

Lecture notes
Math 4377/6308 – Advanced Linear Algebra I

Vaughn Climenhaga

The primary text for this course is “Linear Algebra and its Applications”, second edition, by Peter D. Lax (hereinafter referred to as **[Lax]**). The lectures will follow the presentation in this book, and many of the homework exercises will be taken from it.

You may occasionally find it helpful to have access to other resources that give a more expanded and detailed presentation of various topics than is available in Lax’s book or in the lecture notes. To this end I suggest the following list of external references, which are freely available online.

- (**Bee**) “A First Course in Linear Algebra”, by Robert A. Beezer, University of Puget Sound. Long and comprehensive (1027 pages). Starts from the very beginning: vectors and matrices as arrays of numbers, systems of equations, row reduction. Organisation of book is a little non-standard: chapters and sections are given abbreviations instead of numbers. <http://linear.ups.edu/>
- (**CDW**) “Linear Algebra”, by David Cherney, Tom Denton, and Andrew Waldron, UC Davis. 308 pages. Covers similar material to **[Bee]**. <https://www.math.ucdavis.edu/~linear/>
- (**Hef**) “Linear Algebra”, by Jim Hefferon, Saint Michael’s College. 465 pages. Again, starts from the very beginning. <http://joshua.smcvt.edu/linearalgebra/>
- (**LNS**) “Linear Algebra as an Introduction to Abstract Mathematics”, by Isaiah Lankham, Bruno Nachtergaele, and Anne Schilling, UC Davis. 247 pages. More focused on abstraction than the previous three references, and hence somewhat more in line with the present course. https://www.math.ucdavis.edu/~anne/linear_algebra/
- (**Tre**) “Linear Algebra Done Wrong”,¹ by Sergei Treil, Brown University. 276 pages. Starts from the beginning but also takes a more abstract view. <http://www.math.brown.edu/~treil/papers/LADW/LADW.html>

The books listed above can all be obtained freely via the links provided. (These links are also on the website for this course.) Another potentially useful resource is the series of video lectures by Gilbert Strang from MIT’s Open CourseWare project: <http://ocw.mit.edu/courses/mathematics/18-06-linear-algebra-spring-2010/video-lectures/>

¹If the title seems strange, it may help to be aware that there is a relatively famous textbook by Sheldon Axler called “Linear Algebra Done Right”, which takes a different approach to linear algebra than do many other books, including the ones here.

Motivation, linear spaces, and isomorphisms

Further reading: [Lax] Ch. 1 (p. 1–4). See also [Bee] p. 317–333; [CDW] Ch. 5 (p. 79–87); [Hef] Ch. 2 (p. 76–87); [LNS] Ch. 4 (p. 36–40); [Tre] Ch. 1 (p. 1–5)

1.1 General motivation

We begin by mentioning a few examples that on the surface may not appear to have anything to do with linear algebra, but which turn out to involve applications of the machinery we will develop in this course. These (and other similar examples) serve as a motivation for many of the things that we do.

1. **Fibonacci sequence.** The Fibonacci sequence is the sequence of numbers 1, 1, 2, 3, 5, 8, 13, \dots , where each number is the sum of the previous two. We can use linear algebra to find an exact formula for the n th term. Somewhat surprisingly, it has the odd-looking form

$$\frac{1}{\sqrt{5}} \left(\left(\frac{1 + \sqrt{5}}{2} \right)^n - \left(\frac{1 - \sqrt{5}}{2} \right)^n \right).$$

We will discuss this example when we talk about eigenvalues, eigenvectors, and diagonalisation.

2. **Google.** Linear algebra and Markov chain methods are at the heart of the PageRank algorithm that was central to Google's early success as an internet search engine. We will discuss this near the end of the course.
3. **Multivariable calculus.** In single-variable calculus, the derivative is a number, while in multivariable calculus it is a matrix. The proper way to understand this is that in both cases, the derivative is a linear transformation. We will reinforce this point of view throughout the course.
4. **Singular value decomposition.** This is an important tool that has applications to image compression, suggestion algorithms such as those used by Netflix, and many other areas. We will mention these near the end of the course, time permitting.

5. **Rotations.** Suppose I start with a sphere, and rotate it first around one axis (through whatever angle I like) and then around a different axis (again through whatever angle I like). How does the final position of the sphere relate to the initial one? Could I have gotten from start to finish by doing a single rotation around a single axis? How would that axis relate to the axes I actually performed rotations around? This and other questions in three-dimensional geometry can be answered using linear algebra, as we will see later.
6. **Partial differential equations.** Many important problems in applied mathematics and engineering can be formulated as partial differential equations; the heat equation and the wave equation are two fundamental examples. A complete theory of PDEs requires functional analysis, which considers vector spaces whose elements are not arrays of numbers (as in \mathbb{R}^n), but rather functions with certain differentiability properties.

There are many other examples: to chemistry (vibrations of molecules in terms of their symmetries), integration techniques in calculus (partial fractions), magic squares, error-correcting codes, etc.

1.2 Background: general mathematical notation and terminology

Throughout this course we will assume a working familiarity with standard mathematical notation and terminology. Some of the key pieces of background are reviewed on the first assignment, which is due at the beginning of the next lecture.

For example, recall that the symbol \mathbb{R} stands for the set of real numbers; \mathbb{C} stands for the set of complex numbers; \mathbb{Z} stands for the integers (both positive and negative); and \mathbb{N} stands for the natural numbers $1, 2, 3, \dots$. Of particular importance will be the use of the quantifiers \exists (“there exists”) and \forall (“for all”), which will appear in many definitions and theorems throughout the course.

- Example 1.1.**
1. The statement “ $\exists x \in \mathbb{R}$ such that $x + 2 = 7$ ” is true, because we can choose $x = 5$.
 2. The statement “ $x + 2 = 7 \forall x \in \mathbb{R}$ ” is false, because $x + 2 \neq 7$ when $x \neq 5$.

3. The statement “ $\forall x \in \mathbb{R} \exists y \in \mathbb{R}$ such that $x + y = 4$ ” is true, because no matter what value of x is chosen, we can choose $y = 4 - x$ and then we have $x + y = x + (4 - x) = 4$.

The last example has *nested* quantifiers: the quantifier “ \exists ” occurs inside the statement to which “ \forall ” applies. You may find it helpful to interpret such nested statements as a game between two players. In this example, Player A has the goal of making the statement $x + y = 4$ (the innermost statement) be true, and the game proceeds as follows: first Player B chooses a number $x \in \mathbb{R}$, and then Player A chooses $y \in \mathbb{R}$. If Player A’s choice makes it so that $x + y = 4$, then Player A wins. The statement in the example is true because Player A can always win.

Example 1.2. The statement “ $\exists y \in \mathbb{R}$ such that $\forall x \in \mathbb{R}, x + y = 4$ ” is false. In the language of the game played just above, Player A is forced to choose $y \in \mathbb{R}$ first, and then Player B can choose any $x \in \mathbb{R}$. Because Player B gets to choose *after* Player A, he can make it so that $x + y \neq 4$, so Player A loses.

To parse such statements it may also help to use parentheses: the statement in Example 1.2 would become “ $\exists y \in \mathbb{R}$ (such that $\forall x \in \mathbb{R} (x + y = 4)$)”. Playing the game described above corresponds to parsing the statement from the outside in. This is also helpful when finding the negation of the statement (formally, its *contrapositive* – informally, its opposite).

Example 1.3. The negations of the three statements in Example 1.1 are

1. $\forall x \in \mathbb{R}$ we have $x + 2 \neq 7$.
2. $\exists x \in \mathbb{R}$ such that $x + 2 \neq 7$.
3. $\exists x \in \mathbb{R}$ such that $\forall y \in \mathbb{R}$ we have $x + y \neq 4$.

Notice the pattern here: working from the outside in, each \forall is replaced with \exists , each \exists is replaced with \forall , and the innermost statement is negated (so $=$ becomes \neq , for example). You should think through this to understand why this is the rule.

1.3 Vector spaces

In your first linear algebra course you studied vectors as rows or columns of numbers – that is, elements of \mathbb{R}^n . This is the most important example

of a vector space, and is sufficient for many applications, but there are also many other applications where it is important to take the lessons from that first course and re-learn them in a more abstract setting.

What do we mean by “a more abstract setting”? The idea is that we should look at vectors in \mathbb{R}^n and the things we did with them, and see exactly what properties we needed in order to use the various definitions, theorems, techniques, and algorithms we learned in that setting.

So for the moment, think of a vector as an element of \mathbb{R}^n . What can we do with these vectors? A moment’s thought recalls several things:

1. we can add vectors together;
2. we can multiply vectors by real numbers (scalars) to get another vector, which in some sense points in the same “direction”;
3. we can multiply vectors by matrices;
4. we can find the length of a vector;
5. we can find the angle between two vectors.

The list could be extended, but this will do for now. Indeed, for the time being we will focus only on the first two items on the last. The others will enter later.

So: vectors are things that we can add together, and that we can multiply by scalars. This motivates the following definition.

Definition 1.4. A *vector space* (or *linear space*) over \mathbb{R} is a set X on which two operations are defined:

- addition, so that given any $x, y \in X$ we can consider their sum $x + y \in X$;
- scalar multiplication, so that given any $x \in X$ and $c \in \mathbb{R}$ we can consider their product $cx \in X$.

The operations of addition and scalar multiplication are required to satisfy certain properties:

1. commutativity: $x + y = y + x$ for every $x, y \in X$;
2. associativity of addition: $x + (y + z) = (x + y) + z$ for every $x, y, z \in X$;
3. identity element: there exists an element $\mathbf{0} \in X$ such that $x + \mathbf{0} = x$ for all $x \in X$;

4. additive inverses: for every $x \in X$ there exists $(-x) \in X$ such that $x + (-x) = \mathbf{0}$;
5. associativity of multiplication: $a(bx) = (ab)x$ for all $a, b \in \mathbb{R}$ and $x \in X$;
6. distributivity: $a(x+y) = ax+ay$ and $(a+b)x = ax+bx$ for all $a, b \in \mathbb{R}$ and $x, y \in X$;
7. multiplication by the unit: $1x = x$ for all $x \in X$.

The properties in the list above are the *axioms* of a vector space. They hold for \mathbb{R}^n with the usual definition of addition and scalar multiplication. Indeed, this is in some sense the motivation for this list of axioms: they formalise the properties that we know and love for the example of row/column vectors in \mathbb{R}^n . We will see that these properties are in fact enough to let us do a great deal of work, and that there are plenty of other things besides \mathbb{R}^n that satisfy them.

Remark 1.5. Some textbooks use different font styles or some other typographic device to indicate that a particular symbol refers to a vector, instead of a scalar. For example, one may write \mathbf{x} or \vec{x} instead of x to indicate an element of a vector space. By and large we will not do this; rather, plain lowercase letters will be used to denote both scalars and vectors (although we will write $\mathbf{0}$ for the zero vector, and 0 for the zero scalar). It will always be clear from context which type of object a letter represents: for example, in Definition 1.4 it is always specified whether a letter represents a vector (as in $x \in X$) or a scalar (as in $a \in \mathbb{R}$). You should be very careful when reading and writing mathematical expressions in this course that you are always aware of whether a particular symbol stands for a scalar, a vector, or something else.

Before moving on to some examples, we point out that one may also consider vector spaces over \mathbb{C} , the set of complex numbers; here the scalars may be any complex numbers. In fact, one may consider any *field* K and do linear algebra with vector spaces over K . This has many interesting applications, particularly if K is taken to be a finite field, but these examples lie beyond the scope of this course, and while we will often say “Let X be a vector space over the field K ”, it will always be the case in our examples that K is either \mathbb{R} or \mathbb{C} . Thus we will not trouble ourselves here with the general abstract notion of a field.

Certain properties follow immediately from the axioms, although they are not explicitly included in them. It is a worthwhile exercise to deduce the following results from the axioms.

1. The identity element is unique: if $\mathbf{0}' \in X$ is such that $x + \mathbf{0}' = x$ for all $x \in X$, then $\mathbf{0}' = \mathbf{0}$.
2. $0x = 0$ for all $x \in X$.
3. $(-1)x = -x$ for all $x \in X$.

1.4 Examples

The most familiar examples are the following.

Example 1.6. Let $X = \{(x_1, \dots, x_n) \mid x_j \in \mathbb{R} \forall j\}$ be the set of row vectors with n real components, and let addition and scalar multiplication be defined coordinate-wise. Then X is a vector space over \mathbb{R} .

Example 1.7. Let $Y = \left\{ \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \mid x_j \in \mathbb{R} \forall j \right\}$ be the set of column vectors with n real components, and let addition and scalar multiplication be defined coordinate-wise. Then Y is a vector space over \mathbb{R} .

Analogously, one can define \mathbb{C}^n as either row vectors or column vectors with components in \mathbb{C} .

The two examples above look very similar, but formally they are different vector spaces; after all, the sets are different, and a row vector is not a column vector. Nevertheless, there is a real and precise sense in which they are “the same example”: namely, they are *isomorphic*. This means that there is a bijective (one-to-one and onto) correspondence between them that maps sums into sums and scalar multiples into scalar multiples: in this case we can consider the transpose map $T: X \rightarrow Y$ given by $T(x_1, \dots, x_n) = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$, which has the properties $T(x + y) = T(x) + T(y)$ and $T(cx) = cT(x)$ for all $x, y \in X$ and $c \in \mathbb{R}$.

Remark 1.8. Recall that a map $T: X \rightarrow Y$ is 1-1 if $T(x) = T(x')$ implies $x = x'$, and onto if for every $y \in Y$ there exists $x \in X$ such that $T(x) = y$.

We will discuss isomorphisms, and other linear transformations, at greater length later in the course. The key point for now is that as far as the tools of linear algebra are concerned, isomorphic vector spaces are indistinguishable from each other, although they may be described in quite different ways.

Example 1.9. Let X be the set of all functions $x(t)$ satisfying the differential equation $\ddot{x} + x = 0$. If x and y are solutions, then so is $x + y$; similarly, if x is a solution then cx is a solution for every $c \in \mathbb{R}$. Thus X is a vector space. If p is the initial position and v is the initial velocity, then the pair (p, v) completely determines the solution $x(t)$. The correspondence between the pair $(p, v) \in \mathbb{R}^2$ and the solution $x(t)$ is an isomorphism between \mathbb{R}^2 and X .

Example 1.10. Let \mathbb{P}_n be the set of all polynomials with coefficients in K and degree at most n : that is, $\mathbb{P}_n = \{a_0 + a_1t + a_2t^2 + \cdots + a_nt^n \mid a_0, \dots, a_n \in K\}$. Then \mathbb{P}_n is a vector space over K .

Example 1.11. Let $F(\mathbb{R}, \mathbb{R})$ be the set of all functions from $\mathbb{R} \rightarrow \mathbb{R}$, with addition and scalar multiplication defined in the natural way (pointwise) by $(f + g)(x) = f(x) + g(x)$ and $(cf)(x) = c(f(x))$. Then $F(\mathbb{R}, \mathbb{R})$ is a vector space. It contains several other interesting vector spaces.

1. Let $C(\mathbb{R})$ be the subset of $F(\mathbb{R}, \mathbb{R})$ that contains all *continuous* functions.
2. Let $L^1(\mathbb{R})$ be the subset of $F(\mathbb{R}, \mathbb{R})$ that contains all *integrable* functions: that is, $L^1(\mathbb{R}) = \{f: \mathbb{R} \rightarrow \mathbb{R} \mid \int_{-\infty}^{\infty} |f(x)| dx < \infty\}$.
3. Let $C^1(\mathbb{R})$ be the subset of $F(\mathbb{R}, \mathbb{R})$ that contains all *differentiable* functions.

Each of $C(\mathbb{R})$, $L^1(\mathbb{R})$, and $C^1(\mathbb{R})$ is a vector space.

Vector spaces of functions, such as those introduced in Example 1.11, play a key role in many areas of mathematics, such as partial differential equations.

Subspaces, linear dependence and independence

Further reading: [Lax] Ch. 1 (p. 4–5); see also [Bee] p. 334–372; [CDW] Ch. 9–10 (p. 159–173); [Hef] Ch. 2 (p. 87–108); [LNS] Ch. 4–5 (p. 40–54); [Tre] Ch. 1 (p. 6–9, 30–31)

2.1 Deducing new properties from axioms

Last time we saw the general definition of a vector space in terms of a list of axioms. We also mentioned certain properties that follow immediately from these axioms: uniqueness of the zero element, and the fact that $0x = \mathbf{0}$ and $(-1)x = -x$. Let us briefly go through the proofs of these, to illustrate the use of the axioms in deriving basic properties.

1. Uniqueness of $\mathbf{0}$. Suppose $\mathbf{0}'$ also has the property that $x + \mathbf{0}' = x$ for all $x \in X$. Then in particular, this is true when $x = \mathbf{0}$, and so $\mathbf{0} + \mathbf{0}' = \mathbf{0}$. On the other hand, because $\mathbf{0}$ has the property that $y + \mathbf{0} = y$ for all $y \in X$, we may in particular choose $y = \mathbf{0}'$ and deduce that $\mathbf{0}' + \mathbf{0} = \mathbf{0}'$. Finally, by commutativity of addition we deduce that $\mathbf{0} = \mathbf{0} + \mathbf{0}' = \mathbf{0}' + \mathbf{0} = \mathbf{0}'$, and so the zero element is unique.
2. We prove that $0 \cdot x = \mathbf{0}$ for all $x \in X$. To this end, let $x \in X$ be arbitrary, and make the following deductions:

$$x = 1 \cdot x = (1 + 0) \cdot x = 1 \cdot x + 0 \cdot x = x + 0 \cdot x. \quad (2.1)$$

The first and last equalities use the final axiom (multiplication by the unit), the second equality uses properties of real numbers, and the third equality uses the axiom of distributivity. Now by the axiom on existence of additive inverses, we can add $(-x)$ to both sides and get

$$\begin{aligned} \mathbf{0} &= x + (-x) = (x + 0 \cdot x) + (-x) = (0 \cdot x + x) + (-x) \\ &= 0 \cdot x + (x + (-x)) = 0 \cdot x + \mathbf{0} = 0 \cdot x, \end{aligned} \quad (2.2)$$

where the first equality is the property of additive inverses, the second is from (2.1), the third is from commutativity of addition, the fourth is from associativity of addition, the fifth is the property of additive inverses again, and the last equality is the property of the zero vector.

3. We prove that $(-1) \cdot x = -x$ for all $x \in X$. To this end, we first observe that the additive inverse is unique: if $x + y = \mathbf{0}$, then $y = -x$. Indeed, adding $(-x)$ to both sides gives

$$\begin{aligned} -x &= \mathbf{0} + (-x) = (x + y) + (-x) \\ &= (y + x) + (-x) = y + (x + (-x)) = y + \mathbf{0} = y, \end{aligned}$$

where the first equality uses the axiom on the zero vector, the second comes from the equality $x + y = \mathbf{0}$, the third uses commutativity of addition, the fourth uses associativity of addition, the fifth uses the property of additive inverses, and the last once again uses the property of the zero vector. Armed with this fact on uniqueness, we can now observe that

$$x + (-1)x = 1 \cdot x + (-1) \cdot x = (1 + (-1)) \cdot x = 0 \cdot x = \mathbf{0},$$

where the first equality is from the axiom on multiplication by the unit, the second equality is from the distributivity axiom, the third is from properties of real numbers, and the fourth is what we proved just a moment ago in (2.2). Because additive inverses are unique, it follows that $(-1) \cdot x = -x$.

The arguments above are rather painstaking and difficult to read, but they illustrate the procedure of deducing other general facts from the small handful of axioms with which we begin. From now on we will not usually give explicit references to which axioms are used in any given computation or argument, but you should always keep in mind that every step of a calculation or proof needs to be justified in terms of previous results, which are ultimately based on these axioms.

2.2 Subspaces

Let's move on to something a little less bland and more concrete. Recalling our examples from the previous lecture, we see that it is often the case that one vector space is contained inside another one. For example, $\mathbb{P}_n \subset \mathbb{P}_{n+1}$. Or recall Example 1.11:

- $F(\mathbb{R}, \mathbb{R}) = \{\text{functions } \mathbb{R} \rightarrow \mathbb{R}\}$
- $C(\mathbb{R}) = \{f \in V \mid f \text{ is continuous}\}$
- $C^1(\mathbb{R}) = \{f \in V \mid f \text{ is differentiable}\}$

We have $C^1(\mathbb{R}) \subset C(\mathbb{R}) \subset F(\mathbb{R}, \mathbb{R})$. More generally, given $d \in \mathbb{N}$, we write $C^d(\mathbb{R})$ for the vector space of functions on \mathbb{R} that can be differentiated d times. Note that $C(\mathbb{R}) \supset C^1(\mathbb{R}) \supset C^2(\mathbb{R}) \supset \dots$.

Definition 2.1. When X and V are vector spaces with $X \subset V$, we say that X is a *subspace* of V .

- Example 2.2.**
1. $X_1 = \{(x, y) \in \mathbb{R}^2 \mid x + y = 0\}$ is a subspace of \mathbb{R}^2
 2. $X_2 = \{(x, y, z) \in \mathbb{R}^3 \mid z = 0\}$ is a subspace of \mathbb{R}^3 (the xy -plane)
 3. The set X_3 of solutions to $\ddot{x} = -x$ is a subspace of $C^2(\mathbb{R})$ – it is also a subspace of $C^1(\mathbb{R})$, $C(\mathbb{R})$, and $F(\mathbb{R}, \mathbb{R})$.
 4. $X_4 = \{f \in C^1(\mathbb{R}) \mid f \text{ is } 2\pi\text{-periodic}\}$ is a subspace of $C^1(\mathbb{R})$ – it is also a subspace of $C(\mathbb{R})$ and $F(\mathbb{R}, \mathbb{R})$. If you know something about solutions to ODEs, you will notice that in fact X_3 is a subspace of X_4 .

In each of these cases one can check that the operations of addition and multiplication from the *ambient* vector space (\mathbb{R}^2 , \mathbb{R}^3 , or $F(\mathbb{R}, \mathbb{R})$) define a vector space structure on the given subset, and so it is indeed a subspace. We omit the details of checking all the axioms, since we are about to learn a general fact that implies them.

Here is a convenient fact. In general, if we have a set X with two binary operations (addition and multiplication), and want to check that this is a vector space, we must verify the list of axioms from the previous lecture. When X is contained in a vector space V , life is easier: to check that a non-empty set $X \subset V$ is a subspace, we only need to check the following two conditions:

1. $x + y \in X$ whenever $x, y \in X$ (closure under addition)
2. $cx \in X$ whenever $x \in X$ and $c \in K$ (closure under scalar multiplication)

If these two conditions are satisfied, then the fact that the axioms from the previous lecture hold for X can be quickly deduced from the fact that they hold for V . For example, since addition is commutative for all pairs of elements in V , it is certainly commutative for all pairs of elements in the subset X . Similarly for associativity of addition and multiplication, distributivity, and multiplication by the unit. The only axioms remaining are existence of the identity element $\mathbf{0}$ and additive inverses. To get these, we recall from the previous lecture that $\mathbf{0} = 0x$ and $-x = (-1)x$ for any

$x \in V$. In particular, for any $x \in X$ the second condition just given implies that $\mathbf{0}, -x \in X$.

In fact, it is often convenient to combine the two conditions given above into the single following condition.

Proposition 2.3. *Let V be a vector space over K . A non-empty set $X \subset V$ is a subspace of V if and only if $cx + y \in X$ whenever $x, y \in X$ and $c \in K$.*

Proof. Exercise. □

Now the fact that the sets in Example 2.2 are subspaces of \mathbb{R}^2 , \mathbb{R}^3 , and V , respectively, can be easily checked by observing that each of these sets is closed under addition and scalar multiplication.

1. If (x, y) and (x', y') are in X_1 and $c \in \mathbb{R}$, then $(cx + x', cy + y')$ has $(cx + x') + (cy + y') = c(x + y) + (x' + y') = c \cdot 0 + 0 = 0$.
2. If $(x, y, z), (x', y', z') \in X_2$ and $c \in \mathbb{R}$, then $(cx + x', cy + y', cz + z')$ has third component equal to $cz + z' = c \cdot 0 + 0 = 0$, so it is in X_2 .
3. If $x, y: \mathbb{R} \rightarrow \mathbb{R}$ are in X_3 and $c \in \mathbb{R}$, then $\frac{d^2}{dt^2}(cx + y) = c\ddot{x} + \ddot{y} = -(cx + y)$, so $cx + y \in X_3$.
4. If $f, g \in X_4$ and $c \in \mathbb{R}$, then we check that $cf + g$ is 2π -periodic:

$$(cf + g)(t + 2\pi) = cf(t + 2\pi) + g(t + 2\pi) = cf(t) + g(t) = (cf + g)(t).$$

The first two examples should be familiar to you from a previous linear algebra course, since both X_1 and X_2 are the solution set of a system of linear equations (in fact, a single linear equation) in \mathbb{R}^n . What may not be immediately apparent is that X_3 and X_4 are in fact examples of exactly this same type: we define a condition which is linear in a certain sense, and then consider all elements of the vector space that satisfy this condition. This will be made precise later when we consider null spaces (or kernels) of linear transformations.

Remark 2.4. All of the subspaces considered above are described *implicitly*; that is, a condition is given, and then the subspace X is the set of all elements of V that satisfy this condition. Thus if I give you an element of V , it is easy for you to check whether or not this element is contained in the subspace X . On the other hand, it is not necessarily so easy for you to give me a concrete example of an element of X , or a method for producing all the elements of X . Such a method would be an *explicit* description of X , and the process of solving linear equations may be thought of as the process of going from an implicit to an explicit description of X .

Plenty of interesting subsets of vector spaces are *not* subspaces.

Example 2.5. 1. $X = \{(x, y) \in \mathbb{R}^2 \mid x, y \in \mathbb{Z}\}$, the integer lattice in \mathbb{R}^2 , is not a subspace of \mathbb{R}^2 , because it is not closed under multiplication: $\frac{1}{2}(1, 1) = (\frac{1}{2}, \frac{1}{2}) \notin X$

2. With V the set of differentiable functions on $(-1, 1)$, the set $X = \{x \in V \mid \dot{x} = x^2\}$ is not a subspace of V : it is not closed under addition or scalar multiplication. Indeed, the function $x(t) = (2-t)^{-1}$ is contained in X , but $2x$ is not.

Every vector space V has at least two subspaces: V itself is a subspace, and so is the set $\{\mathbf{0}\}$ that contains only the zero vector (check this!) – this is called the *trivial subspace*.

2.3 New subspaces from old ones

Above we pointed out that many examples of subspaces are obtained as solution sets of systems of linear equations. If $X \subset \mathbb{R}^n$ is the solution set of a system of linear equations, and $Y \subset \mathbb{R}^n$ is the solution set of another system, then $X \cap Y$ is the solution set of the system obtained by taking all equations in these two systems. In particular, $X \cap Y$ is a subspace. This is a general fact that is true in any vector space, not just \mathbb{R}^n .

Exercise 2.6. Prove that if X and Y are subspaces of V , then so is $X \cap Y$.

Thus we can form new subspaces from old ones by taking intersections. We can also form new subspaces by taking sums.

Definition 2.7. If X, Y are two subsets of a vector space V (not necessarily subspaces), then the *sum* of X and Y is $X + Y = \{x + y \mid x \in X, y \in Y\}$.

Example 2.8. 1. Let $X = \mathbb{Z}^2 \subset \mathbb{R}^2$ and $Y = \{(x, y) \mid x^2 + y^2 \leq 1/9\}$. Then $X + Y$ is the set containing a ball of radius $1/3$ around every point with integer coordinates.

2. Let $X = \{(x, 0, 0) \in \mathbb{R}^3 \mid x \in \mathbb{R}\}$ be the x -axis and $Y = \{(0, y, 0) \mid y \in \mathbb{R}\}$ be the y -axis. Then $X + Y = \{(x, y, 0) \mid x, y \in \mathbb{R}\}$ is the xy -plane.

Exercise 2.9. Prove that if X and Y are subspaces of V , then so is $X + Y$.

So intersections and sums of subspaces are subspaces. *The same is not true of unions.*

Exercise 2.10. Give an example of subspaces X and Y of \mathbb{R}^2 such that $X \cup Y$ is not a subspace.

You have already encountered the idea of taking sums of subspaces in a very specific guise, namely taking linear combinations of vectors in \mathbb{R}^n . The notion of linear combination extends naturally to an arbitrary vector space.

Definition 2.11. Given vectors $x_1, \dots, x_k \in V$, a *linear combination* of x_1, \dots, x_k is a vector of the form

$$\sum_{j=1}^k c_j x_j = c_1 x_1 + \dots + c_k x_k,$$

where $c_1, \dots, c_k \in K$ are the *coefficients* of the linear combination. The set of all linear combinations of x_1, \dots, x_k is the *span* of x_1, \dots, x_k , denoted $\text{span}(x_1, \dots, x_k)$, or sometimes $\langle x_1, \dots, x_k \rangle$.

Exercise 2.12. Show that for every choice of x_1, \dots, x_k , the set $\text{span}(x_1, \dots, x_k)$ is a subspace of V .

You should notice that the mechanics of proving Exercises 2.9 and 2.12 are very similar to each other. Indeed, one can observe the following.

1. Given $x_j \in V$, the set $X_j = \text{span}(x_j) = \{cx_j \mid c \in K\}$ is a subspace of V .
2. $\text{span}(x_1, \dots, x_k) = X_1 + \dots + X_k$.

2.4 Spanning sets and linear dependence

We say that the set $\{x_1, \dots, x_k\}$ *spans* V , or *generates* V , if $\text{span}(x_1, \dots, x_k) = V$. In particular, a set spans V if every element of V can be written as a linear combination of elements from that set.

Remark 2.13. If X is a subspace of V and $X = \text{span}(x_1, \dots, x_k)$, then it is reasonable to think of the set $\{x_1, \dots, x_k\}$ as giving an *explicit* description of the subspace X , since it gives us a concrete method to produce all the elements of X : just consider the vectors $c_1 x_1 + \dots + c_k x_k$, where the coefficients c_1, \dots, c_k take arbitrary values in the scalar field K . The process of solving a system of linear equations by row reduction gives a way to find a ‘reasonable’ spanning set for the subspace of solutions to that system. The question of finding a ‘reasonable’ spanning set for the solution space of a linear differential equation, such as X_3 in Example 2.2, can be rather more challenging to carry out, but from the point of view of linear algebra is a very similar task.

What does the word ‘reasonable’ mean in the above remark? Well, consider the following example. The vector space \mathbb{R}^2 is spanned by the set $\{(1, 0), (0, 1), (1, 1)\}$. But this set is somehow redundant, since the smaller set $\{(1, 0), (0, 1)\}$ also spans \mathbb{R}^2 . So in finding spanning sets, it makes sense to look for the smallest one, which is somehow the most efficient. You should recall from your previous linear algebra experience that in the case of \mathbb{R}^n , the following condition is crucial.

Definition 2.14. A set $\{x_1, \dots, x_k\}$ is *linearly dependent* if some non-trivial linear combination yields the zero vector: that is, if there are scalars $c_1, \dots, c_k \in K$ such that not all of the c_j are 0, and $c_1x_1 + \dots + c_kx_k = \mathbf{0}$. A set is *linearly independent* if it is not linearly dependent.

Recall how the notions of span and linear dependence manifest themselves in \mathbb{R}^n . Given a set $S = \{x_1, \dots, x_k\}$, the task of checking whether or not $x \in \text{span } S$ reduces to solving the non-homogeneous system of linear equations $c_1x_1 + \dots + c_kx_k = x$. If this system has a solution, then $x \in \text{span } S$. If it has no solution, then $x \notin \text{span } S$. (Note that $\text{span } S$ is an *explicitly* described subspace, and now the difficult task is to check whether or not a specific vector is included in the subspace, which was the easy task when the subspace was implicitly defined.)

Similarly, you can check for linear dependence in \mathbb{R}^n by writing the condition $c_1x_1 + \dots + c_kx_k = \mathbf{0}$ as a system of linear equations, and using row reduction to see if it has a non-trivial solution. A non-trivial solution corresponds to a non-trivial representation of $\mathbf{0}$ as a linear combination of vectors in S , and implies that S is linearly dependent. If there is only the trivial solution, then S is linearly independent.

Example 2.15. Consider the polynomials $f_1(x) = x + 1$, $f_2(x) = x^2 - 2$, and $f_3(x) = x + 3$ in \mathbb{P}_2 . These span \mathbb{P}_2 and are linearly independent. Indeed, given any polynomial $g \in \mathbb{P}_2$ given by $g(x) = a_2x^2 + a_1x + a_0$, we can try to write $g = c_1f_1 + c_2f_2 + c_3f_3$ by solving

$$a_2x^2 + a_1x + a_0 = c_1(x + 1) + c_2(x^2 - 2) + c_3(x + 3).$$

Comparing coefficients we see that this is equivalent to the system

$$\begin{aligned} a_2 &= c_2, \\ a_1 &= c_1 + c_3, \\ a_0 &= c_1 - 2c_2 + 3c_3, \end{aligned} \tag{2.3}$$

which can be easily checked to have a solution (c_1, c_2, c_3) for every choice of a_1, a_2, a_3 . Similarly, the homogeneous version of this system has only

the trivial solution, which shows that the polynomials f_1, f_2, f_3 are linearly independent.

The previous example could be done by reducing it to the familiar case of systems of linear equations. This is not always possible.

Example 2.16. Let V be the vector space of differentiable functions $\mathbb{R} \rightarrow \mathbb{R}$, and let $f(x) = \sin^2(x)$, $g(x) = \cos(2x)$, and $h(x) = 1$. Then $\{f, g, h\}$ is linearly dependent, because the trigonometric identity $\cos(2x) = 1 - 2\sin^2(x)$ implies the non-trivial linear representation of the zero function as $\mathbf{0} = h - 2f - g$.

Bases.

Further reading: [Lax] Ch. 1 (p. 5–7); see also [Bee] p. 373–398; [CDW] Ch. 11 (p. 175–182); [Hef] Ch. 2 (p. 109–137); [LNS] Ch. 5 (p. 54–59); [Tre] Ch. 1,2 (p. 6–7, 54–56)

3.1 More on spanning sets and linear dependence

The definitions of linear dependence, independence, and spanning in the previous lecture are only made for *finite* sets. It is useful to extend the definition to infinite sets, and this is done as follows.

Definition 3.1. A set $S \subset V$ is *linearly dependent* if there are $x_1, \dots, x_n \in S$ and non-zero scalars $c_1, \dots, c_n \in K$ such that $\sum_{j=1}^n c_j x_j = \mathbf{0}$. If S is not linearly dependent, it is said to be *linearly independent*. The *span* of a set $S \subset V$ is the set of all linear combinations of any finite set of elements of S .

Exercise 3.2. Show that an infinite set S is linearly dependent if and only if it has a finite subset S' that is linearly dependent. Show that an infinite set S is linearly independent if and only if every finite subset of S is linearly independent.

Exercise 3.3. Let V be a vector space, let $S \subset V$ be spanning, and let $L \subset V$ be linearly independent.

1. Show that if $S \subset S' \subset V$, then S' is spanning.
2. Show that if $L' \subset L$, then L' is linearly independent.

The next result is quite important: it says that linearly independent sets cannot be bigger than spanning sets.

Proposition 3.4. Let V be a vector space, let $S = \{x_1, \dots, x_n\} \subset V$ span V , and let $L = \{y_1, \dots, y_k\} \subset V$ be linearly independent. Then $k \leq n$.

Proof. We show by induction that for every $0 \leq j \leq \min(n, k)$, we can replace j elements of S with elements of L to obtain a new spanning set S' . In particular, if $k > n$ then this statement with $j = n$ implies that L contains a spanning subset L' of size n , so that the elements of $L \setminus L'$ can be written as linear combinations of L' , contradicting linear independence of L . Thus it suffices to complete the induction just described.

The base case of the induction is $j = 0$, which is trivial. For the induction step, relabel the elements of S and L so that $S' = \{y_1, \dots, y_j, x_{j+1}, \dots, x_n\}$. Because S' spans V , we can write

$$y_{j+1} = c_1 y_1 + \dots + c_j y_j + c_{j+1} x_{j+1} + \dots + c_n x_n$$

for some coefficients c_i . If $c_i = 0$ for all $j + 1 \leq i \leq n$, then we have a contradiction to the linear independence of $\{y_1, \dots, y_{j+1}\} \subset L$. Thus there is some i with $j < i \leq n$ such that x_i can be written as a linear combination of $\{y_1, \dots, y_{j+1}\}$ and the other x_ℓ . In particular, replacing x_i with y_{j+1} in S' yields the desired spanning set. This completes the induction step, and hence completes the proof by the discussion above. \square

In addition to the definition given above, there are other characterisations of linear dependence that are often useful.

