unitary equivalence preserves singular values.

§78.13 Singular value decomposition

If $\text{rank } A = r$, then

$$A = U \begin{bmatrix} \Sigma & 0 \\ 0 & 0 \end{bmatrix} V^H, \quad \Sigma = \text{diag}(\sigma_1, \dots, \sigma_r).$$

SVD says that A rotates, stretches along orthogonal axes, and rotates again.

§78.14 SVD geometry




Figure 78.1: SVD describes the image of the unit ball as an ellipsoid.

§78.15 Which decomposition matches the matrix

Decomposition	Hypothesis or signal	What it gives
LU	square system without severe pivot trouble	triangular solves
$PA = LU$	row exchanges needed	stable Gaussian elimination
$A = LL^H$	Hermitian positive definite	fast symmetric solve
$A = FG$	$\text{rank } A = r$	factorization through an r -dimensional space
$A = U\Sigma V^H$	arbitrary rectangular matrix	singular directions and rank profile

The choice should be driven by structure rather than by habit. Cholesky is inappropriate without positive definiteness; ordinary LU may need pivoting when a small or zero pivot appears; a full-rank decomposition records the actual dimension through which the linear map factors. The SVD is the most geometric option: it says that the unit ball is rotated, stretched by the singular values, and rotated again. This is why the nonzero singular values detect rank and why small singular values signal directions in which solving or inversion will be unstable.

Generalized Inverses, Projections, and Least-Squares Geometry

N-20260603

§79.1 Why ordinary inverse is not enough

For $A \in \mathbb{C}^{m \times n}$, the system $Ax = b$ may be overdetermined, underdetermined, singular, or rank deficient. The Moore-Penrose inverse A^+ gives a unified optimal answer.

§79.2 Four regimes

<i>I</i>	$m = n$	$r = n$	ordinary inverse
<i>II</i>	$m > n$	$r = n$	least squares
<i>III</i>	$m < n$	$r = m$	minimum norm
<i>IV</i>	any	$r < \min(m, n)$	rank deficient

§79.3 Penrose equations

The Moore-Penrose inverse is the unique X satisfying

$$AXA = A, \quad XAX = X, \quad (AX)^H = AX, \quad (XA)^H = XA.$$

It is denoted A^+ .

§79.4 Fifteen generalized inverses

Choosing any nonempty subset of the four Penrose equations gives one of 15 generalized-inverse classes. Only the Moore-Penrose inverse is unique for every matrix.

§79.5 Projection transformations

If $\mathbb{C}^n = L \oplus M$, the projection onto L along M maps $x = y + z$ to y , where $y \in L, z \in M$.

§79.6 Idempotence and orthogonality

A matrix is a projection iff

$$P^2 = P.$$

It is an orthogonal projection iff

$$P^2 = P, \quad P^H = P.$$

§79.7 Column and row-space projections

AA^+ is the orthogonal projection onto $R(A)$. A^+A is the orthogonal projection onto $R(A^H)$.




Figure 79.1: The vector AA^+b is the orthogonal projection of b onto the column space.

§79.8 SVD computation

If

$$A = U \begin{bmatrix} \Sigma & 0 \\ 0 & 0 \end{bmatrix} V^H,$$

then

$$A^+ = V \begin{bmatrix} \Sigma^{-1} & 0 \\ 0 & 0 \end{bmatrix} U^H.$$

Only nonzero singular values are inverted.

§79.9 Full-rank computation

If $A = FG$ is a full-rank decomposition, then

$$A^+ = G^H(GG^H)^{-1}(F^H F)^{-1}F^H.$$

§79.10 Greville recursion

Greville's method updates A_k^+ when a new column is appended. It is useful for column-by-column or online data, although numerical stability still requires care.

§79.11 Iterative approximations

Iterative methods for A^+ often require a step size satisfying a condition such as

$$0 < \alpha < \frac{2}{\sigma_{\max}^2}.$$

Convergence is again a spectral-radius issue.

§79.12 Solvability

The system $Ax = b$ is solvable iff

$$AA^+b = b.$$

If it is solvable, the general solution is

$$x = A^+b + (I - A^+A)y.$$

§79.13 Least squares and minimum norm

For every b , A^+b is the least-squares solution minimizing $\|Ax - b\|_2$. Among all minimizers, it has minimum 2-norm.

§79.14 Identities

$$(A^+)^+ = A, \quad (A^H)^+ = (A^+)^H, \quad \text{rank } A^+ = \text{rank } A.$$

In general, $(AB)^+ \neq B^+A^+$ without additional assumptions.

§79.15 Reading the Moore–Penrose inverse geometrically

expression	geometric meaning	use
A^+	inverse on nonzero singular directions	canonical generalized inverse
AA^+	orthogonal projection onto $R(A)$	test the reachable part of b
A^+A	orthogonal projection onto $R(A^H)$	remove the null-space ambiguity
$AA^+b = b$	$b \in R(A)$	exact solvability
A^+b	minimum-norm least-squares vector	canonical approximate solution

The SVD formula makes the construction explicit: invert the nonzero singular values and leave the zero singular directions untouched. If b is in the column space, A^+b solves $Ax = b$ and chooses the solution orthogonal to $N(A)$. If b is not in the column space, AA^+b is the closest reachable vector in $R(A)$, so A^+b solves the least-squares problem with minimum norm among all least-squares solutions. This also explains why $(AB)^+ = B^+A^+$ is not generally safe: the projection geometry of the two factors need not line up.

Solving Linear Systems: Direct, Iterative, and Optimization Methods

N-20260604

§80.1 The computational problem

Solving $Ax = b$ can be done by direct methods, iterative methods, or optimization methods. Structure decides which method is appropriate.

§80.2 Triangular solves

For lower triangular L , forward substitution gives

$$y_i = \frac{1}{l_{ii}} \left(b_i - \sum_{j < i} l_{ij} y_j \right).$$

For upper triangular U , backward substitution gives

$$x_i = \frac{1}{u_{ii}} \left(y_i - \sum_{j > i} u_{ij} x_j \right).$$

§80.3 Gauss elimination

At step k , with pivot $a_{kk}^{(k)} \neq 0$, set

$$l_{ik} = -\frac{a_{ik}^{(k)}}{a_{kk}^{(k)}}.$$

Then update

$$a_{ij}^{(k+1)} = a_{ij}^{(k)} + l_{ik} a_{kj}^{(k)}, \quad b_i^{(k+1)} = b_i^{(k)} + l_{ik} b_k^{(k)}.$$

After $n - 1$ steps one solves an upper triangular system.

§80.4 Pivoting

Small pivots amplify rounding error. Partial pivoting swaps the largest available entry in the current column into the pivot position. Algebraically,

$$PA = LU.$$

§80.5 LU solution

If $A = LU$, solve

$$Ly = b, \quad Ux = y.$$

If $PA = LU$, solve

$$Ly = Pb, \quad Ux = y.$$

§80.6 Cholesky and LDL

For Hermitian positive definite A , use

$$A = LL^H.$$

The improved square-root method uses

$$A = LDL^H,$$

where L is unit lower triangular and D is positive diagonal. Then solve $Ly = b$, $Dz = y$, $L^Hx = z$.

§80.7 Tridiagonal chasing

A tridiagonal system can be solved in linear time by a specialized LU factorization. The forward sweep computes modified coefficients; the backward sweep recovers $x_n, x_{n-1}, \dots, x_1$.

§80.8 Stationary iteration

Rewrite $Ax = b$ as

$$x = Bx + f, \quad x^{(k+1)} = Bx^{(k)} + f.$$

The error satisfies

$$e^{(k+1)} = Be^{(k)}.$$

Thus convergence for every starting vector is equivalent to

$$\rho(B) < 1.$$

§80.9 Jacobi iteration

With $A = D + L + U$, Jacobi uses

$$x^{(k+1)} = -D^{-1}(L + U)x^{(k)} + D^{-1}b.$$

Componentwise,

$$x_i^{(k+1)} = \frac{1}{a_{ii}} \left(b_i - \sum_{j \neq i} a_{ij} x_j^{(k)} \right).$$

§80.10 Gauss-Seidel iteration

Gauss-Seidel uses the newest available values:

$$x^{(k+1)} = -(D + L)^{-1}Ux^{(k)} + (D + L)^{-1}b.$$

Componentwise,

$$x_i^{(k+1)} = \frac{1}{a_{ii}} \left(b_i - \sum_{j < i} a_{ij} x_j^{(k+1)} - \sum_{j > i} a_{ij} x_j^{(k)} \right).$$

§80.11 Diagonal dominance

If

$$|a_{ii}| > \sum_{j \neq i} |a_{ij}|$$

for every row, then Jacobi and Gauss-Seidel converge. Irreducible weak diagonal dominance gives another standard sufficient condition.

§80.12 SOR iteration

SOR introduces a relaxation parameter α :

$$B_\alpha = (D + \alpha L)^{-1}((1 - \alpha)D - \alpha U).$$

For real symmetric positive definite A , SOR converges iff

$$0 < \alpha < 2.$$

§80.13 Optimization viewpoint

If A is Hermitian positive definite, solving $Ax = b$ is equivalent to minimizing

$$\phi(x) = \frac{1}{2}x^H Ax - \operatorname{Re}(b^H x).$$

Its gradient is $Ax - b$.

§80.14 Steepest descent

Let $r_k = Ax_k - b$. Steepest descent uses

$$x_{k+1} = x_k - \alpha_k r_k, \quad \alpha_k = \frac{(r_k, r_k)}{(Ar_k, r_k)}.$$

It decreases the quadratic energy, but may be slow when $\kappa(A)$ is large.

§80.15 Method map




Figure 80.1: Linear systems can be solved by factorization, iteration, or minimization.

§80.16 How the linear-system method is selected

Method	Structure used	Warning
Gaussian elimination	general dense system	watch pivot size
Partial pivoting	row exchange freedom	controls unstable pivots
Cholesky	Hermitian positive definite A	fails without positivity
Chasing method	tridiagonal sparsity	do not destroy the band structure
Jacobi/Gauss–Seidel	$x^{(k+1)} = Bx^{(k)} + c$	need $\rho(B) < 1$
SOR	relaxed Gauss–Seidel	$0 < \alpha < 2$ for SPD convergence
Steepest descent	$\min \frac{1}{2}x^T Ax - b^T x$	zigzags when $\kappa(A)$ is large

Direct methods spend work to factor the matrix and then solve triangular systems; they are preferred when the factorization is affordable or reused for many right-hand sides. Iterative methods are chosen when the matrix is large or sparse, but their convergence is a spectral question about the iteration matrix. The optimization viewpoint is available for symmetric positive definite systems because solving $Ax = b$ is equivalent to minimizing a quadratic. Thus the method map at the end of the note is not a menu of names: it asks what structure the matrix has and which failure mode is most dangerous.

VIII

Calculus

Part VIII: Contents

N-20220315	531
81.1 A scattered algebraic note	531
81.2 A quadratic with parameters	531
81.3 A parabola and a hyperbola	532
81.4 A constrained extremum	532
81.5 Tangency as a double-root condition	533
81.6 Completing the square	533
81.7 A wider frame	533
81.8 Tangency as multiplicity	534
81.9 Constrained extrema and normal directions	534
81.10 Degeneracy and second-order data	534
81.11 Multiplicity, normality, and local geometry	534
N-20220319	535
82.1 Sequences and convergence	535
82.2 Limits of functions	535
82.3 Derivatives	536
82.4 Small-o and big-O	536
82.5 Integration	537
82.6 Why the epsilon definition is shaped this way	537
82.7 Continuity by sequences	538
82.8 Derivative as best linear approximation	538
82.9 Integral as accumulation	538
82.10 Convergence as topology	538
82.11 Derivative as a linear map	539
82.12 Returning to the original examples	539
N-20220320	541
83.1 From the mean value theorem to Cauchy's form	541
83.2 Continuity at a point	541
83.3 How this note fits the previous one	542
83.4 Rolle's theorem as the engine	542
83.5 Cauchy's theorem and L'Hopital's rule	543
83.6 Uniform continuity	543
83.7 Further direction	543
83.8 Mean value theorems as compactness plus local linearity	543
83.9 Cauchy's theorem and ratios of differentials	544
83.10 Uniform continuity and why closed intervals matter	544
83.11 From local slopes to global control	544
83.12 Mean value theorem via compactness	544
83.13 Lipschitz estimates and uniqueness	545
83.14 Taylor formula with remainder	545
N-20220409	547
84.1 Motivation	547
84.2 Faà di Bruno formula	547
84.3 Leibniz's differential notation	547
84.4 Parametric derivatives	548
84.5 Differentiating determinants	548
84.6 Matrix multiplication	548
84.7 Schur complement	549
84.8 What the pattern suggests	549
84.9 Faà di Bruno as chain rule for higher jets	549
84.10 Leibniz rules and derivations	549
84.11 Matrix multiplication as composition	549
84.12 Schur complements	550
84.13 Composition, derivations, and elimination	550

N-20220628	551
85.1 Linearity	551
85.2 Partial fractions	551
85.3 Integration by parts	552
85.4 Substitution	552
85.5 Partial fractions as algebra before calculus	553
85.6 Integration by parts as product rule reversed	553
85.7 A second lens	554
85.8 Partial fractions as algebra of poles	554
85.9 Integration by parts as adjointness	554
85.10 Substitution as pullback of one-forms	554
85.11 Integration techniques as structural transformations	554
85.12 Integration as reverse differentiation	554
85.13 Substitution as pullback	555
85.14 Integration by parts and adjoints	555
85.15 Partial fractions and residues	555
N-20221107	557
86.1 Trigonometric limits	557
86.2 Derivative of sine and cosine	557
86.3 Tangent, cotangent, secant, and cosecant	557
86.4 Inverse trigonometric derivatives	558
86.5 Examples from the handwritten page	558
86.6 A closing lens	558
86.7 The circle as a Lie group	558
86.8 Trigonometric derivatives from the exponential	559
86.9 Inverse trigonometric branches	559
86.10 Trigonometry from circle symmetry and branches	559
N-20221211	561
87.1 Average velocity and instantaneous velocity	561
87.2 Derivative notation	561
87.3 Power functions	561
87.4 Exponential and logarithmic derivatives	562
87.5 Trigonometric derivatives	562
87.6 Inverse trigonometric derivatives	562
87.7 Examples	563
87.8 Context and outlook	563
87.9 Derivative as first-order linearization	563
87.10 Inverse trigonometric derivatives and the inverse function theorem	563
87.11 Endpoint blow-up and failure of local invertibility	564
87.12 Logarithm as inverse and as group homomorphism	564
87.13 Automatic differentiation and computation graphs	564
87.14 Jets and higher-order local data	565
87.15 Differentials and one-forms	565
87.16 Derivatives as linearization and pullback	565
N-20230310	567
88.1 The derivative from secants to tangents	567
88.2 Derivative notation and local definition	568
88.3 Rules of differentiation	568
88.4 Elementary derivative formulas	568
88.5 Extrema, Fermat's lemma, and critical points	569
88.6 A derivative proof of a basic inequality	569
88.7 Young's inequality	569
88.8 Holder's inequality	570
88.9 Minkowski's inequality	570
88.10 Convex functions and Jensen's inequality	570
88.11 Typical exercise themes	571
88.12 The reusable pattern	572
88.13 Derivative as local linear approximation	572

88.14	Extrema and first variation	572
88.15	Young, Holder, and convex duality	572
88.16	Local linearity and convex control	572
N-20251112		573
89.1	Cauchy–Schwarz and Minkowski	573
89.2	Neighborhood notation	573
89.3	Trigonometric identities	573
89.4	Hyperbolic functions	574
89.5	Sequence limits: epsilon–N language	574
89.6	Infinite limits and one-sided limits	575
89.7	Heine criterion for function limits	575
89.8	Infinitesimals and infinitesimal orders	575
89.9	Squeeze theorem and monotone convergence	576
89.10	Cauchy criteria	576
89.11	Continuity	576
89.12	Derivatives and differentials	577
89.13	Taylor-type expansions	577
89.14	Where the calculation points	577
89.15	Limits through filters and sequences	578
89.16	Hyperbolic functions and exponential coordinates	578
89.17	Cauchy–Schwarz, Minkowski, and normed spaces	578
89.18	Calculus review through limits and norm geometry	578
N-20251210		579
90.1	Strict monotonicity from the derivative	579
90.2	First derivative test for extrema	579
90.3	Second derivative test	579
90.4	Concavity and inflection points	580
90.5	Asymptotes	580
90.6	General procedure for sketching a curve	581
90.7	Basic antiderivative table	581
90.8	Substitution	581
90.9	Integration by parts	582
90.10	Rational functions and partial fractions	582
90.11	Quadratic denominator reduction	582
90.12	Trigonometric rational functions	583
90.13	Definite integrals and Riemann sums	583
90.14	Basic properties of definite integrals	584
90.15	Integral mean value theorem	584
90.16	Applications: area, volume, and arc length	584
90.17	Improper integrals	585
90.18	Broader perspective	585
90.19	Monotonicity and derivative sign	585
90.20	Convexity and supporting lines	586
90.21	Integration methods as inverse rules	586
90.22	Error terms and numerical integration	586
90.23	Local derivative data and global shape	586
N-20260106		587
91.1	Convexity and the Hermite–Hadamard inequality	587
91.2	Kantorovich-type integral inequality	587
91.3	Integral mean value with odd symmetry	588
91.4	A lower bound involving f double-prime over f	589
91.5	Symmetry of integrals about the midpoint	589
91.6	Midpoint and trapezoidal error terms	590
91.7	First special integral	590
91.8	Positivity of a Fresnel-type integral	591
91.9	A symmetry integral with arcsine	591
91.10	A logarithmic sine integral	592
91.11	Integrals of $1/(a + b \cos x)$	592

91.12	Wallis integral and Wallis product	593
91.13	A wider frame	593
91.14	Convexity as second-order positivity	593
91.15	Integral inequalities as weighted averages	593
91.16	Symmetry and cancellation	594
91.17	Wallis formula and asymptotic products	594
91.18	Integral inequalities through convexity and asymptotics	594

Quadratic Parameters, Tangency, and Constrained Extrema

N-20220315

§81.1 A scattered algebraic note

The original page is titled “some scattered things.” The common theme is that elementary-looking algebraic problems often hide a parameter, a tangency condition, or an extremum under constraint. I rewrite the problems in a more systematic form.

§81.2 A quadratic with parameters

Let

$$f(x) = ax^2 + bx + c, \quad a < b < c,$$

and suppose the graph passes through $A = (1, 0)$ and $B = (m, -a)$. The condition $f(1) = 0$ gives

$$a + b + c = 0.$$

The condition $f(m) = -a$ is equivalent to saying that $f(x) + a$ has $x = m$ as a root:

$$am^2 + bm + c + a = 0.$$

Using $c = -a - b$, this becomes

$$am^2 + bm - b = 0, \quad am^2 = -b(m - 1).$$

The axis of symmetry is

$$x_0 = -\frac{b}{2a}.$$

Thus the parameter m and the axis x_0 are tied by

$$x_0 = \frac{m^2}{2m - 2}.$$

Equivalently,

$$m^2 = (2m - 2)x_0, \quad m^2 - 2x_0m + 2x_0 = 0.$$

Solving for m gives

$$m = x_0 \pm \sqrt{x_0^2 - 2x_0}.$$

The original calculation uses this relation to examine endpoints and stationary points of m as a function of x_0 . The conclusion is that the extremal parameter occurs at

$$x_0 = 0 \quad \text{or} \quad x_0 = -\frac{1}{2},$$

and the relevant value is

$$m = \frac{\sqrt{5} - 1}{2}.$$

This is a small example of a recurring trick: convert the algebraic parameter into a geometric parameter, usually the axis, vertex, or tangency point.

§81.3 A parabola and a hyperbola

The second problem considers the parabola

$$y = x^2 - 2ax + a^2 - 4 = (x - a)^2 - 4$$

and the rectangular hyperbola

$$y = \frac{k}{x}, \quad k > 0,$$

in the first quadrant. The intersection equation is

$$x^2 - 2ax + a^2 - 4 = \frac{k}{x}, \quad x > 0,$$

or

$$x^3 - 2ax^2 + (a^2 - 4)x - k = 0.$$

Questions about “only one intersection” are therefore tangency or root-counting questions for this cubic, but the geometry is usually cleaner: one restricts to the right side of the parabola, because the branch k/x lies in the first quadrant.

Remark 81.3.1 — The hidden principle is that a graph-intersection problem may become either a polynomial root-counting problem or a tangency problem. The latter is often more geometric and less computational.

§81.4 A constrained extremum

The third problem asks for extrema of

$$x + y$$

under the constraint

$$x^2 + y^3 = 2.$$

Using Lagrange multipliers for

$$F(x, y, \lambda) = x + y + \lambda(x^2 + y^3 - 2),$$

we get

$$1 + 2\lambda x = 0, \quad 1 + 3\lambda y^2 = 0, \quad x^2 + y^3 = 2.$$

Eliminating λ gives

$$2x = 3y^2.$$

Hence

$$x = \frac{3}{2}y^2.$$

Substitution into the constraint gives

$$\frac{9}{4}y^4 + y^3 = 2,$$

so the stationary points come from the quartic

$$9y^4 + 4y^3 - 8 = 0.$$

The numerical plot in the original note marks approximately

$$A = (1.15195, 2.02829), \quad B = (1.82947, 0.72509),$$

for the function

$$h(x) = x + (2 - x^2)^{1/3}.$$

These are the local maximum and local minimum along the real branch.




Figure 81.1: Reconstruction of the graph used for the constrained extremum problem.

§81.5 Tangency as a double-root condition

A parameter problem involving a quadratic often asks when a line is tangent to a parabola or when an equation has exactly one solution. Algebraically, this is a discriminant condition:

$$\Delta = b^2 - 4ac = 0.$$

Geometrically, the line touches the parabola at one point. The same idea appears in constrained extrema: a level curve just touches a constraint curve at an optimum.

§81.6 Completing the square

For a quadratic

$$ax^2 + bx + c, \quad a \neq 0,$$

completing the square gives

$$a \left(x + \frac{b}{2a} \right)^2 + c - \frac{b^2}{4a}.$$

This form reveals the vertex, minimum or maximum value, and symmetry axis. It is usually more informative than the expanded form.

§81.7 A wider frame

The three examples are connected by the same idea: reduce a complicated statement to a structural condition. The axis of a quadratic, the tangency of two graphs, and the Lagrange multiplier equations are all ways of saying that an extremum occurs when a first-order freedom disappears.

§81.8 Tangency as multiplicity

A line is tangent to a curve when the intersection equation has a repeated root. If a family is written as

$$F(x, t) = 0,$$

then a tangency or transition value often satisfies

$$F(x, t) = 0, \quad \frac{\partial F}{\partial x}(x, t) = 0.$$

For a quadratic, this is the discriminant condition $\Delta = 0$. The same idea persists for higher-degree equations: a repeated root is detected by a common factor of F and F_x .

§81.9 Constrained extrema and normal directions

For a constraint $g(x, y) = 0$, the Lagrange multiplier condition

$$\nabla f = \lambda \nabla g$$

says that at an interior constrained extremum, the gradient of f has no component tangent to the constraint curve. Equivalently, the differential df vanishes on the tangent space of the constraint.

This is the coordinate-free meaning of the method: extrema occur where the objective's first-order change is normal to every allowed first-order motion.

§81.10 Degeneracy and second-order data

After first derivatives vanish, the next question is second-order behavior. For an unconstrained two-variable function, the Hessian quadratic form determines the local model:

$$f(a + h) = f(a) + \frac{1}{2}h^T H_f(a)h + o(\|h\|^2).$$

For constrained problems, the Hessian must be restricted to tangent directions. Degenerate cases occur when this quadratic form has zero directions, forcing higher-order analysis.

§81.11 Multiplicity, normality, and local geometry

Quadratic parameter problems, tangency, and constrained extrema are one theme: detect where the number or type of solutions changes. Algebra sees this as multiplicity or discriminant zero; geometry sees it as tangency; optimization sees it as vanishing first variation plus second-order testing.

A Review of Calculus: Sequences, Limits, Derivatives, and Integrals

N-20220319

§82.1 Sequences and convergence

After defining real numbers, the first analytic object to understand is a sequence. Intuitively,

$$1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \dots$$

should converge to 0, while

$$1, 2, 3, 4, \dots$$

should diverge. The difficulty is to turn this intuition into a definition that is precise.

Definition 82.1.1. A sequence is a function

$$a : \mathbb{N} \rightarrow \mathbb{R}$$

or more generally $a : \mathbb{N} \rightarrow \mathbb{C}$. We write a_n instead of $a(n)$, and denote the sequence by (a_n) or $(a_n)_{n=1}^{\infty}$.

Definition 82.1.2. Let (a_n) be a real sequence and let $\ell \in \mathbb{R}$. We say $a_n \rightarrow \ell$ if for every $\varepsilon > 0$ there exists $N \in \mathbb{N}$ such that whenever $n \geq N$,

$$|a_n - \ell| < \varepsilon.$$

Equivalently,

$$(\forall \varepsilon > 0)(\exists N)(\forall n \geq N) |a_n - \ell| < \varepsilon.$$

Example 82.1.3

Let

$$x_n = \frac{n}{n+1}.$$

Then $x_n \rightarrow 1$. Indeed,

$$|x_n - 1| = \frac{1}{n+1}.$$

Given $\varepsilon > 0$, choose $N > 1/\varepsilon - 1$. Then $n > N$ implies $|x_n - 1| < \varepsilon$.

§82.2 Limits of functions

Definition 82.2.1. Let $A \subseteq \mathbb{R}$, $f : A \rightarrow \mathbb{R}$, and a be a limit point of A . We write

$$\lim_{x \rightarrow a} f(x) = \ell$$

if

$$(\forall \varepsilon > 0)(\exists \delta > 0)(\forall x \in A)(0 < |x - a| < \delta \Rightarrow |f(x) - \ell| < \varepsilon).$$

The condition $0 < |x - a|$ is important: the value of $f(a)$ is irrelevant, and f does not even have to be defined at a .

§82.3 Derivatives

Definition 82.3.1. The derivative of f at x is

$$\frac{df}{dx}(x) = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h},$$

provided the limit exists.

Theorem 82.3.2 (Chain rule)

If $f(x) = F(g(x))$, then

$$\frac{df}{dx} = \frac{dF}{dg} \frac{dg}{dx}.$$

Theorem 82.3.3 (Product rule)

If $f(x) = u(x)v(x)$, then

$$f'(x) = u'(x)v(x) + u(x)v'(x).$$

Theorem 82.3.4 (Quotient rule)

If $f(x) = u(x)/v(x)$, then

$$f'(x) = \frac{u'(x)v(x) - u(x)v'(x)}{v(x)^2}.$$

Theorem 82.3.5 (Leibniz rule)

If $f = uv$, then

$$f^{(n)}(x) = \sum_{r=0}^n \binom{n}{r} u^{(r)}(x) v^{(n-r)}(x).$$

§82.4 Small-o and big-O

Definition 82.4.1. As $x \rightarrow x_0$, write $f(x) = o(g(x))$ if

$$\lim_{x \rightarrow x_0} \frac{f(x)}{g(x)} = 0.$$

Write $f(x) = O(g(x))$ if $f(x)/g(x)$ remains bounded near x_0 .

Thus $o(g)$ means much smaller than g , while $O(g)$ means no larger than a constant multiple of g . Clearly,

$$f = o(g) \Rightarrow f = O(g).$$

§82.5 Integration

An integral is a limit of sums. The usual Riemann integral is modeled by

$$\int_a^b f(x) dx = \lim_{\Delta x \rightarrow 0} \sum_{n=0}^N f(x_n) \Delta x.$$

For an interval partition with $\Delta x = (b-a)/N$ and $x_n = a + n\Delta x$, this approximates the area under the graph.

If f is differentiable, then over a small interval the area is

$$f(x_n)\Delta x + O((\Delta x)^2).$$

Summing and letting $N \rightarrow \infty$ gives the usual integral.

Theorem 82.5.1 (Fundamental theorem of calculus)

Let

$$F(x) = \int_a^x f(t) dt.$$

Then, under the usual continuity hypothesis on f ,

$$F'(x) = f(x).$$

Remark 82.5.2 — This is a review note rather than a full proof-based analysis text. The next revision should add precise hypotheses for every theorem, especially differentiability and continuity assumptions.

§82.6 Why the epsilon definition is shaped this way

The definition of convergence has the order of quantifiers

$$\forall \varepsilon > 0 \exists N \forall n \geq N.$$

This order is essential. We are not allowed to choose one fixed tolerance; the sequence must eventually enter every tolerance window around the limit. The number N may depend on ε , because smaller error demands later terms.

For example, to prove $1/n \rightarrow 0$, given $\varepsilon > 0$, choose $N > 1/\varepsilon$. Then $n \geq N$ implies

$$\left| \frac{1}{n} - 0 \right| < \varepsilon.$$

The proof is not a calculation of the limit; it is a demonstration that every requested precision can eventually be met.

§82.7 Continuity by sequences

A function f is continuous at a if

$$x_n \rightarrow a \implies f(x_n) \rightarrow f(a)$$

for every sequence (x_n) in the domain. This sequence criterion is often easier to use than epsilon-delta directly. To prove discontinuity, it suffices to find one sequence $x_n \rightarrow a$ such that $f(x_n) \not\rightarrow f(a)$.

For example, the Dirichlet function is discontinuous everywhere. Near any a , choose rational and irrational sequences both converging to a . Their function values approach different limits.

§82.8 Derivative as best linear approximation

The derivative is often introduced as a slope:

$$f'(a) = \lim_{h \rightarrow 0} \frac{f(a+h) - f(a)}{h}.$$

A more structural reading is that differentiability means

$$f(a+h) = f(a) + f'(a)h + o(h).$$

Thus the derivative is the coefficient of the best linear approximation near a . This viewpoint generalizes directly to multivariable calculus, where the derivative becomes a linear map.

§82.9 Integral as accumulation

The definite integral can be understood as the limit of accumulated small contributions. For continuous f , Riemann sums

$$\sum_i f(\xi_i) \Delta x_i$$

converge to the area-signed accumulation. The fundamental theorem of calculus says that differentiation and integration are inverse operations in the sense that

$$\frac{d}{dx} \int_a^x f(t) dt = f(x).$$

This is not merely an area formula; it says local rate and global accumulation determine one another.

§82.10 Convergence as topology

The elementary epsilon definition in this note is the first appearance of a topology. On the real line, saying

$$x_n \rightarrow x$$

means that every interval around x eventually contains all x_n . In a metric space (X, d) the same sentence becomes

$$\forall \varepsilon > 0, \quad \exists N, \quad n > N \Rightarrow d(x_n, x) < \varepsilon.$$

Thus the elementary sequence definition is not a special trick for numbers; it is the definition of convergence once a notion of distance has been chosen.

This immediately explains why different spaces have different convergence. In $C[0, 1]$ with the sup norm,

$$\|f_n - f\|_\infty = \max_{x \in [0, 1]} |f_n(x) - f(x)| \rightarrow 0$$

means uniform convergence. With the L^1 norm,

$$\|f_n - f\|_1 = \int_0^1 |f_n(x) - f(x)| dx \rightarrow 0,$$

thin spikes may disappear in integral size while still being tall. The same elementary sentence “goes to zero” therefore depends on the surrounding space.

§82.11 Derivative as a linear map

In one-variable calculus, the derivative is introduced as a number:

$$f'(a) = \lim_{h \rightarrow 0} \frac{f(a+h) - f(a)}{h}.$$

The higher-level meaning is that $f'(a)h$ is the best first-order linear approximation to the change in f :

$$f(a+h) = f(a) + f'(a)h + o(h).$$

For a map $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$, the derivative at a is a linear map $DF(a)$ satisfying

$$F(a+h) = F(a) + DF(a)h + o(|h|).$$

The Jacobian matrix is only the coordinate representation of this linear map.

§82.12 Returning to the original examples

The elementary limit and derivative computations in the original note should be read in two layers. At the calculation layer, they train epsilon estimates and algebraic simplification. At the structural layer, they introduce topology and linearization. A limit asks whether a function is continuous with respect to the chosen topology; a derivative asks whether the function is locally a linear map plus a smaller error.

A Review of Calculus: Mean Value Theorems and Continuity

N-20220320

§83.1 From the mean value theorem to Cauchy's form

This note continues the calculus review. The main topic is the mean value theorem and its Cauchy form.

Theorem 83.1.1 (Mean value theorem)

Let $f : [a, b] \rightarrow \mathbb{R}$ be continuous on $[a, b]$ and differentiable on (a, b) . Then there exists $c \in (a, b)$ such that

$$f'(c) = \frac{f(b) - f(a)}{b - a}.$$

Geometrically, there is a point at which the tangent slope equals the slope of the secant line through the endpoints.

Theorem 83.1.2 (Cauchy's mean value theorem)

Let $f, g : [a, b] \rightarrow \mathbb{R}$ be continuous on $[a, b]$ and differentiable on (a, b) . If $g'(x) \neq 0$ on (a, b) , then there exists $c \in (a, b)$ such that

$$\frac{f'(c)}{g'(c)} = \frac{f(b) - f(a)}{g(b) - g(a)}.$$

Equivalently,

$$(f(b) - f(a))g'(c) = (g(b) - g(a))f'(c).$$

Remark 83.1.3 — This is the version sometimes called Cauchy's mean value theorem. It is the natural two-function generalization of the ordinary mean value theorem and is often the cleanest route to l'Hospital-type arguments.

§83.2 Continuity at a point

The note then returns to continuity.

Definition 83.2.1. Let f be a real-valued function defined in a neighborhood of $a \in \mathbb{R}$. The function f is continuous at a if

$$\lim_{x \rightarrow a} f(x) = f(a).$$

Equivalently, for every $\varepsilon > 0$ there is $\delta > 0$ such that

$$|x - a| < \delta \Rightarrow |f(x) - f(a)| < \varepsilon.$$

Remark 83.2.2 — The intuitive statement is that when x is close to a , the value $f(x)$ is close to $f(a)$. The epsilon-delta form is the precise version of that intuition.

§83.3 How this note fits the previous one

The previous review note introduced limits, derivatives, big-O/small-o notation, and integration. This continuation adds the mean value theorem and continuity, which are the bridge between local derivative information and global information about a function on an interval.

The original text layer for this note was short and duplicated in places, so the missing proof details are filled in explicitly below. The key point is that both the ordinary mean value theorem and Cauchy's mean value theorem are consequences of Rolle's theorem: one subtracts a secant line in the first case, and a suitable linear combination of two functions in the second.

§83.4 Rolle's theorem as the engine

Theorem 83.4.1 (Rolle's theorem)

Let f be continuous on $[a, b]$, differentiable on (a, b) , and suppose $f(a) = f(b)$. Then there exists $c \in (a, b)$ such that

$$f'(c) = 0.$$

The mean value theorem is usually proved from Rolle's theorem. Rolle says that if $f(a) = f(b)$, and f is continuous on $[a, b]$ and differentiable on (a, b) , then there is $c \in (a, b)$ such that

$$f'(c) = 0.$$

Geometrically, a curve beginning and ending at the same height must have a horizontal tangent somewhere.

To prove the mean value theorem, subtract the secant line. Define

$$h(x) = f(x) - \frac{f(b) - f(a)}{b - a}(x - a).$$

Then $h(a) = h(b)$. Rolle's theorem gives $h'(c) = 0$, which is exactly

$$f'(c) = \frac{f(b) - f(a)}{b - a}.$$

This proof explains the theorem better than the formula alone: the secant line is removed so that the endpoints line up.

§83.5 Cauchy's theorem and L'Hopital's rule

Theorem 83.5.1 (Cauchy's mean value theorem)

Let f, g be continuous on $[a, b]$ and differentiable on (a, b) . If $g'(x) \neq 0$ on (a, b) , then there exists $c \in (a, b)$ such that

$$\frac{f(b) - f(a)}{g(b) - g(a)} = \frac{f'(c)}{g'(c)}.$$

Cauchy's mean value theorem applies Rolle's theorem to a linear combination of two functions. It is the natural source of L'Hopital's rule. Suppose

$$\lim_{x \rightarrow a} f(x) = \lim_{x \rightarrow a} g(x) = 0$$

and $g'(x) \neq 0$. Applying Cauchy's theorem to f and g on $[a, x]$ gives a point c between a and x such that

$$\frac{f(x) - f(a)}{g(x) - g(a)} = \frac{f'(c)}{g'(c)}.$$

If $f'(c)/g'(c)$ tends to a limit as $x \rightarrow a$, then so does $f(x)/g(x)$. This derivation shows why L'Hopital's rule has hypotheses; it is not a symbolic cancellation rule.

§83.6 Uniform continuity

Continuity at each point allows δ to depend on both ε and the point. Uniform continuity allows δ to depend only on ε . On a compact interval, every continuous function is uniformly continuous.

This fact is often proved by contradiction using sequences, or by compactness using finite subcovers. It is the hidden reason many limiting arguments on closed intervals work smoothly.

§83.7 Further direction

Rolle, the mean value theorem, and Cauchy's mean value theorem are not isolated formulas. They are mechanisms for turning global information on an interval into local derivative information at some point. This is the central bridge of one-variable calculus: endpoint data forces a tangent slope somewhere in between.

§83.8 Mean value theorems as compactness plus local linearity

Theorem 83.8.1 (Mean value theorem)

If f is continuous on $[a, b]$ and differentiable on (a, b) , then there exists $c \in (a, b)$ such that

$$f(b) - f(a) = f'(c)(b - a).$$

Rolle's theorem and the mean value theorem combine two ingredients: continuity on a compact interval and differentiability inside. Compactness gives a maximum and minimum. Differentiability says that an interior extremum has derivative zero.

The result

$$f(b) - f(a) = f'(c)(b - a)$$

means that some instantaneous slope equals the average slope. It is not a numerical accident; it is forced by the existence of extrema on a closed interval.

§83.9 Cauchy's theorem and ratios of differentials

Theorem 83.9.1 (Cauchy's theorem, ratio form)

Under the hypotheses of Cauchy's mean value theorem, some $c \in (a, b)$ satisfies

$$(f(b) - f(a))g'(c) = (g(b) - g(a))f'(c).$$

Cauchy's mean value theorem says that for suitable f, g , there is c with

$$\frac{f(b) - f(a)}{g(b) - g(a)} = \frac{f'(c)}{g'(c)}.$$

This is the rigorous ancestor of many quotient-limit rules. L'Hopital's rule is a limiting version in which the interval endpoint moves toward a singular point.

Thus quotient estimates should be read as comparing two functions through their first-order changes.

§83.10 Uniform continuity and why closed intervals matter

Continuity at each point gives a local δ . Uniform continuity asks for one δ that works everywhere. Heine–Cantor says that a continuous function on a compact set is uniformly continuous.

The theorem explains why many epsilon proofs on closed intervals feel easier: compactness lets finitely many local controls become one global control.

§83.11 From local slopes to global control

The mean value theorem is the bridge from local derivative information to global estimates. Rolle gives the zero-slope engine, Cauchy compares two functions, L'Hopital studies singular ratios, and uniform continuity records when local control can be made global.

§83.12 Mean value theorem via compactness

Remark 83.12.1 — The compact interval hypothesis is not decorative: continuity on $[a, b]$ gives maxima and minima, and differentiability inside the interval turns an extremum into a zero derivative through Rolle's theorem.

Rolle's theorem and the mean value theorem are often memorized as calculus facts, but their proof rests on two deeper one-dimensional principles. First, a continuous function on a closed interval attains a maximum and a minimum; this is compactness. Second, differentiability turns information at an extremum into the equation $f'(\xi) = 0$.

For the usual mean value theorem, define

$$g(x) = f(x) - \frac{f(b) - f(a)}{b - a}(x - a).$$

Then $g(a) = g(b)$, so Rolle's theorem gives $g'(\xi) = 0$, hence

$$f(b) - f(a) = f'(\xi)(b - a).$$

The theorem is therefore a statement that a function and its secant line must have a parallel tangent somewhere.

§83.13 Lipschitz estimates and uniqueness

A powerful corollary is the Lipschitz estimate. If $|f'(x)| \leq M$ on an interval, then

$$|f(x) - f(y)| \leq M|x - y|.$$

This turns derivative bounds into metric bounds.

This idea reappears in differential equations. Suppose $y' = F(t, y)$ and F is Lipschitz in y :

$$|F(t, u) - F(t, v)| \leq L|u - v|.$$

If y_1, y_2 are two solutions with the same initial value, then

$$|y_1(t) - y_2(t)| \leq \int_{t_0}^t L|y_1(s) - y_2(s)| ds.$$

Gronwall's inequality forces $y_1 = y_2$. Thus a tangent-slope theorem becomes the analytic mechanism behind uniqueness of solutions.

§83.14 Taylor formula with remainder

The mean value theorem also generates Taylor's theorem. For example,

$$f(x) = f(a) + f'(a)(x - a) + \frac{f''(\xi)}{2}(x - a)^2$$

for some ξ between a and x . The remainder measures the error of replacing a function by its local polynomial model. In numerical analysis this error term controls approximation; in geometry it is the beginning of the language of jets.

Faà di Bruno, Leibniz Rules, and Matrix Multiplication

N-20220409

§84.1 Motivation

This note jumps among several topics: high-order derivatives of composite functions, Leibniz-type determinant differentiation, and the beginning of matrix multiplication. The unifying theme is that complicated formulas become manageable once the indexing convention is chosen correctly.

§84.2 Faà di Bruno formula

For one-variable functions, the n -th derivative of a composite $f \circ g$ is governed by Faà di Bruno's formula:

$$(f \circ g)^{(n)}(x) = \sum_{m=1}^n \sum_{(k_1, \dots, k_n) \in \Gamma_{m,n}} \frac{n!}{k_1! \cdots k_n!} f^{(m)}(g(x)) \prod_{j=1}^n \left(\frac{g^{(j)}(x)}{j!} \right)^{k_j},$$

where

$$\Gamma_{m,n} = \{(k_1, \dots, k_n) : k_j \in \mathbb{Z}_{\geq 0}, k_1 + \cdots + k_n = m, k_1 + 2k_2 + \cdots + nk_n = n\}.$$

This is the high-order chain rule. The indexing records how the n derivatives are partitioned among the derivatives of g .

§84.3 Leibniz's differential notation

If $y = f(x)$, then

$$dy = f'(x) dx.$$

Taking another differential gives

$$d^2y = d(f'(x)dx) = f''(x)dx^2 + f'(x)d^2x.$$

If $x = x(t)$, this corresponds to the ordinary chain rule

$$\frac{dy}{dt} = f'(x(t))x'(t).$$

For the third differential, one obtains

$$d^3y = f'''(x)dx^3 + 3f''(x)dx d^2x + f'(x)d^3x.$$

This is already a visible shadow of Faà di Bruno.

§84.4 Parametric derivatives

If a curve is given parametrically by

$$x = x(t), \quad y = y(t),$$

then

$$y' = \frac{dy}{dx} = \frac{dy/dt}{dx/dt}.$$

The second derivative is

$$y'' = \frac{x'(t)y''(t) - y'(t)x''(t)}{(x'(t))^3}.$$

This can be written as a determinant:

$$y'' = \frac{\begin{vmatrix} x' & y' \\ x'' & y'' \end{vmatrix}}{(x')^3}.$$

§84.5 Differentiating determinants

For an $n \times n$ matrix of functions $A(x) = (f_{ij}(x))$, the derivative of the determinant is obtained by differentiating one row at a time:

$$\frac{d}{dx} \det A(x) = \sum_{i=1}^n \det \begin{pmatrix} f_{11} & \cdots & f_{1n} \\ \vdots & & \vdots \\ f'_{i1} & \cdots & f'_{in} \\ \vdots & & \vdots \\ f_{n1} & \cdots & f_{nn} \end{pmatrix}.$$

This is simply the multilinearity of the determinant.

§84.6 Matrix multiplication

If

$$A \in M_{m \times n}, \quad B \in M_{n \times p},$$

then their product $C = AB \in M_{m \times p}$ has entries

$$C_{ij} = \sum_{k=1}^n a_{ik} b_{kj}.$$

The note emphasizes that in general $AB \neq BA$, while associativity still holds:

$$A(BC) = (AB)C.$$

For block matrices,

$$\begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix} = \begin{pmatrix} A_{11}B_{11} + A_{12}B_{21} & A_{11}B_{12} + A_{12}B_{22} \\ A_{21}B_{11} + A_{22}B_{21} & A_{21}B_{12} + A_{22}B_{22} \end{pmatrix}.$$

§84.7 Schur complement

The block elimination identity in the note is

$$\begin{pmatrix} I & 0 \\ -CA^{-1} & I \end{pmatrix} \begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} A & B \\ 0 & D - CA^{-1}B \end{pmatrix}.$$

The matrix

$$D - CA^{-1}B$$

is the Schur complement of A in the block matrix. It is the block-matrix version of Gaussian elimination.

§84.8 What the pattern suggests

The note moves from high-order calculus to matrix algebra, but the same mathematical taste appears in both places: the right notation makes the formula intelligible. Multi-indices organize Faà di Bruno, determinants organize parametric derivatives, and block matrices organize elimination.

§84.9 Faà di Bruno as chain rule for higher jets

The ordinary chain rule controls first derivatives of a composition. Faà di Bruno controls all higher derivatives and is naturally expressed through partitions:

$$\frac{d^n}{dx^n} f(g(x)) = \sum f^{(k)}(g(x)) \prod_j \left(\frac{g^{(j)}(x)}{j!} \right)^{m_j} \frac{n!}{\prod_j m_j!}.$$

The combinatorics records how n differentiations are distributed among the inner function.

§84.10 Leibniz rules and derivations

A derivation D satisfies

$$D(ab) = D(a)b + aD(b).$$

Higher Leibniz formulas are binomial expansions of repeated derivations:

$$D^n(ab) = \sum_{k=0}^n \binom{n}{k} D^k(a) D^{n-k}(b).$$

§84.11 Matrix multiplication as composition

Matrix multiplication is the coordinate expression of composing linear maps. The index formula

$$(AB)_{ij} = \sum_k A_{ik} B_{kj}$$

contracts the output coordinate of B with the input coordinate of A .

§84.12 Schur complements

For a block matrix, the Schur complement

$$D - CA^{-1}B$$

records what remains after eliminating the variables controlled by A . It is block Gaussian elimination, and it appears in determinants, inverses, covariance matrices, and constrained quadratic forms.

§84.13 Composition, derivations, and elimination

The note links differentiation and linear algebra through composition. Faà di Bruno is higher chain rule, Leibniz is derivation algebra, matrix multiplication is composition, and Schur complements are elimination in block form.

Techniques of Integration: Partial Fractions, Parts, and Substitution

N-20220628

§85.1 Linearity

The note begins with the linearity of the indefinite integral:

$$\int (f(x) + g(x) - h(x)) dx = \int f(x) dx + \int g(x) dx - \int h(x) dx.$$

This is the integration analogue of the linearity of differentiation.

§85.2 Partial fractions

The first worked example is

$$\int \frac{dx}{x^2 - a^2}.$$

Since

$$\frac{1}{x^2 - a^2} = \frac{1}{(x - a)(x + a)} = \frac{1}{2a} \left(\frac{1}{x - a} - \frac{1}{x + a} \right),$$

we get

$$\int \frac{dx}{x^2 - a^2} = \frac{1}{2a} \ln \left| \frac{x - a}{x + a} \right| + C.$$

A general quadratic denominator is handled by completing the square. For

$$\int \frac{mx + n}{x^2 + px + q} dx,$$

put

$$t = x + \frac{p}{2}, \quad x^2 + px + q = t^2 + q - \frac{p^2}{4}.$$

Then write $mx + n = At + B$. The integral splits into

$$A \int \frac{t}{t^2 + a^2} dt + B \int \frac{dt}{t^2 + a^2}$$

or, if the constant term is negative,

$$A \int \frac{t}{t^2 - a^2} dt + B \int \frac{dt}{t^2 - a^2}.$$

This produces logarithms and arctangents depending on the sign of $q - p^2/4$.

A page warns about a common partial-fraction mistake: repeated factors must be treated separately. For example,

$$\int \frac{3x + 5}{x^3 - x^2 - x + 1} dx$$

has denominator

$$x^3 - x^2 - x + 1 = (x + 1)(x - 1)^2.$$

Thus the decomposition must be

$$\frac{3x + 5}{(x + 1)(x - 1)^2} = \frac{A}{x + 1} + \frac{B}{x - 1} + \frac{C}{(x - 1)^2}.$$

Solving gives

$$A = \frac{1}{2}, \quad B = -\frac{1}{2}, \quad C = 4,$$

so

$$\int \frac{3x + 5}{x^3 - x^2 - x + 1} dx = \frac{1}{2} \ln |x + 1| - \frac{1}{2} \ln |x - 1| - \frac{4}{x - 1} + C.$$

§85.3 Integration by parts

The integration-by-parts formula is

$$\int u dv = uv - \int v du.$$

For example,

$$\int \ln x dx = x \ln x - x + C.$$

A higher-order version repeatedly integrates by parts:

$$\int u v^{(n+1)} dx = uv^{(n)} - u'v^{(n-1)} + \cdots + (-1)^n u^{(n)}v + (-1)^{n+1} \int u^{(n+1)}v dx.$$

This is especially useful for

$$\int P(x)e^x dx, \quad \int P(x) \sin x dx, \quad \int P(x) \cos x dx,$$

where P is a polynomial.

One worked example is

$$\int (2x^3 + 3x^2 + 4x + 5)e^x dx.$$

Repeated integration by parts gives

$$e^x(2x^3 - 3x^2 + 10x - 5) + C.$$

§85.4 Substitution

The substitution rule is

$$x = \rho(t), \quad dx = \rho'(t) dt, \quad \int f(x) dx = \int f(\rho(t))\rho'(t) dt.$$

For instance,

$$\int \sin^3 x \cos x dx$$

is solved by setting $t = \sin x$, giving

$$\int t^3 dt = \frac{t^4}{4} + C = \frac{\sin^4 x}{4} + C.$$

Similarly,

$$\int \sin^2 x \, dx$$

can be done by the trigonometric identity

$$\sin^2 x = \frac{1 - \cos 2x}{2},$$

so

$$\int \sin^2 x \, dx = \frac{x}{2} - \frac{\sin 2x}{4} + C.$$

A trigonometric substitution example is

$$\int \sqrt{a^2 - x^2} \, dx.$$

Let $x = a \sin t$. Then $dx = a \cos t \, dt$, and the integral becomes

$$a^2 \int \cos^2 t \, dt.$$

Hence

$$\int \sqrt{a^2 - x^2} \, dx = \frac{x}{2} \sqrt{a^2 - x^2} + \frac{a^2}{2} \arcsin \frac{x}{a} + C.$$

The standard substitutions are:

$$\sqrt{a^2 - x^2} : x = a \sin t, \quad \sqrt{a^2 + x^2} : x = a \tan t, \quad \sqrt{x^2 - a^2} : x = a \sec t.$$

§85.5 Partial fractions as algebra before calculus

Partial fractions are not an integration trick alone. They first require algebraic decomposition. For example, if

$$\frac{P(x)}{(x-a)(x-b)} = \frac{A}{x-a} + \frac{B}{x-b},$$

then multiplying by the denominator gives an identity of polynomials. The coefficients A, B are determined before any integration happens.

This is why one should factor the denominator and compare degrees first. If the rational function is improper, divide polynomials before decomposing.

§85.6 Integration by parts as product rule reversed

Integration by parts comes from

$$(uv)' = u'v + uv'.$$

Integrating gives

$$\int u \, dv = uv - \int v \, du.$$

The choice of u and dv is guided by which part simplifies after differentiation and which part is easy to integrate. Logarithms and inverse trigonometric functions often become simpler after differentiation.

§85.7 A second lens

The note is an inventory of the classical calculus toolkit. Partial fractions exploit algebraic factorization; integration by parts exploits the product rule; substitution exploits the chain rule. In a more advanced language, all three are ways of choosing coordinates and decompositions so that the differential form becomes elementary.

§85.8 Partial fractions as algebra of poles

Partial fractions decompose a rational function according to its poles. For example, a factor $(x - a)^m$ contributes terms

$$\frac{c_1}{x - a} + \frac{c_2}{(x - a)^2} + \cdots + \frac{c_m}{(x - a)^m}.$$

This is the real-variable shadow of meromorphic function theory, where principal parts record local behavior near poles.

§85.9 Integration by parts as adjointness

The formula

$$\int_a^b u v' = uv|_a^b - \int_a^b u' v$$

transfers a derivative from one factor to another. In operator language, differentiation is skew-adjoint up to boundary terms. This viewpoint becomes central in Fourier analysis, weak derivatives, and PDE.

§85.10 Substitution as pullback of one-forms

For $x = \phi(t)$, substitution says

$$\int f(x) dx = \int f(\phi(t)) \phi'(t) dt.$$

The invariant object is the one-form $f(x) dx$. Under the map ϕ , it pulls back to $f(\phi(t)) \phi'(t) dt$.

This explains why the derivative factor is necessary: integration is coordinate-independent only when the differential transforms correctly.

§85.11 Integration techniques as structural transformations

The standard integration techniques are structural transformations. Partial fractions isolate poles, integration by parts moves derivatives with boundary terms, and substitution pulls back differential forms. The goal is not formula memorization but recognizing which structure simplifies the integrand.

§85.12 Integration as reverse differentiation

The integration techniques in this note are best remembered as reversed differentiation rules. Partial fractions reverse the derivative of logarithms and arctangents. Integration by parts reverses the product rule. Substitution reverses the chain rule. Trigonometric

substitution is a geometric way of replacing a radical by an identity such as $1 - \sin^2 t = \cos^2 t$.

A good workflow is therefore: first simplify the algebraic shape of the integrand; next ask which differentiation rule could have produced it; only then choose a technique. This prevents the common mistake of trying substitution, parts, and partial fractions randomly.

§85.13 Substitution as pullback

The elementary substitution rule

$$\int f(g(t))g'(t) dt = \int f(x) dx$$

is the one-dimensional form of pullback. If $x = g(t)$, then $dx = g'(t)dt$, so the differential form $f(x)dx$ pulls back to

$$g^*(f(x)dx) = f(g(t))g'(t)dt.$$

The derivative is not an artificial correction factor; it is how lengths are transformed by the map g .

In several variables the same principle becomes the Jacobian determinant:

$$dx dy = \left| \frac{\partial(x, y)}{\partial(u, v)} \right| du dv.$$

Thus elementary substitution is the first instance of the change-of-variables theorem.

§85.14 Integration by parts and adjoints

Integration by parts comes from the product rule:

$$\int_a^b u'v dx = [uv]_a^b - \int_a^b uv' dx.$$

This formula is also the origin of adjoint differential operators. For $D = d/dx$,

$$\int_a^b (Du)v dx = [uv]_a^b - \int_a^b u(Dv) dx.$$

Ignoring boundary terms, D is formally skew-adjoint. This observation is central in Sturm–Liouville theory, Fourier analysis, and PDEs.

§85.15 Partial fractions and residues

Partial fractions decompose a rational function into simple local pieces around its poles. In complex analysis, the coefficient of $(z - a)^{-1}$ is the residue, and the residue theorem says

$$\int_{\gamma} f(z) dz = 2\pi i \sum_{a \text{ inside } \gamma} \text{Res}(f, a).$$

Thus the elementary partial-fraction method is a real-variable shadow of a deeper principle: rational functions are controlled by their singularities.

Trigonometric Limits, Derivatives, and Inverse Trigonometric Examples

N-20221107

§86.1 Trigonometric limits

The note collects examples about trigonometric functions and derivatives. The first basic limit is

$$\lim_{x \rightarrow 0} \frac{\sin x}{x} = 1.$$

From this one obtains

$$\lim_{x \rightarrow 0} \frac{1 - \cos x}{x} = 0, \quad \lim_{x \rightarrow 0} \frac{1 - \cos x}{x^2} = \frac{1}{2}.$$

The second follows from

$$1 - \cos x = 2 \sin^2 \frac{x}{2}.$$

The handwritten page uses these limits repeatedly when deriving derivatives of trigonometric functions.

§86.2 Derivative of sine and cosine

Using the addition formula,

$$\sin(x + h) - \sin x = \sin x(\cos h - 1) + \cos x \sin h.$$

Therefore

$$\frac{\sin(x + h) - \sin x}{h} = \sin x \frac{\cos h - 1}{h} + \cos x \frac{\sin h}{h}.$$

Letting $h \rightarrow 0$ gives

$$(\sin x)' = \cos x.$$

Similarly,

$$\cos(x + h) - \cos x = \cos x(\cos h - 1) - \sin x \sin h,$$

so

$$(\cos x)' = -\sin x.$$

§86.3 Tangent, cotangent, secant, and cosecant

Using quotient rule,

$$\tan x = \frac{\sin x}{\cos x} \implies (\tan x)' = \frac{1}{\cos^2 x} = \sec^2 x.$$

Similarly,

$$(\cot x)' = -\csc^2 x, \quad (\sec x)' = \sec x \tan x, \quad (\csc x)' = -\csc x \cot x.$$

The note includes examples where the algebra is simplified by rewriting everything in sine and cosine first.

§86.4 Inverse trigonometric derivatives

For $y = \arcsin x$, we have

$$x = \sin y, \quad 1 = \cos y y'.$$

Since $y \in [-\pi/2, \pi/2]$, $\cos y = \sqrt{1 - x^2}$, so

$$(\arcsin x)' = \frac{1}{\sqrt{1 - x^2}}.$$

Similarly,

$$(\arccos x)' = -\frac{1}{\sqrt{1 - x^2}}, \quad (\arctan x)' = \frac{1}{1 + x^2}.$$

The page also works through reciprocal-trigonometric inverse derivatives, using triangles to determine the missing side length.

§86.5 Examples from the handwritten page

A typical example differentiates a composite expression such as

$$y = \frac{\sin x}{1 + \cos x}.$$

One can use quotient rule directly, but the note also points out the simplification

$$\frac{\sin x}{1 + \cos x} = \tan \frac{x}{2}.$$

Thus

$$y' = \frac{1}{2} \sec^2 \frac{x}{2}.$$

This illustrates a useful habit: before differentiating, look for a trigonometric identity that simplifies the expression.

Another example uses implicit differentiation for equations involving $\sin x$ and $\cos x$. The method is always the same: differentiate both sides, then solve for the desired derivative.

§86.6 A closing lens

The derivatives of trigonometric functions are not isolated formulas. They all come from the same two sources: the fundamental limit $\sin h/h \rightarrow 1$ and the addition formulas. Inverse trigonometric derivatives then come from implicit differentiation and right-triangle identities. The note is mainly about learning to move between these representations fluently.

§86.7 The circle as a Lie group

The unit circle can be written as

$$S^1 = \{e^{it} : t \in \mathbb{R}\}.$$

Multiplication of complex numbers adds angles:

$$e^{is}e^{it} = e^{i(s+t)}.$$

Thus trigonometric addition formulas are group laws in coordinates. The identities for sine and cosine are the real and imaginary parts of this multiplication rule.

§86.8 Trigonometric derivatives from the exponential

Euler's formula

$$e^{ix} = \cos x + i \sin x$$

turns differentiation of trigonometric functions into differentiation of e^{ix} :

$$\frac{d}{dx} e^{ix} = i e^{ix} = i \cos x - \sin x.$$

Comparing real and imaginary parts gives

$$(\cos x)' = -\sin x, \quad (\sin x)' = \cos x.$$

§86.9 Inverse trigonometric branches

Inverse trigonometric functions require branch choices. For arcsin, sine is restricted to $[-\pi/2, \pi/2]$ so that it is one-to-one. The derivative

$$(\arcsin x)' = \frac{1}{\sqrt{1-x^2}}$$

blows up at ± 1 because the original sine curve has horizontal tangent there.

§86.10 Trigonometry from circle symmetry and branches

Trigonometric limits and derivatives are not isolated formulas. They come from the geometry and group structure of the circle, while inverse functions require branch choices and the inverse function theorem.

Derivatives: Definition, Rules, Elementary Functions, and Inverse Trigonometric Derivatives

N-20221211

§87.1 Average velocity and instantaneous velocity

The first page starts from a graph and the idea of velocity. The average rate of change of f on $[x_0, x_0 + \Delta x]$ is

$$\frac{f(x_0 + \Delta x) - f(x_0)}{\Delta x}.$$

As $\Delta x \rightarrow 0$, the secant line approaches the tangent line if the limit exists. The instantaneous rate of change is the derivative:

$$f'(x_0) = \lim_{\Delta x \rightarrow 0} \frac{f(x_0 + \Delta x) - f(x_0)}{\Delta x}.$$

Geometrically, this is the slope of the tangent line.

§87.2 Derivative notation

The note lists the notations

$$f'(x), \quad y', \quad \frac{dy}{dx}, \quad \frac{df}{dx}.$$

It also distinguishes the derivative at a point from the derivative function. The derivative at x_0 is a number; the derivative function assigns such a number to each point where the derivative exists.

The tangent line formula is

$$y - f(x_0) = f'(x_0)(x - x_0).$$

§87.3 Power functions

For positive integer n ,

$$(x^n)' = nx^{n-1}.$$

This follows from

$$(x + h)^n - x^n = \sum_{k=1}^n \binom{n}{k} x^{n-k} h^k,$$

and after dividing by h and taking $h \rightarrow 0$, only the $k = 1$ term remains.

The note extends the formula to rational powers by implicit differentiation. If

$$y = x^{m/n},$$

then

$$y^n = x^m.$$

Differentiating gives

$$ny^{n-1}y' = mx^{m-1},$$

so

$$y' = \frac{m}{n}x^{m/n-1}.$$

§87.4 Exponential and logarithmic derivatives

The page records

$$(a^x)' = a^x \ln a, \quad (e^x)' = e^x.$$

For logarithms,

$$(\ln x)' = \frac{1}{x}, \quad (\log_a x)' = \frac{1}{x \ln a}.$$

A useful derivation uses inverse functions: since $y = \ln x$ means $x = e^y$, differentiating gives

$$1 = e^y y' = xy',$$

so $y' = 1/x$.

§87.5 Trigonometric derivatives

Using the standard limits and addition formulas,

$$(\sin x)' = \cos x, \quad (\cos x)' = -\sin x.$$

Then quotient rule gives

$$(\tan x)' = \sec^2 x, \quad (\cot x)' = -\csc^2 x.$$

The reciprocal functions satisfy

$$(\sec x)' = \sec x \tan x, \quad (\csc x)' = -\csc x \cot x.$$

§87.6 Inverse trigonometric derivatives

For $y = \arcsin x$,

$$x = \sin y$$

and so

$$1 = \cos y y'.$$

Since $\cos y = \sqrt{1-x^2}$ on the principal branch,

$$(\arcsin x)' = \frac{1}{\sqrt{1-x^2}}.$$

Similarly,

$$(\arccos x)' = -\frac{1}{\sqrt{1-x^2}}, \quad (\arctan x)' = \frac{1}{1+x^2}.$$

The note also derives inverse secant and inverse cosecant derivatives by drawing a right triangle and expressing the missing side length in terms of x .

§87.7 Examples

A typical example differentiates

$$y = \sqrt{x^2 + 1}.$$

By chain rule,

$$y' = \frac{x}{\sqrt{x^2 + 1}}.$$

Another example differentiates a quotient of trigonometric functions. The handwritten solution usually first rewrites the expression using identities, then differentiates. This is often simpler than applying quotient rule blindly.

§87.8 Context and outlook

This note builds the derivative table from a few principles: definition by difference quotient, binomial expansion for powers, inverse-function differentiation for logarithms and inverse trigonometric functions, and quotient/chain rules for combinations. The goal is not memorizing isolated formulas but recognizing which rule creates each formula.

§87.9 Derivative as first-order linearization

The derivative at a is best read as a first-order approximation:

$$f(a + h) = f(a) + f'(a)h + o(h).$$

The number $f'(a)$ is the linear map that best approximates the change of f near a . In one dimension, linear maps are multiplication by a number, so this viewpoint collapses to the familiar tangent slope.

This explains why the chain rule has the form it does. If

$$f(a + h) = f(a) + f'(a)h + o(h), \quad g(b + k) = g(b) + g'(b)k + o(k), \quad b = f(a),$$

then

$$g(f(a + h)) = g(f(a)) + g'(f(a))f'(a)h + o(h).$$

Thus

$$(g \circ f)'(a) = g'(f(a))f'(a).$$

The chain rule is composition of linear approximations.

§87.10 Inverse trigonometric derivatives and the inverse function theorem

The inverse-function formula

$$(f^{-1})'(y) = \frac{1}{f'(x)}, \quad y = f(x),$$

is a one-dimensional form of the inverse function theorem. If f is continuously differentiable near x_0 and $f'(x_0) \neq 0$, then f has a local differentiable inverse near $f(x_0)$.

For $y = \arcsin x$, the equation $x = \sin y$ is inverted only after restricting sine to

$$[-\pi/2, \pi/2].$$

On this branch, $\cos y \geq 0$, so

$$1 = \cos y \frac{dy}{dx}, \quad \frac{dy}{dx} = \frac{1}{\cos y} = \frac{1}{\sqrt{1-x^2}}.$$

The square root sign is not arbitrary; it comes from the chosen branch of the inverse.

§87.11 Endpoint blow-up and failure of local invertibility

The derivative

$$(\arcsin x)' = \frac{1}{\sqrt{1-x^2}}$$

blows up as $x \rightarrow \pm 1$. Geometrically, near $y = \pm\pi/2$, the graph of $\sin y$ has horizontal tangent:

$$\cos y = 0.$$

The inverse function theorem fails exactly where $f'(y) = 0$. The inverse may still exist globally on a restricted interval, but it is no longer differentiable with a finite derivative at the endpoint.

This is the conceptual reason inverse-trigonometric formulas have domain restrictions and singular denominators.

§87.12 Logarithm as inverse and as group homomorphism

The derivative of $\ln x$ follows from $x = e^y$:

$$1 = e^y y' = xy', \quad y' = \frac{1}{x}.$$

But $\ln$ also has a structural meaning:

$$\ln(xy) = \ln x + \ln y.$$

It converts multiplication in $(0, \infty)$ into addition in $\mathbb{R}$. Its derivative $1/x$ is the infinitesimal version of multiplicative scaling.

In Lie-group language, $(0, \infty)$ is a one-dimensional Lie group under multiplication, and $\ln$ identifies it with the additive group $\mathbb{R}$. The elementary derivative formula is the local shadow of this group isomorphism.

§87.13 Automatic differentiation and computation graphs

The product, quotient, and chain rules are the mathematical basis of automatic differentiation. A complicated expression is a computation graph. Each node knows its local derivative, and the chain rule propagates derivative information through the graph.

For example, if

$$y = \sqrt{x^2 + 1}, \quad u = x^2 + 1, \quad y = \sqrt{u},$$

then

$$\frac{dy}{dx} = \frac{dy}{du} \frac{du}{dx} = \frac{1}{2\sqrt{u}} \cdot 2x = \frac{x}{\sqrt{x^2 + 1}}.$$

The ordinary chain-rule computation is already reverse or forward mode automatic differentiation in miniature.

§87.14 Jets and higher-order local data

The k -jet of a function at a records its derivatives up to order k :

$$j_a^k f = (f(a), f'(a), \dots, f^{(k)}(a)).$$

The Taylor polynomial is the coordinate expression of this finite local data:

$$T_k f(a + h) = \sum_{j=0}^k \frac{f^{(j)}(a)}{j!} h^j.$$

The first derivative is therefore not an isolated operation; it is the first nontrivial layer of a tower of local approximations.

§87.15 Differentials and one-forms

The notation dy/dx is sometimes taught as a fraction-like symbol. A safer advanced interpretation is through differentials. For $y = f(x)$,

$$dy = f'(x) dx.$$

Here dx is the basic one-form on the line, and dy is its pullback through f . The chain rule becomes the rule for pulling back one-forms:

$$(g \circ f)^*(dz) = f^*(g^*(dz)).$$

This gives a structural meaning to the notation without treating infinitesimals as informal numbers.

§87.16 Derivatives as linearization and pullback

The derivative table is built from a few structural principles: linear approximation, inverse functions, branch choices, composition, and local Taylor data. The formulas for powers, logarithms, trigonometric functions, and inverse trigonometric functions are not separate facts; they are examples of how local linear behavior is transported through algebraic operations and inverses.

Derivatives, Tangents, Extremal Problems, Young–Holder–Minkowski, and Jensen

N-20230310

§88.1 The derivative from secants to tangents

The note begins with the geometric idea of derivative. Given two nearby points on the curve,

$$(x_0, f(x_0)), \quad (x_0 + \Delta x, f(x_0 + \Delta x)),$$

the secant slope is

$$\frac{\Delta y}{\Delta x} = \frac{f(x_0 + \Delta x) - f(x_0)}{\Delta x}.$$

If, as $\Delta x \rightarrow 0$, this quotient tends to a finite limit, then the curve has a definite tangent line at x_0 . The derivative is

$$f'(x_0) = \lim_{\Delta x \rightarrow 0} \frac{f(x_0 + \Delta x) - f(x_0)}{\Delta x}.$$


Equivalently,

$$f'(x_0) = \lim_{x \rightarrow x_0} \frac{f(x) - f(x_0)}{x - x_0}.$$

The tangent line at x_0 is therefore

$$y - f(x_0) = f'(x_0)(x - x_0).$$

The source page stresses the visual transition: the secant line approaches a unique limiting line. If that limiting line is not vertical and the slope is finite, we call the slope the derivative.



§88.2 Derivative notation and local definition

The definition requires f to be defined on a punctured neighborhood of x_0 , and usually on a small open interval around it. The derivative at x_0 may be denoted by any of

$$f'(x_0), \quad \left. \frac{df}{dx} \right|_{x=x_0}, \quad y'(x_0), \quad \left. \frac{dy}{dx} \right|_{x=x_0}.$$

If f is differentiable at every point of an open interval (a, b) , then $x \mapsto f'(x)$ is the derivative function on that interval.

The note emphasizes a small logical point: before asking whether f is differentiable at a point, one must first know that the function is defined near that point. Differentiability is a local property.

§88.3 Rules of differentiation

If f and g are differentiable at x_0 , then sums, scalar multiples, products, and quotients obey:

$$(af + bg)' = af' + bg', \quad (fg)' = f'g + fg', \quad \left(\frac{f}{g} \right)' = \frac{f'g - fg'}{g^2}, \quad g(x_0) \neq 0.$$

The chain rule says

$$(f \circ g)'(x_0) = f'(g(x_0))g'(x_0).$$

The inverse-function rule says that if f and f^{-1} are inverse functions, $f'(x_0) \neq 0$, and $y_0 = f(x_0)$, then

$$(f^{-1})'(y_0) = \frac{1}{f'(x_0)}.$$

These rules are not simply a table to memorize. Each rule is a disguised statement about first-order linear approximation. For example,

$$f(x_0 + h) = f(x_0) + f'(x_0)h + o(h), \quad g(x_0 + h) = g(x_0) + g'(x_0)h + o(h),$$

and multiplying these expansions gives the product rule.

§88.4 Elementary derivative formulas

The table in the note includes the standard formulas:

$$C' = 0, \quad (x^a)' = ax^{a-1}, \quad (a^x)' = a^x \ln a, \quad (\log_a x)' = \frac{1}{x \ln a}.$$

For trigonometric functions,

$$\begin{aligned} (\sin x)' &= \cos x, & (\cos x)' &= -\sin x, & (\tan x)' &= \sec^2 x, & (\cot x)' &= -\csc^2 x, \\ (\sec x)' &= \sec x \tan x, & (\csc x)' &= -\csc x \cot x. \end{aligned}$$

For inverse trigonometric functions,

$$(\arcsin x)' = \frac{1}{\sqrt{1-x^2}}, \quad (\arccos x)' = -\frac{1}{\sqrt{1-x^2}}, \quad (\arctan x)' = \frac{1}{1+x^2}.$$

§88.5 Extrema, Fermat's lemma, and critical points

A point x_0 is a local maximum if $f(x) \leq f(x_0)$ near x_0 , and a local minimum if $f(x) \geq f(x_0)$ near x_0 . If $f'(x_0) = 0$, then x_0 is called a stationary or critical point.

Fermat's lemma says: if f is differentiable at x_0 and has a local extremum there, then

$$f'(x_0) = 0.$$

The note writes a warning beside this theorem: the converse is false. For example,

$$f(x) = x^3$$

has $f'(0) = 0$, but 0 is not an extremum. Fermat's lemma only provides candidates for extrema. One must then check sign changes, second derivatives, endpoints, or non-differentiable points.

If $f''(x_0) > 0$ and $f'(x_0) = 0$, then x_0 is a local minimum. If $f''(x_0) < 0$, it is a local maximum. The reason is that f'' describes whether f' is increasing or decreasing near the critical point.

§88.6 A derivative proof of a basic inequality

The note proves, for $x > 0$,

$$x^\alpha - \alpha x + \alpha - 1 \leq 0, \quad 0 < \alpha < 1,$$

and

$$x^\alpha - \alpha x + \alpha - 1 \geq 0, \quad \alpha < 0 \text{ or } \alpha > 1.$$

Let

$$F(x) = x^\alpha - \alpha x + \alpha - 1.$$

Then

$$F'(x) = \alpha(x^{\alpha-1} - 1), \quad F(1) = 0.$$

For $0 < \alpha < 1$, F increases on $(0, 1)$ and decreases on $(1, \infty)$, so $x = 1$ is a maximum and $F(x) \leq 0$. For $\alpha < 0$ or $\alpha > 1$, the monotonicity is reversed, and $x = 1$ is a minimum, so $F(x) \geq 0$.

This small inequality is the engine behind Young's inequality.

§88.7 Young's inequality

Let $a, b > 0$ and let p, q satisfy

$$\frac{1}{p} + \frac{1}{q} = 1.$$

For $p > 1$, Young's inequality says

$$a^{1/p} b^{1/q} \leq \frac{a}{p} + \frac{b}{q},$$

with equality iff $a = b$. The proof substitutes

$$x = \frac{a}{b}, \quad \alpha = \frac{1}{p}$$

into the previous derivative inequality, then multiplies by b .

The source also records the reversed version for $p < 1$, where the inequality sign changes. The important conceptual point is that an inequality between weighted geometric and arithmetic means is reduced to the shape of one single-variable function.

§88.8 Holder's inequality

For $x_i, y_i \geq 0$ and $1/p + 1/q = 1$, Holder's inequality is

$$\sum_{i=1}^n x_i y_i \leq \left(\sum_{i=1}^n x_i^p \right)^{1/p} \left(\sum_{i=1}^n y_i^q \right)^{1/q}, \quad p > 1.$$

The proof in the source normalizes

$$a_i = \frac{x_i^p}{\sum_j x_j^p}, \quad b_i = \frac{y_i^q}{\sum_j y_j^q}$$

and applies Young's inequality termwise. After adding the inequalities, the normalized sum becomes 1, giving Holder.

A handwritten note observes that the integral form follows the same idea:

$$\int_a^b f(x)g(x) dx \leq \left(\int_a^b |f(x)|^p dx \right)^{1/p} \left(\int_a^b |g(x)|^q dx \right)^{1/q}.$$

The comment says this feels almost trivial once the finite-sum proof is understood: one replaces sums by integrals and uses the same normalization logic.

For $p = q = 2$, Holder becomes the Cauchy–Bunyakovsky–Schwarz inequality.

§88.9 Minkowski's inequality

For $p > 1$, Minkowski's inequality says

$$\left(\sum_{i=1}^n (x_i + y_i)^p \right)^{1/p} \leq \left(\sum_{i=1}^n x_i^p \right)^{1/p} + \left(\sum_{i=1}^n y_i^p \right)^{1/p}.$$

The proof applies Holder to the two sums

$$\sum_i x_i (x_i + y_i)^{p-1}, \quad \sum_i y_i (x_i + y_i)^{p-1}.$$

Because $q = p/(p-1)$, we have

$$((x_i + y_i)^{p-1})^q = (x_i + y_i)^p.$$

After factoring the common term and dividing by it, one obtains Minkowski.

The handwritten margin notes that for $p = 2$ this is exactly the triangle inequality for Euclidean length:

$$\sqrt{(x_1 + y_1)^2 + \cdots + (x_n + y_n)^2} \leq \sqrt{x_1^2 + \cdots + x_n^2} + \sqrt{y_1^2 + \cdots + y_n^2}.$$

§88.10 Convex functions and Jensen's inequality

The note defines a convex function by the inequality

$$f(\alpha_1 x_1 + \alpha_2 x_2) \leq \alpha_1 f(x_1) + \alpha_2 f(x_2), \quad \alpha_1 + \alpha_2 = 1, \quad \alpha_i \geq 0.$$

The handwritten annotation explains the geometry: the point on the chord between $(x_1, f(x_1))$ and $(x_2, f(x_2))$ lies above the graph.

Equivalently, for $x_1 < x < x_2$,

$$\frac{f(x) - f(x_1)}{x - x_1} \leq \frac{f(x_2) - f(x)}{x_2 - x}.$$

This slope inequality implies that if f is differentiable, then f' is increasing; if f is twice differentiable, then convexity corresponds to $f'' \geq 0$.

Jensen's inequality generalizes the two-point definition:

$$f\left(\sum_{i=1}^n \alpha_i x_i\right) \leq \sum_{i=1}^n \alpha_i f(x_i), \quad \alpha_i \geq 0, \quad \sum_i \alpha_i = 1.$$

The source proves this by induction. Let $\beta = \alpha_2 + \cdots + \alpha_n = 1 - \alpha_1$ and normalize the remaining weights by α_i/β . Then use the two-point convexity inequality followed by the induction hypothesis.

§88.11 Typical exercise themes

The last pages list exercises. They are not random; they are organized around the derivative as a tool for inequality and extrema.

Examples include:

- (i) determining when a line and cubic have three distinct intersections;
- (ii) finding monotonicity intervals of $f(x) = \sqrt{x} - \ln(x + a)$;
- (iii) maximizing $\sin^2 x \cos x$ on $(0, \pi/2)$;
- (iv) comparing numbers such as $0.1e^{0.1}$, $1/9$, and $-\ln 0.9$ using auxiliary functions;
- (v) proving $e^x \geq 1 + x$ and $\ln x \leq x - 1$;
- (vi) using Jensen or logarithms to prove entropy-like inequalities such as

$$\sum_{i=1}^{2^n} p_i \log_2 p_i \geq -n, \quad \sum_i p_i = 1;$$

- (vii) proving product inequalities of the form

$$\prod_{i=1}^n a_i^{b_i} \leq 1$$

under weighted-sum hypotheses;

- (viii) minimizing expressions such as

$$\frac{9}{\sin^2 x} + \frac{1}{\cos x \cos y \sin^2 y}.$$

The common method is to build an auxiliary function, compute derivatives, locate monotonicity or convexity, and then read off the desired inequality.

§88.12 The reusable pattern

This note shows how the derivative becomes more than a tangent slope. It gives a language for extrema, monotonicity, convexity, and inequalities. Young, Holder, Minkowski, and Jensen all look like different inequalities, but the note reveals a chain:

$$\begin{aligned} \text{single-variable derivative inequality} &\rightarrow \text{Young} \rightarrow \text{Holder} \\ &\rightarrow \text{Minkowski}, \quad \text{convexity} \rightarrow \text{Jensen}. \end{aligned}$$

Thus calculus becomes a proof technology for analysis and inequalities.

§88.13 Derivative as local linear approximation

The derivative is the coefficient of the best first-order model:

$$f(a+h) = f(a) + f'(a)h + o(h).$$

This viewpoint explains the product, quotient, and chain rules: they are rules for composing and multiplying first-order approximations.

In several variables the same idea becomes the Fréchet derivative, a linear map rather than a single number.

§88.14 Extrema and first variation

Fermat's lemma says that if f has a local extremum at an interior point and is differentiable there, then $f'(a) = 0$. The derivative is the first variation; if a small positive and negative move are both allowed, the linear term must vanish.

Second derivative tests then study the next nonzero term in the local expansion.

§88.15 Young, Holder, and convex duality

Young's inequality

$$ab \leq \frac{a^p}{p} + \frac{b^q}{q}, \quad \frac{1}{p} + \frac{1}{q} = 1,$$

is a convex-duality statement. It is the pointwise engine behind Holder's inequality. Minkowski then gives the triangle inequality for L^p norms.

Thus the inequality part of the note connects differential calculus to convex analysis and normed spaces.

§88.16 Local linearity and convex control

The chapter links two central ideas: derivatives provide local linear models, and inequalities describe convex or norm geometry. Tangents, extrema, Young–Holder–Minkowski, and Jensen all express how local or averaged information controls global behavior.

Calculus Review I: Inequalities, Limits, Derivatives, and More

N-20251112

§89.1 Cauchy–Schwarz and Minkowski

The opening page recalls two inequalities:

$$\left(\sum_{i=1}^n a_i b_i\right)^2 \leq \left(\sum_{i=1}^n a_i^2\right) \left(\sum_{i=1}^n b_i^2\right),$$

and

$$\left(\sum_{i=1}^n (a_i + b_i)^2\right)^{1/2} \leq \left(\sum_{i=1}^n a_i^2\right)^{1/2} + \left(\sum_{i=1}^n b_i^2\right)^{1/2}.$$

The first is the Cauchy–Schwarz inequality; the second is Minkowski’s inequality, the triangle inequality in Euclidean norm.

The review function of these formulas is clear: they are often used to control error terms when proving limits.

§89.2 Neighborhood notation

For $x_0 \in \mathbb{R}$ and $\delta > 0$, the δ -neighborhood is

$$N_\delta(x_0) = \{x : |x - x_0| < \delta\}.$$

The punctured neighborhood is

$$\mathring{N}_\delta(x_0) = \{x : 0 < |x - x_0| < \delta\}.$$

The right and left neighborhoods are

$$N_\delta^+(x_0) = (x_0, x_0 + \delta), \quad N_\delta^-(x_0) = (x_0 - \delta, x_0).$$

For large variables, the notes use

$$U(\infty) = \{x : |x| > M\}, \quad U(+\infty) = \{x : x > M\}, \quad U(-\infty) = \{x : x < -M\}.$$

§89.3 Trigonometric identities

The first page lists a full table of trigonometric identities. The basic ones are

$$\sin^2 \alpha + \cos^2 \alpha = 1, \quad 1 + \tan^2 \alpha = \sec^2 \alpha, \quad 1 + \cot^2 \alpha = \csc^2 \alpha.$$

The addition formulas are

$$\sin(\alpha \pm \beta) = \sin \alpha \cos \beta \pm \cos \alpha \sin \beta,$$

$$\cos(\alpha \pm \beta) = \cos \alpha \cos \beta \mp \sin \alpha \sin \beta,$$

$$\operatorname{an}(\alpha \pm \beta) = \frac{\operatorname{an}\alpha \pm \operatorname{an}\beta}{1 \mp \operatorname{an}\alpha \operatorname{an}\beta}.$$

The sum-to-product and product-to-sum formulas appear as well, for example

$$\sin \alpha + \sin \beta = 2 \sin \frac{\alpha + \beta}{2} \cos \frac{\alpha - \beta}{2},$$

$$\sin \alpha \sin \beta = \frac{1}{2} [\cos(\alpha - \beta) - \cos(\alpha + \beta)].$$

The page includes double-angle and power-reduction formulas:

$$\sin 2\alpha = 2 \sin \alpha \cos \alpha, \quad \cos 2\alpha = 2 \cos^2 \alpha - 1 = 1 - 2 \sin^2 \alpha,$$

$$\sin^2 \alpha = \frac{1}{2}(1 - \cos 2\alpha), \quad \cos^2 \alpha = \frac{1}{2}(1 + \cos 2\alpha).$$

§89.4 Hyperbolic functions

The next page introduces functions generated by e^x :

$$\sinh x = \frac{1}{2}(e^x - e^{-x}), \quad \cosh x = \frac{1}{2}(e^x + e^{-x}),$$

$$\operatorname{anh} x = \frac{e^x - e^{-x}}{e^x + e^{-x}}, \quad \coth x = \frac{e^x + e^{-x}}{e^x - e^{-x}}.$$

Their identities mirror trigonometric identities but with sign changes:

$$\cosh^2 x - \sinh^2 x = 1,$$

$$\sinh(x + y) = \sinh x \cosh y + \cosh x \sinh y,$$

$$\cosh(x + y) = \cosh x \cosh y + \sinh x \sinh y.$$

The note treats them as a useful parallel system rather than as isolated special functions.

§89.5 Sequence limits: epsilon–N language

A sequence $\{x_n\}$ converges to A if for every $\varepsilon > 0$ there exists $N \in \mathbb{N}$ such that

$$n > N \implies |x_n - A| < \varepsilon.$$

We write

$$\lim_{n \rightarrow \infty} x_n = A.$$

The note highlights the uniqueness theorem. If a sequence had two limits A and B , choose

$$\varepsilon = \frac{|A - B|}{2}.$$

For large n ,

$$|A - B| \leq |A - x_n| + |x_n - B| < 2\varepsilon = |A - B|,$$

a contradiction.

§89.6 Infinite limits and one-sided limits

A sequence tends to $+\infty$ if for every $G > 0$ there is N such that

$$n > N \implies x_n > G.$$

For functions, the epsilon-delta definition is

$$\lim_{x \rightarrow x_0} f(x) = A$$

if for every $\varepsilon > 0$ there exists $\delta > 0$ such that

$$0 < |x - x_0| < \delta \implies |f(x) - A| < \varepsilon.$$

One-sided limits are defined by replacing the punctured neighborhood with a right or left punctured neighborhood. A two-sided limit exists iff the two one-sided limits exist and are equal.

§89.7 Heine criterion for function limits

The highlighted theorem says:

$$\lim_{x \rightarrow x_0} f(x) = A$$

if and only if for every sequence $x_n \rightarrow x_0$ with $x_n \neq x_0$, one has

$$f(x_n) \rightarrow A.$$

This is often used to prove that a limit does not exist. For example,

$$\lim_{x \rightarrow 0} \cos \frac{1}{x}$$

does not exist because along suitable sequences the values tend to different limits. The scanned page uses the sequences

$$x_n = \frac{1}{2n\pi}, \quad y_n = \frac{1}{2n\pi + \pi/2}.$$

Then $x_n, y_n \rightarrow 0$, but $f(x_n) \rightarrow 1$ and $f(y_n) \rightarrow 0$.

§89.8 Infinitesimals and infinitesimal orders

An infinitesimal is a variable whose limit is zero. The note warns that a fixed small number such as e^{-100} is not an infinitesimal; infinitesimal means a variable in a limiting process.

The pages record the usual rules:

- a finite sum of infinitesimals is infinitesimal;
- a bounded variable times an infinitesimal is infinitesimal;
- a constant times an infinitesimal is infinitesimal;
- the product of two infinitesimals is infinitesimal;
- the product of infinitely many infinitesimals need not be handled by a naive rule.

The handwritten orange comment gives examples such as

$$1, \frac{1}{2}, \frac{1}{3}, \dots, \frac{1}{n}, \dots$$

and warns that infinitely many factors require separate analysis. The problem is not philosophical; it is about whether one is taking a finite algebraic combination or a changing infinite product.

For two infinitesimals α, β , one writes

$$\alpha = o(\beta) \iff \lim \frac{\alpha}{\beta} = 0,$$

and

$$\alpha \sim \beta \iff \lim \frac{\alpha}{\beta} = 1.$$

Equivalent infinitesimals may be substituted in products and quotients when the relevant limits exist.

§89.9 Squeeze theorem and monotone convergence

The squeeze theorem is recorded as:

$$X \leq Y \leq Z, \quad \lim X = \lim Z = A \implies \lim Y = A.$$

The monotone convergence theorem says that a monotone bounded sequence converges. For example, an increasing sequence with an upper bound converges to its least upper bound; a decreasing sequence with a lower bound converges to its greatest lower bound.

The page gives the standard proof for the decreasing case: let A be the infimum. Since A is the greatest lower bound, for every $\varepsilon > 0$ there exists some $x_N < A + \varepsilon$, and monotonicity then gives $A \leq x_n \leq x_N < A + \varepsilon$ for $n > N$.

§89.10 Cauchy criteria

The sequence Cauchy criterion is

$$\{x_n\} \text{extconverges} \iff \forall \varepsilon > 0, \exists N, m, n > N \Rightarrow |x_m - x_n| < \varepsilon.$$

For function limits, the analogous criterion is: $f(x)$ has a finite limit at x_0 iff for every $\varepsilon > 0$ there is a punctured neighborhood such that

$$x_1, x_2 \in \mathring{N}_\delta(x_0) \implies |f(x_1) - f(x_2)| < \varepsilon.$$

The note states that the proof of sufficiency uses completeness of the real numbers, so it is often skipped in elementary courses.

§89.11 Continuity

A function is continuous at x_0 if

$$\lim_{x \rightarrow x_0} f(x) = f(x_0).$$

Equivalent sequentially:

$$x_n \rightarrow x_0 \implies f(x_n) \rightarrow f(x_0).$$

The pages list the usual operations: sums, products, quotients with nonzero denominator, and compositions of continuous functions are continuous where defined.

The note uses continuity to justify passing limits through elementary functions, for example

$$\lim \sin u_n = \sin(\lim u_n)$$

when u_n has a limit.

§89.12 Derivatives and differentials

The later pages review derivative definitions:

$$f'(x_0) = \lim_{\Delta x \rightarrow 0} \frac{f(x_0 + \Delta x) - f(x_0)}{\Delta x}.$$

The differential form is

$$dy = f'(x) dx.$$

The derivative rules are:

$$(u + v)' = u' + v', \quad (uv)' = u'v + uv', \quad (u/v)' = \frac{u'v - uv'}{v^2},$$

$$(f \circ g)'(x) = f'(g(x))g'(x).$$

The scanned pages also include logarithmic differentiation and implicit differentiation examples. The practical method is to take logarithms when products, quotients, and powers are entangled, and to differentiate both sides when the function is defined implicitly.

§89.13 Taylor-type expansions

The final pages include expansion examples and reminders about higher derivatives. The standard Taylor form near $x = 0$ is

$$f(x) = f(0) + f'(0)x + \frac{f''(0)}{2!}x^2 + \cdots.$$

Frequently used expansions include

$$e^x = 1 + x + \frac{x^2}{2!} + \cdots,$$

$$\sin x = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \cdots, \quad \cos x = 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \cdots.$$

The handwritten derivations use these expansions to compare infinitesimals and compute limits.

§89.14 Where the calculation points

This note is a calculus foundation review. The chain of ideas is:

neighborhoods $\rightarrow$ epsilon definitions $\rightarrow$ sequence criteria
 $\rightarrow$ infinitesimals $\rightarrow$ continuity $\rightarrow$ derivatives and expansions.

The handwritten comments are mostly warnings: do not confuse small constants with infinitesimals; do not multiply infinitely many infinitesimals naively; do not prove nonexistence of a limit without constructing incompatible approaches; do not use equivalent infinitesimals outside their valid algebraic context.

§89.15 Limits through filters and sequences

The epsilon- N definition says that eventually all terms lie in every prescribed neighborhood. In metric spaces, sequences are enough to test limits. In general topological spaces, nets and filters replace sequences.

This explains the Heine criterion: for functions between metric spaces, a limit exists iff every convergent sequence approaching the point has the corresponding image limit.

§89.16 Hyperbolic functions and exponential coordinates

The hyperbolic functions

$$\cosh x = \frac{e^x + e^{-x}}{2}, \quad \sinh x = \frac{e^x - e^{-x}}{2}$$

parametrize the hyperbola

$$\cosh^2 x - \sinh^2 x = 1.$$

They are to Lorentzian geometry what sine and cosine are to Euclidean circle geometry.

§89.17 Cauchy–Schwarz, Minkowski, and normed spaces

Cauchy–Schwarz gives inner-product control, while Minkowski gives the triangle inequality for norms. These inequalities are the finite-dimensional entrance to Banach and Hilbert spaces, where convergence, projection, and approximation become geometric.

§89.18 Calculus review through limits and norm geometry

The review topics are connected by structure: limits are neighborhood statements, trigonometric and hyperbolic identities come from group/exponential parametrizations, and inequalities express norm or inner-product geometry.

Calculus Review II:

Monotonicity, Integration

Methods, Applications, and

More

N-20251210

§90.1 Strict monotonicity from the derivative

If f is differentiable on an interval I and

$$f'(x) > 0 \quad (x \in I),$$

then f is strictly increasing on I . If

$$f'(x) < 0 \quad (x \in I),$$

then f is strictly decreasing.

The proof uses the Lagrange mean value theorem. For $x_2 > x_1$,

$$f(x_2) - f(x_1) = f'(\xi)(x_2 - x_1)$$

for some $\xi \in (x_1, x_2)$. If $f'(\xi) > 0$, then the difference is positive.

§90.2 First derivative test for extrema

Suppose f is continuous near x_0 and differentiable in the punctured neighborhood. If the sign of f' changes from negative to positive at x_0 , then $f(x_0)$ is a local minimum. If the sign changes from positive to negative, then $f(x_0)$ is a local maximum.

The page records the compact criterion:

$$(x - x_0)f'(x) > 0 \implies f(x_0) \text{ is a local minimum,}$$

$$(x - x_0)f'(x) < 0 \implies f(x_0) \text{ is a local maximum.}$$

§90.3 Second derivative test

If

$$f'(x_0) = 0, \quad f''(x_0) \neq 0,$$

then

$$f''(x_0) > 0 \implies x_0 \text{ is a local minimum,}$$

$$f''(x_0) < 0 \implies x_0 \text{ is a local maximum.}$$

The handwritten note adds the mnemonic that $f'' > 0$ looks like a smile and $f'' < 0$ looks like a frown.

§90.4 Concavity and inflection points

A curve segment is concave up if it lies above its tangent lines, and concave down if it lies below its tangent lines. In the usual calculus language,

$$f''(x) > 0$$

indicates concave up, while

$$f''(x) < 0$$

indicates concave down.

An inflection point is a point where the concavity changes. A necessary condition, when f'' exists, is

$$f''(x_0) = 0.$$

But this is not sufficient. The note highlights that $f''(x_0) = 0$ or nonexistence of $f''(x_0)$ only gives candidates; one must check whether the sign of f'' changes across x_0 .

The scanned page distinguishes several types: smooth horizontal inflection, smooth non-horizontal inflection, and sharp or vertical tangent behavior. The visual test is concavity change, not merely whether the tangent exists.

§90.5 Asymptotes

The notes classify asymptotes.

A vertical asymptote occurs when

$$\lim_{x \rightarrow x_0} f(x) = \pm\infty.$$

Then $x = x_0$ is a vertical asymptote.

A horizontal asymptote occurs when

$$\lim_{x \rightarrow +\infty} f(x) = c \quad \text{or} \quad \lim_{x \rightarrow -\infty} f(x) = c.$$

Then $y = c$ is a horizontal asymptote.

For an oblique asymptote

$$y = ax + b,$$

we require

$$\lim_{x \rightarrow \infty} [f(x) - (ax + b)] = 0.$$

The page derives

$$a = \lim_{x \rightarrow \infty} \frac{f(x)}{x},$$

and then

$$b = \lim_{x \rightarrow \infty} (f(x) - ax).$$

These formulas are highlighted in the source.

§90.6 General procedure for sketching a curve

The pages give a curve-sketching checklist:

- (i) determine the domain;
- (ii) check symmetry and periodicity;
- (iii) compute $f'(x)$ and $f''(x)$;
- (iv) find critical points and possible inflection points;
- (v) determine intervals of monotonicity and concavity;
- (vi) find asymptotes;
- (vii) locate intercepts or special points if useful;
- (viii) draw the graph according to this information.

This is not a mechanical ritual. The derivative table tells how the curve moves; the second derivative table tells how it bends; asymptotes determine the large-scale shape.

§90.7 Basic antiderivative table

The integration section records standard formulas:

$$\int 0 \, dx = C, \quad \int x^\mu \, dx = \frac{x^{\mu+1}}{\mu+1} + C \quad (\mu \neq -1),$$

$$\int \frac{1}{x} \, dx = \ln |x| + C, \quad \int a^x \, dx = \frac{a^x}{\ln a} + C.$$

Trigonometric formulas include

$$\int \sin x \, dx = -\cos x + C, \quad \int \cos x \, dx = \sin x + C,$$

$$\int \sec^2 x \, dx = \tan x + C, \quad \int \csc^2 x \, dx = -\cot x + C,$$

$$\int \frac{dx}{\sqrt{1-x^2}} = \arcsin x + C, \quad \int \frac{dx}{1+x^2} = \arctan x + C.$$

The page also lists logarithmic forms such as

$$\int \frac{dx}{\sqrt{x^2+a^2}} = \ln |x + \sqrt{x^2+a^2}| + C.$$

§90.8 Substitution

For a composite function, the substitution rule is

$$\int f(\arphi(v))\arphi'(v) \, dv = \int f(u) \, du, \quad u = \arphi(v).$$

Equivalently,

$$\int f(u) \, du = \int f(\arphi(v))\arphi'(v) \, dv.$$

The scanned page calls the two directions the first and second substitution methods. The first recognizes a derivative already present; the second introduces a new variable to simplify the integral.

§90.9 Integration by parts

From

$$d(uv) = u dv + v du,$$

one obtains

$$\int u dv = uv - \int v du.$$

The page presents integration by parts as the other major method after substitution. Its use depends on choosing u so that differentiating it simplifies the expression and choosing dv so that it can be integrated.

§90.10 Rational functions and partial fractions

A rational function is

$$\frac{P(x)}{Q(x)},$$

where P, Q are polynomials and $Q \neq 0$. If $\deg P \geq \deg Q$, first divide polynomials to write it as a polynomial plus a proper rational function.

