

Honors Abstract Algebra

Course Notes
Math 55a, Harvard University

Contents

1	Introduction	1
2	Set Theory	2
3	Vector spaces	4
4	Polynomials	8
5	Linear Operators	9
6	Inner product spaces	16
7	Bilinear forms	24
8	Trace and determinant	28
9	Introduction to Group Theory	34
10	Symmetry	38
11	Finite group theory	51
12	Representation theory	55
13	Group presentations	65
14	Knots and the fundamental group	69

1 Introduction

This course will provide a rigorous introduction to abstract algebra, including group theory and linear algebra.

Topics include:

1. Set theory. Formalization of $\mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$.
2. Linear algebra. Vector spaces and transformations over \mathbb{R} and \mathbb{C} . Other ground fields. Eigenvectors. Jordan form.
3. Multilinear algebra. Inner products, quadratic forms, alternating forms, tensor products, determinants.
4. Abstract groups.
5. Groups, symmetry and representations.

2 Set Theory

Halmos reading. Read Halmos, *Naive Set Theory*, sections 1–15 to learn the foundations of mathematics from the point of view of set theory, and its use in formalizing the integers. Most of this should be review, although the systematic use of a small set of axioms to rigorously establish set theory may be new to you. We will also formalize the rational, real and complex numbers.

Then read 22–23 to learn about cardinality and countable sets.

Finally, read 16–21 and 24–25 to learn about other versions of the axiom of choice, ordinals and cardinals.

Axioms.

1. (Extension.) $A = B$ iff they have the same elements.
2. (Specification.) $\{x \in A : P(x)\}$ exists.
3. (Pairs.) $\{A, B\}$ exists.
4. (Union.) $\bigcup A$ exists.
5. (Infinity.) \mathbb{N} exists. (A set containing zero and closed under $n + 1 = n \cup \{n\}$).
6. (Power set.) $\mathcal{P}(A)$ exists.
7. (Axiom of Choice.) There exists a choice function $c : \mathcal{P}(A) - \{\emptyset\} \rightarrow A$ such that $c(A) \in A$.
8. (Well-foundedness.) There is no infinite sequence $x_1 \ni x_2 \ni x_3 \dots$
9. (Replacement.) If $f(x)$ is a function defined by a logical expression, and A is a set, then $f(A)$ is a set.

Power set and algebra. It is useful to introduce the ring \mathbb{Z}/n (also as an example of an equivalence relation). Then we have a natural bijection between $\mathcal{P}(A)$ and $(\mathbb{Z}/2)^A$. But the latter space is a ring! Note that $+$ becomes symmetric difference, while $*$ is just intersection.

If $f : A \rightarrow B$ is a map, then we get a natural map $f^{-1} : \mathcal{P}(B) \rightarrow \mathcal{P}(A)$. This respects the ring operations. Note that $E \mapsto f(E)$ respects some but not all! Of course if f is a bijection, it gives an isomorphism.

Cardinality. We write $|A| = |B|$ if there is a bijection $f : A \rightarrow B$. A set is *finite* if $|A| = |n|$ for some $n \in \mathbb{N}$. This n is unique and is called the cardinality of A .

Finite sets satisfy the pigeon-hole principle: a map $f : A \rightarrow A$ is 1 – 1 iff it is onto. Infinite sets are *exactly* the sets which do not (Hilbert’s Hotel). Using AC, one can show that any infinite set contains a copy of \mathbb{N} , and then $n \mapsto n + 1$ is 1-1 but not onto.

A set A is *countable* if it is finite or $|A| = |\mathbb{N}|$. Note that \mathbb{N}^2 is countable; you can then show \mathbb{N}^n is countable for all n . Similarly a countable union of countable sets is countable. For examples, $\mathbb{Z}[x]$ is countable.

On the other hand, \mathbb{R} is *uncountable* (diagonalization). So we see there are (at least) 2 kinds of infinity! This also gives a proof that:

There exists a transcendental number.

This proof is *somewhat* nonconstructive. An explicit such number is also easy to give, e.g. $\sum 1/10^{n!}$.

Let us also say $|A| \leq |B|$ if there is an injective map $f : A \hookrightarrow B$. Using the Axiom of Choice, we find:

If $g : B \rightarrow A$ is surjective, then there exists an injective map $f : A \rightarrow B$.

(Note: although $|0| < |1|$, there is no surjective map $1 \rightarrow 0$.)

Theorem 2.1 *We have $|\mathcal{P}(A)| > |A|$.*

Proof. Clearly $A \hookrightarrow \mathcal{P}(A)$ so $|\mathcal{P}(A)| \geq |A|$. To show $|\mathcal{P}(A)|$ is strictly bigger, suppose $f : A \rightarrow \mathcal{P}(A)$ is a surjection, and let $B = \{a \in A : a \notin f(a)\}$. Then we cannot have $B = f(a)$, for if we did, then $a \in B \iff a \in f(a) \iff a \notin f(a)$. ■

Question. (CH) Is there a set with $|\mathbb{N}| < |A| < |\mathcal{P}(\mathbb{N})|$? This famous problem — the continuum hypothesis — is now known to have *no solution*.

(How can you prove something cannot be proved? Compare the parallel postulate, and the axiom of infinity. Note that you need to assume your existing assumptions are consistent.)

Theorem 2.2 (Schröder–Bernstein) *If $|A| \leq |B|$ and $|B| \leq |A|$, then $|A| = |B|$.*

Proof. We may assume A and B are disjoint. Let $C = A \cup B$. Then $F = f \cup g : C \rightarrow C$ is injective. By following F backwards, we can partition C into three sets: $C_\infty = \bigcap f^i(C)$, those with an infinite chain of ancestors; $C_A =$ those born in A , and $C_B =$ those born in B . Taking intersection with A , we get a bijection $h : A \rightarrow B$ by setting $h(x) = f(x)$ for $x \in A \cap C_A$, setting $h(x) = g^{-1}(x)$ for $x \in A \cap C_B$, and setting $h(x) = f(x)$ for $x \in A \cap C_\infty$. ■

Theorem 2.3 *The Axiom of Choice is equivalent to the assertion that every set can be well-ordered.*

Proof. If $(A, <)$ is well-ordered, we can define $c(B)$ to be the least element of B .

For the converse, let c be a choice function for A . Let us say a well-ordering $(B, <)$ of a subset of A is *guided by c* if for all $x \in B$, we have $x = c(A - \{y \in B : y < x\})$. It is easy to see that if orderings on B and B' are both guided by c , then $B \subset B'$ or vice-versa, and the orderings agree on $B \cap B'$. Taking the union of order relations compatible with c , we obtain a well-ordering of A . ■

3 Vector spaces

Axler reading. We will discuss groups, rings, fields, and vector spaces over arbitrary fields. Our main text Axler [Ax] discusses only the fields \mathbb{R} and \mathbb{C} . For more general definitions, see [Ar, Ch. 2,3].

Note also that Axler discusses the direction sum $S \oplus T$ of two *subspaces* of a given vector space V . In general one also uses the same notation, $S \oplus T$, to construct a *new* vector space from two given ones, whose elements are ordered pairs $(s, t) \in S \times T$ with the obvious coordinate-wise operations.

Groups and fields. An *abelian group* is a set G with an operation $+$ and an element 0 such that for all $a, b, c \in G$, $0 + a = a$, $a + b = b + a$, $a + (b + c) = (a + b) + c$, and there exists an element d (denoted $-a$) such that $a + d = 0$.

A *field* K is a set with two operations, plus and times, such that $(K, +)$ is an abelian group with identity 0 , and (K^*, \cdot) is an abelian group with identity 1 (where $K^* = K - \{0\}$), and $a(b + c) = ab + ac$.

Note that under plus, \mathbb{N} is not a group but \mathbb{Z} is. Times makes \mathbb{Q} , \mathbb{R} and \mathbb{C} into fields, but not \mathbb{Z} .

Finite fields. When p is a prime, the ring \mathbb{Z}/p is actually a field (usually denoted \mathbb{F}_p). It is the unique field (up to isomorphism) containing p elements.

To see \mathbb{Z}/p is a field, just note that if $xy = 0 \pmod p$ then $p|xy$, which (by unique factorization) means $p|x$ or $p|y$, and hence $x = 0$ or $y = 0 \pmod p$. So if $x \neq 0$, the map $\mathbb{Z}/p \rightarrow \mathbb{Z}/p$ given by $y \mapsto xy$ is $1 - 1$. By the pigeonhole principle, it is onto, so there is a y such that $xy = 1 \pmod p$.

Inversion in \mathbb{F}_p can be carried out by the Euclidean algorithm. For example, to compute $1/10$ in \mathbb{F}_{37} , we write $37 = 3 \cdot 10 + 7$, $10 = 7 \cdot 1 + 3$, $7 = 2 \cdot 3 + 1$, and then back substitute:

$$1 = 7 - 2 \cdot 3 = 7 - 2 \cdot (10 - 7) = 3 \cdot 7 - 2 \cdot 10 = 3 \cdot (37 - 3 \cdot 10) - 2 \cdot 10 = 3 \cdot 37 - 11 \cdot 10;$$

and thus $1/10 = 11 \pmod{37}$.

The Euclidean algorithm can be conceived of geometrically as using a small ruler to measure a big one and create a still smaller ruler with the remainder. If one starts with the two sides of a golden rectangle, the process never ends. In this way the Greeks (including the secretive Pythagoreans) knew the existence of irrational numbers.

Vector spaces. A vector space over a field K is an abelian group $(V, +)$ equipped with a multiplication map $K \times V \rightarrow V$ such that $(\alpha + \beta)v = \alpha v + \beta v$, $(\alpha\beta)v = \alpha(\beta v)$, and (especially!) $1v = v$.

Examples: K , K^n . The reals form a vector space over \mathbb{Q} .

Matrices. The $n \times n$ matrices $M_n(K)$ with entries in a field K form a vector space under dimension. It looks just like K^{n^2} . More important, they form a *ring*; they can be multiplied by the usual rule

$$(ab)_{ij} = \sum_{k=1}^n a_{ik}b_{kj}.$$

The study of this ring will be central to the theory of vector spaces.

We will later show that $GL_n(K)$, the group of matrices with nonzero determinant, is a group under multiplication. For $n = 2$ this is easy: we have

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}^{-1} = \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix}.$$

This group is (for $n > 1$) never commutative, e.g. $a = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and $b = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$ satisfy

$$ab = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}, \quad ba = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}.$$

(Although there is a field where $2 = 0$, there is no field where $2 = 1$!)

\mathbb{R}^4 . It is sometimes useful to appreciate the geometry and topology of higher-dimensional vector spaces. Here is an example. In \mathbb{R}^2 , a circle can ‘enclose’ a point. The two objects are linked. But in \mathbb{R}^3 you can move the point out of the plane and then transport it to the outside of the circle, without ever crossing it.

As a test of visualization: show that two circles which are linked in \mathbb{R}^3 can be separated in \mathbb{R}^4 .

Linear interdependence. A basic feature of linear dependence is the following. Suppose

$$0 = \sum_{i=1}^n a_i x_i,$$

and all $a_i \neq 0$. (One might then say the x_i are *interdependent*). Then the span of $(x_1, \dots, \widehat{x}_i, \dots, x_n)$ is the same as the span of (x_1, \dots, x_n) , for all i . In other words, *any one* of the x_i can be eliminated, without changing their span.

From this we get the main fact regarding bases.

Theorem 3.1 *Let A be a linear independent set and B a finite spanning set for a vector space V . Then $|A| \leq |B|$.*

Proof. Write $A_i = (a_1, \dots, a_i)$ (so $A_0 = \emptyset$). We inductively construct a sequence of spanning sets of form $A_i \cup B_i$, as follows. Let $B_0 = B$; then $A_0 \cup B_0$ spans. Assuming $A_i \cup B_i$ spans, we can express a_{i+1} as a linear combination of elements in $A_i \cup B_i$. These interdependent vectors must include at least one from B_i , since A is an independent set. Thus we can remove one element of B_i , to obtain a set B_{i+1} such that $A_{i+1} \cup B_{i+1}$ still spans. Note that $|B_i| = |B| - i$.

The induction can proceed until i reaches the minimum n of $|A|$ and $|B|$. If $n = |B| < |A|$ then $B_n = \emptyset$, while A_n is a proper subset of A that spans V . This contradicts the linear independence of A . Thus $|B| \geq |A|$. ■

Assuming V is finite-dimensional, we obtain:

Corollary 3.2 *Any linearly independent set be extended to a basis.*

Corollary 3.3 *All bases have the same number of elements.*

When applied to the case where A and B are both bases, the proof gives more: it gives a sequence of bases $A_i \cup B_i$ that interpolates between A and B . This can be expressed as a factorization theorem for general automorphisms of V .

Dimension. If V is a finite-dimensional vector space, we define $\dim V$ to be the number of elements in a basis. We have just shown $\dim V$ is well-defined.

An *isomorphism* between vector space is a bijection map $T : V_1 \rightarrow V_2$ such that $T(x + y) = T(x) + T(y)$ and $T(\lambda x) = \lambda T(x)$. It is easily verified that T^{-1} is also an isomorphism.

Theorem 3.4 *Any finite-dimensional vector space over K is isomorphic to K^n for a unique n .*

Note that there may be many different isomorphisms $T : K^n \rightarrow V$. In fact, an isomorphism is *the same* as a choice of basis for V . Thus for any two bases $(e_i), (f_i)$ of V , there is a unique isomorphism $T : V \rightarrow V$ such that $T(e_i) = f_i$.

More generally, a linear independent set $(e_i)_1^n$ for V is *the same* as the choice of an *injective linear map* $T : \mathbb{R}^n \rightarrow V$, defined by $T(x) = \sum x_i e_i$; the the choice of a spanning set is the same as the choice of a *surjective linear map* $T : \mathbb{R}^n \rightarrow V$.

Infinite-dimensional vector spaces. Here is a result that conveys the power of the Axiom of Choice.

Theorem 3.5 *Any vector space V has a basis. For example, \mathbb{R} has a basis as a vector space over \mathbb{Q} .*

Proof. Choose a well-ordering $<$ for V (using the Axiom of choice). Let S be the set of $v \in V$ such that v is not in the linear span of $\{w \in V : w < v\}$. It is easy to see that the elements of S are linearly independent. Suppose the span S' of S is not all of V . Then, by well-ordering, there is a *least* element $v \in V - S'$. But then $v = \sum_1^n a_i v_i$ with $v_i < v$, else we would have $v \in S$. And each v_i lies in S' , since $v_i < v$. But then $v \in S'$. ■

Theorem 3.6 *There is a map $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfying $f(x + y) = f(x) + f(y)$, $f(x) = 0$ if x is rational, and $f(\sqrt{2}) = 1$.*

Proof. Let $B_0 = \{1, \sqrt{2}\}$. Using a small variation of the proof above, we can extend B_0 to a basis B for \mathbb{R} over \mathbb{Q} . Then any $x \in \mathbb{R}$ can be written uniquely as $x = \sum_B x_b \cdot b$ with $x_b \in \mathbb{Q}$ and $x_b = 0$ for all but finitely many $b \in B$. This implies $(x + y)_b = x_b + y_b$. Now define $f(x) = x_{\sqrt{2}}$. ■

4 Polynomials

The polynomials $K[x]$ form a vector space over K . The elements of $K[x]$ are formal sums $\sum_0^n a_i x^i$ where $a_i \in K$. Thus the polynomials of degree d or less form a vector space of dimension $d + 1$.

Axler defines polynomials (p.10) as certain *functions* $f : K \rightarrow K$, namely those of the form $f(x) = \sum_0^n a_i x^i$. This is fine for fields like \mathbb{Q} and \mathbb{R} , but it is not the right definition in general. For example, if $K = \mathbb{F}_p$ is the field with p elements, there are infinitely many polynomials in $\mathbb{F}_p[x]$, but only finitely many maps $f : \mathbb{F}_p \rightarrow \mathbb{F}_p$. An important special case is the polynomial $f(x) = x^p - x$, which vanishes for every $x \in \mathbb{F}_p$ but is not the zero polynomial.

Division of polynomials. A ratio of polynomials p/q , $q \neq 0$, can always be written as a ‘proper fraction’, $p/q = s + r/q$, where $\deg(r) < \deg(q)$. Equivalently, we have:

Theorem 4.1 *Given $p, q \in K[x]$, $q \neq 0$, there exist unique polynomials $s, r \in K[x]$ such that $p = sq + r$ and $\deg(r) < \deg(q)$.*

Using this fact one can show that polynomials have unique factorization into irreducibles and have gcd’s. In particular we have this theorem:

Theorem 4.2 *If $p, q \in K[x]$ have no common factor, then there exist $r, s \in K[x]$ such that $sp + rq = 1$.*

Complex polynomials. We will often use:

Theorem 4.3 (Fundamental theorem of algebra) *Every polynomial p in $\mathbb{C}[z]$ of positive degree has at least one complex zero. Thus p is a product of linear polynomials.*

This makes it easy to see when p and q have a common factor.

Theorem 4.4 *If $p, q \in \mathbb{C}[x]$ have no common zeros, then there exist $r, s \in \mathbb{C}[x]$ such that $sp + rq = 1$.*

See [Ar, Ch. 11] for more details on polynomials.

5 Linear Operators

Theorem 5.1 (Conservation of dimension) *For any linear map $T : V \rightarrow W$, we have $\dim V = \dim \text{Ker } T + \dim \text{Im } T$.*

Proof. By lifting a basis for $\text{Im } T$ we get a subspace $S \subset V$ mapping bijectively to the image, and with $V = \text{Ker } T \oplus S$. ■

Corollary 5.2 *There exists a basis for V and W such that T has the form of a projection followed by an inclusion: $T : \mathbb{R}^n \rightarrow \mathbb{R}^i \subset \mathbb{R}^j$.*

This result shows that not only is the theory of finite-dimensional vector spaces trivial (they are classified by their dimension), but the theory of maps between *different* vector spaces V and W is also trivial. It is for this reason that we will concentrate on the theory of *operators*, that is the (dynamical) theory of maps $T : V \rightarrow V$ from a vector space to itself.

Rank. The dimension of $\text{Im } T$ is also known as the *rank* of T . When T is given by a matrix, the columns of the matrix span the image. Thus the rank of T is the maximal number of linearly independent columns.

Clearly the rank of ATB is the same as the rank of T , if A and B are automorphisms.

The row rank and the column rank of a matrix are equal. This is clear when the matrix of T is a projection of \mathbb{R}^n onto \mathbb{R}^i ; it then follows from the Corollary above, by invariance of both quantities under composition with automorphisms.

We will later see a more functorial explanation for this, in terms of vector spaces and their duals.

The rank can be found by repeated row and/or column reduction. Row reduction can be made into an algorithm to compute the inverse of T , as well as a basis for the kernel when T is not invertible.




Figure 1. The quotient space $\mathbb{R}^3/\mathbb{R}(1, 1, 1)$

Example. Row operations on $T = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 2 & 3 \\ 1 & 4 & 9 \end{pmatrix}$ lead to $\begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \\ 0 & 3 & 8 \end{pmatrix}$ and then $\begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \\ 0 & 0 & 2 \end{pmatrix}$, showing this matrix has rank 3.

Quotient spaces. If $U \subset V$ is a subspace, the *quotient space* V/U is defined by $v_1 \sim v_2$ if $v_1 - v_2 \in U$. It is a vector space in its own right, whose elements can be regarded as the parallel translates $v + U$ of the subspace U . There is a natural surjective linear map $V \rightarrow V/U$. We have

$$\dim(V/U) = \dim V - \dim U.$$

For more background see [Hal].

If $T \in \text{Hom}(V, V)$ and $T(U) \subset U$ then T induces a *quotient map* on V/U by $v + U \mapsto T(v + U) + U = T(v) + U$.