Proposition 3.5. *Let V be a vector space over K , and let $S \subset V$. The following are equivalent.*

1. S is linearly dependent.
2. There exists $v \in S$ such that $v \in \text{span}(S \setminus \{v\})$.
3. There exists $v \in S$ such that $\text{span}(S \setminus \{v\}) = \text{span } S$.

Proof. (1) \Rightarrow (2). If S is linearly dependent, then there exist $x_1, \dots, x_n \in S$ and $c_1, \dots, c_n \in K$ such that $c_1 \neq 0$ and

$$c_1 x_1 + c_2 x_2 + \dots + c_n x_n = \mathbf{0}.$$

Subtracting $c_1 x_1$ from both sides and multiplying by c_1^{-1} (since $c_1 \neq 0$) gives

$$x_1 = -\frac{c_2}{c_1} x_2 - \dots - \frac{c_n}{c_1} x_n,$$

which proves (2) with $v = x_1$.

(2) \Rightarrow (3). The inclusion $\text{span}(S \setminus \{v\}) \subset \text{span } S$ is immediate, so we only need to prove the other inclusion. If $v = c_1 v_1 + \dots + c_n v_n$ for $v_1, \dots, v_n \in S \setminus \{v\}$, then for every $x \in \text{span } S$, there are $a, b_1, \dots, b_n \in K$ and $w_1, \dots, w_n \in S \setminus \{v\}$ such that

$$x = av + b_1 w_1 + \dots + b_n w_n = ac_1 v_1 + \dots + ac_n v_n + b_1 w_1 + \dots + b_n w_n \in \text{span}(S \setminus \{v\}).$$

(3) \Rightarrow (1). If $\text{span } S = \text{span}(S \setminus \{v\})$, then $v \in \text{span } S$ implies that $v \in \text{span}(S \setminus \{v\})$ (so (2) holds), and in particular there exist $x_1, \dots, x_n \in S \setminus \{v\}$ and $c_1, \dots, c_n \in K$ such that $v = c_1x_1 + \dots + c_nx_n$. Rearranging gives

$$c_1x_1 + \dots + c_nx_n - v = \mathbf{0},$$

which is a non-trivial linear representation of $\mathbf{0}$ in terms of the elements of S , so S is linearly dependent. \square

Proposition 3.6. *Let V be a vector space over K , and let $S \subset V$. The following are equivalent.*

1. S is linearly independent.
2. Every element of $\text{span } S$ can be represented in a unique way as a linear combination of elements of S .

Proof. (1) \Rightarrow (2). Suppose there is $x \in \text{span } S$ such that

$$x = a_1v_1 + \dots + a_nv_n = b_1w_1 + \dots + b_mw_m$$

for some $a_i, b_j \in K$ and $v_i, w_j \in S$. Without loss of generality (adding some zero coefficients if need be), assume that $m = n$ and $v_i = w_i$ for all $1 \leq i \leq n$, so

$$x = a_1v_1 + \dots + a_nv_n = b_1v_1 + \dots + b_nv_n.$$

This implies that $(a_1 - b_1)v_1 + \dots + (a_n - b_n)v_n = \mathbf{0}$, and by linear independence of S we get $a_i = b_i$ for each i , so the two representations of x as a linear combination of elements of S were in fact the same representation.

(2) \Rightarrow (1). Observe that $\mathbf{0} \in \text{span } S$ can be represented as $\mathbf{0} = 0 \cdot v$ for any $v \in S$, and by uniqueness there is no non-trivial representation of $\mathbf{0}$ as a linear combination of elements of S , so S is linearly independent. \square

Exercise 3.7. Let $L \subset V$ be linearly independent and let $v \in V \setminus L$. Show that $L \cup \{v\}$ is linearly independent if and only if $v \notin \text{span } L$.

3.2 Bases and dimension

As a preview of the next lecture, we can now give two of the fundamental definitions for the course.

Definition 3.8. A set $\{x_1, \dots, x_n\} \subset V$ is a *basis* for V if it spans V and is linearly independent.

Bases are far from unique – there are many choices of basis possible for a finite-dimensional space, as illustrated in Theorem 4.7 below. Sometimes, as in \mathbb{R}^n , there is a natural one, which is usually referred to as the *standard basis* for that vector space: in \mathbb{R}^n , we write \mathbf{e}_i for the vector whose i th component is 1, with all other entries 0.

Theorem 3.9. *If V has a finite basis, then all bases for V have the same number of vectors.*

Proof. Let $B, B' \subset V$ be bases for V . Then B is linearly independent and B' is spanning, so Proposition 3.4 implies that $\#B \leq \#B'$. Similarly, B is spanning and B' is independent, so $\#B' \leq \#B$. \square

Dimension, direct sums, and isomorphisms

Further reading: [Lax] Ch. 1 (p. 5–7); see also [Bee] p. 373–398; [CDW] Ch. 11 (p. 175–182); [Hef] Ch. 2 (p. 109–137); [LNS] Ch. 5 (p. 54–59); [Tre] Ch. 1,2 (p. 6–7, 54–56)

4.1 Dimension

In the last lecture, we hinted that it is useful to have not just spanning sets, but spanning sets that contain a minimum amount of redundancy. The definition of basis makes this precise. The spanning property guarantees that every $v \in V$ can be written as a linear combination $v = c_1x_1 + \cdots + c_nx_n$ of the basis elements, while the linear independence property together with Proposition 3.6 guarantees that the coefficients c_1, \dots, c_n are determined *uniquely* by v and x_1, \dots, x_n . We will come back to this important point later in the lecture..

Exercise 4.1. Show that if $\{x, y\}$ is a basis for X , then so is $\{x + y, x - y\}$.

Definition 4.2. A vector space with a finite basis is called *finite-dimensional*.

Example 4.3. \mathbb{P} , the vector space of all polynomials with coefficients in \mathbb{R} , is not finite-dimensional. Indeed, if $\{f_1, \dots, f_n\}$ is any finite collection of polynomials, then we may let $d = \max_{1 \leq j \leq n} \deg(f_j)$ and consider the polynomial $g(x) = x^{d+1}$. Because every polynomial in $\text{span}\{f_1, \dots, f_n\}$ has degree at most d , we see that g is not in this span, and hence $\{f_1, \dots, f_n\}$ does not span \mathbb{P} .

Remark 4.4. It is possible to consider infinite bases: there is a sense in which the infinite set $\{1, x, x^2, x^3, \dots\}$ is a basis for \mathbb{P} . However, we will not consider these in this course, and for our purposes all bases will be finite, and all results concerning bases of vector spaces will be given for finite-dimensional vector spaces.

Lemma 4.5. *If V has a finite spanning set then it has a finite basis.*

Proof. Let S be a finite spanning set for V . If S is linearly independent then it is a basis, and we are done. If S is linearly dependent then by Proposition 3.5 there is a proper subset $S_1 \subset S$ such that $\text{span } S_1 = \text{span } S = V$. If S_1 is linearly independent then it is our basis, otherwise we repeat the procedure

to obtain S_2, S_3, \dots . Because $\#S_k \leq \#S - k$, the procedure must terminate after at most $\#S$ steps (recall that S is finite), and so we eventually obtain a finite basis. \square

Definition 4.6. The numbers of vectors in a basis for V is the *dimension* of V , denoted $\dim V$. By convention, $\dim\{\mathbf{0}\} = 0$.

Theorem 4.7. *If V is finite-dimensional, then every linearly independent set y_1, \dots, y_n can be completed to a basis for V .*

Proof. If $\{y_1, \dots, y_n\}$ spans V , then it is a basis and we are done. If it does not span, then there is $v \in V \setminus \text{span}\{y_1, \dots, y_n\}$, and by Exercise 3.7 the set $\{y_1, \dots, y_n, v\}$ is linearly independent. If this set spans then we are done, otherwise we repeat the procedure. Eventually we have a linearly independent set with $\dim V$ elements. If this set does not span, the same procedure would give us a linearly independent set with $1 + \dim V$ elements, contradicting Theorem 3.9. \square

4.2 Direct sums

Proposition 3.6 related linear independence to uniqueness of representation for elements in the span of a set; this is the key distinguishing feature between a spanning set and a basis. A similar distinction is drawn when we consider the sum of subspaces. Recall that if $Y_1, \dots, Y_m \subset V$ are subspaces such that every element $x \in V$ can be written as $x = y_1 + \dots + y_m$ for some $y_i \in Y_i$, then we say that V is the *sum* of Y_1, \dots, Y_m and write $V = Y_1 + \dots + Y_m$. If y_1, \dots, y_m are determined *uniquely* by x , then we say that V is the *direct sum* of Y_1, \dots, Y_m , and write

$$V = Y_1 \oplus \dots \oplus Y_m. \quad (4.1)$$

Exercise 4.8. Show that V is the direct sum of Y_1, Y_2 if and only if $V = Y_1 + Y_2$ and $Y_1 \cap Y_2 = \{\mathbf{0}\}$.

Exercise 4.9. Give an example of three subspaces $Y_1, Y_2, Y_3 \subset \mathbb{R}^2$ such that $Y_1 + Y_2 + Y_3 = \mathbb{R}^2$ and $Y_i \cap Y_j = \{\mathbf{0}\}$ for every $i \neq j$, but \mathbb{R}^2 is not the direct sum of Y_1, Y_2, Y_3 .

Exercise 4.10. Show that if $V = Y_1 + \dots + Y_m$, then the sum is a direct sum if and only if the subspaces Y_1, \dots, Y_m are “linearly independent” in the following sense: whenever $y_j \in Y_j$ are such that $\sum_{j=1}^m y_j = \mathbf{0}$, we have $y_j = \mathbf{0}$ for every j .

Theorem 4.11. *Let V be a finite-dimensional vector space and let $W \subset V$ be a subspace. Then W is finite-dimensional and has a complement: that is, another subspace $X \subset V$ such that $V = W \oplus X$.*

Proof. We can build a finite basis for W as follows: take any nonzero vector $w_1 \in W$, so $\{w_1\}$ is linearly independent. If $\{w_1\}$ spans W then we have found our basis; otherwise there exists $w_2 \in W \setminus \text{span}\{w_1\}$. By Exercise 3.7, the set $\{w_1, w_2\}$ is once again linearly independent. We continue in this manner, obtaining linearly independent sets $\{w_1, \dots, w_k\}$ until we find a basis.

As it stands, the above argument is *not yet a proof*, because it is possible that the procedure continues indefinitely: a priori, it could be the case that we get linearly independent sets $\{w_1, \dots, w_k\}$ for every $k = 1, 2, 3, \dots$, without ever obtaining a spanning set. However, because the ambient subspace V is finite-dimensional, we see that the procedure must terminate: by Proposition 3.4, $\{w_1, \dots, w_k\}$ cannot be linearly independent when $k > \dim V$, and so there exists k such that the procedure terminates and gives us a basis for W .

This shows that W is finite-dimensional. To show that it has a complement, we observe that by Theorem 4.7, the set $\{w_1, \dots, w_k\}$, which is a basis for W and hence linearly independent, can be completed to a basis $\beta = \{w_1, \dots, w_k, x_1, \dots, x_m\}$ for V . Let $X = \text{span}\{x_1, \dots, x_m\}$. We claim that every element $v \in V$ can be written in a unique way as $v = w + x$, where $w \in W$ and $x \in X$. To show this, observe that because β is a basis for V , there are unique scalars $a_1, \dots, a_k, b_1, \dots, b_m \in K$ such that

$$v = (a_1w_1 + \dots + a_kw_k) + (b_1x_1 + \dots + b_mx_m).$$

Let $w = \sum_j a_j w_j$ and $x = \sum_i b_i x_i$. Then $v = w + x$. Moreover, by Exercise 4.8 it suffices to check that $W \cap X = \{\mathbf{0}\}$. To see this, suppose that there are coefficients a_j, b_i such that $\sum a_j w_j = \sum b_i x_i$. Bringing everything to one side gives $\sum_j a_j w_j + \sum_i (-b_i) x_i = \mathbf{0}$, and by linear independence of β we have $a_j = 0$ and $b_i = 0$ for each i, j . \square

Exercise 4.12. Show that V is the direct sum of Y_1, \dots, Y_m if and only if $V = Y_1 + \dots + Y_m$ and $\dim V = \sum_{j=1}^m \dim Y_j$.

Quotient spaces and dual spaces

Further reading: [Lax] Ch. 1–2 (p. 7–15); see also [Tre] Ch. 8 (p. 207–214)

5.1 Quotient spaces

Let V be a vector space over K , and $Y \subset V$ a subspace. We say that $v_1, v_2 \in V$ are *congruent modulo Y* if $v_1 - v_2 \in Y$, and write $v_1 \equiv v_2 \pmod{Y}$.

Example 5.1. Let $V = \mathbb{R}^2$ and $Y = \{(x, y) \mid x + y = 0\}$. Then $(x_1, y_1) \equiv (x_2, y_2) \pmod{Y}$ if and only if $(x_1 - x_2, y_1 - y_2) \in Y$; that is, if and only if $x_1 - x_2 + y_1 - y_2 = 0$. We can rewrite this condition as $x_1 + y_1 = x_2 + y_2$.

Exercise 5.2. Show that congruence mod Y is an equivalence relation – that is, it satisfies the following three properties.

1. Symmetric: if $v_1 \equiv v_2 \pmod{Y}$, then $v_2 \equiv v_1 \pmod{Y}$.
2. Reflexive: $v \equiv v \pmod{Y}$ for all $v \in V$.
3. Transitive: if $v_1 \equiv v_2 \pmod{Y}$ and $v_2 \equiv v_3 \pmod{Y}$, then $v_1 \equiv v_3 \pmod{Y}$.

Furthermore, show that if $v_1 \equiv v_2 \pmod{Y}$, then $cv_1 \equiv cv_2 \pmod{Y}$ for all $c \in K$.

Given $v \in V$, let $[v]_Y = \{w \in V \mid w \equiv v \pmod{Y}\}$ be the *congruence class* of v modulo Y . This is also sometimes called the *coset* of Y corresponding to v .

Exercise 5.3. Given $v_1, v_2 \in V$, show that $[v_1]_Y = [v_2]_Y$ if $v_1 \equiv v_2 \pmod{Y}$, and $[v_1]_Y \cap [v_2]_Y = \emptyset$ otherwise.

In example 5.1, we see that $[v]_Y$ is the line through v with slope -1 . Thus every congruence class modulo Y is a line parallel to Y . In fact we see that $[v]_Y = v + Y = \{v + x \mid x \in Y\}$. This fact holds quite generally.

Proposition 5.4. *Let $Y \subset V$ be a subspace, then for every $v \in V$ we have $[v]_Y = v + Y$.*

Proof. (\subset): If $w \in [v]_Y$, then $w - v \in Y$, and so $w = v + (w - v) \in v + Y$.

(\supset): If $w \in v + Y$, then $w = v + y$ for some $y \in Y$, and so $w - v = y \in Y$, whence $w \equiv v \pmod{Y}$, so $w \in [v]_Y$. \square

We see from this result that every congruence class modulo a subspace is just a copy of that subspace shifted by some vector. Such a set is called an *affine subspace*.

Recall once again our definition of setwise addition, and notice what happens if we add two congruence classes of Y : given any $v, w \in V$, we have

$$[v]_Y + [w]_Y = (v+Y) + (w+Y) = \{v+w+y_1+y_2 \mid y_1, y_2 \in Y\} = (v+w) + (Y+Y)$$

But since Y is a subspace, we have $Y + Y = \{y_1 + y_2 \mid y_1, y_2 \in Y\} = Y$, and so

$$[v]_Y + [w]_Y = (v + w) + Y = [v + w]_Y.$$

Similarly, we can multiply a congruence class by a scalar and get

$$c \cdot [v]_Y = c \cdot \{v + y \mid y \in Y\} = \{cv + cy \mid y \in Y\} = cv + Y = [cv]_Y.$$

(If these relations are not clear, spend a minute thinking about what they look like in the example above, where Y is the line through the origin with slope -1 .)

We have just defined a way to add two congruence classes together, and a way to multiply a congruence class by a scalar. One can show that these operations satisfy the axioms of a vector space (commutativity, associativity, etc.), and so the set of congruence classes forms a vector space over K when equipped with these operations. We denote this vector space by V/Y , and refer to it as the *quotient space* of $V \bmod Y$.

Exercise 5.5. What plays the role of the zero vector in V/Y ?

In Example 5.1, V/Y is the space of lines in \mathbb{R}^2 with slope -1 . Notice that any such line is uniquely specified by its y -intercept, and so this is a 1-dimensional vector space, since it is spanned by the single element $[(0, 1)]_Y$.

Example 5.6. Let $V = \mathbb{R}^n$ and $Y = \{(0, 0, 0, x_4, \dots, x_n) \mid x_4, \dots, x_n \in \mathbb{R}\}$. Then two vectors in V are congruent modulo Y if and only if their first three components are equal. Thus a congruence class is specified by its first three components, and in particular $\{[\mathbf{e}_1]_Y, [\mathbf{e}_2]_Y, [\mathbf{e}_3]_Y\}$ is a basis for V/Y .

The previous example also satisfies $\dim Y + \dim(V/Y) = \dim V$. In fact this result always holds.

Theorem 5.7. *Let Y be a subspace of a finite-dimensional vector space V . Then $\dim Y + \dim(V/Y) = \dim V$.*

Proof. Let y_1, \dots, y_k be a basis for Y . By Theorem 4.7, this can be completed to a basis for V by adjoining some vectors v_{k+1}, \dots, v_n . We claim that $\{[v_j]_Y \mid k+1 \leq j \leq n\}$ is a basis for V/Y .

First we show that it spans V/Y . Given $[v]_Y \in V/Y$, because $\{y_1, \dots, y_k, v_{k+1}, \dots, v_n\}$ is a basis for V , there exist $a_1, \dots, a_k, b_{k+1}, \dots, b_n \in K$ such that

$$v = \sum a_i y_i + \sum b_j v_j.$$

Adding the subspace Y to both sides gives

$$[v]_Y = v + Y = Y + \sum b_j v_j = \sum b_j [v_j]_Y.$$

This shows the spanning property. Now we check for linear independence by supposing that $c_{k+1}, \dots, c_n \in K$ are such that

$$\sum c_j [v_j]_Y = [\mathbf{0}]_Y.$$

This is true if and only if $\sum c_j v_j \in Y$, in which case there are d_i such that $\sum c_j v_j = \sum d_i y_i$. But now linear independence of the basis constructed above shows that $c_j = 0$ for all j . Thus $\{[v_j]_Y \mid k+1 \leq j \leq n\}$ is a basis for V/Y .

From this we conclude that $\dim(V/Y) = n - k$ and $\dim Y = k$, so that their sum is $(n - k) + k = n = \dim V$. \square

5.2 Dual spaces

The quotient spaces introduced in the previous section give one way to construct new vector spaces from old ones. Here is another.

Let V be a vector space over K . A scalar valued function $\ell: V \rightarrow K$ is called *linear* if

$$\ell(cx + y) = c\ell(x) + \ell(y) \tag{5.1}$$

for all $c \in K$ and $x, y \in V$. Note that repeated application of this property shows that

$$\ell\left(\sum_{i=1}^k c_i x_i\right) = \sum_{i=1}^k c_i \ell(x_i) \tag{5.2}$$

for every collection of scalars $c_i \in K$ and vectors $x_i \in V$.

The sum of two linear functions is defined by pointwise addition:

$$(\ell + m)(x) = \ell(x) + m(x).$$

Similarly, scalar multiplication is defined by

$$(c\ell)(x) = c(\ell(x)).$$

With these operations, the set of linear functions on V becomes a vector space, which we denote by V' and call the *dual* of V .

Example 5.8. Let $V = C(\mathbb{R})$ be the space of continuous real-valued functions on \mathbb{R} . For any $s \in \mathbb{R}$, define the function

$$\begin{aligned} \ell_s: V &\rightarrow \mathbb{R} \\ f &\mapsto f(s), \end{aligned}$$

which takes a function f to its value at a point s . This is a linear function, and so $\ell_s \in V'$. However, not every element of V' arises in this way. For example, one can define $\ell: V \rightarrow \mathbb{R}$ by

$$\ell(f) = \int_0^1 f(x) dx.$$

Integration over other domains gives other linear functionals, and still there are more possibilities: $C(\mathbb{R})'$ is a very big vector space.

Example 5.9. Let $V = C^1(\mathbb{R})$ be the space of all continuously differentiable real-valued functions on \mathbb{R} . Fix $s \in \mathbb{R}$ and let $\ell_s(f) = f'(s)$. Then $\ell_s \in V'$.

These examples are infinite-dimensional, and in this broader setting the notion of dual space becomes very subtle and requires more machinery to handle properly. In the finite-dimensional case, life is a little more manageable.

Theorem 5.10. *If V is a finite-dimensional vector space, then $\dim(V') = \dim(V)$.*

Proof. Let $\{v_1, \dots, v_n\}$ be a basis for V . Then every $v \in V$ has a unique representation as $v = \sum_i a_i v_i$, where $a_i \in K$. For each $1 \leq i \leq n$, define $\ell_i \in V'$ by

$$\ell_i(v) = \ell_i\left(\sum a_i v_i\right) = a_i.$$

(Alternately, we can define ℓ_i by $\ell_i(v_i) = 1$ and $\ell_i(v_j) = 0$ for $j \neq i$, then extend by linearity to all of V .) We leave as an exercise the fact that $\{\ell_1, \dots, \ell_n\}$ is a basis for V' . \square

Linear maps, nullspace and range

Further reading: [Lax] p. 19–20. See also [Bee] p. 513–581, [CDW] Ch. 6 (p. 89–96), [Hef] Ch. 3 (p. 173–179), [LNS] Ch. 6 (p. 62–81), [Tre] Ch. 1 (p. 12–18)

6.1 Linear maps, examples

Let V and W be vector spaces over the same field K , and $T: V \rightarrow W$ a map that assigns to each input $v \in V$ an output $T(v) \in W$. The vector space V in which the inputs live is the *domain* of the map T , and the vector space W in which the outputs live is the *codomain*, or *target space*.

Definition 6.1. A map $T: V \rightarrow W$ is a *linear transformation* (or a *linear map*, or a *linear operator*) if

1. $T(x + y) = T(x) + T(y)$ for every $x, y \in V$ (so T is *additive*), and
2. $T(cx) = cT(x)$ for every $x \in V$ and $c \in K$ (so T is *homogeneous*).

Exercise 6.2. Show that every linear map T has the property that $T(\mathbf{0}_V) = \mathbf{0}_W$.

Exercise 6.3. Show that T is linear if and only if $T(cx + y) = cT(x) + T(y)$ for every $x, y \in V$ and $c \in K$.

Exercise 6.4. Show that T is linear if and only if for every $x_1, \dots, x_n \in V$ and $a_1, \dots, a_n \in K$, we have

$$T\left(\sum_{i=1}^n a_i x_i\right) = \sum_{i=1}^n a_i T(x_i). \quad (6.1)$$

Example 6.5. The linear functionals discussed in the previous lecture are all examples of linear maps – for these examples we have $W = K$.

Example 6.6. The map $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by $T(x_1, x_2) = (x_2, -x_1 + 2x_2)$ is linear: given any $c \in \mathbb{R}$ and any $x = (x_1, x_2), y = (y_1, y_2) \in \mathbb{R}^2$, we have

$$\begin{aligned} T(cx + y) &= T(cx_1 + y_1, cx_2 + y_2) \\ &= (cx_2 + y_2, -(cx_1 + y_1) + 2(cx_2 + y_2)) \\ &= c(x_2, -x_1 + 2x_2) + (y_2, -y_1 + 2y_2) = cT(x) + T(y). \end{aligned}$$

Example 6.7. The map $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by $T(x_1, x_2) = (x_1 + x_2, x_2^2)$ is *not* linear. To see this, it suffices to observe that $T(0, 1) = (1, 1)$ and $T(0, 2) = (2, 4) \neq 2T(0, 1)$.

Example 6.8. The map $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by $T(x_1, x_2) = (1 + x_1, x_2)$ is *not* linear, because $T\mathbf{0} \neq \mathbf{0}$.

The above examples all look like the sort of thing you will have seen in your first linear algebra course: defining a function or transformation in terms of its coordinates, in which case checking for linearity amounts to confirming that the formula has a specific form, which in turn guarantees that it can be written in terms of matrix multiplication. The world is much broader than these examples, however: there are many examples of linear transformations which are most naturally defined using something other than a coordinate representation.

Example 6.9. Let $V = C^1(\mathbb{R})$ be the vector space of all continuously differentiable functions, and $W = C(\mathbb{R})$ the vector space of all continuous functions. Then $T: V \rightarrow W$ defined by $(Tf)(x) = \frac{d}{dx}f(x)$ is a linear transformation: recall from calculus that $(cf + g)' = cf' + g'$.

Example 6.10. Let $V = W = C(\mathbb{R})$, and define $T: V \rightarrow W$ by

$$(Tf)(x) = \int_0^1 f(y)(x - y)^2 dy.$$

Then $T(cf + g)(x) = \int_0^1 (cf + g)(y)(x - y)^2 dy = c \int_0^1 f(y)(x - y)^2 dy + \int_0^1 g(y)(x - y)^2 dy = (c(Tf) + (Tg))(x)$, and so T is linear.

The previous example illustrates the idea at the heart of the operation of *convolution*, a linear transformation that is used in many applications, including image processing, acoustics, data analysis, electrical engineering, and probability theory. A somewhat more complete description of this operation (which we will not pursue further in this course) is as follows: let $V = W = L^2(\mathbb{R}) = \{f: \mathbb{R} \rightarrow \mathbb{R} \mid \int_{-\infty}^{\infty} f(x)^2 dx < \infty\}$, and fix $g \in C(\mathbb{R})$. Then define $T: V \rightarrow W$ by $(Tf)(x) = \int_a^b f(y)g(x - y) dy$.

Example 6.11. Let $V = X = \mathbb{R}^2$, fix a number $\theta \in \mathbb{R}$, and let $T: V \rightarrow W$ be the map that rotates the input vector v by θ radians around the origin. The homogeneity property $T(cv) = cT(v)$ is easy to see, so to prove linearity of T it suffices to check the additivity property $T(v + w) = T(v) + T(w)$. This property can be proved in two different ways:

1. using the geometric definition of vector addition in \mathbb{R}^2 , where $v + w$ is defined as the diagonal of the parallelogram generated by v and w ;
2. using the observation that if $x = (x_1, x_2)$ makes an angle α with the positive x -axis, then $T(x)$ makes an angle $\alpha + \theta$, and so writing $x_1 = r \cos \alpha$, $x_2 = r \sin \alpha$, we get

$$\begin{aligned} T(x) &= (r \cos(\alpha + \theta), r \sin(\alpha + \theta)) \\ &= (r \cos \alpha \cos \theta - r \sin \alpha \sin \theta, r \sin \alpha \cos \theta + r \cos \alpha \sin \theta) \\ &= (x_1 \cos \theta - x_2 \sin \theta, x_2 \cos \theta + x_1 \sin \theta), \end{aligned}$$

which can be checked to define a linear map. Recall that the formulas for $\sin(\alpha + \theta)$ and $\cos(\alpha + \theta)$ can be recovered from the identities $e^{i\theta} = \cos \theta + i \sin \theta$ and $e^{i(\alpha+\theta)} = e^{i\alpha} e^{i\theta}$.

Example 6.12. Fix $m, n \in \mathbb{N}$ and $\{A_{ij} \in K \mid 1 \leq i \leq m, 1 \leq j \leq n\}$. Let $V = K^n$ and $W = K^m$. Define a map $T: V \rightarrow W$ by $T(v) = w$, where

$$w_i = \sum_{j=1}^n A_{ij} v_j \text{ for each } 1 \leq i \leq m. \quad (6.2)$$

The last example is, of course, the description of a linear transformation as multiplication by a matrix. We will come back to this in due course, and in particular will give a motivation for why matrix multiplication should be defined by a strange-looking formula like (6.2). First we discuss some general properties of linear transformations.

Proposition 6.13. *Suppose that the vectors $v_1, \dots, v_n \in V$ are linearly dependent, and that $T: V \rightarrow W$ is linear. Then $T(v_1), \dots, T(v_n) \in W$ are linearly dependent.*

Proof. By the hypothesis, there are $a_1, \dots, a_n \in K$ such that $\sum_i a_i v_i = \mathbf{0}_V$ and not all the a_i are zero. By Exercises 6.2 and 6.4, this implies that $\mathbf{0}_W = T(\mathbf{0}_V) = T(\sum_i a_i v_i) = \sum_i a_i T(v_i)$, and so $T(v_1), \dots, T(v_n)$ are linearly dependent. \square

Exercise 6.14. Suppose that the vectors $v_1, \dots, v_n \in V$ are linearly independent, and that $T: V \rightarrow W$ is 1-1 and linear. Show that the vectors $T(v_1), \dots, T(v_n) \in W$ are linearly independent.

Exercise 6.15. Let $S \subset V$ be a spanning set, and suppose that $T: V \rightarrow W$ is onto and linear. Show that $T(S) \subset W$ is a spanning set for W .

6.2 Null space and range

Let $T: V \rightarrow W$ be a linear map. Given any subset $X \subset V$, recall that $T(X) = \{T(x) \mid x \in X\}$ is the *image* of X under the action of T . Similarly, given any subset $Y \subset W$, recall that $T^{-1}(Y) = \{v \in V \mid T(v) \in Y\}$ is the *preimage* (or *inverse image*) of Y under T .

We stress that the notation $T^{-1}(Y)$ makes sense even when the map T is not invertible. If T is not 1-1, then T^{-1} should not be thought of as a map from W to V ; rather, it is a map from *subsets of W* to *subsets of V* . Thus given $w \in W$, $T^{-1}(w) = \{v \in V \mid T(v) = w\}$ is in general a *subset* of V , possibly with many elements.

Theorem 6.16. *Let $T: V \rightarrow W$ be a linear map.*

1. *If $X \subset V$ is a subspace of V , then $T(X)$ is a subspace of W .*
2. *if $Y \subset W$ is a subspace of W , then $T^{-1}(Y)$ is a subspace of V .*

Proof. Exercise. □

Definition 6.17. The *range* of T is $T(V) = \{T(v) \mid v \in V\}$. The *nullspace* (or *kernel*) of T is $T^{-1}(\mathbf{0}) = \{v \in V \mid T(v) = \mathbf{0}\}$. We denote the range of T by R_T and the nullspace by N_T .

By Theorem 6.16, both the range and nullspace of T are subspaces (of W and V , respectively – be very careful to remember where these things live). You have encountered both of these subspaces before, in your introduction to linear algebra: if $T: \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a linear transformation defined by coefficients A_{ij} as in Example 6.12, then the equation $T(v) = w$ describes a linear system of m equations (one equation for each w_i) in n variables (v_1, \dots, v_n) . The nullspace of T is the set of vectors v for which $T(v) = \mathbf{0}$: that is, the set of solutions of the homogeneous system of linear equations given by A_{ij} . The range R_T is the set of $w \in \mathbb{R}^m$ for which the non-homogeneous system $T(v) = w$ has a solution. This subspace of \mathbb{R}^m was also characterised as the column space of the matrix A .

Isomorphisms

Further reading: [Lax] p. 3, 7. See also [Bee] p. 544–603, [CDW] Ch. 6, 16 (p. 89–96, 237–246), [Hef] Ch. 3 (p. 157–172, 180–190), [LNS] Ch. 6 (p. 62–81), [Tre] Ch. 1 (p. 12–18)

7.1 More on nullspace and range

The discussion at the end of the previous lecture relates to the following observations. Nullspace is defined *implicitly*, meaning that given $v \in V$, it is easy to test whether $v \in N_T$: just compute $T(v)$ and see whether or not it is $\mathbf{0}_W$. On the other hand, it can be difficult to give an explicit description of all the elements of the nullspace: it was to accomplish precisely this task that you learned to use row reduction to solve homogeneous systems of linear equations. For the range R_T , the situation is reversed: this subspace is defined *explicitly*, and it is easy to produce all the elements of the range, by taking $T(v)$ for each $v \in V$. On the other hand, it can be difficult to determine whether or not a given $w \in W$ is in R_T : this is the same problem as determining whether or not the non-homogeneous system $T(v) = w$ has a solution, or whether w is in the span of the columns of the matrix A .

In the concrete setting of \mathbb{R}^n , problems regarding nullspace and range can be reduced to the familiar problems regarding systems of linear equations. This is not necessarily possible in the abstract setting.

Example 7.1. Consider the linear map $T: C^1(\mathbb{R}) \rightarrow C(\mathbb{R})$ given by differentiation. The nullspace of T is the space of all constant functions, while the range of T is all of $C(\mathbb{R})$. On the other hand, if we restrict our attention to the subspace $\mathbb{P}_n \subset C^1(\mathbb{R})$, then $T: \mathbb{P}_n \rightarrow \mathbb{P}_n$ has $N_T = \mathbb{P}_0$, the space of constant polynomials, while $R_T = \mathbb{P}_{n-1}$.

Example 7.2. Consider $T: C(\mathbb{R}) \rightarrow C(\mathbb{R})$ given by integration: $(Tf)(x) = \int_0^x f(y) dy$. Then N_T is trivial, and $R_T = \{g \in C^1(\mathbb{R}) \mid g(0) = 0\}$.

Example 7.3. Define a map $T: C^2(\mathbb{R}) \rightarrow C(\mathbb{R})$ by $(Tf) = f'' - \alpha f' + f$, where $\alpha \in \mathbb{R}$ is a damping parameter. The nullspace of T is the solution space of the unforced ODE $\ddot{x} - \alpha \dot{x} + x = 0$, and the range of T is the set of all forcing terms $g \in C(\mathbb{R})$ for which the forced ODE $\ddot{x} - \alpha \dot{x} + x = g(t)$ has a solution.

Exercise 7.4. Show that a linear map T is 1-1 if and only if $N_T = \{\mathbf{0}_V\}$.

7.2 Isomorphisms

An important kind of linear transformation is an *isomorphism*: we say that $T: V \rightarrow W$ is an isomorphism if it is linear, 1-1, and onto. If there is an isomorphism between V and W , we say that V and W are *isomorphic*.

Thus we may characterise isomorphisms in terms of nullspace and range as those linear maps for which $N_T = \{\mathbf{0}_V\}$ and $R_T = W$. (Null space is trivial and range is everything.)

Exercise 7.5. Show that if $T: V \rightarrow W$ is an isomorphism, then the inverse map $T^{-1}: W \rightarrow V$ is also an isomorphism.

The notion of isomorphism gives a sense in which various examples of vector spaces that we have seen so far are “the same”. For instance, let V be the vector space of column vectors with 2 real components, and let W be the vector space of row vectors with 2 real components. It seems clear that V and W are in some sense “the same”, in that both can be described by prescribing two numbers. More precisely, we can associate to each column vector $\begin{pmatrix} x \\ y \end{pmatrix}$ the corresponding row vector $(x \ y)$. This association respects vector addition and scalar multiplication:

1. $\begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} x+x' \\ y+y' \end{pmatrix}$ is associated to $(x+x' \ y+y') = (x \ y) + (x' \ y')$;
2. $c \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} cx \\ cy \end{pmatrix}$ is associated to $(cx \ cy) = c(x \ y)$.

Thus writing $T \begin{pmatrix} x \\ y \end{pmatrix} = (x \ y)$, we see that T is linear. Moreover, it is easy to show that T is an isomorphism by checking that $N_T = \{\mathbf{0}\}$ and $R_T = W$.

The key utility of this fact is that any statement concerning properties of V as a vector space can be translated into an equivalent statement concerning properties of W . For example, if we have a set of column vectors and want to know if they are linearly independent, it would suffice to answer the same question for the corresponding row vectors.

In fact, we have already used this principle in Example 2.15. In that example, we considered the polynomials $f_1(x) = x + 1$, $f_2(x) = x^2 - 2$, and $f_3(x) = x + 3$ in \mathbb{P}_2 . We observed that linear dependence or independence of the set $\{f_1, f_2, f_3\}$ could be determined by solving the homogeneous system of linear equations (2.3), which is equivalent to determining linear dependence or independence of the set $\left\{ \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ -2 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 3 \end{pmatrix} \right\}$ in \mathbb{R}^3 . The reason that this works is because we can associate to each column vector $\begin{pmatrix} a \\ b \\ c \end{pmatrix} \in \mathbb{R}^3$ the polynomial $ax^2 + bx + c$ in \mathbb{P}_2 . This association is 1-1, onto, and linear, hence it is an isomorphism.

So we have two examples of isomorphisms so far: the vector space of row vectors is isomorphic to the vector space of column vectors (with the same number of entries); and the vector space \mathbb{P}_n of polynomials with degree $\leq n$ is isomorphic to \mathbb{R}^{n+1} in a natural way, by associating $a_n x^n + \cdots + a_1 x + a_0$ to the vector with entries a_n, \dots, a_0 .