The pages list the elementary partial fraction forms. A real linear factor $(x - a)^k$ contributes

$$\frac{A_1}{x - a} + \frac{A_2}{(x - a)^2} + \cdots + \frac{A_k}{(x - a)^k}.$$

An irreducible quadratic factor $(x^2 + px + q)^k$ with $p^2 - 4q < 0$ contributes

$$\frac{M_1x + N_1}{x^2 + px + q} + \cdots + \frac{M_kx + N_k}{(x^2 + px + q)^k}.$$

Thus integration of rational functions reduces to two basic integral types:

$$\int \frac{dx}{(x - a)^k}, \quad \int \frac{Mx + N}{(x^2 + px + q)^k} dx.$$

§90.11 Quadratic denominator reduction

For the second type, split the numerator as a derivative part plus a remainder:

$$Mx + N = \frac{M}{2}(2x + p) + \left(N - \frac{M}{2}p\right).$$

Then

$$\int \frac{Mx + N}{(x^2 + px + q)^k} dx = \frac{M}{2} I_k^1 + \left(N - \frac{M}{2}p\right) I_k^2,$$

where

$$I_k^1 = \int \frac{d(x^2 + px + q)}{(x^2 + px + q)^k},$$

and

$$I_k^2 = \int \frac{dx}{(x^2 + px + q)^k}.$$

Complete the square by setting

$$t = x + \frac{p}{2}, \quad a^2 = q - \frac{p^2}{4} > 0.$$

Then

$$x^2 + px + q = t^2 + a^2.$$

For $k = 1$,

$$I_1^2 = \frac{1}{a} \arctan \frac{t}{a} + C.$$

For higher k , the recurrence shown on the page is

$$I_{k+1}^2 = \frac{1}{2ka^2} \left(\frac{t}{(t^2 + a^2)^k} + (2k - 1)I_k^2 \right).$$

§90.12 Trigonometric rational functions

The notes next handle rational functions of $\sin x$ and $\cos x$. Use the tangent half-angle substitution

$$t = an \frac{x}{2}, \quad -\pi < x < \pi.$$

Then

$$\sin x = \frac{2t}{1+t^2}, \quad \cos x = \frac{1-t^2}{1+t^2}, \quad an x = \frac{2t}{1-t^2},$$

and

$$dx = \frac{2 dt}{1+t^2}.$$

Thus

$$\int R(\sin x, \cos x) dx = \int R\left(\frac{2t}{1+t^2}, \frac{1-t^2}{1+t^2}\right) \frac{2 dt}{1+t^2},$$

which becomes a rational integral in t .

§90.13 Definite integrals and Riemann sums

The page defines the definite integral. Divide $[a, b]$ by

$$a = x_0 < x_1 < \cdots < x_n = b,$$

let

$$\Delta x_i = x_i - x_{i-1},$$

and choose sample points $\xi_i \in [x_{i-1}, x_i]$. The Riemann sum is

$$\sum_{i=1}^n f(\xi_i) \Delta x_i.$$

If the limit exists as the mesh

$$\lambda = \max_i \Delta x_i$$

tends to zero, then

$$\int_a^b f(x) dx$$

is the definite integral.

The pages list sufficient conditions for integrability: continuous functions, monotone bounded functions, and bounded functions with finitely many discontinuities are integrable on closed intervals.

§90.14 Basic properties of definite integrals

The standard properties are:

$$\int_a^b [\alpha f(x) + \beta g(x)] dx = \alpha \int_a^b f(x) dx + \beta \int_a^b g(x) dx,$$

$$\int_a^b f(x) dx = - \int_b^a f(x) dx, \quad \int_a^a f(x) dx = 0,$$

$$\int_a^b f(x) dx = \int_a^c f(x) dx + \int_c^b f(x) dx.$$

If $f(x) \geq 0$ on $[a, b]$, then

$$\int_a^b f(x) dx \geq 0.$$

If $f \leq g$, then

$$\int_a^b f(x) dx \leq \int_a^b g(x) dx.$$

§90.15 Integral mean value theorem

Theorem 90.15.1 (Integral mean value theorem)

If f is continuous on $[a, b]$, then there exists $c \in [a, b]$ such that

$$\int_a^b f(x) dx = f(c)(b - a).$$

If f is continuous on $[a, b]$, then there exists $\xi \in [a, b]$ such that

$$\int_a^b f(x) dx = f(\xi)(b - a).$$

If g does not change sign, the weighted form is

$$\int_a^b f(x)g(x) dx = f(\xi) \int_a^b g(x) dx.$$

The note uses this to interpret integrals as average values.

§90.16 Applications: area, volume, and arc length

Area between curves is computed by

$$A = \int_a^b [f(x) - g(x)] dx$$

when $f \geq g$. If horizontal slicing is better,

$$A = \int_c^d [x_{right}(y) - x_{left}(y)] dy.$$

The volume of a solid of revolution around the x -axis is

$$V = \pi \int_a^b [f(x)]^2 dx.$$

The shell method has the form

$$V = 2\pi \int_a^b x f(x) dx$$

when rotating a region around the y -axis.

Arc length of $y = f(x)$ is

$$L = \int_a^b \sqrt{1 + [f'(x)]^2} dx.$$

For a parametric curve $x = x(t), y = y(t)$,

$$L = \int_\alpha^\beta \sqrt{(x'(t))^2 + (y'(t))^2} dt.$$

§90.17 Improper integrals

The last pages include reminders about improper integrals. If the interval is unbounded,

$$\int_a^{+\infty} f(x) dx = \lim_{R \rightarrow +\infty} \int_a^R f(x) dx.$$

If the integrand has a singularity at b , then

$$\int_a^b f(x) dx = \lim_{t \rightarrow b^-} \int_a^t f(x) dx.$$

Convergence depends on the existence of these limits.

A standard comparison is the p -integral:

$$\int_1^\infty \frac{dx}{x^p}$$

converges iff $p > 1$, while

$$\int_0^1 \frac{dx}{x^p}$$

converges iff $p < 1$.

§90.18 Broader perspective

This note connects derivative information with the global shape of a curve, then connects antiderivatives with accumulation. Derivatives decide monotonicity, extrema, concavity, inflection points, and asymptotes. Integration methods reduce complicated antiderivatives to known forms. Definite integrals turn limits of sums into area, volume, length, and average value. The same calculus ideas are being used both for qualitative geometry and quantitative computation.

§90.19 Monotonicity and derivative sign

If $f' \geq 0$ on an interval, the mean value theorem gives

$$f(y) - f(x) = f'(c)(y - x) \geq 0 \quad (x < y).$$

So monotonicity is not a separate rule; it is a global consequence of local derivative sign.

Strict monotonicity needs care: $f' > 0$ everywhere is sufficient, but not necessary. Functions can be strictly increasing even when the derivative vanishes at isolated or many points.

§90.20 Convexity and supporting lines

Convexity can be read geometrically: the graph lies below chords and above supporting tangent lines. If f is differentiable and convex, then

$$f(y) \geq f(x) + f'(x)(y - x).$$

This inequality is the foundation of many optimization methods.

§90.21 Integration methods as inverse rules

Substitution reverses the chain rule, integration by parts reverses the product rule, and partial fractions use algebraic decomposition before integration. These methods are not unrelated tricks; each integration technique recognizes which differentiation rule produced the integrand.

§90.22 Error terms and numerical integration

Midpoint and trapezoidal error formulas depend on second derivatives. They express how curvature controls the difference between an integral and a simple area approximation. In numerical analysis, this becomes a systematic study of quadrature rules and convergence rates.

§90.23 Local derivative data and global shape

Curve sketching, monotonicity, extrema, concavity, asymptotes, integration methods, and numerical errors all use derivatives to convert local information into global shape or accumulated quantity.

Calculus Review III: Convexity, Integral Inequalities, Wallis Formula, and More

N-20260106

§91.1 Convexity and the Hermite–Hadamard inequality

The first blackboard block considers a function f continuous on $[a, b]$ and twice differentiable on (a, b) with

$$f''(x) > 0.$$

This means the graph is convex. The visible statement is the Hermite–Hadamard inequality:

$$f\left(\frac{a+b}{2}\right) \leq \frac{1}{b-a} \int_a^b f(x) dx \leq \frac{f(a) + f(b)}{2}.$$

The left inequality says that the value at the midpoint is no larger than the average value of a convex function. The right inequality says that the average value is no larger than the average of the endpoint values.

The geometric proof shown on the board uses two facts. First, a convex curve lies below the chord joining its endpoints:

$$f(x) \leq \frac{b-x}{b-a}f(a) + \frac{x-a}{b-a}f(b).$$

Integrating this gives the upper bound. Second, symmetry around the midpoint and convexity imply

$$f(x) + f(a+b-x) \geq 2f\left(\frac{a+b}{2}\right).$$

Integrating from a to b gives the lower bound.

The handwritten comment points to the geometric picture: the tangent line is below a convex graph, while the chord is above it.

§91.2 Kantorovich-type integral inequality

Another framed problem assumes that f is continuous and bounded between two positive constants:

$$0 < \lambda_1 \leq f(x) \leq \lambda_2, \quad x \in [a, b].$$

The goal is the estimate

$$\left(\int_a^b f(x) dx\right) \left(\int_a^b \frac{dx}{f(x)}\right) \leq \frac{(\lambda_1 + \lambda_2)^2}{4\lambda_1\lambda_2} (b-a)^2.$$

The board proof starts from

$$(f(x) - \lambda_1)(f(x) - \lambda_2) \leq 0.$$

After expanding and dividing by $f(x) > 0$, one obtains

$$f(x) + \frac{\lambda_1 \lambda_2}{f(x)} \leq \lambda_1 + \lambda_2.$$

Integrating gives

$$A + \lambda_1 \lambda_2 B \leq (\lambda_1 + \lambda_2)(b - a),$$

where

$$A = \int_a^b f(x) dx, \quad B = \int_a^b \frac{dx}{f(x)}.$$

Since for positive numbers u, v one has $uv \leq ((u + v)/2)^2$, we get

$$\lambda_1 \lambda_2 AB \leq \frac{(\lambda_1 + \lambda_2)^2}{4} (b - a)^2.$$

This is the desired inequality.

The note's important warning is that the usual Cauchy inequality gives the lower bound

$$\left(\int_a^b f \right) \left(\int_a^b \frac{1}{f} \right) \geq (b - a)^2,$$

whereas this problem asks for an upper bound under a two-sided bound on f .

§91.3 Integral mean value with odd symmetry

A lower-left blackboard block considers a function defined on a symmetric interval and transforms two integrals. The visible pattern is

$$\int_0^x f(t) dt + \int_0^{-x} f(t) dt = \int_0^x [f(t) - f(-t)] dt.$$

If f is differentiable near 0, then the integral mean value theorem gives a number $heta \in (0, 1)$ such that

$$\int_0^x [f(t) - f(-t)] dt = x [f(hetax) - f(-hetax)].$$

When $x \rightarrow 0$, the quotient comparison uses

$$f(hetax) - f(-hetax) \approx 2hetax f'(0),$$

while

$$\int_0^x [f(t) - f(-t)] dt \approx \int_0^x 2t f'(0) dt = x^2 f'(0).$$

Thus the natural limiting value is

$$heta \rightarrow \frac{1}{2}.$$

The board notes this as a typical use of the integral mean value theorem plus local linearization.

§91.4 A lower bound involving f''/f

Another problem assumes

$$f \in C^2[a, b], \quad f(x) > 0 \text{ on } (a, b), \quad f(a) = f(b) = 0.$$

The blackboard proves an inequality of the form

$$\int_a^b \left| \frac{f''(x)}{f(x)} \right| dx > \frac{4}{b-a}.$$

The idea is to let x_0 be a point where f attains its positive maximum. By the mean value theorem, there exist

$$\xi_1 \in (a, x_0), \quad \xi_2 \in (x_0, b)$$

such that

$$f'(\xi_1) = \frac{f(x_0)}{x_0 - a}, \quad f'(\xi_2) = -\frac{f(x_0)}{b - x_0}.$$

Since $f(x) \leq f(x_0)$, one has

$$\int_a^b \left| \frac{f''(x)}{f(x)} \right| dx \geq \frac{1}{f(x_0)} \int_a^b |f''(x)| dx.$$

The total variation of f' is at least the drop from $f'(\xi_1)$ to $f'(\xi_2)$, so

$$\frac{1}{f(x_0)} \int_a^b |f''(x)| dx \geq \frac{1}{x_0 - a} + \frac{1}{b - x_0}.$$

Finally,

$$\frac{1}{x_0 - a} + \frac{1}{b - x_0} \geq \frac{4}{b - a}.$$

The proof is a good example of turning a global integral estimate into a maximum-point argument.

§91.5 Symmetry of integrals about the midpoint

The second page records a common symmetry fact. If

$$f(x) = f(a + b - x),$$

then

$$\int_a^b f(x) dx = 2 \int_a^{(a+b)/2} f(x) dx.$$

Indeed, substitute

$$u = a + b - x$$

in the right half of the integral. This midpoint symmetry is used repeatedly in the board calculations.

More generally, for any integrable f ,

$$\int_a^b f(x) dx = \int_a^b f(a + b - x) dx,$$

and therefore

$$\int_a^b f(x) dx = \frac{1}{2} \int_a^b [f(x) + f(a + b - x)] dx.$$

This averaging trick is useful whenever $f(x) + f(a + b - x)$ simplifies.

§91.6 Midpoint and trapezoidal error terms

The board also derives the error term for the midpoint rule. Let

$$m = \frac{a+b}{2}.$$

Taylor's formula around m gives

$$f(x) = f(m) + f'(m)(x-m) + \frac{1}{2}f''(\xi_x)(x-m)^2.$$

After integrating, the odd term vanishes, and the second derivative can be pulled out by the integral mean value theorem. Thus there exists $\xi \in (a, b)$ such that

$$\int_a^b f(x) dx = (b-a)f(m) + \frac{(b-a)^3}{24}f''(\xi).$$

Similarly, the trapezoidal rule error has the form

$$\int_a^b f(x) dx = \frac{b-a}{2}[f(a) + f(b)] - \frac{(b-a)^3}{12}f''(\xi)$$

for some $\xi \in (a, b)$. The board comments that confusing these constants is a common mistake: the midpoint coefficient is $1/24$, while the trapezoidal coefficient is $1/12$ with the opposite sign.

§91.7 First special integral

One board example evaluates

$$I = \int_0^1 \frac{\ln(1+x)}{1+x^2} dx.$$

Set

$$x = ant, \quad 0 \leq t \leq \frac{\pi}{4}.$$

Then $dx/(1+x^2) = dt$, and

$$I = \int_0^{\pi/4} \ln(1+ant) dt.$$

Use the substitution

$$t \mapsto \frac{\pi}{4} - t.$$

Since

$$1 + an\left(\frac{\pi}{4} - t\right) = \frac{2}{1+ant},$$

one obtains

$$I = \int_0^{\pi/4} [\ln 2 - \ln(1+ant)] dt.$$

Adding the two expressions for I gives

$$2I = \frac{\pi}{4} \ln 2,$$

hence

$$I = \frac{\pi}{8} \ln 2.$$

The trick is not the tangent substitution alone; it is pairing the interval by the involution $t \leftrightarrow \pi/4 - t$.

§91.8 Positivity of a Fresnel-type integral

Another board example proves

$$\int_0^{\sqrt{2\pi}} \sin x^2 dx > 0.$$

Let

$$t = x^2, \quad dx = \frac{dt}{2\sqrt{t}}.$$

Then

$$\int_0^{\sqrt{2\pi}} \sin x^2 dx = \frac{1}{2} \int_0^{2\pi} \frac{\sin t}{\sqrt{t}} dt.$$

Split the interval into $[0, \pi]$ and $[\pi, 2\pi]$. In the second part use $t = u + \pi$; since $\sin(u + \pi) = -\sin u$, the sum becomes

$$\frac{1}{2} \int_0^\pi \sin u \left(\frac{1}{\sqrt{u}} - \frac{1}{\sqrt{u + \pi}} \right) du.$$

For $u \in (0, \pi)$, both factors are positive, so the integral is positive.

The board circles the comparison of the two weights. The oscillation alone is not enough; the positive half has larger weight because $1/\sqrt{t}$ decreases.

§91.9 A symmetry integral with arcsine

The lower-left board example evaluates

$$I = \int_0^1 \frac{\arcsin \sqrt{x}}{\sqrt{x(1-x)}} dx.$$

There are two neat ways. First,

$$d(\arcsin \sqrt{x}) = \frac{dx}{2\sqrt{x(1-x)}}.$$

Hence

$$I = 2 \int_0^1 \arcsin \sqrt{x} d(\arcsin \sqrt{x}) = [\arcsin \sqrt{x}]^2 \Big|_0^1 = \frac{\pi^2}{4}.$$

The board also uses symmetry:

$$\arcsin \sqrt{x} + \arcsin \sqrt{1-x} = \frac{\pi}{2}.$$

Averaging x and $1-x$ gives

$$I = \frac{1}{2} \int_0^1 \frac{\arcsin \sqrt{x} + \arcsin \sqrt{1-x}}{\sqrt{x(1-x)}} dx = \frac{\pi}{4} \int_0^1 \frac{dx}{\sqrt{x(1-x)}}.$$

The last integral equals π , so again $I = \pi^2/4$.

§91.10 A logarithmic sine integral

The lower-right board example is

$$I = \int_0^{\pi/2} \sin x \ln(\sin x) dx.$$

Let $u = \cos x$. Then $du = -\sin x dx$ and

$$I = \int_0^1 \ln \sqrt{1-u^2} du = \frac{1}{2} \int_0^1 \ln(1-u^2) du.$$

Since

$$\ln(1-u^2) = \ln(1-u) + \ln(1+u),$$

we use

$$\int_0^1 \ln(1-u) du = -1,$$

and

$$\int_0^1 \ln(1+u) du = 2 \ln 2 - 1.$$

Therefore

$$I = \ln 2 - 1.$$

This example trains the habit of choosing $u = \cos x$ when the factor $\sin x dx$ appears.

§91.11 Integrals of $1/(a + b \cos x)$

The final page records two standard results. If $a > 0$ and $|b| < a$, then on the interval $[-\pi/2, \pi/2]$,

$$\int_{-\pi/2}^{\pi/2} \frac{dx}{a + b \cos x} = \frac{2}{\sqrt{a^2 - b^2}} \arccos \frac{b}{a}.$$

Using

$$\arccos \frac{b}{a} = 2 \arctan \sqrt{\frac{a-b}{a+b}},$$

this can also be written as

$$\frac{4}{\sqrt{a^2 - b^2}} \arctan \sqrt{\frac{a-b}{a+b}}.$$

On the interval $[0, \pi]$, the standard formula is

$$\int_0^\pi \frac{dx}{a + b \cos x} = \frac{\pi}{\sqrt{a^2 - b^2}}, \quad a > 0, |b| < a.$$

The note remarks that this follows directly from the universal tangent half-angle substitution, and that symmetry over the full half-period makes the final result simpler.

§91.12 Wallis integral and Wallis product

The last block recalls the Wallis integral

$$I_n = \int_0^{\pi/2} \sin^n x \, dx.$$

Integration by parts gives the recurrence

$$I_n = \frac{n-1}{n} I_{n-2} \quad (n \geq 2).$$

Indeed,

$$I_n = \int_0^{\pi/2} \sin^{n-1} x \sin x \, dx,$$

choose

$$u = \sin^{n-1} x, \quad dv = \sin x \, dx.$$

The boundary term vanishes and the remaining integral becomes

$$\frac{n-1}{n} I_{n-2}.$$

Comparing even and odd subsequences I_{2n} and I_{2n+1} and letting $n \rightarrow \infty$ gives the Wallis product.

§91.13 A wider frame

This note is about integral problem-solving as pattern recognition. Convexity gives chord and midpoint inequalities. Bounded positive functions give product estimates. Symmetry turns an integral into an average with its reflected version. Taylor's formula gives quadrature error terms. Special substitutions, especially $x = ant$ and $t = x^2$, turn difficult integrals into symmetric or sign-controlled expressions. Wallis' recurrence shows that integration by parts can convert a family of integrals into a product formula.

§91.14 Convexity as second-order positivity

For twice differentiable functions, convexity on an interval follows from

$$f''(x) \geq 0.$$

The Hermite–Hadamard inequality

$$f\left(\frac{a+b}{2}\right) \leq \frac{1}{b-a} \int_a^b f(x) \, dx \leq \frac{f(a)+f(b)}{2}$$

compares midpoint value, average value, and endpoint average for convex functions.

§91.15 Integral inequalities as weighted averages

Many integral inequalities are Jensen-type statements with a probability measure. If $w \geq 0$ and $\int w = 1$, then

$$f\left(\int xw(x) \, dx\right) \leq \int f(x)w(x) \, dx$$

for convex f . This is Jensen's inequality in measure form.

§91.16 Symmetry and cancellation

Integrals over symmetric intervals often vanish because an odd part cancels:

$$\int_{-a}^a f_{\text{odd}}(x) dx = 0.$$

More generally, symmetry lets one project a function onto invariant and anti-invariant parts. This is the same idea behind Fourier even/odd decompositions.

§91.17 Wallis formula and asymptotic products

Wallis-type integrals encode asymptotic behavior of products and beta functions. The recurrence for

$$I_n = \int_0^{\pi/2} \sin^n x dx$$

leads to ratios of products and ultimately to approximations involving π . This is an early bridge from calculus to asymptotic analysis.

§91.18 Integral inequalities through convexity and asymptotics

The integral inequalities in this review are not isolated estimates. They are consequences of convexity, weighted averages, symmetry, error analysis, and asymptotic product behavior. The advanced habit is to identify which structural feature of the integrand controls the inequality.

IX

Multivariable Calculus

Part IX: Contents

N-20220405	599
92.1 Fréchet differentiability and complex differentiability	599
92.2 Fréchet derivative	599
92.3 Visualizing complex maps	600
92.4 Hessian matrices and Taylor expansion	600
92.5 Bordered Hessian	600
92.6 Example	601
92.7 Finite fields and Galois theory	602
92.8 Why complex analysis meets Hessians	602
92.9 Frechet derivative as linear approximation	602
92.10 Hessian as second derivative	602
N-20220501	603
93.1 Why this note was too compressed before	603
93.2 Lagrange multipliers: the basic idea	603
93.3 Example: minimizing distance to a line	603
93.4 Why the method is valid	604
93.5 Hessian matrix and the second derivative test	604
93.6 Sylvester criterion and eigenvalues	605
93.7 Bordered Hessian	605
93.8 Example revisited with a bordered Hessian	605
93.9 Degenerate cases and higher-order tests	606
93.10 Saddle points and geometric visualization	606
93.11 Optimization beyond equality constraints	606
93.12 Global topology and extrema	607
93.13 Computation and MATLAB	607
93.14 Second variation on the tangent space	607
93.15 A useful afterword	608
93.16 Lagrange multipliers as cotangent annihilators	608
93.17 Bordered Hessian and restricted second variation	608
93.18 Degenerate extrema and higher order	608
93.19 First and second variations under constraints	608
N-20220728	609
94.1 Spatial vectors	609
94.2 Coordinate formulas	609
94.3 Cross product	610
94.4 From products to universal properties	610
94.5 Categories and groupoids	610
94.6 Cross product geometry	611
94.7 Categorical product versus vector product	611
94.8 What changes next	611
94.9 Cross product from the Hodge star	611
94.10 Scalar triple product and volume	611
94.11 Universal products versus vector products	611
94.12 Metric, orientation, and universal products	612
94.13 Cross product and Hodge duality	612
94.14 Scalar triple product as volume form	612
94.15 Lie algebra of rotations	612
94.16 Categorical versus vector products	613
94.17 Why the cross product is exceptional	613
94.18 Angular velocity and infinitesimal rotations	613
94.19 Universal property versus geometric product	613
N-20230205	615
95.1 The vector triple product	615
95.2 Scalar triple product	615

95.3	From vector identities to exterior algebra	615
95.4	Levi-Civita notation	616
95.5	Lagrange identity	616
95.6	Why tensors enter naturally	617
95.7	Coordinate-free moral	617
95.8	The hidden structure	617
95.9	Levi-Civita symbol as a tensor	618
95.10	Contraction identities	618
95.11	Coordinate-free meaning of vector identities	618
95.12	Returning to the original vector identities	618
95.13	From vector identities to index-free proofs	619
95.14	Why tensor notation matters	619
N-20240413		621
96.1	The mapmaker's problem	621
96.2	Derivative first, partial derivatives second	621
96.3	The chain rule is not a formula trick	622
96.4	A small laboratory: polar coordinates	622
96.5	Tangent planes from equations	623
96.6	A manifold is a space with local Euclidean windows	623
96.7	Two pictures of tangent vectors	623
96.8	Regular values: one clean machine for making manifolds	624
96.9	One chart in detail: stereographic projection	624
96.10	Vector fields: tangent spaces glued along a space	625
96.11	Pushforward in a concrete example	625
96.12	What curvature will later add	625
96.13	How the local-linear picture is used	626
N-20260304		627
97.1	Why point-set language appears	627
97.2	Balls and neighborhoods	627
97.3	Interior, exterior, and boundary	627
97.4	Open and closed sets	628
97.5	Accumulation points and closure	629
97.6	Boundedness and compactness	629
97.7	Connectedness and path-connectedness	630
97.8	What this vocabulary controls	630
N-20260306		631
98.1	What changes from one variable to many	631
98.2	How to prove and disprove limits	631
98.3	Continuity	632
98.4	Partial and directional derivatives	632
98.5	Differentiability as linear approximation	633
98.6	Tangent planes and normal vectors	633
98.7	Mixed partial derivatives	634
98.8	The main lesson	635
N-20260313		637
99.1	The chain rule is composition of linear maps	637
99.2	One input variable, two intermediate variables	637
99.3	Two input variables	638
99.4	Total differentials	638
99.5	Implicit differentiation in one equation	639
99.6	Implicit differentiation for systems	639
99.7	A shape table	640
99.8	What to carry forward	640
N-20260314		641
100.1	Purpose of this guide	641
100.2	Classifying points relative to a set	641

100.3	Limit templates	641
100.4	Degree comparison for rational expressions	642
100.5	Differentiability checklist	642
100.6	Tangent and normal templates	643
100.7	Local extrema template	643
100.8	Common failure modes	644
N-20260325		645
101.1	Local extrema	645
101.2	First-order necessary condition	645
101.3	Quadratic model and the Hessian	645
101.4	Second derivative test	646
101.5	A picture of the Hessian test	647
101.6	Constrained extrema	647
101.7	Endpoint and boundary warnings	648

Fréchet Differentiability, Hessians, and Constrained Tests

N-20220405

§92.1 Fréchet differentiability and complex differentiability

The first part of the note asks why complex differentiability is so much stronger than ordinary real differentiability. In one real direction, a limit may exist; in the complex plane, however, the difference quotient must have the same limit along every direction.

Consider

$$f(z) = \bar{z}, \quad z = a + bi.$$

The complex derivative would be

$$f'(z) = \lim_{h \rightarrow 0} \frac{\overline{z+h} - \bar{z}}{h} = \lim_{h \rightarrow 0} \frac{\bar{h}}{h}.$$

Along the real direction $h = c \rightarrow 0$,

$$\frac{\bar{h}}{h} = \frac{c}{c} = 1.$$

Along the imaginary direction $h = di \rightarrow 0$,

$$\frac{\bar{h}}{h} = \frac{-di}{di} = -1.$$

Along the diagonal direction $h = c + ci$,

$$\frac{\bar{h}}{h} = \frac{c - ci}{c + ci} = -i.$$

The directional limits disagree, so $f(z) = \bar{z}$ is not complex differentiable. This is the basic reason complex differentiability is rigid.

If $f(z) = u(x, y) + iv(x, y)$, the Cauchy–Riemann equations are

$$u_x = v_y, \quad u_y = -v_x.$$

When these hold with suitable regularity, the complex derivative is

$$f'(z_0) = u_x(x_0, y_0) + iv_x(x_0, y_0).$$

§92.2 Fréchet derivative

Let V, W be normed vector spaces, and let $U \subseteq V$ be open. A function $f : U \rightarrow W$ is Fréchet differentiable at $x \in U$ if there exists a bounded linear operator $A : V \rightarrow W$ such that

$$\lim_{\|h\| \rightarrow 0} \frac{\|f(x+h) - f(x) - Ah\|_W}{\|h\|_V} = 0.$$

Equivalently,

$$f(x+h) = f(x) + Ah + o(\|h\|).$$

The operator A is unique and is denoted

$$Df(x) = A.$$

The note compares this with the Gâteaux differential. If $k = \varepsilon h$, then

$$f(x + \varepsilon h) = f(x) + \varepsilon(Df)h + h o(\varepsilon),$$

and, in the limit, one gets a directional derivative. But the Fréchet derivative is stronger: it requires one linear map to approximate the function uniformly in all directions.

§92.3 Visualizing complex maps

The page shows conformal grid pictures: a rectangular grid is sent to curved orthogonal families under maps such as

$$f(z) = z^2, \quad f(z) = \frac{1}{z}, \quad f(z) = \sin z.$$

The key geometric point is that holomorphic maps with nonzero derivative preserve angles locally. They may bend the grid, but they preserve the orthogonality of the two coordinate families.




Figure 92.1: Schematic reconstruction: a rectangular grid is deformed into orthogonal curve families by a holomorphic map.

§92.4 Hessian matrices and Taylor expansion

The second main topic is the relationship between Hessians, Taylor expansion, and extrema. For a smooth function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, Taylor expansion at a critical point p gives

$$f(p + h) = f(p) + \frac{1}{2}h^T H_f(p)h + O(\|h\|^3).$$

Thus the quadratic form associated to the Hessian controls the local behavior.

For unconstrained extrema, positive definiteness of $H_f(p)$ gives a local minimum, negative definiteness gives a local maximum, and indefiniteness gives a saddle. For constrained extrema the situation is subtler, because one must restrict the second variation to the tangent space of the constraint.

§92.5 Bordered Hessian

For a constrained problem

$$\text{extremize } f(x) \quad \text{subject to } g_1(x) = \cdots = g_m(x) = 0,$$

one introduces the Lagrangian

$$\mathcal{L}(x, \lambda) = f(x) - \sum_{i=1}^m \lambda_i g_i(x).$$

The bordered Hessian packages the second derivatives of $\mathcal{L}$ together with the first derivatives of the constraints:

$$H_B = \begin{pmatrix} 0 & Dg \\ Dg^T & D_x^2 \mathcal{L} \end{pmatrix}.$$

The sign pattern of certain principal minors then diagnoses constrained maxima and minima.

§92.6 Example

The note considers the inequality problem

$$a, b, c > 0, \quad ab + bc + ca = a + b + c,$$

and asks to prove

$$\sqrt{ab} + \sqrt{bc} + \sqrt{ca} - \sqrt{abc} \geq 2.$$

Let

$$f(a, b, c) = \sqrt{ab} + \sqrt{bc} + \sqrt{ca} - \sqrt{abc},$$

and

$$g(a, b, c) = a + b + c - ab - bc - ca.$$

The Lagrangian is

$$F(a, b, c, \lambda) = f(a, b, c) + \lambda g(a, b, c).$$

The first-order equations are

$$\begin{cases} \frac{\sqrt{abc} - \sqrt{ab} + \sqrt{ac}}{2a} + \lambda(1 - b - c) = 0, \\ \frac{\sqrt{abc} + \sqrt{ab} + \sqrt{bc}}{2b} + \lambda(1 - a - c) = 0, \\ \frac{\sqrt{abc} + \sqrt{ac} + \sqrt{bc}}{2c} + \lambda(1 - a - b) = 0, \\ a + b + c - ab - bc - ca = 0. \end{cases}$$

The symmetric point

$$a = b = c = 1$$

is a critical point. At that point the bordered Hessian has the schematic form

$$H_B = \begin{pmatrix} 0 & -1 & -1 & -1 \\ -1 & -\frac{1}{4} & 0 & 0 \\ -1 & 0 & -\frac{1}{4} & 0 \\ -1 & 0 & 0 & -\frac{1}{4} \end{pmatrix}.$$

The relevant minors have the required sign pattern, so the symmetric point gives the constrained extremum. Since

$$f(1, 1, 1) = 2,$$

we obtain the desired inequality.

§92.7 Finite fields and Galois theory

A final aside in the note jumps to finite fields and Galois theory. The basic facts listed are:

- (i) the finite field $\mathbb{F}_p$ has arithmetic modulo p ;
- (ii) polynomial factorization over $\mathbb{F}_p$ can differ dramatically from factorization over $\mathbb{Q}$;
- (iii) adjoining roots creates finite extensions such as $\mathbb{F}_{p^n}$;
- (iv) quotient rings $\mathbb{F}_p[x]/(q(x))$ construct finite fields when q is irreducible.

Thus the same algebraic idea reappears: to solve an equation, enlarge the field, then study the symmetries of the roots.

§92.8 Why complex analysis meets Hessians

The original pages feel discontinuous because they jump from complex differentiability to Hessians and bordered Hessians. There is nevertheless a shared theme: first-order information tells us what linear approximation is allowed, and second-order information tells us what happens after the first-order obstruction vanishes.

For $f(z) = \bar{z}$, the real derivative exists as a real-linear map, but complex differentiability asks for multiplication by one complex number. The failure is exactly the failure of a real-linear map to commute with multiplication by i . In modern language, complex differentiability means the derivative is $\mathbb{C}$ -linear, not merely $\mathbb{R}$ -linear.

For constrained extrema, the first-order condition removes the tangent freedom, so the Hessian test must be performed only on tangent directions to the constraint. This is why the bordered Hessian appears: it is not a mysterious determinant recipe, but a bookkeeping device for second variation under first-order constraints.

§92.9 Frechet derivative as linear approximation

For a map $F : V \rightarrow W$ between normed spaces, differentiability at a means there is a bounded linear map L such that

$$F(a + h) = F(a) + Lh + o(\|h\|).$$

This L is the Frechet derivative. It is stronger and cleaner than checking directional derivatives separately, because it requires one linear map to approximate the function uniformly in all directions.

§92.10 Hessian as second derivative

For scalar-valued functions, the Hessian is the derivative of the gradient. It records the quadratic part of the Taylor expansion:

$$f(a + h) = f(a) + Df(a)h + \frac{1}{2}h^T H_f(a)h + o(\|h\|^2).$$

At a critical point, the linear term vanishes, so the Hessian controls the first visible local shape.

Lagrange Multipliers, Hessians, Bordered Hessians, and Degenerate Extrema

N-20220501

§93.1 Why this note was too compressed before

The note is not just “Lagrange multipliers plus Hessian.” It is a long reflection on how to decide extrema of multivariable functions. It contains the method, the derivation, examples, doubts, bordered Hessian, degeneracy, higher-order tests, numerical methods, MATLAB, and links to optimization theory. The repaired version preserves that exploration.

§93.2 Lagrange multipliers: the basic idea

Suppose we want to find extrema of

$$f(x_1, \dots, x_n)$$

under the constraint

$$g(x_1, \dots, x_n) = 0.$$

The Lagrange function is

$$L(x_1, \dots, x_n, \lambda) = f(x_1, \dots, x_n) - \lambda g(x_1, \dots, x_n).$$

The stationarity equations are

$$\frac{\partial L}{\partial x_i} = 0, \quad i = 1, \dots, n, \quad \frac{\partial L}{\partial \lambda} = 0.$$

Equivalently,

$$\nabla f = \lambda \nabla g, \quad g = 0.$$

Geometrically, at a constrained extremum, the level surface of f is tangent to the constraint surface. Therefore their normal directions are parallel. This explains why gradients are proportional.

The original Chinese note says that this method is both surprising and confusing: surprising because it transforms a constrained problem into a stationarity system, confusing because the multiplier seems to appear out of nowhere. The geometric explanation removes much of the mystery: λ is the scalar measuring how the two normal vectors compare.

§93.3 Example: minimizing distance to a line

Consider

$$f(x, y) = x^2 + y^2$$

under the constraint

$$x + y = 1.$$

Set

$$L(x, y, \lambda) = x^2 + y^2 - \lambda(x + y - 1).$$

The equations are

$$2x - \lambda = 0, \quad 2y - \lambda = 0, \quad x + y = 1.$$

The first two give $x = y$, and therefore $x = y = 1/2$. This is the closest point of the line $x + y = 1$ to the origin.

The note remarks that this example is simple but educational. For nonlinear constraints, the stationarity system can become hard to solve; in practice one may need numerical methods such as Newton iteration. The Lagrange method gives equations, not automatic closed forms.

§93.4 Why the method is valid

A more rigorous proof uses the implicit function theorem. If $\nabla g(p) \neq 0$, then near p the constraint $g = 0$ is a smooth hypersurface. A tangent vector v to the constraint satisfies

$$\nabla g(p) \cdot v = 0.$$

If p is a constrained extremum of f , then the derivative of f along every tangent direction must vanish:

$$\nabla f(p) \cdot v = 0 \quad \text{for all tangent } v.$$

Thus $\nabla f(p)$ is perpendicular to the tangent space. The normal space is spanned by $\nabla g(p)$, so

$$\nabla f(p) = \lambda \nabla g(p).$$

For several constraints $g_1 = \cdots = g_m = 0$, the condition becomes

$$\nabla f = \sum_{i=1}^m \lambda_i \nabla g_i,$$

assuming the constraint gradients are independent.

This is where the note's later warning about singular constraints enters: if the gradients of constraints are linearly dependent, the usual bordered-Hessian test may fail.

§93.5 Hessian matrix and the second derivative test

For a twice differentiable function $f(x_1, \dots, x_n)$, the Hessian matrix at x_0 is

$$H_f(x_0) = \left(\frac{\partial^2 f}{\partial x_i \partial x_j}(x_0) \right)_{i,j}.$$

If the mixed partials are continuous, the Hessian is symmetric. This lets us use the linear algebra of symmetric matrices.

At an unconstrained critical point $\nabla f(x_0) = 0$:

- (i) if $H_f(x_0)$ is positive definite, x_0 is a local minimum;

- (ii) if $H_f(x_0)$ is negative definite, x_0 is a local maximum;
- (iii) if $H_f(x_0)$ is indefinite, x_0 is a saddle point;
- (iv) if $H_f(x_0)$ is singular or semidefinite, the second derivative test is inconclusive.

The proof comes from Taylor expansion:

$$f(x_0 + h) = f(x_0) + \frac{1}{2}h^T H_f(x_0)h + o(\|h\|^2)$$

when the gradient term vanishes. Thus the sign of the quadratic form controls the local shape.

§93.6 Sylvester criterion and eigenvalues

For a real symmetric matrix, positive definiteness can be tested by eigenvalues or by leading principal minors. Sylvester's criterion says that A is positive definite iff all leading principal minors are positive. Negative definiteness has alternating signs. Eigenvalues give an equivalent test: positive definite means all eigenvalues are positive; indefinite means there are both positive and negative eigenvalues.

The note explicitly connects this to linear algebra. A Hessian is not just an array of second derivatives; it is a quadratic form. The eigenvectors are principal directions of curvature, and the eigenvalues measure curvature along those directions.

§93.7 Bordered Hessian

For constrained extrema, the ordinary Hessian of f alone is not enough because we are only allowed to move along directions tangent to the constraint. The note introduces the bordered Hessian as a way to add constraint-gradient information to the second derivative test.

For constraints $g_i(x) = 0$, define the Lagrangian

$$L = f - \sum_{i=1}^m \lambda_i g_i.$$

A common bordered Hessian has the block form

$$B = \begin{pmatrix} 0_{m \times m} & (Dg)^T \\ Dg & H_x L \end{pmatrix},$$

where Dg records gradients of the constraints and $H_x L$ is the Hessian with respect to the original variables. Different texts use slightly different sign conventions and ordering of rows, so one must check the convention before applying a sign rule.

The intuitive meaning is: the border records which directions are forbidden by the constraints, and the Hessian block records curvature of the Lagrangian in the remaining directions.

§93.8 Example revisited with a bordered Hessian

For the simple constraint $x + y = 1$, the constraint gradient is

$$\nabla g = (1, 1).$$

If $f(x, y) = x^2 + y^2$, then

$$H_f = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}.$$

The tangent direction to $x + y = 1$ is spanned by $(1, -1)$. Along that direction,

$$(1, -1)H_f(1, -1)^T = 4 > 0,$$

so the constrained critical point is a constrained minimum.

This tangent-space calculation is often clearer than memorizing signs of bordered determinants. The bordered Hessian is a compact determinant method for doing the same restriction systematically.

§93.9 Degenerate cases and higher-order tests

The note explicitly worries about cases where the Hessian is singular or the bordered Hessian test is inconclusive. In such cases, second-order information is not enough. One must examine higher-order Taylor terms:

$$f(x_0 + h) = f(x_0) + P_k(h) + o(\|h\|^k),$$

where P_k is the first nonzero homogeneous term after the constant term. If P_k has a fixed sign, one may still obtain an extremum; if it changes sign, one gets a saddle-type point.

A standard example is $f(x, y) = x^4 + y^4$, whose Hessian at the origin is zero, but the origin is still a strict local minimum. Conversely, $f(x, y) = x^4 - y^4$ also has zero Hessian at the origin but has a saddle. Thus degeneracy forces us to look beyond quadratic approximation.

§93.10 Saddle points and geometric visualization

A saddle point is a critical point that is neither a local maximum nor a local minimum. If the Hessian has both positive and negative eigenvalues, then some directions curve upward and others downward. The note suggests using contour plots to detect saddle behavior. This is a good habit: level curves reveal local shape far more clearly than raw formulas.

The classical picture is $f(x, y) = x^2 - y^2$. Along the x -axis it increases; along the y -axis it decreases. The origin is therefore a saddle.

§93.11 Optimization beyond equality constraints

The note mentions Kuhn–Tucker conditions, penalty methods, and interior-point methods. These belong to constrained optimization with inequalities. If constraints include

$$h_j(x) \leq 0,$$

then Lagrange multipliers are supplemented by nonnegativity and complementary slackness:

$$\mu_j \geq 0, \quad \mu_j h_j(x) = 0.$$

This says that an inequality constraint contributes a force only when it is active. This is the natural extension of the Lagrange multiplier idea.

Penalty methods instead replace constraints by large penalty terms, for example

$$f(x) + \rho \sum_j h_j(x)^2.$$

Interior-point methods keep the iterate inside the feasible region using barrier functions. All of these methods share the same philosophy: convert a constrained problem into a system that can be treated with unconstrained or nearly unconstrained tools.

§93.12 Global topology and extrema

The note speculates about global topology. This is not random. Local Hessian information tells us what happens near a critical point, but global existence of maxima and minima depends on compactness. A continuous function on a compact set attains maximum and minimum. Thus topology enters before calculus even begins.

There are also deeper links: Morse theory studies how critical points of smooth functions reflect the topology of a manifold. The Hessian at a nondegenerate critical point gives the Morse index, and the collection of critical points helps reconstruct the space. This is the high-level version of the note's intuition that extrema are connected with global structure.

§93.13 Computation and MATLAB

The note ends with computational reflection. For high-dimensional functions, one can use symbolic or numerical software to compute gradients, Hessians, eigenvalues, and bordered Hessians. A typical workflow is:

- (i) define f and constraints g_i ;
- (ii) solve $\nabla f = \sum \lambda_i \nabla g_i$ together with $g_i = 0$;
- (iii) compute the Hessian of the Lagrangian;
- (iv) restrict to the tangent space or use a bordered Hessian;
- (v) inspect eigenvalues or determinant signs.

The note's MATLAB idea is therefore mathematically meaningful: computation is not replacing theory, but automating the linear algebra after the theoretical conditions have been derived.

§93.14 Second variation on the tangent space

For constrained extrema, the cleanest conceptual test is to restrict the Hessian of the Lagrangian to tangent directions. If the constraint is $g(x) = 0$, then tangent vectors v satisfy

$$\nabla g(x_0) \cdot v = 0.$$

The second variation is

$$v^T H_x L(x_0, \lambda) v.$$

If this quadratic form is positive definite on the tangent space, the point is a constrained local minimum; if negative definite, a constrained local maximum.

This explains what the bordered Hessian is doing: it is a determinant device for testing the same restricted quadratic form.

§93.15 A useful afterword

This note is about the transition from calculus to optimization and differential geometry. Lagrange multipliers are normal-vector equations. Hessians are quadratic forms. Bordered Hessians restrict second variation to tangent directions. Degenerate cases require higher-order terms. Inequality constraints lead to KKT theory. Compactness guarantees global extrema, and Morse theory relates critical points to topology. The original confusion is productive: it points exactly toward the deeper structure behind the computational rules.

§93.16 Lagrange multipliers as cotangent annihilators

For constraints $g_1 = \cdots = g_k = 0$, the tangent space consists of directions annihilated by all dg_i . At a constrained extremum,

$$df \in \text{span}(dg_1, \dots, dg_k).$$

This is the coordinate-free form of Lagrange multipliers.

§93.17 Bordered Hessian and restricted second variation

The second-order test should be applied only to tangent directions of the constraint. The bordered Hessian is a coordinate device for computing the quadratic form of the Lagrangian on that tangent space.

If the restricted quadratic form is positive definite, the point is a constrained local minimum; if negative definite, a constrained local maximum.

§93.18 Degenerate extrema and higher order

When the second variation has zero directions, the Hessian test is inconclusive. One must examine higher-order terms or use normal forms. This is the optimization analogue of a multiple root in algebra: first-order and second-order information may not separate the cases.

§93.19 First and second variations under constraints

Constrained extrema are governed by tangent and cotangent spaces. Multipliers express first-order normality, bordered Hessians test second-order behavior along constraints, and degeneracy signals the need for higher-order local geometry.

Spatial Vectors, Cross Products, and Categorical Products

N-20220728

§94.1 Spatial vectors

The note begins with spatial vectors. A vector has magnitude and direction. The zero vector has length 0, and a unit vector has length 1. Two vectors are parallel if they lie on parallel or coincident directed lines. In particular, the zero vector is conventionally regarded as parallel to every vector.

If two nonzero vectors $\vec{a}, \vec{b}$ are collinear, then there exists a nonzero scalar λ such that

$$\vec{a} = \lambda \vec{b}.$$

If $\vec{a}, \vec{b}$ are not collinear, then every vector in their plane can be written as

$$\vec{p} = \lambda \vec{a} + \mu \vec{b}.$$

If $\vec{i}, \vec{j}, \vec{k}$ are three non-coplanar vectors, every spatial vector can be written uniquely as

$$\vec{p} = x\vec{i} + y\vec{j} + z\vec{k}.$$

This is the vector-space decomposition theorem in geometric form.

§94.2 Coordinate formulas

In an orthonormal basis, if

$$\vec{a} = (x_a, y_a, z_a), \quad \vec{b} = (x_b, y_b, z_b),$$

then

$$\vec{a} + \vec{b} = (x_a + x_b, y_a + y_b, z_a + z_b),$$

$$\lambda \vec{a} = (\lambda x_a, \lambda y_a, \lambda z_a),$$

$$\vec{a} \cdot \vec{b} = x_a x_b + y_a y_b + z_a z_b.$$

The angle formula is

$$\cos(\vec{a}, \vec{b}) = \frac{\vec{a} \cdot \vec{b}}{|\vec{a}| |\vec{b}|}.$$

The note marks the dot product as especially important. The reason is that it turns geometry into algebra: perpendicularity becomes

$$\vec{a} \cdot \vec{b} = 0,$$

and projection becomes a scalar multiple determined by a dot product.

§94.3 Cross product

The cross product is introduced as a vector perpendicular to both $\vec{a}$ and $\vec{b}$. In coordinates,

$$\vec{a} \times \vec{b} = \begin{vmatrix} \vec{i} & \vec{j} & \vec{k} \\ x_a & y_a & z_a \\ x_b & y_b & z_b \end{vmatrix}.$$

It satisfies

$$|\vec{a} \times \vec{b}| = |\vec{a}| |\vec{b}| \sin \theta,$$

so its length is the area of the parallelogram spanned by $\vec{a}$ and $\vec{b}$.

The cross product is anti-commutative:

$$\vec{a} \times \vec{b} = -(\vec{b} \times \vec{a}).$$

This is not an algebraic nuisance; it records orientation.

§94.4 From products to universal properties

The second half of the note suddenly moves into category theory. It asks what a product really is. In set theory, the product $X \times Y$ is the set of ordered pairs. But category theory characterizes it by its mapping property.

A product of objects X, Y is an object P with maps


$$\pi_X : P \rightarrow X, \quad \pi_Y : P \rightarrow Y$$

such that for every object Z and maps $f : Z \rightarrow X, g : Z \rightarrow Y$, there is a unique map

$$\langle f, g \rangle : Z \rightarrow P$$

with

$$\pi_X \circ \langle f, g \rangle = f, \quad \pi_Y \circ \langle f, g \rangle = g.$$



This explains why products are unique up to unique isomorphism. The point is not the elements of P , but the role it plays.

§94.5 Categories and groupoids

A category consists of objects and morphisms, with identity morphisms and associative composition. Examples include:

$$\mathbf{Set}, \quad \mathbf{Vect}_k, \quad \mathbf{Grp}.$$

A group can be viewed as a category with one object in which every morphism is invertible. A groupoid is a category in which every morphism is invertible, but there may be many objects.

This is why the note places groupoids near groups and categories: a groupoid is a many-object group. It remembers symmetry, but allows different objects connected by reversible arrows.

§94.6 Cross product geometry

For vectors $a, b \in \mathbb{R}^3$, the cross product $a \times b$ is perpendicular to both a and b , and its magnitude is

$$\|a \times b\| = \|a\| \|b\| \sin \theta.$$

Thus it measures the area of the parallelogram spanned by a and b . The direction is determined by the right-hand rule.

The cross product is not merely a computational formula. It combines metric information, orientation, and area. That is why it is special to oriented three-dimensional Euclidean space.

§94.7 Categorical product versus vector product

The word “product” appears in different senses. A Cartesian or categorical product is defined by a universal property. A dot product is a scalar-valued bilinear form. A cross product is a vector-valued alternating operation. Keeping these meanings separate prevents confusion: the same English word does not mean the same mathematical structure.

§94.8 What changes next

This note connects spatial vectors and categorical products through the word “structure.” In vector geometry, a basis decomposes every vector into coordinates. In category theory, a product decomposes maps into pairs of maps. Coordinates answer “what are the components of this vector?” Universal properties answer “what data is equivalent to a map into this object?”

§94.9 Cross product from the Hodge star

In oriented Euclidean three-space, the cross product is defined by

$$u \times v = *(u \wedge v).$$

The wedge product records the oriented parallelogram spanned by u, v , and the Hodge star turns that oriented area element into a vector perpendicular to the plane.

§94.10 Scalar triple product and volume

The scalar triple product

$$u \cdot (v \times w)$$

is the oriented volume of the parallelepiped. Algebraically it is a determinant; geometrically it measures whether three vectors form a positively oriented basis.

§94.11 Universal products versus vector products

The categorical product $X \times Y$ is characterized by projections and a universal property. The vector cross product is not a categorical product; it is a metric-and-orientation-dependent operation special to dimension three.

§94.12 Metric, orientation, and universal products

Spatial vector operations separate three structures: affine addition of vectors, Euclidean metric via dot product, and orientation via determinant/Hodge star. Categorical products belong to universal mapping properties, not to the cross product operation.

§94.13 Cross product and Hodge duality

The cross product in $\mathbb{R}^3$ is often introduced by the determinant formula. The deeper reason it exists is the combination of an inner product, an orientation, and exterior algebra. First form the wedge product

$$u \wedge v \in \Lambda^2 \mathbb{R}^3.$$

This is an oriented area element, not yet a vector. The Hodge star operator determined by the Euclidean metric and orientation gives an isomorphism

$$*: \Lambda^2 \mathbb{R}^3 \rightarrow \Lambda^1 \mathbb{R}^3 \cong \mathbb{R}^3.$$

Then

$$u \times v = *(u \wedge v).$$

This explains why the cross product is special to three dimensions. In general dimensions, the wedge product exists, but there is not usually a vector-valued binary cross product.

§94.14 Scalar triple product as volume form

The scalar triple product is

$$u \cdot (v \times w).$$

In exterior algebra this is the coefficient of the top-degree form:

$$u \wedge v \wedge w = (u \cdot (v \times w)) e_1 \wedge e_2 \wedge e_3.$$

Thus the scalar triple product is oriented volume. The vector triple product identities are coordinate expressions of algebraic identities involving wedge product, contraction, and the metric.

§94.15 Lie algebra of rotations

A vector $\omega \in \mathbb{R}^3$ defines a skew-symmetric matrix $[\omega]_{\times}$ by

$$[\omega]_{\times} v = \omega \times v.$$

Explicitly,

$$[\omega]_{\times} = \begin{pmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{pmatrix}.$$

The space of skew-symmetric 3×3 matrices is the Lie algebra $\mathfrak{so}(3)$ of the rotation group $SO(3)$. The cross product therefore encodes infinitesimal rotations.

The elementary right-hand rule is the visible geometric form of orientation in this Lie algebra.

§94.16 Categorical versus vector products

The original note contrasts categorical products with vector products. In a category, a product $A \times B$ is characterized by projection maps and a universal property: for every object X with maps $X \rightarrow A$ and $X \rightarrow B$, there is a unique map $X \rightarrow A \times B$ making the diagram commute.

The cross product $u \times v$ is not a categorical product. It is an operation depending on dimension, metric, and orientation. This contrast is useful: the same word “product” can mean either a universal construction or a geometric bilinear operation.

§94.17 Why the cross product is exceptional

The cross product is a bilinear map

$$\mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$$

that is perpendicular to both inputs and has magnitude equal to the area of the parallelogram they span. Such a product exists in this familiar form essentially because

$$\Lambda^2 \mathbb{R}^3 \cong \mathbb{R}^3$$

once a metric and orientation are chosen.

In $\mathbb{R}^2$, a wedge product of two vectors is an area scalar. In $\mathbb{R}^4$, a wedge product of two vectors is a genuine 2-vector with six independent components, not an ordinary vector. This explains why formulas involving $u \times v$ should not be blindly generalized to higher dimensions.

§94.18 Angular velocity and infinitesimal rotations

If a rigid body rotates with angular velocity ω , the instantaneous velocity of a point at position r is

$$v = \omega \times r.$$

In matrix form, this is $v = [\omega]_{\times} r$, where $[\omega]_{\times}$ is skew-symmetric. The exponential

$$e^{t[\omega]_{\times}}$$

is a rotation matrix. Thus the cross product is not just a static area operation; it is the infinitesimal generator of rotations in three-dimensional mechanics.

§94.19 Universal property versus geometric product

The categorical product $A \times B$ is determined by projections and a universal mapping property. The vector cross product is determined by metric, orientation, and dimension. Putting them side by side is pedagogically useful because it prevents a common mistake: the same word “product” does not mean the same structure. In advanced mathematics, understanding a construction means knowing which data it depends on and which universal property or invariance characterizes it.

Vector Triple Product, Exterior Algebra, and the Hodge Star

N-20230205

§95.1 The vector triple product

The note begins with a proof of the identity

$$u \times (v \times w) = (u \cdot w)v - (u \cdot v)w.$$

The key observation is that $u \times (v \times w)$ lies in the plane spanned by v and w , because it is perpendicular to $v \times w$. Therefore

$$u \times (v \times w) = Av + Bw.$$

Taking dot products with suitable vectors determines A and B , giving

$$A = u \cdot w, \quad B = -(u \cdot v).$$

This proof is better than expanding a determinant because it explains the geometry: the result must live in the v, w -plane, and only the two scalar projections of u onto v, w can appear.

§95.2 Scalar triple product

The scalar triple product is

$$[u, v, w] = (u \times v) \cdot w.$$

It equals the signed volume of the parallelepiped spanned by u, v, w . It is alternating: swapping two vectors changes the sign. It vanishes exactly when the three vectors are coplanar.

The note uses this to justify identities involving cross products and dot products. The recurring geometric idea is that determinants measure oriented volume.

§95.3 From vector identities to exterior algebra

This note should not try to become a general tensor-calculus note. Its natural center is narrower and cleaner: the familiar three-dimensional vector identities are shadows of the metric, the volume form, wedge product, contraction, and the Hodge star.

Definition 95.3.1 (Bivectors and the wedge product). For vectors u, v in a vector space V , the wedge product $u \wedge v$ is an alternating two-vector. It records the oriented area element spanned by u and v , not a third vector perpendicular to them.

In a three-dimensional oriented Euclidean space, the metric and orientation identify bivectors with vectors. This identification is the Hodge star:

$$\star : \Lambda^2 V \longrightarrow V.$$

The cross product is then

$$u \times v = \star(u \wedge v).$$

Thus the cross product is not the primitive operation. The primitive operation is the oriented area element $u \wedge v$; the cross product appears only after choosing both a metric and an orientation.

Remark 95.3.2 (Why dimension three is special) — In dimension 3, bivectors and vectors both have dimension 3, so the Hodge star can identify them. In other dimensions, $\Lambda^2 V$ and V usually have different dimensions, and there is no ordinary vector-valued cross product with the same properties.

§95.4 Levi-Civita notation

The coordinate proof becomes much shorter with the Levi-Civita symbol. Write

$$(u \times v)_i = \varepsilon_{ijk} u_j v_k,$$

using the Einstein summation convention. Then

$$(u \times (v \times w))_i = \varepsilon_{ijk} u_j \varepsilon_{klm} v_l w_m.$$

The contraction identity

$$\varepsilon_{ijk} \varepsilon_{klm} = \delta_{il} \delta_{jm} - \delta_{im} \delta_{jl}$$

gives

$$(u \times (v \times w))_i = (u_j w_j) v_i - (u_j v_j) w_i.$$

This is exactly

$$u \times (v \times w) = (u \cdot w) v - (u \cdot v) w.$$

This notation is useful because it separates the combinatorics of signs from the geometry of the identity.

§95.5 Lagrange identity

Another important identity is

$$(u \times v) \cdot (w \times z) = (u \cdot w)(v \cdot z) - (u \cdot z)(v \cdot w).$$

Setting $w = u$ and $z = v$ gives

$$\|u \times v\|^2 = \|u\|^2 \|v\|^2 - (u \cdot v)^2.$$

This is the area formula for the parallelogram spanned by u and v . It also shows that Cauchy–Schwarz is equivalent to the nonnegativity of a squared area.

§95.6 Why tensors enter naturally

The cross product is special to three dimensions, but the ideas behind it are not. The determinant, alternating forms, metrics, and contractions are tensorial constructions. A tensor is not a multidimensional array by itself; the array is a coordinate representation of a multilinear object.

For example, the metric tensor lets us convert vectors to covectors. The volume form lets us express oriented volume. In three-dimensional Euclidean space, the cross product is essentially produced by combining the metric with the volume form. This explains why the cross product behaves so specially and why it does not generalize naively to every dimension.

§95.7 Coordinate-free moral

The note's vector identities are training for coordinate-free thinking. Coordinates are useful for checking signs, but the real objects are projections, oriented areas, volumes, and multilinear maps. Once this is understood, tensor notation stops looking like decoration: it is a way to write formulas that remain meaningful after a change of basis.

§95.8 The hidden structure

The vector triple product is a projection identity; the scalar triple product is a volume form; the cross product is a Hodge-star translation of an oriented area element. These three statements give the structural content of the note.

Proposition 95.8.1 (Cross product from wedge product and Hodge star)

In an oriented Euclidean three-dimensional vector space,

$$a \times b = \star(a \wedge b).$$

Consequently,

$$c \cdot (a \times b) = \text{vol}(c, a, b),$$

where vol is the oriented volume form.

Proof. The Hodge star is defined so that pairing a vector c with $\star(a \wedge b)$ gives the oriented volume of c, a, b . In the standard oriented orthonormal basis, this is exactly the usual determinant formula for $c \cdot (a \times b)$. Therefore $\star(a \wedge b)$ has the same defining inner products as $a \times b$. $\square$

This is the right high-level stopping point for the present note. General basis-change rules and tensor representations belong more naturally to the following linear-map note.

§95.9 Levi-Civita symbol as a tensor

The vector identities in this note become systematic with the Levi-Civita symbol. In an oriented orthonormal basis of $\mathbb{R}^3$,

$$\varepsilon_{ijk} = \begin{cases} 1, & (i, j, k) \text{ even permutation of } (1, 2, 3), \\ -1, & (i, j, k) \text{ odd permutation}, \\ 0, & \text{otherwise.} \end{cases}$$

Then

$$(u \times v)_i = \sum_{j,k} \varepsilon_{ijk} u_j v_k.$$

The scalar triple product is

$$u \cdot (v \times w) = \sum_{i,j,k} \varepsilon_{ijk} u_i v_j w_k.$$

This notation reveals the tensorial structure behind determinant formulas.

§95.10 Contraction identities

Many vector identities reduce to contractions of ε_{ijk} . The key identity is

$$\sum_k \varepsilon_{ijk} \varepsilon_{mnk} = \delta_{im} \delta_{jn} - \delta_{in} \delta_{jm}.$$

Using this, one obtains the vector triple product formula

$$u \times (v \times w) = v(u \cdot w) - w(u \cdot v).$$

What looks like a clever mnemonic is actually an index contraction identity.

§95.11 Coordinate-free meaning of vector identities

A tensor is not its array of components. Components change when the basis changes, while the tensorial object remains. A bilinear form, for instance, is a map

$$B : V \times V \rightarrow k.$$

Its matrix depends on a basis, but the operation $B(u, v)$ does not.

The elementary formulas of dot product, cross product, and triple product are therefore best read as components of coordinate-free objects: metrics, alternating forms, and contractions.

§95.12 Returning to the original vector identities

The original note's identities are not separate facts to memorize. Dot products come from a metric, cross products come from wedge product plus Hodge star, and triple products come from the volume form. The tensor notation gives a uniform way to prove and remember all of them.

§95.13 From vector identities to index-free proofs

The vector triple product

$$a \times (b \times c) = b(a \cdot c) - c(a \cdot b)$$

can be derived by index contraction, but it also has a geometric meaning. The vector $b \times c$ is perpendicular to the plane spanned by b, c , so $a \times (b \times c)$ lies in that plane. The right-hand side expresses it as the unique linear combination of b and c with coefficients determined by inner products with a .

In exterior algebra, this identity is a special case of relationships between wedge product, contraction, and Hodge star. If ι_a denotes contraction by a , then expressions of the form

$$\iota_a(b \wedge c)$$

encode projections of a against an oriented area element. The elementary formula is therefore a low-dimensional coordinate shadow of a general calculus of multivectors.

§95.14 Why tensor notation matters

Einstein notation is not merely shorter writing. It separates free indices from contracted indices. Free indices describe the type of the output; repeated indices describe summation. For example,

$$(a \times b)_i = \varepsilon_{ijk} a_j b_k$$

has one free index i , so the result is a vector. The repeated indices j, k are summed.

This discipline prevents type errors. A scalar has no free indices, a vector has one, a matrix has two, and higher tensors have more. The original vector computations become much easier to audit once every expression has a tensor type.

Local Linear Geometry: A Mapmaker's View of Derivatives and Tangent Spaces

N-20240413

§96.1 The mapmaker's problem

Suppose one wants to draw a map of a curved surface. A small enough region of the Earth can be treated as flat, even though the Earth is not flat. This is not a lie; it is a local approximation.

Differential geometry begins from the same idea. Curved objects are understood by replacing them, near one point, with linear objects. The derivative is the machine that performs this replacement.




Figure 96.1: The tangent line is the first linear shadow of a curve. Higher-order details are deliberately ignored.

The guiding principle for this chapter is:

to see a curved object clearly, first learn its best linear shadow.

This principle sounds elementary, but it is the starting point for manifolds, tangent spaces, differential forms, and eventually curvature.

§96.2 Derivative first, partial derivatives second

Let

$$f : \mathbb{R}^n \rightarrow \mathbb{R}^m.$$

At a point p , the derivative should be thought of as a linear map

$$Df_p : \mathbb{R}^n \rightarrow \mathbb{R}^m$$

satisfying

$$f(p+h) = f(p) + Df_p(h) + o(\|h\|).$$

This formula says that $Df_p(h)$ is the best first-order prediction of how f changes when the input moves by h .

The Jacobian matrix is only a coordinate display of this linear map. If $f = (f_1, \dots, f_m)$, then

$$J_f(p) = \begin{bmatrix} \partial f_1/\partial x_1 & \cdots & \partial f_1/\partial x_n \\ \vdots & \ddots & \vdots \\ \partial f_m/\partial x_1 & \cdots & \partial f_m/\partial x_n \end{bmatrix}_p.$$

The matrix is useful, but the linear map is the concept.

A useful correction. Students often learn multivariable calculus as a list of partial derivatives. Geometry reverses the priority: the total derivative is the object, and partial derivatives are its coordinates.

§96.3 The chain rule is not a formula trick

If

$$\mathbb{R}^n \xrightarrow{f} \mathbb{R}^m \xrightarrow{g} \mathbb{R}^k,$$

then

$$D(g \circ f)_p = Dg_{f(p)} \circ Df_p.$$

In coordinates this becomes matrix multiplication. Geometrically it says something simpler: if f is locally approximated by Df_p , and g is locally approximated by $Dg_{f(p)}$, then their composition is locally approximated by the composition of the two linear maps.

This is why the chain rule is one of the first genuinely structural theorems in calculus. It says that differentiation respects composition.

§96.4 A small laboratory: polar coordinates

Consider

$$F(r, \theta) = (r \cos \theta, r \sin \theta).$$

Then

$$DF_{(r,\theta)} = \begin{bmatrix} \cos \theta & -r \sin \theta \\ \sin \theta & r \cos \theta \end{bmatrix}.$$

The two columns are

$$\frac{\partial F}{\partial r} = (\cos \theta, \sin \theta), \quad \frac{\partial F}{\partial \theta} = (-r \sin \theta, r \cos \theta).$$

The first vector points outward. The second vector points around the circle. Thus a coordinate change is not only a substitution rule; it turns coordinate motion into tangent motion.

Notice the factor r in the angular direction. Moving by $d\theta$ near the origin is physically shorter than moving by the same $d\theta$ far away. The derivative records this distortion.

§96.5 Tangent planes from equations

A graph $z = f(x, y)$ has tangent plane

$$z = f(a, b) + f_x(a, b)(x - a) + f_y(a, b)(y - b).$$

A level surface

$$F(x, y, z) = c$$

has tangent plane

$$\nabla F(p) \cdot (X - p) = 0.$$

Both formulas are versions of the same principle: differentiate the defining data and keep only the first-order part.

For the sphere

$$S^2 = \{x \in \mathbb{R}^3 : \|x\|^2 = 1\},$$

let $p \in S^2$. If $\gamma(t) \subset S^2$ and $\gamma(0) = p$, then

$$\gamma(t) \cdot \gamma(t) = 1.$$

Differentiating gives

$$p \cdot \gamma'(0) = 0.$$

So

$$T_p S^2 = \{v \in \mathbb{R}^3 : p \cdot v = 0\}.$$

The tangent plane is perpendicular to the radius vector. This is the most concrete form of the slogan “differentiate the constraint.”

§96.6 A manifold is a space with local Euclidean windows

A smooth manifold is a space that can be covered by coordinate windows

$$\varphi : U \rightarrow \mathbb{R}^n$$

whose overlaps are smoothly compatible. A chart is not the manifold itself. It is a local language for talking about it.

The circle, sphere, torus, and projective spaces all share this feature: no single coordinate system describes them perfectly, but many local coordinate systems can be patched together.

Why charts matter. Charts let us import calculus into spaces that are not globally Euclidean. Differentiation, tangent vectors, and smooth maps are all defined by moving into a chart, doing the Euclidean calculation, and checking that the answer does not depend on the chart in a bad way.

§96.7 Two pictures of tangent vectors

There are two useful definitions of a tangent vector at p .

The first is dynamic: a tangent vector is the velocity of a curve through p :

$$v = \gamma'(0), \quad \gamma(0) = p.$$

The second is operational: a tangent vector differentiates functions:

$$v(f) = \left. \frac{d}{dt} \right|_{t=0} f(\gamma(t)).$$

The second definition is more intrinsic. It does not require the manifold to be sitting inside a larger Euclidean space.

A smooth map $F : M \rightarrow N$ sends tangent vectors forward:

$$dF_p : T_p M \rightarrow T_{F(p)} N.$$

This is the derivative on manifolds.

§96.8 Regular values: one clean machine for making manifolds

Suppose

$$F : \mathbb{R}^n \rightarrow \mathbb{R}^k$$

is smooth. If c is a regular value, meaning DF_p is surjective for every $p \in F^{-1}(c)$, then

$$F^{-1}(c)$$

is a smooth manifold of dimension $n - k$.

This theorem is useful because many geometric objects are naturally described by equations. Spheres, matrix groups, and many surfaces arise this way. The theorem says that if the equations cut cleanly, the solution set is a manifold.

§96.9 One chart in detail: stereographic projection

The sphere is the first place where one really feels why a single coordinate system is not enough. Remove the north pole

$$N = (0, 0, 1)$$

from S^2 . Stereographic projection sends a point $(X, Y, Z) \in S^2 \setminus \{N\}$ to the point where the line through N and (X, Y, Z) meets the plane $Z = 0$:

$$(X, Y, Z) \mapsto \left(\frac{X}{1 - Z}, \frac{Y}{1 - Z} \right).$$

The inverse map is

$$(u, v) \mapsto \left(\frac{2u}{u^2 + v^2 + 1}, \frac{2v}{u^2 + v^2 + 1}, \frac{u^2 + v^2 - 1}{u^2 + v^2 + 1} \right).$$

This formula is worth seeing once because it makes the word “chart” concrete. The sphere is not being flattened globally; one point is missing, and the remaining part is described by ordinary coordinates (u, v) .

A second chart, obtained by projecting from the south pole, covers the missing point. The transition map between the two charts is smooth away from the origin. This is the atlas idea in its most visible form.

§96.10 Vector fields: tangent spaces glued along a space

A tangent vector lives at one point. A vector field chooses a tangent vector at every point:

$$p \longmapsto V(p) \in T_p M.$$

On the plane, one may draw a vector field as little arrows attached to points. On a manifold, this picture is still correct, but each arrow belongs to a different tangent space.

For example, on the unit circle $S^1 \subset \mathbb{R}^2$, the vector field

$$V(x, y) = (-y, x)$$

is tangent to the circle because

$$(x, y) \cdot (-y, x) = 0.$$

It rotates counterclockwise around the circle. This example shows a useful habit: to check whether a vector field is tangent to a constraint, check whether it preserves the constraint to first order.

§96.11 Pushforward in a concrete example

Let

$$F : \mathbb{R}^2 \rightarrow \mathbb{R}^2, \quad F(x, y) = (x^2 - y^2, 2xy).$$

This is the complex squaring map $z \mapsto z^2$ written in real coordinates. Its derivative is

$$DF_{(x,y)} = \begin{bmatrix} 2x & -2y \\ 2y & 2x \end{bmatrix}.$$

At the point $(1, 0)$, a tangent vector (a, b) is sent to

$$(2a, 2b).$$

At the point $(0, 1)$, it is sent to

$$(-2b, 2a).$$

The same coordinate vector can be stretched and rotated differently depending on the base point. This is why the derivative of a nonlinear map is a family of linear maps, one for each point.

§96.12 What curvature will later add

This chapter only keeps first-order information. A tangent plane cannot distinguish a sphere from a plane at the point of tangency. Curvature begins when one asks how tangent spaces change from point to point.

So the local-linear story is not the end of geometry. It is the first layer. The later layers ask:

How do tangent spaces vary?

How do we compare vectors at different points?

How does a manifold bend?

Those questions require connections, metrics, and curvature. But without the tangent-space layer, none of them can even be stated cleanly.

§96.13 How the local-linear picture is used

The usable chain from the chapter is

derivative $\longrightarrow$ linear approximation $\longrightarrow$ tangent space $\longrightarrow$ manifold calculus.

In a calculation this chain has a definite order. First choose a chart or a defining equation. Then replace the nonlinear object by its first-order part: for a parametrized surface this means differentiating the parametrization; for a level set $F^{-1}(c)$ this means studying dF_p . The tangent space is the kernel of the linear constraint when the manifold is cut out by equations, and it is the span of coordinate velocity vectors when the manifold is parametrized.

The stereographic-projection example and the regular-value criterion are two versions of the same habit. A chart lets one compute on $\mathbb{R}^n$; a regular value tells one that the kernel of a derivative has the expected dimension. Curvature, connections, and metrics will later add second-order and comparison data, but the first check remains the same: before asking how a manifold bends, identify the tangent space in which the infinitesimal question lives.

Point-Set Language in $\mathbb{R}^n$: Neighborhoods, Boundaries, and Compactness

N-20260304

§97.1 Why point-set language appears

In one-variable calculus, the domain is often an interval, so the topology is mostly hidden. In several variables, a domain may have holes, corners, isolated points, missing boundary pieces, or several connected components. Pictures help, but pictures can lie. The reliable language is the language of balls, neighborhoods, boundary points, and sequences.

Throughout this note, points are in $\mathbb{R}^n$, and distance means Euclidean distance:

$$\rho(P, Q) = \|P - Q\|_2 = \sqrt{\sum_{i=1}^n (p_i - q_i)^2}.$$

§97.2 Balls and neighborhoods

Definition 97.2.1. For $P_0 \in \mathbb{R}^n$ and $\delta > 0$, the open ball of radius δ around P_0 is

$$B_\delta(P_0) = \{P \in \mathbb{R}^n : \rho(P, P_0) < \delta\}.$$

A set U is a **neighborhood** of P_0 if it contains some open ball $B_\delta(P_0)$.

A neighborhood is the formal version of "all points sufficiently close to P_0 ." This phrase is the engine behind limits, continuity, differentiability, and local extrema.

Example 97.2.2

In $\mathbb{R}^2$, the disk

$$D = \{(x, y) : x^2 + y^2 < 1\}$$

is a neighborhood of $(0, 0)$, because it contains $B_{1/2}(0, 0)$. The punctured disk

$$D \setminus \{(0, 0)\}$$

is not a neighborhood of $(0, 0)$, although every small ball around the origin meets it.

§97.3 Interior, exterior, and boundary

Definition 97.3.1. Let $E \subseteq \mathbb{R}^n$. A point P is

- (i) an **interior point** of E if some ball around P is contained in E ;
- (ii) an **exterior point** of E if some ball around P is disjoint from E ;
- (iii) a **boundary point** of E if every ball around P meets both E and $\mathbb{R}^n \setminus E$.

The set of all boundary points is denoted ∂E .

These three classes partition $\mathbb{R}^n$. The boundary is where local membership cannot be decided by zooming in.

Example 97.3.2

For

$$E = \{(x, y) : x^2 + y^2 < 1\},$$

the interior is E itself, the exterior is $\{(x, y) : x^2 + y^2 > 1\}$, and the boundary is the circle

$$\partial E = \{(x, y) : x^2 + y^2 = 1\}.$$

For the closed disk $\{x^2 + y^2 \leq 1\}$, the boundary is the same circle. Boundary is not the same as "points excluded from the set."




Figure 97.1: Interior points have a small ball inside the set; boundary points never do.

§97.4 Open and closed sets

Definition 97.4.1. A set $E \subseteq \mathbb{R}^n$ is **open** if every point of E is an interior point. It is **closed** if its complement is open.

The two words are not opposites in the everyday sense. A set can be both open and closed, or neither open nor closed.

Example 97.4.2

In $\mathbb{R}$, the interval $(0, 1)$ is open but not closed, $[0, 1]$ is closed but not open, $[0, 1)$ is neither open nor closed, and both $\emptyset$ and $\mathbb{R}$ are both open and closed.

Proposition 97.4.3

A set $E \subseteq \mathbb{R}^n$ is closed if and only if it contains all its boundary points.

Proof skeleton. If E is closed, its complement is open, so no boundary point can lie in the complement; hence every boundary point lies in E . Conversely, if E contains all boundary points, then any point outside E is not an interior point of E and not a boundary point, so it is exterior. Thus the complement is open. $\square$

§97.5 Accumulation points and closure

Definition 97.5.1. A point P is an **accumulation point** of E if every punctured ball

$$B_\delta(P) \setminus \{P\}$$

contains a point of E . The **closure** of E , denoted $\overline{E}$, is the union of E with all its accumulation points.

Proposition 97.5.2 (Sequential test)

For $E \subseteq \mathbb{R}^n$, a point P lies in $\overline{E}$ if and only if there is a sequence (P_k) in E such that $P_k \rightarrow P$.

Proof skeleton. If $P \in \overline{E}$, choose $P_k \in E \cap B_{1/k}(P)$, using $P_k = P$ if $P \in E$ and necessary. Then $P_k \rightarrow P$. Conversely, if $P_k \in E$ and $P_k \rightarrow P$, then every ball around P contains some P_k , so P is in the closure. $\square$

This gives the most useful closed-set criterion:

$$E \text{ is closed} \iff \text{every convergent sequence in } E \text{ has its limit in } E.$$

§97.6 Boundedness and compactness

Definition 97.6.1. A set $E \subseteq \mathbb{R}^n$ is **bounded** if it is contained in some ball $B_R(0)$. It is **compact** if every open cover of E has a finite subcover.

In $\mathbb{R}^n$, compactness has a simpler working form.

Theorem 97.6.2 (Heine–Borel)

A subset $K \subseteq \mathbb{R}^n$ is compact if and only if it is closed and bounded.

For computation, compactness is an existence machine. It makes "largest," "smallest," and "convergent subsequence" statements true.

Theorem 97.6.3 (Extreme value theorem)

If $K \subseteq \mathbb{R}^n$ is compact and $f : K \rightarrow \mathbb{R}$ is continuous, then f attains a maximum and a minimum on K .

Example 97.6.4

The function $f(x, y) = x^2 + y^2$ has a maximum and a minimum on the closed disk $x^2 + y^2 \leq 1$. It has a minimum but no maximum on the open disk $x^2 + y^2 < 1$. The missing boundary is exactly where the supremum would be reached.

§97.7 Connectedness and path-connectedness

Definition 97.7.1. A set E is **connected** if it cannot be written as the union of two nonempty separated pieces. It is **path-connected** if for any two points $P, Q \in E$, there is a continuous curve $\gamma : [0, 1] \rightarrow E$ with $\gamma(0) = P$ and $\gamma(1) = Q$.

Path-connectedness implies connectedness. In the elementary domains of multivariable calculus, the two often coincide, but they are not the same in general.

A **region** in many calculus texts usually means an open connected subset of $\mathbb{R}^n$, sometimes together with a piecewise smooth boundary when integration is involved.

§97.8 What this vocabulary controls

The words in this note are not decorative definitions. They decide which theorems are legal:

open set	local perturbations stay inside the domain
closed set	limits of internal sequences remain inside
compact set	continuous functions attain extrema
boundary	where constraints and integration surfaces live
connectedness	where one can move without jumping

When later notes discuss limits, tangent planes, extrema, or multiple integrals, these are the hidden hypotheses that make the calculations legitimate.

Multivariable Limits, Continuity, Differentiability, and the Total Derivative

N-20260306

§98.1 What changes from one variable to many

In one variable, approaching a point means moving from the left or from the right. In several variables, there are infinitely many ways to approach a point: lines, curves, spirals, parabolas, and irregular paths. This is why multivariable limits are stricter than they first appear.

Definition 98.1.1. For $f : D \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^m$, the statement

$$\lim_{x \rightarrow a} f(x) = L$$

means that for every $\varepsilon > 0$ there exists $\delta > 0$ such that

$$0 < \|x - a\| < \delta, \ x \in D \implies \|f(x) - L\| < \varepsilon.$$

The definition requires all sufficiently close points, not only points on convenient lines.

§98.2 How to prove and disprove limits

There are two basic strategies.

To prove a limit, dominate the expression by a function of $r = \|x - a\|$ that tends to zero:

$$|f(x) - L| \leq g(\|x - a\|), \quad g(r) \rightarrow 0.$$

To disprove a limit, find two paths approaching the same point that give different limiting values.

Example 98.2.1 (Two paths are enough to disprove)

Let

$$f(x, y) = \begin{cases} \frac{xy}{x^2 + y^2}, & (x, y) \neq (0, 0), \\ 0, & (x, y) = (0, 0). \end{cases}$$

Along the x -axis, $f(x, 0) = 0$. Along the line $y = x$,

$$f(x, x) = \frac{x^2}{2x^2} = \frac{1}{2}.$$

Thus the limit at the origin does not exist.

Example 98.2.2 (Lines are not enough)

Let

$$f(x, y) = \begin{cases} \frac{x^2 y}{x^4 + y^2}, & (x, y) \neq (0, 0), \\ 0, & (x, y) = (0, 0). \end{cases}$$

Along every line $y = mx$, the value tends to 0. But along $y = x^2$,

$$f(x, x^2) = \frac{x^4}{2x^4} = \frac{1}{2}.$$

So checking all straight lines still does not prove a two-variable limit.

§98.3 Continuity

Definition 98.3.1. A map $f : D \rightarrow \mathbb{R}^m$ is **continuous** at $a \in D$ if

$$\lim_{x \rightarrow a} f(x) = f(a).$$

It is continuous on D if it is continuous at every point of D .

Continuity is not just a pointwise property; it interacts strongly with compactness and connectedness. A continuous real-valued function on a compact set attains maxima and minima; a continuous image of a connected set is connected.

§98.4 Partial and directional derivatives

For $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, the partial derivatives at (x_0, y_0) are

$$\begin{aligned} f_x(x_0, y_0) &= \lim_{h \rightarrow 0} \frac{f(x_0 + h, y_0) - f(x_0, y_0)}{h}, \\ f_y(x_0, y_0) &= \lim_{h \rightarrow 0} \frac{f(x_0, y_0 + h) - f(x_0, y_0)}{h}. \end{aligned}$$

For a unit vector u , the directional derivative is

$$D_u f(a) = \lim_{t \rightarrow 0} \frac{f(a + tu) - f(a)}{t}.$$

Partial derivatives measure coordinate-direction changes. Directional derivatives measure changes along chosen lines. Neither, by itself, guarantees differentiability.

Example 98.4.1

For the function

$$f(x, y) = \begin{cases} \frac{xy}{x^2 + y^2}, & (x, y) \neq (0, 0), \\ 0, & (x, y) = (0, 0), \end{cases}$$

both partial derivatives at the origin exist and equal 0, because the function is zero on both axes. But the function is not continuous at the origin, hence cannot be differentiable there.

§98.5 Differentiability as linear approximation

The correct higher-dimensional analogue of differentiability is not "all partial derivatives exist." It is "there is a good linear approximation."

Definition 98.5.1. A map $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is **differentiable** at a if there is a linear map $L : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that

$$\lim_{h \rightarrow 0} \frac{\|F(a+h) - F(a) - Lh\|}{\|h\|} = 0.$$

The linear map L is the **total derivative** or **Fréchet derivative**, denoted $DF(a)$.

In coordinates, $DF(a)$ is the Jacobian matrix

$$DF(a) = \left(\frac{\partial F_i}{\partial x_j}(a) \right)_{i,j}.$$

For scalar-valued f , the derivative is represented by the row vector $\nabla f(a)^T$, so

$$f(a+h) = f(a) + \nabla f(a) \cdot h + o(\|h\|).$$

Proposition 98.5.2

If F is differentiable at a , then F is continuous at a . Moreover, for every direction u , the directional derivative exists and is

$$D_u F(a) = DF(a)u.$$

Proof. Differentiability gives

$$F(a+h) - F(a) = DF(a)h + r(h), \quad \frac{\|r(h)\|}{\|h\|} \rightarrow 0.$$

Both $DF(a)h$ and $r(h)$ tend to 0, so continuity follows. For the directional derivative, put $h = tu$ and divide by t . $\square$

Theorem 98.5.3 (A useful sufficient condition)

If all first partial derivatives of F exist in a neighborhood of a and are continuous at a , then F is differentiable at a .

The safe hierarchy is therefore

$$C^1 \implies \text{differentiable} \implies \text{continuous} \implies \text{partial derivatives need not exist}.$$

Existence of partial derivatives alone sits outside this chain.

§98.6 Tangent planes and normal vectors

For a graph $z = f(x, y)$, differentiability at (x_0, y_0) gives the tangent plane

$$z - z_0 = f_x(x_0, y_0)(x - x_0) + f_y(x_0, y_0)(y - y_0),$$

where $z_0 = f(x_0, y_0)$. This is just the linear approximation written geometrically.

For an implicit surface $F(x, y, z) = 0$, if $\nabla F(P_0) \neq 0$, the tangent plane is

$$\nabla F(P_0) \cdot (P - P_0) = 0.$$

The gradient is normal to the level surface.

Example 98.6.1

For the sphere $F(x, y, z) = x^2 + y^2 + z^2 - 1 = 0$, at $P_0 = (0, 0, 1)$,

$$\nabla F(P_0) = (0, 0, 2).$$

The tangent plane is therefore

$$(0, 0, 2) \cdot (x, y, z - 1) = 0,$$

so $z = 1$.

§98.7 Mixed partial derivatives

For a surface $z = f(x, y)$, the mixed partial derivative

$$f_{xy} = \partial_x(\partial_y f)$$

asks how the y -direction slope changes when one moves in the x -direction. Under the usual smoothness hypotheses, the order does not matter.

Theorem 98.7.1 (Clairaut)

If f_{xy} and f_{yx} are continuous in a neighborhood of a point, then

$$f_{xy} = f_{yx}$$

at that point.




Figure 98.1: A mixed partial derivative measures how one coordinate slope changes as we move in the other coordinate direction.

§98.8 The main lesson

The word "derivative" in several variables means "linear map that captures first-order change." Partial derivatives are entries of that map, directional derivatives are evaluations of that map, tangent planes are pictures of that map, and the Jacobian is its matrix representation.

Multivariable Chain Rule, Total Differentials, and Implicit Differentiation

N-20260313

§99.1 The chain rule is composition of linear maps

The chain rule is not a collection of unrelated formulas. It says that first-order approximations compose.

Theorem 99.1.1 (Chain rule)

Let $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be differentiable at x , and let $G : \mathbb{R}^m \rightarrow \mathbb{R}^k$ be differentiable at $F(x)$. Then $G \circ F$ is differentiable at x , and

$$D(G \circ F)(x) = DG(F(x)) DF(x).$$

Proof skeleton. Write

$$F(x+h) = F(x) + DF(x)h + r(h), \quad \|r(h)\| = o(\|h\|).$$

Then apply the differentiability expansion for G at $F(x)$:

$$G(F(x) + u) = G(F(x)) + DG(F(x))u + s(u), \quad \|s(u)\| = o(\|u\|).$$

Substitute $u = DF(x)h + r(h)$. Since $u = O(\|h\|)$, the remainder is still $o(\|h\|)$. The linear part is $DG(F(x))DF(x)h$. $\square$

This theorem is the source of all ordinary partial-derivative chain rules.

§99.2 One input variable, two intermediate variables

Suppose

$$z = f(u, v), \quad u = u(x), \quad v = v(x).$$

Then

$$\frac{dz}{dx} = f_u \frac{du}{dx} + f_v \frac{dv}{dx}.$$

In matrix form, this is

$$\frac{dz}{dx} = \begin{pmatrix} f_u & f_v \end{pmatrix} \begin{pmatrix} u_x \\ v_x \end{pmatrix}.$$

The row vector is Df , the column vector is the derivative of (u, v) , and their product is the derivative of the composite.

Example 99.2.1

Let $z = u^2 + uv$, $u = x^2$, and $v = \sin x$. Then

$$z_u = 2u + v, \quad z_v = u, \quad u' = 2x, \quad v' = \cos x.$$

Hence

$$\frac{dz}{dx} = (2u + v)2x + u \cos x.$$

Substituting $u = x^2$, $v = \sin x$ gives

$$\frac{dz}{dx} = 2x(2x^2 + \sin x) + x^2 \cos x.$$

§99.3 Two input variables

Suppose

$$z = f(u, v), \quad u = u(x, y), \quad v = v(x, y).$$

Then

$$\begin{pmatrix} z_x & z_y \end{pmatrix} = \begin{pmatrix} f_u & f_v \end{pmatrix} \begin{pmatrix} u_x & u_y \\ v_x & v_y \end{pmatrix}.$$

Equivalently,

$$z_x = f_u u_x + f_v v_x, \quad z_y = f_u u_y + f_v v_y.$$

Example 99.3.1

Let $z = e^u \cos v$, where $u = x^2 + y^2$ and $v = x - y$. Then

$$f_u = e^u \cos v, \quad f_v = -e^u \sin v,$$

and

$$\begin{pmatrix} u_x & u_y \\ v_x & v_y \end{pmatrix} = \begin{pmatrix} 2x & 2y \\ 1 & -1 \end{pmatrix}.$$

Thus

$$z_x = e^u \cos v \cdot 2x - e^u \sin v, \quad z_y = e^u \cos v \cdot 2y + e^u \sin v.$$

§99.4 Total differentials

The notation

$$dz = f_x dx + f_y dy$$

is best read as a compact way to write a linear approximation. If (x, y) changes by (dx, dy) , then the first-order change in $z = f(x, y)$ is

$$dz \approx f_x dx + f_y dy.$$

For a vector-valued map F , the same statement is

$$dF = J_F dx.$$

The differential notation is safe when one remembers that it represents a linear map, not ordinary fraction arithmetic.

§99.5 Implicit differentiation in one equation

Suppose a relation

$$F(x, y) = 0$$

defines y locally as a differentiable function of x . Differentiating the identity $F(x, y(x)) = 0$ gives

$$F_x + F_y \frac{dy}{dx} = 0.$$

If $F_y \neq 0$, then

$$\frac{dy}{dx} = -\frac{F_x}{F_y}.$$

Example 99.5.1

For the circle $F(x, y) = x^2 + y^2 - 1 = 0$, where $y \neq 0$,

$$\frac{dy}{dx} = -\frac{2x}{2y} = -\frac{x}{y}.$$

This is the slope of the tangent line to the circle.

§99.6 Implicit differentiation for systems

The same idea becomes matrix algebra. Suppose

$$F(x, u) = 0, \quad F : \mathbb{R}^p \times \mathbb{R}^q \rightarrow \mathbb{R}^q,$$

and u is locally a function of x . Differentiating gives

$$F_x dx + F_u du = 0.$$

If F_u is invertible, then

$$du = -F_u^{-1} F_x dx,$$

so

$$Du(x) = -F_u^{-1} F_x.$$

Theorem 99.6.1 (Implicit function theorem, working form)

Let $F : \mathbb{R}^p \times \mathbb{R}^q \rightarrow \mathbb{R}^q$ be C^1 , and suppose

$$F(a, b) = 0, \quad \det F_u(a, b) \neq 0.$$

Then near (a, b) , the equation $F(x, u) = 0$ determines u as a C^1 function of x , and

$$Du(a) = -F_u(a, b)^{-1} F_x(a, b).$$

The theorem says when the formal linearized computation is legitimate. The condition $\det F_u \neq 0$ means the constraint can be solved in the u -directions.

§99.7 A shape table

Keeping track of shapes prevents many mistakes:

map	derivative shape	meaning
$f : \mathbb{R}^n \rightarrow \mathbb{R}$	$1 \times n$	row gradient
$F : \mathbb{R}^n \rightarrow \mathbb{R}^m$	$m \times n$	Jacobian
$G \circ F$	$k \times n$	$(k \times m)(m \times n)$
$F(x, u) = 0$	$q \times p, \quad q \times q$	F_x, F_u

The formulas work because the linear maps have compatible domains and codomains.

§99.8 What to carry forward

The chain rule says that local linear approximations compose. Total differentials are the notation for those approximations. Implicit differentiation says that a constraint can be solved by inverting the linearized constraint in the unknown directions.

A Problem-Solving Guide for Multivariable Limits and Local Geometry

N-20260314

§100.1 Purpose of this guide

This note is a compact problem-solving companion to the previous multivariable calculus notes. It should not replace the definitions. Its job is to say what to try first, what each method proves, and where each method fails.

§100.2 Classifying points relative to a set

For a set $G \subseteq \mathbb{R}^n$, classify a point P by inspecting small balls around it:

classification	ball test
interior	$\exists r > 0, B_r(P) \subseteq G$
exterior	$\exists r > 0, B_r(P) \cap G = \emptyset$
boundary	$\forall r > 0, B_r(P) \cap G \neq \emptyset \text{ and } B_r(P) \cap G^c \neq \emptyset$

Example 100.2.1

For

$$G = \{(x, y) : 0 < x^2 + y^2 < 1\},$$

the origin is not in G , but it is a boundary point. Every ball around the origin contains points with $0 < x^2 + y^2 < 1$, and also contains the origin itself, which lies outside G .

§100.3 Limit templates

To prove a limit at the origin, try to convert the expression into polar form

$$x = r \cos \theta, \quad y = r \sin \theta.$$

If the expression becomes

$$r^\alpha g(\theta)$$

where $\alpha > 0$ and g is bounded, then the limit is 0.

Example 100.3.1

For

$$f(x, y) = \frac{x^2 y}{x^2 + y^2},$$

we have

$$f(r \cos \theta, r \sin \theta) = r \cos^2 \theta \sin \theta.$$

Since $|\cos^2 \theta \sin \theta| \leq 1$, the limit is 0.

To disprove a limit, find paths giving different values. The common path choices are:

$$y = 0, \quad x = 0, \quad y = mx, \quad y = x^2, \quad x = y^2.$$

Caution

Path testing can disprove a limit, but it usually cannot prove one. Even checking every straight line is not enough, because a curved path may still detect different behavior.

§100.4 Degree comparison for rational expressions

For rational expressions near the origin, compare the lowest total degrees of numerator and denominator only after checking whether the denominator is comparable to a power of r . The estimate

$$x^2 + y^2 = r^2$$

is reliable. The expression $x^4 + y^2$, however, is anisotropic: along $y = x^2$, the two terms have the same size.

Example 100.4.1

The function

$$f(x, y) = \frac{x^2 y}{x^4 + y^2}$$

looks as if the numerator has degree 3 and the denominator has degree 2, but the denominator changes scale depending on direction. Along $y = x^2$, the limit is 1/2. Degree-counting alone fails.

§100.5 Differentiability checklist

For a scalar function $f(x, y)$, use the following hierarchy:

$$C^1 \implies \text{differentiable} \implies \text{continuous}.$$

Partial derivatives alone do not imply continuity. Directional derivatives in every direction alone do not imply differentiability.

To prove differentiability at a , the definition asks for

$$f(a + h) = f(a) + L(h) + o(\|h\|).$$

The candidate linear map is usually

$$L(h) = \nabla f(a) \cdot h.$$

Then one must estimate the remainder.

§100.6 Tangent and normal templates

For a graph $z = f(x, y)$, the tangent plane at (x_0, y_0, z_0) is

$$z - z_0 = f_x(x_0, y_0)(x - x_0) + f_y(x_0, y_0)(y - y_0).$$

For an implicit surface $F(x, y, z) = 0$, the normal vector is ∇F , and the tangent plane is

$$\nabla F(P_0) \cdot (P - P_0) = 0.$$

For a space curve given by the intersection

$$F(x, y, z) = 0, \quad G(x, y, z) = 0,$$

a tangent direction is

$$\nabla F(P_0) \times \nabla G(P_0),$$

provided the two normals are not parallel.

Example 100.6.1

The curve

$$x^2 + y^2 + z^2 = 1, \quad z = x$$

at a point P_0 has tangent direction

$$\nabla(x^2 + y^2 + z^2 - 1)(P_0) \times \nabla(z - x)(P_0).$$

This works because the curve lies on both surfaces, so its tangent vector must be perpendicular to both surface normals.

§100.7 Local extrema template

For an unconstrained interior extremum of a differentiable function, first solve

$$\nabla f = 0.$$

Then classify candidates using the Hessian when possible. For two variables,

$$H = \begin{pmatrix} A & B \\ B & C \end{pmatrix}, \quad D = AC - B^2.$$

At a critical point,

condition	classification
$D > 0, A > 0$	local minimum
$D > 0, A < 0$	local maximum
$D < 0$	saddle
$D = 0$	inconclusive

For a constraint $g(x, y) = 0$, first try

$$\nabla f = \lambda \nabla g.$$

Geometrically, the level curve of f is tangent to the constraint curve at a constrained extremum.

§100.8 Common failure modes

- (i) Checking axes proves almost nothing about a two-variable limit.
- (ii) Checking all lines still may not prove a limit.
- (iii) Existence of partial derivatives is not differentiability.
- (iv) The Hessian test classifies only critical points, and it is inconclusive when $D = 0$.
- (v) Lagrange multipliers require checking regularity and comparing all candidate points, including boundary or endpoint cases when the constraint has them.

The common idea is local replacement: balls replace vague closeness, radial estimates replace path guessing, and linear maps replace nonlinear functions to first order.

Extrema of Two-Variable Functions and the Hessian Test

N-20260325

§101.1 Local extrema

Let $f : G \subseteq \mathbb{R}^2 \rightarrow \mathbb{R}$. A point $P_0 \in G$ is a **local maximum** if there is a neighborhood U of P_0 such that

$$f(P) \leq f(P_0) \quad (P \in U \cap G).$$

It is a **local minimum** if the inequality is reversed. If the inequality is strict for all nearby $P \neq P_0$, the extremum is strict.

The word "local" matters. A point can be best among nearby points without being best on the whole domain.