Exact sequences. A pair of linear maps

$$U \xrightarrow{S} V \xrightarrow{T} W$$

is said to be *exact* at V if $\text{Im } S = \text{Ker } T$. A sequence of linear maps $V_1 \rightarrow V_2 \rightarrow V_3 \rightarrow \dots$ is said to be *exact* if it is exact at every V_i . Any linear map $T : V \rightarrow W$ gives rise to a *short exact sequence* of the form

$$0 \rightarrow \text{Ker}(T) \rightarrow V \xrightarrow{T} \text{Im}(T) \rightarrow 0.$$

Block diagonal and block triangular matrices. Suppose $V = A \oplus B$ with a basis (e_i) that runs first through a basis of A , then through a basis of B .

If $T(A) \subset A$, then the matrix of T is *block triangular*. It has the form $T = \begin{pmatrix} T_{AA} & T_{BA} \\ 0 & T_{BB} \end{pmatrix}$. The T_{AA} block gives the matrix for $T|_A$. The T_{BB} block gives the matrix for $T|(V/A)$. So quotient spaces occur naturally when we

consider block triangular matrices. Note that if $S(A) = A$ as well, then the diagonal blocks for TS can be computed from those for S and those for T .

Finally if $T(B) \subset B$ then $T_{BA} = 0$ and the matrix is block diagonal. This means $T = (T|_A) \oplus (T|_B)$.

Eigenvectors. Now let V be a finite-dimensional vector space over \mathbb{C} . An *eigenvector* is a nonzero vector $e \in V$ such that $Te = \lambda e$ for some $\lambda \in \mathbb{C}$, called the *eigenvalue* of e .

Theorem 5.3 *Any linear operator $T : V \rightarrow V$ has an eigenvector.*

Proof. Let $n = \dim V$ and choose $v \in V$, $v \neq 0$. Then the vectors $(v, Tv, \dots, T^n v)$ must be linearly dependent. In other words, there is a polynomial of degree n such that $p(T)v = 0$. Factor this polynomial; we then have

$$(T - \lambda_1) \cdots (T - \lambda_n)v = 0.$$

It follows that $\text{Ker}(T - \lambda_i)$ is nontrivial for some i . ■

Theorem 5.4 *Every operator on a finite-dimensional vector space over \mathbb{C} has a basis such that its matrix is upper-triangular.*

Proof. Since \mathbb{C} is algebraically closed, T has an eigenvector v_1 , and hence an invariant 1-dimensional subspace $S = \mathbb{C}v_1$. Now proceed by induction. Consider the quotient map $T : V/S \rightarrow V/S$. This map has a basis $(\bar{v}_2, \dots, \bar{v}_n)$ making it upper-triangular, where $\bar{v}_i = \pi(v_i)$. Then (v_1, v_2, \dots, v_n) put T into upper-triangular form. ■

Theorem 5.5 *An upper-triangular matrix is invertible iff there are no zeros on its diagonal.*

Proof. If there are no zeros on the diagonal, we can in fact explicitly invert $Ta = b$ by first solving $\lambda_n a_n = b_n$, and then working our way back to a_1 .

For the converse, suppose $\lambda_{i+1} = 0$. Then T maps the span V_{i+1} of (e_1, \dots, e_{i+1}) into V_i . So it has a nontrivial kernel. ■

Corollary 5.6 *The eigenvalues of T coincide with the diagonal values of any upper triangular matrix representing T .*

Corollary 5.7 *The number of eigenvalues of T is at most $\dim V$.*

Flags. A *flag* is an ascending sequence of subspaces, $0 = V_0 \subset V_1 \subset \cdots \subset V_n = V$. A flag is *maximal* if $\dim V_i = i$.

Assume V is a vector space over \mathbb{C} (or any algebraically closed field). The theorem on upper triangular form can be re-phrased and proved as follows.

Theorem 5.8 *Any linear operator $T : V \rightarrow V$ leaves invariant a maximal flag.*

Diagonalization. The simplest possible form for an operator is that it is diagonal: it has a basis of eigenvectors, $Te_i = \lambda_i e_i$. Unfortunately this cannot always be achieved: consider the operator $T(x, y) = (y, 0)$. (Its only eigenvalue is zero, but it is not the zero map!) However we do have:

Theorem 5.9 *If e_1, \dots, e_m are eigenvectors for T with distinct eigenvalues, then they are linearly independent. In particular, if the number of eigenvalues of T is equal to $\dim V$, then T is diagonalizable.*

Proof. Suppose $\sum_1^m a_i e_i = 0$. We can assume all $a_i \neq 0$ and that no smaller set of vectors is linearly dependent. But then $\sum_1^m \lambda_i a_i e_i = 0$, so $\sum_1^m (\lambda_i - \lambda_j) a_i e_i = 0$ for any j with $1 \leq j \leq m$. This gives a smaller linearly dependent set. ■

Statement in terms of matrices. The preceding results show that for any $T \in M_n(\mathbb{C})$, we can find an $A \in GL_n(\mathbb{C})$ such that ATA^{-1} is upper triangular, and even diagonal if T has n distinct eigenvalues.

Generalized kernels. If $T^i v = 0$ for some $i > 0$, we say v is in the *generalized kernel* of T . Since $\text{Ker}(T^i)$ can only increase with i , it must stabilize after at most $n = \dim V$ steps. Thus the generalized kernel of T coincides with $\text{Ker}(T^n)$, $n = \dim V$.

Similarly we say v is a *generalized eigenvector* of T if v is in the generalized kernel of $T - \lambda$, i.e. $(T - \lambda)^i v = 0$ for some i . The set of all such v forms the *generalized eigenspace* V_λ of T , and is given by $V_\lambda = \text{Ker}(T - \lambda)^n$. The number $m(\lambda) = \dim V_\lambda$ is the *multiplicity* of V_λ as an eigenvalue of T .

Proposition 5.10 *We have $T(V_\lambda) = V_\lambda$.*

Example. Let $T(x, y) = (x + y, y)$. The matrix for this map, $T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$, is already upper triangular. Clearly $v_1 = (1, 0)$ is an eigenvector with eigenvalue 1. But there is no second eigenvector to form a basis for V . This is clearly seen geometrically in \mathbb{R}^2 , where the map is a shear along horizontal lines (fixing the x -axis, and moving the line at height y distance y to the right.)

However if we try $v_2 = (0, 1)$ we find that $T(v_2) - v_2 = v_1$. In other words, taking $\lambda = 1$, $v_1(T - \lambda)v_2 \neq 0$ is an eigenvector of T . And indeed $(T - \lambda)^2 = 0$ so $V_\lambda = V$. This is the phenomenon that is captured by generalized eigenspaces.

Here is another way to think about this: if we conjugate T by $S(x, y) = (ax, y)$, then it becomes $T(x + (1/a)y, y)$. So the conjugates of T converge to the identity. The eigenvalues don't change under conjugacy, but in the limit the map is diagonal.

The main theorem about a general linear map over \mathbb{C} is the following:

Theorem 5.11 *For any $T \in \text{Hom}_{\mathbb{C}}(V, V)$, we have $V = \bigoplus V_\lambda$, where the sum is over the eigenvalues of V .*

This means T can be put into block diagonal form, where there is one $m(\lambda) \times m(\lambda)$ block for each eigenvalue λ , which is upper triangular and has λ 's on the diagonal.

The characteristic polynomial. This is given by $p(x) = \prod (x - \lambda)^{m(\lambda)}$. Since $T(V_\lambda) = V_\lambda$, we have $(T - \lambda)^{m(\lambda)}|_{V_\lambda} = 0$. This shows:

Corollary 5.12 (Cayley-Hamilton) *A linear transformation satisfies its characteristic polynomial: we have $p(T) = 0$.*

Determinants. Although we have not official introduced determinants, we would be remiss if we did not mention that $p(x) = \det(xI - T)$.

Lemma 5.13 *If T_{ij} is upper-triangular, then the dimension of the generalized kernel of T ($\text{Ker } T^n$, $n = \dim V$) is the same as the number of zero on the diagonal.*

Proof. The proof is by induction on $\dim V = n$. Put T in upper triangular form with diagonal entries $\lambda_1, \dots, \lambda_n$ and invariant subspaces $V_1 \subset \dots \subset V_n =$

V , with induced operators T_1, \dots, T_n with ‘eventual images’ $S_i = \text{Im } T_i^n$ and generalized kernels K_i . We have $\dim K_i + \dim S_i = i$.

If $\lambda_n = 0$ then $T(V_n) \subset V_{n-1}$ so $S_n = S_{n-1}$ and hence the dimension of the generalized kernel increases by 1. While if $\lambda_n \neq 0$ then $S_n \neq S_{n-1}$, so $\dim S_n = 1 + \dim S_{n-1}$ and hence the dimension of the generalized kernel remains the same. ■

Corollary 5.14 *The dimension of V_λ coincides with the number of times λ occurs on the diagonal of T_{ij} in upper triangular form.*

Corollary 5.15 *We have $\sum \dim V_\lambda = \dim V$.*

Corollary 5.16 *We have $V = \bigoplus V_\lambda$.*

Proof. Let S be the span of $\bigcup V_\lambda$. It remains only to show that $S = V$. But since $T(S) \subset S$, if we decompose S into generalized eigenspaces, then we have $S_\lambda = V_\lambda$. Thus $\dim S = \sum \dim S_\lambda = \sum \dim V_\lambda = \dim V$ and so $S = V$. ■

Note that $T|_{V_\lambda} = \lambda I + N$ where N is nilpotent. Since $N^m = 0$ with $m = m(\lambda)$ we have

$$T^k|_{V_\lambda} = \sum_{i=0}^{\min(k,m)} \binom{k}{k-i} \lambda^{k-i} N^i. \quad (5.1)$$

Note that this is the sum of at most m terms, each of which is a polynomial in k times λ^k .

Spectral radius. The following result is often useful in application with a dynamical flavor. Let $\|T\|$ denote any reasonable measurement of the size of T , for example $\sup |T_{ij}|$, such that $\|\lambda T\| = |\lambda| \cdot \|T\|$.

Theorem 5.17 *We have $\lim \|T^n\|^{1/n} = \sup |\lambda|$, where λ ranges over the eigenvalues of T .*

The eigenvalues of T are also known as its *spectrum*, and $\sup |\lambda| = \rho(T)$ as the *spectral radius* of T . This follows easily from (5.1).

Jordan blocks and similarity. A *Jordan block* is an upper triangular matrix such as

$$T = \begin{pmatrix} \lambda & 1 & & & \\ & \lambda & 1 & & \\ & & \lambda & 1 & \\ & & & \lambda & 1 \\ & & & & \lambda \end{pmatrix}$$

(where the missing entries are zero). It has the form $T = \lambda I + N$, where N is an especially simple nilpotent operator: it satisfies $N(e_i) = e_{i-1}$, and $N(e_1) = 0$.

Two operators on V are *similar* if $ST_1S^{-1} = T_2$ for some $S \in \text{GL}(V)$. This means T_1 and T_2 have the same matrix, for suitable choice of bases.

Theorem 5.18 *Every $T \in M_n(\mathbb{C})$ is similar to a direct sum of Jordan blocks. Two operators are similar iff their Jordan blocks are the same.*

Nilpotence. The *exponent* of a nilpotent operator is the least $q \geq 0$ such that $T^q = 0$.

By considering $(T - \lambda)|_{V_\lambda}$, the proof of the existence of Jordan form follows quickly from the nilpotent case, i.e. the case $\lambda = 0$. In this case the Jordan block is a nilpotent matrix, representing the transitions a graph such as the one shown for $q = 5$ in Figure 2.

Theorem 5.19 *Let T be nilpotent with exponent q , and suppose $v \notin \text{Im}(T)$. Then V admits a T -invariant decomposition of the form $V = A \oplus B$, where A is spanned by $(v, Tv, \dots, T^{q-1}v)$.*

Proof. Let $A_0 = A$ and let $A_1 = T(A_0)$. Note that $T|_{\text{Im}(T)}$ has exponent $q - 1$; thus by induction we can have a T -invariant decomposition $\text{Im}(T) = A_1 \oplus B_1$. Let $B_0 = T^{-1}(B_1)$.

We claim that (i) $A_0 + B_0 = V$ and (ii) $A_0 \cap B_1 = (0)$. To see (i) suppose $v \in V$; then $T(v) = a_1 + b_1 \in A_1 \oplus B_1$ and $a_1 = T(a_0)$ for some $a_0 \in A_0$; and since $T(v - a_0) = b_1 \in B_1$, we have $v - a_0 \in B_0$.

To see (ii) just note that $B_1 \subset \text{Im}(T)$, so $A_0 \cap B_1 = A_1 \cap B_1 = (0)$.

Because of (i) and (ii) we can choose B such that $B_1 \subset B \subset B_0$ and $V = A_0 \oplus B$. Then $T(C) \subset T(B_0) \subset B_1 \subset C$ and we are done. ■

Graphs. Suppose G is a directed graph with vertices $1, 2, \dots, n$. Let $T_{ij} = 1$ if i is connected to j , and 0 otherwise. Then $(T^k)_{ij}$ is just the *number of directed paths* of length k that run from i to j .

Theorem 5.20 *The matrix T is nilpotent if and only if T has no cycle.*




Figure 2. A nilpotent graph.

Jordan canonical form: an algorithm? It is tempting to think that a basis to put a matrix into Jordan canonical form can be obtained by first taking a basis for the kernel of T , then by taking preimages of these base elements, etc. This need not work! For example, we might have $\dim \text{Im}(T) = 1$ and $\dim \text{Ker}(T) = 2$; then the original basis for the kernel had better contain an element from the image, otherwise we will be stuck after the first stage, when we are still short of a basis.

The minimal polynomial. This is the unique monic polynomial $q(x)$ of minimal degree such that $q(T) = 0$. For example, the characteristic and minimal polynomials of the identity operator are given by $p(x) = (x - 1)^n$ and $q(x) = (x - 1)$. It is straightforward to verify that

$$q(x) = \prod (x - \lambda)^{M(\lambda)},$$

where $M(\lambda)$ is the exponent of $(T - \lambda)|_{V_\lambda}$. This is the same as the maximal dimension of a Jordan block of T with eigenvalue λ .

Classification over \mathbb{R} : an example.

Theorem 5.21 *Every $T \in \text{SL}_2(\mathbb{R})$ is either elliptic, parabolic, or hyperbolic.*

6 Inner product spaces

In practice there are often special elements in play that make it possible to diagonalize a linear transformation. Here is one of the most basic.

Theorem 6.1 (Spectral theorem) *Let $T \in M_n(\mathbb{R})$ be a symmetric matrix. Then \mathbb{R}^n has a basis of orthogonal eigenvectors for T .*

In particular, T is diagonalizable, all its eigenvalues are real, and $V_\lambda = \text{Ker}(T - \lambda I)$ (there are no nilpotent phenomena).

But what does it mean for a matrix to be symmetric? In fact, a matrix which is symmetric in one basis need not be symmetric in another. For example, if T has distinct eigenvalues (the typical case) then there is a basis where it is symmetric (even diagonal), even if it did not start that way.

Inner products. Let V be a vector space over \mathbb{R} . With just this structure, there is no notion of lengths or angles. This is added by specifying an *inner product* $V \times V \rightarrow \mathbb{R}$, written $\langle x, y \rangle$ or $x \cdot y$.

An inner product is an example of a *bilinear form* $B(x, y)$. This means a map $B : V \times V \rightarrow \mathbb{R}$ such that B is linear in each variable individually. An inner product satisfies the additional requirements that it is *symmetric* and *positive definite*. This means $\langle x, y \rangle = \langle y, x \rangle$, $\langle x, x \rangle \geq 0$ and $\langle x, x \rangle = 0$ iff $x = 0$.

We will *define* the length of a vector in V by $|x|^2 = \langle x, x \rangle$. (Geometrically, it turns out that $\langle x, y \rangle = |x||y| \cos \theta$, where θ is the angle between x and y .)

The main example is of course \mathbb{R}^n with $\langle x, y \rangle = \sum x_i y_i$.

Basic facts:

1. If $\langle x, y \rangle = 0$ then $|x + y|^2 = |x|^2 + |y|^2$ (Pythagorean rule).
2. We have $|\langle x, y \rangle| \leq |x||y|$ (Cauchy-Schwarz inequality).
3. We have $|x + y| \leq |x| + |y|$ (Triangle inequality).

The Pythagorean rule follows from

$$\langle x + y, x + y \rangle = |x|^2 + |y|^2 + 2\langle x, y \rangle,$$

and the same computation shows the triangle inequality follows from Cauchy-Schwarz.

Projections. The proof of Cauchy-Schwarz uses the idea of projections and orthogonality.

For any subspace S , we let

$$S^\perp = \{v \in V : \langle v, s \rangle = 0 \forall s \in S\}.$$

Now, suppose $|x| = 1$, let $S = \mathbb{R}x \subset V$, and define $P : V \rightarrow S$ by $Py = \langle x, y \rangle x$. Then $P^2 = P$, so P is a projection. Quite generally, this

implies $V = \text{Ker } P \oplus \text{Im } P$, as can be seen by writing $y = Py + (I - P)y$. In the case at hand, clearly $\text{Ker } P = S^\perp$. So we have $V = S \oplus S^\perp$.

Consider any y , and write $y = Py + z$ where $Py \in S$ and $z \in S^\perp$. Then we have

$$|y|^2 = |Py|^2 + |z|^2 \geq |Py|^2 = |\langle x, y \rangle|^2,$$

which gives Cauchy-Schwarz. Note that equality in Cauchy-Schwarz iff $y \in \mathbb{R}x$ or $x = 0$.

Theorem 6.2 *For any subspace S , we have $V = S \oplus S^\perp$.*

Proof. Since S^\perp is defined by $\dim S$ linear conditions, we have $\dim S^\perp + \dim S \geq \dim V$. On the other hand, the Pythagorean rule implies the map $S \oplus S^\perp \rightarrow V$ has no kernel. ■

Corollary 6.3 *There is a natural ‘orthogonal projection’ $P : V \rightarrow S$, characterized by $\langle y, s \rangle = \langle Py, s \rangle$ for all $s \in S$. Its kernel is S^\perp .*

Orthonormal bases. Let us say $(e_i)_1^m$ is an *orthonormal basis* for S if $\langle e_i, e_j \rangle = \delta_{ij}$ and the span of (e_i) is S . If S has an orthonormal basis, then we get an explicit formula for P :

$$P(v) = \sum \langle v, e_i \rangle e_i.$$

Theorem 6.4 *Any finite-dimensional inner product space has an orthonormal basis.*

Proof. By induction on $\dim V$. Clear if $\dim V \leq 1$, so assume $\dim V > 1$. Pick any $e_1 \in V$ with $|e_1| = 1$. Then (e_1) is an orthonormal basis for $S = \mathbb{R}e_1$, so the projection to S is defined and $V = S \oplus S^\perp$. By induction, S^\perp has an orthonormal basis (e_2, \dots, e_n) , which together with e_1 gives one for V . ■

Corollary 6.5 *Any finite-dimensional inner product space is isomorphic to \mathbb{R}^n with the usual inner product.*

Sequence spaces. The Cauchy-Schwarz inequality is already a powerful tool for bounded finite sums:

$$\left(\sum a_n b_n\right)^2 \leq \left(\sum a_n^2\right) \left(\sum b_n^2\right).$$

This works for infinite sums as well, and makes $\ell^2(\mathbb{N}) = \{a_n : \sum a_n^2 < \infty\}$ into an inner product space. For example, we have

$$\left|\sum a_n/n\right|^2 \leq \left(\sum a_n^2\right) \left(\sum 1/n^2\right) = (\pi^2/6) \sum a_n^2.$$

Function spaces. Another typical inner-product space is the space of polynomials $\mathbb{R}[x]$ with

$$\langle f, g \rangle = \int_a^b f(x)g(x) dx.$$

Note that the monomials $(1, x, x^2, \dots)$ form a basis for this space but not an *orthonormal* basis. The *Gram-Schmidt* process gives an inductive procedure for generating an orthonormal basis from a basis.