A less obvious example is the subspace $X_3 \subset C^2(\mathbb{R})$ from Example 2.2, defined as the set of all C^2 functions $x: \mathbb{R} \rightarrow \mathbb{R}$ such that $\ddot{x} + x = 0$. If you have taken a course in ODEs, you will recognise that the solutions are all of the form $x(t) = a \cos(t) + b \sin(t)$, where $a, b \in \mathbb{R}$ are chosen based on the initial conditions. In particular, by associating $(a, b) \in \mathbb{R}^2$ to the function $x(t) = a \cos(t) + b \sin(t)$, we obtain an isomorphism between X_3 and \mathbb{R}^2 .

All of the above examples of isomorphisms come about in the same way: given two vector spaces V and W of the same dimension n , produce a basis $\beta = \{v_1, \dots, v_n\}$ for V , and a basis $\gamma = \{w_1, \dots, w_n\}$ for W ; then given coefficients c_1, \dots, c_n , associate the vector $\sum c_j v_j \in V$ to the vector $\sum c_j w_j \in W$. In fact, this is a very general phenomenon.

Proposition 7.6. *If V, W are vector spaces over K and $T: V \rightarrow W$ is an isomorphism, then $\beta \subset V$ is a basis for V if and only if $T\beta$ is a basis for W .*

Proof. This follows from the exercises following Proposition 6.13: if β is a basis for V , then it is linearly independent, and since T is 1-1, $T\beta$ is linearly independent as well; similarly, β spans V , and since T is onto, $T\beta$ spans W , thus $T\beta$ is a basis. The reverse implication is proved similarly, using the result from Exercise 7.5 that T^{-1} is also an isomorphism. \square

A consequence of Proposition 7.6 is that isomorphic finite-dimensional vector spaces have the same dimension. The converse is also true. We need the following lemma.

Lemma 7.7. *Let V, W be finite-dimensional vector spaces over the same field K , and let $\{v_1, \dots, v_n\}$ be a basis for V . Then given any vectors $w_1, \dots, w_n \in W$, there is a unique linear transformation $T: V \rightarrow W$ such that $T(v_i) = w_i$ for each $1 \leq i \leq n$.*

Proof. Define T as follows: if $v \in V$ can be written as $v = \sum_i a_i v_i$ for $a_i \in K$, then $T(v) = \sum_i a_i w_i$. We must show that T is well-defined and linear. By *well-defined* we mean that there is only one possibility for $T(v)$; since the definition of $T(v)$ involved some coefficients $a_i \in K$, there could be more than one possible definition of $T(v)$ if there was more than one possible choice of a_i . However, because $\{v_1, \dots, v_n\}$ is a basis, the coefficients a_i are

determined *uniquely* by v , and so there is only one possible value of $T(v)$ that satisfies $T(v) = \sum_i a_i w_i$.

For linearity, we observe that if $v = \sum_i a_i v_i$ and $v' = \sum_i a'_i v_i$, then $cv + v' = \sum_i (ca_i + a'_i)v_i$ for all $c \in K$, and so

$$T(cv + v') = \sum_{i=1}^n (ca_i + a'_i)w_i = c \sum a_i w_i + \sum a'_i w_i = cT(v) + T(v').$$

Thus T is linear. Finally, we show that T is unique: indeed, if $U: V \rightarrow W$ is *any* linear transformation with $U(v_i) = w_i$ for every i , then by Exercise 6.4 we have $U(\sum_i a_i v_i) = \sum_i a_i U(v_i) = \sum_i a_i w_i = T(\sum_i a_i v_i)$, and so $U = T$. \square

Theorem 7.8. *Two finite-dimensional vector spaces over the same field K are isomorphic if and only if they have the same dimension.*

Proof. The forward implication (isomorphism implies equal dimension) follows from Proposition 7.6. For the reverse implication, we suppose that V and W are vector spaces over K with $\dim(V) = \dim(W) = n$, so that we can choose bases $\{v_1, \dots, v_n\}$ for V and $\{w_1, \dots, w_n\}$ for W . By Lemma 7.7, there is a unique linear transformation $T: V \rightarrow W$ with $T(v_i) = w_i$ for each i .

We claim that T is an isomorphism. Indeed, if $T(v) = \mathbf{0}_W$ then $v = \sum_i a_i v_i$ has $T(v) = \sum_i a_i w_i = \mathbf{0}_W$, and by linear independence of $\{w_1, \dots, w_n\}$ this implies $a_i = 0$ for each i , so $v = \mathbf{0}_V$, which shows that T is 1-1. To see that T is onto, note that any $w \in W$ can be written as $w = \sum_i a_i w_i$ for some $a_i \in K$, because of the spanning property, and so writing $v = \sum_i a_i v_i$, we have $T(v) = w$. \square

One way of looking at Theorem 7.8 is to say that up to isomorphism, K^n is the only n -dimensional vector space over K . If this is so, then why do we bother with the abstract point of view that has been presented so far? Why not just work with row vectors or as column vectors, since every finite-dimensional example can be reduced to these?

One answer is that there are also very interesting and important examples of infinite-dimensional vector spaces, which we cannot understand by comparing them to K^n . We have already seen some of these: $C(X)$, $C^1(X)$, and so on. Another answer is that there is often some benefit to having a *coordinate-free* presentation of a vector space, since there is not always a natural, readily-available basis to work with. For example, one may consider the set of solutions to a linear ODE of degree n ; this is a vector space of

dimension n , but there is no canonical basis with which to work. Similarly, if one considers all vectors in K^n whose components sum to 0, one obtains a vector space V of dimension $n - 1$, for which there are various choices of basis, but none that is obviously superior to all the others.

Indeed, given any two isomorphic vector spaces, there are *many* isomorphisms between them, and similarly every vector space has many different bases. This will become more and more clear as we go on.

More on linear transformations

Further reading: [Lax] p. 24–26

8.1 The algebra of linear maps

Let V and W be vector spaces over a common field K , and let $\mathbb{L}(V, W)$ be the set of linear maps from V to W . We can define addition and scalar multiplication on $\mathbb{L}(V, W)$ in the natural way:

$$(T + S)(v) = T(v) + S(v), \quad (cT)(v) = c(T(v))$$

for every $c \in K$ and $S, T \in \mathbb{L}(V, W)$. It is not hard to check that $T + S$ and cT are also linear maps, and that these operations make $\mathbb{L}(V, W)$ into a vector space.

There is more to linear maps than just a vector space structure, however, because in addition to adding linear maps and multiplying them by scalars, we can *compose* them, at least under certain conditions. Recall that if V, W, X are any sets (not necessarily vector spaces) and $f: V \rightarrow W$ and $g: W \rightarrow X$ are any maps (not necessarily linear), then the *composition* $g \circ f: V \rightarrow X$ is defined by $(g \circ f)(v) = g(f(v))$ – that is, the output of f is fed into g as its input. Schematically we can write

$$X \xleftarrow{g} W \xleftarrow{f} V. \quad (8.1)$$

Composition is always *associative*, if we compose three maps with compatible inputs/outputs, such as

$$Y \xleftarrow{h} X \xleftarrow{g} W \xleftarrow{f} V, \quad (8.2)$$

then $h \circ (g \circ f) = (h \circ g) \circ f$.

Now we consider composition of *linear* maps. If V, W, X are vector spaces over the same field and we have linear maps $T \in \mathbb{L}(V, W)$ and $S \in \mathbb{L}(W, X)$, then $S \circ T$ is the composition defined above.

Exercise 8.1. Show that $S \circ T$ is linear, and so $S \circ T \in \mathbb{L}(V, X)$.

Exercise 8.2. Show that composition is distributive with respect to addition of linear maps: that is, given any $Q, R \in \mathbb{L}(V, W)$ and $S, T \in \mathbb{L}(W, X)$, we have

$$(T + S) \circ Q = T \circ Q + S \circ Q, \quad T \circ (Q + R) = T \circ Q + T \circ R.$$

In light of the distributive property and the associative law, we will abbreviate our notation and denote the composition of linear maps by juxtaposing them, in a notation suggestive of *multiplication*:

$$ST \stackrel{\text{def}}{=} S \circ T.$$

It is very important to keep in mind that this product (composition) is only defined when the target space of T is the domain of S ; if the inputs and outputs do not match up, then ST is not defined. In particular, it is often the case that TS is not defined, even when ST is. In order to keep track of which maps go between which vector spaces (and hence which can be composed with each other), it is useful to draw diagrams like those in (8.1) and (8.2). For instance, in Exercise 8.2, we have the diagram

$$X \xleftarrow{S,T} W \xleftarrow{Q,R} V.$$

One important example is when $W = V$, and we consider linear maps from V to itself. In this case we will sometimes write $\mathbb{L}(V)$ instead of $\mathbb{L}(V, V)$ for the space of all linear operators on V . Then any two maps in $\mathbb{L}(V)$ have matching inputs and outputs, whichever order we try to compose them in, and products are always defined.

Even in this setting, however, it is important to remember that multiplication is generally not commutative: it is often the case that $ST \neq TS$, even when both are defined.

Example 8.3. Let $V = \mathbb{P}$ be the space of all polynomials in x , let T be differentiation, and let S be integration starting at 0, so $(Sf)(x) = \int_0^x f(t) dt$. (Note that this is always a polynomial whenever f is.) Then it is easy to check that $TS = I$, the identity map, but ST is not the identity map, since $(ST)(f) = \mathbf{0}$ whenever $f(x) = c$ is a constant polynomial.

Exercise 8.4. Let $V = \mathbb{R}^3$, let S be rotation around the x -axis by 90 degrees, and let T be rotation around the y -axis by 90 degrees. Show that $S, T \in \mathbb{L}(V)$ and that $ST \neq TS$.

As usual, when T is 1-1 and onto we say that it is *invertible* and write T^{-1} for its inverse. The same argument as for the inverse of a composition of two functions (not necessarily linear) shows that

$$(ST)^{-1} = T^{-1}S^{-1}$$

whenever ST is defined and S, T are both invertible.

8.2 Dual spaces and transposes

Given a linear map $T: V \rightarrow W$ and a linear functional $\ell \in W'$, we can view ℓ as a linear map from $W \rightarrow K$. Then we have the diagram

$$K \xleftarrow{\ell} W \xleftarrow{T} V,$$

and in particular we can form the composition $m = \ell T = \ell \circ T: V \rightarrow K$, which maps $v \in V$ to the scalar $\ell(Tv) \in K$. As explained in the previous section, m is linear, so $m \in \mathbb{L}(V, K) = V'$, the dual space of V .

Here is another way of thinking about the above construction. Fix a linear map $T \in \mathbb{L}(V, W)$. Then to every $\ell \in W'$ we can assign an element $m \in V'$, defined by $m = \ell \circ T$. That is, T defines a map $T': W' \rightarrow V'$, where $T'\ell = m = \ell \circ T$. In diagram form, the relation is like this:

$$\begin{array}{ccc} V & \xrightarrow{T} & W \\ & \searrow & \swarrow \\ & & K \end{array}$$

$\begin{array}{c} \text{---} T' \text{---} \\ \text{---} \ell \in W' \text{---} \\ \text{---} m \in V' \text{---} \end{array}$

Example 8.5. Let $V = \mathbb{P}_n$, $W = \mathbb{P}_{n-1}$, and let $T \in \mathbb{L}(V, W)$ be differentiation. Let $\ell \in W'$ be the linear functional that evaluates a polynomial g at a specific input t_0 , so $\ell(g) = g(t_0)$. Then $T'\ell \in V'$ is a linear functional on \mathbb{P}_n , given by $(T'\ell)(f) = \ell(Tf) = \ell(f') = f'(t_0)$.

Example 8.6. Let $V = W = \mathbb{R}^2$. Fix $a, b, c, d \in \mathbb{R}$ and consider the linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by

$$T \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

Let $\ell_1(x, y) = x$ and $\ell_2(x, y) = y$ be two linear functionals on \mathbb{R}^2 . The proof of Theorem 5.10 shows that $\{\ell_1, \ell_2\}$ is a basis for V' , so V' is also isomorphic to \mathbb{R}^2 . Thus we can represent any linear functional $\ell \in V'$ as $\ell = s\ell_1 + t\ell_2$, where $s, t \in \mathbb{R}$ are coordinates that depend on ℓ . Indeed, s and t are determined by

$$s(\ell) = \ell(1, 0), \quad t(\ell) = \ell(0, 1), \quad (8.3)$$

since then linearity implies that

$$\ell(x, y) = x\ell(1, 0) + y\ell(0, 1) = \ell_1(x, y)\ell(1, 0) + \ell_2(x, y)\ell(0, 1) = (s\ell_1 + t\ell_2)(x, y).$$

Given such an $\ell = s\ell_1 + t\ell_2$, what is $T'\ell$? We want to write $T'\ell = s'\ell_1 + t'\ell_2$ and then understand how (s, t) and (s', t') are related. Using (8.3), we have

$$\begin{aligned} s' &= (T'\ell)(1, 0) = \ell(T(1, 0)) = \ell(a, c) = as + ct, \\ t' &= (T'\ell)(0, 1) = \ell(T(0, 1)) = \ell(b, d) = bs + dt. \end{aligned}$$

Using matrices, we see that

$$\begin{pmatrix} s' \\ t' \end{pmatrix} = \begin{pmatrix} a & c \\ b & d \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix},$$

so that the linear map T' is given in matrix form (relative to the basis ℓ_1, ℓ_2) by the transpose of the matrix that defined T .

We will eventually see that the phenomenon in Example is very general, and thus we will refer to $T' \in \mathbb{L}(W', V')$ as the *transpose* of $T \in \mathbb{L}(V, W)$; it is also sometimes called the *adjoint*.

The following notation is sometimes useful when thinking of dual spaces: given $v \in V$ and $\ell \in V'$, we write

$$(\ell, v) \stackrel{\text{def}}{=} \ell(v).$$

This notation is suggestive of the inner product (or scalar product, or dot product) that is sometimes used on \mathbb{R}^n , but is more general. (Although if \mathbb{R}^n is equipped with an inner product, then identifying the two notations gives a natural isomorphism between \mathbb{R}^n and its dual space.) Using this notation, we see that the transpose satisfies $(T'\ell, v) = (T'\ell)(v) = \ell(Tv) = (\ell, Tv)$, or more succinctly,

$$(T'\ell, v) = (\ell, Tv). \tag{8.4}$$

If you have worked with the dot product on \mathbb{R}^n before, you may recognise (8.4) as one of the identities satisfied by the matrix transpose.

Nullity and rank

Further reading: [Lax] p. 20–22. See also [Bee] p. 544–603, [CDW] Ch. 6, 16 (p. 89–96, 237–246), [Hef] Ch. 3 (p. 157–172, 180–190), [LNS] Ch. 6 (p. 62–81), [Tre] Ch. 1 (p. 12–18)

9.1 Dimension theorem

The *nullity* of a linear transformation T is the dimension of its nullspace, and the *rank* of a linear transformation is the dimension of its range. These are related by the following fundamental result.

Theorem 9.1. *Let $T: V \rightarrow W$ be a linear map. Then $\dim N_T + \dim R_T = \dim V$.*

Proof. If $v_1 \equiv v_2 \pmod{N_T}$, then there is $x \in N_T$ such that $v_1 = v_2 + x$, and in particular $T(v_1) = T(v_2 + x) = T(v_2) + T(x) = T(v_2) + \mathbf{0}_W = T(v_2)$. Thus we can define a map $U: V/N_T \rightarrow R_T$ by $U([v]_{N_T}) = T(v)$; the previous sentence shows that U is well-defined.

In fact, U is an isomorphism: clearly it is onto, and if $U[v]_{N_T} = \mathbf{0}_W$, we get $T(v) = \mathbf{0}_W$, so $v \in N_T$, hence $[v]_{N_T} = \mathbf{0}_{V/N_T}$. Because U is an isomorphism, it follows from Theorem 9.1 that $\dim(V/N_T) = \dim(R_T)$. Moreover, we saw in Theorem 5.7 that $\dim(Y) + \dim(V/Y) = \dim(V)$ for every subspace $Y \subset V$. Putting $Y = N_T$, this implies the result. \square

Remark 9.2. Note that it is the dimension of the *domain* V , and not the dimension of the codomain W , that enters Theorem 9.1. One way of remembering this is the following: if X is some larger vector space containing W , then we can consider T as a linear map $V \rightarrow X$, and so we would need to replace $\dim W$ by $\dim X$ in any formula like the one in Theorem 9.1. However, $\dim N_T$ and $\dim R_T$ do not change when we replace W with X .

Corollary 9.3. *If $\dim(V) > \dim(W)$, then T is not 1-1.*

Proof. $\dim(R_T) \leq \dim(W)$, and so $\dim N_T = \dim V - \dim R_T \geq \dim V - \dim W \geq 1$. \square

Corollary 9.4. *If $\dim(V) = \dim(W)$, then T is 1-1 if and only if T is onto.*

Proof. Exercise. \square

Matrices

Further reading: [Lax] p. 32–36

10.1 Representing linear maps with matrices

[Bee] p. 518–522, [CDW] p. 97–125, [Hef] p. 191–208, [LNS] p. 71–75, [Tre] p. 13–21

Write $\mathbb{M}_{m \times n}(K)$ for the vector space of $m \times n$ matrices with entries in the field K . We saw in Example 6.12 that every $A \in \mathbb{M}_{m \times n}(K)$ determines a linear transformation $T: K^n \rightarrow K^m$ by the formula

$$(Tv)_i = \sum_{j=1}^n A_{ij}v_j \text{ for each } 1 \leq i \leq m. \quad (10.1)$$

In fact, this gives *all* the linear transformations from K^n to K^m .

Theorem 10.1. *For every $T \in \mathbb{L}(K^n, K^m)$, there is $A \in \mathbb{M}_{m \times n}(K)$ such that $T(v)$ is given by (10.1) for every $v \in K^n$. Moreover, the entries of A are the unique coefficients satisfying*

$$T(\mathbf{e}_j) = \sum_{i=1}^m A_{ij}\mathbf{d}_i \quad (10.2)$$

for every $1 \leq j \leq n$.

Proof. Write $\mathbf{e}_1, \dots, \mathbf{e}_n$ for the standard basis vectors in K^n , and $\mathbf{d}_1, \dots, \mathbf{d}_m$ for the standard basis vectors in K^m . Given $v \in K^n$, let $v_1, \dots, v_n \in K$ be the unique coefficients such that $v = \sum_{j=1}^n v_j\mathbf{e}_j$. Then by linearity,

$$T(v) = T\left(\sum_{j=1}^n v_j\mathbf{e}_j\right) = \sum_{j=1}^n v_jT(\mathbf{e}_j). \quad (10.3)$$

For each j , the image $T(\mathbf{e}_j)$ is a vector in K^m , and so it can be written as a linear combination of the standard basis vectors $\mathbf{d}_1, \dots, \mathbf{d}_m$. Let $A_{ij} \in K$ be the unique coefficients realising this – that is, define A_{ij} by (10.2). Putting (10.3) and (10.2) together gives

$$T(v) = \sum_{j=1}^n v_jT(\mathbf{e}_j) = \sum_{j=1}^n \sum_{i=1}^m A_{ij}v_j\mathbf{d}_i.$$

In other words, $T(v)$ is the vector whose coordinates are given by (10.1). \square

We see from Theorem 10.1 that the formula (10.1) is not arbitrary, but follows naturally from linearity of T and our decision to use the entries of the matrix A to record the partial coefficients of T . The only choice we made was to record the coefficients of the different vectors $T(\mathbf{e}_j)$ as the columns of the matrix, rather than the rows. This corresponds to a decision to treat vectors in K^n and K^m as column vectors, and to have A act by multiplication on the left. If we had made the other choice, we would work with row vectors, and multiply from the right.

Given our choice to work with column vectors, it also makes sense to think of a matrix A as being a “row of column vectors”. Another way of thinking of (10.2) is as the observation that multiplying A by the standard basis vector \mathbf{e}_j returns the j th column of A .

A similar argument to the one above shows that the usual formula for matrix multiplication is actually a very natural one, as it is the proper way to encode composition of linear transformations. To see this, consider the vector spaces K^n, K^m , and K^ℓ , and let $S \in \mathbb{L}(K^n, K^m)$, $T \in \mathbb{L}(K^m, K^\ell)$, so that we have the diagram

$$K^\ell \xleftarrow{T} K^m \xleftarrow{S} K^n,$$

and the composition is $TS \in \mathbb{L}(K^n, K^\ell)$. By Theorem 10.1 there are matrices $A \in \mathbb{M}_{\ell \times m}(K)$ and $B \in \mathbb{M}_{m \times n}(K)$ that encode the transformations T and S , respectively. Theorem 10.1 also shows that TS is encoded by a matrix $C \in \mathbb{M}_{\ell \times n}$. The next result shows that in fact $C = AB$ using the usual definition of matrix multiplication; this fact provides motivation for this definition.

Theorem 10.2. *With T, S and A, B, C as above, the matrix C is given by*

$$C_{ij} = \sum_{k=1}^m A_{ik} B_{kj} \text{ for every } 1 \leq i \leq \ell, 1 \leq j \leq n. \quad (10.4)$$

Proof. Let \mathbf{e}_i be the standard basis vectors for K^ℓ , let \mathbf{d}_j be the standard basis vectors for K^n , and let \mathbf{c}_k be the standard basis vectors for K^m . The key is to write $TS(\mathbf{e}_j)$ as a linear combination of the vectors \mathbf{c}_k . This can be done using the formulae for T and S in terms of A and B . First note that since B is the matrix for S , (10.2) gives

$$S(\mathbf{e}_j) = \sum_{k=1}^m B_{kj} \mathbf{c}_k.$$

Now applying T and using linearity followed by (10.2) for A and T , we get

$$TS(\mathbf{e}_j) = \sum_{k=1}^m B_{kj} T(\mathbf{c}_k) = \sum_{k=1}^m \sum_{i=1}^{\ell} B_{kj} A_{ik} \mathbf{d}_i. \quad (10.5)$$

Finally, using (10.2) for C and TS , we also have

$$TS(\mathbf{e}_j) = \sum_{i=1}^{\ell} C_{ij} \mathbf{d}_i, \quad (10.6)$$

and comparing this with (10.5) proves (10.4). \square

A nice corollary of the above result is the fact that matrix multiplication is associative – that is, $(AB)C = A(BC)$ whenever the products are defined. This is immediate from the fact that composition of linear maps is associative (which is easy to see), and while it could be proved directly using (10.4), that proof is messier and not as illuminating.

Now we have seen that linear maps from K^n to K^m can always be understood in terms of matrix multiplication. What about other vector spaces? Looking back at Theorems 10.1 and 10.2, we see that the proofs do not use any special properties of K^n apart from the fact that we work with the standard basis $\mathbf{e}_1, \dots, \mathbf{e}_n$ (and similarly for K^m, K^ℓ). In particular, if we let V, W be *any* finite-dimensional vector spaces and let $\mathbf{e}_1, \dots, \mathbf{e}_n$ be *any* basis for V , and $\mathbf{d}_1, \dots, \mathbf{d}_m$ *any* basis for W , the proof of Theorem 10.1 shows that the matrix $A \in \mathbb{M}_{m \times n}(K)$ with entries given by (10.2) determines the linear transformation T via the following version of (10.1): if $v = \sum_{j=1}^n v_j \mathbf{e}_j$, then

$$T(v) = \sum_{i=1}^m \left(\sum_{j=1}^n A_{ij} v_j \right) \mathbf{d}_i.$$

There is a very important fact to keep in mind here – the matrix A depends on our choice of basis for V and W . If we choose a different basis, then the vectors appearing in (10.2) will change, and so the entries of A will change as well.

Example 10.3. Let $V = W = \mathbb{P}_2$ and let T be differentiation. We first find the matrix of T relative to the basis $\mathbf{e}_1 = 1$, $\mathbf{e}_2 = x$, $\mathbf{e}_3 = x^2$. We see that

$$T\mathbf{e}_1 = \mathbf{0}, \quad T\mathbf{e}_2 = 1 = \mathbf{e}_1, \quad T\mathbf{e}_3 = 2x = 2\mathbf{e}_2,$$

and so using (10.2) the matrix of T relative to this basis is

$$A = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{pmatrix}.$$

What if we use the basis $\mathbf{d}_1 = 2$, $\mathbf{d}_2 = 2x$, and $\mathbf{d}_3 = x^2$? Then we still have $T\mathbf{d}_1 = \mathbf{0}$ and $T\mathbf{d}_2 = \mathbf{d}_1$, but now $T\mathbf{d}_3 = \mathbf{d}_2$, and so the matrix of T relative to this new basis is

$$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}.$$

Of course, there are many other bases available for \mathbb{P}_2 . If we write $\mathbf{c}_1 = 2+x$, $\mathbf{c}_2 = 1+x$, $\mathbf{c}_3 = \frac{1}{2}x^2$, we see that

$$T\mathbf{c}_1 = 1 = \mathbf{c}_1 - \mathbf{c}_2, \quad T\mathbf{c}_2 = 1 = \mathbf{c}_1 - \mathbf{c}_2, \quad T\mathbf{c}_3 = x = 2\mathbf{c}_2 - \mathbf{c}_1,$$

and so the matrix representing T becomes

$$\begin{pmatrix} 1 & 1 & 2 \\ -1 & -1 & -1 \\ 0 & 0 & 0 \end{pmatrix}.$$

Even for linear maps in K^n , where it seems most natural to use the standard basis $\mathbf{e}_1, \dots, \mathbf{e}_n$, it is often more useful to choose a different basis. For example, if T is the linear transformation represented by the matrix $A = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ in the standard basis, we see that taking $\mathbf{d}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\mathbf{d}_2 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ gives

$$T\mathbf{d}_1 = \begin{pmatrix} 2 \\ 2 \end{pmatrix} = 2\mathbf{d}_1, \quad T\mathbf{d}_2 = \mathbf{0},$$

so the relative to the basis $\mathbf{d}_1, \mathbf{d}_2$, the map T is given by the rather simpler matrix $\begin{pmatrix} 2 & 0 \\ 0 & 0 \end{pmatrix}$.

Two natural questions present themselves. How do we determine the relationship between the matrices that appear for different choices of basis? Given a linear transformation, can we choose a basis in which its matrix takes on a “nice” form that is easier to work with? These questions will motivate a great deal of what we do in the rest of this course.

Changing bases

Further reading:

11.1 Commutative diagrams

Let V be a vector space and let β be a set of vectors $v_1, \dots, v_n \in V$. There is a natural linear map $I_\beta: K^n \rightarrow V$ defined by

$$I_\beta(x_1, \dots, x_n) = \sum_{i=1}^n x_i v_i.$$

The map I_β is 1-1 if and only if β is linearly independent, and onto if and only if β spans V . In particular, it is an isomorphism if and only if β is a basis.

Let V, W be finite-dimensional vector spaces over K , and choose bases β for V and γ for W . Let $n = \dim V$ and $m = \dim W$, so that

$$I_\beta: K^n \rightarrow V, \quad I_\gamma: K^m \rightarrow W$$

are isomorphisms. Write $\beta = \{v_1, \dots, v_n\}$ and $\gamma = \{w_1, \dots, w_m\}$, and let A be the matrix of T relative to β and γ , as defined in the previous section: then A is given by (10.2), which in this case takes the form

$$T(v_j) = \sum_{i=1}^m A_{ij} w_i. \quad (11.1)$$

Now let $\mathbf{e}_j, \mathbf{d}_i$ be the standard basis vectors in K^n, K^m , respectively. By the definition of I_β and I_γ , we have $v_j = I_\beta(\mathbf{e}_j)$ and $w_i = I_\gamma(\mathbf{d}_i)$, so (11.1) implies

$$TI_\beta \mathbf{e}_j = \sum_{i=1}^m A_{ij} I_\gamma \mathbf{d}_i.$$

Applying I_γ^{-1} to both sides gives

$$(I_\gamma^{-1} T I_\beta) \mathbf{e}_j = \sum_{i=1}^m A_{ij} \mathbf{d}_i,$$

and we see that the composition $I_\gamma^{-1}TI_\beta$, which is a map from $K^n \rightarrow K^m$, is given by multiplication by the matrix A . The diagram below illustrates the situation.

$$\begin{array}{ccc} V & \xrightarrow{T} & W \\ I_\beta \uparrow & & \downarrow I_\gamma^{-1} \\ K^n & \xrightarrow{A} & K^m \end{array} \quad (11.2)$$

The above picture is an example of a *commutative diagram* – the word “commutative” here refers to the fact that whether we go from K^n to K^m directly, by following the arrow labeled “ A ”, or via another route, by following the other three sides of the square, we are still defining the same map from K^n to K^m , because $A = I_\gamma^{-1}TI_\beta$.

This discussion gives another way to think of the matrix representing a linear transformation: fix bases β and γ for the domain and codomain, respectively, which give isomorphisms I_β and I_γ with K^n and K^m , so that $I_\gamma^{-1}TI_\beta$ is a linear map from K^n to K^m , which is given by matrix multiplication thanks to Theorem 10.1. This point of view also highlights the role of the bases β and γ ; if we choose different bases, then the isomorphisms I_β and I_γ are different so we get a different map $K^n \rightarrow K^m$, and a different matrix. To emphasise the role of the bases, we will sometimes write $[T]_\beta^\gamma$ for the matrix representing the linear transformation T relative to the bases β and γ .

11.2 Changing bases

[Lax] p. 37–38, [Bee] p. 603–615, [CDW] p. 202–206, [Hef] p. 237–241, [LNS] p. 136–138, [Tre] p. 68–73

Let V be a finite-dimensional vector space and β a basis for V . Given $v \in V$ we will write $[v]_\beta = I_\beta^{-1}(v)$ for the vector in K^n whose components are the coefficients in the unique representation of v as a linear combination of elements of β . That is, $[v]_\beta = x \in K^n$ iff $v = \sum_{i=1}^n x_i v_i$.

Exercise 11.1. In the setting of the previous section, show that the commutative diagram (11.2) can be expressed as the equation $[Tv]_\gamma = [T]_\beta^\gamma[v]_\beta$.

An important special case of Exercise 11.1 comes when $T \in \mathbb{L}(V)$ is a linear operator on V – that is a linear map from V to itself – and we write down the matrix of T with respect to the same basis β in the domain and codomain. In this case we will simply write $[T]_\beta$ for the matrix of T , and

(11.2) becomes

$$\begin{array}{ccc} V & \xrightarrow{T} & V \\ I_\beta \uparrow & & \uparrow I_\beta \\ K^n & \xrightarrow{[T]_\beta} & K^n \end{array}$$

The result of Exercise 11.1 in this case says that $[Tv]_\beta = [T]_\beta[v]_\beta$.

Let β, γ be two bases for V . It is important to understand how $[T]_\beta$ and $[T]_\gamma$ are related, and to do this we must first understand how $[v]_\beta$ and $[v]_\gamma$ are related for $v \in V$. These are related to v by

$$\begin{array}{ccc} K^n & & K^n \\ & \searrow I_\beta & \swarrow I_\gamma \\ & V & \end{array} \qquad \begin{array}{ccc} [v]_\beta & & [v]_\gamma \\ & \searrow I_\beta & \swarrow I_\gamma \\ & v & \end{array} \qquad (11.3)$$

We see that $[v]_\gamma = I_\gamma^{-1}I_\beta[v]_\beta$. Let $I_\beta^\gamma = I_\gamma^{-1}I_\beta$: this is a linear transformation from K^n to K^n , and so is represented by an $n \times n$ matrix (relative to the standard basis). We refer to this matrix as the *change-of-coordinates matrix from β to γ* because it has the property

$$[v]_\gamma = I_\beta^\gamma[v]_\beta \text{ for all } v \in V. \qquad (11.4)$$

The following example illustrates how all this works. Let $V = \mathbb{P}_2$, let $\beta = \{f_1, f_2, f_3\}$ be the basis given by $f_1(x) = 2 + 3x$, $f_2(x) = 1 + x$, and $f_3(x) = x + x^2$, and let $\gamma = \{g_1, g_2, g_3\}$ be the standard basis $g_1(x) = 1$, $g_2(x) = x$, $g_3(x) = x^2$. Let $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ be the standard basis for \mathbb{R}^3 . In order to write down the matrix I_β^γ , we need to compute $I_\beta^\gamma \mathbf{e}_j$ for $j = 1, 2, 3$.

First note that $I_\beta^\gamma \mathbf{e}_j = I_\gamma^{-1}I_\beta \mathbf{e}_j$, and that $I_\beta \mathbf{e}_j = f_j$. (This is the definition of I_β .) Thus $I_\beta^\gamma \mathbf{e}_j = I_\gamma^{-1}(f_j) = [f_j]_\gamma$. That is, $I_\beta^\gamma \mathbf{e}_j$ is given by the coordinate representation of f_j in terms of the basis γ . In particular, since

$$\begin{aligned} f_1(x) &= 2 + 3x = 2g_1(x) + 3g_2(x), \\ f_2(x) &= 1 + x = g_1(x) + g_2(x), \\ f_3(x) &= x + x^2 = g_2(x) + g_3(x), \end{aligned}$$

we see that

$$[f_1]_\gamma = \begin{pmatrix} 2 \\ 3 \\ 0 \end{pmatrix}, \quad [f_2]_\gamma = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}, \quad [f_3]_\gamma = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix},$$

and so

$$I_{\beta}^{\gamma} = \begin{pmatrix} 2 & 1 & 0 \\ 3 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}.$$

Thus for example, the polynomial $p(x) = (2+3x) - 2(1+x) - (x+x^2) = -x^2$ has

$$[p]_{\beta} = \begin{pmatrix} 1 \\ -2 \\ -1 \end{pmatrix}, \quad [p]_{\gamma} = \begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix},$$

and indeed we can verify that

$$I_{\beta}^{\gamma}[p]_{\beta} = \begin{pmatrix} 2 & 1 & 0 \\ 3 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ -2 \\ -1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix} = [p]_{\gamma}.$$

To go the other way and find I_{γ}^{β} one needs to compute $I_{\gamma}^{\beta} \mathbf{e}_j = [g_j]_{\beta}$. That is, one must represent elements of β in terms of γ , and the coefficients of this representation will give the columns of the matrix. So we want to find coefficients a_i, b_i, c_i such that

$$\begin{aligned} g_1 &= a_1 f_1 + a_2 f_2 + a_3 f_3, \\ g_2 &= b_1 f_1 + b_2 f_2 + b_3 f_3, \\ g_3 &= c_1 f_1 + c_2 f_2 + c_3 f_3, \end{aligned} \tag{11.5}$$

and then we will have

$$I_{\gamma}^{\beta} = \begin{pmatrix} a_1 & b_1 & c_1 \\ a_2 & b_2 & c_2 \\ a_3 & b_3 & c_3 \end{pmatrix},$$

using the fact that $I_{\gamma}^{\beta} \mathbf{e}_1 = [g_1]_{\beta} = a_1 \mathbf{e}_1 + a_2 \mathbf{e}_2 + a_3 \mathbf{e}_3$, and similarly for \mathbf{e}_2 and \mathbf{e}_3 . The system of equations (11.5) becomes

$$\begin{aligned} 1 &= a_1(2+3x) + a_2(1+x) + a_3(x+x^2), \\ x &= b_1(2+3x) + b_2(1+x) + b_3(x+x^2), \\ x^2 &= c_1(2+3x) + c_2(1+x) + c_3(x+x^2). \end{aligned}$$

Comparing coefficients yields nine equations in nine variables, which can (eventually) be solved to produce

$$1 = -1(2+3x) + 3(1+x),$$

$$\begin{aligned}x &= (2 + 3x) - 2(1 + x), \\x^2 &= -(2 + 3x) + 2(1 + x) + (x + x^2).\end{aligned}$$

We conclude that

$$I_\gamma^\beta = \begin{pmatrix} -1 & 1 & -1 \\ 3 & -2 & 2 \\ 0 & 0 & 1 \end{pmatrix}.$$

We note that the product $I_\gamma^\beta I_\beta^\gamma$ corresponds to changing from β -coordinates to γ -coordinates, and then back to β -coordinates, so we expect that the product is the identity matrix, and indeed a direct computation verifies that

$$\begin{pmatrix} -1 & 1 & -1 \\ 3 & -2 & 2 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 2 & 1 & 0 \\ 3 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

The procedure described for the above example works very generally. Given two bases $\beta = \{v_1, \dots, v_n\}$ and $\gamma = \{w_1, \dots, w_n\}$ for a vector space V , let $A_{ij} \in K$ be the unique coefficients such that

$$v_j = \sum_{i=1}^n A_{ij} w_i \text{ for all } j = 1, \dots, n. \quad (11.6)$$

Then I_β^γ is the $n \times n$ matrix whose coefficients are given by A_{ij} – that is, the j th column of I_β^γ is given by the coefficients needed to express v_j in terms of the basis γ .