§101.2 First-order necessary condition

Theorem 101.2.1 (Fermat theorem in several variables)

If f is differentiable at an interior point $P_0 \in G$ and P_0 is a local extremum, then

$$\nabla f(P_0) = 0.$$

Proof. For every direction u , the one-variable function

$$\varphi(t) = f(P_0 + tu)$$

has a local extremum at $t = 0$. Hence $\varphi'(0) = 0$. But

$$\varphi'(0) = \nabla f(P_0) \cdot u.$$

This vanishes for every u , so $\nabla f(P_0) = 0$. □

A point with $\nabla f = 0$ is called a **critical point**. Critical points are candidates, not conclusions.

§101.3 Quadratic model and the Hessian

At a critical point P_0 , the Taylor expansion begins with the quadratic term:

$$f(P_0 + h) = f(P_0) + \frac{1}{2}h^T H(P_0)h + o(\|h\|^2),$$

where

$$H(P_0) = \begin{pmatrix} f_{xx}(P_0) & f_{xy}(P_0) \\ f_{yx}(P_0) & f_{yy}(P_0) \end{pmatrix}.$$

When the second partial derivatives are continuous, $f_{xy} = f_{yx}$, so the Hessian is symmetric.

Thus the classification problem becomes a problem about a quadratic form.

§101.4 Second derivative test

Let

$$A = f_{xx}(P_0), \quad B = f_{xy}(P_0), \quad C = f_{yy}(P_0), \quad D = AC - B^2.$$

Theorem 101.4.1 (Second derivative test)

Assume $\nabla f(P_0) = 0$ and the second partial derivatives are continuous near P_0 . Then:

condition	conclusion
$D > 0, A > 0$	P_0 is a strict local minimum
$D > 0, A < 0$	P_0 is a strict local maximum
$D < 0$	P_0 is a saddle point
$D = 0$	the test is inconclusive

Proof skeleton. The quadratic part is

$$q(h, k) = Ah^2 + 2Bhk + Ck^2.$$

A symmetric 2×2 matrix is positive definite exactly when $A > 0$ and $D = AC - B^2 > 0$. It is negative definite exactly when $A < 0$ and $D > 0$. If $D < 0$, the quadratic form takes both positive and negative values. The Taylor expansion shows that the sign of the quadratic term controls the sign of $f(P_0 + h) - f(P_0)$ for sufficiently small h , except in the inconclusive case $D = 0$. $\square$

Example 101.4.2

For

$$f(x, y) = x^2 + 3y^2,$$

the only critical point is $(0, 0)$, and

$$H = \begin{pmatrix} 2 & 0 \\ 0 & 6 \end{pmatrix}.$$

Here $D = 12 > 0$ and $A = 2 > 0$, so $(0, 0)$ is a strict local minimum.

Example 101.4.3

For

$$f(x, y) = x^2 - y^2,$$

the origin is critical, but

$$H = \begin{pmatrix} 2 & 0 \\ 0 & -2 \end{pmatrix}, \quad D = -4 < 0.$$

The origin is a saddle point. Indeed, along the x -axis the function is positive, while along the y -axis it is negative.

Example 101.4.4 (Inconclusive case)

For $f(x, y) = x^4 + y^4$, the Hessian at the origin is the zero matrix, so $D = 0$. The test is inconclusive, but the origin is still a strict local minimum. For $g(x, y) = x^3 - y^3$, the Hessian at the origin is also zero, but the origin is not an extremum. This is why $D = 0$ really means "use another method."


§101.5 A picture of the Hessian test

Figure 101.1: Near a nondegenerate critical point, the Hessian determines the local quadratic shape.

§101.6 Constrained extrema

Suppose one wants to optimize $f(x, y)$ subject to a constraint

$$g(x, y) = 0.$$

At a regular point of the constraint curve, $\nabla g \neq 0$, so ∇g is normal to the constraint. At a constrained extremum, moving along the constraint produces no first-order change in f . Thus ∇f must also be normal to the constraint.

Theorem 101.6.1 (Lagrange multiplier condition)

If P_0 is a constrained local extremum of f on $g = 0$, and $\nabla g(P_0) \neq 0$, then there exists λ such that

$$\nabla f(P_0) = \lambda \nabla g(P_0).$$

For several constraints $g_1 = \cdots = g_k = 0$, the gradient of f lies in the span of the constraint gradients:

$$\nabla f = \lambda_1 \nabla g_1 + \cdots + \lambda_k \nabla g_k.$$

Example 101.6.2

Find extrema of $f(x, y) = x + y$ on the circle $x^2 + y^2 = 1$. The equations are

$$(1, 1) = \lambda(2x, 2y), \quad x^2 + y^2 = 1.$$

Thus $x = y$, and $2x^2 = 1$, so

$$(x, y) = \left(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}} \right) \quad \text{or} \quad (x, y) = \left(-\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}} \right).$$

The corresponding values are $\sqrt{2}$ and $-\sqrt{2}$, the maximum and minimum.

§101.7 Endpoint and boundary warnings

The equation $\nabla f = 0$ applies only to interior unconstrained extrema. The equation $\nabla f = \lambda \nabla g$ applies only at regular constrained points. If the domain has corners, endpoints, singular constraint points, or boundary pieces, they must be checked separately.

A reliable workflow is:

1. find interior critical points;
2. find constrained or boundary candidates;
3. check singular points and endpoints;
4. compare values when the problem asks for absolute extrema.

X

Vector Calculus

Part X: Contents

N-20230123	655
102.1 Indefinite and definite integrals	655
102.2 Area by slicing	655
102.3 Work as an integral	655
102.4 Line integrals	656
102.5 Green-type intuition	656
102.6 Polar coordinates	656
102.7 Spherical coordinates	657
102.8 Dot product and cross product	657
102.9 After the examples: what object is being integrated?	657
102.10 Line integrals as pullbacks	658
102.11 Exact forms and potentials	658
102.12 Conservation laws	658
102.13 De Rham preview	659
102.14 Work integrals and gauge freedom	659
102.15 Polar and spherical adapted charts	659
 N-20230201	 661
103.1 First orientation	661
103.2 Statement of Newton's shell theorem	661
103.3 Why symmetry alone is not enough	661
103.4 Solid-angle proof in full detail	662
103.5 Potential proof: doing the integral	663
103.6 The direct force integral	664
103.7 Gauss-law viewpoint	665
103.8 Harmonic potential	665
103.9 Why the inverse-square law is special	666
103.10 Application: a uniform solid sphere	666
103.11 Multipole viewpoint	666
103.12 A broader reading	667
 N-20230311	 669
104.1 Why two theories of integration exist	669
104.2 Riemann sums	669
104.3 Darboux upper and lower sums	669
104.4 A simple integrable discontinuity	670
104.5 Dirichlet's function	670
104.6 The Lebesgue idea	670
104.7 Characteristic functions and sets	671
104.8 Convergence theorems	671
104.9 Improper Riemann integrals versus Lebesgue integrals	672
104.10 Comparison table	672
104.11 Source repair: the measure-space starting point	672
104.12 Source repair: simple functions as building blocks	672
104.13 Source repair: probability viewpoint	673
104.14 A note on scope	673
104.15 Riemann integration partitions the domain	673
104.16 Lebesgue integration partitions the values	673
104.17 Dirichlet function as separator	673
104.18 Convergence theorems as exchange principles	673
104.19 Integration theory through measurable structure	675
104.20 Riemann sums and measure spaces	675
104.21 Dominated and monotone convergence in practice	676
104.22 L_p spaces and duality	676
104.23 Returning to the original Riemann examples	676

N-20240427	677
105.1 A theorem looking for its natural language	677
105.2 One-forms: measuring motion	677
105.3 The pullback experiment	677
105.4 Two-forms: measuring oriented area	678
105.5 Exterior derivative as one operator	678
105.6 A familiar theorem reappears	678
105.7 Orientation is not decoration	679
105.8 Why arc length is a different kind of geometry	679
105.9 Closed forms can detect holes	679
105.10 A table of classical shadows	680
105.11 Flux as a form	680
105.12 The moral of pullback	680
105.13 Exit line	680
N-20260403	681
106.1 From area sums to double integrals	681
106.2 Geometric meanings	681
106.3 Basic properties	682
106.4 Existence of the integral	682
106.5 Symmetry before computation	683
106.6 Mean value theorem for double integrals	683
106.7 A worked average-value example	684
106.8 Relation with product measure	684
106.9 A real Fubini warning: conditional two-dimensional accumulation	684
106.10 Product measure behind rectangular sums	685
106.11 Symmetry as projection onto invariant functions	685
106.12 Average value as a pushforward distribution	686
106.13 How the double-integral setup is checked	686
N-20260405	687
107.1 The real task: describe the region	687
107.2 Changing the order of integration	687
107.3 Polar coordinates	688
107.4 Jacobian as local area scaling	688
107.5 Why the absolute value appears	689
107.6 A linear change example	689
107.7 Choosing coordinates	689
107.8 Slicing as pushforward of area	689
107.9 Exterior algebra proof of the Jacobian factor	690
107.10 Polar coordinates as a singular chart	690
107.11 Coarea formula in the simplest radial case	690
107.12 When a clever substitution is actually a quotient	691
107.13 How a change of variables earns its bounds	691
N-20260406	693
108.1 Triple integrals as weighted volume	693
108.2 Rectangular-coordinate example	693
108.3 Cylindrical coordinates	693
108.4 Spherical coordinates	694
108.5 Coordinate-choice table	694
108.6 Mass, center of mass, and moments	694
108.7 Symmetry	695
108.8 Three ways to understand volume factors	695
108.9 Toward divergence theorem	695
108.10 Metric tensor derivation of spherical volume	696
108.11 Divergence as infinitesimal volume expansion	696
108.12 Pushforward of densities and probability	696
108.13 Volume forms and orientation	697
108.14 How the solid determines the coordinates	697

N-20260408	699
109.1 Volume as an integral	699
109.2 Describing spatial regions	699
109.3 A reconstruction workflow	700
109.4 Changing order of integration	700
109.5 Cross-sections	701
109.6 Top minus bottom	701
109.7 Cylindrical and spherical descriptions	701
109.8 Absolute extrema over solids	702
109.9 Fibre integration: top-minus-bottom as a pushforward	702
109.10 Cavalieri principle and layer-cake formula	702
109.11 Why regions must be split: projections can change topology	703
109.12 Order of integration as choosing a foliation	703
109.13 Reconstructing the solid before integrating	703
N-20260416	705
110.1 Two meanings of integration on curves	705
110.2 Line integrals with respect to arc length	705
110.3 Line integrals of vector fields	705
110.4 Surface integrals of scalar fields	706
110.5 Flux integrals	706
110.6 First kind versus second kind	706
110.7 Parametrization invariance	707
110.8 Differential forms viewpoint	707
110.9 Orientability warning	707
110.10 Densities versus differential forms	707
110.11 Flux and the Hodge star	707
110.12 Boundary orientation is not a sign convention	708
110.13 Non-orientability as a real obstruction	708
110.14 What the parametrization is responsible for	708
N-20260417	709
111.1 Green's theorem	709
111.2 Proof idea: rectangles and cancellation	709
111.3 Conservative fields and path independence	710
111.4 The punctured-plane warning	711
111.5 Area by Green's theorem	711
111.6 Differential forms notation	711
111.7 Closed forms, exact forms, and topology	712
111.8 Connection with de Rham cohomology	712
111.9 Winding number hidden in the punctured-plane example	712
111.10 From Green's theorem to residues	713
111.11 Cohomology as bookkeeping for failed potentials	713
111.12 Boundary cancellation and the algebra $\partial^2 = 0$	713
111.13 What Green's theorem checks in practice	714
N-20260615	715
112.1 Why this note is split this way	715
112.2 Four similar-looking integrals	715
112.3 First-kind surface integrals	715
112.4 Template surfaces	716
112.5 Example pattern: spherical cap	717
112.6 Second-kind surface integrals as flux	718
112.7 Model flux computations: closed, capped, singular, and sheeted	718
112.8 Gauss theorem	720
112.9 Singular fields and missing points	720
112.10 Stokes theorem	720
112.11 Stokes model computations: area vector and spanning-surface freedom	721
112.12 Decision table	722
112.13 Conceptual endpoint	722
112.14 The de Rham complex behind the vector formulas	722

112.15 General Stokes theorem	722
112.16 Singular sources and distributions	723
112.17 Conservation laws	723
112.18 Currents	724
112.19 Gauss and Stokes as dual local-to-global laws	724
112.20 Weak forms and PDE	724

Definite Integrals, Line Integrals, Work, Green-Type Thinking, and Vector Operations

N-20230123

§102.1 Indefinite and definite integrals

The note begins by distinguishing indefinite and definite integrals. An indefinite integral is an antiderivative family:

$$\int f(x) dx = F(x) + C, \quad F'(x) = f(x).$$

A definite integral is a number:

$$\int_a^b f(x) dx.$$

Geometrically, it is signed area under the graph. Analytically, it is the limit of Riemann sums.

The fundamental theorem of calculus connects the two:

$$\int_a^b f(x) dx = F(b) - F(a)$$

when $F' = f$.

§102.2 Area by slicing

The page includes diagrams of shaded regions. The standard method is to slice a region into thin strips. If vertical strips are convenient, then

$$A = \int_a^b (y_{top}(x) - y_{bottom}(x)) dx.$$

If horizontal strips are more convenient, then

$$A = \int_c^d (x_{right}(y) - x_{left}(y)) dy.$$

The note's sketches show that choosing the direction of slicing is part of the solution, not a fixed rule.

§102.3 Work as an integral

The physical meaning of a definite integral appears in work. If a force $F(x)$ acts along a line, then the work from $x = a$ to $x = b$ is

$$W = \int_a^b F(x) dx.$$

If the force is constant, this reduces to $W = F(b - a)$. If the force changes with position, integration adds the small contributions

$$dW = F(x) dx.$$

This is the same accumulation idea as area, but with a physical interpretation.

§102.4 Line integrals

The note then introduces line integrals. If a curve is parametrized by

$$r(t) = (x(t), y(t)), \quad a \leq t \leq b,$$

then the scalar line element is

$$ds = \sqrt{(x'(t))^2 + (y'(t))^2} dt.$$

A scalar line integral is

$$\int_C f ds = \int_a^b f(x(t), y(t)) \sqrt{(x'(t))^2 + (y'(t))^2} dt.$$

For a vector field $F = (P, Q)$, the work integral is

$$\int_C F \cdot dr = \int_C P dx + Q dy = \int_a^b [P(x(t), y(t))x'(t) + Q(x(t), y(t))y'(t)] dt.$$

§102.5 Green-type intuition

The handwritten diagrams show a small oriented boundary and a vector field. The idea is that circulation around a boundary can be related to a double integral over the region. In the plane, Green's theorem says

$$\oint_{\partial D} P dx + Q dy = \iint_D \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} \right) dA.$$

Even if the note does not prove the theorem fully, it records the key intuition: local rotation accumulated over an area produces circulation on the boundary.

§102.6 Polar coordinates

The note uses polar coordinates in examples. The change of variables is

$$x = r \cos \theta, \quad y = r \sin \theta.$$

The area element becomes

$$dA = r dr d\theta.$$

The factor r is the Jacobian. Geometrically, a small polar rectangle has side lengths approximately dr and $r d\theta$.

Thus a disk of radius R has area

$$\int_0^{2\pi} \int_0^R r dr d\theta = \pi R^2.$$

§102.7 Spherical coordinates

The later page contains a spherical-coordinate diagram. A common convention is

$$x = \rho \sin \varphi \cos \theta, \quad y = \rho \sin \varphi \sin \theta, \quad z = \rho \cos \varphi.$$

The volume element is

$$dV = \rho^2 \sin \varphi \, d\rho \, d\varphi \, d\theta.$$

The factor $\rho^2 \sin \varphi$ comes from the three small lengths:

$$d\rho, \quad \rho \, d\varphi, \quad \rho \sin \varphi \, d\theta.$$

This explains the diagram on the final page: the small spherical patch has dimensions determined by radius and angle.

§102.8 Dot product and cross product

The note also summarizes vector operations. The dot product is

$$a \cdot b = |a||b| \cos \theta.$$

It measures projection and appears in work:

$$dW = F \cdot dr.$$

The cross product has magnitude

$$|a \times b| = |a||b| \sin \theta$$

and direction given by orientation. It measures the area of the parallelogram spanned by a and b .

In coordinates,

$$a \times b = \begin{vmatrix} i & j & k \\ a_1 & a_2 & a_3 \\ b_1 & b_2 & b_3 \end{vmatrix}.$$

§102.9 After the examples: what object is being integrated?

The examples develop integration as accumulation over different geometric objects. The clean way to organize them is not by dimension alone, but by the type of object being integrated.

Integral	Object being integrated	Geometric meaning
$\int_a^b f(x) \, dx$	density on an oriented interval	signed accumulation
$\int_C f \, ds$	scalar density along a curve	mass/length-weighted accumulation
$\int_C P \, dx + Q \, dy$	one-form	work or circulation
$\iint_D f \, dA$	area density	accumulation over a region
$\iiint_\Omega f \, dV$	volume density	weighted volume or mass

This table prevents a common confusion. A line integral of $f ds$ and a work integral $P dx + Q dy$ both run along a curve, but they integrate different geometric objects. The first uses arclength and ignores orientation; the second is a pullback of a one-form and changes sign when the orientation is reversed.

Green's theorem then becomes a boundary-to-interior conversion:

$$\oint_{\partial D} P dx + Q dy = \iint_D (\partial_x Q - \partial_y P) dA.$$

It says that circulation around the boundary is curl accumulated over the interior.

§102.10 Line integrals as pullbacks

For a parametrized curve $\gamma : [a, b] \rightarrow \mathbb{R}^n$, a line integral of a one-form

$$\omega = P_1 dx_1 + \cdots + P_n dx_n$$

is

$$\int_{\gamma} \omega = \int_a^b \sum_i P_i(\gamma(t)) \gamma'_i(t) dt.$$

This is the pullback of the one-form to the interval:

$$\gamma^* \omega = \sum_i P_i(\gamma(t)) \gamma'_i(t) dt.$$

Thus the elementary formula $P dx + Q dy$ is already differential-form notation.

§102.11 Exact forms and potentials

A one-form ω is exact if $\omega = df$ for some potential f . Then

$$\int_{\gamma} df = f(\gamma(b)) - f(\gamma(a)).$$

This is the fundamental theorem of line integrals.

A form is closed if $d\omega = 0$. In the plane, for $\omega = P dx + Q dy$, this means

$$Q_x - P_y = 0.$$

On simply connected domains, closed one-forms are exact. On domains with holes, this may fail.

§102.12 Conservation laws

In mechanics, work is a line integral

$$W = \int_C F \cdot dr.$$

If $F = -\nabla V$ is conservative, then

$$W = V(A) - V(B)$$

depends only on endpoints. Conservation of energy is therefore a statement about exact one-forms.

This connects elementary work integrals to topology: path independence is not only about derivatives; it also depends on whether the domain has holes.

§102.13 De Rham preview

De Rham cohomology measures the failure of closed forms to be exact:

$$H_{dR}^1(M) = \frac{\{\text{closed 1-forms}\}}{\{\text{exact 1-forms}\}}.$$

For the punctured plane, the form

$$\omega = \frac{-y dx + x dy}{x^2 + y^2}$$

is closed but not exact. Its integral around the unit circle is 2π .

The elementary line integral around a circle is therefore a measurement of a topological hole.

§102.14 Work integrals and gauge freedom

If a force field is conservative, $F = -\nabla V$, then work depends only on endpoints. The potential V is not unique: adding a constant changes nothing. This is the simplest form of gauge freedom. The physically meaningful object is the differential dV , not the absolute zero of potential.

In electromagnetism, potentials have richer gauge freedom. A vector potential A may be changed by a gradient without changing the magnetic field $B = \nabla \times A$. The elementary fact that potentials are determined only up to constants is the first example of this broader principle.

§102.15 Polar and spherical adapted charts

The original integral examples using polar or spherical coordinates should be read as choosing charts adapted to symmetry. A chart is a coordinate map from a parameter domain to the geometric region. The integration element is obtained by pulling back the metric or volume form.

Thus polar coordinates are not merely a trick for circular regions. They are a coordinate chart on the punctured plane adapted to the action of rotations. Spherical coordinates are adapted to radial symmetry in three dimensions.

Newton's Shell Theorem: Solid Angle, Potential, and Gauss Law

N-20230201

§103.1 First orientation

The elementary statement is familiar: the gravitational field produced by a uniform spherical shell is zero at every point inside the shell. But this result is not a trivial symmetry statement. If the field point is not the center, the near side of the shell is closer and should pull more strongly, while the far side has more area visible in the same direction. The theorem says that these two effects cancel exactly.

This note rewrites the old solution into three coherent languages:

solid-angle cancellation,
scalar potential computation,
Gauss law and harmonic potential.

The point is not only to remember that the answer is zero. The point is to understand why the inverse-square law, spherical geometry, and potential theory fit together so perfectly.

§103.2 Statement of Newton's shell theorem

Let a spherical shell of radius R have uniform surface density σ . Its total mass is

$$M = 4\pi R^2 \sigma.$$

Let $\mathbf{x} \in \mathbb{R}^3$, and write $r = |\mathbf{x}|$, with the shell centered at the origin.

Theorem 103.2.1 (Newton's shell theorem)

The gravitational field of the shell is

$$\mathbf{g}(\mathbf{x}) = \begin{cases} \mathbf{0}, & 0 \leq r < R, \\ -\frac{GM}{r^2} \frac{\mathbf{x}}{r}, & r > R. \end{cases}$$

Equivalently, outside the shell the field is the same as if all mass were concentrated at the center, while inside the shell the field is zero.

The old note focused only on the interior part. It is better to prove both halves together, because the same potential calculation gives both at once.

§103.3 Why symmetry alone is not enough

At the center of the shell, cancellation is immediate: every mass element has an opposite mass element at the same distance. Away from the center, this argument fails. A point P inside the shell is closer to one side than the other.

The correct local cancellation uses three facts at once:

the same narrow cone cuts two opposite surface patches,
patch area grows like distance squared,
gravity decays like distance squared.

There is also a small geometric correction: the patch may be tilted relative to the cone. This is the obliquity factor. Omitting it makes the proof feel like a slogan rather than a proof.

§103.4 Solid-angle proof in full detail

Let P be an interior point. Draw a line through P , meeting the shell at two points A and B . Around this line take a very narrow cone with vertex P and solid angle $d\Omega$. It cuts two small patches of the shell near A and B .




Figure 103.1: The geometric cancellation proof: the same narrow cone from an interior point cuts two opposite shell patches, and the equal angle geometry supplies the common obliquity factor.

Let

$$s_A = PA, \quad s_B = PB.$$

Let β_A and β_B be the angles between the ray and the outward normal to the shell, taken with absolute cosine. For a small surface patch, solid angle and area satisfy

$$d\Omega = \frac{|\cos \beta|}{s^2} dS,$$

so

$$dS = \frac{s^2}{|\cos \beta|} d\Omega.$$

Therefore

$$dS_A = \frac{s_A^2}{|\cos \beta_A|} d\Omega, \quad dS_B = \frac{s_B^2}{|\cos \beta_B|} d\Omega.$$

The missing geometric point is that the two obliquity factors are equal. Since $OA = OB = R$, the triangle OAB is isosceles. Thus the base angles at A and B are equal. Because P lies on the chord AB , the angles between the chord direction and the radii at A and B have the same absolute cosine:

$$|\cos \beta_A| = |\cos \beta_B|.$$

The mass of a patch is $dm = \sigma dS$. Its gravitational field magnitude at P is

$$|d\mathbf{g}| = G \frac{dm}{s^2} = G\sigma \frac{dS}{s^2}.$$

Substituting the area formula gives

$$|d\mathbf{g}_A| = G\sigma \frac{d\Omega}{|\cos \beta_A|}, \quad |d\mathbf{g}_B| = G\sigma \frac{d\Omega}{|\cos \beta_B|}.$$

These magnitudes are equal, and the directions are opposite along the same line through P . Hence the two contributions cancel.

Pairing all such cones through P gives

$$\mathbf{g}(P) = \mathbf{0}.$$

The proof is now precise: the inverse-square law cancels the square growth of apparent area, and spherical geometry cancels the obliquity factor.

§103.5 Potential proof: doing the integral

The force integral is a vector integral. The potential is scalar, so it is cleaner. Let the field point have distance r from the center, and put it on the positive z -axis. A point of the shell at polar angle θ has distance

$$\ell(\theta) = \sqrt{R^2 + r^2 - 2Rr \cos \theta}$$

from the field point.




Figure 103.2: Geometry of the shell integral: the distance ℓ from the field point to a surface element and the projection angle used when passing from scalar potential to radial force.

The gravitational potential is

$$V(r) = -G \int \frac{dm}{\ell}.$$

For the shell,

$$dm = \sigma R^2 \sin \theta d\theta d\phi.$$

Hence

$$V(r) = -G\sigma R^2 \int_0^{2\pi} \int_0^\pi \frac{\sin \theta d\theta d\phi}{\sqrt{R^2 + r^2 - 2Rr \cos \theta}}.$$

Integrating over ϕ and putting $u = \cos \theta$, we obtain

$$V(r) = -2\pi G\sigma R^2 \int_{-1}^1 \frac{du}{\sqrt{R^2 + r^2 - 2Rru}}.$$

Now compute the elementary integral:

$$\begin{aligned} \int_{-1}^1 \frac{du}{\sqrt{R^2 + r^2 - 2Rru}} &= \left[-\frac{1}{Rr} \sqrt{R^2 + r^2 - 2Rru} \right]_{-1}^1 \\ &= \frac{(R+r) - |R-r|}{Rr}. \end{aligned}$$

Therefore

$$V(r) = -2\pi G\sigma R^2 \cdot \frac{(R+r) - |R-r|}{Rr}.$$

Since $M = 4\pi R^2\sigma$, this becomes

$$V(r) = \begin{cases} -\frac{GM}{R}, & 0 \leq r < R, \\ -\frac{GM}{r}, & r > R. \end{cases}$$

Now recover the field by

$$\mathbf{g} = -\nabla V.$$

For a radial potential, this means

$$\mathbf{g}(r) = -\frac{dV}{dr} \hat{\mathbf{r}}.$$

Inside the shell, $V(r)$ is constant, so

$$\mathbf{g} = \mathbf{0}.$$

Outside the shell,

$$V(r) = -\frac{GM}{r},$$

so

$$\mathbf{g}(r) = -\frac{GM}{r^2} \hat{\mathbf{r}}.$$

This proves both halves of Newton's shell theorem.

§103.6 The direct force integral

The direct force calculation is possible, but it is less elegant. By symmetry only the radial component survives. The radial component from a ring at polar angle θ is proportional to the projection of the force onto the radial axis.

One obtains

$$g_r(r) = 2\pi G\sigma R^2 \int_0^\pi \frac{(R \cos \theta - r) \sin \theta}{(R^2 + r^2 - 2Rr \cos \theta)^{3/2}} d\theta.$$

The sign convention here takes the positive direction away from the center. This integral is exactly

$$-\frac{dV}{dr}.$$

So the potential method does not avoid the physics; it packages the same computation in a scalar form and differentiates only at the end.

This explains why potential theory is the natural language for shell problems. Instead of separately tracking force directions, one integrates a scalar quantity whose gradient gives the field.

§103.7 Gauss-law viewpoint

There is a still shorter proof once the field is known to be spherically symmetric. The gravitational Gauss law is

$$\nabla \cdot \mathbf{g} = -4\pi G\rho,$$

or, in integral form,

$$\oint_S \mathbf{g} \cdot d\mathbf{S} = -4\pi G M_{\text{enc}}.$$

For a spherical shell, symmetry implies that the field at radius r is radial and has constant magnitude on each sphere:

$$\mathbf{g}(r) = g(r)\hat{\mathbf{r}}.$$

Taking S to be the sphere of radius r , the flux is

$$\oint_S \mathbf{g} \cdot d\mathbf{S} = 4\pi r^2 g(r).$$

Thus

$$4\pi r^2 g(r) = -4\pi G M_{\text{enc}}(r),$$

so

$$g(r) = -\frac{G M_{\text{enc}}(r)}{r^2}.$$

Inside the shell, $M_{\text{enc}}(r) = 0$, hence $g(r) = 0$. Outside the shell, $M_{\text{enc}}(r) = M$, hence

$$g(r) = -\frac{GM}{r^2}.$$

The nuance is important: Gauss law alone gives the result quickly only because spherical symmetry has already reduced the field to one radial function. The uniform spherical shell provides that symmetry.

§103.8 Harmonic potential

In an empty region, the gravitational potential satisfies Laplace's equation:

$$\Delta V = 0.$$

Inside the shell there is no mass. Moreover the shell is spherically symmetric, so the interior potential must be radial:

$$V = V(r).$$

For a radial function in $\mathbb{R}^3$, Laplace's equation becomes

$$\frac{1}{r^2} \frac{d}{dr} \left(r^2 \frac{dV}{dr} \right) = 0.$$

Integrating once gives

$$r^2 V'(r) = C.$$

If $C \neq 0$, then $V'(r) = C/r^2$, which is singular at $r = 0$. The potential inside the shell should be regular at the center, so $C = 0$. Therefore

$$V'(r) = 0,$$

and the potential is constant inside.

This is the PDE version of the theorem. The absence of mass gives harmonicity; spherical symmetry reduces the harmonic equation to an ODE; regularity at the center forces the constant solution.

§103.9 Why the inverse-square law is special

The solid-angle proof shows exactly why the exponent 2 matters. For a patch cut out by a fixed solid angle,

$$dS \sim s^2 d\Omega$$

up to the common obliquity factor. Gravity contributes a factor $1/s^2$. These cancel.

If the force law were $1/s^p$, the patch contribution would scale like

$$\frac{dS}{s^p} \sim s^{2-p} d\Omega.$$

The near and far patches would have equal magnitudes for all distances only when

$$p = 2.$$

Thus the shell theorem is not a generic fact about attractive forces. It is a special feature of inverse-square fields.

The other hypotheses also matter. If the shell is not uniform, the two paired patches do not have masses proportional only to area. If the surface is not spherical, the obliquity factors do not match in the same way. The theorem is the result of three matching structures:

$$\text{uniform density} + \text{extspherical geometry} + \text{inverse-square law}.$$

§103.10 Application: a uniform solid sphere

A uniform solid sphere can be regarded as a stack of thin spherical shells. At a point of radius r , all shells outside radius r contribute zero field. Only the mass enclosed inside radius r matters.

If the solid sphere has volume density ρ , then

$$M(r) = \frac{4\pi}{3} \rho r^3.$$

Thus the interior field is

$$g(r) = -\frac{GM(r)}{r^2} = -\frac{4\pi G\rho}{3} r.$$

So the field inside a uniform ball grows linearly with distance from the center.

This has a famous dynamical consequence. If one imagines a straight tunnel through a uniform spherical planet and ignores friction and rotation, a particle moving along the tunnel satisfies

$$\ddot{r} = -\frac{4\pi G\rho}{3} r.$$

This is simple harmonic motion. The shell theorem turns a gravitational problem into the equation of a spring.

§103.11 Multipole viewpoint

There is also a useful advanced memory aid. The Newtonian kernel has the expansion

$$\frac{1}{|x - y|} = \sum_{\ell=0}^{\infty} \frac{r_{<}^{\ell}}{r_{>^{\ell+1}}} P_{\ell}(\cos \gamma),$$

where

$$r_{<} = \min(|x|, |y|), \quad r_{>} = \max(|x|, |y|).$$

For a uniformly distributed spherical shell, integrating over the angles kills all higher spherical harmonics. Only the monopole term survives outside. That is why the exterior field sees only total mass.

Inside the shell, the same expansion leaves a constant potential. A constant potential has zero gradient, hence zero field. This is the harmonic-analysis version of the same theorem.

§103.12 A broader reading

The shell theorem should not be remembered as “obvious by symmetry.” For an off-center interior point, ordinary pairwise symmetry is not enough. The real mechanism is:

spherical geometry controls apparent area and obliquity,
the inverse-square law cancels the distance-squared factor,
potential theory turns cancellation into harmonicity,
and Gauss law turns it into flux conservation.

The geometric proof explains the local cancellation. The potential proof computes the whole field. Gauss law explains the flux structure. Laplace's equation explains why the interior potential must be constant.

Together these four views make Newton's shell theorem more than a trick: it is a meeting point of geometry, analysis, and field theory.

Riemann Integration and Lebesgue Integration: From Partitions of the Domain to Partitions of Values

N-20230311

§104.1 Why two theories of integration exist

This note compares Riemann integration and Lebesgue integration. A compressed summary such as “Riemann cuts the x -axis; Lebesgue cuts the y -axis” is memorable, but it is not enough. The real issue is what kinds of limiting processes the integral is supposed to tolerate.

Riemann integration works beautifully for continuous functions and many functions with mild discontinuities. But once functions oscillate too much, especially when rational and irrational points behave differently, the Riemann approach becomes rigid. Lebesgue integration was designed to handle limits of functions more naturally, especially in analysis, probability, and Fourier theory.

§104.2 Riemann sums

Let $f : [a, b] \rightarrow \mathbb{R}$ be bounded. A partition is a finite sequence

$$P : a = x_0 < x_1 < \cdots < x_n = b.$$

Choose a sample point $\xi_i \in [x_{i-1}, x_i]$ in each subinterval. The Riemann sum is

$$\sum_{i=1}^n f(\xi_i) \Delta x_i, \quad \Delta x_i = x_i - x_{i-1}.$$

If all such sums approach the same limit as the mesh

$$\|P\| = \max_i \Delta x_i$$

tends to zero, that limit is the Riemann integral.

The geometric picture is rectangles under a graph. But the definition is not simply about drawing rectangles; it says that the total contribution should stabilize no matter how the sample points are chosen, once the partition is sufficiently fine.

§104.3 Darboux upper and lower sums

A cleaner equivalent formulation is Darboux integration. On each interval $[x_{i-1}, x_i]$, define

$$M_i = \sup_{x \in [x_{i-1}, x_i]} f(x), \quad m_i = \inf_{x \in [x_{i-1}, x_i]} f(x).$$

The upper and lower sums are

$$U(P, f) = \sum_i M_i \Delta x_i, \quad L(P, f) = \sum_i m_i \Delta x_i.$$

Then

$$f \text{ is Riemann integrable} \iff \inf_P U(P, f) = \sup_P L(P, f).$$

The common value is the integral.

This formulation explains why discontinuities matter. If a function oscillates wildly on every interval, then $M_i - m_i$ remains large no matter how small the interval is, and the upper and lower sums cannot meet.

§104.4 A simple integrable discontinuity

Consider

$$f(x) = \begin{cases} 1, & x = c, \\ 0, & x \neq c \end{cases}$$

on $[a, b]$. This function is discontinuous at one point, but it is Riemann integrable with integral zero. Indeed, the lower sums are all zero. For any partition, only one subinterval can contain c , and the contribution to the upper sum from that subinterval can be made arbitrarily small.

Thus Riemann integration tolerates finite discontinuities. More generally, a bounded function is Riemann integrable iff its set of discontinuities has Lebesgue measure zero. This theorem already hints that measure theory is hiding behind Riemann integration.

§104.5 Dirichlet's function

The Dirichlet function is

$$1_{\mathbb{Q}}(x) = \begin{cases} 1, & x \in \mathbb{Q}, \\ 0, & x \notin \mathbb{Q}. \end{cases}$$

On every nonempty interval there are rational and irrational numbers. Therefore on every subinterval,

$$M_i = 1, \quad m_i = 0.$$

Hence

$$U(P, f) = b - a, \quad L(P, f) = 0$$

for every partition. The upper and lower integrals never meet, so the function is not Riemann integrable.

Lebesgue integration handles this example immediately: since $\mathbb{Q}$ is countable and has measure zero,

$$\int_{[a,b]} 1_{\mathbb{Q}} dx = 0.$$

This one example explains much of the motivation for Lebesgue's theory.

§104.6 The Lebesgue idea

Riemann integration partitions the domain into small intervals and asks how the function behaves on each interval. Lebesgue integration partitions according to values. For a nonnegative simple function

$$s = \sum_{j=1}^m a_j 1_{E_j},$$

where the E_j are measurable and disjoint, define

$$\int s \, d\mu = \sum_{j=1}^m a_j \mu(E_j).$$

For a nonnegative measurable function f , define

$$\int f \, d\mu = \sup \left\{ \int s \, d\mu : s \leq f, \, s \text{ simple} \right\}.$$

For a general integrable function, write

$$f = f^+ - f^-.$$

If both $\int f^+$ and $\int f^-$ are finite, define

$$\int f = \int f^+ - \int f^-.$$

The construction is less pictorial than Riemann sums, but more robust. It counts how much space is occupied by each range of values.

§104.7 Characteristic functions and sets

The Lebesgue integral turns set measure into function integration:

$$\int 1_E \, d\mu = \mu(E).$$

This is conceptually important. Sets become functions, and measuring a set becomes integrating its indicator. Probability theory uses exactly this viewpoint: expectation is integration, and probabilities are integrals of indicators.

For the Dirichlet function,

$$1_{\mathbb{Q} \cap [a,b]}$$

has integral zero because the rationals are countable. For the Thomae function, which is discontinuous at every rational and continuous at every irrational, the function is still Riemann integrable because its discontinuity set is countable. Lebesgue theory makes such examples easier to organize.

§104.8 Convergence theorems

The major advantage of Lebesgue integration is not merely that it integrates more functions. It has powerful limit theorems.

The Monotone Convergence Theorem says: if

$$0 \leq f_1 \leq f_2 \leq \cdots, \quad f_n \rightarrow f,$$

then

$$\int f \, d\mu = \lim_{n \rightarrow \infty} \int f_n \, d\mu.$$

The Dominated Convergence Theorem says: if $f_n \rightarrow f$ pointwise, and $|f_n| \leq g$ for an integrable g , then

$$\lim_{n \rightarrow \infty} \int f_n \, d\mu = \int f \, d\mu.$$

These are theorems Riemann integration cannot match in such generality. They explain why Lebesgue integration becomes the standard language of modern analysis.

§104.9 Improper Riemann integrals versus Lebesgue integrals

A subtle point: some improper Riemann integrals converge conditionally but are not Lebesgue integrable. For example,

$$\int_1^\infty \frac{\sin x}{x} dx$$

converges as an improper integral, but

$$\int_1^\infty \left| \frac{\sin x}{x} \right| dx$$

diverges. Lebesgue integrability of a signed function requires integrability of the absolute value. Thus Lebesgue integration is not simply “more permissive”; it is more structurally consistent with normed function spaces such as L^1 .

§104.10 Comparison table

Aspect	Riemann	Lebesgue
Partition	Domain intervals	Value levels and measurable sets
Basic object	Bounded functions on intervals	Measurable functions
Bad example	$1_{\mathbb{Q}}$ not integrable	Integral is 0
Limit theorems	Limited	MCT, Fatou, DCT
Natural space	Continuous functions and elementary calculus	L^p spaces, probability, Fourier analysis

§104.11 Source repair: the measure-space starting point

The uploaded note explicitly starts the Lebesgue construction from a measure space

$$(X, \mathcal{M}, \mu),$$

where $\mathcal{M}$ is a σ -algebra and μ is a measure. A function $f : X \rightarrow \mathbb{R}$ is measurable if sets such as

$$\{x \in X : f(x) > \alpha\}$$

belong to $\mathcal{M}$ for every real α . This is the point at which the theory leaves elementary calculus: before integrating, one must decide which subsets are measurable.

§104.12 Source repair: simple functions as building blocks

The uploaded text defines a nonnegative simple function by

$$\varphi(x) = \sum_{i=1}^n a_i \chi_{A_i}(x),$$

and sets

$$\int_X \varphi d\mu = \sum_{i=1}^n a_i \mu(A_i).$$

For a nonnegative measurable function,

$$\int_X f d\mu = \sup \left\{ \int_X \varphi d\mu : 0 \leq \varphi \leq f, \varphi \text{ simple} \right\}.$$

This is the formal version of cutting by values. Instead of approximating the domain by small intervals, we approximate the function from below by simple measurable functions.

§104.13 Source repair: probability viewpoint

The source also points toward probability. A random variable is a measurable function, and expectation is an integral:

$$\mathbb{E}[X] = \int_{\Omega} X \, d\mathbb{P}.$$

Events are measured by probabilities, and indicators satisfy

$$\mathbb{E}[1_A] = \mathbb{P}(A).$$

This makes the Lebesgue framework the natural language of modern probability, not only a generalization of calculus.

§104.14 A note on scope

The slogan is useful:

Riemann cuts the domain; Lebesgue cuts by values.

But the deeper difference is categorical: Riemann integration is built around approximating graphs by rectangles, while Lebesgue integration is built around measurable sets and stable limiting operations. That is why Lebesgue integration becomes the foundation for modern analysis, probability, PDE, and functional analysis.

§104.15 Riemann integration partitions the domain

Riemann integration controls oscillation on subintervals of the domain. A bounded function is Riemann integrable when upper and lower sums can be made arbitrarily close.

§104.16 Lebesgue integration partitions the values

Lebesgue integration instead measures level sets. Simple functions approximate by values, and measurable sets replace intervals as the basic pieces. This is why characteristic functions become fundamental.

§104.17 Dirichlet function as separator

The function $1_{\mathbb{Q}}$ is discontinuous everywhere and not Riemann integrable, but it is Lebesgue integrable with integral zero because $\mathbb{Q}$ has measure zero.

§104.18 Convergence theorems as exchange principles

The central question behind the convergence theorems is not only whether a limit exists. It is whether the following exchange is legal:

$$\int \lim_{n \rightarrow \infty} f_n \, d\mu \stackrel{?}{=} \lim_{n \rightarrow \infty} \int f_n \, d\mu.$$

Lebesgue theory gives precise hypotheses under which this exchange is safe.

Theorem 104.18.1 (Monotone convergence, informal reading)

If $0 \leq f_1 \leq f_2 \leq \cdots$ and $f_n \rightarrow f$, then

$$\int f_n d\mu \longrightarrow \int f d\mu.$$

The intuition is that positive mass is only accumulating upward; no cancellation or escaping negative part is hidden.

Theorem 104.18.2 (Dominated convergence, informal reading)

If $f_n \rightarrow f$ pointwise and there is an integrable function g such that

$$|f_n| \leq g$$

for all n , then

$$\int f_n d\mu \longrightarrow \int f d\mu.$$

The dominating function prevents mass from escaping to infinity or concentrating in narrower and narrower spikes.

Fatou's lemma is the one-sided safety rule:

$$\int \liminf_{n \rightarrow \infty} f_n d\mu \leq \liminf_{n \rightarrow \infty} \int f_n d\mu \quad (f_n \geq 0).$$

It says that, for nonnegative functions, the integral of the limiting lower mass cannot exceed the limiting lower integrals.

Example 104.18.3 (Escaping mass)

On $\mathbb{R}$, let

$$f_n = \mathbf{1}_{[n, n+1]}.$$

Then $f_n(x) \rightarrow 0$ for every fixed x , but

$$\int_{\mathbb{R}} f_n dx = 1$$

for all n . Therefore

$$\int \lim f_n dx = 0 \neq 1 = \lim \int f_n dx.$$

The mass has not disappeared; it has escaped to infinity. Dominated convergence fails because no single integrable g can dominate all these translated bumps.

Example 104.18.4 (Concentrating mass)

On $[0, 1]$, let

$$f_n = n \mathbf{1}_{[0, 1/n]}.$$

Then $f_n(x) \rightarrow 0$ for every $x > 0$, but

$$\int_0^1 f_n(x) dx = 1.$$

The mass concentrates into a narrower spike near 0. Again, pointwise convergence alone is too weak to justify exchanging limit and integral.

Remark 104.18.5 (Probability language) — In probability, the same question is whether

$$\mathbb{E}[\lim X_n] = \lim \mathbb{E}[X_n]$$

is legal. Monotone convergence, Fatou's lemma, and dominated convergence are therefore not abstract measure-theory decorations; they are rules for passing limits through expectations.

§104.19 Integration theory through measurable structure

Riemann integration manages oscillation on intervals; Lebesgue integration measures sets of values. The shift from partitions of the domain to measurable level sets is the conceptual upgrade.

§104.20 Riemann sums and measure spaces

The Riemann integral begins with partitions of intervals and rectangles. Lebesgue integration changes the organizing principle: instead of cutting the domain into small intervals first, it studies the size of sets on which the function takes certain values. A measure space is a triple

$$(X, \mathcal{A}, \mu),$$

where $\mathcal{A}$ is a sigma-algebra of measurable sets and μ assigns sizes to those sets.

A simple function has the form

$$s = \sum_{j=1}^m c_j \mathbf{1}_{E_j}.$$

Its integral is defined by

$$\int s d\mu = \sum_{j=1}^m c_j \mu(E_j).$$

Nonnegative measurable functions are integrated by approximating from below by simple functions, and general integrable functions are handled by positive and negative parts.

This preserves the elementary intuition of area, but it makes convergence theorems much stronger.

§104.21 Dominated and monotone convergence in practice

The three convergence tools should be read as three different controls on a limiting process.

tool	hypothesis	what it prevents
MCT	$0 \leq f_n \uparrow f$	loss by cancellation
Fatou	$f_n \geq 0$	overestimating the limiting lower mass
DCT	$ f_n \leq g \in L^1$	escape or concentration of mass

This is why Lebesgue integration is the natural language of probability, Fourier analysis, and PDEs. These subjects constantly form limiting objects: limits of random variables, limits of Fourier partial sums, weak or strong limits of approximate solutions. The integral must be stable enough to survive those passages to the limit.

§104.22 L_p spaces and duality

For $1 \leq p < \infty$, define

$$L^p(X) = \{f : \int |f|^p d\mu < \infty\} / \sim,$$

where functions equal almost everywhere are identified. The norm is

$$\|f\|_p = \left(\int |f|^p d\mu \right)^{1/p}.$$

Hölder's inequality says that if $1/p + 1/q = 1$, then

$$\int |fg| d\mu \leq \|f\|_p \|g\|_q.$$

Thus integration becomes a pairing between function spaces. The elementary integral $\int fg$ is now a duality bracket, and inequalities become statements about bounded linear functionals.

§104.23 Returning to the original Riemann examples

Darboux sums, Riemann sums, and step-function pictures remain valuable because they teach integral as accumulation. The advanced reconstruction does not erase them. It explains why later mathematics replaces intervals by measurable sets, pointwise equality by almost-everywhere equality, and ordinary convergence by modes of convergence suitable for passing limits through integrals.

Differential Forms and Stokes: The Geometry of Boundary Measurements

N-20240427

§105.1 A theorem looking for its natural language

The fundamental theorem of calculus says

$$\int_a^b f'(x) dx = f(b) - f(a).$$

Green's theorem, the divergence theorem, and Stokes' theorem look more complicated, but they share the same shape: integrate a derivative over a region, and the answer lives on the boundary.

Differential forms are the language in which this pattern becomes one theorem:

$$\int_M d\omega = \int_{\partial M} \omega.$$

The purpose of this chapter is not to glorify notation. It is to make clear why dx , dy , pullback, orientation, and boundary signs belong together.

§105.2 One-forms: measuring motion

A tangent vector tells us a direction of motion. A covector measures such a direction. At a point p , a covector is a linear map

$$\alpha : T_p M \rightarrow \mathbb{R}.$$

A 1-form is a smoothly varying covector.

In the plane, a typical 1-form is

$$\omega = P(x, y) dx + Q(x, y) dy.$$

If

$$v = a \frac{\partial}{\partial x} + b \frac{\partial}{\partial y},$$

then

$$\omega(v) = Pa + Qb.$$

Thus a 1-form is not a tiny arrow. It is a measuring rule for arrows.

§105.3 The pullback experiment

Let $\gamma(t) = (x(t), y(t))$ be a parametrized curve. Pulling back the form

$$\omega = P(x, y) dx + Q(x, y) dy$$

to the interval gives

$$\gamma^* \omega = (P(x(t), y(t))x'(t) + Q(x(t), y(t))y'(t)) dt.$$

Therefore

$$\int_{\gamma} \omega = \int_a^b (P(x(t), y(t))x'(t) + Q(x(t), y(t))y'(t)) dt.$$

This is the usual line integral. The conceptual advantage is that the curve integral is reduced to ordinary integration after pullback.

§105.4 Two-forms: measuring oriented area

The wedge product is antisymmetric:

$$dx \wedge dy = -dy \wedge dx, \quad dx \wedge dx = 0.$$

A 2-form such as

$$f(x, y) dx \wedge dy$$

measures signed area with density f . The sign changes if the orientation is reversed.

This is why forms are naturally tied to oriented geometry. A form does not merely ask “how large?” It asks “how large, with which orientation?”

§105.5 Exterior derivative as one operator

For a function f ,

$$df = f_x dx + f_y dy.$$

For a 1-form $\omega = P dx + Q dy$,

$$d\omega = (Q_x - P_y) dx \wedge dy.$$

The rule

$$d^2 = 0$$

is the compact source of several familiar identities: gradients have zero curl, curls have zero divergence, and exact forms are closed.

§105.6 A familiar theorem reappears

Apply Stokes to a planar region D and a 1-form $\omega = P dx + Q dy$:

$$\int_D d\omega = \int_{\partial D} \omega.$$

Writing this in coordinates gives

$$\iint_D (Q_x - P_y) dx dy = \int_{\partial D} P dx + Q dy.$$

This is Green’s theorem. The point is not that Green’s theorem becomes shorter, but that its boundary nature becomes unavoidable.

§105.7 Orientation is not decoration




Figure 105.1: The orientation of a patch forces an induced orientation on its boundary.

The boundary of an oriented region inherits an orientation. If this convention is changed, Stokes' theorem changes sign. So orientation is not a drawing preference; it is part of the data needed for the theorem to be true.

§105.8 Why arc length is a different kind of geometry

It is tempting to treat ds as a 1-form, but arc length is not linear in the velocity:

$$ds(\gamma') = \|\gamma'\| dt.$$

A 1-form is linear in the tangent vector. A norm is not. This distinction separates two kinds of geometry:

differential forms	oriented measurement and integration
Riemannian metrics	lengths, angles, and distances

Both are important, but they should not be confused.

§105.9 Closed forms can detect holes

The distinction between closed and exact forms is one of the first places where analysis notices topology. On $\mathbb{R}^2 \setminus \{0\}$, consider

$$\omega = \frac{-y dx + x dy}{x^2 + y^2}.$$

One computes that

$$d\omega = 0.$$

So ω is closed. But it is not exact on $\mathbb{R}^2 \setminus \{0\}$. Indeed, along the unit circle $\gamma(t) = (\cos t, \sin t)$, $0 \leq t \leq 2\pi$, we get

$$\gamma^*\omega = dt,$$

and therefore

$$\int_{\gamma} \omega = 2\pi.$$

If $\omega = df$, the integral around a closed loop would be 0. The nonzero integral detects the missing origin.

This is a beautiful lesson: a differential form can carry topological information.

§105.10 A table of classical shadows

The same theorem appears in several costumes:

dimension	form	classical theorem
1	f	$\int_a^b df = f(b) - f(a)$
2	$P dx + Q dy$	Green's theorem
3	flux form	divergence theorem
3	circulation form	classical Stokes theorem

The table is not meant as a replacement for calculation. It is meant to prevent the common mistake of treating these theorems as unrelated facts.

§105.11 Flux as a form

In $\mathbb{R}^3$, a vector field

$$\mathbf{F} = (P, Q, R)$$

corresponds to the 2-form

$$\omega = P dy \wedge dz + Q dz \wedge dx + R dx \wedge dy.$$

Then

$$d\omega = (P_x + Q_y + R_z) dx \wedge dy \wedge dz.$$

Thus the divergence theorem is exactly Stokes' theorem for this 2-form:

$$\int_M d\omega = \int_{\partial M} \omega.$$

This is a good example of how vector calculus becomes cleaner when translated into forms.

§105.12 The moral of pullback

Integration should not depend on a lucky parametrization. If a surface is parametrized by

$$\Phi : U \rightarrow M,$$

then an integral over M is computed by pulling the form back to U :

$$\int_M \omega = \int_U \Phi^* \omega.$$

The pullback automatically inserts the Jacobian-like factors and orientation signs. This is why forms are the natural objects to integrate: they know how to move backward along maps.

§105.13 Exit line

Differential forms are a language for quantities that can be pulled back and integrated. Stokes' theorem says that the integral of a derivative over a body is a boundary measurement.

The short memory is:

forms measure; pullbacks transport;
 d differentiates;
 Stokes moves to the boundary.

Double Integrals: Riemann Sums, Area, Symmetry, and Average Value

N-20260403

§106.1 From area sums to double integrals

A double integral is the two-dimensional version of adding many small contributions. If a quantity has density $f(x, y)$ over a plane region D , then a small cell of area ΔA_i contributes approximately

$$f(\xi_i, \eta_i) \Delta A_i.$$

Adding all cells and refining the partition gives the integral.

Definition 106.1.1 (Riemann double integral). Let $D \subseteq \mathbb{R}^2$ be a bounded region and let $f : D \rightarrow \mathbb{R}$ be bounded. Partition D into small subregions ΔD_i , choose sample points $(\xi_i, \eta_i) \in \Delta D_i$, and form

$$S(P, f) = \sum_i f(\xi_i, \eta_i) \Delta A_i.$$

If $S(P, f)$ tends to a limit independent of the partition and sample points as the mesh tends to 0, then the limit is denoted

$$\iint_D f(x, y) dA.$$

The notation dA or $d\sigma$ represents an infinitesimal area element. It is not an independent variable; it tells us what geometric measure is being used.

§106.2 Geometric meanings

The meaning depends on the integrand.

integrand	meaning of $\iint_D f dA$
$f = 1$	area of D
$f \geq 0$	volume under $z = f(x, y)$
f sign-changing	signed volume
$f = \rho$	mass of a lamina with density ρ

Example 106.2.1 (Area as an integral)

For the disk $D = \{(x, y) : x^2 + y^2 \leq R^2\}$,

$$\text{area}(D) = \iint_D 1 \, dA.$$

Using polar coordinates later gives

$$\int_0^{2\pi} \int_0^R r \, dr \, d\theta = \pi R^2.$$

§106.3 Basic properties

The following properties are the main algebra of double integrals.

Proposition 106.3.1

Assume the displayed integrals exist.

(i) Linearity:

$$\iint_D (af + bg) \, dA = a \iint_D f \, dA + b \iint_D g \, dA.$$

(ii) Additivity over regions: if $D = D_1 \cup D_2$ and the overlap has area zero, then

$$\iint_D f \, dA = \iint_{D_1} f \, dA + \iint_{D_2} f \, dA.$$

(iii) Monotonicity: if $f \leq g$ on D , then

$$\iint_D f \, dA \leq \iint_D g \, dA.$$

(iv) Estimate: if $m \leq f \leq M$ on D , then

$$m \, \text{area}(D) \leq \iint_D f \, dA \leq M \, \text{area}(D).$$

Proof idea. Each statement is first true for finite Riemann sums. Passing to the limit preserves addition, scalar multiplication, and inequalities. The estimate comes from multiplying $m \leq f(\xi_i, \eta_i) \leq M$ by ΔA_i and summing. $\square$

§106.4 Existence of the integral

For most calculus computations, the following sufficient condition is enough.

Theorem 106.4.1

If f is continuous on a bounded closed elementary region D , then f is Riemann integrable on D .

The theorem is often used silently. It is the reason that ordinary continuous examples can be computed by choosing convenient coordinates and orders of integration.

Remark 106.4.2 — The statement is not saying that only continuous functions are integrable. Piecewise continuous functions and many functions with mild discontinuities are also integrable. Continuity is simply the safest elementary hypothesis.

§106.5 Symmetry before computation

Symmetry is one of the fastest ways to evaluate an integral.

Proposition 106.5.1 (Odd cancellation)

Let D be symmetric under a measure-preserving transformation $T : D \rightarrow D$. If

$$f(T(x, y)) = -f(x, y),$$

then

$$\iint_D f \, dA = 0.$$

If $f \circ T = f$, then the integral over D can often be reduced to a fundamental half or sector of D .

Example 106.5.2

Let D be the disk centered at the origin. Since D is symmetric under $(x, y) \mapsto (-x, y)$,

$$\iint_D x e^{y^2} \, dA = 0.$$

The integrand is odd in x , while the region is unchanged by reflecting across the y -axis.

§106.6 Mean value theorem for double integrals

Theorem 106.6.1 (Mean value theorem)

Let D be compact and connected, and let $f : D \rightarrow \mathbb{R}$ be continuous. Then there exists a point $P_0 \in D$ such that

$$\iint_D f \, dA = f(P_0) \operatorname{area}(D).$$

Equivalently, the average value

$$f_{\text{avg}} = \frac{1}{\operatorname{area}(D)} \iint_D f \, dA$$

is attained somewhere on D .

Proof skeleton. By compactness, f attains a minimum m and maximum M on D . The

estimate property gives

$$m \leq f_{\text{avg}} \leq M.$$

Since D is connected and f is continuous, the image $f(D)$ is an interval. Therefore the intermediate value theorem gives a point P_0 where $f(P_0) = f_{\text{avg}}$. $\square$

§106.7 A worked average-value example

Example 106.7.1

Let $D = [0, 1] \times [0, 2]$ and $f(x, y) = x + y$. Then

$$\iint_D (x + y) dA = \int_0^1 \int_0^2 (x + y) dy dx = \int_0^1 (2x + 2) dx = 3.$$

Since $\text{area}(D) = 2$, the average value is $3/2$. The mean-value theorem says there is a point in the rectangle with $x + y = 3/2$, for example $(1/2, 1)$.

§106.8 Relation with product measure

The elementary Riemann picture divides a region into small rectangles. In measure-theoretic language, this is integration with respect to two-dimensional Lebesgue measure:

$$\iint_D f(x, y) dA = \int_{\mathbb{R}^2} f(x, y) \mathbf{1}_D(x, y) d\lambda^2.$$

This viewpoint separates two choices:

- (i) the function being accumulated;
- (ii) the measure that tells how large each piece of the region is.

For nonnegative measurable functions, Tonelli's theorem justifies iterated integration even when the value may be infinite. For sign-changing functions, Fubini's theorem requires integrability of $|f|$. The elementary rule "switch the order of integration" is therefore backed by real hypotheses.

§106.9 A real Fubini warning: conditional two-dimensional accumulation

It is tempting to treat order-switching as a harmless algebraic move. The advanced point is sharper: for sign-changing functions, order-switching is licensed by absolute integrability. Without that hypothesis, different limiting procedures can see different cancellations.

A standard improper example on the square near the origin is

$$f(x, y) = \frac{x^2 - y^2}{(x^2 + y^2)^2}, \quad (x, y) \in (0, 1]^2.$$

For fixed $x > 0$,

$$\int_0^1 \frac{x^2 - y^2}{(x^2 + y^2)^2} dy = \left[\frac{y}{x^2 + y^2} \right]_0^1 = \frac{1}{1 + x^2}.$$

Hence

$$\int_0^1 \int_0^1 f(x, y) dy dx = \frac{\pi}{4}.$$

But for fixed $y > 0$,

$$\int_0^1 \frac{x^2 - y^2}{(x^2 + y^2)^2} dx = \left[-\frac{x}{x^2 + y^2} \right]_0^1 = -\frac{1}{1 + y^2},$$

so

$$\int_0^1 \int_0^1 f(x, y) dx dy = -\frac{\pi}{4}.$$

The two answers differ because the singularity at $(0, 0)$ is not absolutely integrable. This is the concrete content of the distinction between Tonelli's theorem and Fubini's theorem.

§106.10 Product measure behind rectangular sums

The phrase "divide the region into rectangles" hides a construction. Rectangles $A \times B$ are assigned area by

$$\lambda^2(A \times B) = \lambda(A)\lambda(B).$$

This rule first defines area on finite unions of rectangles. One then extends it to the sigma-algebra generated by rectangles. The double integral is integration with respect to this product measure.

This explains why iterated integrals work first for indicator functions of rectangles:

$$\int \mathbf{1}_{A \times B}(x, y) d(\mu \times \nu) = \mu(A)\nu(B) = \int_A \left(\int_B 1 d\nu \right) d\mu.$$

Linearity handles simple functions, monotone convergence handles nonnegative measurable functions, and absolute integrability handles signed functions. Thus the elementary Riemann-sum picture is not abandoned in measure theory; it is completed.

§106.11 Symmetry as projection onto invariant functions

If a finite group G acts on D by area-preserving transformations, averaging over the group defines the Reynolds operator

$$\mathcal{R}f = \frac{1}{|G|} \sum_{g \in G} f \circ g.$$

Because the action preserves area,

$$\iint_D f dA = \iint_D \mathcal{R}f dA.$$

So integration only sees the invariant part of the function. Odd functions cancel because their invariant part is zero. This is the same idea that later appears in representation theory as projection onto the trivial representation, but here it is already visible in a disk or rectangle symmetry problem.

For example, on the disk, the rotation group kills all angular Fourier modes except the constant mode:

$$f(r, \theta) = a_0(r) + \sum_{n \geq 1} (a_n(r) \cos n\theta + b_n(r) \sin n\theta)$$

gives

$$\iint_D f \, dA = \int_0^R \int_0^{2\pi} a_0(r) r \, d\theta \, dr.$$

The elementary symmetry shortcut is the first case of harmonic decomposition.

§106.12 Average value as a pushforward distribution

For continuous $f : D \rightarrow \mathbb{R}$, the average value can be viewed probabilistically. Put the uniform probability measure on D :

$$\mathbb{P}(E) = \frac{\text{area}(E)}{\text{area}(D)}.$$

Then

$$f_{\text{avg}} = \mathbb{E}[f].$$

The pushforward measure $f_*\mathbb{P}$ on $\mathbb{R}$ records how much area of D lies at each range of values of f . The mean value theorem says that if D is connected and f is continuous, this expectation lies inside the interval $f(D)$, hence equals $f(P_0)$ for some point.

This interpretation is useful later: integration can be studied either by integrating over the original region or by pushing measure forward through a function and integrating over its values.

§106.13 How the double-integral setup is checked

A double integral begins as a limit of weighted area sums, so the first data to identify are the region D , the area element, and the function being averaged over that region. Linearity and additivity come directly from cutting the region into pieces; this is why splitting a domain along a change in boundary formula is legitimate rather than cosmetic.

The examples in this note give three practical tests. If the region has a symmetry and the integrand is odd under that symmetry, the corresponding contribution vanishes because the area measure is preserved. If the question asks for an average value, rewrite

$$\iint_D f \, dA = f_{\text{avg}} \text{area}(D)$$

and check connectedness and continuity before invoking the mean value theorem. If the computation is rearranged by Fubini or by a change of variables, the real issue is not the order of symbols but whether the slices describe the same region and whether the Jacobian accounts for area distortion. This is the point at which Tonelli, Fubini, and change of variables turn the geometric setup into a justified calculation.

Computing Double Integrals: Slicing, Polar Coordinates, and Change of Variables

N-20260405

§107.1 The real task: describe the region

In double-integral problems, the hard part is usually not the antiderivative. It is translating a plane region into inequalities. Once the region is described correctly, Fubini's theorem turns the double integral into ordinary one-variable integration.

Definition 107.1.1. A region D is **vertically simple** if it has the form

$$D = \{(x, y) : a \leq x \leq b, \varphi_1(x) \leq y \leq \varphi_2(x)\}.$$

It is **horizontally simple** if it has the form

$$D = \{(x, y) : c \leq y \leq d, \psi_1(y) \leq x \leq \psi_2(y)\}.$$

For a vertically simple region,

$$\iint_D f(x, y) dA = \int_a^b \int_{\varphi_1(x)}^{\varphi_2(x)} f(x, y) dy dx.$$

For a horizontally simple region,

$$\iint_D f(x, y) dA = \int_c^d \int_{\psi_1(y)}^{\psi_2(y)} f(x, y) dx dy.$$

§107.2 Changing the order of integration

Changing order means describing the same set with the other variable outermost.

Example 107.2.1

Consider

$$I = \int_0^1 \int_{x^2}^x f(x, y) dy dx.$$

The region is

$$0 \leq x \leq 1, \quad x^2 \leq y \leq x.$$

For a horizontal slice, $0 \leq y \leq 1$, and the inequalities become

$$y \leq x \leq \sqrt{y}.$$

Thus

$$I = \int_0^1 \int_y^{\sqrt{y}} f(x, y) dx dy.$$

A reliable order-changing workflow is:

1. write the original region as inequalities;
2. draw or infer its projection on the new outer axis;
3. solve the old inequalities for the new inner variable;
4. split the region if one pair of bounds is not enough.

§107.3 Polar coordinates

Polar coordinates are useful when circles, disks, sectors, or expressions $x^2 + y^2$ appear:

$$x = r \cos \theta, \quad y = r \sin \theta, \quad r \geq 0.$$

The area element is

$$dA = r \, dr \, d\theta.$$

The factor r is not a decoration; it is the Jacobian determinant of the map $(r, \theta) \mapsto (x, y)$.

Example 107.3.1

For the upper half disk $x^2 + y^2 \leq 4$, $y \geq 0$, the polar description is

$$0 \leq r \leq 2, \quad 0 \leq \theta \leq \pi.$$

Therefore

$$\iint_D (x^2 + y^2) \, dA = \int_0^\pi \int_0^2 r^2 \cdot r \, dr \, d\theta = \pi \cdot 4 = 4\pi.$$

§107.4 Jacobian as local area scaling

Let

$$T(u, v) = (x(u, v), y(u, v)).$$

The derivative matrix is

$$DT = \begin{pmatrix} x_u & x_v \\ y_u & y_v \end{pmatrix}.$$

A tiny rectangle in the (u, v) -plane is sent approximately to a parallelogram spanned by the two columns of DT . Its area is multiplied by

$$|\det DT| = \left| \frac{\partial(x, y)}{\partial(u, v)} \right|.$$

Theorem 107.4.1 (Change of variables)

Let $T : D' \rightarrow D$ be a one-to-one C^1 change of variables, except possibly on harmless boundary pieces, and suppose its Jacobian does not vanish in the interior. Then

$$\iint_D f(x, y) \, dx \, dy = \iint_{D'} f(x(u, v), y(u, v)) \left| \frac{\partial(x, y)}{\partial(u, v)} \right| \, du \, dv.$$

§107.5 Why the absolute value appears

For scalar area, orientation is irrelevant. A coordinate map may preserve or reverse orientation, but both cases scale ordinary area by a positive factor. Hence the absolute value.

For oriented forms, the sign is kept:

$$dx \wedge dy = \det \frac{\partial(x, y)}{\partial(u, v)} du \wedge dv.$$

This explains the difference between scalar area integrals and oriented circulation or flux formulas.

§107.6 A linear change example

Example 107.6.1

Let

$$x = u + v, \quad y = u - v.$$

Then

$$\frac{\partial(x, y)}{\partial(u, v)} = \det \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} = -2.$$

Thus $dA = 2 du dv$. If a diamond-shaped region in the (x, y) -plane is hard to describe but becomes a rectangle in (u, v) -coordinates, this coordinate change turns geometric complexity into a constant Jacobian.

§107.7 Choosing coordinates

The best coordinate system simplifies both the integrand and the boundary.

signal in problem	candidate coordinates	reason
$x^2 + y^2$	polar	radial expression
circle or sector	polar	simple bounds
parallelogram	linear change	constant Jacobian
curved family of level sets	adapted variables	region straightens

§107.8 Slicing as pushforward of area

Slicing a region vertically means integrating over fibers $x = \text{constant}$ and then integrating over x . In measure language, the map

$$\pi_x : D \rightarrow \mathbb{R}, \quad (x, y) \mapsto x$$

pushes area measure on D down to a measure on the x -axis. The inner integral computes the contribution of each fiber, and the outer integral adds the fibers.

This viewpoint separates three related operations:

operation	map used	geometric effect
Fubini slicing	$(x, y) \mapsto x$	integrate over fibers
change of variables	$(u, v) \mapsto (x, y)$	pull back area by a chart
coarea	$F(x, y) \mapsto t$	integrate over level sets

Ordinary order-changing is the classroom version of this general fiber-integration principle.

§107.9 Exterior algebra proof of the Jacobian factor

The Jacobian determinant is not merely a remembered formula. It is the coefficient by which a coordinate change pulls back the oriented area form.

For

$$x = x(u, v), \quad y = y(u, v),$$

we have

$$dx = x_u du + x_v dv, \quad dy = y_u du + y_v dv.$$

Taking wedge products,

$$\begin{aligned} dx \wedge dy &= (x_u du + x_v dv) \wedge (y_u du + y_v dv) \\ &= (x_u y_v - x_v y_u) du \wedge dv. \end{aligned}$$

Thus

$$dx \wedge dy = \det \frac{\partial(x, y)}{\partial(u, v)} du \wedge dv.$$

Scalar area forgets orientation, so it uses the absolute value. Oriented integration keeps the sign. This is why the same determinant appears both in change of variables and in Green/Stokes type formulas.

§107.10 Polar coordinates as a singular chart

The polar map

$$T(r, \theta) = (r \cos \theta, r \sin \theta)$$

has Jacobian r . It is not one-to-one globally: all values of θ represent the origin when $r = 0$, and θ is periodic. Calculus usually ignores this because the bad set has area zero, but geometrically it matters.

This is the first hint that coordinates are local. On a manifold, one generally cannot describe everything by one global coordinate chart. Spherical coordinates have the same issue at the poles; longitude is not defined there. Integration survives because measure-zero coordinate singularities do not affect scalar integrals, while topology remembers them.

§107.11 Coarea formula in the simplest radial case

Changing variables is not the only way to decompose an integral. The coarea formula decomposes a domain by level sets. For the distance function

$$r(x, y) = \sqrt{x^2 + y^2}, \quad |\nabla r| = 1 \quad ((x, y) \neq 0),$$

coarea says, informally,

$$\int_{\mathbb{R}^2} g(x, y) dA = \int_0^\infty \left(\int_{r^{-1}(t)} g dS \right) dt.$$

If $g(x, y) = h(r)$ is radial, the inner integral over the circle of radius t is $2\pi th(t)$, so

$$\int_{\mathbb{R}^2} h(r) dA = \int_0^\infty 2\pi th(t) dt.$$

This is the polar-coordinate formula reinterpreted as integration over level curves of r . The elementary formula $dA = r dr d\theta$ and the advanced coarea formula are describing the same geometry from two sides.

§107.12 When a clever substitution is actually a quotient

Sometimes a change of variables collapses a family of curves into a coordinate. For example, the variables

$$u = x + y, \quad v = x - y$$

are adapted to diagonal lines. A region bounded by lines $x + y = \text{constant}$ and $x - y = \text{constant}$ becomes a rectangle in (u, v) -space. This is not just a trick: the new variables are invariants of two transverse foliations of the plane.

In harder problems, one chooses variables because level sets of the integrand or boundary are natural. The computation succeeds when the coordinate fibers match the geometry already present in the problem.

§107.13 How a change of variables earns its bounds

Fubini should be read as a slicing statement: a two-dimensional region is decomposed into one-dimensional fibres over a projection. A change of variables is stronger: it replaces the region by a coordinate domain and pays for the replacement with a Jacobian determinant. In polar coordinates this payment is the familiar factor r , because a small rectangle in (r, θ) -space becomes an annular sector whose area is approximately $r dr d\theta$.

The examples show how to decide whether a substitution is natural. Linear maps are useful when the boundary lines are already expressed by linear combinations such as $x + y$ and $x - y$. Polar coordinates are useful when the boundary or integrand is organized by distance from a point. More generally, the variables should be chosen so that level curves of the new coordinates match the visible families of curves in the problem. After that choice, the bounds are obtained by drawing the image region in the new variables; they are not copied from the old coordinates by syntax.

Triple Integrals: Solids, Cylindrical Coordinates, and Spherical Coordinates

N-20260406

§108.1 Triple integrals as weighted volume

A triple integral accumulates a scalar field over a spatial region $\Omega \subseteq \mathbb{R}^3$:

$$\iiint_{\Omega} f(x, y, z) dV.$$

If $f = 1$, the integral is volume. If $f = \rho$ is a mass density, the integral is mass. If f is source density, the integral is total source.

Definition 108.1.1. An xy -simple solid has the form

$$\Omega = \{(x, y, z) : (x, y) \in D, z_1(x, y) \leq z \leq z_2(x, y)\}.$$

Then

$$\iiint_{\Omega} f dV = \iint_D \int_{z_1(x,y)}^{z_2(x,y)} f(x, y, z) dz dA.$$

There are analogous yz -simple and zx -simple descriptions. If no single description works, split the solid.

§108.2 Rectangular-coordinate example

Example 108.2.1

Let Ω be the tetrahedron

$$x \geq 0, \quad y \geq 0, \quad z \geq 0, \quad x + y + z \leq 1.$$

One description is

$$0 \leq x \leq 1, \quad 0 \leq y \leq 1 - x, \quad 0 \leq z \leq 1 - x - y.$$

Thus

$$\text{vol}(\Omega) = \int_0^1 \int_0^{1-x} \int_0^{1-x-y} 1 dz dy dx = \frac{1}{6}.$$

§108.3 Cylindrical coordinates

Cylindrical coordinates keep the vertical coordinate and use polar coordinates in the horizontal plane:

$$x = r \cos \theta, \quad y = r \sin \theta, \quad z = z.$$

The volume element is

$$dV = r \, dr \, d\theta \, dz.$$

They are natural for cylinders, cones around the z -axis, and solids where $x^2 + y^2$ appears.

Example 108.3.1

For the cylinder $x^2 + y^2 \leq R^2$, $0 \leq z \leq h$,

$$V = \int_0^{2\pi} \int_0^R \int_0^h r \, dz \, dr \, d\theta = \pi R^2 h.$$

§108.4 Spherical coordinates

Spherical coordinates are

$$x = \rho \sin \varphi \cos \theta, \quad y = \rho \sin \varphi \sin \theta, \quad z = \rho \cos \varphi,$$

where usually $\rho \geq 0$, $0 \leq \varphi \leq \pi$, and $0 \leq \theta < 2\pi$. The volume element is

$$dV = \rho^2 \sin \varphi \, d\rho \, d\varphi \, d\theta.$$

The factor has a geometric interpretation: small coordinate lengths are approximately

$$d\rho, \quad \rho \, d\varphi, \quad \rho \sin \varphi \, d\theta.$$

Multiplying them gives $\rho^2 \sin \varphi$.

Example 108.4.1

The ball $x^2 + y^2 + z^2 \leq R^2$ has volume

$$\int_0^{2\pi} \int_0^\pi \int_0^R \rho^2 \sin \varphi \, d\rho \, d\varphi \, d\theta = \frac{4}{3} \pi R^3.$$

§108.5 Coordinate-choice table

geometry	coordinates	typical bounds
box or graph over a plane	rectangular	$z_1(x, y) \leq z \leq z_2(x, y)$
cylinder, cone around z -axis	cylindrical	r, θ, z
ball, spherical shell, cone from origin	spherical	ρ, φ, θ

Coordinates should be chosen because they simplify the solid, not because they are fashionable.

§108.6 Mass, center of mass, and moments

If $\rho(x, y, z)$ is density, then

$$M = \iiint_{\Omega} \rho \, dV.$$

The center of mass is

$$\bar{x} = \frac{1}{M} \iiint_{\Omega} x \rho \, dV, \quad \bar{y} = \frac{1}{M} \iiint_{\Omega} y \rho \, dV, \quad \bar{z} = \frac{1}{M} \iiint_{\Omega} z \rho \, dV.$$

Moments of inertia are obtained by integrating squared distance to an axis. For example, about the z -axis,

$$I_z = \iiint_{\Omega} (x^2 + y^2) \rho \, dV.$$

§108.7 Symmetry

Before computing a triple integral, check symmetry. If Ω is symmetric under $x \mapsto -x$ and f is odd in x , then

$$\iiint_{\Omega} f(x, y, z) \, dV = 0.$$

If f is even, integrate over half the solid and double. In three dimensions, symmetry can remove a whole family of terms.

Example 108.7.1

On the ball centered at the origin,

$$\iiint_{x^2+y^2+z^2 \leq 1} xz^2 \, dV = 0,$$

because the region is symmetric under $x \mapsto -x$, while the integrand changes sign.

§108.8 Three ways to understand volume factors

The Jacobian factors in cylindrical and spherical coordinates can be read in three compatible ways.

Viewpoint	What it explains
linear approximation	DT sends a small box to a parallelepiped
volume form	$dx \wedge dy \wedge dz$ pulls back by $\det DT$
metric tensor	volume density is $\sqrt{\det g}$ in curvilinear coordinates

For scalar volume, orientation is forgotten and one uses $|\det DT|$. For oriented integration, the sign is kept. Cylindrical and spherical volume elements are special cases of these principles.

§108.9 Toward divergence theorem

Triple integrals of divergence appear in Gauss' theorem:

$$\iiint_{\Omega} \nabla \cdot F \, dV = \iint_{\partial\Omega} F \cdot n \, dS.$$

Before learning the theorem, $\iiint_{\Omega} \nabla \cdot F \, dV$ can be read as total source in the volume. After the theorem, the same total source is measured by boundary flux.

§108.10 Metric tensor derivation of spherical volume

The spherical factor $\rho^2 \sin \varphi$ can be derived from the metric. In spherical coordinates,

$$ds^2 = d\rho^2 + \rho^2 d\varphi^2 + \rho^2 \sin^2 \varphi d\theta^2.$$

Thus the metric matrix is diagonal:

$$g = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \rho^2 & 0 \\ 0 & 0 & \rho^2 \sin^2 \varphi \end{pmatrix}.$$

The Riemannian volume density is $\sqrt{\det g}$, hence

$$\sqrt{\det g} = \rho^2 \sin \varphi.$$

So

$$dV = \rho^2 \sin \varphi d\rho d\varphi d\theta.$$

This is the same factor as the Jacobian determinant, but the metric viewpoint generalizes directly to curved surfaces and Riemannian manifolds.

§108.11 Divergence as infinitesimal volume expansion

Divergence is not just a formal sum of partial derivatives. If a vector field F generates a flow Φ_t , then

$$\mathcal{L}_F(dV) = (\nabla \cdot F)dV,$$

where $\mathcal{L}_F$ is the Lie derivative. In words: divergence is the infinitesimal rate at which the flow expands volume.

For $F(x, y, z) = (x, y, z)$, the flow is $\Phi_t(p) = e^t p$. A region is scaled by e^t in each of three directions, so its volume is scaled by e^{3t} . The infinitesimal expansion rate is 3, and indeed

$$\nabla \cdot F = 1 + 1 + 1 = 3.$$

This makes the divergence theorem conceptually natural: total infinitesimal volume production inside Ω should equal flux through $\partial\Omega$.

§108.12 Pushforward of densities and probability

The triple-integral change-of-variables formula is also the density transformation formula in probability. If X has density f_X and $Y = T(X)$, then for an invertible smooth T ,

$$f_Y(y) = f_X(T^{-1}y) \left| \det DT^{-1}(y) \right|.$$

For example, in $\mathbb{R}^3$, a radially symmetric density can be pushed forward by the radius map $R = |X|$. The shell of radius r has area $4\pi r^2$, so the radial density contains the same factor that appears in spherical coordinates:

$$f_R(r) = 4\pi r^2 f_X(r) \quad \text{for a radial density } f_X.$$

Thus the spherical Jacobian is not just a calculus artifact; it is the volume growth of shells.

§108.13 Volume forms and orientation

The standard oriented volume form is

$$dx \wedge dy \wedge dz.$$

Under a coordinate map $T(u, v, w)$, it pulls back as

$$T^*(dx \wedge dy \wedge dz) = \det DT \, du \wedge dv \wedge dw.$$

Scalar volume uses $|\det DT|$, but oriented integration keeps $\det DT$. This distinction becomes essential in Stokes' theorem, where signs encode boundary orientation.

Triple integrals are therefore the scalar shadow of a more structured object: integration of top-degree forms.

§108.14 How the solid determines the coordinates

A triple integral is weighted volume, but the coordinate system decides how that volume is sliced. Rectangular coordinates are best when the solid is bounded by graphs over coordinate planes. Cylindrical coordinates are best when the projection to the xy -plane is circular or when the integrand depends on $x^2 + y^2$. Spherical coordinates are best when the boundary is organized by distance from the origin or by cones through the origin.

The volume factors are not separate formulas to memorize. They are Jacobians:

$$dV = dx \, dy \, dz, \quad dV = r \, dr \, d\theta \, dz, \quad dV = \rho^2 \sin \varphi \, d\rho \, d\varphi \, d\theta.$$

The metric-tensor and volume-form derivations explain why these factors appear and why orientation matters once one passes from scalar volume to Stokes-type formulas. In a computation the safe order is therefore: describe the solid geometrically, choose coordinates whose coordinate surfaces match that description, compute or recall the Jacobian, and only then write the iterated bounds.

Volumes, Spatial Regions, and Order of Integration

N-20260408

§109.1 Volume as an integral

The volume of a solid is the integral of the constant density 1. If a solid lies between two graphs

$$z = z_1(x, y), \quad z = z_2(x, y)$$

over a plane region D , then

$$V = \iint_D [z_2(x, y) - z_1(x, y)] dA.$$

If the solid is better described directly in three dimensions, then

$$V = \iiint_{\Omega} 1 dV.$$

Example 109.1.1

The solid under $z = 4 - x^2 - y^2$ and above the disk $x^2 + y^2 \leq 4$ has volume

$$V = \iint_{x^2+y^2 \leq 4} (4 - x^2 - y^2) dA.$$

In polar coordinates this becomes

$$V = \int_0^{2\pi} \int_0^2 (4 - r^2)r dr d\theta = 8\pi.$$

§109.2 Describing spatial regions

A spatial region can be simple with respect to different coordinate planes. An xy -simple description has the form

$$\Omega = \{(x, y, z) : (x, y) \in D, z_1(x, y) \leq z \leq z_2(x, y)\}.$$

Then

$$\iiint_{\Omega} f dV = \iint_D \int_{z_1(x, y)}^{z_2(x, y)} f(x, y, z) dz dA.$$

The same solid may be yz -simple or zx -simple instead. The best description is the one in which the inner bounds are simplest.




Figure 109.1: A composite solid with a hemispherical part and a conical part. The picture is used to decide cross-sections and height bounds before writing the integral.

§109.3 A reconstruction workflow

When an integral or picture is given, reconstruct the solid before integrating.

- (i) Identify the bounding surfaces.
- (ii) Find the projection or shadow on a coordinate plane.
- (iii) Decide which variable has clear lower and upper surfaces.
- (iv) Split the shadow when the top, bottom, left, or right boundary changes formula.
- (v) Only then write the integral.

Most wrong triple integrals have correct antiderivatives but wrong regions.

§109.4 Changing order of integration

Changing the order of integration means choosing a new projection and new fibers. For example, an integral written as

$$\int_a^b \int_{\varphi_1(x)}^{\varphi_2(x)} \int_{z_1(x,y)}^{z_2(x,y)} f \, dz \, dy \, dx$$

uses vertical segments above an xy -region. Reordering may require projecting onto the xz - or yz -plane instead.

Caution

Do not change the order by simply permuting differentials. The bounds describe a region. A new order requires a new description of the same region.




Figure 109.2: A rotational solid with circular cross-sections at different heights. Such diagrams are read by projecting to the base plane and then recovering the vertical bounds.

§109.5 Cross-sections

Sometimes the cleanest volume method is to integrate cross-sectional area:

$$V = \int_a^b A(t) dt.$$

Here $A(t)$ is the area of the slice of the solid by a plane such as $z = t$. This is the one-variable version of slicing.

Example 109.5.1

For a right circular cone of height h and base radius R , the slice at height z has radius $R(1 - z/h)$, so

$$A(z) = \pi R^2 \left(1 - \frac{z}{h}\right)^2.$$

Thus

$$V = \int_0^h \pi R^2 \left(1 - \frac{z}{h}\right)^2 dz = \frac{1}{3} \pi R^2 h.$$

§109.6 Top minus bottom

For a solid between surfaces, the double-integral formula

$$V = \iint_D (z_{\text{top}} - z_{\text{bottom}}) dA$$

is often faster than a full triple integral. But this formula is only legal after the top and bottom are correctly identified over the same base region D .

§109.7 Cylindrical and spherical descriptions

If the solid has rotational symmetry, rectangular projections may be awkward. Cylindrical coordinates describe a solid by inequalities in r, θ, z , and spherical coordinates describe it by inequalities in ρ, φ, θ . The cost is the Jacobian:

$$dV = r dr d\theta dz, \quad dV = \rho^2 \sin \varphi d\rho d\varphi d\theta.$$

The coordinate system should match the surface equations. A sphere centered at the origin asks for spherical coordinates; a cylinder around the z -axis asks for cylindrical coordinates.

§109.8 Absolute extrema over solids

Volume setup also appears in average-value problems. If f is continuous on a solid Ω , its average value is

$$f_{\text{avg}} = \frac{1}{\text{vol}(\Omega)} \iiint_{\Omega} f \, dV.$$

When Ω is compact and connected, this average value is attained somewhere in Ω . This is the three-dimensional analogue of the mean-value theorem for integrals.

§109.9 Fibre integration: top-minus-bottom as a pushforward

The formula

$$V = \iint_D (z_{\text{top}} - z_{\text{bottom}}) \, dA$$

has a precise structural meaning. Let

$$\pi : \Omega \rightarrow D, \quad \pi(x, y, z) = (x, y)$$

be projection onto the base. The fibre over (x, y) is the vertical segment

$$\pi^{-1}(x, y) = \{z : z_{\text{bottom}}(x, y) \leq z \leq z_{\text{top}}(x, y)\}.$$

Its length is

$$\ell(x, y) = z_{\text{top}}(x, y) - z_{\text{bottom}}(x, y).$$

The volume is the integral of fibre length over the base:

$$\text{vol}(\Omega) = \int_D \ell(x, y) \, dA.$$

This is the elementary version of integration along fibres. In differential geometry and topology, fibre integration pushes a form on a total space down to a form on the base.

§109.10 Cavalieri principle and layer-cake formula

Cross-sectional integration

$$V = \int_a^b A(t) \, dt$$

is Cavalieri's principle in analytic form. A more advanced cousin is the layer-cake formula. For a nonnegative function f ,

$$\int_D f \, dA = \int_0^\infty \text{area}(\{(x, y) \in D : f(x, y) > t\}) \, dt.$$

Geometrically, the volume under $z = f(x, y)$ can be computed by stacking horizontal slices instead of vertical columns.

For example, for $f(x, y) = 1 - x^2 - y^2$ on the unit disk, the level set $f > t$ is the disk $x^2 + y^2 < 1 - t$, so

$$\int_D f \, dA = \int_0^1 \pi(1 - t) \, dt = \frac{\pi}{2}.$$

This equals the usual polar-coordinate computation. The point is not the answer but the shift: one can integrate by columns, by horizontal layers, or by fibres of a map.

§109.11 Why regions must be split: projections can change topology

A projection can make a simple solid look complicated. Suppose a vertical line over a base point intersects the solid in two disjoint intervals rather than one. Then there is no single pair $z_1(x, y) \leq z \leq z_2(x, y)$ describing the fibre. The correct response is to split the base into pieces on which the fibre structure is stable.

This is a low-dimensional version of a general principle: maps are easiest to integrate over where their fibres vary regularly. Critical values and singular fibres force decompositions. In advanced language, one tries to partition the base so that the projection behaves like a fibration on each piece.

§109.12 Order of integration as choosing a foliation

Writing

$$\iiint_{\Omega} f \, dz \, dy \, dx$$

means that Ω is being decomposed into one-dimensional fibres parallel to the z -axis, then those fibres are organized over a region in the xy -plane. A different order chooses a different family of fibres. Thus changing order is choosing a different foliation of the same solid.

This language sounds advanced, but it is exactly what the pictures are doing. A good diagram shows which family of slices has the simplest geometry.

§109.13 Reconstructing the solid before integrating

For the problems in this note, the integral sign is less important than the solid it describes. An order such as

$$dz \, dy \, dx$$

means that vertical line fibres are being placed over a region in the xy -plane. Changing the order means choosing a different family of fibres, so the projection may change and the projected region may need to be split.

A reliable setup has four checks. First draw or describe the solid independently of the given bounds. Second choose the projection plane whose shadow has the simplest boundary. Third write the inner bounds as the entrance and exit points of a fibre through the solid. Fourth check whether the shadow changes formula across the projection; if it does, split the integral. The layer-cake and Cavalieri interpretations are useful because they keep the same discipline: volume is assembled from slices, and each slice must correspond to an actual cross-section of the solid.

Line Integrals and Surface Integrals of the First and Second Kind

N-20260416

§110.1 Two meanings of integration on curves

There are two common line integrals, and they answer different questions.

A first-kind line integral integrates a scalar density along length:

$$\int_C f \, ds.$$

A second-kind line integral integrates a vector field along an oriented curve:

$$\int_C P \, dx + Q \, dy + R \, dz = \int_C F \cdot dr.$$

The first is about amount distributed along a curve. The second is about work or circulation.

§110.2 Line integrals with respect to arc length

Let C be parametrized by $r(t) = (x(t), y(t), z(t))$, $a \leq t \leq b$. Then

$$ds = \|r'(t)\| \, dt,$$

and

$$\int_C f \, ds = \int_a^b f(r(t)) \|r'(t)\| \, dt.$$

This integral does not depend on orientation.

Example 110.2.1

For the circle $r(t) = (R \cos t, R \sin t)$, $0 \leq t \leq 2\pi$, and $f(x, y) = x^2 + y^2$,

$$\int_C f \, ds = \int_0^{2\pi} R^2 \cdot R \, dt = 2\pi R^3.$$

§110.3 Line integrals of vector fields

For $F = (P, Q, R)$,

$$\int_C F \cdot dr = \int_a^b F(r(t)) \cdot r'(t) \, dt.$$

In coordinates,

$$\int_C P \, dx + Q \, dy + R \, dz = \int_a^b (Px'(t) + Qy'(t) + Rz'(t)) \, dt.$$

Reversing orientation changes the sign.

Example 110.3.1

Let $F = (-y, x)$ and let C be the unit circle traversed counterclockwise, $r(t) = (\cos t, \sin t)$. Then

$$F(r(t)) = (-\sin t, \cos t), \quad r'(t) = (-\sin t, \cos t),$$

so

$$\int_C F \cdot dr = \int_0^{2\pi} 1 \, dt = 2\pi.$$

This is circulation around the origin.

§110.4 Surface integrals of scalar fields

Let a surface S be parametrized by $r(u, v)$, $(u, v) \in D$. The surface area element is

$$dS = \|r_u \times r_v\| \, du \, dv.$$

The first-kind surface integral is

$$\iint_S f \, dS = \iint_D f(r(u, v)) \|r_u \times r_v\| \, du \, dv.$$

It integrates scalar density over area and does not require an orientation.

Example 110.4.1

For the graph $z = g(x, y)$, use $r(x, y) = (x, y, g(x, y))$. Then

$$r_x \times r_y = (-g_x, -g_y, 1), \quad dS = \sqrt{1 + g_x^2 + g_y^2} \, dx \, dy.$$

§110.5 Flux integrals

A flux integral measures how much a vector field crosses an oriented surface. If S is parametrized by $r(u, v)$, then

$$\iint_S F \cdot n \, dS = \iint_D F(r(u, v)) \cdot (r_u \times r_v) \, du \, dv,$$

where the direction of $r_u \times r_v$ determines the orientation.

Reversing the normal reverses the sign. This is the surface analogue of reversing a line integral of a vector field.

§110.6 First kind versus second kind

object	first kind	second kind
curve	$\int_C f \, ds$	$\int_C F \cdot dr$
surface	$\iint_S f \, dS$	$\iint_S F \cdot n \, dS$

First-kind integrals use metric size: length or area. Second-kind integrals use orientation: direction along a curve or normal direction through a surface.

§110.7 Parametrization invariance

A geometric integral should not depend on the particular parametrization. For $\int_C f ds$, reparametrizing a curve changes $\|r'(t)\| dt$ in exactly the compensating way. For $\int_C F \cdot dr$, orientation-preserving reparametrizations leave the value unchanged, while orientation reversal changes sign.

For surfaces, changing parameters multiplies the area element by the absolute Jacobian in first-kind integrals. In flux integrals, the sign of the Jacobian records whether the orientation is preserved or reversed.

§110.8 Differential forms viewpoint

A vector-field line integral in $\mathbb{R}^3$ can be read as the integral of a one-form

$$\omega = P dx + Q dy + R dz.$$

A flux integral corresponds to the two-form

$$\omega_F = P dy \wedge dz + Q dz \wedge dx + R dx \wedge dy.$$

Parametrization formulas are pullback formulas for these forms.

This viewpoint explains why orientation is built into second-kind integrals but not into scalar length or area integrals.

§110.9 Orientability warning

Flux through a surface requires a consistent choice of normal. Spheres and tori are orientable; a Möbius strip is not. On a non-orientable surface, a globally defined flux integral needs extra care or must be treated locally.

§110.10 Densities versus differential forms

The first-kind line integral and the second-kind line integral look similar only because both are written with an integral sign. Their transformation laws are different.

For a curve $r(t)$, arclength is a density:

$$ds = \|r'(t)\| dt.$$

The absolute value makes it insensitive to orientation. By contrast, a one-form pulls back as

$$r^*(P dx + Q dy + R dz) = (P(r(t))x'(t) + Q(r(t))y'(t) + R(r(t))z'(t)) dt.$$

If t is reversed, dt changes sign, so the integral changes sign. This is the exact mathematical distinction between density and form.

§110.11 Flux and the Hodge star

In Euclidean three-space, a vector field $F = (P, Q, R)$ can be converted into a two-form by the Hodge star:

$$F \cdot dS \longleftrightarrow \omega_F = P dy \wedge dz + Q dz \wedge dx + R dx \wedge dy.$$

If $r(u, v)$ parametrizes a surface, then

$$\iint_S F \cdot n \, dS = \iint_D r^* \omega_F.$$

Expanding this pullback gives exactly the cross-product formula

$$F(r(u, v)) \cdot (r_u \times r_v) \, du \, dv.$$

Thus the vector-calculus formula is not arbitrary: it is the coordinate expression of pulling back a two-form.

§110.12 Boundary orientation is not a sign convention

For Stokes-type theorems, the orientation of a boundary is forced by the orientation of the interior. In the plane, positive orientation means the region stays on the left as one traverses the boundary. On an oriented surface, the boundary orientation is determined by the right-hand rule.

This matters because the theorem

$$\int_{\partial M} \omega = \int_M d\omega$$

can only be correct when the two orientations are compatible. A wrong sign in a flux or circulation problem is often not a computational accident; it means the boundary orientation and interior orientation were not matched.

§110.13 Non-orientability as a real obstruction

A scalar surface integral $\iint_S f \, dS$ makes sense on a Möbius strip because area does not need a global normal. A flux integral $\iint_S F \cdot n \, dS$ needs a consistent normal direction. The Möbius strip has no such global choice.

This is a topological obstruction, not a defect of parametrization. Locally one can always choose a normal; globally the choices may return reversed after going once around the strip. Vector calculus is already detecting orientability.

§110.14 What the parametrization is responsible for

First-kind integrals integrate scalar densities against geometric measure. For a curve this means that the parametrization contributes the speed $\|r'(t)\|$; for a surface it contributes the area factor $\|r_u \times r_v\|$. Changing parameters should not change the value, because these factors exactly compensate for stretching or compressing the parameter domain.

Second-kind integrals keep orientation. A line integral of work uses the directed tangent displacement dr , and a flux integral uses the oriented normal vector $r_u \times r_v$ or a chosen unit normal n . This is why reversing orientation changes the sign of work and flux but not the sign of scalar length or area. The Möbius-strip example marks the boundary of the usual formula: scalar surface area is still meaningful, but a global flux integral is not available until a consistent normal direction has been chosen.

Green's Theorem, Conservative Fields, and the Planar Stokes Principle

N-20260417

§111.1 Green's theorem

Green's theorem relates circulation around a boundary to local rotation inside the region.

Theorem 111.1.1 (Green's theorem)

Let $D \subseteq \mathbb{R}^2$ be a positively oriented region with piecewise smooth simple closed boundary $C = \partial D$. If P and Q have continuous first partial derivatives on a neighborhood of D , then

$$\oint_C P dx + Q dy = \iint_D \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} \right) dA.$$

The left side is circulation along the boundary. The right side is total scalar curl over the region.




Figure 111.1: Simple closed curves and non-simple closed curves in the statement of Green's theorem.

§111.2 Proof idea: rectangles and cancellation

For a rectangle $[a, b] \times [c, d]$, the theorem follows from the fundamental theorem of calculus. The $Q dy$ part detects horizontal change in Q , and the $P dx$ part detects vertical change in P with the opposite sign.

For a region subdivided into many rectangles, internal edges are traversed twice in opposite directions. Their line integrals cancel. Only the outer boundary remains.

对于平面区域 D 的边界曲线 C , 我们规定 C 的正向如下: 当观察者沿 C 的这个方向行走时, 区域 D 总在他的左手边. 例如, 圆环 $R = \{(x, y) | 1 < x^2 + y^2 < 2\}$ 的边界是圆周 $L: x^2 + y^2 = 2$ 及 $l: x^2 + y^2 = 1$. 作为 R 的正向边界, L 的正向是逆时针方向, 而 l 的正向是顺时针方向 (见图 57).




Figure 111.2: For a vertically simple region, Green's theorem is proved by reducing to one-variable fundamental theorem computations.




Figure 111.3: Subdivision by auxiliary arcs: internal boundary contributions cancel, leaving only the outer boundary.

§111.3 Conservative fields and path independence

A vector field $F = (P, Q)$ is **conservative** on a region U if there is a potential function ϕ such that

$$\nabla \phi = F.$$

Then

$$\int_C P dx + Q dy = \phi(B) - \phi(A)$$

for every path C from A to B . Thus the integral depends only on endpoints.