Hermitian inner products. We mention that similar definitions and results hold for vector spaces over \mathbb{C} . In this setting the standard example is $V = \mathbb{C}^n$ with $\langle x, y \rangle = \sum x_i \bar{y}_i$. Note that $\langle y, x \rangle = \overline{\langle x, y \rangle}$ and we only have (complex) linearity in the first coordinate. A positive-definite bilinear form with these properties is called a *Hermitian* inner product.

Example: the space $V = \mathbb{C}[z, z^{-1}]$ of function on the circle $S^1 \subset \mathbb{C}$ that are polynomials in z and $1/z$ forms a Hermitian vector space with respect to the inner product

$$\langle f, g \rangle = \frac{1}{2\pi} \int_{S^1} f(z) \overline{g(z)} |dz|.$$

More on projections. Let $P : V \rightarrow S$ be orthogonal projection to S . We remark that Pv is the unique point in S closest to v . To see this, recall that $t = v - Pv \in S^\perp$. So if $s' \in S$ we have

$$|v - s'|^2 = |s - s' + t|^2 = |s - s'|^2 + |t|^2 \geq |t|^2 = |Pv - v|^2.$$

Signal processing. For example, let $V = P[0, 2\pi]$ be the vector space of piecewise continuous functions $f : [0, 2\pi] \rightarrow \mathbb{R}$ with the inner product $\langle f, g \rangle = \pi^{-1} \int_0^{2\pi} f(x)g(x) dx$. Let $e_0(x) = 2^{-1/2}$, $c_n = \cos(nx)$ and $s_n(x) = \sin(nx)$, and let V_n the span of e_0 and $c_i, s_i, i = 1, \dots, n$. One can show that




Figure 3. Square wave approximated by terms up to $\sin(9x)$.

this is an orthogonal basis, so it can be used to compute the projection g_n of f to V_n . For $f(x) = 1$ on $[0, \pi]$ and 0 on $[\pi, 2\pi]$ we find:

$$g_n(x) = \frac{1}{2} + \sum_{\substack{k=1 \\ k \text{ odd}}}^n \frac{2}{\pi k} \sin kx,$$

where only *odd* values of k are included in the sum.

The dual space. For a more functorial discussion, we introduce the *dual space* $V^* = \text{Hom}(V, \mathbb{R})$. This is a new vector space built naturally from V ; its elements consists of linear maps or *functionals* $\phi : V \rightarrow \mathbb{R}$.

Theorem 6.6 *For all $\phi \in V^*$ there exists a unique $y \in V$ such that $\phi(x) = \langle x, y \rangle$ for all $x \in V$.*

Proof. Choose an orthonormal basis (e_i) and let $y = \sum \phi(e_i)e_i$. ■

Corollary 6.7 *We have $\dim V = \dim V^*$.*

Adjoints. Let T be a linear operator on V . Then for any $y \in V$, the map $x \mapsto \langle Tx, y \rangle$ is a linear functional on V . Thus there is a unique vector z such that $\langle Tx, y \rangle = \langle x, z \rangle$ for all $x \in V$. We write $z = T^*(y)$ and use it to define a new linear operator, the *adjoint* of T . It satisfies

$$\langle Tx, y \rangle = \langle x, T^*y \rangle$$

for all $x, y \in V$.

Theorem 6.8 *Let e_i be an orthonormal basis for V . Then the matrix for T^* with respect to (e_i) is the transpose of the matrix for T .*

Corollary 6.9 *An operator $T \in M_n(\mathbb{R})$ is self-adjoint iff $T_{ij} = T_{ji}$.*

Note that $(ST)^* = T^*S^*$, $T^{**} = T$, and if $T(S) \subset S$ then $T(S^\perp) \subset S^\perp$.

The spectral theorem. We can now prove that a self-adjoint operator is diagonal with respect to a suitable orthonormal basis. In other words, if $T = T^*$ then V has an orthonormal basis of eigenvectors for T . Note that we are working over \mathbb{R} , not \mathbb{C} , and we are getting honest eigenvectors, not generalized eigenspaces.

The proof is based on some preliminary good properties for T . Suppose T is self-adjoint.

Lemma 6.10 *If $T(S) \subset S$ then $T(S^\perp) \subset S^\perp$.*

Lemma 6.11 *We have $\langle v, T^2v \rangle \geq 0$ for all v .*

(In other words, whenever T is self-adjoint (or ‘real’), T^2 is a *positive operator*).

Corollary 6.12 *If $a > 0$ then $T^2 + a$ is invertible. More generally, $T^2 + aT + b$ is invertible if $p(x) = x^2 + ax + b$ has no real root.*

Corollary 6.13 *The operator T has a (real) eigenvector.*

Proof. As usual, by considering $v, Tv, \dots, T^n v$ as find a nontrivial monic polynomial $p \in \mathbb{R}[x]$ such that $\dim \text{Ker } p(T) > 0$. Now over \mathbb{R} , $p(x)$ factors into linear terms $(x - r_i)$ (with real roots) and quadratic terms $q_i(x)$ (with no real roots). By the preceding Corollary, $q_i(T)$ is invertible for each i , so $(T - r_i I)$ must be noninvertible for some i . ■

Proof of the spectral theorem. By induction on $\dim V$. Let $S = \mathbb{R}e_1$ be the span of a unit eigenvector. Then $T(S^\perp) \subset S^\perp$, and $T|_{S^\perp}$ has a basis of eigenvectors by induction. ■

(Note: one should verify that $T^*(S) \subset S$.)

The orthogonal group. The symmetries of V which preserve its inner product form the *orthogonal group* $O(V)$. These are exactly the linear transformations T such that

$$\langle Tx, Ty \rangle = \langle x, y \rangle$$

for all x, y . Equivalently, we have

$$\langle x, T^*Ty \rangle = \langle x, y \rangle.$$

But this happens iff $T^*T = I$. Thus:

$$O(V) = \{T \in \text{Hom}(V, V) : T^*T = I\}.$$

Equivalently, we have $T^{-1} = T^*$. So these transformations are sort of ‘anti-self adjoint’.

The terminology is explained by the following observation: *a matrix T_{ij} belongs to $O(\mathbb{R}^n)$ iff its columns form an orthonormal basis for \mathbb{R}^n .*

Example: the orthogonal group of the plane. A matrix belongs to $O_2(\mathbb{R})$ iff it has the form $\begin{pmatrix} a & -b \\ b & a \end{pmatrix}$ or $\begin{pmatrix} a & b \\ b & -a \end{pmatrix}$, with $a^2 + b^2 = 1$. The rotations of \mathbb{R}^2 are given by $T = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$. The others give reflections.

Theorem 6.14 *For any orthogonal transformation T , there exists an orthogonal splitting $V = \oplus V_i$ where each V_i has dimension 1 or 2. When V_i is two-dimensional, $T|_{V_i}$ is a rotation; when V_i is one-dimensional, $T|_{V_i} = \pm I$.*

Proof. Recall that every real operator has an invariant subspace V_1 of dimension one or two. This is the span of v and Tv where v is in the kernel of $T - a$ or $T^2 - bT - c$. Since T is an isometry, $T(V_1) = V_1$ and $T(V_1^\perp) = V_1^\perp$. The proof then proceeds by induction. ■

Corollary 6.15 *Every isometry of the 2-sphere that preserves its orientation is a rotation about an axis.*

Interaction of orthogonal and self-adjoint transformations. Note that if $R \in O(V)$, then $(RTR^{-1})^* = RT^*R^{-1}$. In particular, the conjugate of a self-adjoint transformation by an orthogonal one is still self-adjoint.

The spectral theorem can be restated as:

Theorem 6.16 *If T is a symmetric matrix, then there is an $R \in O_n(\mathbb{R})$ such that RTR^{-1} is diagonal.*

Polar decomposition. An operator T is *positive* if $T = T^*$ and $\langle Tv, v \rangle \geq 0$ for all v . This is equivalent to its eigenvalues satisfying $\lambda_i \geq 0$. (Q. Is the condition that $T = T^*$ redundant?)

Proposition 6.17 *A positive operator has a unique positive square-root.*

Clearly A^*A is positive for any operator A .

Theorem 6.18 (Polar decomposition) *Any invertible operator A on V can be unique expressed as $A = TR$, where T is positive and $R \in O_n(\mathbb{R})$.*

Proof. Let $T = \sqrt{A^*A}$; then $R = AT^{-1}$ satisfies $RR^* = AT^{-2}A^* = A(A^*A)^{-1}A = I$, so $R \in O_n(\mathbb{R})$. ■

Functional calculus. More generally, if $T = T^*$ and $f : \mathbb{R} \rightarrow \mathbb{R}$ is any *continuous* function, then $f(T)$ makes sense; and if T is positive, then f need only be defined on $[0, \infty)$. This ‘functional calculus’ takes the obvious values on polynomials and satisfies $f_n(T) \rightarrow f(T)$ if $f_n \rightarrow f$; and $f(g(T)) = (f \circ g)(T)$.

Normal operators. Here is a useful extension to a vector space V over \mathbb{C} equipped with a Hermitian form. We say an operator is *normal* if $TT^* = T^*T$.

Theorem 6.19 *A transformation $T \in GL_n(V)$ has an orthonormal basis of eigenvectors iff T is normal.*

Proof. Clearly diagonal matrices are normal. For the converse, let $V_1 = \text{Ker}(T - \lambda_1 I) \subset V$ be a nontrivial eigenspace. We claim T and T^* both preserve the splitting $V = V_1 \oplus V_1^\perp$.

First, we always have $T^*(V_1^\perp) \subset V_1^\perp$ (whether or not T is normal). But we also have $T^*(V_1) \subset V_1$, since for $v \in V_1$ we have

$$TT^*(v_1) = T^*T(v_1) = T^*(\lambda_1 v_1) = \lambda_1 T^*(v_1).$$

This implies $T(V_1^\perp) \subset V_1^\perp$. Then $T|_{V_1^\perp}$ is also normal and the proof proceeds by induction. ■

7 Bilinear forms

Let $B : V \times V \rightarrow \mathbb{R}$ be a bilinear form on a finite-dimensional real vector space.

We say B is *degenerate* if there is an $x \neq 0$ such that $B(x, y) = 0$ for all y ; otherwise it is *nondegenerate*. Any bilinear form is the sum of a symmetric and an antisymmetric one. Our aim is to classify these.

It is often useful to give V an inner product, which we can then compare to B .

Proposition 7.1 *If V has an inner product, then there is a unique operator A such that*

$$B(x, y) = \langle x, Ay \rangle.$$

The form B is symmetric iff A is self-adjoint; antisymmetric iff $A^ = -A$; and B is nondegenerate iff A is invertible.*

Quadratic forms. Suppose B is symmetric. Associated to $B(x, y)$ is the *quadratic form*

$$Q(x) = B(x, x).$$

The form Q determines B , since

$$2B(x, y) = Q(x + y) - Q(x) - Q(y).$$

(Note that for an antisymmetric form, Q would be zero.)

(Aside: this formula is related to the famous equation

$$(x^2 + y^2) - (x^2 - y^2)^2 = (2xy)^2,$$

which allows one to generate all Pythagorean triples. The geometric meaning of this equation is that the rational solutions to $a^2 + b^2 = 1$, projected from $(a, b) = (0, 1)$, are in bijection with the rational numbers on the a -axis.)

A quadratic form on \mathbb{R}^n is nothing more than a homogeneous quadratic polynomial in the coordinates on \mathbb{R}^n :

$$Q(x) = \sum a_{ij} x_i x_j,$$

and the coefficients a_{ij} are unique if we require $a_{ij} = a_{ji}$. If we say $A(x) = \sum a_{ij} x_j$, then

$$B(x, y) = \langle x, Ay \rangle.$$

Examples with mixed signatures. For the usual inner product, $Q(x) = |x|^2$. But there are other important examples! The standard quadratic form of *signature* (p, q) on \mathbb{R}^n , $n \geq p + q$, is given by

$$Q(x) = \sum_1^p x_i^2 - \sum_{p+1}^q x_i^2.$$

This form is nondegenerate iff $p + q = n$.

Note that even when Q is nondegenerate, the ‘lengths’ of vectors can have both signs, or even be zero.

Another important example is provided by $Q(x, y) = xy$ on \mathbb{R}^2 . This is in fact equivalent to the standard form of signature $(1, 1)$. In fact, if we let $e_1 = (1, 1)/\sqrt{2}$ and $e_2 = (1, -1)/\sqrt{2}$, then $B(e_i, e_j) = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$.

Theorem 7.2 *For any real quadratic form on a finite-dimensional vector space V , there exists a basis such that $Q(x)$ is the standard quadratic form of signature (p, q) with $p + q \leq n$.*

Proof. Choose an inner product on V and write the associated bilinear form as $B(x, y) = \langle x, Ay \rangle$ where A is self-adjoint. Then choose an orthonormal basis e_i of eigenvectors for A with eigenvalues λ_i . Replacing e_i with $|\lambda_i|^{-1/2}e_i$, we can assume $B(e_i, e_i) = \pm 1$ or 0. ■

Conics and quadrics. A *real quadric* $X \subset \mathbb{R}^n$ is the locus defined by an equation of the form

$$f(x) = Q(x) + L(x) + C = 0,$$

where $Q(x)$ is a quadratic form, $L(x)$ is linear and C is a constant. We say two quadrics are equivalent if there is an (invertible) affine transformation $T(x) = Ax + b$ such that $T(X_1) = X_2$. A quadric is *degenerate* if Q is zero, or it is equivalent to one where f does not depend on all the coordinates, or where the linear and constant term both vanish.

To classify nondegenerate quadrics, we first arrange that Q has signature (p, q) , $1 \leq p + q \leq n$. Then we can complete the square to eliminate all but $r = n - p - q$ coordinates from $L(x)$. If $r > 0$ then we can translate in these coordinates to absorb the constant. If $r > 1$ then we can get rid of one of the linear coordinates completely, so the quadric is degenerate. And if there is no linear term, we can scale Q so the constant becomes 1. So in the end we find:

Theorem 7.3 Any nondegenerate quadric is equivalent to one given by a purely quadratic equation of the form

$$Q_{pq}(x) = 1,$$

with $p + q = n$, $p \geq 1$; or by a parabolic equation

$$x_n = Q_{pq}(x),$$

where $p \geq q$, $p + q = n - 1$.

Corollary 7.4 Any nondegenerate conic is equivalent to the circle $x^2 + y^2 = 1$, the hyperbola $x^2 - y^2 = 1$, or the parabola $y = x^2$.

Theorem 7.5 Any nondegenerate quadric in \mathbb{R}^3 is equivalent to either:

1. The sphere $x^2 + y^2 + z^2 = 1$; or
2. The 1-sheeted hyperboloid $x^2 + y^2 - z^2 = 1$; or
3. The 2-sheeted hyperboloid $x^2 - y^2 - z^2 = 1$; or
4. The elliptic paraboloid $z = x^2 + y^2$; or
5. The hyperbolic paraboloid $z = x^2 - y^2$.

Alternating forms. We now handle the alternating case.

Suppose $B(x, y) = -B(y, x)$. The main example of such a form is

$$B(x, y) = \det \begin{pmatrix} x_1 & x_2 \\ y_1 & y_2 \end{pmatrix} = x_1y_2 - x_2y_1.$$

This form measures the signed area of the parallelogram with sides x and y . By taking the sum of n such forms, we obtain a nondegenerate alternating form on \mathbb{R}^{2n} . It is given by $B(x, y) = \langle x, J_{2n}y \rangle$ where J_{2n} has n blocks of the form $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ on the diagonal.

A nondegenerate alternating form is also called a *symplectic form*

Theorem 7.6 Let $B(x, y)$ be a symplectic form on finite-dimensional real vector space V . Then $\dim V = 2n$ is even, and there exists a basis such that $B(e_i, e_j) = (J_{2n})_{ij}$.

Note. The prototypical example of such a form arises on the homology of a closed orientable surface of genus g .

Proof. Choose any $e_1 \neq 0$. Then there exists an e_2 such that $B(e_1, e_2) = 1$ (else B is degenerate). Let V_1 be the span of e_1 and e_2 ; this is 2-dimensional, else we would have $B(e_1, e_2) = 0$. The matrix of $B|_{V_1}$ is now J_2 .

Let $V_1^\perp = \{v \in V : B(v, e_1) = B(v, e_2) = 0\}$. It is easy to see that $V = V_1 \oplus V_1^\perp$. First, these two subspaces meet only at 0 since $B|_{V_1}$ is non-degenerate. And second, V_1^\perp is the kernel of a map to \mathbb{R}^2 , so its dimension is (at least) $n - 2$.

It follows that $B|_{V_1^\perp}$ is non-degenerate, and we proceed by induction. ■

Change of basis. The matrix of $T : V \rightarrow V$ relative to a given basis e_i is determined by the condition:

$$T(e_j) = \sum_i T_{ij} e_i.$$

What is the matrix of T in another basis (f_i) ? To find this we form the matrix S satisfying $S(e_j) = f_j$, so the columns of S give the new basis.

Proposition 7.7 *The matrix of T relative to the new basis (f_i) is the same as the matrix of $U = S^{-1}TS$ relative to the old basis.*

Proof. We have $U(e_j) = \sum U_{ij} e_i$ and $SU = TS$, so

$$T(f_j) = TS(e_j) = SU(e_j) = \sum U_{ij} S(e_i) = \sum U_{ij} f_j. \quad \blacksquare$$

On the other hand, the matrix of a bilinear form B relative to (e_i) is given by

$$B_{ij} = B(e_i, e_j).$$

It satisfies $B(v, w) = \sum v_i B_{ij} w_j$.

Proposition 7.8 *The matrix of B relative to the new basis (f_i) is given by $S^t B S$.*

Proof. We have

$$B(f_i, f_j) = B(S(e_i), S(e_j)) = \sum_{k,l} S_{ki} B_{kl} S_{lj} = \sum S_{ik}^t B_{kl} S_{lj} = (S^t B S)_{ij}. \quad \blacksquare$$

Note that these operations are very different: if $S = \lambda I$, $\lambda \neq 0$, then $S^{-1}TS = T$ but $S^tBS = |\lambda|^2B$. They agree if S is orthogonal (or unitary), i.e. if $SS^t = I$.

8 Trace and determinant

The trace and determinant can be approached in many different ways: using bases, using eigenvalues, and using alternating forms. It is useful to understand all approaches and their interrelations.

Trace. Let V be a finite-dimensional space over K , and let $T : V \rightarrow V$ be a linear map. The trace of T is defined by choosing a basis and then setting $\text{tr}(T) = \sum T_{ii}$. We then find:

$$\text{tr}(ST) = \sum_{i,j} S_{ij}T_{ji} = \text{tr}(TS).$$

This shows $\text{tr}(S^{-1}TS) = \text{tr}(T)$ so in fact the trace doesn't depend on the choice of basis.

Now suppose $K = \mathbb{R}$ or \mathbb{C} . By putting T into upper-triangular form, we see:

$$\text{tr}(T) = \sum m(\lambda)\lambda.$$

This gives another, basis-free definition of the trace. And finally we see that characteristic polynomial satisfies

$$p(t) = \prod (t - \lambda)^{m(\lambda)} = t^n - \text{Tr}(T)t^{n-1} + \dots$$

This shows that any convenient definition of $p(t)$ also gives a definition of the trace.

Theorem 8.1 *If $T : V \rightarrow V$ is linear, then $T^* : V^* \rightarrow V^*$ satisfies $\text{tr}(T) = \text{tr}(T^*)$.*

Determinant. We begin with the case $V = \mathbb{R}^n$. We can associate to a matrix T its column vectors $(v_1, \dots, v_n) = (T(e_1), \dots, T(e_n))$. We then wish to define a quantity which measures the volume of the parallelepiped with edges v_1, \dots, v_n . We will denote this quantity by $\det(v_1, \dots, v_n)$ and *define*

it by induction on the dimension using expansion by minors in the first row. That is, writing $v_i = (a_i, w_i)$, we set

$$\det(v_1, \dots, v_n) = a_1 \det(w_2, \dots, w_n) - a_2 \det(w_1, w_3, \dots, w_n) + \dots \pm a_n \det(w_1, a_2, \dots, a_{n-1}).$$

By induction on n , we find this function has three key properties:

1. It is linear in each v_i .
2. If $v_i = v_{i+1}$ then the determinant is zero.
3. It is normalized so that $\det(e_1, \dots, e_n) = 1$. (Volume of a cube).

