

LINEAR ALGEBRA

ANTON BERNSHTEYN

CONTENTS

1	Fields	3
1.A	A problem to think about	3
1.B	Groups, rings, fields	3
1.C	Examples of fields	5
1.D	Finite fields	6
1.E	Matrices over rings	8
1.F	An application of linear algebra over an unusual field	9
	Extra exercises for Section 1.	10
2	Vector spaces.	11
2.A	The definition of a vector space	11
2.B	Subspaces	13
2.C	Linear functions	15
2.D	Quotient spaces	16
	Extra exercises for Section 2.	18
3	Bases	18
3.A	A “better” version of Theorem 2.31	18
3.B	Spanning sets and independent sets	19
3.C	The first fundamental theorem of linear algebra and its ramifications	21
3.D	Proof of Theorem 3.15.	23
3.E	Proof of Zorn’s lemma.	25
3.F	Applying Theorem 3.15	28
	Extra exercises for Section 3.	30
4	Dimension	31
4.A	The second fundamental theorem of linear algebra.	31
4.B	Dimension and finite-dimensional spaces.	33
4.C	Using dimension: algebraic numbers	34
	Extra exercises for Section 4.	37

DEPARTMENT OF MATHEMATICAL SCIENCES, CARNEGIE MELLON UNIVERSITY, PA, USA

Date: September 2, 2019.

The author is grateful to Zichao Dong and Kehao Zeng for their comments on earlier versions of these notes.

5	Spaces of linear functions	37
5.A	The dual space	37
5.B	The double dual	39
5.C	Annihilators	40
5.D	Dual functions	42
5.E	Representation of linear functions by matrices	44
5.F	Fast matrix multiplication	47
	Extra exercises for Section 5.	49
6	Exterior algebra	50
6.A	Multilinear functions	50
6.B	Alternating maps	51
6.C	Discussion of Theorem 6.15	53
6.D	Parity of permutations	54
6.E	Determinants	57
6.F	Proof of Theorem 6.15	60
6.G	Quantifier elimination	63
6.H	A combinatorial application: the skew set pairs inequality	64
	Extra exercises for Section 6.	65
7	Polynomials	65
7.A	Basic properties of polynomials	65
7.B	Polynomial division	67
7.C	The resultant	68
7.D	Multiple roots and derivatives	70
7.E	Algebraically closed fields	72
7.F	The Schwartz–Zippel lemma	72
7.G	Application: identity testing	74
	Extra exercises for Section 7.	76
8	Classification of linear transformations	78
8.A	Conjugacy	78
8.B	Eigenvectors and eigenvalues	79
8.C	Diagonalizable transformations	81
8.D	The Cayley–Hamilton theorem	82
8.E	Generalized eigenspaces	85
8.F	Structure of nilpotent transformations	86
8.G	The Jordan normal form	88
	Extra exercises for Section 8.	90
	References	90

1. FIELDS

1.A. A problem to think about

Linear algebra, by and large, is concerned with systems of linear equations of the form

$$a_1x_1 + \cdots + a_nx_n = b.$$

Ponder the following two questions:

(Q1) Where do the coefficients a_1, \dots, a_n, b and the variables x_1, \dots, x_n come from?

(Q2) How many equations and variables can such a system include? Can there be *infinitely* many?

Most introductory linear algebra courses answer (Q1) by requiring the coefficients and the variables to be real numbers, and (Q2) by prohibiting infinite systems. However, it turns out that a lot of linear-algebraic techniques can be used in a much broader context, as we will soon discover.

The following problem gives an example of a situation when there are infinitely many variables and equations:

Problem 1.1. Let a_1, \dots, a_n be distinct integers and suppose that $f: \mathbb{Z} \rightarrow \mathbb{R}$ is a function such that for all $k \in \mathbb{Z}$ and $\ell \in \mathbb{Z}^+$, we have

$$f(k + a_1\ell) + f(k + a_2\ell) + \cdots + f(k + a_n\ell) = 0. \tag{1.2}$$

Must it be that $f(m) = 0$ for all $m \in \mathbb{Z}$?

We can think of the values $f(m)$, $m \in \mathbb{Z}$, as real variables indexed by the integers. Then (1.2) is a homogeneous linear equation in these variables, and what Problem 1.1 is asking is whether the infinite system formed by these equations has a nontrivial solution. Here is how one can go about answering this question for a specific choice of a_1, \dots, a_n :

Example 1.3. Suppose that $n = 3$ and $a_1 = 0, a_2 = 1, a_3 = 2$. Then $f(m) = 0$ for all $m \in \mathbb{Z}$, as the following calculation shows:

$$\begin{aligned} f(m) &= \frac{2}{3} \cdot \underbrace{(f(m) + f(m+1) + f(m+2))}_{k=m, \ell=1} - \frac{2}{3} \cdot \underbrace{(f(m+1) + f(m+2) + f(m+3))}_{k=m+1, \ell=1} \\ &\quad + \frac{1}{3} \cdot \underbrace{(f(m+3) + f(m+4) + f(m+5))}_{k=m+3, \ell=1} - \frac{1}{3} \cdot \underbrace{(f(m+4) + f(m+5) + f(m+6))}_{k=m+4, \ell=1} \\ &\quad + \frac{1}{3} \cdot \underbrace{(f(m) + f(m+3) + f(m+6))}_{k=m, \ell=3} \\ &= 0. \end{aligned} \tag{1.4}$$

In principle, we might hope to combine different instances of (1.2) in a manner similar to (1.4) to obtain the equality $f(m) = 0$. However, doing this explicitly for arbitrary a_1, \dots, a_n is tricky. We will eventually be able to sidestep this difficulty and solve Problem 1.1 almost effortlessly. To really appreciate the power of the general theory that we will develop, the reader is encouraged to try their hand on some concrete instances of Problem 1.1, such as the following:

Exercise 1.5. Solve Problem 1.1 for $n = 3$ and $a_1 = 0, a_2 = 1, a_3 = 3$.

1.B. Groups, rings, fields

For now we will focus on question (Q1). It is clear that for a linear equation

$$a_1x_1 + \cdots + a_nx_n = b$$

to make sense, there has to be a way to *multiply* and *add* entities such as the a_i 's and the x_i 's. For instance, the a_i 's and the x_i 's might be real numbers. But they could also be *complex* numbers, or

rational numbers, or even *integers*. (The problem of finding integer solutions to linear equations is part of the subject called *integer programming*, which has numerous applications in computer science.) It turns out that linear algebra works best when the coefficients and the variables are elements of an algebraic structure called a *field*.

To define fields, we will need a brief recap of some basic notions from abstract algebra. The reader should be warned: The rest of this subsection is a tiresome journey through a sea of boring-sounding terms and technical definitions, but, trust me, this evil is a necessary one, and soon enough we will reap the bountiful fruits of our labors.

A **binary operation** on a set S is a map

$$\star: S \times S \rightarrow S.$$

In other words, \star takes an (ordered) pair of elements of S as an input and outputs a single element of S . It is customary to write $a \star b$ instead of $\star(a, b)$ (for instance, we write $a + b$ instead of $+(a, b)$). A binary operation \star is **associative** if for all $a, b, c \in S$, we have

$$(a \star b) \star c = a \star (b \star c).$$

Associativity of \star means that it makes sense to write $a \star b \star c$, since the placement of parentheses, doesn't affect the outcome. This observation generalizes to more than three elements:

Exercise 1.6. Show that if \star is an associative binary operation on a set S , then for all $a_1, \dots, a_n \in S$, the value $a_1 \star a_2 \star \dots \star a_n$ is well-defined and independent of the placement of parentheses.

A binary operation \star on a set S is **commutative** if for all $a, b \in S$, we have

$$a \star b = b \star a.$$

An element $e \in S$ is an **identity** of \star if for all $a \in S$,

$$a \star e = e \star a = a.$$

Lemma 1.7. *If \star is a binary operation with identity, then the identity of \star is unique.*

PROOF. If e, e' are identities of \star , then

$$e = e \star e' = e'. \quad \blacksquare$$

Let \star be a binary operation on a set S and let e be an identity of \star . An **inverse** of $a \in S$ is an element $b \in S$ such that

$$a \star b = b \star a = e.$$

One might hope that, by analogy with Lemma 1.7, every element $a \in S$ can have at most one inverse. Unfortunately, this hope is false in general:

Exercise 1.8. Give an example of a binary operation \star with identity for which there is an element with more than one inverse.

Nevertheless, if \star is *associative*, then inverses are unique:

Lemma 1.9. *If \star is an associative binary operation with identity e , then every element has at most one inverse with respect to \star .*

PROOF. Suppose that b and b' are both inverses of a . Then

$$b = b \star e = b \star (a \star b') = (b \star a) \star b' = e \star b' = b'. \quad \blacksquare$$

The next definition identifies what can perhaps be called the most important class of algebraic structures:

Definition 1.10. A **group** is a set G equipped with an associative binary operation \star such that \star has an identity e and every element $g \in G$ has an inverse.

A group (G, \star) is called **commutative**, or **Abelian**, if the operation \star is commutative.

Example 1.11. $(\mathbb{R}, +)$ is an Abelian group. The identity of this group is 0, and the inverse of an element $a \in \mathbb{R}$ is $(-a)$.

Example 1.12. (\mathbb{R}, \cdot) is *not* a group. The multiplication operation is associative and commutative; it also has an identity, namely the number 1. But the element 0 has no inverse.

Exercise 1.13. Is $(\mathbb{R} \setminus \{0\}, \cdot)$ an Abelian group?

Now we define a class of structures with two operations that resemble the usual addition and multiplication:

Definition 1.14. A **ring** is a set R equipped with a pair of binary operations $+$ and \cdot , called **addition** and **multiplication** respectively, such that:

- (R1) $(R, +)$ is an Abelian group, whose identity is denoted 0;
- (R2) multiplication is an associative operation with identity, which is denoted 1;
- (R3) for all $a, b, c \in R$, we have

$$a \cdot (b + c) = (a \cdot b) + (a \cdot c) \quad \text{and} \quad (a + b) \cdot c = (a \cdot c) + (b \cdot c).$$

A ring $(R, +, \cdot)$ is called **commutative** if \cdot is a commutative operation.¹

Example 1.15. $(\mathbb{R}, +, \cdot)$ and $(\mathbb{Z}, +, \cdot)$ (where $+$ and \cdot are the usual addition and multiplication) are commutative rings.

Example 1.16. Let $M_{n \times n}(\mathbb{R})$ denote the set of all n -by- n matrices with real entries. Then $M_{n \times n}(\mathbb{R})$ is a ring under the usual operations of matrix addition and multiplication, but when $n \geq 2$, this ring is not commutative. For instance, for $n = 2$, we have

$$\begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix} \neq \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}.$$

Exercise 1.17 (important). Show that in any ring R , the following identities hold for all $a, b \in R$:

$$\begin{aligned} a \cdot 0 &= 0 \cdot a = 0; \\ a \cdot (-b) &= (-a) \cdot b = -(a \cdot b). \end{aligned}$$

Finally, we can define fields, which are particularly nice rings:

Definition 1.18. A **field** is a commutative ring F in which $0 \neq 1$ and every element $a \in F \setminus \{0\}$ has a multiplicative inverse.

Exercise 1.19. Why is it necessary to require $0 \neq 1$ in the definition of a field?

1.C. Examples of fields

Example 1.20. \mathbb{R} , \mathbb{C} , and \mathbb{Q} , equipped with the usual addition and multiplication operations, are fields. On the other hand, \mathbb{Z} is a commutative ring but not a field, since some elements (actually, *all* nonzero elements apart from 1 and -1) have no multiplicative inverses. The set of nonnegative integers \mathbb{N} is not even a ring, since most elements of \mathbb{N} have no additive inverses.

Example 1.21. Consider the set

$$\mathbb{Q}(i) := \{a + bi : a, b \in \mathbb{Q}\} \subset \mathbb{C}.$$

We claim that $\mathbb{Q}(i)$ is a field under the usual addition and multiplication operations; in other words, $\mathbb{Q}(i)$ is a *subfield* of \mathbb{C} . First, we have to show that $\mathbb{Q}(i)$ is *closed* under addition and multiplication;

¹Because $+$ is commutative by definition.

i.e., if $x, y \in \mathbb{Q}(i)$, then $x + y, xy \in \mathbb{Q}(i)$ as well. For addition, this is almost trivial ([exercise!](#)), while for multiplication, we have

$$(a + bi)(c + di) = (ac - bd) + (ad + bc)i,$$

and if a, b, c , and d are rational, then so are $(ac - bd)$ and $(ad + bc)$. Now it is fairly easy to see that $\mathbb{Q}(i)$ is a ring (another [exercise!](#)), so it remains to verify that every nonzero element of $\mathbb{Q}(i)$ has a multiplicative inverse in $\mathbb{Q}(i)$. But

$$\frac{1}{a + bi} = \frac{a}{a^2 + b^2} + \frac{-b}{a^2 + b^2}i,$$

and if $a, b \in \mathbb{Q}$, then $a/(a^2 + b^2), -b/(a^2 + b^2) \in \mathbb{Q}$ as well.

Exercise 1.22. Let n be a positive integer and define

$$\mathbb{Q}(\sqrt{n}) := \{a + b\sqrt{n} : a, b \in \mathbb{Q}\} \subset \mathbb{R}.$$

Prove that $\mathbb{Q}(\sqrt{n})$ is a subfield of \mathbb{R} .

Example 1.23. A complex number $a \in \mathbb{C}$ is **algebraic** if there is a nonzero polynomial $p(x)$ with rational coefficients such that $p(a) = 0$. Some examples of algebraic numbers are:

- rational numbers (if $a \in \mathbb{Q}$, then it is a root of the polynomial $x - a$ with rational coefficients);
- $\sqrt{2}$, which is a root of $x^2 - 2$;
- the imaginary unit i , which is a root of $x^2 + 1$;
- $\sqrt[3]{5}$, which is a root of $x^3 - 5$;
- the golden ratio $(1 + \sqrt{5})/2$, which is a root of $x^2 - x - 1$;
- the five complex roots of $x^5 - 4x + 2$;
- &etc.

Exercise 1.24. Show that every element of $\mathbb{Q}(i)$ is algebraic. (See Example 1.21 for the definition of $\mathbb{Q}(i)$.) Show that if n is a positive integer, then every element of $\mathbb{Q}(\sqrt{n})$ is algebraic. (See Exercise 1.22 for the definition of $\mathbb{Q}(\sqrt{n})$.)

Some numbers that are *not* algebraic are $\pi = 3.1415\dots$ and $e = 2.7182\dots$ ² Denote the set of all algebraic numbers by $\overline{\mathbb{Q}}$. Then $\overline{\mathbb{Q}}$ is a subset of \mathbb{C} . It turns out that $\overline{\mathbb{Q}}$ is a field:

Theorem 1.25. $\overline{\mathbb{Q}}$ is a subfield of \mathbb{C} .

Note that this theorem is far from obvious, because it is not clear from the definition that $\overline{\mathbb{Q}}$ is closed under addition and multiplication:

Exercise 1.26 (hard). Suppose that a is a root of $x^5 - 4x + 2$ and b is a root of $2x^5 - 5x^4 + 5$. Find a nonzero polynomial $p(x)$ with rational coefficients such that $p(a + b) = 0$.

We will prove Theorem 1.25 later on, when we have enough machinery to attack it efficiently.

1.D. Finite fields

By definition, every field must contain at least two distinct elements, namely 0 and 1. It turns out, there is a field with *only* two elements. Namely, let $\mathbb{F}_2 := \{0, 1\}$ and define $+$ and \cdot on \mathbb{F}_2 as follows:

$$\begin{array}{c|cc} + & 0 & 1 \\ \hline 0 & 0 & 1 \\ 1 & 1 & 0 \end{array} \qquad \begin{array}{c|cc} \cdot & 0 & 1 \\ \hline 0 & 0 & 0 \\ 1 & 0 & 1 \end{array}$$

²The fact that e is not algebraic was first established by Charles Hermite in 1873, while the non-algebraicity of π was proved in 1882 by Ferdinand von Lindemann.

It is not hard to check that these two operations turn \mathbb{F}_2 into a field.

More generally, let n be an integer ≥ 2 . Let \mathbb{Z}_n be the set $\{0, 1, \dots, n-1\}$, equipped with addition and multiplication modulo n . In other words, to add/multiply two elements $a, b \in \mathbb{Z}_n$, we first add/multiply them as integers and then compute the remainder of the result after division by n . For instance, $2 + 2 = 1$ in \mathbb{Z}_3 , because, if we add 2 and 2 as integers, we get 4, and the remainder of 4 after division by 3 is 1. We usually write $2 + 2 = 1 \pmod{3}$ to emphasize that the operation is performed in \mathbb{Z}_3 instead of \mathbb{Z} .

Exercise 1.27 (tedious but straightforward). Show that \mathbb{Z}_n is a commutative ring.

Theorem 1.28. *Let n be an integer ≥ 2 . \mathbb{Z}_n is a field if and only if n is a prime number.*

PROOF. First, suppose that n is not prime. To show that \mathbb{Z}_n is not a field, we will use the following lemma:

Lemma 1.29. *If F is a field and a, b are nonzero elements of F , then $ab \neq 0$.*

Proof. Assume, towards a contradiction, that a, b are nonzero elements such that $ab = 0$. Since $a \neq 0$ and F is a field, a has a multiplicative inverse a^{-1} . Then

$$b = 1 \cdot b = (a^{-1}a)b = a^{-1}(ab) = a^{-1} \cdot 0 = 0,$$

where the last equality is due to Exercise 1.17. —

If n is not a prime number, then $n = k\ell$ for some integers $2 \leq k, \ell < n$. Then k, ℓ are nonzero in \mathbb{Z}_n , yet $k\ell = n = 0 \pmod{n}$. In other words, \mathbb{Z}_n contradicts Lemma 1.29, and thus it is not a field.

Now suppose that n is a prime number. Since we already know that \mathbb{Z}_n is a commutative ring, and it is clear that $0 \neq 1 \pmod{n}$, we only need to show that every nonzero element of \mathbb{Z}_n has a multiplicative inverse. To that end, take any $a \in \mathbb{Z}_n \setminus \{0\}$. We need to find some $b \in \mathbb{Z}_n$ such that $ab = 1 \pmod{n}$. Consider the following function:

$$f_a: \mathbb{Z}_n \rightarrow \mathbb{Z}_n: b \mapsto ab.$$

We claim that f_a is *injective*, i.e., if $f_a(b_1) = f_a(b_2)$, then $b_1 = b_2$. Indeed, $f_a(b_1) = f_a(b_2)$ means that $ab_1 = ab_2$ in \mathbb{Z}_n , i.e., $a(b_1 - b_2)$ is divisible by n . Since n is a prime number and $a \neq 0 \pmod{n}$, it must be that $(b_1 - b_2)$ is divisible by n . In other words, $b_1 - b_2 = 0 \pmod{n}$, or, equivalently, $b_1 = b_2 \pmod{n}$, as claimed. The set \mathbb{Z}_n is finite, so if the map $f_a: \mathbb{Z}_n \rightarrow \mathbb{Z}_n$ is injective, then it must also be *surjective*, i.e., for every $c \in \mathbb{Z}_n$ there is some $b \in \mathbb{Z}_n$ such that $f_a(b) = c$. Taking $c = 1$, we obtain $b \in \mathbb{Z}_n$ such that $ab = f_a(b) = c = 1$, as desired. Therefore, every nonzero element of \mathbb{Z}_n has a multiplicative inverse, and so \mathbb{Z}_n is indeed a field. ■

Remark. There are other ways to show that \mathbb{Z}_n is a field when n is prime. A very general approach involves the so-called Euclidean algorithm. We will discuss it later on in the context of polynomial division.

For a prime number p , we write \mathbb{F}_p instead of \mathbb{Z}_p to emphasize that it is a field. (Another common notation for this field is $GF(p)$, standing for “the Galois field of order p .”)

Exercise 1.30. Suppose that p is a prime number and F is a finite field of size p . Show that F is isomorphic to \mathbb{F}_p .

There exist other finite fields. In fact, the following is true:

Theorem 1.31. *Let n be an integer ≥ 2 . There exists a finite field of size n if and only if n is a power of a prime number, in which case all the fields of size n are isomorphic to each other.*

If q is a prime power, then the unique (up to isomorphism) field of size q is denoted \mathbb{F}_q (or $GF(q)$). Note that unless q is itself prime, \mathbb{F}_q is *not* the same as \mathbb{Z}_q . We will not prove Theorem 1.31 in these notes.

1.E. Matrices over rings

Matrices are very important objects in linear algebra. Let R be a ring (not necessarily commutative). An **m -by- n matrix** over R is a rectangular array of elements of R indexed by the pairs (i, j) with $1 \leq i \leq m$, $1 \leq j \leq n$. For a matrix A , we write $A(i, j)$ for the entry of A in the position (i, j) .^{3,4} Thus, a typical m -by- n matrix A looks like this:

$$A = \begin{bmatrix} A(1,1) & A(1,2) & \cdots & A(1,n) \\ A(2,1) & A(2,2) & \cdots & A(2,n) \\ \vdots & \vdots & \ddots & \vdots \\ A(m,1) & A(m,2) & \cdots & A(m,n) \end{bmatrix}.$$

The set of all m -by- n matrices over R is denoted by $M_{m \times n}(R)$. For a pair of matrices $A, B \in M_{m \times n}(R)$, their **sum** $A + B \in M_{m \times n}(R)$ is the matrix given by

$$(A + B)(i, j) := A(i, j) + B(i, j) \quad (1.32)$$

for all $1 \leq i \leq m$, $1 \leq j \leq n$. Note that “+” on the right-hand side of (1.32) indicates the addition operation in the ring R .

Exercise 1.33. Show that $(M_{m \times n}(R), +)$ is an Abelian group.

To define matrix multiplication, we first need a convenient piece of notation. Let I be a finite set, say of size n , and let $(a_i)_{i \in I}$ be a sequence of elements of R indexed by I . Choose an arbitrary ordering i_1, \dots, i_n of I and define

$$\sum_{i \in I} a_i := a_{i_1} + \cdots + a_{i_n}. \quad (1.34)$$

The expression on the right-hand side of (1.34) makes sense (because the addition in R is associative, see Exercise 1.6) and, crucially, its value is independent of the particular ordering i_1, \dots, i_n (because addition is commutative), so (1.34) is a valid definition. If $I = \emptyset$, then, by convention, $\sum_{i \in \emptyset} a_i := 0$. When $I = \{1, \dots, n\}$ for some $n \in \mathbb{N}$, we also use the notation $\sum_{i=1}^n a_i := \sum_{i \in \{1, \dots, n\}} a_i$.

Now let $A \in M_{m \times n}(R)$ and $B \in M_{n \times r}(R)$. Note that we require the number of columns of A to match the number of rows of B . The **product** of A and B is the m -by- r matrix $AB \in M_{m \times r}(R)$ given by

$$(AB)(i, j) := \sum_{k=1}^n A(i, k)B(k, j) \quad (1.35)$$

for all $1 \leq i \leq m$, $1 \leq j \leq r$. Again, note that the addition and multiplication on the right-hand side of (1.35) are the corresponding operations in the ring R .

Exercise 1.36. Let $A \in M_{m \times n}(R)$, $B \in M_{n \times r}(R)$, and $C \in M_{r \times s}(R)$. Show that $(AB)C = A(BC)$.

While Exercise 1.36 can be solved by a direct calculation, we will eventually establish a conceptual reason for the equality $(AB)C = A(BC)$ as well as for the specific way matrix multiplication is defined.

By definition, the set $M_{n \times n}(R)$ of all n -by- n matrices is closed under matrix addition and multiplication. Furthermore, these operations, restricted to $M_{n \times n}(R)$, have identities: The additive identity is the **zero matrix**, i.e., the n -by- n matrix all of whose entries are zero; while the multiplicative

³If you wish to be pedantic, you could say that an m -by- n matrix over R is a function $A: \{1, \dots, m\} \times \{1, \dots, n\} \rightarrow R$.

⁴Actually, if you wish to be *even more* pedantic, you should note that the definition in ³ doesn't quite work, since it identifies the (unique) m -by-0 matrix with the (unique) n -by-0 matrix, even if $m \neq n$, which breaks down the rules of matrix multiplication.

identity is the **identity matrix** $I_n(R)$ (or simply I_n if the ring R is understood from the context), which looks like this:

$$I_n := \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{bmatrix}.$$

In other words,

$$I_n(i, j) = \begin{cases} 1 & \text{if } i = j; \\ 0 & \text{if } i \neq j. \end{cases}$$

In fact, $M_{n \times n}(R)$ is a ring:

Exercise 1.37. Show that $M_{n \times n}(R)$ is a ring under the matrix addition and multiplication.

Another useful matrix operation is transposition. The **transpose** of a matrix $A \in M_{m \times n}(R)$ is the matrix $A^\top \in M_{n \times m}(R)$ given by $A^\top(i, j) := A(j, i)$ for all $1 \leq i \leq n$, $1 \leq j \leq m$.

Exercise 1.38. Let $A \in M_{m \times n}(R)$ and $B \in M_{n \times r}(R)$. Prove that $(A^\top)^\top = A$ and if R is commutative, then $(AB)^\top = B^\top A^\top$.

1.F. An application of linear algebra over an unusual field

One of the results that we will prove is that for any field F , there is a way to assign to each matrix A over F a natural number $\text{rank}(A)$, called the *rank* of A , with the following properties:

- (r1) $\text{rank}(I_n) = n$ for all $n \in \mathbb{N}$;
- (r2) if $A \in M_{m \times n}(F)$, then $\text{rank}(A) \leq \min\{m, n\}$;
- (r3) if $A \in M_{m \times n}(F)$ and $B \in M_{n \times r}(F)$, then $\text{rank}(AB) \leq \min\{\text{rank}(A), \text{rank}(B)\}$.

The reader is probably familiar with at least some of these properties in the case of matrices with real entries; what makes this result particularly striking is that the same holds for matrices over an arbitrary field. The proof of the next theorem shows the power of a judicious choice of a field:

Theorem 1.39 (Babai–Frankl?). *Let n, m be positive integers and suppose that S_1, \dots, S_m are subsets of the set $\{1, \dots, n\}$ such that:*

- for each i , the size of S_i is odd;
- for all distinct i, j , the size of the intersection $S_i \cap S_j$ is even.

Then $m \leq n$.

Two quick remarks before we start the proof: First, there are exponentially many (2^{n-1} , to be precise) distinct subsets of $\{1, \dots, n\}$ of odd size, and it is remarkable how drastically this number decreases when we add the requirement that the pairwise intersections of the sets must be even. Second, the bound $m \leq n$ is sharp, since the n sets $\{1\}, \{2\}, \dots, \{n\}$ all have odd size (namely 1) while their pairwise intersections have even size (namely 0).

PROOF. Form an m -by- n matrix A according to the formula

$$A(i, j) := \begin{cases} 1 & \text{if } j \in S_i; \\ 0 & \text{if } j \notin S_i. \end{cases} \tag{1.40}$$

We view A as a matrix over the field \mathbb{F}_2 . Consider the matrix $B := AA^\top$ (where the matrix multiplication is performed over \mathbb{F}_2). Then B is an m -by- m matrix and, by definition,

$$B(i, j) = \sum_{k=1}^n A(i, k)A(j, k).$$

The product $A(i, k)A(j, k)$ is either 0 or 1, and it is equal to 1 only if $A(i, k) = A(j, k) = 1$. By (1.40), we have $A(i, k) = A(j, k) = 1$ if and only if $k \in S_i \cap S_j$. Therefore,

$$B(i, j) = \sum_{k \in S_i \cap S_j} 1 = |S_i \cap S_j| \pmod{2}.$$

(Addition is performed modulo 2, since we are working in \mathbb{F}_2 .) By the assumptions of the theorem,

$$|S_i \cap S_j| \text{ is } \begin{cases} \text{odd} & \text{if } i = j; \\ \text{even} & \text{if } i \neq j. \end{cases}$$

Hence,

$$B(i, j) = \begin{cases} 1 & \text{if } i = j; \\ 0 & \text{if } i \neq j; \end{cases}$$

in other words, B is the identity matrix I_m . And now we are done, since

$$m \stackrel{\text{(r1)}}{=} \text{rank}(I_m) = \text{rank}(B) = \text{rank}(AA^\top) \stackrel{\text{(r3)}}{\leq} \text{rank}(A) \stackrel{\text{(r2)}}{\leq} n. \quad \blacksquare$$

The proof of Theorem 1.39 given above uses only a minimum amount of linear algebra, and all the missing details can be filled in quite easily. For matrices over \mathbb{F}_2 , the rank of an m -by- n matrix $A \in M_{m \times n}(\mathbb{F}_2)$ can be defined by

$$\text{rank}(A) := \log_2 |\{Ax : x \in M_{n \times 1}(\mathbb{F}_2)\}|. \quad (1.41)$$

It is fairly straightforward to check that this definition fulfills conditions (r1), (r2), and (r3). For instance, note that if $A \in M_{m \times n}(\mathbb{F}_2)$, then $\{Ax : x \in M_{n \times 1}(\mathbb{F}_2)\} \subseteq M_{m \times 1}(\mathbb{F}_2)$, so

$$\text{rank}(A) = \log_2 |\{Ax : x \in M_{n \times 1}(\mathbb{F}_2)\}| \leq \log_2 |M_{m \times 1}(\mathbb{F}_2)| = m.$$

On the other hand, the number of m -by-1 matrices that can be expressed in the form Ax with $x \in M_{n \times 1}(\mathbb{F}_2)$ cannot exceed the number of different choices for x , and so

$$\text{rank}(A) = \log_2 |\{Ax : x \in M_{n \times 1}(\mathbb{F}_2)\}| \leq \log_2 |M_{n \times 1}(\mathbb{F}_2)| = n.$$

This proves (r2).

Exercise 1.42. Verify that the notion of rank given by (1.41) satisfies (r1) and (r3).

A definition similar to (1.41) works over any *finite* field. We will show that even over an infinite field, there is a way to perform analogous “counting” arguments. Another difficulty that we will tackle is finding simple ways for computing or estimating the rank of a given matrix.

Extra exercises for Section 1

Exercise 1.43. An **integral domain** is a commutative ring R in which $0 \neq 1$ and for all $a, b \in R \setminus \{0\}$, we have $ab \neq 0$. By Lemma 1.29, every field is an integral domain.

- Give an example of an integral domain that is not a field.
- Prove that every *finite* integral domain is a field.

Exercise 1.44 (Characteristic). Let F be a field. For $n \in \mathbb{N}$ and $a \in F$, define

$$n \cdot a := \underbrace{a + a + \cdots + a}_{n \text{ summands}}.$$

The **characteristic** of a field F is the natural number $\text{char}(F)$ that is equal to the smallest positive integer n satisfying $n \cdot 1 = 0$ if such n exists, and 0 otherwise. Show that if $\text{char}(F) > 0$, then $\text{char}(F)$ is a prime number.

Exercise 1.45. We say that a field K **contains a copy** of a field F if K has a subfield that is isomorphic to F . Let K be a field.

- (a) Show that if $\text{char}(K) = 0$, then K contains a copy of \mathbb{Q} .
 (b) Show that if $\text{char}(K) = p > 0$, then K contains a copy of \mathbb{F}_p .

2. VECTOR SPACES

2.A. The definition of a vector space

A central notion in linear algebra is that of a *vector space*. It is an abstract concept that captures the properties of sets of solutions to systems of homogeneous linear equations.

Definition 2.1. A **vector space** over a field F (also called an F -**vector space**) is a set V , whose elements are referred to as **vectors**, equipped with a binary operation $+$, called **(vector) addition**, and a function $\cdot : F \times V \rightarrow V$, called **scalar multiplication**, or **scaling**, such that:

- (V1) $(V, +)$ is an Abelian group;
- (V2) for all $a, b \in F$ and $v \in V$, $(ab) \cdot v = a \cdot (b \cdot v)$;
- (V3) for all $a, b \in F$ and $v \in V$, $(a + b) \cdot v = (a \cdot v) + (b \cdot v)$;
- (V4) for all $a \in F$ and $v, w \in V$, $a \cdot (v + w) = (a \cdot v) + (a \cdot w)$;
- (V5) for all $v \in V$, $1 \cdot v = v$.

A vector space is an object that is heavily laden with operations (and the fact that many of them share a name makes things even more confusing!). If V is a vector space over a field F , then the following operations are around:

- the field addition on F , denoted $+$;
- the vector addition on V , *also* denoted $+$;⁵
- the field multiplication on F , denoted \cdot or by juxtaposition; and
- the scalar multiplication, which allows one to multiply a vector $v \in V$ by a field element $a \in F$, and is *also* denoted \cdot or by juxtaposition.