Proposition 111.3.1

If $F = \nabla \phi$, then $P_y = Q_x$.

Proof. If $P = \phi_x$ and $Q = \phi_y$, then under the usual continuity hypotheses,

$$P_y = \phi_{xy} = \phi_{yx} = Q_x.$$

□

The converse requires a topological hypothesis.

Theorem 111.3.2 (Conservative-field test)

If $U \subseteq \mathbb{R}^2$ is simply connected and $P_y = Q_x$ on U , then $F = (P, Q)$ is conservative on U .




Figure 111.4: Simply connected and multiply connected plane regions. The distinction matters for path independence and conservative fields.

§111.4 The punctured-plane warning

The standard example is

$$F(x, y) = \left(-\frac{y}{x^2 + y^2}, \frac{x}{x^2 + y^2} \right)$$

on $\mathbb{R}^2 \setminus \{0\}$. A direct computation gives $P_y = Q_x$ away from the origin, but around the unit circle $r(t) = (\cos t, \sin t)$,

$$F(r(t)) = (-\sin t, \cos t), \quad r'(t) = (-\sin t, \cos t),$$

so

$$\oint_{S^1} F \cdot dr = \int_0^{2\pi} 1 \, dt = 2\pi.$$

Therefore no global potential exists on the punctured plane. The hole is detected by circulation.

§111.5 Area by Green's theorem

Green's theorem can compute area. Choose

$$P = -\frac{y}{2}, \quad Q = \frac{x}{2}.$$

Then $Q_x - P_y = 1$, so

$$\text{area}(D) = \iint_D 1 \, dA = \frac{1}{2} \oint_{\partial D} x \, dy - y \, dx.$$

Example 111.5.1

For the unit circle $x = \cos t$, $y = \sin t$,

$$\frac{1}{2} \int_0^{2\pi} (xy' - yx') \, dt = \frac{1}{2} \int_0^{2\pi} 1 \, dt = \pi.$$

§111.6 Differential forms notation

Let

$$\omega = P \, dx + Q \, dy.$$

Then

$$d\omega = (Q_x - P_y) \, dx \wedge dy.$$

Green's theorem becomes

$$\int_{\partial D} \omega = \int_D d\omega.$$

This is Stokes' theorem in the plane.

The algebraic slogan is: integrate a form over the boundary, or integrate its derivative over the interior.

§111.7 Closed forms, exact forms, and topology

A one-form ω is **closed** if $d\omega = 0$. It is **exact** if $\omega = d\phi$ for some function ϕ . Exact forms are always closed, because $d^2 = 0$. The converse depends on the domain.

On simply connected plane regions, closed one-forms are exact. On the punctured plane, the form

$$\omega = \frac{-y dx + x dy}{x^2 + y^2}$$

is closed but not exact, because

$$\int_{S^1} \omega = 2\pi.$$

§111.8 Connection with de Rham cohomology

The first de Rham cohomology group is

$$H_{\text{dR}}^1(M) = \frac{\ker(d : \Omega^1(M) \rightarrow \Omega^2(M))}{\text{im}(d : \Omega^0(M) \rightarrow \Omega^1(M))}.$$

It measures closed one-forms modulo exact one-forms. In elementary vector calculus language, it measures the obstruction to finding potentials for curl-free fields.

§111.9 Winding number hidden in the punctured-plane example

For the one-form

$$\omega = \frac{-y dx + x dy}{x^2 + y^2},$$

write $x = r \cos \theta$, $y = r \sin \theta$. Then formally

$$\omega = d\theta.$$

The word "formally" matters: θ is not a globally single-valued function on $\mathbb{R}^2 \setminus \{0\}$. Around a closed loop γ ,

$$\frac{1}{2\pi} \int_{\gamma} \omega$$

is the winding number of γ about the origin.

For the unit circle, the winding number is 1. For a loop going twice around, it is 2. For a loop not enclosing the origin, it is 0. Thus the line integral is not merely detecting that a potential is missing; it measures how the curve winds around the hole.

§111.10 From Green's theorem to residues

The same punctured-plane form is connected to complex analysis. Since

$$\frac{dz}{z} = \frac{dx + i dy}{x + iy},$$

one computes

$$\operatorname{Im} \frac{dz}{z} = \frac{x dy - y dx}{x^2 + y^2} = \omega.$$

Thus

$$\int_{S^1} \omega = 2\pi$$

is the imaginary part of

$$\int_{S^1} \frac{dz}{z} = 2\pi i.$$

The residue theorem is a complex-analytic refinement of the same topological fact: a closed integral can detect a puncture and its winding number.

§111.11 Cohomology as bookkeeping for failed potentials

The statement "curl-free implies conservative" is not universally true; it is true when the relevant cohomology vanishes. In de Rham notation,

$$H_{\text{dR}}^1(M) = \frac{\text{closed one-forms}}{\text{exact one-forms}}.$$

A nonzero class is represented by a closed form that is not the differential of any global function.

For $M = \mathbb{R}^2 \setminus \{0\}$, the class of ω above generates $H_{\text{dR}}^1(M) \cong \mathbb{R}$. The integral over S^1 is a period pairing:

$$([\omega], [S^1]) \mapsto \int_{S^1} \omega.$$

So the elementary warning about holes is really the first example of pairing cohomology with cycles.

§111.12 Boundary cancellation and the algebra $\partial^2 = 0$

The rectangle-subdivision proof of Green's theorem says that every internal edge appears twice with opposite orientations. Algebraically this is the principle

$$\partial^2 = 0 : \quad \text{the boundary of a boundary cancels.}$$

De Rham theory has the dual identity

$$d^2 = 0.$$

Stokes' theorem connects them:

$$\int_{\partial C} \omega = \int_C d\omega.$$

If one applies this again, both sides become zero for structural reasons. Thus Green's theorem is not an isolated vector-calculus trick; it is the two-dimensional visible face of a chain/cochain duality.

§111.13 What Green's theorem checks in practice

Green's theorem converts a boundary integral into an area integral:

$$\oint_{\partial D} P dx + Q dy = \iint_D \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} \right) dA.$$

In a computation this requires more than recognizing the formula. The boundary must be positively oriented, the vector field must be defined on a region containing D , and holes in the domain must be accounted for by extra boundary components with the correct orientation.

The punctured-plane example explains why this check matters. The form

$$\omega = \frac{-y dx + x dy}{x^2 + y^2}$$

is closed away from the origin, but its integral around S^1 is nonzero, so it is not globally exact on $\mathbb{R}^2 \setminus \{0\}$. Green's theorem, the identity $d^2 = 0$, and the cancellation of internal boundaries are therefore three views of the same mechanism: local derivative data controls boundary circulation only after the topology of the region has been included.

Surface Integrals, Gauss Theorem, and Stokes Theorem

N-20260615

§112.1 Why this note is split this way

Surface integration is where several notations begin to look dangerously similar. The guiding principle of this note is that every integral should be identified by the geometric object it integrates over and by the degree of the quantity being integrated. A scalar surface integral measures density against area; a flux integral measures a two-dimensional oriented flow; Gauss turns a two-dimensional boundary flux into a three-dimensional divergence integral; Stokes turns a one-dimensional boundary circulation into a two-dimensional curl integral.

This is not a taxonomy for its own sake. It is the first place where the general formula

$$\int_{\partial M} \omega = \int_M d\omega$$

starts to become visible behind vector calculus.

§112.2 Four similar-looking integrals

The first task is to separate four objects:

$$\begin{aligned} \iint_S f(x, y, z) dS, & \quad \iint_S P dy dz + Q dz dx + R dx dy, \\ \iiint_\Omega \nabla \cdot F dV, & \quad \oint_{\partial S} P dx + Q dy + R dz. \end{aligned}$$

The first is a scalar surface integral: it integrates a scalar density over area. The second is a flux integral: it integrates a vector field through an oriented surface. The third is the volume integral of divergence. The fourth is circulation along an oriented boundary curve.

The common mistake is to treat all of them as “surface integral formulas.” They are not the same mathematical object. Their relation is supplied by Gauss’ theorem and Stokes’ theorem.

§112.3 First-kind surface integrals

Let S be a smooth surface parametrized by

$$r(u, v) = (x(u, v), y(u, v), z(u, v)), \quad (u, v) \in D.$$

Then

$$dS = |r_u \times r_v| du dv, \quad \iint_S f dS = \iint_D f(r(u, v)) |r_u \times r_v| du dv.$$

This is the surface analogue of the first-kind line integral $\int_C f ds$.

For a graph $z = z(x, y)$,

$$r(x, y) = (x, y, z(x, y)), \quad r_x \times r_y = (-z_x, -z_y, 1),$$

so

$$dS = \sqrt{1 + z_x^2 + z_y^2} dx dy.$$

For a graph $x = x(y, z)$ or $y = y(z, x)$, analogous formulas hold:

$$dS = \sqrt{1 + x_y^2 + x_z^2} dy dz, \quad dS = \sqrt{1 + y_z^2 + y_x^2} dz dx.$$




Figure 112.1: A surface area element and its projection. The factor between dS and the projected area element is determined by the angle between the surface normal and the projection direction.

§112.4 Template surfaces

The original note lists the common templates, and they should be kept because they prevent repeated derivations.

For a plane $z = ax + by + c$,

$$dS = \sqrt{1 + a^2 + b^2} dx dy.$$

For a sphere $x^2 + y^2 + z^2 = R^2$, spherical coordinates give

$$dS = R^2 \sin \varphi d\varphi d\theta.$$

For a cylinder $x = R \cos \theta$, $y = R \sin \theta$, $z = z$,

$$dS = R d\theta dz.$$

For a cone $z = \sqrt{x^2 + y^2}$, use $x = r \cos \theta$, $y = r \sin \theta$, $z = r$, so

$$|r_r \times r_\theta| = \sqrt{2} r.$$




Figure 112.2: A torus as a standard example of a parametrized surface.

§112.5 Example pattern: spherical cap

For a spherical cap on $x^2 + y^2 + z^2 = R^2$, the scalar integral of f should normally be written in spherical coordinates. If the cap is cut by $z = h$, then $z = R \cos \varphi$ gives

$$0 \leq \varphi \leq \arccos(h/R).$$

Thus

$$\iint_S f \, dS = \int_0^{2\pi} \int_0^{\arccos(h/R)} f(R \sin \varphi \cos \theta, R \sin \varphi \sin \theta, R \cos \varphi) R^2 \sin \varphi \, d\varphi \, d\theta.$$

The important part of the example is not the final antiderivative. It is recognizing the right parameter domain.




Figure 112.3: A spherical surface patch and its projection domain. This is the geometric model behind many spherical-cap surface integrals.




Figure 112.4: A related spherical surface cut out by a vertical boundary. The dashed curves indicate the hidden projection and bounding geometry.

§112.6 Second-kind surface integrals as flux

Let $F = (P, Q, R)$. The flux through an oriented surface S is

$$\iint_S F \cdot n \, dS.$$

In differential notation this is

$$\iint_S P \, dy \, dz + Q \, dz \, dx + R \, dx \, dy.$$

For a parametrization $r(u, v)$,

$$\iint_S F \cdot n \, dS = \iint_D F(r(u, v)) \cdot (r_u \times r_v) \, du \, dv,$$

where the order of u, v determines orientation.

For a graph $z = z(x, y)$ with upward orientation,

$$n \, dS = (-z_x, -z_y, 1) \, dx \, dy,$$

so

$$\iint_S F \cdot n \, dS = \iint_D [-Pz_x - Qz_y + R] \, dx \, dy.$$

For downward orientation, change the sign.

§112.7 Model flux computations: closed, capped, singular, and sheeted

The useful flux examples are not a random list of surfaces; they represent four different decisions about orientation and theorem choice.

§112.7.i Closed sphere: Gauss beats parametrization

Let $F(x, y, z) = (x, y, z)$ and let S be the sphere $x^2 + y^2 + z^2 = R^2$ with outward orientation. Directly, the outward unit normal is $n = x/R$, so

$$F \cdot n = R, \quad dS = R^2 \sin \varphi \, d\varphi \, d\theta.$$

Hence

$$\iint_S F \cdot n \, dS = R \cdot 4\pi R^2 = 4\pi R^3.$$

Gauss gives the same answer faster:

$$\nabla \cdot F = 3, \quad \iiint_{B_R} 3 \, dV = 3 \cdot \frac{4}{3}\pi R^3 = 4\pi R^3.$$

The theorem choice is dictated by closedness of the surface.

§112.7.ii Open surface: close it and subtract the cap

For the upper hemisphere S of radius R , outward orientation, Gauss cannot be applied to S alone because it is not closed. Close it by adding the disk D in the plane $z = 0$. For $F = (x, y, z)$, the flux through D is zero because $F \cdot n = 0$ there. Therefore the hemisphere flux equals the full closed half-ball flux:

$$\iint_S F \cdot n \, dS = \iiint_{z \geq 0, x^2 + y^2 + z^2 \leq R^2} 3 \, dV = 2\pi R^3.$$

For another field, the cap flux might not vanish; subtracting it is the whole point of the method.

§112.7.iii Cone: orientation is the computation

For the cone $z = r$, $0 \leq r \leq a$, use

$$r(r, \theta) = (r \cos \theta, r \sin \theta, r).$$

Then

$$r_r \times r_\theta = (-r \cos \theta, -r \sin \theta, r),$$

which has positive z -component. Thus it is the upward orientation. If a problem asks for outward orientation of the cone as the side of a solid below its tip, the sign may be the opposite. Here the formula is easy; choosing the correct normal is the real work.

§112.7.iv Two sheets over one projection

For a sphere cut by projection onto the xy -plane, upper and lower sheets have opposite normal components. A flux written as

$$\iint_S R \, dx \, dy$$

really means integrating the two-form $R \, dx \wedge dy$. On the upper sheet, upward orientation contributes $+R \, dx \, dy$; on the lower sheet, the same geometric outward orientation may contribute $-R \, dx \, dy$. Thus sheets may add or cancel depending on the chosen orientation and the component of the flux form.

This is the differential-form reason that projection formulas are dangerous when treated as scalar area formulas.

§112.8 Gauss theorem

Theorem 112.8.1 (Gauss divergence theorem)

For a compact region Ω with oriented boundary $\partial\Omega$,

$$\int_{\partial\Omega} F \cdot n \, dS = \int_{\Omega} \nabla \cdot F \, dV.$$

Let Ω be a solid with outward boundary $\partial\Omega$. For $F = (P, Q, R)$,

$$\iint_{\partial\Omega} F \cdot n \, dS = \iiint_{\Omega} (P_x + Q_y + R_z) \, dV.$$

The divergence

$$\nabla \cdot F = P_x + Q_y + R_z$$

measures local source strength. Gauss' theorem says that total source inside equals outward flux through the boundary.

Use Gauss when the surface is closed, or when it can be closed by adding a simple cap. If the required surface is not closed, compute the closed flux and subtract the flux across the added part. If the problem asks for inward orientation, multiply the outward result by -1 .

§112.9 Singular fields and missing points

The note includes the important example type where the field is singular at a point. Gauss' theorem cannot be applied across a singularity. The usual remedy is to remove a small sphere around the singular point, apply Gauss' theorem to the punctured region, and then let the radius tend to zero. For inverse-square radial fields, the flux through the small sphere often remains nonzero. This is the analytic origin of point charges and delta-type sources.

§112.10 Stokes theorem

Theorem 112.10.1 (Stokes theorem)

For an oriented surface S with boundary ∂S ,

$$\int_{\partial S} F \cdot dr = \int_S (\nabla \times F) \cdot n \, dS.$$

Let S be an oriented surface with boundary ∂S . Then

$$\oint_{\partial S} F \cdot dr = \iint_S (\nabla \times F) \cdot n \, dS.$$

Here

$$\nabla \times F = \begin{vmatrix} i & j & k \\ \partial_x & \partial_y & \partial_z \\ P & Q & R \end{vmatrix}.$$

The right-hand rule fixes orientation: if the thumb points along n , the fingers curl in the positive boundary direction.

Use Stokes when the boundary curve is complicated but some spanning surface is simple, or when the curl is much simpler than the original vector field. The surface can be changed as long as the boundary and orientation are preserved.

§112.11 Stokes model computations: area vector and spanning-surface freedom

§112.11.i Triangle in a plane

Suppose ∂S is the positively oriented boundary of a triangle with vertices A, B, C , and suppose $\nabla \times F = c$ is constant. Stokes gives

$$\oint_{\partial S} F \cdot dr = \iint_S c \cdot n \, dS = c \cdot \vec{A}_S,$$

where the oriented area vector is

$$\vec{A}_S = \frac{1}{2}(B - A) \times (C - A).$$

So the whole line integral collapses to one vector dot product. The triangle is not merely a convenient region; it packages orientation and area into $\vec{A}_S$.

§112.11.ii Ellipse as a chosen spanning surface

If the boundary curve is an ellipse lying in a plane and $\nabla \times F$ is simple, Stokes lets us choose the planar elliptical disk rather than parameterizing the boundary. If the ellipse has semiaxes a, b and unit normal n , then its oriented area vector is

$$\vec{A} = \pi ab n.$$

When $\nabla \times F = c$ is constant,

$$\oint_{\partial S} F \cdot dr = c \cdot (\pi ab n).$$

This illustrates the true power of Stokes: the line integral depends on the boundary, but the chosen surface can be any surface spanning it, provided orientation is consistent.

§112.11.iii Why surface replacement is legal

If two oriented surfaces S_1 and S_2 have the same boundary, then applying Stokes to the closed surface $S_1 \cup (-S_2)$ gives

$$\iint_{S_1} (\nabla \times F) \cdot n \, dS - \iint_{S_2} (\nabla \times F) \cdot n \, dS = 0,$$

as long as the field is smooth in the region between them. If there is a singularity between the two surfaces, this argument can fail. This is the Stokes-side analogue of the singular-source warning in Gauss' theorem.

§112.12 Decision table

Object	Meaning	Best tool
$\iint_S f dS$	scalar density on surface	parametrization or graph dS
$\iint_S F \cdot n dS$	flux through oriented surface	graph formula or Gauss
closed flux	total outward flow	Gauss theorem
line circulation	work/circulation along boundary	direct line integral or Stokes

§112.13 Conceptual endpoint

The high-level statement is the boundary principle:

$$\int_{\partial M} \omega = \int_M d\omega.$$

Green, Gauss, and Stokes are all versions of this formula. The old vector-calculus notation hides the unity; differential forms reveal it. The derivative $d\omega$ records local change inside a region, and the boundary integral records the accumulated effect at the edge.

§112.14 The de Rham complex behind the vector formulas

The vector identities become less mysterious if the objects are placed by degree:

degree	form in $\mathbb{R}^3$	vector-calculus name
0	f	scalar potential
1	$P dx + Q dy + R dz$	work/circulation form
2	$P dy \wedge dz + Q dz \wedge dx + R dx \wedge dy$	flux form
3	$g dx \wedge dy \wedge dz$	volume density

The exterior derivative moves one row down:

$$d : \Omega^0 \rightarrow \Omega^1 \rightarrow \Omega^2 \rightarrow \Omega^3.$$

In vector notation, these three moves are gradient, curl, and divergence. The identities

$$\nabla \times (\nabla f) = 0, \quad \nabla \cdot (\nabla \times F) = 0$$

are both the single identity

$$d^2 = 0.$$

This is why conservative fields, curl-free forms, flux through closed surfaces, and singular sources are connected rather than separate topics.

§112.15 General Stokes theorem

Theorem 112.15.1 (General Stokes theorem)

For a compact oriented manifold M with boundary and a differential form ω ,

$$\int_M d\omega = \int_{\partial M} \omega.$$

The vector-calculus theorems in this note are all cases of the general Stokes theorem. If M is an oriented smooth manifold with boundary and ω is a compactly supported differential form of degree one less than $\dim M$, then

$$\int_{\partial M} \omega = \int_M d\omega.$$

Green's theorem, Stokes' theorem, and Gauss' theorem correspond to different degrees and dimensions:

Green	$M \subseteq \mathbb{R}^2$	$\omega = Pdx + Qdy,$
Stokes	M a surface in $\mathbb{R}^3$	$\omega = Pdx + Qdy + Rdz,$
Gauss	$M \subseteq \mathbb{R}^3$	$\omega = Pdy \wedge dz + Qdz \wedge dx + Rdx \wedge dy.$

The apparently different vector formulas are one theorem written in different degrees.

§112.16 Singular sources and distributions

The note mentions singular fields and missing points. The advanced language for point sources is distribution theory. The Dirac delta δ_0 is not an ordinary function but satisfies

$$\int \delta_0 \varphi = \varphi(0)$$

for test functions φ .

For the inverse-square field in $\mathbb{R}^3$,

$$F(x) = \frac{x}{|x|^3},$$

one has classically $\nabla \cdot F = 0$ away from the origin, but distributionally

$$\nabla \cdot F = 4\pi\delta_0.$$

This reconciles the nonzero flux through spheres with zero ordinary divergence away from the singularity.

§112.17 Conservation laws

Many PDEs have conservation-law form

$$\partial_t \rho + \nabla \cdot J = 0.$$

Integrating over a region Ω and applying Gauss' theorem gives

$$\frac{d}{dt} \int_{\Omega} \rho dV = - \iint_{\partial\Omega} J \cdot n dS.$$

Change inside equals negative outward flux. Thus Gauss' theorem is the geometric mechanism behind conservation of mass, charge, and probability.

§112.18 Currents

A current generalizes an oriented submanifold by acting on differential forms. An oriented curve C defines a current

$$T_C(\omega) = \int_C \omega.$$

The boundary of a current is defined by

$$\partial T(\omega) = T(d\omega).$$

This extends Stokes' theorem to singular chains, weak surfaces, and geometric measure theory.

The elementary surface and curve integrals in the note are therefore the smooth case of a much more flexible integration theory over generalized geometric objects.

§112.19 Gauss and Stokes as dual local-to-global laws

Gauss' theorem and Stokes' theorem are both local-to-global principles. Gauss says that local source density accumulates to boundary flux. Stokes says that local rotation accumulates to boundary circulation. In differential-form notation, both are one formula:

$$\int_{\partial M} \omega = \int_M d\omega.$$

The difference lies only in the degree of the form and the dimension of the manifold.

This unification matters because it prevents memorizing vector identities as separate facts. The object being integrated determines the theorem: a one-form over a boundary curve gives curl through a surface; a two-form over a closed surface gives divergence through a volume.

§112.20 Weak forms and PDE

In PDE, integration by parts and Stokes-type formulas define weak solutions. Instead of requiring derivatives to exist pointwise, one tests against smooth functions. For example, a weak divergence equation is read through

$$\int_{\Omega} F \cdot \nabla \varphi \, dV = - \int_{\Omega} (\nabla \cdot F) \varphi \, dV$$

up to boundary terms.

Thus vector-calculus integral theorems are not merely computational shortcuts; they are the foundation of weak formulations, conservation laws, and modern PDE theory.

XI

Complex Analysis

Part XI: Contents

N-20220321	729
113.1 The complex exponential as motion	729
113.2 Roots and angles	730
113.3 A divisibility result	730
113.4 Closing perspective	732
113.5 Why multiplication by i is rotation	732
113.6 Roots of unity as symmetry	732
113.7 Context and outlook	732
113.8 Roots of unity as a finite symmetry group	733
113.9 Cyclotomic factorization	733
113.10 Fourier viewpoint	733
113.11 Cyclic symmetry on the unit circle	733
113.12 Roots of unity as characters	733
113.13 Discrete Fourier transform	734
113.14 Line-to-circle covering map	734
113.15 Cyclotomic fields	734
113.16 Finite Fourier analysis from roots of unity	734
113.17 Cyclotomic fields and constructibility	735
N-20220323	737
114.1 Why complex analysis feels different	737
114.2 Complex derivative	737
114.3 Basic examples	738
114.4 Closure properties	739
114.5 The main intuition	739
114.6 Complex versus real differentiability	739
114.7 Local power series intuition	739
N-20220813	741
115.1 Why complex integration is different	741
115.2 Real integrals as a warm-up	741
115.3 Contours and parametrization	741
115.4 Riemann sums and midpoint sums	742
115.5 Complex Riemann sums	743
115.6 Integrals of powers on the unit circle	743
115.7 Cauchy–Goursat theorem	743
115.8 Contour integrals depend on paths	744
115.9 Primitive functions and path independence	744
115.10 Why singularities matter	744
115.11 A bridge outward	744
115.12 Contour integrals as integrals of one-forms	745
115.13 Cauchy theorem and exactness	745
115.14 Residues as obstruction	745
115.15 Residues, topology, and exactness	745
N-20260616	747
116.1 Why these topics belong together	747
116.2 Numerical series	747
116.3 Positive-term series	747
116.4 Ratio and root tests	748
116.5 Alternating and arbitrary-sign series	748
116.6 Function series and uniform convergence	749
116.7 Power series	749
116.8 Taylor series	749
116.9 Improper integrals	750
116.10 Unified viewpoint	750
116.11 Uniform convergence as norm convergence	750

116.12	Analytic functions and convergence radius	751
116.13	Dominated convergence and improper integrals	751
116.14	Tauberian flavor	751
N-20260617		753
117.1	Why Fourier series is a separate theme	753
117.2	Orthogonality of trigonometric functions	753
117.3	Fourier coefficients	753
117.4	Dirichlet convergence theorem	753
117.5	Even and odd functions	754
117.6	Half-range expansions	754
117.7	Concrete coefficient models	754
117.8	Dirichlet kernel: what partial sums really do	756
117.9	Fejer means: averaging partial sums improves convergence	756
117.10	Hilbert-space viewpoint	756
117.11	Fourier series in Hilbert space	757
117.12	Parseval identity	757
117.13	Gibbs phenomenon	757
117.14	Sturm–Liouville and Fourier bases	757
117.15	Heat and wave equations	758
117.16	Convergence modes of Fourier series	758
117.17	Smoothness is decay of coefficients	758
117.18	Sobolev norm from Fourier coefficients	759
117.19	Fourier transform as the non-periodic limit	759
117.20	Distributions and weak Fourier series	759
N-20260618		761
118.1	Why the differential-equation part is one note	761
118.2	Basic concepts	761
118.3	Separable equations	761
118.4	Homogeneous first-order equations	761
118.5	First-order linear equations	762
118.6	Bernoulli equations	762
118.7	Exact equations	762
118.8	Reducible higher-order equations	763
118.9	Second-order linear equations	763
118.10	Constant-coefficient homogeneous equations	763
118.11	Constant-coefficient nonhomogeneous equations	764
118.12	Initial-value problems	764
118.13	Functional-analytic viewpoint	764
118.14	Differential equations as operators	764
118.15	Green functions	765
118.16	Semigroups and evolution	765
118.17	Sturm–Liouville and Fourier expansion	765
N-20260619		767
119.1	What a first-order PDE asks for	767
119.2	The unforced transport equation	767
119.3	Example: translating a parabola	768
119.4	Forced transport	769
119.5	What can go wrong	770
119.6	General linear first-order equations	770
119.7	Quasilinear equations	770
119.8	Crossing characteristics and shocks	771
119.9	Initial curves and well-posedness	772
119.10	Conservation-law viewpoint	772
N-20260620		773
120.1	The equation and its physical meaning	773
120.2	Factorization into transport operators	773
120.3	Initial-value problem on the whole line	774

120.4	Finite strings and Fourier series	775
120.5	Energy conservation	775
120.6	Inhomogeneous wave equation and Duhamel's principle	776
120.7	Boundary conditions on a finite interval	776
120.8	Deriving the Fourier coefficients	776
120.9	Standing waves and travelling waves	777
120.10	The wave equation with damping	777
N-20260621		779
121.1	Diffusion versus wave motion	779
121.2	Heat kernel on the whole line	779
121.3	Similarity solution on a semi-infinite rod	779
121.4	Fourier series for a finite rod	780
121.5	Maximum principle	781
121.6	Energy decay on an interval	781
121.7	The maximum principle with proof idea	781
121.8	Smoothing and the heat semigroup	782
121.9	Fourier transform solution on the real line	782
121.10	Boundary conditions and conservation	783
121.11	Steady states and long-time behavior	783
121.12	Comparison with the wave equation	784

Some Pre-Complex Analysis: Motion on the Unit Circle and Roots of Unity

N-20220321

§113.1 The complex exponential as motion

The first example to keep in mind is

$$Z(t) = e^{it}.$$

Here t is time and $Z(t)$ is the position of a moving particle in the complex plane.

Differentiating gives

$$\frac{d}{dt}e^{it} = ie^{it}.$$

Thus the velocity is

$$V(t) = iZ(t).$$

Multiplication by i rotates a complex number by a right angle, so the velocity is always perpendicular to the position vector.

At $t = 0$, the particle starts at

$$Z(0) = 1,$$

and the velocity is

$$V(0) = i.$$

So the particle initially moves upward. A moment later the position vector has turned slightly, and the velocity is again obtained by rotating that new position vector by 90° . Continuing this construction gives motion around the unit circle.




Figure 113.1: For $Z(t) = e^{it}$, multiplication by i rotates the position vector by 90° , giving the tangent velocity.




Figure 113.2: Original geometric picture for $Z'(t) = iZ(t)$: the velocity is obtained by multiplying the current position by i , so the motion follows the unit circle from the initial position 1.

§113.2 Roots and angles

The same geometric viewpoint is useful for roots in the complex plane. A root of unity is best understood not merely as a solution of an equation, but as a point on the unit circle with a specified angle.

Remark 113.2.1 — This is why complex numbers make polynomial divisibility questions visual. If a polynomial is built from powers of x , its roots often arrange themselves by angles on the unit circle.

§113.3 A divisibility result

The note records a result analogous to a previous observation about roots and angles.

Theorem 113.3.1

The polynomial

$$x^{2n} + x^n + 1$$

has the factor

$$x^2 + x + 1$$

if and only if

$$n \not\equiv 0 \pmod{3}.$$

Proof. The roots of $x^2 + x + 1$ are the primitive third roots of unity, say

$$\omega = e^{2\pi i/3}, \quad \omega^2 = e^{4\pi i/3}.$$

Thus $x^2 + x + 1 \mid x^{2n} + x^n + 1$ exactly when

$$\omega^{2n} + \omega^n + 1 = 0.$$

Let $k \equiv n \pmod{3}$. If $k = 0$, then $\omega^n = 1$, so the value is

$$1 + 1 + 1 = 3 \neq 0.$$

If $k = 1$, then $\omega^n = \omega$, and

$$\omega^{2n} + \omega^n + 1 = \omega^2 + \omega + 1 = 0.$$

If $k = 2$, then $\omega^n = \omega^2$, and

$$\omega^{2n} + \omega^n + 1 = \omega^4 + \omega^2 + 1 = \omega + \omega^2 + 1 = 0.$$

Therefore divisibility holds precisely when $n \not\equiv 0 \pmod{3}$. □




Figure 113.3: Primitive third roots of unity on the unit circle. Divisibility by $x^2 + x + 1$ is checked by evaluating at ω and ω^2 .




Figure 113.4: Original root-of-unity plots: increasing the number of roots fills the unit circle more densely, while highlighted roots show how special powers or congruence classes are read geometrically.

§113.4 Closing perspective

The two ideas in this note are connected by geometry. The derivative of e^{it} says that motion on the unit circle is governed by rotation by i . The divisibility theorem says that evaluating a polynomial at roots of unity turns algebra into angle arithmetic.

§113.5 Why multiplication by i is rotation

Write a complex number as $z = x + iy$. Then

$$iz = i(x + iy) = -y + ix.$$

As a vector, this sends

$$(x, y) \mapsto (-y, x),$$

which is rotation by 90° counterclockwise. Therefore the differential equation

$$Z'(t) = iZ(t)$$

says that the velocity is always obtained by rotating the radius vector by a right angle. This is exactly uniform circular motion.

The equation also explains why exponentials appear in complex analysis: the function e^{it} is the solution of a linear differential equation whose coefficient is multiplication by i .

§113.6 Roots of unity as symmetry

The equation

$$z^n = 1$$

has solutions

$$z_k = e^{2\pi i k/n}, \quad k = 0, 1, \dots, n-1.$$

These points form a regular n -gon on the unit circle. Multiplying by $e^{2\pi i/n}$ rotates the polygon by one step. Thus roots of unity are not merely algebraic solutions; they are the rotational symmetries of the circle.

For $n = 3$, if $\omega = e^{2\pi i/3}$, then

$$1 + \omega + \omega^2 = 0.$$

Geometrically, the three vectors form an equilateral triangle centered at the origin, so their sum is zero. Algebraically, the same identity comes from

$$x^3 - 1 = (x - 1)(x^2 + x + 1).$$

§113.7 Context and outlook

This note is a prelude to complex analysis because it already contains three central ideas: complex multiplication is geometry, exponentials solve differential equations, and roots of unity connect algebraic equations with symmetry. Later topics such as residues, Fourier analysis, cyclotomic polynomials, and covering maps all reuse this same unit-circle picture.

§113.8 Roots of unity as a finite symmetry group

The n -th roots of unity form a cyclic group

$$\mu_n = \{z \in \mathbb{C} : z^n = 1\} = \{e^{2\pi i k/n} : 0 \leq k < n\}.$$

Multiplication corresponds to addition of exponents modulo n . Thus roots of unity are not just solutions of an equation; they are the rotational symmetries of a regular n -gon.

§113.9 Cyclotomic factorization

The factorization

$$x^n - 1 = \prod_{d|n} \Phi_d(x)$$

separates roots of unity by exact order. The polynomial Φ_n collects primitive n -th roots. This is the algebraic form of decomposing rotations by their period.

Such factorization explains why divisibility problems involving $x^n - 1$ often reduce to order and gcd computations.

§113.10 Fourier viewpoint

Roots of unity also underlie the discrete Fourier transform. For a sequence $(a_0, \dots, a_{n-1})$, the transform evaluates the polynomial

$$A(x) = \sum_{k=0}^{n-1} a_k x^k$$

at roots of unity. Orthogonality of characters gives inversion.

§113.11 Cyclic symmetry on the unit circle

Motion on the unit circle, roots of unity, and divisibility of cyclotomic expressions are one story: the circle group has finite cyclic subgroups, and algebra records their periods through cyclotomic polynomials and Fourier characters.

§113.12 Roots of unity as characters

The n -th roots of unity form a cyclic group

$$\mu_n = \{z \in \mathbb{C} : z^n = 1\}.$$

If $\zeta = e^{2\pi i/n}$, then every root is ζ^k . The functions

$$\chi_m(k) = \zeta^{mk}$$

are characters of the finite cyclic group $\mathbb{Z}/n\mathbb{Z}$.

Orthogonality of roots of unity becomes character orthogonality:

$$\sum_{k=0}^{n-1} \zeta^{mk} = 0$$

unless $n \mid m$, in which case the sum is n .

§113.13 Discrete Fourier transform

For a function $f : \mathbb{Z}/n\mathbb{Z} \rightarrow \mathbb{C}$, its discrete Fourier transform is

$$\widehat{f}(m) = \sum_{k=0}^{n-1} f(k) e^{-2\pi i m k / n}.$$

The inverse formula is

$$f(k) = \frac{1}{n} \sum_{m=0}^{n-1} \widehat{f}(m) e^{2\pi i m k / n}.$$

Thus roots of unity are the basis functions for finite periodic data.

This connects elementary complex-number geometry to signal processing, finite harmonic analysis, and representation theory of finite abelian groups.

§113.14 Line-to-circle covering map

The exponential map

$$\mathbb{R} \rightarrow S^1, \quad t \mapsto e^{it}$$

is a covering map with kernel $2\pi\mathbb{Z}$. It wraps the real line around the circle infinitely many times.

Angles modulo 2π are therefore elements of the quotient group

$$\mathbb{R}/2\pi\mathbb{Z}.$$

The elementary fact that angles differing by 2π are equivalent is a quotient-space statement.

§113.15 Cyclotomic fields

The field $\mathbb{Q}(\zeta_n)$ generated by a primitive n -th root of unity is a cyclotomic field. Its Galois group is

$$\text{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \cong (\mathbb{Z}/n\mathbb{Z})^\times.$$

The elementary root powers ζ_n^a become field automorphisms when a is invertible modulo n .

Thus roots of unity are simultaneously geometry of the circle, Fourier basis functions, and algebraic numbers with arithmetic symmetries.

§113.16 Finite Fourier analysis from roots of unity

The orthogonality identity

$$\sum_{k=0}^{n-1} \zeta^{mk} = 0$$

for $n \nmid m$ is the engine of finite Fourier analysis. It implies that the matrix

$$F_{mk} = \zeta^{mk}$$

is unitary up to normalization. Therefore every function on a cyclic group decomposes uniquely into frequency components.

This gives an advanced interpretation of elementary root-of-unity sums. They are not only tools for solving polynomial equations; they are projection formulas onto irreducible characters of a finite abelian group.

§113.17 Cyclotomic fields and constructibility

Roots of unity also control constructible regular polygons. A regular n -gon is constructible by straightedge and compass exactly when the relevant cyclotomic field can be built by successive quadratic extensions. Gauss' theorem says this occurs when

$$n = 2^k p_1 \cdots p_m$$

where the p_i are distinct Fermat primes.

Thus elementary geometry of regular polygons connects to field extensions and Galois theory through roots of unity.

Complex Differentiability and the Rigidity of Holomorphic Functions

N-20220323

§114.1 Why complex analysis feels different

Real analysis contains many pathological examples:

- (i) the Cantor function is continuous, has derivative zero almost everywhere, yet is not constant;
- (ii) the Weierstrass function is continuous everywhere and differentiable nowhere;
- (iii) a function may have many derivatives but fail at the next derivative;
- (iv) a smooth real function may have a Taylor series that does not recover the function.

For example,

$$f(x) = \begin{cases} e^{-1/x}, & x > 0, \\ 0, & x \leq 0 \end{cases}$$

is smooth, but its Taylor series at 0 is identically zero and therefore does not equal f for $x > 0$.

Complex analysis is much more rigid. Holomorphic functions behave as if a tiny amount of local data can determine a large amount of global behavior.

§114.2 Complex derivative

Definition 114.2.1. Let $U \subseteq \mathbb{C}$ and $f : U \rightarrow \mathbb{C}$. The derivative of f at $z_0 \in U$ is

$$f'(z_0) = \lim_{h \rightarrow 0} \frac{f(z_0 + h) - f(z_0)}{h},$$

provided the limit exists.

The key difference from real differentiability is that h approaches 0 through the complex plane, not merely from the left and right along a line. The quotient must have the same limit from every direction.

Example 31.2.1 (Important: conjugation is *not* holomorphic)

Let $f(z) = \bar{z}$ be complex conjugation, $f : \mathbb{C} \rightarrow \mathbb{C}$. This function, despite its simple nature, is not holomorphic! Indeed, at $z = 0$ we have,

$$\frac{f(h) - f(0)}{h} = \frac{\bar{h}}{h}.$$

This does not have a limit as $h \rightarrow 0$, because depending on “which direction” we approach zero from we have different values.




Figure 114.1: In complex differentiability, the increment approaches from all directions in the plane.

Definition 114.2.2. A function $f : U \rightarrow \mathbb{C}$ is holomorphic if it is complex differentiable at every point of U . If $U = \mathbb{C}$, then f is called entire.

§114.3 Basic examples

Example 114.3.1

Every polynomial

$$a_n z^n + a_{n-1} z^{n-1} + \cdots + a_0$$

is holomorphic.

Example 114.3.2

The complex exponential is holomorphic. In real and imaginary parts,

$$e^{x+iy} = e^x (\cos y + i \sin y).$$

Example 114.3.3

The trigonometric functions extend holomorphically to the complex plane by

$$\cos z = \frac{e^{iz} + e^{-iz}}{2}, \quad \sin z = \frac{e^{iz} - e^{-iz}}{2i}.$$

§114.4 Closure properties

Holomorphic functions are closed under the usual operations:

- (i) sums of holomorphic functions are holomorphic;
- (ii) products of holomorphic functions are holomorphic;
- (iii) quotients are holomorphic wherever the denominator is nonzero;
- (iv) compositions of holomorphic functions are holomorphic.

Remark 114.4.1 — This resembles real calculus formally, but the consequences are stronger. Complex differentiability is a much more restrictive condition than real differentiability.

§114.5 The main intuition

The main point of this note is that complex differentiability is not merely ordinary differentiability with complex numbers inserted. Requiring the difference quotient to approach the same value from every complex direction forces strong compatibility conditions, later expressed by the Cauchy–Riemann equations.

Remark 114.5.1 — The next step should be to derive the Cauchy–Riemann equations and then explain why holomorphic functions are analytic. This will make precise the slogan that complex functions are rigid.

§114.6 Complex versus real differentiability

A complex function $f(z) = u(x, y) + iv(x, y)$ is real differentiable as a map $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ if it has a real linear approximation. It is complex differentiable if that linear approximation is multiplication by a complex number. This extra requirement forces the Cauchy–Riemann equations:

$$u_x = v_y, \quad u_y = -v_x.$$

Thus holomorphic functions are far more rigid than arbitrary differentiable two-variable functions.

§114.7 Local power series intuition

Once a function is holomorphic on a disk, it is represented by a power series there. This is the source of many surprising consequences: holomorphic functions are infinitely differentiable, zeros are isolated unless the function is identically zero, and values on small sets can determine the function globally.

Complex Contour Integrals and the Geometry of Integration

N-20220813

§115.1 Why complex integration is different

The note begins by comparing real functions with complex functions. If a complex function is regarded merely as a map

$$\mathbb{R}^2 \rightarrow \mathbb{R}^2,$$

then it is just a two-dimensional real function. But the complex plane has more structure: it is a field. We can multiply and divide complex numbers, and this extra structure makes holomorphic functions far more rigid than ordinary differentiable maps.

The printed text on the first page highlights two key facts:

- (i) contour integrals around closed loops often vanish;
- (ii) holomorphic functions are essentially determined by their Taylor series.

The handwritten remark correctly notices that this explains earlier surprising facts: knowing a holomorphic function on a boundary, or on a sequence with a limit point, can determine it inside.

§115.2 Real integrals as a warm-up

For a complex-valued function of one real variable,

$$f(t) = u(t) + iv(t),$$

we define

$$\int_a^b f(t) dt = \int_a^b u(t) dt + i \int_a^b v(t) dt.$$

This is just componentwise integration. Nothing mysterious has happened yet.

The conceptual jump comes when the input is not just a real interval but a path in the complex plane.

§115.3 Contours and parametrization

A contour is a path

$$\alpha : [a, b] \rightarrow \mathbb{C}.$$

If f is complex-valued, its contour integral along α is

$$\int_{\alpha} f(z) dz = \int_a^b f(\alpha(t)) \alpha'(t) dt.$$

This formula is exactly the substitution rule. The factor $\alpha'(t)$ records both speed and direction along the curve.


For the unit circle, the following parametrizations describe the same geometric path run once counterclockwise:

$$\begin{aligned}\alpha_1(t) &= e^{it}, & 0 \leq t \leq 2\pi, \\ \alpha_2(t) &= e^{2\pi it}, & 0 \leq t \leq 1.\end{aligned}$$

They give the same contour integral. A path such as

$$\alpha_3(t) = e^{2\pi it^5}$$

also traces the circle once, but at a different speed. The integral is unchanged because the parametrization is merely a way of walking along the same curve.



The unit circle as a contour.

Figure 115.1: A contour integral depends on the oriented geometric path, not on the accidental speed of parametrization.

§115.4 Riemann sums and midpoint sums

The note inserts a real-calculus reminder about Riemann sums. Divide an interval into small pieces Δ_i , choose sample points x_i , and form

$$R = \sum_{i=1}^n f(x_i) \Delta_i.$$

As the mesh tends to zero, this approximates the area under the curve.

The trapezoidal rule and midpoint rule are refinements of this same idea. The note's diagrams compare rectangles, trapezoids, and midpoint rectangles. The moral is that integration is not a formula first; it is a limiting process of adding many tiny contributions.




Figure 115.2: A schematic midpoint Riemann sum. The exact picture is less important than the limiting idea.

§115.5 Complex Riemann sums

For a complex path K , divide the path into small directed pieces Δ_i . Choose sample points z_i on the path. The complex Riemann sum is

$$\sum_i f(z_i) \Delta_i.$$

Each term means: take the tiny tangent displacement Δ_i , multiply it by the complex number $f(z_i)$, thereby scaling and rotating it, and then add all these transformed displacements.

This gives the geometric meaning of

$$\int_K f(z) dz.$$

A complex integral is not area in the usual sense. It is a limit of transformed tangent vectors.

§115.6 Integrals of powers on the unit circle

The note states the key example:

$$\gamma(t) = e^{it}, \quad 0 \leq t \leq 2\pi.$$

Then

$$\int_{\gamma} z^m dz = \int_0^{2\pi} e^{imt} i e^{it} dt = i \int_0^{2\pi} e^{i(m+1)t} dt.$$

If $m \neq -1$, then the exponential completes full oscillations and cancels out:

$$\int_{\gamma} z^m dz = 0.$$

If $m = -1$, then the integrand is simply i , so

$$\int_{\gamma} \frac{1}{z} dz = 2\pi i.$$

Thus

$$\int_{\gamma} z^m dz = \begin{cases} 2\pi i, & m = -1, \\ 0, & m \neq -1. \end{cases}$$

This is the seed of residue theory. The function $1/z$ is special because it detects the winding of the loop around the origin.

§115.7 Cauchy–Goursat theorem

Theorem 115.7.1 (Cauchy–Goursat theorem)

If f is holomorphic on a simply connected domain D , then for every closed contour $\gamma \subset D$,

$$\int_{\gamma} f(z) dz = 0.$$

The last page states the Cauchy–Goursat theorem: if f is holomorphic on a simply connected domain Ω , then for every closed contour γ in Ω ,

$$\int_{\gamma} f(z) dz = 0.$$

The note compares two paths with the same endpoints and suggests path independence in simply connected domains. This is the right intuition. If every loop can contract without crossing a singularity, holomorphic integrals do not accumulate circulation.

§115.8 Contour integrals depend on paths

For a complex function, the integral

$$\int_{\gamma} f(z) dz$$

depends on both the function and the path γ . A parametrization $z = \gamma(t)$ gives

$$\int_{\gamma} f(z) dz = \int_a^b f(\gamma(t)) \gamma'(t) dt.$$

Thus complex integration is line integration in the plane with complex multiplication built in.

§115.9 Primitive functions and path independence

If $f = F'$ on a domain, then

$$\int_{\gamma} f(z) dz = F(\gamma(b)) - F(\gamma(a)).$$

So integrals over closed curves vanish. Conversely, on nice domains, vanishing around closed curves is closely related to the existence of a primitive. This is the beginning of Cauchy's theorem.

§115.10 Why singularities matter

The function $1/z$ has no primitive on $\mathbb{C} \setminus \{0\}$. Around the unit circle,

$$\int_{|z|=1} \frac{1}{z} dz = 2\pi i.$$

The nonzero integral detects the hole at the origin. This is the geometric heart of residue theory.

§115.11 A bridge outward

The important insight is that complex integration is simultaneously analytic and topological. Analytic regularity says holomorphic functions have local antiderivative-like behavior. Topology says singularities and holes can obstruct cancellation. The integral of $1/z$ around the unit circle is the prototype: the local function looks harmless away from 0, but the loop winds around a missing point and records $2\pi i$.

§115.12 Contour integrals as integrals of one-forms

A complex contour integral

$$\int_{\gamma} f(z) dz$$

integrates the complex one-form $f(z) dz$ along a parametrized curve. Under $z = \gamma(t)$, it becomes

$$\int_a^b f(\gamma(t)) \gamma'(t) dt.$$

This is exactly the pullback rule for one-forms.

§115.13 Cauchy theorem and exactness

If f has a primitive F , then

$$\int_{\gamma} f(z) dz = F(\gamma(b)) - F(\gamma(a)).$$

Closed-loop integrals vanish. Cauchy–Goursat says holomorphic functions have this local exactness behavior on suitable domains.

§115.14 Residues as obstruction

For $f(z) = 1/z$, the integral around the unit circle is

$$\int_{|z|=1} \frac{dz}{z} = 2\pi i.$$

There is no global primitive on $\mathbb{C} \setminus \{0\}$. The residue measures the obstruction created by the puncture.

§115.15 Residues, topology, and exactness

Complex integration differs from real integration because holomorphicity, topology, and orientation interact. Contour integrals are one-form integrals; Cauchy theory is exactness; residues measure the failure of exactness around singularities.

Numerical Series, Power Series, Taylor Series, and Improper Integrals

N-20260616

§116.1 Why these topics belong together

The long Markdown note moves from numerical series to function series, then to power series and improper integrals. The common theme is convergence: when does an infinite summation or an infinite accumulation define a finite object, and when may we manipulate it like a finite expression?

§116.2 Numerical series

A numerical series is

$$\sum_{n=1}^{\infty} a_n, \quad s_N = \sum_{n=1}^N a_n.$$

It converges if the partial sums s_N converge. The first necessary test is

$$a_n \rightarrow 0.$$

If this fails, the series diverges. If it holds, nothing is decided.

The geometric series is the base model:

$$\sum_{n=0}^{\infty} ar^n = \frac{a}{1-r}, \quad |r| < 1.$$

The harmonic series shows why terms tending to zero is insufficient:

$$\sum_{n=1}^{\infty} \frac{1}{n} = \infty.$$

A standard proof groups terms:

$$1 + \frac{1}{2} + \left(\frac{1}{3} + \frac{1}{4}\right) + \left(\frac{1}{5} + \cdots + \frac{1}{8}\right) + \cdots.$$

Each block after the first contributes at least $1/2$.

§116.3 Positive-term series

For $a_n \geq 0$, convergence is controlled by size.

Comparison test: if $0 \leq a_n \leq b_n$ and $\sum b_n$ converges, then $\sum a_n$ converges. If $a_n \geq b_n \geq 0$ and $\sum b_n$ diverges, then $\sum a_n$ diverges.

Limit comparison: if

$$\lim_{n \rightarrow \infty} \frac{a_n}{b_n} = c, \quad 0 < c < \infty,$$

then $\sum a_n$ and $\sum b_n$ have the same convergence behavior.

The p -series is the main comparison model:

$$\sum_{n=1}^{\infty} \frac{1}{n^p}$$

converges iff $p > 1$.

The examples in the original note include

$$\sum \sin \frac{1}{n}, \quad \sum \ln \left(1 + \frac{1}{n^2}\right), \quad \sum (\sqrt{n+1} - \sqrt{n}).$$

The first behaves like $\sum 1/n$ and diverges. The second behaves like $\sum 1/n^2$ and converges. The third behaves like $\sum 1/(2\sqrt{n})$ and diverges.

§116.4 Ratio and root tests

The ratio test is suited to factorials and exponentials:

$$L = \lim_{n \rightarrow \infty} \left| \frac{a_{n+1}}{a_n} \right|.$$

If $L < 1$, the series converges absolutely; if $L > 1$, it diverges; if $L = 1$, the test is inconclusive.

The root test is suited to expressions of the form $(\dots)^n$:

$$L = \limsup_{n \rightarrow \infty} \sqrt[n]{|a_n|}.$$

Again $L < 1$ gives absolute convergence, $L > 1$ gives divergence, and $L = 1$ is inconclusive.

Typical original examples include $\sum a^n/n^s$, factorials in the numerator or denominator, and powers with variable bases. The rule is not to memorize which test belongs to which problem, but to inspect which part of a_n changes fastest.

§116.5 Alternating and arbitrary-sign series

Absolute convergence means

$$\sum |a_n| < \infty.$$

Conditional convergence means $\sum a_n$ converges but $\sum |a_n|$ diverges.

The alternating-series test says that if $b_n \downarrow 0$, then

$$\sum_{n=1}^{\infty} (-1)^{n-1} b_n$$

converges. The alternating harmonic series is the standard example.

Dirichlet's test: if partial sums $A_N = \sum_{n=1}^N a_n$ are bounded and $b_n \downarrow 0$, then $\sum a_n b_n$ converges. This handles examples such as

$$\sum_{n=1}^{\infty} \frac{\cos n}{n}.$$

Abel's test is a neighboring result where one factor has a convergent series and the other is monotone bounded.

§116.6 Function series and uniform convergence

A function series is

$$\sum_{n=1}^{\infty} u_n(x).$$

Pointwise convergence asks for each fixed x . Uniform convergence asks whether the tail becomes small simultaneously for all x in a set:

$$\sup_{x \in E} |s_N(x) - s(x)| \rightarrow 0.$$

Uniform convergence is what allows continuity and integration to pass to the limit. Without it, a limit of continuous functions may fail to be continuous.

The Weierstrass M-test says that if

$$|u_n(x)| \leq M_n, \quad \sum M_n < \infty,$$

then $\sum u_n$ converges uniformly and absolutely.

§116.7 Power series

A power series centered at x_0 is

$$\sum_{n=0}^{\infty} a_n(x - x_0)^n.$$

There exists a radius of convergence R such that the series converges absolutely for $|x - x_0| < R$ and diverges for $|x - x_0| > R$. Endpoints require separate testing.

A common formula is

$$R = \lim_{n \rightarrow \infty} \left| \frac{a_n}{a_{n+1}} \right|,$$

when the limit exists. Missing powers require care: if the series contains $(x - x_0)^{2n}$, one should apply the ratio/root test directly rather than blindly reading a coefficient sequence as though every power were present.

Inside the interval of convergence, power series may be differentiated and integrated term by term:

$$\left(\sum a_n(x - x_0)^n \right)' = \sum n a_n(x - x_0)^{n-1}.$$

The radius of convergence remains the same, though endpoints may change.

§116.8 Taylor series

The Taylor series at a is

$$\sum_{n=0}^{\infty} \frac{f^{(n)}(a)}{n!} (x - a)^n.$$

At $a = 0$, it is a Maclaurin series. Standard expansions include

$$e^x = \sum_{n=0}^{\infty} \frac{x^n}{n!}, \quad \sin x = \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n+1}}{(2n+1)!}, \quad \cos x = \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n}}{(2n)!},$$

$$\frac{1}{1-x} = \sum_{n=0}^{\infty} x^n, \quad |x| < 1.$$

Many exam examples are indirect expansions: rewrite the function in terms of known series, substitute, integrate, differentiate, or multiply series.

§116.9 Improper integrals

Improper integrals appear when the interval is infinite or the integrand is unbounded:

$$\int_a^\infty f(x) dx = \lim_{R \rightarrow \infty} \int_a^R f(x) dx,$$

$$\int_a^b f(x) dx = \lim_{t \rightarrow b^-} \int_a^t f(x) dx$$

if f is singular at b .

The two fundamental p -templates are

$$\int_1^\infty \frac{dx}{x^p} \quad \text{converges iff } p > 1,$$

$$\int_0^1 \frac{dx}{x^p} \quad \text{converges iff } p < 1.$$

Logarithmic refinements include

$$\int_e^\infty \frac{dx}{x(\ln x)^p},$$

which converges iff $p > 1$.

For sign-changing improper integrals, absolute convergence means $\int |f| < \infty$. Conditional convergence often comes from oscillation, as in Dirichlet-type integrals.

§116.10 Unified viewpoint

Series and improper integrals ask the same question: can infinitely many small contributions accumulate to a finite value? Positive objects are controlled by size; sign-changing objects are controlled by cancellation plus size. Power series add a new idea: convergence depends on the point x , producing a domain where infinite algebraic manipulation becomes legitimate.

§116.11 Uniform convergence as norm convergence

For a function series $\sum u_n(x)$, uniform convergence on E is convergence in the sup norm:

$$\|s_N - s\|_{\infty, E} = \sup_{x \in E} |s_N(x) - s(x)| \rightarrow 0.$$

The Weierstrass M-test is simply comparison in this norm. If $|u_n(x)| \leq M_n$ and $\sum M_n < \infty$, then the tails satisfy

$$\sup_x \left| \sum_{n=N}^\infty u_n(x) \right| \leq \sum_{n=N}^\infty M_n \rightarrow 0.$$

Thus the elementary test is really completeness of a normed function space at work.

§116.12 Analytic functions and convergence radius

A power series

$$\sum a_n(z - z_0)^n$$

defines a holomorphic function inside its disk of convergence. The radius is controlled by the nearest complex singularity, not merely by real-variable behavior. This explains why real Taylor series sometimes stop converging at a point where the real function looks harmless: the obstruction may live off the real axis.

For example,

$$\frac{1}{1+x^2}$$

has real Taylor series at 0

$$1 - x^2 + x^4 - \dots,$$

with radius 1, because the complex singularities are at $x = \pm i$.

§116.13 Dominated convergence and improper integrals

Improper integrals and parameter-dependent integrals become cleaner in measure language. If $f_n \rightarrow f$ pointwise and $|f_n| \leq g$ with g integrable, dominated convergence gives

$$\int f_n \rightarrow \int f.$$

This justifies many limiting operations that elementary calculus handles case by case.

A useful contrast is between absolute convergence and conditional cancellation. The integral

$$\int_1^\infty \frac{\sin x}{x} dx$$

converges conditionally by oscillation, but

$$\int_1^\infty \left| \frac{\sin x}{x} \right| dx$$

diverges. This is the integral analogue of conditional series.

§116.14 Tauberian flavor

Power series also encode sequences. Given a_n , define

$$F(r) = \sum_{n=0}^{\infty} a_n r^n, \quad 0 < r < 1.$$

Abelian theorems say that ordinary convergence of $\sum a_n$ often implies limiting behavior of $F(r)$ as $r \rightarrow 1^-$. Tauberian theorems ask for converses under additional hypotheses. This is the research-level continuation of the elementary question: how does convergence of coefficients relate to behavior of the generating function near the boundary?

Fourier Series and Orthogonal Expansion of Functions

N-20260617

§117.1 Why Fourier series is a separate theme

The Fourier part of the original long note deserves its own chapter. It is not merely another convergence test. It is the beginning of representing functions by an orthogonal basis, which is the infinite-dimensional analogue of decomposing a vector in an orthonormal basis.

§117.2 Orthogonality of trigonometric functions

On $[-\pi, \pi]$,

$$\int_{-\pi}^{\pi} \cos mx \cos nx \, dx = 0 \quad (m \neq n),$$
$$\int_{-\pi}^{\pi} \sin mx \sin nx \, dx = 0 \quad (m \neq n), \quad \int_{-\pi}^{\pi} \sin mx \cos nx \, dx = 0.$$

Also,

$$\int_{-\pi}^{\pi} \cos^2 nx \, dx = \int_{-\pi}^{\pi} \sin^2 nx \, dx = \pi.$$

This is why the Fourier coefficients are projections.

§117.3 Fourier coefficients

For a 2π -periodic function,

$$f(x) \sim \frac{a_0}{2} + \sum_{n=1}^{\infty} (a_n \cos nx + b_n \sin nx),$$

where

$$a_n = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \cos nx \, dx, \quad b_n = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \sin nx \, dx.$$

For a_0 ,

$$a_0 = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \, dx.$$

§117.4 Dirichlet convergence theorem

Theorem 117.4.1 (Dirichlet convergence theorem)

If a periodic function is piecewise smooth, its Fourier series converges at each point to the average of the left and right limits.

Under standard piecewise smooth hypotheses, the Fourier series converges at x to

$$\frac{f(x+0) + f(x-0)}{2}.$$

At a point of continuity, this equals $f(x)$. At a jump, it equals the midpoint of the jump. This is why Fourier series may represent a function correctly except at discontinuities while still showing endpoint averaging.

§117.5 Even and odd functions

If f is even, then all sine coefficients vanish and the Fourier series is a cosine series. If f is odd, all cosine coefficients vanish and the series is a sine series. This is not just a computational shortcut; it says symmetry determines which basis vectors can appear.

For example, $f(x) = x$ on $[-\pi, \pi]$ is odd, so only sine terms occur. The standard result is

$$x \sim 2 \sum_{n=1}^{\infty} (-1)^{n+1} \frac{\sin nx}{n}.$$

For $f(x) = |x|$, only cosine terms occur, and evaluating the series at special points gives sums of reciprocal squares over odd integers.

§117.6 Half-range expansions

On $[0, l]$, a function may be extended oddly or evenly to $[-l, l]$. Odd extension gives a sine series; even extension gives a cosine series. The choice should match the boundary behavior one wants. In boundary-value problems for PDEs, sine series naturally encode zero boundary values, while cosine series encode even reflection or Neumann-type behavior.

For period $2l$, the basis becomes

$$\cos \frac{n\pi x}{l}, \quad \sin \frac{n\pi x}{l}.$$

The coefficients are

$$a_n = \frac{1}{l} \int_{-l}^l f(x) \cos \frac{n\pi x}{l} dx, \quad b_n = \frac{1}{l} \int_{-l}^l f(x) \sin \frac{n\pi x}{l} dx.$$

§117.7 Concrete coefficient models

The standard examples are valuable only when they are computed, not merely named.

§117.7.i The sawtooth function $f(x) = x$

On $[-\pi, \pi]$, the function $f(x) = x$ is odd, so only sine terms occur. The coefficient is

$$b_n = \frac{1}{\pi} \int_{-\pi}^{\pi} x \sin nx dx = \frac{2}{\pi} \int_0^{\pi} x \sin nx dx.$$

Integration by parts gives

$$\int_0^{\pi} x \sin nx dx = -\frac{\pi(-1)^n}{n},$$

so

$$b_n = 2 \frac{(-1)^{n+1}}{n}.$$

Hence

$$x \sim 2 \sum_{n=1}^{\infty} (-1)^{n+1} \frac{\sin nx}{n}.$$

At $x = \pi$, the periodic extension jumps from π to $-\pi$, so the Fourier series converges to 0, not to π .

§117.7.ii The even function $f(x) = |x|$

Since $|x|$ is even, only cosine terms occur. One has

$$a_0 = \frac{1}{\pi} \int_{-\pi}^{\pi} |x| dx = \pi,$$

and for $n \geq 1$,

$$a_n = \frac{2}{\pi} \int_0^{\pi} x \cos nx dx = \frac{2}{\pi} \frac{(-1)^n - 1}{n^2}.$$

Thus only odd n contribute:

$$|x| \sim \frac{\pi}{2} - \frac{4}{\pi} \sum_{k=0}^{\infty} \frac{\cos(2k+1)x}{(2k+1)^2}.$$

Putting $x = 0$ gives

$$0 = \frac{\pi}{2} - \frac{4}{\pi} \sum_{k=0}^{\infty} \frac{1}{(2k+1)^2},$$

so

$$\sum_{k=0}^{\infty} \frac{1}{(2k+1)^2} = \frac{\pi^2}{8}.$$

This is a typical Fourier-series trick: evaluate a function expansion at a point where the convergence value is known.

§117.7.iii A square wave

For the 2π -periodic square wave equal to 1 on $(0, \pi)$ and -1 on $(-\pi, 0)$, the function is odd, so

$$b_n = \frac{2}{\pi} \int_0^{\pi} \sin nx dx = \frac{2}{\pi} \frac{1 - (-1)^n}{n}.$$

Thus

$$f(x) \sim \frac{4}{\pi} \left(\sin x + \frac{\sin 3x}{3} + \frac{\sin 5x}{5} + \cdots \right).$$

At the jump points the series converges to 0, the midpoint of the jump.

§117.8 Dirichlet kernel: what partial sums really do

The N -th Fourier partial sum is not a vague approximation; it is convolution with the Dirichlet kernel:

$$S_N f(x) = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x-t) D_N(t) dt,$$

where

$$D_N(t) = \sum_{n=-N}^N e^{int} = \frac{\sin((N+1/2)t)}{\sin(t/2)}.$$

The kernel has total mass 2π , but its L^1 -norm grows. This growth is one reason Fourier partial sums can behave badly near discontinuities.

The Dirichlet convergence theorem is therefore not magic: near a point x , the oscillatory kernel samples values of f from both sides. At a jump, the symmetric sampling yields the midpoint $(f(x+0) + f(x-0))/2$.

§117.9 Fejer means: averaging partial sums improves convergence

The Cesaro average of Fourier partial sums is

$$\sigma_N f = \frac{1}{N+1} \sum_{k=0}^N S_k f.$$

It is convolution with the Fejer kernel

$$F_N(t) = \frac{1}{N+1} \left(\frac{\sin((N+1)t/2)}{\sin(t/2)} \right)^2.$$

Unlike the Dirichlet kernel, $F_N \geq 0$ and has total mass 2π . This positivity makes Fejer's theorem work: for continuous periodic f , $\sigma_N f$ converges uniformly to f .

This is a concrete lesson about convergence methods. The same Fourier coefficients can be summed by raw partial sums or by averaged partial sums, and averaging changes the analytic behavior without changing the underlying frequency data.

§117.10 Hilbert-space viewpoint

The conceptual endpoint is that Fourier series is linear algebra in the Hilbert space $L^2[-\pi, \pi]$. The inner product is

$$\langle f, g \rangle = \int_{-\pi}^{\pi} f(x)g(x) dx.$$

The trigonometric system is an orthogonal basis in an appropriate sense. Fourier coefficients are coordinates, Parseval's identity is the Pythagorean theorem, and convergence questions ask which topology of function space is being used.

§117.11 Fourier series in Hilbert space

Let

$$L^2[-\pi, \pi]$$

be the space of square-integrable functions with inner product

$$\langle f, g \rangle = \int_{-\pi}^{\pi} f(x) \overline{g(x)} dx.$$

The functions

$$e_n(x) = \frac{1}{\sqrt{2\pi}} e^{inx}, \quad n \in \mathbb{Z},$$

form an orthonormal basis in the Hilbert-space sense. The Fourier coefficient is

$$\hat{f}(n) = \langle f, e_n \rangle.$$

Thus Fourier series is orthogonal projection onto an infinite orthonormal basis.

§117.12 Parseval identity

For $f \in L^2[-\pi, \pi]$, Parseval's identity says

$$\|f\|_2^2 = \sum_{n \in \mathbb{Z}} |\hat{f}(n)|^2.$$

This is the Pythagorean theorem in infinite dimensions. In real sine-cosine notation, it becomes a statement relating integrals of f^2 to sums of squares of Fourier coefficients.

The elementary coefficient computations in the note are therefore coordinates of a vector in Hilbert space.

§117.13 Gibbs phenomenon

At jump discontinuities, Fourier series converges to the midpoint of the jump, but partial sums overshoot near the discontinuity. The overshoot does not vanish in height as the number of terms grows; it only concentrates closer to the jump. This is the Gibbs phenomenon.

This explains why Fourier series can converge correctly pointwise while still showing visible oscillations near discontinuities. It is a warning that convergence mode matters: pointwise, uniform, and L^2 convergence behave differently.

§117.14 Sturm–Liouville and Fourier bases

The functions $\sin nx$ and $\cos nx$ are eigenfunctions of

$$-\frac{d^2}{dx^2}.$$

For example,

$$-\frac{d^2}{dx^2} \sin nx = n^2 \sin nx.$$

With suitable boundary conditions, this operator is self-adjoint, and its eigenfunctions form an orthogonal basis.

Fourier series is therefore one example of Sturm–Liouville expansion. Other boundary conditions and intervals produce other orthogonal systems.

§117.15 Heat and wave equations

For the heat equation on an interval,

$$u_t = u_{xx},$$

expand

$$u(x, t) = \sum_n a_n(t) \sin nx.$$

Then each coefficient satisfies

$$a'_n(t) = -n^2 a_n(t), \quad a_n(t) = a_n(0) e^{-n^2 t}.$$

High-frequency modes decay faster. For the wave equation,

$$u_{tt} = u_{xx},$$

Fourier modes oscillate in time.

Thus Fourier series is not only a representation of functions; it diagonalizes differential operators and solves PDEs.

§117.16 Convergence modes of Fourier series

Fourier series can converge in different senses. For $f \in L^2$, partial sums converge to f in the L^2 norm:

$$\|S_N f - f\|_2 \rightarrow 0.$$

For smoother functions, convergence may be pointwise or uniform. At jumps, Dirichlet's theorem gives midpoint convergence, while Gibbs oscillation persists near the discontinuity.

This distinction is essential. A Fourier series may be excellent for energy approximation while not uniformly accurate near jumps. Thus the elementary endpoint rules and half-range expansions should always be paired with a statement of the convergence mode.

§117.17 Smoothness is decay of coefficients

Fourier coefficients measure regularity. If f is smooth, repeated integration by parts gives decay. For example, in complex notation,

$$\widehat{f}(n) = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) e^{-inx} dx.$$

If f is C^1 and periodic with matching endpoint values, then

$$\widehat{f}(n) = \frac{1}{in} \frac{1}{2\pi} \int_{-\pi}^{\pi} f'(x) e^{-inx} dx,$$

so $|\widehat{f}(n)| = O(1/|n|)$ if f' is integrable. More derivatives give faster decay.

This explains the difference between the examples above. The square wave has jumps, so coefficients decay like $1/n$. The function $|x|$ is continuous but has a corner, so its nonzero coefficients decay like $1/n^2$. Smoothness in physical space becomes decay in frequency space.

§117.18 Sobolev norm from Fourier coefficients

On the circle, Sobolev spaces can be described by Fourier coefficients:

$$\|f\|_{H^s}^2 = \sum_{n \in \mathbb{Z}} (1 + n^2)^s |\hat{f}(n)|^2.$$

Large positive s demands fast coefficient decay and therefore smoothness. Negative s permits rough objects, even distributions.

This is the precise functional-analytic version of the heuristic: high frequencies measure oscillation and roughness. Fourier series turns differential regularity into weighted square summability.

§117.19 Fourier transform as the non-periodic limit

Fourier series decomposes periodic functions into discrete frequencies. On the real line, the Fourier transform decomposes non-periodic functions into a continuum of frequencies:

$$\hat{f}(\xi) = \int_{-\infty}^{\infty} f(x) e^{-ix\xi} dx.$$

The inverse transform reconstructs f under suitable hypotheses. The derivative becomes multiplication by frequency:

$$\hat{f}'(\xi) = i\xi \hat{f}(\xi).$$

This is why Fourier methods solve constant-coefficient differential equations: they diagonalize differentiation.

§117.20 Distributions and weak Fourier series

Some important objects, such as the Dirac comb

$$\sum_{k \in \mathbb{Z}} \delta(x - 2\pi k),$$

are not ordinary functions but distributions. Fourier analysis extends to them by duality against test functions. This viewpoint explains why formal series sometimes work beyond classical convergence: they are meaningful in a weak or distributional sense.

First-Order and Higher-Order Differential Equations

N-20260618

§118.1 Why the differential-equation part is one note

The last part of the original long Markdown note covers first-order equations and higher-order linear equations. These belong together because the central idea is the same: transform an equation for an unknown function into a form that can be integrated, linearized, or solved as an operator equation.

§118.2 Basic concepts

A differential equation relates an unknown function to its derivatives. The order is the highest derivative appearing. A general first-order equation has the form

$$F(x, y, y') = 0.$$

A solution is a function $y(x)$ satisfying the equation on an interval. A general solution contains arbitrary constants; an initial condition selects a particular solution.

§118.3 Separable equations

A separable equation has

$$\frac{dy}{dx} = g(x)h(y).$$

If $h(y) \neq 0$, then

$$\frac{dy}{h(y)} = g(x) dx.$$

Integrating gives an implicit solution. One must also check equilibrium solutions coming from $h(y) = 0$, because division by $h(y)$ can lose them.

§118.4 Homogeneous first-order equations

A homogeneous first-order equation has the form

$$y' = F\left(\frac{y}{x}\right).$$

Set $v = y/x$, so $y = vx$ and

$$y' = v + xv'.$$

Then

$$v + x \frac{dv}{dx} = F(v),$$

which is separable:

$$\frac{dv}{F(v) - v} = \frac{dx}{x}.$$

Some equations become homogeneous after translating the origin. If the linear expressions in x, y meet at a point, shift that point to the origin.

§118.5 First-order linear equations

A first-order linear equation is

$$y' + p(x)y = q(x).$$

The integrating factor is

$$\mu(x) = e^{\int p(x) dx}.$$

Multiplying the equation by μ gives

$$(\mu y)' = \mu q.$$

Thus

$$y = e^{-\int p} \left(\int q e^{\int p} dx + C \right).$$

The integrating factor is not a trick; it conjugates the operator $D + p(x)$ into the ordinary derivative D .

§118.6 Bernoulli equations

A Bernoulli equation has

$$y' + p(x)y = q(x)y^n.$$

For $n \neq 0, 1$, set

$$u = y^{1-n}.$$

Then

$$u' = (1 - n)y^{-n}y',$$

and the equation becomes linear in u . Again, check special solutions lost by division.

§118.7 Exact equations

An equation

$$M(x, y) dx + N(x, y) dy = 0$$

is exact if there is a potential Φ with

$$d\Phi = M dx + N dy.$$

The test is

$$M_y = N_x$$

on a suitable simply connected region. Then the solution is

$$\Phi(x, y) = C.$$

If the equation is not exact, one may look for an integrating factor, often depending only on x or only on y .

§118.8 Reducible higher-order equations

If a second-order equation does not explicitly contain y , set $p = y'$ and reduce the order. If it does not explicitly contain x , set $p = p(y) = y'$, so

$$y'' = \frac{dp}{dx} = \frac{dp}{dy} \frac{dy}{dx} = p \frac{dp}{dy}.$$

These substitutions are not random: they use a missing variable as a symmetry.

§118.9 Second-order linear equations

A second-order linear equation has

$$y'' + p(x)y' + q(x)y = f(x).$$

The associated homogeneous equation has a solution space. If y_1, y_2 are independent homogeneous solutions, then

$$y_h = C_1 y_1 + C_2 y_2.$$

A general nonhomogeneous solution is

$$y = y_h + y_p.$$

This is the affine-space principle: one particular solution plus the kernel of a linear operator.

The Wronskian

$$W(y_1, y_2) = y_1 y_2' - y_1' y_2$$

measures independence. If it is nonzero at one point for solutions of a linear homogeneous equation, the two solutions form a fundamental system.

§118.10 Constant-coefficient homogeneous equations

For

$$y'' + ay' + by = 0,$$

try $y = e^{rx}$. The characteristic equation is

$$r^2 + ar + b = 0.$$

Distinct real roots r_1, r_2 give

$$y = C_1 e^{r_1 x} + C_2 e^{r_2 x}.$$

A repeated root r gives

$$y = (C_1 + C_2 x) e^{rx}.$$

Complex roots $\alpha \pm i\beta$ give

$$y = e^{\alpha x} (C_1 \cos \beta x + C_2 \sin \beta x).$$

§118.11 Constant-coefficient nonhomogeneous equations

For

$$L[y] = f(x), \quad L = D^2 + aD + b,$$

the method of undetermined coefficients works when f is built from polynomials, exponentials, sines, and cosines. If the trial function overlaps with a homogeneous solution, multiply by a sufficient power of x .

Variation of parameters works more generally. If y_1, y_2 are fundamental solutions, seek

$$y_p = u_1(x)y_1(x) + u_2(x)y_2(x).$$

The auxiliary condition $u_1'y_1 + u_2'y_2 = 0$ produces a solvable linear system for u_1', u_2' .

§118.12 Initial-value problems

An initial-value problem specifies

$$y(x_0) = a, \quad y'(x_0) = b.$$

For regular linear equations, existence and uniqueness say that these data select exactly one solution. This is conceptually important: a second-order equation needs two pieces of initial data because its solution space is two-dimensional.

§118.13 Functional-analytic viewpoint

Let

$$L[y] = y'' + p(x)y' + q(x)y.$$

Then a differential equation is an operator equation

$$L[y] = f.$$

The homogeneous solutions form $\ker L$. A nonhomogeneous solution set, if nonempty, is an affine translate of $\ker L$. Initial conditions are linear functionals that choose one element from that affine space.

This viewpoint explains why linear differential equations behave like linear algebra in function space. The derivative operator replaces a matrix; the Wronskian replaces a determinant; fundamental solutions replace a basis; Green functions, when introduced later, are kernels for inverse operators.

§118.14 Differential equations as operators

The elementary note classifies equations by solvable forms: separable, homogeneous, linear, Bernoulli, exact, constant-coefficient. The higher-level unity is that a differential equation is an equation in a function space. For example,

$$L[y] = f, \quad L = \frac{d^2}{dx^2} + p(x)\frac{d}{dx} + q(x).$$

The homogeneous solutions form the kernel $\ker L$. A nonhomogeneous solution set, when nonempty, is an affine translate

$$y_p + \ker L.$$

This is the same algebraic structure as a linear system $Ax = b$.

§118.15 Green functions

For a linear operator L , a Green function $G(x, \xi)$ is a kernel that represents the inverse problem. Formally,

$$L_x G(x, \xi) = \delta(x - \xi),$$

where δ is the Dirac delta distribution. Then a solution of

$$L[y] = f$$

can often be written as

$$y(x) = \int G(x, \xi) f(\xi) d\xi$$

with boundary conditions built into G .

This connects elementary particular-solution methods to integral operators. Undetermined coefficients works for special forcing terms; Green functions describe the response to arbitrary forcing as a superposition of responses to point sources.

§118.16 Semigroups and evolution

For first-order evolution equations, one often writes

$$\frac{du}{dt} = Au.$$

If A were a number, the solution would be $u(t) = e^{tA}u(0)$. In a Banach or Hilbert space, A may be an unbounded operator, and the solution is described by a semigroup

$$T(t) = e^{tA}, \quad T(t+s) = T(t)T(s).$$

The heat equation is the model example: the Laplacian generates a smoothing semigroup. The elementary exponential solution of constant-coefficient ODEs is the finite-dimensional shadow of this theory.

§118.17 Sturm–Liouville and Fourier expansion

A second-order operator of the form

$$L[y] = -(p(x)y')' + q(x)y$$

with suitable boundary conditions is often self-adjoint. Its eigenvalue problem

$$L[y] = \lambda w(x)y$$

produces orthogonal eigenfunctions. Fourier sine and cosine series are special cases of this principle.

Thus the elementary method of solving $y'' + \lambda y = 0$ under boundary conditions is not an isolated trick. It is the beginning of spectral theory for differential operators.

First-Order Partial Differential Equations and Characteristics

N-20260619

§119.1 What a first-order PDE asks for

A partial differential equation is an equation for an unknown function of several variables and its partial derivatives. A first-order scalar equation for $u = u(x, t)$ has the schematic form

$$F(x, t, u, u_x, u_t) = 0.$$

The simplest useful class is the linear transport equation

$$u_t + au_x = b(x, t),$$

where a is a constant. It is called a transport equation because it moves the initial profile without changing its shape, except for the forcing term b .

The geometric method for first-order equations is the method of characteristics. Instead of trying to understand u on the whole (x, t) -plane at once, one looks for curves along which the PDE becomes an ordinary differential equation.

Definition 119.1.1 (Characteristic curves). For the constant-coefficient transport equation

$$u_t + au_x = b(x, t),$$

a characteristic is a curve $x = x(t)$ satisfying

$$\frac{dx}{dt} = a.$$

Along such a curve,

$$\frac{d}{dt}u(x(t), t) = u_x(x(t), t)x'(t) + u_t(x(t), t) = au_x + u_t = b(x(t), t).$$

Thus a first-order PDE has been replaced by two ordinary differential equations:

$$\frac{dx}{dt} = a, \quad \frac{du}{dt} = b(x(t), t).$$

The first equation describes the characteristic line, and the second describes the value of the solution along that line.

§119.2 The unforced transport equation

Consider

$$u_t - cu_x = 0.$$

This is the sign convention used by the first-order wave example in the original notes. It has the form $u_t + au_x = 0$ with $a = -c$, so the characteristics satisfy

$$\frac{dx}{dt} = -c.$$

Hence

$$x(t) = x_0 - ct, \quad x_0 = x + ct.$$

Along each such line,

$$\frac{d}{dt}u(x(t), t) = u_t - cu_x = 0.$$

Therefore u is constant on the line $x + ct = x_0$.

Theorem 119.2.1 (Solution of the constant-speed transport equation)

Let f be a differentiable function. The solution of

$$u_t - cu_x = 0, \quad u(x, 0) = f(x)$$

is

$$u(x, t) = f(x + ct).$$

Proof. The initial point of the characteristic through (x, t) is $(x + ct, 0)$. Since u is constant along the characteristic,

$$u(x, t) = u(x + ct, 0) = f(x + ct).$$

A direct check gives

$$u_t = cf'(x + ct), \quad u_x = f'(x + ct),$$

so $u_t - cu_x = 0$. □




Figure 119.1: Characteristics for $u_t - cu_x = 0$. The solution is constant along each red line and the initial profile moves to the left as t increases.

§119.3 Example: translating a parabola

Let

$$u_t - cu_x = 0, \quad u(x, 0) = x^2 - 3.$$

The theorem gives

$$u(x, t) = (x + ct)^2 - 3.$$

For each fixed time t , the graph in the (x, u) -plane is the same parabola shifted left by ct .




Figure 119.2: Time slices of $u(x, t) = (x + ct)^2 - 3$, shown as translated copies of the initial profile.

§119.4 Forced transport

Now consider

$$u_t + au_x = b(x, t), \quad u(x, 0) = f(x).$$

The characteristic through (x, t) has the form

$$x(s) = x - a(t - s), \quad 0 \leq s \leq t,$$

because $x(t) = x$. Along this path,

$$\frac{d}{ds}u(x(s), s) = b(x(s), s).$$

Integrating from 0 to t gives the formula

$$u(x, t) = f(x - at) + \int_0^t b(x - a(t - s), s) ds.$$

Example 119.4.1 (A forced first-order wave equation)

Solve

$$u_t + 5u_x = e^{-t}, \quad u(x, 0) = e^{-x^2}.$$

Here $a = 5$, so

$$x(s) = x - 5(t - s).$$

The forcing depends only on time, and therefore

$$\int_0^t e^{-s} ds = 1 - e^{-t}.$$

The initial point is $x - 5t$. Hence

$$u(x, t) = e^{-(x-5t)^2} + 1 - e^{-t}.$$

Equivalently, if one writes the characteristic label as $x_0 = x - 5t$, then $u = f(x_0) + 1 - e^{-t}$. The sign of the last term is fixed by differentiating the proposed solution back into the PDE.

§119.5 What can go wrong

For first-order equations, the initial curve must meet the characteristic curves transversely. If the data are prescribed on a characteristic, then the PDE may fail to determine a unique solution. For constant-speed transport, prescribing $u(x, 0) = f(x)$ is safe because the initial line $t = 0$ intersects each characteristic once. This is the first place where PDEs differ from ordinary differential equations: the geometry of the set on which data are given is part of the problem.

§119.6 General linear first-order equations

The constant-coefficient equation is only the first model. A more general linear first-order equation is

$$a(x, t)u_x + b(x, t)u_t = c(x, t)u + d(x, t).$$

The characteristic curves now satisfy

$$\frac{dx}{ds} = a(x, t), \quad \frac{dt}{ds} = b(x, t).$$

Along such a curve, the value of u satisfies

$$\frac{du}{ds} = c(x(s), t(s))u + d(x(s), t(s)).$$

Thus even when the coefficients vary, the method is still the same: solve a characteristic system in the independent variables, then solve an ordinary differential equation for the dependent variable along the characteristic.

Example 119.6.1 (Variable speed transport)

Consider

$$u_t + xu_x = 0, \quad u(x, 0) = f(x).$$

The characteristics solve

$$\frac{dx}{dt} = x,$$

so $x(t) = x_0 e^t$. The initial footpoint through (x, t) is therefore $x_0 = x e^{-t}$. Since u is constant along characteristics,

$$u(x, t) = f(x e^{-t}).$$

A direct check gives

$$u_t = -x e^{-t} f'(x e^{-t}), \quad x u_x = x e^{-t} f'(x e^{-t}),$$

and hence $u_t + x u_x = 0$.

§119.7 Quasilinear equations

A first-order equation is called quasilinear if it is linear in the derivatives but the coefficients may depend on u :

$$a(x, t, u)u_x + b(x, t, u)u_t = c(x, t, u).$$

The characteristic system is now

$$\frac{dx}{ds} = a(x, t, u), \quad \frac{dt}{ds} = b(x, t, u), \quad \frac{du}{ds} = c(x, t, u).$$

This is already a nonlinear system of ordinary differential equations. In this sense, first-order nonlinear PDEs are not an entirely new object: they are families of coupled ODEs whose curves fill the plane.

Example 119.7.1 (Inviscid Burgers equation)

The equation

$$u_t + uu_x = 0$$

has characteristic speed u itself. Along a characteristic,

$$\frac{du}{dt} = u_t + uu_x = 0,$$

so u is constant. If $u(x, 0) = f(x)$, then the characteristic starting from x_0 has

$$x = x_0 + f(x_0)t, \quad u = f(x_0).$$

Thus the implicit solution is obtained from

$$x_0 = x - ut, \quad u = f(x - ut).$$

This formula is useful only as long as the map $x_0 \mapsto x_0 + f(x_0)t$ remains one-to-one.

§119.8 Crossing characteristics and shocks

For linear transport with constant speed, characteristics never cross. For nonlinear equations such as Burgers' equation, faster characteristics can overtake slower ones. When this happens, the classical solution becomes multi-valued if one continues the characteristic construction naively.

Differentiate the characteristic map for Burgers' equation:

$$x = x_0 + f(x_0)t.$$

The map stops being locally invertible when

$$\frac{\partial x}{\partial x_0} = 1 + f'(x_0)t = 0.$$

Thus if $f'(x_0) < 0$ somewhere, characteristic crossing occurs at time

$$t_* = -\frac{1}{\min f'(x_0)}$$

provided the minimum is negative. This is the first appearance of shock formation.

In an introductory PDE course, one often stops here and says that classical differentiable solutions break down. A more advanced treatment replaces them by weak solutions and adds an entropy condition to select the physically relevant shock. The point for this note is more elementary: first-order PDEs are controlled by the geometry of characteristic curves, and loss of uniqueness in that geometry is visible before one writes any distribution theory.

§119.9 Initial curves and well-posedness

Let initial data be prescribed on a curve Γ in the (x, t) -plane. For the characteristic method to determine a solution near Γ , the curve must not be tangent to the characteristic direction. For

$$a(x, t)u_x + b(x, t)u_t = c(x, t, u),$$

the characteristic direction in the (x, t) -plane is (a, b) . If Γ has tangent vector τ , then the non-characteristic condition is

$$\tau \nparallel (a, b).$$

Equivalently, the characteristic curves must cross the data curve rather than slide along it.

Example 119.9.1 (Bad data curve)

For $u_t + u_x = 0$, the characteristics are lines $x - t = \text{const.}$. If data are prescribed on the line $x - t = 0$, then the data are given along a characteristic. The equation says the solution is constant on that line, but it does not determine what happens on neighboring characteristic lines. Therefore the problem is not locally well-posed from that data curve.

§119.10 Conservation-law viewpoint

Some first-order equations are written in conservation form:

$$u_t + \partial_x F(u) = 0.$$

For smooth solutions, this is the same as

$$u_t + F'(u)u_x = 0.$$

The characteristic speed is $F'(u)$. Burgers' equation corresponds to $F(u) = u^2/2$. Conservation form becomes important when shocks appear, because it still has meaning after differentiability is lost.

Integrating over an interval $[a, b]$, we obtain

$$\frac{d}{dt} \int_a^b u(x, t) dx = F(u(a, t)) - F(u(b, t)).$$

Thus the amount of u in the interval changes only by flux through the boundary. This is the conceptual reason that conservation laws are the natural nonlinear continuation of transport equations.

The Second-Order Wave Equation and d'Alembert's Formula

N-20260620

§120.1 The equation and its physical meaning

The one-dimensional wave equation is

$$u_{tt} = c^2 u_{xx},$$

where $u(x, t)$ is the vertical displacement of a string and $c > 0$ is the wave speed. The equation says that acceleration is proportional to curvature. This is physically reasonable: translating the whole string upward produces no restoring force, and a straight slanted string has no net restoring force; only bending produces acceleration.




Figure 120.1: A vibrating string. The vertical displacement $u(x, t)$ is measured from the equilibrium line.

The equation is called hyperbolic because its highest-order part has the sign pattern of $\partial_t^2 - c^2 \partial_x^2$. This name is about the algebraic type of the differential operator, not about hyperbolas as curves.

§120.2 Factorization into transport operators

The wave operator factors as

$$\partial_t^2 - c^2 \partial_x^2 = (\partial_t + c \partial_x)(\partial_t - c \partial_x).$$

Since the two constant-coefficient operators commute, the wave equation can be read as a pair of first-order transport equations. Any function of $x + ct$ solves the equation, and any function of $x - ct$ also solves it.

Theorem 120.2.1 (d'Alembert form)

Every twice differentiable solution of

$$u_{tt} = c^2 u_{xx}$$

on a simply connected region in the (x, t) -plane can locally be written as

$$u(x, t) = F(x + ct) + G(x - ct).$$

The term $F(x + ct)$ travels left, and $G(x - ct)$ travels right.

Proof. Set

$$\xi = x + ct, \quad \eta = x - ct.$$

By the chain rule,

$$\partial_x = \partial_\xi + \partial_\eta, \quad \partial_t = c\partial_\xi - c\partial_\eta.$$

A direct calculation gives

$$u_{tt} - c^2 u_{xx} = -4c^2 u_{\xi\eta}.$$

Thus the wave equation is $u_{\xi\eta} = 0$. Integrating once in η gives $u_\xi = A(\xi)$, and integrating in ξ gives $u = F(\xi) + G(\eta)$. $\square$




Figure 120.2: The two characteristic families of the wave equation. A point receives information from both left-moving and right-moving characteristic directions.

§120.3 Initial-value problem on the whole line

Suppose

$$u(x, 0) = \phi(x), \quad u_t(x, 0) = \psi(x).$$

From $u = F(x + ct) + G(x - ct)$, the first condition gives

$$F(x) + G(x) = \phi(x).$$

The second condition gives

$$cF'(x) - cG'(x) = \psi(x).$$

Solving these two equations yields the d'Alembert formula

$$u(x, t) = \frac{1}{2} [\phi(x + ct) + \phi(x - ct)] + \frac{1}{2c} \int_{x-ct}^{x+ct} \psi(s) ds.$$

This formula shows finite propagation speed: the value at (x, t) depends only on the initial data in the interval $[x - ct, x + ct]$.

Example 120.3.1 (Zero initial velocity)

Let

$$\phi(x) = \frac{1}{1 + x^2}, \quad \psi(x) = 0.$$

Then

$$u(x, t) = \frac{1}{2} \left[\frac{1}{1 + (x + ct)^2} + \frac{1}{1 + (x - ct)^2} \right].$$

The initial bump splits into two half-sized bumps travelling in opposite directions.

§120.4 Finite strings and Fourier series

For a string fixed at $x = 0$ and $x = L$, the boundary conditions are

$$u(0, t) = u(L, t) = 0.$$

The spatial eigenfunctions satisfying these boundary conditions are

$$\sin \frac{n\pi x}{L}, \quad n = 1, 2, 3, \dots$$

Separation of variables gives normal modes

$$u_n(x, t) = \sin \frac{n\pi x}{L} \left(A_n \cos \frac{n\pi ct}{L} + B_n \sin \frac{n\pi ct}{L} \right).$$

A general solution is a Fourier sine series built from these modes:

$$u(x, t) = \sum_{n=1}^{\infty} \sin \frac{n\pi x}{L} \left(A_n \cos \frac{n\pi ct}{L} + B_n \sin \frac{n\pi ct}{L} \right).$$

This is the first bridge from the Fourier-series note to PDEs: boundary conditions choose the spatial basis, while the PDE determines the time oscillation of each coefficient.

§120.5 Energy conservation

The wave equation has a conserved energy. On the whole line, assume u and its derivatives decay sufficiently fast as $|x| \rightarrow \infty$. Define

$$E(t) = \frac{1}{2} \int_{-\infty}^{\infty} \left(u_t(x, t)^2 + c^2 u_x(x, t)^2 \right) dx.$$

Then

$$\frac{dE}{dt} = \int_{-\infty}^{\infty} (u_t u_{tt} + c^2 u_x u_{xt}) dx.$$

Using $u_{tt} = c^2 u_{xx}$,

$$\frac{dE}{dt} = c^2 \int_{-\infty}^{\infty} (u_t u_{xx} + u_x u_{xt}) dx.$$

Integrating the first term by parts gives

$$\int u_t u_{xx} dx = - \int u_{tx} u_x dx,$$

so the two terms cancel and $E'(t) = 0$.

Theorem 120.5.1 (Uniqueness by energy)

For the wave equation on the whole line with sufficiently decaying smooth solutions, the initial data $u(x, 0)$ and $u_t(x, 0)$ determine at most one solution.

Proof. If u and v are two solutions with the same initial data, set $w = u - v$. Then w solves the homogeneous wave equation with zero initial displacement and zero initial velocity. Its energy at $t = 0$ is zero. Since energy is conserved, $E_w(t) = 0$ for all t . Hence $w_t = 0$ and $w_x = 0$, so w is constant in space and time. The initial data give $w = 0$. $\square$

Energy is often more robust than an explicit formula. On domains where d'Alembert's formula is unavailable, energy still gives uniqueness and stability.

§120.6 Inhomogeneous wave equation and Duhamel's principle

Consider

$$u_{tt} - c^2 u_{xx} = F(x, t), \quad u(x, 0) = 0, \quad u_t(x, 0) = 0.$$

The forcing F can be viewed as continuously injecting infinitesimal initial velocities. Duhamel's principle gives

$$u(x, t) = \frac{1}{2c} \int_0^t \int_{x-c(t-s)}^{x+c(t-s)} F(\xi, s) d\xi ds.$$

This formula is the forced analogue of d'Alembert's formula. The triangular integration region is the past light cone of (x, t) .

Example 120.6.1 (Localized forcing)

If F is supported near a point (x_*, s_*) , then it can influence (x, t) only if

$$|x - x_*| \leq c(t - s_*), \quad s_* \leq t.$$

This is finite propagation speed in another form. Information does not instantly appear everywhere; it travels along the characteristic cone.

§120.7 Boundary conditions on a finite interval

For a finite string $0 < x < L$, the physical boundary condition depends on how the endpoint is constrained.

- Fixed endpoints give Dirichlet conditions:

$$u(0, t) = u(L, t) = 0.$$

- Free endpoints give Neumann conditions:

$$u_x(0, t) = u_x(L, t) = 0.$$

- A damped or elastic endpoint can lead to mixed conditions such as

$$\alpha u + \beta u_x = 0$$

at an endpoint.

The boundary condition selects the spatial eigenfunctions. Dirichlet conditions produce sine modes, while Neumann conditions produce cosine modes. This explains why half-range Fourier expansions are not only computational tricks: they encode boundary behavior.

§120.8 Deriving the Fourier coefficients

For the fixed string, write

$$u(x, t) = \sum_{n=1}^{\infty} q_n(t) \sin \frac{n\pi x}{L}.$$

Substituting into $u_{tt} = c^2 u_{xx}$ gives

$$q_n''(t) + \left(\frac{n\pi c}{L}\right)^2 q_n(t) = 0.$$

Hence

$$q_n(t) = A_n \cos \frac{n\pi ct}{L} + B_n \sin \frac{n\pi ct}{L}.$$

If

$$u(x, 0) = \phi(x), \quad u_t(x, 0) = \psi(x),$$

then

$$A_n = \frac{2}{L} \int_0^L \phi(x) \sin \frac{n\pi x}{L} dx,$$

and

$$\frac{n\pi c}{L} B_n = \frac{2}{L} \int_0^L \psi(x) \sin \frac{n\pi x}{L} dx.$$

Thus

$$B_n = \frac{2}{n\pi c} \int_0^L \psi(x) \sin \frac{n\pi x}{L} dx.$$

The role of orthogonality is explicit: it decouples one PDE into infinitely many harmonic oscillators.

§120.9 Standing waves and travelling waves

The functions

$$\sin \frac{n\pi x}{L} \cos \frac{n\pi ct}{L}$$

are standing waves. Their nodes are fixed in space. By contrast, d'Alembert's formula writes the solution as travelling waves. These are not contradictory descriptions. A standing wave is the superposition of a left-moving and a right-moving travelling wave of the same frequency.

For example,

$$\sin(kx) \cos(kct) = \frac{1}{2} [\sin(k(x + ct)) + \sin(k(x - ct))].$$

The Fourier description is best adapted to bounded intervals and boundary conditions; the travelling-wave description is best adapted to the whole line and propagation.

§120.10 The wave equation with damping

A simple damped string model is

$$u_{tt} + \gamma u_t = c^2 u_{xx}, \quad \gamma > 0.$$

The energy is no longer conserved. Multiplying by u_t and integrating gives

$$\frac{dE}{dt} = -\gamma \int u_t^2 dx \leq 0.$$

Thus damping removes kinetic energy. This example shows how a small change in the PDE changes the qualitative behavior: the undamped wave equation transports energy, while the damped wave equation dissipates it.

The Heat Equation, Diffusion, and Similarity Solutions

N-20260621

§121.1 Diffusion versus wave motion

The one-dimensional heat equation is

$$T_t = \kappa T_{xx}, \quad \kappa > 0.$$

Here $T(x, t)$ is temperature and κ is the thermal diffusivity. The equation says that the rate of change of temperature is proportional to curvature in the temperature profile. Peaks have negative curvature and decay; troughs have positive curvature and rise. Unlike the wave equation, diffusion is dissipative rather than oscillatory.

The heat equation is parabolic. One practical consequence is infinite propagation speed in the ideal model: a localized heat distribution immediately becomes nonzero everywhere, although it may be extremely small far away.