From these properties one easily deduces the behavior of the determinant under *elementary operations*:

1. The determinant is unchanged if a multiple of one vector is added to another,
2. The determinant changes sign when 2 adjacent entries are reversed;
3. The determinant scales by a if one of the v_i is scaled by a .

Theorem 8.2 *There is a unique function $\det(v_1, \dots, v_n)$ with these properties.*

Proof. Observe that if the (v_i) form a basis, we can repeatedly apply the operations above to convert (v_i) to (e_i) , and hence compute \det . Similarly if there is a linear relation, we can apply these operations to get $v_i = 0$ for some i , and then by linearity $\det(v_i) = 0$. Thus the properties above uniquely determine $\det(v_i)$. ■

Elementary matrices. Let E^{ij} denote the matrix with 1 in position (i, j) and 0 elsewhere. The *elementary matrices* are those of the form

1. $E = I + aE^{ij}$, $i \neq j$; which satisfies $E(e_j) = e_j + ae_i$;
2. $I + E^{ij} + E^{ji} - E^{ii} - E^{jj}$; which satisfies $E(e_i) = e_j$ and vice-versa; and
3. $I + (a - 1)E^{ii}$; which satisfies $E(e_i) = ae_i$.

In each case the vectors $(w_i) = TE(e_i)$ are obtained from $(v_i) = T(e_i)$ by performing an elementary operation. The argument just given shows:

Theorem 8.3 *The elementary matrices generate $\mathrm{GL}_n(\mathbb{R})$.*

We have also seen that the elementary operation changes the determinant in a predictable way: $\det(TE) = \det(T)$, $-\det(T)$ or $a\det(T)$. But $\det(E)$ is also easily computed, allowing us to verify:

Theorem 8.4 *We have $\det(TE) = \det(T)\det(E)$ for any elementary matrix.*

Corollary 8.5 *We have $\det(ST) = \det(S)\det(T)$ for all $S, T \in \mathrm{M}_n(\mathbb{R})$.*

(Here we use the fact that $\det(T) = 0$ if T is not invertible.)

Eigenvalues and determinants. Since $\det(S^{-1}TS) = \det(S)$, we also find

$$\det(T) = \prod \lambda^{m(\lambda)}$$

and hence $p(T) = T^n + \dots + (-1)^n \det(T)$.

Rotations versus reflections. Recall that $T \in \mathrm{O}_3(\mathbb{R})$ is either a rotation about an axis, or a reflection, since it must have a real eigenvector. We can distinguish these 2 cases by determinant: $\det(T) = 1$ for rotations and $\det(T) = -1$ for reflections.

The group of orientation-preserving orthogonal transformations, in any dimension, is denoted $\mathrm{SO}_n(\mathbb{R})$.

Alternating forms. Here is a more functorial approach to the determinant. Recall that V^* denote the space of linear forms on V , i.e. linear maps $\phi : V \rightarrow K$.

Let $V^* \otimes V^*$ denote the space of *bilinear* forms. The notation is meant to suggest that $B(v_1, v_2) = \phi(v_1)\psi(v_2)$ is always such a form, so there is a natural map $V^* \oplus V^* \rightarrow V^* \otimes V^*$. A bilinear form is determined by its values $B_{ij} = B(e_i, e_j)$ on pairs of basis elements.

Similarly let $\otimes^k V^*$ denotes the space of k -linear forms. Such a form is determined by a ‘matrix’ with k indices, and hence $\dim \otimes^k V^* = n^k$.

A k -form is *alternating* if it vanishes whenever two entries are equal. In characteristic zero, this is equivalent to saying that ϕ changes sign whenever any two adjacent entries are interchanged. We let

$$\wedge^k V^* \subset \otimes^k V^*$$

denote the vector space of alternating k -linear forms $\phi(v_1, \dots, v_k)$ on V .

For any linear map $T : V \rightarrow W$ we get a natural map $T^* : W^* \rightarrow V^*$. (How is this related to the adjoint in an inner product space?) Similarly, by ‘pullback’, we get a natural map

$$\wedge^k T^* : \wedge^k W^* \rightarrow \wedge^k V^*,$$

defined by

$$\wedge^k T^*(\phi) = \phi(Tv_1, \dots, Tv_k).$$

Sign of a permutation. A *permutation* $\sigma \in S_n$ is a bijection from $\{1, \dots, n\}$ to itself. Every permutation is a product of transpositions, indeed, adjacent transpositions would suffice.

The *sign* function $\text{sign} : S_n \rightarrow (\pm 1)$ is the function that is 1 if σ is a product of an even number of transpositions, and -1 if it is a product of an odd number of transpositions. The fact that this function is well-defined can be seen as follows. Consider the polynomial

$$P(t_1, \dots, t_n) = \prod_{i < j} (t_i - t_j).$$

Given σ we can form a new polynomial

$$Q = Q(t_1, \dots, t_n) = P(t_{\sigma 1}, \dots, t_{\sigma n}) = \pm P,$$

and we set $\text{sign}(\sigma) = P/Q$. Evidently $\text{sign}(\sigma\sigma') = \text{sign}(\sigma)\text{sign}(\sigma')$ and $\text{sign}(\tau) = -1$ for a transposition, so sign gives the parity. In particular this definition of π shows that no permutation is both even and odd.

Theorem 8.6 *For any $\phi \in \wedge^k V^*$, and any $\sigma \in S_k$, we have*

$$\phi(v_{\sigma 1}, \dots, v_{\sigma k}) = \text{sign}(\sigma)\phi(v_1, \dots, v_k).$$

Theorem 8.7 *If $n = \dim V$, the space $\wedge^n V^*$ is 1-dimensional.*

Proof. Any ϕ is determined by its value on (e_1, \dots, e_n) , and there is a unique ϕ satisfying $\phi(e_{\sigma 1}, \dots, e_{\sigma n}) = \text{sign}(\sigma)$. ■

To compute the action of T on $\wedge^n V^*$ it then suffices to compute the single number

$$\phi(Te_1, \dots, Te_n) / \phi(e_1, \dots, e_n).$$

Using the multilinear and alternating properties of ϕ it follows that this number obeys the same axioms as the determinant, which shows:

Theorem 8.8 *We have $\wedge^n T^*(\phi) = \det(T)\phi$ for all $\phi \in \wedge^n V^*$.*

Corollary 8.9 *We have $\det T = \sum_{S_n} \text{sign}(\sigma) \prod_{i=1}^n T_{i,\sigma i}$.*

Theorem 8.10 *The characteristic polynomial is given by $p(t) = \det(tI - T)$.*

These result can also be read backwards, as alternative, basis-free/eigenvector-free definitions of $\det(T)$ and of the characteristic polynomial. For example we immediately see that the constant term in $p(t)$ is given by $p(0) = \det(-T) = (-1)^n \det(T)$.

Theorem 8.11 *We have $\det(T) = \det(T^*)$.*

Proof. This is clear in the one-dimensional case; now use the fact that $\wedge^n T^* = (\wedge^n T)^*$. ■

Theorem 8.12 *The coefficients of the characteristic polynomial*

$$p(t) = \det(tI - T) = t^n + a_1 t^{n-1} + \dots + a_n$$

are given by

$$a_i = (-1)^i \text{Tr } \wedge^i T^*.$$

Cramer's rule. Let M'_{ij} be the $(n-1) \times (n-1)$ (minor) matrix obtained by deleting the i th row and j th column from T_{ij} , and let

$$T'_{ij} = (-1)^{i+j} \det M_{ij}.$$

Then $\sum T_{i1} T'_{i1} = \det(T)$. Similarly $\sum T_{i1} T'_{ij} = 0$ for $j \neq 1$, since this gives an expansion for the determinant of a matrix with two equal columns. The same reasoning shows

$$\sum_i T_{ij} T'_{ik} = (\det T) \delta_{ik},$$

or equivalently $T^{-1} = (T')^t$.

Corollary 8.13 *A matrix in $M_n(K)$ is invertible iff its determinant is nonzero, in which case the inverse is also in $M_n(K)$.*

If $A \subset K$ is a subring (e.g. $\mathbb{Z} \subset \mathbb{R}$), $T \in M_n(A)$ and $\det(T)^{-1} \in A$, then $T^{-1} \in M_n(A)$.

Intrinsically, T' gives the matrix of $\wedge^{n-1}T^*$ acting on $\wedge^{n-1}V^*$.

Tensor products. In general, the *tensor product* $V \otimes W$ of two vector spaces is a new vector space with a bilinear map $V \times W \rightarrow V \otimes W$, written $(v, w) \mapsto v \otimes w$, such that for any other bilinear map $\phi : V \times W \rightarrow Z$ there is a *unique* compatible linear map $\Phi : V \otimes W \rightarrow Z$.

If (v_i) and (w_j) are bases for V and W then $(v_i \otimes w_j)$ is a basis for $V \otimes W$, and $\Phi(\sum a_{ij}v_i \otimes w_j) = \sum a_{ij}\phi(v_i, w_j)$. In particular, we have

$$\dim V \otimes W = (\dim V)(\dim W).$$

Theorem 8.14 *There is a natural isomorphism between $\text{Hom}(V, W)$ and $V^* \otimes W$.*

Proof. Define a bilinear map $V^* \times W \rightarrow \text{Hom}(V, W)$ by sending (ϕ, w) to the map $T(v) = \phi(v)w$, then take its canonical extension to $V^* \otimes W$. ■

Functoriality. If $T : V \rightarrow W$ then we have a natural map $T^* : V^* \rightarrow W^*$. If $S : V \rightarrow V'$ and $T : W \rightarrow W'$, then we get a map $S \otimes T : V \otimes W \rightarrow V' \otimes W'$ satisfying

$$(S \otimes T)(v \otimes w) = S(v) \otimes T(w).$$

Reciprocal polynomials. Let us say the *reciprocal* of a monic, degree n polynomial $p(t) = t^n + a_1t^{n-1} + \cdots + a_n$ is the monic polynomial $p^*(t)$ whose roots are the reciprocals of the roots of $p(t)$. It is given explicitly by $p^*(t) = t^n p(1/t)/a_n$. We say p is a *reciprocal polynomial* if $p(t) = p^*(t)$.

Example: if $p(t) = \det(tI - T)$ is the characteristic polynomial of T , and T is invertible, then $p^*(t)$ is the characteristic polynomial of T^{-1} .

Theorem 8.15 *If $T \in \text{Hom}(V, V)$ preserves a nondegenerate quadratic form B — that is, if $T \in \text{O}(V, B)$ — then its characteristic polynomial satisfies $p^*(t) = p(t)$.*

Proof. To say V preserves B is to say that the isomorphism $B : V \rightarrow V^*$ conjugates T to T^* , reversing the direction of the arrow; i.e. T^* is conjugate to T^{-1} . Since T and T^* have the same characteristic polynomial, the result follows. In terms of equations, we have $T^*BT = B$ and so $BTB^{-1} = (T^*)^{-1}$. ■

Corollary 8.16 *A rotation of \mathbb{R}^{2n+1} preserves an axis.*

Proof. It must have an eigenvalue which satisfies $\lambda = 1/\lambda$, so $\lambda = \pm 1$. ■

Diagonalization. One might hope that any $T \in O(V, B)$ is diagonalizable, at least over \mathbb{C} . This is in fact true for symmetric forms, using the fact that if S is an eigenspace then $V = S \oplus S^\perp$. But it is false for antisymmetric forms; for example $T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ preserves area on \mathbb{R}^2 , so it lies in $\text{Sp}_2(\mathbb{R})$, but it is not diagonalizable.

9 Introduction to Group Theory

A group is a G set with an associative operation $G \times G \rightarrow G$ that has a (two-sided) identity and inverses.

A *subgroup* $H \subset G$ is a subset that is closed under multiplication (i.e. $HH = H$) and, with respect to this operation, is a group in its own right.

The *intersection* $\bigcap H_i$ of any collection of subgroups is a subgroup. (Compare subspaces.) For example, if $S \subset G$ is a finite set, the subgroup $\langle S \rangle$ generated by S is the intersection of all subgroups containing S . (It is also the space of words in S and S^{-1} .)

Maps. A homomorphism of groups is just a map $f : G \rightarrow H$ such that $f(xy) = f(x)f(y)$. Its kernel and image are subgroups of G and H respectively. If f is a bijection, then its inverse is a homomorphism, and we say G and H are *isomorphic*.

Examples. The integers \mathbb{Z} and the integers mod n , \mathbb{Z}/n . For any field (e.g. \mathbb{R} , \mathbb{Q} , \mathbb{F}_p , \mathbb{C}) we have (abelian) subgroups K and K^* . The group $(\mathbb{Z}/n)^*$.

Linear groups: $\text{GL}_n(K)$; $\text{O}_n(\mathbb{R})$; $\text{SO}_n(\mathbb{R})$; $\text{U}_n(\mathbb{C})$; $\text{SU}_n(\mathbb{C})$; $\text{Sp}_{2g}(\mathbb{R})$; $\text{SO}(p, q)$; $\text{GL}_n(\mathbb{Z})$; $\text{O}_n(\mathbb{Z})$.

The dihedral group. What is $\text{O}_2(\mathbb{Z})$? It is the same as the symmetries of a diamond or square, allowing flips. We can record its 8 elements as a subgroup




Figure 4. The dihedral group D_4 .

of S_4 . It is generated by $r = (1234)$ and $f = (24)$. We have $rf = f^3r$. More generally, $D_{2n} = \langle r, f : r^n = f^2 = e, rf = fr^{-1} \rangle$.

The symmetries of a cube. What is $G = \text{SO}_3(\mathbb{Z})$? It is the orientation-preserving symmetric group of a cube or octahedron. Its order is 24. In fact G is isomorphic to S_4 . (Consider the pairs of antipodal vertices).

The quaternion group. This is a nonabelian group of order 8: $\langle i, j : i^2 = j^2 = (ij)^2 = -e \rangle$.

Finite fields. What is $\text{GL}_2(\mathbb{F}_2)$? It permutes the nonzero vectors in \mathbb{F}_2^2 in any way; so G is isomorphic to S_3 .

The symmetric group. The symmetric group S_n . We write (a_1, \dots, a_r) for the cyclic permutation with $\sigma(a_i) = a_{i+1}$. We have $|S_n| = n!$

The (adjacent) transpositions generate S_n . Any permutation $\sigma \in S_n$ can be written uniquely as a product of disjoint cycles. The parity of a k -cycle is $k + 1 \pmod 2$. Thus it is easy to calculate the parity of an given element:

$$\text{sign}(\sigma) = (\text{number of cycles}) + \sum (\text{lengths of each}) \pmod 2.$$

Proof that this is correct: if we multiply a product of cycles by a transposition (ij) , we either meld or break apart the cycle(s) containing i and j .

Examples of subgroups. $\text{GL}_2(\mathbb{R})$ contains $\text{SL}_2(\mathbb{R})$ which in turn contains AN , A and N (isomorphic to \mathbb{R} and \mathbb{R}^*). $\text{GL}_2(\mathbb{R})$ also contains the $f(x) = ax + b$ group as the matrices $\begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix}$ acting on the line $y = 1$.

Subgroups of \mathbb{Z} . Every subgroup of \mathbb{Z} has the form $H = a\mathbb{Z}$, $a \geq 0$. We have $a\mathbb{Z} \subset b\mathbb{Z}$ iff $b|a$. Thus the lattice of subgroups is the same as the lattice \mathbb{N} under divisibility.

It follows that

$$a\mathbb{Z} + b\mathbb{Z} = \text{gcd}(a, b)\mathbb{Z}.$$

In particular this shows $\gcd(a, b) = ar + bs$. Similarly $a\mathbb{Z} \cap b\mathbb{Z} = \text{lcm}(a, b)\mathbb{Z}$. Thus the additive theory of \mathbb{Z} knows the multiplicative theory.

Cyclic subgroups. Given an $a \in G$ there is a map $\phi : \mathbb{Z} \rightarrow G$ given by $\phi(n) = a^n$. Its kernel is $b\mathbb{Z}$ where b is the order of a (or zero).

Isomorphisms. Example: $\phi : \mathbb{R} \rightarrow \text{SO}_2(\mathbb{R})$ given by $\phi(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$ is an isomorphism (verification includes the addition laws for sine and cosine).

For any group G , $\text{Aut}(G)$ is also a group. Does \mathbb{Z} know the difference between ± 1 ? Between 1 and 2? What about $\mathbb{Z}/7$? To answer these questions, one can first show $\text{Aut}(\mathbb{Z}) \cong \mathbb{Z}/2$; and $\text{Aut}(\mathbb{Z}/n) \cong (\mathbb{Z}/n)^*$.

We also have $\text{Aut}(\mathbb{Z}^n) \cong \text{GL}_n(\mathbb{Z})$. This shows, for example, that \mathbb{Z}^2 has an automorphism of order 6.

Inner automorphisms. There is always a map $G \rightarrow \text{Aut}(G)$ sending g to $\phi_g(x) = gxg^{-1}$. The placement of g^{-1} is important; check: $\phi_{gh}(x) = g(hxh^{-1})g^{-1} = \phi_g(\phi_h(x))$. These are called *inner automorphisms* of G .

The center. The natural map $G \rightarrow \text{Aut}(G)$ has kernel $Z(G)$, the *center* of G . Example: $Z(\text{GL}_n(\mathbb{R})) \cong \mathbb{R}^*I$; $Z(S_n)$ is trivial for $n \geq 3$. (Proof: use the fact that $\text{Fix}(ghg^{-1}) = g(\text{Fix}(h))$ and the fact that there are permutations with single fixed-points.)

Two elements of G are *conjugate* if $gx_1g^{-1} = x_2$ for some $g \in G$. Compare similarity: two matrices in $\text{GL}_n(\mathbb{R})$ are similar iff they are conjugate. Thus conjugate elements are the same up to a ‘change of basis’. For example, x_1 and x_2 have the same order.

Example. In D_8 (or D_{4n}), the diagonal flips are not conjugate to the vertical and horizontal flips. That is, f and rf are not conjugate.

Homomorphisms $\phi : G \rightarrow H$. Image and kernel are subgroups; $\phi(e) = e$, $\phi(x)^{-1} = \phi(x^{-1})$.

Examples:

1. $\det : \text{GL}_n(\mathbb{R}) \rightarrow \mathbb{R}^*$, kernel is $\text{SL}_n(\mathbb{R})$.
2. $\text{sign} : S_n \rightarrow \mathbb{Z}/2$, kernel is A_n .
3. $\mathbb{Z} \rightarrow \mathbb{Z}/n$, kernel is $n\mathbb{Z}$.

Normal subgroups. Problem: is there a homomorphism $\phi : \text{SL}_2(\mathbb{R}) \rightarrow H$ whose kernel is the group A of diagonal matrices? No. The kernel is always a *normal subgroup*. It is a *union of conjugacy classes*. A *group homomorphism can only kill properties that are conjugacy invariant*.

Example: the center $Z(G)$ is always normal.

Cosets. Let H be a subgroup of G . The sets $aH \subset G$ are called the (left) *cosets* of G . Any two cosets have the same cardinality, since we can multiply by a^{-1} to get back to H . If aH and bH both contain $x = ah_1 = bh_2$ then $a = bh_2h_1^{-1}$ so $a \in bH$ and $b \in aH$; but $aHb = aH$ so $aH = bH$. Thus any 2 cosets are equal or disjoint. In fact the cosets are the equivalence classes of the relation $a \equiv b$ iff $a^{-1}b \in H$. In particular we have:

Theorem 9.1 (Lagrange) *The order $|H|$ of any subgroup of G divides $|G|$.*

Corollary 9.2 *The only group of prime order p is \mathbb{Z}/p .*

Corollary 9.3 *For any $\phi : G \rightarrow H$ we have $|G| = |\text{Ker } \phi| \cdot |\text{Im } \phi|$.*

Compare this to the dimension addition formula for linear maps. Over finite fields, they are the same theorem.