The principle of “express elements of one basis in terms of the other basis, then use those coefficients to make a matrix” is easy enough to remember. In principle one might try to build a matrix by putting the coefficients found above into a row of the matrix, instead of a column, or by expressing elements of γ in terms of β , instead of what we did above. To keep things straight it is useful to remember the following.

1. The coefficients we find will go into a column of the matrix, because we always work with column vectors – $[v]_\beta$ and $[v]_\gamma$ are expressed as column vectors, not row vectors.
2. The purpose of I_β^γ is to turn $[v]_\beta$ into $[v]_\gamma$ – that is, the output, the information that is needed, is on how to express vectors as linear combinations of elements of γ . Thus to determine I_β^γ , we should express elements of β in terms of γ , and not vice versa.

We end this section by remarking that (11.6) is actually a specific case of (10.2). Indeed, we defined I_β^γ as $I_\gamma^{-1}I_\beta$, using (11.3). We can draw the diagrams in (11.3) slightly differently, as

$$\begin{array}{ccc}
 K^n & \xrightarrow{I_\beta^\gamma} & K^n \\
 \downarrow I_\beta & & \downarrow I_\gamma \\
 V & \xrightarrow{I} & V
 \end{array}
 \qquad
 \begin{array}{ccc}
 [v]_\beta & \xrightarrow{I_\beta^\gamma} & [v]_\gamma \\
 \downarrow I_\beta & & \downarrow I_\gamma \\
 v & \xrightarrow{I} & v
 \end{array}
 \qquad (11.7)$$

Comparing this to (11.2), we see that I_β^γ is exactly $[I]_\beta^\gamma$, the matrix of the identity transformation relative to the bases β (in the domain) and γ (in the codomain).

Conjugacy, types of operators, dual spaces

Further reading: [Lax] p. 37–43.

12.1 Conjugacy

[Lax] p. 37–38, [Bee] p. 616–676, [CDW] p. 202–206, [Hef] p. 241–248, [LNS] p. 138–143, [Tre] p. 68–73

The previous section lets us describe the relationship between $[v]_\beta$ and $[v]_\gamma$ when $v \in V$ and β, γ are two bases for V . Now we are in a position to describe the relationship between $[T]_\beta$ and $[T]_\gamma$ when $T \in \mathbb{L}(V)$. The following commutative diagram relates these matrices to T itself, and hence to each other:

$$\begin{array}{ccc}
 K^n & \xrightarrow{[T]_\gamma} & K^n \\
 \downarrow I_\gamma & & \downarrow I_\gamma \\
 V & \xrightarrow{T} & V \\
 \uparrow I_\beta & & \uparrow I_\beta \\
 K^n & \xrightarrow{[T]_\beta} & K^n
 \end{array}$$

Recall that $I_\beta^\gamma = I_\gamma^{-1}I_\beta$ and $I_\gamma^\beta = I_\beta^{-1}I_\gamma$ be obtained by following the right-hand side from the top to the bottom. Then because the diagram commutes, we have

$$[T]_\beta = I_\gamma^\beta [T]_\gamma I_\beta^\gamma = (I_\beta^\gamma)^{-1} [T]_\gamma I_\beta^\gamma. \quad (12.1)$$

We could also have gotten to (12.1) by observing that $[T]_\beta$ satisfies the relationship $[Tv]_\beta = [T]_\beta[v]_\beta$, and that using the change-of-coordinates properties from the previous section, we have

$$[Tv]_\beta = I_\gamma^\beta [Tv]_\gamma = I_\gamma^\beta [T]_\gamma [v]_\gamma = I_\gamma^\beta [T]_\gamma I_\beta^\gamma [v]_\beta.$$

Thus the matrix $[T]_\gamma$ is obtained from $[T]_\beta$ by multiplying on the right and left by I_β^γ and its inverse. This motivates the following definition.

Definition 12.1. Let V, W be vector spaces and consider two linear maps $S \in \mathbb{L}(V)$ and $T \in \mathbb{L}(W)$. We say that S and T are *conjugate* if there is an isomorphism $J: V \rightarrow W$ such that $T = J^{-1}SJ$. That is, the following

diagram commutes:

$$\begin{array}{ccc} V & \xrightarrow{S} & V \\ \downarrow J & & \downarrow J \\ W & \xrightarrow{T} & W \end{array}$$

Similarly, we say that two $n \times n$ matrices A, B are conjugate, or *similar*, if there is an invertible $n \times n$ matrix C such that $B = C^{-1}AC$.

Exercise 12.2. Show that if two linear transformations $S, T \in \mathbb{L}(V)$ are conjugate, then there are bases β, γ for V such that $[S]_{\beta} = [T]_{\gamma}$.

Exercise 12.3. Let A, B, C be invertible $n \times n$ matrices. Show that AB and BA are always conjugate. Must ABC and BAC always be conjugate?

12.2 Nilpotent operators and projections

[Lax] p. 30–31. [Bee] p. 687–704, [Hef] p. 373–377, [Tre] p. 256–262

Certain linear transformations $T \in \mathbb{L}(V)$ have particular properties that warrant giving them special names. Here we quickly go over two such classes.

Definition 12.4. If there is $k \in \mathbb{N}$ such that $T^k = 0$, then T is a *nilpotent* operator.

Example 12.5. Let $T \in \mathbb{L}(\mathbb{R}^2)$ be the linear operator whose matrix (in the standard basis, or indeed, in any basis) is $A = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$. Then the matrix of T^2 is $A^2 = 0$, and so T^2 is the zero transformation.

Example 12.5 can be generalised. Say that a matrix $A \in \mathbb{M}_{n \times n}$ is *strictly upper-triangular* if $A_{ij} = 0$ whenever $i \geq j$. That is, all entries on the main diagonal, and below it, are equal to 0.

Exercise 12.6. Prove that if $T \in \mathbb{L}(V)$ can be represented by a strictly upper-triangular matrix, then T is nilpotent.

Exercise 12.7. Use the following steps to prove that the converse is also true – that is, if T is nilpotent, then there is a basis β such that $[T]_{\beta}$ is strictly upper-triangular.

1. Show that the null spaces of T, T^2, \dots satisfy the following relationship: $N_T \subset N_{T^2} \subset N_{T^3} \subset \dots \subset N_{T^k} = V$.
2. Let $j_i = \dim(N_{T^i})$, for $1 \leq i \leq k$. Note that $j_k = \dim V$.

3. Let v_1, \dots, v_{j_1} be a basis for N_T ; extend this to a basis v_1, \dots, v_{j_2} for N_{T^2} , and so on.
4. Let $\beta = \{v_1, \dots, v_{\dim V}\}$ be the basis for V obtained in this way, and show that $[T]_\beta$ is strictly upper-triangular.

Another important class are the *projections*. We say that $P \in \mathbb{L}(V)$ is a *projection* if there are subspace $W, X \subset V$ such that $V = W \oplus X$ and $P(w + x) = w$ for every $w \in W$ and $x \in X$. In this case we say that P is *the projection onto W along X* .

The definition given here of a projection is fairly geometric: to compute $P(v)$, one first needs to write v in the form $v = w + x$, where $w \in W$ and $x \in X$. There is a unique way to do this because $V = W \oplus X$. This amounts to taking “the part of v in the direction of W ” and “the part of v in the direction of X ” – then one throws away the second and keeps the first. Another way of thinking of it is that one starts at v and then moves along the congruence class $[v]_X = v + X = \{v + x \mid x \in X\}$. This congruence class intersects W exactly once (again, this is because of the direct sum property), and this unique point of intersection is $P(v)$.

Projections can also be characterised algebraically.

Exercise 12.8. Show that every projection P satisfies the equation $P^2 = P$.

Exercise 12.9. Show that the converse is also true: if $P \in \mathbb{L}(V)$ satisfies $P^2 = P$, then there are subspaces $W, X \subset V$ such that $V = W \oplus X$ and P is the projection onto W along X . (*Hint: take $W = R_P$ and $X = N_P$.*)

In order to represent a projection as a matrix, we need to choose a basis. The previous exercise shows that if $P \in \mathbb{L}(V)$ is a projection, then $V = R_P \oplus N_P$, and P is the projection onto R_P along N_P . Let v_1, \dots, v_k be a basis for R_P , and let w_1, \dots, w_ℓ be a basis for N_P . Note that $k + \ell = n = \dim V$. We see that

- $\beta = \{v_1, \dots, v_k, w_1, \dots, w_\ell\}$ is a basis for V ;
- $Pv_i = v_i$ for every $1 \leq i \leq k$;
- $Pw_i = \mathbf{0}$ for every $1 \leq i \leq \ell$.

This implies that $[Pv_i]_\beta = \mathbf{e}_i$ and $[Pw_i]_\beta = \mathbf{0}$, so the matrix representation of P relative to β can be written in *block form* as

$$[P]_\beta = \begin{pmatrix} I_{k \times k} & \mathbf{0}_{k \times \ell} \\ \mathbf{0}_{\ell \times k} & \mathbf{0}_{\ell \times \ell} \end{pmatrix},$$

where $I_{k \times k}$ is the $k \times k$ identity matrix, $\mathbf{0}_{k \times \ell}$ is the $k \times \ell$ matrix of all 0s, and so on. In other words,

$$[P]_{\beta} = \text{diag}(\overbrace{1, \dots, 1}^{k \text{ times}}, \overbrace{0, \dots, 0}^{\ell \text{ times}}),$$

where $\text{diag}(d_1, \dots, d_n)$ is the diagonal matrix D whose entries are given by $D_{ij} = 0$ when $i \neq j$ and $D_{ii} = d_i$. Using the change-of-coordinates procedure described earlier, this allows us to compute $[P]_{\gamma}$ for any basis γ .

12.3 Dual spaces

The representation of linear transformations via matrices gives us a concrete way to look at the dual space. Let V be an n -dimensional vector space. Then once we have fixed a basis β for V , an element of v is represented by the column vector $[v]_{\beta} \in K^n$. Recall that the dual space V' is the space of linear functionals on V – that is, linear transformations from V to the scalar field K . Thus $V' = \mathbb{L}(V, K)$ fits into the scheme we have been discussing, and elements of V' can be represented by matrices.

Recall that if $\dim V = n$ and $\dim W = m$, then $T \in \mathbb{L}(V, W)$ is represented by an $m \times n$ matrix once we have fixed bases $\beta = \{v_1, \dots, v_n\}$ for V and $\gamma = \{w_1, \dots, w_m\}$ for w . This is because T is characterised by the equation

$$[T]_{\beta}^{\gamma} [v]_{\beta} = [T(v)]_{\gamma},$$

and $[v]_{\beta} \in K^n$, $[T(v)]_{\gamma} \in K^m$.

Applying this to linear functionals and using the natural basis $\{1\}$ for K , we see that $\ell \in V'$ is represented by a matrix $[\ell]_{\beta}$ satisfying

$$[\ell]_{\beta} [v]_{\beta} = \ell(v).$$

The right hand side is a number, which is a 1×1 matrix, so $[\ell]_{\beta}$ must be a $1 \times n$ matrix in order for things to match up. In other words, a linear functional $\ell \in V'$ is represented by a row vector $[\ell]_{\beta}$, and $\ell(v)$ can be computed by multiplying the row and column vectors of ℓ and v relative to β .

More on nullity and rank

Further reading:

13.1 Row and column rank via annihilators

[Lax] p. 15–17, 27–28, 37.

Recall that the transpose of $T \in \mathbb{L}(V, W)$ is the map $T' \in \mathbb{L}(W', V')$ given by

$$T'(\ell)(v) = \ell(Tv) \quad (13.1)$$

for all $\ell \in W'$ and $v \in V$. Suppose $V = K^n$ and $W = K^m$, and let $A \in \mathbb{M}_{m \times n}$ be the matrix representing T relative to the standard bases. Then as described above, every linear functional $\ell \in (K^m)'$ is represented by a row vector, which we also denote ℓ . From (13.1), we see that

$$(T'\ell)(v) = \ell(Tv) = \ell Av = (\ell A)v.$$

In other words, the linear functional $T'\ell \in (K^n)'$ is represented by the row vector ℓA . Thus the transpose of T can also be represented via matrix multiplication, but multiplication of *row* vectors from the *right*, instead of column vectors from the left.

To represent T' as left multiplication of column vectors, we can do the following. Let $A^t \in \mathbb{M}_{n \times m}$ be the regular matrix transpose, that is, $(A^t)_{ij} = A_{ji}$. By the properties of this matrix transpose, we get $(\ell A)^t = A^t \ell^t$. Since ℓ^t is the column vector associated to the linear functional ℓ , we see that A^t is the matrix representing T' .

An important object in the dual space is the *annihilator*, defined as follows. Let V be any vector space and $S \subset V$ be any subset. The set

$$S^\perp = \{\ell \in V' \mid \ell(v) = 0 \text{ for all } v \in S\}$$

is called the *annihilator* of S .

Exercise 13.1. Show that S^\perp is always a subspace of V .

Annihilators satisfy a dimension theorem similar to Theorems 5.7 and 9.1.

Theorem 13.2. *If V is finite-dimensional and $W \subset V$ is a subspace of V , then $\dim W + \dim W^\perp = \dim V$.*

Proof. Let $\{v_1, \dots, v_m\}$ be a basis for W , and define a linear transformation $T: V' \rightarrow K^m$ by

$$T(\ell) = \begin{pmatrix} \ell(v_1) \\ \vdots \\ \ell(v_m) \end{pmatrix}.$$

Then $W^\perp = N_T$. One can show that T is onto (this is an exercise), and thus $R_T = K^m$. In particular, we have $\dim W^\perp = \dim N_T$ and $\dim W = m = \dim R_T$, and so by Theorem 9.1 we have $\dim W + \dim W^\perp = \dim R_T + \dim N_T = \dim V'$. By Theorem 5.10, this is equal to $\dim V$. \square

Theorem 13.3. *For any linear transformation $T \in \mathbb{L}(V, W)$, the annihilator of the range of T is the nullspace of its transpose: $R_T^\perp = N_{T'}$.*

Proof.

$$\begin{aligned} R_T^\perp &= \{\ell \in W' \mid \ell(w) = 0 \text{ for all } w \in R_T\} && \text{def. of annihilator} \\ &= \{\ell \in W' \mid \ell(Tv) = 0 \text{ for all } v \in V\} && \text{def. of range} \\ &= \{\ell \in W' \mid (T'\ell)(v) = 0 \text{ for all } v \in V\} && \text{def. of transpose} \\ &= \{\ell \in W' \mid T'\ell = \mathbf{0}\} && \text{def. of zero functional} \\ &= N_{T'} && \text{def. of nullspace} \end{aligned}$$

\square

The definition of annihilator also goes from the dual space V' to V itself: given $S \subset V'$, we write

$$S^\perp = \{v \in V \mid \ell(v) = 0 \text{ for all } \ell \in S\}.$$

Theorem 13.4. *If V is finite-dimensional and $X \subset V$ is a subspace of V , then $(X^\perp)^\perp = X$.*

Proof. By definition of X^\perp , every $v \in X$ satisfies $\ell(v) = 0$ for all $\ell \in X^\perp$, thus $X \subset (X^\perp)^\perp$. On the other hand, Theorem 13.2 implies that

$$\dim(X^\perp)^\perp = \dim V' - \dim X^\perp = \dim V' - (\dim V - \dim X) = \dim X$$

using the fact that $\dim V' = \dim V$. The result follows. \square

Together with Theorem 13.3, this result immediately implies that the range of T is the annihilator of the nullspace of T' :

$$R_T = (N_{T'})^\perp \tag{13.2}$$

We can put these results together to conclude that T and T' have the same rank for any finite-dimensional V, W and any $T \in \mathbb{L}(V, W)$. Indeed, Theorems 13.2 and 9.1 yield

$$\begin{aligned}\dim R_T^\perp + \dim R_T &= \dim W, \\ \dim N_{T'} + \dim R_{T'} &= \dim W'.\end{aligned}$$

The last terms are equal by Theorem 5.10, and the first terms are equal by Theorem 13.3. Thus we conclude that

$$\dim R_T = \dim R_{T'}. \quad (13.3)$$

This can be interpreted as the statement that row and column rank of a matrix are the same. Indeed, if T is an $m \times n$ matrix, then $R_T \subset K^m$ is the column space of T , and $R_{T'} \subset K^n$ is the row space of T . The above result shows that these spaces have the same dimension.

Exercise 13.5. Show that if $\dim V = \dim W$ and $T \in \mathbb{L}(V, W)$, then $\dim N_T = \dim N_{T'}$.

13.2 Existence of solutions

[Lax] p. 23–24

The key relationship between nullity and rank is Theorem 9.1, which states that dimension of the domain of a linear transformation is the dimension of the range plus the dimension of the null space. In the concrete language of a first course in linear algebra, Corollary 9.3 is the statement that every undetermined system (m equations in n variables, where $n > m$) has a non-trivial solution, and Corollary 9.4 is the statement that for a system of n equations in n variables, the following two statements are equivalent:

1. the homogeneous system $Ax = 0$ has only the trivial solution;
2. the non-homogeneous system $Ax = b$ has a solution for every value of the data b .

This fundamental fact has many applications. Here is one. Given a bounded domain $G \subset \mathbb{R}^2$, the *Laplace equation* on G is the PDE $\Delta u = u_{xx} + u_{yy} = 0$ in G , with the value of u prescribed on the boundary of G ; this describes heat distributions and many other important things in physics. As with any PDE, it is important to understand whether a solution always exists, for given boundary conditions.

Of course we are not studying PDEs in this course, but there is a useful finite-dimensional analogue of this problem. Suppose we discretise the problem by filling the region G with a square lattice and recording the value of u at every point of the lattice. Then each point has four neighbours, which we think of as lying to the north, south, east, and west, at a distance of h . Writing u_0 for the value of u at the central point, and u_N, u_S, u_E, u_W for the values at its neighbours, we have the approximations

$$\begin{aligned} u_{xx} &\approx \frac{1}{h^2}(u_W - 2u_0 + u_E), \\ u_{yy} &\approx \frac{1}{h^2}(u_N - 2u_0 + u_S), \end{aligned}$$

and so the approximate version of the Laplace equation $u_{xx} + u_{yy} = 0$ is the linear equation

$$u_0 = \frac{1}{4}(u_W + u_E + u_N + u_S). \quad (13.4)$$

That is, u solves the discretised equation if and only if its value at every lattice point is the average of its values at the four neighbours of that point. We require this to hold whenever all four neighbours $u_W, u_E, u_N,$ and u_S are in G . If any of those neighbours are outside of G , then we require u_0 to take the value prescribed at the nearest boundary point of G .

Writing n for the number of lattice points in G , we see that the discretised version of the Laplace equation $\Delta u = 0$ is a system of n equations in n unknowns. We want to answer the following question: for a given choice of boundary conditions, is there a solution of the discretised Laplace equation on the interior of G ? If so, is that solution unique?

Because the discretised Laplace equation is a system of n linear equations in n unknowns, this question is equivalent to the following: does the given non-homogeneous system of equations have a unique solution? We know that if the *homogeneous* system has only the trivial solution, then the non-homogeneous system has a unique solution for any choice of data – in particular, the discretised Laplace equation has a unique solution on the interior of G for any choice of boundary conditions.

The homogeneous system for the Laplace equation corresponds to having the zero boundary condition ($u = 0$ on the boundary of G). Let M be the maximum value that u takes at any point in G . We claim that $M = 0$. Indeed, if $M > 0$ then there is some lattice point at the interior of G for which $u_0 = M$. By (13.4) and the fact that each of the neighbouring points has $u \leq M$, we see that $u = M$ at every neighbouring point. Continuing in

this way, we conclude that $u = M$ on *every* point, contradicting the fact that it vanishes on the boundary. Thus $u \leq 0$ on all of G . A similar argument (considering the minimum value) shows that $u \geq 0$ on all of G , and so $u = 0$ is the only solution to the homogeneous problem. This implies that there is a unique solution to the non-homogeneous problem for every choice of boundary conditions.

In fact, we have proved something slightly stronger. Laplace's equation is a special case of *Poisson's equation* $\Delta u = f$, where $f: G \rightarrow \mathbb{R}$ is some function – examples where $f \neq 0$ arise in various physical applications. Upon discretising, (13.4) becomes $u_0 = \frac{1}{4}(u_W + u_E + u_N + u_S - f_0)$, where f_0 is the value of f at the given point. Thus the situation is the following: if G contains n lattice points and k of these are on the boundary in the sense that one of their neighbours lies outside of G , then the non-homogeneous part of k of the equations specifies the boundary condition, while the non-homogeneous part of the remaining $n - k$ equations specifies the function f . The above argument permits both parts to be non-zero, and so it shows that the discretised version of Poisson's equation has a unique solution for every choice of boundary condition and every choice of f .

Powers of matrices

Further reading:

14.1 Iterating linear transformations

[Lax] p. 58–59

Suppose we have some system which is modeled by a map (not necessarily linear) from \mathbb{R}^n into itself. That is, a point in \mathbb{R}^n represents a particular state of the system, and if the system is in state x at the present time, then its state one unit of time into the future is given by $F(x)$. The map $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$ encodes the *dynamics* of the system; it may be given by integrating an ODE, or by some other method.

Now suppose that the motion of this system starting from the initial condition $x = 0$ is periodic – that is, there is $p \in \mathbb{N}$ such that $F^p(0) = 0$. Then writing $G = F^p$, we see that 0 is a fixed point of G – that is, $G(0) = 0$. If the map G is differentiable (which is usually the case for physically interesting systems), then for states starting near this fixed point we have

$$G(x) = DG(0)x + O(x^2),$$

where $DG(0) \in \mathbb{M}_{n \times n}$ is the derivative of G at 0, and the remainder term $O(x^2)$ is negligible when $\|x\|$ is small. Let $A = DG(0)$ – then determining the behaviour of a trajectory starting at $x \approx 0$ amounts to computing $A^k(x)$ as k increases. That is, we want to study the behaviour of the sequence of vectors x, Ax, A^2x, A^3x, \dots

There are other motivations for studying this sequence of vectors. Consider the Fibonacci sequence, defined by $x_0 = x_1 = 1$ and then iteratively by $x_n = x_{n-1} + x_{n-2}$ for $n \geq 2$. Let v_n be the vector $\begin{pmatrix} x_{n-1} \\ x_n \end{pmatrix}$; then we see that v_{n+1} and v_n are related by

$$v_{n+1} = \begin{pmatrix} x_n \\ x_{n+1} \end{pmatrix} = \begin{pmatrix} x_n \\ x_n + x_{n-1} \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} x_{n-1} \\ x_n \end{pmatrix} = Av_n,$$

where $A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$, and thus we can write down the formula

$$v_n = A^{n-1}v_1 = A^{n-1} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = A^{n-1}A \begin{pmatrix} 0 \\ 1 \end{pmatrix} = A^n \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

That is, v_n is the second column of the matrix A^n , and x_n is the second entry of this column, so we have the explicit formula

$$x_n = (A^n)_{22}.$$

This gives us an explicit formula for the n th Fibonacci number, if only we can find a formula for A^n .

The difficulty, of course, is that it is not clear how to write down a formula for A^k for an arbitrary matrix A , or even to predict whether the entries of A^k will grow, shrink, or oscillate in absolute value. For example, consider the matrices

$$A = \begin{pmatrix} 5 & 7.1 \\ -3 & -4 \end{pmatrix} \quad B = \begin{pmatrix} 5 & 6.9 \\ -3 & -4 \end{pmatrix} \quad (14.1)$$

It turns that

$$A^{50} \approx \begin{pmatrix} -1476 & -3176 \\ 1342 & 2549 \end{pmatrix} \quad B^{50} \approx \begin{pmatrix} .004 & .008 \\ -.003 & -.006 \end{pmatrix} \quad (14.2)$$

Thus despite the fact that A and B have entries that are very close to each other, the iterates $A^n v$ and $B^n v$ behave very differently: we will eventually be able to show that if $v \in \mathbb{R}^2 \setminus \{\mathbf{0}\}$ is any non-zero vector, then $\|A^n v\| \rightarrow \infty$ as $n \rightarrow \infty$, while $\|B^n v\| \rightarrow \mathbf{0}$.

We could formulate a similar example for a *one-dimensional* linear map, which is just multiplication by a real number. Let $a = 1.1$ and $b = 0.9$. Then we see immediately that for any non-zero number x , we have $a^n x \rightarrow \infty$ and $b^n x \rightarrow 0$ as $n \rightarrow \infty$. In this case, it is clear what the issue is: $|a|$ is larger than 1, and $|b|$ is smaller than 1.

Can we find a similar criterion for matrices? Is there some property of a matrix that we can look at and determine what kind of behaviour the iterates $A^n v$ will exhibit? This is one question that will motivate our discussion of eigenvalues and eigenvectors, determinants, and eventually spectral theory.

14.2 Eigenvectors and eigenvalues

[Lax] p. 59–60. [Bee] p. 453–479, [CDW] p. 187–200, [Hef] p. 364–373, [LNS] p. 82–88, [Tre] p. 99–105

First we point out that there is one situation in which it is easy to determine the behaviour of the iterates $A^n v$. If $v \in \mathbb{R}^n$ has the property that Av is a scalar multiple of v – that is, there exists $\lambda \in \mathbb{R}$ such that $Av = \lambda v$,

then the linearity property of multiplication by A makes computing the iterates $A^n v$ very easy: we just observe that

$$A^n v = A^{n-1}(Av) = A^{n-1}(\lambda v) = \lambda(A^{n-1}v) = \cdots = \lambda^n v.$$

Then $A^n v$ grows exponentially in magnitude if $|\lambda| > 1$, and decreases exponentially if $|\lambda| < 1$.

If $v \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ is a non-zero vector such that $Av = \lambda v$ for some $\lambda \in \mathbb{R}$, then v is called an *eigenvector* of the matrix A , and the corresponding value of λ is called an *eigenvalue* of A . We use the same terminology for linear transformations: given a vector space V and a linear transformation $T \in \mathbb{L}(V)$, an eigenvector v is a non-zero vector such that $T(v) = \lambda v$, and λ is called an eigenvalue of T .

Note that if v is an eigenvector for T , then so is cv for every $c \in K$, with the same eigenvalue: indeed, $T(cv) = cT(v) = c(\lambda v) = \lambda(cv)$.

Example 14.1. Let $A = \begin{pmatrix} 2 & 0 \\ 0 & 3 \end{pmatrix}$. Then \mathbf{e}_1 is an eigenvector with eigenvalue 2, and \mathbf{e}_2 is an eigenvector with eigenvalue 3. On the other hand, $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is not an eigenvector, because $A \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 2 \\ 3 \end{pmatrix}$ is not a scalar multiple of $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$.

The following elementary result gives an important characterisation of eigenvalues.

Proposition 14.2. *A scalar $\lambda \in K$ is an eigenvalue of a linear transformation $T \in \mathbb{L}(V)$ if and only if $T - \lambda I$ is non-invertible. The set of eigenvectors corresponding to λ is precisely the set of non-trivial vectors in the nullspace of $T - \lambda I$. (This nullspace is often called the eigenspace of λ .)*

Proof. Note that v is an eigenvector for λ if and only if $Tv = \lambda v$, which is equivalent to $(T - \lambda I)v = Tv - \lambda v = \mathbf{0}$. Thus a non-zero vector v is in the nullspace of $T - \lambda I$ if and only if it is an eigenvector for T with eigenvalue λ . This proves the second part of the proposition, and the first part follows because $T - \lambda I$ is invertible if and only if the nullspace is trivial. \square

For eigenvectors at least, the question of computing $A^n v$ is tractable. But how restrictive is the condition that v is an eigenvector of A ? The example above shows that in general, not every vector is an eigenvector. In fact, when $K = \mathbb{R}$, there may not be *any* eigenvectors.

Exercise 14.3. Let $A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$, and show that A does not have any eigenvectors in \mathbb{R}^2 .

If we are willing to work over the complex numbers (that is, with $K = \mathbb{C}$ instead of $K = \mathbb{R}$), then there is always at least one eigenvector.

Proposition 14.4. *Let $A \in \mathbb{M}_{n \times n}(\mathbb{C})$. Then there is $v \in \mathbb{C}^n \setminus \{\mathbf{0}\}$ and $\lambda \in \mathbb{C}$ such that $Av = \lambda v$.*

Proof. Let w be any nonzero vector in \mathbb{C}^n , and consider the vectors

$$w, Aw, A^2w, \dots, A^n w.$$

There are $n + 1$ of these vectors, and $\dim(\mathbb{C}^n) = n$, thus these vectors are linearly dependent. That is, there exist coefficients $c_0, \dots, c_n \in \mathbb{C}$ such that

$$\sum_{j=0}^n c_j A^j w = \mathbf{0}. \quad (14.3)$$

Let $p \in \mathbb{P}_n(\mathbb{C})$ be the polynomial with coefficients c_j – that is, $p(t) = \sum_{j=0}^n c_j t^j$. We will write $p(A) = \sum_{j=0}^n c_j A^j$, so that in particular (14.3) can be written as $p(A)w = \mathbf{0}$.

By the fundamental theorem of algebra, p can be written as a product of linear factors – that is, there are $a_1, \dots, a_n \in \mathbb{C}$ and $c \in \mathbb{C} \setminus \{0\}$ such that

$$p(t) = c \prod_{j=1}^n (t - a_j).$$

This factorisation works for $p(A)$ as well, and we have

$$p(A) = c(A - a_1 I)(A - a_2 I) \cdots (A - a_n I).$$

Thus (14.3) can be rewritten as

$$c(A - a_1 I)(A - a_2 I) \cdots (A - a_n I)w = \mathbf{0}.$$

Because $w \neq \mathbf{0}$, this means that the product $(A - a_1 I) \cdots (A - a_n I)$ has non-trivial nullspace, hence is non-invertible. But a product of invertible matrices is invertible, so we conclude that at least one of the matrices $A - a_j I$ is non-invertible. In particular, $A - a_j I$ has a non-trivial nullspace, and taking v to be a non-trivial element of this nullspace, we get $(A - a_j I)v = \mathbf{0}$, so $Av = a_j v$, and v is an eigenvector with eigenvalue a_j . \square

Proposition 14.4 shows that over \mathbb{C} , every matrix has at least one eigenvalue and eigenvector, but does not show how to compute them in practice, or how many of them there are. To do this, we need to return to the characterisation of eigenvalues and eigenvectors in Proposition 14.2. If we know that λ is an eigenvalue, then eigenvectors for λ are elements of the nullspace

of $A - \lambda I$, which can be found via row reduction. Thus the primary challenge is to find the eigenvalues – that is, to find λ such that $A - \lambda I$ is non-invertible.

Thus we have reduced the problem of determining eigenvalues and eigenvectors to the problem of determining invertibility of a matrix. This motivates the introduction of the *determinant*, a quantity which we will spend the next few lectures studying. We end this lecture with the following observation: if $\beta = \{v_1, \dots, v_n\}$ is a basis for V consisting of eigenvectors of a linear operator T , then for each $j = 1, \dots, n$ we have $Tv_j = \lambda_j v_j$ for some $\lambda_j \in K$, so that $[Tv_j]_\beta = \lambda_j \mathbf{e}_j$. This implies that $[T]_\beta$ is the diagonal matrix with entries $\lambda_1, \dots, \lambda_n$. Thus finding a basis of eigenvectors has very strong consequences.

Exercise 14.5. Consider the matrix $A = \begin{pmatrix} -4 & -6 & 3 \\ 2 & 4 & -2 \\ -2 & -2 & 1 \end{pmatrix}$. Show that 0, -1, and 2 are all eigenvalues of A , and find eigenvectors associated to each. Check that these eigenvectors form a basis β , and find the change-of-coordinates matrix $Q = I_\beta^\alpha$, where α is the standard basis for \mathbb{R}^3 . Verify by direct computation that $Q^{-1}AQ = \text{diag}(0, -1, 2)$.

Introduction to determinants

Further reading:

15.1 Determinant of 2×2 matrices

Given a matrix A , it is important to be able to determine whether or not A is invertible, and if it is invertible, how to compute A^{-1} . This arises, for example, when trying to solve the non-homogeneous equation $Ax = y$, or when trying to find eigenvalues by determining whether or not $A - \lambda I$ is invertible.

In this and the coming lectures, we will study determinants, which provide a powerful tool for answering this question. It would be customary at this point to introduce the determinant either via a rather complicated formula or via a list of properties; we opt to postpone these for a little while, beginning rather with a discussion that leads us through the process that one might follow in order to ultimately develop that formula and those properties, beginning with the question of invertibility posed above.

Let us first consider the case where $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ is a 2×2 matrix. We know that A is invertible if and only if it row reduces to $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, and that in this case A^{-1} can be computed by row reducing $[A \mid I]$ to $[I \mid A^{-1}]$. So let's try to row reduce

$$[A \mid I] = \left(\begin{array}{cc|cc} a & b & 1 & 0 \\ c & d & 0 & 1 \end{array} \right). \quad (15.1)$$

If $c = 0$, then our task is easy. Recall that the following are equivalent:

- A is invertible
- The columns of A form a basis for \mathbb{R}^2 .
- The rows of A form a basis for \mathbb{R}^2 .

If A is invertible and $c = 0$, we must have $a \neq 0$ and $d \neq 0$, since otherwise one of the rows or columns of A vanishes. On the other hand, if $c = 0$ and a, d are non-zero, we can row reduce (15.1) as follows:

$$[A \mid I] \rightarrow \left(\begin{array}{cc|cc} 1 & \frac{b}{a} & \frac{1}{a} & 0 \\ 0 & 1 & 0 & \frac{1}{d} \end{array} \right) \rightarrow \left(\begin{array}{cc|cc} 1 & 0 & \frac{1}{a} & -\frac{b}{ad} \\ 0 & 1 & 0 & \frac{1}{d} \end{array} \right)$$

We deduce that in the case $c = 0$, invertibility of a is equivalent to the statement that $ad \neq 0$, and that in this case we have

$$\begin{pmatrix} a & b \\ 0 & d \end{pmatrix}^{-1} = \begin{pmatrix} \frac{1}{a} & -\frac{b}{ad} \\ 0 & \frac{1}{d} \end{pmatrix} = \frac{1}{ad} \begin{pmatrix} d & -b \\ 0 & a \end{pmatrix}$$

Exercise 15.1. A matrix $A \in \mathbb{M}_{n \times n}$ is *upper triangular* if $A_{ij} = 0$ whenever $i > j$. Show that an upper triangular matrix A is invertible if and only if $A_{jj} \neq 0$ for all $1 \leq j \leq n$. Equivalently, A is invertible if and only if $\prod_{j=1}^n A_{jj} \neq 0$ (the product of the diagonal entries is non-zero). Note that upper triangular is a different criterion from *strictly* upper triangular. Show that every strictly upper triangular matrix is non-invertible.

Let's go back to the 2×2 matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ and consider now the case $c \neq 0$. In this case we can row reduce (15.1) by multiplying the first row by c and then subtracting a times the second row:

$$\left(\begin{array}{cc|cc} a & b & 1 & 0 \\ c & d & 0 & 1 \end{array} \right) \rightarrow \left(\begin{array}{cc|cc} ac & bc & c & 0 \\ c & d & 0 & 1 \end{array} \right) \rightarrow \left(\begin{array}{cc|cc} 0 & bc - ad & c & -a \\ c & d & 0 & 1 \end{array} \right)$$

At this point we can observe that one of two things happens: either $ad = bc$ and it is impossible to row reduce this to $[I \mid A^{-1}]$, hence A is non-invertible; or $ad \neq bc$ and we can get to that form via four more row operations (divide first row by the non-zero value $bc - ad$, divide second row by the non-zero value c , swap the rows, and finally subtract a multiple of a row to eliminate the remaining non-zero term in the left half). In particular, we conclude that A is invertible if and only if $ad - bc \neq 0$. Note that this criterion works for the case $c = 0$ as well, because in this case it reduces to $ad \neq 0$, which was the criterion found in the previous argument.

Exercise 15.2. Complete the row reduction above to show that the augmented matrix $[A \mid I]$ row reduces to

$$\left(\begin{array}{cc|cc} ad - bc & 0 & d & -b \\ 0 & ad - bc & -c & a \end{array} \right)$$

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

Thus for 2×2 matrices, invertibility can be determined via a formula — one simply tests whether or not the *determinant* $ad - bc$ is zero or non-zero. It is natural to ask if there is a similar quantity for larger matrices, and we will address this shortly. First we give a geometric interpretation of the determinant for a 2×2 matrix. Recall again that A is invertible if and only

if the columns of A form a basis for \mathbb{R}^2 , which is true if and only if they are linearly independent (because there are 2 columns). Let $v = \begin{pmatrix} a \\ c \end{pmatrix}$ and $w = \begin{pmatrix} b \\ d \end{pmatrix}$ be the columns of A . Plotting v and w in the plane, we can make the following geometric observation: v, w are linearly dependent if and only if they point in the same direction or in opposite directions. Let θ be the angle between v and w , then the vectors are linearly dependent if and only if $\theta = 0$ or $\theta = \pi$, which is equivalent to $\sin \theta = 0$.