These operations bring in with them some extra notation:

- the additive identity in F , denoted 0 ;
- the additive identity in V , *also* denoted 0 ;⁶
- the additive inverse of an element $a \in F$, denoted $-a$;
- the additive inverse of an element $v \in V$, denoted $-v$;⁷
- the multiplicative identity in F , denoted 1 ;
- the multiplicative inverse of a nonzero element $a \in F$, denoted a^{-1} or $1/a$.

It is very important to keep in mind, however, that scalar multiplication is *not* a binary operation on V , so there is no such thing as a “multiplicative identity in V .”

Exercise 2.2. Let V be a vector space over a field F . Show that for all $a \in F$ and $v \in V$,

$$0_F \cdot v = a \cdot 0_V = 0_V \quad \text{and} \quad (-1) \cdot v = -v.$$

(Cf. Exercise 1.17.)

The definition of a vector space is a true Goliath. Thankfully, with few examples in hand, one very rarely actually needs to apply it directly.⁸

⁵To make it absolutely clear which $+$ is meant, people sometimes write things like $+_F$ and $+_V$, but only occasionally.

⁶You can, of course, clarify the matters by writing something like 0_F and 0_V . Also, people sometimes use boldface zero: $\mathbf{0}$, or a zero with an arrow: $\vec{0}$, for the additive identity in V , but both these symbols are uncommon.

⁷One might think that there is no way to confuse the notation for the additive inverses in F and in V since the first one applies to the elements of the field F and the second one to the elements of the vector space V . However, nothing prevents an object x to be *both* an element of F and of V and to have different additive inverses there. Then, the expression “ $-x$ ” is ambiguous.

⁸If you’ve seen one vector space, you’ve seen them all!

Example 2.3. This is the prototypical example of a vector space. Let F be any field and consider the set F^n of all n -tuples of the elements of F . We can add and scale them as follows:

$$\begin{aligned}(x_1, \dots, x_n) + (y_1, \dots, y_n) &:= (x_1 + y_1, \dots, x_n + y_n); \\ a \cdot (x_1, \dots, x_n) &:= (ax_1, \dots, ax_n).\end{aligned}$$

It is easy to check that these definitions give F^n the structure of an F -vector space. (Most of the required properties follow simply because F is a field.)




Figure 1. Vector addition in \mathbb{R}^2 .

Example 2.4. The set $M_{m \times n}(F)$ of m -by- n matrices over a field F is a vector space over F under matrix addition and entry-wise scaling: $(a \cdot A)(i, j) := a(A(i, j))$. This is a special case of Example 2.3, since an m -by- n matrix can be viewed as a tuple of elements of F of length mn , and thus, as a vector space over F , $M_{m \times n}(F)$ is essentially the same as F^{mn} .

Example 2.5. Let F be a field. Then F is already equipped with addition and multiplication, and these operations make F a vector space over itself. (This is a special case of Example 2.3, since F can be identified with F^1 .) More generally, suppose that F is a subfield of a field K . Then the elements of K can be added to each other and multiplied by the elements of F , making K a vector space over F . Thus, for example, \mathbb{C} is a vector space over \mathbb{R} and \mathbb{R} is a vector space over \mathbb{Q} . Note that \mathbb{R} is at the same time *also* a vector space over \mathbb{R} —it is important to remember which field you are working over.

Example 2.6. A vector space must contain *at least one* element, namely its additive identity, and it is possible to construct a vector space with *only one* element. To that end, consider a one-element set $\{0\}$, whose only element is denoted 0 . We can equip $\{0\}$ with the structure of a vector space over any given field F by setting $0 + 0 := 0$ and $a \cdot 0 := 0$ for all $a \in F$. It is trivial to check that this is indeed a vector space. (It is sometimes convenient to think that this is also a special case of Example 2.3 with $n = 0$.)

Example 2.7. Let F be a field and let $F^{\mathbb{N}}$ denote the set of all infinite sequences (x_0, x_1, \dots) of elements of F . By analogy with Example 2.3, define

$$\begin{aligned}(x_0, x_1, \dots) + (y_0, y_1, \dots) &:= (x_0 + y_0, x_1 + y_1, \dots); \\ a \cdot (x_0, x_1, \dots) &:= (ax_0, ax_1, \dots).\end{aligned}$$

This makes $F^{\mathbb{N}}$ a vector space over F .

Example 2.8. This example is a common generalization of Examples 2.3 and 2.7. Let F be a field and let X be an arbitrary set. The set F^X of all functions from X to F is an F -vector space under the operations of pointwise addition and multiplication. That is, for each $f, g \in F^X$, we let $f + g$ be the function such that for all $x \in X$,

$$(f + g)(x) := f(x) + g(x).$$

Similarly, for each $f \in F^X$ and $a \in F$, we let $a \cdot f$ be the function given by

$$(a \cdot f)(x) := a(f(x)).$$

Example 2.3 is the special case of this construction for $X = \{1, \dots, n\}$, while Example 2.7 corresponds to the case $X = \mathbb{N}$.

Example 2.9. For a set X , let $\mathcal{P}(X)$ denote the **powerset** of X , i.e., the set of all subsets of X . We equip $\mathcal{P}(X)$ with the addition given by

$$A + B := A \triangle B,$$

where \triangle is the symmetric difference operation.⁹ Also, let $0 \cdot A := \emptyset$ and $1 \cdot A := A$ for all $A \in \mathcal{P}(X)$. These operations make $\mathcal{P}(X)$ a vector space over \mathbb{F}_2 . The easiest way to see this is as follows:

Exercise 2.10. Show that the \mathbb{F}_2 -vector space $\mathcal{P}(X)$ is isomorphic to \mathbb{F}_2^X .

2.B. Subspaces

Let V be an F -vector space. When is a subset $W \subseteq V$ a **subspace** of V , i.e., a vector space in its own right under the operations inherited from V ? First of all, W must be nonempty (recall that a vector space can't be empty—it must have an additive identity). Second, W must be closed under the vector space operations; that is, for all $x, y \in W$ and $a \in F$, the elements $x + y$ and $a \cdot x$ must be in W as well. Additionally, W must satisfy the axioms for being a vector space—but, as the following very useful lemma asserts, we get the vector space axioms for free:

Lemma 2.11. *Let V be a vector space over a field F . A nonempty subset $W \subseteq V$ is a subspace of V if and only if W is closed under addition and scaling by the elements of F .*

PROOF. If W is a subspace of V , then it is closed under addition and scaling by definition. Now suppose that $W \subseteq V$ is a nonempty subset of V that is closed under addition and scaling. Most of the properties required of a vector space hold in W simply because they hold in V (since they say something about *all* elements of a vector space). The only things to check are:

- Addition, restricted to W , has an additive identity. To that end, we will show that $0_V \in W$.
- Every element $x \in W$ has an additive inverse in W . We will show that $-x \in W$, where $-x$ is the additive inverse of x in V .

Since W is nonempty, there is at least one element $x \in W$. But W is closed under scaling, so

$$0_V = 0_F \cdot x \in W,$$

as claimed. (Here we use Exercise 2.2.) Also, for any $x \in W$, we have

$$-x = (-1) \cdot x \in W,$$

where we again use Exercise 2.2. ■

Example 2.12. The set

$$\{(x_1, x_2, x_3) : x_1 + 2x_2 + 3x_3 = 0 \text{ and } x_1 - x_2 + x_3 = 0\}$$

is a subspace of \mathbb{R}^3 (as a vector space over \mathbb{R}). More generally, the set of solutions to a system of homogeneous linear equations is a vector space. In some sense, the entire theory of vector spaces is a way to generalize this example.

Example 2.13. The set $\mathcal{C}([0; 1])$ of all continuous functions $f: [0; 1] \rightarrow \mathbb{R}$ is a subspace of $\mathbb{R}^{[0;1]}$ (considered as an \mathbb{R} -vector space).

Example 2.14. The set

$$\left\{ f \in \mathcal{C}([0; 1]) : \int_0^1 f(x) \, dx = 0 \right\}$$

is a subspace of $\mathcal{C}([0; 1])$. This is analogous to Example 2.12, with the expression

$$\int_0^1 f(x) \, dx = 0$$

⁹The **symmetric difference** of two sets A and B is the set consisting of all elements that belong to exactly one of the sets A, B . In other words, $A \triangle B = (A \cup B) \setminus (A \cap B)$.

playing the role of a homogeneous linear equation.

Example 2.15. The set of all twice-differentiable functions $f: [0; 1] \rightarrow \mathbb{R}$ satisfying

$$f'' - f' + f = 0 \tag{2.16}$$

is another subspace of $\mathcal{C}([0; 1])$. Here, (2.16) is playing the role of a homogeneous linear equation.

Example 2.17. The set \mathcal{A} of all functions $f: [0; 1] \rightarrow \mathbb{R}$ that have an antiderivative (i.e., a function $F: [0; 1] \rightarrow \mathbb{R}$ such that $F' = f$) is a subspace of $\mathbb{R}^{[0;1]}$. Notice that $\mathcal{C}([0; 1])$ is a subspace of \mathcal{A} .

Exercise 2.18. Show that $\mathcal{A} \neq \mathcal{C}([0; 1])$. *Hint:* Consider the derivative of $x^2 \sin(1/x)$.

Example 2.19. The set of all sequences (x_0, x_1, \dots) of real numbers such that for every $n \in \mathbb{N}$,

$$x_{n+2} = x_{n+1} + x_n,$$

is a subspace of $\mathbb{R}^{\mathbb{N}}$. One of the points in this subspace is the Fibonacci sequence $(1, 1, 2, 3, 5, 8, \dots)$.

Exercise 2.20. Show that the following sets are subspaces of $\mathbb{R}^{\mathbb{N}}$:

$$\ell^\infty(\mathbb{N}) := \left\{ (x_0, x_1, \dots) \in \mathbb{R}^{\mathbb{N}} : \sup_{n \in \mathbb{N}} |x_n| < \infty \right\};$$

$$\ell^1(\mathbb{N}) := \left\{ (x_0, x_1, \dots) \in \mathbb{R}^{\mathbb{N}} : \sum_{n=0}^{\infty} |x_n| < \infty \right\};$$

$$\ell^2(\mathbb{N}) := \left\{ (x_0, x_1, \dots) \in \mathbb{R}^{\mathbb{N}} : \sum_{n=0}^{\infty} x_n^2 < \infty \right\}.$$

(The last one is a bit tricky.)

Example 2.21. Let X be any set. Recall how in Example 2.9 we equipped the powerset $\mathcal{P}(X)$ with the structure of an \mathbb{F}_2 -vector space. Now let $[X]^{<\infty}$ denote the set of all *finite* subsets of X . Then $[X]^{<\infty}$ is a subspace of $\mathcal{P}(X)$ (because the symmetric difference of two finite sets is finite).

Example 2.22. This is an extension of Example 2.21. Let X be a set and let F be a field. For a function $f: X \rightarrow F$, let the **support** of f be the set

$$\text{supp}(f) := \{x \in X : f(x) \neq 0\}.$$

Denote by $[X \rightarrow F]^{<\infty}$ the set of all functions $f: X \rightarrow F$ whose support is finite. Then $[X \rightarrow F]^{<\infty}$ is a subspace of F^X . This example will become surprisingly important later.




Figure 2. The graph of a function with finite support.

2.C. Linear functions

The notion of a linear equation over an abstract vector space is captured in the concept of a *linear function*.

Definition 2.23. Let V and W be vector spaces over the same field F . A function $\varphi: V \rightarrow W$ is called **linear** (or, sometimes, **F -linear**) if:

- (L1) for all $x, y \in V$, $\varphi(x + y) = \varphi(x) + \varphi(y)$;
- (L2) for all $a \in F$ and $x \in V$, $\varphi(a \cdot x) = a \cdot \varphi(x)$.

If $\varphi: V \rightarrow W$ is a linear function, then the **kernel** of φ is the set

$$\ker(\varphi) := \{x \in V : \varphi(x) = 0\} \subseteq V,$$

and the **image** of φ is the set

$$\text{im}(\varphi) := \{\varphi(x) : x \in V\} \subseteq W.$$

Exercise 2.24. Show that if $\varphi: V \rightarrow W$ is a linear function, then $\varphi(0_V) = 0_W$.

You should think of the kernel of φ as the set of solutions to the “generalized linear equation”

$$\varphi(x) = 0.$$

The next lemma justifies this attitude by showing that $\ker(\varphi)$ is a vector space:

Lemma 2.25. Let V and W be vector spaces over a field F and let $\varphi: V \rightarrow W$ be a linear function. Then $\ker(\varphi)$ is a subspace of V and $\text{im}(\varphi)$ is a subspace of W .

PROOF. Since $0 \in \ker(\varphi)$ (see Exercise 2.24), the set $\ker(\varphi)$ is nonempty. To check that $\ker(\varphi)$ is closed under addition, consider any $x, y \in \ker(\varphi)$. We have

$$\varphi(x + y) = \varphi(x) + \varphi(y) = 0 + 0 = 0,$$

so $x + y \in \ker(\varphi)$, as desired. Similarly, if $x \in \ker(\varphi)$ and $a \in F$, then

$$\varphi(a \cdot x) = a \cdot \varphi(x) = a \cdot 0 = 0,$$

hence $a \cdot x \in \ker(\varphi)$. By Lemma 2.11, we conclude that $\ker(\varphi)$ is a subspace of V . Showing that $\text{im}(\varphi)$ is a subspace of W is left as an **exercise** (see Exercise 2.26). ■

Exercise 2.26. Prove that if V and W are vector spaces over a field F and $\varphi: V \rightarrow W$ is a linear function, then $\text{im}(\varphi)$ is a subspace of W .

Several of the examples in §2.B are naturally described as kernels or images of linear functions.

Example 2.27. The function

$$\mathbb{R}^3 \rightarrow \mathbb{R}^2: (x_1, x_2, x_3) \mapsto (x_1 + 2x_2 + 3x_3, x_1 - x_2 + x_3)$$

is \mathbb{R} -linear. (Cf. Example 2.12.)

Example 2.28. The function

$$\mathcal{C}([0; 1]) \rightarrow \mathbb{R}: f \mapsto \int_0^1 f(x) \, dx$$

is \mathbb{R} -linear. (Cf. Example 2.14.)

Example 2.29. Let $\mathcal{D}([0; 1])$ denote the set of all differentiable functions $f: [0; 1] \rightarrow \mathbb{R}$. It is a subspace of $\mathbb{R}^{[0; 1]}$, and the map

$$\partial: \mathcal{D}([0; 1]) \rightarrow \mathbb{R}^{[0; 1]}: f \mapsto f'$$

is linear. The image of ∂ is the space \mathcal{A} from Example 2.17.

2.D. Quotient spaces

In the light of the examples in §2.C, the following question appears natural:

Question 2.30. Can every subspace of a given vector space be expressed as $\ker(\varphi)$ or $\text{im}(\varphi)$ for some linear function φ ?

Recall that if $W = \ker(\varphi)$, then W is the solution set for the “generalized linear equation”

$$\varphi(x) = 0.$$

Thus, if the answer to Question 2.30 were positive, it would mean that linear equations are a general way of identifying subspaces. The next theorem asserts that this is indeed the case:

Theorem 2.31. *Let V be a vector space over a field F and let W be a subspace of V . Then there exist F -vector spaces X and Y and linear functions*

$$\varphi: X \rightarrow V \quad \text{and} \quad \psi: V \rightarrow Y$$

such that $W = \text{im}(\varphi) = \ker(\psi)$.

PROOF. For X we can just take the space W itself with $\varphi: W \rightarrow V$ being the identity map on W :

$$\varphi(x) := x \quad \text{for all } x \in W.$$

Clearly, φ is linear and $\text{im}(\varphi) = W$ by definition.

The construction of the space Y and the map $\psi: V \rightarrow Y$ such that $\ker(\psi) = W$ is somewhat more subtle. To motivate it, imagine that we are already given a linear map ψ such that $W = \ker(\psi)$. Then, for all $x \in V$ and $w \in W$, we must have

$$\psi(x + w) = \psi(x) + \psi(w) = \psi(x) + 0 = \psi(x). \quad (2.32)$$

For $x \in V$, let $x + W$ denote the following subset of V :

$$x + W := \{x + w : w \in W\}.$$

The set $x + W$ is called the W -**coset** of x , or the **translate** of W by x . The set of all W -cosets is denoted by V/W (so V/W is a *set of sets*). Note that $W = 0 + W \in V/W$. Observation (2.32) can be summarized as, “ ψ is constant on the W -cosets.” Thus, for a coset $S \in V/W$, we can set $\bar{\psi}(S)$ to be the common value of $\psi(x)$ for all $x \in S$.

The linearity of ψ imposes some restrictions on the relationship between the values $\bar{\psi}(S)$ for different cosets $S \in V/W$. Indeed, for all $x, y \in V$,

$$\bar{\psi}((x + y) + W) = \psi(x + y) = \psi(x) + \psi(y) = \bar{\psi}(x + W) + \bar{\psi}(y + W). \quad (2.33)$$

Similarly, if $x \in V$ and $a \in F$, then we have

$$\bar{\psi}((ax) + W) = \psi(ax) = a\psi(x) = a \cdot \bar{\psi}(x + W). \quad (2.34)$$

The idea now is to equip the set V/W *itself* with the structure of a vector space such that the map ψ given by $\psi(x) := x + W$ satisfies (2.33) and (2.34) by virtue of the definition.

Specifically, we endow V/W with addition and scaling operations defined by the formulas

$$(x + W) + (y + W) := (x + y) + W \quad \text{and} \quad a(x + W) := (ax) + W. \quad (2.35)$$

Some explanation is necessary here. The above expressions are really shortcuts for more technical definitions. A more precise way to define, say, the addition on V/W would be as follows: Given two cosets $S, T \in V/W$, choose any $x, y \in V$ such that $S = x + W$ and $T = y + W$ and set

$$S + T := (x + y) + W. \quad (2.36)$$

This leads to the question, why is the right-hand side of (2.36) independent of the choice of x and y ? If we chose some *other* elements $x', y' \in V$ such that $S = x' + W$ and $T = y' + W$, then $S + T$ would be defined as $(x' + y') + W$, and we must make sure that

$$(x + y) + W = (x' + y') + W,$$

for otherwise (2.36) is not a proper definition of $S + T$.




Figure 3. Adding W -cosets S and T in two ways.

To deal with this issue, we will use the following observation:

Claim 2.31.1. Let $x, y \in V$. We have $x + W = y + W$ if and only if $x - y \in W$.

Proof. Note that $x = x + 0 \in x + W$. Hence, if $x + W = y + W$, then $x \in y + W$, i.e., $x = y + w$ for some $w \in W$. But then $w = x - y$, and hence $x - y \in W$, as desired.

Now assume that $x - y \in W$. For each $w \in W$, we have

$$x + w = y + (x - y) + w.$$

Since W is closed under addition, we conclude that $(x - y) + w \in W$, so $x + w \in y + W$. Therefore,

$$x + W \subseteq y + W.$$

On the other hand, if $x - y \in W$, then $y - x = (-1) \cdot (x - y)$ is also in W (as W is closed under scaling), and the same argument as before shows that

$$y + W \subseteq x + W.$$

Hence, $x + W = y + W$, as claimed. ◻

To show that the right-hand side of (2.36) is independent of the choice of x and y , suppose that $S = x + W = x' + W$ and $T = y + W = y' + W$. By Claim 2.31.1, the differences $x - x'$ and $y - y'$ belong to W . Therefore,

$$(x + y) - (x' + y') = (x - x') + (y - y') \in W,$$

since W is closed under addition. But this means, by Claim 2.31.1 again, that

$$(x + y) + W = (x' + y') + W,$$

which is what we wanted.

Exercise 2.37. Show that the operation of scalar multiplication on V/W is similarly well-defined.

It is now easy to see that V/W is a vector space.

Exercise 2.38. Show that V/W , equipped with the operations given by (2.35), is a vector space.

The vector space V/W is called the **quotient** of V by W . The additive identity of V/W is W (recall that $W = 0 + W \in V/W$). Consider the function $\psi: V \rightarrow V/W$ given by

$$\psi(x) := x + W.$$

It is called the **quotient map** corresponding to W . The vector space structure on V/W is defined precisely so as to make the quotient map linear; furthermore,

$$\ker(\psi) = \{x \in V : x + W = W\} = W.$$

This completes the proof of Theorem 2.31. ■

Exercise 2.39. Let V be a vector space over a field F and let W be a subspace of V . Let $x \in V$ and $S \in V/W$. Show that $S = x + W$ if and only if $x \in S$.

Exercise 2.40 (First isomorphism theorem). Let V and W be vector spaces over a field F and let $\varphi: V \rightarrow W$ be a linear function. Show that the space $\text{im}(\varphi)$ is isomorphic to $V/\ker(\varphi)$.

Extra exercises for Section 2

Exercise 2.41. For a real number $x \in \mathbb{R}$, let $x \pmod{1}$ denote the **fractional part** of x , i.e., the unique number $\alpha \in [0; 1)$ such that $x - \alpha$ is an integer. For $\alpha, \beta \in [0; 1)$ and $r \in \mathbb{R}$, define

$$\alpha \oplus \beta := (\alpha + \beta) \pmod{1} \quad \text{and} \quad r \odot \alpha := (r\alpha) \pmod{1}.$$

Does this definition make $[0; 1)$ into an \mathbb{R} -vector space?

Exercise 2.42 (Direct sums). Fix a field F . The **direct sum** of two F -vector spaces V and W is the F -vector space $V \oplus W$ defined as follows. As a set, $V \oplus W$ is equal to $V \times W$, and addition and scalar multiplication on $V \oplus W$ are defined component-wise:

$$(v_1, w_1) + (v_2, w_2) := (v_1 + v_2, w_1 + w_2) \quad \text{and} \quad a \cdot (v, w) := (a \cdot v, a \cdot w).$$

Prove that a function $f: V \rightarrow W$ is linear if and only if its graph is a subspace of $V \oplus W$.

3. BASES

3.A. A “better” version of Theorem 2.31

Recall that the purported goal of Theorem 2.31 was to show that every subspace of a vector space can be defined by a “generalized system of linear equations.” In that regard, Theorem 2.31 is not entirely satisfactory. Imagine that W is a subspace of, say, \mathbb{R}^5 , viewed as a vector space over \mathbb{R} . Then, according to Theorem 2.31, W is the kernel of the quotient map $\mathbb{R}^5 \rightarrow \mathbb{R}^5/W$. This is not particularly illuminating, because the quotient space \mathbb{R}^5/W is *defined* to have precisely this property; in some sense, it is not a very “natural” space.

This situation is remedied by the following fact:

Theorem 3.1. *Let V be a vector space over a field F and let W be a subspace of V . Then there exist linear functions*

$$\varphi: V \rightarrow V \quad \text{and} \quad \psi: V \rightarrow V$$

such that $W = \text{im}(\varphi) = \ker(\psi)$.

If W is a subspace of \mathbb{R}^5 , Theorem 3.1 asserts that W is the kernel of some linear function $\mathbb{R}^5 \rightarrow \mathbb{R}^5$. It is not hard to see that this is equivalent to identifying W with the solution set of an “honest-to-goodness” system of five homogeneous linear equations in five variables.

Let us ponder on how we could try to prove Theorem 3.1 in a specific case: Suppose that $V = \mathbb{R}$, viewed as a vector space over \mathbb{Q} (so only scaling by the rationals is allowed), and $W = \mathbb{Q}$. To prove that W is the image of a linear map from \mathbb{R} to \mathbb{R} , we have to find a \mathbb{Q} -linear function $\varphi: \mathbb{R} \rightarrow \mathbb{Q}$ that is not identically zero.

Since we don't know what else to do, let's take an arbitrary real number and decide where φ should send it. Say, take the number 1 and declare $\varphi(1) := 1$. Since φ must be \mathbb{Q} -linear, we then also have

$$\varphi(a) = a \cdot \varphi(1) = a, \quad \text{for all } a \in \mathbb{Q}.$$

But the values of $\varphi(x)$ for *irrational* x are still undefined. So take some irrational number, say $\sqrt{2}$, and declare $\varphi(\sqrt{2}) := 1$. Again, φ must be \mathbb{Q} -linear, and we are forced to have

$$\varphi(a + b\sqrt{2}) = a \cdot \varphi(1) + b \cdot \varphi(\sqrt{2}) = a + b, \quad \text{for all } a, b \in \mathbb{Q}.$$

Thus, the values of φ are determined for all real numbers of the form $a + b\sqrt{2}$ with a and b rational. But, for instance, $\varphi(\sqrt{3})$ is still undefined, so we can declare $\varphi(\sqrt{3}) := 1$, which forces

$$\varphi(a + b\sqrt{2} + c\sqrt{3}) = a + b + c, \quad \text{for all } a, b, c \in \mathbb{Q}.$$

Yet, there are still infinitely many real numbers x for which $\varphi(x)$ is undefined. . .

In the above attempted construction, we were building a sequence of real numbers $1, \sqrt{2}, \sqrt{3}, \dots$. To each of the numbers in the sequence, we could assign the value of φ arbitrarily, but these arbitrary decisions were forcing particular values at some other real numbers. If the values at *all* real numbers were determined, we would have our desired φ . For this, the sequence $1, \sqrt{2}, \sqrt{3}, \dots$ must have the properties of a *basis* for \mathbb{R} over \mathbb{Q} . In this section, we will define what a basis is formally and show that every vector space has one.

3.B. Spanning sets and independent sets

Exercise 3.2 (important). Let V and W be F -vector spaces and let $\varphi: V \rightarrow W$ be a linear map. Show that:

- φ is surjective $\iff \text{im}(\varphi) = W$;
- φ is injective $\iff \ker(\varphi) = \{0\}$.

Exercise 3.3. Let V be a vector space over a field F and let \mathcal{W} be a nonempty family of subspaces of V . Show that $\bigcap \mathcal{W}$, the intersection of all $W \in \mathcal{W}$, is also a subspace of V .

Remark 3.4. Note that the family \mathcal{W} in Exercise 3.3 is allowed to be infinite.

Using the result of Exercise 3.3, we can make the following important definition:

Definition 3.5. Let V be a vector space over a field F and let $X \subseteq V$ be a subset of V . Let \mathcal{W}_X be the set of all subspaces $W \subseteq V$ such that $X \subseteq W$. Since $V \in \mathcal{W}_X$, the family \mathcal{W}_X is nonempty, and hence we can define the **span** of X to be the space

$$\text{Span}(X) := \bigcap \mathcal{W}_X.$$

In other words, $\text{Span}(X)$ is the smallest subspace of V that contains X . Sometimes, we write $\text{Span}_F(X)$ instead of $\text{Span}(X)$ to explicitly indicate that we are working with vector spaces over F (for instance, if we want to make a distinction between $\text{Span}_{\mathbb{R}}(X)$ and $\text{Span}_{\mathbb{Q}}(X)$ for $X \subseteq \mathbb{R}$).

Exercise 3.6. Verify the following properties of span:

- $\text{Span}(\emptyset) = \{0\}$ and $\text{Span}(V) = V$;
- $X \subseteq \text{Span}(X)$;
- $Y \subseteq X \implies \text{Span}(Y) \subseteq \text{Span}(X)$;
- $\text{Span}(\text{Span}(X)) = \text{Span}(X)$.

Definition 3.5 describes the span of X abstractly and does not provide a concrete way of determining whether a given vector y is in $\text{Span}(X)$. Such a concrete description is supplied by the next lemma:

Lemma 3.7. Let V be a vector space over a field F and let $X \subseteq V$ be a subset of V . Define $\text{Span}^*(X)$ to be the set of all **linear combinations** of elements of X , i.e., vectors of the form

$$a_1x_1 + \cdots + a_nx_n, \quad (3.8)$$

where $a_1, \dots, a_n \in F$ and $x_1, \dots, x_n \in X$ (this includes the case $n = 0$, in which, by convention, expression (3.8) evaluates to 0_V). Then $\text{Span}^*(X) = \text{Span}(X)$.

PROOF. We have $\text{Span}^*(X) \subseteq \text{Span}(X)$ because $\text{Span}(X)$ contains X and 0_V and is closed under addition and scaling. Conversely, $\text{Span}(X) \subseteq \text{Span}^*(X)$ because, as can be easily checked, $\text{Span}^*(X)$ is a subspace of V containing X . ■

Note that (3.8) involves only finitely many elements of X , even if X itself is infinite (this is to be expected since there is no way to define infinite sums in an arbitrary vector space). Nevertheless, it is convenient to think of linear combinations of elements of X as sums running over the entire set:

$$\sum_{x \in X} c(x) \cdot x,$$

where all but finitely many of the coefficients $c(x)$ are zero, which, in effect, makes the sum finite. Recall from Example 2.22 that $[X \rightarrow F]^{<\infty}$ is the set of all functions $c: X \rightarrow F$ whose **support** $\text{supp}(c) := \{x \in X : c(x) \neq 0\}$ is finite. For each such $c \in [X \rightarrow F]^{<\infty}$, it makes sense to define

$$\ell_X(c) := \sum_{x \in X} c(x) \cdot x \in V.$$

This gives a linear function $\ell_X: [X \rightarrow F]^{<\infty} \rightarrow V$. With this notation, Lemma 3.7 can be stated as:

$$\text{Span}(X) = \text{im}(\ell_X).$$

Definition 3.9. Let V be a vector space over a field F . We say that a set $X \subseteq V$ is:

- **spanning** if $\text{im}(\ell_X) = \text{Span}(X) = V$, i.e., if the function ℓ_X is surjective;
- **independent** if $\ker(\ell_X) = \{0\}$, i.e., if the function ℓ_X is injective;
- a **basis** if it is both spanning and independent, i.e., if the function ℓ_X is bijective.

Remark 3.10. Explicitly, $\ker(\ell_X) = \{0\}$ means that for any $c \in [X \rightarrow F]^{<\infty}$,

$$\sum_{x \in X} c(x) \cdot x = 0 \iff c(x) = 0 \text{ for all } x \in X.$$

Remark 3.11. Since the map $\ell_X: [X \rightarrow F]^{<\infty} \rightarrow V$ is linear, if X is a basis for V , then ℓ_X is an *isomorphism* of vector spaces, and thus V is isomorphic to $[X \rightarrow F]^{<\infty}$.

Exercise 3.12. Show that a subset of an independent set is independent. Show that a superset of a spanning set is spanning.

Once we have found a basis for a vector space, we have full control over the linear functions on V , as explained by the following theorem:

Theorem 3.13 (Linear functions in terms of a basis). *Let V and W be F -vector spaces. Suppose that $B \subseteq V$ is a basis for V . Then for each function $\varphi: B \rightarrow W$, there exists a unique linear function $\hat{\varphi}: V \rightarrow W$ such that $\hat{\varphi}(x) = \varphi(x)$ for all $x \in B$, and this $\hat{\varphi}$ is given by the formula*

$$\hat{\varphi} \left(\sum_{x \in B} c(x) \cdot x \right) := \sum_{x \in B} c(x) \cdot \varphi(x). \quad (3.14)$$

PROOF. This is a theorem that “proves itself,” meaning that once we “unwrap” its statement, it becomes almost tautological. Suppose that $\hat{\varphi}: V \rightarrow W$ is a linear map that extends φ . Since B is

spanning, every element of V can be written as a linear combination of elements of B , i.e., in the form $\sum_{x \in B} c(x) \cdot x$ for some $c \in [X \rightarrow F]^{<\infty}$. Since $\hat{\varphi}$ is linear, it respects addition and scaling, so

$$\hat{\varphi} \left(\sum_{x \in B} c(x) \cdot x \right) = \sum_{x \in B} \hat{\varphi}(c(x) \cdot x) = \sum_{x \in B} c(x) \cdot \hat{\varphi}(x) = \sum_{x \in B} c(x) \cdot \varphi(x).$$

This shows that $\hat{\varphi}$ *must* be defined by (3.14), which proves its uniqueness. Furthermore, since B is independent, every element of V can be expressed in the form $\sum_{x \in B} c(x) \cdot x$ in *only one* way, and thus (3.14) is a valid definition. It remains to verify that the function $\hat{\varphi}$ given by (3.14) is linear, but that is a straightforward [exercise](#). ■

3.C. The first fundamental theorem of linear algebra and its ramifications

Theorem 3.15 (First fundamental theorem). *Every vector space has a basis.*

Moreover, if V is an F -vector space, $I \subseteq V$ is an independent set, $S \subseteq V$ is a spanning set, and $I \subseteq S$, then there exists a basis B such that $I \subseteq B \subseteq S$.

We will prove Theorem 3.15 in the next subsection. For now, let us consider some of its consequences. For instance, we can now easily deduce Theorem 3.1:

Theorem 3.1. *Let V be a vector space over a field F and let W be a subspace of V . Then there exist linear functions*

$$\varphi: V \rightarrow V \quad \text{and} \quad \psi: V \rightarrow V$$

such that $W = \text{im}(\varphi) = \text{ker}(\psi)$.