We can similarly form the right cosets $\{Ha : a \in G\} = H \backslash G$.

Theorem 9.4 *A subgroup H is normal iff $aH = Ha$ for all $a \in G$.*

Internal and external products. For any pair of groups we can form the (external) product $A \times B$. In this product, elements of A commute with elements of B . In particular, A and B are normal subgroups of their product.

We can also make an internal version.

Theorem 9.5 *If A and B are normal subgroups of G meeting only at the identity, and $AB = G$, then $G \cong A \times B$.*

Proof. We claim elements of A and B commute. To see this we consider the commutator $[a, b] = aba^{-1}b^{-1}$. Parenthesizing one way, we see this lies in A ; the other way, we see it lies in B ; hence it is the identity.

Define a map $\phi : A \times B \rightarrow G$ by forming the product of two elements. By commutativity this map is a homomorphism, and our assumptions imply it is injective and surjective. ■

Compare this to the decomposition of a vector space as $V = A \oplus B$. Note that it is enough to require that A and B generate G , since the proof shows AB is a subgroup.

Example: $\mathbb{Z}/10 \cong \mathbb{Z}/2 \times \mathbb{Z}/5$.

Quotient groups.

Theorem 9.6 *If $N \subset G$ is a normal subgroup, then G/N is itself a group with the group law $(aN) * (bN) = (aN)(bN) = abN$, and the map $\phi(a) = aN$ is a homomorphism with kernel N .*

This leads to the ‘first isomorphism theorem’:

Theorem 9.7 *If $\phi : G \rightarrow H$ is a surjective homomorphism, then H is naturally isomorphic to $G/\text{Ker } \phi$.*

Example: let $\phi(z) = \exp(z) : \mathbb{C} \rightarrow \mathbb{C}^*$. Then $\mathbb{C} \cong \mathbb{C}^*/(2\pi i\mathbb{Z})$.

10 Symmetry

Group actions. Group theory emerged as an abstraction of the notion of symmetry. This notation is made systematic by the idea of a group acting on a set. Such an action is given by a map $G \times S \rightarrow S$, written $g(s)$ or gs or $g \cdot s$, satisfying $(gh)(s) = g(h(s))$ and $e(s) = s$. Equivalently, a group action is given by a homomorphism $G \rightarrow \text{Sym}(S)$.

The *orbit* of $s \in S$ is the set Gs .

The *stabilizer* $\text{Stab}(s) = G^s = \{g \in G : gs = s\}$ is a subgroup of G .

If s and t are in the same orbit, then G^s is conjugate to G^t . The following result is useful for counting orbits.

Theorem 10.1 *We have $|G| = |Gs| \cdot |\text{Stab}(s)|$.*

We also define, for any set $A \subset S$, $\text{Stab}(A) = \{g \in G : g(A) = A\}$.

Example: binomial coefficients. The group S_n acts transitively on $\mathcal{P}_k(n)$, the k -element subsets of n ; and the stabilizer of a given set is $S_k \times S_{n-k}$. Thus $|\mathcal{P}_k| = n!/(k!(n-k)!)$.

How many conjugates does $\sigma = (12 \cdots k)$ have in S_n ? Its centralizer is $\mathbb{Z}/k \times S_{n-k}$, so the number is $n!/(k(n-k)!) = (k-1)! \binom{n}{k}$. You have to

choose the points in the cycle, and then give them an order, but k different orders give the same cycle.

Euclidean motions. As our first example we let $G = \text{Isom}(\mathbb{R}^n)$. An isometry fixing the origin is a linear map (check this!) and hence an orthogonal transformation. Thus any $g \in G$ has the form

$$g(x) = T(x) + a$$

where $T \in O_n(\mathbb{R})$ and $a \in \mathbb{R}^n$. The group G can be represented by matrices of the form $\begin{pmatrix} T & a \\ 0 & 1 \end{pmatrix} \in \text{GL}_{n+1}(\mathbb{R})$. Alternatively, we can let (T, a) represent g . Then we have the group law

$$(S, b)(T, a)x = S(Tx + a) + b = STx + Sa + b = (ST, Sa + b).$$

This presents G as the *semidirect product* of $O_n(\mathbb{R})$ and \mathbb{R}^n .

Motions of the plane. Now let us focus on the case of \mathbb{R}^2 . In addition to rotations and translations, we have reflections in a line L and *glide reflections*: reflection in L followed by translation along L . It is crucial to notice that these classes are *conjugacy invariant*.

Theorem 10.2 *Every isometry of \mathbb{R}^2 is of one of these 4 types.*

Proof. Let $g(x) = Tx + a$. The 4 classes are distinguished by (i) whether or not $\det(T) = 1$ and (ii) whether or not g has a fixed point. When g has a fixed point the classification is clear. Otherwise, $(T - I)x = -a$ has no solution. If $\det T = 1$ this implies T is a translation.

Now assume $\det(T) = -1$. Then $x \mapsto Tx$ is reflection, say through a line L ; and g is the composition of a reflection and a translation. It is useful to write $a = b + c$ where b is orthogonal to L and c is parallel to L . Then $Tx + b$ is a reflection through $L' = L + b/2$, and thus g is the composition of reflection through L' with translation by c along L' . This is an ordinary reflection if $c = 0$ and a glide reflection if $c \neq 0$. ■

Theorem 10.3 *Every finite subgroup of $O_2(\mathbb{R})$ is cyclic or dihedral.*

The same is true for $\text{Isom}(\mathbb{R}^2)$.

Proof. For the first statement, look at the smallest rotation in the group. For the second, observe that the barycenter $\beta(S) = |S|^{-1} \sum_S s$ satisfies $\beta(gS) = \beta(S)$ for any $g \in \text{Isom}(\mathbb{R}^n)$; so for any x , $p = \beta(Gx)$ is a fixed-point of G . ■

Finiteness and orthogonality. By the same token, if $G \subset \text{GL}_n(\mathbb{R})$ is a finite group, we can sum over G to get an inner product B which is invariant under G . Since all inner products are equivalent, this shows:

Theorem 10.4 *Any finite subgroup of $\text{GL}_n(\mathbb{R})$ is conjugate to a subgroup of $\text{O}_n(\mathbb{R})$. Any finite subgroup of $\text{GL}_n(\mathbb{C})$ is conjugate to a subgroup of $\text{U}_n(\mathbb{C})$.*

Recalling that unitary transformations are diagonalizable, this implies:

Corollary 10.5 *Any finite order automorphism of a finite-dimensional vector space over \mathbb{C} is diagonalizable.*

Discrete groups. Next we analyze discrete subgroups $G \subset \text{Isom}(\mathbb{R}^2)$.

Associated to any group G of isometries of \mathbb{R}^2 , we have a translational subgroup $L \subset G$ and a rotational quotient group $DG \subset \text{O}_2(\mathbb{R})$.

Theorem 10.6 *If $L \subset \mathbb{R}^n$ is discrete, then it is generated by a set of linearly independent vectors, and L is isomorphic to \mathbb{Z}^k for $1 \leq k \leq n$. Conversely, any set of linearly independent vectors generates a discrete subgroup of \mathbb{R}^n .*

Proof for $n = 2$. Consider the shortest vector, and normalize it to be e_1 ; then consider the lowest vector $v = (x, y)$ with $y > 0$ and $|x| \leq 1/2$; and then observe that these span. ■

Theorem 10.7 *There are 2, 4 or 6 shortest vectors in a lattice $L \subset \mathbb{R}^2$.*

Proof. The lowest vector (x, y) must satisfy $1/4 + y^2 \geq |v|^2 \geq 1$ and hence $y \geq \sqrt{3}/4 > 1/2$. This shows any ties for the shortest vector must lie at the same height as e_1 or $\pm v$. Then use the fact that $|x| \leq 1/2$. ■

Corollary 10.8 *The symmetry group of a lattice is typically $\mathbb{Z}/2$, but it can also be D_4 (rectangular or checkerboard case), D_6 (hexagonal case) or D_8 (square case).*

Theorem 10.9 *We have $g(L) = L$ for all $g \in DG$. Thus the rotational part DG of an infinite discrete subgroup of $\text{Isom}(\mathbb{R}^2)$ has order 1, 2, 3, 4 or 6, and DG is cyclic or dihedral.*

Proof. If $Dg = A$ then $g = SA$ where S is a translation; consequently for any $T_a(x) = x + a \in L$ we have

$$gkg^{-1} = SAT_aA^{-1}S^{-1} = ST_{A(a)}S^{-1} = T_{A(a)}.$$

For the second part, let P the convex hull of the shortest vectors in L . Then P is either an interval, a rectangle/square or a regular hexagon. Then DG must be contained in the symmetry group $\text{Isom}(P)$, which is D_4 , D_8 or D_{12} . ■

Example. You cannot order wallpaper with order 5 symmetry. (See however Penrose tilings!)

Constraints on lattices. Every lattice L has a rotational symmetry of order 2, but only the square and hexagonal lattices have symmetries of order 3, 4 or 6.

A lattice with a reflection symmetry must be rectangular *or* rhombic (the center of each rectangle is added).

The 17 types of wallpaper groups. A discrete subgroup $G \subset \text{Isom}(\mathbb{R}^2)$ whose translational part L is a lattice is called a *wallpaper group*. There are 17 types of these. They can be classified according to their ‘point groups’.

The group G is generated by L and by the lifts of the one or two generators of $DG \cong D_{2n}$. If a given generator is a rotation we can always normalize so it fixes the original. But if it is a reflection, its lift might be a reflection or a glide reflection (by an element of $L/2$). And when DG contains both a rotation and a reflection, they might or might not have a common fixed point. These consideration make the classification more intricate.

1. (1) $DG = \langle 1 \rangle$. Then $G = L$.
2. (1) $DG = \mathbb{Z}/2$ acting as a rotation. Any L has such an extension. The points of $L/2$ are the poles of rotation. There are 4 types of poles.
3. (3) $DG = \mathbb{Z}/2$ acting as a reflection. In this case, L must be invariant by a reflection, say $(x, y) \mapsto (x, -y)$. Then L might be rectangular or rhomboidal, e.g. L might be \mathbb{Z}^2 or $D_2 = \{(a, b) \in \mathbb{Z}^2 : a = b \pmod{2}\}$. (Put differently, there are 2 reflections in $\text{GL}_2(\mathbb{Z})$, one of which fixes the line $|z| = 1$ and one of which fixes $\text{Re } z = 0$.) And a generate of DG might lift to a reflection or to a glide reflection.

This gives 3 possibilities for G . (i) rectangular/reflection: G has 2 conjugacy classes of reflections. (ii) rectangular/glide reflection: G has 2 types of glide reflections. (iii) rhomboidal, reflection: then we get a reflection and a glide reflection.

4. (3) $DG = \mathbb{Z}/3, \mathbb{Z}/4, \mathbb{Z}/6$. There are unique lattices that fit with these (rotational) symmetries.
5. (4) $DG = D_4 \cong \mathbb{Z}/2 \times \mathbb{Z}/2$. We can lift so the order 2 rotation fixes 0. Then the points of L are poles of rotation. The lattice can be either rectangular or rhomboidal. The reflections may or may not pass through the pole. This gives 4 possibilities.
6. (2) $DG = D_6$. The reflection group of the (60,60,60) triangle; and this example with face-center rotations adjoined.
7. (2) $DG = D_8$. The reflection group of the (45,45,90) triangle and this example with long-edge center rotations adjoined.
8. (1) $DG = D_{12}$. This is the reflection group of the (30, 60, 90) triangle.

The wallpaper groups can also be classified topologically by forming the orbifold $X = \mathbb{R}^2/G$. The underlying space can be:

1. (4) The sphere; we get the (2, 2, 2, 2), (2, 4, 4), (3, 3, 3) and (2, 3, 6) orbifolds.
2. (8) The disk with mirror boundary. Then its double is a sphere as above. We write $(a_i|b_i)$ for a_i -type points in the interior and b_i on the boundary. The possible signatures are:

$$\begin{aligned}
 &(|2, 2, 2, 2); (2|2, 2); (2, 2|) \\
 &(|2, 4, 4); (4|2) \\
 &(|3, 3, 3); (3|3) \\
 &(|2, 3, 6).
 \end{aligned}$$

3. (1) The projective plane with signature (2, 2).
4. (4) The torus, Klein bottle, cylinder and Möbius band.

Platonic solids. The symmetry group of a polyhedron $S \subset \mathbb{R}^3$ is its stabilizer in $\text{Isom}^+(\mathbb{R}^3)$. (Note that we have required determinant one, so the symmetry can be seen by a rigid motion of the polyhedron.)

We say S is a *Platonic solid* if its symmetry group acts transitively on its vertices, edges and faces. As a consequence, the number of vertices, faces and edges must divide the order of the symmetry group.

There are exactly 5 Platonic solids: the tetrahedron, the cube, the octahedron, the dodecahedron and the icosahedron.

1. *The tetrahedron.* The symmetry group of a tetrahedron is A_4 ; it can be described as the orientation-preserving permutations of the vertices.
2. *The cube.* The symmetry group of a cube has 24 elements, since there are 6 faces each with stabilizer of order 4.

In fact G is isomorphic to S_4 , acting on the long diagonals! To see this, note that a rotation fixing a face gives the permutation $\sigma = (1234)$, and a rotation fixing an edge gives the permutation (12) . These two elements together generate S_4 .

3. The cube is dual to the octahedron.
4. *The dodecahedron.* How large is the symmetry group of a dodecahedron? A face has stabilizer of order 5, and there are 12 faces, so $|G| = t \times 12 = 60$. Similarly there are 30 edges (since each has stabilizer 2) and 20 vertices (since 5 faces come together at each).

It turns out we have $G \cong A_5$. To see this, one can find 5 *cubes* whose vertices lie on the vertices of a dodecahedron. There are 20 vertices all together, and each belongs to two cubes — which works out, since 5 cubes have $5 \cdot 8 = 40$ vertices all together.

It is important to note that not every symmetry of an inscribed cube extends to a symmetry of the dodecahedron. In fact we have $S_4 \cap A_5 = A_4$ under the embedding. A given cube has ‘clockwise’ and ‘counterclockwise’ vertices: see Figure 5.

Put differently, every cube contains a unique left-handed tetrahedron and a unique right-handed tetrahedron. Thus the 10 inscribed tetrahedrons fall into 2 types, again giving the isomorphism $G \cong A_5$.

5. The dodecahedron is dual to the icosahedron.




Figure 5. A cube inscribed in a dodecahedron has 2 types of vertices.

A non-Platonic solid. The rhombic dodecahedron is *not* a Platonic solid. All its 12 faces are equivalent, and their stabilizer is of order 2, so $|G| = 24$. There are 14 vertices, but they are *not* all equivalent! In fact they fall into two classes of sizes $6 + 8 = 14$, and each of those divides 24.

	G	$ G $	V	E	F	$V - E + F$
Tetrahedron	A_4	12	4	6	4	2
Cube	S_4	24	8	12	6	2
Octahedron	S_4	24	6	12	8	2
Dodecahedron	A_5	60	20	30	12	2
Icosahedron	A_5	60	12	30	20	2
Rhombic Dodecahedron	S_4	24	$6+8=14$	12	12	2

Kepler's Cosmology. Kepler believed that the orbits of the planets were determined by the Platonic solids. Each eccentric orbit determines a thickened sphere or orb, centered at the Sun, that just encloses it. The 5 Platonic solids thus correspond exactly to the gaps between the 6 planets known at that time. Between the orbits of Saturn and Jupiter you can just fit a cube; between Jupiter and Mars, a tetrahedron; between Mars and Earth, a dodecahedron; between Earth and Venus an icosahedron, and between Venus and Mercury, an octahedron.

This theory is discussed in the *Mysterium Cosmographicum*, 1596. Using the astronomical data of Copernicus, Kepler found:

	Predicted	Actual
Jupiter/Saturn	577	635
Mars/Jupiter	333	333
Earth/Mars	795	757
Venus/Earth	795	794
Mercury/Venus	577	723

See ‘Kepler’s Geometrical Cosmology’, J. V. Field, University of Chicago Press, 1988.

Finite subgroups of $\mathrm{SO}_3(\mathbb{R})$. A modern, and more abstract version of the classification of Platonic solids is the classification of finite groups $G \subset \mathrm{SO}_3(\mathbb{R})$.

Theorem 10.10 *Any finite subgroup of $\mathrm{SO}_3(\mathbb{R})$ is abstractly isomorphic to \mathbb{Z}/n , D_{2n} , A_4 , S_4 or A_5 . These subgroups are rigid: any 2 finite subgroups of $\mathrm{SO}_3(\mathbb{R})$ which are isomorphic abstractly are conjugate in $\mathrm{SO}_3(\mathbb{R})$.*

Compare the case of lattices $\mathbb{Z}^2 \subset \mathrm{Isom}(\mathbb{R}^2)$: these are *not* rigid, but they become rigid when more rotational symmetries are added.

Here is part of the proof.

Theorem 10.11 *If $G \subset \mathrm{SO}_3(\mathbb{R})$ is a finite group, then G is \mathbb{Z}/n , D_{2n} or a group of order 12, 24 or 60.*

Proof. Let $P \subset S^2$ be the set of poles of G , i.e. the endpoints of the axes of nontrivial rotations $G' \subset G$. Clearly G acts on P , say with orbits P_1, \dots, P_m . Let N_i be the order of the stabilizer of a point in P_i . Now every $p \in P_i$ is the pole of $N_i - 1$ elements in G' , and every element of G' has 2 poles. Thus if we count up the number S of pairs (g, p) in $G' \times P$ such that p is a pole of g , we get

$$S = 2|G'| = 2|G|(1 - 1/|G|) = \sum |P_i|(N_i - 1) = |G| \sum (1 - 1/N_i).$$

In other words, we have

$$\sum (1 - 1/N_i) = 2 - 2/N < 2,$$

where $N = |G|$.

Now each term on the left is at least $1/2$, so the number of terms is 1, 2 or 3. And in fact (unless G is trivial) there must be at least 2 terms, because $1 - 1/N_1 < 1$ and $2 - 2/N \geq 1$.

If there are exactly 2 terms we get $2/N = 1/N_1 + 1/N_2$, and N_i divide N . The only possibility is $N_1 = N_2 = N$ and then G is cyclic.

If there are exactly 3 terms then we get $1 + 2/N = 1/N_1 + 1/N_2 + 1/N_3$. The possible solutions are $(N_1, N_2, N_3) = (2, 2, N/2), (2, 3, 3), (2, 3, 4)$ and $(2, 3, 5)$. In the first case G is dihedral, and in the last 3 cases we get $N = 12, 24$ or 60 . ■

Euler characteristic. From a topological point of view, what's going is that $\chi(S^2/G) = 2/|G| = 2 - \sum(1 - 1/N_i)$.

Abstract group actions. Any transitive action of G on a set S is isomorphic to the action by left translation of G on G/H , where $H = \text{Stab}(x)$.

This identification *depends* on the choice of x . The stabilizer of $y = gx$ is the conjugate gHg^{-1} of H ; this observation applies in particular to the stabilizer of gH .

The stabilizers of the points of S are the conjugates of H .

Example: D_3 acts on the 3 vertices of a triangle with $H = \langle f \rangle$. The other vertices correspond to $\langle rf \rangle$ and $\langle r^2f \rangle$, the conjugates of H .

Note however that H may fix many points in S , so there is not a bijection between the points of S and the conjugates of H in general.

The class equation and its consequences. The fact that G acts on G by translation (on the right or the left) is not very useful, since there is only one orbit.

But the fact that G also acts by *conjugation* is very useful. In this case the orbit of $x \in G$ is its *conjugacy class* $C(x)$, and its *stabilizer* is its *centralizer*. We have $|C(x)| \cdot |Z(x)| = |G|$.