So, can we compute $\sin \theta$ in terms of a, b, c, d ? Let $r = \|v\| = \sqrt{a^2 + c^2}$ and $s = \|w\| = \sqrt{b^2 + d^2}$ be the lengths of the vectors v and w , so that in polar coordinates we have

$$v = \begin{pmatrix} a \\ c \end{pmatrix} = \begin{pmatrix} r \cos \alpha \\ r \sin \alpha \end{pmatrix} \quad w = \begin{pmatrix} b \\ d \end{pmatrix} = \begin{pmatrix} s \cos \beta \\ s \sin \beta \end{pmatrix}.$$

Then $\theta = \beta - \alpha$, and we have

$$\sin \theta = \sin(\beta - \alpha) = \sin \beta \cos \alpha - \cos \beta \sin \alpha = \frac{1}{rs}(ad - bc). \quad (15.2)$$

As long as neither of the columns of A is zero (in which case A is non-invertible), we see that the angle between the column vectors is proportional to the determinant $ad - bc$.

It is still not clear how to generalise this to higher dimensions. What angle should we be computing for a 3×3 matrix, since there are 3 column vectors to work with?

To this end, we make the following observation. Let $P = P(v, w)$ be the parallelogram with vertices at $0, v, w$, and $v + w$. Then taking the edge from 0 to v as the base of the parallelogram, the area of P is (base) \times (height), where the base has length $r = \|v\|$ and the height h is given by the formula $\sin \theta = h/\|w\|$, so $h = \|w\| \sin \theta$. Using (15.2), we conclude that the area of the parallelogram is

$$\text{area}(P) = \|v\| \|w\| \sin \theta = rs \frac{ad - bc}{rs} = ad - bc,$$

so that the area is equal to the determinant of A .

There is a small caveat we must insert at this point – the previous paragraph was not entirely honest. If w lies on the other side of v , then $\sin \theta$ will be negative, and the formula will return a negative number. But after all, the determinant $ad - bc$ may be negative as well. The correct interpretation here is that the determinant gives the *signed* area of the parallelogram, and that reversing the order in which we input the parallelogram's edges as the columns of the matrix will switch the sign of the determinant, while preserving the absolute value.

We observe that the formula for determinant of a 2×2 matrix immediately shows that

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

That is, both A and its transpose A^t have the same determinant. We will later see that this property continues to hold for larger matrices. For now, note that this means that the determinant of A gives the area of the parallelogram spanned by the rows of A , as well as the one spanned by the columns.

Finally, we observe that this geometric interpretation can also be interpreted in terms of row reduction. Given two row vectors $v, w \in \mathbb{R}^2$, it is easy to see that $P(v + tw, w)$ and $P(v, w)$ have the same area for every $t \in \mathbb{R}$. This corresponds to the fact that we may treat w as the base of the parallelogram, and the height does not change if we replace v with $v + tw$ for any t . Thus assuming v and w are in “general position”, we can make one row reduction to go from $\begin{pmatrix} v \\ w \end{pmatrix}$ to $\begin{pmatrix} v' \\ w \end{pmatrix}$, where $v' = v + tw$ is of the form $(a, 0)$, and then we can make one more row reduction to go from $\begin{pmatrix} v' \\ w \end{pmatrix}$ to $\begin{pmatrix} v' \\ w' \end{pmatrix}$, where $w' = w + sv'$ is of the form $(0, b)$. Then $P(v, w) = ab$, since the parallelogram spanned by v' and w' is just a rectangle with sides a and b .

The previous paragraph shows that if we row reduce a 2×2 matrix A to a diagonal form D using only operations of the form “add a multiple of a row to a different row”, then the product of the diagonal entries of D gives the area of the parallelogram spanned by the rows of A . A completely analogous result is true for $n \times n$ matrices. In the next section, we consider the 3×3 case and find a formula for the product of these diagonal entries in terms of the entries of A itself.

15.2 Determinant of 3×3 matrices

Let $v, w, x \in \mathbb{R}^3$ be the row vectors of a 3×3 matrix $A = \begin{pmatrix} v \\ w \\ x \end{pmatrix}$. Following the previous section and the 2×2 case, it is reasonable to ask how we compute the (signed) volume of the parallelepiped $P(v, w, x)$ spanned by v, w, x , since this number tells us how far A is from being invertible, and may reasonably be expected to appear in formulas for A^{-1} . As at the end of the previous section, we want to row reduce A to $D = \begin{pmatrix} v' \\ w' \\ x' \end{pmatrix}$, where $v' = (a, 0, 0)$, $w' = (0, b, 0)$ and $x' = (0, c, 0)$, so that $\det A = \det D = abc$ gives the signed volume of $P(v, w, x)$. (Note that we have not yet given a

formal definition of $\det A$, but our working definition is that at least up to sign, it should be the volume of this parallelepiped.)

We go in three steps: first go from $\begin{pmatrix} v \\ w \\ x \end{pmatrix}$ to $\begin{pmatrix} v' \\ w \\ x \end{pmatrix}$, where $v' = v + sw + tx = (a, 0, 0)$ for some $s, t \in \mathbb{R}$; then row reduce to $\begin{pmatrix} v' \\ w' \\ x \end{pmatrix}$, where $x = (0, b, 0)$; and finally, row reduce to $\begin{pmatrix} v' \\ w' \\ x' \end{pmatrix}$.

To find v' , we must find s, t such that $v_2 + sw_2 + tx_2 = 0$ and $v_3 + sw_3 + tx_3 = 0$. This amounts to solving the following non-homogeneous system of linear equations in two variables:

$$\begin{pmatrix} s & t \end{pmatrix} \begin{pmatrix} w_2 & w_3 \\ x_2 & x_3 \end{pmatrix} = - \begin{pmatrix} v_2 & v_3 \end{pmatrix}.$$

Of course, this system may not have any solutions, depending on the values of $w_2, w_3, x_2, x_3, v_2, v_3$. For now we assume that these are such that there is always a unique solution, and worry about degenerate cases later. Although we work with row vectors here instead of column vectors, we can still solve this system by using Exercise 15.2, which tells us that

$$\begin{pmatrix} w_2 & w_3 \\ x_2 & x_3 \end{pmatrix}^{-1} = \frac{1}{w_2x_3 - w_3x_2} \begin{pmatrix} x_3 & -w_3 \\ -x_2 & w_2 \end{pmatrix}.$$

Thus

$$\begin{aligned} \begin{pmatrix} s & t \end{pmatrix} &= \frac{- \begin{pmatrix} v_2 & v_3 \end{pmatrix}}{w_2x_3 - w_3x_2} \begin{pmatrix} x_3 & -w_3 \\ -x_2 & w_2 \end{pmatrix} \\ &= \frac{\begin{pmatrix} v_3x_2 - v_2x_3 & v_2w_3 - v_3w_2 \end{pmatrix}}{w_2x_3 - w_3x_2} \end{aligned}$$

We conclude that

$$\begin{aligned} a &= v_1 + sw_1 + tx_1 = v_1 + \frac{v_3w_1x_2 - v_2w_1x_3 + v_2w_3x_1 - v_3w_2x_1}{w_2x_3 - w_3x_2} \\ &= \frac{v_1w_2x_3 - v_1w_3x_2 + v_3w_1x_2 - v_2w_1x_3 + v_2w_3x_1 - v_3w_2x_1}{w_2x_3 - w_3x_2}. \end{aligned} \quad (15.3)$$

At this point we have successfully row reduced

$$A = \begin{pmatrix} v_1 & v_2 & v_3 \\ w_1 & w_2 & w_3 \\ x_1 & x_2 & x_3 \end{pmatrix} \rightarrow \begin{pmatrix} a & 0 & 0 \\ w_1 & w_2 & w_3 \\ x_1 & x_2 & x_3 \end{pmatrix}.$$

Note that the expression for a has the following properties: every term in the numerator is the product of exactly one entry from each of v, w, x , and the denominator is the determinant of the 2×2 matrix $\begin{pmatrix} w_2 & w_3 \\ x_2 & x_3 \end{pmatrix}$, which is the matrix lying directly below the two entries of A that we transformed to 0 in our row reduction.

We need to repeat the same procedure with w and x to find b, c , but now it turns out to be simpler. Indeed, w' is determined by the requirement that $w' = w + sv' + tx = (0, b, 0)$ for some $s, t \in \mathbb{R}$ (not necessarily the same as before). This gives

$$\begin{aligned} w_1 + sa + tx_1 &= 0 \\ w_2 + tx_2 &= b \\ w_3 + tx_3 &= 0. \end{aligned}$$

The third equation determines t , the second determines b , and the first determines s . We get $t = -\frac{w_3}{x_3}$, hence

$$b = w_2 - \frac{w_3}{x_3}x_2 = \frac{w_2x_3 - w_3x_2}{x_3}.$$

Note that the numerator of b matches the denominator of a . Now we have row reduced A to $\begin{pmatrix} a & 0 & 0 \\ 0 & b & 0 \\ x_1 & x_2 & x_3 \end{pmatrix}$, and we see immediately that row reducing to diagonal form does not change the bottom right entry, so $c = x_3$, and the original matrix A is row equivalent to the diagonal matrix $\text{diag}(a, b, c)$. In particular, the volume of the parallelepiped spanned by v, w, x is the same as the volume of the rectangular prism with side lengths a, b, c – that is, it is given by

$$\begin{aligned} abc &= \frac{v_1w_2x_3 - v_1w_3x_2 + v_3w_1x_2 - v_2w_1x_3 + v_2w_3x_1 - v_3w_2x_1}{w_2x_3 - w_3x_2} \cdot \frac{w_2x_3 - w_3x_2}{x_3} \cdot x_3 \\ &= v_1w_2x_3 - v_1w_3x_2 + v_3w_1x_2 - v_2w_1x_3 + v_2w_3x_1 - v_3w_2x_1. \end{aligned} \quad (15.4)$$

So far we cannot say anything rigorously about this quantity, because the discussion above did not properly account for the possibility that we may have $x_3 = 0$ or $w_2x_3 - w_3x_2 = 0$. Nevertheless, we may tentatively call this quantity the determinant of A , and see if it does what we expect it to do – that is, tell us whether or not A is invertible, and help in computing the inverse. Thus we make the following definition: given a 3×3 matrix A with entries A_{ij} , the *determinant* of A is

$$\begin{aligned} \det(A) &= A_{11}A_{22}A_{33} - A_{11}A_{23}A_{32} \\ &\quad + A_{13}A_{21}A_{32} - A_{12}A_{21}A_{33} + A_{12}A_{23}A_{31} - A_{13}A_{22}A_{31}. \end{aligned} \quad (15.5)$$

More about determinants

Further reading:

16.1 Two interpretations of determinant

If we look carefully at (15.5) or the equivalent formula (15.4), we see that the formula for the determinant can be analysed in two different ways. First we note that it can be expressed in terms of 2×2 determinants: writing

$$A = \begin{pmatrix} v_1 & v_2 & v_3 \\ w_1 & w_2 & w_3 \\ x_1 & x_2 & x_3 \end{pmatrix} = \begin{pmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \end{pmatrix},$$

we see from (15.5) and (15.4) that

$$\begin{aligned} \det A &= v_1(w_2x_3 - w_3x_2) - w_1(v_2x_3 - v_3x_2) + x_1(v_2w_3 - v_3w_2) \\ &= v_1 \det \begin{pmatrix} w_2 & w_3 \\ x_2 & x_3 \end{pmatrix} - w_1 \det \begin{pmatrix} v_2 & v_3 \\ x_2 & x_3 \end{pmatrix} + x_1 \det \begin{pmatrix} v_2 & v_3 \\ w_2 & w_3 \end{pmatrix} \\ &= A_{11} \det \begin{pmatrix} A_{22} & A_{23} \\ A_{32} & A_{33} \end{pmatrix} - A_{21} \det \begin{pmatrix} A_{12} & A_{13} \\ A_{32} & A_{33} \end{pmatrix} + A_{31} \det \begin{pmatrix} A_{12} & A_{13} \\ A_{22} & A_{23} \end{pmatrix} \\ &= A_{11} \det \tilde{A}_{11} - A_{21} \det \tilde{A}_{21} + A_{31} \det \tilde{A}_{31}, \end{aligned}$$

where \tilde{A}_{ij} denotes the 2×2 matrix obtained from A by deleting the i th row and j th column. This formula, defining the determinant of a 3×3 matrix in terms of the determinants of the 2×2 matrices \tilde{A}_{ij} , is called *cofactor expansion*, or *Laplace expansion*. We will later see that a generalisation of this can be used to define the determinant of an $n \times n$ matrix in terms of determinants of $(n-1) \times (n-1)$ matrices.

Another interpretation of the formula for a 3×3 determinant is illustrated by the following diagram, which relates each term of (15.5) to a choice of 3

entries on a 3×3 grid, one from each row and column.

$$\begin{array}{ccc}
 \begin{array}{|c|c|c|} \hline \bullet & & \\ \hline & \bullet & \\ \hline & & \bullet \\ \hline \end{array} & \begin{array}{|c|c|c|} \hline \bullet & & \\ \hline & & \bullet \\ \hline & \bullet & \\ \hline \end{array} & \begin{array}{|c|c|c|} \hline & & \bullet \\ \hline \bullet & & \\ \hline & & \bullet \\ \hline \end{array} \\
 A_{11}A_{22}A_{33} & -A_{11}A_{23}A_{32} & A_{13}A_{21}A_{32} \\
 \\
 \begin{array}{|c|c|c|} \hline & \bullet & \\ \hline \bullet & & \\ \hline & & \bullet \\ \hline \end{array} & \begin{array}{|c|c|c|} \hline & \bullet & \\ \hline & & \bullet \\ \hline \bullet & & \\ \hline \end{array} & \begin{array}{|c|c|c|} \hline & & \bullet \\ \hline & \bullet & \\ \hline \bullet & & \\ \hline \end{array} \\
 -A_{12}A_{21}A_{33} & A_{12}A_{23}A_{31} & -A_{13}A_{22}A_{31}
 \end{array} \tag{16.1}$$

Each of the six terms in (15.5) is a product of one entry from the first row of A , one entry from the second row, and one entry from the third row. Moreover, a similar statement is true of columns: each term involves exactly one matrix entry from each column. For a 3×3 matrix, there are exactly six ways to select three entries that represent each row and each column exactly once; the diagram above shows these six ways, together with the terms from $\det A$ that they correspond to.

Note that the first two grids correspond to the term $A_{11} \det \tilde{A}_{11}$ in the cofactor expansion; the next two correspond to $-A_{21} \det \tilde{A}_{21}$; and the final two correspond to $A_{31} \det \tilde{A}_{31}$. By pairing them differently, we see that the cofactor expansion can be carried out along any row or any column:

$$\begin{aligned}
 \det A &= A_{11} \det \tilde{A}_{11} - A_{21} \det \tilde{A}_{21} + A_{31} \det \tilde{A}_{31} \\
 &= -A_{12} \det \tilde{A}_{12} + A_{22} \det \tilde{A}_{22} - A_{32} \det \tilde{A}_{32} \\
 &= A_{13} \det \tilde{A}_{13} - A_{23} \det \tilde{A}_{23} + A_{33} \det \tilde{A}_{33} \\
 &= A_{11} \det \tilde{A}_{11} - A_{12} \det \tilde{A}_{12} + A_{13} \det \tilde{A}_{13} \\
 &= -A_{21} \det \tilde{A}_{21} + A_{22} \det \tilde{A}_{22} - A_{23} \det \tilde{A}_{23} \\
 &= A_{31} \det \tilde{A}_{31} - A_{32} \det \tilde{A}_{32} + A_{33} \det \tilde{A}_{33}.
 \end{aligned}$$

What's going on with the signs of the terms? Which ones are positive and which are negative? For the cofactor expansion, observe that the sign of the term $A_{ij} \tilde{A}_{ij}$ depends on whether $i + j$ is even or odd: it is positive if $i + j$ is even, and negative if $i + j$ is odd. We can capture this in a formula by writing the sign as $(-1)^{i+j}$, so that the cofactor expansion along the i th row is

$$\det A = \sum_{j=1}^3 (-1)^{i+j} A_{ij} \det \tilde{A}_{ij}, \tag{16.2}$$

and along the j th column we have

$$\det A = \sum_{i=1}^3 (-1)^{i+j} A_{ij} \det \tilde{A}_{ij}, \quad (16.3)$$

Pictorially, the positive and negative signs associated to the i, j -term follow the chessboard pattern shown:

$$\begin{array}{c|c|c} + & - & + \\ \hline - & + & - \\ \hline + & - & + \end{array} \quad (16.4)$$

What about (16.1) itself? How do we determine which terms are positive and which are negative?

16.2 Permutations on three symbols

To each of the configurations in (16.1) we can associate the 3×3 matrix which has 1s in the marked positions and 0s everywhere else: thus the first configuration, with markers on the main diagonal, corresponds to the identity matrix, the second configuration (that goes with $-A_{11}A_{23}A_{32}$) corresponds to $\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$, and so on.

Definition 16.1. A *permutation matrix* is an $n \times n$ matrix in which every row and column has exactly one entry equal to 1, and the rest equal to 0.

Writing $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ for the standard basis of row vectors in \mathbb{R}^3 , we can write the six 3×3 permutation matrices from (16.1), together with their corresponding terms in $\det A$, as

$$\begin{array}{l} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} \mathbf{e}_1 \\ \mathbf{e}_2 \\ \mathbf{e}_3 \end{pmatrix} \quad v_1 w_2 x_3 = A_{11} A_{22} A_{33} \\ \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} = \begin{pmatrix} \mathbf{e}_1 \\ \mathbf{e}_3 \\ \mathbf{e}_2 \end{pmatrix} \quad -v_1 w_2 x_3 = -A_{11} A_{22} A_{33} \\ \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} = \begin{pmatrix} \mathbf{e}_3 \\ \mathbf{e}_1 \\ \mathbf{e}_2 \end{pmatrix} \quad v_3 w_1 x_2 = A_{13} A_{21} A_{32} \end{array}$$

$$\begin{aligned}
\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} &= \begin{pmatrix} \mathbf{e}_2 \\ \mathbf{e}_1 \\ \mathbf{e}_3 \end{pmatrix} && -v_2w_1x_3 = -A_{12}A_{21}A_{33} \\
\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} &= \begin{pmatrix} \mathbf{e}_2 \\ \mathbf{e}_3 \\ \mathbf{e}_1 \end{pmatrix} && v_2w_3x_1 = A_{12}A_{23}A_{31} \\
\begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} &= \begin{pmatrix} \mathbf{e}_3 \\ \mathbf{e}_2 \\ \mathbf{e}_1 \end{pmatrix} && -v_3w_2x_1 = -A_{13}A_{22}A_{31}
\end{aligned}$$

The three permutation matrices corresponding to negative terms are $\begin{pmatrix} \mathbf{e}_1 \\ \mathbf{e}_3 \\ \mathbf{e}_2 \end{pmatrix}$, $\begin{pmatrix} \mathbf{e}_2 \\ \mathbf{e}_1 \\ \mathbf{e}_3 \end{pmatrix}$, and $\begin{pmatrix} \mathbf{e}_3 \\ \mathbf{e}_2 \\ \mathbf{e}_1 \end{pmatrix}$. Notice that each of these can be obtained from the identity matrix via a single row operation of swapping two rows. The identity matrix, of course, requires zero swaps, while the other two permutation matrices corresponding to positive terms require two swaps.

Exercise 16.2. Check this last claim, that $\begin{pmatrix} \mathbf{e}_3 \\ \mathbf{e}_1 \\ \mathbf{e}_2 \end{pmatrix}$ and $\begin{pmatrix} \mathbf{e}_2 \\ \mathbf{e}_3 \\ \mathbf{e}_1 \end{pmatrix}$ can be obtained from I by two row operations of exchanging rows.

The above discussion suggests the following criterion:

$$\begin{aligned}
\text{even number of swaps} &\Leftrightarrow \text{positive term} \\
\text{odd number of swaps} &\Leftrightarrow \text{negative term}
\end{aligned}$$

Let us formalise this notion. A *permutation* of $\{1, 2, 3\}$ is a bijection $\pi: \{1, 2, 3\} \rightarrow \{1, 2, 3\}$ – that is, a 1-1 and onto map from the set of indices $\{1, 2, 3\}$ to itself. This may be thought of as a rearrangement of these indices, so that in the new ordering $\pi(1)$ comes first, $\pi(2)$ comes second, and $\pi(3)$ third. Then each of the six permutation matrices listed above is given by a permutation π : for example,

$$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} = \begin{pmatrix} \mathbf{e}_2 \\ \mathbf{e}_3 \\ \mathbf{e}_1 \end{pmatrix} = \begin{pmatrix} \mathbf{e}_{\pi(1)} \\ \mathbf{e}_{\pi(2)} \\ \mathbf{e}_{\pi(3)} \end{pmatrix}, \quad \begin{aligned} \pi(1) &= 2 \\ \pi(2) &= 3 \\ \pi(3) &= 1 \end{aligned} \quad (16.5)$$

Definition 16.3. Let S_3 be the set of all permutations of three symbols. Given a permutation $\pi \in S_3$, the permutation matrix associated to π is

$$P(\pi) = \begin{pmatrix} \mathbf{e}_{\pi(1)} \\ \mathbf{e}_{\pi(2)} \\ \mathbf{e}_{\pi(3)} \end{pmatrix} \quad P(\pi)_{ij} = \begin{cases} 1 & \text{if } j = \pi(i) \\ 0 & \text{otherwise} \end{cases}$$

Each permutation matrix corresponds to the term in $\det A$ which is a product of the entries of A where the permutation matrix has a 1 – in terms of the permutation π , we see that π corresponds to the product $A_{1,\pi(1)}A_{2,\pi(2)}A_{3,\pi(3)}$. Moreover, we can determine the sign associated with this term as follows.

Describe π by the list $(\pi(1), \pi(2), \pi(3))$, which gives the rearrangement of $(1, 2, 3)$ described by π . Thus for example we write the permutation in (16.5) as $\pi = (2\ 3\ 1)$.

Observe that exchanging two rows of the permutation matrix corresponds to exchanging two symbols in this list. Let

$$\operatorname{sgn}(\pi) = \begin{cases} +1 & \text{if } \pi \text{ is reached in an even number of exchanges} \\ -1 & \text{if } \pi \text{ is reached in an odd number of exchanges} \end{cases} \quad (16.6)$$

be the *sign* of the permutation π . Then the determinant of the 3×3 matrix A is

$$\det A = \sum_{\pi \in S_3} \operatorname{sgn}(\pi) A_{1,\pi(1)} A_{2,\pi(2)} A_{3,\pi(3)}. \quad (16.7)$$

There is a potential problem with the definition in (16.6): the same permutation can be reached in different ways. For example, to reach the permutation $\pi = (2\ 3\ 1)$ from (16.5) from the original arrangement $(1\ 2\ 3)$, we can either swap the first two symbols to get $(2\ 1\ 3)$, and then the last two to get π , or we can swap the first and third to get $(3\ 2\ 1)$, and then the first two to get π . In this case both procedures involve the same number of swaps, but will we always be so lucky? What if there is some π such that one way of getting π involves an even number of swaps, and the other involves an odd number? What should $\operatorname{sgn}(\pi)$ be?

In fact, it turns out that for any permutation π , the number of exchanges in *any* process leading to π is either always even (in which case $\operatorname{sgn}(\pi) = 1$) or always odd (in which case $\operatorname{sgn}(\pi) = -1$), as we will see in the next section.

16.3 Oriented volumes in \mathbb{R}^n

Consider two vectors $v_1, v_2 \in \mathbb{R}^2$, and assume that they form a basis. Let $D(v_1, v_2)$ be the determinant of the matrix $A = \begin{pmatrix} v_1 \\ v_2 \end{pmatrix}$, which as we have seen is the signed area of the parallelogram spanned by v_1, v_2 . Note that this region can be described as

$$P(v_1, v_2) = \{t_1 v_1 + t_2 v_2 \mid t_1, t_2 \in [0, 1]\}.$$

In other words, $P(v_1, v_2)$ is the image of the square with vertices $(0, 0)$, $(1, 0)$, $(0, 1)$, and $(1, 1)$ under the linear transformation $x \mapsto xA$.

More generally, given $v_1, \dots, v_n \in \mathbb{R}^n$, the *parallelepiped spanned by* v_1, \dots, v_n is

$$\begin{aligned} P(v_1, \dots, v_n) &= \left\{ \sum_{i=1}^n t_i v_i \mid t_i \in [0, 1] \text{ for all } 1 \leq i \leq n \right\} \\ &= \{xA \mid x \in C\}, \end{aligned}$$

where $A = \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix}$ is the linear transformation sending \mathbf{e}_i to v_i , and C is the n -dimensional cube with side lengths 1 and vertices at $\mathbf{0}$ and \mathbf{e}_i for $1 \leq i \leq n$. (Note that C also has vertices at $\mathbf{e}_{i_1} + \dots + \mathbf{e}_{i_k}$ for every $1 \leq i_1 < i_2 < \dots < i_k \leq n$.)

For 2×2 matrices, recall that $\det \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \|v_1\| \|v_2\| \sin \theta$, where θ is the angle from v_1 to v_2 . In particular, $\sin \theta$ is positive if the shortest way from v_1 to v_2 is to go counterclockwise, and is negative if the shortest way is clockwise. Thus the sign of the determinant tells us what the *orientation* of the vectors v_1, v_2 is relative to each other. We write $S(v_1, v_2) = \pm 1$ depending on whether this orientation is positive or negative.

What about 3×3 matrices, or larger? Given a basis $v_1, v_2, v_3 \in \mathbb{R}^3$, we can no longer use the words “clockwise” and “counterclockwise” to describe where these vectors are in relation to each other. Instead we note the following: in \mathbb{R}^2 , the standard basis $\mathbf{e}_1, \mathbf{e}_2$ is positively oriented, and any positively oriented basis v_1, v_2 can be *continuously deformed* into the standard basis without becoming dependent. More precisely, if v_1, v_2 is a positively oriented basis, then there are continuous functions $w_1, w_2: [0, 1] \rightarrow \mathbb{R}^2$ such that

- $w_i(0) = v_i$ and $w_i(1) = \mathbf{e}_i$ for $i = 1, 2$;
- $w_1(t), w_2(t)$ is a basis for every $t \in [0, 1]$.

On the other hand, if v_1, v_2 is negatively oriented – for example, if $v_1 = \mathbf{e}_2$ and $v_2 = \mathbf{e}_1$ – then there is no way to accomplish this deformation without the vectors $w_1(t), w_2(t)$ becoming linearly dependent at some point. Thus for a basis $v_1, \dots, v_n \in \mathbb{R}^n$, we define orientation as follows: $S(v_1, \dots, v_n) = 1$ if there exist continuous functions $w_1, \dots, w_n: [0, 1] \rightarrow \mathbb{R}^n$ such that

- $w_i(0) = v_i$ and $w_i(1) = \mathbf{e}_i$ for $i = 1, \dots, n$;
- $w_1(t), \dots, w_n(t)$ is a basis for every $t \in [0, 1]$.

If no such continuous functions exist, then $S(v_1, \dots, v_n) = -1$. If v_1, \dots, v_n is not a basis, we put $S(v_1, \dots, v_n) = 0$.

Exercise 16.4. Show that $S(\mathbf{e}_{\pi(1)}, \mathbf{e}_{\pi(2)}, \mathbf{e}_{\pi(3)}) = \text{sgn}(\pi)$ for every $\pi \in S_3$.

Now we make the key definition: given vectors $v_1, \dots, v_n \in \mathbb{R}^n$, let

$$D(v_1, \dots, v_n) = S(v_1, \dots, v_n) \text{Vol}(P(v_1, \dots, v_n)) \quad (16.8)$$

be the signed volume of the parallelepiped spanned by v_1, \dots, v_n . Given a matrix $A \in \mathbb{M}_{n \times n}$, let v_1, \dots, v_n be the row vectors of A , and define the determinant of A to be

$$\det A = D(v_1, \dots, v_n). \quad (16.9)$$

Earlier, we saw that this definition of determinant implies the formula (15.5) for 3×3 matrices, which we rewrote in two different ways, (16.2)–(16.3) and (16.7). Now we will show that (16.8) and (16.9) force the determinant to have several properties that in turn force it to be given by a generalisation of (16.7), the sum over permutations, and that this is equivalent to a recursive definition following (16.2)–(16.3).

A general approach to determinants

Further reading:

17.1 Properties of determinants

[Lax] p. 44–46

The definition of $D(v_1, \dots, v_n)$ in (16.8) forces the function $D: (\mathbb{R}^n)^n \rightarrow \mathbb{R}$ to have the following properties.

1. If $v_i = v_j$ for some $i \neq j$, then $D(v_1, \dots, v_n) = 0$.
2. $D(v_1, \dots, v_n)$ is linear in each of its arguments – that is, for every $1 \leq j \leq n$ and every $v_1, \dots, v_{j-1}, v_{j+1}, \dots, v_n$, the function $T: x \mapsto D(v_1, \dots, v_{j-1}, x, v_{j+1}, \dots, v_n)$ is a linear function of x .
3. $D(\mathbf{e}_1, \dots, \mathbf{e}_n) = 1$.

To see the first property, note that if $v_i = v_j$ for $i \neq j$, then the parallelepiped $P(v_1, \dots, v_n)$ lies in an $(n - 1)$ -dimensional subspace of \mathbb{R}^n , and hence has zero volume. To see the third property, note that $P(\mathbf{e}_1, \dots, \mathbf{e}_n)$ is the unit n -cube, and this basis is positively oriented by definition. It only remains to prove the second property, linearity in each argument.

Given $1 \leq j \leq n$ and v_1, \dots, v_n , let V be the $(n - 1)$ -dimensional volume of the parallelepiped spanned by $v_1, \dots, v_{j-1}, v_{j+1}, \dots, v_n$. Given $x \in \mathbb{R}^n$, let $h(x)$ be the signed distance from x to the subspace in which this parallelepiped lies – that is, $h(x)$ is positive if $S(v_1, \dots, v_{j-1}, x, v_{j+1}, \dots, v_n) = 1$, and negative otherwise. Then h is a linear function of x , and we have

$$T(x) = D(v_1, \dots, v_{j-1}, x, v_{j+1}, \dots, v_n) = h(x)V,$$

so T is linear in x .

There are two more important properties of determinants that follow from the three listed above.

4. D is an *alternating* function of the vectors v_1, \dots, v_n ; that is, if v_i and v_j are interchanged, then D changes by a factor of -1 .
5. If v_1, \dots, v_n are linearly dependent, then $D(v_1, \dots, v_n) = 0$.

To prove that D is alternating, we fix v_k for $k \neq i, j$ and write

$$T(w, x) = D(v_1, \dots, v_{i-1}, w, v_{i+1}, \dots, v_{j-1}, x, v_{j+1}, \dots, v_n),$$

so that we are trying to show $T(x, w) = -T(w, x)$. Using property 1, we have

$$T(w, x) + T(x, w) = T(w, x) + T(w, w) + T(x, x) + T(x, w),$$

and by property 2 (multilinearity), this is equal to

$$T(w, x + w) + T(x, x + w) = T(x + w, x + w),$$

which vanishes by property 1. Thus D is alternating.

To prove that D vanishes on a linearly dependent set, we note that by linear dependence, one of the v_i s can be written as a linear combination of the others, so there are coefficients $c_j \in \mathbb{R}$ such that

$$v_i = \sum_{j \neq i} c_j v_j.$$

Let $T(w) = D(v_1, \dots, v_{i-1}, w, v_{i+1}, \dots, v_n)$, then we have

$$D(v_1, \dots, v_n) = T(v_i) = T\left(\sum_{j \neq i} c_j v_j\right) = \sum_{j \neq i} c_j T(v_j) = 0,$$

where the last two equalities use property 2 followed by property 1.

17.2 A formula for determinant

[Lax] p. 46–49

Now we will see that properties 1–3 can be used to derive a formula for the function D , and hence for the determinant. Given vectors $v_1, \dots, v_n \in \mathbb{R}^n$, write $v_i = (v_{i1}, \dots, v_{in})$, so that $v_i = \sum_{j=1}^n v_{ij} \mathbf{e}_j$. Now by property 2 (multilinearity), we can expand along the coefficients of v_1 to get

$$\begin{aligned} D(v_1, \dots, v_n) &= D\left(\sum_{j=1}^n v_{1j} \mathbf{e}_j, v_2, \dots, v_n\right) \\ &= \sum_{j=1}^n v_{1j} D(\mathbf{e}_j, v_2, \dots, v_n). \end{aligned}$$

Expanding this along the coefficients of $v_2 = v_{21}\mathbf{e}_1 + \cdots + v_{2n}\mathbf{e}_n$ gives

$$\begin{aligned} D(v_1, \dots, v_n) &= \sum_{j=1}^n v_{1j} D\left(\mathbf{e}_j, \sum_{k=1}^n v_{2k}\mathbf{e}_k, v_3, \dots, v_n\right) \\ &= \sum_{j=1}^n \sum_{k=1}^n v_{1j} v_{2k} D(\mathbf{e}_j, \mathbf{e}_k, v_3, \dots, v_n). \end{aligned}$$

Continuing, we eventually get

$$D(v_1, \dots, v_n) = \sum_{j_1=1}^n \sum_{j_2=1}^n \cdots \sum_{j_n=1}^n v_{1j_1} v_{2j_2} \cdots v_{nj_n} D(\mathbf{e}_{j_1}, \dots, \mathbf{e}_{j_n}). \quad (17.1)$$

Let's consider this sum for a moment. Every term in the sum is a product of n entries v_{ij} multiplied by the factor $D(\mathbf{e}_{j_1}, \dots, \mathbf{e}_{j_n})$. The n entries v_{ij} in any given term have the property that one of them is an entry in v_1 , one is an entry in v_2 , and so on. If we consider the $n \times n$ matrix A whose rows are the vectors v_1, \dots, v_n , then this corresponds to taking exactly one entry from each row of A .

We can rewrite (17.1) as

$$D(v_1, \dots, v_n) = \sum_f v_{1,f(1)} \cdots v_{n,f(n)} D(\mathbf{e}_{f(1)}, \dots, \mathbf{e}_{f(n)}), \quad (17.2)$$

where the sum is taken over all functions $f: \{1, \dots, n\} \rightarrow \{1, \dots, n\}$. If f is not 1-1, then there are indices $i \neq j$ such that $\mathbf{e}_{f(i)} = \mathbf{e}_{f(j)}$, and hence by property 1, $D(\mathbf{e}_{f(1)}, \dots, \mathbf{e}_{f(n)}) = 0$. Thus the only non-zero terms in the above sum are those corresponding to 1-1 functions f . A function from an n -element set to itself is 1-1 if and only if it is onto; in this case it is called a *permutation* of the set, as in the 3×3 case we discussed earlier. Let S_n denote the set of all permutations on n symbols; then since the only functions that contribute to the sum in (17.2) are permutations, we have

$$D(v_1, \dots, v_n) = \sum_{\pi \in S_n} D(\mathbf{e}_{\pi(1)}, \dots, \mathbf{e}_{\pi(n)}) \prod_{i=1}^n v_{i,\pi(i)} \quad (17.3)$$

Now we will use property 3 of the function D (normalisation), together with property 5 (alternating). Property 3 says that $D(\mathbf{e}_1, \dots, \mathbf{e}_n) = 1$. Property 5 says that every time we interchange two vectors, we reverse the sign of D . Thus $D(\mathbf{e}_2, \mathbf{e}_1, \mathbf{e}_3, \mathbf{e}_4, \dots, \mathbf{e}_n) = -1$, and so on. In particular, if it is possible to reach the arrangement $\mathbf{e}_{\pi(1)}, \dots, \mathbf{e}_{\pi(n)}$ in an even number of swaps, then

$D(\mathbf{e}_{\pi(1)}, \dots, \mathbf{e}_{\pi(n)}) = 1$, whereas if it takes an odd number of swaps, then $D(\mathbf{e}_{\pi(1)}, \dots, \mathbf{e}_{\pi(n)}) = -1$.

This is exactly the way we defined the sign of a permutation earlier: $\text{sgn } \pi = 1$ if $(\pi(1), \dots, \pi(n))$ can be obtained from $(1, \dots, n)$ in an even number of swaps, and $\text{sgn } \pi = -1$ if it takes an odd number. Thus (17.3) can be rewritten as

$$D(v_1, \dots, v_n) = \sum_{\pi \in S_n} \text{sgn } \pi \prod_{i=1}^n v_{i, \pi(i)}. \quad (17.4)$$

Given an $n \times n$ matrix A , let v_1, \dots, v_n be the row vectors of A , so $A_{ij} = v_{ij}$. Then using the definition of determinant given by $\det(A) = D(v_1, \dots, v_n)$, we can write (17.4) as

$$\det A = \sum_{\pi \in S_n} \text{sgn } \pi \prod_{i=1}^n A_{i, \pi(i)}, \quad (17.5)$$

which generalises (16.7) from the 3×3 case.