PROOF. By Theorem 3.15, W has a basis B_W . Since B_W is, by definition, an independent set, we may apply the “moreover” part of Theorem 3.15 with $I = B_W$ and $S = V$ to obtain a basis B for V such that $B_W \subseteq B$. By Theorem 3.13, there exist linear functions $\varphi: V \rightarrow V$ and $\psi: V \rightarrow V$ such that for all $x \in B$,

$$\varphi(x) = \begin{cases} x & \text{if } x \in B_W; \\ 0 & \text{if } x \in B \setminus B_W, \end{cases} \quad \text{and} \quad \psi(x) = \begin{cases} 0 & \text{if } x \in B_W; \\ x & \text{if } x \in B \setminus B_W. \end{cases}$$

We claim that $W = \text{im}(\varphi) = \text{ker}(\psi)$, as desired. Indeed, consider any vector $v \in V$. Since B is a basis for V , there is a unique way to express v as a linear combination of the elements of B :

$$v = \sum_{x \in B} c(x) \cdot x,$$

where $c \in [B \rightarrow F]^{<\infty}$. Separating the terms corresponding to the basis vectors in B_W and in $B \setminus B_W$, we can write

$$v = \sum_{x \in B_W} c(x) \cdot x + \sum_{x \in B \setminus B_W} c(x) \cdot x.$$

Then

$$\varphi(v) = \sum_{x \in B_W} c(x) \cdot \varphi(x) + \sum_{x \in B \setminus B_W} c(x) \cdot \varphi(x) = \sum_{x \in B_W} c(x) \cdot x. \quad (3.16)$$

This shows that $\varphi(v)$ is a linear combination of elements of B_W , and hence $\varphi(v) \in W$ and $\text{im}(\varphi) \subseteq W$. On the other hand, *any* linear combination of the elements of B_W can appear as the last expression in (3.16), and, since B_W spans W , this means that $W \subseteq \text{im}(\varphi)$.

Exercise 3.17. Show that $\varphi(w) = w$ for all $w \in W$.

Proving that $W = \text{ker}(\psi)$ is left as an [exercise](#). ■

Theorem 3.15 is much less obvious than might seem at first. We are used to thinking about vector spaces such as \mathbb{R}^n (over \mathbb{R}), where a basis is easy to find:

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

On the other hand, consider the space $\mathbb{R}^{\mathbb{N}}$ of all infinite sequences of real numbers. It is tempting to guess that the following infinite set should be a basis for this space:

$$\begin{aligned} e_0 &:= (1, 0, 0, 0, \dots), \\ e_1 &:= (0, 1, 0, 0, \dots), \\ e_2 &:= (0, 0, 1, 0, \dots), \\ &\dots \end{aligned}$$

However, this guess is wrong! Indeed, the sequence

$$(1, 1, 1, 1, \dots),$$

all of whose entries are equal to 1, is not in the span of $\{e_0, e_1, e_2, \dots\}$ (because linear combinations only involve finite sums). In fact, the span of $\{e_0, e_1, e_2, \dots\}$ is the space $[\mathbb{N} \rightarrow \mathbb{R}]^{<\infty}$ of all sequences with finite support (i.e., with only finitely many nonzero entries). Actually, any basis for $\mathbb{R}^{\mathbb{N}}$ is much larger than the set $\{e_0, e_1, e_2, \dots\}$ —it is necessarily *uncountable*.¹⁰

The following corollary is another indication of how surprising Theorem 3.15 is:

Corollary 3.18 (to Theorem 3.15). *Every F -vector space is isomorphic to a space of the form $[X \rightarrow F]^{<\infty}$ for some set X .*

PROOF. See Remark 3.11. ■

So, for example, \mathbb{R} , viewed as a vector space over \mathbb{Q} , is isomorphic to $[X \rightarrow \mathbb{Q}]^{<\infty}$ for some set X . In fact, one can show that \mathbb{R} is isomorphic to $[\mathbb{R} \rightarrow \mathbb{Q}]^{<\infty}$. What's more, \mathbb{R}^2 , viewed as a vector space over \mathbb{Q} , is *also* isomorphic to $[\mathbb{R} \rightarrow \mathbb{Q}]^{<\infty}$, and hence, \mathbb{R}^2 and \mathbb{R} are isomorphic as \mathbb{Q} -vector spaces!¹¹ In other words, there exists a \mathbb{Q} -linear bijection $f: \mathbb{R}^2 \rightarrow \mathbb{R}$. Let us split this statement into two parts:

- there is a *bijection* $f: \mathbb{R}^2 \rightarrow \mathbb{R}$;
- such a bijection can be made *\mathbb{Q} -linear*.

Without the second part, constructing an arbitrary bijection $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ is actually not difficult. Below we describe an injection $\mathbb{R}^2 \rightarrow \mathbb{R}$; making it into a bijection is left as an exercise.¹²

Lemma 3.19. *There exists an injective function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$.*

PROOF SKETCH. We have to describe a way to “encode” a pair of real numbers (a, b) into a single real number c so that a and b can be recovered from c uniquely. One way to achieve this is to write a and b in decimal, adding leading zeros if necessary to ensure that the integer parts of a and b are of the same length, and then assemble c as follows:

$$c := \begin{pmatrix} 1 & \text{if } a \geq 0 \\ 2 & \text{if } a < 0 \end{pmatrix} \begin{pmatrix} \text{integer} \\ \text{part of } a \end{pmatrix} \begin{pmatrix} 1 & \text{if } b \geq 0 \\ 2 & \text{if } b < 0 \end{pmatrix} \begin{pmatrix} \text{integer} \\ \text{part of } b \end{pmatrix} \cdot \begin{pmatrix} \text{intersperse} \\ \text{the digits of the} \\ \text{fractional parts} \\ \text{of } a \text{ and } b \end{pmatrix}.$$

So, for example, if $a = 314.1592\dots$ and $b = -1.2345\dots$, then $c = 13142001.12539425\dots$ ■

¹⁰A set X is **uncountable** if there is no surjection $\mathbb{N} \rightarrow X$, in other words, if there is no way to list all the elements of X in a sequence x_0, x_1, x_2, \dots .

¹¹They are *not* isomorphic as \mathbb{R} -vector spaces though.

¹²Although one might say that an injection from \mathbb{R}^2 to \mathbb{R} that also leaves some of the elements of \mathbb{R} uncovered is even more counter-intuitive than a bijection!

Exercise 3.20. Describe a bijective function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$.

It turns out that there is no “explicit” way, like in the above proof of Lemma 3.19, to describe a \mathbb{Q} -linear injection $\mathbb{R}^2 \rightarrow \mathbb{R}$. It follows from results in the area of mathematics called **descriptive set theory** that every \mathbb{Q} -linear function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ that one can “explicitly write down” (the technical term is “Borel”) must be continuous. (This phenomenon is known as *automatic continuity*.) And it is fairly easy to see that a continuous function $\mathbb{R}^2 \rightarrow \mathbb{R}$ cannot be injective.

How can this be? How can we prove that a \mathbb{Q} -linear bijection $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ exists without being able to explicitly describe it? This apparent paradox is a consequence of the fact that the proof of Theorem 3.15 relies on the so-called *Axiom of Choice*. It is one of the generally accepted axioms of set theory that form the foundation of mathematics. While most other axioms assert the existence of “concrete” sets, such as the empty set or the powerset $\mathcal{P}(X)$ of a given set X , the Axiom of Choice postulates the existence of a certain function without explicitly stating what it is:

Axiom 3.21 (Axiom of Choice). *Let \mathcal{F} be a set of nonempty sets. Then there exists a function $\text{ch}: \mathcal{F} \rightarrow \bigcup \mathcal{F}$ such that for all $X \in \mathcal{F}$, we have $\text{ch}(X) \in X$.*

The function ch in the Axiom of Choice is called a **choice function**, because it “chooses” one element $\text{ch}(X)$ from each set $X \in \mathcal{F}$. Notice that the Axiom of Choice does not specify how the chosen element $\text{ch}(X)$ is determined; it merely claims that *some* choice is possible.

Finally, let us point out that Theorem 3.15 crucially relies on the fact that F is a field. If we replace the field F by a commutative ring R in the definition of a vector space, we obtain a structure called a **module** over R (so, a vector space is a module over a field). The definition of a basis makes sense for modules as well as for vector spaces, but if R is not a field, then a module over R may not have a basis. Here’s a simple example:

Example 3.22. Let n be an integer ≥ 2 . The set \mathbb{Z}_n of residues modulo n is naturally a module over \mathbb{Z} , but this module does not have a basis. Indeed, for every $x \in \mathbb{Z}_n$, we have

$$n \cdot x = 0 \pmod{n},$$

so the set $\{x\}$ is not independent. Hence, the only independent set in \mathbb{Z}_n is \emptyset , and it is certainly not spanning. Note, however, that if n is prime, then $\mathbb{Z}_n = \mathbb{F}_n$ *does* have a basis as a vector space over \mathbb{F}_n (any one-element set $\{x\}$ with $x \neq 0$ is a basis).

3.D. Proof of Theorem 3.15

Lemma 3.23. *Let V be a space over a field F . Let $I \subseteq V$ be an independent set and let $y \in V \setminus I$. The following statements are equivalent:*

- (1) $I \cup \{y\}$ is not an independent set;
- (2) $y \in \text{Span}(I)$.

PROOF. (1) \implies (2). If the set $I \cup \{y\}$ is not independent, then we can write

$$a_0 y + a_1 x_1 + \cdots + a_n x_n = 0,$$

where $x_1, \dots, x_n \in I$ and not all of the coefficients a_0, a_1, \dots, a_n are zero. Since I is independent, we must have $a_0 \neq 0$; thus,

$$y = -\frac{a_1}{a_0} x_1 - \cdots - \frac{a_n}{a_0} x_n \in \text{Span}(I).$$

(2) \implies (1). Since $y \in \text{Span}(I)$, y can be expressed as a linear combination of elements of I :

$$y = a_1 x_1 + \cdots + a_n x_n.$$

But then y can be expressed in two *distinct* ways as a linear combination of elements of $I \cup \{y\}$, meaning that the set $I \cup \{y\}$ is not independent. ■

Lemma 3.23 gives a convenient criterion for when a given independent set is a basis. Say that an independent set $I \subseteq V$ is **maximal** if there is no independent set $J \subseteq V$ with $J \supsetneq I$. Similarly, a spanning set S is **minimal** if there is no spanning set T such that $T \subsetneq S$.

Lemma 3.24. *Let V be an F -vector space and let $X \subseteq V$. The following statements are equivalent:*

- (1) X is a basis;
- (2) X is a maximal independent set;
- (3) X is a minimal spanning set.

PROOF. We will prove the equivalence (1) \iff (2), while (1) \iff (3) is left as an **exercise**.

(1) \implies (2). Assume that X is a basis. Consider any $y \in V \setminus X$. Since X is spanning, $y \in \text{Span}(X)$. By Lemma 3.23, the set $X \cup \{y\}$ is not independent, and hence X is a maximal independent set.

(2) \implies (1). Assume that X is a maximal independent set. We need to argue that X is spanning. To that end, consider any $y \in V$. If $y \in X$, then $y \in \text{Span}(X)$ by definition, so assume $y \notin X$. But then the set $X \cup \{y\}$ is not independent, which, by Lemma 3.23, means that $y \in \text{Span}(X)$. ■

Definition 3.25. We say that sets A and B are **comparable** if $A \subseteq B$ or $B \subseteq A$. A **chain** is a set \mathcal{C} of pairwise comparable sets.

Recall that $\bigcup \mathcal{C}$ denotes the union of all the sets in \mathcal{C} ; i.e.,

$$\bigcup \mathcal{C} := \{x : x \in A \text{ for some } A \in \mathcal{C}\}.$$

Example 3.26. The set $\{\{0\}, \{0, 1\}, \{0, 1, 2\}, \dots\}$ is a chain. The union of this chain is \mathbb{N} .

Example 3.27. More generally, if A_0, A_1, A_2, \dots are sets such that

$$A_0 \subset A_1 \subset A_2 \subset \dots,$$

then the set $\{A_0, A_1, A_2, \dots\}$ is a chain, whose union is $A_0 \cup A_1 \cup A_2 \cup \dots$. Similarly, if

$$A_0 \supset A_1 \supset A_2 \supset \dots,$$

then $\{A_0, A_1, A_2, \dots\}$ is a chain, whose union is A_0 .

Example 3.28. The set $\{(-\infty; \alpha) : \alpha \in \mathbb{R}\}$ is a chain. The union of this chain is \mathbb{R} .

Observe that if \mathcal{C} is a chain and $A_1, \dots, A_n \in \mathcal{C}$ (when n is finite), then there is an index i such that $A_i = A_1 \cup \dots \cup A_n$.

Lemma 3.29. *Let V be a vector space over a field F . If \mathcal{C} is a chain of independent subsets of V , then the set $\bigcup \mathcal{C}$ is also independent.*

PROOF. Suppose, towards a contradiction, that \mathcal{C} is not independent. This means that there exist some $x_1, \dots, x_n \in \bigcup \mathcal{C}$ and nonzero $a_1, \dots, a_n \in F$ such that

$$a_1x_1 + \dots + a_nx_n = 0.$$

Since $x_1, \dots, x_n \in \bigcup \mathcal{C}$, there are sets $A_1, \dots, A_n \in \mathcal{C}$ such that $x_1 \in A_1, \dots, x_n \in A_n$. But since \mathcal{C} is a chain, there is an index i such that $A_i = A_1 \cup \dots \cup A_n$. Therefore, $x_1, \dots, x_n \in A_i$, and hence A_i is not independent. This is a contradiction. ■

Remark 3.30. It is important in Lemma 3.29 to assume that \mathcal{C} is a chain, since a union of independent sets is not, in general, independent. For instance, the sets

$$\{(1, 1)\} \quad \text{and} \quad \{(2, 2)\}$$

are independent in \mathbb{R}^2 , but their union $\{(1, 1), (2, 2)\}$ is not.

With Lemma 3.29 in hand, we can deduce Theorem 3.15 from the following general fact, known as **Zorn's lemma**:

Theorem 3.31 (Zorn’s lemma). *Let \mathcal{F} be a family of sets with the following properties:*

- (Z1) $\emptyset \in \mathcal{F}$;
- (Z2) if $A \in \mathcal{F}$ and $B \subseteq A$, then $B \in \mathcal{F}$;
- (Z3) if $\mathcal{C} \subseteq \mathcal{F}$ is a chain, then $\bigcup \mathcal{C} \in \mathcal{F}$.

Then \mathcal{F} has a **maximal element**; i.e., there is a set $A \in \mathcal{F}$ such that there is no $B \in \mathcal{F}$ with $B \supsetneq A$.

Remark 3.32. Theorem 3.31 is not the most general form of Zorn’s lemma. In particular, assumption (Z2) can be removed, while (Z1) can be weakened to $\mathcal{F} \neq \emptyset$ (see Theorem 4.3).

PROOF THAT EVERY VECTOR SPACE HAS A BASIS. Let V be an F -vector space and let \mathcal{F} be the set of all independent subsets of V . Then \mathcal{F} satisfies the assumptions of Zorn’s lemma (where (Z3) is given by Lemma 3.29), and hence \mathcal{F} has a maximal element. By Lemma 3.24, a maximal independent set is a basis, and hence we are done. ■

Exercise 3.33. Prove the “moreover” part of Theorem 3.15. *Hint:* Consider the family \mathcal{F} of all sets $X \subseteq S \setminus I$ such that $I \cup X$ is independent.

What remains is to prove Theorem 3.31; this will be done in the next subsection.

3.E. Proof of Zorn’s lemma

The Axiom of Choice is obviously true, the well-ordering principle obviously false, and who can tell about Zorn’s lemma?

Jerry L. Bona

This is likely the most challenging proof in these notes, but coming to grips with its logic can be immensely beneficial, as it involves several fundamental ideas that play an important role throughout mathematics.

We argue by contradiction. Let \mathcal{F} be a family of sets satisfying (Z1), (Z2), and (Z3) and assume that \mathcal{F} has no maximal element. In other words, for each $A \in \mathcal{F}$, there is some $B \in \mathcal{F}$ such that $A \subsetneq B$. Due to (Z2), we may in fact assume that $B = A \cup \{x\}$ for a single element $x \notin A$. (This is the only reason for including (Z2) in the list of assumptions.)

3.E.1. The plan of attack.—Before we proceed to the proof, let’s try to build some intuition about the structure of \mathcal{F} . By (Z1), $\emptyset \in \mathcal{F}$. Since \emptyset is not a maximal element of \mathcal{F} , there is some a_0 such that $\{a_0\} \in \mathcal{F}$. But since $\{a_0\}$ is also not maximal, there is some $a_1 \neq a_0$ such that $\{a_0, a_1\} \in \mathcal{F}$. Repeating this argument, we obtain an infinite sequence of sets

$$\emptyset \subset \{a_0\} \subset \{a_0, a_1\} \subset \{a_0, a_1, a_2\} \subset \dots,$$

all of which belong to \mathcal{F} . Now notice that

$$\{\emptyset, \{a_0\}, \{a_0, a_1\}, \{a_0, a_1, a_2\}, \dots\}$$

is a chain, so we can apply (Z3) and conclude that the union of this chain is in \mathcal{F} ; that is, we have $\{a_0, a_1, a_2, \dots\} \in \mathcal{F}$. Thus, we’ve found an *infinite* set in \mathcal{F} . But that’s not all. The set $\{a_0, a_1, \dots\}$ is *also* not maximal—hence, there is some b_0 such that $\{a_0, a_1, \dots\} \cup \{b_0\} \in \mathcal{F}$. As before, we can repeat this argument to obtain an infinite sequence of sets

$$\{a_0, a_1, \dots\} \subset \{a_0, a_1, \dots\} \cup \{b_0\} \subset \{a_0, a_1, \dots\} \cup \{b_0, b_1\} \subset \{a_0, a_1, \dots\} \cup \{b_0, b_1, b_2\} \subset \dots,$$

and then use (Z3) to conclude that $\{a_0, a_1, \dots\} \cup \{b_0, b_1, \dots\} \in \mathcal{F}$. In other words, we are able to add infinitely many new elements to any set in \mathcal{F} . Iterating this construction, we now get a sequence

$$\{a_0, a_1, \dots\} \subset \{a_0, a_1, \dots\} \cup \{b_0, b_1, \dots\} \subset \{a_0, a_1, \dots\} \cup \{b_0, b_1, \dots\} \cup \{c_0, c_1, \dots\} \subset \dots$$

But this sequence is again a chain, so we can apply (Z3) to infer that

$$\{a_0, a_1, \dots\} \cup \{b_0, b_1, \dots\} \cup \{c_0, c_1, \dots\} \cup \dots \in \mathcal{F}.$$

And so on—this process can be continued indefinitely.

Our strategy now is to identify the “final stage” of this process: the collection of *all* sets that will ever appear in this construction. We will denote this collection \mathcal{T}_0 (this choice of notation will make sense soon). So, \mathcal{T}_0 should look like this:

$$\begin{aligned} \mathcal{T}_0 = & \{ \emptyset, \{a_0\}, \{a_0, a_1\}, \{a_0, a_1, a_2\}, \dots \dots \dots \\ & \{a_0, a_1, \dots\}, \{a_0, a_1, \dots\} \cup \{b_0\}, \{a_0, a_1, \dots\} \cup \{b_0, b_1\}, \{a_0, a_1, \dots\} \cup \{b_0, b_1, b_2\}, \dots \dots \dots \\ & \{a_0, a_1, \dots\} \cup \{b_0, b_1, \dots\}, \{a_0, a_1, \dots\} \cup \{b_0, b_1, \dots\} \cup \{c_0\}, \dots \dots \dots \\ & \{a_0, a_1, \dots\} \cup \{b_0, b_1, \dots\} \cup \{c_0, c_1, \dots\}, \dots \dots \dots \\ & \dots \dots \dots \\ & \{a_0, a_1, \dots\} \cup \{b_0, b_1, \dots\} \cup \{c_0, c_1, \dots\} \cup \dots, \dots \dots \dots \\ & \dots \dots \dots \}. \end{aligned}$$

By definition, for every set $A \in \mathcal{T}_0$, \mathcal{T}_0 *also* contains a set $A \cup \{x\}$ with $x \notin A$. As well, for every chain $\mathcal{C} \subseteq \mathcal{T}_0$, the union $\bigcup \mathcal{C}$ of this chain is in \mathcal{T}_0 . And here’s the punchline:

The set \mathcal{T}_0 itself is a chain!

This means that $\bigcup \mathcal{T}_0 \in \mathcal{T}_0$. But then there is some *new* element $x \notin \bigcup \mathcal{T}_0$ such that $(\bigcup \mathcal{T}_0) \cup \{x\} \in \mathcal{T}_0$, which is, of course, impossible.

This is all nice and well, I hear you say, but what *is* this \mathcal{T}_0 , exactly? “All the sets that will ever appear in this construction” is far from a precise definition. For that matter, how is “this construction” defined? To address these questions, and make the above intuition into a rigorous proof, note that on each step of the construction we do one of the following two things:

- (i) either we add a new element to a set that has already been constructed;
- (ii) or we take the union of a chain of constructed sets.

Thus, we could say that \mathcal{T}_0 is “the set of all sets that can be built from \emptyset by repeatedly applying operations (i) and (ii).” However, this is also not fully satisfactory, since it is not clear what “repeatedly” means here (you should keep in mind that we have to repeat (i) and (ii) *infinitely many times*). What we will actually do is let \mathcal{T}_0 be the *smallest set that contains \emptyset and is closed under operations (i) and (ii)*. (This is reminiscent of the definition of the span of a subset $X \subseteq V$: It can be both defined as the set of all vectors that can be obtained from X by taking linear combinations, and as the smallest subspace of V containing X .)

3.E.2. Towers and the Induction Principle.—Recall that, since \mathcal{F} has no maximal element, for each $A \in \mathcal{F}$, there is some $x \notin A$ such that $A \cup \{x\} \in \mathcal{F}$. Pick one such x and denote it $f(A)$. Define

$$A' := A \cup \{f(A)\}.$$

By definition, $A' \in \mathcal{F}$, $A \subset A'$, and A' contains precisely one element that is not in A , namely $f(A)$. We call A' the **successor** of A . (This is where we use the Axiom of Choice.)

Call a subset $\mathcal{T} \subseteq \mathcal{F}$ a **tower** if it has the following properties:

- (T1) $\emptyset \in \mathcal{T}$;
- (T2) if $A \in \mathcal{T}$, then $A' \in \mathcal{T}$;
- (T3) if $\mathcal{C} \subseteq \mathcal{T}$ is a chain, then $\bigcup \mathcal{C} \in \mathcal{T}$.

Note that there is at least one tower, namely \mathcal{F} . Let \mathcal{T}_0 be the intersection of all towers; that is, $A \in \mathcal{T}_0$ if and only if $A \in \mathcal{T}$ for every tower \mathcal{T} .

Exercise 3.34. Show that \mathcal{T}_0 is a tower.

Thus, \mathcal{T}_0 is the smallest tower.¹³ By definition, \mathcal{T}_0 contains those and only those sets that must belong to every tower; morally speaking, this should mean that \mathcal{T}_0 is the set of all sets that can be obtained from \emptyset by repeatedly taking successors and unions of chains. To vindicate this intuition, we have to prove that \mathcal{T}_0 is a chain.

Claim. *If we can show that \mathcal{T}_0 is a chain, then we can finish the proof of Zorn's lemma.*

PROOF. If \mathcal{T}_0 is a chain, then, applying (T3) with \mathcal{T}_0 in place of \mathcal{C} , we get $\bigcup \mathcal{T}_0 \in \mathcal{T}_0$. By (T2), $(\bigcup \mathcal{T}_0)' \in \mathcal{T}_0$ as well, and thus $f(\bigcup \mathcal{T}_0) \in (\bigcup \mathcal{T}_0)' \subseteq \bigcup \mathcal{T}_0$. But $f(\bigcup \mathcal{T}_0) \notin \bigcup \mathcal{T}_0$ by definition, and this contradiction completes the proof of Zorn's lemma. ■

If P is a property of sets, then we write $P(A)$ to mean that the set A has P .

Lemma 3.35 (Induction Principle). *Let P be a property of sets. Suppose that:*

- (I1) $P(\emptyset)$;
- (I2) for all $A \in \mathcal{T}_0$, if $P(A)$, then $P(A')$;
- (I3) if $\mathcal{C} \subseteq \mathcal{T}_0$ is a chain of sets that have P , then $P(\bigcup \mathcal{C})$.

Then $P(A)$ for all $A \in \mathcal{T}_0$.

PROOF. The assumptions of the lemma mean that the set

$$\mathcal{U} := \{A \in \mathcal{T}_0 : P(A)\}$$

is a tower contained in \mathcal{T}_0 . But \mathcal{T}_0 is the smallest tower, and hence $\mathcal{U} = \mathcal{T}_0$. ■

The name “Induction Principle” is due to the analogy between Lemma 3.35 and the principle of mathematical induction, which says that if P is a property of natural numbers such that:

- (1) $P(0)$; and
- (2) for all $n \in \mathbb{N}$, if $P(n)$, then $P(n + 1)$,

then $P(n)$ for all $n \in \mathbb{N}$. Condition (I1) is analogous to (1), while (I2) plays the role of the induction step (2). However, in Lemma 3.35, there is one more assumption, namely (I3), which handles the induction after “more than \mathbb{N} steps.”

3.E.3. *The proof.*—We say that a set $A \in \mathcal{T}_0$ is \mathcal{T}_0 -**comparable** if A is comparable with every $B \in \mathcal{T}_0$. Our goal is to show that *every* set $A \in \mathcal{T}_0$ is \mathcal{T}_0 -comparable, since this would mean that \mathcal{T}_0 is a chain. To this end, we will use the Induction Principle.

- (I1) The empty set \emptyset is \mathcal{T}_0 -comparable, since $\emptyset \subseteq B$ for all $B \in \mathcal{T}_0$.
- (I2) This step is somewhat complicated, and we will come back to it later.
- (I3) Suppose that $\mathcal{C} \subseteq \mathcal{T}_0$ is a chain of \mathcal{T}_0 -comparable sets. Take any $B \in \mathcal{C}$. We have to show that

$$B \subseteq \bigcup \mathcal{C} \quad \text{or} \quad \bigcup \mathcal{C} \subseteq B.$$

There are two cases to consider.

Case 1: *There is some $A \in \mathcal{C}$ such that $B \subseteq A$.* In this case $B \subseteq \bigcup \mathcal{C}$.

Case 2: *There is no $A \in \mathcal{C}$ such that $B \subseteq A$.* Since every $A \in \mathcal{C}$ is \mathcal{T}_0 -comparable, and hence comparable with B , this means that for all $A \in \mathcal{C}$, we have $A \subseteq B$. But then $\bigcup \mathcal{C} \subseteq B$ as well.

Now let's return to (I2). Let $A \in \mathcal{T}_0$ be \mathcal{T}_0 -comparable. We need to show that A' is \mathcal{T}_0 -comparable as well. To that end, take any $B \in \mathcal{T}_0$. Our goal is to prove that $B \subseteq A'$ or $A' \subseteq B$.

Case 1: $B \subseteq A$. Then $B \subseteq A'$ as well.

Case 2: $B \not\subseteq A$. Since A is \mathcal{T}_0 -comparable, this means that $A \not\subseteq B$. In this case, we want to conclude that $A' \subseteq B$. To achieve this, we will prove the following lemma:

¹³Perhaps calling it the *slimmest* tower would be even more accurate.

Lemma 3.36. For $A \in \mathcal{T}_0$, define

$$\mathcal{U}_A := \{B \in \mathcal{T}_0 : B \subseteq A \text{ or } A' \subseteq B\}.$$

If A is \mathcal{T}_0 -comparable, then $\mathcal{U}_A = \mathcal{T}_0$.

Once Lemma 3.36 is established, we may conclude that if $B \not\subseteq A$, then $A' \subseteq B$, as desired.

It now remains to prove Lemma 3.36.

PROOF OF LEMMA 3.36. The proof uses the Induction Principle again.

Exercise 3.37. Show that (I1) and (I3) hold; that is, prove that:

- $\emptyset \in \mathcal{U}_A$;
- if $\mathcal{C} \subseteq \mathcal{U}_A$ is a chain, then $\bigcup \mathcal{C} \in \mathcal{U}_A$.

To verify (I2), let $B \in \mathcal{U}_A$. We have to show that $B' \in \mathcal{U}_A$. Since $B \in \mathcal{U}_A$, there are three cases:

Case 1: $A' \subseteq B$. Then $A' \subseteq B'$, and hence $B' \in \mathcal{U}_A$.


Case 2: $B = A$. Then $B' = A'$, and hence $B' \in \mathcal{U}_A$.

Case 3: $B \subsetneq A$. This is the most interesting case, and the crux of the entire proof of Zorn's lemma. We wish to show that $B' \subseteq A$. Suppose, towards a contradiction, that $B' \not\subseteq A$. Since A is \mathcal{T}_0 -comparable, it is, in particular, comparable with B' . Thus, if $B' \not\subseteq A$, then $A \subsetneq B'$. This means that B' contains all the elements of $A \setminus B$ (of which there is at least one since $B \subsetneq A$) plus *also* at least one element not in A . Therefore, $|B' \setminus B| \geq 2$. But, by definition, the only element of B' that is not in B is $f(B)$, i.e., $|B' \setminus B| = 1$. This contradiction completes the proof. ■

3.F. Applying Theorem 3.15

We already saw a corollary of Theorem 3.15, namely Theorem 3.1. Another immediate consequence of Theorem 3.15 is the existence of a large family of functions $\mathbb{R} \rightarrow \mathbb{R}$ that are \mathbb{Q} -linear but not \mathbb{R} -linear. Indeed, let B be a basis for \mathbb{R} as a \mathbb{Q} -vector space. Then every function $\varphi: B \rightarrow \mathbb{R}$ has a (unique) \mathbb{Q} -linear extension $\hat{\varphi}: \mathbb{R} \rightarrow \mathbb{R}$; but unless there is a real number $c \in \mathbb{R}$ such that $\varphi(x) = cx$ for all $x \in B$, the function $\hat{\varphi}$ is not going to be \mathbb{R} -linear.

Exercise 3.38 (\mathbb{Q} -linear functions are weird). A set $S \subseteq \mathbb{R}^2$ is **dense** in \mathbb{R}^2 if for every point $p \in \mathbb{R}^2$ and for every real $\varepsilon > 0$, there is a point $q \in S$ such that the distance between p and q is less than ε . In other words, S is dense in \mathbb{R}^2 if S intersects every disc $D \subset \mathbb{R}^2$ of positive radius:



If we were to draw a picture of a dense subset of the plane giving each point an arbitrarily small positive thickness, it would look like this:



Suppose that $f: \mathbb{R} \rightarrow \mathbb{R}$ is a function that is \mathbb{Q} -linear but not \mathbb{R} -linear. Show that the graph of f is dense in \mathbb{R}^2 . *Hint.* What is $\text{Span}_{\mathbb{R}}(\Gamma_f)$?

With Theorem 3.15 in hand, we can now solve Problem 1.1. The proof given below uses a basis for \mathbb{R} as a \mathbb{Q} -vector space to reduce a question about real numbers to a question about finite fields.

Theorem 3.39. *Let $0 \leq a_1 < \dots < a_n$ be distinct integers and suppose that $f: \mathbb{Z} \rightarrow \mathbb{R}$ is a function such that for all $k \in \mathbb{Z}$ and $\ell \in \mathbb{Z}^+$, we have*

$$f(k + a_1\ell) + f(k + a_2\ell) + \dots + f(k + a_n\ell) = 0. \quad (3.40)$$

Then $f(m) = 0$ for all $m \in \mathbb{Z}$.

PROOF. This elegant argument is due to the user **grobber** on *The Art of Problem Solving*.

Suppose that f satisfies (3.40). Plugging in $k = m - a_n$ and $\ell = 1$, we get

$$f(m) = -f(m - a_n + a_1) - f(m - a_n + a_2) - \dots - f(m - a_n + a_{n-1}). \quad (3.41)$$

The main consequence of (3.41) is that if we know the values

$$f(m - a_n), \quad f(m - a_n + 1), \quad \dots, \quad f(m - 1),$$

then we can apply (3.41) repeatedly to compute the values $f(m)$, $f(m + 1)$, $f(m + 2)$, and so on. Similarly, plugging in $k = m - a_1$ and $\ell = 1$ gives

$$f(m) = -f(m - a_1 + a_2) - f(m - a_1 + a_3) - \dots - f(m - a_1 + a_n),$$

which means that knowing the values

$$f(m + 1), \quad f(m + 2), \quad \dots, \quad f(m + a_n)$$

is enough to also compute the values $f(m)$, $f(m - 1)$, $f(m - 2)$, &tc. To summarize, f is completely determined by its values at any a_n consecutive integers.