Now recall that the partition of G into orbits of G yields the *class equation*

$$|G| = \sum C_i,$$

where C_i are the cardinalities of the conjugacy classes in G . Each C_i divides $|G|$.

We also know that we can take $C_1 = 1$, i.e. the conjugacy class of the identity is trivial. More generally, the set of elements with $|C(x)| = 1$ forms the center $Z(G)$.

Examples: the symmetries G of the cube are comprised of the identity (1), elements of order 4 (6), elements of order 3 (8), and elements of order 2 flipping edges (6), and elements of order 2 spinning faces (3). We have

$$|G| = 24 = 1 + 6 + 8 + 6 + 3.$$

Note that an element of order 4 is determined by the choice of a face, which is then rotated $1/4$ of a turn. The opposite face is rotated $3/4$.

For the dodecahedron, the conjugacy class is almost determined by the order; we have elements of orders 1, 2, 3 and 5, but the last fall into 2 types. The class equation is

$$|G| = 60 = 1 + 15 + 20 + 12 + 12.$$

Note that an element of order 5 is determined by a face plus a choice of whether to rotate $1/5$ or $2/5$ of a turn.

Note that $|C_i|$ must divide the order of G . Indeed, $|C_i| = |Z(g)|$ for any $g \in C_i$.

Simple groups. The class equation also helps us to recognize *simple groups*, i.e. groups with no normal subgroups. A general finite group is built up out of finite simple groups, so their classification is particularly important. By considering the conjugacy classes of G contained in N , we find:

Theorem 10.12 *Let $|G| = \sum C_i$ be the class equation. If G has a proper normal subgroup, then there is a proper sub-sum containing $C_1 = 1$ such that $\sum C'_i$ divides $|G|$.*

Corollary 10.13 *The icosahedral group A_5 has no normal subgroup.*

Proof. Recall its class equation is $60 = 1 + 15 + 20 + 12 + 12$. There is no sub-sum of the required type. ■

Riddle. What is the next term in the sequence

$$2, 3, 5, 7, 11, 13, 17, 19, 23, 29, 31, 37, 41, 43, 47, 53, 59?$$

(Ans: 60.)

Normal subgroups. By the same token, the class equation $24 = 1 + 6 + 8 + 6 + 3$ for S_4 allows one to show: S_4 contains a unique nontrivial normal subgroup, namely the group of order 4 generated by 180 degree face rotations.

Burnside counting theorem. We now return to the setting of G acting on a set S . If the action of G on S is not necessarily transitive, then S breaks up into orbits. We write the space of orbits as S/G , although $G \backslash S$ would be more accurate. Let $S^g = \{x \in S : gx = x\}$.

We then have the following result of Burnside:

Theorem 10.14 *The number of orbits $|S/G|$ is given by $|G|^{-1} \sum_G |S^g|$.*

Equivalently, the number of orbits is the same as the average number of fixed points.

Proof. We consider the function on $G \times S$ which is 1 if $gx = x$ and otherwise 0. Summing one way gives $\sum_G |S^g|$; the other way gives $\sum_S |\text{Stab}(x)|$. Now $|\text{Stab}(x)|$ is constant on each orbit, so the second sum can be rewritten as

$$\sum_S |\text{Stab}(x)| = \sum_{S/G} |\text{Stab}(x)| |Gx| = \sum_{S/G} |G| = |S/G| |G|.$$

■

What makes this formula useful is that $S^g = S^h$ whenever g and h are conjugate. So we can organize the sum over G as a sum over conjugacy classes.

Example: coloring the cube. (The Heinz problem.) How many rotationally distinct ways are there to color the faces of a cube with 3 colors? Let S be the set of all colorings and let $G \cong S_4$ be the symmetries of the cube. Then $|S^g| = 3^n$ where n is the number of orbits of faces under g .

There are 5 cases to consider. The unique element of order 1 contributes 3^6 . The 6 elements of order 4 each contribute 3^3 . The 8 elements of order 3 each contribute 3^2 . The 6 edge-type elements of order 2 each contribute 3^3 . And the 3 face-type elements of order 2 each contribute 3^4 as well. So we get:

$$|S/G| = (1/24)(1 \cdot 3^6 + 6 \cdot 3^3 + 8 \cdot 3^2 + 6 \cdot 3^3 + 3 \cdot 3^4) = 1368/24 = 57.$$

Semidirect products (This is a supplement to Artin, which does not treat this topic.)

Here is a very important generalization of the notion of a product of groups. It uses the idea that both G and $\text{Aut}(G)$ are subgroups of the permutation group $\text{Sym}(G)$.

Let G and H be groups, and let $\phi : H \rightarrow \text{Aut}(G)$ be a homomorphism. With ϕ understood, it is convenient to use the notation

$$g^h = \phi(h)(g);$$

the fact that $\phi(h)$ is an automorphism of G implies $(g_1g_2)^h = g_1^hg_2^h$.

We can now form a new group $P = G \rtimes_{\phi} H$, the *semidirect product* of G and H . (Often the ϕ is understood and suppressed.)

As a set, P consists of all pairs (g, h) . We interpret such a pair as the product gh ; thus any element of P can be uniquely expressed as such a product.

What about products in the other order? The crucial relation is that we define:

$$hg = g^hh.$$

Using this relation, any product of elements of P can be rewritten uniquely in the form gh again. For example:

$$g_1h_1g_2h_2 = g_1g_2^{h_1}h_1h_2.$$

An equivalent and more formal definition is that P consists of the ordered pairs (g, h) with the group law $(g_1, h_1)(g_2, h_2) = (g_1g_2^{h_1}, h_1h_2)$.

We can also think of P as the set of ‘affine transformations’ $T : G \rightarrow G$ of the form $T(x) = gh(x)$. Then

$$g_1h_1(g_2h_2(x)) = g_1g_2^{h_1}h_1h_2(x).$$

Note that G is always a *normal* subgroup of $P = G \rtimes H$, and $P/G \cong H$. As a universal example, one can always take $G \rtimes \text{Aut}(G)$.

Recognizing semidirect products. The basic result is:

Theorem 10.15 *Let P be a group containing subgroups G and H such that:*

1. G is normal,
2. $GH = \{gh : g \in G, h \in H\} = P$, and
3. $G \cap H = \{e\}$.

Then P is isomorphic to $G \rtimes_{\phi} H$, where $\phi(h)(g) = hgh^{-1} \in G$.

Remark. Since G is normal, GH is always a *subgroup* of S . This often helps in verifying that $GH = S$.

Splitting exact sequences. We say a surjective homomorphism $\pi : A \rightarrow B$ *splits* if it has a ‘section’ $\xi : B \rightarrow A$, i.e. a homomorphism satisfying $\pi\xi(b) = b$ for all $b \in B$.

Theorem 10.16 *Let N be a normal subgroup of G . Then G is the semidirect product of N and G/N iff the exact sequence*

$$0 \rightarrow N \rightarrow G \rightarrow G/N \rightarrow 0$$

splits.

Examples.

1. The group \mathbb{Z}/n admits an automorphism of order 2 given by $\alpha(k) = -k$. Thus there is a map $\phi : \mathbb{Z}/2 \rightarrow \text{Aut}(\mathbb{Z}/n)$. The resulting semidirect product is the usual dihedral group:

$$D_n = \langle r, \rho : r^2 = \rho^n = 1, r\rho r^{-1} = \rho^{-1} \rangle.$$

2. If ϕ is trivial, so $g^h = g$, then the semidirect product becomes the ordinary product of groups $G \times H$.
3. The group of motions of the plane is given by $M = \mathbb{R}^2 \rtimes \text{O}_2(\mathbb{R})$.
4. The orthogonal group itself is the semidirect product $\text{O}_2 = \text{SO}_2(\mathbb{R}) \rtimes \mathbb{Z}/2$, since we have $r\rho_\theta r^{-1} = \rho_{-\theta}$.
5. The group $AN \subset \text{SL}_2(\mathbb{R})$ is the semidirect product of A and N .
6. The group $\mathbb{Z}/7$ admits an automorphism of the form $\alpha(k) = 2k$. We have $\alpha(\alpha(\alpha(k))) = 8k = k \pmod{7}$, so α has order 3 in $\text{Aut}(\mathbb{Z}/7)$. Thus there is a *nonabelian* group of order 21 given by $S = \mathbb{Z}/7 \rtimes_\phi \mathbb{Z}/3$.
7. The sequence $0 \rightarrow \mathbb{Z} \xrightarrow{2} \mathbb{Z} \rightarrow \mathbb{Z}/2$ does not split.

We recall from Theorem 9.5 that as a special case: if G is generated by a pair of normal subgroups A and B meeting only at the identity, then G is isomorphic to $A \times B$.

11 Finite group theory

Classification of finitely-generated abelian groups. We begin with the classification of finite abelian groups. The finitely-generated case is no more difficult. (Note that infinitely-generated abelian groups like \mathbb{Q} can have a complicated structure.)

Theorem 11.1 *Any finitely-generated abelian group is isomorphic to a product of cyclic groups.*

Proof. Consider a presentation matrix $T : \mathbb{Z}^r \rightarrow \mathbb{Z}^g$ such that $G \cong \mathbb{Z}^g/T(\mathbb{Z}^r)$, and then apply row and column operations to diagonalize it.

In more detail, first arrange that T_{11} is the smallest nonzero element of T . Then make the other entries in the first row and column of T equal to zero. If a remainder results, the size of the smallest element in T has decreased: start over. When done, proceed by induction of the rank of T . ■

Note that $\mathbb{Z}/a \times \mathbb{Z}/b \cong \mathbb{Z}/ab$ whenever $\gcd(a, b) = 1$. Because of this, the factorization of G into cyclic groups is not quite unique.

There are 2 standard forms for the torsion part $\mathbb{Z}/a_1 \times \dots \times \mathbb{Z}/a_n$ of the product. In one form, $a_1|a_2|\dots|a_n$. These integers are called the *elementary divisors* of G . In the other form, we have $a_i = p_i^{e_i}$ where p_i are prime. The same prime can appear multiple times.

In both cases, the numbers (a_i) are invariants of G . For example, in the elementary divisor case, a_n is the largest order of an element in G .

See Artin, Chapter 12.6 for a treatment of this classification in the context of modules.

Classification of small groups. We will next explore some methods for classifying groups of ‘small order’. We already understand the groups whose order is prime. The first ‘open’ case is $|G| = 4$.

Lemma 11.2 *Any group G in which every element has order 1 or 2 is abelian. In fact, G is a vector space over \mathbb{F}_2 .*

For $|G| = 4$ we either have a cyclic group $\mathbb{Z}/4$ or, by the preceding Lemma, the field $\mathbb{Z}/2 \times \mathbb{Z}/2$.

A similar argument shows for $|G| = 6$ there are 2 just groups, $\mathbb{Z}/6$ and S_3 . E.g. we know G must have an element of order 3 or 6, else it would be abelian.

The classification of groups of order 8 is left as an exercise. We will next show any group of order 9 (or order p^2) is abelian. (There are 2 such groups for each prime.)

Nilpotent groups. The class of nilpotent groups is the smallest class of groups containing the abelian groups and with the property that G is nilpotent if $G/Z(G)$ is nilpotent. For example, S_3 is *not* nilpotent (its center is trivial) but D_8 is nilpotent (the ‘antipodal map’ r^2 lies in the center and yields as quotient $\mathbb{Z}/2 \times \mathbb{Z}/2$). In the same way the quaternion group Q is nilpotent.

(The terminology is explained by examples of linear groups where every element has the form $I + N$, N a nilpotent matrix).

p -groups. If $|G| = p^e$ is a power of a prime, G is a p -group. The class equation enables one to analyze p -groups rather well:

Theorem 11.3 *Every p -group is nilpotent.*

Proof. The center of G must be nontrivial, else the class equation would give $|G| = 1 \pmod{p}$. And then $G/Z(G)$ is again a p -group. ■

Theorem 11.4 *Every group of order p^2 is abelian.*

As we have seen, this implies G is isomorphic to $\mathbb{Z}/p \times \mathbb{Z}/p$ or \mathbb{Z}/p^2 .

Proof. If $|Z(G)| = p^2$ we are done, so we may assume $|Z(G)| = p$. Now we use a basic fact: for any $x \in G$, x and $Z(G)$ together generated an abelian subgroup. So taking any $x \notin Z(G)$, we get an abelian subgroup with $|A| > p$, and hence $|A| = p^2 = |G|$. ■

Note that D_8 and Q are nonabelian groups of order p^3 . The upper-triangular matrices in $\text{GL}_n(\mathbb{F}_p)$ with 1’s on the diagonal also give nice examples of nonabelian p -groups.

The Sylow theorems. Recall that a p -group is nilpotent. Remarkably, every finite group contains a p -group that is as large as possible, and this p -subgroup is unique up to conjugacy.

Theorem 11.5 (First Sylow) *Let $|G| = p^e m$ where p is prime and $\gcd(p, m) = 1$. Then G contains a subgroup H with $|H| = p^e$.*

Such an H is called a *p-Sylow subgroup* of G .

Example. The ‘prime’ example is the general linear group $G = \text{GL}_n(\mathbb{F}_p)$. By considering bases it is easy to see that

$$|G| = (p^n - 1)(p^n - p) \cdots (p^n - p^{n-1}) = p^e m,$$

where $e = 0 + 1 + 2 + \cdots + (n - 1)$. A p -Sylow subgroup of G is given by the group of upper-triangular matrices with 1’s along the diagonal.

Theorem 11.6 (Second Sylow) *Any p -subgroup is contained in a p -Sylow subgroup, and all p -Sylow subgroups are conjugate.*

Theorem 11.7 (Third Sylow) *Suppose $|G| = p^e m$, $\gcd(p, m) = 1$. Then the number of p -Sylow subgroups in G divides m , and is congruent to 1 mod p .*

Proof of Sylow I. Let $\mathcal{S} = \mathcal{P}_{p^e}(G)$ denote the set of all $S \subset G$ such that $|S| = p^e$. Then

$$|\mathcal{S}| = \binom{p^e m}{p^e} = \frac{(mp^e)(mp^e - 1) \cdots ((m - 1)p^e + 1)}{p^e(p^e - 1) \cdots 1}$$

is not divisible by p , because corresponding factors in the numerator and denominator are congruent modulo p^e . Now let G act on \mathcal{S} by $g(S) = g \cdot S$ (*not* by conjugation!). Since \mathcal{S} is partitioned into orbits of G , some orbit has size $|G \cdot S| = |G| / \text{Fix}(S)$ which is not divisible by p . It follows that $p^e \mid |H|$, where $H = \text{Fix}(S)$.

On the other hand, by decomposing S into orbits of the *pointwise* action of H , we see $|H|$ divides $|S| = p^e$. It follows that $|H| = p^e$. ■

The next proof exploits the following general observation, already used in the proof that p -groups are nilpotent. Namely if a p -group G acts on a set S , and S^G denotes the set of points fixed by every element of G , then we have

$$|S^G| \equiv |S| \pmod{p}.$$

Indeed, we can decompose S into G -orbits, and any orbit which is not a single point has size divisible by p . The orbits which are single points, on the other hand, are exactly S^G .

Proof of Sylow II. Let P be a Sylow subgroup of G , and let H be a p -subgroup. We will show H is contained in some conjugate of P . This is the same as showing that H fixes some point when it acts on G/P . But $|G/P|$ is not divisible by p , so this follows from the basic lemma on orbits. ■

There are 2 ideas in the final proof. The first is that, by II, the set of all p -Sylow subgroups is in bijection with $N_G(P)/P$ — this gives $n_p|m$. The second is that P is the *unique* p -Sylow subgroup of $N_G(P)$ — since all p -Sylow subgroups of $N_G(P)$ are conjugate.

Proof of Sylow III. The calculation of the number of Sylow subgroups is the same as the calculation of the number of conjugates of P .

Let $N_G(P)$ denote the normalizer of P , i.e. the elements of G that conjugate P to itself. (E.g. $N_G(P) = G$ iff P is a normal subgroup.) It is a general fact, true for any subgroup, that the number of conjugates is just $|G|/|N_G(P)|$.

In the case at hand, $|P| = p^e$ divides $|N_G(P)|$, and thus the number of conjugates of P divides m .

For the last part, we consider the action of P by conjugation on the set S of all p -Sylow subgroups. We will show that $S^P = \{P\}$. This shows by the lemma that $|S| = 1 \pmod{p}$.

To finish the proof suppose Q is a p -Sylow subgroup and $xQx^{-1} = Q$ for all $x \in P$. Then P is contained in $N_G(Q)$. But then Q and P are both subgroups of $N_G(Q)$, so P and Q are conjugate by Sylow II. But Q is normal in $N_G(Q)$, so this means $P = Q$. ■

Groups of order 15. A well-known application of the Sylow theory is to show that the only group of order 15 is $\mathbb{Z}/3 \times \mathbb{Z}/5$. For this we let n_p denote the number of p -Sylow subgroups. We have $n_3 = 1(3)$ and $n_3|5$ so $n_3 = 1$. Similarly $n_5 = 1(5)$ and $n_5|3$ so $n_5 = 1$. Thus we have normal cycle subgroups A and B of size 3 and 5, which can only meet in the identity. This implies (as we have seen) that $G \cong A \times B$.

More generally we have:

Theorem 11.8 *Let $|G| = pq$ be a product of distinct primes with $p < q$. Then either G is abelian — and hence isomorphic to $\mathbb{Z}/p \times \mathbb{Z}/q$ — or $p|q-1$ and G is the unique nonabelian group of this order.*

Proof. This is a mild generalization of the case $pq = 15$. We have $n_q = 1$ or p , but also $n_q = 1(q)$, so $n_q = 1$. Thus there exists a normal subgroup $N \cong \mathbb{Z}/q$. Similarly $n_p = 1$ or q . If $n_p = 1$ then we have a product group as before.

If $n_p = q$ then $q = 1(p)$, i.e. $p|q-1$. In this case we still have an exact sequence

$$0 \rightarrow \mathbb{Z}/q \rightarrow G \rightarrow \mathbb{Z}/p \rightarrow 0.$$

Since G has an element of order p , this sequence splits and we have a semidirect product, easily seen to be unique up to isomorphism. ■

12 Representation theory

In this section we discuss the classification of representations of finite groups using characters.

Basic definitions. Let G be a finite group and let V be a finite-dimensional complex vector space. A *representation* of G is a homomorphism

$$\rho : G \rightarrow \mathrm{GL}(V).$$

We write $\rho(g)(v) = gv$ with ρ understood. Of course representations of larger groups, on infinite-dimensional spaces, are also interesting, but for the purposes of this section we restrict to the case above.

Two representations are *equivalent* (consider the same) if there is a linear isomorphism $T : V \rightarrow V'$ such that $T(gv) = g'T(v)$.

The regular representation. The symmetric group S_n has a natural action on \mathbb{C}^n by permuting the basis vectors. Since any finite G embeds as a subgroup of S_n , $n = |G|$, we also get a faithful linear representation of G in this way.

To make this representation more functorial, let $L^2(G)$ denote the space of all functions $\phi : G \rightarrow \mathbb{C}$ with the inner product

$$\langle \phi(g), \psi(g) \rangle = |G|^{-1} \sum_G \phi(g) \overline{\psi(g)}.$$

The action of G on $V = L^2(G)$ is called the *regular representation* of G . Formally it is given by

$$(g \cdot \phi)(x) = \phi(g^{-1}x),$$

or

$$\left(g \cdot \sum a_h h \right) = \sum a_h gh = \sum a_{g^{-1}h} h.$$

This representation is *unitary*.

The group ring. We can extend the definition above to make $L^2(G)$ into a *ring* $\mathbb{C}[G]$. Any representation of G extends to a homomorphism of rings

$\rho : \mathbb{C}[G] \rightarrow \text{End}(V)$. Conversely, a representation of G is the same as a module over the group ring.

Direct sums; reducibility. There is a natural notion of direct sum of representations of G : $V = V_1 \oplus V_2$. More generally we can write $V = \sum m_i V_i$ for integers $m_i > 0$.