There is still an issue to be addressed, which is the problem of well-definedness: given a permutation π , what happens if we choose two different ways to get to the arrangement $(\pi(1), \dots, \pi(n))$ from the original arrangement $(1, \dots, n)$? Can they give two different answers for $\text{sgn } \pi$? We will address this in the next section. First we observe that if we write $P(\pi)$ for the permutation matrix corresponding to π – that is, the $n \times n$ matrix whose rows are given by $\mathbf{e}_{\pi(i)}$, so that $P(\pi)_{ij} = 1$ if $j = \pi(i)$ and 0 otherwise – then

$$\det P(\pi) = D(\mathbf{e}_{\pi(1)}, \dots, \mathbf{e}_{\pi(n)}) = \text{sgn } \pi.$$


Thus the sign of a permutation is given by the determinant of the corresponding permutation matrix, and vice versa.

Yet more on determinants

Further reading:

18.1 Signs of permutations


Let π be a permutation on n symbols. To settle the question of whether $\text{sgn } \pi$ is well defined, we describe another way of determining whether π is even ($\text{sgn } \pi = 1$) or odd ($\text{sgn } \pi = -1$). Write two rows of numbers: the first row is the numbers $1, \dots, n$, and the second is $\pi(1), \dots, \pi(n)$. For each i , draw an arrow from the top row to the bottom row connecting the two places where i occurs. For example, the permutation $\pi = (2, 4, 1, 3)$,¹ which has $\pi(1) = 2$, $\pi(2) = 4$, $\pi(3) = 1$, and $\pi(4) = 3$, corresponds to the following picture.



Now count the number of times the arrows cross (in the above example, there are three crossings). We claim that the number of crossings has the same parity (even or odd) as the number of transpositions it takes to obtain π . To see this, choose any way of obtaining π via transpositions: in the above example, one possibility would be

$$(1, 2, 3, 4) \rightarrow (2, 1, 3, 4) \rightarrow (2, 4, 3, 1) \rightarrow (2, 4, 1, 3).$$

Draw a similar diagram to (18.1), where each transposition gets its own row:



¹For any readers who are familiar with the representation of a permutation in terms of its cycle structure, notice that we are using a similar-looking notation here to mean a different thing.

Note that we are careful to draw the arrows in such a way that no more than two arrows intersect at a time.² (This is why the second row has a curved arrow.) This diagram has more intersections than (18.1), but the parity of intersections is the same. To see that this is true in general, say that i and j are *inverted* by π if $\pi(i), \pi(j)$ occur in a different order than i, j . That is, i, j are inverted if $i < j$ and $\pi(i) > \pi(j)$, or if $i > j$ and $\pi(i) < \pi(j)$.

Proposition 18.1. *For any choice of how to draw the arrows in the above diagram, and any $i \neq j$, the arrows corresponding to i and j cross an odd number of times if i, j are inverted by π , and an even number of times if they are not inverted.*

Proof. Every time the arrows for i, j cross, the order of i, j is reversed. If this order is reversed an odd number of times, then i, j are inverted. If it is reversed an even number of times, then i, j are not inverted. \square

Example 18.2. The permutation $(2, 4, 1, 3)$ has three inversions, corresponding to the pairs of arrows starting out at $(1, 2)$, at $(1, 4)$, and at $(3, 4)$.

Using Proposition 18.1, we see that the number of intersections always has the same parity as the number of inversions (or “backwards pairs”). This is useful because the number of inversions, unlike the number of transpositions or the number of intersections in the diagram, does not depend on any choices we make – it is completely determined by the permutation. In particular, its parity is determined by the permutation, and since this parity matches the parity of the number of transpositions or the number of intersections, we have proved most of the following result.

Proposition 18.3. *If k is the number of inversions in a permutation π – that is, the number of pairs $i < j$ for which $\pi(i) > \pi(j)$ – then $\operatorname{sgn} \pi = (-1)^k$.*

If π can be reached in two different ways by a sequence of k transpositions and a sequence of ℓ transpositions, then $\operatorname{sgn} \pi = (-1)^k = (-1)^\ell$.

Similarly, if π leads to a diagram such as (18.1) with k intersections, and another such diagram with ℓ intersections, then $\operatorname{sgn} \pi = (-1)^k = (-1)^\ell$.

Proof. The only thing left to prove is that the number of intersections has the same parity as the number of transpositions. This follows because every diagram as in (18.1) can be expanded to a diagram as in (18.2), in which every row corresponds to a single transposition, and every row has an odd number of intersections. \square

²For the sake of completeness, we also specify that two arrows do not touch unless they actually cross each other, and that arrows are not allowed to go above the first row or below the last row.

Exercise 18.4. Let $S(v_1, \dots, v_n)$ be the orientation function defined earlier, and show that $S(\mathbf{e}_{\pi(1)}, \dots, \mathbf{e}_{\pi(n)}) = \text{sgn } \pi$ for every permutation π .

We close this section by remarking that Proposition 18.3 is the tool we need to complete the definition of determinant. Before this point, we had shown that any function $D(v_1, \dots, v_n)$ satisfying the three properties – multilinearity, normalisation, and vanishing when $v_i = v_j$ – must be given by (17.5), but we did not know if (17.5) made sense until we could guarantee that $\text{sgn } \pi$ was well-defined.

18.2 Determinant of transpose

Let π be a permutation on n symbols, and $P(\pi)$ the associated permutation matrix, so that the i th row of $P(\pi)$ is $\mathbf{e}_{\pi(i)}$; equivalently, $P(\pi)_{ij} = 1$ if $j = \pi(i)$, and 0 otherwise. The inverse permutation π^{-1} is such that $\pi^{-1}(i) = j$ whenever $\pi(j) = i$, and so its associated permutation matrix is given by

$$\begin{aligned} P(\pi^{-1})_{ij} &= \begin{cases} 1 & \text{if } \pi(j) = i \\ 0 & \text{otherwise} \end{cases} \\ &= P(\pi)_{ji}. \end{aligned}$$

We conclude that $P(\pi^{-1}) = P(\pi)^t$. Moreover, an arrowed diagram as in (18.1) for π^{-1} can be obtained from such a diagram for π by reversing the directions of the arrows, which does not change the number of intersections, and thus $\text{sgn}(\pi^{-1}) = \text{sgn } \pi$. This will be important in the proof of the following result.

Proposition 18.5. *For every $A \in \mathbb{M}_{n \times n}$, we have $\det(A^t) = \det A$.*

Proof. Using (17.5), we have

$$\det(A^t) = \sum_{\pi \in S_n} \text{sgn } \pi \prod_{i=1}^n A_{i, \pi(i)}^t = \sum_{\pi \in S_n} \text{sgn } \pi \prod_{i=1}^n A_{\pi(i), i}.$$

Note that $\prod_{i=1}^n A_{\pi(i), i} = \prod_{j=1}^n A_{j, \pi^{-1}(j)}$, and since $\text{sgn}(\pi^{-1}) = \text{sgn } \pi$, we have

$$\det(A^t) = \sum_{\pi \in S_n} \text{sgn}(\pi^{-1}) \prod_{i=1}^n A_{j, \pi^{-1}(j)}.$$

As π ranges over all permutations in S_n , so does π^{-1} , and so this is equal to $\det A$. \square

Corollary 18.6. *Let A be an $n \times n$ matrix with rows v_1, \dots, v_n and columns w_1, \dots, w_n . Then $\det A = D(v_1, \dots, v_n) = D(w_1, \dots, w_n)$, where $D: (K^n)^n \rightarrow K$ has the properties listed in §17.1.*

Proof. The first equality is the definition of the determinant, and the second follows from the first because w_1, \dots, w_n are the rows of A^t , which has the same determinant as A by Proposition 18.5. \square

18.3 Determinant of product

[Lax] p. 49–51

An important additional property of determinants is the following.

Theorem 18.7. *For all $A, B \in \mathbb{M}_{n \times n}$, we have $\det(AB) = (\det A)(\det B)$.*

Proof. Recall that the j th column of AB is $(AB)\mathbf{e}_j$, and the j th column of B is $B\mathbf{e}_j$. Thus writing b_1, \dots, b_n for the columns of B , we have $(AB)\mathbf{e}_j = A(B\mathbf{e}_j) = Ab_j$, and by Corollary 18.6 we have

$$\det(AB) = D(Ab_1, \dots, Ab_n). \quad (18.3)$$

Define $C: (K^n)^n \rightarrow K$ by

$$C(b_1, \dots, b_n) = D(Ab_1, \dots, Ab_n) \quad (18.4)$$

We claim that the function C has the first two properties described in §17.1, and that $C(\mathbf{e}_1, \dots, \mathbf{e}_n) = \det A$. Then it will follow from the discussion in §17.2 that $C(b_1, \dots, b_n) = (\det A)D(b_1, \dots, b_n)$, which will suffice to prove the theorem.

To verify these properties, first observe that if $b_i = b_j$ for some $b_i \neq b_j$, then $Ab_i = Ab_j$, and so the numerator in (18.4) vanishes. This takes care of property 1. Property 2 follows because Ab_i depends linearly on b_i , and the right-hand side of (18.4) depends linearly on Ab_i .

Now we observe that if a_1, \dots, a_n are the columns of A , then

$$C(\mathbf{e}_1, \dots, \mathbf{e}_n) = D(A\mathbf{e}_1, \dots, A\mathbf{e}_n) = D(a_1, \dots, a_n) = \det A,$$

where the second equality follows from matrix multiplication, and the third from the definition of determinant.

As indicated above, §17.2 now gives

$$C(b_1, \dots, b_n) = (\det A)D(b_1, \dots, b_n) = (\det A)(\det B)$$

which together with (18.3) completes the proof. \square

Corollary 18.8. *An $n \times n$ matrix A is invertible if and only if $\det A \neq 0$. If it is invertible, then $\det(A^{-1}) = \det(A)^{-1}$.*

Proof. (\Leftarrow): If A is non-invertible then its columns are not a basis for K^n , and so they are linearly dependent. Thus Property 5 of the function D implies that $\det A = 0$.

(\Rightarrow): If A is invertible, then there is $B \in \mathbb{M}_{n \times n}$ such that $AB = I$, and so by Theorem 18.7 we have

$$1 = \det I = \det(AB) = (\det A)(\det B),$$

implying that $\det A \neq 0$, and that $\det B = 1/\det A$. □

Corollary 18.9. *If two matrices A and B are similar, then $\det A = \det B$.*

Proof. By similarity there is an invertible $Q \in \mathbb{M}_{n \times n}$ such that $B = Q^{-1}AQ$. By Corollary 18.8 we have

$$\det B = \det(Q^{-1}AQ) = \frac{1}{\det Q}(\det A)(\det Q) = \det A. \quad \square$$

Other ways of computing determinant

Further reading:

19.1 Cofactor expansion

[Lax]p. 51–53

Recall that one of the ways we were able to interpret the expression for a 3×3 determinant was the Laplace expansion, or cofactor expansion, which gave the determinant of a 3×3 matrix in terms of the 2×2 determinants of matrices obtained by deleting a row and a column. It is natural to ask if we can interpret an $n \times n$ determinant in terms of $(n-1) \times (n-1)$ determinants in the same way, and indeed we can. As before, we write $\tilde{A}_{ij} \in \mathbb{M}_{(n-1) \times (n-1)}$ for the matrix obtained from A by deleting the i th row and j th column.

Theorem 19.1. *For any $A \in \mathbb{M}_{n \times n}$, the determinant is given by*

$$\det A = \sum_{i=1}^n (-1)^{i+1} A_{i1} \det \tilde{A}_{i1}. \quad (19.1)$$

Proof. Let $a_1, \dots, a_n \in \mathbb{R}^n$ be the columns of A . Then $a_1 = \sum_{i=1}^n A_{i1} \mathbf{e}_i$, and thus

$$\det A = D(a_1, \dots, a_n) = \sum_{i=1}^n A_{i1} D(\mathbf{e}_i, a_2, \dots, a_n),$$

where the first equality is the definition of determinant, and the second follows from multilinearity of D . To prove (19.1), we need to show that

$$D(\mathbf{e}_i, a_2, \dots, a_n) = (-1)^{i+1} \det \tilde{A}_{i1}. \quad (19.2)$$

Note that by Properties 1 and 2 of the function D , the value of D does not change if we add a multiple of one vector to another: in particular, for every $2 \leq j \leq n$ and every $\lambda \in K$, we have

$$\begin{aligned} D(\mathbf{e}_i, a_2, \dots, a_j + \lambda \mathbf{e}_i, \dots, a_n) &= D(\mathbf{e}_i, a_2, \dots, a_j, \dots, a_n) \\ &\quad + \lambda D(\mathbf{e}_i, a_2, \dots, \mathbf{e}_i, \dots, a_n) \\ &= D(\mathbf{e}_i, a_2, \dots, a_n) \end{aligned}$$

In particular, we can add $-A_{ij}\mathbf{e}_i$ to a_j for each j and obtain

$$D(\mathbf{e}_i, a_2, \dots, a_n) = \det \begin{pmatrix} 0 & A_{12} & \cdots & A_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & A_{i-1,2} & \cdots & A_{i-1,n} \\ 1 & 0 & \cdots & 0 \\ 0 & A_{i+1,2} & \cdots & A_{i+1,n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & A_{n2} & \cdots & A_{nn} \end{pmatrix}. \quad (19.3)$$

Notice that we have set every entry in the i th row and 1st column to 0, except for the entry in the place $(i, 1)$, which is now equal to 1. The remaining entries determine the $(n-1) \times (n-1)$ matrix \tilde{A}_{i1} . It follows from multilinearity of determinant that the right-hand side of (19.3) is linear in each column of \tilde{A}_{i1} , and hence $D(\mathbf{e}_i, a_2, \dots, a_n)$ is as well. Similarly, $D(\mathbf{e}_i, a_2, \dots, a_n) = 0$ if any two columns of \tilde{A}_{i1} are equal. Because these properties characterise determinant up to a constant normalising factor, we conclude that

$$D(\mathbf{e}_i, a_2, \dots, a_n) = D(\mathbf{e}_i, \mathbf{e}_1, \dots, \mathbf{e}_{i-1}, \mathbf{e}_{i+1}, \dots, \mathbf{e}_n) \det \tilde{A}_{i1}.$$

We can go from the order $(i, 1, \dots, i-1, i+1, \dots, n)$ to the order $(1, 2, \dots, n)$ in $i-1$ swaps, and (19.2) follows. \square

This gives the cofactor (Laplace) expansion along the first column: the only difference from the 3×3 case is that the sum has n terms, instead of 3, and the matrix A_{i1} is an $(n-1) \times (n-1)$ matrix, instead of a 2×2 matrix.

Because interchanging two columns reverses the sign of the determinant, we immediately obtain from Theorem 19.1 the formula for cofactor expansion along the j th column as

$$\det A = \sum_{i=1}^n (-1)^{i+j} A_{ij} \det \tilde{A}_{ij}.$$

Similarly, because determinant is preserved under taking transposes, we can use cofactor expansion along the i th row, getting

$$\det A = \sum_{j=1}^n (-1)^{i+j} A_{ij} \det \tilde{A}_{ij}.$$

19.2 Determinants and row reduction

The formulas for determinant in terms of a sum over permutation and in terms of a cofactor expansion are very important theoretically, and will continue to play a key role for us when we return to the study of eigenvalues and eigenvectors. However, as the size of the matrix increases, these formulas quickly become inefficient from the point of view of computation. Indeed, in each case the number of calculations to be performed in the evaluation of an $n \times n$ determinant is on the order of $n!$, which grows very quickly.

A more efficient way to compute determinant is via row reduction. Recall that there are three types of row operations:

1. add a multiple of a row to another row;
2. interchange two rows;
3. multiply a row by a nonzero scalar.

We saw in the proof of Theorem 19.1 that

$$\begin{aligned} D(v_1, \dots, v_i + \lambda v_1, \dots, v_n) &= D(v_1, \dots, v_i, \dots, v_n) \\ &\quad + \lambda D(v_1, \dots, v_{i-1}, v_1, v_{i+1}, \dots, v_n) \\ &= D(v_1, \dots, v_n), \end{aligned}$$

and a similar computation shows that for *any* $i \neq j$, we can replace v_i with $v_i + \lambda v_j$ without changing the value of D . In particular, if A row reduces to B using only the first type of row operation – adding a multiple of a row to another row – then $\det A = \det B$.

Furthermore, we know that interchanging two rows multiplies the determinant by -1 . Thus we have the following result.

Proposition 19.2. *Suppose A can be row reduced to B using only row operations of Types 1 and 2. Let k be the number of row operations of Type 2 required to carry out this row reduction. Then $\det A = (-1)^k \det B$.*

This gives a more efficient way of computing determinants. If A is non-invertible, then it row reduces to a matrix with an empty row, and we conclude that $\det A = 0$. If A is invertible, then it row reduces to a diagonal matrix $\text{diag}(\lambda_1, \dots, \lambda_n)$, and we conclude that

$$\det A = (-1)^k \det \text{diag}(\lambda_1, \dots, \lambda_n), \quad (19.4)$$

where k is the number of transpositions involved in the row reduction.

Exercise 19.3. Show that $\det \text{diag}(\lambda_1, \dots, \lambda_n) = \prod_{i=1}^n \lambda_i$.

Using this exercise, we see that (19.4) reduces to

$$\det A = (-1)^k \lambda_1 \cdots \lambda_n.$$

It can be shown that the number of row reductions involved in carrying out the above process for an $n \times n$ matrix is on the order of n^3 , which gives a much more manageable number of computations than a direct application of the formulas via summing over permutations or cofactor expansion. Indeed, for $n = 10$ we have $n^3 = 1000$, while $n! \approx 3.6 \times 10^6$, so the two differ by three orders of magnitude. For $n = 20$ we have $n^3 = 8000$, while $n! \approx 2.4 \times 10^{18}$. At a billion operations a second, it would take 1/100,000 of a second to carry out 20^3 operations, but about 76 years to carry out $20!$ operations. Thus row reduction is a vastly more efficient way of computing large determinants than the previous formulas we saw.

Exercise 19.4. Suppose A can be row reduced to B via *any* sequence of row operations; let k be the number of transpositions (Type 2) involved, and let p_1, \dots, p_m be the non-zero factors by which rows are multiplied in the m steps of Type 3. Use multilinearity of determinant to show that

$$\det B = (-1)^k p_1 \cdots p_m \det A,$$

and in particular, if A row reduces to the identity through this process, then

$$\det A = (-1)^k (p_1 \cdots p_m)^{-1}.$$

In fact, it suffices to row reduce to an upper triangular matrix, thanks to the following exercise.

Exercise 19.5. Let A be upper triangular, and show that $\det A = \prod_{i=1}^n A_{ii}$.

Example 19.6. Consider the matrix

$$A = \begin{pmatrix} 1 & -2 & 3 & -12 \\ -5 & 12 & -14 & 19 \\ -9 & 22 & -20 & 31 \\ -4 & 9 & -14 & 15 \end{pmatrix}.$$

Row operations of the first type reduce this to

$$\begin{pmatrix} 1 & -2 & 3 & -12 \\ 0 & 2 & 1 & -41 \\ 0 & 4 & 7 & -77 \\ 0 & 1 & -2 & -33 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & -2 & 3 & -12 \\ 0 & 0 & 5 & 25 \\ 0 & 0 & 15 & 55 \\ 0 & 1 & -2 & -33 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & -2 & 3 & -12 \\ 0 & 0 & 5 & 25 \\ 0 & 0 & 0 & -20 \\ 0 & 1 & -2 & -33 \end{pmatrix}.$$

With two more transpositions, this becomes an upper triangular matrix with diagonal entries 1, 1, 5, -20 , and so

$$\det A = (-1)^2(1)(1)(5)(-20) = -100.$$

Inverses, trace, and characteristic polynomial

Further reading:

20.1 Matrix inverses

[Lax] p. 53–54

We can use determinants to solve non-homogeneous systems of linear equations. Suppose $A \in \mathbb{M}_{n \times n}$ is invertible, and consider the equation $Ax = u$, where $u \in K^n$ is given and $x \in K^n$ is unknown. Because A is invertible, this has a unique solution for every u . But how to express it?

Let $a_j \in K^n$ be the j th column of A , so that $Ae_j = a_j$. The unknown vector x can be expressed as $x = \sum_{j=1}^n x_j e_j$, and thus

$$u = Ax = \sum_{j=1}^n x_j (Ae_j) = \sum_{j=1}^n x_j a_j. \quad (20.1)$$

To find x_k , we need to find something to do to both sides of the above equation that eliminates all the terms $x_j a_j$ for $j \neq k$, so that the only remaining unknown is x_k . This can be accomplished by using the fact that $D(a_1, \dots, a_{k-1}, a_j, a_{k+1}, \dots, a_n) = 0$ for all $j \neq k$. More precisely, let $B_k = [a_1 \mid \dots \mid a_{k-1} \mid u \mid a_{k+1} \mid \dots \mid a_n]$ be the matrix obtained from A by replacing the k th column with the vector u . Then we have

$$\begin{aligned} \det B_k &= D(a_1, \dots, a_{k-1}, u, a_{k+1}, \dots, a_n) \\ &= D\left(a_1, \dots, a_{k-1}, \sum_{j=1}^n x_j a_j, a_{k+1}, \dots, a_n\right) \\ &= \sum_{j=1}^n x_j D(a_1, \dots, a_{k-1}, a_j, a_{k+1}, \dots, a_n) \\ &= x_k \det A, \end{aligned}$$

where the first equality is the definition of determinant, the second equality uses (20.1), the third uses multilinearity of determinant, and the fourth uses the fact that D vanishes whenever two of its arguments coincide. We conclude that $x_k = (\det B_k)/(\det A)$, since A is invertible and hence $\det A \neq 0$.

To compute $\det B_k$ in terms of A , we use cofactor expansion along the k th column, and observe that $(\tilde{B}_k)_{ik} = A_{ik}$ since deleting the k th column of B_k yields the same matrix as deleting the k th column of A . Thus

$$\det B_k = \sum_{i=1}^n (-1)^{i+k} u_i \det \tilde{A}_{ik},$$

since $(B_k)_{ik} = u_i$ by the definition of B_k . Together with the computation above, this gives

$$x_k = \sum_{i=1}^n (-1)^{i+k} \frac{\det \tilde{A}_{ik}}{\det A} u_i, \quad (20.2)$$

which is known as *Cramer's rule*.

We can use Cramer's rule to give a formula for the inverse of a matrix. Given an invertible matrix A , let $B = A^{-1}$ and let b_1, \dots, b_n be the columns of B . Then the condition $AB = I$ can be rewritten as the system of equations

$$Ab_1 = \mathbf{e}_1, \dots, Ab_n = \mathbf{e}_n.$$

Applying Cramer's rule to the equations $Ab_j = \mathbf{e}_j$ for $1 \leq j \leq n$, we get

$$B_{kj} = (b_j)_k = \sum_{i=1}^n (-1)^{i+k} \frac{\det \tilde{A}_{ik}}{\det A} (\mathbf{e}_j)_i = (-1)^{j+k} \frac{\det \tilde{A}_{jk}}{\det A}.$$

Notice that the order of the indices is reversed between the left and right-hand sides of the equation. We conclude that the inverse of a matrix A is given by the formula

$$(A^{-1})_{ij} = (-1)^{i+j} \frac{\det \tilde{A}_{ji}}{\det A}. \quad (20.3)$$

For large matrices, this is not a practical method for computing the inverse, because it involves computing n^2 determinants, and the computation of inverse via row reduction of $[A \mid I]$ to $[I \mid A^{-1}]$ is significantly faster. Nevertheless, (20.3) has important theoretical significance.

20.2 Trace

[Lax] p. 55–56

The determinant is a function that assigns a scalar value to every $n \times n$ matrix. Another important scalar-valued function is the *trace*, defined as

$$\operatorname{Tr} A = \sum_{i=1}^n A_{ii}. \quad (20.4)$$

That is, the trace of an $n \times n$ matrix A is the sum of the diagonal entries of A . (This looks like the formula for a determinant of an upper-triangular matrix, where determinant is the product of the diagonal entries, but this definition of trace is for *any* square matrix A , regardless of whether or not it is upper-triangular.)

Determinant is a multilinear function – it depends linearly on each row/column of A . Trace, on the other hand, is a linear function of the entire matrix A . Indeed, given any $A, B \in \mathbb{M}_{n \times n}$ and $c \in K$, we have

$$\operatorname{Tr}(cA + B) = \sum_{i=1}^n (cA + B)_{ii} = \sum_{i=1}^n cA_{ii} + B_{ii} = c \operatorname{Tr} A + \operatorname{Tr} B.$$

Trace is not multiplicative in the same way determinant is: as a general rule, $\operatorname{Tr}(AB) \neq \operatorname{Tr}(A) \operatorname{Tr}(B)$. (Indeed, $\operatorname{Tr}(I) = n$ for the $n \times n$ identity matrix, so $\operatorname{Tr}(IB) = \operatorname{Tr}(B) \neq \operatorname{Tr}(I) \operatorname{Tr}(B)$ for every $B \in \mathbb{M}_{n \times n}$ with non-zero trace.) Nevertheless, trace has the following important property with respect to multiplication.

Theorem 20.1. *Given any $A, B \in \mathbb{M}_{n \times n}$, we have $\operatorname{Tr}(AB) = \operatorname{Tr}(BA)$.*

Proof. By the formula for matrix multiplication, we have $(AB)_{ii} = \sum_{k=1}^n A_{ik} B_{ki}$, and so

$$\operatorname{Tr}(AB) = \sum_{i=1}^n \sum_{k=1}^n A_{ik} B_{ki}.$$

Similarly, $(BA)_{ii} = \sum_{k=1}^n B_{ik} A_{ki}$, and so

$$\operatorname{Tr}(BA) = \sum_{i=1}^n \sum_{k=1}^n B_{ik} A_{ki}.$$

The two sums are equal (just interchange the names of the indices). \square

We saw in Corollary 18.9 that similar (conjugate) matrices have the same determinant; this followed from the fact that determinant is multiplicative. Even though trace is not multiplicative, Theorem 20.1 leads to the analogous result for similar matrices.

Corollary 20.2. *Similar matrices have the same trace.*

Proof. Let $A, B \in \mathbb{M}_{n \times n}$ be such that $B = QAQ^{-1}$ for some invertible $Q \in \mathbb{M}_{n \times n}$. Then applying Theorem 20.1 to the matrices QA and Q^{-1} , we have

$$\operatorname{Tr} B = \operatorname{Tr}((QA)(Q^{-1})) = \operatorname{Tr}((Q^{-1})(QA)) = \operatorname{Tr} A. \quad \square$$

Exercise 20.3. Corollary 20.2 used the fact that $\text{Tr}(ABC) = \text{Tr}(BCA)$. Is it always true that $\text{Tr}(ABC) = \text{Tr}(BAC)$? If so, prove it. If not, provide a counterexample.

20.3 Characteristic polynomial

[Lax] p. 60–63

Now that we have equipped ourselves with the tools of determinant and trace, we return to the question of finding eigenvalues and eigenvectors for a square matrix A . That is, given $A \in \mathbb{M}_{n \times n}$, for which $\lambda \in K$ and $v \in K^n \setminus \{\mathbf{0}\}$ do we have $Av = \lambda v$?

Recall that this can be rephrased as the equation $(A - \lambda I)v = \mathbf{0}$, and that consequently $\lambda \in K$ is an eigenvalue for A if and only if the matrix $A - \lambda I$ is non-invertible. By Theorem 18.8, $A - \lambda I$ is non-invertible if and only if $\det(\lambda I - A) = 0$. This gives us a method of finding the eigenvalues.

Example 20.4. Let $A = \begin{pmatrix} 3 & 1 \\ 2 & 4 \end{pmatrix}$. Then we have

$$\begin{aligned} \det(\lambda I - A) &= \det \begin{pmatrix} \lambda - 3 & -1 \\ -2 & \lambda - 4 \end{pmatrix} \\ &= (\lambda - 3)(\lambda - 4) - (-1)(-2) = \lambda^2 - 7\lambda + 10. \end{aligned}$$

Thus the eigenvalues of A are precisely the roots of $\lambda^2 - 7\lambda + 10 = 0$. The quadratic polynomial factors as $(\lambda - 5)(\lambda - 2)$, so the eigenvalues of A are 2 and 5. The null spaces of $A - 2I$ and $A - 5I$ give us eigenvectors:

$$A - 2I = \begin{pmatrix} 1 & 1 \\ 2 & 2 \end{pmatrix}, \quad A - 5I = \begin{pmatrix} -2 & 1 \\ 2 & -1 \end{pmatrix},$$

thus $\lambda_1 = 2$ has a corresponding eigenvector $v_1 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$, and $\lambda_2 = 5$ has a corresponding eigenvector $v_2 = \begin{pmatrix} 1 \\ 2 \end{pmatrix}$.

Note that v_1, v_2 are linearly independent, hence they form a basis for \mathbb{R}^2 . In particular, any $v \in \mathbb{R}^2$ can be written as $v = a_1 v_1 + a_2 v_2$ for some $a_1, a_2 \in \mathbb{R}$, and then we have

$$A^N v = a_1 (A^N v_1) + a_2 (A^N v_2) = a_1 \lambda_1^N v_1 + a_2 \lambda_2^N v_2 = a_1 2^N v_1 + a_2 5^N v_2.$$

The above example illustrates a powerful procedure – if we can find a basis for K^n that consists of eigenvectors for the matrix A , then we can write down an explicit formula for the iterates $A^N v$ of any initial vector v .

Example 20.5. Recall that the Fibonacci numbers defined by $x_1 = x_2 = 1$ and $x_n = x_{n-1} + x_{n-2}$ can be defined in terms of matrix multiplication by putting $w_n = \begin{pmatrix} x_n \\ x_{n+1} \end{pmatrix}$ and $A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$, so that $w_n = A^n \begin{pmatrix} 0 \\ 1 \end{pmatrix}$. To find the eigenvalues of the matrix A , we look for $\lambda \in \mathbb{R}$ such that

$$0 = \det(\lambda I - A) = \det \begin{pmatrix} \lambda & -1 \\ -1 & \lambda - 1 \end{pmatrix} = \lambda^2 - \lambda - 1. \quad (20.5)$$

By the quadratic equation, the eigenvalues are $\lambda_1 = \frac{1}{2}(1 + \sqrt{5})$ and $\lambda_2 = \frac{1}{2}(1 - \sqrt{5})$. To find eigenvectors for λ_i , we find the null space of $A - \lambda_i I = \begin{pmatrix} -\lambda_i & 1 \\ 1 & 1 - \lambda_i \end{pmatrix}$. We see that $v_i = \begin{pmatrix} 1 \\ \lambda_i \end{pmatrix}$ is an eigenvector for λ_i : the fact that $(A - \lambda_i I)v_i$ has zero first coordinate is clear, and we see that the second coordinate is $1 + (1 - \lambda_i)\lambda_i = 1 + \lambda_i - \lambda_i^2$, which is 0, because λ_i is a root of (20.5).

Now to find $w_n = A^n \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, we find coefficients $a_1, a_2 \in \mathbb{R}$ such that

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} = a_1 v_1 + a_2 v_2 = a_1 \begin{pmatrix} 1 \\ \lambda_1 \end{pmatrix} + a_2 \begin{pmatrix} 1 \\ \lambda_2 \end{pmatrix}.$$

That is, we find the coordinate representation of $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ relative to the basis v_1, v_2 . From the first coordinate of the above equation we have $a_2 = -a_1$, and so the equation in the second coordinate becomes $1 = a_1(\lambda_1 - \lambda_2) = a_1\sqrt{5}$. We conclude that

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \frac{v_1 - v_2}{\sqrt{5}} \quad \Rightarrow \quad w_n = A^n \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \frac{A^n v_1 - A^n v_2}{\sqrt{5}} = \frac{\lambda_1^n v_1 - \lambda_2^n v_2}{\sqrt{5}}.$$

In particular, since x_n is the first component of w_n , and since v_1, v_2 both have a 1 in their first component, we get the following formula for the n th Fibonacci number:

$$x_n = \frac{\lambda_1^n - \lambda_2^n}{\sqrt{5}}. \quad (20.6)$$

Notice that λ_1 is equal to the golden ratio $\varphi = \frac{1+\sqrt{5}}{2} \approx 1.618$. Observe that the sum of the two eigenvalues is 1, and so the second eigenvalue is $1 - \varphi$. Dividing the quadratic equation $\varphi^2 - \varphi - 1 = 0$ by φ , we get $\varphi - 1 - 1/\varphi = 0$, and in particular we see that $\lambda_2 = 1 - \varphi = -1/\varphi$. Thus (20.6) can be rewritten as

$$x_n = \frac{\varphi^n - (-1)^n \varphi^{-n}}{\sqrt{5}} = \frac{\varphi^n}{\sqrt{5}} - (-1)^n \frac{\varphi^{-n}}{\sqrt{5}}$$

Note that $\varphi > 1$ and $\varphi^{-1} \approx .618$, so the second term in the above formula is always $< \frac{1}{2}$, and goes to 0 as $n \rightarrow \infty$. We can thus make the following

two conclusions: first, the ratio between successive Fibonacci numbers is

$$\begin{aligned} \frac{x_{n+1}}{x_n} &= \frac{\varphi^{n+1} - (-1)^{n+1}\varphi^{-(n+1)}}{\varphi^n - (-1)^n\varphi^{-n}} \\ &= \varphi \left(\frac{\varphi^n - (-1)^n\varphi^{-n} + (-1)^n\varphi^{-n} - (-1)^{n+1}\varphi^{-(n+2)}}{\varphi^n - (-1)^n\varphi^{-n}} \right) \\ &= \varphi \left(1 + (-1)^n \frac{\varphi^{-n} + \varphi^{-(n+2)}}{\varphi^n - (-1)^n\varphi^{-n}} \right), \end{aligned}$$

which converges to the golden ratio φ as $n \rightarrow \infty$; second, the n th Fibonacci number is the closest integer to $\varphi^n/\sqrt{5} = \frac{1}{\sqrt{5}} \left(\frac{1+\sqrt{5}}{2} \right)^n$.

In both of the above examples, the quantity $\det(\lambda I - A)$ turned out to be a quadratic polynomial in the variable λ . This was because A was a 2×2 matrix. The following example shows what happens when we consider a 3×3 matrix.

Example 20.6. Let $A = \begin{pmatrix} 2 & 0 & 0 \\ 1 & 3 & -1 \\ 2 & 2 & 0 \end{pmatrix}$. Then the condition $\det(\lambda I - A) = 0$ becomes

$$\begin{aligned} 0 &= \det(\lambda I - A) = \det \begin{pmatrix} \lambda - 2 & 0 & 0 \\ -1 & \lambda - 3 & 1 \\ -2 & -2 & \lambda \end{pmatrix} \\ &= (\lambda - 2)(\lambda - 3)\lambda + (0)(1)(-2) + (0)(-1)(-2) \\ &\quad - (\lambda - 2)(-2)(1) - (-1)(0)\lambda - (-2)(\lambda - 3)(0) \\ &= (\lambda^3 - 5\lambda^2 + 6\lambda) + (2\lambda - 4) \\ &= \lambda^3 - 5\lambda^2 + 8\lambda - 4. \end{aligned}$$

Thus we must find the roots of a cubic polynomial. Observe that $\lambda = 1$ is a root of this polynomial, and we have

$$\lambda^3 - 5\lambda^2 + 8\lambda - 4 = (\lambda - 1)(\lambda^2 - 4\lambda + 4) = (\lambda - 1)(\lambda - 2)^2.$$

Thus the eigenvalues of A are $\lambda = 1$ and $\lambda = 2$. To find the eigenspace for $\lambda = 1$, we row reduce

$$A - I = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 2 & -1 \\ 2 & 2 & -1 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & -1 \\ 0 & 0 & 0 \end{pmatrix}$$

and see that the eigenspace is spanned by $v_1 = \begin{pmatrix} 0 \\ 1 \\ 2 \end{pmatrix}$. The eigenspace for $\lambda = 2$ is found by row reducing

$$A - 2I = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 1 & -1 \\ 2 & 2 & -2 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 1 & -1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

so the eigenspace is spanned by $v_2 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$ and $v_3 = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}$. We see that v_1, v_2, v_3 is a basis for \mathbb{R}^3 consisting of eigenvectors of A , even though A only has two eigenvalues.