After these preliminary observations, we proceed in four steps.

STEP 1: *Let p be a prime number $> n$. If $f: \mathbb{Z} \rightarrow \mathbb{F}_p$ satisfies (3.40), then $f(m) = 0$ for all m .*

Proof. Indeed, let p be a prime number $> n$ and suppose that $f: \mathbb{Z} \rightarrow \mathbb{F}_p$ satisfies (3.40). There are only finitely many (namely p^{a_n}) distinct sequences of elements of \mathbb{F}_p of length a_n , while there are infinitely many integers. Therefore, for some $i < j$, the sequences

$$(f(i + 1), f(i + 2), \dots, f(i + a_n)) \quad \text{and} \quad (f(j + 1), f(j + 2), \dots, f(j + a_n))$$

coincide. But then $f(m) = f(m + j - i)$ for all $m \in \mathbb{Z}$, i.e., f is periodic with period $t := j - i$. Applying (3.40) with $k = m$ and $\ell = t$ gives

$$0 = f(m + a_1t) + \dots + f(m + a_nt) = n \cdot f(m).$$

Since $p > n$, this implies $f(m) = 0$, as desired. \dashv

STEP 2: *If $f: \mathbb{Z} \rightarrow \mathbb{Z}$ satisfies (3.40), then $f(m) = 0$ for all m .*

Proof. For a prime number $p > n$, define $f_p: \mathbb{Z} \rightarrow \mathbb{F}_p$ by

$$f_p(m) := f(m) \pmod{p}.$$

By Step 1, for every $m \in \mathbb{Z}$, we have $f_p(m) = 0$, i.e., $f(m)$ is divisible by p . The only integer that is divisible by every prime number bigger than n is 0, so $f(m) = 0$, as claimed. \dashv

STEP 3: *If $f: \mathbb{Z} \rightarrow \mathbb{Q}$ satisfies (3.40), then $f(m) = 0$ for all m .*

Proof. Let d be a common denominator of $f(1)$, $f(2)$, \dots , $f(a_n)$, and let $g(m) := df(m)$. Then $g(1)$, $g(2)$, \dots , $g(a_n)$ are integers. Moreover, g still satisfies (3.40). From (3.41), we conclude that $g(m)$ is an integer for all $m \in \mathbb{Z}$, and, by Step 2, $g(m) = 0$, and hence $f(m) = 0$, for all $m \in \mathbb{Z}$. \dashv

Finally, consider a function $f: \mathbb{Z} \rightarrow \mathbb{R}$ that satisfies (3.40). Let $B \subset \mathbb{R}$ be a basis for \mathbb{R} as a vector space over \mathbb{Q} . For each $m \in \mathbb{Z}$ and $x \in B$, let $f_x(m)$ be the coefficient of x in the representation of $f(m)$ as a linear combination of the elements of B ; in other words, write

$$f(m) = \sum_{x \in B} f_x(m) \cdot x,$$

where the values $f_x(m)$, $x \in B$, are rational numbers, only finitely many of which are nonzero. We claim that for each $x \in B$, the function $f_x: \mathbb{Z} \rightarrow \mathbb{Q}$ satisfies (3.40). Indeed, we have

$$\begin{aligned} f(k + a_1\ell) + \cdots + f(k + a_n\ell) &= \sum_{x \in B} f_x(k + a_1\ell) \cdot x + \cdots + \sum_{x \in B} f_x(k + a_n\ell) \cdot x \\ &= \sum_{x \in B} (f_x(k + a_1\ell) + \cdots + f_x(k + a_n\ell)) \cdot x = 0, \end{aligned}$$

and, since B is a basis, this is only possible when

$$f_x(k + a_1\ell) + \cdots + f_x(k + a_n\ell) = 0 \quad \text{for all } x \in B.$$

Hence, by Step 3, $f_x(m) = 0$ for all m . But then $f(m) = \sum_{x \in B} 0 \cdot x = 0$ as well, and we are done. ■

Extra exercises for Section 3

Exercise 3.42. Let V be a vector space over a field F and let $W \subseteq V$ be a subspace of V .

- Show that there is a subspace $W' \subseteq V$ such that every vector $v \in V$ can be *uniquely* expressed as a sum $v = w + w'$ with $w \in W$ and $w' \in W'$.
- Show that every subspace $W' \subseteq V$ as in (a) is isomorphic to V/W .
- Conclude that V is isomorphic to $W \oplus (V/W)$.

Exercise 3.43. Consider the \mathbb{R} -vector space $\mathbb{R}^{\mathbb{N}}$ of all infinite sequences of reals. For each $\alpha \in \mathbb{R}$, let

$$e_\alpha := (1, \alpha, \alpha^2, \alpha^3, \dots).$$

Show that the set $\{e_\alpha : \alpha \in \mathbb{R}\}$ is independent. This means that you can find as many independent vectors in $\mathbb{R}^{\mathbb{N}}$ as there are real numbers!

Exercise 3.44. This exercise outlines a proof of the following theorem:

Theorem 3.45. *Let R be a rectangle with side lengths 1 and x . If x is irrational, then R cannot be tiled by finitely many squares (so that the squares have disjoint interiors and cover all of R).*

Given $f: \mathbb{R} \rightarrow \mathbb{R}$, define the f -**area** of a rectangle R with side lengths a and b by the formula

$$A_f(R) := f(a) \cdot f(b).$$

- Let $I, J \subset \mathbb{R}$ be two intervals and consider the rectangle $R := I \times J$. Suppose that the intervals I and J are tiled by finitely many smaller intervals:

$$I = I_1 \cup \cdots \cup I_n \quad \text{and} \quad J = J_1 \cup \cdots \cup J_m.$$


Then R is tiled by the rectangles $R_{ij} := I_i \times J_j$, $1 \leq i \leq n$, $1 \leq j \leq m$, in a grid-like fashion:

$$\begin{array}{c} J_m \\ \vdots \\ J_1 \end{array} \begin{array}{|c|c|c|c|} \hline R_{1m} & & & R_{nm} \\ \hline \vdots & & & \\ \hline R_{11} & R_{21} & \cdots & R_{n1} \\ \hline \end{array} \begin{array}{c} I_1 \\ I_2 \\ \cdots \\ I_n \end{array}$$

Prove that if $f: \mathbb{R} \rightarrow \mathbb{R}$ is \mathbb{Q} -linear, then

$$A_f(R) = \sum_{i=1}^n \sum_{j=1}^m A_f(R_{ij}).$$

(b) Suppose that a rectangle R is tiled arbitrarily by finitely many rectangles Q_1, \dots, Q_k :



Prove that if $f: \mathbb{R} \rightarrow \mathbb{R}$ is \mathbb{Q} -linear, then

$$A_f(R) = \sum_{i=1}^k A_f(Q_i).$$

(c) Show that if $x \in \mathbb{R} \setminus \mathbb{Q}$, then there is a \mathbb{Q} -linear map $f: \mathbb{R} \rightarrow \mathbb{R}$ such that

$$f(1) = 1 \quad \text{and} \quad f(x) = -1.$$

(d) Deduce Theorem 3.45. *Hint:* What can you say about the f -area of a square?

4. DIMENSION

4.A. The second fundamental theorem of linear algebra

The goal of this subsection is to show that any two bases in a vector space have the same size:

Theorem 4.1 (Second fundamental theorem). *Let V be a vector space over a field F . If $B_1, B_2 \subseteq V$ are bases for V , then there is a bijection $f: B_1 \rightarrow B_2$.*

In particular, if B_1 is finite, then B_2 is also finite and $|B_1| = |B_2|$.

According to Theorem 4.1, the F -vector spaces

$$\{0\} = F^0, \quad F = F^1, \quad F^2, \quad F^3, \quad \dots, \quad F^n, \quad \dots$$

are pairwise non-isomorphic. Also, the vector spaces $[\mathbb{N} \rightarrow F]^{<\infty}$ and $[\mathbb{R} \rightarrow F]^{<\infty}$ are not isomorphic to each other, because there is no bijection $\mathbb{N} \rightarrow \mathbb{R}$ (even though both \mathbb{N} and \mathbb{R} are infinite sets). Together, Theorems 3.15 and 4.1 give a complete characterization of all F -vector spaces up to isomorphism.

Lemma 4.2 (Exchange lemma). *Let V be an F -vector space and let $I, B \subseteq V$. Suppose that the set I is independent, while B is a basis for V . Then for every $x \in I \setminus B$, there is some $y \in B \setminus I$ such that the set $(I \setminus \{x\}) \cup \{y\}$ is independent. Furthermore, if I is a basis for V , then so is $(I \setminus \{x\}) \cup \{y\}$.*

PROOF. First, we find $y \in B \setminus I$ such that the set $S_y := (I \setminus \{x\}) \cup \{y\}$ is independent. Suppose, towards a contradiction, that no such y exists. This means that for all $y \in B \setminus I$, the set S_y is not independent, i.e., $y \in \text{Span}(I \setminus \{x\})$ (see Lemma 3.23). Thus, $B \setminus I \subseteq \text{Span}(I \setminus \{x\})$. Since $x \notin B$, we also have $B \cap I \subseteq I \setminus \{x\}$; hence, $B \subseteq \text{Span}(I \setminus \{x\})$. Since $\text{Span}(B) = V$, we conclude that the set $I \setminus \{x\}$ is spanning, and in particular $x \in \text{Span}(I \setminus \{x\})$, which contradicts the independence of I .

To prove the “furthermore” part of the lemma, assume that I is a basis for V and let $y \in B \setminus I$ be an arbitrary element such that the set S_y is independent. We claim that S_y is also spanning, and hence it is a basis, as desired. By definition, $I \setminus \{x\} \subseteq S_y$, so we only need to show that $x \in \text{Span}(S_y)$. But $S_y \cup \{x\} = I \cup \{y\}$ is not an independent set, so $x \in \text{Span}(S_y)$ by Lemma 3.23. ■

Using the exchange lemma, it is easy to derive Theorem 4.1 in the case one of B_1, B_2 is finite. Indeed, suppose that B_1 contains n elements:

$$B_1 = \{x_1, x_2, \dots, x_n\}.$$

Applying the exchange lemma repeatedly, we can replace the elements of B_1 , one by one, by elements of B_2 , producing a sequence of bases, each containing n distinct elements:

$$\begin{aligned} &\{y_1, x_2, \dots, x_n\}, \\ &\{y_1, y_2, \dots, x_n\}, \\ &\quad \dots \\ &\{y_1, y_2, \dots, y_n\}. \end{aligned}$$

Since $\{y_1, \dots, y_n\} \subseteq B_2$ and B_2 itself is a basis, this means that $B_2 = \{y_1, \dots, y_n\}$.

It might seem that when both B_1 and B_2 are infinite, we need a more powerful tool than the exchange lemma, which only treats one element of $B_1 \setminus B_2$ at a time. It turns out, however, that just dealing with one element at a time is enough, even when the sets B_1 and B_2 are infinite, thanks to Zorn's lemma. Similar arguments based on Zorn's lemma are often used in different parts of mathematics.

We will need a form of Zorn's lemma that is slightly stronger than Theorem 3.31:

Theorem 4.3 (Zorn's lemma #2). *Let \mathcal{F} be a nonempty family of sets such that if $\mathcal{C} \subseteq \mathcal{F}$ is a chain, then $\bigcup \mathcal{C} \in \mathcal{F}$. Then \mathcal{F} has a maximal element.*

Exercise 4.4. For a family \mathcal{F} of sets, let \mathcal{F}^+ denote the set of all chains $\mathcal{C} \subseteq \mathcal{F}$.

- Show that for every family \mathcal{F} of sets, the set \mathcal{F}^+ satisfies the assumptions (Z1), (Z2), and (Z3) of Zorn's lemma (in the form of Theorem 3.31).
- Conclude that \mathcal{F}^+ has a maximal element; in other words, every family \mathcal{F} of sets contains a maximal chain $\mathcal{C} \subseteq \mathcal{F}$.
- Deduce Theorem 4.3.

Another tool we will need is a classical result known as the **Bernstein–Cantor–Schröder theorem**, which allows one to construct a *bijection* between two sets out of a pair of *injections*:

Theorem 4.5 (Bernstein–Cantor–Schröder). *Let A and B be sets. If $f: A \rightarrow B$ and $g: B \rightarrow A$ are injective functions, then there also exists a bijection $h: A \rightarrow B$.*

PROOF. Define a sequence of sets $B_0, A_0, B_1, A_1, B_2, A_2, \dots$ as follows:

$$B_0 := B \setminus \text{im}(f); \quad A_n := g(B_n); \quad B_{n+1} := f(A_n).$$

By definition, $A_n \subseteq A$ and $B_n \subseteq B$ for all $n \in \mathbb{N}$. Let

$$\begin{aligned} A' &:= A_0 \cup A_1 \cup A_2 \cup \dots & \text{and} & & A'' &:= A \setminus A'; \\ B' &:= B_0 \cup B_1 \cup B_2 \cup \dots & \text{and} & & B'' &:= B \setminus B'. \end{aligned}$$

Note that $A' \subseteq \text{im}(g)$, so we can define a function $h: A \rightarrow B$ by

$$h(a) := \begin{cases} f(a) & \text{if } a \in A''; \\ g^{-1}(a) & \text{if } a \in A'. \end{cases} \quad (4.6)$$

This function h is a desired bijection ([exercise!](#)). ■

PROOF OF THEOREM 4.1. Let $B_1, B_2 \subseteq V$ be two bases of V . In view of Theorem 4.5, we just have to show that there is an injective function from B_1 to B_2 .

An **exchange** is a function f with the following properties:

- $\text{dom}(f) \subseteq B_1$ and $\text{im}(f) \subseteq B_2$;
- f is injective;
- the sets $\text{im}(f)$ and $B_1 \setminus \text{dom}(f)$ are disjoint;
- the set $I_f := \text{im}(f) \cup (B_1 \setminus \text{dom}(f))$ is independent.

Let \mathcal{E} be the set of all exchanges. Note that $\mathcal{E} \neq \emptyset$, since the *empty function* $\emptyset: \emptyset \rightarrow \emptyset$, whose domain and image are both empty, is an exchange. (In particular, $I_\emptyset = B_1$, so I_\emptyset is independent.) We wish to apply Zorn's lemma to \mathcal{E} to obtain a *maximal* exchange, and then to show that it must be an injection $B_1 \rightarrow B_2$. Of course, Zorn's lemma applies to families of sets, while \mathcal{E} is a family of functions, but we can view each $f \in \mathcal{E}$ as a set by identifying f with its graph, i.e., with the set

$$\Gamma_f := \{(x, y) : f(x) = y\}.$$

Thus, $f \subseteq g$ means that $\text{dom}(f) \subseteq \text{dom}(g)$ and $g(x) = f(x)$ for all $x \in \text{dom}(f)$.

Exercise 4.7. Show that if $\mathcal{C} \subseteq \mathcal{E}$ is a chain of exchanges, then $\bigcup \mathcal{C}$ is an exchange.

With Exercise 4.7 in hand, we can apply Theorem 4.3 to conclude that there is a maximal exchange $f \in \mathcal{E}$. Suppose, towards a contradiction, that $B_1 \setminus \text{dom}(f) \neq \emptyset$ and let x be an arbitrary element of $B_1 \setminus \text{dom}(f)$. We will show how to extend f to x ; i.e., we shall construct an exchange f' such that $f' \supset f$ and $\text{dom}(f') = \text{dom}(f) \cup \{x\}$. This would contradict the maximality of f , thus proving that $\text{dom}(f) = B_1$, as desired. There are two cases to consider:

Case 1: $x \in B_2$. Then we can set $f'(x) := x$, and it is not hard to check that f' is an exchange.

Case 2: $x \notin B_2$. Then we have $x \in I_f \setminus B_2$, so we may apply the exchange lemma to the basis B_2 and the independent set I_f to obtain $y \in B_2 \setminus I_f$ such that $(I_f \setminus \{x\}) \cup \{y\}$ is independent. Then we can set $f'(x) := y$.

Hence, $\text{dom}(f) = B_1$ and $f: B_1 \rightarrow B_2$ is a desired injection. ■

4.B. Dimension and finite-dimensional spaces

Definition 4.8. A vector space V is called **finite-dimensional** if it has a finite basis, and **infinite-dimensional** otherwise. The size of any basis in a finite-dimensional vector space V is called the **dimension** of V , denoted $\dim V$. Sometimes, we write $\dim_F V$ instead of $\dim V$ to explicitly indicate that we are working over F .

Example 4.9. Let F be a field and let $n \in \mathbb{N}$. Then the dimension of F^n is n .

Example 4.10. Let F be a field and let $m, n \in \mathbb{N}$. Then the dimension of $M_{m \times n}(F)$, as a vector space over F , is mn .

Example 4.11. The dimension of \mathbb{C} as a vector space over \mathbb{R} is $\dim_{\mathbb{R}} \mathbb{C} = 2$, since the set $\{1, i\}$ is a basis for \mathbb{C} over \mathbb{R} . On the other hand, \mathbb{C} is also a vector space over \mathbb{C} , and $\dim_{\mathbb{C}} \mathbb{C} = 1$. (In general, every field is a one-dimensional vector space over itself.)

Example 4.12. As a vector space over \mathbb{Q} , \mathbb{R} is infinite-dimensional. The most straightforward way to see this is by using basic set theory, which, unfortunately, falls outside the scope of this course.¹⁴ It is fairly easy to guess an infinite subset of \mathbb{R} that is independent over \mathbb{Q} , but it is usually surprisingly hard to *prove* that it is independent. For instance, the set

$$\{\sqrt{p} : p \text{ is a prime number}\}$$

is \mathbb{Q} -linearly independent, but the proof of this fact is quite complicated. However, the independence of the following set is easy to verify:

Exercise 4.13. Show that the set $\{\ln p : p \text{ is a prime number}\}$ is \mathbb{Q} -linearly independent.

Lemma 4.14. Let V be a finite-dimensional vector space and let $W \subseteq V$ be a subspace. Then W is also finite-dimensional and $\dim W \leq \dim V$. Furthermore, if $W \neq V$, then $\dim W < \dim V$.

¹⁴For the initiated: every finite-dimensional \mathbb{Q} -vector space is isomorphic to \mathbb{Q}^n for some $n \in \mathbb{N}$; in particular, it is countable. On the other hand, \mathbb{R} is well-known to be uncountable.

PROOF. Let $B_W \subseteq W$ be any basis for W . Since B_W is an independent subset of V , there is a basis B for V such that $B \supseteq B_W$. Then

$$\dim W = |B_W| \leq |B| = \dim V,$$

as desired. Furthermore, if $\dim W = \dim V$, then $B_W = B$, and hence $W = V$. ■

Exercise 4.15. Let V be a finite-dimensional vector space. Show that the size of every independent set $I \subseteq V$ is at most $\dim V$, while the size of every spanning set $S \subseteq V$ is at least $\dim V$.

Exercise 4.16. Let V and W be finite-dimensional vector spaces over a field F . Show that

$$\dim(V \oplus W) = \dim V + \dim W.$$

Theorem 4.17 (Rank–nullity). *Let V and W be vector spaces over a field F and let $\varphi: V \rightarrow W$ be a linear function. Suppose that V is finite-dimensional. Then the spaces $\ker(\varphi)$, $\text{im}(\varphi)$ are also finite-dimensional, and we have*

$$\dim V = \dim \ker(\varphi) + \dim \text{im}(\varphi).$$

PROOF. By Exercise 3.42, V is isomorphic to $\ker(\varphi) \oplus (V/\ker(\varphi))$. By Exercise 2.40, $V/\ker(\varphi)$ is isomorphic to $\text{im}(\varphi)$. Thus, V is isomorphic to $\ker(\varphi) \oplus \text{im}(\varphi)$, and we are done by Exercise 4.16. ■

Corollary 4.18. *Let V be a finite-dimensional vector space and let $\varphi: V \rightarrow V$ be a linear function. The following statements are equivalent:*

- (1) φ is injective;
- (2) φ is surjective;
- (3) $\dim \ker(\varphi) = 0$;
- (4) $\dim \text{im}(\varphi) = \dim V$.

PROOF. The equivalences (1) \iff (3) and (2) \iff (4) follow from Exercise 3.2 (the second of these also relies on Lemma 4.14). By Theorem 4.17, $\dim \ker(\varphi) = \dim V - \dim \text{im}(\varphi)$, so $\dim \ker(\varphi) = 0$ if and only if $\dim V = \dim \text{im}(\varphi)$, which proves (3) \iff (4). ■

Remark 4.19. The equivalence (1) \iff (2) in Corollary 4.18 may fail for infinite-dimensional V . For example, consider the linear function $\varphi: F^{\mathbb{N}} \rightarrow F^{\mathbb{N}}$ given by $\varphi(x_0, x_1, x_2, \dots) := (x_1, x_2, x_3, \dots)$. Then φ is surjective but not injective. Similarly, the function $\psi: F^{\mathbb{N}} \rightarrow F^{\mathbb{N}}$ given by $\psi(x_0, x_1, x_2, \dots) := (0, x_0, x_1, \dots)$ is injective but not surjective.

4.C. Using dimension: algebraic numbers

Recall that a complex number $a \in \mathbb{C}$ is **algebraic** if there is a nonzero polynomial $p(x)$ with rational coefficients such that $p(a) = 0$. (For more details, see Example 1.23.) Denote the set of all algebraic numbers by $\overline{\mathbb{Q}}$. We have now the tools to prove the following theorem, which was stated in §1.C:

Theorem 1.25. $\overline{\mathbb{Q}}$ is a subfield of \mathbb{C} .

Lemma 4.20. *For $\alpha \in \mathbb{C}$, let $S_\alpha := \{\alpha^k : k \in \mathbb{N}\}$ and $V_\alpha := \text{Span}_{\mathbb{Q}}(S_\alpha)$. Then α is algebraic if and only if the \mathbb{Q} -vector space V_α is finite-dimensional.*

PROOF. Suppose that V_α is finite-dimensional. Then S_α is not an infinite independent set, and hence there exists a nontrivial linear combination of elements of S_α with rational coefficients that evaluates to zero; i.e., we can write

$$a_0 \cdot 1 + a_1 \cdot \alpha + a_2 \cdot \alpha^2 + \cdots + a_n \cdot \alpha^n = 0,$$

for some $n \in \mathbb{N}$, $a_0, \dots, a_n \in \mathbb{Q}$, and $a_n \neq 0$. Thus, α is a root of a nonzero polynomial with rational coefficients, as desired. Conversely, suppose that

$$a_0 + a_1\alpha + a_2\alpha^2 + \cdots + a_n\alpha^n = 0,$$

for some $n \in \mathbb{N}$, $a_0, \dots, a_n \in \mathbb{Q}$, and $a_n \neq 0$. We claim that then

$$V_\alpha = \text{Span}_{\mathbb{Q}}(\{1, \alpha, \alpha^2, \dots, \alpha^{n-1}\}),$$

and, in particular, $\dim_{\mathbb{Q}} V_\alpha \leq n$. It suffices to show that $\alpha^k \in \text{Span}_{\mathbb{Q}}(\{1, \alpha, \dots, \alpha^{n-1}\})$ for all $k \in \mathbb{N}$. The proof is by induction on k . For $k \leq n-1$ the statement is clear; furthermore, we have

$$\alpha^n = -\frac{a_0}{a_n} - \frac{a_1}{a_n}\alpha - \frac{a_2}{a_n}\alpha^2 - \dots - \frac{a_{n-1}}{a_n}\alpha^{n-1} \in \text{Span}_{\mathbb{Q}}(\{1, \alpha, \dots, \alpha^{n-1}\}). \quad (4.21)$$

We need to show that if $\alpha^k \in \text{Span}_{\mathbb{Q}}(\{1, \alpha, \dots, \alpha^{n-1}\})$, then $\alpha^{k+1} \in \text{Span}_{\mathbb{Q}}(\{1, \alpha, \dots, \alpha^{n-1}\})$. Write

$$\alpha^k = c_0 + c_1\alpha + c_2\alpha^2 + \dots + c_{n-1}\alpha^{n-1}.$$

Then $\alpha^{k+1} \in \text{Span}_{\mathbb{Q}}(\{1, \alpha, \dots, \alpha^{n-1}\})$ since

$$\alpha^{k+1} = \alpha \cdot \alpha^k = \underbrace{c_0\alpha + c_1\alpha^2 + \dots + c_{n-2}\alpha^{n-1}}_{\in \text{Span}_{\mathbb{Q}}(\{1, \alpha, \dots, \alpha^{n-1}\})} + \underbrace{c_{n-1}\alpha^n}_{\in \text{Span}_{\mathbb{Q}}(\{1, \alpha, \dots, \alpha^{n-1}\}) \text{ by (4.21)}}. \quad \blacksquare$$

Lemma 4.22. *Let $S, T \subseteq \mathbb{C}$ and let $ST := \{st : s \in S, t \in T\}$. If the \mathbb{Q} -vector spaces $\text{Span}_{\mathbb{Q}}(S)$ and $\text{Span}_{\mathbb{Q}}(T)$ are finite-dimensional, then so is $\text{Span}_{\mathbb{Q}}(ST)$; moreover,*

$$\dim \text{Span}_{\mathbb{Q}}(ST) \leq (\dim \text{Span}_{\mathbb{Q}}(S)) \cdot (\dim \text{Span}_{\mathbb{Q}}(T)).$$

PROOF. Suppose that $\{s_1, \dots, s_n\} \subseteq S$ and $\{t_1, \dots, t_m\} \subseteq T$ are bases for $\text{Span}_{\mathbb{Q}}(S)$ and $\text{Span}_{\mathbb{Q}}(T)$, respectively. We claim that

$$\text{Span}_{\mathbb{Q}}(ST) = \text{Span}_{\mathbb{Q}}(\{s_it_j : 1 \leq i \leq n, 1 \leq j \leq m\}),$$

and thus $\dim \text{Span}_{\mathbb{Q}}(ST) \leq nm$, as desired. It suffices to show that for all $s \in S$ and $t \in T$,

$$st \in \text{Span}_{\mathbb{Q}}(\{s_it_j : 1 \leq i \leq n, 1 \leq j \leq m\}).$$

To that end, write

$$s = \sum_{i=1}^n a_i s_i \quad \text{and} \quad t = \sum_{j=1}^m b_j t_j,$$

where the coefficients a_i , $1 \leq i \leq n$, and b_j , $1 \leq j \leq m$, are rational. Then

$$st = \left(\sum_{i=1}^n a_i s_i \right) \left(\sum_{j=1}^m b_j t_j \right) = \sum_{i=1}^n \sum_{j=1}^m (a_i b_j) (s_i t_j) \in \text{Span}_{\mathbb{Q}}(\{s_it_j : 1 \leq i \leq n, 1 \leq j \leq m\}). \quad \blacksquare$$

PROOF OF THEOREM 1.25. The only things that require verification are:

- $\overline{\mathbb{Q}}$ is closed under addition;
- $\overline{\mathbb{Q}}$ is closed under multiplication;
- $\overline{\mathbb{Q}}$ is closed under taking additive inverses;
- $\overline{\mathbb{Q}}$ is closed under taking multiplicative inverses of nonzero elements.

Perhaps somewhat surprisingly, the latter two bullet points are relatively straightforward to check, while the former two are somewhat tricky. Indeed, suppose that $\alpha \in \overline{\mathbb{Q}} \setminus \{0\}$. Then we can write

$$a_0 + a_1\alpha + a_2\alpha^2 + \dots + a_n\alpha^n = 0, \quad (4.23)$$

for rational $a_0, \dots, a_n \in \mathbb{Q}$, not all of which are zero. Dividing both sides of (4.23) by α^n , we obtain

$$a_n + a_{n-1}\alpha^{-1} + a_{n-2}(\alpha^{-1})^2 + \dots + a_0(\alpha^{-1})^n = 0,$$

and hence α^{-1} is also algebraic. Additive inverses can be treated in a similar fashion.

Exercise 4.24. Show that $\overline{\mathbb{Q}}$ is closed under taking additive inverses.

Now we proceed to show that $\overline{\mathbb{Q}}$ is closed under addition. Let $\alpha, \beta \in \overline{\mathbb{Q}}$. By Lemma 4.20, this means that the associated \mathbb{Q} -vector spaces

$$V_\alpha = \text{Span}_{\mathbb{Q}}(S_\alpha) \quad \text{and} \quad V_\beta = \text{Span}_{\mathbb{Q}}(S_\beta)$$

are finite-dimensional. We wish to show that $\alpha + \beta \in \overline{\mathbb{Q}}$, i.e., that the space

$$V_{\alpha+\beta} = \text{Span}_{\mathbb{Q}}(S_{\alpha+\beta})$$

is also finite-dimensional. Observe that for each $k \in \mathbb{N}$, we have

$$(\alpha + \beta)^k = \sum_{i=0}^k \binom{k}{i} \alpha^i \beta^{k-i} \in \text{Span}_{\mathbb{Q}}(\{\alpha^i \beta^j : i, j \in \mathbb{N}\}) = \text{Span}_{\mathbb{Q}}(S_\alpha S_\beta),$$

where the product $S_\alpha S_\beta$ is defined as in Lemma 4.22. Therefore, $V_{\alpha+\beta} \subseteq \text{Span}_{\mathbb{Q}}(S_\alpha S_\beta)$; but the space $\text{Span}_{\mathbb{Q}}(S_\alpha S_\beta)$ is finite-dimensional by Lemma 4.22, and hence we are done.

Finally, the proof that $\overline{\mathbb{Q}}$ is closed under multiplication is left as an **exercise**.

Exercise 4.25. Show that $\overline{\mathbb{Q}}$ is closed under multiplication. ■

It is worthwhile to go over a concrete example to see what the proof of Theorem 1.25 means computationally. Suppose that $\alpha, \beta \in \mathbb{C}$ satisfy

$$1 + \alpha + \alpha^2 = 0 \quad \text{and} \quad -1 - 2\beta + \beta^2 = 0. \quad (4.26)$$

Then α and β are algebraic, and by Theorem 1.25, so is their sum $\alpha + \beta$. How do we actually find a polynomial $p(x)$ with rational coefficients such that $p(\alpha + \beta) = 0$? The strategy is to express the powers of $\alpha + \beta$ as linear combinations of the four monomials $1, \alpha, \beta,$ and $\alpha\beta$, using (4.26) to eliminate all the higher powers of α and β . Then we will be able to find a nontrivial linear relation between the five expressions $1, \alpha + \beta, (\alpha + \beta)^2, (\alpha + \beta)^3,$ and $(\alpha + \beta)^4$, showing that $\alpha + \beta$ is a root of a polynomial of degree at most 4.

To begin with, we repeatedly apply (4.26) to express the first five powers of α (resp. β) as linear combinations of 1 and α (resp. 1 and β):

$$\begin{array}{ll} \alpha^0 = 1 & \beta^0 = 1 \\ \alpha^1 = \alpha & \beta^1 = \beta \\ \alpha^2 = -1 - \alpha & \beta^2 = 1 + 2\beta \\ \alpha^3 = -\alpha - (-1 - \alpha) = 1 & \beta^3 = \beta + 2(1 + 2\beta) = 2 + 5\beta \\ \alpha^4 = \alpha & \beta^4 = 2\beta + 5(1 + 2\beta) = 5 + 12\beta \end{array}$$

Now we compute:

$$\begin{aligned} (\alpha + \beta)^0 &= 1, & (\alpha + \beta)^1 &= \alpha + \beta, \\ (\alpha + \beta)^2 &= \alpha^2 + 2\alpha\beta + \beta^2 = (-1 - \alpha) + 2\alpha\beta + (1 + 2\beta) = -\alpha + 2\beta + 2\alpha\beta, \\ (\alpha + \beta)^3 &= \alpha^3 + 3\alpha^2\beta + 3\alpha\beta^2 + \beta^3 \\ &= 1 + 3(-1 - \alpha)\beta + 3\alpha(1 + 2\beta) + (2 + 5\beta) \\ &= 3 + 3\alpha + 2\beta + 3\alpha\beta, \\ (\alpha + \beta)^4 &= \alpha^4 + 4\alpha^3\beta + 6\alpha^2\beta^2 + 4\alpha\beta^3 + \beta^4 \\ &= \alpha + 4\beta + 6(-1 - \alpha)(1 + 2\beta) + 4\alpha(2 + 5\beta) + (5 + 12\beta) \\ &= -1 + 3\alpha + 4\beta + 8\alpha\beta. \end{aligned}$$

We can now form a matrix whose rows correspond to the monomials $1, \alpha, \beta, \alpha\beta$ and whose columns correspond to the powers of $\alpha + \beta$, from 0 to 4, that contains the coefficients of the above expressions:

$$A := \begin{bmatrix} 1 & 0 & 0 & 3 & -1 \\ 0 & 1 & -1 & 3 & 3 \\ 0 & 1 & 2 & 2 & 4 \\ 0 & 0 & 2 & 3 & 8 \end{bmatrix}.$$

To find a desired polynomial p , we just need to find a nontrivial linear combination of the columns of A that evaluates to zero, which requires solving a system of 4 homogeneous linear equations in 5 variables. This is a computationally tractable problem; for instance, we find that

$$7 \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + 2 \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} - \begin{bmatrix} 0 \\ -1 \\ 2 \\ 2 \end{bmatrix} - 2 \begin{bmatrix} 3 \\ 3 \\ 2 \\ 3 \end{bmatrix} + \begin{bmatrix} -1 \\ 3 \\ 4 \\ 8 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix},$$

and hence

$$7 + 2(\alpha + \beta) - (\alpha + \beta)^2 - 2(\alpha + \beta)^3 + (\alpha + \beta)^4 = 0.$$

Exercise 4.27. Show that if a complex number $a \in \mathbb{C}$ is a root of a nonzero polynomial with algebraic coefficients, then a is itself algebraic.