We say V is *irreducible* if it has no nontrivial invariant subspace. We say V is *semisimple* or *completely reducible* if it is a direct sum of irreducible representations. By putting together isomorphic representations, this means we can write $V = \oplus m_i V_i$ as a sum of inequivalent representations with multiplicities. (These generalize the multiplicities of eigenvalues.)

A representation is *unitary* if there is a (Hermitian) inner product $\langle v, w \rangle$ on V preserved by ρ ; i.e. satisfying $\langle gv, gw \rangle = \langle v, w \rangle$.

Theorem 12.1 *Any representation of a finite group preserves an inner product on V . Thus we can think of the target of ρ as the unitary group U_n .*

Corollary 12.2 *Every representation of a finite group is a direct sum of irreducible representations.*

A related result: every element $g \in \text{GL}_n(\mathbb{R})$ of finite order is conjugate into $O_n(\mathbb{R})$.

Characters. The *character* of (ρ, V) is the function

$$\chi(g) = \text{tr } \rho(g).$$

If V is irreducible, we say χ is an *irreducible character*.

Evidently isomorphic representations have the same character. One of the main points of representation theory is the converse. Thus characters give a concrete way to record and distinguish representations.

Here are some general facts about characters.

1. We have $\chi(e) = \dim V$.
2. We have $\chi(g^{-1}) = \overline{\chi(g)}$.
3. Similarly $\overline{\chi(g)}$ is the character of the adjoint representation, $g \mapsto \rho(g^{-1})^*$. (Using duality or an invariant Hermitian form.)
4. If $\dim V = 1$, the character is a homomorphism: we have $\chi(gh) = \chi(g)\chi(h)$.

5. In general, χ is a *class function*: it satisfies

$$\chi(ghg^{-1}) = \chi(h),$$

so it is constant on conjugacy classes.

6. The character of $V_1 \oplus V_2$ is given by $\chi_1 + \chi_2$.

Examples of characters.

1. For the regular representation, $\chi(e) = |G|$ and $\chi(g) = 0$ for $g \neq e$.
2. For the action of S_n on \mathbb{C}^n , $\chi(g) = |\text{Fix}(g)|$.
3. For the action of \mathbb{Z}/n by rotations on $\mathbb{R}^2 \subset \mathbb{C}^2$, with k rotating by k/n , we have

$$\chi(k) = 2 \cos(2\pi k/n).$$

It is useful to note that the trace is 1, 0 and -1 for rotations by 60° , 90° and 120° ; and also that any reflection on \mathbb{R}^2 has $\chi(g) = 0$.

4. For the action A_4 on \mathbb{R}^3 by the symmetries of the cube we have:

A_4	(1)	(3)	(4)	(4)	
	e	f	r	r^2	
χ	3	-1	0	0	

Note that we have given one entry for each conjugacy class, and recorded the order of the conjugacy class along the top.

Duality for abelian groups. Now suppose G is a finite abelian group.

Theorem 12.3 *Every irreducible representation of G is 1-dimensional.*

Proof. It suffices to show that every $g \in G$ acts by a multiple of the identity. For this consider any $g \in G$ and, using the fact that $g^n = e$, decompose V into eigenspaces $\oplus V_\lambda$ of g . These spaces must be preserved by every $h \in G$ because of commutativity. Thus g has only one eigenvalue, so it acts by a multiple of the identity. ■

The dual group. We define the *dual* of G by

$$\widehat{G} = \text{Hom}(G, S^1).$$

Note that this is indeed a group, and in fact $\widehat{\widehat{G}} \cong G$, as can be checked with cyclic groups. There is no *canonical* isomorphism between G and \widehat{G} ; however, there is a natural isomorphism

$$G \cong \widehat{\widehat{G}},$$

sending g to the 1-dimensional character of \widehat{G} defined by $\alpha_g(\chi) = \chi(g)$.

Theorem 12.4 *The set $\widehat{G} \subset L^2(G)$ coincides with the characters of 1-dimensional representations of G .*

Theorem 12.5 *If $\chi : G \rightarrow S^1$ is a homomorphism of a finite Abelian group, then either $\chi(g) = 1$ for all g , or $S = \sum_G \chi(g) = 0$.*

Proof. In the second case, pick an $h \in G$ such that $\chi(h) \neq 1$, and observe that $\chi(h)S = S$. ■

Corollary 12.6 (Orthogonality relations) *The irreducible characters of a finite Abelian G form an orthonormal basis for $L^2(G)$.*

Proof. For any $\alpha, \beta \in \widehat{G}$ we have $\langle \alpha, \beta \rangle = (1/|G|) \sum_G (\alpha/\beta)(g)$, which gives orthonormality. Thus we have an injective map $L^2(\widehat{G}) \rightarrow L^2(G)$. To see it is surjective, one can use the fact that $|G| = |\widehat{G}|$. For a better proof, decompose $L^2(G)$ into 1-dimensional irreducible subspaces $\oplus \mathbb{C}\chi_i$, and observe that if we normalize so $\chi_i(e) = 1$, then it is a character. ■

Theorem 12.7 *If G is an abelian group, then the regular representation decomposes as a direct sum of $|G|$ distinct 1-dimensional representations, one for each character of G .*

Fourier transform. The preceding discussion shows that every function on G can be written as a sum of characters:

$$f(g) = \sum_{\widehat{G}} \widehat{f}(\chi) \chi(g).$$

The coefficients are given by

$$\widehat{f}(\chi) = \langle f, \chi \rangle$$

and the map $L^2(G) \rightarrow L^2(\widehat{G})$ is called the *Fourier transform* of G .

Decomposing abelian representations. In the abelian case, we know that any representation can be decomposed as $V \cong \sum m_i V_i$ where V_i is the 1-dimensional representation associated to the character χ_i . Consequently:

$$\chi(g) = \sum m_i \chi_i.$$

This shows:

Theorem 12.8 *The decomposition of a representation V of an abelian group into its irreducible components is given by the Fourier transform of its trace character. That is, if $f(g) = \text{Tr } \rho(g)$, then $\widehat{f}(\chi) = m_\chi$ is a integer for each $\chi \in \widehat{G}$, and $V \cong \oplus m_\chi V_\chi$.*

Corollary 12.9 *Two representations of a finite abelian group are equivalent iff they have the same character.*

Example. The regular representation V_n of $G = \mathbb{Z}/n$ has character $\chi_n(0) = n$, $\chi_n(k) = 0$ for $k \neq 0$.

If $d|n$ then the regular representation V_d of \mathbb{Z}/d can also be regarded as a representation of \mathbb{Z}/n , with character $\chi_d(k) = d$ if $k = 0 \pmod{d}$, and $\chi_d(k) = 0$ otherwise.

How does the representation $W = \text{Sym}^2(V_n)$ decompose when n is even? This is the same as the permutation representation on pairs of elements of G . Thus $\chi(k)$ is the number of fixed pairs. We find $\chi(0) = n(n+1)/2 = n/2 + n(n/2)$, $\chi(n/2) = n/2$, and $\chi(k) = 0$ otherwise. Thus $\chi = (n/2)\chi_n + \chi_{n/2}$, and thus $\text{Sym}^2(V_n)$ is the sum of $n/2$ copies of V_n and one copy of $V_{n/2}$.

E.g. when $n = 4$ the orbits of $\mathbb{Z}/4$ on $(\mathbb{Z}/4)^2$ consist of $(00, 11, 22, 33)$, $(01, 12, 23, 03)$ and $(02, 13)$, giving 2 copies of V_4 and one copy of V_2 .

More general groups. In general every locally compact abelian group G has a dual group \widehat{G} of the same time, consisting of the continuous characters $\chi : G \rightarrow S^1$. The Fourier transform sends $L^2(G)$ to $L^2(\widehat{G})$.

Examples of dual pairs: S^1 and \mathbb{Z} ; \mathbb{R} and \mathbb{R} . For the first the Fourier transform goes back and forth between functions on S^1 and the coefficients of their powers series $f(z) = \sum a_n z^n$. For the second, the characters are $\exp(itx)$.

The Γ -function. Note that e^{-t} and t^s are homomorphism from \mathbb{R} and \mathbb{R}_+ into \mathbb{C}^* . The Γ -function gives the Fourier transform of e^{-t} as a function on \mathbb{R}^* : it is given by

$$\Gamma(s) = \int_0^\infty e^{-t} t^s \frac{dt}{t}.$$

Its analogue for \mathbb{Z}/p is a Gauss sum,

$$\sigma(\chi) = \sum_{(n,p)=1} \chi(n) e^{2\pi i n/p},$$

where $\chi : (\mathbb{Z}/p)^* \rightarrow S^1$.

Finite, nonabelian groups. Now assume G is a general finite group.

We will shortly see that G has only a finite number of isomorphism classes of irreducible representations. Let us enumerate them V_1, \dots, V_m with V_1 the trivial representation, and let χ_1, \dots, χ_m be the corresponding characters.

The following is the most fundamental result in representations of finite groups.

Theorem 12.10 *The irreducible characters χ_i form an orthonormal basis for the class functions in $L^2(G)$ (functions constant on conjugacy classes).*

Corollary 12.11 *The number of irreducible representations of G is the same as the number of conjugacy classes of G .*

Corollary 12.12 *Any representation V with character χ decomposes as $V \cong \sum m_i V_i$ where $m_i = \langle \chi, \chi_i \rangle$.*

Corollary 12.13 *A representation is irreducible iff its character satisfies $\langle \chi, \chi \rangle = 1$.*

Corollary 12.14 *The dimensions $d_i = \dim V_i$ satisfy*

$$|G| = \sum d_i^2.$$

More precisely, the regular representation decomposes as

$$\mathbb{C}[G] \cong \sum d_i V_i.$$

Proof. The character of the regular representation is $\chi(e) = |G|$ and otherwise $\chi(g) = 0$. Thus $\langle \chi, \chi_i \rangle = \chi_i(e) = d_i$. ■

It can also be shown that d_i divides $|G|$ but this is more difficult.

Corollary 12.15 *Two representations V and V' of G are equivalent iff $\chi = \chi'$.*

Character tables. Note that the theorem gives (i) a quick way to check if a representation is irreducible and (ii) a quick way to check if we have a complete list of irreducible representations. This makes it easy to generate a character table in simple examples, as illustrated below for $G = \mathbb{Z}/3$ and $G = D_3$. The numbers on the top row give the size of each conjugacy class.

These weights must be taken in to account when checking the orthogonality relations.

$\mathbb{Z}/3$	(1)	(1)	(1)	D_6	(1)	(3)	(2)
	e	r	r^2		e	f	r
χ_1	1	1	1	χ_1	1	1	1
χ_2	1	ω	ω^2	χ_2	1	-1	1
χ_3	1	ω^2	ω	χ_3	2	0	-1

The group A_4 has a 3-dimensional geometric representation as the symmetries of a cube. It also admits a map $A_4 \rightarrow \mathbb{Z}/3$, so the representations of $\mathbb{Z}/3$ must occur inside those of A_4 . Using this fact, its character table can be easily computed below.

A_4	(1)	(3)	(4)	(4)
	e	f	r	r^2
χ_1	1	1	1	1
χ_2	1	1	ω	ω^2
χ_3	1	1	ω^2	ω
χ_4	3	-1	0	0

A lot of information is encoded in these tables. For example, the first column gives the *dimension* of the representations. The *commutator subgroup* is the intersection of the kernels of all 1-dimensional characters. More generally, the kernel of any character is *normal*, and the intersections of these give all normal subgroups.

Schur's lemma. As usual in linear algebra, the central idea is to work in a basis-independent fashion. Since we implicitly identify representations that are isomorphic, it is important to have a good method to produce these isomorphisms.

Theorem 12.16 *If $T : V \rightarrow V'$ is a G -map between irreducible representations, then $T = 0$ unless $V \cong V'$. If $V = V'$ then T is a multiple of the identity.*

Proof. Since $\text{Ker } T$ and $\text{Im } T$ are invariant, by irreducibility T is an isomorphism or 0. If $V = V'$ then the eigenspaces of T give further invariant subspaces, so there can be only one. ■

Averaging and matrix coefficients. Here are 2 good averaging procedures to produce G -invariant maps. Suppose $S : V \rightarrow V'$ is any linear map between G -spaces. Then we can set

$$T = \frac{1}{|G|} \sum_G g' S g^{-1}$$

and it is easily verified that T is a G -map. So if V and V' are irreducible, T is a multiple of the identity, which can only be nonzero if $V \cong V'$. In this case we can assume $V = V'$ and then we find

$$\text{tr}(T) = \text{tr}(\lambda I) = \text{tr}(S) = \dim(V)\lambda,$$

so we can compute the value of T from $\text{tr}(S)$.

Matrix coefficients. If V is a unitary representation of G , then (for any basis of V) the functions $g \mapsto \langle ge_i, e_j \rangle$ are called the *matrix coefficients* of V . They span a subspace $c(V) \subset L^2(G)$ that is well-defined independent of the choice of basis or inner product. The span contains $\langle gv, w \rangle$ for any v and w of course.

Within $c(V)$ we have the character χ of V , since this is a sum of matrix coefficients. Thus the orthogonality of different characters follows from:

Theorem 12.17 *If V and V' are distinct irreducible representations of G , then $c(V)$ and $c(V')$ are orthogonal in $L^2(G)$.*

Proof. Let $S : V \rightarrow V'$ be any linear map. Then we have $\sum (g')^{-1} Sg = 0$, and thus applying to v and taking the inner product with v' we get

$$0 = \sum_G \langle S(gv), g'v' \rangle$$

using the unitarity to move g' over to v' . Now suppose $S(x) = \langle x, w \rangle w'$. Then this gives

$$\sum_G \langle gv, w \rangle \langle w', g'v' \rangle = \sum_G \langle gv, w \rangle \overline{\langle g'v', w' \rangle} = 0.$$

But this is just the inner product of 2 matrix coefficients, considered as functions on G . ■

Theorem 12.18 *If V is irreducible then $\langle \chi, \chi \rangle = 1$.*

Proof. Let $(e_i)_1^d$ be an orthonormal basis for V . Then we can write $\chi = \sum \langle ge_i, e_i \rangle = \sum \xi_i$. To complete the proof is it enough to show that $\langle \xi_i, \xi_j \rangle = \delta_{ij}/d$.

To this end, let $S(v) = \langle v, e_i \rangle e_j$. Then $\text{tr}(S) = 1$ and so $\text{tr}(T) = 1$, where $T = |G|^{-1} \sum_G g^{-1} Sg$. Consequently $Tv = v/d$ by Schur's lemma. This shows:

$$\langle Te_i, e_j \rangle = \delta_{ij}/d = |G|^{-1} \sum_G \langle ge_i, e_i \rangle \overline{\langle ge_j, e_j \rangle} = \langle \xi_i, \bar{\xi}_j \rangle$$

as desired. ■

Corollary 12.19 *The number of irreducible characters is at most the number of conjugacy classes in G .*

Class functions from G -maps. We have seen that the character χ of any representation (ρ, V) is a positive integral sum of class functions. Where do the other class functions come from?

A natural source is to consider G -maps $T : V \rightarrow V$ where $V \cong \sum m_i V_i$. That is, $Tg = gT$. Suppose we set

$$\phi(g) = \text{tr}(Tg).$$

Then we have

$$\phi(gh) = \text{Tr}(Tgh) = \text{Tr}(gTh) = \text{Tr}(Thg) = \phi(hg),$$

so ϕ is a class function. Moreover, by Schur's lemma, one can verify that

$$\text{tr}(Tg|m_i V_i) = \lambda_i \chi_i(g)$$

(this is immediate when $m_i = 1$). Consequently

$$\phi(g) = \sum \lambda_i \chi_i$$

(and conversely all sums of characters can be obtained by this construction).

Averaging and class functions. We can also go backwards in a certain sense: given any G -space V and class function ϕ , we can form

$$T = \sum_G \phi(g)g.$$

This is a G -map because:

$$\begin{aligned} T(hv) &= \sum_G \phi(g)ghv = \sum_G \phi(g)h(h^{-1}gh)v \\ &= \sum_G h\phi(hgh^{-1})(hgh^{-1})v = h(Tv) \end{aligned}$$

for all $h \in G$. Using this averaging procedure we find:

Theorem 12.20 *The irreducible characters span the class functions.*

Proof. Let $\phi(g)$ be a class function, let $V = L^2(G)$, and form T as above. Then $\psi(g) = \text{tr}(Tg)$ is a linear combination of characters. But we have

$$\psi(h) = \sum_G \phi(g) \text{Tr}(gh) = |G|\phi(g^{-1}),$$

by properties of the regular representation. Since any function can be expressed as $\phi(g^{-1})$, the sums of characters span the class functions. ■

The Monster. The largest sporadic finite simple group is the *monster*, M , with $|M| \approx 10^{54}$. One of the few ways of grappling with this group is via its character table, which is given in [CCN].

13 Group presentations

In this section we present a topological perspective on group theory: the Cayley graph, free groups, and group presentations.

The Cayley graph. Given generators a_i for G , we draw a directed graph with a vertex for each $g \in G$ and an edge from g to $a_i g$, colored by a_i . If a_i has order two, the arrow is dropped.

Examples: $\langle \mathbb{Z}, 1 \rangle$; $\langle \mathbb{Z}/n, 1 \rangle$; $\langle V_4, a, b \rangle$; generators i, j ; the star, i.e. $\mathbb{Z}/5$ with generator 2.

Examples: (S_3, f, r) vs. $(\mathbb{Z}/6, 2, 3)$. Two triangles running opposite directions in one case, the same direction in the other. Visualizing commutativity.

Example: (S_4, a, b) where $a^3 = b^2 = e$. The graph is a cube with the corners trimmed off.

Groups of order 8 as graphs. The dihedral group (D_4, f, r) . The group $\mathbb{Z}/2 \times \mathbb{Z}/4$.

The quaternion group Q with generators i, j ; 8 points on S^3 ! (Put $\pm 1, \pm i, \pm j$ on the vertices of an octahedron. Then put k in the center of the octahedron and $-k$ at infinity!)

Example: the Cayley graph of $\mathbb{Z}/a \times \mathbb{Z}/b$ can be visualized as a torus.

The free group on 2 generators. Any one generator G group admits a *surjective* map $\phi: \mathbb{Z} \rightarrow G$. Thus \mathbb{Z} is the ‘biggest’ group with one generator — every other cyclic group is a quotient of it.

What is the biggest group you can make with two generators?

Answer: Consider all finite words in the letters a, a', b and b' . A word is *reduced* if it does not contain $aa', a'a, bb'$ or $b'b$. An arbitrary word can be reduced by repeatedly canceling out (removing) these 2-letter subwords whenever they occur. Example: $ab'aa'ba \rightarrow ab'ba \rightarrow aa$.

Let G be the set of all reduced words. Multiplication is defined by concatenation followed by reduction. The inverse of a word is obtained by writing it down backwards and exchanging the primed and unprimed variables.

The result is the *free group* $\mathbb{Z} * \mathbb{Z} = \langle a, b \rangle$.

Trees. The Cayley graph of the free group is an infinite tree.

Theorem 13.1 *Let G' be a group that can be generated by two elements x and y . Then there is a unique surjective map $\phi : \mathbb{Z} * \mathbb{Z} \rightarrow G'$ with $\phi(a) = x$ and $\phi(b) = y$.*

This means G' can be described as $G/\text{Ker } \phi$.

Example. Let $H \subset G = \mathbb{Z} * \mathbb{Z} = \langle a, b \rangle$ be the subgroup generated by all commutators, $[g, h] = ghg^{-1}h^{-1}$. Since the set of commutators is closed under conjugation, so is H . Therefore H is *normal*.