The examples we have done so far all have the property that there is a basis for K^n consisting of eigenvectors for the matrix A . Whenever this is the case, we can write any vector $v \in K^n$ as $v = \sum_i a_i v_i$, where $v_i \in K^n$ is an eigenvector for $\lambda_i \in K$, and $a_i \in K$ are scalar coefficients. Note that some eigenvalues λ_i may be repeated, as in Example 20.6; however, the eigenvectors v_i must be linearly independent. Then we have

$$\begin{aligned} Av &= \sum_{i=1}^n a_i (Av_i) = \sum_{i=1}^n (a_i \lambda_i) v_i, \\ A^N v &= \sum_{i=1}^n a_i (A^N v_i) = \sum_{i=1}^n (a_i \lambda_i^N) v_i, \end{aligned}$$

which allows us to quickly determine how A and A^N act on any vector v .

Thus we have the following important question: Does every matrix A have the property that there is a basis for K^n consisting of eigenvectors for A ? If not, how do we determine which matrices have this property?

The first of these questions is answered by the following exercise. The second will be answered later.

Exercise 20.7. Consider the matrix $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$. Show that $\lambda = 1$ is the only eigenvalue of A . Show that \mathbb{R}^2 does *not* have a basis consisting of eigenvectors for A .

We conclude this lecture with the following observation. When we found the eigenvalues of a 2×2 matrix, we saw that $\det(\lambda I - A)$ was a quadratic polynomial in λ . For a 3×3 matrix, we saw that $\det(\lambda I - A)$ was a cubic polynomial in λ . More generally, we observe that for an $n \times n$ matrix A ,

$$\det(\lambda I - A) = \sum_{\pi \in S_n} \operatorname{sgn} \pi \prod_{i=1}^n (\lambda I - A)_{i, \pi(i)}.$$

We see that $(\lambda I - A)_{ij} = \lambda - A_{ij}$ if $i = j$, and $-A_{ij}$ otherwise; consequently, the product $\prod_{i=1}^n (\lambda I - A)_{i, \pi(i)}$ is a polynomial in λ , whose degree is equal to the number of indices i for which $\pi(i) = i$. In particular, it is a polynomial with degree $\leq n$, and thus $\det(\lambda I - A)$ is also a polynomial in λ with degree $\leq n$.

In fact, a little more thought shows that there is exactly one permutation for which $\prod_{i=1}^n (\lambda I - A)_{i, \pi(i)}$ has degree n – this is the identity permutation, $\pi(i) = i$ for all i . The leading coefficient here is equal to 1 – that is, the product is equal to $\lambda^n + \text{something with smaller degree}$, and so we see that $\det(\lambda I - A)$ is a polynomial of degree exactly n , and the coefficient on the highest order term (λ^n) is equal to 1. We call this polynomial the *characteristic polynomial* of the matrix A . We will sometimes write it as $p_A(\lambda)$, so that

$$p_A(\lambda) = \det(\lambda I - A). \quad (20.7)$$

Exercise 20.8. Show that $p_A(\lambda) = (-1)^n \det(A - \lambda I)$.

Diagonalisability

Further reading:

21.1 Spectrum of a matrix

Recall that $\lambda \in K$ is an eigenvalue of an $n \times n$ matrix A if and only if $A - \lambda I$ is non-invertible (has non-trivial nullspace) – this is equivalent to the condition that $p_A(\lambda) = \det(\lambda I - A) = 0$, and we see that the eigenvalues are the roots of the characteristic polynomial.

Definition 21.1. The *spectrum* of A is the set of its eigenvalues. The *algebraic multiplicity* of an eigenvalue is the number of times it appears as a root of the characteristic polynomial. The *geometric multiplicity* of an eigenvalue is the dimension of the eigenspace corresponding to it.

Example 21.2. Let A be the 3×3 matrix from Example 20.6, then the spectrum of A is $\{1, 2\}$. The eigenvalue $\lambda = 1$ has algebraic and geometric multiplicities equal to 1, and the eigenvalue $\lambda = 2$ has algebraic and geometric multiplicities equal to 2.

Exercise 21.3. Let A be an upper triangular matrix. Show that the eigenvalues of A are precisely the diagonal entries of A , and that the algebraic multiplicity of an eigenvalue is the number of times it appears on the diagonal of A .

In Exercise 20.7, the eigenvalue $\lambda = 1$ has algebraic multiplicity 2 but geometric multiplicity 1, so the two numbers are not always the same. We will come back to the relationship between these two numbers later on. First we establish some general results, beginning with the fact that eigenvectors for different eigenvalues are linearly independent.

21.2 Bases of eigenvectors

Theorem 21.4. *If $\lambda_1, \dots, \lambda_k$ are distinct eigenvalues for a matrix A ($\lambda_i \neq \lambda_j$ for $i \neq j$), and v_1, \dots, v_k are eigenvectors for these eigenvalues, then v_1, \dots, v_k are linearly independent.*

Proof. By the definition of eigenvector, we have $Av_i = \lambda_i v_i$ for each $1 \leq i \leq k$. Suppose that v_1, \dots, v_k are linearly dependent. Then there is an m

such that $v_m \in \text{span}\{v_1, \dots, v_{m-1}\}$. Consider the smallest such m with this property. We will use the fact that v_1, \dots, v_m are eigenvectors for distinct eigenvalues to show that the same property holds for some $1 \leq \ell < m$, which will be a contradiction. To this end, let $a_1, \dots, a_{m-1} \in K$ be such that

$$v_m = a_1 v_1 + \dots + a_{\ell} v_{\ell}, \quad (21.1)$$

where $1 \leq \ell < m$ and $a_{\ell} \neq 0$. Consider two equations derived from this one; multiplying both sides of (21.1) by A and using the eigenvector property gives

$$\lambda_m v_m = a_1 \lambda_1 v_1 + \dots + a_{\ell} \lambda_{\ell} v_{\ell}, \quad (21.2)$$

while multiplying both sides of (21.1) by λ_m gives

$$\lambda_m v_m = a_1 \lambda_m v_1 + \dots + a_{\ell} \lambda_m v_{\ell}. \quad (21.3)$$

Subtracting (21.3) from (21.2) gives

$$\mathbf{0} = a_1(\lambda_1 - \lambda_m)v_1 + \dots + a_{\ell}(\lambda_{\ell} - \lambda_m)v_{\ell}.$$

Because $\lambda_{\ell} \neq \lambda_m$ and $a_{\ell} \neq 0$, this implies that $v_{\ell} \in \text{span}\{v_1, \dots, v_{\ell-1}\}$, which contradicts our assumption that m was the smallest index with this property. We conclude that v_1, \dots, v_k are linearly independent. \square

Corollary 21.5. *If $A \in \mathbb{M}_{n \times n}$ and the characteristic polynomial p_A has n distinct roots, then A has n linearly independent eigenvectors, which form a basis for K^n .*

Proof. Suppose $p_A(\lambda) = (\lambda - \lambda_1) \cdots (\lambda - \lambda_n)$, where the λ_i are distinct. Let v_i be an eigenvector for λ_i . By Theorem 21.4, v_1, \dots, v_n are linearly independent. Because $\dim K^n = n$, they form a basis. \square

We saw earlier that similar matrices have the same determinant and trace. They also have the same characteristic polynomial, eigenvalues, and there is a clear relationship between their eigenspaces.

Theorem 21.6. *If A and B are similar via a change-of-coordinates matrix Q – that is, if $B = Q A Q^{-1}$ – then A and B have the same characteristic polynomial. Thus every eigenvalue of A is also an eigenvalue of B , with the same algebraic multiplicity. Moreover, the geometric multiplicities agree, and if $E_{\lambda}^A = N_{A - \lambda I}$ is the eigenspace for A corresponding to λ , then E_{λ}^A and E_{λ}^B are related by $E_{\lambda}^B = Q(E_{\lambda}^A)$.*

Proof. First observe that

$$\begin{aligned} p_B(\lambda) &= \det(\lambda I - B) = \det(\lambda I - Q^{-1}AQ) \\ &= \det(Q^{-1}(\lambda I - A)Q) = \det(\lambda I - A) = p_A(\lambda), \end{aligned}$$

using the fact that similar matrices have the same determinant. This immediately implies that A and B have the same eigenvalues, with the same algebraic multiplicities. Now given an eigenvalue λ for A and the corresponding eigenspace E_λ^A , we see that for any $v \in E_\lambda^A$, we have

$$B(Qv) = (QAQ^{-1})(Qv) = QAv = Q(\lambda v) = \lambda(Qv),$$

so $Qv \in E_\lambda^B$. This shows that $Q(E_\lambda^A) \subset E_\lambda^B$. Similarly, if $w \in E_\lambda^B$, then writing $v = Q^{-1}w$, we see that

$$Av = AQ^{-1}w = Q^{-1}BQQ^{-1}w = Q^{-1}Bw = Q^{-1}\lambda w = \lambda(Q^{-1}w) = \lambda v,$$

so $v \in E_\lambda^A$ and hence $w = Qv \in Q(E_\lambda^A)$. This proves the theorem. \square

We described earlier how a basis of eigenvectors can be used to compute $A^N v$ for arbitrary vectors v . It can also be used to perform a change of coordinates in which A takes a particularly simple form.

Definition 21.7. A matrix A is *diagonalisable* if there is an invertible matrix Q such that $D = Q^{-1}AQ$ is diagonal. Similarly, a linear transformation $T \in \mathbb{L}(V)$ is diagonalisable if there is a basis β for V such that $[T]_\beta$ is diagonal.

Theorem 21.8. A is diagonalisable if and only if K^n has a basis consisting of eigenvectors for A . The entries of the corresponding diagonal matrix are the eigenvalues of A .

Proof. (\Rightarrow) Suppose there is an invertible matrix Q such that $D = Q^{-1}AQ$ is diagonal. Then $A = QDQ^{-1}$. Note that a diagonal matrix $D = \text{diag}(\lambda_1, \dots, \lambda_n)$ has eigenvalues λ_j with eigenvectors \mathbf{e}_j . Let $v_j = Q\mathbf{e}_j$ be the j th column of Q , then $\mathbf{e}_j = Q^{-1}v_j$, and so

$$Av_j = QDQ^{-1}v_j = QD\mathbf{e}_j = Q\lambda_j\mathbf{e}_j = (\lambda_j)Q\mathbf{e}_j = \lambda_j v_j.$$

Because v_1, \dots, v_n are the columns of an invertible matrix, they form a basis.

(\Leftarrow) Let $\beta = (v_1, \dots, v_n)$ be such a basis, with eigenvalues $\lambda_1, \dots, \lambda_n$. Let Q be the matrix with column vectors v_1, \dots, v_n . Then Q is invertible because v_1, \dots, v_n is a basis. Note that if α is the standard ordered basis,

then $Q = I_\beta^\alpha$. This is because Q has the property that $Q\mathbf{e}_j = v_j$ for each j , and similarly $Q^{-1}v_j = \mathbf{e}_j$. Let $D = Q^{-1}AQ$. Then

$$D\mathbf{e}_j = Q^{-1}AQ\mathbf{e}_j = Q^{-1}Av_j = Q^{-1}(\lambda_j v_j) = \lambda_j(Q^{-1}v_j) = \lambda_j\mathbf{e}_j.$$

Because $D\mathbf{e}_j$ gives the j th column of D , we see that D is diagonal with entries $\lambda_1, \dots, \lambda_n$. \square

By Theorem 21.8, the matrix in Exercise 20.7 is not diagonalisable. It is an important question what to do with non-diagonalisable matrices, and we will come back to it.

Corollary 21.9. *If the characteristic polynomial of A factors as a product of linear polynomials with distinct roots, then A is diagonalisable.*

Proof. If $p_A(\lambda) = \prod_{j=1}^n (\lambda - \lambda_j)$, where $\lambda_i \neq \lambda_j$ for $i \neq j$, then each λ_j is an eigenvalue since $\det(\lambda_j I - A) = 0$, and hence has an eigenvector $v_j \in K^n$. The eigenvectors v_1, \dots, v_n are linearly independent by Theorem 21.4. Thus they form a basis, so A is diagonalisable by Theorem 21.8. \square

21.3 Multiplicities

We observed above that the algebraic and geometric multiplicities can differ. However, there is a universal relationship between them: the geometric multiplicity is always less than or equal to the algebraic multiplicity. We need the following exercise.

Exercise 21.10. Let A be a square matrix with the block form $A = \begin{pmatrix} X & Y \\ \mathbf{0} & Z \end{pmatrix}$, where X, Z are square matrices and $Y, \mathbf{0}$ have the appropriate dimensions. Show that $\det(A) = \det(X)\det(Z)$.

Theorem 21.11. *Given an $n \times n$ matrix A and an eigenvalue λ of A , let $m_a(\lambda)$ and $m_g(\lambda)$ be the algebraic and geometric multiplicities, respectively. Then $1 \leq m_g(\lambda) \leq m_a(\lambda)$.*

Proof. Let E_λ^A be the eigenspace for A corresponding to λ , and let $m = m_g(\lambda) = \dim E_\lambda^A$ be the geometric multiplicity of λ . We must show that $m_a(\lambda) \geq m$. Let v_1, \dots, v_m be a basis for E_λ^A . Extend it to a basis v_1, \dots, v_n for K^n . Let $Q \in \mathbb{M}_{n \times n}$ be the invertible matrix with column vectors v_1, \dots, v_n , and let $B = Q^{-1}AQ$. Then for each $1 \leq j \leq m$, we have

$$B\mathbf{e}_j = Q^{-1}AQ\mathbf{e}_j = Q^{-1}Av_j = Q^{-1}\lambda v_j = \lambda(Q^{-1}v_j) = \lambda\mathbf{e}_j,$$

and thus B has the block form

$$B = \begin{pmatrix} \lambda I_m & X \\ \mathbf{0} & Y \end{pmatrix},$$

where I_m is the $m \times m$ identity matrix, $\mathbf{0}$ is the $(n - m) \times m$ zero matrix, X is a $m \times (n - m)$ matrix, and Y is an $(n - m) \times (n - m)$ matrix. Using Exercise 21.10 and the fact that similar matrices have the same characteristic polynomial, we have

$$\begin{aligned} p_A(t) = p_B(t) &= \det \begin{pmatrix} (t - \lambda)I_m & -X \\ \mathbf{0} & tI_{n-m} - Y \end{pmatrix} \\ &= \det((t - \lambda)I_m) \det(tI_{n-m} - Y) = (t - \lambda)^m p_Y(t). \end{aligned}$$

This shows that λ appears at least m times as a root of $p_A(t)$, and so $m_a(t) \geq m = m_g(t)$. \square

More spectral theory

Further reading:

22.1 More on multiplicities

Given a matrix $A \in \mathbb{M}_{n \times n}$, let $M_g(A)$ be the maximum number of linearly independent eigenvectors of A . That is, $M_g(A)$ is the largest number m such that there are $v_1, \dots, v_m \in K^n$ for which each v_j is an eigenvector of A , and v_1, \dots, v_m are linearly independent. Theorem 21.8 shows that A is diagonalisable if and only if $M_g(A) = n$. Let $\sigma(A)$ be the spectrum of A – that is, the set of eigenvalues. The following result says that $M_g(A)$ is given by the sum of the geometric multiplicities of the eigenvalues of A .

Proposition 22.1. $M_g(A) = \sum_{\lambda \in \sigma(A)} m_g(\lambda)$.

Proof. For each eigenvalue $\lambda \in \sigma(A)$, let $v_1^\lambda, \dots, v_{m_g(\lambda)}^\lambda$ be a basis for the eigenspace E_λ^A . Then $\{v_j^\lambda \mid \lambda \in \sigma(A), 1 \leq j \leq m_g(\lambda)\}$ has $\sum_{\lambda \in \sigma(A)} m_g(\lambda)$ elements. Moreover, it is linearly independent by Theorem 21.4 and the fact that each $v_1^\lambda, \dots, v_{m_g(\lambda)}^\lambda$ is linearly independent. Thus $M_g(A) \geq \sum_{\lambda \in \sigma(A)} m_g(\lambda)$.

Now we prove the other inequality. Suppose v_1, \dots, v_m are linearly independent eigenvectors for A . We must show that $m \leq \sum_{\lambda \in \sigma(A)} m_g(\lambda)$. Because every v_j must be contained in some E_λ^A , it suffices to show that any linearly independent set $w_1, \dots, w_k \in E_\lambda^A$ has $k \leq m_g(\lambda)$. But this follows immediately from the fact that $m_g(\lambda) = \dim(E_\lambda^A)$. \square

Proposition 22.1 has the following important consequence.

Corollary 22.2. *A matrix A is diagonalisable if and only if $m_g(\lambda) = m_a(\lambda)$ for every $\lambda \in \sigma(A)$.*

Proof. By Theorem 21.8 and Proposition 22.1, A is diagonalisable if and only if $\sum_{\lambda \in \sigma(A)} m_g(\lambda) = n$. By the definition of algebraic multiplicity, we see that $\sum_{\lambda} m_a(\lambda) = n$. Thus A is diagonalisable if and only if $\sum_{\lambda} m_g(\lambda) = \sum_{\lambda} m_a(\lambda)$. By Theorem 21.11, we have $m_g(\lambda) \leq m_a(\lambda)$ for every λ , and so equality occurs in the sum if and only if $m_g(\lambda) = m_a(\lambda)$ for every λ . \square

Informally, we can think of Theorem 21.11 and Corollary 22.2 as saying the following: the algebraic multiplicity of an eigenvalue says how many linearly independent eigenvectors that eigenvalue *should* have (according to the

characteristic polynomial), while the geometric multiplicity says how many it *does* have, and a matrix is diagonalisable if and only if every eigenvalue has all the eigenvectors it is “supposed to”.

22.2 Trace and determinant

From now on, unless otherwise stated, we will always consider our scalar field to be \mathbb{C} , the complex numbers. The reason for doing this is the fundamental theorem of algebra, which states that every polynomial factors into linear terms over \mathbb{C} . In particular, given any $A \in \mathbb{M}_{n \times n}(\mathbb{C})$, we have

$$p_A(\lambda) = \det(\lambda I - A) = (\lambda - \lambda_1) \cdots (\lambda - \lambda_n)$$

for some $\lambda_1, \dots, \lambda_n \in \mathbb{C}$. The eigenvalues λ_j may not be distinct, so Corollary 21.9 does not necessarily apply. However, if A is diagonalisable, we can make the following observations: A is similar to a diagonal matrix D , hence it has the same eigenvalues, trace, and determinant as D . Since the diagonal entries of a diagonal matrix D are the eigenvalues, we see that the trace of D is the sum of the eigenvalues, and the determinant is the product of the eigenvalues. Because A and D have the same eigenvalues, trace, and determinant, this continues to be true for A .

In fact, this relationship continues to hold even when the matrix A is not diagonalisable.

Theorem 22.3. *Let $A \in \mathbb{M}_{n \times n}$ have eigenvalues $\lambda_1, \dots, \lambda_n$, where the number of times each value appears in the list is equal to its algebraic multiplicity. Then $\text{Tr}(A) = \sum_{j=1}^n \lambda_j$ and $\det(A) = \prod_{j=1}^n \lambda_j$.*

Proof. Let the characteristic polynomial of A be given as

$$p_A(t) = t^n + a_{n-1}t^{n-1} + a_{n-2}t^{n-2} + \cdots + a_1t + a_0.$$

Because $(t - \lambda_j)$ is a factor of p_A for every $1 \leq j \leq n$, we have

$$p_A(t) = \prod_{j=1}^n (t - \lambda_j) = t^n - \left(\sum_{j=1}^n \lambda_j \right) t^{n-1} + \cdots + (-1)^n \prod_{j=1}^n \lambda_j.$$

That is $a_{n-1} = -\sum_{j=1}^n \lambda_j$ and $a_0 = (-1)^n \prod_{j=1}^n \lambda_j$. Thus to prove the theorem it suffices to show that $a_{n-1} = -\text{Tr}(A)$ and $a_0 = (-1)^n \det A$.

Using the formula for determinant as a sum over permutations, we get

$$\begin{aligned}
 p_A(t) &= \det(tI - A) = \det \begin{pmatrix} t - A_{11} & -A_{12} & \cdots & -A_{1n} \\ -A_{21} & t - A_{22} & \cdots & \vdots \\ \vdots & \vdots & \cdots & \vdots \\ -A_{n1} & \cdots & \cdots & t - A_{nn} \end{pmatrix} \\
 &= \sum_{\pi \in S_n} \operatorname{sgn}(\pi) \prod_{j=1}^n (tI - A)_{j, \pi(j)} \\
 &= \sum_{\pi \in S_n} \operatorname{sgn}(\pi) \prod_{j=1}^n (t\delta_{j, \pi(j)} - A_{j, \pi(j)}),
 \end{aligned} \tag{22.1}$$

where δ_{jk} is 1 if $j = k$ and 0 otherwise. (This is the *Kronecker delta*.) Every permutation other than the identity ($\pi(j) = j$ for all j) has at least two values of j for which $\delta_{j, \pi(j)} = 0$, and hence does not contribute to the terms of $p_A(t)$ with degree $\geq n - 1$. Thus

$$\begin{aligned}
 p_A(t) &= \left(\prod_{j=1}^n (t - A_{jj}) \right) + (\text{something with degree at most } n - 2) \\
 &= t^n - \left(\sum_{j=1}^n A_{jj} \right) t^{n-1} + (\text{something with degree at most } n - 2).
 \end{aligned}$$

This gives the desired form for a_{n-1} . For a_0 , we observe that the only term in (22.1) without any factor of t comes from choosing $-A_{j, \pi(j)}$ in every term of the product, and so

$$a_0 = \sum_{\pi \in S_n} \operatorname{sgn} \pi \prod_{j=1}^n (-A_{j, \pi(j)}) = (-1)^n \det A.$$

This completes the proof. \square

22.3 Spectral mapping and Cayley–Hamilton theorems

Recall that if $q(t) = c_0 + c_1t + \cdots + c_d t^d$ is a polynomial in t , then given a matrix $A \in \mathbb{M}_{n \times n}$, we write

$$q(A) = c_0 I + c_1 A + c_2 A^2 + \cdots + c_d A^d,$$

so that $q(A)$ is also an $n \times n$ matrix. There is a relationship between the eigenvalues of A and the eigenvalues of $q(A)$.

Theorem 22.4 (Spectral mapping theorem). *Let q be a polynomial and A a square matrix. Then $\lambda \in \mathbb{C}$ is an eigenvalue of $q(A)$ if and only if $\lambda = q(\mu)$ for some eigenvalue μ of A .*

Proof. (\Leftarrow). This direction is a straightforward computation: if v is an eigenvector for A and μ , so that $Av = \mu v$, then we have

$$q(A)v = \left(\sum_{i=0}^d c_i A^i \right) v = \sum_{i=0}^d c_i (A^i v) = \sum_{i=0}^d c_i \mu^i v = q(\mu)v,$$

so $\lambda = q(\mu)$ is an eigenvalue of $q(A)$.

(\Rightarrow). Suppose λ is an eigenvalue of $q(A)$; thus $q(A) - \lambda I$ is not invertible. By the Fundamental Theorem of Algebra, we may factor the polynomial $q(t) - \lambda$ as

$$q(t) - \lambda = a \prod_{i=1}^d (t - \nu_i)$$

for some $\nu_1, \dots, \nu_d \in \mathbb{C}$, and observe that this implies

$$q(A) - \lambda I = a \prod_{i=1}^d (A - \nu_i I).$$

Since the left-hand side is not invertible, it follows that one of the factors $A - \nu_i I$ is non-invertible – in particular, some ν_i is an eigenvalue of A . Because ν_i is a root of $q(t) - \lambda$, we see that $q(\nu_i) - \lambda = 0$, so $\lambda = q(\nu_i)$. \square

Cayley–Hamilton and 2×2 matrices

Further reading:

23.1 Cayley–Hamilton theorem

It follows from Theorem 22.4 that if $p_A(t)$ is the characteristic polynomial of A , then all eigenvalues of $p_A(A)$ are zero – indeed, every eigenvalue μ of A has the property that $p_A(\mu) = 0$, and by Theorem 22.4 every eigenvalue of $p_A(A)$ is of the form $p_A(\mu)$ for some eigenvalue μ of A , so every eigenvalue of $p_A(A)$ is zero.

Recall that every nilpotent matrix has the property that all eigenvalues are zero. Thus the previous paragraph suggests that $p_A(A)$ is nilpotent.¹ In fact, even more is true.

Theorem 23.1 (Cayley–Hamilton theorem). *If $A \in \mathbb{M}_{n \times n}$ is a square matrix and $p_A(t)$ is the characteristic polynomial of A , then $p_A(A) = \mathbf{0}$.*

It is tempting to give the following “proof”:

$$p_A(A) = \det(AI - A) = \det \mathbf{0} = 0.$$

This is wrong. In the first place, notice that the output of determinant is a scalar, while $p_A(A)$ is a matrix – indeed, writing $p_A(t) = a_0 + a_1t + \cdots + a_{n-1}t^{n-1} + t^n$, we have

$$p_A(A) = a_0I + a_1A + a_2A^2 + \cdots + a_{n-1}A^{n-1} + A^n.$$

This is an $n \times n$ matrix, not a determinant of anything, and it is this expression that we need to show vanishes.

We first observe that if A happens to be diagonalisable (in particular, if A has n distinct eigenvalues), then writing any $v \in \mathbb{C}^n$ as a linear combination of eigenvectors v_1, \dots, v_n , we get

$$p_A(A)v = p_A(A) \sum_{j=1}^n c_j v_j = \sum_{j=1}^n c_j p_A(A)v_j = \sum_{j=1}^n c_j p_A(\lambda_j)v_j = \mathbf{0},$$

showing that $p_A(A) = \mathbf{0}$. However, if A is not diagonalisable then we need to use a different argument.

¹We have not yet proved that all eigenvalues zero implies nilpotence, but this can be shown.

Proof of Theorem 23.1. First we recall (20.3), which stated that for an invertible matrix A , we have $(A^{-1})_{ij} = (-1)^{i+j} \frac{\det \tilde{A}_{ji}}{\det A}$. By using the same computations as we used in the proof of this formula, we can show that the matrix B defined by $B_{ij} = \det \tilde{A}_{ji}$ has the property that

$$BA = AB = (\det A)I. \quad (23.1)$$

(B is sometimes called the *adjugate* of A .) We will apply this fact to the matrices $Q(t) = tI - A$, where $t \in K$. Let $P(t)$ be the matrix whose entries are given by

$$P(t)_{ij} = (-1)^{i+j} \det \tilde{Q}(t)_{ji};$$

then (23.1) shows that

$$P(t)Q(t) = (\det Q(t))I = p_A(t)I. \quad (23.2)$$

This is a statement about multiplication of polynomials with matrix coefficients: indeed, $P(t)_{ij}$ is a polynomial of degree $n - 1$ in t , so we can write

$$P(t)_{ij} = \sum_{k=0}^{n-1} C_{k,ij} t^k$$

for some coefficients $C_{k,ij} \in K$, and then letting $C_k \in \mathbb{M}_{n \times n}$ be the matrix with coefficients $C_{k,ij}$, we get

$$P(t) = \sum_{k=0}^{n-1} C_k t^k.$$

Thus (23.2) implies

$$P(t)Q(t) = \left(\sum_{k=0}^{n-1} C_k t^k \right) (tI - A) = p_A(t)I.$$

Expanding the product in the middle of the above equation gives

$$\begin{aligned} p_A(t)I &= (C_0 + C_1 t + \cdots + C_{n-1} t^{n-1})(tI - A) \\ &= -C_0 A + (C_0 - C_1 A)t + \cdots + (C_{n-2} - C_{n-1} A)t^{n-1} + C_{n-1} t^n. \end{aligned}$$

Let a_0, \dots, a_{n-1} be the coefficients of the characteristic polynomial of A , so that

$$p_A(t)I = a_0 I + a_1 I t + a_2 I t^2 + \cdots + a_{n-1} I t^{n-1} + I t^n.$$

(Recall that the top coefficient is always 1.) Then since the previous two equations hold for all $t \in K$, the coefficients of t^k must agree, giving

$$\begin{aligned} C_{n-1} &= I, \\ C_{n-2} - C_{n-1}A &= a_{n-1}I, \\ C_{n-3} - C_{n-2}A &= a_{n-2}I, \\ &\dots \\ C_0 - C_1A &= a_1I \\ -C_0A &= a_0I. \end{aligned}$$

Now multiply both sides of the first line by A^n (from the right), both sides of the second by A^{n-1} , and so on. We get

$$\begin{aligned} C_{n-1}A^n &= A^n, \\ C_{n-2}A^{n-1} - C_{n-1}A^n &= a_{n-1}A^{n-1}, \\ C_{n-3}A^{n-2} - C_{n-2}A^{n-1} &= a_{n-2}A^n, \\ &\dots \\ C_0A - C_1A^2 &= a_1A \\ -C_0A &= a_0I. \end{aligned}$$

Adding up all these equations, we see that the left side is $\mathbf{0}$ and the right side is $p_A(A)$. \square

23.2 Working over \mathbb{C}

Consider the 2×2 matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$, where $a, b, c, d \in \mathbb{C}$. The characteristic polynomial of A is

$$p_A(t) = \det(tI - A) = t^2 - (\operatorname{Tr} A)t + \det A,$$

which has roots given by

$$\begin{aligned} \lambda_{1,2} &= \frac{1}{2} \left(\operatorname{Tr} A \pm \sqrt{(\operatorname{Tr} A)^2 - 4 \det A} \right) \\ &= \frac{1}{2} \left(a + d \pm \sqrt{a^2 + 2ad + d^2 - 4(ad - bc)} \right) \\ &= \frac{1}{2} \left(a + d \pm \sqrt{(a - d)^2 + 4bc} \right). \end{aligned} \tag{23.3}$$

If $(a-d)^2 + 4bc \neq 0$, then we have $\lambda_1 \neq \lambda_2$, and thus the corresponding eigenvectors v_1, v_2 are linearly independent. Thus \mathbb{C}^2 has a basis of eigenvectors for A , and so A is diagonalisable.

On the other hand, we may have $(a-d)^2 + 4bc = 0$, in which case $p_A(t) = (t - \lambda)^2$, so A has only a single eigenvalue λ , with algebraic multiplicity 2. Now one of two things happens: either A is diagonalisable, or it is not.

If A is diagonalisable, then the geometric multiplicity of λ is equal to the algebraic multiplicity, by Corollary 22.2. That is, the eigenspace E_λ^A has dimension 2, which means that $E_\lambda^A = \mathbb{C}^2$, since the whole space has dimension 2. Thus every vector in \mathbb{C}^2 is an eigenvector: $Av = \lambda v$ for every $v \in \mathbb{C}^2$, which means that $A = \lambda I$. Thus the only way that a 2×2 matrix A with a repeated eigenvalue can be diagonalisable is if it is already a scalar multiple of the identity – this means that $b = c = 0$. There are plenty of 2×2 matrices with $(a-d)^2 + 4bc = 0$ (and hence a repeated eigenvalue) for which b and c are not both 0, so these matrices are not diagonalisable.

Suppose A is such a matrix – 2×2 , repeated eigenvalue, not diagonalisable. This means that $m_g(\lambda) = 1 < 2 = m_a(\lambda)$, so that λ has an eigenvector v , but there is no second linearly independent eigenvector. Thus $A \neq \lambda I$, so in particular $A - \lambda I \neq \mathbf{0}$. However, by the Cayley–Hamilton theorem (Theorem 23.1), A satisfies its own characteristic polynomial $p_A(t) = (t - \lambda)^2$, which means that $(A - \lambda I)^2 = \mathbf{0}$.

Now choose $w \in \mathbb{C}^2 \setminus E_\lambda^A$, and let $v = (A - \lambda I)w$. Because $(A - \lambda I)^2 = \mathbf{0}$, we see that

$$(A - \lambda I)v = (A - \lambda I)^2 w = \mathbf{0},$$

so $v \in E_\lambda^A$. In particular, this means that v, w are linearly independent, hence a basis. Moreover, we have $v = Aw - \lambda w$, so

$$Aw = \lambda w + v, \quad Av = \lambda v.$$

Thus although w is not actually an eigenvector for A , it has some behaviour that is reminiscent of eigenvectors: multiplying w by A has the effect of scaling it by the eigenvalue λ and then adding v .

We can use the above to find a matrix that is similar to A and has a reasonably nice structure. Let Q be the 2×2 matrix with columns v, w ; that is, $Q\mathbf{e}_1 = v$ and $Q\mathbf{e}_2 = w$. Let $B = Q^{-1}AQ$. Then

$$\begin{aligned} B\mathbf{e}_1 &= Q^{-1}AQ\mathbf{e}_1 = Q^{-1}Av = Q^{-1}\lambda v = \lambda\mathbf{e}_1, \\ B\mathbf{e}_2 &= Q^{-1}AQ\mathbf{e}_2 = Q^{-1}Aw = Q^{-1}(\lambda w + v) = \lambda\mathbf{e}_2 + \mathbf{e}_1, \end{aligned}$$

and we see that

$$B = Q^{-1}AQ = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}. \quad (23.4)$$

We have proved the following theorem.

Theorem 23.2. *Let A be a 2×2 matrix. Then there is an invertible 2×2 matrix Q with entries in \mathbb{C} such that $Q^{-1}AQ$ is either a diagonal matrix or has the form in (23.4).*

Note that (23.4) is the sum of the diagonal matrix λI and the nilpotent matrix $\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$. It can be shown that this holds for matrices of any size – any matrix can be written as the sum of a diagonalisable matrix and a nilpotent matrix.

Theorem 23.2 has an important application: it lets us compute powers of 2×2 matrices with relatively little fuss. Indeed, suppose $A \in \mathbb{M}_{2 \times 2}$ is such that $D = Q^{-1}AQ$ is diagonal for some invertible Q . Then we have $A = QDQ^{-1}$, and so

$$A^N = (QDQ^{-1})^N = \overbrace{(QDQ^{-1})(QDQ^{-1}) \cdots (QDQ^{-1})}^{N \text{ times}} = QD^N Q^{-1}$$

for every $N \in \mathbb{N}$. It is easy to compute powers of $D = \begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix}$, since we get $D^N = \begin{pmatrix} \lambda^N & 0 \\ 0 & \mu^N \end{pmatrix}$, and thus

$$A^N = Q \begin{pmatrix} \lambda^N & 0 \\ 0 & \mu^N \end{pmatrix} Q^{-1}.$$

Similarly, if $J = Q^{-1}AQ = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}$, then we get $A = QJ^N Q^{-1}$ by the same argument as above, and can use the fact that $J = \lambda I + B$, where $B = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ has $B^2 = \mathbf{0}$, to compute

$$\begin{aligned} J^2 &= (\lambda I + B)(\lambda I + B) = \lambda^2 I + 2\lambda B + B^2 = \lambda^2 I + 2\lambda B, \\ J^3 &= (\lambda^2 I + 2\lambda B)(\lambda I + B) = \lambda^3 I + 3\lambda^2 B, \end{aligned}$$

which suggests the general formula $J^N = \lambda^N I + N\lambda^{N-1}B$. And indeed, if this formula holds for N , we have

$$J^{N+1} = J^N J = (\lambda^N I + N\lambda^{N-1}B)(\lambda I + B) = \lambda^{N+1} I + (N+1)\lambda^N B,$$

using the fact that $B^2 = \mathbf{0}$. By induction, we have for every N that

$$A^N = QJ^N Q^{-1} = Q(\lambda^N I + N\lambda^{N-1}B)Q^{-1} = Q \begin{pmatrix} \lambda^N & N\lambda^{N-1} \\ 0 & \lambda^N \end{pmatrix} Q^{-1}.$$

Scaled rotations, etc.

Further reading:

24.1 Working over \mathbb{R}

In general, the eigenvalues of A may be complex, even if all the entries of A are real. This is because real polynomials always factor over the complex numbers, but not necessarily over the reals. Thus while Theorem 23.2 guarantees that every 2×2 matrix is similar over \mathbb{C} to a diagonal matrix or something in the form (23.4), this may not always be true if we want to only consider similarity over \mathbb{R} – that is, if we require the similarity matrix Q to have real entries.

For example, the matrix $A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ has characteristic polynomial $p_A(t) = t^2 + 1$ and eigenvalues $\pm i$, so its eigenvectors are also complex. Indeed, if $v \in \mathbb{R}^2$ then $Av \in \mathbb{R}^2$, which shows that we cannot have $Av = \lambda v$ for any $\lambda \in \mathbb{C} \setminus \mathbb{R}$.

More generally, given $\theta \in \mathbb{R}$, the rotation matrix $R_\theta = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$ has the effect of rotating a vector in \mathbb{R}^2 counterclockwise by the angle θ , and we immediately see that R_θ has no eigenvectors in \mathbb{R}^2 when θ is not a multiple of π . Similarly, given any $q > 0$, the matrix

$$qR_\theta = \begin{pmatrix} q \cos \theta & -q \sin \theta \\ q \sin \theta & q \cos \theta \end{pmatrix} \quad (24.1)$$

has the effect of rotating by θ and then scaling by q , so it has no real eigenvectors when θ is not a multiple of π .