Exercise 4.28. When K is a field and $F \subseteq K$ is a subfield of K , we say that K is an **extension** of F . Recall that if K is an extension of F , then K can be naturally viewed as a vector space over F . A field extension $K \supseteq F$ is called **finite** if the dimension of K , as an F -vector space, is finite. A field extension $K \supseteq F$ is **algebraic** if for each element $a \in K$, there is a nonzero polynomial $p(x)$ with coefficients in F such that $p(a) = 0$. Thus, \mathbb{C} is a finite extension of \mathbb{R} and $\overline{\mathbb{Q}}$ is an algebraic extension of \mathbb{Q} .

Let F be a field. Show that every finite extension of F is algebraic.

Extra exercises for Section 4

Exercise 4.29. Fix distinct $a_1, \dots, a_n \in \mathbb{R}$ and let $P_{n-1}(\mathbb{R})$ denote the set of all polynomials $p(x)$ with real coefficients in a single variable x of degree at most $n - 1$.

- (a) Show that $P_{n-1}(\mathbb{R})$ is an \mathbb{R} -vector space. What is its dimension?
- (b) For each $1 \leq i \leq n$, let $q_i(x)$ be the polynomial

$$q_i(x) := (x - a_1) \cdots (x - a_{i-1})(x - a_{i+1}) \cdots (x - a_n) \quad (n - 1 \text{ factors}).$$

Let $Q := \{q_i : 1 \leq i \leq n\}$. Show that Q is an independent subset of $P_{n-1}(\mathbb{R})$. *Hint:* What would happen to a linear combination of elements of Q if we plug in a_i instead of x ?

- (c) Conclude that for every $p \in P_{n-1}(\mathbb{R})$, there exist coefficients $c_1, \dots, c_n \in \mathbb{R}$ such that

$$\frac{p(x)}{(x - a_1) \cdots (x - a_n)} = \frac{c_1}{x - a_1} + \cdots + \frac{c_n}{x - a_n}.$$

This fact is used in calculus to find antiderivatives of rational functions.

Exercise 4.30. Show that if F is a finite field, then the size of F is a power of a prime number. *Hint:* Use the result of Exercise 1.45.

5. SPACES OF LINEAR FUNCTIONS

5.A. The dual space

Definition 5.1. Let V, W be vector spaces over a field F . We use $\text{Lin}(V, W)$ to denote the set of all linear functions $f: V \rightarrow W$. Note that $\text{Lin}(V, W)$ is a subspace of W^V , viewed as an F -vector space

under pointwise addition and scaling. In the special case when $W = F$, we set $V^* := \text{Lin}(V, F)$ and call V^* the **dual space** of V .

Let $B \subseteq V$ be a basis for an F -vector space V . For each $x \in B$, let $x_B^* : V \rightarrow F$ be the unique linear function such that for all $y \in B$,

$$x_B^*(y) = \begin{cases} 1 & \text{if } y = x; \\ 0 & \text{if } y \neq x. \end{cases}$$

Explicitly, for each $v \in V$, the value $x_B^*(v)$ is determined as follows. Write v as a linear combination of the elements of the basis B :

$$v = \sum_{y \in B} c(y) \cdot y,$$

where $c \in [B \rightarrow F]^{<\infty}$. Then we have

$$x_B^*(v) = \sum_{y \in B} c(y) \cdot x_B^*(y) = c(x),$$

that is, $x_B^*(v)$ is equal to the coefficient of x in the unique expansion of v as a linear combination of the elements of B . In other words, we can write

$$v = \sum_{x \in B} x_B^*(v) \cdot x. \quad (5.2)$$

Lemma 5.3. *Let V be a finite-dimensional vector space over a field F and let $B \subseteq V$ be a basis for V . Then $B^* := \{x_B^* : x \in B\}$ is a basis for V^* , called the **dual basis** corresponding to B .*

PROOF. First, we show that B^* is independent; this is true regardless of whether V is finite-dimensional or not. Let $c \in [B \rightarrow F]^{<\infty}$ and suppose that

$$\sum_{x \in B} c(x) \cdot x_B^* = 0. \quad (5.4)$$

Take any $y \in B$ and plug it into (5.4); we then obtain

$$0 = \sum_{x \in B} c(x) \cdot x_B^*(y) = c(y),$$

i.e., $c(y) = 0$ for all $y \in B$, as desired.

Now we show that $\text{Span}(B^*) = V^*$. To that end, let $f \in V^*$. We claim that

$$f = \sum_{x \in B} f(x) \cdot x_B^*. \quad (5.5)$$

Since the set B is finite, the right-hand side of (5.5) is a valid linear combination, and hence (5.5) implies $f \in \text{Span}(B^*)$. To prove (5.5), we have to show that f agrees with the right-hand side of (5.5) when applied to each vector $v \in V$. And indeed, by (5.2), we have

$$f(v) = f\left(\sum_{x \in B} x_B^*(v) \cdot x\right) = \sum_{x \in B} x_B^*(v) \cdot f(x),$$

which precisely coincides with the right-hand side of (5.5) applied to v . ■

Remark 5.6. The conclusion of Lemma 5.3 fails if V is infinite-dimensional. Consider, for example, the space $[\mathbb{N} \rightarrow F]^{<\infty}$ of all infinite sequences of elements of F with only finitely many nonzero entries. The set $B = \{e_n : n \in \mathbb{N}\}$ is a basis for $[\mathbb{N} \rightarrow F]^{<\infty}$, where $e_n = (e_n(0), e_n(1), \dots)$ denotes the sequence such that

$$e_n(i) = \begin{cases} 1 & \text{if } i = n; \\ 0 & \text{if } i \neq n. \end{cases}$$

The corresponding elements of $([\mathbb{N} \rightarrow F]^{<\infty})^*$ are the functions $e_n^*: [\mathbb{N} \rightarrow F]^{<\infty} \rightarrow F$ given by

$$e_n^*(x_0, x_1, \dots) := x_n.$$

Now consider the linear function $\sigma: [\mathbb{N} \rightarrow F]^{<\infty} \rightarrow F$ given by

$$\sigma(x_0, x_1, \dots) := x_0 + x_1 + \dots. \tag{5.7}$$

The right-hand side of (5.7) is a well-defined summation, since only finitely many of the entries x_0, x_1, \dots are nonzero. It may be tempting to write

$$\sigma = e_0^* + e_1^* + \dots, \tag{5.8}$$

but the expression on the right-hand side of (5.8) is not a valid linear combination; in fact, we claim that $\sigma \notin \text{Span}(B^*)$. Indeed, if $f \in \text{Span}(B^*)$, then f can be written as a *finite* linear combination of elements of B^* , so there is $m \in \mathbb{N}$ such that

$$f = a_0 \cdot e_0^* + \dots + a_m \cdot e_m^*.$$

Then for all $n > m$, we have

$$f(e_n) = a_0 \cdot e_0^*(e_n) + \dots + a_m \cdot e_m^*(e_n) = 0.$$

On the other hand, we have $\sigma(e_n) = 1$ for all $n \in \mathbb{N}$.

Exercise 5.9. Show that if V is an *infinite*-dimensional vector space and $B \subseteq V$ is a basis for V , then the set $B^* := \{x_B^* : x \in B\}$ is *never* a basis for V^* .

Corollary 5.10. *If V is a finite-dimensional vector space, then V is isomorphic to V^* .* ■

Again, Corollary 5.10 fails for infinite-dimensional vector spaces.

Exercise 5.11. Let F be a field and let X be an arbitrary set. Consider the vector space $[X \rightarrow F]^{<\infty}$ over F . Show that the dual space $([X \rightarrow F]^{<\infty})^*$ is isomorphic to F^X .

Example 5.12. Consider the n -dimensional F -vector space F^n . By Corollary 5.10, the dual space $(F^n)^*$ is also isomorphic to F^n . An explicit isomorphism can be obtained by assigning to each tuple $(a_1, \dots, a_n) \in F^n$ the linear map $F^n \rightarrow F$ given by $(x_1, \dots, x_n) \mapsto a_1x_1 + \dots + a_nx_n$.

5.B. The double dual

Let V be a vector space over a field F . The space $V^{**} := (V^*)^*$ is called the **double dual** of V . Let $\iota: V \rightarrow V^{**}$ be the map that sends each $x \in V$ to the linear function $\iota(x): V^* \rightarrow F$ given by

$$(\iota(x))(f) := f(x) \quad \text{for all } f \in V^*.$$

Exercise 5.13. Show that the function $\iota(x): V^* \rightarrow F$ is indeed linear.

Exercise 5.14. Show that the map $\iota: V \rightarrow V^{**}$ is an *embedding*; that is, it is a linear injection.

If V is finite-dimensional, then, by Corollary 5.10, we have a chain of equalities

$$\dim V = \dim V^* = \dim V^{**}.$$

Hence, the injective linear map $\iota: V \rightarrow V^{**}$ must also be surjective, and thus, it is an isomorphism:

Lemma 5.15. *If V is a finite-dimensional vector space, then $\iota: V \rightarrow V^{**}$ is an isomorphism.* ■

There is an important distinction between Corollary 5.10 and Lemma 5.15. While Corollary 5.10 asserts the *existence* of an isomorphism $V \rightarrow V^*$, constructing such an isomorphism requires choosing a basis for V , and the result essentially depends on this choice. On the contrary, Lemma 5.15 provides a “canonical” isomorphism $V \rightarrow V^{**}$, namely ι , independent of any auxiliary choices.

5.C. Annihilators

Definition 5.16. Let V be an F -vector space and let $S \subseteq V$. The **annihilator** of S is the set

$$\text{Ann}(S) := \{f \in V^* : S \subseteq \ker(f)\}.$$

Exercise 5.17. Show that $\text{Ann}(S)$ is a subspace of V^* .

Lemma 5.18. Let V be a vector space over a field F and let $W \subseteq V$ be a subspace. Then

$$W = \{x \in V : f(x) = 0 \text{ for all } f \in \text{Ann}(W)\}. \quad (5.19)$$

PROOF. Let W' denote the right-hand side of (5.19). It is clear that $W \subseteq W'$. To establish the opposite inclusion, consider any $x \in V \setminus W$. Let $B_W \subseteq W$ be a basis for W . Since the set $B_W \cup \{x\}$ is independent, it can be extended to a basis for V . This means that any assignment $B_W \cup \{x\} \rightarrow F$ can be extended to a linear function $V \rightarrow F$. In particular, there is $f \in V^*$ such that $f(x) = 1$ while $f(w) = 0$ for all $w \in B_W$. Then $f \in \text{Ann}(W)$ and $f(x) \neq 0$, and thus $x \notin W'$. ■

Corollary 5.20 (Capelli–Fontené–Frobenius–Kronecker–Rouché–...). Let V be a vector space over a field F and let $X \subseteq V$. The following statements are equivalent for a vector $v \in V$:

- (1) $v \notin \text{Span}(X)$;
- (2) there is $f \in V^*$ such that $f(x) = 0$ for all $x \in X$, while $f(v) \neq 0$.

Exercise 5.21. Deduce Corollary 5.20 from Lemma 5.18.


Corollary 5.20 is important enough to have a name; furthermore, it has many *different* names depending on where you are. In particular, it is called

- Rouché–Capelli theorem in Italy and English-speaking countries;
- Kronecker–Capelli theorem in Russia and Poland;
- Rouché–Fontené theorem in France;
- Rouché–Frobenius theorem in Spain and Latin America.

The following result is a simple (and somewhat facetious) application of Corollary 5.20. The game **Lights Out** is played as follows. Let $G = (V, E)$ be a finite graph¹⁵. Suppose that at each vertex of G , there is a light bulb that can be in one of the two states: on or off. Each vertex of G is equipped with a light switch. Flicking the switch at a vertex $u \in V$ simultaneously changes the states of the light bulbs at u and all the vertices adjacent to u . At the start of the game, all the lights are off. The goal is to turn all the lights on in finitely many moves.¹⁶

Theorem 5.22 (Lights Out). For every finite graph $G = (V, E)$, it is possible to turn all the lights on using the rules of **Lights Out**.

Remark 5.23. It is important that our goal is to turn *all* the lights on. For example, if G is the 4-vertex graph shown below, then it is impossible to turn on *precisely one* light bulb:



Exercise 5.24. Show that if G is the graph from Remark 5.23, then it is impossible to turn on precisely one light bulb. More generally, show that if every vertex of G has odd degree, then it is impossible to turn on an odd number of light bulbs. (Solving this exercise may be easier after reading the proof of Theorem 5.22 given below.)

¹⁵By a *graph* here we mean a **simple** graph, i.e., one in which no vertex is adjacent to itself and every pair of vertices is joined by at most one edge.

¹⁶As the name suggests, the original version involves turning all the lights *off*, but the two versions are equivalent.

Before proving Theorem 5.22, let us quickly review some graph-theoretic notation. Let $G = (V, E)$ be a finite graph. For a vertex $v \in V$, let $N_G(v)$ denote the **neighborhood** of v in G , i.e., the set of all vertices $u \in V$ that are adjacent to v :

$$N_G(v) := \{u \in V : uv \in E\}.$$

The **degree** of a vertex v is the size of its neighborhood: $\deg_G(v) := |N_G(v)|$.

Exercise 5.25 (Handshake lemma). Let $G = (V, E)$ be a finite graph. Show that

$$\sum_{v \in V} \deg_G(v) = 2|E|.$$

Conclude that the number of vertices of odd degree in G is even.

Finally, for a subset $U \subseteq V$, the subgraph of G **induced** by U , denoted $G[U]$, is the graph with vertex set U in which two vertices are adjacent if and only if they are adjacent in G (see Fig. 4).




Figure 4. Induced subgraphs

PROOF OF THEOREM 5.22. To turn this into a linear algebra problem, we represent each state of the lights by a vector $x \in \mathbb{F}_2^V$ such that for all $u \in V$,

$$x(u) = \begin{cases} 1 & \text{if the light at } u \text{ is on;} \\ 0 & \text{if the light at } u \text{ is off.} \end{cases}$$

For each $v \in V$, let $s_v \in \mathbb{F}_2^V$ be the vector given by

$$s_v(u) := \begin{cases} 1 & \text{if } u = v \text{ or } u \in N_G(v); \\ 0 & \text{otherwise.} \end{cases}$$

Also, let $e := (1, \dots, 1) \in \mathbb{F}_2^V$ be the vector all of whose entries are 1. The key observation is that, since we are working modulo 2, flicking the switch at a vertex $v \in V$ results in adding s_v to the vector $x \in \mathbb{F}_2^V$ representing the current state of the lights. This has two consequences, which are otherwise not entirely obvious:

- the order of switches does not affect the resulting arrangement of lights; and
- it is never necessary to flick the same switch more than once.

Our goal is to show that for some $v_1, \dots, v_k \in V$, we have $e = s_{v_1} + \dots + s_{v_k}$, or, equivalently, e is in the span of $\{s_v : v \in V\}$. Thanks to Corollary 5.20, it suffices to prove that whenever $f \in (\mathbb{F}_2^V)^*$ satisfies $f(s_v) = 0$ for all $v \in V$, then $f(e) = 0$ as well. Using the theory developed in §5.A, we can describe all linear functions $f \in (\mathbb{F}_2^V)^*$ explicitly:

Exercise 5.26. Let $f: \mathbb{F}_2^V \rightarrow \mathbb{F}_2$ be a linear function. Show that there is a subset $U \subseteq V$ such that

$$f(x) = \sum_{u \in U} x(u) \quad \text{for all } x \in \mathbb{F}_2^V.$$

Therefore, our problem can now be restated as follows: Suppose that $U \subseteq V$ is a subset such that

$$\sum_{u \in U} s_v(u) = 0 \quad \text{for all } v \in V. \quad (5.27)$$

From this, we wish to deduce that

$$\sum_{u \in U} e(u) = \sum_{u \in U} 1 = |U| = 0 \pmod{2},$$

i.e., the size of U is even. To this end, consider the induced subgraph $G' := G[U]$. For each $v \in U$,

$$\sum_{u \in U} s_v(u) = \underbrace{1}_{\text{the contribution of } v} + \underbrace{|U \cap N_G(v)|}_{\text{the contribution of the neighbors of } v} = 1 + |N_{G'}(v)| = 1 + \deg_{G'}(v),$$

and hence, by (5.27), the degree of every vertex in G' is odd. Hence, by the handshake lemma, the number of vertices of G' —i.e., the size of U —is even, as desired. ■

5.D. Dual functions

Definition 5.28. Let V, W be vector spaces over a field F and let $\varphi: V \rightarrow W$ be a linear function. The **dual** of φ is the function $\varphi^*: W^* \rightarrow V^*$ given by

$$\varphi^*(f) := f \circ \varphi \quad \text{for all } f \in W^*.$$

Here \circ denotes composition of functions, as shown on the diagram below:

$$\begin{array}{ccccc} V & \xrightarrow{\varphi} & W & \xrightarrow{f} & F \\ & \searrow & & \nearrow & \\ & & & & f \circ \varphi = \varphi^*(f) \end{array}$$

Example 5.29. If $\varphi: \mathbb{R} \rightarrow \mathbb{R}$ is given by $\varphi(x) = 2x$, then for every $f \in \mathbb{R}^*$, $\varphi^*(f)$ is the function

$$\mathbb{R} \rightarrow \mathbb{R}: x \mapsto f(2x).$$

Exercise 5.30. Let V, W be vector spaces over a field F .

- Show that if $\varphi \in \text{Lin}(V, W)$ and $f \in W^*$, then the function $\varphi^*(f): V \rightarrow F$ is linear (and hence it indeed belongs to V^* , as asserted in Definition 5.28).
- Show that if $\varphi \in \text{Lin}(V, W)$, then the dual map $\varphi^*: W^* \rightarrow V^*$ is linear.
- Show that the function $\text{Lin}(V, W) \rightarrow \text{Lin}(W^*, V^*): \varphi \mapsto \varphi^*$ is linear.

Exercise 5.31. Fix a field F . For an F -vector space V , let $\text{id}_V: V \rightarrow V$ denote the corresponding identity map.

- Let V be an F -vector space. Show that $(\text{id}_V)^* = \text{id}_{V^*}$.
- Let U, V , and W be F -vector spaces and let $\varphi: U \rightarrow V$ and $\psi: V \rightarrow W$ be linear maps. Show that $(\psi \circ \varphi)^* = \varphi^* \circ \psi^*$. This situation is illustrated by the following diagram:

$$\begin{array}{ccccc}
 & & \psi \circ \varphi & & \\
 & \curvearrowright & & \curvearrowleft & \\
 U & \xrightarrow{\varphi} & V & \xrightarrow{\psi} & W \\
 & & & & \\
 U^* & \xleftarrow{\varphi^*} & V^* & \xleftarrow{\psi^*} & W^* \\
 & \curvearrowright & & \curvearrowleft & \\
 & & (\psi \circ \varphi)^* = \varphi^* \circ \psi^* & &
 \end{array}$$

If you wish to impress your friends at a party, you can summarize the above statements as follows:

Taking duals is a contravariant functor from the category of F -vector spaces and linear maps to itself.

Exercise 5.32. Let V, W be vector spaces over a field F .

- (a) Show that the function $\text{Lin}(V, W) \rightarrow \text{Lin}(W^*, V^*): \varphi \mapsto \varphi^*$ is injective.
- (b) Show that if W is finite-dimensional, then the function $\text{Lin}(V, W) \rightarrow \text{Lin}(W^*, V^*): \varphi \mapsto \varphi^*$ is surjective, and hence it is an isomorphism of F -vector spaces.

Lemma 5.33. Let V, W be vector spaces over a field F and let $\varphi: V \rightarrow W$ be a linear function. Then $\ker(\varphi^*) = \text{Ann}(\text{im}(\varphi))$.

PROOF. For each $f \in W^*$, we have

$$\begin{aligned}
 f \in \ker(\varphi^*) &\iff \varphi^*(f) = 0 \iff (\varphi^*(f))(v) = 0 \text{ for all } v \in V \\
 &\iff f(\varphi(v)) = 0 \text{ for all } v \in V \\
 &\iff f(w) = 0 \text{ for all } w \in \text{im}(\varphi) \\
 &\iff f \in \text{Ann}(\text{im}(\varphi)). \quad \blacksquare
 \end{aligned}$$

Theorem 5.34. Let V, W be vector spaces over a field F and let $\varphi: V \rightarrow W$ be a linear function. Then the F -vector spaces $\text{im}(\varphi^*)$ and $(\text{im}(\varphi))^*$ are isomorphic.

PROOF. From Lemma 5.33 and the first isomorphism theorem, it follows that

$$\text{im}(\varphi^*) \cong W^*/\ker(\varphi^*) = W^*/\text{Ann}(\text{im}(\varphi)).$$

Let $\rho: W^* \rightarrow (\text{im}(\varphi))^*$ be the function given by

$$\rho(f) := f|_{\text{im}(\varphi)} \quad \text{for all } f \in W^*,$$

where $f|_{\text{im}(\varphi)}$ denotes the restriction of f to the subset $\text{im}(\varphi)$ of its domain. It is easy to verify that ρ is linear. Also, ρ is surjective, since every linear function $\text{im}(\varphi) \rightarrow F$ can be extended to a linear function $W \rightarrow F$ (exercise!). Finally, $\ker(\rho) = \text{Ann}(\text{im}(\varphi))$ (this is essentially the definition of the annihilator). Hence, by the first isomorphism theorem again,

$$(\text{im}(\varphi))^* = \text{im}(\rho) \cong W^*/\ker(\rho) = W^*/\text{Ann}(\text{im}(\varphi)),$$

and we are done. \blacksquare

Corollary 5.35. Let V, W be vector spaces over a field F and let $\varphi: V \rightarrow W$ be a linear function. If the space $\text{im}(\varphi)$ is finite-dimensional, then $\dim \text{im}(\varphi) = \dim \text{im}(\varphi^*)$.

PROOF. By Corollary 5.10 and Theorem 5.34, $\dim \text{im}(\varphi) = \dim(\text{im}(\varphi))^* = \dim \text{im}(\varphi^*)$. \blacksquare

5.E. Representation of linear functions by matrices

The introduction of numbers as coordinates ... is an act of violence.

Hermann Weyl

In this subsection we discuss a very convenient way of encoding linear functions via matrices. To start with, we need to introduce the somewhat technical notion of an *ordered basis*. Let V be a vector space over a field F . We say that a tuple $(x_1, \dots, x_n) \in V^n$ is **independent** if the set $\{x_1, \dots, x_n\}$ is independent *and* the vectors x_1, \dots, x_n are pairwise distinct. (So, for example, if $x \in V \setminus \{0\}$, then the set $\{x, x\} = \{x\}$ is independent, while the pair (x, x) is not.) An **ordered basis** for V is an independent tuple (x_1, \dots, x_n) such that $\text{Span}(\{x_1, \dots, x_n\}) = V$; equivalently, $\{x_1, \dots, x_n\}$ is a basis and x_1, \dots, x_n are distinct.

Exercise 5.36. Show that a tuple (x_1, \dots, x_n) is independent if and only if for all $a_1, \dots, a_n \in F$,

$$a_1x_1 + \dots + a_nx_n = 0 \iff a_1 = \dots = a_n = 0.$$

Let $X = (x_1, \dots, x_n)$ be an ordered basis for V (in particular, $\dim V = n$). Then for each $v \in V$, there is a unique sequence of coefficients $a_1, \dots, a_n \in F$ such that

$$v = a_1x_1 + \dots + a_nx_n.$$

We put these coefficients together in a column matrix and define $[v]_X \in M_{n \times 1}(F)$ by

$$[v]_X := \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix}.$$

The function $V \rightarrow M_{n \times 1}(F): v \mapsto [v]_X$ is an isomorphism of F -vector spaces.

Now let V and W be two F -vector spaces and let $\varphi \in \text{Lin}(V, W)$. If $X = (x_1, \dots, x_n)$ is an ordered basis for V and $Y = (y_1, \dots, y_m)$ is an ordered basis for W , then we let $[\varphi]_{X,Y} \in M_{m \times n}(F)$ be the m -by- n matrix such that for each $1 \leq i \leq n$, the i -th column of $[\varphi]_{X,Y}$ is $[\varphi(x_i)]_Y$; in symbols,

$$[\varphi]_{X,Y} := \begin{bmatrix} [\varphi(x_1)]_Y & \cdots & [\varphi(x_n)]_Y \end{bmatrix}.$$

Again, the map $\text{Lin}(V, W) \rightarrow M_{m \times n}(F): \varphi \mapsto [\varphi]_{X,Y}$ is an isomorphism of F -vector spaces.

Example 5.37. Let $P_2(\mathbb{R})$ be the \mathbb{R} -vector space of polynomials with real coefficients in a single variable x of degree at most 2. Then $X = (1, x, x^2)$ is an ordered basis for this space. Consider the linear function $\partial: P_2(\mathbb{R}) \rightarrow P_2(\mathbb{R})$ that sends each polynomial $p \in P_2(\mathbb{R})$ to its derivative p' . Since $\partial(1) = 0$, $\partial(x) = 1$, and $\partial(x^2) = 2x$, we obtain that

$$[\partial]_{X,X} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix}.$$

Exercise 5.38. If $\dim V = n$ and $\dim W = m$, then what is $\dim \text{Lin}(V, W)$?

Exercise 5.39. Let V be a vector space over a field F and let $X = (x_1, \dots, x_n)$ be an ordered basis for V . Show that if $\text{id}_V: V \rightarrow V$ denotes the identity function, then $[\text{id}_V]_{X,X} = I_n(F)$.

Exercise 5.40 (important). Let V and W be vector spaces over a field F and let $\varphi \in \text{Lin}(V, W)$. Let $X = (x_1, \dots, x_n)$ and $Y = (y_1, \dots, y_m)$ be ordered bases for V and W , respectively. Show that

$$[\varphi(v)]_Y = [\varphi]_{X,Y} [v]_X \quad \text{for all } v \in V.$$

(Juxtaposition on the right-hand side indicates matrix multiplication.) This is the *reason* why matrix multiplication is defined the way it is.

In view of Exercise 5.40, it makes sense to introduce the following notational convention. Let $A \in M_{m \times n}(F)$. Then A defines a linear map $T_A: M_{n \times 1}(F) \rightarrow M_{m \times 1}(F): v \mapsto Av$, and, according to Exercise 5.40, every linear function between two finite-dimensional vector spaces is “essentially” of this form. We often conflate T_A with A ; in particular, we write $\text{im}(A)$ and $\text{ker}(A)$ to indicate the image and the kernel of T_A , respectively.

Exercise 5.41. Let $A \in M_{m \times n}(F)$ be an m -by- n matrix over a field F . Let the columns of A be x_1, \dots, x_n . Show that $\text{im}(A) = \text{Span}(\{x_1, \dots, x_n\})$.

Lemma 5.42. Let U, V, W be F -vector spaces and let $\varphi: U \rightarrow V$ and $\psi: V \rightarrow W$ be linear maps. Let $X = (x_1, \dots, x_n)$, $Y = (y_1, \dots, y_m)$, and $Z = (z_1, \dots, z_k)$ be ordered bases for U, V , and W , respectively. Then

$$[\psi \circ \varphi]_{X,Z} = [\psi]_{Y,Z}[\varphi]_{X,Y}.$$

PROOF. Note that for all $A \in M_{k \times m}(F)$, $B \in M_{m \times n}(F)$ and for each $1 \leq i \leq n$, the i -th column of the matrix product AB is equal to A times the i -th column of B . In particular, we have

$$\begin{aligned} \text{the } i\text{-th column of } [\psi]_{Y,Z}[\varphi]_{X,Y} &= [\psi]_{Y,Z} \cdot (\text{the } i\text{-th column of } [\varphi]_{X,Y}) \\ &= [\psi]_{Y,Z}[\varphi(x_i)]_Y \\ \text{[by Exercise 5.40]} &= [(\psi(\varphi(x_i)))]_Z \\ &= [(\psi \circ \varphi)(x_i)]_Z \\ &= \text{the } i\text{-th column of } [\psi \circ \varphi]_{X,Z}. \quad \blacksquare \end{aligned}$$

Lemma 5.43. Let V and W be F -vector spaces and let $\varphi \in \text{Lin}(V, W)$. Let $X = (x_1, \dots, x_n)$ and $Y = (y_1, \dots, y_m)$ be ordered bases for V and W , respectively, and let $X^* = (x_1^*, \dots, x_n^*)$ and $Y^* = (y_1^*, \dots, y_m^*)$ be the corresponding dual bases for V^* and W^* . Then

$$[\varphi^*]_{Y^*,X^*} = ([\varphi]_{X,Y})^\top.$$

Example 5.44. Before proving Lemma 5.43, let us consider a simple concrete example. Let $\partial: P_2(\mathbb{R}) \rightarrow P_2(\mathbb{R})$ be the linear function that sends each polynomial $p \in P_2(\mathbb{R})$ to its derivative p' (as in Example 5.37). Using the ordered basis $X = (1, x, x^2)$, we obtain

$$[\partial]_{X,X} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix}.$$

Let the dual basis be $X^* = (f_0, f_1, f_2)$. For a polynomial $a + bx + cx^2$, we have

$$\begin{aligned} f_0(a + bx + cx^2) &= a; \\ f_1(a + bx + cx^2) &= b; \\ f_2(a + bx + cx^2) &= c. \end{aligned}$$

Thus,

$$\begin{aligned} (\partial^*(f_0))(a + bx + cx^2) &= (f_0 \circ \partial)(a + bx + cx^2) = f_0(b + 2cx) = b; \\ (\partial^*(f_1))(a + bx + cx^2) &= (f_1 \circ \partial)(a + bx + cx^2) = f_1(b + 2cx) = 2c; \\ (\partial^*(f_2))(a + bx + cx^2) &= (f_2 \circ \partial)(a + bx + cx^2) = f_2(b + 2cx) = 0. \end{aligned}$$

Therefore, we conclude that

$$\partial^*(f_0) = f_1, \quad \partial^*(f_1) = 2f_2, \quad \text{and} \quad \partial^*(f_2) = 0,$$

and hence

$$[\partial^*]_{X^*, X^*} = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 2 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix}^\top = ([\partial]_{X, X})^\top,$$

as predicted by Lemma 5.43.

PROOF OF LEMMA 5.43. Our aim is to show that $[\varphi]_{X, Y}(i, j) = [\varphi^*]_{Y^*, X^*}(j, i)$ for all $1 \leq i \leq m$, $1 \leq j \leq n$. Let us first compute $[\varphi]_{X, Y}(i, j)$. By definition, the j -th column of $[\varphi]_{X, Y}$ is $[\varphi(x_j)]_Y$. Recall that for each $w \in W$, we have

$$w = \sum_{i=1}^m y_i^*(w) \cdot y_i \quad (\text{see (5.2)}), \quad \text{hence} \quad [w]_Y = \begin{bmatrix} y_1^*(w) \\ \vdots \\ y_m^*(w) \end{bmatrix}.$$

Therefore, the j -th column of $[\varphi]_{X, Y}$ is

$$[\varphi(x_j)]_Y = \begin{bmatrix} y_1^*(\varphi(x_j)) \\ \vdots \\ y_m^*(\varphi(x_j)) \end{bmatrix},$$

and thus $[\varphi]_{X, Y}(i, j) = y_i^*(\varphi(x_j)) = (y_i^* \circ \varphi)(x_j)$. Now we need to compute $[\varphi^*]_{Y^*, X^*}(j, i)$. Again, by definition, the i -th column of $[\varphi^*]_{Y^*, X^*}$ is $[\varphi^*(y_i^*)]_{X^*}$. For each $f \in V^*$, we have

$$f = \sum_{j=1}^n f(x_j) \cdot x_j^* \quad (\text{see (5.5)}), \quad \text{hence} \quad [f]_{X^*} = \begin{bmatrix} f(x_1) \\ \vdots \\ f(x_n) \end{bmatrix},$$

and so the i -th column of $[\varphi^*]_{Y^*, X^*}$ is

$$[\varphi^*(y_i^*)]_{X^*} = \begin{bmatrix} (\varphi^*(y_i^*))(x_1) \\ \vdots \\ (\varphi^*(y_i^*))(x_n) \end{bmatrix}.$$

Thus, $[\varphi^*]_{Y^*, X^*}(j, i) = (\varphi^*(y_i^*))(x_j) = (y_i^* \circ \varphi)(x_j) = [\varphi]_{X, Y}(i, j)$, and we are done. \blacksquare

Definition 5.45. Let F be a field and let $A \in M_{m \times n}(F)$. The **rank** of A is $\text{rank}(A) := \dim \text{im}(A)$.