What is G/H ? Construct a map $\phi : G \rightarrow \mathbb{Z}^2$ sending a to $(1, 0)$ and b to $(0, 1)$. Since every commutator is in the kernel, we actually get a map $\phi : G/H \rightarrow \mathbb{Z}^2$. Now construct a map $\psi : \mathbb{Z}^2 \rightarrow G/H$, sending (i, j) to $a^i b^j$. Since G/H is abelian, this map is a homomorphism. Clearly the compositions are the identity, so G/H is isomorphic to \mathbb{Z}^2 .

Group presentations. By $G = \langle g_1, \dots, g_n : r_1, \dots, r_m \rangle$ we denote F/N where F is the free group on generators g_1, \dots, g_n , and $N \subset F$ is the smallest normal subgroup containing $R = \{r_1, \dots, r_m\}$.

Another way to put it is that N is the set of *consequences* of the relations R . Here a consequence of R means either

1. the identity e (the trivial consequence),
2. an element $r \in R$,
3. the product of two consequences,
4. the inverse of a consequence, or
5. fxf^{-1} where $f \in F$ and x is a consequence.

Clearly if we have a group in which the elements of R represent the identity, so do the consequences of R . Now it is easy to see that the consequences form a normal subgroup, and in fact they are exactly N , the smallest normal subgroup containing R .

Theorem 13.2 *To give a homomorphism $\psi : G \rightarrow H$, where $G = \langle g_1, \dots, g_n : r_1, \dots, r_m \rangle$, it suffices to give values $h_i = \psi(g_i) \in H$ for each generator, and check that $\psi(r_i) = e$ for each relations.*

Proof. Let F be the free group on generators $\langle g_1, \dots, g_n \rangle$. The h_i determine a homomorphism $\phi : F \rightarrow H$. By assumption, $r_i \in \text{Ker}(\phi)$ for each relation. Since the kernel is a normal subgroup, it *contains* the smallest normal subgroup N containing the relations r_i . But by definition, $F/N \cong G$, so ϕ descends to a map $\psi : F/N \rightarrow H$. ■

Presenting \mathbb{Z}^2 . Let $G = \langle a, b : ab = ba \rangle$. Then G is abelian, and there is a map to \mathbb{Z}^2 , which is obviously surjective. There is also an inverse map, which is a homomorphism because G is abelian. Both compositions give the identity, so $G \cong \mathbb{Z}^2$.

The checkerboard. It is useful to compare the Cayley graphs of the free group $\langle a, b \rangle$ and of the free abelian group $\langle a, b : ab = ba \rangle$. The relation gives loops in the second graph.

Dihedral groups. As an example, let's consider the group

$$G = \langle f, r : f^2 = r^n = e, rf = fr^{-1} \rangle.$$

Using the relations, every word can be represented in the form $r^i f^j$ where $0 \leq i < n$ and $j = 0, 1$. Thus G has *at most* $2n$ elements.

Define the obvious homomorphism $\phi : G \rightarrow D_n$. Then ϕ is onto, so by counting we see it is an isomorphism.

The infinite dihedral group. Let $D_\infty = \langle f, r : rf = fr^{-1} \rangle$. We can think of D_∞ as acting on \mathbb{R} as the boundary of an 'infinite polygon', with vertices at the integers, by $f(x) = -x$, $r(x) = x + 1$. Then $frf(x) = -((-x) + 1) = x - 1 = r^{-1}(x)$ as required.

Another presentation for D_∞ . Let $G = \langle a, b : a^2 = b^2 = e \rangle = \mathbb{Z}/2 * \mathbb{Z}/2$. It is easy to draw the Cayley graph of G ; it's a straight line, just like the boundary of an infinite polygon.

Theorem 13.3 *D_∞ and G are isomorphic.*

Proof. Define a map $\phi : G \rightarrow D_\infty$ by $\phi(a) = f$, $\phi(b) = rf$. Then clearly $\phi(a^2) = e$ and $\phi(b^2) = rfrf = rr^{-1} = e$, so ϕ is a homomorphism.

Now define a map $\psi : D_\infty \rightarrow G$ by $\psi(f) = a$ and $\psi(r) = ba$. Then $\psi(f^2) = a^2 = e$ and

$$\psi(fr^{-1}) = a(ba)' = aa'b' = b' = b = (ba)a = \psi(rf),$$

so ψ is a homomorphism. We then compute $\psi \circ \phi(a) = a$,

$$\psi \circ \phi(b) = \psi(rf) = baa = b,$$

so $\psi \circ \phi$ is the identity. Similarly $\phi \circ \psi$ is the identity, so these two groups are isomorphic. ■

Generators and relations of S_n .

Theorem 13.4 S_n has generators $\tau_i = (i, i + 1)$, $i = 1, \dots, n - 1$, with relations

$$\begin{aligned} \tau_i^2 &= e; \\ \tau_i \tau_{i+1} \tau_i &= \tau_{i+1} \tau_i \tau_{i+1}; \text{ and} \\ \tau_i \tau_j &= \tau_j \tau_i \text{ if } |i - j| > 1. \end{aligned}$$

Proof. To check the main relation, let $(i, i + 1, i + 2) = (i, j, k)$; then we have: $(ij)(jk)(ij) = (ik) = (jk)(ij)(jk)$. So there is a map of the group above to S_n , and since adjacent permutations generate, it is *onto*.

Now by the picture of changing crossings, it is clear that any two diagrams of the same permutation differ by these relations. ■

Cor. The parity of an element in S_n is well-defined.

Proof. The relations preserve parity. Alternatively, define a map from S_n to $\mathbb{Z}/2$ by sending each τ_i to one, and observe that the relations are satisfied. ■

Trivial groups. It is not always easy to tell whether or not a presentation is just giving the trivial group. For example, $\langle a : a^{12} = e, a^{25} = e \rangle$ is trivial.

14 Knots and the fundamental group

What is a knot? It is a smooth closed curve in 3-space. A knot is not allowed to cross itself. A knot can be moved a little so it is a polygon. We do not allow wild knots.

Two knots K_0 and K_1 are *equivalent* if you can make a smoothly moving family of knots K_t that connects them. You can imagine this motion taking place in discrete steps, K_0, \dots, K_n , where K_i and K_{i+1} differ by a triangle move.

A *link* is defined similarly as a finite number of disjoint closed loops.

Knot projections. A useful way to discuss knots is by projections: you put the knot almost in a plane, with pairs of strands meeting at crossings.

Any knot can be given a knot projection; in fact a generic projection will work. You just have to avoid the directions tangent to the knot, and the directions of lines passing through the knot in 3 points (taken with multiplicities). Each forms a one-dimensional set.

Examples.

1. The unknot. There are several projections. Any knot projection with 0, 1 or 2 crossings is the unknot. Any knot projection you can draw without anticipating crossings is the unknot.
2. The trefoil knot, 3_1 .
3. The figure-eight knot, 4_1 .
4. The (p, q) -torus knot. You start with q strands and form a braid of the form β^p , where β is a cyclic permutation; then close. If p and q are relatively prime, you get a knot. The $(1, 3)$ torus knot is the unknot; $(2, 3)$ is the trefoil.
5. The unlink on two components.
6. The Hopf link.
7. The Borromean rings, after the Renaissance family crest of the Borromeas.

History. Lord Kelvin conjectured that atoms are knots in ether. Tait and Little undertook the tabulation of knots. In recent times biologists have

discovered that DNA is often knotted. The classification of 3-dimensional spaces is intimately tied up with knot theory.

Showing two knots are the same. Suppose K_1 and K_2 are projections that happen to correspond to the same knot. Then you can transform K_1 to K_2 be a sequence of *Reidemeister moves* (or their inverses). These moves are:

- I Transform one strand into a loop with one crossing. The singularity is a cusp.
- II Transform two parallel strands by adding two crossings. The singularity is a tangency.
- III Transform three strands preserving the number of crossings. The singularity is a triple-point.

The Reidemeister moves can be remembered by the number of strands they involve. Planar isotopy is also allowed. The Reidemeister moves also work for links.

Proof. One way to approach the proof is to consider what kinds of singularities can arise as you view a generic knot projection during isotopy. The generic singularities are cusps, tangencies and triple points, accounting for the 3 moves.

Example: Draw a trefoil with one crossing wrong. This can be undone by the sequence III, I, II.

Example: Tangle up the trefoil.

Example: For each crossing of 6_3 , change it and simplify the result.

Orientation, linking number and tricoloring. A knot or link is *oriented* if we have chosen a direction (usually indicated by an arrow) to traverse each component. A link with n components has 2^n possible orientations.

Showing two links are different. Let $L = K_1 \cup K_2$ be a two component oriented link. The *linking number* $\ell(K_1, K_2)$ is defined as follows: at each crossing between K_1 and K_2 , count +1 if it is a right-hand turn to get onto the overpass, otherwise -1. Add up and divide by two; this is $\ell(K_1, K_2)$.

Theorem 14.1 *The linking number is an invariant of L . Even though it is defined using a projection, the answer for two different projections is the same.*

Proof. Type I moves don't involve both components. A type two moves creates a pair of crossings of opposite sign. And type III doesn't really change any pair of strands, only the configuration of all three. ■

Examples: the unlink, Hopf link and Whitehead link.

Unoriented links. The *absolute value* of the linking number is independent of orientation.

Tricoloring. We still haven't shown there are any real knots. To do this, let's say a *tricoloring* of a knot or link projection is an assignment of colors R, G, B to the arcs such that:

- 1) at least 2 colors are used; and at each crossing either
- 2a) all three strands have the same color or
- 2b) all three strands have different colors.

Theorem 14.2 *If one projection of a knot or link can be tricolored, then they all can.*

Proof. We must check the Reidemeister moves.

- (I) The entire loop must be a single color.
- (II) If the parallel strands are colored R and G , then we color the underloop B .
- (III) If 3 colors appear, then a crossing where 3 colors are the same either appears or disappears.

In all 3 cases, if 2 colors were used before, then 2 are used after (see especially 2, which can increase the number of colors from 2 to 3).

Example. The unknot cannot be tricolored.

Example. The trefoil *can* be tricolored. Thus the trefoil is really knotted!

Example. The figure 8 knot cannot be tricolored. Thus tricoloring is not powerful enough to detect all knots.

Connect sum and colorings. Example. The connect sum of 2 trefoils can be tricolored in 11 different ways. (!) The point is that we can use a 1-coloring on one factor so long as it hooks up with a 3-coloring on the other factor.

The *number* of tricolorings is also an invariant of the knot or link.

Example. Let L be the unlink on 2 components. Then L can be tricolored with 2 colors — if drawn with no crossings — and with 3 colors — if drawn with 2 crossings.

Fundamental groups. *Examples of $\pi_1(X)$.* The disk, the annulus, the circle, the figure eight, the torus, the torus with a hole, a surface of genus two.

The knot group. This is the group of flight plans for airplanes leaving from a base located outside the knot and flying around it.

Examples: For the unknot, $G(K) = \mathbb{Z}$. For the unlink on two components, $G(L) = \mathbb{Z} * \mathbb{Z}$.

Presenting a knot or link group. There is one generator for each arc of the knot projection. To each arc we associate a generator of $\pi_1(\mathbb{R}^3 - K)$ that makes a right-handed underpass for that arc. (This means as one moves along the generator, one can make a safe right-hand turn to get on the superhighway.)

Now write ab for the loop that does a first, then b . Then when c crosses over $a - b$, so that we have a right-hand underpass, we get $ac = cb$. At a left-handed underpass, we get $ca = bc$.




Figure 6. Proof that $ac = cb$ at safe underpass.

The trefoil knot group. We can now finally return to the trefoil knot. In the picture with all safe crossings, the group is

$$G(K) = \langle a, b, c : ab = bc, bc = ca, ca = ab \rangle.$$

Mapping $G(K)$ to \mathbb{Z} . This is easy: define $\phi(a) = \phi(b) = \phi(c) = 1$. It works for every knot group.

Mapping $G(K)$ to S_3 . Notice that S_3 has 3 transpositions, call them A, B, C ; they are all the odd permutations. They must satisfy $AB = CA$, since ABA is odd, it can't be A (else $A = B$) and it can't be B (else S_3 is commutative). Mapping (a, b, c) to (A, B, C) we see G admits S_3 as a quotient!

Corollary 14.3 *The trefoil knot really is knotted.*

Theorem 14.4 *The tricolorings of a knot correspond one-to-one to the surjective homomorphisms $\phi : G(K) \rightarrow S_3$.*

Proof. Let A, B, C be the three flips in S_3 as before. Since the generators of $G(K)$ correspond to arcs, we can use the three colors to define a map $\phi : G(K) \rightarrow S_3$ on the generators. Then the tricoloring condition shows that each relation in $G(K)$ is satisfied. So we can map the generators for strands of color A to flip a , etc. Since at least two colors are used, we get a surjective map.

Similarly, if we have a surjection, then the generators (all being conjugate) must go to flips, else the image would be in the abelian subgroup of rotations. We then obtain a coloring. ■

Note: for a link the statement above is not quite true. A homomorphism that kills one component of a link does not correspond to a coloring. That is the tricolorings correspond to maps to S_3 that send all generators of the Wirtinger presentation to flips.

Changing presentations. To prove $G(K)$ is a knot invariant, not just an invariant of the knot projection, it is important to understand elementary (or Tietze) moves on a presentation. There are just two:

- (1) $\langle g_i : r_i \rangle \iff \langle g_i : r_i, s \rangle$, where s is a consequence of the given relations. That means s is a product of conjugates of the r_i .
- (2) $\langle g_i : r_i \rangle \iff \langle g_i, h : r_i, h = w(g_1, \dots, g_n) \rangle$, where $w(\cdot)$ is a word in the generators g_i . This means we can add a new generator so long as we add a relation putting it in the group generated by the (g_i) .

Example. The trefoil group can be simplified to 2 generators, one relation.
Invariance of $G(K)$.

Theorem 14.5 *Equivalent projections of a knot (or link) give isomorphic groups $G(K)$.*

Proof. We must check the Reidemeister moves. (I) A loop gives $\langle a, b : aa = ba \rangle$, so we can use Tietze move (2) to eliminate b .

(II) Suppose the arc a is underpassed by (b, c, d) . Then we get from the 2 crossings the relations $ba = ac, ac = da$. From this we derive $b = d$ (Tietze

1), then eliminate c, d (Tietze 2). We are left with a, b and no relations, which is the contribution of two parallel arcs.

(III) Let the topmost arc be c , NW to SE, over (a, d) , SW to NE, with finally (b, f, e) passing W to E under a and c . The 3 crossings give the relations

$$R = \langle ac = cd, ab = fa, fc = ce \rangle.$$

See Figure 7.

We can obtain the relation $aba' = cec' = f$, and then eliminate f , giving the new set of relations

$$R' = \langle ac = cd, aba' = cec' \rangle$$




Figure 7. Reidemeister move III.

After the Reidemeister move we get a new arc g and relations

$$S = \langle ac = cd, bc = cg, dg = ed \rangle.$$

We can obtain the relation $c'bc = d'ed = g$, and eliminate g , obtaining

$$S' = \langle ac = cd, c'bc = d'ed \rangle.$$

We now must show the two sets of relations R' and S' are equivalent. To see this, we will use the relation $ac = cd$, present in each group. Then the relation $c'bc = d'ed$ in S' , implies the relation $(ac)c'bc(ac)' = (cd)d'ed(cd)'$, which simplifies to the relation $aba' = cec'$ in R' . The process can be reversed, so indeed the groups are isomorphic. ■

Search for cycles. If we want to find a quotient $\phi : G(K) \rightarrow F$, it is *very useful* to note that the standard generators of a knot group are all conjugates. Thus the images of these generators must also be conjugate, and generate F (if ϕ is to be surjective.)

With this in mind, we can see that a natural choice for the trefoil group is S_3 , the symmetries of a triangle, since it has three edges. We map each generator to the flip of the edge, and this gives ϕ .

The figure 8 group. We have the presentation

$$G = \langle a, b, c, d : da = bd, ba = ac, bc = db, dc = ca \rangle.$$

Now the elements a, b, c, d are *all conjugate* in G . So if we try to find a map $\phi : G \rightarrow F$, where F is a finite group, then the images of a, b, c, d all have the same order — and indeed they are all conjugate. So we need a group with 4 conjugate elements, that generate the group.

A nice candidate is $F = A_4$, the symmetries of a tetrahedron T . The motivation is that T has 4 faces, corresponding to the generators of G , just like the edges of the triangle did for the trefoil knot.

The clockwise rotations of a face are the permutations (123), (134), (243) and (142). Any two faces stand in the same relation, so we may as well send a to (123) and d to (134).

Then the first relations gives $b = dad' = (432)$, and the second gives $c = a'ba = (421)$. We can then check that the last two relations are satisfied.

Geometric check. Another way to think about the relations in the figure 8 group is to let a, b, c, d correspond to faces A, B, C, D of the tetrahedron. Then an equation like $b = dad'$ is equivalent to $B = d(A)$. So we get the relations $B = d(A), C = a'(B), D = b(C)$ and $D = c(A)$. Once A and D are chosen, the first two relations specify C and B , and then we can check the last two.

Trefoil \neq Figure Eight.

Theorem 14.6 *There is no surjection from the trefoil group to A_4 . Thus the unknot, the trefoil and the figure eight knot are all distinct.*

Proof. Let $\phi : G \rightarrow A_4$ be a surjective homomorphism from the trefoil group to A_4 . Since the generators are all conjugate, they all have the same order, which must be 1, 2 or 3. But the elements of order 2 in A_4 generate a proper V_4 subgroup, so the order must be 3. Moreover two generators must map to different elements, else the image would be $\mathbb{Z}/3$.

Up to symmetry we can then take the images of the generators to be $a = (123)$ and $b = (134)$. Then the relations $ab = bc = ca$ imply $c = b'ab = (432)$, but we then find $bc = (321) \neq ca = (431)$. ■

Hopf link. The Hopf link L is *less knotted* than the unlink, since $G(L) \cong \mathbb{Z} \oplus \mathbb{Z}$. As a trick, one can weave the commutator through two unlinked carabiner, in such a way that the loop comes free when the carabiner are linked! (Cf. homework on computing $G(L)$.)

Theorem 14.7 *The linking number for $L = K_1 \cup K_2$ corresponds to the abelianization of the element of $\pi_1(\mathbb{R}^3 - K_1)$ represented by K_2 .*

Proof. The proof is a little tricky. Working with a link projection, one can first change crossings of K_2 with itself so K_2 becomes the unknot. This change does not change our projection-based calculation of the linking number (obviously), nor does it change the image of K_2 in $\pi_1(\mathbb{R} - K_1)$ (obviously).

Now re-arrange the projection so K_2 is a counter-clockwise round circle. It is then clear that the number of outward-heading crossings of K_1 with K_2 is the same as the number of inward-heading crossings.

Count the crossings in 4 groups, TO, BO, TI, BI , where T/B means K_1 is on top/bottom, and O/I means it is headed out/in. Letting P be the link number in π_1 , and L the linking number from the diagram, we have

$$\begin{aligned} TO + BO &= TI + BI \\ L &= (TO + BI - TI - BO)/2 \quad \text{and} \\ P &= TO - TI. \end{aligned}$$

Using the first equation we have $P = BI - BO$; averaging these two expressions for P , we obtain $P = L$. ■

Theorem 14.8 (Gordon-Luecke) *Let $(G(K), H(K))$ be a knot group and a subgroup $H(K) \cong \mathbb{Z}$ generated by a meridian. Then K is equivalent to K' , or its mirror image, iff there is an isomorphism $G(K) \cong G(K')$ sending $H(K)$ to $H(K')$.*

The $H(K)$ is needed for the square and granny knots, which have isomorphic groups. Often the $H(K)$ is unnecessary.

References

- [Ar] M. Artin. *Algebra*. Prentice-Hall, 1991.
- [Ax] S. Axler. *Linear Algebra Done Right*. Springer, 1997.
- [CCN] J. H. Conway, R. T. Curtis, S. P. Norton, R. A. Parker, and R. A. Wilson. *Atlas of Finite Groups*. Oxford University Press, 1985.
- [Hal] P. Halmos. *Finite-Dimensional Vector Spaces*. Springer-Verlag, 1987.