§121.2 Heat kernel on the whole line

The fundamental solution on the real line is the heat kernel

$$G(x, t) = \frac{1}{\sqrt{4\pi\kappa t}} \exp\left(-\frac{x^2}{4\kappa t}\right), \quad t > 0.$$

If $T(x, 0) = f(x)$, then the solution is the convolution

$$T(x, t) = (G(\cdot, t) * f)(x) = \int_{-\infty}^{\infty} G(x - s, t) f(s) ds.$$

The kernel has total mass one:

$$\int_{-\infty}^{\infty} G(x, t) dx = 1,$$

so heat is redistributed rather than created. The width of the Gaussian is of order $\sqrt{\kappa t}$, which is the natural diffusion length scale.

§121.3 Similarity solution on a semi-infinite rod

Consider a semi-infinite rod $x > 0$ initially at temperature zero, with the end held at temperature one after $t = 0$:

$$T(x, 0) = 0, \quad T(0, t) = 1, \quad T(x, t) \rightarrow 0 \quad (x \rightarrow \infty).$$

The heat equation is invariant under the scaling

$$x \mapsto \lambda x, \quad t \mapsto \lambda^2 t.$$

This suggests the similarity variable

$$\eta = \frac{x}{2\sqrt{\kappa t}}, \quad T(x, t) = \theta(\eta).$$

The derivatives are

$$T_x = \frac{1}{2\sqrt{\kappa t}}\theta'(\eta), \quad T_{xx} = \frac{1}{4\kappa t}\theta''(\eta), \quad T_t = -\frac{\eta}{2t}\theta'(\eta).$$

Substitution into $T_t = \kappa T_{xx}$ gives

$$\theta'' + 2\eta\theta' = 0.$$

Let $v = \theta'$. Then

$$v' + 2\eta v = 0, \quad (e^{\eta^2} v)' = 0,$$

so

$$\theta' = Ae^{-\eta^2}, \quad \theta = A \int_0^\eta e^{-s^2} ds + B.$$

Using

$$\operatorname{erf}(\eta) = \frac{2}{\sqrt{\pi}} \int_0^\eta e^{-s^2} ds, \quad \operatorname{erfc}(\eta) = 1 - \operatorname{erf}(\eta),$$

the boundary conditions give

$$T(x, t) = \operatorname{erfc}\left(\frac{x}{2\sqrt{\kappa t}}\right).$$




Figure 121.1: A qualitative picture of the semi-infinite rod solution. As time grows, the transition layer spreads over length scale $\sqrt{\kappa t}$.

§121.4 Fourier series for a finite rod

For a finite rod $0 < x < L$ with fixed zero temperature at the endpoints,

$$T(0, t) = T(L, t) = 0,$$

separation of variables gives

$$T(x, t) = X(x)A(t).$$

Substitution yields

$$\frac{A'}{\kappa A} = \frac{X''}{X} = -\lambda.$$

The boundary value problem

$$X'' + \lambda X = 0, \quad X(0) = X(L) = 0$$

has eigenfunctions

$$X_n(x) = \sin \frac{n\pi x}{L}, \quad \lambda_n = \left(\frac{n\pi}{L}\right)^2.$$

Thus each Fourier mode decays exponentially:

$$T(x, t) = \sum_{n=1}^{\infty} b_n e^{-\kappa(n\pi/L)^2 t} \sin \frac{n\pi x}{L}.$$

High-frequency modes have larger n^2 and therefore decay faster. This is the analytic meaning of smoothing by the heat equation.

§121.5 Maximum principle

A basic qualitative theorem for the heat equation is the maximum principle.

Theorem 121.5.1 (Maximum principle, informal form)

For a smooth solution of $T_t = \kappa T_{xx}$ on a closed space-time rectangle, the maximum and minimum occur either at the initial time or on the spatial boundary, unless the solution is constant.

This theorem matches the physical intuition that heat does not create a new interior temperature peak. It also explains why the heat equation behaves differently from the wave equation: diffusion forgets fine details, while waves transport and superpose them.

§121.6 Energy decay on an interval

For the heat equation on $0 < x < L$ with zero boundary values, define the thermal energy norm

$$\|T(\cdot, t)\|_{L^2}^2 = \int_0^L T(x, t)^2 dx.$$

Then

$$\frac{d}{dt} \frac{1}{2} \int_0^L T^2 dx = \int_0^L T T_t dx = \kappa \int_0^L T T_{xx} dx.$$

Integrating by parts and using $T(0, t) = T(L, t) = 0$,

$$\frac{d}{dt} \frac{1}{2} \int_0^L T^2 dx = -\kappa \int_0^L T_x^2 dx \leq 0.$$

Thus the L^2 -size of the temperature distribution decreases in time. This is the heat-equation analogue of the energy identity for the wave equation, but with a crucial difference: wave energy is conserved, while heat energy decays.

§121.7 The maximum principle with proof idea

The maximum principle says that heat cannot create a new interior maximum. Here is the local calculation behind it. Suppose T has a strict interior maximum at (x_0, t_0) with $t_0 > 0$. At that point,

$$T_x(x_0, t_0) = 0, \quad T_{xx}(x_0, t_0) \leq 0.$$

The heat equation gives

$$T_t(x_0, t_0) = \kappa T_{xx}(x_0, t_0) \leq 0.$$

Thus the value cannot be increasing at an interior spatial maximum. A rigorous proof perturbs the function by a small term such as $-\varepsilon t$ to rule out equality cases.

Theorem 121.7.1 (Uniqueness for the heat equation)

For the heat equation on a bounded interval with fixed boundary and initial data, there is at most one smooth solution.

Proof. Let T_1 and T_2 be two solutions and set $w = T_1 - T_2$. Then $w_t = \kappa w_{xx}$ with zero initial data and zero boundary data. By the maximum principle, $w \leq 0$. Applying the same argument to $-w$ gives $w \geq 0$. Hence $w = 0$. $\square$

§121.8 Smoothing and the heat semigroup

On the whole line, the solution operator is convolution with the heat kernel:

$$T(\cdot, t) = G_t * f.$$

The family G_t satisfies the semigroup identity

$$G_t * G_s = G_{t+s}.$$

Therefore the solution operators satisfy

$$S(t)S(s) = S(t+s), \quad S(t)f = G_t * f.$$

This is the infinite-dimensional analogue of the exponential law $e^{tA}e^{sA} = e^{(t+s)A}$. Formally,

$$S(t) = e^{t\kappa\partial_x^2}.$$

The heat equation is the basic example of a smoothing semigroup: even if f is rough, convolution with a Gaussian produces a smooth function for every $t > 0$.

§121.9 Fourier transform solution on the real line

Let

$$\hat{T}(\xi, t) = \int_{-\infty}^{\infty} T(x, t) e^{-ix\xi} dx.$$

Taking the Fourier transform of $T_t = \kappa T_{xx}$ gives

$$\partial_t \hat{T}(\xi, t) = -\kappa \xi^2 \hat{T}(\xi, t).$$

For each fixed frequency ξ , this is an ordinary differential equation:

$$\hat{T}(\xi, t) = e^{-\kappa \xi^2 t} \hat{f}(\xi).$$

High frequencies are multiplied by a rapidly decaying factor. This is the frequency-space explanation of smoothing.

Inverting the Fourier transform gives the heat-kernel formula. Thus the Gaussian kernel, the semigroup law, and the Fourier multiplier $e^{-\kappa \xi^2 t}$ are three descriptions of the same solution operator.

§121.10 Boundary conditions and conservation

Boundary conditions change what is conserved.

For insulated endpoints, one uses Neumann conditions:

$$T_x(0, t) = T_x(L, t) = 0.$$

Then the total heat is conserved:

$$\frac{d}{dt} \int_0^L T(x, t) dx = \kappa \int_0^L T_{xx} dx = \kappa [T_x]_0^L = 0.$$

For fixed-temperature endpoints, total heat may enter or leave through the boundary. The PDE describes local diffusion; the boundary condition describes the external thermal reservoir.

§121.11 Steady states and long-time behavior

A steady state satisfies $T_t = 0$, so

$$T_{xx} = 0.$$

On an interval, the steady states are linear functions

$$T_\infty(x) = Ax + B.$$

If the endpoints are held at temperatures $T(0, t) = A_0$ and $T(L, t) = A_L$, then the steady state is

$$T_\infty(x) = A_0 + \frac{A_L - A_0}{L}x.$$

The transient part $T - T_\infty$ satisfies zero boundary conditions and decays by the Fourier sine expansion. This separates the problem into two pieces: first solve the static boundary problem, then solve a homogeneous heat equation for the decaying error.

Example 121.11.1 (Nonzero boundary temperatures)

Let $T(0, t) = 0$, $T(L, t) = 1$, and $T(x, 0) = f(x)$. The steady state is x/L . Set

$$v(x, t) = T(x, t) - \frac{x}{L}.$$

Then v satisfies

$$v_t = \kappa v_{xx}, \quad v(0, t) = v(L, t) = 0,$$

with initial data $v(x, 0) = f(x) - x/L$. Hence

$$v(x, t) = \sum_{n=1}^{\infty} b_n e^{-\kappa(n\pi/L)^2 t} \sin \frac{n\pi x}{L},$$

where

$$b_n = \frac{2}{L} \int_0^L \left(f(x) - \frac{x}{L} \right) \sin \frac{n\pi x}{L} dx.$$

§121.12 Comparison with the wave equation

The heat and wave equations both use the second spatial derivative, but they use it in different ways:

$$T_t = \kappa T_{xx}, \quad u_{tt} = c^2 u_{xx}.$$

For heat, curvature determines velocity; for waves, curvature determines acceleration. This is why heat smooths and decays, while waves oscillate and propagate. Fourier modes make the difference transparent:

$$e^{-\kappa n^2 t} \quad \text{versus} \quad \cos(nct), \sin(nct).$$

The same spatial eigenfunctions occur, but their time factors encode the qualitative nature of the PDE.

XII

Classical Geometry

Part XII: Contents

N-20220615	791
122.1 What the note contains	791
122.2 Problem 1: two distances under a linear constraint	791
122.3 Problem 2: parabola and distance	791
122.4 Problem 3: packing unit circles	792
122.5 Problem 4: a symmetric Diophantine equation	792
122.6 The larger pattern	793
122.7 Coordinate choice and invariants	793
122.8 Packing and density	793
122.9 Symmetric Diophantine equations	793
122.10 Olympiad structure behind coordinates and symmetry	793
122.11 What an olympiad sheet trains	793
122.12 Coordinate geometry as optimization	794
122.13 Packing and convex bodies	794
122.14 Group actions and symmetry reduction	794
122.15 Convexity and supporting hyperplanes	795
122.16 Returning to the original coordinate methods	795
122.17 Coordinate choice as quotienting symmetry	795
122.18 Packing bounds as volume plus local obstruction	795
122.19 Equality cases as rigid configurations	795
 N-20220708	 797
123.1 Reflecting the orthocenter	797
123.2 The incenter–excenter lemma	797
123.3 Directed angles modulo 180°	798
123.4 Miquel point of a triangle	799
123.5 Directed angles reduce case distinctions	799
123.6 Miquel point proof pattern	799
123.7 Further context	799
123.8 Directed angles as projective bookkeeping	800
123.9 Reflections and isogonal structure	800
123.10 Miquel point as a compatibility condition	800
123.11 Angle invariants and cyclic compatibility	800
123.12 Why directed angles are worth learning	800
123.13 Directed angles and invariants	800
123.14 Circle pencils	801
123.15 Cross ratio	801
123.16 Moduli of configurations	801
123.17 Miquel configurations and spiral similarity	801
123.18 Inversion as normalization of circles	802
 N-20220715	 803
124.1 Conic sections	803
124.2 Latus rectum and tangent geometry	803
124.3 Rational and irrational density	804
124.4 Vector spaces	804
124.5 Dedekind cuts and order	804
124.6 Addition of Dedekind cuts	805
124.7 Additive inverses	805
124.8 Conics through focus and directrix	805
124.9 Dedekind cuts and addition	806
124.10 A compact bridge	806
124.11 Conics through affine and projective lenses	806
124.12 Density and completion	806
124.13 Dedekind addition as forced operation	806
124.14 Coordinates, conics, and the completed line	806

N-20220716	807
125.1 Line and ellipse	807
125.2 Tangents from a point to an ellipse	807
125.3 Line and hyperbola	808
125.4 Parabola tangent through a point	808
125.5 A parabola and a circle	809
125.6 An ellipse with prescribed quadrilateral area	809
125.7 The next layer	810
125.8 Tangency via discriminants	810
125.9 Polarity with respect to a conic	810
125.10 Area constraints and affine invariance	811
125.11 Tangency, polarity, and affine normalization	811
N-20220719	813
126.1 Why solid geometry feels different	813
126.2 Basic axioms of incidence	813
126.3 Relations between two lines	813
126.4 A line and a plane	814
126.5 Parallel planes and parallel lines	814
126.6 Perpendicularity	814
126.7 Projection theorem	814
126.8 A coordinate viewpoint	815
126.9 Skew lines	815
126.10 Planes by normal vectors	815
126.11 The organizing idea	815
126.12 Incidence geometry	816
126.13 Orthogonality as inner-product structure	816
126.14 Projection as least-distance decomposition	816
126.15 Affine incidence plus metric projection	816
126.16 Affine geometry of points, lines, and planes	816
126.17 Projective completion	817
126.18 Grassmannians	817
126.19 Exterior algebra and incidence	817
126.20 Returning to solid geometry	817
126.21 Incidence via ranks	817
126.22 Projective duality of points and planes	818
N-20220802	819
127.1 Angles and radians	819
127.2 Sine, cosine, and tangent	819
127.3 Special angles	820
127.4 Addition formulas	820
127.5 Double-angle and half-angle formulas	820
127.6 Graphs of trigonometric functions	820
127.7 Triangle formulas	821
127.8 Radians are not a convention only	821
127.9 Addition formulas as rotation composition	821
127.10 From technique to structure	821
127.11 Radians and the exponential map	822
127.12 Addition formulas from rotation matrices	822
127.13 Triangle formulas and invariant quantities	822
127.14 Trigonometry as rotation geometry	822
127.15 Trigonometry and rotations	822
127.16 Complex exponentials	823
127.17 Spherical trigonometry	823
127.18 Harmonic analysis	823
127.19 Trigonometric identities as algebra	823
127.20 Curvature and spherical trigonometry	824
N-20220911	825
128.1 Line formulas	825

128.2	Parallel and perpendicular lines	825
128.3	Direction and normal vectors	826
128.4	Angles, distances, and parallel lines	826
128.5	Circle equations	827
128.6	A fractional-linear range problem	827
128.7	A graph-slope counting problem	828
128.8	A moving-line maximum	828
128.9	Triangle coordinate problem	829
128.10	Coordinate tricks as changes of viewpoint	829
128.11	Power of a point	829
128.12	The structural moral	829
128.13	Lines as affine equations	830
128.14	Circles and quadratic forms	830
128.15	Fractional-linear transformations	830
128.16	Analytic geometry through algebraic invariants	830
128.17	Conics as quadratic forms	830
128.18	Affine transformations	831
128.19	Projective viewpoint on conics	831
128.20	Duality	831
128.21	Returning to analytic geometry	831
128.22	Distance formulas as dual norms	831
128.23	Conics and degenerations	832
N-20230119		833
129.1	Why this geometry chapter matters	833
129.2	Ceva's theorem	833
129.3	Menelaus' theorem	834
129.4	Ptolemy's theorem	834
129.5	Generalized Ptolemy inequality	834
129.6	Simson line	835
129.7	Euler line and Euler formula	835
129.8	Why the pattern matters	836
129.9	An invariant map for the theorem list	836
129.10	Barycentric coordinates	837
129.11	Projective transformations	837
129.12	Ptolemy and inversive geometry	837
129.13	Elliptic and hyperbolic analogues	837
129.14	Returning to the theorem list	837
129.15	Ceva and Menelaus as dual affine statements	838
129.16	Euler line and affine covariance	838
129.17	Simson line through projection	838
N-20230220		839
130.1	Motivation	839
130.2	Sine difference formula by geometry	839
130.3	Rotation-matrix proof	839
130.4	Euler formula proof	840
130.5	Tangent formula	840
130.6	The method behind the trick	840
130.7	Rotation proof as the structural proof	840
130.8	Euler formula and characters	840
130.9	Tangent formula as projective coordinate	841
130.10	Addition formulas as rotation composition	841
N-20230223		843
131.1	Correcting a trigonometric mistake	843
131.2	A trigonometric system	843
131.3	A rational-parametrization idea	843
131.4	A direct algebraic route	844
131.5	Multiple-angle formulas	844
131.6	The determinant recurrence	845

131.7	From technique to structure	845
131.8	Trigonometric systems as algebraic systems	845
131.9	Rational parametrization of the circle	845
131.10	Multiple-angle formulas and Chebyshev polynomials	845
131.11	Determinant recurrences	846
131.12	Trigonometry through algebraic parametrization	846
N-20240518		847
132.1	The annoyance of parallel lines	847
132.2	Homogeneous coordinates are ratios	847
132.3	A coordinate scene: where parallel lines meet	847
132.4	Duality: points and lines speak the same language	848
132.5	Transformations and what survives	848
132.6	A first invariant: the cross-ratio	848
132.7	Conics as one projective family	848
132.8	The projective plane as all lines through the origin	849
132.9	Desargues as a projective-flavored theorem	849
132.10	Projective transformations in coordinates	849
132.11	The conic as a disguised projective line	849
132.12	Final checkpoint	850
N-20250211		851
133.1	The organizing question	851
133.2	Affine transformations	851
133.3	Conics in matrix form	851
133.4	Projective geometry	853
133.5	Cross-ratio	853
133.6	Projective duality and conics	854
133.7	Inversion	854
133.8	Möbius transformations	854
133.9	Conic sections	855
133.10	The Erlangen program	855
133.11	A hierarchy of invariants	855
133.12	Quadratic forms and matrix language	855
133.13	Affine classification versus projective classification	856
133.14	Polar lines	856
133.15	The structural moral	856
133.16	How to choose the right geometry	856
133.17	Using the transformation group in a conic problem	857

A Russian-Style Early Olympiad Game: Coordinates, Packing, and Symmetry

N-20220615

§122.1 What the note contains

This note consists of a short mock or translated olympiad-style problem sheet. The cover page imitates an elementary mathematics olympiad booklet and states the usual rules: a multiple-choice test, one correct answer per problem, no calculators, and figures not necessarily drawn to scale. The mathematically relevant part is a collection of four problems about coordinate geometry, optimization, packing, and symmetric equations.

§122.2 Problem 1: two distances under a linear constraint

In the coordinate plane let

$$A = (0, -2), \quad B = (2, -2), \quad M = (x, y), \quad x + y = 1.$$

The problem asks for the maximum and minimum of

$$AM^2 + BM^2.$$

Since $y = 1 - x$,

$$\begin{aligned} AM^2 + BM^2 &= x^2 + (y + 2)^2 + (x - 2)^2 + (y + 2)^2 \\ &= x^2 + (3 - x)^2 + (x - 2)^2 + (3 - x)^2 \\ &= 4x^2 - 16x + 22 \\ &= 4(x - 2)^2 + 6. \end{aligned}$$

Thus the minimum is 6, achieved at $x = 2, y = -1$. There is no maximum on the whole line $x + y = 1$, because the expression tends to infinity as $|x| \rightarrow \infty$.

Remark 122.2.1 — The page asks for both maximum and minimum. This is a useful trap: on an unbounded line, a positive quadratic has a minimum but no maximum.

§122.3 Problem 2: parabola and distance

Let $M = (x, y)$ lie on the parabola

$$y = ax^2 + bx, \quad a < 0, \quad b > 0.$$

The quantity to maximize is

$$OM^2 = x^2 + y^2 = x^2 + (ax^2 + bx)^2.$$

The problem asks for the maximum on intervals such as

$$0 \leq x \leq -\frac{b}{a}$$

and a larger interval involving b . The first interval is exactly the interval between the two x -intercepts of the downward-opening parabola. Hence the endpoints lie on the x -axis, while the vertex lies at

$$x = -\frac{b}{2a}, \quad y = -\frac{b^2}{4a} > 0.$$

The derivative method is direct. Put

$$F(x) = x^2 + (ax^2 + bx)^2.$$

Then

$$F'(x) = 2x + 2(ax^2 + bx)(2ax + b).$$

The maximum on a closed interval occurs at an endpoint or at a critical point of this derivative.

§122.4 Problem 3: packing unit circles

The third problem asks about packing n circles of unit diameter into a circle of smallest possible diameter, and into a square of smallest possible side length. This is a circle-packing problem. For small n , the optimal configurations are concrete and geometric; for general n , exact formulas are difficult and often unknown.

A few elementary bounds are immediate. If n unit-diameter circles are packed into a container circle of diameter D , then area gives

$$n \cdot \pi \left(\frac{1}{2}\right)^2 \leq \pi \left(\frac{D}{2}\right)^2,$$

so

$$D \geq \sqrt{n}.$$

This bound is not usually sharp because circles leave gaps. For a square of side s , the area bound gives

$$n \cdot \frac{\pi}{4} \leq s^2, \quad s \geq \frac{\sqrt{n\pi}}{2}.$$

Again, the actual packing problem is subtler.

§122.5 Problem 4: a symmetric Diophantine equation

The final problem asks for the least integer solution of

$$\frac{a}{b+c} + \frac{b}{a+c} + \frac{c}{a+b} = 4.$$

This is a symmetric rational equation. A standard approach is to clear denominators:

$$a(a+c)(a+b) + b(b+c)(a+b) + c(b+c)(a+c) = 4(a+b)(a+c)(b+c).$$

After expansion and symmetrization, one may search for integer solutions or reduce by substituting elementary symmetric polynomials

$$s = a + b + c, \quad p = ab + bc + ca, \quad q = abc.$$

The problem belongs to the family of olympiad equations where symmetry is more important than direct expansion.

§122.6 The larger pattern

This short problem sheet moves between three worlds: quadratic optimization, geometric packing, and symmetric algebra. The common olympiad skill is to recognize the structure before calculating: a quadratic may have no maximum on an unbounded domain, a packing problem is governed first by area bounds and then by geometry, and a symmetric equation should be attacked through symmetric functions rather than arbitrary expansion.

§122.7 Coordinate choice and invariants

The olympiad-style coordinate problems are not about coordinates for their own sake. A good coordinate system preserves the invariant that matters and simplifies the constraint. For distance problems, squared distance often removes radicals:

$$d^2 = (x - a)^2 + (y - b)^2.$$

For symmetry problems, placing the origin at the center makes odd terms cancel.

§122.8 Packing and density

Packing unit circles is an elementary entrance to discrete geometry. A packing argument usually has two sides: a construction giving a lower bound and an area or distance obstruction giving an upper bound. In the plane, the optimal infinite equal-circle packing density is

$$\frac{\pi}{\sqrt{12}}.$$

Even when this theorem is far beyond the exercise, the same local logic appears: disks cannot overlap, so centers must be separated.

§122.9 Symmetric Diophantine equations

A symmetric equation in integers should first be rewritten in symmetric variables, such as

$$u = x + y, \quad v = xy.$$

This reduces the number of independent expressions and often exposes divisibility or parity obstructions.

§122.10 Olympiad structure behind coordinates and symmetry

The sheet trains the olympiad habit of choosing structure before computation: coordinates for constraints, area/distance for packing, symmetry for Diophantine equations, and equality cases for sharpness.

§122.11 What an olympiad sheet trains

This sheet looks elementary, but it trains several habits that reappear later. The distance problem asks us to parametrize a line and reduce geometry to a quadratic. The parabola problem asks us to decide whether a maximum exists on a specified interval. The packing

problem teaches the difference between lower bounds and exact solutions. The symmetric equation teaches that expanding blindly is rarely the best first move.

Thus the real notebook value is in the decision process: before computing, decide whether the domain is bounded, whether the object is symmetric, whether a crude bound is sharp, and whether a problem is actually asking for an exact classification.

§122.12 Coordinate geometry as optimization

The original olympiad-style coordinate problems often reduce a geometric claim to an algebraic inequality or an extremum. This is the beginning of optimization geometry. For example, placing points in coordinates turns distances into quadratic functions:

$$|x - y|^2 = (x - y) \cdot (x - y).$$

Then many geometric constraints become quadratic inequalities.

A modern viewpoint is to treat a geometric configuration as a feasible set and a desired length, angle, or area as an objective function. Convexity, symmetry, and compactness then explain why extrema exist and why equality cases are structured.

§122.13 Packing and convex bodies

Packing problems ask how many objects can fit without overlap. In the plane, disk packing leads to density questions. A convex body $K \subseteq \mathbb{R}^n$ is a compact convex set with nonempty interior. A packing by translates of K is a family

$$K + x_i$$

with disjoint interiors.

The density of a lattice packing is

$$\Delta = \frac{\text{vol}(K)}{\text{covol}(\Lambda)}$$

when the translates are indexed by a lattice Λ . Elementary circle or sphere packing problems are finite visible versions of questions about optimal densities.

This connects school geometry to research-level topics such as lattice packing, sphere packing, and the geometry of numbers.

§122.14 Group actions and symmetry reduction

Many coordinate solutions become shorter after using symmetry. A group action formalizes this. If a group G acts on a set X , then points in the same orbit

$$Gx = \{gx : g \in G\}$$

are equivalent under the symmetry.

For a symmetric geometric configuration, one can choose coordinates adapted to the group. For example, rotational symmetry suggests polar coordinates; reflection symmetry suggests placing axes on symmetry lines. The elementary trick of “put the origin at the center” is the coordinate form of quotienting by translations or rotations.

§122.15 Convexity and supporting hyperplanes

A line tangent to a convex curve or a plane tangent to a convex body is a supporting hyperplane. A hyperplane H supports a convex set K if K lies in one closed half-space determined by H and $H \cap K \neq \emptyset$.

Many olympiad inequalities about distances to lines or extreme points can be rephrased using supporting lines. This is the finite Euclidean beginning of convex analysis.

§122.16 Returning to the original coordinate methods

The elementary coordinate transformations, distance formulas, and symmetry placements should be kept exactly because they teach the practical side. The advanced layer explains what they are doing: reducing a geometric problem by symmetry, encoding constraints algebraically, and optimizing over a convex or compact feasible region.

§122.17 Coordinate choice as quotienting symmetry

When an olympiad solution places the origin at a center or aligns an axis with a symmetry line, it is not merely making the algebra shorter. It is quotienting out irrelevant degrees of freedom. Translations remove the absolute position of the configuration; rotations remove absolute direction; scaling removes absolute size when the problem is similarity-invariant.

A clean coordinate setup therefore identifies the true moduli of the configuration. For a triangle up to similarity, for instance, three vertex coordinates contain redundant information because translation, rotation, and scaling do not change the intrinsic shape. The remaining parameters describe shape, not placement.

§122.18 Packing bounds as volume plus local obstruction

Elementary packing arguments often begin with area or volume comparison. But pure volume comparison is rarely sharp because local geometry imposes additional constraints. In circle packing, density depends on how neighboring circles can touch; the hexagonal packing is optimal in the plane.

The higher-dimensional sphere-packing problem asks for the maximal density of non-overlapping equal balls in $\mathbb{R}^n$. In special dimensions, such as 8 and 24, exceptional lattices provide optimal packings. Thus a simple finite packing exercise belongs to a long research line connecting geometry, lattices, modular forms, and optimization.

§122.19 Equality cases as rigid configurations

In geometry problems, equality cases often correspond to symmetric or rigid configurations. If a bound is proved by Cauchy, AM–GM, or convexity, equality conditions translate back into geometric constraints: equal lengths, parallel directions, tangency, or regularity. The advanced habit is to track equality throughout the proof, because it reveals the extremal geometry.

Directed Angles, Reflections, and the Miquel Point

N-20220708

§123.1 Reflecting the orthocenter

Let H be the orthocenter of triangle ABC . Let X be the reflection of H over BC , and let Y be the reflection of H over the midpoint of BC . The lemma in the note states:

- (i) X lies on the circumcircle (ABC) ;
- (ii) AY is a diameter of (ABC) .

For the first claim, reflecting across BC preserves angles with BC . Since $BH \perp AC$ and $CH \perp AB$, angle chasing gives

$$\angle BXC = \angle BAC.$$

Thus A, B, C, X are concyclic.

For the second claim, the reflection of H through the midpoint of BC has the property that

$$\angle ABY = 90^\circ,$$

so AY is a diameter of the circumcircle by Thales' theorem.




Figure 123.1: Schematic reconstruction of the orthocenter reflection lemma.

§123.2 The incenter–excenter lemma

Let ABC have incenter I . Ray AI meets the circumcircle again at L . Let I_A be the reflection of I over L . Then:

- (i) I, B, C, I_A lie on a circle with diameter II_A and center L ;
- (ii) in particular,

$$LI = LB = LC = LI_A;$$

(iii) the rays BI_A and CI_A bisect the exterior angles of triangle ABC .

The geometry is governed by the well-known fact that the midpoint of the arc BC not containing A is equidistant from B, C, I . Reflecting I across this point gives the corresponding excenter configuration.

§123.3 Directed angles modulo 180°

The note then explains why ordinary angles create too many cases in cyclic geometry. The cure is to use directed angles modulo 180° . I write

$$\angle ABC$$

for the directed angle from line BA to line BC , modulo 180° .

The basic rules are:

- (i) Oblivion: $\angle APA = 0$.
- (ii) Anti-reflexivity: $\angle ABC = -\angle CBA$.
- (iii) Replacement: if A, B, C are collinear, then

$$\angle PBA = \angle PBC.$$

- (iv) Right angles: if $AP \perp BP$, then

$$\angle APB = 90^\circ.$$

- (v) Addition:

$$\angle APB + \angle BPC = \angle APC.$$

- (vi) Triangle sum:

$$\angle ABC + \angle BCA + \angle CAB = 0.$$

- (vii) Isosceles triangles:

$$AB = AC \iff \angle ACB = \angle CBA.$$

- (viii) Inscribed angle theorem:

if (ABC) has center P , then

$$\angle APB = 2\angle ACB.$$

- (ix) Parallel lines: if $AB \parallel CD$, then

$$\angle ABC + \angle BCD = 0.$$

One warning is important: because angles are modulo 180° , the statement

$$2\angle ABC = 2\angle XYZ$$

does not necessarily imply

$$\angle ABC = \angle XYZ.$$

Therefore one should not casually divide a directed angle congruence by 2.

§123.4 Miquel point of a triangle

Let D, E, F lie on the lines BC, CA, AB , respectively. The Miquel point theorem says that there is a point P lying on all three circles

$$(AEF), \quad (BFD), \quad (CDE).$$

Directed-angle proof. Let P be the second intersection of (AEF) and (BFD) . Since A, E, F, P are concyclic,

$$\angle EPF = \angle EAF.$$

Since B, F, D, P are concyclic,

$$\angle FPD = \angle FBD.$$

Adding gives

$$\angle EPD = \angle EAF + \angle FBD.$$

Using the collinearities E, A, C , F, A, B , and D, B, C , the right side becomes the angle condition for C, D, E, P to be cyclic. Hence $P \in (CDE)$. $\square$

§123.5 Directed angles reduce case distinctions

Using directed angles modulo 180° allows one to avoid separating many configurations. The cyclicity criterion becomes

$$A, B, C, D \text{ concyclic} \iff \angle ABC \equiv \angle ADC \pmod{180^\circ}.$$

This remains valid even when points move to extensions of lines. That is why olympiad geometry often uses directed angles once configurations become crowded.

§123.6 Miquel point proof pattern

For a triangle with points on the three sides, one proves the existence of the Miquel point by showing that two circumcircles meet at a point and then checking that this point lies on the third circle. The check is an angle chase: cyclicity in the first two circles gives angle equalities, and the directed-angle sum forces cyclicity in the third.

The proof is not about guessing the point. It is about showing that once two circles meet, the third circle is forced by angle consistency.

§123.7 Further context

The page is an excellent example of why directed angles are not merely notation. They turn many separate configuration cases into one proof. The deeper geometric message is inversive: circles and lines should be treated as one family of generalized circles, and Miquel configurations are natural incidence phenomena in that family.

§123.8 Directed angles as projective bookkeeping

Directed angles modulo 180° turn many cyclic and collinearity statements into algebraic angle sums. A statement like

$$\angle(AB, AC) = \angle(DB, DC)$$

encodes concyclicity without separating acute and obtuse cases.

This is a first step toward projective thinking: incidence and cross-ratio-type data matter more than metric length.

§123.9 Reflections and isogonal structure

Reflecting the orthocenter or a line across an angle bisector preserves angle data. In triangle geometry, many constructions are isogonal: two lines are mirror images with respect to an angle bisector. This explains why incenter/excenter lemmas naturally produce paired angle equalities.

§123.10 Miquel point as a compatibility condition

The Miquel point appears when several cyclic conditions are compatible. Each quadrilateral supplies a circle; directed angle chasing proves that the remaining circle passes through the same point. The deeper structure is consistency of local cyclic constraints.

§123.11 Angle invariants and cyclic compatibility

The geometry here is controlled by angle invariants. Directed angles remove cases, reflections preserve angle data, and the Miquel point records when cyclic conditions glue into one global configuration.

§123.12 Why directed angles are worth learning

The directed-angle pages are a turning point in geometry study. Before directed angles, one often proves the same cyclic statement four times because points move across arcs or outside the triangle. After directed angles modulo 180° , all these cases collapse into one algebraic identity of angles.

The Miquel point theorem is the perfect example. The theorem is not mainly about a mysterious point; it is about converting three cyclicity conditions into a chain of directed-angle equalities. Once the notation is accepted, the proof becomes almost linear.

§123.13 Directed angles and invariants

Directed angles modulo π are a way to make angle chasing stable under orientation and cyclic order. In projective geometry, angles themselves are not generally preserved, but incidence and cross ratios are. The bridge is inversive geometry: circles and angles are preserved by Möbius transformations.

A Möbius transformation has the form

$$z \mapsto \frac{az + b}{cz + d}, \quad ad - bc \neq 0.$$

It maps generalized circles to generalized circles and preserves oriented angles. Many Miquel-point and cyclic quadrilateral configurations simplify after a Möbius transformation sends one circle or line to a convenient position.

§123.14 Circle pencils

A pencil of circles is a one-parameter family of circles passing through the same two points, tangent at a point, or sharing limiting points. Algebraically, if two circles are given by equations $C_1 = 0$ and $C_2 = 0$, then the pencil is

$$C_1 + \lambda C_2 = 0.$$

Miquel configurations can often be interpreted as statements about intersecting pencils of circles.

This replaces repeated angle chasing by a structural object: the family of all circles compatible with part of the configuration.

§123.15 Cross ratio

For four points on a line or circle, the cross ratio is

$$[a, b; c, d] = \frac{(c - a)(d - b)}{(c - b)(d - a)}$$

in affine coordinates. It is invariant under projective transformations. Many cyclicity and concurrence statements can be reformulated by equality or reality of cross ratios.

The elementary theorem that equal angles imply concyclicity has a projective/inversive continuation: concyclicity and cross-ratio reality are two descriptions of the same configuration under suitable coordinates.

§123.16 Moduli of configurations

A geometric configuration can be studied up to a symmetry group. For instance, four distinct points on a projective line modulo projective transformations are classified by one parameter, the cross ratio. This is a simple example of a moduli problem: classify objects up to equivalence.

The Miquel point in the note is therefore not merely a special point. It belongs to a broader practice of understanding which features of a configuration are invariant under allowed transformations.

§123.17 Miquel configurations and spiral similarity

A Miquel point is often best understood through spiral similarity. If two segments AB and CD are seen under equal oriented angles from a point P , then P is the center of a spiral similarity sending one segment to the other. This combines a rotation and a dilation.

In complex coordinates, a direct similarity has the form

$$z \mapsto az + b, \quad a \in \mathbb{C}^\times.$$

A spiral similarity center can be solved algebraically from equations relating two pairs of points. Thus angle chasing and complex affine transformations are two languages for the same configuration.

§123.18 Inversion as normalization of circles

Inversion centered at O with radius r sends a point P to P' satisfying

$$OP \cdot OP' = r^2.$$

It sends circles and lines to circles and lines, preserving angles. If a configuration contains many circles through a common point, inversion at that point turns them into lines, often reducing Miquel-type statements to elementary parallel or intersection statements.

This is why inversive geometry is not an optional decoration. It is a normalization technique for circle configurations, just as choosing coordinates is a normalization technique for analytic geometry.

Conics and the Dedekind Construction of Addition

N-20220715

§124.1 Conic sections

The note starts with the geometry of conic sections. A cone cut by a plane produces the classical conics:

circle, ellipse, parabola, hyperbola.

An ellipse can be defined as the locus of points P for which the sum of distances to two foci is constant:

$$PF_1 + PF_2 = 2a.$$

With foci $F_1 = (-c, 0)$, $F_2 = (c, 0)$, this gives the standard equation

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1, \quad c^2 = a^2 - b^2.$$

The eccentricity is

$$e = \frac{c}{a}, \quad 0 < e < 1.$$

A hyperbola is defined by a constant difference of distances to the two foci:

$$|PF_1 - PF_2| = 2a,$$

with standard equations

$$\frac{x^2}{a^2} - \frac{y^2}{b^2} = 1$$

or

$$\frac{y^2}{a^2} - \frac{x^2}{b^2} = 1,$$

depending on the transverse axis.

A parabola is the locus of points equidistant from a focus and a directrix. Its standard equations are

$$y^2 = 2px, \quad y^2 = -2px, \quad x^2 = 2py, \quad x^2 = -2py.$$

For $y^2 = 2px$, the focus is $(p/2, 0)$, the directrix is $x = -p/2$, and the eccentricity is 1.

§124.2 Latus rectum and tangent geometry

For an ellipse, a vertical line through a focus intersects the ellipse in endpoints of a latus rectum. The length is

$$\frac{2b^2}{a}.$$

The tangent at $P = (x_0, y_0)$ to

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$$

is

$$\frac{xx_0}{a^2} + \frac{yy_0}{b^2} = 1.$$

The note also records the focus-directrix relation. For an ellipse with eccentricity $e = c/a$, the directrices are

$$x = \pm \frac{a^2}{c}.$$

The ratio

$$\frac{PF}{\text{dist}(P, \text{directrix})} = e$$

characterizes the conic.

§124.3 Rational and irrational density

The page then shifts to real analysis and proves that rational and irrational numbers are dense in $\mathbb{R}$. For rationals, given $x \in \mathbb{R}$ and $\varepsilon > 0$, choose $q \in \mathbb{N}$ with

$$\frac{1}{q} < \varepsilon.$$

Then choose an integer p such that

$$p \leq qx < p + 1.$$

Thus

$$\left| x - \frac{p}{q} \right| < \frac{1}{q} < \varepsilon.$$

For irrationals, take a rational approximation r to $x - \sqrt{2}/N$, then

$$r + \frac{\sqrt{2}}{N}$$

is irrational and can be made arbitrarily close to x .

§124.4 Vector spaces

The note includes the vector space axioms. Let F be a field and V a set equipped with addition and scalar multiplication:

$$+ : V \times V \rightarrow V, \quad \cdot : F \times V \rightarrow V.$$

The axioms are associativity and commutativity of addition, existence of 0, additive inverses, compatibility of scalar multiplication, distributivity, and $1v = v$. With the usual coordinate operations,

$$\mathbb{R}^n$$

is an $\mathbb{R}$ -linear space.

§124.5 Dedekind cuts and order

The note returns to Dedekind cuts. We denote by $\mathbb{R}$ the set of Dedekind cuts. For cuts X, Y , define an order by

$$X < Y \iff X \subset Y, \quad X \neq Y.$$

This order is total because if $X \not\subseteq Y$, choose $x \in X \setminus Y$; the downward-closed property of cuts forces $Y \subset X$.

§124.6 Addition of Dedekind cuts

Define

$$X + Y = \{x + y : x \in X, y \in Y\}.$$

The zero cut is

$$\bar{0} = \{x \in \mathbb{Q} : x < 0\}.$$

One must check that $X + Y$ is again a Dedekind cut:

- (i) it is nonempty and not all of $\mathbb{Q}$;
- (ii) it is downward closed;
- (iii) it has no greatest element.

The page gives the standard proof. For downward closure, if $z < x + y$, then $z = x + (z - x)$ and $z - x < y$, so $z - x \in Y$. For no greatest element, choose larger elements $x' > x$ in X and $y' > y$ in Y ; then $x' + y' > x + y$.

Associativity and commutativity follow directly:

$$(X + Y) + Z = X + (Y + Z), \quad X + Y = Y + X.$$

Order compatibility is subtler. One needs the lemma: if X is a cut and $n \in \mathbb{N}$, then there exist $x \in X$ and $x' \in X'$ such that

$$0 < x' - x < \frac{1}{n}.$$

This approximation lemma lets one prove

$$X < Y \Rightarrow X + Z < Y + Z.$$

§124.7 Additive inverses

For a cut X , define

$$-X = \{y - x' : y \in \bar{0}, x' \in X'\}.$$

Then $-X$ is the unique cut Z such that

$$X + Z = \bar{0}.$$

This reproduces the usual negative of a real number. In particular, if $p/q \in \mathbb{Q}$, then

$$-X_{p/q} = X_{-p/q}.$$

§124.8 Conics through focus and directrix

A conic can be defined by the ratio

$$e = \frac{\text{distance to focus}}{\text{distance to directrix}}.$$

If $e < 1$, the conic is an ellipse. If $e = 1$, it is a parabola. If $e > 1$, it is a hyperbola. This unifies the three classical curves by one parameter: eccentricity.

§124.9 Dedekind cuts and addition

The Dedekind construction treats a real number as a cut of rational numbers. Addition is then defined by adding representatives from cuts:

$$A + B = \{a + b : a \in A, b \in B\}.$$

The hard part is not writing the formula, but proving that it again satisfies the properties of a cut. This illustrates how arithmetic operations must be reconstructed after the real numbers are built.

§124.10 A compact bridge

This note is a striking mixture: classical conics and the construction of real addition by Dedekind cuts. The hidden connection is that both topics turn geometry into algebra. A conic becomes an equation and a focus-directrix ratio; a real number becomes a cut with algebraic operations.

§124.11 Conics through affine and projective lenses

A conic can be represented by a quadratic equation

$$ax^2 + bxy + cy^2 + dx + ey + f = 0.$$

Affine transformations preserve conic type up to degeneracy, while projective transformations unify ellipses, parabolas, and hyperbolas as sections of a cone.

The elementary focus-directrix and latus-rectum formulas are metric shadows of a more invariant projective object.

§124.12 Density and completion

The density of $\mathbb{Q}$ in $\mathbb{R}$ says every real interval contains rationals. Dedekind cuts use order rather than decimal expansions to construct real numbers. Addition of cuts must be defined so that order and completeness are preserved.

This is the structural reason the note can move from conics to real-number construction: analytic geometry needs a complete ordered field.

§124.13 Dedekind addition as forced operation

If A and B are lower cuts, their sum is

$$A + B = \{a + b : a \in A, b \in B\}.$$

The definition is not arbitrary. It is forced by wanting the embedding of rationals into cuts to preserve addition and order.

§124.14 Coordinates, conics, and the completed line

Conic geometry depends on the coordinate field. Dedekind cuts explain what the real line is, and projective/affine transformations explain which features of conics survive coordinate changes.

Conic Exercises: Lines, Ellipses, Parabolas, and Tangents

N-20220716

§125.1 Line and ellipse

The note continues the conic-section exercises. A line and an ellipse can have three possible positional relationships:

$$\begin{aligned}\Delta > 0 &: \text{two intersections,} \\ \Delta = 0 &: \text{tangent,} \\ \Delta < 0 &: \text{no real intersection.}\end{aligned}$$

This is obtained by substituting the line equation into the conic equation and examining the discriminant of the resulting quadratic.

For example, consider

$$y = 2x - 2, \quad \frac{x^2}{5} + \frac{y^2}{4} = 1.$$

Substitution gives

$$\frac{x^2}{5} + \frac{(2x - 2)^2}{4} = 1,$$

so

$$\frac{6}{5}x^2 - 2x = 0.$$

Thus

$$x = 0, \quad x = \frac{5}{3},$$

and the intersection points are

$$(0, -2), \quad \left(\frac{5}{3}, \frac{4}{3}\right).$$

§125.2 Tangents from a point to an ellipse

Let

$$C: \frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$$

and let $P = (0, 2)$. A line through P has equation

$$y = kx + 2.$$

Substituting into the ellipse gives a quadratic in x . Tangency occurs exactly when the discriminant is zero.

For the example on the page, this produces

$$k = \pm \frac{\sqrt{6}}{3},$$

so the two tangents are

$$y = \frac{\sqrt{6}}{3}x + 2, \quad y = -\frac{\sqrt{6}}{3}x + 2.$$

§125.3 Line and hyperbola

For the hyperbola

$$\frac{x^2}{2} - \frac{y^2}{2} = 1,$$

and the line

$$y = x - 1,$$

substitution gives

$$x^2 - (x - 1)^2 = 2.$$

Hence

$$2x - 1 = 2, \quad x = \frac{3}{2}, \quad y = \frac{1}{2}.$$

The line is not tangent; it meets the hyperbola at one point because the equation degenerates to a linear equation after substitution. This is a useful warning: for hyperbolas, “one intersection” need not mean tangency, because the line may be parallel to an asymptotic direction.

§125.4 Parabola tangent through a point

For the parabola

$$x^2 = \frac{1}{2}y,$$

we have

$$y = 2x^2.$$

A line through $P = (2, 6)$ may be written

$$y = kx + 6 - 2k.$$

Tangency means that

$$2x^2 = kx + 6 - 2k$$

has a repeated root. Thus

$$2x^2 - kx - 6 + 2k = 0,$$

and the discriminant condition is

$$k^2 - 4 \cdot 2(-6 + 2k) = 0.$$

Solving gives

$$k = 12, \quad k = -4.$$

The tangent lines are therefore

$$y = 12x - 18, \quad y = -4x + 14.$$

§125.5 A parabola and a circle

The second page considers the parabola

$$C : x^2 = 2py, \quad p > 0,$$

with focus F , and a circle

$$M : x^2 + (y + 4)^2 = 1.$$

The problem asks for p when the shortest distance between F and the circle is 4, and then asks for a maximal triangle area using two tangents from a point P on the circle.

Since the focus of $x^2 = 2py$ is

$$F = \left(0, \frac{p}{2}\right),$$

and the circle has center

$$O = (0, -4)$$

and radius 1, the distance from F to the circle is

$$FO - 1 = \left(\frac{p}{2} + 4\right) - 1 = \frac{p}{2} + 3.$$

Setting this equal to 4 gives

$$p = 2.$$

Hence the parabola is

$$x^2 = 4y.$$

For a tangent point $A = (x_1, x_1^2/4)$ on $x^2 = 4y$, the tangent has equation

$$xx_1 = 2(y + x_1^2/4),$$

or

$$y = \frac{x_1}{2}x - \frac{x_1^2}{4}.$$

Similarly for a second tangent point $B = (x_2, x_2^2/4)$. The two tangents meet at

$$P = \left(\frac{x_1 + x_2}{2}, \frac{x_1 x_2}{4}\right).$$

Using the condition that P lies on the circle gives an optimization problem for the area of triangle PAB . The computation in the page reduces it to a one-variable expression and obtains the maximum

$$[PAB]_{\max} = 20\sqrt{5}.$$

§125.6 An ellipse with prescribed quadrilateral area

The final exercise considers an ellipse

$$E : \frac{x^2}{a^2} + \frac{y^2}{b^2} = 1, \quad a > b > 0,$$

passing through $A = (0, -2)$. Four vertices form a quadrilateral of area $4\sqrt{5}$. Since $A = (0, -2)$ lies on the ellipse, $b = 2$. The area condition gives

$$2ab = 4\sqrt{5},$$

so

$$a = \sqrt{5}.$$

Therefore

$$E : \frac{x^2}{5} + \frac{y^2}{4} = 1.$$

For the line through $P = (0, -3)$

$$\ell : y = kx - 3,$$

intersecting the ellipse at B, C , the equation is

$$\frac{x^2}{5} + \frac{(kx - 3)^2}{4} = 1.$$

This gives a quadratic in x . The condition that the line meet the ellipse in two points is

$$k > 1 \quad \text{or} \quad k < -1.$$

The problem then studies where the chords meet the line $y = -3$ and imposes

$$|PM| + |PN| \leq 15.$$

The computation on the page reduces this to

$$|k| \leq 3.$$

Combining conditions gives

$$k \in [-3, -1) \cup (1, 3].$$

§125.7 The next layer

The page is a concentrated exercise in the algebraic method for conics: substitute a line into a conic, read the discriminant, and interpret the result geometrically. The only subtlety is that different conics behave differently: for ellipses, one intersection usually signals tangency; for hyperbolas, a line parallel to an asymptotic direction can produce a single finite intersection without being tangent.

§125.8 Tangency via discriminants

A line intersecting a conic usually produces a quadratic equation in the line parameter. Tangency occurs when the quadratic has a repeated root:

$$\Delta = 0.$$

This converts a geometric contact condition into algebra. It is the same multiplicity idea that appears in polynomial factorization.

§125.9 Polarity with respect to a conic

For a nondegenerate conic, a point outside the conic determines a polar line through its contact points with tangents. In matrix form, if the conic is

$$X^T A X = 0,$$

then the polar of P is

$$P^T A X = 0.$$

This gives a projective explanation of many tangent-chord relations.

§125.10 Area constraints and affine invariance

Some ellipse problems impose an area or quadrilateral condition. Area changes by the determinant of an affine map. Therefore one can often transform an ellipse to a circle, solve the symmetric problem, then transform back with the determinant factor.

§125.11 Tangency, polarity, and affine normalization

Line-conic exercises are controlled by three ideas: intersection multiplicity, polar duality, and affine normalization. The computations are coordinate expressions of these invariant geometric mechanisms.

Solid Geometry: Points, Lines, Planes, Parallelism, and Perpendicularity

N-20220719

§126.1 Why solid geometry feels different

The note starts with a warning: solid geometry is not just plane geometry with one more coordinate. The difficulty is that objects can miss each other in space. Two lines in the plane either meet or are parallel; in space, they may also be skew. A line can be parallel to a plane without being parallel to every line in that plane. These small distinctions are exactly why the axioms matter.

§126.2 Basic axioms of incidence

Axiom 126.2.1 (Plane incidence axioms)

Two distinct points determine a unique line, and two distinct lines intersect in at most one point. These axioms encode the basic incidence structure before metric notions such as distance and angle are introduced.

The fundamental incidence principles recorded in the note are:

- (i) Through three non-collinear points there is one and only one plane.
- (ii) If two points of a line lie in a plane, then the whole line lies in that plane.
- (iii) If two distinct planes have a common point, then their intersection is exactly one line through that point.

From these one derives several useful consequences:

- (i) Through a line and a point not on that line, there is one and only one plane.
- (ii) Through two intersecting lines, there is one and only one plane.
- (iii) Through two parallel lines, there is one and only one plane.

The proofs are good training in mathematical language. One repeatedly shows existence by choosing suitable points, and uniqueness by reducing to the uniqueness of a plane through three non-collinear points.

§126.3 Relations between two lines

In space, two distinct lines can be:

- (i) coplanar and intersecting;
- (ii) coplanar and parallel;

(iii) non-coplanar, called skew lines.

The first page underlines that skew lines are the genuinely new case. Drawings can be misleading: two projected lines may cross on paper even when the actual spatial lines do not intersect.

§126.4 A line and a plane

A line l and a plane α have three possible relations:

- (i) $l \subset \alpha$;
- (ii) l intersects α at exactly one point;
- (iii) $l \parallel \alpha$, meaning there is no common point.

The note stresses a common mistake: a line parallel to a plane is not necessarily parallel to every line in the plane. It is parallel to some line in the plane, but it may be skew to many others.

§126.5 Parallel planes and parallel lines

Two planes are parallel if they have no common point. The standard theorem is: if two intersecting lines in one plane are respectively parallel to two intersecting lines in another plane, then the two planes are parallel. Conversely, if two parallel planes are cut by a third plane, the intersection lines are parallel.

This explains many of the little diagrams in the note: the right move is to find the auxiliary plane that contains the relevant lines, then reduce the spatial statement to a planar one.

§126.6 Perpendicularity

A line l is perpendicular to a plane α if l is perpendicular to every line in α through the foot point. The useful test is stronger-looking but easier to check:

Theorem 126.6.1

If a line is perpendicular to two intersecting lines in a plane, then it is perpendicular to the plane.

Once a line is perpendicular to a plane, it is perpendicular to every line in that plane through the foot. This is the basic tool behind many distance and projection problems.

§126.7 Projection theorem

Theorem 126.7.1 (Projection theorem)

In an inner-product space, the closest point in a subspace W to a vector v is the vector $p \in W$ such that

$$v - p \perp W.$$

The note also touches the three-perpendicular theorem. In one common form: let $PO \perp \alpha$, let $AB \subset \alpha$, and let OQ be the projection of PQ onto α . Then perpendicularity of PQ to a line in space can be detected by perpendicularity of its projection in the plane. Informally, right angles survive projection when the projection is taken along a perpendicular direction.

§126.8 A coordinate viewpoint

Although the note is synthetic, it is useful to connect it with vectors. A plane in $\mathbb{R}^3$ can be written

$$ax + by + cz = d,$$

with normal vector

$$n = (a, b, c).$$

A line with direction vector v is perpendicular to the plane iff

$$v \parallel n.$$

A line lies in or is parallel to the plane iff

$$v \cdot n = 0.$$

Thus the synthetic tests are shadows of dot-product computations.

§126.9 Skew lines

In space, two lines may be neither intersecting nor parallel. Such lines are skew. This is the first major difference from plane geometry. A diagram can be misleading because a projection of skew lines may appear to intersect.

To prove two lines are skew, one usually shows they are not parallel and that the system of equations for intersection has no solution. In vector form, lines

$$p + tu, \quad q + sv$$

intersect iff there exist s, t such that $p + tu = q + sv$.

§126.10 Planes by normal vectors

A plane can be described by a point p_0 and a normal vector n :

$$n \cdot (x - p_0) = 0.$$

This equation is often more useful than a picture. Parallelism of planes becomes parallelism of normal vectors. Perpendicularity of a line and a plane becomes parallelism between the line direction and the plane normal.

§126.11 The organizing idea

Solid geometry trains an important habit: always ask what ambient plane contains the objects you are comparing. Many false arguments in space come from treating skew lines as if they were coplanar. Later, vector algebra solves this with span and dot products; projective geometry solves it with incidence; linear algebra solves it with subspaces and dimensions.

§126.12 Incidence geometry

Solid geometry begins with incidence: which points lie on which lines and planes. Many theorems are consequences of dimension counting. Two distinct points determine a line; three non-collinear points determine a plane.

This is the three-dimensional analogue of affine geometry, where parallelism and intersection are primary relations.

§126.13 Orthogonality as inner-product structure

Perpendicularity is not merely visual; it depends on an inner product. In coordinates,

$$u \perp v \iff u \cdot v = 0.$$

A line perpendicular to a plane is perpendicular to every direction vector in that plane. This is why normal vectors encode planes efficiently.

§126.14 Projection as least-distance decomposition

Orthogonal projection decomposes a vector into parallel and perpendicular parts:

$$v = \text{proj}_W v + (v - \text{proj}_W v).$$

The perpendicular part is the shortest error. This geometric fact later becomes least squares in linear algebra.

§126.15 Affine incidence plus metric projection

Solid geometry has two layers: affine incidence/parallelism and metric orthogonality/projection. Confusing the layers causes many errors. Lines and planes are affine objects; perpendicularity and distance require inner-product structure.

§126.16 Affine geometry of points, lines, and planes

Solid geometry often begins with points, lines, and planes. In affine geometry, a line through points p, q is

$$p + t(q - p), \quad t \in k,$$

and a plane through p with independent directions u, v is

$$p + su + tv.$$

Incidence questions become linear-dependence questions among direction vectors.

This explains the elementary tests: parallelism is dependence of direction vectors; coplanarity is rank at most two; intersection of planes is solving a linear system.

§126.17 Projective completion

Parallel lines do not meet in affine geometry, but they meet at points at infinity in projective geometry. Projective space $\mathbb{P}^n$ consists of one-dimensional subspaces of k^{n+1} . A point is written in homogeneous coordinates

$$[x_0 : x_1 : \cdots : x_n].$$

Affine space sits inside projective space as the chart $x_0 \neq 0$.

Adding points at infinity makes many incidence theorems cleaner because parallel and intersecting cases become one case.

§126.18 Grassmannians

The set of all k -dimensional linear subspaces of an n -dimensional vector space is the Grassmannian

$$Gr(k, n).$$

Lines through the origin in k^3 form $Gr(1, 3) = \mathbb{P}^2$. Planes through the origin form $Gr(2, 3)$.

Thus the collection of all possible directions of lines and planes is itself a geometric space. The elementary act of choosing a line direction is a point in a Grassmannian.

§126.19 Exterior algebra and incidence

A line direction spanned by u and a plane direction spanned by v, w can be encoded by wedge products:

$$u, \quad v \wedge w.$$

Incidence conditions become wedge equations. For example, u lies in the plane spanned by v, w iff

$$u \wedge v \wedge w = 0.$$

This gives a coordinate-free version of determinant tests for coplanarity.

§126.20 Returning to solid geometry

The original point-line-plane exercises are exactly the practical computations needed for affine and projective geometry. Direction vectors become points in Grassmannians; determinant tests become wedge-product incidence; parallel cases become ordinary intersections after projective completion.

§126.21 Incidence via ranks

In three-dimensional analytic geometry, incidence conditions are rank conditions. Points P_0, P_1, P_2, P_3 are coplanar iff the direction vectors

$$P_1 - P_0, \quad P_2 - P_0, \quad P_3 - P_0$$

are linearly dependent. Equivalently, the determinant formed by these three vectors is zero.

A line intersects a plane when the linear system consisting of the line equation and plane equation is solvable. A line is parallel to a plane when its direction vector lies in the kernel of the plane's normal functional. These are all linear-algebraic conditions.

§126.22 Projective duality of points and planes

In projective three-space, a plane is given by a linear equation

$$a_0X_0 + a_1X_1 + a_2X_2 + a_3X_3 = 0.$$

Thus planes correspond to covectors up to scale, i.e. points in the dual projective space. Point-plane incidence is the pairing

$$\ell(P) = 0.$$

This duality explains why many theorems about points have dual theorems about planes.

The elementary distinction between direction vectors and normal vectors is the affine shadow of vector-covector duality.

Trigonometric Functions, Radians, Identities, and Triangle Formulas

N-20220802

§127.1 Angles and radians

The note begins with angles. An angle has an initial side and a terminal side. Rotating counterclockwise gives a positive angle; rotating clockwise gives a negative angle. Angles differing by a full revolution represent the same terminal direction:

$$\theta \sim \theta + 2k\pi, \quad k \in \mathbb{Z}.$$

Radian measure is defined by arc length:

$$\theta = \frac{s}{r}.$$

Thus one radian is the angle subtending an arc whose length equals the radius. Since the circumference of a circle is $2\pi r$, one full turn is

$$2\pi \text{ radians} = 360^\circ.$$

Therefore

$$1 \text{ radian} \approx 57.3^\circ.$$

§127.2 Sine, cosine, and tangent

For a point $P = (x, y)$ on the terminal side of an angle α , put

$$r = |OP| = \sqrt{x^2 + y^2}.$$

Then

$$\sin \alpha = \frac{y}{r}, \quad \cos \alpha = \frac{x}{r}, \quad \tan \alpha = \frac{y}{x}.$$

On the unit circle, $r = 1$, so

$$P = (\cos \alpha, \sin \alpha).$$

The basic identity is

$$\sin^2 \alpha + \cos^2 \alpha = 1,$$

and

$$\tan \alpha = \frac{\sin \alpha}{\cos \alpha}.$$

The note's sign diagrams for $\sin, \cos, \tan$ are exactly quadrant information.

§127.3 Special angles

The page computes values such as

$$\sin \frac{2\pi}{3} = \frac{\sqrt{3}}{2}, \quad \cos \frac{2\pi}{3} = -\frac{1}{2}, \quad \tan \frac{2\pi}{3} = -\sqrt{3},$$

and

$$\sin \frac{3\pi}{2} = -1, \quad \cos \frac{3\pi}{2} = 0.$$

The method is not memorization alone: reduce to a reference angle and then use the quadrant sign.

§127.4 Addition formulas

The central identity is the cosine subtraction formula:

$$\cos(\alpha - \beta) = \cos \alpha \cos \beta + \sin \alpha \sin \beta.$$

From it one obtains

$$\cos(\alpha + \beta) = \cos \alpha \cos \beta - \sin \alpha \sin \beta,$$

$$\sin(\alpha + \beta) = \sin \alpha \cos \beta + \cos \alpha \sin \beta,$$

$$\sin(\alpha - \beta) = \sin \alpha \cos \beta - \cos \alpha \sin \beta.$$

The tangent formula follows by dividing sine by cosine:

$$\tan(\alpha + \beta) = \frac{\tan \alpha + \tan \beta}{1 - \tan \alpha \tan \beta}.$$

The note writes these not as isolated formulas but as relations among angles on the unit circle. That is the best way to remember them.

§127.5 Double-angle and half-angle formulas

From the addition formulas:

$$\sin 2\alpha = 2 \sin \alpha \cos \alpha,$$

$$\cos 2\alpha = \cos^2 \alpha - \sin^2 \alpha = 1 - 2 \sin^2 \alpha = 2 \cos^2 \alpha - 1.$$

The half-angle identities are

$$\sin^2 \frac{\alpha}{2} = \frac{1 - \cos \alpha}{2}, \quad \cos^2 \frac{\alpha}{2} = \frac{1 + \cos \alpha}{2}.$$

These formulas are especially useful in integrals and triangle geometry.

§127.6 Graphs of trigonometric functions

The note draws the graphs of $y = \sin x$, $y = \cos x$, and $y = \tan x$. The key facts are:

$$\sin(x + 2\pi) = \sin x, \quad \cos(x + 2\pi) = \cos x,$$

$$\tan(x + \pi) = \tan x.$$

Sine and cosine are bounded between -1 and 1 ; tangent has vertical asymptotes at

$$x = \frac{\pi}{2} + k\pi.$$

§127.7 Triangle formulas

For a triangle ABC with sides a, b, c opposite the corresponding angles, the law of sines is

$$\frac{a}{\sin A} = \frac{b}{\sin B} = \frac{c}{\sin C} = 2R,$$

where R is the circumradius.

The area formulas include

$$[ABC] = \frac{1}{2}bc \sin A$$

and Heron's formula

$$[ABC] = \sqrt{s(s-a)(s-b)(s-c)}, \quad s = \frac{a+b+c}{2}.$$

The note also records trigonometric transformations involving sums of angles in a triangle, especially $A + B + C = \pi$. For example,

$$\sin(A + B) = \sin C, \quad \cos(A + B) = -\cos C.$$

§127.8 Radians are not a convention only

The radian measure of an angle is defined by arc length divided by radius:

$$\theta = \frac{s}{r}.$$

This makes the derivative formulas clean:

$$\frac{d}{dx} \sin x = \cos x, \quad \frac{d}{dx} \cos x = -\sin x.$$

If degrees were used inside calculus, extra conversion constants would appear everywhere. Thus radians are the natural unit for analysis.

§127.9 Addition formulas as rotation composition

The matrix

$$R(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$$

rotates the plane by θ . Since rotations compose by adding angles,

$$R(\alpha + \beta) = R(\alpha)R(\beta).$$

Multiplying the matrices gives the sine and cosine addition formulas. This explains why those identities are not arbitrary memorized formulas but the algebra of rotations.

§127.10 From technique to structure

Trigonometry is not just a table of identities. It is the coordinate algebra of the circle. Radians measure angle by arc length; sine and cosine are coordinates on the unit circle; addition formulas express rotation composition; triangle formulas translate between angle data and side data. This is why trigonometry becomes natural again in complex numbers, Fourier analysis, rotations, and differential equations.

§127.11 Radians and the exponential map

Radians are natural because arc length on the unit circle equals angle. The map

$$t \mapsto e^{it}$$

wraps the real line around the circle, and addition of angles becomes multiplication of complex numbers.

This is the structural source of the addition formulas for sine and cosine.

§127.12 Addition formulas from rotation matrices

The rotation matrix

$$R(t) = \begin{pmatrix} \cos t & -\sin t \\ \sin t & \cos t \end{pmatrix}$$

satisfies

$$R(s+t) = R(s)R(t).$$

Multiplying matrices gives the sine and cosine addition formulas. Thus the identities express composition of rotations.

§127.13 Triangle formulas and invariant quantities

The sine rule and cosine rule relate side lengths and angles. The cosine rule is an inner-product identity:

$$c^2 = a^2 + b^2 - 2ab \cos C.$$

It is the law of dot products written in triangle language.

§127.14 Trigonometry as rotation geometry

Trigonometry is geometry of the circle and rotations. Radians come from arc length, addition formulas from group composition, graphs from periodicity, and triangle formulas from inner-product geometry.

§127.15 Trigonometry and rotations

The matrix

$$R_\theta = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$$

represents rotation by angle θ . The identity

$$R_\alpha R_\beta = R_{\alpha+\beta}$$

contains the addition formulas for sine and cosine.

Thus trigonometric identities are representation theory of the group $SO(2)$ in coordinates.

§127.16 Complex exponentials

Euler's formula

$$e^{i\theta} = \cos \theta + i \sin \theta$$

turns trigonometric identities into exponential algebra. For example,

$$e^{i(\alpha+\beta)} = e^{i\alpha} e^{i\beta}$$

gives both addition formulas at once.

The multiple-angle formulas are simply

$$(e^{i\theta})^n = e^{in\theta}.$$

Taking real and imaginary parts yields Chebyshev-type identities:

$$\cos(n\theta) = T_n(\cos \theta).$$

§127.17 Spherical trigonometry

On a sphere, triangles are bounded by great-circle arcs. Their angles and side lengths satisfy different laws from Euclidean triangles. The spherical law of cosines is

$$\cos c = \cos a \cos b + \sin a \sin b \cos C.$$

This reduces to the Euclidean law of cosines in a small-scale limit.

The elementary triangle formulas in the plane are therefore the zero-curvature case of trigonometry on constant-curvature geometries.

§127.18 Harmonic analysis

Trigonometric functions form an orthogonal system. On $[-\pi, \pi]$,

$$\int_{-\pi}^{\pi} \sin mx \sin nx \, dx = 0$$

for $m \neq n$, and similarly for cosines. This is why Fourier series can decompose functions into sine and cosine modes.

The original trigonometric identities are therefore also algebra of Fourier modes. Products of sines and cosines correspond to sums of frequencies.

§127.19 Trigonometric identities as algebra

The functions $e^{in\theta}$ form characters of the circle group. Products of trigonometric functions are products of characters, which add frequencies:

$$e^{im\theta} e^{in\theta} = e^{i(m+n)\theta}.$$

The product-to-sum formulas are just this identity rewritten in real and imaginary parts.

This explains why trigonometric simplification is often frequency bookkeeping. A complicated expression may be simplified by converting to exponentials, combining frequencies, and returning to real functions.

§127.20 Curvature and spherical trigonometry

In Euclidean geometry, the angle sum of a triangle is π . On a unit sphere, the spherical excess satisfies

$$A + B + C - \pi = \text{area}(\triangle).$$

Thus trigonometry detects curvature. Plane triangle formulas are the zero-curvature limit of spherical and hyperbolic formulas.

This provides a genuine high-level extension of elementary trigonometry: identities involving sides and angles are not merely computational; they reflect the geometry of the underlying space.

Analytic Geometry: Lines, Circles, and Coordinate Tricks

N-20220911

§128.1 Line formulas

This note is a formula sheet and worked exercise collection for analytic geometry. The first part summarizes equations of a line.

If a line has inclination angle α , then its slope is

$$k = \tan \alpha, \quad \alpha \neq \frac{\pi}{2}.$$

If it passes through two points

$$P_1 = (x_1, y_1), \quad P_2 = (x_2, y_2),$$

then

$$k = \frac{y_2 - y_1}{x_2 - x_1}.$$

The point-slope form is

$$y - y_1 = k(x - x_1),$$

the slope-intercept form is

$$y = kx + b,$$

the two-point form is

$$\frac{y - y_1}{y_2 - y_1} = \frac{x - x_1}{x_2 - x_1},$$

the intercept form is

$$\frac{x}{a} + \frac{y}{b} = 1,$$

and the general form is

$$Ax + By + C = 0.$$

The note's diagram of intercepts reminds us that the intercept form is convenient only when the line meets both coordinate axes away from the origin.

§128.2 Parallel and perpendicular lines

For

$$l_1 : y = k_1x + b_1, \quad l_2 : y = k_2x + b_2,$$

we have

$$l_1 \parallel l_2 \iff k_1 = k_2, \quad b_1 \neq b_2,$$

and

$$l_1 \perp l_2 \iff k_1 k_2 = -1.$$

For lines in general form,

$$l_1 : A_1x + B_1y + C_1 = 0, \quad l_2 : A_2x + B_2y + C_2 = 0,$$

parallelism is detected by proportional normal vectors:

$$\frac{A_1}{A_2} = \frac{B_1}{B_2} \neq \frac{C_1}{C_2},$$

and perpendicularity is

$$A_1 A_2 + B_1 B_2 = 0.$$

This is just the dot product of the normal vectors.

§128.3 Direction and normal vectors

For a line

$$Ax + By + C = 0,$$

a normal vector is

$$n = (A, B),$$

and a direction vector is

$$d = (-B, A).$$

Thus the vector form through $P_0 = (x_0, y_0)$ is

$$\frac{x - x_0}{u} = \frac{y - y_0}{v}$$

if $d = (u, v)$, and

$$a(x - x_0) + b(y - y_0) = 0$$

if $n = (a, b)$.

This vector viewpoint is the easiest way to remember the formulas. The line is the set of points whose displacement from P_0 is perpendicular to a fixed normal vector.

§128.4 Angles, distances, and parallel lines

The angle between two nonvertical lines is given by

$$\tan \theta = \left| \frac{k_1 - k_2}{1 + k_1 k_2} \right|.$$

Equivalently, using direction vectors d_1, d_2 ,

$$\cos \theta = \frac{|d_1 \cdot d_2|}{|d_1| |d_2|}.$$

The distance between points is

$$|AB| = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}.$$

The distance from a point $P(x_0, y_0)$ to a line

$$Ax + By + C = 0$$

is

$$d = \frac{|Ax_0 + By_0 + C|}{\sqrt{A^2 + B^2}}.$$

For parallel lines

$$Ax + By + C_1 = 0, \quad Ax + By + C_2 = 0,$$

the distance is

$$d = \frac{|C_1 - C_2|}{\sqrt{A^2 + B^2}}.$$

§128.5 Circle equations

A circle with center $C(a, b)$ and radius r has equation

$$(x - a)^2 + (y - b)^2 = r^2.$$

The general equation is

$$x^2 + y^2 + Dx + Ey + F = 0.$$

Completing the square gives

$$\left(x + \frac{D}{2}\right)^2 + \left(y + \frac{E}{2}\right)^2 = \frac{D^2 + E^2 - 4F}{4}.$$

Hence the center is

$$\left(-\frac{D}{2}, -\frac{E}{2}\right),$$

and it is a real circle exactly when

$$D^2 + E^2 - 4F > 0.$$

The handwritten note underlines this discriminant-like quantity because forgetting it leads to treating imaginary circles as real ones.

§128.6 A fractional-linear range problem

One example asks for the range of

$$\mu = \frac{x}{y - 3}$$

under

$$|x| + |y| \leq 1.$$

The idea is geometric. The equation

$$\mu = \frac{x}{y - 3}$$

can be rewritten as

$$x = \mu(y - 3).$$

For fixed μ , this is a line through the point $(0, 3)$. The range of μ is determined by which slopes make this line intersect the diamond

$$|x| + |y| \leq 1.$$

Definition 128.6.1 (Pencil of lines). A pencil of lines through a fixed point P is the family of all lines passing through P . Choosing a slope, or another one-dimensional parameter, gives a coordinate on most of this pencil.

Here the fixed point is $P = (0, 3)$. The parameter μ is not a random fraction; it is a coordinate on the pencil of lines through P . The range problem asks:

which lines in this pencil meet the diamond?

Because the diamond is convex and P lies outside it, the extreme values of μ occur at the two supporting lines from P to the diamond. These supporting lines touch the diamond at its upper-left and upper-right vertices:

$$(-1, 0), \quad (1, 0).$$

Their parameters are

$$\mu = \frac{-1}{0-3} = \frac{1}{3}, \quad \mu = \frac{1}{0-3} = -\frac{1}{3}.$$

Therefore

$$\mu \in \left[-\frac{1}{3}, \frac{1}{3}\right].$$

Remark 128.6.2 (Why this is not only a trick) — The algebraic fraction becomes transparent after dualizing the problem: points (x, y) in the diamond are replaced by lines through $P = (0, 3)$. The range is an interval because the diamond is convex, and the endpoints are supporting lines. This is a first glimpse of projective and convex dual thinking inside analytic geometry.

§128.7 A graph-slope counting problem

Another exercise asks for the number of x_i such that

$$\frac{f(x_i)}{x_i}$$

has a fixed value. Geometrically, the ratio $f(x)/x$ is the slope of the line from the origin to the point $(x, f(x))$. So the problem asks how many times a line through the origin can be tangent or secant to the drawn curve with a given slope. The answer is read from the graph, not from algebraic manipulation.

This is a useful change of viewpoint: ratios involving y/x often mean slopes of rays from the origin.

§128.8 A moving-line maximum

A later problem has moving lines through fixed points $A = (-3, 0)$ and $B = (1, 6)$, intersecting at P , and asks for the maximum of

$$|PA| \cdot |PB|.$$

The solution uses a right-angle configuration. Since the two moving lines are perpendicular, P lies on the circle with diameter AB . Then

$$|PA|^2 + |PB|^2 = |AB|^2.$$

Thus

$$|PA||PB| \leq \frac{|PA|^2 + |PB|^2}{2} = \frac{|AB|^2}{2}.$$

This is equality when $PA = PB$. With

$$|AB|^2 = (1 + 3)^2 + 6^2 = 52,$$

the maximum is

$$26.$$

§128.9 Triangle coordinate problem

The final worked example gives $A = (5, 1)$, the median line CM , and the perpendicular bisector BN , then asks for B and the equation of BC . The handwritten solution introduces unknown coordinates and uses two facts:

- (i) a median means a midpoint condition;
- (ii) a perpendicular bisector means equal distances or perpendicular slopes.

Coordinate geometry often works this way: translate each geometric phrase into an equation, then solve the resulting system.

§128.10 Coordinate tricks as changes of viewpoint

Analytic geometry often becomes simple after choosing better coordinates. A line is easiest when written in point-direction form,

$$p + tv,$$

while a circle is easiest when written by center and radius,

$$(x - a)^2 + (y - b)^2 = r^2.$$

Expanding too early hides the geometry. Completing the square recovers it.

§128.11 Power of a point

For a circle with center O and radius r , the power of a point P is

$$OP^2 - r^2.$$

If a line through P meets the circle at A, B , then

$$PA \cdot PB$$

is constant. This is a standard way to convert circle intersection geometry into algebraic length equations.

§128.12 The structural moral

Analytic geometry is not merely a bag of formulas. Each formula has a vector meaning: slope is direction, $Ax + By + C = 0$ encodes a normal vector, distance to a line is projection onto that normal vector, and circle equations are completed squares. The more one sees these meanings, the less one needs to memorize.

§128.13 Lines as affine equations

A line in the plane is an affine equation

$$ax + by + c = 0.$$

The vector (a, b) is normal to the line. Parallelism means proportional normal vectors; perpendicularity means the direction vector of one line is proportional to the normal vector of the other.

This turns many line problems into vector algebra.

§128.14 Circles and quadratic forms

A circle has equation

$$(x - a)^2 + (y - b)^2 = r^2.$$

Completing the square reveals the center and radius. More general second-degree equations are conics, and their classification depends on the quadratic part.

Thus circle problems are the simplest non-linear quadratic geometry.

§128.15 Fractional-linear transformations

A fractional-linear expression

$$y = \frac{ax + b}{cx + d}$$

is a Möbius transformation on the projective line when $ad - bc \neq 0$. Range problems for such expressions are often solved by inverting the equation and detecting excluded values.

§128.16 Analytic geometry through algebraic invariants

Analytic geometry translates geometric relations into algebraic invariants: normal vectors for lines, quadratic forms for circles/conics, and fractional-linear maps for rational range problems.

§128.17 Conics as quadratic forms

A conic in the plane has equation

$$ax^2 + 2bxy + cy^2 + dx + ey + f = 0.$$

The quadratic part is the symmetric matrix

$$Q = \begin{pmatrix} a & b \\ b & c \end{pmatrix}.$$

Orthogonal diagonalization of Q removes the xy term. Translation then attempts to remove linear terms. This is the linear-algebraic classification behind the elementary conic formulas.

The discriminant

$$b^2 - ac$$

distinguishes ellipse, parabola, and hyperbola types in the nondegenerate real case. In projective geometry, these distinctions are understood through rank and intersection with the line at infinity.

§128.18 Affine transformations

An affine transformation has the form

$$x \mapsto Ax + b, \quad A \in GL(n).$$

It sends lines to lines and preserves ratios along a line, but not necessarily distances or angles. Many analytic-geometry simplifications are affine changes of coordinates.

Under affine transformations, ellipses may be transformed into circles, and general nondegenerate conics may be brought to standard forms. The elementary completion-of-squares method is one coordinate expression of this broader equivalence.

§128.19 Projective viewpoint on conics

Using homogeneous coordinates $[X : Y : Z]$, a conic is given by a homogeneous quadratic equation

$$\begin{pmatrix} X & Y & Z \end{pmatrix} A \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = 0.$$

Projective transformations act by

$$A \mapsto P^T A P.$$

This unifies ellipses, parabolas, and hyperbolas over the complex projective plane; the affine differences come from how the conic meets the line at infinity $Z = 0$.

§128.20 Duality

In projective geometry, points and lines have a dual relationship. A nondegenerate conic identifies points with polar lines. If a conic is represented by a symmetric matrix A , the polar line of a point p is

$$p^T A x = 0.$$

This gives an advanced interpretation of tangents and chord relations.

§128.21 Returning to analytic geometry

The original formulas for lines, circles, and conics remain the computational core. The advanced layer says: lines are affine/projective objects, conics are quadratic forms, coordinate changes are group actions, and tangents/polars come from bilinear algebra.

§128.22 Distance formulas as dual norms

The distance from a point x_0 to a line

$$ax + by + c = 0$$

is

$$\frac{|ax_0 + by_0 + c|}{\sqrt{a^2 + b^2}}.$$

The numerator evaluates an affine functional; the denominator is the norm of its linear part. In normed-space language, distance to a hyperplane uses the dual norm of the defining functional.

For a hyperplane $\varphi(x) = c$ in an inner product space,

$$\text{dist}(x_0, H) = \frac{|\varphi(x_0) - c|}{\|\varphi\|_*}.$$

In Euclidean coordinates, $\|\varphi\|_*$ is the length of the normal vector.

§128.23 Conics and degenerations

A conic represented by a symmetric homogeneous matrix A may be nondegenerate or degenerate. Degenerate conics include pairs of lines, double lines, or points over the real affine plane. Algebraically, degeneracy corresponds to rank drop of A .

This gives a unified view of many analytic-geometry cases: a circle, ellipse, parabola, hyperbola, or pair of lines is controlled by rank, signature, and the line at infinity. Completing squares is a coordinate method for revealing those invariants.

Plane Geometry I: Ceva, Menelaus, Ptolemy, Simson Line, and Euler Line

N-20230119

§129.1 Why this geometry chapter matters

The note is a chapter-style plane geometry note. It does not merely list named theorems; it gives statements, pictures, proofs, and comments about when each theorem is useful. The main pattern is:

Ceva detects concurrence, Menelaus detects collinearity,
Ptolemy detects cyclicity.

Then Simson line and Euler line enter as richer configurations involving circles, projections, centers, and ratios.

§129.2 Ceva's theorem

Theorem 129.2.1 (Ceva's theorem)

For points on the sides of a triangle, the corresponding cevians are concurrent exactly when the signed product of the three side ratios is 1.

Let D, E, F lie on BC, CA, AB or their extensions in triangle ABC . Ceva's theorem says that AD, BE, CF are concurrent iff

$$\frac{BD}{DC} \cdot \frac{CE}{EA} \cdot \frac{AF}{FB} = 1$$

using directed lengths.

The proof in the note uses areas. Suppose the cevians meet at O . Since triangles with bases on the same line and the same altitude have areas proportional to their bases,

$$\frac{BD}{DC} = \frac{[ABD]}{[ACD]} = \frac{[OBD]}{[OCD]}.$$

Subtracting the smaller triangles gives

$$\frac{[ABO]}{[ACO]} = \frac{BD}{DC}.$$

Similarly,

$$\frac{[BCO]}{[BAO]} = \frac{CE}{EA}, \quad \frac{[CAO]}{[CBO]} = \frac{AF}{FB}.$$

Multiplying the three identities cancels all area factors and gives the Ceva product 1.

The inverse theorem is often the useful direction: if the product of ratios is 1, then the three cevians are concurrent. In olympiad geometry, this is the standard tool for proving three lines meet at one point.

§129.3 Menelaus' theorem

Theorem 129.3.1 (Menelaus' theorem)

For points on the three sidelines of a triangle, the points are collinear exactly when the signed product of the three side ratios is -1 .

Menelaus is the collinearity counterpart of Ceva. If a line meets the sidelines of triangle ABC at D, E, F , then

$$\frac{AE}{EC} \cdot \frac{CD}{DB} \cdot \frac{BF}{FA} = 1$$

again with directed lengths.

The proof in the original note draws a parallel through a vertex. For instance, draw through B a line parallel to the transversal and compare similar triangles. The ratios on the sides of the triangle become ratios on a pair of parallel lines. After substitution, the product becomes 1.

The memory rule is useful but not enough: Ceva and Menelaus look similar, so one must remember what they prove. Ceva: three cevians concurrent. Menelaus: three points collinear.

§129.4 Ptolemy's theorem

Theorem 129.4.1 (Ptolemy's theorem)

For a cyclic quadrilateral with side lengths a, b, c, d and diagonals e, f ,

$$ef = ac + bd.$$

For a cyclic quadrilateral $ABCD$, Ptolemy's theorem states

$$AC \cdot BD = AB \cdot CD + AD \cdot BC.$$

The note also records the converse: if a convex quadrilateral satisfies this equality, then it is cyclic.

A classical proof chooses a point on diagonal construction so that two pairs of triangles become similar. The products $AB \cdot CD$ and $AD \cdot BC$ then appear as pieces of $AC \cdot BD$. What matters is that cyclicity converts equal angles into similarity.

Ptolemy is metric: it relates lengths, not only ratios. That is why it is more rigid than Ceva or Menelaus.

§129.5 Generalized Ptolemy inequality

For any convex quadrilateral $ABCD$, the generalized Ptolemy inequality says

$$AC \cdot BD \leq AB \cdot CD + AD \cdot BC,$$

with equality iff the quadrilateral is cyclic.

The note gives a complex-plane proof. Assign complex numbers a, b, c, d to the vertices. The identity

$$(a - b)(c - d) + (a - d)(b - c) = (a - c)(b - d)$$

is exact. Taking moduli and applying the triangle inequality gives

$$|a - c||b - d| \leq |a - b||c - d| + |a - d||b - c|.$$

This is precisely Ptolemy's inequality. Equality in the triangle inequality corresponds to the relevant complex numbers having the same argument, which geometrically means the points are concyclic in the required order.

This proof is important because it changes the language of geometry: a quadrilateral becomes four complex numbers, and a metric inequality becomes the ordinary triangle inequality.

§129.6 Simson line

Simson's theorem says: if P lies on the circumcircle of triangle ABC , and perpendiculars are dropped from P to the three sides (or their extensions), then the three feet are collinear. The common line is the Simson line of P .

The proof in the note uses cyclic quadrilaterals. Let the feet be D, E, F . Because $PE \perp AB$ and $PF \perp AC$, quadrilaterals such as $AEPF$ are cyclic. Similar cyclic-angle relations show that the angle between two of the foot segments is supplementary to the angle between another pair. Hence the three feet lie on a line.

The inverse theorem is also true. The practical use is: when a problem contains many perpendicular feet from a point to the sides of a triangle, look for a circumcircle point and a Simson line.

§129.7 Euler line and Euler formula

For triangle ABC , let O be the circumcenter, G the centroid, and H the orthocenter. Euler's theorem says that O, G, H are collinear and

$$OG : GH = 1 : 2,$$

equivalently $GH = 2OG$. This common line is the Euler line.

One proof uses midpoints and similar triangles. Let D be the midpoint of BC . Since G is the centroid, $AG : GD = 2 : 1$. Using the facts that the circumcenter lies on perpendicular bisectors and the orthocenter lies on altitudes, one constructs similar triangles showing that the point on OH with the centroid ratio is exactly G .

The note also states Euler's formula relating the circumradius R , inradius r , and the distance $d = OI$ between circumcenter and incenter:

$$R^2 - d^2 = 2Rr.$$

This immediately gives Euler's inequality

$$R \geq 2r,$$

with equality iff the triangle is equilateral. The proof in the note uses the angle bisector through the incenter, intersections with the circumcircle, and similar triangles to derive a product relation. The key point is that the inequality is not a random radius estimate; it is the nonnegativity of d^2 .

§129.8 Why the pattern matters

This chapter is a small map of classical geometry. Ceva and Menelaus are ratio theorems, almost projective in flavor. Ptolemy is metric and cyclic. Simson line turns perpendicular projections into collinearity. Euler line organizes triangle centers, and Euler's formula links local radii with the global placement of centers. The proofs constantly switch languages: area, similarity, complex numbers, cyclic angles, and centers. That switching is the real skill of plane geometry.

§129.9 An invariant map for the theorem list

The theorems in this chapter are not all of the same type. Each belongs to a different invariance principle.

Theorem	Main structure	What is invariant or detected
Ceva	affine/projective incidence	concurrence of cevians through side ratios
Menelaus	affine/projective incidence	collinearity of side points through signed ratios
Ptolemy	cyclic metric geometry	cyclicity through a length identity
Simson line	cyclic projection geometry	collinearity from perpendicular projections
Euler line	mixed affine/metric geometry	alignment of centroid, circumcenter, orthocenter

Remark 129.9.1 (Why this classification matters) — A geometric theorem becomes easier to use when one knows which transformations preserve it. Incidence statements survive projective transformations. Ratio statements survive affine changes on a line or projective normalization after coordinates are chosen. Perpendicularity, circumcenters, and orthocenters depend on the Euclidean metric and therefore do not survive arbitrary affine or projective transformations.

Proposition 129.9.2 (Mixed nature of the Euler line)

The centroid is affine in the vertices, while the circumcenter and orthocenter depend on perpendicularity. Therefore the Euler line is not a purely affine theorem and not a purely projective theorem.

Proof. The centroid has the coordinate expression

$$G = \frac{A + B + C}{3},$$

so affine maps send the centroid of ABC to the centroid of the image triangle. By contrast, the circumcenter is defined using perpendicular bisectors, and the orthocenter is defined using altitudes. Perpendicularity is not preserved by a general affine transformation. Hence the Euler line mixes an affine point with metric points. $\square$

This map keeps the elementary proofs in place while explaining why different tools appear. Area ratios prove Ceva, similar triangles prove Menelaus, the triangle inequality in the complex plane proves Ptolemy's inequality, cyclic angles prove Simson, and metric center relations prove Euler-type formulas.

§129.10 Barycentric coordinates

For a triangle ABC , barycentric coordinates write a point P as

$$P = \frac{\alpha A + \beta B + \gamma C}{\alpha + \beta + \gamma}.$$

The ratios in Ceva and Menelaus are naturally expressed through barycentric coordinates. For instance, a point on side BC has coordinates $(0 : \beta : \gamma)$, and the side ratio is encoded by β/γ .

Ceva's theorem becomes a product condition on ratios because three cevians are concurrent exactly when their barycentric coordinate equations have a common solution.

§129.11 Projective transformations

Ceva, Menelaus, and many incidence theorems are projective: they remain true under projective transformations. A projective transformation sends lines to lines and preserves incidence, but not lengths or angles.

This explains why ratio and cross-ratio methods are so powerful. They use quantities that survive projective changes. One may move a complicated triangle configuration to a simpler projective position, prove the statement there, and transfer it back.

§129.12 Ptolemy and inversive geometry

Ptolemy's theorem characterizes cyclic quadrilaterals:

$$AC \cdot BD = AB \cdot CD + BC \cdot AD$$

for a cyclic quadrilateral $ABCD$. Inversive geometry explains why cyclic configurations are flexible. Inversion in a circle sends generalized circles to generalized circles and preserves angles.

Many circle theorems, including Simson line configurations, become simpler after an inversion sends one circle to a line or normalizes the circumcircle.

§129.13 Elliptic and hyperbolic analogues

Classical Euclidean triangle geometry has analogues in spherical and hyperbolic geometry. Ceva-type and Menelaus-type theorems survive with trigonometric replacements. For example, in spherical geometry, side lengths and angles interact through sine laws rather than linear ratios.

This shows which parts of theorems are affine/projective and which depend on Euclidean metric. Ratio statements tend to be projective; angle and length statements depend more strongly on the geometry.

§129.14 Returning to the theorem list

The original list of Ceva, Menelaus, Ptolemy, Simson, and Euler-line facts is not just a catalogue. It spans three geometries: affine ratio geometry, projective incidence geometry, and inversive circle geometry. The advanced reconstruction should keep the elementary proofs while identifying which invariance principle makes each theorem natural.

§129.15 Ceva and Menelaus as dual affine statements

Ceva concerns concurrence of cevians; Menelaus concerns collinearity of points on sides. In barycentric coordinates, these become dual linear-algebraic conditions. Concurrence means three lines have a common solution. Collinearity means three points satisfy one linear equation.

This duality is why the product of ratios appears in both theorems, with reciprocal-looking signs. The elementary ratio-chasing proof is a coordinate proof in the affine-/projective plane.

§129.16 Euler line and affine covariance

The centroid, circumcenter, orthocenter, and nine-point center belong to metric triangle geometry, but some constructions behave well under affine maps and others do not. The centroid is affine: it is the average of vertices. Orthocenter and circumcenter depend on perpendicularity, so they are metric.

The Euler line therefore mixes affine and metric structures. Understanding which points survive affine transformations clarifies why some triangle centers are more universal than others.

§129.17 Simson line through projection

The Simson line theorem says that projections of a point on the circumcircle onto the three sides of a triangle are collinear. This can be interpreted through oriented angles and cyclic quadrilaterals, but also through projective degeneration: as a point moves on the circumcircle, a family of pedal triangles degenerates exactly when the point lies on the circumcircle.

Thus the theorem is not just a miracle of angle chasing; it records a degeneracy condition in a moving geometric construction.

Special Proofs of Trigonometric Addition Formulas

N-20230220

§130.1 Motivation

This handwritten note tries several proofs of trigonometric addition formulas. The author first tried a mixed line-slope proof for tangent, then tried reciprocal bases, then settled on more reliable methods for sine and cosine. This history is important: it records why some apparently clever approaches become awkward.

§130.2 Sine difference formula by geometry

Place two unit vectors at angles α and β :

$$u = (\cos \alpha, \sin \alpha), \quad v = (\cos \beta, \sin \beta).$$

The magnitude of the two-dimensional cross product is

$$|u \times v| = \sin |\beta - \alpha|.$$

But the determinant gives

$$u \times v = \cos \alpha \sin \beta - \sin \alpha \cos \beta.$$

Hence, with orientation taken into account,

$$\sin(\beta - \alpha) = \sin \beta \cos \alpha - \cos \beta \sin \alpha.$$

This proof is clean because the determinant already measures signed area, and signed area of two unit vectors is sine of the included angle.

§130.3 Rotation-matrix proof

The rotation matrix is

$$R(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}.$$

Because rotating by α and then by β is the same as rotating by $\alpha + \beta$,

$$R(\alpha + \beta) = R(\beta)R(\alpha).$$

Multiplying matrices gives

$$\cos(\alpha + \beta) = \cos \alpha \cos \beta - \sin \alpha \sin \beta,$$

$$\sin(\alpha + \beta) = \sin \alpha \cos \beta + \cos \alpha \sin \beta.$$

This is probably the most conceptual elementary proof: addition formulas are simply composition laws for rotations.

§130.4 Euler formula proof

Euler's formula gives

$$e^{i\theta} = \cos \theta + i \sin \theta.$$

Then

$$e^{i(\alpha+\beta)} = e^{i\alpha} e^{i\beta}.$$

Expanding the right-hand side gives

$$\begin{aligned} & (\cos \alpha + i \sin \alpha)(\cos \beta + i \sin \beta) \\ &= (\cos \alpha \cos \beta - \sin \alpha \sin \beta) + i(\sin \alpha \cos \beta + \cos \alpha \sin \beta). \end{aligned}$$

Equating real and imaginary parts yields the addition formulas.

§130.5 Tangent formula

Once sine and cosine are known,

$$\tan(\alpha + \beta) = \frac{\sin(\alpha + \beta)}{\cos(\alpha + \beta)} = \frac{\tan \alpha + \tan \beta}{1 - \tan \alpha \tan \beta},$$

provided the denominators are nonzero. This is why trying to prove tangent first is often less natural: tangent is a quotient, while sine and cosine are coordinates of rotation.

§130.6 The method behind the trick

The same formulas have three interpretations: geometry says sine is signed area, linear algebra says rotations compose by matrix multiplication, and complex analysis says rotations are multiplication by unit complex numbers. The formulas are memorable once they are seen as structure-preservation laws rather than isolated identities.

§130.7 Rotation proof as the structural proof

The cleanest proof of trigonometric addition formulas is the identity

$$R(a + b) = R(a)R(b)$$

for plane rotations. Matrix multiplication gives

$$\cos(a + b) = \cos a \cos b - \sin a \sin b,$$

$$\sin(a + b) = \sin a \cos b + \cos a \sin b.$$

This proof shows that the formulas express composition of rotations.

§130.8 Euler formula and characters

Euler's formula

$$e^{i(a+b)} = e^{ia} e^{ib}$$

gives the same identities by comparing real and imaginary parts. The map $t \mapsto e^{it}$ is a character from the additive group of real angles to the multiplicative circle group.

§130.9 Tangent formula as projective coordinate

Tangent is the slope coordinate on the unit circle, where defined. The formula

$$\tan(a + b) = \frac{\tan a + \tan b}{1 - \tan a \tan b}$$

is the induced fractional-linear rule for slopes under rotation.

§130.10 Addition formulas as rotation composition

The many proofs of addition formulas are different coordinate expressions of one fact: rotations form a group. Geometry, matrices, complex exponentials, and tangent slopes all encode the same composition law.

Trigonometric Systems, Multiple-Angle Formulas, and a Determinant Recurrence

N-20230223

§131.1 Correcting a trigonometric mistake

The first handwritten line says that an earlier calculation from February 20 contained a mistake. The attempted determinant manipulation involved vectors built from angles α and β , and a factor such as $\tan(\alpha - \beta)$. The note records that the determinant-style computation did not work and that the reason for the error was not immediately clear.

This is worth preserving. A failed determinant manipulation is not useless; it tells us that the problem probably has hidden dependencies among the trigonometric variables. Before choosing a linear algebra method, one must check whether the equations are genuinely linear in the chosen unknowns.

§131.2 A trigonometric system

The main problem asks for $\cos \alpha \cos \beta$ under conditions of the form

$$\sin \alpha + \sin \beta = \frac{1}{2}, \quad \cos \alpha - \cos \beta = \frac{1}{3}.$$

The note introduces

$$a = \sin \alpha, \quad b = \sin \beta, \quad c = \cos \alpha, \quad d = \cos \beta.$$

Then the equations become

$$a + b = \frac{1}{2}, \quad c - d = \frac{1}{3}, \quad a^2 + c^2 = 1, \quad b^2 + d^2 = 1.$$

This transforms a trigonometric problem into an algebraic system.

§131.3 A rational-parametrization idea

The handwritten solution tries to express the four variables as sums of rational parts and radicals:

$$a = q_1 + r_1, \quad b = q_2 - r_1, \quad c = q_3 + r_2, \quad d = q_4 + r_2,$$

with

$$q_1 + q_2 = \frac{1}{2}, \quad q_3 - q_4 = \frac{1}{3}.$$

The two circle equations become quadratic equations in the radical parts. Comparing rational and irrational parts can force relations among the q_i and r_i .

The note finds relations such as

$$q_1^2 + q_3^2 = q_2^2 + q_4^2,$$

and a linear relation between r_1 and r_2 . One conclusion written on the page is that the signs may be controlled by introducing one radical and then solving back for c, d . The method is experimental, and the note explicitly says that another method was not successful.

§131.4 A direct algebraic route

A cleaner derivation is to use sum-to-product formulas. We have

$$\sin \alpha + \sin \beta = 2 \sin \frac{\alpha + \beta}{2} \cos \frac{\alpha - \beta}{2} = \frac{1}{2},$$

and

$$\cos \alpha - \cos \beta = -2 \sin \frac{\alpha + \beta}{2} \sin \frac{\alpha - \beta}{2} = \frac{1}{3}.$$

Dividing gives

$$-\tan \frac{\alpha - \beta}{2} = \frac{2}{3},$$

provided the common factor is nonzero. This determines the half-difference angle. Then one can recover products such as

$$\cos \alpha \cos \beta = \frac{\cos(\alpha + \beta) + \cos(\alpha - \beta)}{2}.$$

The direct method shows why the algebraic system has hidden trigonometric structure.