Exercise 5.46. Let $A \in M_{m \times n}(F)$ be an m -by- n matrix over a field F . Show that the rank of A is equal to the largest number of independent columns of A .

Corollary 5.47. Let $A \in M_{m \times n}(F)$ be an m -by- n matrix over a field F . Then $\text{rank}(A) = \text{rank}(A^\top)$.

PROOF. Follows immediately from Corollary 5.35 and Lemma 5.43. \blacksquare

Example 5.48. Here's a simple application of Corollary 5.47. Let $\alpha_1, \dots, \alpha_n \in \mathbb{R}$ be distinct real numbers and define

$$v_i := [1 \quad \alpha_i \quad \alpha_i^2 \quad \cdots \quad \alpha_i^{n-1}]^\top \in M_{n \times 1}(\mathbb{R}).$$

We claim that the tuple (v_1, \dots, v_n) is independent. If we were to show this directly, we would have to consider the equation

$$c_1 v_1 + \cdots + c_n v_n = 0,$$

which is equivalent to

$$c_1 \alpha_1^k + \cdots + c_n \alpha_n^k = 0 \quad \text{for all } 0 \leq k \leq n-1. \quad (5.49)$$

From this, we have to deduce that $c_1 = \cdots = c_n = 0$. Unfortunately, equations (5.49) involve all of the variables $\alpha_1, \dots, \alpha_n$ at once and are hard to tackle. To simplify the problem, we first apply Corollary 5.47. Let A be the n -by- n matrix whose columns are v_1, \dots, v_n and let the rows

of A be w_0, \dots, w_{n-1} , so $w_k = [\alpha_1^k \ \alpha_2^k \ \dots \ \alpha_n^k]$. By Corollary 5.47, the tuple (v_1, \dots, v_n) is independent if and only if the tuple (w_0, \dots, w_{n-1}) is. Thus, we need to consider the equation

$$c_0 w_0 + \dots + c_{n-1} w_{n-1} = 0,$$

which is equivalent to

$$c_0 + c_1 \alpha_i + \dots + c_{n-1} \alpha_i^{n-1} = 0 \quad \text{for all } 1 \leq i \leq n. \quad (5.50)$$

Notice that each equation in (5.50) involves only one variable α_i . This means that the polynomial $p(x) := c_0 + c_1 x + \dots + c_{n-1} x^{n-1}$ has n distinct roots. But the degree of p is at most $n - 1$, so it must be the zero polynomial and hence $c_0 = \dots = c_{n-1} = 0$, as desired.

Exercise 5.51. Show that there exists an infinite sequence of vectors $x_0, x_1, \dots \in \mathbb{R}^n$ such that for all $i_1 < i_2 < \dots < i_n$, the tuple $(x_{i_1}, x_{i_2}, \dots, x_{i_n})$ is independent.

Exercise 5.52. Verify the following claims (which were advertised back in §1.F):

- (r1) $\text{rank}(I_n) = n$ for all $n \in \mathbb{N}$;
- (r2) if $A \in M_{m \times n}(F)$, then $\text{rank}(A) \leq \min\{m, n\}$;
- (r3) if $A \in M_{m \times n}(F)$ and $B \in M_{n \times r}(F)$, then $\text{rank}(AB) \leq \min\{\text{rank}(A), \text{rank}(B)\}$.

5.F. Fast matrix multiplication

In view of the correspondence between matrix multiplication and composition of linear functions established in §5.E, an important and natural question arises:

How quickly can we multiply two matrices?

Let R be a ring and let $A, B \in M_{n \times n}(R)$ be two n -by- n matrices over R . We shall view addition and multiplication in the ring R itself to be elementary operations. By definition,

$$(AB)(i, j) = \sum_{k=1}^n A(i, k)B(k, j).$$

Thus, computing a single entry of AB requires adding n terms, each a product of two elements of R . Since the matrix AB has n^2 entries, computing AB directly using the definition takes roughly n^3 steps. Somewhat surprisingly, there exist clever algorithms for matrix multiplication that take *fewer* than a cubic number of steps. The first such algorithm was developed by Volker Strassen in 1969. Instead of n^3 , Strassen's algorithm requires only $O(n^{2.80\dots})$ operations, where the **big-O** symbol O means that the exact number of operations is upper bounded by $n^{2.80\dots}$ times some constant independent of n (which may vary depending on the particular implementation of the algorithm). The exact value of the exponent is $\log_2 7 = 2.80\dots$; we will soon see where this value comes from. After Strassen's seminal work, several improved algorithms have been proposed. In 1990, Don Coppersmith and Shmuel Winograd introduced an algorithm with running time $O(n^{2.375\dots})$. The Coppersmith–Winograd algorithm was unbeaten until 2010; the best currently known matrix multiplication algorithm is a modification of the Coppersmith–Winograd algorithm due to François Le Gall from 2014, with running time $O(n^{2.372\dots})$. Finding the best possible matrix multiplication algorithm is an important open problem in computer science; in particular, the following is an open question:

Open Problem 5.53. Is it true that for every $\varepsilon > 0$, there is an algorithm for multiplying two n -by- n matrices with running time $O(n^{2+\varepsilon})$?

While the later approaches are quite technical and fall beyond the scope of these notes, Strassen's algorithm is actually rather easy to explain.¹⁷ For convenience, assume that $n = 2^k$ is a power of

¹⁷Another advantage of Strassen's algorithm is that it is actually used in practice, whereas the other methods, such as the Coppersmith–Winograd algorithm, only become superior for values of n that are too large to appear in practical applications.

2. The strategy is to reduce computing the product of two n -by- n matrices to multiplying several pairs of $(n/2)$ -by- $(n/2)$ matrices and then proceed recursively. Let $A, B \in M_{n \times n}(R)$. Split each of A and B into four $(n/2)$ -by- $(n/2)$ “blocks,” as follows:

$$A =: \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \quad \text{and} \quad B =: \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix}.$$

If we also write

$$AB =: \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix},$$

then

$$\begin{aligned} C_{11} &= A_{11}B_{11} + A_{12}B_{21}; & C_{12} &= A_{11}B_{12} + A_{12}B_{22}; \\ C_{21} &= A_{21}B_{11} + A_{22}B_{21}; & C_{22} &= A_{21}B_{12} + A_{22}B_{22}. \end{aligned} \tag{5.54}$$

Equations (5.54) reduce computing AB to multiplying 8 pairs of $(n/2)$ -by- $(n/2)$ matrices (matrix addition takes only $O(n^2)$ steps, so its contribution to the total running time of the algorithm is negligible). This means that if we apply formulas (5.54) recursively, then doubling n would increase the running time of the algorithm approximately by a factor of 8, which means that the running time is of the order n^3 . So far, we have not improved on the naïve approach that simply uses the definition of matrix multiplication.

Strassen’s key insight is that to compute C_{11} , C_{12} , C_{21} , and C_{22} , one can get away with only *seven* multiplications instead of eight. Namely, consider the following matrices:

$$\begin{aligned} M_1 &:= (A_{11} + A_{22})(B_{11} + B_{22}); \\ M_2 &:= (A_{21} + A_{22})B_{11}; \\ M_3 &:= A_{11}(B_{12} - B_{22}); \\ M_4 &:= A_{22}(B_{21} - B_{11}); \\ M_5 &:= (A_{11} + A_{12})B_{22}; \\ M_6 &:= (A_{21} - A_{11})(B_{11} + B_{12}); \\ M_7 &:= (A_{12} - A_{22})(B_{21} + B_{22}). \end{aligned}$$

Each of M_1, \dots, M_7 is computed using a simple matrix multiplication (and a few matrix additions), and an easy direct calculation shows that

$$\begin{aligned} C_{11} &= M_1 + M_4 - M_5 + M_7; & C_{12} &= M_3 + M_5; \\ C_{21} &= M_2 + M_4; & C_{22} &= M_1 - M_2 + M_3 + M_6. \end{aligned}$$

Using these formulas recursively, we obtain an algorithm whose running time increases by a factor of 7 each time n doubles, and hence the running time is $O(n^{\log_2 7})$, as desired.

A nice application of fast matrix multiplication is to the following problem:

*Given a graph $G = (V, E)$ on n vertices, how quickly can we determine whether it contains a **triangle**, i.e., a triple or pairwise adjacent vertices?*

A naïve approach is to simply check every possible triple of vertices to see whether they form a triangle. Since there are $\binom{n}{3} \approx n^3/6$ triples to check, this gives an algorithm with running time $O(n^3)$. It turns out that we can do better using matrix multiplication. Denote the vertices of G by v_1, \dots, v_n . Let $A \in M_{n \times n}(\mathbb{Z})$ be the **adjacency matrix** of G , i.e., the n -by- n matrix given by

$$A(i, j) := \begin{cases} 1 & \text{if } v_i v_j \in E; \\ 0 & \text{otherwise.} \end{cases}$$

Compute $B := AA = A^2$. Using fast matrix multiplication, this can be done in sub-cubic time (the precise running time depends on the chosen matrix multiplication algorithm). Notice that


$$B(i, j) = \sum_{k=1}^n A(i, k)A(k, j) = |N_G(v_i) \cap N_G(v_j)|,$$

i.e., $B(i, j)$ is the number of common neighbors of v_i and v_j . Therefore, G contains a triangle if and only if there exist some $1 \leq i, j \leq n$ such that $A(i, j) = 1$ and $B(i, j) > 0$, and whether such a pair (i, j) exists can be determined in time $O(n^2)$. Using the fastest known matrix multiplication algorithm in this procedure leads to the fastest known algorithm for checking whether a graph contains a triangle.

Extra exercises for Section 5

Exercise 5.55. For a set V , let $K(V)$ denote the graph with vertex set V in which there is an edge between every pair of distinct vertices. The graph $K(V)$ is called the **complete graph** on V . Similarly, if U, W are disjoint sets, then $K(U, W)$ is the graph with vertex set $V := U \cup W$ and edge set $\{uv : u \in U, v \in W\}$. The graph $K(U, W)$ is called the **complete bipartite graph with bipartition** (U, W) .

Let G be a graph with edge set E . We say that graphs H_1, \dots, H_k form an **edge decomposition** of G if the edge sets of H_1, \dots, H_k are pairwise disjoint and their union is E . The figure below shows an edge decomposition of a complete graph on 5 vertices into 4 complete bipartite graphs:



- (a) For every $n \geq 1$, show that a complete graph on n vertices admits an edge decomposition into $n - 1$ complete bipartite graphs.

The goal of this exercise is to establish the following result:

Theorem 5.56. Let G be a complete graph on n vertices and suppose that H_1, \dots, H_k are complete bipartite graphs forming an edge decomposition of G . Then $k \geq n - 1$.

For concreteness, assume that the vertex set of G is $\{1, \dots, n\}$. Suppose, towards a contradiction, that $k < n - 1$. For each $1 \leq j \leq k$, let (U_j, W_j) be the bipartition of H_j .

- (b) Show that for every vector $(x_1, \dots, x_n) \in \mathbb{R}^n$,

$$\sum_{i=1}^n x_i^2 = \left(\sum_{i=1}^n x_i \right)^2 - 2 \sum_{j=1}^k \left(\sum_{i \in U_j} x_i \right) \left(\sum_{i \in W_j} x_i \right). \tag{5.57}$$

- (c) Prove that there is a nonzero vector $(x_1, \dots, x_n) \in \mathbb{R}^n$ for which the right-hand side of (5.57) is equal to 0. *Hint:* Use linear algebra.
 (d) Deduce Theorem 5.56.

Exercise 5.58. Let $P_n(\mathbb{R})$ be the \mathbb{R} -vector space of all polynomials with real coefficients in a single variable x of degree at most n . For $q \in P_n(\mathbb{R})$, let

$$\varphi(q) : P_n(\mathbb{R}) \rightarrow \mathbb{R} : p \mapsto \int_0^1 p(x)q(x) dx.$$

- (a) Check that $\varphi(q) \in (P_n(\mathbb{R}))^*$ and that the function $\varphi: P_n(\mathbb{R}) \rightarrow (P_n(\mathbb{R}))^*$ is linear.
 (b) Show that $\ker(\varphi) = \{0\}$ and conclude that the function $\varphi: P_n(\mathbb{R}) \rightarrow (P_n(\mathbb{R}))^*$ is surjective.
 (c) Deduce that for every $\alpha \in \mathbb{R}$, there exists a unique polynomial $q \in P_n(\mathbb{R})$ such that

$$p(\alpha) = \int_0^1 p(x)q(x) dx \quad \text{for all } p \in P_n(\mathbb{R}).$$

Exercise 5.59. Let K be a field and let $F \subseteq K$ be a subfield of K . (Or, in the terminology of Exercise 4.28, K is an extension of F .) Suppose that K is finite-dimensional as an F -vector space and let $n := \dim_F K$. Show that there is a set $R \subseteq M_{n \times n}(F)$ of n -by- n matrices over F such that K , as a field, is isomorphic to R with matrix addition and multiplication. *Hint:* For each $a \in K$, consider the F -linear function $\varphi_a: K \rightarrow K$ given by $\varphi_a(x) := ax$ for all $x \in K$.

6. EXTERIOR ALGEBRA

6.A. Multilinear functions

Definition 6.1. Let V_1, \dots, V_k, W be vector spaces over a field F . A function $f: V_1 \times \dots \times V_k \rightarrow W$ is k -**linear** (or **multilinear** if k is implicit) if for all $1 \leq i \leq k$ and for all $v_1 \in V_1, \dots, v_{i-1} \in V_{i-1}, v_{i+1} \in V_{i+1}, \dots, v_k \in V_k$, the function

$$V_i \rightarrow W: v \mapsto f(v_1, \dots, v_{i-1}, v, v_{i+1}, \dots, v_k)$$

is linear. Informally, $f: V_1 \times \dots \times V_k \rightarrow W$ is k -linear if it is linear separately in each variable.

A function $f: V \rightarrow W$ is 1-linear if and only if it is linear. A function $f: V_1 \times V_2 \rightarrow W$ is 2-linear (or **bilinear**) if the following equations hold for all $x_1, y_1 \in V_1, x_2, y_2 \in V_2$, and $a \in F$:

$$\begin{aligned} f(x_1 + y_1, x_2) &= f(x_1, x_2) + f(y_1, x_2); & f(x_1, x_2 + y_2) &= f(x_1, x_2) + f(x_1, y_2); \\ f(ax_1, x_2) &= f(x_1, ax_2) = af(x_1, x_2). \end{aligned}$$

These equations show that bilinear functions are “multiplication-like”; and indeed, most natural examples of bilinear functions are various multiplication operations.

Example 6.2. Let F be a field. Then the multiplication operation $F \times F \rightarrow F: (a, b) \mapsto ab$ is bilinear. More generally, if V is an F -vector space, then the scalar multiplication $F \times V \rightarrow V: (a, v) \mapsto a \cdot v$ is a bilinear function. Generalizing this even further, if $n, m, k \in \mathbb{N}$, then the matrix multiplication

$$M_{k \times m}(F) \times M_{m \times n}(F) \rightarrow M_{k \times n}(F): (A, B) \mapsto AB$$

is a bilinear function.

Example 6.3. Let V be an F -vector space. Then the function

$$V \times V^* \rightarrow F: (v, f) \mapsto f(v) \tag{6.4}$$

is bilinear. Note that in the finite-dimensional case, this can be viewed as a special case of the previous example, since if $\dim V = n$, then the map (6.4) can be represented by the matrix multiplication

$$M_{1 \times n}(F) \times M_{n \times 1}(F) \rightarrow M_{1 \times 1}(F) \cong F.$$

Example 6.5. In general, it is fairly easy to write down many bilinear functions that do not have any special “meaning” attached to them. For instance, any function of the form

$$F^2 \times F^2 \rightarrow F: ((x_1, x_2), (y_1, y_2)) \mapsto ax_1y_1 + bx_1y_2 + cx_2y_1 + dx_2y_2,$$

where a, b, c, d are fixed parameters from F , is bilinear.

The next useful lemma states that a multilinear map is determined by its values on basis vectors:

Lemma 6.6. Let V_1, \dots, V_k, W be vector spaces over a field F . For each $1 \leq i \leq k$, let $B_i \subseteq V_i$ be a basis for V_i . Suppose that $f, g: V_1 \times \dots \times V_k \rightarrow W$ are k -linear functions such that for all $x_1 \in B_1, \dots, x_k \in B_k$, we have $f(x_1, \dots, x_k) = g(x_1, \dots, x_k)$. Then $f = g$.

PROOF. We will give a proof for $k = 2$, the general case being left as an [exercise](#). Take any $v \in V_1$ and $w \in V_2$. Our goal is to show that $f(v, w) = g(v, w)$. To this end, write

$$\begin{aligned} v &= a_1x_1 + \cdots + a_nx_n, & a_1, \dots, a_n &\in F, & x_1, \dots, x_n &\in B_1; \\ w &= b_1y_1 + \cdots + b_my_m, & b_1, \dots, b_m &\in F, & y_1, \dots, y_m &\in B_2. \end{aligned}$$

Then we have

$$\begin{aligned} f(v, w) &= f(a_1x_1 + \cdots + a_nx_n, w) \\ \text{[since } f \text{ is linear in the first variable]} &= \sum_{i=1}^n a_i f(x_i, w) \\ &= \sum_{i=1}^n a_i f(x_i, b_1y_1 + \cdots + b_my_m) \\ \text{[since } f \text{ is linear in the second variable]} &= \sum_{i=1}^n \sum_{j=1}^m a_i b_j f(x_i, y_j) \\ \text{[since } f \text{ and } g \text{ agree on the basis vectors]} &= \sum_{i=1}^n \sum_{j=1}^m a_i b_j g(x_i, y_j) \\ \text{[since } g \text{ is bilinear]} &= g(v, w). \quad \blacksquare \end{aligned}$$

Exercise 6.7 (important). Let F be a field and let $A \in M_{m \times n}(F)$ be an m -by- n matrix over F . Define a function

$$f_A: M_{m \times 1}(F) \times M_{n \times 1}(F) \rightarrow M_{1 \times 1}(F): (x, y) \mapsto x^\top Ay.$$

Notice that since $M_{m \times 1}(F) \cong F^m$, $M_{n \times 1}(F) \cong F^n$, and $M_{1 \times 1}(F) \cong F$, we could think of f_A as a function from $F^m \times F^n$ to F .

- (a) Show that for each $A \in M_{m \times n}(F)$, the function f_A is bilinear.
- (b) Use Lemma 6.6 to prove that *every* bilinear function

$$f: M_{m \times 1}(F) \times M_{n \times 1}(F) \rightarrow M_{1 \times 1}(F)$$

is equal to f_A for some matrix $A \in M_{m \times n}(F)$.

Example 6.8. For a 2-by-2 matrix $A \in M_{2 \times 2}(F)$, the function f_A acts as follows:

$$\begin{aligned} f_A \left(\left[\begin{array}{c} x_1 \\ x_2 \end{array} \right], \left[\begin{array}{c} y_1 \\ y_2 \end{array} \right] \right) &= \left[\begin{array}{cc} x_1 & x_2 \end{array} \right] \left[\begin{array}{cc} A(1,1) & A(1,2) \\ A(2,1) & A(2,2) \end{array} \right] \left[\begin{array}{c} y_1 \\ y_2 \end{array} \right] \\ &= \left[A(1,1)x_1y_1 + A(1,2)x_1y_2 + A(2,1)x_2y_1 + A(2,2)x_2y_2 \right]. \end{aligned}$$

6.B. Alternating maps

Definition 6.9. Let V and W be F -vector spaces. A k -linear map $f: V^k \rightarrow W$ is **alternating** if for all $x_1, \dots, x_k \in V$, we have $f(x_1, \dots, x_k) = 0$ whenever the vectors x_1, \dots, x_k are *not* pairwise distinct, i.e., when there exist $1 \leq i < j \leq k$ such that $x_i = x_j$.

Example 6.10. The function $V^k \rightarrow W$ that sends every tuple (x_1, \dots, x_k) to zero is alternating. (This example illustrates that an alternating function *can* be zero even if the inputs are distinct.)

Example 6.11. Every linear (i.e., 1-linear) function is, vacuously, alternating (since it accepts only a single input).

Example 6.12. Let F be a field and consider the function $f: F^2 \times F^2 \rightarrow F$ given by

$$f((x_1, x_2), (y_1, y_2)) := x_1y_2 - x_2y_1.$$

This function is clearly bilinear. Furthermore, it is alternating, since

$$f((x_1, x_2), (x_1, x_2)) = x_1x_2 - x_2x_1 = 0.$$

The name “alternating” is explained by the following lemma:

Lemma 6.13. *Let V and W be F -vector spaces and suppose that $f: V^k \rightarrow W$ is an alternating k -linear map. Then for any $x_1, \dots, x_k \in V$ and $1 \leq i < j \leq k$,*

$$f(x_1, \dots, x_k) = -f(x_1, \dots, x_{i-1}, x_j, x_{i+1}, \dots, x_{j-1}, x_i, x_{j+1}, \dots, x_k);$$

in other words, switching any two inputs changes the sign of f .

PROOF. We will give a proof for $k = 2$, leaving the general case as an [exercise](#). Our goal is to show that for all $x, y \in V$, $f(x, y) = -f(y, x)$. To that end, consider the value $f(x + y, x + y)$. Since f is alternating, $f(x + y, x + y) = 0$. On the other hand,

$$f(x + y, x + y) = f(x, x) + f(x, y) + f(y, x) + f(y, y) = f(x, y) + f(y, x),$$

where we used that $f(x, x) = f(y, y) = 0$. Hence, $f(x, y) + f(y, x) = 0$, as desired. ■

The importance of alternating functions lies in the following fact:

Lemma 6.14. *Let V and W be F -vector spaces and let $f: V^k \rightarrow W$ be an alternating k -linear map. Then for all $x_1, \dots, x_k \in V$, we have $f(x_1, \dots, x_k) = 0$ whenever the tuple (x_1, \dots, x_k) is not independent.*

PROOF. Suppose $x_1, \dots, x_k \in V$ are vectors such that the tuple (x_1, \dots, x_k) is not independent. This means that one of the vectors x_1, \dots, x_k can be expressed as a linear combination of the other ones ([exercise!](#)). For concreteness, assume that $x_k = a_1x_1 + \dots + a_{k-1}x_{k-1}$. Then we have

$$f(x_1, \dots, x_k) = f(x_1, \dots, x_{k-1}, a_1x_1 + \dots + a_{k-1}x_{k-1}) = \sum_{i=1}^{k-1} a_i f(x_1, \dots, x_{k-1}, x_i) = 0. \quad \blacksquare$$

Our goal is to construct multilinear functions for which the *converse* of Lemma 6.14 holds, i.e., such that $f(x_1, \dots, x_k) = 0$ *if and only if* the tuple (x_1, \dots, x_k) is not independent. Furthermore, the multilinear functions we construct will be given by iteratively applying a certain associative binary operation. Without further ado, let us state the main result of this section:

Theorem/Definition 6.15 (Exterior products). *Let F be a field and let V be an F -vector space of dimension n . Then there exist disjoint F -vector spaces V_1, V_2, \dots and an associative binary operation \wedge on $V_1 \cup V_2 \cup \dots$, called the **exterior product** (or **wedge product**) such that:*

- (EP1) $V_1 = V$.
- (EP2) For each $k \geq 1$, $\dim V_k = \binom{n}{k}$. In particular, $\dim V_m = 0$ for all $m > n$.
- (EP3) If $x \in V_k$ and $y \in V_\ell$, then $x \wedge y \in V_{k+\ell}$.
- (EP4) For each $k \geq 1$, we have $V_k = \text{Span}\{x_1 \wedge \dots \wedge x_k : x_1, \dots, x_k \in V\}$.
- (EP5) For all $k, \ell \geq 1$, the map $V_k \times V_\ell \rightarrow V_{k+\ell}: (x, y) \mapsto x \wedge y$ is bilinear.
- (EP6) The map $V \times V \rightarrow V_2: (x, y) \mapsto x \wedge y$ is alternating.
- (EP7) For all $x_1, \dots, x_k \in V$, the tuple (x_1, \dots, x_k) is independent if and only if $x_1 \wedge \dots \wedge x_k \neq 0$.
- (EP8) For every F -vector space W and an alternating k -linear map $f: V^k \rightarrow W$, there is a unique linear function $\varphi: V_k \rightarrow W$ such that for all $x_1, \dots, x_k \in V$, we have

$$f(x_1, \dots, x_k) = \varphi(x_1 \wedge \dots \wedge x_k).$$

For each $k \geq 1$, the space V_k is referred to as the **k -th exterior power** of V and is denoted by $\bigwedge^k V$.

6.C. Discussion of Theorem 6.15

For now, let's take the existence of the structure described in Theorem 6.15 for granted and see what consequences we can draw from it. We start with two basic remarks:

Remark 6.16. It follows from (EP5) and (EP6) that for each $k \geq 1$, the map

$$V^k \rightarrow V_k: (x_1, \dots, x_k) \mapsto x_1 \wedge \dots \wedge x_k$$

is k -linear and alternating. ([Exercise](#): prove this!)

Remark 6.17. The structure $(V_1, V_2, \dots; \wedge)$ is “essentially” unique, meaning that if W_1, W_2, \dots is another sequence of disjoint F -vector spaces equipped with an associative binary operation $\tilde{\wedge}$ satisfying (EP1)–(EP8), then there exist unique isomorphisms $\varphi_k: V_k \rightarrow W_k$ that send \wedge to $\tilde{\wedge}$. (This is why we can refer to *the* k -th exterior power of V .) Indeed, the function

$$V^k \rightarrow W_k: (x_1, \dots, x_k) \mapsto x_1 \tilde{\wedge} \dots \tilde{\wedge} x_k$$

is k -linear and alternating, so, by (EP8), there is a unique linear map $\varphi_k: V_k \rightarrow W_k$ such that

$$x_1 \tilde{\wedge} \dots \tilde{\wedge} x_k = \varphi_k(x_1 \wedge \dots \wedge x_k).$$

Similarly, the map

$$V^k \rightarrow V_k: (x_1, \dots, x_k) \mapsto x_1 \wedge \dots \wedge x_k$$

is also alternating, so, by (EP8) applied to $\tilde{\wedge}$, there is a unique linear map $\psi_k: W_k \rightarrow V_k$ with

$$x_1 \wedge \dots \wedge x_k = \psi_k(x_1 \tilde{\wedge} \dots \tilde{\wedge} x_k).$$

Since, by (EP4),

$$V_k = \text{Span}\{x_1 \wedge \dots \wedge x_k : x_1, \dots, x_k \in V\} \quad \text{and} \quad W_k = \text{Span}\{x_1 \tilde{\wedge} \dots \tilde{\wedge} x_k : x_1, \dots, x_k \in V\},$$

we conclude that the functions φ_k and ψ_k are inverses of each other, and hence they are isomorphisms.

Let us now consider some low-dimensional examples.

Example 6.18 (The 2-dimensional case). Suppose that V is 2-dimensional; for concreteness, let $V = M_{2 \times 1}(F)$. Then $\dim V_2 = \binom{2}{2} = 1$, so we can take $V_2 = F$. Now we need to define a bilinear function $\wedge: V \times V \rightarrow F$ such that a pair of vectors $(x, y) \in V \times V$ is independent if and only if $x \wedge y \neq 0$. We claim that the following function works:

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \wedge \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} := x_1 y_2 - x_2 y_1.$$

Indeed, with this definition, we have $x \wedge y = 0$ if and only if $x_1 y_2 = x_2 y_1$. If $y_1, y_2 \neq 0$, then this means that $x_1/y_1 = x_2/y_2 := c$, and hence $x = cy$, so the pair (x, y) is not independent.

Exercise 6.19. Finish this argument (i.e., consider the cases when one or both of y_1, y_2 are zero).

Example 6.20 (The 3-dimensional case). Suppose that $\dim V = 3$; for concreteness, assume that $V = M_{3 \times 1}(F)$. Now both V_2 and V_3 are nontrivial, with $\dim V_2 = \binom{3}{2} = 3$ and $\dim V_3 = \binom{3}{3} = 1$. Take $V_2 = M_{3 \times 1}(F)$ and $V_3 = F$. We have to define a bilinear function $\wedge: V \times V \rightarrow V_2$ such that a pair of vectors $(x, y) \in V \times V$ is independent if and only if $x \wedge y \neq 0$. Here's the idea: Suppose that the pair (x, y) is *not* independent, where

$$x = \begin{bmatrix} x_1 & x_2 & x_3 \end{bmatrix}^\top \quad \text{and} \quad y = \begin{bmatrix} y_1 & y_2 & y_3 \end{bmatrix}^\top.$$

Then the following three pairs of vectors in $M_{2 \times 1}(F)$ are also not independent:

$$\left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}, \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} \right), \quad \left(\begin{bmatrix} x_1 \\ x_3 \end{bmatrix}, \begin{bmatrix} y_1 \\ y_3 \end{bmatrix} \right), \quad \text{and} \quad \left(\begin{bmatrix} x_2 \\ x_3 \end{bmatrix}, \begin{bmatrix} y_2 \\ y_3 \end{bmatrix} \right). \quad (6.21)$$

It turns out that the converse implication also holds:

Exercise 6.22. Show that if the three pairs of vectors in (6.21) are not independent, then neither is the pair (x, y) .

Note that we do need all three pairs in (6.21): For instance, if

$$x = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}^\top \quad \text{and} \quad y = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^\top,$$

then the pair (x, y) is independent, even though the two pairs

$$\left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}, \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} \right) = \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \end{bmatrix} \right) \quad \text{and} \quad \left(\begin{bmatrix} x_1 \\ x_3 \end{bmatrix}, \begin{bmatrix} y_1 \\ y_3 \end{bmatrix} \right) = \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right)$$

are not. Combining Exercise 6.22 with Example 6.18, we conclude that the following definition gives a bilinear function that correctly “detects” independence:

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \wedge \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} := \begin{bmatrix} x_1y_2 - x_2y_1 \\ x_1y_3 - x_3y_1 \\ x_2y_3 - x_3y_2 \end{bmatrix}. \quad (6.23)$$

Now we *also* have to find bilinear maps $\wedge: V \times V_3 \rightarrow F$ and $\wedge: V_2 \times V \rightarrow F$ so that for all $x, y, z \in V$, $(x \wedge y) \wedge z = x \wedge (y \wedge z) := x \wedge y \wedge z$, and $x \wedge y \wedge z \neq 0$ if and only if (x, y, z) is a basis for V . It is not at all obvious how to achieve this (and whether it is even *possible*) by writing an explicit formula such as (6.23). In the proof of Theorem 6.15 that we will give, we shall try to completely avoid such “numerical” expressions and instead construct exterior products “abstractly.” Nevertheless, *after* proving Theorem 6.15, we will be able to derive such explicit formulas. Here are the ones for the 3-dimensional case: If

$$x = \begin{bmatrix} x_1 & x_2 & x_3 \end{bmatrix}^\top \in V \quad \text{and} \quad u = \begin{bmatrix} u_1 & u_2 & u_3 \end{bmatrix}^\top \in V_2,$$

then we can set

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \wedge \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} := x_1u_3 - x_2u_2 + x_3u_1 \quad \text{and} \\ \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} \wedge \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} := u_1x_3 - u_2x_2 + u_3x_1.$$

6.D. Parity of permutations

Let V and W be vector spaces over a field F and suppose that $f: V^3 \rightarrow W$ is an alternating 3-linear function. For any $x_1, x_2, x_3 \in V$, there are six ways to plug them into f :

$$f(x_1, x_2, x_3), \quad f(x_1, x_3, x_2), \quad f(x_2, x_1, x_3), \quad f(x_2, x_3, x_1), \quad f(x_3, x_1, x_2), \quad \text{and} \quad f(x_3, x_2, x_1).$$

Using Lemma 6.13, we can express all six of these in terms of $f(x_1, x_2, x_3)$:

$$\begin{aligned} & f(x_1, x_2, x_3); \\ & f(x_1, x_3, x_2) = -f(x_1, x_2, x_3); \\ & f(x_2, x_1, x_3) = -f(x_1, x_2, x_3); \\ & f(x_2, x_3, x_1) = -f(x_1, x_3, x_2) = f(x_1, x_2, x_3); \\ & f(x_3, x_1, x_2) = -f(x_1, x_3, x_2) = f(x_1, x_2, x_3); \\ & f(x_3, x_2, x_1) = -f(x_1, x_2, x_3). \end{aligned}$$

To generalize this to alternating k -linear maps for arbitrary $k \geq 1$, we need a quick review of permutations and their parity.