On the other hand, qR_θ has characteristic polynomial

$$t^2 - 2q \cos \theta + q^2,$$

and thus its eigenvalues (in \mathbb{C}) are

$$q \cos \theta \pm \sqrt{q^2 \cos^2 \theta - q^2} = q(\cos \theta \pm i \sin \theta) = qe^{\pm i\theta}.$$

Observe that for $\lambda = q(\cos \theta + i \sin \theta)$ we have

$$qR_\theta - \lambda I = \begin{pmatrix} -qi \sin \theta & -q \sin \theta \\ q \sin \theta & -qi \sin \theta \end{pmatrix} = q \sin \theta \begin{pmatrix} -i & -1 \\ 1 & -i \end{pmatrix},$$

and the corresponding complex eigenvector is $\begin{pmatrix} i \\ 1 \end{pmatrix} = i\mathbf{e}_1 + \mathbf{e}_2$. It turns out that any real 2×2 matrix with complex eigenvalues is similar to qR_θ , and the key is to write its complex eigenvector as $iv_1 + v_2$, where $v_1, v_2 \in \mathbb{R}^2$. Let us explain this a little more carefully; first we need the following elementary result.

Proposition 24.1. *Let $A \in \mathbb{M}_{n \times n}$ have real entries, and suppose that $v \in \mathbb{C}^n$ is an eigenvector for A with eigenvalue $\lambda \in \mathbb{C}$. Then $\bar{v} \in \mathbb{C}^n$ is an eigenvector for A with eigenvalue $\bar{\lambda} \in \mathbb{C}$, where \bar{v} and $\bar{\lambda}$ denote complex conjugation.*

Proof. By properties of complex conjugation and the fact that A is real, so $\bar{A} = A$, we have

$$A\bar{v} = \bar{A}v = \overline{Av} = \overline{\lambda v} = \bar{\lambda}\bar{v}. \quad \square$$

Let A be a 2×2 real matrix with a complex eigenvalue $a + ib$, where $b \neq 0$. Let $w \in \mathbb{C}^2$ be a corresponding eigenvector, and write $w = iv_1 + v_2$, where $v_1, v_2 \in \mathbb{R}^2$. (This may seem like an unnatural order to write the real and imaginary parts in, but turns out to work better.) Then by Proposition 24.1, $a - ib$ is also an eigenvalue of A , with eigenvector $\bar{w} = -iv_1 + v_2$.

By the assumption that $b \neq 0$, we see that w and \bar{w} are linearly independent, hence form a basis for \mathbb{C}^2 . This means that the 2×2 complex matrix $P = [w \mid \bar{w}]$ is invertible, and diagonalises A to the 2×2 complex matrix

$$P^{-1}AP = \begin{pmatrix} a + ib & 0 \\ 0 & a - ib \end{pmatrix}.$$

Now P is related to the 2×2 real matrix $S = [v_1 \mid v_2]$ by

$$P = [w \mid \bar{w}] = [v_1 \mid v_2] \begin{pmatrix} i & -i \\ 1 & 1 \end{pmatrix} = ST,$$

where $T = \begin{pmatrix} i & -i \\ 1 & 1 \end{pmatrix}$ is invertible. Hence S is invertible as well, so v_1, v_2 form a basis for \mathbb{R}^2 . Note that $v_1 = \frac{1}{2i}(w - \bar{w})$ and $v_2 = \frac{1}{2}(w + \bar{w})$.

We claim that $B = S^{-1}AS$ has the form (24.1). Indeed, observing that $S = PT^{-1}$ and that $T^{-1} = \frac{1}{2} \begin{pmatrix} -i & 1 \\ i & 1 \end{pmatrix}$, we have

$$\begin{aligned} B &= S^{-1}AS = TP^{-1}APT^{-1} = T \begin{pmatrix} a + ib & 0 \\ 0 & a - ib \end{pmatrix} T^{-1} \\ &= \frac{1}{2} \begin{pmatrix} i & -i \\ 1 & 1 \end{pmatrix} \begin{pmatrix} a + ib & 0 \\ 0 & a - ib \end{pmatrix} \begin{pmatrix} -i & 1 \\ i & 1 \end{pmatrix} \\ &= \frac{1}{2} \begin{pmatrix} ai - b & -ai - b \\ a + ib & a - ib \end{pmatrix} \begin{pmatrix} -i & 1 \\ i & 1 \end{pmatrix} \end{aligned}$$

$$= \begin{pmatrix} a & -b \\ b & a \end{pmatrix}$$

and writing $a + ib = q(\cos \theta + i \sin \theta) = qe^{i\theta}$ for some $q > 0$ and $\theta \in \mathbb{R}$, we have

$$B = \begin{pmatrix} a & -b \\ b & a \end{pmatrix} = \begin{pmatrix} q \cos \theta & -q \sin \theta \\ q \sin \theta & q \cos \theta \end{pmatrix} = qR_\theta.$$

We have proved the following theorem.

Theorem 24.2. *Given any real 2×2 matrix A , there is an invertible 2×2 real matrix Q such that QAQ^{-1} has one of the following forms:*

$$\begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix} \quad \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix} \quad q \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}, \quad (24.2)$$

where $\lambda, \mu, q, \theta \in \mathbb{R}$.

In fact, one can determine which of the three cases in (24.2) occurs with a quick look at $p_A(t)$ and A itself: the first case occurs if and only if $p_A(t)$ has distinct real roots or is a scalar multiple of the identity matrix; the second case occurs if and only if $p_A(t)$ has a repeated real root and is not a scalar multiple of the identity; the third case occurs if and only if $p_A(t)$ has complex roots.

The situation for larger matrices is a more complicated, but in some sense mimics the behaviour seen here. We observe that Theorem 24.2 can also be formulated for linear transformations: if V is a vector space over \mathbb{R} with dimension 2 and $T \in \mathbb{L}(V)$ is any linear operator, then there is a basis β for V such that $[T]_\beta$ has one of the forms in (24.2).

We saw in the previous section how to compute powers of A when it is similar to one of the first two cases in (24.2). In fact, if we want to compute A^N for a real matrix then even when A has complex eigenvalues we can still use the diagonal form $\begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix}$, because the complex parts of Q , D^N , and Q^{-1} will ultimately all cancel out in the product $A^N = QD^NQ^{-1}$. Nevertheless, it is worth pointing out that powers of the scaled rotation matrix also have a nice formula.

Proposition 24.3. *Writing $R_\theta = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$, we have*

$$(qR_\theta)^N = q^N R_{N\theta} \text{ for all } N \in \mathbb{Z}. \quad (24.3)$$

Proof. $N = 0, 1$ is immediate. To prove the rest observe that

$$\begin{aligned} R_\alpha R_\beta &= \begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} \cos \beta & -\sin \beta \\ \sin \beta & \cos \beta \end{pmatrix} \\ &= \begin{pmatrix} \cos \alpha \cos \beta - \sin \alpha \sin \beta & -\cos \alpha \sin \beta - \sin \alpha \cos \beta \\ \sin \alpha \cos \beta + \cos \alpha \sin \beta & \cos \alpha \cos \beta - \sin \alpha \sin \beta \end{pmatrix} \quad (24.4) \\ &= \begin{pmatrix} \cos(\alpha + \beta) & -\sin(\alpha + \beta) \\ \sin(\alpha + \beta) & \cos(\alpha + \beta) \end{pmatrix} = R_{\alpha+\beta} \end{aligned}$$

This gives $N > 1$ by induction, since $R_{(N+1)\theta} = R_{N\theta}R_\theta = R_\theta^N R_\theta = R_\theta^{N+1}$. For $N < 0$ it suffices to observe that $R_{-N\theta}R_{N\theta} = R_0 = I$, and so $R_{-N\theta} = R_{N\theta}^{-1} = R_\theta^{-N}$. \square

An alternate proof of (24.4) that avoids trigonometric identities can be given by diagonalising R_θ .

Exercise 24.4. Show that R_θ has eigenvalues $e^{\pm i\theta} = \cos \theta \pm i \sin \theta$ with eigenvectors $\begin{pmatrix} \pm i \\ 1 \end{pmatrix}$. Conclude that

$$R_\theta = Q \begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{pmatrix} Q^{-1}, \quad Q = \begin{pmatrix} i & -i \\ 1 & 1 \end{pmatrix},$$

and use this to show (24.4).

24.2 Application to ODEs

Consider the second order linear ODE $\ddot{x} + a\dot{x} + bx = 0$. Recall that this can be turned into a first order system by putting $y = \dot{x}$ and $\mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}$, so that

$$\dot{\mathbf{x}} = \begin{pmatrix} y \\ \ddot{x} \end{pmatrix} = \begin{pmatrix} y \\ -ay - bx \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ -b & -a \end{pmatrix} \mathbf{x} = A\mathbf{x}.$$

This is a system of two first order ODEs, which are coupled. If we decouple them by diagonalising A , then we can solve each independently. The characteristic polynomial of A is

$$p_A(s) = s^2 - \text{Tr}(A)s + \det(A) = s^2 + as + b. \quad (24.5)$$

If A diagonalises as $A = QDQ^{-1}$, then putting $\mathbf{z} = Q^{-1}\mathbf{x}$ gives

$$\dot{\mathbf{z}} = Q^{-1}\dot{\mathbf{x}} = Q^{-1}A\mathbf{x} = DQ^{-1}\mathbf{x} = D\mathbf{z},$$

and so writing $D = \begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix}$, this becomes

$$\dot{z}_1 = \lambda z_1, \quad \dot{z}_2 = \mu z_2.$$

These can be solved independently. (Note that if the roots are complex, then we will need to take appropriate linear combinations of the complex solutions $z_{1,2}$ to obtain real solutions.)

But what if A does not diagonalise? If (24.5) has a repeated root, then because A is not a scalar multiple of the identity we know that we will get $A = QJQ^{-1}$, where $J = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}$, and so the change of coordinates $\mathbf{z} = Q^{-1}\mathbf{x}$ gives $\dot{\mathbf{z}} = J\mathbf{z}$, that is,

$$\begin{pmatrix} \dot{z}_1 \\ \dot{z}_2 \end{pmatrix} = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix},$$

so that we must solve the system

$$\begin{aligned} \dot{z}_1 &= \lambda z_1 + z_2, \\ \dot{z}_2 &= \lambda z_2. \end{aligned}$$

This is not fully decoupled, but because J is upper triangular, it is decoupled enough to solve. Indeed, we get $z_2(t) = e^{\lambda t} z_2(0)$, and then the first equation reduces to

$$\dot{z}_1 = \lambda z_1 + g(t),$$

where $g(t) = z_2(0)e^{\lambda t}$ is a forcing term. This can be solved via standard ODE methods: in the absence of the forcing term, the solution would be $z_1(t) = e^{\lambda t} z_1(0)$, so the quantity $z_1 e^{-\lambda t}$ would be constant. Differentiating this quantity and using the actual ODE gives

$$\frac{d}{dt}(z_1 e^{-\lambda t}) = \dot{z}_1 e^{-\lambda t} - \lambda z_1 e^{-\lambda t} = \lambda z_1 e^{-\lambda t} + g(t) e^{-\lambda t} - \lambda z_1 e^{-\lambda t} = g(t) e^{-\lambda t}.$$

Thus we have

$$z_1(t) e^{-\lambda t} = z_1(0) + \int_0^t g(s) e^{-\lambda s} ds,$$

and so the solution is

$$z_1(t) = z_1(0) e^{\lambda t} + e^{\lambda t} \int_0^t g(s) e^{\lambda(-s)} ds.$$

In the particular case $g(t) = z_2(0)e^{\lambda t}$, we see that $g(s)e^{-\lambda s} = z_2(0)$ is constant, and so

$$z_1(t) = z_1(0) e^{\lambda t} + z_2(0) t e^{\lambda t}.$$

Thus the source of the solution $te^{\lambda t}$ when there is a repeated root is the fact that a eigenvalue with $m_g < m_a$ leads to a Jordan block. Another way of viewing this is the following. Let V be the vector space of smooth (infinitely many times differentiable) functions from \mathbb{R} to \mathbb{R} , and let $D \in \mathbb{L}(V)$ be the differentiation operator. Then the ODE $\ddot{x} + a\dot{x} + bx = 0$ becomes $(D^2 + aD + bI)(x) = 0$. When there is a repeated root we have

$$s^2 + as + b = (s - \lambda)^2 \Rightarrow D^2 + aD + bI = (D - \lambda I)^2,$$

and so solving the ODE becomes a question of finding the null space of $(D - \lambda I)^2$. The null space of $D - \lambda I$ is eigenfunctions of the differentiation operator, which is all scalar multiples of $e^{\lambda t}$, but then we must also consider generalised eigenfunctions, which are in the null space of $(D - \lambda I)^2$ but not $(D - \lambda I)$. We see that

$$\frac{d}{dt}(te^{\lambda t}) = \lambda(te^{\lambda t}) + e^{\lambda t}.$$

Writing $x(t) = e^{\lambda t}$ and $y(t) = te^{\lambda t}$, this can be rewritten as

$$Dy = \lambda y + x,$$


which we saw earlier as the equation characterising a generalised eigenvector of order 2, since x is a genuine eigenvector.

Markov chains

Further reading:


We can introduce Markov chains via the following simplified example. Suppose that people move back and forth between Dallas and Houston according to the following rule: every year, 10% of the people in Dallas move to Houston (and the rest remain in Dallas), while 5% of the people in Houston move to Dallas. In the long run, how do the populations of the two cities relate? To put it another way, if we pick a person at random, what is the probability that they live in Houston?

The situation can be described by the following graph, in which there are two “states” each person can be in: “living in Houston”, or “living in Dallas”.



The arrows in the graph illustrate the four “transitions” that can occur: “staying in Houston”; “moving from Houston to Dallas”; “moving from Dallas to Houston”; “staying in Dallas”. The number with each arrow gives the probability of that transition occurring.

The circles are called “vertices” of the graph, and the arrows are called “edges”. If we label the vertices 1 and 2, then the graph drawn above has the following structure:



Here P_{ij} is the probability that a person in city i this year will be in city j next year. We can represent all the information in the graph by writing the 2×2 matrix P whose entries are P_{ij} :

$$P = \begin{pmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{pmatrix} = \begin{pmatrix} .95 & .05 \\ .1 & .9 \end{pmatrix}.$$

Suppose $v \in \mathbb{R}^2$ is a vector that encodes the populations of the two cities this year – that is, v_1 is the population of Houston, and v_2 is the population of Dallas. Let w be the vector encoding *next* year’s populations. Then we see that

$$\begin{aligned}w_1 &= .95v_1 + .1v_2 \\w_2 &= .05v_1 + .9v_2,\end{aligned}$$

which can be written in terms of matrix multiplication as

$$w = \begin{pmatrix} w_1 & w_2 \end{pmatrix} = \begin{pmatrix} .95v_1 + .1v_2 \\ .05v_1 + .9v_2 \end{pmatrix} = \begin{pmatrix} v_1 & v_2 \end{pmatrix} \begin{pmatrix} .95 & .05 \\ .1 & .9 \end{pmatrix} = vP.$$

Thus next year’s population distribution can be obtained from this year’s by right multiplication by P . Iterating this procedure, the distribution two years from now will be $wP = vP^2$, three years from now will be vP^3 , and more generally, the population distribution after n years will be vP^n .

A system like the one just described is called a *Markov chain*. More precisely, a Markov chain is a collection of states together with a collection of transition probabilities. The states are usually encoded with the labels $1, 2, \dots, d$, and then the transition probabilities P_{ij} are viewed as the entries of a $d \times d$ matrix P , where P_{ij} is the probability of making a transition from state i to state j . A Markov chain can also be encoded by a directed graph such as the one above – that is, a collection of vertices and edges, where the vertices correspond to states of the system, and each edge is an arrow from one vertex to another labelled with the probability of the corresponding transition.

Note that every entry P_{ij} lies in the interval $[0, 1]$. In fact, something more can be said about the transition probabilities P_{ij} . If we fix i , then $P_{i1} + P_{i2}$ represents the probability of going from state i to either of states 1 or 2; similarly, $P_{i1} + P_{i2} + P_{i3}$ is the probability of going from state i to any of the states 1, 2, 3, and so on. In particular, $\sum_{j=1}^d P_{ij} = 1$, since we go *somewhere* with probability 1.

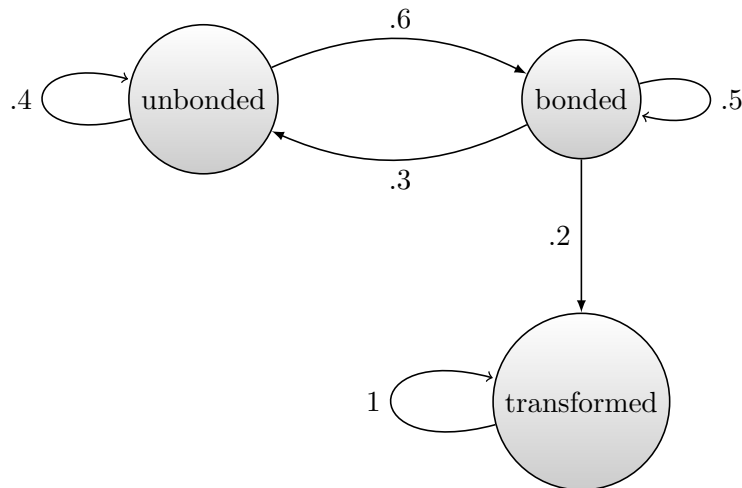
Definition 25.1. A vector $v \in \mathbb{R}^d$ is a *probability vector* if $v_j \geq 0$ for every j and $\sum_{j=1}^d v_j = 1$. A matrix $P \in \mathbb{M}_{d \times d}$ is a *stochastic matrix* if every row vector of P is a probability vector.

The previous discussion shows that the matrix of transition probabilities is a stochastic matrix. This matrix encodes all the important information about a Markov chain.

We remark that sometimes the definition of a stochastic matrix is given in terms of column vectors instead of row vectors. This would be the way to do things if P_{ij} represented the probability of going *to* state i *from* state j , instead of the other way around.

We often think of a Markov chain as describing a system about which we have incomplete information, so that at time t the best we can hope to do is say what probability the system has of being in a given state. The following example illustrates this.

Example 25.2. A simple model of enzyme activity is given by Michaelis–Menten kinetics. Consider the following version of this model: a certain molecule (call it M) undergoes a particular reaction only when it has bonded to an enzyme. At each time step, the molecule can be in one of three states: bonded, unbonded, and transformed. It can pass between bonded and unbonded states (catching and releasing enzyme molecules without undergoing the reaction); it cannot pass between unbonded and transformed (the reaction only occurs when the enzyme is present); it can pass from bonded to transformed but not vice versa (the reaction only runs one way). This can be modeled as a Markov chain using the following graph, which also includes some transition probabilities for the sake of argument.



The graph is interpreted as follows.

- If the untransformed molecule is not bonded to an enzyme, then it has a 60% chance of bonding at the next time step, and a 40% chance of remaining unbonded.

- If the untransformed molecule is currently bonded to an enzyme, then it has a 30% chance of unbonded, a 20% chance of undergoing the transformative reaction, and a 50% chance of staying as it is (bonded but untransformed).
- One the molecule has been transformed, it remains transformed.

Notice that the arrows leading away from any given vertex have associated probabilities that sum up to 1. If we assign the unbonded state the label 1, the bonded state the label 2, and the transformed state the label 3, then the transition matrix for this Markov chain is

$$P = \begin{pmatrix} .4 & .6 & 0 \\ .3 & .5 & .2 \\ 0 & 0 & 1 \end{pmatrix}. \quad (25.1)$$

Returning now to our general discussion of Markov chains, we say that a probability vector v represents the current distribution of the system if the system is currently in state 1 with probability v_1 , in state 2 with probability v_2 , and so on. In particular, if the system is currently in state i with probability 1, then it is represented by the probability vector \mathbf{e}_i , the i th standard basis vector. In this case the distribution of the system at the next time step is given by the i th row of P , which is $\mathbf{e}_i P$, since

$$P_{ij} = \mathbb{P}(\text{state } j \text{ tomorrow given state } i \text{ today}),$$

and so at the next time step the system is in state 1 with probability P_{i1} , state 2 with probability P_{i2} , and so on. More generally, if the current distribution of the system is represented by a probability vector v , and the distribution at the next time step is represented by a probability vector w , then we have

$$\begin{aligned} w_j &= \mathbb{P}(\text{state } j \text{ tomorrow}) \\ &= \sum_{i=1}^d \mathbb{P}(\text{state } j \text{ tomorrow given state } i \text{ today}) \cdot \mathbb{P}(\text{state } i \text{ today}) \\ &= \sum_{i=1}^d v_i P_{ij} = (vP)_j. \end{aligned}$$

That is, v and w are related by $w = vP$. Thus, just as before, if v represents the probability distribution of the system at the present time, then vP^n gives the probability distribution after n time steps have passed.

Proposition 25.3. *If v is a probability vector and P is a stochastic matrix, then vP is a probability vector.*

Proof. Since $v_i \geq 0$ and $P_{ij} \geq 0$ for all i, j , we immediately get $(vP)_j \geq 0$ for all j . Moreover,

$$\sum_{j=1}^d (vP)_j = \sum_{j=1}^d \sum_{i=1}^d v_i P_{ij} = \sum_{i=1}^d v_i \left(\sum_{j=1}^d P_{ij} \right) = \sum_{i=1}^d v_i = 1,$$

where the first equality is matrix multiplication, the second is just rearranging terms, the third is because P is a stochastic matrix, and the fourth is because v is a probability vector. \square

Example 25.4. Consider the Michaelis–Menten kinetics described above, with the transition matrix P given in (25.1). The probability that an initially unbonded molecule has bonded and transformed by time n is $(\mathbf{e}_1 P^n)_3 = (P^n)_{13}$. That is, this probability is the third coordinate of the first row of the matrix

$$P^n = \begin{pmatrix} .4 & .6 & 0 \\ .3 & .5 & .2 \\ 0 & 0 & 1 \end{pmatrix}^n.$$

Note that P has $\lambda_1 = 1$ as an eigenvalue, with corresponding left eigenvector $v_1 = (0 \ 0 \ 1)$. Numerical computation shows that it has two other eigenvalues $\lambda_2 \approx 0.87$ and $\lambda_3 \approx 0.023$; letting v_2, v_3 be left eigenvectors for these, we see that v_1, v_2, v_3 is a basis for \mathbb{R}^3 , and so in particular there are coefficients a_1, a_2, a_3 such that

$$\begin{aligned} \mathbf{e}_1 &= a_1 v_1 + a_2 v_2 + a_3 v_3 \\ \mathbf{e}_1 P^n &= a_1 v_1 P^n + a_2 v_2 P^n + a_3 v_3 P^n \\ &= a_1 \lambda_1^n v_1 + a_2 \lambda_2^n v_2 + a_3 \lambda_3^n v_3 \\ &= a_1 \mathbf{e}_3 + (a_2 \lambda_2^n v_2 + a_3 \lambda_3^n v_3). \end{aligned}$$

As $n \rightarrow \infty$, the quantity inside the brackets goes to 0 because $|\lambda_2|, |\lambda_3| < 1$, and so $\mathbf{e}_1 P^n \rightarrow a_1 \mathbf{e}_3$. Because $\mathbf{e}_1 P^n$ is a probability vector for all n , we see that $a_1 = 1$, and so $\mathbf{e}_1 P^n \rightarrow \mathbf{e}_3$. In terms of the original example, this is the statement that at large times, the molecule has almost certainly been transformed.

The phenomenon that occurred in the previous example is in fact quite a general one: the stochastic matrix representing a Markov chain has 1 as

an eigenvalue, and all other eigenvalues λ have $|\lambda| < 1$. Thus writing the initial probability distribution as $v = w + x$, where w is a left eigenvector for 1 and x is a sum of generalised eigenvectors for other eigenvalues, one gets

$$vP^n = wP^n + xP^n = w + xP^n \rightarrow w;$$

the term xP^n goes to $\mathbf{0}$ because all other eigenvalues have absolute value less than 1. Consequently, it is the eigenvector corresponding to the largest eigenvalue 1 that governs the long-term behaviour of the system. Moreover, this eigenvector w has the property that $w = wP$, meaning that w is a *stationary probability distribution* for the Markov chain: if w describes the probability distribution of the system at the present time, then it will also describe the probability distribution of the system at all future times.

Perron–Frobenius and Google

Further reading:

To turn the informal discussion from the end of the last lecture into a precise result, we need a little more terminology. Consider a Markov chain with states $\{1, \dots, d\}$ and transition matrix $P \in \mathbb{M}_{d \times d}$. Let G be the corresponding graph, where we only draw edges that correspond to a transition with a non-zero probability of occurring. (Thus in the example of the Michaelis–Menten kinetics, we do not draw the edge from “unbonded” to “transformed”, or the edges from “transformed” to “bonded” or “unbonded”.)

Definition 26.1. A *path* in a directed graph is a sequence of edges such that each edge starts at the vertex where the previous one terminated. A Markov chain is *irreducible* if given any two states i and j , there is a path that starts at vertex i and terminates at vertex j .

Example 26.2. The Markov chain for the Dallas–Houston example is irreducible, since every transition happens with a non-zero probability, and so the graph contains all possible edges. The Markov chain for the Michaelis–Menten kinetics is not irreducible, since there is no path from “transformed” to either of the other vertices.

Proposition 26.3. A Markov chain is irreducible if and only if its transition matrix P has the property that for every $1 \leq i, j \leq d$, there is some $n \in \mathbb{N}$ such that $(P^n)_{ij} > 0$.

Proof. (\Rightarrow). Given irreducibility, for every i, j there is a path in the graph that connects i to j . Let the vertices of this path be i_0, i_1, \dots, i_n , where $i_0 = i$ and $i_n = j$. Then because each $i_k \rightarrow i_{k+1}$ is required to be an edge in the graph, we must have $P_{i_k i_{k+1}} > 0$, and so

$$\begin{aligned} (P^n)_{ij} &= \sum_{j_1=1}^d \sum_{j_2=2}^d \cdots \sum_{j_{n-1}=1}^d P_{ij_1} P_{j_1 j_2} \cdots P_{j_{n-1} j} \\ &\geq P_{i_0 i_1} P_{i_1 i_2} \cdots P_{i_{n-1} i_n} > 0. \end{aligned} \tag{26.1}$$

(\Leftarrow). Given any i, j and n such that $(P^n)_{ij} > 0$, the first line of (26.1) shows that there exist j_1, \dots, j_{n-1} such that $P_{ij_1} > 0$, $P_{j_1 j_2} > 0$, \dots , $P_{j_{n-1} j} > 0$. Thus the graph contains an edge from i to j_1 , from j_1 to j_2 , and so on through to j_{n-1} to j . This gives a path from i to j , and since this holds for all i, j , we conclude that the Markov chain is irreducible. \square

Exercise 26.4. In the definition of irreducibility, the value of n for which $(P^n)_{ij}$ is allowed to depend on i, j , and the same value of n may not work for all i, j . Give an example of a Markov chain for which there is no single value of n that works for all i, j .

Exercise 26.5. Say that a Markov chain is *primitive* if there is a single value of n such that $(P^n)_{ij} > 0$ for all i, j . Describe a condition on the graph of the Markov chain that is equivalent to this property.

The primary result concerning Markov chains (which we will not have time to prove completely) is the following theorem.

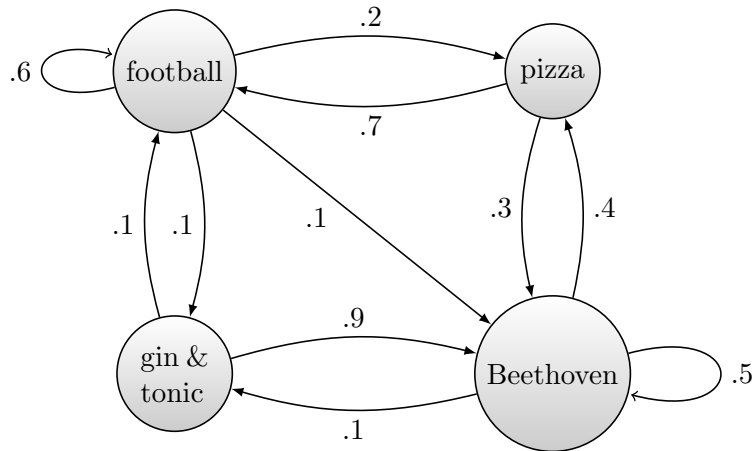
Theorem 26.6 (Perron–Frobenius Theorem for stochastic matrices). *Let P be a primitive stochastic matrix. Then the following statements are true.*

1. 1 is an eigenvalue of P , with algebraic and geometric multiplicity 1.
2. All other eigenvalues of P have $|\lambda| < 1$.
3. The eigenvalue 1 has a left eigenvector v which is a probability vector and whose components are all positive.
4. Let Q be the $d \times d$ matrix whose row vectors are all equal to v . Then $\lim_{n \rightarrow \infty} P^n = Q$. Equivalently, $P^n w \rightarrow v$ for every probability vector w .

Exercise 26.7. Show that for any stochastic matrix, the column vector whose entries are all equal to 1 is a right eigenvector for the eigenvalue 1.

There is also a version of this theorem that works for irreducible matrices (it is slightly more complicated in this case because P may have other eigenvalues with $|\lambda| = 1$). The Perron–Frobenius theorem tells us that for a primitive (or more generally, irreducible) Markov chain, the long-term behaviour is completely governed by the eigenvector corresponding to the eigenvalue 1.

Example 26.8. Consider a mildly inebriated college student who is watching the Super Bowl on TV, but doesn't really care for football, and so at any given point in time might decide to leave the game and go get pizza, or go make another gin and tonic (his drink of choice, which he thinks goes really well with mushroom and anchovy pizza), or step over for a little while to the outdoor theatre next to his apartment, where the symphony is performing Beethoven's seventh. Suppose that every ten minutes, he either moves to a new location or stays where he is, with transition probabilities as shown in the following graph.



If we order the states as $\{1, 2, 3, 4\} = \{\text{football}, \text{pizza}, \text{gin\&tonic}, \text{Beethoven}\}$, the corresponding stochastic transition matrix is

$$P = \begin{pmatrix} .6 & .2 & .1 & .1 \\ .7 & 0 & 0 & .3 \\ .1 & 0 & 0 & .9 \\ 0 & .4 & .1 & .5 \end{pmatrix}$$

Notice that some transition probabilities are 0: he never stays at the pizza place for more than ten minutes, and it doesn't take longer than ten minutes to make his next gin & tonic; similarly, after either of those he always goes to either football or Beethoven, and he never goes from Beethoven directly to football without first getting either pizza or a drink. Nevertheless, one can check that P^2 has all positive entries, and so P is primitive.

Numerical computations show that the largest eigenvalue (in absolute value) is 1, and that there are three other (distinct) eigenvalues, each with $|\lambda| < \frac{1}{2}$. The left eigenvector corresponding to the largest eigenvalue is $v \approx (.39 \ .21 \ .07 \ .33)$, so that in the long run the student will spend 39% of his time watching football, 21% of his time getting pizza, 7% of his time making himself a gin & tonic, and 33% of his time at the symphony.

It is also instructive to compute some powers of P and compare them to the matrix Q whose row vectors are all equal to v . We find that

$$P^3 \approx \begin{pmatrix} .42 & .20 & .08 & .3 \\ .45 & .17 & .06 & .32 \\ .31 & .20 & .05 & .44 \\ .32 & .24 & .08 & .36 \end{pmatrix}, \quad P^6 \approx \begin{pmatrix} .39 & .21 & .07 & .33 \\ .39 & .21 & .07 & .33 \\ .38 & .21 & .07 & .34 \\ .38 & .21 & .07 & .34 \end{pmatrix}$$

These can be interpreted as follows: the row vectors of P^3 represent possible probability distributions after 3 transitions (half an hour), while the rows of P^6 represent the corresponding distributions after 6 transitions (one hour). The first row represents the distribution if he started out watching football, the second if he started at the pizza place, the third if he started making a drink, and the fourth if he started at the symphony. We see that after half an hour, there is still some dependence on where he started out, while after a full hour, the dependence is down to under one percentage point, and all the rows of P^6 are nearly identical to the eigenvector v .

One important application of all this occurs if we want to “rank” the different states somehow. In the previous example, we might want to know which of the four activities the student spends most of his time doing. Upon looking at the graph, or the matrix P , we might guess that Beethoven will be the most common activity: after all, if we sum the probabilities in the j th column of P , we are adding up all the probabilities of making a transition to state j , and this number is higher in the fourth column (Beethoven) than in the first (football).

Exercise 26.9. Explain why football ends up being more likely than Beethoven, despite the above fact.

A rather more serious example of this process lies behind the PageRank algorithm, which was an important part of Google’s early success. The problem confronting Google, or any search engine, is that it is relatively easy for a computer to produce a list of websites that are relevant to a given search phrase, but then one is faced with the problem of ranking these websites. For example, although the poets may have been mysteriously silent on the subject of cheese (if Chesterton is to be believed), the internet is certainly not, and Google currently turns up about 450 million websites related to “cheese” in one way or another. Not all of these are created equal, though – out of the resulting list of shops, restaurants, news articles, recipes, encyclopedia entries, and advertisements for Wisconsin, how does one choose the ten most important websites to appear on the first page of the search results? The word “important” here can be replaced with “interesting” or “valuable” or whatever we want to use to describe the qualities of a website we are concerned with ranking.

At this point one observes that the internet has the structure of a directed graph, such as the ones we have been studying here: each website is represented by a vertex, and there is an edge from vertex A to vertex B if website A has a link to website B . One expects that the most important websites are those that get linked to a lot; thus perhaps one should count the

number of incoming links, or in terms of the graph, the number of incoming edges.

Upon further thought, though, this needs some refinement. Suppose you start a cheese-making operation and want to have a web presence, so you put up a website touting your gouda (or curd of choice). Now if I post a link to your website from my UH faculty webpage, you have one more incoming link, which marginally increases your website's importance. On the other hand, if CNN does a story about you and posts a link to your website from their homepage, that should increase your website's visibility and importance rather more than the link from my website did.

This leads to the basic principle followed by the PageRank algorithm: *A site is important if important sites link to it.* To put it another way, the importance of a site is proportional to the importance of the sites that link to it.

Let's make this notion precise. Suppose we have a list of N websites, and we want to rank them. Label the sites $1, 2, \dots, N$, and let L_{ij} be the number of times that website i links to website j . Let M_i be the total number of outgoing links from website i , and let $P_{ij} = L_{ij}/M_i$. Then we see that P is a stochastic $N \times N$ matrix, because


$$\sum_{j=1}^N P_{ij} = \sum_{j=1}^N \frac{L_{ij}}{M_i} = \frac{M_i}{M_i} = 1.$$

Now let $v \in \mathbb{R}^n$ be a probability vector that encodes our current best guess as to the relative "importance" of the websites: that is, website i has a number v_i between 0 and 1 attached to it, and $\sum_i v_i = 1$. According to the above principle (importance being proportional to importance of linking sites), the "correct" v should have the property that

$$v_j = \sum_{i=1}^N v_i P_{ij},$$

where we think of $v_i P_{ij}$ as the amount of the site i 's importance that it passes to the site j . This equation is just the statement that $v = vP$, that is, that v is an eigenvector for P with eigenvalue 1. By the Perron–Frobenius theorem, there is a unique such eigenvector (as long as P is primitive), and we can use this to rank the websites: the highest-ranked website is the one for which v_i is largest, and so on.

As an illustration, consider the following very small internet, which only has five websites.



The matrix L described above, where L_{ij} is the number of links from i to j , is

$$L = \begin{pmatrix} 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

To get the stochastic matrix P , we have to normalise each row of L to be a probability vector, obtaining

$$P = \begin{pmatrix} 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ \frac{1}{3} & 0 & \frac{1}{3} & 0 & \frac{1}{3} \\ 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Numerical computation shows that the eigenvector v satisfying $v = vP$ is

$$v \approx (.29, .26, .09, .15, .21).$$

Thus A is the most important website, then B , then E , then D , then C . Note that this happens even though A only has two incoming edges, while B and E have three each; the reason is that the incoming edges for A come from more important sources (B and E).

The above process can be interpreted in terms of a “random walk”. Suppose you put a monkey on the internet and have him navigate at random – whatever website he is currently on, he follows a random link from that page to get to the next website. Then his location is given by a Markov chain with transition matrix P ; if w describes the probability distribution

of where the monkey currently is, then wP describes the probability distribution of where he will be after he follows the next link. After n clicks, his likely location is described by wP^n , which by the Perron–Frobenius theorem converges to the eigenvector v .

In practice, some tweaks to the process described here are needed to deal with the real internet - for example, if there is a website that doesn't link anywhere, then P is not primitive and the monkey would simply get trapped at that site, even if it's not really that important a website. Thus one has to add a small probability of “resetting” at any given time. And there are other issues as well – but this approach via Markov chains is a very powerful one and illustrates the basic idea.