§131.5 Multiple-angle formulas

The second page records standard multiple-angle formulas. One form is

$$\sin(nx) = \sum_{k=0}^n \binom{n}{k} \cos^k x \sin^{n-k} x \sin\left(\frac{(n-k)\pi}{2}\right),$$

and

$$\cos(nx) = \sum_{k=0}^n \binom{n}{k} \cos^k x \sin^{n-k} x \cos\left(\frac{(n-k)\pi}{2}\right).$$

These come from expanding

$$(\cos x + i \sin x)^n = \cos(nx) + i \sin(nx).$$

Thus De Moivre's formula is the conceptual source of the binomial-looking expressions.

The recurrence formulas are

$$\sin(nx) = 2 \sin((n-1)x) \cos x - \sin((n-2)x),$$

$$\cos(nx) = 2 \cos((n-1)x) \cos x - \cos((n-2)x),$$

and

$$\tan(nx) = \frac{\tan((n-1)x) + \tan x}{1 - \tan((n-1)x) \tan x}.$$

The sine and cosine recurrences are Chebyshev-type recurrences.

§131.6 The determinant recurrence

The final page asks one to prove a determinant formula. The displayed matrix is tridiagonal, with diagonal entries involving $2 \cos \theta$ and off-diagonal entries equal to 1, with a modified first entry involving $\cos \theta$. The intended proof is expansion along the last row or last column.

If D_n denotes the determinant of the $n \times n$ tridiagonal matrix with diagonal $2 \cos \theta$ and off-diagonal 1, then it satisfies

$$D_n = 2 \cos \theta D_{n-1} - D_{n-2}.$$

This is the same recurrence as the multiple-angle formulas. Therefore D_n can be expressed in terms of $\sin((n+1)\theta)/\sin \theta$ or a related Chebyshev polynomial, depending on the exact boundary convention.

The note's blue comment says that proving the determinant is a cosine expression can be done by the recurrence. This is exactly the right idea: the matrix determinant and the trigonometric formula are the same recurrence with the same initial values.

§131.7 From technique to structure

This note is about recognizing when a problem should be treated as trigonometry, algebra, or linear algebra. A failed determinant trick reveals hidden constraints. Sum-to-product formulas solve the angle system more naturally. Multiple-angle identities come from De Moivre's formula. Tridiagonal determinants satisfy the same recurrence as Chebyshev polynomials. The common theme is that recurrence relations are a bridge between trigonometry and matrices.

§131.8 Trigonometric systems as algebraic systems

A trigonometric system can often be converted into algebra by setting

$$u = \cos x, \quad v = \sin x, \quad u^2 + v^2 = 1.$$

This turns identities into polynomial relations on the unit circle. The constraint must always be kept; otherwise extraneous algebraic solutions appear.

§131.9 Rational parametrization of the circle

The substitution

$$t = \tan \frac{x}{2}, \quad \cos x = \frac{1-t^2}{1+t^2}, \quad \sin x = \frac{2t}{1+t^2}$$

rationalizes the unit circle except one point. It turns many trigonometric equations into rational equations.

§131.10 Multiple-angle formulas and Chebyshev polynomials

The identity

$$\cos(nx) = T_n(\cos x)$$

defines the Chebyshev polynomial T_n . Thus multiple-angle formulas are polynomial dynamics on the interval $[-1, 1]$.

§131.11 Determinant recurrences

A determinant recurrence often comes from expanding along a row or column. The resulting sequence is commonly governed by a second-order linear recurrence, whose characteristic roots explain closed forms.

§131.12 Trigonometry through algebraic parametrization

The note joins trigonometry and algebra: systems become polynomial constraints, half-angle substitution gives rational parametrization, multiple angles give Chebyshev polynomials, and determinant recurrences reveal linear recurrence structure.

Projective Geometry as the Art of Removing Exceptions

N-20240518

§132.1 The annoyance of parallel lines

Euclidean geometry has a small but persistent exception: two lines meet unless they are parallel. Projective geometry removes the exception by adding points at infinity.

This is not a cosmetic choice. It changes the grammar of geometry. Statements about incidence become cleaner, and perspective drawing becomes natural.

The guiding sentence is:

projective geometry sacrifices distance to save incidence.

§132.2 Homogeneous coordinates are ratios

The projective line over a field k is

$$\mathbb{P}^1(k) = \{[x : y] : (x, y) \neq (0, 0)\} / \sim,$$

where

$$[x : y] = [\lambda x : \lambda y] \quad (\lambda \neq 0).$$

The pair $[x : y]$ is not a point in k^2 . It is a line through the origin in k^2 .

When $y \neq 0$, the point becomes $[x/y : 1]$. The missing point $[1 : 0]$ is the point at infinity.

Similarly,

$$\mathbb{P}^2(k) = \{[x : y : z] : (x, y, z) \neq (0, 0, 0)\} / \sim.$$

The chart $z \neq 0$ is the ordinary affine plane, and the equation $z = 0$ is the line at infinity.

§132.3 A coordinate scene: where parallel lines meet

The affine lines

$$y = 0, \quad y = 1$$


are parallel. Homogenizing them gives

$$y = 0, \quad y - z = 0.$$

On the line at infinity $z = 0$, both equations force $y = 0$, so the intersection is

$$[1 : 0 : 0].$$

Thus projective geometry does not say that parallel lines visibly meet in the affine plane. It says that the affine plane was missing one direction point.



The diagram is symbolic: the point at infinity is not an ordinary point far to the right. It is a direction added to complete the incidence relation.

§132.4 Duality: points and lines speak the same language

A line in $\mathbb{P}^2$ has equation

$$ax + by + cz = 0.$$

The coefficients $[a : b : c]$ are homogeneous. Thus lines themselves have projective coordinates.

Incidence is the bilinear equation

$$ax + by + cz = 0.$$

This symmetry produces projective duality: many theorems remain true after swapping point and line.

For instance:

$$\text{two points determine a line} \quad \longleftrightarrow \quad \text{two lines determine a point.}$$

Projective geometry is full of such paired statements.

§132.5 Transformations and what survives

An invertible linear map $A : k^3 \rightarrow k^3$ induces a projective transformation

$$[v] \longmapsto [Av].$$

It preserves incidence but not length or angle. A circle may become an ellipse under a perspective projection, but a line remains a line.

This is the geometric lesson: once one changes the allowed transformations, one changes the invariants. Euclidean geometry cares about distances and angles. Projective geometry cares about incidence, alignment, cross-ratio, and projective equivalence.

§132.6 A first invariant: the cross-ratio

For four points a, b, c, d on a projective line, the cross-ratio is written in an affine coordinate as

$$[a, b; c, d] = \frac{(c - a)(d - b)}{(c - b)(d - a)}.$$

It is preserved by projective transformations of the line.

This formula is less important here than the role it plays. It is a warning that projective geometry is not invariant-free. It simply has different invariants from Euclidean geometry.

§132.7 Conics as one projective family

A projective conic is given by a homogeneous quadratic equation

$$X^T A X = 0, \quad X = (x, y, z)^T.$$

In the affine chart $z = 1$, it becomes an ordinary quadratic equation in x and y . Ellipses, parabolas, and hyperbolas are different affine views of projective conics.

This explains why conics look artificially split in elementary analytic geometry. The affine chart cuts a projective object and then classifies the slice.

§132.8 The projective plane as all lines through the origin

There is a concrete model that removes much of the mystery. A point $[x : y : z] \in \mathbb{P}^2$ is a line through the origin in k^3 . A projective line in $\mathbb{P}^2$ is a plane through the origin in k^3 .

Then incidence becomes containment:

$$\begin{array}{c} \text{projective point on projective line} \\ \iff \\ \text{line in } k^3 \text{ inside plane in } k^3. \end{array}$$

This model is often the easiest way to remember why two projective lines meet. Two planes through the origin in k^3 usually meet in a line through the origin, which is a projective point.

§132.9 Desargues as a projective-flavored theorem

Projective geometry is naturally comfortable with theorems about incidence. Desargues' theorem says roughly: if two triangles are perspective from a point, then the intersections of corresponding sides lie on a line.

The exact diagram is less important here than the type of statement. It uses points lying on lines, lines meeting at points, and no measurements. This is projective geometry's native habitat.

One reason projective geometry feels more elegant than Euclidean geometry is that many such incidence theorems survive projection. A perspective drawing may distort lengths wildly, but it preserves collinearity and concurrence.

§132.10 Projective transformations in coordinates

A projective transformation of the projective line is represented by a matrix

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

with nonzero determinant. In affine coordinate $t = x/y$, it acts by a fractional linear transformation

$$t \mapsto \frac{at + b}{ct + d}.$$

This explains why the cross-ratio is built from ratios of differences: it is designed to survive these fractional linear transformations.

The point at infinity is not an exception in this formula. It is exactly what is needed to make the transformation defined as a map of the whole projective line.

§132.11 The conic as a disguised projective line

A nonsingular conic is projectively equivalent to $\mathbb{P}^1$. For example, the conic

$$xz - y^2 = 0$$

has the parametrization

$$[s : t] \mapsto [s^2 : st : t^2].$$

This formula sends a point of $\mathbb{P}^1$ to a point on the conic, because

$$s^2t^2 - (st)^2 = 0.$$

So a conic is not merely a quadratic curve; projectively, it behaves like a line wrapped into the plane by quadratic coordinates.

This fact is a useful bridge to later algebraic geometry: parametrizations, homogeneous equations, and projective spaces start to work together.

§132.12 Final checkpoint

Projective geometry is not the study of a strange faraway infinity. It is the study of incidence after removing affine exceptions.

The three core moves are:

- replace coordinates by ratios,
- add direction points at infinity,
- study transformations preserving incidence.

Transformations of Conics and the Erlangen View of Geometry

N-20250211

§133.1 The organizing question

This chapter studies conic sections through transformations. The guiding question is not only how an ellipse, parabola, or hyperbola is described by an equation, but what kind of geometry regards two such objects as the same.

The hierarchy is:

Euclidean $\subset$ similarity $\subset$ affine $\subset$ projective $\subset$ conformal/inversive phenomena.

Each enlargement forgets some invariants and preserves others. This is the spirit of Klein's Erlangen program: a geometry is the study of properties invariant under a chosen transformation group.

§133.2 Affine transformations

Definition 133.2.1. An affine transformation of $\mathbb{R}^n$ is a map

$$T(x) = Ax + b,$$

where A is a linear map and $b \in \mathbb{R}^n$. If $A \in GL(n, \mathbb{R})$, then T is an affine automorphism.

Affine transformations preserve straight lines, parallelism, ratios along a line, convexity, and barycentric combinations. They do not preserve lengths, angles, or circles.

Example 133.2.2

An ellipse can be sent to a circle by an affine transformation. If

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1,$$

then the affine map

$$(x, y) \mapsto (x/a, y/b)$$

sends it to the unit circle.

§133.3 Conics in matrix form

A plane conic may be written as

$$Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0.$$

In homogeneous coordinates $[x : y : z]$, this becomes

$$Ax^2 + Bxy + Cy^2 + Dxz + Eyz + Fz^2 = 0.$$

Equivalently,

$$X^T Q X = 0, \quad X = (x, y, z)^T,$$

where Q is a symmetric 3×3 matrix.

Under a projective transformation $X \mapsto MX$, the conic transforms by

$$Q \mapsto M^{-T} Q M^{-1}.$$

Thus projective classification of conics becomes a question about symmetric bilinear forms up to change of coordinates.

A concrete matrix is the best antidote to the vague sentence that all nondegenerate conics are projectively equivalent. The parabola

$$C_1: \quad y = x^2$$

has homogeneous equation

$$X^2 - YZ = 0,$$

so it is represented by

$$A_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & -\frac{1}{2} \\ 0 & -\frac{1}{2} & 0 \end{pmatrix}.$$

The unit circle

$$C_2: \quad X^2 + Y^2 = Z^2$$

is represented by

$$A_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}.$$

Now take

$$M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & -1 & 1 \end{pmatrix}.$$

A direct multiplication gives

$$M^T A_1 M = A_2.$$

Indeed, if $U = (u, v, w)^T$ and $X = MU = (u, v + w, -v + w)^T$, then

$$X^T A_1 X = u^2 - (v + w)(-v + w) = u^2 + v^2 - w^2 = U^T A_2 U.$$

Thus the projective change of coordinates $U = M^{-1}X$ sends the parabola equation to the circle equation.

This is not an affine transformation. The inverse matrix is

$$M^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{2} & -\frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}.$$

The unique real point at infinity of the parabola,

$$[0 : 1 : 0],$$

is sent to

$$M^{-1}[0 : 1 : 0] = [0 : 1 : 1],$$

a finite real point on the circle. Meanwhile the usual circular points at infinity of the real circle are the complex points

$$[1 : i : 0], \quad [1 : -i : 0].$$

So the projective transformation has moved the affine line at infinity itself; this is exactly why a parabola and a circle can be the same projective conic while remaining different in affine geometry. The difference between them is absolute for affine geometry, but only a coordinate choice for projective geometry.

§133.4 Projective geometry

Projective geometry adds points at infinity. In the real projective plane,

$$\mathbb{P}^2(\mathbb{R}) = \{[x : y : z] \mid (x, y, z) \neq 0\} / \sim,$$

where

$$[x : y : z] = [\lambda x : \lambda y : \lambda z], \quad \lambda \neq 0.$$

The affine plane is the chart $z \neq 0$, and the line at infinity is $z = 0$.

Definition 133.4.1. A projective transformation of $\mathbb{P}^2$ is induced by an invertible matrix

$$M \in GL(3, \mathbb{R}),$$

acting by

$$[X] \mapsto [MX].$$

Scalar multiples of M induce the same projective transformation. Thus the transformation group is

$$PGL(3, \mathbb{R}).$$

Projective transformations preserve incidence and cross-ratio, but not parallelism, distance, or angle. In projective geometry, all nondegenerate conics over an algebraically closed field are projectively equivalent.

§133.5 Cross-ratio

For four collinear points A, B, C, D , the cross-ratio is denoted

$$(A, B; C, D).$$

In an affine coordinate it has the form

$$(a, b; c, d) = \frac{(c - a)(d - b)}{(c - b)(d - a)}.$$

Theorem 133.5.1

Projective transformations preserve cross-ratio.

Remark 133.5.2 — The cross-ratio is the prototype of a projective invariant. Once one understands it, projective geometry stops looking like distorted Euclidean geometry and becomes an invariant theory of incidence.

§133.6 Projective duality and conics

Projective duality interchanges points and lines. A nondegenerate conic defines a duality: to a point p , one associates its polar line with respect to the conic. In matrix form, if the conic is $X^T Q X = 0$, then the polar line of p is

$$p^T Q X = 0.$$

Remark 133.6.1 — The same quadratic form that defines the conic also turns points into lines. The conic is therefore not just a curve; it is a device for dualizing projective geometry.

§133.7 Inversion

Definition 133.7.1. Inversion in the circle of radius r centered at the origin is the map

$$x \mapsto \frac{r^2}{\|x\|^2} x.$$

In the complex plane, inversion in the unit circle is essentially

$$z \mapsto \frac{1}{\bar{z}}.$$

Inversion sends lines and circles to lines and circles, with the convention that lines are circles through infinity. It is conformal away from the center of inversion: it preserves angles but not distances.

§133.8 Möbius transformations

A Möbius transformation has the form

$$z \mapsto \frac{az + b}{cz + d}, \quad ad - bc \neq 0.$$

It acts naturally on the Riemann sphere

$$\hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}.$$

The group of Möbius transformations is

$$PGL(2, \mathbb{C}).$$

Möbius transformations send generalized circles to generalized circles and preserve cross-ratios. They are simultaneously projective transformations of $\mathbb{P}^1(\mathbb{C})$ and conformal automorphisms of the Riemann sphere.

§133.9 Conic sections

Classically, conics arise as plane sections of a cone. Analytically, they are degree-two plane curves. Projectively, a conic is a homogeneous quadratic equation in $\mathbb{P}^2$.

Geometry	Transformations	Conic classification
Euclidean	Rigid motions	Metric shape matters
Similarity	Rotations, translations, scalings	Eccentricity-like shape remains
Affine	Invertible affine maps	Ellipses are equivalent to circles
Projective	$PGL(3)$	Nondegenerate conics are equivalent over algebraically closed fields

§133.10 The Erlangen program

Klein's Erlangen program says that geometry should be understood through a group of transformations and the invariants under that group. Euclidean geometry studies invariants under Euclidean motions; affine geometry studies invariants under affine transformations; projective geometry studies invariants under projective transformations.

Remark 133.10.1 — This viewpoint explains why the same object can look different in different geometries. A circle is special in Euclidean geometry, but not in affine geometry. Parallel lines are meaningful in affine geometry, but not in projective geometry. Cross-ratio is invisible in elementary Euclidean geometry, but central in projective geometry.

§133.11 A hierarchy of invariants

Transformation group	Preserved	Forgotten
Euclidean	Distance, angle, incidence	Coordinates
Similarity	Angle, ratios of lengths	Absolute scale
Affine	Incidence, parallelism, ratios on a line	Angle and length
Projective	Incidence, cross-ratio	Parallelism and affine ratio
Möbius	Generalized circles, angles, cross-ratio on $\hat{\mathbb{C}}$	Euclidean size

§133.12 Quadratic forms and matrix language

A conic in the projective plane can be written as

$$X^T Q X = 0,$$

where $X = (x : y : z)^T$ is a homogeneous coordinate vector and Q is a symmetric matrix. Passing to homogeneous coordinates is not cosmetic. It lets ellipses, parabolas, and hyperbolas be treated uniformly and allows points at infinity to participate in the geometry.

Under a projective change of coordinates $X = MX'$, the conic matrix transforms as

$$Q' = M^T Q M$$

or equivalently $Q \mapsto M^{-T} Q M^{-1}$, depending on whether one transforms coordinates or points. The equation changes, but the geometric conic is the same object in new coordinates.

§133.13 Affine classification versus projective classification

In affine geometry, ellipses, parabolas, and hyperbolas are genuinely different because the line at infinity is fixed. The way a conic meets the line at infinity distinguishes the types: an ellipse has no real points at infinity, a parabola is tangent to the line at infinity, and a hyperbola meets it in two real points.

In projective geometry over an algebraically closed field, all nonsingular conics are projectively equivalent. Projective geometry forgets metric shape and focuses on incidence and cross-ratio-type data.

§133.14 Polar lines

For a nonsingular conic $X^T Q X = 0$, a point P has a polar line

$$P^T Q X = 0.$$

If P lies on the conic, the polar line is the tangent line at P . If P lies outside, the polar encodes the chord of contact of tangents from P .

The pole-polar relation is one of the best examples of why projective conic theory is richer than coordinate classification. It converts points to lines and lines to points in a way controlled by the conic.

§133.15 The structural moral

The transformation-centered view naturally leads to modern geometry:

- (i) algebraic geometry studies varieties and morphisms, with projective space as a basic ambient object;
- (ii) differential geometry studies invariants under diffeomorphisms or isometries;
- (iii) complex geometry studies holomorphic maps;
- (iv) representation theory studies symmetries through linear actions;
- (v) category theory abstracts the entire pattern by treating structure-preserving maps as primary.

§133.16 How to choose the right geometry

The Erlangen viewpoint becomes useful only when it tells us which tools are allowed. For conic problems, a practical guide is:

Question being asked	Natural geometry	Reason
length, angle, focus, eccentricity midpoints, parallelism, area ratios	Euclidean/similarity affine	metric data is essential affine maps preserve barycentric structure
incidence, tangency, conic equivalence	projective	the line at infinity may move
cross-ratio or four points on a line	projective	cross-ratio is projectively invariant
angle with generalized circles	conformal/inversive	inversion preserves angles and generalized circles

Remark 133.16.1 (A common mistake) — Saying that two conics are projectively equivalent does not mean they are the same in Euclidean or affine geometry. A parabola and a circle may be projectively equivalent, but the projective map can move the line at infinity. Once that happens, affine notions such as parallelism and the distinction between ellipse/parabola/hyperbola need not be preserved.

§133.17 Using the transformation group in a conic problem

A conic in homogeneous coordinates is

$$X^T A X = 0.$$

If the problem asks only for incidence or tangency, one may use a projective change of coordinates; algebraically this replaces A by a congruent matrix $P^T A P$. The tangent at a point X is read from the bilinear form as

$$X^T A Y = 0,$$

so the same matrix controls the conic, its tangent lines, and the pole–polar correspondence.

The restrictions become important when the question contains metric words. A projective map may send a circle to a parabola because it can move the line at infinity, so it does not preserve Euclidean length, angle, focus, or eccentricity. An affine map keeps parallelism and ratios on parallel lines but may change angles. Inversion and conformal maps preserve angles and generalized circles, but not distances. Thus in an actual problem the first step is to mark which data must survive: cross-ratio and incidence suggest projective coordinates; parallelism suggests affine coordinates; angles with circles suggest inversive geometry; lengths force a Euclidean or similarity argument.

XIII

Curves and Structures

Part XIII: Contents

N-20220212	863
134.1 Structure beyond topology	863
134.2 From smooth structure to complex structure	864
134.3 Definition of a Riemann surface	864
134.4 Examples	864
134.5 Holomorphic maps between Riemann surfaces	866
134.6 Meromorphic functions as maps to the sphere	867
134.7 Degree and multiplicity	868
134.8 Local power form is not global	869
134.9 Holomorphic transition maps	869
134.10 Examples beyond the plane	870
134.11 Meromorphic functions	870
134.12 Conceptual synthesis	870
134.13 Complex structure is stronger than topology	870
134.14 Meromorphic functions and branched coverings	871
134.15 Riemann–Hurwitz	871
134.16 Complex curves through topology and ramification	871
N-20221214	873
135.1 The route of the note	873
135.2 Equidecomposability and dissection congruence	873
135.3 Hilbert’s third problem and the Dehn invariant	875
135.4 Measure-theoretic background	875
135.5 Tarski’s circle-squaring problem	875
135.6 Scissors congruence and why nice pieces fail	876
135.7 Laczkovich’s positive theorem	876
135.8 First idea: work in the torus	876
135.9 Random translations and a free action	877
135.10 The graph and bounded-distance bijection	877
135.11 Orbit matching	878
135.12 Discrepancy estimates	879
135.13 From Laczkovich to Borel circle squaring	879
135.14 Bridge to later ideas	880
135.15 Scissors congruence and invariants	880
135.16 Measurable versus nonmeasurable decompositions	880
135.17 Circle squaring and group actions	880
135.18 Equidecomposition under measure and group action	881
N-20230801	883
136.1 Starting point	883
136.2 Connections and curvature	883
136.3 Higgs bundles	883
136.4 The moduli space	884
136.5 The Hitchin fibration	884
136.6 Spectral curves	884
136.7 Geometric Langlands perspective	886
136.8 Synthesis	886
136.9 Moment-map reading of the equations	886
136.10 Spectral curves as global eigenvalues	886
136.11 Four descriptions to keep together	886
136.12 Source repair: Yang–Mills background	887
136.13 Source repair: dimensional reduction	887
136.14 Source repair: gauge invariance	887
136.15 Source repair: Higgs bundles and stability	888
136.16 Source repair: spectral data and Hitchin fibration	888
136.17 Source repair: geometric Langlands and S-duality	888
136.18 Source repair: local work still missing	889

N-20240608	891
137.1 The curve is not only the drawing	891
137.2 The function ring of a curve	891
137.3 When factorization becomes geometry	892
137.4 Smooth points and singular points	892
137.5 Tangent cones: the first nonzero shadow	892
137.6 The cusp as a ring	893
137.7 Intersections: one visible point may count twice	893
137.8 Projective closure repairs missing intersections	893
137.9 Bezout as a promise of clean accounting	893
137.10 Normalization: repairing a singular curve	894
137.11 Resultants: eliminating a variable	894
137.12 A projective calculation with a line and a conic	894
137.13 Real pictures and algebraically closed accounting	895
137.14 How to read a plane curve from its equation	895
N-20240914	897
138.1 A ledger for meromorphic functions	897
138.2 Principal divisors: the ledger of one function	897
138.3 Reading a divisor as permission	897
138.4 Line bundles: permission becomes geometry	898
138.5 The canonical divisor	898
138.6 Riemann–Roch without intimidation	899
138.7 Hurwitz as topology with branch points	899
138.8 A divisor game on the sphere	899
138.9 Degree as a first estimate	899
138.10 Canonical divisors through examples	900
138.11 Why Riemann–Roch has a correction term	900
138.12 Final caution	900
N-20241026	901
139.1 Starting from functions, not pictures	901
139.2 Coordinate rings: a space through its functions	901
139.3 Nullstellensatz: the dictionary becomes reliable	901
139.4 The Zariski topology is coarse on purpose	902
139.5 Distinguished opens and localization	902
139.6 Sheaves: local data with a gluing rule	902
139.7 Stalks: looking through a microscope	902
139.8 A local example: two axes crossing	903
139.9 Why the strange topology helps	903
139.10 Products and unions through rings	903
139.11 Radical ideals and forgotten nilpotents	903
139.12 Regular functions versus continuous functions	904
139.13 Morphisms as pullback of functions	904
139.14 A compact dictionary	904
N-20250118	905
140.1 A new kind of point	905
140.2 The affine line spectrum	905
140.3 Closed sets and distinguished opens	905
140.4 Localization is zooming in	906
140.5 The structure sheaf	906
140.6 Nilpotents: invisible but geometric	906
140.7 The arithmetic curve	906
140.8 Maximal ideals are not enough	907
140.9 Specialization as movement from generic to concrete	907
140.10 Residue fields	907
140.11 A tiny nonreduced calculation	907
140.12 Reading $\text{Spec } R$ in examples	908

N-20250301	909
141.1 Projective geometry returns with functions attached	909
141.2 The projective line as the model	909
141.3 Graded rings and homogeneous data	909
141.4 Projective space as Proj	909
141.5 The projective plane through three windows	910
141.6 A conic across charts	910
141.7 Projective closure with scheme memory	910
141.8 What Proj adds	911
141.9 Why the irrelevant ideal is excluded	911
141.10 Standard opens as spectra	911
141.11 A second conic chart calculation	911
141.12 Line bundles are already nearby	912
141.13 Closing the geometry arc	912
 N-20250705	 913
142.1 First orientation	913
142.2 The topological object and its planar shadow	913
142.3 Some basic knot-theory language	914
142.4 The pieces of a diagram	916
142.5 Nikonov's bridge: diagram elements as probes	917
142.6 Invariants, coinvariants, and labels	918
142.7 What this note is not trying to do	919
 N-20250707	 921
143.1 Reading only one small part of a long paper	921
143.2 Where this sits inside Nikonov's paper	921
143.3 A region is not just a blank patch	922
143.4 At a crossing, four regions want to talk to each other	923
143.5 Why the operation is partial	924
143.6 Probe clearing and color rules	925
143.7 What about crossings themselves?	926
143.8 The paper in one playful sentence	927

Riemann Surfaces as One-Dimensional Complex Manifolds

N-20220212

§134.1 Structure beyond topology

The starting question is: what does it mean for a space to carry a structure? A topological manifold has a topology; a smooth manifold has a smooth structure; a complex manifold has a complex structure.

A topology tells us which sets are open, which sequences or nets converge, and which maps are continuous. A smooth structure adds the ability to speak about differentiability in local coordinates. A complex structure is stricter: coordinate changes must be holomorphic.

Example 134.1.1

The interval $(0, 2) \subset \mathbb{R}$ is a topological 1-manifold. One chart may be the identity map

$$\varphi_1(x) = x.$$

Another chart can still be a homeomorphism while failing to be smooth. This shows why topology alone is not enough to define smoothness.




Figure 134.1: Two coordinate descriptions of the interval used to separate topology from smooth structure.

Remark 134.1.2 — The lesson is that the same underlying topological space may admit different additional structures. Smoothness is not merely a property of the set; it is encoded by charts and transition maps.

§134.2 From smooth structure to complex structure

A smooth surface embedded in $\mathbb{R}^3$ can be described locally by diffeomorphisms to open subsets of $\mathbb{R}^2$. But this embedded definition is not sufficiently general: for example, the Klein bottle cannot be embedded in $\mathbb{R}^3$.

The abstract definition of a smooth manifold therefore uses charts

$$\varphi_i : U_i \rightarrow E_i \subseteq \mathbb{R}^n$$

whose transition maps are smooth. The same philosophy gives complex manifolds by replacing $\mathbb{R}^n$ with $\mathbb{C}^n$ and smooth transition maps with holomorphic transition maps.

Remark 134.2.1 — Because holomorphic functions are automatically analytic, a complex structure is very rigid. This is why a Riemann surface is much more than a real two-dimensional smooth manifold.

§134.3 Definition of a Riemann surface

Definition 134.3.1. A Riemann surface is a connected, Hausdorff, second-countable topological space X equipped with an atlas $\{(U_i, \varphi_i)\}$, where

$$\varphi_i : U_i \xrightarrow{\sim} E_i \subseteq \mathbb{C}$$

is a homeomorphism onto an open subset of $\mathbb{C}$, and all transition maps

$$\varphi_j \circ \varphi_i^{-1}$$

are holomorphic on their domains.

Each φ_i is a complex chart. A local coordinate near $p \in X$ is the complex number

$$z = \varphi_i(x).$$

If $\varphi_i(p) = 0$, the coordinate is centered at p .

Definition 134.3.2. More generally, a complex n -manifold is a Hausdorff space locally modeled on open subsets of $\mathbb{C}^n$, with holomorphic transition maps.

Every complex n -manifold is naturally a smooth real $2n$ -manifold.

§134.4 Examples

Example 134.4.1 (Open subsets of the plane)

Every connected open subset $U \subseteq \mathbb{C}$ is a Riemann surface, with the identity chart.

Example 134.4.2 (The Riemann sphere)

The Riemann sphere $\mathbb{C}_\infty$ is the complex plane together with a point at infinity. One way to construct its complex structure is to cover it by two charts: one chart with coordinate z , and another chart near infinity with coordinate

$$t = 1/z.$$

On the overlap, the transition function is

$$z = 1/t,$$

which is holomorphic away from 0. Thus the two charts glue into a Riemann surface.




Figure 134.2: Stereographic charts used to define the complex structure on the Riemann sphere.

Caution

The sign choice in the stereographic chart matters. If the transition map became $z = 1/\bar{t}$, it would be antiholomorphic rather than holomorphic. The complex structure depends on getting this orientation right.




Figure 134.3: The second stereographic chart and the transition coordinate near infinity.

Example 134.4.3 (A complex torus)

Let $L = \mathbb{Z}[i] \subset \mathbb{C}$, a lattice under addition. The quotient

$$\mathbb{C}/L$$

is a complex torus. Locally, the quotient map from $\mathbb{C}$ gives complex coordinates, and different choices of lifts differ by translations

$$z \mapsto z + a, \quad a \in L,$$

which are holomorphic.




Figure 134.4: The quotient picture of a complex torus, with opposite sides identified.

Remark 134.4.4 — The complex torus is compact. Any holomorphic function on it is constant, but meromorphic functions on it are much more interesting.

Example 134.4.5

A disjoint union of two Riemann spheres is not a Riemann surface under the convention used here, because a Riemann surface is required to be connected. If one drops connectedness, each connected component is a Riemann surface.

§134.5 Holomorphic maps between Riemann surfaces

Definition 134.5.1. Let X, Y be Riemann surfaces. A map $f : X \rightarrow Y$ is holomorphic at $p \in X$ if, for charts

$$\varphi_1 : U_1 \rightarrow E_1, \quad \varphi_2 : U_2 \rightarrow E_2$$

with $p \in U_1$ and $f(p) \in U_2$, the local expression

$$\varphi_2 \circ f \circ \varphi_1^{-1}$$

is holomorphic near $\varphi_1(p)$. The map is holomorphic if this holds at every point.

Example 134.5.2 (The cubic map)

The map

$$f : \mathbb{C} \rightarrow \mathbb{C}, \quad f(z) = z^3$$

is holomorphic. In the usual global coordinate on $\mathbb{C}$, this is just a polynomial, hence holomorphic everywhere. At a point $a \neq 0$, the derivative

$$f'(a) = 3a^2$$

is nonzero, so the inverse function theorem gives a local holomorphic inverse near a . At 0, however, the derivative vanishes. The local model is exactly

$$z \mapsto z^3,$$

so the map has multiplicity 3 at the origin. This is the first concrete place where “holomorphic map” and “branched covering” begin to touch.

Example 134.5.3 (Inclusion into the Riemann sphere)

The inclusion

$$\mathbb{C} \hookrightarrow \mathbb{C}_\infty$$

is a holomorphic map of Riemann surfaces. On the finite chart of $\mathbb{C}_\infty$, this map is simply

$$z \mapsto z,$$

so it is holomorphic in the most literal possible sense. The point of the example is not the formula itself, but the target: the Riemann sphere adds one extra point ∞ , and the ordinary plane sits inside it as the finite chart. Thus a map into $\mathbb{C}_\infty$ may be checked by using the coordinate z away from infinity and the coordinate $w = 1/z$ near infinity.

§134.6 Meromorphic functions as maps to the sphere

The Riemann sphere lets us treat poles as honest values. A meromorphic function

$$f : X \rightarrow \mathbb{C}$$

can be regarded as a holomorphic map

$$g : X \rightarrow \mathbb{C}_\infty$$

by sending poles to ∞ .

Proposition 134.6.1

Nonconstant meromorphic functions $X \rightarrow \mathbb{C}$ correspond to holomorphic maps $X \rightarrow \mathbb{C}_\infty$ not identically equal to ∞ .

Idea. Away from a pole this is clear. Near a pole, use the coordinate $t = 1/z$ near infinity on $\mathbb{C}_\infty$. If f has a pole, then $1/f$ is holomorphic and vanishes there, so the map into the sphere is holomorphic in the t -coordinate. $\square$

§134.7 Degree and multiplicity

For nonconstant holomorphic maps between compact Riemann surfaces, fibres behave surprisingly well when counted with multiplicity.

Proposition 134.7.1

Let X, Y be compact Riemann surfaces and let $f : X \rightarrow Y$ be nonconstant holomorphic. For each $y \in Y$, count the points in $f^{-1}(y)$ with multiplicity. The resulting number is finite and independent of y . This number is the degree of f , denoted $\deg(f)$.

Remark 134.7.2 — This is much stronger than what one expects from arbitrary smooth maps. The rigidity of holomorphic maps makes the fibre count locally constant, and connectedness makes it globally constant.

Example 134.7.3 (The power map has degree k)

The map

$$z \mapsto z^k$$

extends to a holomorphic map $\mathbb{C}_\infty \rightarrow \mathbb{C}_\infty$ of degree k . On the finite chart this is a polynomial. Near infinity, use the coordinate $w = 1/z$. Then

$$z^k = \frac{1}{w^k}, \quad \frac{1}{z^k} = w^k,$$

so the extended map is also holomorphic at ∞ . For a generic value $a \in \mathbb{C}^*$, the equation

$$z^k = a$$

has k distinct solutions. Counting multiplicity, the same number remains true over the exceptional values 0 and ∞ . This is why the degree is k , not merely because the formula contains the exponent k .

Definition 134.7.4. Let $f : X \rightarrow Y$ be a nonconstant holomorphic map and $p \in X$. The multiplicity of f at p is the unique integer $m \geq 1$ such that, after choosing suitable local coordinates centered at p and $f(p)$, the map is locally of the form

$$z \mapsto z^m.$$

Remark 134.7.5 — The guiding slogan is: every nonconstant holomorphic map locally looks like a power map. This is one of the reasons Riemann surfaces are so much cleaner than arbitrary smooth surfaces.

Idea of the local normal form. Choose local coordinates centered at p and $f(p)$, so the local expression becomes a nonconstant holomorphic function F with $F(0) = 0$. Its Taylor expansion has the form

$$F(z) = a_m z^m + a_{m+1} z^{m+1} + \cdots, \quad a_m \neq 0.$$

Thus

$$F(z) = z^m u(z),$$

where u is holomorphic and $u(0) \neq 0$. After shrinking the coordinate neighborhood, u has a holomorphic m -th root v . Replacing the source coordinate z by $zv(z)$, the map becomes exactly

$$z \mapsto z^m.$$

The integer m is uniquely determined by the order of vanishing of F at 0, hence does not depend on the chosen coordinates. $\square$

§134.8 Local power form is not global

The local normal form $z \mapsto z^m$ must be read with the word *local* emphasized. A useful test case is

$$U = \mathbb{C}^*, \quad F : U \rightarrow \mathbb{C}^*, \quad F(z) = z^2.$$

At $p = 1$, the derivative is nonzero, so in the local-normal-form sense the multiplicity is $m = 1$. Nearby, one can choose a local inverse

$$z = \sqrt{w}$$

with $w = F(z)$. But this square-root coordinate cannot be chosen globally on $\mathbb{C}^*$. If the target point travels once around the loop

$$w(t) = e^{2\pi it}, \quad 0 \leq t \leq 1,$$

then analytic continuation of the local branch starting at $\sqrt{1} = 1$ ends at -1 , not 1. This sign change is monodromy.

There is a small trap here. If one instead moves the source point as

$$p(t) = e^{4\pi it},$$

then $F(p(t)) = e^{4\pi it}$ winds twice around the target origin, so the square root changes sign twice and returns to the original branch. The monodromy is seen by following a local inverse over a loop in the target, not by confusing the covering coordinate with a single global coordinate downstairs.

Thus the map $z \mapsto z^2$ is an unbranched two-fold covering of $\mathbb{C}^*$, while the missing point 0 is where the corresponding map on $\mathbb{C}$ would branch. To make $\sqrt{w}$ single-valued one must either cut the base, for example by removing a ray, or pass to a covering space. This is the global meaning of the failure: local coordinate charts glue with a monodromy representation, and on a non-simply-connected region those gluings need not collapse to one global coordinate.

§134.9 Holomorphic transition maps

A Riemann surface is locally modeled on open subsets of $\mathbb{C}$, but the crucial condition is on overlaps. If two charts

$$\varphi : U \rightarrow V \subseteq \mathbb{C}, \quad \psi : U' \rightarrow V' \subseteq \mathbb{C}$$

overlap, then the transition map

$$\psi \circ \varphi^{-1} : \varphi(U \cap U') \rightarrow \psi(U \cap U')$$

must be holomorphic. This is what makes complex analysis possible on the surface without choosing one global coordinate.

The condition is stronger than smoothness. Holomorphic maps are rigid: they preserve angles, satisfy Cauchy–Riemann equations, and are locally power series. Thus a Riemann surface is not just a two-dimensional smooth surface; it is a surface with complex-analytic coordinate changes.

§134.10 Examples beyond the plane

The complex plane $\mathbb{C}$ is a Riemann surface with one global chart. The Riemann sphere

$$\widehat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$$

requires two charts: one coordinate z near finite points and another coordinate $w = 1/z$ near infinity. The transition map $w = 1/z$ is holomorphic away from zero, so the charts are compatible.

A complex torus is obtained from $\mathbb{C}$ by quotienting by a lattice

$$\Lambda = \mathbb{Z}\omega_1 + \mathbb{Z}\omega_2.$$

The quotient

$$\mathbb{C}/\Lambda$$

is topologically a torus but also carries a complex structure inherited from the plane. This example is the bridge from elementary complex analysis to elliptic curves.

§134.11 Meromorphic functions

On a compact Riemann surface, holomorphic functions are often forced to be constant by the maximum principle. Meromorphic functions are therefore more flexible. A meromorphic function is holomorphic except at isolated poles. Locally near a pole, it has a Laurent expansion with finitely many negative terms.

Thinking of meromorphic functions as holomorphic maps to the Riemann sphere is useful:

$$f : X \rightarrow \widehat{\mathbb{C}}.$$

A pole is simply a point mapped to ∞ . This viewpoint turns singularities into ordinary values on a compact target.

§134.12 Conceptual synthesis

The main transition in the note is from topology to complex geometry. A topological surface only knows neighborhoods up to homeomorphism. A smooth surface knows differentiable coordinate changes. A Riemann surface knows holomorphic coordinate changes, so complex analysis can be done intrinsically. This is why Riemann surfaces are simultaneously one-dimensional complex manifolds and two-dimensional real manifolds.

§134.13 Complex structure is stronger than topology

A Riemann surface is topological dimension two but complex dimension one. The complex charts have holomorphic transition maps, so local coordinates remember angle and orientation. This is why every holomorphic map between Riemann surfaces is far more rigid than a continuous map.

§134.14 Meromorphic functions and branched coverings

A nonconstant meromorphic function on a compact Riemann surface is a holomorphic map to the Riemann sphere:

$$f : X \rightarrow \mathbb{P}^1(\mathbb{C}).$$

Its degree counts preimages with multiplicity. Branch points are where local behavior looks like

$$z \mapsto z^e.$$

The integer e records ramification.

§134.15 Riemann–Hurwitz

For a holomorphic map $f : X \rightarrow Y$ of degree d , Riemann–Hurwitz says

$$2g_X - 2 = d(2g_Y - 2) + \sum_{p \in X} (e_p - 1).$$

This formula ties topology, degree, and branching into one invariant statement.

§134.16 Complex curves through topology and ramification

Riemann surfaces are where topology, complex analysis, and algebraic geometry first coincide. Holomorphic charts give rigidity, meromorphic functions give maps to the sphere, divisors record zeros and poles, and ramification translates analytic multiplicity into topological genus.

Tarski's Circle-Squaring Problem: From Polygon Dissection to Borel Equidecomposition

N-20221214

§135.1 The route of the note

This note is about Tarski's circle-squaring problem, but the source pages do not treat it as a single isolated puzzle. They move through a sequence of increasingly subtle notions:

Tarski circle-squaring $\longrightarrow$ Hilbert's third problem
 $\longrightarrow$ measure and paradoxical decomposition
 $\longrightarrow$ Laczkovich's theorem $\longrightarrow$ Borel circle squaring.

The figures below are kept only when the visual configuration itself matters. The surrounding theorem statements, comments, formulas, and handwritten annotations have been incorporated into the prose.

Tarski asked in 1925 whether a disk and a square of equal area in the plane are equidecomposable: can the disk be partitioned into finitely many pieces, and can those pieces be moved by isometries to partition the square?

§135.2 Equidecomposability and dissection congruence

Let a group G act on a set X . Two subsets $A, B \subseteq X$ are G -equidecomposable if there are finite partitions

$$A = A_1 \sqcup \cdots \sqcup A_m, \quad B = B_1 \sqcup \cdots \sqcup B_m,$$

and elements $g_i \in G$ such that

$$g_i(A_i) = B_i.$$

For plane polygons, one usually lets G be the group of Euclidean isometries. In that setting the classical language is dissection congruence: cut one polygon into finitely many polygonal pieces and rearrange them to obtain the other.

The first example is the familiar parallelogram-to-rectangle cut.




Figure 135.1: A parallelogram dissected and rearranged into a rectangle.

Area is obviously preserved by dissection congruence. The Wallace–Bolyai–Gerwien theorem says that, for plane polygons, area is also sufficient.

Theorem 135.2.1 (Wallace–Bolyai–Gerwien)

Any two plane polygons of the same area are dissection congruent.

One small point in the note is transitivity. If polygon P is dissection congruent to Q , and Q is dissection congruent to R , then P is dissection congruent to R . One refines the two dissections so that the intermediate polygon Q is cut compatibly, then composes the rearrangements.




Figure 135.2: The transitivity intuition: a dissection from P to Q and one from Q to R compose.

The proof of the theorem is constructive. First chop a polygon into triangles.





Figure 135.3: The first reduction: chop a polygon into triangles.

Then dissect each triangle into a parallelogram and then a rectangle; dissect each rectangle into a square; combine squares by the Pythagorean dissection.

Dissect each triangle into a parallelogram and then a rectangle



Dissect each rectangle into a square



Combine these squares using our Pythagorean proof.




Figure 135.4: The visual chain: triangle to parallelogram/rectangle, rectangles to squares, then combine squares.

The theorem establishes the baseline: for tame plane polygons, equal area exactly characterizes finite dissection congruence.

§135.3 Hilbert's third problem and the Dehn invariant

Hilbert's third problem asks whether equal volume similarly characterizes dissection congruence of three-dimensional polytopes. Dehn's theorem says no: a cube and a regular tetrahedron of the same volume are not dissection congruent.

The extra obstruction is the Dehn invariant. If a polyhedron has edge lengths ℓ_i and corresponding dihedral angles θ_i , the invariant has the form

$$D(P) = \sum_i \ell_i \otimes \theta_i$$

with values in a tensor product such as $\mathbb{R} \otimes_{\mathbb{Z}} \mathbb{R}/(2\pi\mathbb{Z})$, depending on conventions. The invariant records length-angle data that volume does not see. The handwritten note correctly marks the tensor product as a point worth noticing: this is not elementary area bookkeeping anymore.

Sydler's theorem says that volume and Dehn invariant together are complete in dimension three. Thus the two-dimensional polygon theorem has no naive three-dimensional analogue.

§135.4 Measure-theoretic background

The source page then discusses measure. Vitali sets imply that for every $n \geq 1$, there is no extension of Lebesgue measure to all subsets of $\mathbb{R}^n$ that is both countably additive and invariant under isometries.

If countable additivity is weakened to finite additivity, the dimension matters. For $n \geq 3$, the Banach–Tarski paradox shows that there is no such finitely additive isometry-invariant measure on all bounded sets compatible with volume. For $n \leq 2$, Banach measures exist. The note emphasizes the reason: in dimension at least three, the isometry group contains enough free-group behavior to allow paradoxical decompositions; in dimensions one and two it does not.

For Tarski's problem in the plane, this means that equal area is still a necessary condition. A disk and square that are equidecomposable by isometries must have the same Lebesgue measure, even if the pieces themselves are nonmeasurable.

§135.5 Tarski's circle-squaring problem

The central question is then:

Are a disk and a square in $\mathbb{R}^2$ of the same area equidecomposable?



Figure 135.5: The central visual question: equal-area disk and square.

At first sight the boundary seems to obstruct the construction. A disk has circular boundary; a square has straight boundary. This intuition is correct for a stricter notion, but not for arbitrary equidecomposition.

§135.6 Scissors congruence and why nice pieces fail

The note distinguishes equidecomposability from scissors congruence. In the scissors-congruence version, the pieces must be tame: each piece is homeomorphic to a disk and has boundary given by curves of finite length. Under such restrictions, the disk and square are not scissors congruent.

The invariant mentioned in the source is:

$$\text{convex circular perimeter} - \text{concave circular perimeter}.$$

When a curved circular arc is created inside a dissection, it is paired with an opposite concave arc, so the internal contributions cancel. The remaining outer boundary contribution is preserved. A disk has a nonzero circular-boundary contribution, while a square has none. Hence a disk and square cannot be scissors congruent in this tame sense.




Figure 135.6: Tame scissors pieces preserve circular-boundary information.

This is the first major subtlety: the positive answer to Tarski's problem requires pieces much wilder than ordinary scissors pieces.

§135.7 Laczkovich's positive theorem

Laczkovich proved in 1990 that Tarski's circle-squaring problem has a positive answer. In 1992 he proved a broader theorem. A simplified statement is:

Theorem 135.7.1 (Laczkovich)

Let $A, B \subseteq \mathbb{R}^k$ be bounded sets with the same positive Lebesgue measure. If the boundaries of A and B have upper Minkowski dimension strictly less than k , then A and B are equidecomposable by translations.

For a disk and square in the plane, the boundaries are one-dimensional, so the hypothesis applies.

The proof does not draw a literal jigsaw. It converts the problem into a bounded matching problem in a graph generated by translations.

§135.8 First idea: work in the torus

Scale and translate A and B so that they lie inside $[0, 1)^k$. Identify this cube with the k -torus

$$\mathbb{T}^k = (\mathbb{R}/\mathbb{Z})^k.$$

Then A and B are equidecomposable by translations as subsets of the torus iff, after possibly using more pieces, they are equidecomposable by translations in $\mathbb{R}^k$.




Figure 135.7: Working on the torus: translations wrap around the fundamental cube.

This torus reduction matters because translations become an action on a compact group. Compactness and periodic boundary conditions make the orbit structure manageable.

§135.9 Random translations and a free action

Choose a sufficiently large integer d and random translations

$$u_1, \dots, u_d \in \mathbb{T}^k.$$

They define an action of $\mathbb{Z}^d$ on $\mathbb{T}^k$ by

$$(n_1, \dots, n_d) \cdot x = n_1 u_1 + \dots + n_d u_d + x.$$

For almost every choice of the u_i , this action is free: no nonzero element of $\mathbb{Z}^d$ fixes a point. Thus each orbit can be visualized as a copy of $\mathbb{Z}^d$.




Figure 135.8: Random translations generate lattice-like orbits on the torus.

§135.10 The graph and bounded-distance bijection

Define a graph G with vertex set $\mathbb{T}^k$. Two points x, y are adjacent if there exists $g \in \mathbb{Z}^d$ with

$$g \cdot x = y, \quad \|g\|_\infty = 1.$$

So edges correspond to single generator moves.




Figure 135.9: The generator graph: locally it looks like a lattice with bounded moves.

To show A and B are equidecomposable, it suffices to find a Borel bijection

$$f : A \rightarrow B$$

of bounded distance in G : there is some fixed N such that

$$d_G(x, f(x)) \leq N \quad \text{for all } x \in A.$$

If such an f exists, then only finitely many group elements g can occur. Put

$$A_g = \{x \in A : f(x) = g \cdot x\}.$$

The sets A_g partition A , and the translated sets $g \cdot A_g$ partition B . This gives the desired finite equidecomposition.

This is the conceptual heart of the proof:

$$\text{finite translation equidecomposition} \iff \text{bounded-distance orbit matching}.$$

§135.11 Orbit matching

Inside a single orbit of the $\mathbb{Z}^d$ action, the problem becomes matching red A -points to blue B -points by bounded paths.




Figure 135.10: Within one orbit, equidecomposition becomes a matching problem.

A bounded matching can exist only if large finite regions of the orbit contain approximately the same number of A -points and B -points. This leads to discrepancy estimates.

§135.12 Discrepancy estimates

For an orbit point x , define a finite box

$$S_N(x) = \{(n_1, \dots, n_d) \cdot x : 0 \leq n_i < N\}.$$

Ergodic intuition suggests

$$|S_N(x) \cap A| \approx \lambda(A)N^d.$$

The handwritten note warns that this is only a heuristic: ordinary ergodic theorems alone are not strong enough. Laczkovich needs a uniform quantitative estimate from Diophantine approximation and discrepancy theory.




Figure 135.11: A finite orbit box $S_N(x)$; discrepancy compares observed counts with expected measure.

The key lemma states that there exist $\varepsilon > 0$ and M such that for all x and N ,

$$\left| |S_N(x) \cap A| - \lambda(A)N^d \right| \leq MN^{d-1-\varepsilon},$$

and similarly for B :

$$\left| |S_N(x) \cap B| - \lambda(B)N^d \right| \leq MN^{d-1-\varepsilon}.$$

Since $\lambda(A) = \lambda(B)$, large boxes contain nearly equal numbers of A and B points. The error term grows more slowly than the volume term N^d , which is exactly what makes matching possible.

§135.13 From Laczkovich to Borel circle squaring

Laczkovich's original proof used the axiom of choice. Later work by Marks and Unger showed that the pieces can be chosen Borel. This is much stronger: the pieces are still extremely complicated, but they are definable in the Borel hierarchy.




Figure 135.12: A visual impression of the Borel circle-squaring pieces.

The final handwritten note says the main components are: Dehn invariant, group action, torus embedding, traversal/matching, and measure/discrepancy theory. That summary is accurate. The proof is not one trick; it is a combination of several fields.

§135.14 Bridge to later ideas

The same question changes answer depending on the category of pieces. Plane polygons are controlled by area. Three-dimensional polyhedra require the Dehn invariant. Tame planar scissors pieces preserve circular-boundary data, so a disk cannot become a square. But arbitrary or Borel pieces are flexible enough for Laczkovich's theorem and Borel circle squaring.

Thus Tarski's circle-squaring problem is a crossroads of geometry, measure theory, group actions, discrepancy estimates, graph matching, and descriptive set theory. The important lesson is not only that the answer is yes, but why the yes lives far beyond ordinary scissors geometry.

§135.15 Scissors congruence and invariants

Polygon dissection problems ask whether one shape can be cut into finitely many pieces and rearranged into another. Area is an invariant, but in higher dimensions it is not enough; Hilbert's third problem needs the Dehn invariant.

§135.16 Measurable versus nonmeasurable decompositions

Tarski-style paradoxes show that arbitrary set-theoretic decompositions can violate geometric intuition. Requiring pieces to be measurable or Borel changes the problem drastically because measure becomes a meaningful invariant.

§135.17 Circle squaring and group actions

Laczkovich's theorem uses translations and carefully controlled pieces to equidecompose a disk and a square of equal area. The modern refinements ask for regularity of the pieces, such as Borel or measurable structure.

§135.18 Equidecomposition under measure and group action

Circle-squaring is not just a geometric puzzle. It sits at the meeting point of group actions, measure theory, regularity of sets, and invariants of dissection.

Reading Hitchin: Self-Duality Equations, Higgs Bundles, and Geometric Structures

N-20230801

§136.1 Starting point

This note is a reading reconstruction around Hitchin's paper on the self-duality equations on a Riemann surface. The main mathematical movement is a dimensional reduction from four-dimensional self-dual Yang–Mills theory to equations on a Riemann surface, producing what are now called Hitchin equations.

Let Σ be a compact Riemann surface, $E \rightarrow \Sigma$ a complex vector bundle, A a connection, and Φ a Higgs field, locally an $\text{End}(E)$ -valued $(1,0)$ -form. The Hitchin equations have the schematic form

$$F_A + [\Phi, \Phi^*] = 0, \quad \bar{\partial}_A \Phi = 0.$$

Here F_A is the curvature of A , and Φ^* is the adjoint of the Higgs field with respect to a Hermitian metric.

§136.2 Connections and curvature

A connection A is a way of differentiating sections of a bundle. Its curvature is

$$F_A = dA + A \wedge A,$$

which measures the failure of covariant derivatives to commute. Gauge transformations act on connections, and the equations are considered modulo this gauge symmetry.

Remark 136.2.1 — This already explains why the moduli space is not just a solution set. One must divide by gauge equivalence, so the actual object is a quotient, often with singularities and rich geometric structure.

§136.3 Higgs bundles

A Higgs bundle over a Riemann surface is roughly a pair

$$(E, \Phi),$$

where E is a holomorphic vector bundle and

$$\Phi : E \rightarrow E \otimes K_\Sigma$$

is a holomorphic K_Σ -twisted endomorphism. The field Φ is the Higgs field.

The equation

$$\bar{\partial}_A \Phi = 0$$

says that Φ is holomorphic with respect to the holomorphic structure determined by A .

§136.4 The moduli space

The moduli space of solutions to Hitchin's equations is deeply structured. It carries a hyperkähler metric and admits several interpretations:

- (i) as a moduli space of solutions to gauge-theoretic PDEs;
- (ii) as a moduli space of stable Higgs bundles;
- (iii) as a moduli space of flat connections or representations of the fundamental group.

This is the beginning of non-abelian Hodge theory. In rank one, Hodge theory relates differential forms, cohomology, and harmonic representatives. In higher rank, the analogous story involves Higgs bundles, flat connections, and representations.

§136.5 The Hitchin fibration

The characteristic polynomial of the Higgs field gives invariant functions. For example, in rank n , one considers

$$\mathrm{tr}(\Phi), \quad \mathrm{tr}(\Phi^2), \quad \dots, \quad \mathrm{tr}(\Phi^n).$$

These define the Hitchin map from the moduli space of Higgs bundles to an affine space of differentials, often called the Hitchin base.

The generic fibers are abelian varieties, related to spectral curves. Thus the Hitchin system is an algebraically completely integrable system.

§136.6 Spectral curves

The spectral curve is defined by the eigenvalue equation of the Higgs field. Informally, if λ denotes a coordinate on the cotangent direction, then the spectral curve is given by

$$\det(\lambda - \Phi) = 0.$$

It lives inside the total space of the canonical bundle K_Σ . A Higgs bundle can often be reconstructed from line bundle data on its spectral curve.

Remark 136.6.1 — This is one of the most attractive parts of the theory: a non-abelian problem about vector bundles and Higgs fields is transformed into abelian geometry on a spectral curve.

A concrete rank-two model makes the slogan much less vague. I take Σ to be the genus-two hyperelliptic curve

$$\Sigma : \quad y^2 = (x^2 - 1)(x^2 - 4)(x^2 - 9).$$

This is the promised genus-two surface in a form one can calculate with. I deliberately do not call it an elliptic curve or a one-dimensional complex torus, since those have genus one. A basis of holomorphic one-forms is

$$\omega_0 = \frac{dx}{y}, \quad \omega_1 = x \frac{dx}{y}.$$

Now set

$$E = \mathcal{O}_\Sigma \oplus \mathcal{O}_\Sigma$$

and choose a local coordinate z so that quadratic differentials are written as $Q(z) dz^2$. In companion-matrix form a rank-two Higgs field may be written locally as

$$\Phi = \begin{pmatrix} 0 & 1 \\ -Q_2 & -Q_1 \end{pmatrix} dz,$$

where $q_1 = Q_1 dz \in H^0(K)$ and $q_2 = Q_2 dz^2 \in H^0(K^2)$. For the simplest trace-free calculation I put $q_1 = 0$ and choose the explicit holomorphic quadratic differential

$$q_2 = (x^2 - 2) \left(\frac{dx}{y} \right)^2 \in H^0(\Sigma, K^2).$$

The zeros occur at $x = \pm\sqrt{2}$. Since these are not branch values of the hyperelliptic map, each value of x gives two points of Σ , so q_2 has four simple zeros, exactly as $\deg K^2 = 4g - 4 = 4$ predicts for $g = 2$.

The characteristic equation is now just a two-by-two determinant:

$$\begin{aligned} \det(\lambda I - \Phi) &= \det \begin{pmatrix} \lambda & -1 \\ Q_2 & \lambda + Q_1 \end{pmatrix} dz^2 \\ &= \lambda^2 + q_1 \lambda + q_2. \end{aligned}$$

With $q_1 = 0$, the spectral curve $\tilde{\Sigma} \subset \text{Tot}(K)$ is therefore

$$\eta^2 + q_2 = 0,$$

where η is the tautological one-form on the total space of K . Equivalently, after absorbing the harmless sign into the name of q_2 , one writes the same calculation as $\eta^2 = q_2(z)$. The projection

$$\pi : \tilde{\Sigma} \rightarrow \Sigma$$

is a double cover, branched over the four zeros of q_2 . Riemann–Hurwitz gives

$$2g(\tilde{\Sigma}) - 2 = 2(2g(\Sigma) - 2) + 4 = 2 \cdot 2 + 4 = 8,$$

so $g(\tilde{\Sigma}) = 5$.

The reconstruction is also concrete. If $\mathcal{L}$ is a line bundle on the spectral curve, then

$$E = \pi_* \mathcal{L}.$$

Multiplication by the tautological eigenvalue η gives

$$\mathcal{L} \longrightarrow \mathcal{L} \otimes \pi^* K,$$

and after applying π_* this becomes

$$\Phi : E = \pi_* \mathcal{L} \longrightarrow \pi_*(\mathcal{L} \otimes \pi^* K) \cong E \otimes K.$$

Thus the Higgs field is literally multiplication by the coordinate on the spectral cover. In the trace-free, fixed-determinant version relevant to this example, the Hitchin base is

$$H^0(K^2), \quad \dim H^0(K^2) = 3g - 3 = 3,$$

and the generic fiber is the Prym variety

$$\text{Prym}(\tilde{\Sigma}/\Sigma) = \ker \left(\text{Nm} : \text{Jac}(\tilde{\Sigma}) \rightarrow \text{Jac}(\Sigma) \right)^0.$$

If the determinant condition is removed one gets the full Jacobian of the spectral curve instead. This distinction is worth keeping: “abelian variety” is the broad slogan, while “Prym” is the sharper answer in the fixed-determinant rank-two case.

§136.7 Geometric Langlands perspective

Hitchin moduli spaces appear naturally in the geometric Langlands program. The Hitchin fibration organizes the moduli space into a structure where duality phenomena become visible. Mirror symmetry for Hitchin systems is one of the bridges between gauge theory, algebraic geometry, and representation theory.

§136.8 Synthesis

Hitchin's equations sit at a crossroads:

gauge theory Riemann surfaces Higgs bundles
integrable systems geometric Langlands.

The central idea is that reducing self-duality equations to a Riemann surface produces Hitchin's equations. Their moduli space remembers differential-geometric structure on one side and algebraic-geometric structure on the other.

§136.9 Moment-map reading of the equations

A useful way to read Hitchin's equations is through symplectic reduction. The space of pairs (A, Φ) is infinite-dimensional, and the unitary gauge group acts on it. The curvature equation

$$F_A + [\Phi, \Phi^*] = 0$$

can be interpreted as a zero moment-map condition for this action.

This viewpoint explains why stability appears. In many moduli problems, solving a differential equation modulo a compact group is equivalent to imposing an algebraic stability condition modulo a complexified group. For Higgs bundles, this is the Hitchin–Kobayashi style bridge.

§136.10 Spectral curves as global eigenvalues

For a single matrix, eigenvalues are roots of its characteristic polynomial. A Higgs field is a matrix-valued one-form varying over Σ . The spectral curve is the global version of taking eigenvalues:

$$\det(\lambda - \Phi) = 0.$$

It lives in the total space of the canonical bundle. Over a generic point of Σ , it records eigenvalues of Φ ; globally, these eigenvalues form a branched cover.

This is why the Hitchin fibration is powerful. The complicated nonabelian data of a Higgs bundle can often be transformed into abelian data: a line bundle on the spectral curve.

§136.11 Four descriptions to keep together

The reader should keep four translations active:

PDE solution $\leftrightarrow$ stable Higgs bundle $\leftrightarrow$ flat connection $\leftrightarrow$ representation of $\pi_1(\Sigma)$.

The difficulty of Hitchin's theory is precisely that none of these descriptions is disposable. The equations are analytic, the moduli problem is algebraic, the spectral curve is geometric, and the flat connection is representation-theoretic.

§136.12 Source repair: Yang–Mills background

The uploaded version begins from Yang–Mills theory. Let P be a principal G -bundle over a compact manifold M , and let A be a connection. Locally, A is a Lie-algebra-valued one-form, and its curvature is

$$F_A = dA + \frac{1}{2}[A, A].$$

The Yang–Mills functional is

$$S_{YM}(A) = \int_M \text{tr}(F_A \wedge *F_A).$$

Its Euler–Lagrange equation is

$$d_A * F_A = 0.$$

In four dimensions, two-forms decompose into self-dual and anti-self-dual parts:

$$\omega = \omega^+ + \omega^-, \quad *\omega^\pm = \pm\omega^\pm.$$

A self-dual connection satisfies

$$F_A = *F_A.$$

Such a connection automatically solves the Yang–Mills equation. This is the four-dimensional origin of the Hitchin system.

§136.13 Source repair: dimensional reduction

Hitchin’s equations can be understood as a dimensional reduction of the four-dimensional self-duality equations. Roughly, one splits the four-dimensional geometry into a two-dimensional Riemann surface direction and two suppressed directions. Components of the original connection in the suppressed directions become a Higgs field.

The resulting equations on a Riemann surface Σ are

$$F_A + [\Phi, \Phi^*] = 0, \quad \bar{\partial}_A \Phi = 0.$$

Here A is a connection on a bundle over Σ , Φ is an $\text{End}(E) \otimes K$ -valued Higgs field, and K is the canonical bundle.

The uploaded note emphasizes that Φ is not an arbitrary extra variable. It is the remnant of the higher-dimensional connection after reduction.

§136.14 Source repair: gauge invariance

Under a gauge transformation g , the Higgs field transforms as

$$\Phi \mapsto g\Phi g^{-1}, \quad \Phi^* \mapsto g\Phi^* g^{-1}.$$

Therefore

$$[\Phi, \Phi^*] \mapsto g[\Phi, \Phi^*]g^{-1}.$$

This explains why the bracket term is globally meaningful and compatible with quotienting by the gauge group. The curvature term transforms in the same conjugation pattern, so the first Hitchin equation is gauge invariant.

§136.15 Source repair: Higgs bundles and stability

A Higgs bundle is a pair

$$(E, \Phi), \quad \Phi \in H^0(\Sigma, \text{End}(E) \otimes K).$$

Stability is measured by slope:

$$\mu(E) = \frac{\deg E}{\text{rank } E}.$$

For an ordinary vector bundle, stability asks that every proper nonzero subbundle $F \subset E$ have

$$\mu(F) < \mu(E).$$

For Higgs bundles one restricts to Φ -invariant subbundles. This condition is essential because the moduli space behaves well only after excluding unstable objects.

This is the algebraic-geometric side of the Hitchin–Kobayashi philosophy: solving differential equations modulo unitary gauge transformations corresponds to imposing stability modulo complex gauge transformations.

§136.16 Source repair: spectral data and Hitchin fibration

The spectral curve is defined by the characteristic polynomial of the Higgs field:

$$\det(\lambda I - \Phi) = \lambda^n + a_1 \lambda^{n-1} + \cdots + a_n = 0.$$

The coefficients a_i are sections of powers of the canonical bundle. This gives the Hitchin map

$$h : \mathcal{M} \rightarrow \bigoplus_{i=1}^n H^0(\Sigma, K^i).$$

For a generic spectral curve, the fiber of this map is an abelian variety; in the rank-two fixed-determinant case above it is more precisely a Prym variety, while in the unconstrained GL_2 version it is the Jacobian of the spectral curve.

This is the mechanism by which a nonabelian moduli problem becomes an algebraically completely integrable system. The base records eigenvalue-like data; the fiber records line-bundle-like data on the spectral curve.

§136.17 Source repair: geometric Langlands and S-duality

The uploaded note connects Hitchin systems to geometric Langlands. The rough slogan is that local systems for the Langlands dual group ${}^L G$ correspond to sheaf-theoretic automorphic data on the moduli of G -bundles.

Hitchin moduli spaces provide the geometric stage on which this duality becomes visible. In the physical picture developed by Kapustin and Witten, S-duality in four-dimensional supersymmetric Yang–Mills theory becomes a geometric Langlands duality after reduction. Hitchin’s equations are the bridge between the gauge-theoretic and algebro-geometric languages.

§136.18 Source repair: local work still missing

The uploaded text names the right objects, but several parts still need to be made local and computable before the story is really understood.

- (i) In a local holomorphic coordinate z , Hitchin's equations should be written in terms of a connection form, a Higgs field φdz , and the adjoint $\varphi^* d\bar{z}$. Without this coordinate calculation, the equation $F_A + [\varphi, \varphi^*] = 0$ remains only symbolic.
- (ii) The tangent space to the moduli problem is controlled by a deformation complex. The missing calculation is to identify infinitesimal gauge transformations, infinitesimal deformations of (A, φ) , and the linearized equations in one complex.
- (iii) Stability is not just a label attached to Higgs bundles. It is the condition that removes bad orbit-closure behavior, so it is the algebro-geometric replacement for asking the gauge quotient to have reasonable points.
- (iv) The spectral curve construction should be checked in low rank by writing

$$\det(\eta - \text{varphi}) = 0$$

inside the total space of the canonical bundle. This makes the Hitchin base concrete: its coordinates are the coefficients of the characteristic polynomial.

A good next pass through this topic should therefore not try to memorize the whole nonabelian Hodge correspondence at once. It should work through one rank-two example, compute the characteristic polynomial of a Higgs field, locate the spectral curve, and then match the deformation complex with the expected dimension of the moduli space.

Plane Curves as Algebraic Creatures: Equations, Rings, and Singularities

N-20240608

§137.1 The curve is not only the drawing

A plane curve often begins as a picture. Algebraic geometry asks for something more durable: an equation.

Given a polynomial $f(x, y)$, define

$$C = V(f) = \{(x, y) : f(x, y) = 0\}.$$

The drawing of C may be helpful, but the equation controls its functions, singularities, and intersections.

A curve is therefore a small laboratory where geometry and algebra constantly translate into each other.

§137.2 The function ring of a curve

For $C = V(f)$, define

$$k[C] = k[x, y]/(f).$$

This is the coordinate ring of the curve. Two polynomials define the same function on C if their difference is a multiple of f .

For the parabola $C = V(y - x^2)$,

$$k[C] = k[x, y]/(y - x^2) \cong k[x].$$

So although the parabola is drawn curved, its polynomial function theory is the same as that of a line.




Figure 137.1: A parabola is geometrically curved in the plane, but its coordinate ring is as simple as $k[x]$.

§137.3 When factorization becomes geometry

If

$$f = gh,$$

then

$$V(f) = V(g) \cup V(h).$$

Thus factorization of equations corresponds to decomposition of curves.

Example:

$$xy = 0$$

is the union of the two coordinate axes. The equation is reducible, and the curve is reducible.

This is the first piece of the algebra-geometry dictionary: algebraic factorization can mean geometric splitting.

§137.4 Smooth points and singular points

A point $p = (a, b) \in V(f)$ is singular if

$$f(a, b) = 0, \quad f_x(a, b) = 0, \quad f_y(a, b) = 0.$$

At such a point, the linear part of the equation vanishes. The usual tangent line no longer comes from the first derivative.

The cusp

$$y^2 = x^3$$

has a singularity at the origin. The node

$$y^2 = x^2(x + 1)$$

also has a singularity at the origin, but it looks like two branches crossing. Both are singular, but they are not singular in the same way.

§137.5 Tangent cones: the first nonzero shadow

When the linear part vanishes, look at the lowest-degree nonzero part of f . If it is f_d , then

$$f_d = 0$$

is the tangent cone.

For

$$f(x, y) = y^2 - x^2 - x^3,$$

the lowest-degree part is

$$y^2 - x^2 = (y - x)(y + x).$$

So the tangent cone at the origin consists of the two lines

$$y = x, \quad y = -x.$$

This detects the two branches of the node.

§137.6 The cusp as a ring

For the cusp $C = V(y^2 - x^3)$, the parametrization

$$x = t^2, \quad y = t^3$$

gives a ring map

$$k[x, y]/(y^2 - x^3) \longrightarrow k[t], \quad x \mapsto t^2, \quad y \mapsto t^3.$$

Its image is $k[t^2, t^3]$, not all of $k[t]$. This failure to become the full polynomial ring reflects the singularity. The cusp is close to a line, but not quite a smooth line.

This example is more memorable than the definition: singularities change the algebra of functions.

§137.7 Intersections: one visible point may count twice

The line $y = 0$ and the parabola $y = x^2$ meet at the origin. Substituting $y = 0$ into $y = x^2$ gives

$$x^2 = 0.$$

The intersection has multiplicity 2. Geometrically, the line is tangent to the parabola.

This explains why intersection theory counts multiplicities. A tangent intersection is not the same as a transverse crossing.

§137.8 Projective closure repairs missing intersections

Affine curves may hide points at infinity. Homogenization fills them in. For

$$f(x, y) = y - x^2,$$

the homogenization is

$$F(x, y, z) = yz - x^2.$$

The projective closure is $V(F) \subset \mathbb{P}^2$. In the chart $z = 1$ it is the original parabola. At infinity $z = 0$, it contains the point

$$[0 : 1 : 0].$$

This is the same philosophy as projective geometry: complete the affine picture so that geometry stops losing information at the boundary of the chart.

§137.9 Bezout as a promise of clean accounting

A first form of Bezout's theorem says that two projective plane curves of degrees m and n meet in mn points, counted with multiplicity, over an algebraically closed field, provided they have no common component.

One should not first treat this as a formula to memorize. It is a statement that projective geometry and multiplicity create the correct accounting system for intersections.

§137.10 Normalization: repairing a singular curve

The cusp $y^2 = x^3$ is parametrized by

$$x = t^2, \quad y = t^3.$$

The parameter line with coordinate t is smooth. The map from the line to the cusp is almost one-to-one, but at the cusp it repairs the singularity by replacing the bad point with a smooth parameter.

Algebraically, the coordinate ring of the cusp is

$$k[t^2, t^3] \subset k[t].$$

Passing to $k[t]$ is a normalization. One can think of it as adding the missing function t , which the singular coordinate ring did not contain.

This is a useful slogan:

normalization separates or smooths hidden branches of a curve.

For a node, normalization separates the two crossing branches. For a cusp, it smooths a single branch with a bad tangent behavior.

§137.11 Resultants: eliminating a variable

Intersections can also be studied by eliminating variables. If two curves are given by

$$f(x, y) = 0, \quad g(x, y) = 0,$$

one may eliminate y to get a polynomial condition on x . This is the idea behind the resultant.

For example, intersect

$$y = x^2, \quad y = mx + b.$$

Substitution gives

$$x^2 - mx - b = 0.$$

The discriminant

$$m^2 + 4b$$

tells whether the line meets the parabola in two distinct points, one tangent point, or no real points. Over an algebraically closed field, the quadratic always has two roots counted with multiplicity.

This elementary calculation is a small preview of Bezout's theorem.

§137.12 A projective calculation with a line and a conic

Take the conic

$$xz - y^2 = 0$$

and the line

$$z = 0.$$

Substituting $z = 0$ into the conic gives

$$-y^2 = 0.$$

Thus $y = 0$, and the intersection point is

$$[1 : 0 : 0].$$

The square indicates multiplicity 2. Geometrically, the line at infinity is tangent to this projective conic.

This example is compact but important: projective closure not only adds the missing point; it also remembers how the curve touches the line at infinity.

§137.13 Real pictures and algebraically closed accounting

Over $\mathbb{R}$, the circle

$$x^2 + y^2 + 1 = 0$$

has no real points. Over $\mathbb{C}$, it has many points. This is why algebraic geometry often works over algebraically closed fields when stating clean intersection theorems.

Real drawings are valuable, but they can hide complex points. Bezout's theorem counts in the algebraic world where equations have all their roots.

§137.14 How to read a plane curve from its equation

For the examples in this note, the useful dictionary is concrete:

equation	zero set being studied
factorization	decomposition into components
coordinate ring	polynomial functions restricted to the curve
vanishing of F_x, F_y	possible singular point
homogenization	points added on the line at infinity

Thus a quick analysis of a plane curve should not stop after drawing its real graph. One first factors the polynomial if possible, then solves $F = F_x = F_y = 0$ to find singularities, and finally homogenizes if an intersection count looks wrong in the affine plane. The line-conic example shows why this last step matters: an intersection can disappear to infinity, and over $\mathbb{R}$ it can also be hidden by the absence of real points. Bezout's theorem becomes reliable only after the curve is viewed projectively and, when necessary, over an algebraically closed field.

Divisors on Curves: The Bookkeeping of Zeros, Poles, and Functions

N-20240914

§138.1 A ledger for meromorphic functions

A meromorphic function on a compact Riemann surface has zeros and poles. Instead of listing them in prose, we record them in a divisor:

$$D = \sum_p n_p [p].$$

Positive coefficients record zeros or permissions. Negative coefficients record poles or debts.

This is the central metaphor of the chapter:

a divisor is a ledger for allowed behavior at points.

§138.2 Principal divisors: the ledger of one function

For a nonzero meromorphic function f , its principal divisor is

$$(f) = \sum_p \text{ord}_p(f) [p].$$

On the Riemann sphere, the function $z - a$ has divisor

$$(z - a) = [a] - [\infty].$$

The function

$$\frac{z - a}{z - b}$$

has divisor

$$\left(\frac{z - a}{z - b} \right) = [a] - [b].$$

Zeros and poles balance:

$$\deg(f) = 0.$$

This is the simplest global constraint on rational functions.

§138.3 Reading a divisor as permission

Given a divisor D , define

$$L(D) = \{f \text{ meromorphic} : (f) + D \geq 0\} \cup \{0\}.$$

This is the vector space of functions whose poles are no worse than D allows.

For example, on $\mathbb{CP}^1$, take

$$D = n[\infty].$$

Then $L(D)$ consists of polynomials of degree at most n :

$$1, z, z^2, \dots, z^n.$$

So

$$\ell(D) = n + 1.$$

This example should be kept close. It is the friendly face of the general theory.

§138.4 Line bundles: permission becomes geometry

A line bundle is a family of one-dimensional vector spaces varying over a curve. Locally it looks like

$$U \times \mathbb{C},$$

but the local pieces may glue nontrivially.

A divisor D determines a line bundle $\mathcal{O}(D)$. Its sections are the geometric version of functions satisfying the permission rule encoded by D .

Thus divisors are not just notation for zeros and poles. They produce spaces of sections, and those spaces are the objects one can count.

§138.5 The canonical divisor

Holomorphic 1-forms also have zeros. Their divisor class is called the canonical divisor K . For a compact Riemann surface of genus g ,

$$\deg K = 2g - 2.$$

The formula already links analysis and topology. A differential form knows something about the number of handles of the surface.




Figure 138.1: A curve is simultaneously analytic and topological; divisors help transfer information between these views.

§138.6 Riemann–Roch without intimidation

The theorem says

$$\ell(D) - \ell(K - D) = \deg D + 1 - g.$$

It counts the functions allowed by D , with a correction term involving the canonical divisor.

On $\mathbb{CP}^1$, this becomes the elementary fact that polynomials of degree at most n form an $(n + 1)$ -dimensional space. In that case $g = 0$, and for $D = n[\infty]$ with $n \geq 0$, the correction term disappears.

So Riemann–Roch is deep, but its first example is not mysterious. It extends a familiar polynomial dimension count to arbitrary compact Riemann surfaces.

§138.7 Hurwitz as topology with branch points

If $f : X \rightarrow Y$ is a nonconstant holomorphic map of compact Riemann surfaces of degree d , then

$$2g_X - 2 = d(2g_Y - 2) + \sum_{p \in X} (e_p - 1).$$

The correction term records ramification. Locally, ramification means the map looks like

$$z \mapsto z^{e_p}.$$

Thus the formula says that global topology changes according to local branching.

§138.8 A divisor game on the sphere

Let us play the simplest possible game. On $\mathbb{CP}^1$, ask for meromorphic functions with at most a double pole at infinity and no other poles. The answer is

$$L(2[\infty]) = \text{span}\{1, z, z^2\}.$$

If we also require a zero at 0, we are asking for

$$L(2[\infty] - [0]).$$

The allowed functions are now

$$az + bz^2,$$

so the dimension drops from 3 to 2.

This is how divisors should feel: adding a pole allowance increases freedom, while imposing a zero condition removes freedom.

§138.9 Degree as a first estimate

The degree of a divisor

$$D = \sum_p n_p [p]$$

is

$$\deg D = \sum_p n_p.$$

A positive degree suggests more allowed poles than required zeros, so one expects more functions. A negative degree usually means the conditions are too strict.

This intuition is made precise in many cases. On a compact Riemann surface, if $\deg D < 0$, then

$$L(D) = 0.$$

Indeed, a nonzero meromorphic function has principal divisor of degree 0, so $(f) + D \geq 0$ would force a nonnegative divisor of negative total degree, impossible.

§138.10 Canonical divisors through examples

On $\mathbb{CP}^1$, the differential dz has a double pole at infinity. In the coordinate $w = 1/z$,

$$dz = -w^{-2} dw.$$

Thus the canonical divisor has degree -2 , matching

$$2g - 2 = -2$$

for genus 0.

On a torus, $g = 1$, so $\deg K = 0$. This matches the existence of a nowhere-vanishing holomorphic differential. The formula $\deg K = 2g - 2$ is therefore not merely symbolic; it reflects the geometry of different surfaces.

§138.11 Why Riemann–Roch has a correction term

The naive guess would be that the dimension of $L(D)$ is roughly

$$\deg D + 1 - g.$$

Riemann–Roch says this guess is corrected by

$$\ell(K - D).$$

This correction measures hidden differentials compatible with the opposite divisor condition. It is the reason the theorem is both a counting formula and a duality statement.

The formula therefore has two personalities:

$$\begin{array}{l|l} \text{counting side} & \ell(D) \text{ counts functions} \\ \text{duality side} & \ell(K - D) \text{ counts differential obstructions} \end{array}$$

§138.12 Final caution

Divisors are useful because they turn local analytic behavior into algebraic data. A zero is not merely a point where a function vanishes; it is an entry in a global accounting system. A pole is not merely a blow-up; it is a controlled permission.

The core question behind this chapter is:

$$\text{given allowed zeros and poles, which functions can exist?}$$

Riemann–Roch is the grand answer to that question.

Zariski Geometry: When Functions Become the Space

N-20241026

§139.1 Starting from functions, not pictures

In elementary geometry, one first draws a space and then studies functions on it. Algebraic geometry often reverses the direction: start with polynomial functions, and let the space be what those functions define.

Given an ideal

$$I \subset k[x_1, \dots, x_n],$$

we define the zero set

$$V(I) = \{a \in k^n : f(a) = 0 \text{ for all } f \in I\}.$$

Conversely, a set $X \subset k^n$ has an ideal of functions vanishing on it:

$$I(X) = \{f : f(x) = 0 \text{ for all } x \in X\}.$$

This back-and-forth is the first algebra-geometry dictionary.

§139.2 Coordinate rings: a space through its functions

If $X = V(I)$, the coordinate ring is

$$k[X] = k[x_1, \dots, x_n]/I(X).$$

It is the ring of polynomial functions on X .

For the parabola $X = V(y - x^2)$,

$$k[X] = k[x, y]/(y - x^2) \cong k[x].$$

For the hyperbola $X = V(xy - 1)$,

$$k[X] = k[x, y]/(xy - 1) \cong k[x, x^{-1}].$$

The hyperbola is algebraically the affine line with the origin removed. The drawing is curved, but the function ring reveals a simpler identity.

§139.3 Nullstellensatz: the dictionary becomes reliable

Over an algebraically closed field, Hilbert's Nullstellensatz says

$$I(V(I)) = \sqrt{I}.$$

This means that polynomial equations and radical ideals determine each other.

The theorem is not just a technical result. It is a guarantee that algebraic information can faithfully describe geometric zero sets, once nilpotent ambiguity is removed.

§139.4 The Zariski topology is coarse on purpose

Closed sets in the Zariski topology are algebraic sets. On the affine line, proper closed sets are finite. Therefore nonempty open sets are very large.

This topology is not designed to model physical closeness. It is designed to model algebraic dependence. If a polynomial vanishes on a large enough open set, it is forced to vanish much more widely than a continuous function would be.




Figure 139.1: Pictures can guide intuition, but the real geometry is carried by equations and functions.

§139.5 Distinguished opens and localization

For $f \in k[X]$, the distinguished open set $D(f)$ is the region where $f \neq 0$. On this region, f should be invertible.

For example, in $\mathbb{A}^1$, the open set $D(x)$ has ring of functions

$$k[x, x^{-1}].$$

This is localization. It is the algebraic operation of allowing denominators that do not vanish on the open set.

This idea is one of the most important technical moves in algebraic geometry: local geometry is controlled by localized rings.

§139.6 Sheaves: local data with a gluing rule

Regular functions are local. A function may be represented by different fractions on different open sets, as long as they agree on overlaps.

A sheaf formalizes this behavior. If a space is covered by opens U_i , and functions f_i on U_i agree on every overlap $U_i \cap U_j$, then they glue to a unique function on the union.

This is not abstraction for its own sake. It is the ordinary behavior of functions, stated carefully enough to work in algebraic geometry.

§139.7 Stalks: looking through a microscope

The stalk at a point stores the local behavior near that point. For regular functions, a stalk consists of germs of functions: two functions are the same germ if they agree on some neighborhood of the point.

In algebraic geometry, stalks are often local rings. They are the algebraic microscopes through which one studies smoothness, singularities, and local equations.

§139.8 A local example: two axes crossing

Consider the curve

$$X = V(xy) \subset \mathbb{A}^2.$$

Geometrically it is the union of the two coordinate axes. Algebraically its coordinate ring is

$$k[X] = k[x, y]/(xy).$$

Away from the origin, each branch looks like an ordinary line. At the origin, the two branches meet, and the local ring remembers both directions.

On the open set $D(x)$, the element x is invertible. Since $xy = 0$, multiplying by x^{-1} gives $y = 0$. Thus on $D(x)$, only the x -axis branch remains. Similarly, on $D(y)$, only the y -axis branch remains.

This is a small but useful lesson: localization is not just algebraic manipulation. It literally focuses attention on the part of the variety where a chosen function does not vanish.

§139.9 Why the strange topology helps

The Zariski topology feels strange because its open sets are large. But large opens are exactly what make polynomial continuation possible. If two regular functions agree on a nonempty open subset of an irreducible variety, they often must agree much more widely.

This is why the topology is useful for algebraic geometry even though it is poor for measuring distance. It is built for rigidity, not for analysis.

§139.10 Products and unions through rings

Algebraic operations on rings often reverse geometric operations. The product of equations gives a union:

$$V(fg) = V(f) \cup V(g).$$

The sum of ideals gives an intersection:

$$V(I + J) = V(I) \cap V(J).$$

This reversal is a persistent theme. Geometry and algebra face opposite directions: more equations mean smaller spaces, while larger ideals mean fewer points.

For instance, in $k[x, y]$, the ideals (x) and (y) define the two coordinate axes. Their sum (x, y) defines the origin, the intersection of the two axes.

§139.11 Radical ideals and forgotten nilpotents

The ideals (x) and (x^2) have the same classical zero set in $\mathbb{A}^1$: both vanish only at 0. But the quotient rings differ:

$$k[x]/(x), \quad k[x]/(x^2).$$

The second ring contains a nonzero nilpotent element.

Classical varieties often pass to radical ideals and forget nilpotents. Schemes later restore them. This is one reason the transition from varieties to schemes is not artificial: nilpotent structure is already waiting in simple equations.

§139.12 Regular functions versus continuous functions

On a usual topological space, continuous functions are flexible. They can be changed locally with bump functions. Regular functions in algebraic geometry are rigid. They are controlled by polynomial or rational formulas.

This rigidity is why algebraic geometry can extract strong global conclusions from local algebra. The price is that the topology becomes coarse and unintuitive from an analytic point of view.

§139.13 Morphisms as pullback of functions

A map of affine varieties

$$F : X \rightarrow Y$$

induces a pullback of coordinate rings

$$F^* : k[Y] \rightarrow k[X].$$

A function on Y becomes a function on X by composition with F .

This reverses arrows:

$$X \longrightarrow Y \quad \rightsquigarrow \quad k[Y] \longrightarrow k[X].$$

This arrow reversal is one of the conceptual gates into modern algebraic geometry. Spaces are studied through their function rings, and maps of spaces become maps of rings in the opposite direction.

§139.14 A compact dictionary

The essential translations are:

geometric phrase	algebraic phrase
zero set	$V(I)$
functions on X	$k[X]$
remove $f = 0$	localize at f
local behavior at p	stalk at p
glue local functions	sheaf condition

The point of Zariski geometry is not to make topology strange. It is to let functions define the correct topology for polynomial geometry.

Schemes Gently: A Guided Tour of Spectrum, Nilpotents, and Local Rings

N-20250118

§140.1 A new kind of point

Classical algebraic geometry studies points satisfying polynomial equations. Scheme theory asks us to enlarge the meaning of point.

For a ring R , define

$$\operatorname{Spec} R = \{\mathfrak{p} \subset R : \mathfrak{p} \text{ is prime}\}.$$

This is the spectrum of R . Its points are prime ideals.

At first this feels like a category mistake: ideals are algebraic objects, not points. But this change is exactly what allows geometry to remember generic behavior, arithmetic information, and infinitesimal thickening.

§140.2 The affine line spectrum

For $R = k[x]$, the maximal ideals


$$(x - a)$$

correspond to ordinary points $a \in k$. But there is also the prime ideal (0) , the generic point.

The closure of the generic point is the whole space:

$$\overline{\{(0)\}} = \operatorname{Spec} k[x].$$

This is not Hausdorff behavior, but it is meaningful. The generic point represents the whole irreducible line, while $(x - a)$ represents a closed point.



The arrows are suggestive: the generic point specializes to closed points.

§140.3 Closed sets and distinguished opens

For an ideal $I \subset R$, define

$$V(I) = \{\mathfrak{p} \in \operatorname{Spec} R : I \subseteq \mathfrak{p}\}.$$

These are closed sets.

For $f \in R$, define

$$D(f) = \{\mathfrak{p} : f \notin \mathfrak{p}\}.$$

This is the region where f is allowed to be inverted. Distinguished opens are the basic open sets of affine schemes.

§140.4 Localization is zooming in

At a prime ideal $\mathfrak{p}$, the local ring is

$$R_{\mathfrak{p}}.$$

It is obtained by inverting every element not in $\mathfrak{p}$. The interpretation is simple: near $\mathfrak{p}$, functions not vanishing there should be units.

For $R = k[x]$ and $\mathfrak{p} = (x - a)$, the ring

$$k[x]_{(x-a)}$$

contains rational functions whose denominators do not vanish at a . This is the algebraic version of functions defined near a .

§140.5 The structure sheaf

A spectrum is not only a topological space. It carries functions. On a distinguished open $D(f)$, the structure sheaf satisfies

$$\mathcal{O}_{\mathrm{Spec} R}(D(f)) = R_f.$$

The stalk at $\mathfrak{p}$ is

$$\mathcal{O}_{\mathrm{Spec} R, \mathfrak{p}} = R_{\mathfrak{p}}.$$

These two formulas are enough to make affine schemes computational rather than mystical.

§140.6 Nilpotents: invisible but geometric

Consider

$$R = k[x]/(x^2).$$

The element x is nonzero but satisfies $x^2 = 0$. Topologically, $\mathrm{Spec} R$ has only one point. Algebraically, it remembers an infinitesimal thickening of that point.

A useful picture is a point with a first-order tangent fuzz. The fuzz cannot be seen as a second classical point, but it can be detected by functions.

This is one of the reasons schemes are powerful: they keep information that ordinary point sets discard.

§140.7 The arithmetic curve

The spectrum of the integers has points

$$(p)$$

for prime numbers p , together with the generic point (0) . An integer n vanishes at exactly the primes dividing it:

$$V(n) = \{(p) : p \mid n\}.$$

Thus divisibility becomes geometry.

This is one of the most charming features of schemes: number theory begins to resemble geometry, not metaphorically but structurally.

§140.8 Maximal ideals are not enough

It is tempting to keep only maximal ideals, because over an algebraically closed field they correspond to ordinary closed points. But prime ideals contain information about irreducible subspaces.

In $k[x, y]$, the prime ideal (x) is not a closed point. It corresponds to the whole y -axis. The maximal ideal $(x, y - a)$ is a closed point on that axis.

Thus prime ideals let a scheme contain both points and subvarieties as points of different levels of generality. This is strange at first, but it is exactly what makes generic points possible.

§140.9 Specialization as movement from generic to concrete

In $\text{Spec } R$, one says that $\mathfrak{q}$ specializes to $\mathfrak{p}$ when

$$\mathfrak{q} \subseteq \mathfrak{p}.$$

The smaller prime is more generic; the larger prime is more special.

For $k[x]$,

$$(0) \subset (x - a),$$

so the generic point specializes to the closed point $(x - a)$. This order relation is the topological shadow of ideal containment.

§140.10 Residue fields

At a prime $\mathfrak{p}$, the local ring $R_{\mathfrak{p}}$ has maximal ideal $\mathfrak{p}R_{\mathfrak{p}}$. The residue field is

$$\kappa(\mathfrak{p}) = R_{\mathfrak{p}}/\mathfrak{p}R_{\mathfrak{p}}.$$

For a closed point $(x - a) \subset k[x]$, the residue field is k . For the generic point (0) , the residue field is $k(x)$, the field of rational functions.

This is another way to see the meaning of generic points: at the generic point of the affine line, the coordinate x is not specialized to a number; it remains a variable.

§140.11 A tiny nonreduced calculation

Let

$$R = k[\epsilon]/(\epsilon^2).$$

A function on $\text{Spec } R$ can have the form

$$a + b\epsilon.$$

The term $b\epsilon$ is invisible if one only looks at ordinary points, because ϵ is nilpotent. But algebraically it records first-order infinitesimal information.

This is the simplest model of a dual number. It is also a first hint of why tangent vectors and infinitesimal thickenings are naturally scheme-theoretic.

§140.12 Reading $\operatorname{Spec} R$ in examples

The dictionary from this note should be used with actual rings in mind:

$\operatorname{Spec} R$	prime ideals, including non-maximal ones
$D(f)$	the part where f becomes invertible
R_f	functions allowed to divide by powers of f
$R_{\mathfrak{p}}$	functions near the prime $\mathfrak{p}$
$\kappa(\mathfrak{p})$	the residue field visible at $\mathfrak{p}$
nilpotents	infinitesimal information not seen by closed points

For $\operatorname{Spec} \mathbb{Z}$, the generic point (0) and the closed points (p) are both needed: localizing at (p) records arithmetic near one prime, while the generic point records information over $\mathbb{Q}$. For $k[\epsilon]/(\epsilon^2)$, there is only one underlying prime, but the function $a + b\epsilon$ still contains a first-order term. These two examples explain why schemes keep more than the set of maximal ideals: they retain specialization, residue fields, and nilpotent thickening.

Projective Schemes and Gluing: Building Projective Space from Affine Windows

N-20250301

§141.1 Projective geometry returns with functions attached

Earlier projective geometry used homogeneous coordinates to add points at infinity and remove affine exceptions. Scheme theory adds another layer: every affine window comes with a ring of functions, and the windows must be glued compatibly.

The construction that does this is Proj. It is to graded rings what Spec is to ordinary rings.

§141.2 The projective line as the model

The best first example is $\mathbb{P}^1$. Take two affine lines:

$$U_0 = \operatorname{Spec} R[t], \quad U_1 = \operatorname{Spec} R[s].$$

On the region where both coordinates are nonzero, glue them by

$$s = \frac{1}{t}.$$

The result is the projective line. The point is not the formula itself, but the construction style:

projective space is made by gluing affine charts.

§141.3 Graded rings and homogeneous data

Let

$$S = S_0 \oplus S_1 \oplus S_2 \oplus \cdots$$

be a graded ring. Projective geometry only respects homogeneous equations, because

$$[x_0 : \cdots : x_n] = [\lambda x_0 : \cdots : \lambda x_n].$$

The construction $\operatorname{Proj} S$ uses homogeneous prime ideals not containing the irrelevant ideal.

This definition sounds heavier than the picture, but it encodes a familiar rule: equations in projective space must be homogeneous.

§141.4 Projective space as Proj

The scheme-theoretic projective space is

$$\mathbb{P}_R^n = \operatorname{Proj} R[x_0, \dots, x_n],$$

where the variables have degree 1.

The standard affine open where $x_i \neq 0$ is written

$$D_+(x_i).$$

On it, the ratios

$$\frac{x_j}{x_i}$$

serve as affine coordinates.

Thus $\mathbb{P}^n$ is covered by $n + 1$ affine schemes. Projective geometry is global, but it is computed through affine windows.

§141.5 The projective plane through three windows

For $\mathbb{P}^2$, the standard opens are

$$D_+(x), \quad D_+(y), \quad D_+(z).$$

On $D_+(z)$, use

$$X = \frac{x}{z}, \quad Y = \frac{y}{z}.$$

On $D_+(x)$, use

$$U = \frac{y}{x}, \quad V = \frac{z}{x}.$$

Each chart is affine. The projective plane is the result of gluing these affine pieces by ratio changes on overlaps.

This is the practical meaning of projective schemes: use affine algebra locally, then glue.

§141.6 A conic across charts

Consider the projective conic

$$xz - y^2 = 0.$$

On the chart $z \neq 0$, set $X = x/z$ and $Y = y/z$. The equation becomes

$$X = Y^2.$$

On the chart $x \neq 0$, set $U = y/x$ and $V = z/x$. The equation becomes

$$V = U^2.$$

The same projective curve has multiple affine faces. A single affine chart may hide what happens at infinity, but the projective object keeps all charts at once.

§141.7 Projective closure with scheme memory

Given an affine equation $f(x, y) = 0$, homogenizing gives a projective equation

$$F(x, y, z) = 0.$$

Classically, this adds points at infinity. Scheme-theoretically, it also keeps track of nilpotent or embedded structure if such structure is present.

So Proj is not merely a notation for old projective varieties. It is the projective version of the local-function viewpoint developed in scheme theory.

§141.8 What Proj adds

The earlier projective note emphasized incidence and transformations. The Proj view-point emphasizes functions and gluing.

homogeneous coordinates	homogeneous primes
affine charts	standard affine opens
projective equations	homogeneous ideals
points at infinity	glued affine schemes

These are not competing languages. They are layers of the same geometry.

§141.9 Why the irrelevant ideal is excluded

For $S = R[x_0, \dots, x_n]$, the ideal

$$S_+ = (x_0, \dots, x_n)$$

is called the irrelevant ideal. It corresponds to the forbidden point where all homogeneous coordinates vanish at once:

$$[0 : \dots : 0],$$

which is not a projective point.

Thus $\text{Proj } S$ uses homogeneous prime ideals that do not contain all positive-degree elements. This technical-looking exclusion is simply the algebraic version of the rule that homogeneous coordinates cannot all be zero.

§141.10 Standard opens as spectra

For a homogeneous element $f \in S$, the standard open $D_+(f)$ is affine. Its coordinate ring is the degree-zero part of the localization:

$$\mathcal{O}(D_+(f)) = (S_f)_0.$$

This formula explains why ratios appear in projective coordinates. Localizing at x_i allows division by powers of x_i , and taking degree zero keeps only scale-invariant expressions such as

$$\frac{x_j}{x_i}.$$

So affine coordinates on projective charts are not arbitrary. They are degree-zero fractions.