Let X be a finite set. A **permutation** of (or on) X is a bijection $\sigma: X \rightarrow X$. The set of all permutations of X is denoted by $\text{Sym}(X)$ and is called the **symmetric group** of X . The **product** of two permutations $\sigma, \pi \in \text{Sym}(X)$ is defined by $\sigma\pi := \sigma \circ \pi$:

$$\begin{array}{ccccc} X & \xrightarrow{\pi} & X & \xrightarrow{\sigma} & X \\ & \searrow & & \nearrow & \\ & & \sigma \circ \pi = \sigma\pi & & \end{array}$$

This operation indeed makes $\text{Sym}(X)$ into a *group*: Composition of permutations is clearly associative; the identity element of $\text{Sym}(X)$ is the identity map $\text{id}_X: X \rightarrow X$; and every $\sigma \in \text{Sym}(X)$ has an inverse $\sigma^{-1} \in \text{Sym}(X)$ because σ a bijection.

Example 6.24. Let $X := \{1, 2, 3\}$ and consider the permutations $\sigma, \pi \in \text{Sym}(X)$ given by

$$\begin{aligned} \sigma &: 1 \mapsto 2, \quad 2 \mapsto 3, \quad 3 \mapsto 1; \\ \pi &: 1 \mapsto 3, \quad 2 \mapsto 2, \quad 3 \mapsto 1. \end{aligned}$$

Then $\sigma\pi$ is the permutation given by

$$\sigma\pi : 1 \mapsto 1, \quad 2 \mapsto 3, \quad 3 \mapsto 2,$$

and σ^{-1} is given by

$$\sigma^{-1} : 1 \mapsto 3, \quad 2 \mapsto 1, \quad 3 \mapsto 2.$$

Definition 6.25. A **transposition** on a finite set X is a permutation $\tau \in \text{Sym}(X)$ that interchanges two elements of X while keeping the rest of the elements fixed. Explicitly, for $i, j \in X, i \neq j$, the **transposition of i and j** is the permutation $\tau_{ij} \in \text{Sym}(X)$ such that

$$\tau_{ij}(i) := j, \quad \tau_{ij}(j) := i, \quad \text{and} \quad \tau_{ij}(x) := x \text{ for all } x \notin \{i, j\}.$$

The key to understanding exterior products lies in the following basic combinatorial fact, which we will leave as an exercise (it is usually covered in most abstract algebra courses):

Exercise 6.26 (important). Let X be a finite set and let $\sigma \in \text{Sym}(X)$ be a permutation. Prove the following statements.

- (a) There is a finite sequence of transpositions τ_1, \dots, τ_k such that $\sigma = \tau_1 \cdots \tau_k$.
- (b) If $\tau_1, \dots, \tau_k, \rho_1, \dots, \rho_\ell$ are transpositions such that

$$\sigma = \tau_1 \cdots \tau_k = \rho_1 \cdots \rho_\ell,$$

then $k = \ell \pmod{2}$.

Example 6.27. Let $X := \{1, 2, 3, 4, 5\}$ and suppose that $\sigma \in \text{Sym}(X)$ is given by

$$\sigma : 1 \mapsto 2, \quad 2 \mapsto 3, \quad 3 \mapsto 1, \quad 4 \mapsto 5, \quad 5 \mapsto 4,$$

then $\sigma = \tau_{13}\tau_{12}\tau_{45} = \tau_{23}\tau_{34}\tau_{13}\tau_{15}\tau_{14}$ (**exercise!**), i.e., σ can be expressed as a product of 3 or 5 transpositions. However, in accordance with Exercise 6.26, it is impossible to write σ as a product of 4, 6, or any other *even* number of transpositions.

Definition 6.28. Let X be a finite set and let $\sigma \in \text{Sym}(X)$. Write σ as a product of transpositions: $\sigma = \tau_1 \cdots \tau_k$. Then the **sign** of σ is $\text{sign}(\sigma) := (-1)^k$. (The sign of σ is well-defined since it only depends on the *parity* of k .) If $\text{sign}(\sigma) = 1$, then we say that σ is **even**; otherwise, σ is **odd**.

Exercise 6.29. Show that for any $\sigma, \pi \in \text{Sym}(X)$, we have $\text{sign}(\sigma\pi) = \text{sign}(\sigma)\text{sign}(\pi)$.

Exercise 6.30. Show that for any $\sigma \in \text{Sym}(X)$, $\text{sign}(\sigma^{-1}) = \text{sign}(\sigma)$.

Exercise 6.31. Show that if X is a finite set of size at least 2, then the number of even permutations in $\text{Sym}(X)$ is equal to the number of odd permutations in $\text{Sym}(X)$. *Hint:* Take any transposition $\tau \in \text{Sym}(X)$ and consider the map $\text{Sym}(X) \rightarrow \text{Sym}(X): \sigma \mapsto \tau\sigma$.

The following statement is an immediate consequence of Lemma 6.13 that generalizes the discussion from the beginning of this subsection:

Exercise 6.32. Let V and W be F -vector spaces and suppose that $f: V^k \rightarrow W$ is an alternating k -linear map. Show that for any $x_1, \dots, x_k \in V$ and $\sigma \in \text{Sym}(\{1, \dots, k\})$,

$$f(x_{\sigma(1)}, \dots, x_{\sigma(k)}) = \text{sign}(\sigma) \cdot f(x_1, \dots, x_k).$$

From Exercise 6.32 and assuming the existence of exterior products (which we haven't proved yet), we can derive a useful corollary that will guide our construction of exterior powers:

Corollary 6.33. Let F be a field and let V be an F -vector space of dimension n . Pick a basis $\{e_1, \dots, e_n\}$ for V . Then the set $\{e_{i_1} \wedge \dots \wedge e_{i_k} : 1 \leq i_1 < \dots < i_k \leq n\}$ is a basis for $\bigwedge^k V$.

PROOF. By (EP4), $\bigwedge^k V = \text{Span}(\{x_1 \wedge \dots \wedge x_k : x_1, \dots, x_k \in V\})$. Consider any element of the form $x_1 \wedge \dots \wedge x_k$. Since $\{e_1, \dots, e_n\}$ is a basis for V , we can express each of x_1, \dots, x_k as a linear combination of e_1, \dots, e_n . If we replace the vectors x_1, \dots, x_k in $x_1 \wedge \dots \wedge x_k$ by these linear combinations and expand use the bilinearity of \wedge , we obtain that

$$x_1 \wedge \dots \wedge x_k \in \text{Span}(\{e_{i_1} \wedge \dots \wedge e_{i_k} : 1 \leq i_1, \dots, i_k \leq n\}).$$

Now consider an expression of the form $e_{i_1} \wedge \dots \wedge e_{i_k}$. If the indices i_1, \dots, i_k are not distinct, then $e_{i_1} \wedge \dots \wedge e_{i_k} = 0$. If, on the other hand, they are distinct, then, by Exercise 6.32, $e_{i_1} \wedge \dots \wedge e_{i_k}$ is equal to plus or minus one times the wedge product of e_{i_1}, \dots, e_{i_k} taken in the increasing order of indices. This shows that the set

$$\{e_{i_1} \wedge \dots \wedge e_{i_k} : 1 \leq i_1 < \dots < i_k \leq n\}$$

spans $\bigwedge^k V$. Since the size of this set is equal to $\binom{n}{k} = \dim \bigwedge^k V$, it must be a basis. ■

Example 6.34. Let $V = M_{3 \times 1}(\mathbb{R})$ and let

$$e_1 = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^\top, \quad e_2 = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}^\top, \quad \text{and} \quad e_3 = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^\top$$

be the standard basis vectors for V . Let's use this basis to compute

$$\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \wedge \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \wedge \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} = (e_1 + e_2) \wedge (e_1 + e_3) \wedge (e_2 + e_3).$$

We have

$$\begin{aligned} (e_1 + e_2) \wedge (e_1 + e_3) &= e_1 \wedge e_1 + e_1 \wedge e_3 + e_2 \wedge e_1 + e_2 \wedge e_3 \\ &= e_1 \wedge e_3 - e_1 \wedge e_2 + e_2 \wedge e_3. \end{aligned}$$

Therefore,

$$\begin{aligned} (e_1 + e_2) \wedge (e_1 + e_3) \wedge (e_2 + e_3) &= (e_1 \wedge e_3 - e_1 \wedge e_2 + e_2 \wedge e_3) \wedge (e_2 + e_3) \\ &= e_1 \wedge e_3 \wedge e_2 - e_1 \wedge e_2 \wedge e_3 = -2e_1 \wedge e_2 \wedge e_3. \end{aligned}$$

The significance of the coefficient -2 will be explained in the next subsection.

6.E. Determinants

Let V be an n -dimensional vector space over a field F and let $\varphi: V \rightarrow V$ be a linear function. Consider the map $V^n \rightarrow \bigwedge^n V$ given by

$$(x_1, \dots, x_n) \mapsto \varphi(x_1) \wedge \dots \wedge \varphi(x_n).$$

This map is n -linear and alternating (**exercise!**), so, by (EP8), there is a unique linear function $f: \bigwedge^n V \rightarrow \bigwedge^n V$ such that $\varphi(x_1) \wedge \dots \wedge \varphi(x_n) = f(x_1 \wedge \dots \wedge x_n)$. Since $\dim \bigwedge^n V = \binom{n}{n} = 1$, the only linear functions from $\bigwedge^n V$ to itself are of the form $v \mapsto a \cdot v$ for a fixed $a \in F$; therefore there exists a unique element of F , called the **determinant** of φ and denoted by $\det(\varphi)$, such that

$$\varphi(x_1) \wedge \dots \wedge \varphi(x_n) = \det(\varphi) \cdot x_1 \wedge \dots \wedge x_n.$$

If $A \in M_{n \times n}(F)$, then A represents a linear map $M_{n \times 1}(F) \rightarrow M_{n \times 1}(F): x \mapsto Ax$, and we denote the determinant of this map by $\det(A)$ and call it the **determinant** of A . Note that if $\{e_1, \dots, e_n\}$ is the standard basis for $M_{n \times 1}(F)$, then, by definition,

$$Ae_1 \wedge \dots \wedge Ae_n = \det(A) \cdot e_1 \wedge \dots \wedge e_n,$$

and thus calculating $\det(A)$ is tantamount to computing the wedge product of the columns of A .

Example 6.35. Recall from Example 6.34 that

$$\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \wedge \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \wedge \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} = -2e_1 \wedge e_2 \wedge e_3.$$

This means that, by definition,

$$\det \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} = -2.$$

Notice that our definition of the determinant is *coordinate-free*: We don't need to pick a basis for V in order to determine $\det(\varphi)$ for $\varphi: V \rightarrow V$. However, by generalizing the calculation from Example 6.34, it is possible to obtain an explicit numerical expression for the determinant:

Theorem 6.36 (Leibniz formula). *Let F be a field and let $A \in M_{n \times n}(F)$. Then*

$$\det(A) = \sum_{\sigma} \text{sign}(\sigma) \prod_{i=1}^n A(i, \sigma(i)), \tag{6.37}$$

where the summation is over all $\sigma \in \text{Sym}(\{1, \dots, n\})$.

PROOF SKETCH. Let $\{e_1, \dots, e_n\}$ be the standard basis for $M_{n \times 1}(F)$. Then the exterior product of the columns of A can be written as

$$\begin{aligned} Ae_1 \wedge \dots \wedge Ae_n &= (A(1,1)e_1 + A(2,1)e_2 + \dots + A(n,1)e_n) \\ &\quad \wedge (A(1,2)e_1 + A(2,2)e_2 + \dots + A(n,2)e_n) \\ &\quad \wedge \dots \\ &\quad \dots \\ &\quad \wedge (A(1,n)e_1 + A(2,n)e_2 + \dots + A(n,n)e_n). \end{aligned}$$

Expanding this product (using that \wedge is bilinear and alternating) gives (6.37). ■

Exercise 6.38. Fill in the details in the proof of Theorem 6.36.

Exercise 6.39. Use Theorem 6.36 to show that for all $A \in M_{n \times n}(F)$, $\det(A^\top) = \det(A)$.

Example 6.40. There are precisely two permutations of the set $\{1, 2\}$, namely $(1 \mapsto 1, 2 \mapsto 2)$ and $(1 \mapsto 2, 2 \mapsto 1)$. The first of these permutations (the identity) is even, while the second one is odd. Therefore, by the Leibniz formula, the determinant of a 2-by-2 matrix A is

$$\det(A) = A(1, 1)A(2, 2) - A(1, 2)A(2, 1).$$

Similarly, there are six permutations of the set $\{1, 2, 3\}$, three of which are even and three odd, and the determinant of a 3-by-3 matrix A is

$$\begin{aligned} \det(A) &= A(1, 1)A(2, 2)A(3, 3) - A(1, 1)A(2, 3)A(3, 2) - A(1, 2)A(2, 1)A(3, 3) \\ &\quad + A(1, 2)A(2, 3)A(3, 1) + A(1, 3)A(2, 1)A(3, 2) - A(1, 3)A(2, 2)A(3, 1). \end{aligned}$$

In principle, the Leibniz formula could be used to *define* the determinant of a matrix. However, our “coordinate-free” definition, which is more abstract, has some significant advantages, as it allows us to derive several useful properties of determinants almost effortlessly.

Theorem 6.41. *Let V be an n -dimensional vector space over a field F and let $\varphi: V \rightarrow V$ be a linear function. Then φ is bijective if and only if $\det(\varphi) \neq 0$.*

PROOF. Let $\{e_1, \dots, e_n\}$ be a basis for V . Then the function φ is bijective if and only if the set $\{\varphi(e_1), \dots, \varphi(e_n)\}$ is a basis for V , i.e., if the tuple $(\varphi(e_1), \dots, \varphi(e_n))$ is independent. By (EP7), this is equivalent to $\varphi(e_1) \wedge \dots \wedge \varphi(e_n) \neq 0$. But $\varphi(e_1) \wedge \dots \wedge \varphi(e_n) = \det(\varphi) \cdot e_1 \wedge \dots \wedge e_n$ by definition, and so $\varphi(e_1) \wedge \dots \wedge \varphi(e_n) \neq 0$ if and only if $\det(\varphi) \neq 0$, as desired. ■

Lemma 6.42. *Let V be an n -dimensional vector space over a field F and let $\varphi, \psi: V \rightarrow V$ be linear functions. Then $\det(\varphi \circ \psi) = \det(\varphi) \det(\psi)$. Hence, if $A, B \in M_{n \times n}(F)$ are n -by- n matrices over F , then $\det(AB) = \det(A) \det(B)$.*

PROOF. Let $x_1, \dots, x_n \in V$ and let $y_i := \psi(x_i)$ for each $1 \leq i \leq n$. We have

$$\begin{aligned} (\varphi \circ \psi)(x_1) \wedge \dots \wedge (\varphi \circ \psi)(x_n) &= \varphi(\psi(x_1)) \wedge \dots \wedge \varphi(\psi(x_n)) \\ &= \varphi(y_1) \wedge \dots \wedge \varphi(y_n) \\ &= \det(\varphi) \cdot y_1 \wedge \dots \wedge y_n \\ &= \det(\varphi) \cdot \psi(x_1) \wedge \dots \wedge \psi(x_n) \\ &= \det(\varphi) \det(\psi) \cdot x_1 \wedge \dots \wedge x_n, \end{aligned}$$

as desired. ■

Our next result is an extension of Lemma 6.42 to non-square matrices. Before we state it, let us fix some notation. For $n \in \mathbb{N}$, let $[n] := \{1, \dots, n\}$. For a set X and $k \in \mathbb{N}$, let $\mathcal{P}_k(X)$ be the set of all k -element subsets of X . Notice that

$$|\mathcal{P}_k([n])| = \binom{n}{k}.$$

For an m -by- n matrix A and a pair of subsets $S \subseteq [m]$, $T \subseteq [n]$, let $A_{S,T}$ be the matrix obtained from A by only keeping the entries in the rows whose indices are in S and the columns whose indices are in T ; more formally, if $S = \{s_1, \dots, s_k\}$ and $T = \{t_1, \dots, t_\ell\}$, where

$$s_1 < \dots < s_k \quad \text{and} \quad t_1 < \dots < t_\ell,$$

then $A_{S,T}$ is the k -by- ℓ matrix such that

$$A_{S,T}(i, j) := A(s_i, t_j) \quad \text{for all } i \in [k] \text{ and } j \in [\ell].$$

For instance,

$$\text{if } A = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 3 & 1 & 4 & 5 \\ 0 & 1 & 0 & 1 \end{bmatrix} \quad \text{and } S = \{1, 3\}, T = \{1, 2, 4\}, \quad \text{then } A_{S,T} = \begin{bmatrix} 1 & 2 & 4 \\ 0 & 1 & 1 \end{bmatrix}.$$

Theorem 6.43 (Binet–Cauchy formula). *Let F be a field and let m, n be positive integers with $m \geq n$. Then, for all $A \in M_{n \times m}(F)$ and $B \in M_{m \times n}(F)$, we have*

$$\det(AB) = \sum_{S \in \mathcal{P}_n([m])} \det(A_{[n],S}) \det(B_{S,[n]}). \quad (6.44)$$

Example 6.45. Working over \mathbb{R} , consider the matrices

$$A = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 1 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 1 & 1 \\ 2 & 2 \\ 3 & 4 \end{bmatrix}.$$

Then

$$\det(AB) = \det \begin{bmatrix} 14 & 17 \\ 5 & 6 \end{bmatrix} = 14 \cdot 6 - 5 \cdot 17 = -1.$$

On the other hand, the right-hand side of (6.44) is

$$\begin{aligned} & \det(A_{[2],\{1,2\}}) \det(B_{\{1,2\},[2]}) + \det(A_{[2],\{1,3\}}) \det(B_{\{1,3\},[2]}) + \det(A_{[2],\{2,3\}}) \det(B_{\{2,3\},[2]}) \\ &= \det \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix} \det \begin{bmatrix} 1 & 1 \\ 2 & 2 \end{bmatrix} + \det \begin{bmatrix} 1 & 3 \\ 0 & 1 \end{bmatrix} \det \begin{bmatrix} 1 & 1 \\ 3 & 4 \end{bmatrix} + \det \begin{bmatrix} 2 & 3 \\ 1 & 1 \end{bmatrix} \det \begin{bmatrix} 2 & 2 \\ 3 & 4 \end{bmatrix} \\ &= (1 \cdot 1 - 0 \cdot 2) \cdot (1 \cdot 2 - 2 \cdot 1) + (1 \cdot 1 - 0 \cdot 3) \cdot (1 \cdot 4 - 3 \cdot 1) + (2 \cdot 1 - 1 \cdot 3) \cdot (2 \cdot 4 - 3 \cdot 2) \\ &= 1 \cdot 0 + 1 \cdot 1 + (-1) \cdot 2 = 0 + 1 - 2 = -1, \end{aligned}$$

in agreement with Theorem 6.43.

To establish Theorem 6.43, we will need two lemmas that are also interesting and useful in their own right. Let V be an n -dimensional vector space over a field F and let $\{e_1, \dots, e_n\}$ be a basis for V . For a set $S = \{s_1, \dots, s_k\} \subseteq [n]$, where $s_1 < \dots < s_k$, write

$$e_S := e_{s_1} \wedge \dots \wedge e_{s_k}. \quad (6.46)$$

By Corollary 6.33, the set $\{e_S : S \in \mathcal{P}_k([n])\}$ is a basis for $\bigwedge^k V$.

Lemma 6.47. *Let F be a field and let m, n be positive integers with $m \geq n$. Set $V := M_{n \times 1}(F)$ and $W := M_{m \times 1}(F)$. Let $\{e_1, \dots, e_n\}$ and $\{f_1, \dots, f_m\}$ be the standard bases for V and W , respectively. Let $A \in M_{n \times m}(F)$ and let $\varphi_A: \bigwedge^n W \rightarrow \bigwedge^n V$ be the unique linear function such that*

$$A y_1 \wedge \dots \wedge A y_n = \varphi_A(y_1 \wedge \dots \wedge y_n) \quad \text{for all } y_1, \dots, y_n \in W.$$

(Such a function φ_A exists due to (EP8).) Then, for all $S \in \mathcal{P}_n([m])$, we have

$$\varphi_A(f_S) = \det(A_{[n],S}) \cdot e_1 \wedge \dots \wedge e_n.$$

PROOF. The proof of this lemma is shorter than its statement. Take any $S \in \mathcal{P}_n([m])$ and let the elements of S be $s_1 < \dots < s_n$. Then $f_S = f_{s_1} \wedge \dots \wedge f_{s_n}$, so

$$\varphi_A(f_S) = \varphi_A(f_{s_1} \wedge \dots \wedge f_{s_n}) = A f_{s_1} \wedge \dots \wedge A f_{s_n}.$$

But the vectors $A f_{s_1}, \dots, A f_{s_n} \in V$ are precisely the columns of $A_{[n],S}$, and hence

$$A f_{s_1} \wedge \dots \wedge A f_{s_n} = \det(A_{[n],S}) \cdot e_1 \wedge \dots \wedge e_n,$$

as desired. ■

Lemma 6.48. *Let F be a field and let m, n be positive integers with $m \geq n$. Set $V := M_{n \times 1}(F)$ and $W := M_{m \times 1}(F)$. Let $\{e_1, \dots, e_n\}$ and $\{f_1, \dots, f_m\}$ be the standard bases for V and W , respectively. Let $B \in M_{m \times n}(F)$ and let $\varphi_B: \bigwedge^n V \rightarrow \bigwedge^n W$ be the unique linear function such that*

$$B x_1 \wedge \dots \wedge B x_n = \varphi_B(x_1 \wedge \dots \wedge x_n) \quad \text{for all } x_1, \dots, x_n \in V.$$

(Such a function φ_B exists due to (EP8).) Then we have

$$\varphi_B(e_1 \wedge \dots \wedge e_n) = \sum_{S \in \mathcal{P}_n([m])} \det(B_{S,[n]}) \cdot f_S.$$

Exercise 6.49. Prove Lemma 6.48.

PROOF OF THEOREM 6.43. Set $V := M_{n \times 1}(F)$ and $W := M_{m \times 1}(F)$, and let $\{e_1, \dots, e_n\}$ and $\{f_1, \dots, f_m\}$ be the standard bases for V and W , respectively. Also, let $\varphi_A: \bigwedge^n W \rightarrow \bigwedge^n V$ and $\varphi_B: \bigwedge^n V \rightarrow \bigwedge^n W$ be as in Lemmas 6.47 and 6.48. For each $1 \leq i \leq n$, let $y_i := Be_i$ be the i -th column of B . Then we have

$$\begin{aligned} ABe_1 \wedge \dots \wedge ABe_n &= Ay_1 \wedge \dots \wedge Ay_n = \varphi_A(y_1 \wedge \dots \wedge y_n) \\ &= \varphi_A(Be_1 \wedge \dots \wedge Be_n) = \varphi_A(\varphi_B(e_1 \wedge \dots \wedge e_n)). \end{aligned}$$

Now, using Lemma 6.48 and the linearity of φ_A , we obtain

$$\varphi_A(\varphi_B(e_1 \wedge \dots \wedge e_n)) = \varphi_A \left(\sum_{S \in \mathcal{P}_n([m])} \det(B_{S,[n]}) \cdot f_S \right) = \sum_{S \in \mathcal{P}_n([m])} \det(B_{S,[n]}) \cdot \varphi_A(f_S).$$

Finally, by Lemma 6.47,

$$\sum_{S \in \mathcal{P}_n([m])} \det(B_{S,[n]}) \cdot \varphi_A(f_S) = \sum_{S \in \mathcal{P}_n([m])} \det(B_{S,[n]}) \det(A_{[n],S}) \cdot e_1 \wedge \dots \wedge e_n,$$

as desired. ■

Exercise 6.50. Let $A \in M_{m \times n}(\mathbb{R})$. Show that $\det(A^\top A) \geq 0$.

Exercise 6.51 (Cramer's rule). Let F be a field and let $A \in M_{n \times n}(F)$ be a matrix with $\det(A) \neq 0$. Fix some $y \in M_{n \times 1}(F)$. Let the columns of A be a_1, \dots, a_n . For each $1 \leq i \leq n$, let A_i be the matrix obtained from A by replacing its i -th column by y ; i.e.,

$$A_1 := [y \ a_2 \ \cdots \ a_n], \quad A_2 := [a_1 \ y \ \cdots \ a_n], \quad \dots, \quad A_n := [a_1 \ a_2 \ \cdots \ y].$$

Show that the unique vector $x \in M_{n \times 1}(F)$ such that $Ax = y$ is given by

$$x = \frac{1}{\det(A)} \begin{bmatrix} \det(A_1) \\ \det(A_2) \\ \vdots \\ \det(A_n) \end{bmatrix}.$$

6.F. Proof of Theorem 6.15

The method of “postulating” what we want has many advantages; they are the same as the advantages of theft over honest toil.

Bertrand Russell

Now it's finally time to construct exterior products. Let V be an n -dimensional vector space over a field F and let $\{e_1, \dots, e_n\}$ be a basis for V . We already *know* (from Corollary 6.33) that the exterior power $\bigwedge^k V$ must have a basis of the form $\{e_S : S \in \mathcal{P}_k([n])\}$. Here's the idea: We will let V_k be *some* (any!) vector space of the right dimension, pick an arbitrary basis in V_k , and simply label the elements of this basis by the sets $S \in \mathcal{P}_k([n])$. Then we will *define* the exterior product so that it has all the required properties.

So, let $V_1 := V$ and for each $k \geq 2$, let V_k be an arbitrary F -vector space of dimension $\binom{n}{k}$. Pick a basis B_k for V_k . Since $\dim V_k = \binom{n}{k}$, B_k has $\binom{n}{k}$ elements, so we can index them by the k -element subsets of $[n]$ (of which there are exactly $\binom{n}{k}$) and let

$$B_k = \{e_S : S \in \mathcal{P}_k([n])\}.$$

Again, “ e_S ” here is just an arbitrary *name* for a vector in V_k . But these names inform our construction, as our plan is to define an operation \wedge that makes equation (6.46) true. For convenience, we also let $e_{\{i\}} := e_i$ to make the set $\{e_S : S \in \mathcal{P}_1([n])\} = \{e_{\{i\}} : i \in [n]\}$ a basis for $V_1 = V$.

Now we need to define \wedge . First, we define it on the basis vectors: For all $\emptyset \neq S, T \subseteq [n]$, set

$$e_S \wedge e_T := \begin{cases} 0 \in V_{k+\ell} & \text{if } S \cap T \neq \emptyset; \\ \text{sign}(S, T)e_{S \cup T} & \text{if } S \cap T = \emptyset, \end{cases} \quad (6.52)$$

where $\text{sign}(S, T)$ is either 1 or -1 and is determined as follows: Let $S = \{s_1, \dots, s_k\}$ and $T = \{t_1, \dots, t_\ell\}$ with $s_1 < \dots < s_k$ and $t_1 < \dots < t_\ell$. Then $\text{sign}(S, T)$ is (-1) to the power of a number of transpositions needed to put the sequence $(s_1, \dots, s_k, t_1, \dots, t_\ell)$ in increasing order.

Remark 6.53. The value $\text{sign}(S, T)$ is well-defined, as it only depends on the *parity* of the number of transpositions that put the sequence $(s_1, \dots, s_k, t_1, \dots, t_\ell)$ in increasing order. (See Exercise 6.26 and Definition 6.28.)

Remark 6.54. Formula (6.52) is the *only* way to define \wedge that is consistent with (6.46) (think why!). For example, if $S = \{1, 3\}$ and $T = \{2, 3\}$, then we *should* end up having $e_S \wedge e_T = e_1 \wedge e_3 \wedge e_2 \wedge e_3 = 0$, as reflected in (6.52). Similarly, if $S = \{1, 3\}$ and $T = \{2, 4\}$, then $\text{sign}(S, T) = -1$ (because the sequence $(1, 3, 2, 4)$ can be put in increasing order by a single transposition), and we *should* get $e_S \wedge e_T = e_1 \wedge e_3 \wedge e_2 \wedge e_4 = -e_1 \wedge e_2 \wedge e_3 \wedge e_4$, which is again what (6.52) ensures.

Notice that if $|S| = k$ and $|T| = \ell$, then $e_S \wedge e_T \in V_{k+\ell}$ (as desired). Since \wedge should be bilinear, there is a unique way to extend (6.52) to arbitrary vectors. Namely, for all $x \in V_k$ and $y \in V_\ell$, we define $x \wedge y \in V_{k+\ell}$ as follows: Write x and y in terms of the corresponding bases:

$$x = \sum_{S \in \mathcal{P}_k([n])} a_S e_S \quad \text{and} \quad y = \sum_{T \in \mathcal{P}_\ell([n])} b_T e_T.$$

Then

$$x \wedge y := \sum_{S \in \mathcal{P}_k([n])} \sum_{T \in \mathcal{P}_\ell([n])} a_S b_T \cdot (e_S \wedge e_T).$$

That’s it! Notice that we had essentially no “freedom” in this construction: the definitions were “forced” on us by the requirements of Theorem 6.15. What’s left to do now is to verify that the structure obtained in this manner actually satisfies all the claims made in Theorem 6.15.

6.F.1. The operation \wedge is associative.—This is the most subtle part of the argument. Since \wedge is defined to be bilinear, it is enough to check associativity on the basis vectors (**exercise!**). So, take $S = \{s_1, \dots, s_k\}$, $T = \{t_1, \dots, t_\ell\}$, and $R = \{r_1, \dots, r_m\}$ with $s_1 < \dots < s_k$, $t_1 < \dots < t_\ell$, and $r_1 < \dots < r_m$, and consider the expressions

$$(e_S \wedge e_T) \wedge e_R \quad \text{and} \quad e_S \wedge (e_T \wedge e_R).$$

If the sets S, T, R are not pairwise disjoint, then both of these expressions are equal to 0. Otherwise, both of them are equal to plus or minus $e_{S \cup T \cup R}$. Furthermore, the coefficient in front of $e_{S \cup T \cup R}$ is equal to (-1) to the power of a number of transpositions that put the sequence

$$(s_1, \dots, s_k, t_1, \dots, t_\ell, r_1, \dots, r_m)$$

in increasing order (**exercise!**). Therefore, $(e_S \wedge e_T) \wedge e_R = e_S \wedge (e_T \wedge e_R)$, as desired.

6.F.2. Properties (EP1), (EP2), (EP3), and (EP5).—These properties are satisfied by construction.

6.F.3. *Property (EP4): We have $V_k = \text{Span}\{x_1 \wedge \dots \wedge x_k : x_1, \dots, x_k \in V\}$.—This holds since the space V_k is spanned by the elements e_S , $S \in \mathcal{P}_k([n])$, and \wedge is defined so that we have*

$$e_S = e_{s_1} \wedge \dots \wedge e_{s_k},$$

where $S = \{s_1, \dots, s_k\}$ and $s_1 < \dots < s_k$.

6.F.4. *Property (EP6): The map $V \times V \rightarrow V_2: (x, y) \mapsto x \wedge y$ is alternating.—Take any $x \in V$ and write it as $x = \sum_{i=1}^n a_i e_i$ with $a_1, \dots, a_n \in F$. Then, by definition,*

$$x \wedge x = \sum_{i=1}^n \sum_{j=1}^n a_i a_j (e_i \wedge e_j).$$

Note that

$$e_i \wedge e_j = \begin{cases} 0 & \text{if } i = j; \\ e_{\{i,j\}} & \text{if } i < j; \\ -e_{\{j,i\}} & \text{if } i > j. \end{cases}$$

Hence,

$$x \wedge x = \sum_{1 \leq i < j \leq n} (a_i a_j - a_j a_i) e_{\{i,j\}} = 0,$$

as claimed.