§141.11 A second conic chart calculation

For the conic

$$xz - y^2 = 0,$$

the chart $y \neq 0$ uses coordinates

$$A = \frac{x}{y}, \quad B = \frac{z}{y}.$$

The equation becomes

$$AB = 1.$$

So in this chart the conic looks like a hyperbola. In another chart it looked like a parabola. This is a concrete reminder that affine appearances depend on the chosen window.

The projective curve is the object obtained by holding all these windows together.

§141.12 Line bundles are already nearby

Projective geometry naturally leads to line bundles. On $\mathbb{P}^n$, the twisting sheaves $\mathcal{O}(d)$ organize homogeneous polynomials of degree d . A homogeneous equation of degree d is better thought of as a section of $\mathcal{O}(d)$ than as an ordinary global function.

This explains why projective varieties have fewer global regular functions than affine varieties. Projective geometry shifts attention from functions to sections of line bundles.

This remark points back to divisors and Riemann–Roch: projective geometry, line bundles, and spaces of sections are parts of the same story.

§141.13 Closing the geometry arc

This chapter is a suitable endpoint for the Geometry volume because it returns to a familiar object, projective space, with a richer language.

The compact summary is:

Proj is projective geometry rebuilt from homogeneous algebra and affine gluing.

It keeps the intuitive projective picture while adding the scheme-theoretic memory of local functions.

Nikonov's Seminar Prelude: Knots, Diagrams, Probes, and the Space Between the Lines

N-20250705

§142.1 First orientation

This note has a slightly unusual origin. It was not written after a finished course or after a fully digested paper. It grew out of a post-Gaokao attempt to understand an online seminar by I. M. Nikonov on partial tribrackets. The seminar was exciting, but the words arrived too quickly: quasigroups, biquandloids, thickened surfaces, crossings, arcs, regions, and probes all appeared before I had a stable picture of what they were supposed to mean.

The first reading memo I made after the seminar was useful as a record of confusion. Its main focus was correct: it noticed that Nikonov's paper is not merely introducing another knot invariant, but trying to explain why the pieces of a knot diagram should have topological meaning. However, that memo was too blog-like and too top-heavy for a mathematical notebook. It mixed motivation, informal analogies, and technical claims in one stream. Here I rewrite the same impulse into a more usable note.

The rule of this note is therefore modest: do not try to read the whole paper. Instead, understand the first conceptual move. Nikonov starts with an old tension in knot theory:

a knot is a topological object, but we calculate using diagrams.

The paper asks whether the diagrammatic pieces themselves—arcs, semiarcs, regions, and crossings—can be described topologically.

§142.2 The topological object and its planar shadow

Topologically, a knot is not first of all a picture on paper. It is an embedded circle in three-dimensional space:

$$K : S^1 \hookrightarrow \mathbb{R}^3,$$

or, more generally, a tangle in the thickening $F \times I$ of a surface F . The equivalence relation is ambient isotopy: the whole surrounding space is allowed to deform, carrying one knot to the other.

The story begins with the familiar planar shadow of a knot. A diagram is what we can draw, but the actual knot lives one dimension higher.




Figure 142.1: From Nikonov’s paper: a knot diagram. It is a planar record of a three-dimensional object.

This distinction matters. The diagram has crossings, but the knot in space has no literal self-intersection. At a crossing, one strand passes over the other. The diagram remembers this by marking over/under information.

§142.3 Some basic knot-theory language

A few standard definitions make the paper much less mysterious.

Definition 142.3.1 (Knot and ambient isotopy). A knot is a tame embedding $K : S^1 \hookrightarrow S^3$, usually drawn in $\mathbb{R}^3 \subset S^3$. Two knots K_0, K_1 are equivalent if there is an ambient isotopy $H_t : S^3 \rightarrow S^3$ with $H_0 = \text{id}$ and $H_1(K_0) = K_1$.

The word “ambient” is important. We do not merely slide points of the curve along the curve. We deform the whole surrounding three-dimensional space. This is why a knot is a topological object and not just a parametrized curve.

Definition 142.3.2 (Tangle in a thickened surface). Let F be an oriented compact surface and $I = [0, 1]$. A tangle in $F \times I$ is a properly embedded compact one-dimensional manifold, possibly with boundary on $\partial F \times I$. A knot is the closed one-component special case.

The thickened-surface language is the natural setting of Nikonov’s paper. Instead of drawing everything in the plane, one works over a surface F , and the knot or tangle floats in the slab $F \times I$.




Figure 142.2: From Nikonov’s paper: a tangle diagram on a surface.

Definition 142.3.3 (Diagram). A diagram of a tangle is the projection of the tangle to F , with crossing data recording which strand passes over the other. Thus a diagram is a two-dimensional shadow plus over/under information.

The basic theorem justifying all diagrammatic knot theory is Reidemeister's theorem.

Theorem 142.3.4 (Reidemeister theorem)

Two tame knot diagrams represent ambient isotopic knots if and only if they are related by a finite sequence of planar isotopies and Reidemeister moves.

This theorem is the bridge between topology and combinatorics. It says that if a construction on diagrams is unchanged under the three Reidemeister moves, then it really defines a knot invariant.

The combinatorial definition says that a knot can be represented by a diagram modulo Reidemeister moves. These are the local moves that change the drawing without changing the underlying knot.




Figure 142.3: From Nikonov's paper: Reidemeister moves. These are the local equivalences that make diagrams represent knots.

So the first dictionary is:

$$\text{knot in space} \longleftrightarrow \text{diagram modulo Reidemeister moves.}$$

This is powerful because diagrams are finite and computable. But it also creates the problem that motivates the paper: if a calculation assigns labels to arcs or regions of a diagram, what exactly are those arcs and regions as topological objects?

A few famous theorems live on the topological side of this dictionary. They are not the main subject of the note, but they explain the background against which the diagrammatic approach became so attractive.

Definition 142.3.5 (Knot complement and knot group). The complement of a knot $K \subset S^3$ is $S^3 \setminus N(K)$, where $N(K)$ is a small tubular neighbourhood of the knot. The knot group is

$$\pi_1(S^3 \setminus K),$$

or equivalently the fundamental group of the complement after removing a tubular neighbourhood.

The complement is often more informative than the curve itself. A loop in the complement asks: can this loop be contracted to a point without crossing the knot? The knot group packages the obstruction algebraically.

Theorem 142.3.6 (Schubert prime decomposition, informal form)

Every nontrivial knot decomposes as a connected sum of prime knots, and this decomposition is unique up to order.

Here the connected sum operation is the knot-theoretic analogue of multiplication: cut out a small arc from each knot and splice the loose ends together. Prime knots are the knots that cannot be decomposed further.

Another landmark is the geometric decomposition picture. Very roughly, many knot complements split into pieces that are torus-like, satellite-like, or hyperbolic. The figure-eight knot is the standard first example of a hyperbolic knot. This is the “God’s-eye” topological world mentioned in the original reading memo: elegant, deep, and hard to compute with directly.

§142.4 The pieces of a diagram

A diagram is not just a curve with crossings. It has several kinds of local pieces. Four basic types are especially important here:

arcs, semiarcs, crossings, regions.

These are the pieces that appear in coloring rules, parity theories, skein relations, and many diagrammatic constructions.




Figure 142.4: From Nikonov’s paper: the elementary pieces of a knot diagram.

The naive view is that these are pieces of ink on the page. An arc is a piece of curve between undercrossings. A semiarc is a smaller piece between consecutive crossings. A region is a component of the complement of the diagram. A crossing is a vertex with over/under data.

This is good enough for drawing, but it is not yet good enough for topology. A Reidemeister move can create or destroy a crossing, split an arc, merge two regions, or change how pieces correspond. If one wants to build invariants by labeling these pieces, one needs a way to say when a piece before the move is the same as a piece after the move.

This is exactly where many elementary knot invariants already give hints. For example, the linking number of a two-component link is computed from signs of mixed crossings. A quandle coloring assigns algebraic labels to arcs. A shadow or region coloring assigns labels to complementary regions. These methods work only because the labels transform correctly under Reidemeister moves.

So there are two levels:

first choose pieces of a diagram
 $\Downarrow$
 then assign algebraic labels to those pieces
 $\Downarrow$
 prove the labels survive Reidemeister moves.

Nikonov's paper attacks the first level. Before asking which labels to assign, ask what the pieces are.

This is the hidden question behind the paper:

Can a diagram element be defined without choosing a particular diagram?

§142.5 Nikonov's bridge: diagram elements as probes

Nikonov's answer is to replace a two-dimensional piece of a diagram by a three-dimensional path in a thickened surface. Work in

$$F \times I, \quad I = [0, 1],$$

where the knot or tangle lies between the bottom surface $F \times \{0\}$ and the top surface $F \times \{1\}$. A diagram is obtained by projecting down to F .

The key idea is that a diagram element can be detected by a probe: a path travelling through the thickened surface and interacting with the knot in a controlled way.




Figure 142.5: From Nikonov's paper: topological interpretation of diagram elements by probes in $F \times I$.

The slogan is:

diagram element	topological probe
arc	a path from the knot to the top surface
semiarc	a bottom-to-top path meeting the knot once
crossing	a bottom-to-top path meeting the knot twice
region	a bottom-to-top path avoiding the knot

This is the first place where the seminar became more understandable to me. The word “arc” is no longer merely a segment in a planar picture. It is represented by a way of reaching the knot from the ambient three-dimensional space. The word “region” is no longer just an empty component on paper. It is represented by a path through the complement of the knot.




Figure 142.6: From Nikonov’s paper: probe diagrams for arc, semiarc, region, and crossing elements.

The great advantage of this formulation is that isotopy becomes natural. If the tangle moves in $F \times I$, the probes can move with it. Thus the correspondence of diagram elements under Reidemeister moves is no longer a purely combinatorial bookkeeping rule; it is the shadow of a topological motion.

§142.6 Invariants, coinvariants, and labels

The first seminar memo paid special attention to the distinction between invariants and coinvariants. This distinction is worth keeping, but it is better to phrase it in a cleaner way.

An invariant is something that gives the same value for every diagram of the same knot. For example, a polynomial invariant or a numerical invariant should not depend on which diagram we draw.

A coinvariant is subtler. It may give different-looking objects for different diagrams, but these objects are canonically isomorphic once we account for the moves between diagrams. A vector space is a good mental model: its dimension is an invariant, while the vector space itself, together with the way it changes under coordinate changes, is closer to a coinvariant.

For diagram elements, the relevant object is often not one number but a whole set of possible classes. The set of arc probes, region probes, or crossing probes should be thought of as a universal object from which more concrete labels can be obtained.

In this language, Nikonov's first main conceptual claim is:

isotopy classes of probes give universal homotopical coinvariants of diagram elements.

This sentence is dense, but the intuition is simple. Instead of assigning a label to an arc immediately, first ask what the arc is intrinsically. The answer is a class of probes. Any later coloring rule is a way of reading algebraic information from these probe classes.

Remark 142.6.1 — The word “universal” should be read in the usual algebraic way. A universal object is not one more label among many labels. It is the object from which all compatible labels should factor. Thus the universal arc object is like the raw space of possible arc labels before choosing a particular quandle, polynomial, or coefficient set.

This is also why the paper talks about coinvariants rather than only invariants. A set of probes may change its concrete presentation when the diagram changes, but the induced object is naturally identified through the isotopy or through the homotopy class of Reidemeister moves. The object is stable, even when its drawing is not literally identical.

§142.7 What this note is not trying to do

This note does not attempt to understand the whole paper. The paper continues into quandles, partial ternary quasigroups, biquandoids, crossoids, and a multicrossing complex. Those structures are not independent tricks. They are different algebraic shadows of the same topological idea:

diagram pieces become probes, and probes have isotopy or homotopy classes.

For a first reading, the main thing to remember is this:

$$\begin{array}{c}
 \text{topological knot or tangle in } F \times I \\
 \Downarrow \\
 \text{planar diagram with arcs, regions, crossings} \\
 \Downarrow \\
 \text{topological reinterpretation of those pieces as probes} \\
 \Downarrow \\
 \text{algebraic colorings extracted from homotopy classes of probes.}
 \end{array}$$

This is why the seminar about partial tribrackets led back to the beginning of knot theory. To understand a ternary operation on regions, one first has to understand what a region is. Nikonov's paper answers: a region is not merely a blank patch in a diagram; it is a topological probe through the complement of the knot.


Nikonov’s Paper at One Crossing: Region Probes, Partial Tribrackets, and What the Seminar Was About

N-20250707

§143.1 Reading only one small part of a long paper

This note is deliberately selective. Nikonov’s paper is long, and it is not reasonable to digest all of it just because one seminar made the subject look interesting. The first seminar memo had the right instinct: the seminar topic, partial tribrackets, is best understood by looking at the part of the paper where regions of diagrams are turned into topological objects and then into algebraic operations.

So here I focus on one small path through the paper:



The point is not to master all definitions. The point is to understand why the word “partial” appears and why a tribracket has anything to do with a knot diagram.

§143.2 Where this sits inside Nikonov’s paper

The paper has two large movements. The first movement is topological: arcs, regions, semiarcs, and crossings of diagrams are reinterpreted as isotopy classes of probes. The second movement is algebraic: homotopy classes of these probes naturally carry coloring structures.

The local dictionary is:

diagram element	algebraic shadow
arc	quandle
region	partial ternary quasigroup / partial tribracket
semiarc	biquandloid
crossing	crossoid

This note chooses the second row. That is the row closest to the seminar topic.




Figure 143.1: From Nikonov's paper: arc and semiarc colorings are the neighbouring examples around the region story.

The figures above are not the main object of this note, but they help orient the reader. Quandles and biquandles are familiar algebraic gadgets for diagram colorings. Nikonov's point is that the less familiar partial tribracket and crossoid fit into the same topological pattern.

§143.3 A region is not just a blank patch

In an ordinary knot diagram, a region is a connected component of the plane or surface after removing the diagram. Informally, it is a blank patch between strands. But this definition depends on the drawing.

Nikonov replaces the patch by a topological path. In a thickened surface $F \times I$, a region probe is a path from the bottom $F \times \{0\}$ to the top $F \times \{1\}$ that does not intersect the knot. In other words, it is a way to pass through the empty space around the knot.




Figure 143.2: From Nikonov's paper: a region probe passes through the complement of the knot.

This is the most important mental shift. In the diagram, a region looks passive. In the thickened surface, a region is represented by an actual route through space. If the knot moves by isotopy, the route can move too. Thus a region can be followed through Reidemeister moves without pretending that it is only a two-dimensional patch of paper.

For a quick analogy, imagine the knot as a collection of wires suspended in a slab of transparent jelly. A region probe is a thin needle pushed from the bottom of the jelly to the top without touching any wire. Moving the knot deforms the possible needle paths. The algebra of regions is really an algebra of these possible needle paths.

§143.4 At a crossing, four regions want to talk to each other

A tribracket is a ternary operation. Instead of combining two inputs, it combines three:

$$[a, b, c] = d.$$

In knot diagrams, the reason for a ternary operation is visual. Around one crossing there are four adjacent regions. If three of their colors are known, the coloring rule determines the fourth.




Figure 143.3: From Nikonov's paper: the tribracket coloring rule near a crossing.

This is the cleanest way to remember the algebra. A quandle is suited to arc colorings because arcs meet at a crossing in a binary-looking way. A tribracket is suited to region colorings because four regions meet around a crossing, so knowing three should determine the fourth.

The operation is designed so that the coloring survives Reidemeister moves. That is the usual diagrammatic requirement: if two diagrams represent the same knot, their coloring data should correspond.

Definition 143.4.1 (Horizontal tribracket, informal version). A tribracket is a set X with a ternary operation

$$[\ , \ , \] : X^3 \rightarrow X$$

such that in an equation

$$[a, b, c] = d$$

any three of a, b, c, d determine the fourth, and such that the operation satisfies the compatibility identity forced by the third Reidemeister move.

The first condition is the analogue of invertibility. If a crossing tells us the fourth region from three known regions, then we must also be able to solve backwards when the diagram is read in another local direction.

The second condition is the real knot-theoretic condition. Around a third Reidemeister move, there are two ways of propagating region colors across three nearby crossings. These two ways must give the same final color. In Nikonov's source, this is expressed by the identity

$$[b, [a, b, c], [a, b, d]] = [c, [a, b, c], [a, c, d]] = [d, [a, b, d], [a, c, d]].$$




Figure 143.4: From Nikonov's paper: the third Reidemeister move is the source of the tribracket identity.

Theorem 143.4.2 (Coloring invariance for tribrackets)

If X is a tribracket, then tribracket colorings of a diagram correspond bijectively under Reidemeister moves.

This theorem is the usual coloring-invariant mechanism. The algebraic axioms are not arbitrary: they are exactly the local rules needed so that the diagram move does not change the set of colorings.

Example 143.4.3 (Two useful models)

An Alexander-type tribracket has the form

$$[a, b, c] = tb + sc - tsa$$

on a module over $\mathbb{Z}[t^{\pm 1}, s^{\pm 1}]$. A group-based Dehn-type model, in the convention used later for the partial construction, has the form

$$[a, b, c] = ba^{-1}c.$$

The first example is linear and computational; the second example is closer to the topology of paths in a complement.

§143.5 Why the operation is partial

In the plane, it is tempting to think that any three region colors can be inserted into the operation. Nikonov's point is that, on a fixed surface and from the topological viewpoint, not every formal triple of regions is genuinely realizable around a crossing.

A region is represented by a probe. Three regions can be composed at a crossing only when the corresponding probes can be arranged in the correct local configuration. If topology prevents that configuration, then the operation should simply not be defined. This is why the algebra is partial.

The paper includes a useful warning example: a formal composition may look legal if one only stares at symbols, but may fail to be realizable on the actual diagram.




Figure 143.5: From Nikonov's paper: an indicated composition of regions that is not realizable on a classical diagram.

This figure is the whole story in miniature. A purely algebraic operation would like to accept any triple. The topology says: no, only triples that can actually occur as neighbouring regions around a crossing should be allowed. The adjective “partial” is not a technical nuisance; it records genuine geometric obstruction.

Definition 143.5.1 (Partial ternary quasigroup, notebook version). A partial ternary quasigroup is a set X , a compatibility relation $\uparrow$, and a ternary operation

$$[a, b, c]$$

defined only when $b \uparrow a$ and $c \uparrow a$. The axioms say that the output is compatible with the neighbouring inputs, that the operation can be solved uniquely for any missing entry whenever the compatibility conditions make sense, and that the third-Reidemeister identity still holds for compatible quadruples.

For this note, the compatibility relation is the essential new ingredient. One should read

$$b \uparrow a$$

as “the region b is allowed to sit next to the region a in the required local way.” In a flat beginner picture this sounds unnecessary, but in a fixed thickened surface it records real topology.

Example 143.5.2 (Alexander numbering)

The paper gives a very small partial example:

$$X = \mathbb{Z}, \quad a \uparrow b \iff b = a - 1, \quad [a, b, c] = a + 2.$$

This is not meant to be the most exciting invariant. It is a memory aid: partiality can be as simple as allowing only neighbouring integer levels to interact.

§143.6 Probe clearing and color rules

The paper also explains how a region probe can be cleared through the diagrammatic situation. The figure below should be read as a topological mechanism behind the coloring rule. We are not just writing $[a, b, c] = d$; we are moving a probe through the thickened surface and watching how its class changes.




Figure 143.6: From Nikonov’s paper: clearing a region probe in the construction of the partial tribracket.

This is the connection with the first note. Diagrammatic elements become probes; homotopy classes of probes become the raw material for algebra; algebraic coloring rules become shadows of probe motion.

The theorem proved in this part of the paper can be read as follows.

Theorem 143.6.1 (Nikonov’s fundamental partial tribracket, informal form)

For a tangle $T \subset F \times I$ with diagram D , colorings of D by a partial ternary quasigroup correspond to homomorphisms out of the topological partial tribracket built from homotopy classes of region probes, with the natural $\pi_1(F)$ -invariance condition.

The exact statement has more notation, but the meaning is very clean. A diagram coloring is not merely a coloring of a planar picture. It is the same thing as a compatible algebraic map from the topological object made of region probes.

In the proof, “clearing a probe” is the operational step. One pulls the diagram away from the probe so that the probe’s value can be read as the color of a region after the transformation. Reidemeister invariance then says that this value does not depend on the particular clearing process.

A small local rule can also be expressed as propagation of colors between neighbouring regions.

$$c \quad \uparrow \quad \phi(c)$$

Figure 143.7: From Nikonov’s paper: a local propagation rule for neighbouring regions.

This is the part I would keep from the seminar if I had to compress it into one sentence:

a partial tribracket is the algebra of how region probes interact near crossings.

§143.7 What about crossings themselves?

The paper does not stop at regions. Crossings also become topological objects: a crossing probe is a bottom-to-top path intersecting the tangle twice, with extra framing data on the segment between the two intersection points.




Figure 143.8: From Nikonov's paper: a crossing probe. Crossings are treated as probe classes rather than merely as vertices of a drawing.

This leads toward crossoids, which generalize parity-type structures on knots. I do not try to unfold crossoids here, but the pattern is now visible:

choose a diagram element
 $\Downarrow$
 replace it by a probe
 $\Downarrow$
 take homotopy classes
 $\Downarrow$
 read off an algebraic coloring structure.

The partial tribracket is simply the region case of this general machine.

§143.8 The paper in one playful sentence

A knot diagram looks like a flat picture, but Nikonov treats its arcs, regions, and crossings as little test paths sent through a three-dimensional slab. Once those test paths are allowed to move, their homotopy classes begin to carry algebra. Arc probes lead toward quandles; region probes lead toward partial tribrackets; semiarc probes lead toward biquandoids; crossing probes lead toward crossoids.

So the paper is not merely saying “here is another coloring invariant.” It is saying something more structural:

the algebraic coloring gadgets are shadows of topological probe spaces.

For a casual post-seminar reading, this is already enough. The rest of the paper builds a much larger machine, but this one crossing explains why the machine is interesting.

XIV

Topology and Homotopy

Part XIV: Contents

N-20210806	933
144.1 Purpose of the note	933
144.2 Group actions	933
144.3 Homotopy	934
144.4 Cofibrations, fibrations, and deformation retracts	935
144.5 Spheres and homotopy groups	938
144.6 The Hopf fibration	939
144.7 References to return to	940
144.8 Group actions as symmetry bookkeeping	941
144.9 Hopf fibration intuition	941
 N-20221218	 943
145.1 Why the note starts with metrics	943
145.2 Boundedness: diameter and radius	943
145.3 Metric spaces	943
145.4 Sequences and continuity	944
145.5 Standard metrics	944
145.6 Norms and inner products	945
145.7 Open and closed balls	945
145.8 Open sets, closed sets, and limit points	945
145.9 Cauchy sequences and completeness	946
145.10 Total boundedness and compactness	946
145.11 Beyond the first calculation	946
145.12 Equivalent metrics and what topology forgets	946
145.13 Completion as adding ideal limit points	947
145.14 Uniform structure hidden inside metric topology	947
145.15 Compactness through total boundedness	947
145.16 Baire category as a completeness phenomenon	948
145.17 Metric spaces through completion and category	948
 N-20221221	 949
146.1 Completion	949
146.2 Let the buyer beware	949
146.3 Homeomorphisms and embeddings	949
146.4 Quotient intuition	950
146.5 A more structural view	950
146.6 Completion and its universal property	950
146.7 Homeomorphism versus uniform equivalence	950
146.8 Continuous bijections that are not homeomorphisms	950
146.9 Quotient topology as final topology	951
146.10 Endpoint gluing and local neighborhoods	951
146.11 Completion, gluing, and quotient control	951
 N-20221222	 953
147.1 Re-defining continuity	953
147.2 Closed sets and closure	953
147.3 Subspaces	953
147.4 Connectedness	953
147.5 Path-connectedness	954
147.6 Continuity: why preimages appear	954
147.7 Closure through neighborhoods	954
147.8 Connectedness by continuous images	954
147.9 Path-connectedness and components	955
147.10 Next viewpoint	955
147.11 Continuity as a contravariant test	955
147.12 Closure as an operator	955
147.13 Why sequences are not enough	955

147.14	Connectedness as maps to a two-point space	956
147.15	Path components versus connected components	956
147.16	Local connectedness as the missing hypothesis	956
147.17	Continuity, closure, and connectedness	956
N-20221223		957
148.1	Homotopy intuition	957
148.2	Simply connected spaces	957
148.3	Semigroups, monoids, and groups	957
148.4	Congruence, quotients, and homomorphisms	957
148.5	Loops and the beginning of the fundamental group	958
148.6	Why the punctured plane has nontrivial loops	958
148.7	Homomorphisms as structure-preserving maps	958
148.8	Quotients in topology and algebra	959
148.9	Higher viewpoint: from holes to algebra	959
148.10	A useful afterword	959
148.11	Concatenation of loops and why homotopy classes form a group	959
148.12	The exponential covering and the integer invariant	959
148.13	Punctured plane deformation retracts to the circle	960
148.14	Functoriality: maps of spaces give homomorphisms	960
148.15	Van Kampen as a computation machine	960
148.16	Group completion and why inverses matter	961
148.17	Loops, coverings, and algebraic invariants	961
N-20221228		963
149.1	Interval modulo endpoints	963
149.2	More quotient spaces	963
149.3	Definition of quotient topology	963
149.4	Product topology	964
149.5	Universal property of quotient topology	964
149.6	Endpoint identification in detail	964
149.7	Product topology via projections	964
149.8	Quotient and product as opposite directions	965
149.9	A second lens	965
149.10	Saturated sets and quotient openness	965
149.11	Quotients can destroy Hausdorffness	965
149.12	Product topology and the mapping property	965
149.13	From squares to surfaces through edge words	966
149.14	Products and quotients need not commute blindly	966
149.15	Products, quotients, and gluing data	966
N-20221229		967
150.1	What this note is really doing	967
150.2	The torus as a parametrized product	967
150.3	Disjoint union	968
150.4	Wedge sum	968
150.5	CW complexes: building spaces by cells	968
150.6	The torus as a quotient of a square	969
150.7	Cosets as algebraic quotients	970
150.8	The Klein bottle	971
150.9	Polygon schemata for surfaces	971
150.10	Projective spaces as quotient spaces	972
150.11	Cellular boundary words and fundamental groups	972
150.12	Why polygon schemata are efficient	973
150.13	Further context	973
150.14	Cellular chain complex of the torus	973
150.15	Klein bottle: same cells, different boundary	974
150.16	Euler characteristic as cell bookkeeping	974
150.17	Orientability from polygon words	974
150.18	Van Kampen and presentations from gluing	974
150.19	Projective spaces through cells	975

150.20 Cellular surfaces through homology and presentations	975
N-20240727	977
151.1 What this chapter is for	977
151.2 The first analytic reflex	977
151.3 Why the Hopf and K-theory hints changed the search	978
151.4 The Eilenberg swindle as the central guess	979
151.5 How this looked like a K-theory statement	980
151.6 The role I imagined for the Hopf fibration	980
151.7 The meme as a learning ladder	981
151.8 What the original interpretation claimed	982
151.9 A slightly uneasy stopping point	982
N-20251217	983
152.1 Why this second note exists	983
152.2 The lecture starts from ordinary equations	984
152.3 The Hopf fibration as the first nontrivial map	984
152.4 Homotopy groups and the first stable class	985
152.5 From natural numbers to finite sets	986
152.6 The space of finite sets	986
152.7 The K-space of finite signed sets	987
152.8 The loop h	987
152.9 Why the first guess was close but reversed	988
152.10 Finite sets and stable spheres	989
152.11 Relation with the classical Hopf map	990
152.12 Cobordism: the movie becomes a circle	990
152.13 Algebraic K-theory of a field	991
152.14 Intrinsic skein-lasagna	992
152.15 The V - k cobordism group	992
152.16 Generalized Pontryagin–Thom	993
152.17 The $K2$ slide and Steinberg relations	994
152.18 The final lasagna diagrams	995
152.19 What the equation finally says	996
152.20 Structural viewpoint	997

Group Actions, Homotopy, and a Rough Road to the Hopf Fibration

N-20210806

§144.1 Purpose of the note

This note was written as preparation for a seminar on homotopy groups and the Hopf fibration. The goal is modest: I only want enough background so that the seminar is not blocked by unfamiliar words. The content is therefore a quick route through group actions, homotopy, cofibrations and fibrations, homotopy groups, and finally the Hopf fibration.

Caution

The original note was written quickly and openly says that errors may remain. I have corrected the most obvious convention problems and phrasing issues, while keeping the exploratory style.

§144.2 Group actions

A group action is one of the first ways in which abstract algebra touches geometry and topology. A group can be understood as acting by symmetries or permutations of a set.

Definition 144.2.1. Let G be a group and X a set. A left action of G on X is a map

$$G \times X \rightarrow X, \quad (g, x) \mapsto g \cdot x,$$

such that

$$e \cdot x = x, \quad (gh) \cdot x = g \cdot (h \cdot x).$$

Equivalently, it is a homomorphism

$$\varphi : G \rightarrow S_X,$$

where S_X is the permutation group of X .

Remark 144.2.2 — One may also use right actions, written $x \cdot g$, satisfying

$$x \cdot 1 = x, \quad (x \cdot g) \cdot h = x \cdot (gh).$$

The original note switches viewpoints, so I fix the left-action convention unless stated otherwise.

Example 144.2.3

The symmetric group S_n acts on $X = \{1, 2, \dots, n\}$. Cayley's theorem says that every group G can be realized as a subgroup of the permutation group of its underlying set, by the action of left multiplication.

Example 144.2.4

The symmetry group of a cube acts on the cube and may also be regarded as a group of linear transformations of $\mathbb{R}^3$ preserving the cube.

Example 144.2.5 (Conjugation)

A group G acts on itself by conjugation:

$$g \cdot x = gxg^{-1}.$$

For a fixed $a \in G$, its orbit is

$$\text{Orb}_G(a) = \{gag^{-1} \mid g \in G\},$$

which is the conjugacy class of a . Its stabilizer is

$$\text{Stab}_G(a) = \{g \in G \mid gag^{-1} = a\},$$

which is the centralizer of a .

Definition 144.2.6. For an action of G on X , the stabilizer and orbit of $x \in X$ are

$$\text{Stab}_G(x) = \{g \in G \mid g \cdot x = x\},$$

$$\text{Orb}_G(x) = \{g \cdot x \mid g \in G\}.$$

Theorem 144.2.7 (Orbit-stabilizer)

If G is finite, then

$$|G| = |\text{Stab}_G(x)| |\text{Orb}_G(x)|.$$

Proof. The map $G \rightarrow \text{Orb}_G(x)$, $g \mapsto g \cdot x$, has fibres equal to left cosets of $\text{Stab}_G(x)$. Hence the orbit has $[G : \text{Stab}_G(x)]$ elements. $\square$

§144.3 Homotopy

Algebraic topology tries to classify topological spaces and continuous maps using algebraic invariants. Before doing that, one needs a flexible notion of sameness. Homotopy is such a notion.

Definition 144.3.1. Let X, Y be topological spaces and $f, g : X \rightarrow Y$ continuous maps. A homotopy from f to g is a continuous map

$$F : X \times I \rightarrow Y, \quad I = [0, 1],$$

such that

$$F(x, 0) = f(x), \quad F(x, 1) = g(x).$$

When such an F exists, write $f \simeq g$.

Equivalently, a homotopy may be viewed as a map

$$F : X \rightarrow Y^I,$$

where $Y^I = \text{Top}(I, Y)$ is the path space, and the endpoint evaluation maps recover f and g .

Proposition 144.3.2

Homotopy is an equivalence relation on maps $X \rightarrow Y$, and it is compatible with composition. If $f \simeq g$, then

$$k \circ f \circ h \simeq k \circ g \circ h$$

whenever the compositions are defined.

Definition 144.3.3. Two spaces X and Y have the same homotopy type if there are maps $f : X \rightarrow Y$ and $g : Y \rightarrow X$ such that

$$g \circ f \simeq \text{id}_X, \quad f \circ g \simeq \text{id}_Y.$$

§144.4 Cofibrations, fibrations, and deformation retracts

The original note introduces cofibrations and fibrations through lifting diagrams. These diagrams are not decorative: each arrow is one piece of the data in the definition. I spell them out because this is exactly where the abstraction tends to hide the simple geometric idea.

§144.4.i Homotopy extension property

A map $j : A \rightarrow X$ is a cofibration if it has the homotopy extension property. The situation is this. We know a map

$$g : X \rightarrow Z,$$

and we also know a homotopy on the smaller space

$$G : A \times I \rightarrow Z.$$

These two pieces of data must agree at time 0: for every $a \in A$,

$$G(a, 0) = g(j(a)).$$

If this compatibility condition holds, the homotopy extension property asks for a map

$$\overline{G} : X \times I \rightarrow Z$$

such that

$$\overline{G}(x, 0) = g(x), \quad \overline{G}(j(a), t) = G(a, t).$$

In other words, a homotopy that has only been prescribed on the subspace A must extend to a homotopy on all of X , while keeping the original time-zero map on X .

The corresponding diagram is

$$\begin{array}{ccccc}
 A & \xrightarrow{j} & X & \xrightarrow{\quad g \quad} & Z \\
 \downarrow i_0 & & \downarrow i_0 & & \parallel \\
 A \times I & \xrightarrow{j \times I} & X \times I & \xrightarrow{\quad \overline{G} \quad} & Z \\
 & \searrow G & & &
 \end{array}$$

Here i_0 denotes insertion at time zero: $a \mapsto (a, 0)$ or $x \mapsto (x, 0)$, according to the source. The bottom arrow $j \times I$ says that the cylinder over A sits inside the cylinder over X . The two solid maps into Z are the information already given: g on X , and G on $A \times I$. The rightmost equality means that both pieces are maps into the same target Z , not two different copies of Z . The dashed arrow is not part of the data. It is exactly the thing whose existence is required.

This is why the word “extension” is literal. We first define a map on the union

$$X \times \{0\} \cup A \times I.$$

On $X \times \{0\}$ it is g . On $A \times I$ it is G . The compatibility condition $G(a, 0) = g(j(a))$ says these two formulas agree on the overlap $A \times \{0\}$. The H.E.P. says this partial map extends across the whole cylinder $X \times I$.

A useful mental picture is: $A \subset X$ is a part of the space whose motion has already been specified. A cofibration is an inclusion good enough that this partial motion can be filled in over the rest of X . Thus cofibrations are the maps along which homotopies extend outward.

§144.4.ii Homotopy lifting property

A fibration is the dual-looking story. Let

$$p : X \rightarrow A$$

be a map. The homotopy lifting property says that if a homotopy has been drawn downstairs in A , and if its starting point has already been lifted upstairs to X , then the whole homotopy can be lifted upstairs.

Using path spaces makes the diagram compact. The notation A^I means the space of paths $I \rightarrow A$, and X^I means the space of paths $I \rightarrow X$. Evaluation at time 0 gives

$$ev_0 : A^I \rightarrow A, \quad ev_0 : X^I \rightarrow X.$$

The map $p : X \rightarrow A$ also induces a map on path spaces

$$p^I : X^I \rightarrow A^I, \quad \gamma \mapsto p \circ \gamma.$$

Now take a parameter space Z . A map

$$F : Z \rightarrow A^I$$

is a Z -indexed family of paths downstairs in A . A map

$$\tilde{f}_0 : Z \rightarrow X$$

chooses a starting lift upstairs. The compatibility condition is

$$p \circ \tilde{f}_0 = ev_0 \circ F.$$

This simply says: the chosen point upstairs really lies over the initial point of the path downstairs.

The lifting diagram is

$$\begin{array}{ccc} X^I & \xrightarrow{ev_0} & X \\ p^I \downarrow & & \downarrow p \\ A^I & \xrightarrow{ev_0} & A \end{array} \quad \begin{array}{ccc} & & X^I \\ & \nearrow \bar{F} & \downarrow (p^I, ev_0) \\ Z & \xrightarrow{(F, \tilde{f}_0)} & A^I \times_A X \end{array}$$

The square on the left is the structural square. The horizontal map $p^I : X^I \rightarrow A^I$ sends an upstairs path to its downstairs projection, and the two ev_0 arrows take initial points. The square commutes because evaluating $p \circ \gamma$ at time 0 gives $p(\gamma(0))$. The diagram on the right packages the actual lifting problem. A pair $(F, \tilde{f}_0)$ lands in the fiber product $A^I \times_A X$ exactly when $p \circ \tilde{f}_0 = ev_0 \circ F$. The dashed arrow $\bar{F} : Z \rightarrow X^I$ is the desired lifted family of paths.

Unpacking the dashed arrow, for each $z \in Z$ it gives a path

$$\bar{F}(z) : I \rightarrow X.$$

The two required equations are

$$p \circ \bar{F}(z) = F(z), \quad \bar{F}(z)(0) = \tilde{f}_0(z).$$

So the lifted path projects to the given downstairs path and starts at the chosen upstairs point. This is the exact formal meaning of “lift upward from the base.”

The relationship with the earlier definition is now transparent. A cofibration problem begins with a homotopy on a subspace and asks to extend it sideways across a larger cylinder. A fibration problem begins with a homotopy downstairs and asks to lift it vertically to a path or homotopy upstairs. Both are extension problems, but in different directions: H.E.P. fills missing points in the domain; H.L.P. fills missing lifts in the total space.

Definition 144.4.1. A map $j : A \rightarrow X$ is a cofibration if it has the homotopy extension property described above.

Definition 144.4.2. A map $p : X \rightarrow A$ is a fibration if it has the homotopy lifting property described above.

Remark 144.4.3 — Cofibration and fibration are dual-looking ideas. A cofibration lets homotopies extend outward from a subspace; a fibration lets homotopies lift upward from the base. In the Hopf fibration $S^3 \rightarrow S^2$, the word “fibration” means exactly that small motions of a point in the base sphere can be lifted to motions in S^3 , once a starting point in the circle fiber has been chosen.

Definition 144.4.4. A deformation retract of X onto a subspace A is a homotopy

$$H : X \times I \rightarrow X$$

such that

$$H(x, 0) = x, \quad H(x, 1) \in A, \quad H(a, t) = a$$

for all $x \in X$, $a \in A$, and $t \in I$.

Proposition 144.4.5

The pair (X, A) has the homotopy extension property if and only if

$$X \times \{0\} \cup A \times I$$

is a retract of $X \times I$.

§144.5 Spheres and homotopy groups

The n -sphere is

$$S^n = \{(x_1, \dots, x_{n+1}) \in \mathbb{R}^{n+1} \mid x_1^2 + \dots + x_{n+1}^2 = 1\}.$$

The n -torus is

$$T^n = (S^1)^n.$$

Thus a point of T^2 may be described by a pair of angles, or equivalently by a point of $S^1 \times S^1$.

Caution

The original note says that a personal translation habit used the wrong word for “equator”. I use the standard word *equator* throughout.

Stereographic projection identifies a sphere minus one point with Euclidean space. For example,

$$S^2 \setminus \{\text{north pole}\} \cong \mathbb{R}^2.$$

The equator is the region where the projection is least distorted; away from it, shapes become increasingly distorted.

Definition 144.5.1. Let X be a based space with basepoint x_0 . The n -th homotopy group or homotopy set is

$$\pi_n(X, x_0) = [(S^n, *), (X, x_0)],$$

the set of based homotopy classes of based maps $S^n \rightarrow X$. For $n \geq 1$, it has a natural group structure.

Using the model

$$S^n \simeq I^n / \partial I^n,$$

the group operation for $n \geq 1$ may be described by pinching the sphere along an equator:

$$S^n \rightarrow S^n \vee S^n.$$

Given two based maps $f, g : S^n \rightarrow X$, the product first pinches S^n into two copies, then applies f on one copy and g on the other.

Remark 144.5.2 — For $n = 0$, $\pi_0(X)$ records path components and is not usually a group. For $n \geq 1$, $\pi_n(X, x_0)$ is a group; for $n \geq 2$, it is abelian.

§144.6 The Hopf fibration

The Hopf fibration is a map

$$S^3 \rightarrow S^2.$$

It is one of the first examples showing that spheres are connected by unexpectedly rich maps.

Describe

$$S^3 = \{(z_1, z_2) \in \mathbb{C}^2 \mid |z_1|^2 + |z_2|^2 = 1\}.$$

A clean formula for the Hopf map is

$$h : S^3 \rightarrow \mathbb{C}P^1 \cong S^2, \quad h(z_1, z_2) = [z_1 : z_2].$$

On the chart $z_1 \neq 0$, this looks like

$$z_2/z_1.$$

The case $z_1 = 0$ is the point at infinity.

Caution

It is tempting to say $z_2/0 = \infty$, and that intuition is useful. The mathematically clean way is to use projective coordinates $[z_1 : z_2]$, which handle infinity automatically.

The fibers over three simple points can be written without any abstraction:

$$F_1 = h^{-1}([1 : 0]) = \{(e^{it}, 0) : 0 \leq t < 2\pi\},$$

$$F_2 = h^{-1}([0 : 1]) = \{(0, e^{it}) : 0 \leq t < 2\pi\},$$

and

$$F_3 = h^{-1}([1 : 1]) = \left\{ \left(\frac{e^{it}}{\sqrt{2}}, \frac{e^{it}}{\sqrt{2}} \right) : 0 \leq t < 2\pi \right\}.$$

To see them inside ordinary three-space, use stereographic projection

$$\Pi(z_1, z_2) = \left(\frac{\operatorname{Re}(z_1)}{1 - \operatorname{Im}(z_2)}, \frac{\operatorname{Im}(z_1)}{1 - \operatorname{Im}(z_2)}, \frac{\operatorname{Re}(z_2)}{1 - \operatorname{Im}(z_2)} \right).$$

Writing $z_1 = x_1 + iy_1$ and $z_2 = x_2 + iy_2$, this is the projection from the point $(0, i) \in S^3$.

For F_1 , substitution gives

$$\Pi(e^{it}, 0) = (\cos t, \sin t, 0),$$

the standard unit circle in the plane $z = 0$. For F_2 ,

$$\Pi(0, e^{it}) = \left(0, 0, \frac{\cos t}{1 - \sin t} \right),$$

so this fiber becomes the vertical line together with the point at infinity; the missing value $t = \pi/2$ is precisely the projection point. For F_3 ,

$$\Pi\left(\frac{e^{it}}{\sqrt{2}}, \frac{e^{it}}{\sqrt{2}}\right) = \left(\frac{\cos t/\sqrt{2}}{1 - \sin t/\sqrt{2}}, \frac{\sin t/\sqrt{2}}{1 - \sin t/\sqrt{2}}, \frac{\cos t/\sqrt{2}}{1 - \sin t/\sqrt{2}} \right).$$

This curve lies in the plane $X = Z$. Eliminating t gives

$$2X^2 + (Y - 1)^2 = 2, \quad Z = X,$$

so it is again a circle, not a random distorted loop.

The important topological fact is that $\Pi(F_1)$ and $\Pi(F_3)$ form a Hopf link. With the orientations induced by increasing t , their linking number is $+1$; reversing one orientation changes the sign. The computation matches the picture: one circle passes once through a spanning disk of the other. The core geometric intuition of the Hopf fibration is that any two distinct circle fibers in S^3 become linked once after stereographic projection to $\mathbb{R}^3$.

There is also a quaternionic picture. Regard S^3 as the unit quaternions and S^2 as the purely imaginary unit quaternions. Fix $p \in S^2$. For a unit quaternion $q \in S^3$, define

$$q \mapsto qpq^{-1}.$$

This gives another description of a map $S^3 \rightarrow S^2$, equivalent to the Hopf fibration.

Remark 144.6.1 — The main fact to remember is simple but powerful: the Hopf fibration has total space S^3 , base S^2 , and fibre S^1 . A full understanding needs pictures, but the formula already shows why complex numbers, quaternions, and topology appear together.

§144.7 References to return to

The original note points toward Hatcher's *Algebraic Topology*, Niles Johnson's visual material on the Hopf fibration, and interactive Hopf-fibration demonstrations. Those should be treated as the real next step, because this note is mainly a vocabulary bridge.

The diagrammatic part of the note should be read as a staged program. The homotopy extension property and the homotopy lifting property have now been drawn explicitly

above, because those two diagrams are the quickest way to connect the abstract definitions with the words “extend” and “lift”. The same visual habit should also be used when returning to the pinch map

$$S^n \rightarrow S^n \vee S^n$$

and to the Hopf fibration itself. The definitions and formulas are already in place, but the geometry becomes much more memorable when the pinch map, circle fibers, and linked fibers are kept visible.

§144.8 Group actions as symmetry bookkeeping

A group action of G on X is a map

$$G \times X \rightarrow X$$

satisfying $e \cdot x = x$ and $(gh) \cdot x = g \cdot (h \cdot x)$. The orbit of x is

$$Gx = \{g \cdot x : g \in G\},$$

and the stabilizer is

$$G_x = \{g : g \cdot x = x\}.$$

The orbit-stabilizer relation says that symmetry and movement are complementary. Large stabilizer means small orbit; small stabilizer means large orbit.

§144.9 Hopf fibration intuition

The Hopf fibration is a map

$$S^3 \rightarrow S^2$$

whose fibers are circles. One way to see it is through complex coordinates: points of $S^3 \subset \mathbb{C}^2$ are pairs (z_1, z_2) with $|z_1|^2 + |z_2|^2 = 1$. Multiplying both coordinates by the same unit complex number moves along a circle fiber. The quotient by this circle action gives $\mathbb{CP}^1 \cong S^2$.

The sentence to remember is concrete: the fibers are not merely many disjoint circles; any two of them are linked once. This is why the formula $S^3 \rightarrow S^2$ is already a three-dimensional knot-theoretic picture.

Basic Topology I: Metric Spaces, Boundedness, Norms, Open Sets, and Completeness

N-20221218

§145.1 Why the note starts with metrics

This note begins a topology sequence, but it does not begin with arbitrary open sets. It begins with metric spaces, because metric spaces keep the analysis intuition alive. In real analysis, convergence means

$$x_n \rightarrow x \iff (\forall \varepsilon > 0)(\exists N)(\forall n > N) |x_n - x| < \varepsilon.$$

The first abstraction is to replace $|x_n - x|$ by a general distance $d(x_n, x)$. Thus a sequence in a metric space converges when

$$(\forall \varepsilon > 0)(\exists N)(\forall n > N) d(x_n, x) < \varepsilon.$$

The hidden question is: which parts of analysis really require a numerical distance, and which only require a notion of neighborhood? Metric spaces are the intermediate language between analysis and topology.

§145.2 Boundedness: diameter and radius

A metric space M is bounded if there is a constant D such that

$$d(p, q) \leq D \quad \text{for all } p, q \in M.$$

One can think of D as a diameter bound. The note also gives an equivalent radius formulation: fixing a point p , if there is R such that

$$d(p, q) \leq R \quad \text{for all } q \in M,$$

then M is bounded. Indeed, for any $q_1, q_2 \in M$,

$$d(q_1, q_2) \leq d(q_1, p) + d(p, q_2) \leq 2R.$$

Conversely, a diameter bound immediately gives a radius bound from any chosen point.

The examples are deliberately chosen to break naive intuition: $[0, 1]$ is bounded, $\mathbb{R}^n$ is not, an infinite set with the discrete metric is bounded, but $\mathbb{N}$ with the usual metric is not. Thus boundedness belongs to a metric, not to a bare set.

§145.3 Metric spaces

A metric space is a pair (X, d) where $d : X \times X \rightarrow \mathbb{R}$ satisfies:

- (i) $d(x, y) \geq 0$;

- (ii) $d(x, y) = 0$ iff $x = y$;
- (iii) $d(x, y) = d(y, x)$;
- (iv) $d(x, z) \leq d(x, y) + d(y, z)$.

A subset $Y \subseteq X$ inherits the subspace metric

$$d_Y(a, b) = d_X(a, b).$$

This is the metric ancestor of subspace topology.

§145.4 Sequences and continuity

A sequence (x_n) converges to x if every small ball around x eventually contains all x_n . In $\mathbb{R}^k$ with the Euclidean metric, convergence is coordinatewise:

$$x_n = (x_n^1, \dots, x_n^k) \rightarrow x = (x^1, \dots, x^k)$$

iff $x_n^i \rightarrow x^i$ for each i . One direction follows from $|x_n^i - x^i| \leq \|x_n - x\|_2$; the other follows by summing the small coordinate errors.

A map $f : X \rightarrow Y$ between metric spaces is continuous at x_0 if

$$(\forall \varepsilon > 0)(\exists \delta > 0) \quad d_X(x, x_0) < \delta \Rightarrow d_Y(f(x), f(x_0)) < \varepsilon.$$

In metric spaces, continuity can also be detected by sequences:

$$x_n \rightarrow x_0 \Rightarrow f(x_n) \rightarrow f(x_0).$$

The later general topology notes warn that sequences no longer detect everything in arbitrary topological spaces.

§145.5 Standard metrics

On $\mathbb{R}^n$ the Euclidean, taxicab, and max metrics are

$$d_2(x, y) = \left(\sum_i |x_i - y_i|^2 \right)^{1/2},$$

$$d_1(x, y) = \sum_i |x_i - y_i|,$$

and

$$d_\infty(x, y) = \max_i |x_i - y_i|.$$

Their balls have different shapes: round disks, diamonds, and squares. Nevertheless they generate the same usual topology on $\mathbb{R}^n$. This is an early hint that topology remembers neighborhoods but forgets exact metric geometry.

The discrete metric on a set X is

$$d(x, y) = \begin{cases} 0, & x = y, \\ 1, & x \neq y. \end{cases}$$

Then $B_{1/2}(x) = \{x\}$. Every subset is open and closed, and a sequence converges iff it is eventually constant.

§145.6 Norms and inner products

A norm on a vector space V satisfies positivity, homogeneity, and the triangle inequality:

$$\|v + w\| \leq \|v\| + \|w\|.$$

Every norm gives a metric by

$$d(v, w) = \|v - w\|.$$

The p -norms are

$$\|x\|_p = \left(\sum_i |x_i|^p \right)^{1/p},$$

and $\|x\|_\infty = \max_i |x_i|$.

An inner product gives a norm by

$$\|v\| = \sqrt{\langle v, v \rangle}.$$

The triangle inequality for this norm follows from Cauchy–Schwarz. So Euclidean metric geometry, linear algebra, and inequalities are already intertwined.

§145.7 Open and closed balls

For $x \in X$ and $r > 0$,

$$B_r(x) = \{y : d(x, y) < r\}, \quad \overline{B}_r(x) = \{y : d(x, y) \leq r\}.$$

Open balls are open. If $y \in B_r(x)$, choose

$$\delta = r - d(x, y) > 0.$$

Then $B_\delta(y) \subseteq B_r(x)$ by the triangle inequality. The proof says: every point inside a ball has some remaining margin before hitting the boundary.

§145.8 Open sets, closed sets, and limit points

A set $U \subseteq X$ is open if every point $x \in U$ has some ball $B_r(x)$ contained in U . A set is closed if its complement is open. In metric spaces, a set C is closed iff it contains the limits of all convergent sequences from C .

A point x is a limit point of A if every ball around x meets A in a point other than x . The closure satisfies

$$x \in \overline{A} \iff (\forall r > 0) B_r(x) \cap A \neq \emptyset.$$

The boundary is

$$\partial A = \overline{A} \cap \overline{X \setminus A}.$$

Examples matter: every real number is a limit point of $\mathbb{Q}$, and also of the irrational numbers; in the discrete metric, points are isolated by small balls.

§145.9 Cauchy sequences and completeness

A sequence is Cauchy if

$$(\forall \varepsilon > 0)(\exists N)(\forall m, n > N) \, d(x_m, x_n) < \varepsilon.$$

A metric space is complete if every Cauchy sequence converges in the space. The standard contrast is

$$\mathbb{R} \text{ is complete,} \quad \mathbb{Q} \text{ is not complete.}$$

A rational sequence converging to $\sqrt{2}$ is Cauchy in $\mathbb{Q}$, but it has no rational limit.

Convergence names a destination; Cauchy only says the sequence internally stabilizes. Completeness says internal stabilization is enough to guarantee a destination inside the space.

§145.10 Total boundedness and compactness

The note also points toward total boundedness and compactness. A metric space is totally bounded if for every $\varepsilon > 0$, finitely many ε -balls cover it. Total boundedness is stronger than boundedness: an infinite discrete metric space is bounded, but for $\varepsilon < 1$ it cannot be covered by finitely many ε -balls.

In $\mathbb{R}^n$, compactness is equivalent to closed and bounded, but that is the Heine–Borel theorem, not the definition in arbitrary metric spaces. A better metric-space memory is:

$$\text{compact} \Rightarrow \text{complete and totally bounded.}$$

Conversely, in metric spaces, complete plus totally bounded implies compact.

§145.11 Beyond the first calculation

This note stages the transition from analysis to topology. Metrics let us talk about boundedness, convergence, continuity, balls, closed sets, limit points, Cauchy sequences, and completeness. But later topology keeps only what can be expressed through open sets. Boundedness and completeness will fail to be topological invariants; continuity and connectedness will survive.

§145.12 Equivalent metrics and what topology forgets

On $\mathbb{R}^n$, the usual norms produce different numerical distances but the same open sets. For instance,

$$\|x\|_\infty \leq \|x\|_2 \leq \sqrt{n} \|x\|_\infty.$$

Therefore a ball for one norm contains a smaller ball for the other norm, and the two norms generate the same topology.

This inequality is the concrete mechanism behind the slogan “topology forgets exact metric geometry.” A square ball and a round ball look different, but each can be squeezed inside the other after changing radius. Thus convergence and continuity agree, while lengths, angles, and diameters do not.

Remark 145.12.1 — The finite-dimensional hypothesis matters. In infinite-dimensional function spaces, norms need not be equivalent. On $C[0, 1]$, the sup norm and the L^1 -norm lead to very different convergence: narrow spikes can have small L^1 -norm but large sup norm.

§145.13 Completion as adding ideal limit points

The completion of a metric space is not just a larger space containing more points. It is a controlled way to turn every Cauchy sequence into a convergent sequence. The construction by Cauchy sequences makes this precise:

$$\widehat{X} = \{\text{Cauchy sequences in } X\} / \sim, \quad (x_n) \sim (y_n) \iff d(x_n, y_n) \rightarrow 0.$$

The distance between two completed points is

$$\widehat{d}([(x_n)], [(y_n)]) = \lim_{n \rightarrow \infty} d(x_n, y_n).$$

This is well-defined because the two sequences are Cauchy.

For $\mathbb{Q}$, a Cauchy sequence of rationals approximating $\sqrt{2}$ becomes a genuine point in the completion. The point is not a preferred decimal expansion; it is the whole equivalence class of rational approximations that close in on the same missing limit.

§145.14 Uniform structure hidden inside metric topology

Continuity only cares about open sets. Cauchy sequences care about more: they need a way to compare two moving points x_m and x_n . This information is called a uniform structure. Two metrics may define the same topology but different Cauchy behavior.

The homeomorphism

$$(0, 1) \rightarrow \mathbb{R}, \quad x \mapsto \tan \pi(x - 1/2)$$

shows the issue. Topologically the spaces are the same. But the Euclidean metric on $(0, 1)$ sees sequences approaching 1 as Cauchy without a limit inside $(0, 1)$, while their images in $\mathbb{R}$ go to infinity and are not Cauchy. Thus completeness is not topological; it is uniform/metric.

§145.15 Compactness through total boundedness

In metric spaces,

$$\text{compact} \iff \text{complete and totally bounded.}$$

This statement is more portable than Heine–Borel. The two hypotheses do different jobs. Total boundedness prevents points from spreading out into infinitely many separated regions; completeness prevents Cauchy subsequences from escaping the space.

The contrast with boundedness is useful. An infinite set with the discrete metric $d(x, y) = 1$ is bounded, because all distances are at most 1. But it is not totally bounded: for $\varepsilon < 1$, every ε -ball contains just one point, so no finite family can cover the space.

§145.16 Baire category as a completeness phenomenon

Completeness has consequences beyond sequences. The Baire category theorem says that a complete metric space cannot be written as a countable union of nowhere dense closed sets. In practice, this means complete spaces are too large to be exhausted by thin pieces.

This theorem explains why complete function spaces are powerful. Many existence theorems in analysis prove that a desired object lies in a countable intersection of dense open sets. Completeness then guarantees that the intersection is nonempty and large. Thus the elementary Cauchy condition eventually becomes a tool for proving generic existence.

§145.17 Metric spaces through completion and category

Metric topology is a bridge. It keeps enough numerical structure to discuss boundedness, Cauchy sequences, total boundedness, and completeness, while already pointing toward open-set topology through balls, neighborhoods, closure, and continuity. The advanced lesson is to separate three layers:

$$\text{topological} \subsetneq \text{uniform/metric} \subsetneq \text{geometric}.$$

Continuity lives in the first layer, completeness in the second, and lengths/angles in the third.

Basic Topology II: Completion, Homeomorphism Warnings, and Quotient Intuition

N-20221221

§146.1 Completion

The note continues with the theorem that every metric space can be completed. The standard example is

$$\mathbb{Q} \leadsto \mathbb{R}.$$

A rational Cauchy sequence should be thought of as trying to name a real number. If the real number is irrational, the sequence has no limit in $\mathbb{Q}$, but it has a limit after completion.

One construction takes equivalence classes of Cauchy sequences. Two Cauchy sequences (x_n) , (y_n) represent the same completed point if

$$d(x_n, y_n) \rightarrow 0.$$

This is an early quotient construction: the new point is not one sequence but a class of mutually converging sequences.

§146.2 Let the buyer beware

The note warns that boundedness and completeness are not topological invariants. The interval $(0, 1)$ and the line $\mathbb{R}$ are homeomorphic, for instance by

$$f(x) = \tan(\pi(x - 1/2)).$$

But $(0, 1)$ is bounded and not complete under the Euclidean metric, while $\mathbb{R}$ is complete and unbounded.

This is the first serious conceptual warning: homeomorphism preserves open-set structure, not distances, diameters, or Cauchy behavior.

§146.3 Homeomorphisms and embeddings

A homeomorphism $f : M \rightarrow N$ is a bijection such that both f and f^{-1} are continuous. The inverse condition is essential. A continuous bijection can fail to be a homeomorphism if the inverse destroys neighborhoods.

An embedding realizes a space as a subspace of another space in a topology-preserving way. It is more than an injection: it must preserve the local structure inherited from the target.

§146.4 Quotient intuition

A quotient space identifies points. For example,

$$[-1, 1]/(-1 \sim 1) \cong S^1.$$

A neighborhood of the glued point corresponds upstairs to two half-neighborhoods at the two endpoints. This explains why quotient topology is defined by requiring a set downstairs to be open exactly when its preimage upstairs is open.

§146.5 A more structural view

Completion adds missing limits. Homeomorphism forgets metric size. Embedding preserves a space as a subspace. Quotient topology glues points together. These are four ways of changing a space while controlling which structure survives.

§146.6 Completion and its universal property

The completion $X \hookrightarrow \widehat{X}$ has a useful universal property. If Y is complete and $f : X \rightarrow Y$ is uniformly continuous, then f extends uniquely to a continuous map

$$\widehat{f} : \widehat{X} \rightarrow Y.$$

On a completed point $[(x_n)]$, define

$$\widehat{f}([(x_n)]) = \lim_{n \rightarrow \infty} f(x_n).$$

Uniform continuity ensures that $(f(x_n))$ is Cauchy in Y , and completeness of Y supplies the limit.

This explains why completion is canonical. It is not merely "adding some missing points"; it is the best possible way to make Cauchy limits available without changing the original metric information.

§146.7 Homeomorphism versus uniform equivalence

The example $(0, 1) \cong \mathbb{R}$ shows that topology does not preserve boundedness or completeness. A sharper statement is that the map is a homeomorphism but not a uniform equivalence for the usual metrics.

A uniformly continuous map controls all small distances by one global δ . The tangent map

$$x \mapsto \tan(\pi(x - 1/2))$$

blows up near the endpoints. Points very close to 1 can be very close in $(0, 1)$ but sent far apart in $\mathbb{R}$. Thus topology sees local neighborhoods, while uniform structure sees global control of closeness.

§146.8 Continuous bijections that are not homeomorphisms

A continuous bijection need not be a homeomorphism. A standard example is

$$f : [0, 2\pi) \rightarrow S^1, \quad f(t) = (\cos t, \sin t).$$

It is continuous and bijective, but its inverse is not continuous at the point $(1, 0)$. Points on the circle approaching $(1, 0)$ from below have parameters approaching 2π , while the point itself has parameter 0.

This example explains why the inverse condition in the definition of homeomorphism is not pedantry. A bijection can preserve points but fail to preserve neighborhoods.

§146.9 Quotient topology as final topology

For a quotient map $q : X \rightarrow Y$, the quotient topology on Y is the strongest topology making q continuous. It is characterized by

$$U \subseteq Y \text{ open} \iff q^{-1}(U) \text{ open in } X.$$

The universal property is:

$$f : Y \rightarrow Z \text{ is continuous} \iff f \circ q : X \rightarrow Z \text{ is continuous.}$$

This is the correct abstract version of endpoint gluing. To define a continuous function out of a glued space, define it before gluing and check that it is constant on equivalence classes.

§146.10 Endpoint gluing and local neighborhoods

In $[0, 1]/(0 \sim 1)$, a neighborhood of the glued point is not a one-sided interval. Its preimage is a union of a small interval near 0 and a small interval near 1:

$$[0, \varepsilon) \cup (1 - \varepsilon, 1].$$

This two-sided local structure is why the quotient is a circle. The topology is doing exactly what the picture suggests: it forces the two endpoints to share neighborhoods after identification.

§146.11 Completion, gluing, and quotient control

Completion, homeomorphism, embedding, and quotient are not interchangeable changes of space. Completion preserves metric Cauchy information while adding limits. Homeomorphism preserves open-set structure and may forget metric size. Embedding realizes a topology inside another space. Quotient topology glues points and controls maps out of the glued space.

Topological Re-definitions, Subspaces, Connectedness, and Path-Connectedness

N-20221222

§147.1 Re-defining continuity

After defining topological spaces, the note rewrites metric definitions using open sets. A function $f : X \rightarrow Y$ is continuous if

$$f^{-1}(V) \text{ is open in } X$$

for every open $V \subseteq Y$. This definition no longer mentions ε , δ , or distances. It says that open information in Y pulls back to open information in X .

A homeomorphism is a bijection preserving open sets in both directions. This is the topological notion of sameness.

§147.2 Closed sets and closure

A set $C \subseteq X$ is closed if $X \setminus C$ is open. The closure of S is the smallest closed set containing S :

$$\overline{S} = \bigcap \{C : C \text{ closed and } S \subseteq C\}.$$

The note warns that complete and totally bounded are metric properties, not topological properties. It also warns that sequences do not work perfectly in general topology; one often needs nets or filters.

§147.3 Subspaces

If $Y \subseteq X$, the subspace topology on Y is

$$\tau_Y = \{U \cap Y : U \in \tau_X\}.$$

The key warning is: open in Y need not be open in X . For example,

$$[0, 1) = [0, 2] \cap (-1, 1)$$

is open in $[0, 2]$, but not open in $\mathbb{R}$.

§147.4 Connectedness

A space X is disconnected if there are nonempty open sets U, V such that

$$X = U \cup V, \quad U \cap V = \emptyset.$$

If no such separation exists, X is connected. Equivalently, X has no nontrivial clopen subset. Separated unions such as $[1, 2] \cup [3, 4]$ are disconnected, while intervals are connected.

§147.5 Path-connectedness

A space X is path-connected if for any $x, y \in X$, there is a continuous map

$$\gamma : [0, 1] \rightarrow X, \quad \gamma(0) = x, \quad \gamma(1) = y.$$

Path-connectedness implies connectedness. The converse can fail. In ordinary regions of $\mathbb{R}^n$, it matches the intuition that one can walk from one point to another without leaving the set.

§147.6 Continuity: why preimages appear

The preimage definition of continuity can feel backward at first. We define continuity by saying that for every open set $V \subseteq Y$, the preimage $f^{-1}(V)$ is open in X , not by saying images of open sets are open. The reason is logical: to know whether $f(x)$ lies in a neighborhood V , we ask which points x are sent into V . That set is the preimage.

Images of open sets need not be open. For example, the map

$$f : \mathbb{R} \rightarrow \mathbb{R}, \quad f(x) = x^2$$

sends the open interval $(-1, 1)$ to $[0, 1)$, which is not open in $\mathbb{R}$. But f is continuous, because preimages of open sets are open.

This example is important because it prevents a common misconception: continuity is about pulling back open tests, not pushing open sets forward.

§147.7 Closure through neighborhoods

The closure $\overline{S}$ can also be described pointwise:

$$x \in \overline{S} \iff \text{every open neighborhood of } x \text{ meets } S.$$

This description is often easier than the intersection-of-closed-sets definition. For instance,

$$\overline{\mathbb{Q}} = \mathbb{R}$$

in the usual topology because every interval contains rational numbers. Similarly,

$$\overline{\mathbb{R} \setminus \mathbb{Q}} = \mathbb{R}.$$

The two dense subsets are disjoint, showing that closure is not about having large measure or containing intervals; it is about meeting every neighborhood.

§147.8 Connectedness by continuous images

A useful theorem is that continuous images of connected spaces are connected. If X is connected and $f : X \rightarrow Y$ is continuous, then $f(X)$ is connected. Otherwise, if $f(X) = U \cup V$ were a separation, then $f^{-1}(U)$ and $f^{-1}(V)$ would separate X .

This theorem proves that intervals in $\mathbb{R}$ are connected. If an interval were split into two nonempty open pieces, one could construct a continuous map to the discrete two-point space, contradicting the intermediate value principle. Conversely, the intermediate value theorem is essentially a connectedness statement about intervals.

§147.9 Path-connectedness and components

Path-connectedness is stronger because it requires an explicit continuous path. The path-connected component of a point x is the set of all points reachable from x by paths. These components partition the space.

Every path-connected space is connected. The proof is simple: if $X = U \cup V$ were a separation and a path $\gamma : [0, 1] \rightarrow X$ started in U and ended in V , then $[0, 1]$ would be separated by $\gamma^{-1}(U)$ and $\gamma^{-1}(V)$, impossible because $[0, 1]$ is connected.

The converse fails in spaces like the topologist's sine curve. That example is not needed for computation, but it explains why topology keeps the two definitions separate.

§147.10 Next viewpoint

Continuity, closure, subspace topology, connectedness, and path-connectedness are open-set notions. Boundedness and completeness are not. The note is teaching what topology remembers and what it deliberately forgets.

§147.11 Continuity as a contravariant test

The preimage definition of continuity has a categorical flavor: open sets pull back. If

$$X \xrightarrow{f} Y \xrightarrow{g} Z$$

are continuous, then for every open $W \subseteq Z$,

$$(g \circ f)^{-1}(W) = f^{-1}(g^{-1}(W))$$

is open in X . This one-line identity is why continuous maps compose.

Images do not behave this cleanly. A continuous map may send open sets to non-open sets and closed sets to non-closed sets. Preimages are the stable operation; topology is built around that stability.

§147.12 Closure as an operator

The closure operation satisfies Kuratowski-style rules:

$$S \subseteq \overline{S}, \quad \overline{\overline{S}} = \overline{S}, \quad \overline{S \cup T} = \overline{S} \cup \overline{T}.$$

It also preserves inclusions: if $S \subseteq T$, then $\overline{S} \subseteq \overline{T}$.

These algebraic properties turn closure into more than a point-set definition. They let one reason with closure abstractly, without drawing neighborhoods each time.

§147.13 Why sequences are not enough

In metric spaces, closure can be detected by sequences. In arbitrary topological spaces, this may fail. The replacement notions are nets and filters. A net is like a sequence whose index set is not necessarily $\mathbb{N}$, but any directed set. This extra flexibility lets nets test all neighborhoods in spaces that are too large or too non-countable for ordinary sequences.

The warning is not meant to make topology harder for its own sake. It explains exactly which metric-space proof habits are safe and which are not. Sequential reasoning works in first-countable spaces, including metric spaces, but general topology requires a broader convergence language.

§147.14 Connectedness as maps to a two-point space

A space X is disconnected iff there is a nonconstant continuous map

$$X \rightarrow \{0, 1\}$$

where $\{0, 1\}$ has the discrete topology. Indeed, the preimages of 0 and 1 are disjoint nonempty clopen sets.

So connectedness can be read as: every continuous discrete-valued function is constant. This perspective is useful because it turns a separation problem into a mapping problem.

§147.15 Path components versus connected components

Path components always lie inside connected components. Equality holds in many familiar spaces, such as open subsets of $\mathbb{R}^n$, because they are locally path-connected. But equality can fail in pathological spaces such as the topologist's sine curve.

The difference matters conceptually. Connectedness says the space cannot be split by open sets. Path-connectedness says points can be joined by a continuous interval. The first is global and negative; the second is constructive.

§147.16 Local connectedness as the missing hypothesis

A space is locally path-connected if every point has a basis of path-connected neighborhoods. In such spaces, connected components and path components agree, and path components are open. This is why ordinary geometric regions behave better than arbitrary topological spaces.

Thus the hierarchy is:

$$\text{path-connected} \Rightarrow \text{connected},$$

but local path-connectedness often lets one recover path information from connectedness.

§147.17 Continuity, closure, and connectedness

This note is about what survives after distances disappear. Continuity is preimage stability, closure is unavoidable nearness, subspace topology is restriction, connectedness is the impossibility of a clopen split, and path-connectedness is reachability by intervals. The advanced warning is that metric intuition remains useful but not universal: sequences may need nets, and connectedness may not mean path-connectedness.

Homotopy, Simply Connected Spaces, Semigroups, Groups, and Homomorphisms

N-20221223

§148.1 Homotopy intuition

The first part explains homotopy through paths in $\mathbb{C}$. If $p, q \in \mathbb{C}$, many paths from p to q can be continuously deformed into one another while the endpoints stay fixed. But in

$$\mathbb{C} \setminus \{0\},$$

the missing origin becomes an obstacle. A path going around the hole one way may not be deformable into a path going the other way without crossing the missing point.

A homotopy between paths $f_0, f_1 : [0, 1] \rightarrow X$ is a continuous map

$$F : [0, 1] \times [0, 1] \rightarrow X$$

with $F(s, 0) = f_0(s)$, $F(s, 1) = f_1(s)$. For fixed-endpoint homotopy, the endpoints remain fixed throughout.

§148.2 Simply connected spaces

A path-connected space is simply connected if every loop can be contracted to a point. The plane $\mathbb{C}$ is simply connected; the punctured plane $\mathbb{C} \setminus \{0\}$ is not. This is the fundamental intuition behind algebraic topology: holes can be detected by loops that refuse to shrink.

§148.3 Semigroups, monoids, and groups

A semigroup is a set with an associative binary operation. A monoid is a semigroup with an identity element. A group is a monoid in which every element has an inverse:

$$\text{semigroup} \longrightarrow \text{monoid} \longrightarrow \text{group}.$$

A group G has associativity, identity, and inverses. If $ab = ba$ for all $a, b \in G$, it is abelian. Examples include $(\mathbb{Z}, +)$, S_n , and symmetry groups of geometric figures.

§148.4 Congruence, quotients, and homomorphisms

A congruence relation on an algebraic structure is an equivalence relation compatible with the operation. This lets operations descend to equivalence classes. For groups, normal subgroups play this role for quotient groups.

A group homomorphism $f : G \rightarrow H$ satisfies

$$f(ab) = f(a)f(b).$$

It automatically gives $f(e_G) = e_H$ and $f(a^{-1}) = f(a)^{-1}$. Its kernel

$$\ker f = \{g \in G : f(g) = e_H\}$$

is a normal subgroup, and the image is a subgroup.

§148.5 Loops and the beginning of the fundamental group

A loop based at $x_0 \in X$ is a path

$$\gamma : [0, 1] \rightarrow X, \quad \gamma(0) = \gamma(1) = x_0.$$

Two loops are considered equivalent if one can be continuously deformed into the other while keeping the basepoint fixed. The set of such equivalence classes becomes a group under concatenation:

$$[\gamma] \cdot [\eta] = [\gamma * \eta].$$

The identity is the constant loop, and the inverse of a loop is the same path traversed backward.

This explains why group theory appears immediately after homotopy in the note. The algebra is not an unrelated topic; it is the algebra of loops.

§148.6 Why the punctured plane has nontrivial loops

In $\mathbb{C} \setminus \{0\}$, the loop

$$\gamma(t) = e^{2\pi it}$$

winds once around the origin. It cannot be contracted to a point without crossing the missing origin. More generally,

$$\gamma_n(t) = e^{2\pi i n t}$$

winds n times. The integer n is the winding number. This gives the fundamental group

$$\pi_1(\mathbb{C} \setminus \{0\}) \cong \mathbb{Z}.$$

The plane itself has trivial fundamental group because every loop can be shrunk through the missing-free interior.

§148.7 Homomorphisms as structure-preserving maps

The algebraic part of the note introduces homomorphisms because topology will soon produce groups naturally. A homomorphism does not preserve elements individually; it preserves the operation:

$$f(ab) = f(a)f(b).$$

From this one derives

$$f(e_G) = e_H, \quad f(a^{-1}) = f(a)^{-1}.$$

The kernel records what the map collapses to the identity:

$$\ker f = \{g \in G : f(g) = e_H\}.$$

In topology, continuous maps induce homomorphisms on fundamental groups. Thus a map of spaces becomes a map of algebraic invariants.

§148.8 Quotients in topology and algebra

Both quotient spaces and quotient groups follow the same pattern: impose an equivalence relation and then ask whether the structure descends. For spaces, open sets in the quotient are defined by preimages under the quotient map. For groups, normal subgroups are exactly what make multiplication of cosets well-defined.

This shared pattern is one of the first conceptual bridges of algebraic topology. A loop space quotient by homotopy becomes a group; a group quotient by a normal subgroup becomes another group. In both cases, equivalence classes are not just bookkeeping but new mathematical objects.

§148.9 Higher viewpoint: from holes to algebra

The important transition is:

$$\text{hole in a space} \quad \leadsto \quad \text{nontrivial loop} \quad \leadsto \quad \text{group element.}$$

This is the beginning of algebraic topology. Instead of trying to draw all possible deformations, we assign algebraic invariants to spaces. If two spaces have different fundamental groups, they cannot be homeomorphic.

§148.10 A useful afterword

The note looks split between topology and algebra, but the split is exactly the point. Homotopy studies continuous deformation of paths; groups study composable reversible operations. The fundamental group combines them: homotopy classes of loops form a group, and that group records holes in the space.

§148.11 Concatenation of loops and why homotopy classes form a group

Given loops $\gamma, \eta : [0, 1] \rightarrow X$ based at x_0 , their concatenation is

$$(\gamma * \eta)(t) = \begin{cases} \gamma(2t), & 0 \leq t \leq \frac{1}{2}, \\ \eta(2t - 1), & \frac{1}{2} \leq t \leq 1. \end{cases}$$

This operation is not strictly associative on the nose, because parametrizations differ. It becomes associative after passing to homotopy classes. The identity is the constant loop, and the inverse is

$$\bar{\gamma}(t) = \gamma(1 - t).$$

Thus the group structure of $\pi_1(X, x_0)$ is already a quotient phenomenon: strict path operations become algebraic only after identifying homotopic paths.

§148.12 The exponential covering and the integer invariant

The map

$$p : \mathbb{R} \rightarrow S^1, \quad p(t) = e^{2\pi it}$$

is a covering map. A loop in S^1 based at 1 lifts uniquely to a path in $\mathbb{R}$ starting at 0. The endpoint of the lift is an integer, and that integer is the winding number.

This gives a concrete isomorphism

$$\pi_1(S^1) \cong \mathbb{Z}.$$

The generator is one counterclockwise turn. The inverse is one clockwise turn. Concatenating loops adds winding numbers.

§148.13 Punctured plane deformation retracts to the circle

The punctured plane $\mathbb{C} \setminus \{0\}$ has the same fundamental group as S^1 because it deformation retracts onto the unit circle:

$$H(z, t) = \left((1 - t) + \frac{t}{|z|} \right) z.$$

At $t = 0$, this is the identity. At $t = 1$, it sends z to $z/|z|$. The missing origin is exactly what makes this formula possible without crossing 0.

Thus

$$\pi_1(\mathbb{C} \setminus \{0\}) \cong \pi_1(S^1) \cong \mathbb{Z}.$$

The hole is measured by how many times a loop winds around it.

§148.14 Functoriality: maps of spaces give homomorphisms

A continuous based map

$$f : (X, x_0) \rightarrow (Y, y_0)$$

sends a loop γ in X to the loop $f \circ \gamma$ in Y . This respects homotopy and concatenation, so it induces a group homomorphism

$$f_* : \pi_1(X, x_0) \rightarrow \pi_1(Y, y_0).$$

This is the real reason group homomorphisms appear in the same note as homotopy. Algebraic topology assigns algebraic invariants to spaces and homomorphisms to continuous maps.

§148.15 Van Kampen as a computation machine

The Seifert–van Kampen theorem says, roughly, that if a space is built from open pieces U and V with path-connected overlap, then $\pi_1(U \cup V)$ is assembled from $\pi_1(U)$ and $\pi_1(V)$, with relations coming from the overlap.

For the figure-eight space $S^1 \vee S^1$, the result is

$$\pi_1(S^1 \vee S^1) \cong F(a, b),$$

the free group on two generators. The two loops are independent because the wedge point imposes no commutation relation.

§148.16 Group completion and why inverses matter

A monoid records composable operations with an identity; a group additionally allows reversal. Loops are naturally reversible by traversing them backward, so the fundamental group is a group rather than merely a monoid. In other contexts, one starts with a monoid and formally adds inverses; $\mathbb{N}$ under addition completes to $\mathbb{Z}$.

This mirrors the topological story: algebraic structures often arise by imposing equivalence relations and adding the formal operations needed to make the intended behavior exact.

§148.17 Loops, coverings, and algebraic invariants

The central bridge is:

continuous deformation of loops $\longrightarrow$ homotopy classes $\longrightarrow$ group elements.

The punctured plane shows the bridge concretely: a missing point creates winding number, winding number is an integer, and concatenation of loops becomes addition of integers.

Quotient Topology and Product Topology

N-20221228

§149.1 Interval modulo endpoints

The note begins with the quotient

$$D^1 = [-1, 1], \quad -1 \sim 1.$$

Identifying the two endpoints gives a circle:

$$D^1/(-1 \sim 1) \cong S^1.$$

A small neighborhood around the glued point corresponds to two half-interval neighborhoods at -1 and 1 upstairs. This is exactly the local reason quotient topology is defined by preimages.

More generally,

$$D^n/S^{n-1} \cong S^n.$$

Collapsing the boundary of a disk to one point produces a sphere.

§149.2 More quotient spaces

The note records the standard square identifications. If a square has opposite vertical sides identified in the same direction, one obtains a cylinder. If the two circular boundary components of that cylinder are then identified in the same direction, one obtains a torus. If one pair of edges is identified with a twist, one obtains a Möbius band or a Klein bottle depending on the full gluing pattern.

Another example is

$$\mathbb{R}/\mathbb{Z} \cong S^1,$$

where $x \sim y$ iff $y - x \in \mathbb{Z}$. This is the real line wrapped around the circle.

§149.3 Definition of quotient topology

Definition 149.3.1 (Quotient topology). Given a surjective map $q : X \rightarrow Y$, the quotient topology on Y declares $U \subseteq Y$ open exactly when $q^{-1}(U)$ is open in X .

Let $q : X \rightarrow X/\sim$ be the quotient map. A set $U \subseteq X/\sim$ is open iff

$$q^{-1}(U) \text{ is open in } X.$$

This is the finest topology on the quotient making q continuous. The note emphasizes that quotient topology is not guessed by drawing; it is forced by the universal property of the quotient map.

§149.4 Product topology

For spaces X, Y , the product topology on $X \times Y$ is generated by basic open rectangles

$$U \times V, \quad U \subseteq X \text{ open}, \quad V \subseteq Y \text{ open}.$$

Thus a set is open in $X \times Y$ if it is a union of such rectangles.

The page's picture of a small rectangle in a grid illustrates the basic rule: local neighborhoods in a product are products of local neighborhoods in the factors.

§149.5 Universal property of quotient topology

The quotient topology is not chosen by visual intuition alone. Let $q : X \rightarrow Y = X/\sim$ be the quotient map. A map $f : Y \rightarrow Z$ is continuous iff

$$f \circ q : X \rightarrow Z$$

is continuous. This is the universal property of the quotient topology.

This property explains the definition. We declare $U \subseteq Y$ open precisely when $q^{-1}(U)$ is open in X , because this is exactly what makes continuity out of the quotient detectable before quotienting.

§149.6 Endpoint identification in detail

For $[0, 1]/(0 \sim 1)$, points near the glued point come from two pieces upstairs: a small interval near 0 and a small interval near 1. This is why the quotient looks like a circle rather than an interval with a strange endpoint. The local neighborhoods have been glued.

Similarly, $D^2/S^1 \cong S^2$: collapse the boundary circle of a disk to a single point. The interior of the disk becomes the sphere with one point removed, and the collapsed boundary becomes the missing point.

§149.7 Product topology via projections

The product topology on $X \times Y$ is the coarsest topology making the projection maps

$$\pi_X : X \times Y \rightarrow X, \quad \pi_Y : X \times Y \rightarrow Y$$

continuous. Its basis consists of rectangles

$$U \times V$$

where $U \subseteq X$ and $V \subseteq Y$ are open.

The word “coarsest” matters. We do not add more open sets than necessary. Once all projection preimages $\pi_X^{-1}(U) = U \times Y$ and $\pi_Y^{-1}(V) = X \times V$ are open, finite intersections give $U \times V$, and arbitrary unions give the full product topology.

§149.8 Quotient and product as opposite directions

Product topology builds a larger space from two spaces while preserving maps out to the factors. Quotient topology builds a smaller space by identifying points while preserving maps out of the quotient. Product is controlled by projections; quotient is controlled by the quotient map.

This contrast is a useful memory:

product: maps out to factors, quotient: maps out after gluing.

Both definitions are forced by continuity requirements rather than by pictures.

§149.9 A second lens

Quotient topology and product topology are dual-looking constructions. A quotient identifies points and defines open sets by pulling back along the quotient map. A product combines spaces and defines open sets by rectangles from the two factors. Both are examples of topology being controlled by maps rather than by drawings alone.

§149.10 Saturated sets and quotient openness

For a quotient map $q : X \rightarrow Y$, a subset $A \subseteq X$ is saturated if it is a union of equivalence classes, equivalently

$$A = q^{-1}(q(A)).$$

Open sets in Y correspond exactly to saturated open sets in X . This is a practical way to test quotient topology: first check that the upstairs set is open, then check that it contains whole equivalence classes.

Endpoint gluing works because neighborhoods near the glued point have saturated preimages containing pieces near both endpoints. A one-sided interval near only one endpoint is not saturated, so it does not define a neighborhood downstairs.

§149.11 Quotients can destroy Hausdorffness

Not every quotient of a nice space is nice. For example, identify all rational numbers in $\mathbb{R}$ to one point and all irrational numbers to another point. The quotient has two points, but the preimage of either singleton is dense in $\mathbb{R}$, so the two quotient points cannot be separated by disjoint open neighborhoods.

This warning matters because quotient pictures can be misleading. The quotient topology is defined by preimages, and sometimes that produces spaces with poor separation properties.

§149.12 Product topology and the mapping property

The product topology is characterized by a universal property dual to the quotient one. A map

$$h : Z \rightarrow X \times Y$$

is continuous iff both coordinate maps

$$\pi_X \circ h : Z \rightarrow X, \quad \pi_Y \circ h : Z \rightarrow Y$$

are continuous. This is why the product topology is the coarsest topology making the projections continuous.

Thus product and quotient are both controlled by maps:

construction	distinguished map	continuity test
quotient	$q : X \rightarrow X/\sim$	f out of quotient via $f \circ q$
product	π_X, π_Y	h into product via coordinates

§149.13 From squares to surfaces through edge words

The quotient square diagrams can be encoded algebraically. A torus has boundary word

$$aba^{-1}b^{-1},$$

while a Klein bottle has a word of the form

$$aba^{-1}b.$$

The arrows in the picture determine whether a generator appears as a or a^{-1} . This is already the language of CW complexes and fundamental-group presentations.

So the square gluing is not merely visual. It tells us how a two-cell is attached to a one-skeleton, and the attaching word becomes a relation in the fundamental group.

§149.14 Products and quotients need not commute blindly

Although $(\mathbb{R}/\mathbb{Z}) \times (\mathbb{R}/\mathbb{Z})$ gives the torus, one should not assume products and quotients always interact without hypotheses. The map

$$X \times Y \rightarrow (X/\sim) \times (Y/\approx)$$

passes through the quotient by the product equivalence relation, but whether the resulting topology behaves as expected depends on the quotient maps and spaces involved.

This is another reason universal properties are safer than pictures: they tell us which maps are actually continuous and which topology is being imposed.

§149.15 Products, quotients, and gluing data

Quotient topology is about gluing while controlling maps out of the glued space. Product topology is about pairing spaces while controlling maps into the product. Endpoint identification, $\mathbb{R}/\mathbb{Z}$, torus gluings, and square diagrams are not separate examples; they are the same principle expressed through equivalence classes, saturated sets, and universal properties.

Algebraic Topology and Algebra: Products, Wedge Sums, CW Complexes, Cosets, and Surface Gluing

N-20221229

§150.1 What this note is really doing

This note is a working page where several threads meet: the torus as $(\mathbb{R}/\mathbb{Z})^2$, the product $S^1 \times S^1$, disjoint union and wedge sum, CW complexes, the algebra of cosets, and the square/polygon gluing pictures for surfaces. The original page should not be compressed into a list of definitions. Its real movement is: topology suggests quotient pictures; quotient pictures resemble congruence classes; congruence classes lead to cosets; then the topological quotient of a square returns as the torus and Klein bottle.

The top margin says that there is another intuitive understanding of

$$S^1 \times S^1 \cong T^2.$$

Think of a torus as a vertical circle carried around a horizontal circle. The path traced by that moving circle is the torus. This is not merely a picture: it is a parametrization.

下面写一下 $S^1 \times S^1 = \mathbb{T}$ 的方程解:

可以把 $\mathbb{T}$ 视为一个竖过来的 S^1 绕一个水平的 S^1 转一周, 其边界的轨迹形成的点集。




Figure 150.1: The original visual analysis of $S^1 \times S^1$: one parameter goes around the big circle, the other around the small circular fiber.

§150.2 $S^1 \times S^1$ as a parametrized torus

Write a point of the first circle as an angle θ , and a point of the second circle as an angle ϕ . For a torus with major radius R and minor radius r , a standard parametrization is

$$F(\theta, \phi) = ((R + r \cos \phi) \cos \theta, (R + r \cos \phi) \sin \theta, r \sin \phi).$$

This realizes $S^1 \times S^1 \rightarrow T^2$. The first angle decides where the circular cross-section is placed around the central hole; the second decides where we are on that cross-section.

The handwritten picture tries to choose a point A and decompose it into two circle parameters. That is exactly the right intuition, but the invariant statement is cleaner: the torus is the space of ordered pairs of circle-points, then realized geometrically by the map above.

§150.3 Disjoint union

For topological spaces X and Y , the disjoint union $X \amalg Y$ keeps the two spaces formally separate. A subset $U \subseteq X \amalg Y$ is open iff

$$U \cap X \text{ is open in } X, \quad U \cap Y \text{ is open in } Y.$$

If X and Y are nonempty, then $X \amalg Y$ is disconnected. Indeed, X and Y are both open in the disjoint union topology, disjoint, nonempty, and their union is the whole space. This proof is more precise than the vague statement that the spaces are simply “side by side”. The topology, not only the picture, gives the separation.

§150.4 Wedge sum

The wedge sum is obtained from a disjoint union by identifying chosen basepoints:

$$X \vee Y = (X \amalg Y) / (x_0 \sim y_0).$$

Thus disjoint union places spaces side by side, while wedge sum fuses them at one point. The prototype is $S^1 \vee S^1$, the figure-eight space.

Example 57.4.4 ($S^1 \vee S^1$ is a figure eight)

Let $X = S^1$ and $Y = S^1$, and let $x_0 \in X$ and $y_0 \in Y$ be any points. Then $X \vee Y$ is a “figure eight”: it is two circles fused together at one point.




Figure 150.2: The wedge sum $S^1 \vee S^1$ as two circles fused at one point.

This construction matters because loops in the two circles remain visibly independent. Algebraically, the fundamental group of the figure eight is the free group on two generators. So the picture is not a toy: it is one of the first spaces where topology produces noncommutative algebra.

§150.5 CW complexes: building spaces by cells

A CW complex is built inductively. Start with a set of 0-cells. Attach 1-cells by gluing endpoints to the 0-skeleton. Attach 2-cells by gluing the boundary circle of a disk to the 1-skeleton. Continue dimension by dimension:

$$X^0 \subseteq X^1 \subseteq X^2 \subseteq \cdots$$

A k -cell is a copy of D^k attached along S^{k-1} by an attaching map

$$\varphi_\alpha : S_\alpha^{k-1} \rightarrow X^{k-1}.$$

Informally: attach disks along their boundaries.

The note compares two CW structures on D^2 . One uses two 0-cells, two 1-cells, and one 2-cell. The other uses one 0-cell, one 1-cell whose two endpoints are glued to that point, and one 2-cell. Both build a disk, but the second is more economical.

Example 57.5.2 (D^2 with $2 + 2 + 1$ and $1 + 1 + 1$ cells)

- (a) First, we start with X^0 having two points e_a^0 and e_b^0 . Then, we join them with two 1-cells D^1 (green), call them e_c^1 and e_d^1 . The endpoints of each 1-cell (the copy of S^0) get identified with distinct points of X^0 ; hence $X^1 \cong S^1$. Finally, we take a single 2-cell e^2 (yellow) and weld it in, with its boundary fitting into the copy of S^1 that we just drew. This gives the figure on the left.
- (b) In fact, one can do this using just $1 + 1 + 1 = 3$ cells. Start with X^0 having a single point e^0 . Then, use a single 1-cell e^1 , fusing its two endpoints into the single point of X^0 . Then, one can fit in a copy of S^1 as before, giving D^2 as on the right.




Figure 150.3: Two CW descriptions of D^2 : one more literal, one more economical.

The lesson is that a CW decomposition is not unique. It is chosen because it makes later topology and algebra computable.

§150.6 The torus as a quotient of a square

The torus can be formed from a square by identifying opposite edges in the same direction. Walking off the right edge returns one at the corresponding point on the left edge; walking off the top returns one at the bottom. Algebraically,

$$T^2 \cong (\mathbb{R}/\mathbb{Z})^2 \cong S^1 \times S^1.$$

The quotient square viewpoint treats $[0, 1] \times [0, 1]$ as a fundamental domain for $(\mathbb{R}/\mathbb{Z})^2$.




Figure 150.4: The quotient square/cylinder/torus picture: $(\mathbb{R}/\mathbb{Z})^2$ as a glued square.

As a CW complex, the torus has one 0-cell, two 1-cells a, b , and one 2-cell. The boundary of the 2-cell is attached along the word

$$aba^{-1}b^{-1}.$$

This gives the familiar presentation

$$\pi_1(T^2) = \langle a, b \mid aba^{-1}b^{-1} = 1 \rangle \cong \mathbb{Z}^2.$$

The commutativity of $\mathbb{Z}^2$ is encoded by the boundary word of the square.

§150.7 Cosets as algebraic quotients

The algebra excerpt in the original pages is not random. It is the algebraic analogue of quotienting. If H is a subgroup of G , right congruence modulo H is

$$a \equiv_r b \pmod{H} \iff ab^{-1} \in H.$$

It is an equivalence relation. Reflexivity uses $aa^{-1} = e \in H$; symmetry uses $(ab^{-1})^{-1} = ba^{-1}$; transitivity uses

$$ac^{-1} = (ab^{-1})(bc^{-1}).$$

The equivalence class of a is the right coset

$$Ha = \{ha : h \in H\}.$$

The map $H \rightarrow Ha$, $h \mapsto ha$, is a bijection, so $|Ha| = |H|$. The group G is partitioned into cosets, and two cosets are either equal or disjoint. This is exactly the algebraic version of cutting a space into equivalence classes.

This is why the note can return from cosets to $(\mathbb{R}/\mathbb{Z})$ naturally. Both are quotient constructions: identify things differing by a subgroup.

§150.8 The Klein bottle

The Klein bottle is built like the torus, except one pair of edges is identified in the opposite direction. If one first glues a pair of edges, one gets a cylinder. But gluing the resulting boundary circles in opposite directions cannot be embedded in ordinary three-dimensional space without self-intersection. The usual picture in $\mathbb{R}^3$ is therefore an immersion, not a faithful embedding.




Figure 150.5: The Klein bottle: one edge identification is reversed, so the usual 3D picture self-intersects.

As a CW complex, the Klein bottle again has one 0-cell, two 1-cells, and one 2-cell, but the attaching word changes. A common presentation is

$$\pi_1(K) = \langle a, b \mid aba^{-1} = b^{-1} \rangle.$$

The relation says that going around one direction reverses the other. Thus the torus and Klein bottle have the same cell counts but different attaching words, and this difference controls orientability and fundamental group.

§150.9 Polygon schemata for surfaces

The last figure records polygon schemata for surfaces. A polygon with labelled, oriented boundary edges can be glued according to those labels. The resulting quotient is a surface. For example, the torus corresponds to

$$aba^{-1}b^{-1},$$

while an orientable surface of genus g has a boundary word of the form

$$a_1b_1a_1^{-1}b_1^{-1} \cdots a_gb_ga_g^{-1}b_g^{-1}.$$




Figure 150.6: Polygon schemata and surface gluings: topology encoded by labelled boundary words.

The handwritten comment says these diagrams took a long time to understand. That is exactly right. The picture compresses several ideas: cut a surface into a polygon, label the boundary, glue equal labels, and read the labels as algebraic generators and relations.

§150.10 Projective spaces as quotient spaces

The section title also mentions real and complex projective spaces. They fit the same quotient philosophy. Real projective space is

$$\mathbb{RP}^n = S^n / (x \sim -x),$$

the space of lines through the origin in $\mathbb{R}^{n+1}$. Complex projective space is the space of complex lines through the origin in $\mathbb{C}^{n+1}$, equivalently

$$\mathbb{CP}^n = S^{2n+1} / S^1.$$

A useful CW fact is that $\mathbb{RP}^n$ has one cell in each dimension $0, 1, \dots, n$, while $\mathbb{CP}^n$ has one cell in each even real dimension $0, 2, 4, \dots, 2n$.

§150.11 Cellular boundary words and fundamental groups

The attaching word of the two-cell is not a decorative label. For the torus, the word

$$aba^{-1}b^{-1}$$

becomes the relation saying that a and b commute. For the Klein bottle, the twist changes the relation to one of the form

$$aba^{-1} = b^{-1}.$$

Thus two spaces can have the same numbers of cells but different topology because the attaching maps differ.

§150.12 Why polygon schemata are efficient

A labelled polygon schema compresses a surface into boundary data. The labels record which edges are glued, and the arrows record orientation. Once this is understood, the diagram becomes an algebraic object: reading around the boundary gives a word in generators and inverses.

This is the reason algebraic topology can turn pictures into group presentations and chain complexes.

§150.13 Further context

The repaired version of this note should be read as one continuous story. First, $S^1 \times S^1$ parametrizes the torus. Second, a square with opposite sides identified gives the same torus as a quotient. Third, disjoint union and wedge sum explain how spaces can be combined before quotienting. Fourth, CW complexes systematize this by attaching cells. Fifth, cosets provide the algebraic analogue of quotienting. Finally, torus, Klein bottle, and polygon schemata show that a whole surface can be encoded by a boundary word.

The central lesson is that algebraic topology is not simply topology plus algebra added later. The topology is already algebraic as soon as we record how pieces are glued: arrows, labels, words, equivalence classes, and quotient maps are the bridge.

§150.14 Cellular chain complex of the torus

The torus CW structure with one 0-cell, two 1-cells a, b , and one 2-cell gives a cellular chain complex

$$0 \rightarrow C_2 \cong \mathbb{Z} \xrightarrow{\partial_2} C_1 \cong \mathbb{Z}^2 \xrightarrow{\partial_1} C_0 \cong \mathbb{Z} \rightarrow 0.$$

Since both ends of each one-cell attach to the single zero-cell, $\partial_1 = 0$. The two-cell is attached by the commutator word

$$aba^{-1}b^{-1}.$$

In cellular homology, the total exponent of a is $1 - 1 = 0$, and the total exponent of b is $1 - 1 = 0$, so $\partial_2 = 0$. Therefore

$$H_0(T^2) \cong \mathbb{Z}, \quad H_1(T^2) \cong \mathbb{Z}^2, \quad H_2(T^2) \cong \mathbb{Z}.$$

The same square diagram that gives the fundamental group also gives homology by linearizing the boundary word.

§150.15 Klein bottle: same cells, different boundary

The Klein bottle has the same cell counts as the torus: one 0-cell, two 1-cells, and one 2-cell. But the attaching word is different. With a common convention, it is

$$aba^{-1}b.$$

The total exponent of a is 0, while the total exponent of b is 2. Hence the cellular boundary has the form

$$\partial_2(1) = 2b.$$

So

$$H_1(K) \cong \mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}, \quad H_2(K) = 0.$$

The absence of $H_2 \cong \mathbb{Z}$ is the homological shadow of non-orientability.

§150.16 Euler characteristic as cell bookkeeping

For a finite CW complex,

$$\chi(X) = \sum_k (-1)^k c_k,$$

where c_k is the number of k -cells. The torus and Klein bottle both have

$$\chi = 1 - 2 + 1 = 0.$$

Thus Euler characteristic cannot distinguish them. Their difference lies not in the number of cells but in the attaching map of the two-cell. This is the core lesson of CW complexes: cell counts are only the first approximation; attaching maps carry the topology.

§150.17 Orientability from polygon words

In polygon schemata, orientability can be read from edge orientations. For an orientable surface, each label appears once in each direction, as in

$$aba^{-1}b^{-1}.$$

For a non-orientable surface, some label appears twice with the same direction after normalizing the word, as in projective-plane-type words.

This explains why the torus is orientable and the Klein bottle is not. The twist in the edge gluing reverses local orientation after going around a loop.

§150.18 Van Kampen and presentations from gluing

The one-skeleton of the torus is $S^1 \vee S^1$, whose fundamental group is the free group $F(a, b)$. Attaching the two-cell along $aba^{-1}b^{-1}$ imposes the relation

$$aba^{-1}b^{-1} = 1.$$

Thus

$$\pi_1(T^2) = \langle a, b \mid aba^{-1}b^{-1} = 1 \rangle \cong \mathbb{Z}^2.$$

For the Klein bottle, the attaching word imposes a different relation, giving

$$\pi_1(K) = \langle a, b \mid aba^{-1} = b^{-1} \rangle.$$

This is the precise algebraic meaning of the arrows on the edge-identification diagram.

§150.19 Projective spaces through cells

The quotient descriptions

$$\mathbb{RP}^n = S^n / (x \sim -x), \quad \mathbb{CP}^n = S^{2n+1} / S^1$$

are not just definitions. They explain the CW structures:

$$\mathbb{RP}^n : \text{one cell in every dimension } 0, 1, \dots, n,$$

while

$$\mathbb{CP}^n : \text{one cell in every even dimension } 0, 2, \dots, 2n.$$

For $\mathbb{CP}^n$, this leads to homology groups $\mathbb{Z}$ in even degrees and 0 in odd degrees. Thus a quotient construction immediately produces computable algebraic invariants.

§150.20 Cellular surfaces through homology and presentations

This note is about turning pictures into algebra. A product picture gives $S^1 \times S^1$. A gluing picture gives a quotient. A CW picture gives chain groups and boundary maps. A labelled boundary word gives a fundamental-group presentation and, after abelianization, homology. The advanced lesson is that algebraic topology is not added after the drawings; it is already encoded in the arrows, labels, and attaching maps.

Hopf Fibration, Homotopy, and the Infinite-Stability Reading

N-20240727

§151.1 What this chapter is for

This note records the first route by which the line

$$0 = 1 - 1 = -1 + 1 = 0$$

became mathematically suggestive to me. It should be read as a working interpretation rather than as the final explanation. The important point is to preserve the original path of thought in full detail.

The triggering background was a lecture poster by Peng Keyao, whose abstract seemed to connect elementary arithmetic, the Hopf fibration, homotopy groups, configuration spaces, algebraic K -theory, and the strange equation above. The poster immediately reminded me of an older lecture on the Hopf fibration. At that time the picture of linked circles in S^3 was one of the first images that made topology feel like something deeper than point-set definitions.

But the line itself first looked analytic and algebraic rather than topological. The string

$$1 - 1 = -1 + 1$$

looks like rearrangement, cancellation, and ambiguity of formal sums. So the first reading I tried to build was:

alternating sums $\longrightarrow$ regularization
 $\longrightarrow$ Eilenberg swindle
 $\longrightarrow$ vanishing under infinite stability.

This note unfolds that reading carefully.

§151.2 The first analytic reflex

The most immediate association is the infinite alternating series

$$1 - 1 + 1 - 1 + \cdots$$

Its partial sums are

$$1, 0, 1, 0, \dots,$$

so the ordinary limit does not exist.

Definition 151.2.1 (Cesàro summation). Let $s_n = a_0 + \cdots + a_n$ be the partial sums of a series. If the averages

$$\sigma_N = \frac{s_0 + s_1 + \cdots + s_N}{N+1}$$

converge to L , then the series is Cesàro summable to L .

For the alternating series, the Cesàro averages of the partial sums tend to

$$\frac{1}{2}.$$

Thus the expression is not convergent in the usual sense, but it has a regularized value in the Cesàro sense.

Definition 151.2.2 (Abel summation). Given a formal series $\sum_{n \geq 0} a_n$, form its power series

$$F(r) = \sum_{n \geq 0} a_n r^n$$

for $0 < r < 1$. If $F(r)$ has a limit L as $r \rightarrow 1^-$, then the original series is Abel summable to L .

For

$$1 - 1 + 1 - 1 + \cdots,$$

one has

$$1 - r + r^2 - r^3 + \cdots = \frac{1}{1+r},$$

so Abel summation again gives $1/2$.

Remark 151.2.3 — The same cultural memory also points toward the analytic continuation of the Dirichlet eta function. At the level of this first note, the important point is not the exact analytic formalism, but the habit of reading a formally divergent expression through a more flexible rule.

This is why the line

$$0 = 1 - 1 = -1 + 1 = 0$$

first felt like a statement about how formal algebra changes when one passes from finite arithmetic to limiting procedures. In ordinary arithmetic, the line is trivial. In the world of infinite expressions, however, the order of grouping and the rule of interpretation may change what the expression seems to mean.

Claim 151.2.4 (First analytic slogan). The equation should perhaps be read not as a numerical equality, but as evidence that cancellation becomes unstable once infinite formal processes are allowed.

§151.3 Why the Hopf and K-theory hints changed the search

The analytic interpretation alone did not explain why the meme included the Hopf fibration and K -theory. Those are not the usual tools for assigning values to divergent series. The Hopf fibration suggests homotopy and stable phenomena. K -theory suggests addition, cancellation, and group completion of mathematical objects.

So the first analytic reflex had to be upgraded. Instead of asking only whether

$$1 - 1 + 1 - 1 + \cdots$$

has a regularized value, I began asking whether there is a categorical or K -theoretic situation in which cancellation becomes legitimate because an object is infinitely large or infinitely stable.

That is exactly the kind of situation where the Eilenberg swindle appears.

§151.4 The Eilenberg swindle as the central guess

The Eilenberg swindle is the mechanism that made the analytic reflex look relevant to K -theory.

Theorem 151.4.1 (Eilenberg swindle, working form)

Suppose an additive category has countable direct sums and an object

$$S = A \oplus A \oplus A \oplus \cdots.$$

Then

$$S \cong A \oplus S.$$

Consequently, any Grothendieck-type invariant K_0 compatible with direct sums satisfies

$$[S] = [A] + [S],$$

and therefore forces

$$[A] = 0.$$

Idea. The isomorphism is the shift of the infinite direct sum:

$$A \oplus A \oplus A \oplus \cdots \cong A \oplus (A \oplus A \oplus A \oplus \cdots).$$

The finite summand A is absorbed by the infinite tail. Passing to an additive invariant converts this absorption into a vanishing statement. $\square$

This is the algebraic version of the same psychological picture as the alternating series. One can group an infinite expression in different ways:

$$(1 - 1) + (1 - 1) + \cdots = 0,$$

but also

$$1 + (-1 + 1) + (-1 + 1) + \cdots = 1.$$

The paradox is not that arithmetic is broken. The point is that infinite formal manipulation can erase the distinction between adding one copy and adding no copy, because the tail of the expression absorbs the finite change.

Remark 151.4.2 — This is the reason I originally took the line $0 = 1 - 1 = -1 + 1 = 0$ as a compact symbolic form of infinite absorption. The two middle expressions are different finite presentations, but after entering an infinitely stable world they both collapse back to the same zero object.

§151.5 How this looked like a K-theory statement

In algebraic K -theory, one often begins with a category of objects equipped with a direct-sum operation. The first invariant records formal additive differences.

Definition 151.5.1 (Grothendieck group). Let M be a commutative monoid. Its Grothendieck group $K_0(M)$ is the universal abelian group receiving a monoid homomorphism from M . Concretely, elements may be represented by formal differences

$$[A] - [B],$$

subject to the relation

$$[A \oplus C] - [B \oplus C] = [A] - [B].$$

The expression $1 - 1$ then resembles a formal difference between one positive copy and one negative copy. The Eilenberg-swindle reading was that some sufficiently infinite environment might force such formal differences to vanish.

Proposition 151.5.2 (Vanishing by absorption, heuristic form)

If a category contains an object S with

$$S \cong S \oplus A,$$

then any additive invariant that sees direct sum as addition tends to identify the class of A with zero.

This is why swindle arguments prove vanishing theorems in categories with infinite co-products or infinite-dimensional stabilization. At the level of intuition, the meme ladder became a short chain: Hopf fibration points toward stable homotopy; stable homotopy points toward group completion in K -theory; group completion suggests infinite stability and the Eilenberg swindle; the final line is read as symbolic infinite cancellation. In this reading, the final line is a “God-tier” arithmetic trick: after enough stabilization, the complicated machinery collapses into the fact that infinity can absorb a finite summand.

§151.6 The role I imagined for the Hopf fibration

The Hopf fibration still needed a place in this interpretation. I imagined it as the geometric entrance into stable homotopy. The Hopf map

$$S^3 \rightarrow S^2$$

is a concrete example of a map whose nontriviality is not visible from ordinary dimension-counting. It teaches that topology can remember a hidden process.

From there, stable homotopy suggests that some phenomena persist after suspending or stabilizing. In the original reading, I loosely connected this with the word “stable” in algebra: perhaps after passing to a sufficiently stabilized category, the same kind of hidden structure becomes algebraic cancellation, and the end result is a vanishing statement.

This was not a precise theorem. It was a ladder of associations:

$$\text{Hopf} \rightsquigarrow \text{stable homotopy} \rightsquigarrow \text{stable algebra} \rightsquigarrow \text{swindle}.$$

But as a first notebook explanation, it made the meme feel less arbitrary. The Hopf fibration stood for the beginning of stable topology; the equation stood for the end of a stabilization trick.

§151.7 The meme as a learning ladder

The meme that triggered the reflection should be kept as part of the record. It should not be read as a proof. It is a ladder of associations.



Figure 151.1: The meme that compressed Hopf fibration, stable homotopy, K -theory, and the equation $0 = 1 - 1 = -1 + 1 = 0$ into one ladder.

In the first reading, the ladder was interpreted as follows: In the first reading, the ladder was interpreted as a sequence of working meanings. The Hopf fibration meant hidden topology; homotopy groups meant invariants of deformation; stable homotopy meant persistence under stabilization; K -theory meant organized addition and cancellation; the Eilenberg swindle meant vanishing by infinite stability. This is why the original explanation was not merely a random guess about divergent series. It was trying to rec-

oncile every visible clue in the meme with a known mechanism in K -theory.

§151.8 What the original interpretation claimed

The original interpretation can be stated cleanly as a working thesis.

Theorem 151.8.1 (Original working interpretation, not yet the final reading)

In elementary arithmetic, $0 = 1 - 1 = -1 + 1 = 0$ is trivial. In an infinitely stable categorical setting, however, finite additions and cancellations can be absorbed by an infinite tail. The Eilenberg swindle expresses this absorption. Therefore the equation may be read as a symbolic way to say that certain K -theoretic invariants vanish once infinite direct-sum operations are allowed.

This reading emphasizes three words:

infinity, stability, vanishing.

The line $0 = 1 - 1 = -1 + 1 = 0$ then becomes an emblem of the idea that a sophisticated invariant may collapse because the ambient category is too large.

§151.9 A slightly uneasy stopping point

Even at this first stage, there was something slightly unstable about the explanation. The lecture poster mentioned finite sets, configuration spaces, and the Hopf fibration, while the swindle explanation relies heavily on infinite objects. The analytic examples also use infinite series and limiting procedures. This suggested that the original route might be pointing toward the right neighbourhood but not yet at the exact object.

Still, the first interpretation is worth preserving. It captured several genuine connections:

- alternating sums explain why $1 - 1$ and $-1 + 1$ feel like more than ordinary arithmetic;
- the Eilenberg swindle explains how infinite stabilization can trivialize K -theoretic information;
- K -theory really is about direct sum, group completion, and cancellation;
- the Hopf fibration really does point toward stable homotopy.

So this note stops at the original working picture. The central claim of that picture is not that 0 is secretly nonzero. It is that in an infinitely stabilized algebraic world, the distinction between adding and not adding a finite piece can disappear. The later note will separate this first interpretation from the finite homotopical mechanism actually used in the lecture.

A Nontrivial Loop at Zero: FinSet K-Theory and Stable Homotopy

N-20251217

§152.1 Why this second note exists

This note is the reflective version of the earlier summer memo. The first reaction to the line

$$0 = 1 - 1 = -1 + 1 = 0$$

was not the answer eventually given by the lecture. A few days before writing this correction, while I was organizing my homepage, my roommate Su Yuyang saw the same line. His immediate thought was the familiar analytic story around

$$\sum_{k \geq 0} (-1)^k = 1 - 1 + 1 - 1 + \cdots,$$

whose Cesàro value and Abel value are $1/2$, and which is also connected with the analytic continuation of the Dirichlet eta function. This was also my own first thought. The string $1 - 1 = -1 + 1$ looks like a rearrangement or regularization trick before it looks like topology.

But the meme and the lecture poster clearly featured the Hopf fibration and K -theory. That made the analytic-summation explanation feel too small. Following the same original line of thought, I then searched for a K -theoretic or categorical mechanism in which an apparently impossible cancellation becomes legitimate. This led naturally to the Eilenberg swindle.

The model I had in mind was this. If an object admits an infinite decomposition

$$S = A \oplus A \oplus A \oplus \cdots,$$

then adding or removing one copy of A may not change S . Formally, one shifts the infinite direct sum:

$$S \cong A \oplus S.$$

In a Grothendieck group this kind of absorption can force

$$[A] = 0.$$

In the language of the meme, I was reading

$$0 = 1 - 1 = -1 + 1 = 0$$

as a symbolic version of infinite stability: by regrouping an infinite alternating expression, one can make the same expression look like 0, 1, or a vanishing object. This was my “Level 6” interpretation: the ultimate trick behind the line is that infinity absorbs finite differences, and therefore certain K -groups or invariants vanish.

That guess was not completely unrelated. It correctly noticed that K -theory is about addition, cancellation, group completion, and stability. It also correctly noticed that

an equality such as $0 = 0$ might be hiding a process rather than an elementary arithmetic fact. But the direction was almost the opposite of Prof. Peng's point. My first interpretation emphasized:

infinite objects,
analytic or algebraic cancellation,
triviality and vanishing.

The lecture's point is instead:

finite sets,
homotopy of a finite creation–swap–annihilation process,
nontriviality of a loop.

So the slogan of this note is:

the important object is not the endpoint 0, but the loop based at 0.

The loop is created by making a $+$ and a $-$, swapping them, and annihilating them again:

$$\emptyset \sim \{+, -\} \sim \{-, +\} \sim \emptyset.$$

Under the K -theory of finite sets, this loop becomes a nontrivial element. Under the Barratt–Priddy–Quillen theorem, it corresponds to the stable Hopf class.

§152.2 The lecture starts from ordinary equations

The slide deck begins with familiar equations: Pythagoras, Euler's identity, Stokes' theorem, the Navier–Stokes equation, the Langlands slogan, and then the strange meme equation. This is pedagogically important. It suggests that an equation is not merely a string of symbols with a truth value. An equation may also encode a method, a structure, or a path.

In elementary arithmetic,

$$0 = 1 - 1 = -1 + 1 = 0$$

is boring. Both sides are zero. But the lecture asks us to stop looking only at the value. In a higher structure, there can be more than one way for an object to equal itself. The identity loop at 0 and the creation–swap–annihilation loop at 0 need not be the same.

Remark 152.2.1 — This is the conceptual bridge to homotopy type theory and higher category theory. A statement $a = b$ can be represented by a path from a to b . Then two proofs of equality may themselves be compared by a homotopy. The lecture uses this philosophy in a concrete topological model rather than as formal type theory.

§152.3 The Hopf fibration as the first nontrivial map

The first mathematical anchor is the Hopf fibration. Define

$$S^3 = \{(z_1, z_2) \in \mathbb{C}^2 : |z_1|^2 + |z_2|^2 = 1\}.$$

The Hopf map is

$$H : S^3 \rightarrow \mathbb{CP}^1 \cong S^2, \quad (z_1, z_2) \mapsto [z_1 : z_2].$$

For every point of S^2 , its preimage is a circle. These circles are linked in a highly organized way.




Figure 152.1: the Hopf fibration visualized by linked fibres.

The Hopf fibration matters because it is not null-homotopic. It represents a nontrivial element of

$$\pi_3(S^2).$$

In beginner language: even though S^3 and S^2 are both spheres, the map $H : S^3 \rightarrow S^2$ cannot be continuously shrunk to a constant map.

Definition 152.3.1 (Null-homotopic map). A map $f : X \rightarrow Y$ is null-homotopic if it is homotopic to a constant map. Equivalently, there is a continuous family of maps $f_t : X \rightarrow Y$ with $f_0 = f$ and f_1 constant.

The Hopf map is the first warning: a process can be topologically nontrivial even if it looks, from far away, like a map between simple spaces.

§152.4 Homotopy groups and the first stable class

For a based space X , the group $\pi_n(X)$ is made from based maps $S^n \rightarrow X$, modulo based homotopy. The lecture reviews the low-dimensional intuition:

$$\pi_0(X) = \text{connected components}, \quad \pi_1(X) = \text{loops modulo homotopy}.$$

Higher homotopy groups generalize this to maps from higher-dimensional spheres.

For spheres, the groups are easy below the diagonal:

$$\pi_i(S^n) = 0 \quad (i < n).$$

On the diagonal, one has the degree:

$$\pi_n(S^n) \cong \mathbb{Z}.$$

But above the diagonal the world becomes subtle. The Hopf map gives the famous example

$$\pi_3(S^2) \neq 0.$$

The stable homotopy groups of spheres are obtained by repeatedly suspending:

$$\pi_i^s(\mathbb{S}) = \lim_{n \rightarrow \infty} \pi_{i+n}(S^n).$$

The Hopf map contributes the first stable nontrivial class

$$\eta \in \pi_1^s(\mathbb{S}) \cong \mathbb{Z}/2.$$

This class has order two: doing it twice becomes trivial, but doing it once is not trivial.

§152.5 From natural numbers to finite sets

The lecture then changes viewpoint. Instead of defining the natural numbers as primitive symbols, define them as isomorphism classes of finite sets:

$$\mathbb{N} \cong \pi_0(\text{FinSet}).$$

Here a finite set of size n represents the number n . Disjoint union represents addition:

$$A + B := A \sqcup B.$$

To get the integers, one takes a group completion. Informally, one allows formal differences

$$A - B.$$

Thus the number $1 - 1$ is represented by a pair consisting of one positive point and one negative point.

Remark 152.5.1 — The corresponding slide phrases this as a move from numbers to finite sets: the natural number n is represented by an n -element set, and addition is represented by disjoint union. This is the first categorification step in the lecture: instead of remembering only the cardinality, one keeps the finite set itself and later also its symmetries.

Definition 152.5.2 (Grothendieck group). Given a commutative monoid M , its Grothendieck group $K_0(M)$ is the universal abelian group receiving a monoid homomorphism from M . For finite sets under disjoint union this gives

$$K_0(\text{FinSet}) \cong \mathbb{Z}.$$

So far, this still only explains the value of $1 - 1$. It does not yet explain why the path

$$0 \rightarrow 1 - 1 \rightarrow -1 + 1 \rightarrow 0$$

can be interesting. For that, we must remember not just the set of isomorphism classes, but the space of finite sets and isomorphisms.

§152.6 The space of finite sets

The slide deck asks us to consider the moduli space

$$\mathcal{F} = \text{FinSet} / \sim.$$

This is not merely a set of cardinalities. It remembers automorphisms of finite sets. A finite set with n elements has automorphism group S_n , the symmetric group.

Remark 152.6.1 — The slide at this point depicts the moduli space of finite sets as a space whose n -point component remembers S_n . In notebook language, the important formula is

$$\mathcal{F} \simeq \bigsqcup_{n \geq 0} BS_n.$$

Thus the component remembers the number n , while loops in that component remember permutations.

One can think of $\mathcal{F}$ as a disjoint union of classifying spaces:

$$\mathcal{F} \simeq \bigsqcup_{n \geq 0} BS_n.$$

This formula is the first place where the hidden loops appear. A point in BS_n remembers a finite set of size n , but loops at that point remember permutations of the set.

Thus the finite set with one positive and one negative point is not just a cardinality-zero object. It can have a nontrivial path coming from swapping the two labelled positions before annihilation.

§152.7 The K -space of finite signed sets

To group-complete finite sets at the level of spaces, the slides introduce a space $K(\mathcal{F})$. Its points are finite ordered configurations whose elements are marked by $+$ or $-$. Adjacent opposite signs may cancel.

Remark 152.7.1 — The slide introduces the K -space by replacing finite sets with signed finite configurations. The picture is simple enough to state in words: points marked $+$ and points marked $-$ may be created and cancelled in opposite-signed pairs, and the empty configuration is used as the base point.

The empty configuration $\emptyset$ is the base point. The configuration $\{+, -\}$ has total signed cardinality zero, but it is not identical to the empty configuration before one performs cancellation. This distinction is the seed of the whole example.

The higher K -groups are defined by

$$K_i(\text{FinSet}) = \pi_i(K(\mathcal{F}), \emptyset).$$

Thus K_0 records components, while K_1 records loops based at the empty configuration.

§152.8 The loop h

Now the lecture constructs the loop

$$h : \emptyset \sim \{+, -\} \sim \{-, +\} \sim \emptyset.$$

It has three stages.

First, create a cancelling pair:

$$\emptyset \longrightarrow \{+, -\}.$$

Second, swap the two points:

$$\{+, -\} \longrightarrow \{-, +\}.$$

Third, annihilate the cancelling pair:

$$\{-, +\} \longrightarrow \emptyset.$$

K-groups of finite sets

Let $K_i(\text{FinSet}) = \pi_i(K(F), \emptyset)$ to be the higher K-groups, we have the following observations:

- $\pi_0 K(F) = \mathbb{Z}$, by $x \mapsto \|x\|$. This coincides with $K_0(\text{FinSet})$.
- We can construct a loop $h \in \pi_1(K(F), \emptyset)$ by using identifications

$$h : \emptyset \sim \{+, -\} \sim \{-, +\} \sim \emptyset$$

- In fact $\pi_1(K(F), \emptyset) = S_n^{ab} = \mathbb{Z}/2$ with h as the generator. That is to say, h is a non-trivial identification.

14/27

◀ ▶ ↺ ↻

Figure 152.2: From the lecture slides: the loop h represented by creation, swap, and annihilation.

This is the corrected meaning of

$$0 = 1 - 1 = -1 + 1 = 0.$$

It is not the claim that zero equals zero. It is the claim that there is a particular loop at zero. The middle step, the swap, is the nontrivial part.

Remark 152.8.1 — If the two points were not moved, the movie would look like a boring creation followed by immediate annihilation. The swap forces a braid-like motion. In configuration space, this motion wraps around the collision locus where the two points would meet.

The slides record the key computation:

$$K_1(\text{FinSet}) \cong \mathbb{Z}/2,$$

and the loop h gives the nontrivial element. Doing the swap twice is null-homotopic, but doing it once is not.

§152.9 Why the first guess was close but reversed

The analytic reading begins with the infinite alternating expression

$$1 - 1 + 1 - 1 + \cdots .$$

Cesàro summation replaces the raw partial sums

$$1, 0, 1, 0, \dots$$

by their averages, which tend to $1/2$. Abel summation studies

$$1 - r + r^2 - r^3 + \dots = \frac{1}{1 + r}$$

and then lets $r \rightarrow 1^-$, again obtaining $1/2$. Analytic continuation of the Dirichlet eta function gives another sophisticated version of the same cultural memory.

This is why Su Yuyang's first reaction and my first reaction were natural. The line

$$0 = 1 - 1 = -1 + 1 = 0$$

looks as if it is inviting a discussion of how formal rearrangements, limiting processes, and regularizations change the meaning of an alternating expression.

The Eilenberg swindle pushes that same instinct from analysis into algebra. Suppose a category admits countable direct sums, and an object S satisfies

$$S \cong A \oplus S.$$

Then in a Grothendieck group one obtains

$$[S] = [A] + [S],$$

so formally

$$[A] = 0.$$

This is the mechanism behind many vanishing arguments: an infinite object absorbs finite additions, and the invariant becomes trivial. In my first interpretation, the meme equation was a compressed version of this phenomenon. The “truth” was supposed to be that infinity makes cancellation too flexible, so the sophisticated K -theoretic machinery collapses into an infinite arithmetic trick.

The actual lecture keeps several words from that interpretation but reverses their role. It still cares about K -theory, addition, cancellation, and stability. But it does not use infinity to erase a finite object. It uses finite configurations to produce a nontrivial loop. The contrast is:

first interpretation	lecture's mechanism
infinite alternating series	finite signed configuration
Cesàro/Abel/eta regularization	creation–swap–annihilation path
Eilenberg swindle	$K(\text{FinSet})$
absorption at infinity	nontrivial loop at the base point
vanishing of invariants	generator of $K_1(\text{FinSet}) \cong \mathbb{Z}/2$
triviality	stable Hopf nontriviality

Thus the first guess was close in vocabulary but opposite in direction. It looked for the triviality of the infinite. Prof. Peng's point is the nontriviality of the finite.

§152.10 Finite sets and stable spheres

The surprising theorem behind the example is the Barratt–Priddy–Quillen theorem. In the form used by the lecture, it says that the K -theory of finite sets is the sphere spectrum:

$$K_i(\text{FinSet}) \cong \pi_i^s(\mathbb{S}).$$

Theorem 152.10.1 (Barratt–Priddy–Quillen, lecture form)

The K -theory of finite sets is the sphere spectrum. Equivalently,

$$K_i(\text{FinSet}) \cong \pi_i^s(\mathbb{S}).$$

In the slide deck this theorem is presented as the conceptual bridge: a finite-set loop can detect a stable homotopy class of spheres.

This theorem explains why the finite-set loop h can be compared to the Hopf class. Under BPQ, the nontrivial element of

$$K_1(\text{FinSet}) \cong \mathbb{Z}/2$$

corresponds to

$$\eta \in \pi_1^s(\mathbb{S}) \cong \mathbb{Z}/2.$$

This is the stable Hopf class.

Another way to say this is that $K(\mathcal{F})$ can be seen as a stable configuration space of framed points:

$$K(\mathcal{F}) \simeq \text{Conf}^{\text{fr}}(\mathbb{R}^\infty).$$

The creation–swap–annihilation loop becomes a framed one-dimensional motion. Its framing carries the $\mathbb{Z}/2$ -twist.

§152.11 Relation with the classical Hopf map

The classical Hopf fibration represents a nontrivial element of $\pi_3(S^2)$. Stabilizing the Hopf map gives the stable element

$$\eta \in \pi_1^s(\mathbb{S}).$$

The finite-set loop h gives the same class under the BPQ theorem.

Thus the meme ladder is not merely free association. The chain is:

$$\begin{array}{c} \text{Hopf fibration } S^3 \rightarrow S^2 \\ \Downarrow \\ \text{stable Hopf element } \eta \in \pi_1^s(\mathbb{S}) \\ \Downarrow \\ K_1(\text{FinSet}) \cong \pi_1^s(\mathbb{S}) \\ \Downarrow \\ \text{the loop } h : \emptyset \sim \{+, -\} \sim \{-, +\} \sim \emptyset. \end{array}$$


This is the clean mathematical sense in which the arithmetic-looking line is the Hopf map in disguise.

§152.12 Cobordism: the movie becomes a circle

The slides then place the motion in three-dimensional space. As time evolves, the moving point configuration sweeps out a one-dimensional manifold. The creation of the pair, their swap, and their annihilation form a closed loop. Because the points carry signs or framings, the resulting circle is a framed one-manifold.

What about cobordism?

If we place this motion in 3-dimensional space, we obtain a circle:



17/27



Figure 152.3: From the lecture slides: the creation–swap–annihilation movie becomes a framed circle.

Framed cobordism asks whether this framed circle is the boundary of a framed disk. The answer is no for the loop h . This is another expression of the same nontrivial $\mathbb{Z}/2$ -class.

Definition 152.12.1 (Framed cobordism, informal). A framed n -manifold is an n -manifold equipped with a trivialization of its normal bundle. Two framed manifolds are framed cobordant if they form the boundary of a framed manifold one dimension higher.

This is the geometric face of the equation. The path starts and ends at the empty configuration, but its trace is not the boundary of a trivial framed disk. The loop is geometrically nonzero.

§152.13 Algebraic K -theory of a field

After finite sets, the lecture repeats the construction with finite-dimensional vector spaces over a field k . Let

$$\mathcal{V}_k = \text{FinVect}_k / \sim$$

be the moduli space of finite-dimensional vector spaces. Under direct sum, it behaves like a categorified version of addition.

The low K -groups are familiar:

$$K_0(k) \cong \mathbb{Z}, \quad K_1(k) \cong k^\times.$$

The first says that a vector space is measured by its dimension after group completion. The second says that the determinant detects the abelianized automorphism information:

$$\text{GL}_\infty(k)^{\text{ab}} \cong k^\times.$$

This parallels the finite-set case:

object	K_0	K_1
FinSet	$\mathbb{Z}$	$\mathbb{Z}/2$
FinVect $_k$	$\mathbb{Z}$	$k^\times$

The lecture now asks for a geometric model of these algebraic K -groups beyond degree one.

§152.14 Intrinsic skein-lasagna

The advanced second half of the slide deck introduces intrinsic skein-lasagna. This is not merely decorative. It is a proposed geometric language for K -groups analogous to how framed cobordism gives a geometric language for stable homotopy.

The first definition is an n -treefold: a space locally homeomorphic to

$$T \times \mathbb{R}^{n-1},$$

where T is a tree. The branching models a controlled kind of singularity.

Given a symmetric monoidal category $(\mathcal{C}, \otimes)$, an n - $(\mathcal{C}, \otimes)$ -lasagna is, roughly, a framed n -treefold with labels from $\mathcal{C}$, satisfying compatibility at the branching loci.

Definition 152.14.1 (Intrinsic skein-lasagna, informal). An intrinsic lasagna is a geometric object built from treefold-like pieces, with labels coming from a symmetric monoidal category. The word “intrinsic” means that the object is studied on its own, rather than as something embedded in a previously chosen ambient manifold.

The slide illustrates n -treefolds as spaces locally modelled on $T \times \mathbb{R}^{n-1}$, where T is a tree. For the notebook, the useful point is that branching is allowed but controlled.

The word “intrinsic” is important. Classical skein-lasagna modules often involve embedded objects inside an ambient manifold. Here the lecture emphasizes that one can study the geometry of the lasagna itself, without first choosing an ambient space. If desired, one may imagine the ambient space as $\mathbb{R}^\infty$, but it is no longer part of the primary data.

§152.15 The $(\mathcal{V}_k, \oplus)$ -cobordism group

The lecture specializes the lasagna story to

$$(\mathcal{V}_k, \oplus),$$

where $\mathcal{V}_k$ is the category or moduli object of finite-dimensional vector spaces under direct sum. One obtains groups denoted

$$\Omega_n^{\mathcal{V}_k, \oplus}.$$

The examples in low degree match algebraic K -theory:

$$\Omega_0^{\mathcal{V}_k, \oplus} \cong K_0(k) \cong \mathbb{Z},$$

$$\Omega_1^{\mathcal{V}_k, \oplus} \cong K_1(k) \cong k^\times.$$

Remark 152.15.1 — The slide summarizes the first two $(\mathcal{V}_k, \oplus)$ -cobordism groups: degree 0 recovers dimension, and degree 1 recovers determinant. In formulas,

$$\Omega_0^{\mathcal{V}_k, \oplus} \cong K_0(k) \cong \mathbb{Z}, \quad \Omega_1^{\mathcal{V}_k, \oplus} \cong K_1(k) \cong k^\times.$$

This is why vector-space-labelled circles are the K_1 analogue of the finite-set swap loop.

The degree-one case is especially helpful. A circle labelled by an automorphism $f \in \mathrm{GL}(V)$ represents a class, and the determinant gives its value in $k^\times$. A commutator becomes null-cobordant, matching the abelianization of $\mathrm{GL}_\infty(k)$.

This is the lasagna analogue of the finite-set story. In finite sets, a swap gives the nontrivial element of $\mathbb{Z}/2$. In vector spaces, an automorphism gives an element of $k^\times$.

§152.16 Generalized Pontryagin–Thom

The slide deck then states the central claim:

$$K_n(k) \xrightarrow{\cong} \Omega_n^{\mathcal{V}_k, \oplus} \quad (n \geq 0).$$

This is presented as a generalized Pontryagin–Thom isomorphism.

Generalized Pontryagin–Thom Isomorphism

Recall that $K_0(k) = \mathbb{Z}$ and $K_1(k) = k^\times$. This is not just a coincidence:

Generalized Pontryagin–Thom Isomorphism (Claim)

For all $n \geq 0$, we have the isomorphism

$$K_n(k) \xrightarrow{\cong} \Omega_n^{\mathcal{V}_k, \oplus}$$

Figure 152.4: From the lecture slides: the generalized Pontryagin–Thom isomorphism identifies $K_n(k)$ with a lasagna-cobordism group.

The analogy is:

classical/stable homotopy	algebraic K -theory side
$\pi_n^s(\mathbb{S})$	$K_n(k)$
Ω_n^{fr}	$\Omega_n^{\mathcal{V}_k, \oplus}$
framed manifolds	vector-space-labelled lasagnas

This is why the lecture moves from elementary arithmetic to Hopf fibration and then to intrinsic skein-lasagna. It is building a geometric model of algebraic K -theory.

§152.17 The K_2 slide and Steinberg relations

In degree two, the lecture recalls the classical presentation

$$K_2(k) \cong k^\times \otimes k^\times / \langle a \otimes (1 - a) \rangle.$$

It then considers matrices $f, g \in \text{GL}(k)$ with commutator

$$[f, g] = 1.$$

Such a commuting pair defines a two-dimensional lasagna $C(f, g)$.

$\Omega_2^{\mathcal{V}_k, \oplus}$ and $K_2(k)$

For matrices $f, g \in \text{GL}(k)$ such that $[f, g] = 1$, one can define the following 2-lasagna $C(f, g)$:




Figure 152.5: From the lecture slides: a two-lasagna $C(f, g)$ associated to a commuting pair.

The next slide gives a more concrete example using elementary matrices. For $i \neq j$, let

$$e_{ij}(a)$$

be the elementary matrix with a in the (i, j) -entry and ones on the diagonal. These matrices satisfy Steinberg relations, such as

$$[e_{ij}(a), e_{kl}(b)] = \begin{cases} e_{il}(ab), & j = k, \\ \mathbf{1}, & \text{otherwise.} \end{cases}$$

The slide then uses Hall–Witt-type commutator identities to produce a nontrivial proof that a certain commutator is the identity. This proof is not just algebraic decoration; it is a blueprint for constructing a bounding three-lasagna.

3-Lasagna bounds $C(e_{12}(a), e_{12}(b))$

For $i \neq j$, let the elementary matrices $e_{ij}(a)$ be


$$(e_{ij}(a))_{kl} = \begin{cases} 1, & \text{if } k = l, \\ a, & \text{if } k = i, l = j, \\ 0, & \text{otherwise.} \end{cases}$$

Figure 152.6: From the lecture slides: elementary matrices, Steinberg relations, and the algebraic pattern behind the three-lasagna.

§152.18 The final lasagna diagrams

The final nontrivial slides display a Heegaard splitting and a boundary relation. They are not necessary for understanding the original meme, but they are essential for understanding where the lecture is going. The creation–swap–annihilation loop was the baby example. The lasagna diagrams are the research-level continuation of the same philosophy: algebraic K -theory classes are represented by geometric/cobordism-like objects.

Following this blueprint, we obtain the following Heegaard splitting to construct a lasagna:




25/27



Figure 152.7: From the lecture slides: a Heegaard-splitting picture used to construct a three-lasagna.

At this point the lecture has travelled very far from the original equation. But the path is coherent:



§152.19 What the equation finally says

The equation

$$0 = 1 - 1 = -1 + 1 = 0$$

should be read as the name of a loop:

$$h : \emptyset \rightarrow \{+, -\} \rightarrow \{-, +\} \rightarrow \emptyset.$$

It is a loop at zero in the K -space of finite sets. It is not the identity loop. It represents the generator of

$$K_1(\text{FinSet}) \cong \mathbb{Z}/2.$$




Figure 152.8: From the lecture slides: the boundary relation for the three-lasagna construction.

By Barratt–Priddy–Quillen, this group is

$$\pi_1^s(\mathbb{S}) \cong \mathbb{Z}/2,$$

and h corresponds to the stable Hopf class η .

This is why the Hopf fibration belongs in the same story as a signed finite-set calculation. The shared object is not an ordinary number. It is a stable homotopy class.

§152.20 Structural viewpoint

The two notes should now be read in chronological order rather than as if the first one already knew the second one's answer. The earlier note preserves the first route: a summer memory of the Hopf fibration, a lecture poster, a meme, and the attempt to explain the line through analytic summation, infinite cancellation, and the Eilenberg swindle. This second note records the later correction: the same vocabulary of K -theory and stability is present, but the mathematical emphasis moves from vanishing to a nontrivial finite loop.

The hierarchy is:

first reaction	later understanding
$0 = 0$	a loop from 0 to 0
$1 - 1 + 1 - 1 + \dots$	one created $+/-$ pair
Cesàro/Abel/eta	finite configuration space
Eilenberg swindle	group completion of finite sets
infinite absorption	creation–swap–annihilation movie
vanishing of invariants	$K_1(\text{FinSet}) \cong \mathbb{Z}/2$
triviality of the infinite	nontriviality of the finite
arithmetic/regularization instinct	stable Hopf class η

This distinction is the real lesson of the reflection. The first idea was not stupid; it was a reasonable search direction because K -theory really does interact with group completion, stability, and swindle arguments. But it placed the weight on the wrong side. It tried to make the equation true by allowing infinitely flexible cancellation. The lecture instead makes the equation interesting by refusing to collapse the finite movie to its endpoint.

The advanced tail of the lecture should also be remembered. The same style of thought continues beyond finite sets. For a field k , algebraic K -groups can be approached through vector-space-labelled lasagnas:

$$K_n(k) \cong \Omega_n^{\mathcal{V}_{k,\oplus}}.$$

The final slides are not random extra material. They show that the playful equation is the doorway into a geometric model of algebraic K -theory.

So the final lesson is not that 0 is secretly nonzero, and it is not that infinity makes all cancellation meaningless. The lesson is sharper: in higher mathematics, the space of ways for 0 to be 0 can contain a nontrivial loop, and in this example that loop is already visible in the smallest possible finite signed configuration.

XV

Machine Learning

Part XV: Contents

N-20250926	1003
153.1 The setting	1003
153.2 The theoretical tensor	1003
153.3 Computing the sparse derivative	1004
153.4 Contraction by the chain rule	1004
153.5 Why the tensor disappears	1005
153.6 Contravariant reading of the same calculation	1005
153.7 Outer-product expression	1006
153.8 Numerical example	1006
153.9 Einstein notation proof	1006
153.10 Kronecker product viewpoint	1007
153.11 Softmax-cross-entropy simplification	1007
153.12 A wider interpretation	1007
153.13 Backpropagation as cotangent pullback	1007
153.14 Why the three-dimensional tensor collapses	1007
153.15 Batch form	1008
153.16 Backpropagation as sparse tensor contraction	1008
 N-20251008	 1009
154.1 Why probability enters generative modeling	1009
154.2 Random variables, distributions, and densities	1009
154.3 Joint, marginal, and conditional distributions	1010
154.4 Bayes' rule as inversion of a generative story	1010
154.5 Gaussian distributions	1010
154.6 Expectation and Monte Carlo approximation	1011
154.7 KL divergence	1011
154.8 Jensen's inequality	1012
154.9 The next abstraction	1012
 N-20251009	 1013
155.1 From autoencoders to latent variable models	1013
155.2 The generative story	1013
155.3 Graphical model notation	1014
155.4 Marginal likelihood	1014
155.5 The posterior problem	1015
155.6 Why naive Monte Carlo is not enough	1015
155.7 The recognition model	1015
155.8 Amortized inference	1016
155.9 The VAE model diagram	1016
155.10 Remaining derivation	1016
 N-20251010	 1017
156.1 The approximate posterior as an optimization problem	1017
156.2 The ELBO identity	1017
156.3 Jensen's inequality derivation	1018
156.4 Reconstruction term and regularization term	1018
156.5 The ELBO gap	1019
156.6 Closed-form KL for diagonal Gaussians	1019
156.7 Negative ELBO as a loss	1020
156.8 Empirical lower bounds	1020
156.9 Information-theoretic reading	1021
156.10 Limits of the ELBO	1021
156.11 What survives later	1021
 N-20251011	 1023
157.1 The gradient problem	1023
157.2 Score-function estimator and pathwise estimator	1023

157.3	Gaussian reparameterization	1024
157.4	The stochastic ELBO estimator	1024
157.5	Why implementations output log variance	1024
157.6	Training as ELBO maximization	1025
157.7	Decoder likelihoods and reconstruction losses	1025
157.8	Generation after training	1025
157.9	Empirical marginal likelihood curves	1026
157.10	Limitations and extensions	1026
157.11	How the pieces fit	1027
 N-20260302		1029
158.1	Backpropagation as cotangent pullback	1029
158.2	Geometric data analysis	1030
158.3	Manifold hypothesis and embeddings	1030
158.4	Optimal transport and stochastic dynamics	1031
158.5	Symmetry, equivariance, and category-like thinking	1032
158.6	Further direction	1032
158.7	Backpropagation and differential geometry	1033
158.8	Manifold hypothesis and representation learning	1033
158.9	Symmetry and equivariance	1033
158.10	Optimal transport and distributions	1033
158.11	Mathematical structures for learning systems	1033
158.12	Rebuilt reading path: what pure mathematics contributes	1033
158.13	Topological data analysis in usable form	1034
158.14	Wasserstein geometry and stochastic dynamics	1034
158.15	Category language without exaggeration	1034

Backpropagation in One MLP Layer: From a Three-Dimensional Tensor to an Outer Product

N-20250926

§153.1 The setting

Consider one middle layer of a multilayer perceptron:

$$z = Wx + b, \quad y = \sigma(z).$$

Here

$$x \in \mathbb{R}^n, \quad W \in \mathbb{R}^{m \times n}, \quad b \in \mathbb{R}^m, \quad z, y \in \mathbb{R}^m.$$

The activation function σ acts elementwise, for example sigmoid or ReLU. The goal of backpropagation is to compute

$$\frac{\partial L}{\partial W}$$

for a loss function L .

The note uses the common softmax-cross-entropy setting as motivation. If the final layer produces logits z_i , then

$$\hat{y}_i = \frac{e^{z_i}}{\sum_j e^{z_j}}, \quad L = - \sum_i t_i \ln \hat{y}_i,$$

where t_i is a one-hot label. Backpropagation asks how this scalar loss changes when each weight $W_{j,k}$ changes.

§153.2 The theoretical tensor

Strictly speaking, the derivative of the vector $y \in \mathbb{R}^m$ with respect to the matrix $W \in \mathbb{R}^{m \times n}$ has three indices:

$$\frac{\partial y_i}{\partial W_{j,k}}.$$

The index i chooses the output component, j chooses the row of W , and k chooses the column of W . So the object has shape

$$m \times m \times n.$$

This is the source of the confusion recorded in the note. The derivative is theoretically a third-order tensor, but in neural-network code one usually sees only a matrix gradient. The matrix appears after chain-rule contraction with the upstream gradient.

§153.3 Computing the sparse derivative

Write the linear preactivation component as

$$z_r = \sum_{t=1}^n W_{r,t} x_t + b_r.$$

Then

$$\frac{\partial z_r}{\partial W_{j,k}} = \begin{cases} x_k, & r = j, \\ 0, & r \neq j, \end{cases} = \delta_{rj} x_k.$$

This is sparse because z_r depends only on row r of W .

For $y_i = \sigma(z_i)$ with elementwise activation,

$$\frac{\partial y_i}{\partial z_r} = 0 \quad (i \neq r), \quad \frac{\partial y_i}{\partial z_i} = \sigma'(z_i).$$

Thus

$$\frac{\partial y_i}{\partial W_{j,k}} = \sum_{r=1}^m \frac{\partial y_i}{\partial z_r} \frac{\partial z_r}{\partial W_{j,k}} = \frac{\partial y_i}{\partial z_j} x_k.$$

For elementwise activation this is nonzero only when $i = j$.

§153.4 Contraction by the chain rule

The loss L is scalar. Therefore

$$\frac{\partial L}{\partial W_{j,k}} = \sum_{i=1}^m \frac{\partial L}{\partial y_i} \frac{\partial y_i}{\partial W_{j,k}}.$$

Substituting the previous formula gives

$$\frac{\partial L}{\partial W_{j,k}} = \sum_{i=1}^m \frac{\partial L}{\partial y_i} \frac{\partial y_i}{\partial z_j} x_k.$$

The sum over i is exactly the j -th component of the upstream gradient with respect to z :

$$\delta_j = \frac{\partial L}{\partial z_j}.$$

Hence

$$\frac{\partial L}{\partial W_{j,k}} = \delta_j x_k.$$

Stacking all j, k gives the outer product

$$\boxed{\frac{\partial L}{\partial W} = \delta x^T}$$

where

$$\delta = \frac{\partial L}{\partial z}.$$

§153.5 Why the tensor disappears

The original derivative $\partial y / \partial W$ is a tensor with shape (m, m, n) . But backpropagation does not need to store it. The chain rule contracts it with the vector $\partial L / \partial y$, leaving a matrix of shape (m, n) .

In the note's language:

$$\text{three-dimensional tensor} \xrightarrow{\text{chain-rule summation}} \text{matrix gradient.}$$

This is not a loss of information for gradient descent, because gradient descent only needs the derivative of the scalar loss with respect to each parameter.

§153.6 Contravariant reading of the same calculation

The same computation can be described as a pullback of covectors. A forward layer is a map

$$f : \text{parameters and activations} \longrightarrow \text{outputs.}$$

Its differential sends small forward perturbations to small output perturbations:

$$df : T_x X \rightarrow T_{f(x)} Y.$$

But the loss supplies a covector at the output,

$$dL \in T_{f(x)}^* Y,$$

and backpropagation sends it backward by precomposition:

$$f^*(dL) = dL \circ df \in T_x^* X.$$

This is why reverse-mode differentiation is naturally contravariant.

For a batched linear layer in PyTorch convention,

$$X \in \mathbb{R}^{b \times n}, \quad W \in \mathbb{R}^{m \times n}, \quad Y = XW^T \in \mathbb{R}^{b \times m},$$

the component formula is

$$Y_{\alpha i} = \sum_{k=1}^n X_{\alpha k} W_{ik}.$$

Hence

$$\frac{\partial Y_{\alpha i}}{\partial W_{jk}} = \delta_{ij} X_{\alpha k},$$

a Jacobian tensor of shape (b, m, m, n) . Contracting with the upstream gradient $G = \partial L / \partial Y$, of shape (b, m) , gives

$$\begin{aligned} \frac{\partial L}{\partial W_{jk}} &= \sum_{\alpha, i} G_{\alpha i} \frac{\partial Y_{\alpha i}}{\partial W_{jk}} \\ &= \sum_{\alpha} G_{\alpha j} X_{\alpha k}, \end{aligned}$$

so

$$\frac{\partial L}{\partial W} = G^T X.$$

If the weights are stored instead as $W_{\text{col}} \in \mathbb{R}^{n \times m}$ and $Y = XW_{\text{col}}$, the same formula appears as

$$\frac{\partial L}{\partial W_{\text{col}}} = X^T G.$$

The two expressions differ only by the convention for storing the weight matrix.

§153.7 Outer-product expression

Let

$$\frac{\partial y}{\partial z}$$

be treated as a column vector for elementwise activation, with entries $\sigma'(z_i)$. Then the derivative of y with respect to W can be represented row by row as

$$\frac{\partial y}{\partial W} \sim \begin{bmatrix} \frac{\partial y_1}{\partial z_1} x^T \\ \frac{\partial y_2}{\partial z_2} x^T \\ \vdots \\ \frac{\partial y_m}{\partial z_m} x^T \end{bmatrix} = \left(\frac{\partial y}{\partial z} \right) x^T.$$

For the scalar loss, replace $\partial y / \partial z$ by the backpropagated vector $\delta = \partial L / \partial z$.

§153.8 Numerical example

The note gives a simple example with

$$m = 2, \quad n = 3, \quad x = [1, 2, 3]^T, \quad z = [0.1, -1.0]^T.$$

For sigmoid,

$$\sigma'(0.1) \approx 0.2494, \quad \sigma'(-1.0) \approx 0.1966.$$

Therefore

$$\frac{\partial y}{\partial W} \sim \begin{bmatrix} 0.2494 \\ 0.1966 \end{bmatrix} \begin{bmatrix} 1 & 2 & 3 \end{bmatrix} = \begin{bmatrix} 0.2494 & 0.4988 & 0.7482 \\ 0.1966 & 0.3932 & 0.5898 \end{bmatrix}.$$

Each row is the input vector scaled by the derivative of the corresponding output.

§153.9 Einstein notation proof

The handwritten proof uses indices. Let $y = f(z)$ and $z = Wx$. For a matrix entry $A_{pq} = W_{pq}$,

$$\frac{\partial z_i}{\partial A_{pq}} = \frac{\partial}{\partial A_{pq}} \sum_j A_{ij} x_j = \delta_{ip} x_q.$$

Then

$$\frac{\partial L}{\partial A_{pq}} = \sum_i \frac{\partial L}{\partial y_i} \frac{\partial y_i}{\partial z_p} x_q.$$

For elementwise activation, this becomes

$$\frac{\partial L}{\partial A_{pq}} = \frac{\partial L}{\partial z_p} x_q.$$

This is the component form of $\partial L / \partial W = \delta x^T$.

§153.10 Kronecker product viewpoint

The note also explores vectorization. If $y = f(Wx)$ and we vectorize W , then the Jacobian with respect to $\text{vec}(W)$ can be written using Kronecker products. The important identity is that changing one matrix entry affects only one row of z , which is encoded by a Kronecker delta.

In batch form, if X is a matrix of inputs and Δ is the matrix of upstream gradients, then the gradient becomes

$$\frac{\partial L}{\partial W} = \Delta X^T.$$

This is the same outer-product rule, summed over the batch dimension.

§153.11 Softmax–cross-entropy simplification

For the final layer with softmax and cross entropy, the derivative simplifies to

$$\frac{\partial L}{\partial z} = \hat{y} - t.$$

Therefore the final-layer weight gradient is

$$\frac{\partial L}{\partial W} = (\hat{y} - t)x^T.$$

This is why implementation code often looks much simpler than the theoretical tensor derivation.

§153.12 A wider interpretation

The key lesson is that matrix calculus is compressed tensor calculus. The full derivative of a vector with respect to a matrix is a higher-order object, but backpropagation contracts it immediately with an upstream gradient. The result is an outer product: error signal times input. This is the local rule behind every dense neural-network layer.

§153.13 Backpropagation as cotangent pullback

For a layer $z = Wx + b$, the forward map sends activations forward, while backpropagation sends covectors backward:

$$\bar{x} = W^T \bar{z}.$$

This is the pullback on cotangent vectors.

§153.14 Why the three-dimensional tensor collapses

The derivative of $z_i = \sum_j W_{ij}x_j + b_i$ with respect to W_{ab} is

$$\frac{\partial z_i}{\partial W_{ab}} = \delta_{ia}x_b.$$

Contracting with upstream gradient $\bar{z}_i$ gives

$$\frac{\partial L}{\partial W_{ab}} = \bar{z}_a x_b.$$

Thus the huge derivative tensor disappears after contraction, leaving an outer product.

§153.15 Batch form

For a batch, gradients add over examples:

$$\nabla_W L = \bar{Z} X^T.$$

This is matrix multiplication as accumulated outer products.

§153.16 Backpropagation as sparse tensor contraction

Backpropagation is not magic tensor cancellation. It is cotangent pullback plus contraction of sparse Jacobians, and the familiar weight gradient is the outer product of upstream error with input activation.

Probability Language for Variational Autoencoders

N-20251008

§154.1 Why probability enters generative modeling

A deterministic autoencoder learns a pair of functions

$$x \longmapsto h \longmapsto \hat{x}.$$

This is useful for compression and representation learning, but it is not yet a probabilistic generative model. A generative model should tell us how likely different data points are and how to sample new ones. Therefore the mathematical object is not merely a map, but a probability distribution.

The VAE begins from the following change of viewpoint:

a data point is not only reconstructed;
it is regarded as a sample from an unknown distribution.

If X denotes the random data point, then the ideal object is the data distribution $p_{\text{data}}(x)$. A model tries to build a tractable family $p_{\theta}(x)$ and choose θ so that p_{θ} resembles the data distribution.

The slogan for this chapter is:

generation means sampling from a learned distribution.

Before discussing encoders, decoders, or neural networks, we first need the probability language in which the VAE objective is written.

§154.2 Random variables, distributions, and densities

A random variable X is a measurable quantity whose value is uncertain. Its distribution tells us the probability of events such as $X \in A$. When X has a density $p(x)$ with respect to Lebesgue measure, probabilities are computed by integration:

$$\mathbb{P}(X \in A) = \int_A p(x) dx.$$

A density is not itself a probability. The value $p(x)$ may even be larger than 1. What matters is the integral over a region.

The expectation of a function $f(X)$ is

$$\mathbb{E}[f(X)] = \int f(x)p(x) dx.$$

For a discrete variable, the corresponding formula is a sum:

$$\mathbb{E}[f(X)] = \sum_x f(x)p(x).$$

The formulas look similar, but the measure matters. In VAEs, latent variables are usually continuous, so integrals over z appear everywhere.

§154.3 Joint, marginal, and conditional distributions

VAEs use latent variables, so we need distributions involving both visible variables and hidden variables. Let X be observed data and Z be a latent variable. Their joint density is

$$p(x, z).$$

The joint density can be factored in two ways:

$$p(x, z) = p(x | z)p(z) = p(z | x)p(x).$$

The marginal density of X is obtained by integrating out Z :

$$p(x) = \int p(x, z) dz.$$

This operation is called marginalization. It means that all possible hidden explanations z are summed or integrated away.

The conditional density is given by Bayes' rule:

$$p(z | x) = \frac{p(x, z)}{p(x)} = \frac{p(x | z)p(z)}{\int p(x | z')p(z') dz'}.$$

The three words should be kept distinct:

joint	describe x and z together
marginal	hide one variable by integration
conditional	update one variable after observing the other

§154.4 Bayes' rule as inversion of a generative story

A latent variable model usually starts with a generative story:

$$z \sim p(z), \quad x \sim p_\theta(x | z).$$

This is the easy direction: first choose a hidden cause z , then generate an observation x .

Inference goes in the opposite direction. Given an observation x , we ask which latent values z could have produced it. The posterior is

$$p_\theta(z | x) = \frac{p_\theta(x | z)p(z)}{p_\theta(x)}.$$

The denominator

$$p_\theta(x) = \int p_\theta(x | z)p(z) dz$$

is called the evidence or marginal likelihood.

This is the first obstacle behind VAEs. The generative direction is easy to write, but the inverse direction requires an integral that is usually not tractable.

§154.5 Gaussian distributions

The standard VAE uses Gaussian latent variables. The one-dimensional Gaussian density is

$$\mathcal{N}(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right).$$

Here μ is the mean and σ^2 is the variance.

The multivariate Gaussian density is

$$\mathcal{N}(z; \mu, \Sigma) = \frac{1}{(2\pi)^{d/2} |\Sigma|^{1/2}} \exp \left(-\frac{1}{2} (z - \mu)^T \Sigma^{-1} (z - \mu) \right).$$

The matrix Σ is the covariance matrix. It measures not only the variance of each coordinate, but also how coordinates co-vary.

The diagonal case is especially important:

$$\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_d^2).$$

Then the coordinates are independent under the Gaussian, and the density factorizes:

$$\mathcal{N}(z; \mu, \text{diag}(\sigma^2)) = \prod_{j=1}^d \mathcal{N}(z_j; \mu_j, \sigma_j^2).$$

This is why VAE encoders often output a vector of means and a vector of variances rather than a full covariance matrix.

§154.6 Expectation and Monte Carlo approximation

Many VAE formulas contain expectations that cannot be computed exactly. If

$$z \sim q(z),$$

then

$$\mathbb{E}_{z \sim q}[f(z)] = \int f(z) q(z) dz.$$

A Monte Carlo approximation uses samples $z^{(1)}, \dots, z^{(L)} \sim q$:

$$\mathbb{E}_{z \sim q}[f(z)] \approx \frac{1}{L} \sum_{\ell=1}^L f(z^{(\ell)}).$$

This is not a deterministic quadrature rule. It is a random approximation whose error decreases statistically. The key benefit is that the dimension of z can be large: the same formula still makes sense, while grid integration becomes impossible.

§154.7 KL divergence

The Kullback–Leibler divergence compares two distributions q and p :

$$D_{\text{KL}}(q \| p) = \mathbb{E}_q \left[\log \frac{q(Z)}{p(Z)} \right].$$

For continuous densities, this means

$$D_{\text{KL}}(q \| p) = \int q(z) \log \frac{q(z)}{p(z)} dz.$$

For discrete distributions, replace the integral by a sum.

The order matters. Usually

$$D_{\text{KL}}(q \| p) \neq D_{\text{KL}}(p \| q).$$

So KL divergence is not a metric distance.

A useful interpretation is coding loss. If data are really drawn from q , but one uses a code optimized for p , then KL measures the expected extra log-loss. In variational inference, we use it to measure how expensive it is to replace the true posterior by an approximate posterior.

The fundamental inequality is

$$D_{\text{KL}}(q\|p) \geq 0,$$

with equality exactly when $q = p$ almost everywhere. This nonnegativity is what turns the ELBO into a lower bound.

§154.8 Jensen's inequality

The other basic ingredient is Jensen's inequality. Since log is concave,

$$\log \mathbb{E}[Y] \geq \mathbb{E}[\log Y]$$

for positive random variables Y .

This inequality is the analytic reason variational lower bounds exist. In the VAE derivation, we will write a difficult likelihood as

$$\log \int q(z) \frac{p(x, z)}{q(z)} dz = \log \mathbb{E}_q \left[\frac{p(x, Z)}{q(Z)} \right]$$

and then move the logarithm inside the expectation:

$$\log \mathbb{E}_q \left[\frac{p(x, Z)}{q(Z)} \right] \geq \mathbb{E}_q \left[\log \frac{p(x, Z)}{q(Z)} \right].$$

This is exactly the shape of the ELBO.

§154.9 The next abstraction

The VAE is built from probability, not merely from neural network architecture. The important vocabulary is:

term	role in VAE
$p(z)$	prior over latent variables
$p_\theta(x z)$	decoder likelihood
$p_\theta(x) = \int p_\theta(x z)p(z) dz$	marginal likelihood
$p_\theta(z x)$	true posterior
$q_\phi(z x)$	approximate posterior / encoder
$D_{\text{KL}}(q\ p)$	penalty for replacing p by q
$\mathbb{E}_q[f]$	average under the approximate posterior

The next note will use this language to describe a latent variable model. The main tension will be simple: the model is easy to sample from, but hard to invert.

Latent Variable Models: Priors, Decoders, and Intractable Posteriors

N-20251009

§155.1 From autoencoders to latent variable models

A deterministic autoencoder is a pair of maps

$$x \xrightarrow{\text{encoder}} h \xrightarrow{\text{decoder}} \hat{x}.$$

The code h is a deterministic representation of x . A VAE keeps the encoder–decoder visual shape, but the mathematical object is different. The hidden representation is a random variable Z , and the decoder defines a conditional distribution, not merely a reconstruction function.

The generative story is

$$z \sim p(z), \quad x \sim p_\theta(x | z).$$

Thus a VAE is best understood as a latent variable model equipped with an efficient approximate inference model.

The key distinction is:

deterministic autoencoder	VAE / latent variable model
$h = f_\phi(x)$	$z \sim q_\phi(z x)$
$\hat{x} = g_\theta(h)$	$x \sim p_\theta(x z)$
reconstruction map	probabilistic generative model

§155.2 The generative story

The model begins with a prior distribution on the latent variable:

$$p(z).$$

In the basic VAE,

$$p(z) = \mathcal{N}(0, I).$$

This prior says that latent variables should live in a simple, continuous, isotropic space.

Given z , the decoder defines the likelihood of an observation:

$$p_\theta(x | z).$$

The joint density is then

$$p_\theta(x, z) = p(z)p_\theta(x | z).$$

This formula is the mathematical core of the generative model.

The decoder network does not directly output a final sample. It outputs parameters of a distribution. For binary data, one often writes

$$p_\theta(x | z) = \text{Bernoulli}(x; \pi_\theta(z)).$$

Here $\pi_\theta(z)$ gives the success probabilities of the Bernoulli coordinates. For real-valued data, one may use

$$p_\theta(x | z) = \mathcal{N}(x; \mu_\theta(z), \sigma^2 I).$$


Here the decoder outputs a mean $\mu_\theta(z)$, and the variance is either fixed or parameterized separately.

§155.3 Graphical model notation

The generative model has one directed probabilistic arrow:

$$z \longrightarrow x.$$

A small diagram is:



For a dataset $x^{(1)}, \dots, x^{(N)}$, the story repeats independently:

$$z^{(i)} \sim p(z), \quad x^{(i)} \sim p_\theta(x | z^{(i)}).$$

The parameters θ are shared across all observations.

§155.4 Marginal likelihood

The model likelihood of one observation is obtained by integrating out the latent variable:

$$p_\theta(x) = \int p_\theta(x, z) dz = \int p(z) p_\theta(x | z) dz.$$

For a dataset, maximum likelihood would choose θ by maximizing

$$\sum_{i=1}^N \log p_\theta(x^{(i)}).$$

Equivalently, one wants the model distribution $p_\theta(x)$ to assign high probability density to the observed data.

The problem is that the integral

$$\int p(z) p_\theta(x | z) dz$$

is usually intractable. If $p_\theta(x | z)$ is parameterized by a nonlinear neural network, there is no closed-form expression for the integral over z .

This is the first reason variational methods enter.

§155.5 The posterior problem

The posterior distribution of the latent variable is

$$p_{\theta}(z | x) = \frac{p_{\theta}(x, z)}{p_{\theta}(x)} = \frac{p(z)p_{\theta}(x | z)}{\int p(z')p_{\theta}(x | z') dz'}.$$

The numerator is easy to evaluate up to normalization. The denominator is the marginal likelihood, which is the same difficult integral.

Therefore the model has a basic asymmetry:

sampling direction	inference direction
$z \sim p(z), x \sim p_{\theta}(x z)$	$p_{\theta}(z x)$ is hard

This is the central probabilistic problem behind VAEs:

to train the model, we want $\log p_{\theta}(x)$,
 to understand $\log p_{\theta}(x)$, we need to reason about z ,
 but the true posterior $p_{\theta}(z | x)$ is intractable.

§155.6 Why naive Monte Carlo is not enough

One might try to estimate the marginal likelihood by drawing samples from the prior:

$$z^{(\ell)} \sim p(z), \quad p_{\theta}(x) \approx \frac{1}{L} \sum_{\ell=1}^L p_{\theta}(x | z^{(\ell)}).$$

This is formally reasonable, but in high dimension it is usually inefficient. Most prior samples z do not explain a fixed observation x well. The average is then dominated by rare samples that land in useful regions of latent space.

For inference about a particular x , we would rather sample z values that are already plausible explanations of x . That is the motivation for introducing an approximate posterior.

§155.7 The recognition model

A VAE introduces a family

$$q_{\phi}(z | x)$$

to approximate the true posterior $p_{\theta}(z | x)$. This is often called the recognition model or encoder.

The standard Gaussian encoder has the form

$$q_{\phi}(z | x) = \mathcal{N}(z; \mu_{\phi}(x), \text{diag}(\sigma_{\phi}^2(x))).$$

The functions $\mu_{\phi}(x)$ and $\sigma_{\phi}(x)$ are outputs of a neural network. Mathematically, they parameterize a distribution over latent variables.

It is important not to confuse the two conditional distributions:

symbol	name	direction
$p_{\theta}(x z)$	decoder likelihood	$z \rightarrow x$
$q_{\phi}(z x)$	approximate posterior / encoder	$x \rightarrow z$

Only $p_{\theta}(x | z)$ belongs to the generative model. The encoder $q_{\phi}(z | x)$ is introduced to make inference and training feasible.

§155.8 Amortized inference

Classical variational inference might introduce separate variational parameters for each data point. For example, for every observation $x^{(i)}$, one could choose its own mean and variance:

$$q_i(z) = \mathcal{N}(z; \mu_i, \text{diag}(\sigma_i^2)).$$

A VAE instead uses one shared inference network:


$$x \mapsto \mu_\phi(x), \quad x \mapsto \sigma_\phi(x).$$

This is called amortized inference. The cost of inference is amortized across the dataset: after learning ϕ , the model can quickly produce an approximate posterior for a new x .

The mathematical tradeoff is that the family is restricted. We do not choose arbitrary variational parameters for each observation; we require them to be outputs of the same function.

§155.9 The VAE model diagram

The two directions can be displayed together:



The upper arrow is the probabilistic generative story. The lower arrow is the learned inference shortcut. The VAE objective will force these two directions to be compatible, but they play different roles.

§155.10 Remaining derivation

At this point, the problem is clear. We want to maximize

$$\log p_\theta(x),$$

but

$$p_\theta(x) = \int p(z) p_\theta(x | z) dz$$

is hard. We also want to use

$$p_\theta(z | x),$$

but it contains the same intractable marginal likelihood.

The solution is to introduce $q_\phi(z | x)$, then build a lower bound on $\log p_\theta(x)$ that can be optimized. That lower bound is the evidence lower bound, or ELBO.

The next note derives it from two facts:

$$D_{\text{KL}}(q||p) \geq 0, \quad \log \mathbb{E}[Y] \geq \mathbb{E}[\log Y].$$

Variational Inference: KL Divergence, Jensen, and the ELBO

N-20251010

§156.1 The approximate posterior as an optimization problem

The true posterior in a latent variable model is

$$p_{\theta}(z | x) = \frac{p_{\theta}(x, z)}{p_{\theta}(x)}.$$

It is usually intractable because the marginal likelihood

$$p_{\theta}(x) = \int p_{\theta}(x, z) dz$$

is hard to compute.

Variational inference introduces a tractable family

$$q_{\phi}(z | x)$$

and tries to make it close to the true posterior. A natural measure of closeness is

$$D_{\text{KL}}(q_{\phi}(z | x) || p_{\theta}(z | x)).$$

This is the reverse direction from what one might first guess: we place expectation under q_{ϕ} , because q_{ϕ} is the distribution from which we can sample and whose density we can evaluate.

The obstacle is that this KL divergence contains $p_{\theta}(z | x)$, which itself contains the intractable $p_{\theta}(x)$. The central algebraic trick is to rewrite it so that $\log p_{\theta}(x)$ appears as a constant plus a lower bound.

§156.2 The ELBO identity

Start from the KL divergence:

$$D_{\text{KL}}(q_{\phi}(z | x) || p_{\theta}(z | x)) = \mathbb{E}_{q_{\phi}(z|x)} \left[\log \frac{q_{\phi}(z | x)}{p_{\theta}(z | x)} \right].$$

Using

$$p_{\theta}(z | x) = \frac{p_{\theta}(x, z)}{p_{\theta}(x)},$$

we get

$$\begin{aligned} D_{\text{KL}}(q_{\phi}(z | x) || p_{\theta}(z | x)) &= \mathbb{E}_{q_{\phi}} [\log q_{\phi}(z | x) - \log p_{\theta}(x, z) + \log p_{\theta}(x)] \\ &= \log p_{\theta}(x) - \mathbb{E}_{q_{\phi}} [\log p_{\theta}(x, z) - \log q_{\phi}(z | x)]. \end{aligned}$$

Define

$$\mathcal{L}(\theta, \phi; x) = \mathbb{E}_{q_{\phi}(z|x)} [\log p_{\theta}(x, z) - \log q_{\phi}(z | x)].$$

Then the identity becomes

$$\log p_\theta(x) = \mathcal{L}(\theta, \phi; x) + D_{\text{KL}}(q_\phi(z | x) \| p_\theta(z | x)).$$

Since KL divergence is nonnegative,

$$\mathcal{L}(\theta, \phi; x) \leq \log p_\theta(x).$$

This is why $\mathcal{L}$ is called the evidence lower bound, or ELBO.

Remark 156.2.1 — The ELBO is not introduced by magic. It is exactly what remains when the intractable posterior KL is separated from the log marginal likelihood.

§156.3 Jensen's inequality derivation

The same lower bound can be derived directly from Jensen's inequality. Insert any density $q_\phi(z | x)$ whose support is compatible with $p_\theta(x, z)$:

$$\begin{aligned} \log p_\theta(x) &= \log \int p_\theta(x, z) dz \\ &= \log \int q_\phi(z | x) \frac{p_\theta(x, z)}{q_\phi(z | x)} dz \\ &= \log \mathbb{E}_{q_\phi(z|x)} \left[\frac{p_\theta(x, z)}{q_\phi(z | x)} \right]. \end{aligned}$$

Because log is concave,

$$\begin{aligned} \log p_\theta(x) &\geq \mathbb{E}_{q_\phi(z|x)} \left[\log \frac{p_\theta(x, z)}{q_\phi(z | x)} \right] \\ &= \mathbb{E}_{q_\phi(z|x)} [\log p_\theta(x, z) - \log q_\phi(z | x)] \\ &= \mathcal{L}(\theta, \phi; x). \end{aligned}$$

The two derivations emphasize different things. The KL derivation shows the gap between the bound and the true log-likelihood. The Jensen derivation shows why a lower bound exists at all.

§156.4 Reconstruction term and regularization term

Using the latent variable factorization

$$p_\theta(x, z) = p(z)p_\theta(x | z),$$

the ELBO becomes

$$\begin{aligned} \mathcal{L}(\theta, \phi; x) &= \mathbb{E}_{q_\phi(z|x)} [\log p_\theta(x | z) + \log p(z) - \log q_\phi(z | x)] \\ &= \mathbb{E}_{q_\phi(z|x)} [\log p_\theta(x | z)] - D_{\text{KL}}(q_\phi(z | x) \| p(z)). \end{aligned}$$

This is the familiar VAE decomposition:

$$\mathcal{L} = \text{expected log-likelihood} - \text{KL to the prior}.$$

The first term asks whether latent samples from $q_\phi(z | x)$ explain x well under the decoder. The second term asks whether the approximate posterior stays close to the prior.

It is common to call the first term the reconstruction term. But mathematically it is an expected log-likelihood. It becomes a squared-error reconstruction loss only under a Gaussian decoder with fixed variance. It becomes a cross-entropy-like loss under a Bernoulli decoder.

§156.5 The ELBO gap

The identity

$$\log p_\theta(x) = \mathcal{L}(\theta, \phi; x) + D_{\text{KL}}(q_\phi(z | x) \| p_\theta(z | x))$$

shows that there are two goals mixed together.

Increasing $\mathcal{L}$ can improve the model likelihood $\log p_\theta(x)$. But it also tries to reduce the posterior approximation gap

$$D_{\text{KL}}(q_\phi(z | x) \| p_\theta(z | x)).$$

If q_ϕ were flexible enough to equal the true posterior, then the bound would be tight:

$$\mathcal{L}(\theta, \phi; x) = \log p_\theta(x).$$

In practice, q_ϕ is restricted, often to a diagonal Gaussian, so the bound is generally not exact.

§156.6 Closed-form KL for diagonal Gaussians

The standard VAE uses

$$q_\phi(z | x) = \mathcal{N}(z; \mu, \text{diag}(\sigma^2)), \quad p(z) = \mathcal{N}(0, I).$$

For this pair, the KL divergence has a closed form:

$$D_{\text{KL}}(q \| p) = \frac{1}{2} \sum_{j=1}^d (\mu_j^2 + \sigma_j^2 - 1 - \log \sigma_j^2).$$

Here is the one-dimensional calculation. Let

$$q(z) = \mathcal{N}(\mu, \sigma^2), \quad p(z) = \mathcal{N}(0, 1).$$

Then

$$\log q(z) = -\frac{1}{2} \log(2\pi\sigma^2) - \frac{(z - \mu)^2}{2\sigma^2},$$

and

$$\log p(z) = -\frac{1}{2} \log(2\pi) - \frac{z^2}{2}.$$

Therefore

$$\begin{aligned} D_{\text{KL}}(q \| p) &= \mathbb{E}_q[\log q(z) - \log p(z)] \\ &= -\frac{1}{2} \log \sigma^2 - \frac{1}{2} \mathbb{E}_q \left[\frac{(z - \mu)^2}{\sigma^2} \right] + \frac{1}{2} \mathbb{E}_q[z^2] \\ &= -\frac{1}{2} \log \sigma^2 - \frac{1}{2} + \frac{1}{2} (\mu^2 + \sigma^2) \\ &= \frac{1}{2} (\mu^2 + \sigma^2 - 1 - \log \sigma^2). \end{aligned}$$

The multivariate diagonal formula follows by summing over independent coordinates.

This is one reason the diagonal Gaussian encoder is practical: the regularization term does not need Monte Carlo estimation.

§156.7 Negative ELBO as a loss

Training usually minimizes the negative ELBO:

$$-\mathcal{L}(\theta, \phi; x) = -\mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x | z)] + D_{\text{KL}}(q_\phi(z | x) \| p(z)).$$

For a Bernoulli decoder with probabilities $\pi_\theta(z)$,

$$p_\theta(x | z) = \prod_i \pi_i(z)^{x_i} (1 - \pi_i(z))^{1-x_i},$$

so

$$-\log p_\theta(x | z) = -\sum_i [x_i \log \pi_i(z) + (1 - x_i) \log(1 - \pi_i(z))].$$

For a Gaussian decoder

$$p_\theta(x | z) = \mathcal{N}(x; \mu_\theta(z), \sigma^2 I),$$

we have

$$-\log p_\theta(x | z) = \frac{1}{2\sigma^2} \|x - \mu_\theta(z)\|^2 + \text{constant}.$$

Thus the common reconstruction losses are not arbitrary engineering choices. They are negative log-likelihoods under different decoder distributions.

§156.8 Empirical lower bounds

The original AEVB experiments plotted the variational lower bound during training. The figure below is useful as a historical reminder: what is being optimized is not raw reconstruction error, but a lower bound $\mathcal{L}$ on marginal likelihood.




Figure 156.1: Variational lower bound curves in the AEVB experiments.

For the purpose of these notes, the important point is mathematical rather than empirical: the plotted quantity is meaningful because of the identity

$$\log p_\theta(x) = \mathcal{L}(\theta, \phi; x) + D_{\text{KL}}(q_\phi(z | x) \| p_\theta(z | x)).$$

§156.9 Information-theoretic reading

The ELBO also resembles a rate–distortion tradeoff. The expected log-likelihood asks the latent variable to preserve enough information to explain x . The KL term penalizes deviations of $q_\phi(z | x)$ from the prior $p(z)$.

Loosely:

term	role
$-\mathbb{E}_q[\log p_\theta(x z)]$	distortion / reconstruction cost
$D_{\text{KL}}(q_\phi(z x) \ p(z))$	rate / information cost

This interpretation is helpful, but it should not replace the probabilistic derivation. The ELBO is first of all a lower bound on marginal likelihood.

§156.10 Limits of the ELBO

The ELBO is powerful, but it is not magic. Several limitations remain.

First, the variational family may be too small. A diagonal Gaussian $q_\phi(z | x)$ cannot represent a complicated multimodal posterior.

Second, the prior may be too simple. The standard normal prior is convenient, but the true aggregated latent structure of data may not be standard Gaussian.

Third, the decoder may be too powerful. If the decoder can model the data without using z , the approximate posterior may collapse toward the prior. This is posterior collapse.

Fourth, maximizing a lower bound is not the same as exactly maximizing $\log p_\theta(x)$. The gap is the posterior KL.

§156.11 What survives later

The VAE objective is not merely a heuristic sum of two losses. It is a variational identity:

$$\log p_\theta(x) = \mathcal{L}(\theta, \phi; x) + D_{\text{KL}}(q_\phi(z | x) \| p_\theta(z | x)).$$

The ELBO can be written as

$$\mathcal{L}(\theta, \phi; x) = \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x | z)] - D_{\text{KL}}(q_\phi(z | x) \| p(z)).$$

The first form explains why it is a lower bound. The second form explains why it looks like reconstruction plus regularization.

The remaining question is computational: how can we differentiate an expectation with respect to parameters that also control the sampling distribution? That is the role of the reparameterization trick.

The Reparameterization Trick and Auto-Encoding Variational Bayes

N-20251011

§157.1 The gradient problem

The ELBO for one data point is

$$\mathcal{L}(\theta, \phi; x) = \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x | z)] - D_{\text{KL}}(q_\phi(z | x) \| p(z)).$$

For the diagonal Gaussian encoder and standard Gaussian prior, the KL term has a closed form. The difficult part is the expectation

$$\mathbb{E}_{z \sim q_\phi(z|x)}[\log p_\theta(x | z)].$$

We need gradients with respect to both θ and ϕ :

$$\nabla_{\theta, \phi} \mathbb{E}_{z \sim q_\phi(z|x)}[\log p_\theta(x | z)].$$

The subtlety is that ϕ controls the distribution from which z is sampled. The sample z is random, but its randomness depends on the encoder parameters.

The reparameterization trick is the mathematical device that turns this stochastic node into a differentiable transformation of parameter-free noise.

§157.2 Score-function estimator and pathwise estimator

There is a general identity for differentiating expectations:

$$\nabla_\phi \mathbb{E}_{q_\phi(z)}[f(z)] = \mathbb{E}_{q_\phi(z)}[f(z) \nabla_\phi \log q_\phi(z)].$$

This is often called a score-function estimator. It is very general because it does not require z to be a differentiable function of ϕ . But it often has high variance.

The reparameterization trick uses a more geometric idea. Instead of sampling

$$z \sim q_\phi(z | x),$$

write

$$z = g_\phi(\epsilon, x), \quad \epsilon \sim p(\epsilon),$$

where the noise distribution $p(\epsilon)$ does not depend on ϕ .

Then

$$\mathbb{E}_{z \sim q_\phi(z|x)}[f(z)] = \mathbb{E}_{\epsilon \sim p(\epsilon)}[f(g_\phi(\epsilon, x))].$$

Now the parameters appear inside an ordinary differentiable function. Gradients can pass through g_ϕ .

§157.3 Gaussian reparameterization

For the standard VAE encoder,

$$q_\phi(z \mid x) = \mathcal{N}(z; \mu_\phi(x), \text{diag}(\sigma_\phi^2(x))).$$

A sample from this distribution can be written as

$$z = \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I).$$

Here $\odot$ denotes coordinatewise multiplication.

Thus

$$\begin{aligned} \mathbb{E}_{z \sim q_\phi(z \mid x)} [\log p_\theta(x \mid z)] \\ = \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} [\log p_\theta(x \mid \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon)]. \end{aligned}$$

The distribution of ϵ is fixed. The dependence on ϕ is now inside μ_ϕ and σ_ϕ , so ordinary backpropagation can be used.

The trick is not a heuristic. It is a change of variables in the sampling procedure.

§157.4 The stochastic ELBO estimator

With L samples $\epsilon^{(1)}, \dots, \epsilon^{(L)} \sim \mathcal{N}(0, I)$, define

$$z^{(\ell)} = \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon^{(\ell)}.$$

Then the Monte Carlo estimator of the ELBO is

$$\hat{\mathcal{L}}(\theta, \phi; x) = \frac{1}{L} \sum_{\ell=1}^L \log p_\theta(x \mid z^{(\ell)}) - D_{\text{KL}}(q_\phi(z \mid x) \parallel p(z)).$$

For minibatches, the dataset objective is approximated by averaging this quantity over the examples in the minibatch.

The original AEVB method is therefore a stochastic-gradient method for maximizing a variational lower bound. The randomness is not removed; it is organized so that differentiation is possible.

§157.5 Why implementations output log variance

Mathematically, the encoder needs $\sigma_\phi(x) > 0$. In practice, one often outputs

$$\log \sigma_\phi^2(x)$$

instead of $\sigma_\phi(x)$ directly. If

$$\ell_\phi(x) = \log \sigma_\phi^2(x),$$

then

$$\sigma_\phi(x) = \exp\left(\frac{1}{2}\ell_\phi(x)\right).$$

This automatically keeps the standard deviation positive.

The Gaussian KL then becomes

$$D_{\text{KL}}(q \parallel p) = \frac{1}{2} \sum_{j=1}^d \left(\mu_j^2 + \exp(\ell_j) - 1 - \ell_j \right).$$

This is the same formula as before, written in terms of log variance.

§157.6 Training as ELBO maximization

The mathematical training loop is:

$$\begin{aligned} \text{encoder:} & \quad x \mapsto \mu_\phi(x), \log \sigma_\phi^2(x), \\ \text{noise:} & \quad \epsilon \sim \mathcal{N}(0, I), \\ \text{latent sample:} & \quad z = \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon, \\ \text{decoder:} & \quad z \mapsto \text{parameters of } p_\theta(x | z), \\ \text{objective:} & \quad \text{maximize } \hat{\mathcal{L}}(\theta, \phi; x). \end{aligned}$$

If one minimizes a loss instead, the loss is the negative ELBO:

$$-\hat{\mathcal{L}} = -\frac{1}{L} \sum_{\ell=1}^L \log p_\theta(x | z^{(\ell)}) + D_{\text{KL}}(q_\phi(z | x) \| p(z)).$$

The first term depends on the decoder likelihood. The second term is analytic for the standard Gaussian encoder.

§157.7 Decoder likelihoods and reconstruction losses

For binary data, a Bernoulli decoder gives

$$p_\theta(x | z) = \prod_i \pi_i(z)^{x_i} (1 - \pi_i(z))^{1-x_i}.$$

The negative log-likelihood is

$$-\log p_\theta(x | z) = -\sum_i [x_i \log \pi_i(z) + (1 - x_i) \log(1 - \pi_i(z))].$$

This is binary cross entropy.

For real-valued data, a Gaussian decoder with fixed variance gives

$$p_\theta(x | z) = \mathcal{N}(x; \mu_\theta(z), \sigma^2 I).$$

Then

$$-\log p_\theta(x | z) = \frac{1}{2\sigma^2} \|x - \mu_\theta(z)\|^2 + \text{constant}.$$

This is squared error up to scaling and constants.

Thus the word “reconstruction loss” hides a probabilistic choice. Different likelihood models produce different reconstruction terms.

§157.8 Generation after training

After learning θ , generation follows the model direction, not the inference direction:

$$z \sim p(z), \quad x \sim p_\theta(x | z).$$

If the decoder likelihood is Gaussian, one may sample x from that Gaussian or use its mean as a representative output. If the decoder likelihood is Bernoulli, one samples binary coordinates or uses probabilities for visualization.

Latent interpolation uses points such as

$$z_t = (1 - t)z_0 + tz_1, \quad 0 \leq t \leq 1.$$

This often produces visually smooth transitions, but it is not a theorem that every linear path in latent space is semantically meaningful. It works when the learned latent geometry is reasonably aligned with the prior and decoder.

§157.9 Empirical marginal likelihood curves

The original experiments compared AEVB with older methods by plotting marginal log-likelihood estimates. The figure is included not as the main point of the theory, but as a reminder that the ELBO machinery was designed to make likelihood-based generative modeling trainable.




Figure 157.1: Marginal log-likelihood curves in the AEVB experiments.

The mathematical object behind these curves is still the same:

$$\log p_{\theta}(x) = \mathcal{L}(\theta, \phi; x) + D_{\text{KL}}(q_{\phi}(z | x) \| p_{\theta}(z | x)).$$

Training improves a tractable stochastic estimate of $\mathcal{L}$, and the hope is that this also improves the marginal likelihood.

§157.10 Limitations and extensions

Several limitations are already visible from the mathematics.

First, the posterior family may be too simple. A diagonal Gaussian cannot capture all posterior shapes. Normalizing flows enrich $q_{\phi}(z | x)$ by transforming a simple base distribution through invertible maps.

Second, the prior may be too simple. The standard normal prior is convenient, but the aggregated posterior may have more complicated geometry.

Third, posterior collapse may occur. If the decoder is too expressive, the model may learn

$$q_{\phi}(z | x) \approx p(z),$$

so that z carries little information about x . Then the KL term is small, but the latent representation is weak.

Fourth, the usual ELBO uses one particular weighting between reconstruction and KL. Variants such as β -VAE modify the objective:

$$\mathbb{E}_q[\log p_{\theta}(x | z)] - \beta D_{\text{KL}}(q_{\phi}(z | x) \| p(z)).$$

This changes the rate–distortion tradeoff and can encourage more disentangled latent coordinates, though not automatically.

§157.11 How the pieces fit


The VAE becomes trainable because the stochastic latent sample is rewritten as

$$z = \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I).$$

This makes the expectation differentiable with respect to the encoder parameters:

$$\mathbb{E}_{z \sim q_\phi(z|x)}[f(z)] = \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)}[f(\mu_\phi(x) + \sigma_\phi(x) \odot \epsilon)].$$

The full mathematical chain is now visible:



This is why the VAE is more than an autoencoder with noise. It is a method for turning approximate Bayesian inference in a latent variable model into differentiable optimization.

Pure Mathematics as a Language for Machine Learning and Data Analysis

N-20260302

§158.1 Backpropagation as cotangent pullback

The previous backpropagation calculation can be read in a categorical way without making it mystical. A forward neural network is a composite of maps between parameter spaces, activation spaces, and finally the one-dimensional loss space. If

$$L : Y \rightarrow \mathbb{R}$$

is the loss, then its differential at the output is a cotangent vector

$$dL_y \in T_y^*Y.$$

Backpropagation pulls this covector backward along the forward computational graph. In this sense the derivative operation used in training is contravariant:

$$f : X \rightarrow Y \implies f^* : T^*Y \rightarrow T^*X.$$

Forward maps push tangent perturbations forward; reverse-mode automatic differentiation pulls covectors backward.

For a concrete PyTorch-shaped layer, take

$$X \in \mathbb{R}^{b \times d}, \quad W \in \mathbb{R}^{n \times d}, \quad Y = XW^T \in \mathbb{R}^{b \times n}.$$

In indices,

$$Y_{\alpha i} = \sum_{k=1}^d X_{\alpha k} W_{ik}.$$

The full Jacobian with respect to the weight matrix has entries

$$\frac{\partial Y_{\alpha i}}{\partial W_{jk}} = \delta_{ij} X_{\alpha k},$$

so its tensor shape is

$$(b, n, n, d).$$

This is the honest tensor object. Neural-network code does not materialize it because the upstream gradient

$$G = \frac{\partial L}{\partial Y} \in \mathbb{R}^{b \times n}$$

immediately contracts it:

$$\begin{aligned} \frac{\partial L}{\partial W_{jk}} &= \sum_{\alpha=1}^b \sum_{i=1}^n \frac{\partial L}{\partial Y_{\alpha i}} \frac{\partial Y_{\alpha i}}{\partial W_{jk}} \\ &= \sum_{\alpha=1}^b G_{\alpha j} X_{\alpha k}. \end{aligned}$$

Therefore

$$\frac{\partial L}{\partial W} = G^T X.$$

If one stores weights in the column convention $W_{\text{col}} \in \mathbb{R}^{d \times n}$ and writes $Y = XW_{\text{col}}$, the same contraction is

$$\frac{\partial L}{\partial W_{\text{col}}} = X^T G.$$

The categorical phrase “pull back the covector” is exactly this matrix multiplication: the high-order Jacobian collapses because it is being paired with a covector from the scalar loss.

§158.2 Geometric data analysis

The original note is an ambitious AI-assisted essay about the role of pure mathematics in machine learning. Its most useful part is not the rhetoric, but the catalogue of mathematical languages that repeatedly appear when data are no longer well described by coordinates alone.

A point cloud is often treated as a finite sample from a hidden space. Topological data analysis makes this idea precise by replacing the cloud with a filtered simplicial complex. For a scale parameter ε , one forms for instance the Vietoris–Rips complex

$$VR_\varepsilon(X) = \{\sigma \subseteq X : d(x, y) \leq \varepsilon \text{ for all } x, y \in \sigma\}.$$

As ε grows, connected components merge, loops appear and disappear, and higher-dimensional holes may be born and die. Persistent homology records this evolution. The chain complex viewpoint is governed by the boundary identity

$$\partial_{k-1} \circ \partial_k = 0,$$

and homology is

$$H_k = \ker \partial_k / \text{im } \partial_{k+1}.$$

This quotient is the mathematical form of the phrase “holes that are cycles but not boundaries.”

§158.3 Manifold hypothesis and embeddings

The manifold hypothesis says that high-dimensional data often concentrate near a lower-dimensional manifold $M \subseteq \mathbb{R}^N$. If M is a smooth d -manifold, Whitney’s embedding theorem says that M can be embedded in $\mathbb{R}^{2d}$, and generically in a slightly larger ambient space. In machine learning this gives a conceptual reason why dimensionality reduction is not only compression: it tries to recover the intrinsic coordinates of M .

The practical warning is that an embedding preserves the manifold as a smooth object, but not necessarily distances, probability densities, or task-relevant semantics. PCA preserves a linear subspace model; manifold learning tries to preserve local neighborhoods; representation learning tries to learn coordinates useful for a loss function.

§158.4 Optimal transport and stochastic dynamics

For probability measures μ, ν on a metric space, the p -Wasserstein distance is

$$W_p(\mu, \nu) = \left(\inf_{\pi \in \Pi(\mu, \nu)} \int d(x, y)^p d\pi(x, y) \right)^{1/p}.$$

Here $\Pi(\mu, \nu)$ is the set of couplings with marginals μ and ν . The dual viewpoint is equally important. For $p = 1$,

$$W_1(\mu, \nu) = \sup_{\|f\|_{\text{Lip}} \leq 1} \left(\int f d\mu - \int f d\nu \right).$$

For $p = 2$ there is a useful one-dimensional closed form. Let

$$\mu = \mathcal{N}(a, \sigma_1^2), \quad \nu = \mathcal{N}(b, \sigma_2^2), \quad \sigma_1, \sigma_2 > 0.$$

In one dimension the optimal map between these two Gaussians is affine and monotone, so write

$$T(x) = \alpha x + \beta.$$

The condition $T_{\#}\mu = \nu$ forces

$$\alpha = \frac{\sigma_2}{\sigma_1}, \quad \beta = b - \alpha a.$$

If $X = a + \sigma_1 Z$ with $Z \sim \mathcal{N}(0, 1)$, then

$$T(X) = b + \sigma_2 Z.$$

Thus

$$\begin{aligned} W_2^2(\mu, \nu) &= \mathbb{E}[(X - T(X))^2] \\ &= \mathbb{E}[((a - b) + (\sigma_1 - \sigma_2)Z)^2] \\ &= (a - b)^2 + (\sigma_1 - \sigma_2)^2. \end{aligned}$$

This formula is small, but it explains a lot of geometric language around generative models: matching Gaussians means matching both their centers and their spreads.

Diffusion models can be read through stochastic differential equations

$$dX_t = b(t, X_t) dt + \sigma(t) dW_t.$$

The density p_t evolves according to a Fokker–Planck equation, and the reverse-time dynamics involve the score $\nabla_x \log p_t(x)$. Thus score-based generative modeling is an attempt to learn the vector field that reverses a noisy transport of probability mass.

For DDPMs the training objective is usually written as a variational lower bound, or in its simplified form as a denoising mean-square loss. It is not literally a Wasserstein objective. Still, each reverse step compares Gaussian conditional distributions, and for Gaussians the geometry above says that a squared displacement of means and a mismatch of variances are exactly the ingredients of W_2^2 . The ELBO uses KL divergences between Gaussian conditionals; when the variance is fixed or separately prescribed, minimizing that KL has the same practical center-matching behavior as minimizing the corresponding Gaussian W_2 term. My preferred reading is therefore: optimal transport gives the geometric picture of local probability movement, while the DDPM ELBO gives the statistically convenient training objective.

§158.5 Symmetry, equivariance, and category-like thinking

If a group G acts on an input space and an output space, a map F is equivariant if

$$F(gx) = gF(x).$$

Convolutional neural networks exploit translation equivariance; graph neural networks exploit permutation equivariance; gauge-equivariant networks try to respect local choices of frame on a bundle.

Here is the matrix calculation behind the slogan. Let the graph be the cycle C_4 , with adjacency matrix

$$A = \begin{pmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{pmatrix},$$

and let P be the permutation matrix rotating the four vertices by 90° . Since P is a graph automorphism,

$$PA = AP.$$

The same is true for the usual normalized adjacency $\hat{A}$, because normalization is built from A and the degree matrix, both invariant under the same rotation:

$$P\hat{A} = \hat{A}P.$$

A standard graph neural-network layer is

$$H^{(l+1)} = \sigma(\hat{A}H^{(l)}W),$$

where σ acts entrywise. If the input features are rotated, then

$$\begin{aligned} F(PH^{(l)}) &= \sigma(\hat{A}PH^{(l)}W) \\ &= \sigma(P\hat{A}H^{(l)}W) \\ &= P\sigma(\hat{A}H^{(l)}W) \\ &= PH^{(l+1)}. \end{aligned}$$

Equivariance is therefore not only a philosophical commitment to symmetry. In this example it is the algebraic fact that the permutation matrix commutes with the graph operator.

The original essay also speculates about topos-theoretic or categorical semantics for AI. A cautious reconstruction is this: category theory is useful when one needs to track structure-preserving maps, functorial changes of representation, and compositional semantics. It does not by itself explain intelligence. Its honest role is to provide a disciplined language for invariance, abstraction, and transport of structure.

§158.6 Further direction

The note is best read as a map of mathematical languages for data: topology sees holes, geometry sees coordinates and curvature, measure theory sees distributions, dynamics sees evolution, algebra sees symmetry, and category theory sees structure-preserving translation. The useful question is not which language is fashionable, but which structure is actually stable under the transformations used by the model.

§158.7 Backpropagation and differential geometry

Backpropagation is reverse-mode differentiation. Geometrically, forward computation pushes tangent vectors forward, while backpropagation pulls cotangent vectors backward through the computation graph.

§158.8 Manifold hypothesis and representation learning

The manifold hypothesis says high-dimensional data may concentrate near lower-dimensional geometric structures. Embeddings, charts, curvature, and local neighborhoods then become relevant to learning representations.

§158.9 Symmetry and equivariance

If a task has a symmetry group G , an equivariant model satisfies

$$f(gx) = gf(x).$$

This reduces sample complexity by building the correct transformation law into the model.

§158.10 Optimal transport and distributions

Optimal transport compares probability measures by the cost of moving mass. It connects geometry, PDE, and machine learning through distances between distributions.

§158.11 Mathematical structures for learning systems

Pure mathematics contributes language for structure: cotangent pullback for gradients, manifolds for data geometry, equivariance for symmetry, optimal transport for distributions, and category-like compositionality for architectures.

§158.12 Rebuilt reading path: what pure mathematics contributes

There are several layers.

- (i) **Geometry** studies the shape of the representation space: distances, curvature, geodesics, embeddings, and coordinate charts.
- (ii) **Topology** studies features stable under continuous deformation: connected components, loops, voids, and persistence across scales.
- (iii) **Measure theory** studies data as distributions rather than isolated points.
- (iv) **Dynamical systems** study learning or generation as evolution of states or densities.
- (v) **Algebra and category theory** study invariance, symmetry, and composition of structure-preserving maps.

Pure mathematics does not magically explain AI. It supplies languages in which different kinds of stability can be stated precisely.

§158.13 Topological data analysis in usable form

Given a finite metric point cloud X , persistent homology constructs a family of complexes K_ε indexed by a scale parameter. For the Vietoris–Rips complex,

$$\{x_0, \dots, x_k\} \in VR_\varepsilon(X) \iff d(x_i, x_j) \leq \varepsilon \text{ for all } i, j.$$

The homology group is

$$H_k(K_\varepsilon) = \ker \partial_k / \operatorname{im} \partial_{k+1}.$$

This quotient is the mathematical form of the phrase “holes that are cycles but not boundaries.”

§158.14 Wasserstein geometry and stochastic dynamics

For probability measures μ, ν on a metric space,

$$W_p(\mu, \nu) = \left(\inf_{\pi \in \Pi(\mu, \nu)} \int d(x, y)^p d\pi(x, y) \right)^{1/p}.$$

The coupling π is the hidden plan. In generative modeling, one often tries to construct a transport from a simple distribution to a complicated one. Normalizing flows do this by invertible maps, diffusion models by stochastic dynamics, and optimal transport by variational minimization.

A diffusion process has the schematic form

$$dX_t = b(t, X_t) dt + \sigma(t) dW_t.$$

The reverse process depends on the score field $\nabla_x \log p_t(x)$. Training a score network is therefore an attempt to approximate a time-dependent vector field on probability density.

§158.15 Category language without exaggeration

A category consists of objects and morphisms, with associative composition and identity morphisms. A functor preserves this structure. Category theory is useful when the same construction can be performed across many spaces in a way compatible with maps. For example, homology is functorial:

$$f : X \rightarrow Y \implies H_k(f) : H_k(X) \rightarrow H_k(Y).$$

The honest categorical contribution is functoriality and compositionality, not a mystical explanation of intelligence.