6.F.5. *Property (EP7): A tuple $(x_1, \dots, x_k) \in V^k$ is independent if and only if $x_1 \wedge \dots \wedge x_k \neq 0$.—We already know from Lemma 6.14 that if (x_1, \dots, x_k) is not independent, then $x_1 \wedge \dots \wedge x_k = 0$. Now suppose that the tuple (x_1, \dots, x_k) is independent. Then we can extend it to an ordered basis (x_1, \dots, x_n) , and it suffices to show that*

$$x_1 \wedge \dots \wedge x_n = (x_1 \wedge \dots \wedge x_k) \wedge (x_{k+1} \wedge \dots \wedge x_n) \neq 0.$$

Since (x_1, \dots, x_n) is a basis, we can express each of e_1, \dots, e_n as a linear combination of x_1, \dots, x_n . Plugging these linear combinations into $e_1 \wedge \dots \wedge e_n$ and expanding, we obtain that

$$e_1 \wedge \dots \wedge e_n \in \text{Span}(\{x_{i_1} \wedge \dots \wedge x_{i_n} : 1 \leq i_1, \dots, i_n \leq n\}).$$

An expression of the form $x_{i_1} \wedge \dots \wedge x_{i_n}$ can only be nonzero if the indices i_1, \dots, i_n are pairwise distinct, in which case $x_{i_1} \wedge \dots \wedge x_{i_n}$ is equal to either plus or minus $x_1 \wedge \dots \wedge x_n$. Therefore, $e_1 \wedge \dots \wedge e_n$ is a scalar multiple of $x_1 \wedge \dots \wedge x_n$. Since $e_1 \wedge \dots \wedge e_n = e_{[n]} \neq 0$ by definition, we conclude that $x_1 \wedge \dots \wedge x_n \neq 0$ as well, as claimed.

6.F.6. *Property (EP8): For every alternating k -linear map $f: V^k \rightarrow W$, there is a unique linear function $\varphi: V_k \rightarrow W$ such that we have $f(x_1, \dots, x_k) = \varphi(x_1 \wedge \dots \wedge x_k)$.—Since $\{e_S : S \in \mathcal{P}_k([n])\}$ is a basis for V_k , there is a unique linear function $\varphi: V_k \rightarrow W$ such that $\varphi(e_S) = f(e_{s_1}, \dots, e_{s_k})$ for all $S = \{s_1, \dots, s_k\}$ with $s_1 < \dots < s_k$. We claim that this function also satisfies $\varphi(x_1 \wedge \dots \wedge x_k) = f(x_1, \dots, x_k)$ for all $x_1, \dots, x_k \in V$. Indeed, consider the map*

$$g: V^k \rightarrow W: (x_1, \dots, x_k) \mapsto \varphi(x_1 \wedge \dots \wedge x_k).$$

It is clear that g is k -linear and alternating, and, by definition,

$$g(e_{s_1}, \dots, e_{s_k}) = f(e_{s_1}, \dots, e_{s_k}) \quad \text{whenever } 1 \leq s_1 < \dots < s_k \leq n. \quad (6.55)$$

Since both f and g are alternating, (6.55) implies that f and g must agree on arbitrary sequences of basis vectors. Hence, $f = g$ by Lemma 6.6.

The proof of Theorem 6.15 is complete.

6.G. Quantifier elimination

Suppose that $A \in M_{m \times n}(F)$. By definition, the columns of A are independent if and only if

for all $x \in M_{n \times 1}(F)$, if $x \neq 0$, then $Ax \neq 0$.

Notice that this definition starts with a *universal quantifier*: Some property must hold for every single vector in the (possibly infinite) space $M_{n \times 1}(F)$. We also know (**exercise!**) that this property can be stated equivalently as follows:

there exists a matrix $B \in M_{n \times m}(F)$ such that $BA = I_n(F)$.

This reformulation starts with an *existential quantifier* that is asking us to find a matrix with certain properties in the (possibly infinite) space $M_{n \times m}(F)$. When $n = m$ (i.e., the matrix A is square), there is a third formulation:

$$\det(A) \neq 0,$$

which is *quantifier-free*: To verify it, we simply need to perform a direct computation that only involves the entries of A . Exterior products provide a similar quantifier free characterization of linear independence in the general case (i.e., when n and m may differ):

Theorem 6.56. *Let F be a field and let $A \in M_{m \times n}(F)$. Then the columns of A are independent if and only if there is a set $S \in \mathcal{P}_n([m])$ such that $\det(A_{S,[n]}) \neq 0$.*

Example 6.57. We've already encountered a special case of this with $n = 2$, $m = 3$ in Exercise 6.22.

PROOF. Set $V := M_{n \times 1}(F)$ and $W := M_{m \times 1}(F)$. Let $\{e_1, \dots, e_n\}$ and $\{f_1, \dots, f_m\}$ be the standard bases for V and W , respectively. The columns of A are independent if and only if their exterior product is nonzero, i.e., if $Ae_1 \wedge \dots \wedge Ae_n \neq 0$. By Lemma 6.48,

$$Ae_1 \wedge \dots \wedge Ae_n = \sum_{S \in \mathcal{P}_n([m])} \det(A_{S,[n]}) \cdot f_S.$$

Since $\{f_S : S \in \mathcal{P}_n([m])\}$ is a basis for $\bigwedge^n W$, the last expression is nonzero if and only if at least one of the coefficients is nonzero, i.e., when $\det(A_{S,[n]}) \neq 0$ for some $S \in \mathcal{P}_n([m])$, as desired. ■

Corollary 6.58 (Rank in terms of determinants). *Let F be a field and let $A \in M_{m \times n}(F)$. Then $\text{rank}(A)$ is equal to the largest integer k such that there exist k -element subsets $S \subseteq [m]$ and $T \subseteq [n]$ with $\det(A_{S,T}) \neq 0$.*

Exercise 6.59. Prove Corollary 6.58.

Remark 6.60. Notice that Corollary 6.58 provides another proof that $\text{rank}(A) = \text{rank}(A^\top)$.

As an application of Theorem 6.56, we shall establish a connection between linear independence over \mathbb{Q} and over \mathbb{F}_p for a prime p . As a motivating example, consider the three vectors

$$x = \begin{bmatrix} 0 \\ 4 \\ 2 \end{bmatrix}, \quad y = \begin{bmatrix} 2 \\ 0 \\ 1 \end{bmatrix}, \quad z = \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix} \in M_{3 \times 1}(\mathbb{Z}).$$

The triple (x, y, z) is independent over \mathbb{Q} . However, since the entries of x , y , and z are integers, we can reduce them modulo a prime p and inquire whether the resulting vectors in $M_{3 \times 1}(\mathbb{F}_p)$ are also independent over \mathbb{F}_p . When $p = 2$, we have $x = 0 \pmod{2}$, and thus the triple (x, y, z) loses its independence in \mathbb{F}_2 . Similarly,

$$x + y + z = \begin{bmatrix} 3 \\ 6 \\ 3 \end{bmatrix} = 0 \pmod{3},$$

which means that the triple (x, y, z) is *not* independent over \mathbb{F}_3 . Yet, it *is* independent in \mathbb{F}_p for all primes $p \geq 5$:

Exercise 6.61. Show that the triple (x, y, z) is independent modulo p for all primes $p \geq 5$.

Our question is, what happens in general? In other words, given vectors $x_1, \dots, x_k \in M_{n \times 1}(\mathbb{Z})$ such that the tuple (x_1, \dots, x_k) is independent over \mathbb{Q} , what can be said about the independence of (x_1, \dots, x_k) modulo a prime p ? From the above example, we know that there can be finitely many primes over which the tuple (x_1, \dots, x_k) loses its independence. It turns out that there can be *only* finitely many such “bad” primes:

Theorem 6.62. Let $x_1, \dots, x_k \in M_{n \times 1}(\mathbb{Z})$. The following statements are equivalent:

- (1) the tuple (x_1, \dots, x_k) is independent over \mathbb{Q} ;
- (2) the tuple (x_1, \dots, x_k) is independent over \mathbb{F}_p for some prime p ;
- (3) the tuple (x_1, \dots, x_k) is independent over \mathbb{F}_p for all but finitely many primes p .

PROOF. (2) \implies (1). We will prove the contrapositive of this implication. Suppose that the tuple (x_1, \dots, x_k) is not independent over \mathbb{Q} . This means that there exist rational numbers $a_1, \dots, a_k \in \mathbb{Q}$, not all of which are zero, such that $a_1x_1 + \dots + a_kx_k = 0$. After clearing the denominators, we may assume that a_1, \dots, a_k are actually integers. Furthermore, we may assume that $\gcd(a_1, \dots, a_k) = 1$, since otherwise we can simply replace each a_i with $a_i/\gcd(a_1, \dots, a_k)$. With these assumptions, we see that for all primes p , at least one of a_1, \dots, a_k is nonzero modulo p , while $a_1x_1 + \dots + a_kx_k = 0 \pmod{p}$, showing that (x_1, \dots, x_k) is not independent over \mathbb{F}_p .

(3) \implies (2). This implication is trivial.

(1) \implies (3). This is the interesting (and somewhat surprising) part of the theorem, and this is where Theorem 6.56 comes in handy. Let A be the n -by- k matrix whose columns are x_1, \dots, x_k and suppose that the tuple (x_1, \dots, x_k) is independent over \mathbb{Q} . By Theorem 6.56, this means that there is a set $S \in \mathcal{P}_k([n])$ with $\det(A_{S,[k]}) \neq 0$. Since $\det(A_{S,[k]})$ is a nonzero integer, there are only finitely many primes p that divide it, and for all *other* p , $\det(A_{S,[k]}) \neq 0 \pmod{p}$, implying, by Theorem 6.56 again, that the columns of A are independent in \mathbb{F}_p , as desired. ■

In the above proof of (1) \implies (3) we tacitly relied on the result of the following exercise:

Exercise 6.63. Let $A \in M_{n \times n}(\mathbb{Z})$ and let p be a prime number. Let $A' \in M_{n \times n}(\mathbb{F}_p)$ be the matrix obtained by reducing each entry of A modulo p . Show that $\det(A) = \det(A') \pmod{p}$.

6.H. A combinatorial application: the skew set pairs inequality

In this subsection, we shall establish the following combinatorial fact, which, at first glance, has little to do with linear algebra (althou it is somewhat reminiscent of Theorem 1.39):

Theorem 6.64 (Skew set pairs inequality; Frankl–Lovász). *Suppose that A_1, \dots, A_n are k -element sets and B_1, \dots, B_n are ℓ -element sets such that:*

- $A_i \cap B_i = \emptyset$ for all $1 \leq i \leq n$; and
- $A_i \cap B_j \neq \emptyset$ for all $1 \leq i < j \leq n$.

Then $n \leq \binom{k+\ell}{k}$.

A few remarks are in order. First, note that in the statement of Theorem 6.64, there are no assumption on the size of the ground set $\bigcup_{i=1}^n (A_i \cup B_i)$; in other words, the same upper bound on n is valid regardless of the total number of elements that the sets in the theorem are allowed to contain. Second, the upper bound $n \leq \binom{k+\ell}{k}$ is best possible. To see this, let

$$A_1, \quad \dots, \quad A_{\binom{k+\ell}{k}}$$

be all the k -element subsets of $[k + \ell]$ and set $B_i := [k + \ell] \setminus A_i$. Third, notice that there is some asymmetry in how the A_i s and the B_i s are treated. Specifically, we are only requiring the intersection $A_i \cap B_j$ to be nonempty when $i < j$, while nothing is said about the case $j < i$. This is why

Theorem 6.64 is referred to as the *skew* set pairs inequality. The non-skew version, in which $A_i \cap B_j \neq \emptyset$ whenever $i \neq j$, was proved earlier by Bollobás with a clever combinatorial argument. However, for the skew version, only algebraic proofs are known!

PROOF. Without loss of generality, we may assume that all the sets $A_1, \dots, A_n, B_1, \dots, B_n$ are subsets of $[N]$ for some $N \in \mathbb{N}$. Using Exercise 5.51, we get a sequence of vectors $x_1, \dots, x_N \in \mathbb{R}^{k+\ell}$ such that every $(k + \ell)$ of them are independent. To each A_i we associate an element $x_{A_i} \in \bigwedge^k \mathbb{R}^{k+\ell}$ given by the formula

$$x_{A_i} := x_{a_1} \wedge \dots \wedge x_{a_k}, \quad \text{where } A_i = \{a_1, \dots, a_k\} \text{ with } a_1 < \dots < a_k.$$

Similarly, let $x_{B_i} \in \bigwedge^\ell \mathbb{R}^{k+\ell}$ be given by

$$x_{B_i} := x_{b_1} \wedge \dots \wedge x_{b_\ell}, \quad \text{where } B_i = \{b_1, \dots, b_\ell\} \text{ with } b_1 < \dots < b_\ell.$$

If $1 \leq i < j \leq n$, then $A_i \cap B_j \neq \emptyset$, and so $x_{A_i} \wedge x_{B_j} = 0$. On the other hand, $x_{A_i} \wedge x_{B_i}$ is the wedge product of $(k + \ell)$ distinct elements of the sequence x_1, \dots, x_N , and therefore $x_{A_i} \wedge x_{B_i} \neq 0$. The crux of the argument is in the following observation:

Claim. *The tuple $(x_{A_1}, \dots, x_{A_n})$ is independent.*

Proof. Suppose not. Then we have $\sum_{i=1}^n c_i x_{A_i} = 0$ for some coefficients $c_1, \dots, c_n \in \mathbb{R}$, not all of which are zero. Let j be the largest index such that $c_j \neq 0$. Then

$$0 = \left(\sum_{i=1}^j c_i x_{A_i} \right) \wedge x_{B_j} = \sum_{i=1}^j c_i (x_{A_i} \wedge x_{B_j}) = c_j (x_{A_j} \wedge x_{B_j}) \neq 0.$$

This contradiction completes the proof of the claim. −

From the above claim, we conclude that $n \leq \dim \bigwedge^k \mathbb{R}^{k+\ell} = \binom{k+\ell}{k}$, as desired. ■

Extra exercises for Section 6

Exercise 6.65. Let V be a finite-dimensional vector space over a field F and let $\text{Alt}_k(V)$ be the set of all alternating k -linear maps $f: V^k \rightarrow F$. Show that $\text{Alt}_k(V)$, viewed as an F -vector space, is isomorphic to $\bigwedge^k V$.

Exercise 6.66 (important). Let V be a finite-dimensional vector space over a field F and let $v_1, \dots, v_k, w_1, \dots, w_k \in V$. Suppose that the tuples (v_1, \dots, v_k) and (w_1, \dots, w_k) are independent. Show that $\text{Span}(\{v_1, \dots, v_k\}) = \text{Span}(\{w_1, \dots, w_k\})$ if and only if $v_1 \wedge \dots \wedge v_k$ is a nonzero scalar multiple of $w_1 \wedge \dots \wedge w_k$.

7. POLYNOMIALS

7.A. Basic properties of polynomials

Polynomials play an extremely important role in algebra in general and in linear algebra in particular.

Example 7.1. If we consider the entries of an n -by- n matrix A as variables, the determinant $\det(A)$ becomes a (multivariate) polynomial. For instance,

$$\det \begin{bmatrix} x_1 & x_3 \\ x_2 & x_4 \end{bmatrix} = x_1 x_4 - x_2 x_3.$$

Even though we have already encountered polynomials on a few occasions in these notes, it will do us good to briefly review the definition of a polynomial. Let R be a commutative ring. A **polynomial** over R in a single variable x is an expression of the form

$$a_0 + a_1 x + a_2 x^2 + \dots + a_n x^n, \tag{7.2}$$

where $n \in \mathbb{N}$ and $a_0, \dots, a_n \in R$. To be more precise, an expression such as (7.2) is a shortcut that stands for

$$a_0 + a_1x + a_2x^2 + \cdots + a_nx^n + 0x^{n+1} + 0x^{n+2} + \cdots \quad (\text{the sum is infinite}).$$

Thus, for example, $1 + x$ and $1 + x + 0x^2$ are two expressions for the same polynomial. If we wanted to be completely formal, we could say that a one-variable polynomial over R is simply an element of the set $[\mathbb{N} \rightarrow R]^{<\infty}$, i.e., a sequence $(a_0, a_1, a_2, \dots) \in R^{\mathbb{N}}$ with only finitely many nonzero entries.

The set of all polynomials over R in a variable x is denoted by $R[x]$. Each polynomial $p \in R[x]$ gives rise to a function $R \rightarrow R$ obtained by evaluating p in the usual way; that is, if

$$p = a_0 + a_1x + a_2x^2 + \cdots + a_nx^n,$$

then for each $c \in R$, we set

$$p(c) := a_0 + a_1c + a_2c^2 + \cdots + a_nc^n,$$

where addition and multiplication are interpreted as the corresponding operations in R , and we use the standard abbreviation $c^k := c \cdot c \cdots c$ (k factors). Crucially, two polynomials are considered equal when they have the same coefficients, and so two *distinct* polynomials may give rise to *the same* function. For example, consider the following two polynomials in $\mathbb{F}_3[x]$:

$$p := x^3 + x + 1 \quad \text{and} \quad q := 2x + 1.$$

We then have

$$0^3 + 0 + 1 = 2 \cdot 0 + 1 = 1 \pmod{3},$$

$$1^3 + 1 + 1 = 2 \cdot 1 + 1 = 0 \pmod{3},$$

$$2^3 + 2 + 1 = 2 \cdot 2 + 1 = 2 \pmod{3},$$

so $p(c) = q(c)$ for all $c \in \mathbb{F}_3$, and yet $p \neq q$ as *polynomials*. Nevertheless, we will soon see that if the ring R is *infinite*, then any two polynomials that give rise to the same function must, in fact, coincide as polynomials (see Exercises 7.11 and 7.12).

If p is a polynomial in a variable x , then we write $[x^k]p$ to indicate the coefficient of p corresponding to the monomial x^k ; in other words,

$$[x^k](a_0 + a_1x + a_2x^2 + \cdots + a_nx^n) := a_k.$$

Polynomials can be added and multiplied together, which makes $R[x]$ into a commutative ring in its own right. Some care has to be taken when defining addition and multiplication of polynomials, since, as explained above, polynomials cannot, in general, be identified with their corresponding functions. This means that, given a pair of polynomials $p, q \in R[x]$, we have to be able to describe the *coefficients* of $p + q$ and pq in terms of the coefficients of p and q . This can be done as follows:

$$[x^k](p + q) := [x^k]p + [x^k]q \quad \text{and} \quad [x^k](pq) := \sum_{i=0}^k ([x^i]p)([x^{k-i}]q).$$

Exercise 7.3. Let R be a commutative ring and let $p, q \in R[x]$ be two polynomials. Show that for all $c \in R$, $(p + q)(c) = p(c) + q(c)$ and $(pq)(c) = p(c) \cdot q(c)$. Explain why the latter equality might fail if the ring R is not assumed to be commutative.

Exercise 7.4. Let R be a commutative ring. Show that $R[x]$ is also a commutative ring.

Exercise 7.5. Let F be a field. Show that $F[x]$ is a vector space over F .

The **degree** of a nonzero polynomial $p \in R[x]$, denoted by $\deg p$, is the largest $k \in \mathbb{N}$ such that $[x^k]p \neq 0$. By definition, $\deg 0 := -\infty$. The set of all polynomials over R in a variable x of degree at most n is denoted by $P_n(R)$. If F is a field, then $P_n(F)$ is a subspace of $F[x]$ of dimension $n + 1$.

Exercise 7.6. Let R be a commutative ring and let $p, q \in R[x]$. Show that $\deg(pq) = \deg p + \deg q$.

Exercise 7.7 (important). Give a definition of the ring of **multivariate polynomials** $R[x_1, \dots, x_k]$.

7.B. Polynomial division

Let F be a field. We say that a polynomial $q \in F[x]$ **divides** another polynomial $p \in F[x]$ if there is a polynomial $f \in F[x]$ such that $p = fq$. If q divides p , then we call q a **divisor** of p . Notice that if $p \neq 0$, then the degree of every divisor of p is at most $\deg p$. A polynomial p is **irreducible** if all its divisors q satisfy $\deg q = 0$ or $\deg q = \deg p$.

Lemma 7.8 (Polynomial division with remainder). *Let F be a field and let $p, q \in F[x]$. If $q \neq 0$, then there exist unique polynomials f (the **quotient**) and r (the **remainder**) such that*

$$\deg r < \deg q \quad \text{and} \quad p = fq + r.$$

PROOF. Call a polynomial $h \in F[x]$ a **potential remainder** if there is $f \in F[x]$ with $p = fq + h$. Note that p itself is a potential remainder since $p = 0 \cdot q + p$. Let r be a potential remainder of the smallest degree. We claim that then $\deg r < \deg q$, proving the existence part of the lemma. Indeed, since r is a potential remainder, there is $f \in F[x]$ with $p = fq + r$, so we only need to show that $\deg r < \deg q$. Suppose, towards a contradiction, that $\deg r \geq \deg q$. Let

$$a := [x^{\deg q}]q \quad \text{and} \quad b := [x^{\deg r}]r,$$

and consider the polynomial $r' := r - (b/a)x^{\deg r - \deg q}q$. By construction, $\deg r' < \deg r$; but

$$p = fq + r = (f + (b/a)x^{\deg r - \deg q})q + r',$$

so r' is a potential remainder, which contradicts our choice of r . For the uniqueness part, suppose that $p = f_1q + r_1 = f_2q + r_2$, where $\deg r_1, \deg r_2 < \deg q$. Then $(f_1 - f_2)q = r_1 - r_2$, and hence q is a divisor of $r_1 - r_2$. Since $\deg q > \deg(r_1 - r_2)$, this is only possible if $r_1 - r_2 = 0$, as desired. ■

A **root** of a polynomial $p \in F[x]$ is any element $c \in F$ such that $p(c) = 0$. The following properties of roots follow easily from Lemma 7.8:

Exercise 7.9. Let F be a field and let $p \in F[x]$. Show that if $c \in F$ is a root of p , then the polynomial $x - c$ is a divisor of p .

Exercise 7.10 (Baby Bézout's theorem). Let F be a field and let $p \in F[x]$ be a nonzero polynomial. Show that p can have at most $\deg p$ distinct roots.

Exercise 7.11. Suppose that F is an infinite field and $p, q \in F[x]$ are polynomials such that $p(c) = q(c)$ for all $c \in F$. Show that $p = q$ as polynomials.

Exercise 7.12. Suppose that R is an infinite commutative ring (but not necessarily a field) and let $p, q \in R[x]$ be polynomials such that $p(c) = q(c)$ for all $c \in R$. Show that $p = q$ as polynomials. *Hint:* Make sure that the conclusion of Lemma 7.8 holds over R when $q = x - c$ for some $c \in R$.

Let $p, q \in F[x]$. A polynomial $s \in F[x]$ is a **greatest common divisor** (or a **gcd**) of p and q if s divides both p and q , and whenever $t \in F[x]$ divides both p and q , t also divides s .

Exercise 7.13. Show that if s_1 and s_2 are gcds of $p, q \in F[x]$, then $s_1 = as_2$ for some $a \in F \setminus \{0\}$.

The following result is one of the most fundamental properties of polynomials over a field:

Theorem 7.14 (Euclidean algorithm for polynomials). *Let F be a field and let $p, q \in F[x]$. If at least one of p and q is nonzero, then p and q have a gcd $s \in F[x]$; furthermore, there exist polynomials $u, v \in F[x]$ such that $up + vq = s$.*

PROOF. Without loss of generality, assume that $p \neq 0$ and $\deg p \geq \deg q$. We argue by induction on $\deg q$. If $\deg q = -\infty$, i.e., if $q = 0$, then we can take $s = p$, $u = 1$, and $v = 0$. Now suppose that $q \neq 0$. By Lemma 7.8, we can write $p = fq + r$ with $f, r \in F[x]$ and $\deg r < \deg q$. By the inductive assumption, q and r have a gcd s , and there exist polynomials g and h such that $gq + hr = s$.

Exercise 7.15. Show that s is also a gcd of p and q .

Hence, p and q have a gcd. Furthermore,

$$s = gq + hr = gq + h(p - fq) = hp + (g - hf)q,$$

so we can take $u = h$ and $v = g - hf$. ■

A polynomial $s \in F[x]$ is **monic** if $s \neq 0$ and $[x^{\deg s}]s = 1$. Given two polynomials $p, q \in F[x]$, at least one of which is nonzero, we write $\gcd(p, q)$ for the unique monic gcd of p and q .

Example 7.16. Let $p := x^3 + x + 1$, $q := 2x + 1$. Viewing p and q as polynomials over \mathbb{R} , we have

$$\gcd(p, q) = 1, \quad \text{and} \quad \frac{8}{3} \cdot p + \left(-\frac{4}{3}x^2 + \frac{2}{3}x - \frac{5}{3}\right) \cdot q = 1.$$

On the other hand, if we think of p and q as polynomials over \mathbb{F}_3 , then $p = (2x^2 + 2x + 1)q$, so

$$\gcd(p, q) = x + 2, \quad \text{and} \quad 0 \cdot p + 2 \cdot q = x + 2.$$

Exercise 7.17. Working over \mathbb{R} , compute $\gcd(x^2 + 1, x^3 + 1)$ and find polynomials u and v such that $u \cdot (x^2 + 1) + v \cdot (x^3 + 1) = \gcd(x^2 + 1, x^3 + 1)$.

Exercise 7.18. Same as Exercise 7.17, but working over \mathbb{F}_2 instead.

The notion of a gcd extends naturally to more than two polynomials. Let $p_1, \dots, p_k \in F[x]$. A polynomial $s \in F[x]$ is a **greatest common divisor** (or a **gcd**) of p_1, \dots, p_k if s divides all of p_1, \dots, p_k , and whenever $t \in F[x]$ divides all of p_1, \dots, p_k , t also divides s .

Exercise 7.19 (Euclidean algorithm for several polynomials). Prove the following extension of Theorem 7.14 to several polynomials: Let F be a field and let $p_1, \dots, p_k \in F[x]$. If at least one of p_1, \dots, p_k is nonzero, then p_1, \dots, p_k have a gcd $s \in F[x]$; furthermore, there exist polynomials $u_1, \dots, u_k \in F[x]$ such that $u_1p_1 + \dots + u_kp_k = s$.

7.C. The resultant

Let F be a field and let $p, q \in F[x]$ be two nonzero polynomials. A **least common multiple** (or an **lcm**) of p and q is a nonzero polynomial $s \in F[x]$ such that p and q both divide s , and whenever p and q both divide some $t \in F[x]$, s also divides t .

Lemma 7.20 (Least common multiples). *Let F be a field and let $p, q \in F[x]$ be two nonzero polynomials. Then p and q have an lcm, namely the polynomial $pq/\gcd(p, q)$.*

PROOF. Clearly,

$$pq/\gcd(p, q) = p \cdot (q/\gcd(p, q)) = q \cdot (p/\gcd(p, q))$$

is divisible by p and q . What remains to show is that if $f \in F[x]$ is a polynomial divisible by p and q , then f is also divisible by $pq/\gcd(p, q)$, or, equivalently, $f \cdot \gcd(p, q)$ is divisible by pq . To that end, write $\gcd(p, q) = up + vq$ for some $u, v \in F[x]$. Then

$$f \cdot \gcd(p, q) = f \cdot (up + vq) = ufp + vfq.$$

Since f is divisible by q , fp is divisible by pq . Similarly, since f is divisible by p , fq is also divisible by pq . Thus, $f \cdot \gcd(p, q)$ is divisible by pq , and we are done. ■

Suppose we are given a pair of nonzero polynomials $p, q \in F[x]$. We shall use linear-algebraic tools to tackle the following general problem:

How can we decide whether $\gcd(p, q) \neq 1$, that is, whether p and q have a nontrivial common divisor?

For concreteness, let $n := \deg p$ and $m := \deg q$ and write

$$p = a_0 + a_1x + \dots + a_nx^n \quad \text{and} \quad q = b_0 + b_1x + \dots + b_mx^m.$$

Exercise 7.21. Deduce from Lemma 7.20 that $\gcd(p, q) \neq 1$ if and only if there exist nonzero polynomials $u, v \in F[x]$ such that

$$\deg u \leq m - 1, \quad \deg v \leq n - 1, \quad \text{and} \quad up + vq = 0.$$

In view of Exercise 7.21, it makes sense to consider the function

$$\varphi_{p,q}: P_{m-1}(F) \oplus P_{n-1}(F) \rightarrow P_{n+m-1}(F): (u, v) \mapsto up + vq.$$

This function is linear, and, by Exercise 7.21, we have $\gcd(p, q) \neq 1$ if and only if $\ker(\varphi_{p,q}) \neq \{0\}$. Notice that since

$$\dim(P_{m-1}(F) \oplus P_{n-1}(F)) = \dim P_{n+m-1}(F) = n + m,$$

we have $\ker(\varphi_{p,q}) \neq \{0\}$ if and only if $\varphi_{p,q}$ is not a bijection.

The next step is to compute a matrix corresponding to $\varphi_{p,q}$. Specifically, let

$$X := ((1, 0), (x, 0), (x^2, 0), \dots, (x^{m-1}, 0), (0, 1), (0, x), (0, x^2), \dots, (0, x^{n-1}))$$

be the “obvious” ordered basis for $P_{m-1}(F) \oplus P_{n-1}(F)$ and let

$$Y := (1, x, x^2, \dots, x^{n+m-1})$$

be the ordered basis for $P_{n+m-1}(F)$. The matrix $[\varphi_{p,q}]_{X,Y}$ that expresses $\varphi_{p,q}$ with respect to these bases is called the **Sylvester matrix**¹⁸ of p and q and is denoted by $\text{Syl}(p, q)$. Note that $\text{Syl}(p, q)$ is an $(n + m)$ -by- $(n + m)$ matrix. It is not hard to describe the entries of the Sylvester matrix explicitly. Indeed, by definition, the first column of $\text{Syl}(p, q)$ is

$$[\varphi_{p,q}(1, 0)]_Y = [1 \cdot p + 0 \cdot q]_Y = [p]_Y = [a_0 \ a_1 \ a_2 \ \cdots \ a_n \ 0 \ \cdots \ 0]^\top,$$

ending with $m - 1$ zeros. Similarly, the second column of $\text{Syl}(p, q)$ is

$$[\varphi_{p,q}(x, 0)]_Y = [x \cdot p + 0 \cdot q]_Y = [xp]_Y = [0 \ a_0 \ a_1 \ \cdots \ a_n \ 0 \ \cdots \ 0]^\top,$$

now ending with only $m - 2$ zeros; and so on. The first m columns of $\text{Syl}(p, q)$ are

$$\begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ \vdots \\ a_n \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \begin{bmatrix} 0 \\ a_0 \\ a_1 \\ \vdots \\ a_{n-1} \\ a_n \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \begin{bmatrix} 0 \\ 0 \\ a_0 \\ \vdots \\ a_{n-2} \\ a_{n-1} \\ a_n \\ \vdots \\ 0 \end{bmatrix}, \quad \dots, \quad \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ a_0 \\ a_1 \\ a_2 \\ \vdots \\ a_n \end{bmatrix}.$$

¹⁸Named after James Joseph Sylvester, a 19th century mathematician who, among other things, laid the foundations of modern linear algebra. In particular, he introduced the term “matrix.”

Analogously, the remaining n columns are

$$\begin{bmatrix} b_0 \\ b_1 \\ b_2 \\ \vdots \\ b_m \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ b_0 \\ b_1 \\ \vdots \\ b_{m-1} \\ b_m \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ b_0 \\ \vdots \\ b_{m-2} \\ b_{m-1} \\ b_m \\ \vdots \\ 0 \end{bmatrix}, \dots, \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ b_0 \\ b_1 \\ b_2 \\ \vdots \\ b_m \end{bmatrix}.$$

Example 7.22. If $p = x^3 + x + 1$ and $q = 2x + 1$, then $\deg p = 3$ and $\deg q = 1$, so $\text{Syl}(p, q)$ is a 4-by-4 matrix; namely

$$\text{Syl}(p, q) = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 2 & 1 & 0 \\ 0 & 0 & 2 & 1 \\ 1 & 0 & 0 & 2 \end{bmatrix}.$$

The **resultant** of p and q is defined by the formula $\text{res}(p, q) := \det(\text{Syl}(p, q))$. It is immediate from this definition that $\varphi_{p,q}$ is bijective if and only if $\text{res}(p, q) \neq 0$. Hence, we have the following:

Theorem 7.23 (Sylvester). *Let F be a field and let $p, q \in F[x]$ be nonzero polynomials. Then we have $\gcd(p, q) \neq 1$ if and only if $\text{res}(p, q) = 0$. ■*

Example 7.24. Continuing Example 7.22, if $p = x^3 + x + 1$ and $q = 2x + 1$, then

$$\text{res}(p, q) = \det \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 2 & 1 & 0 \\ 0 & 0 & 2 & 1 \\ 1 & 0 & 0 & 2 \end{bmatrix} = 3.$$

Hence, $\gcd(p, q) = 1$ when p and q are viewed as polynomials over \mathbb{R} . On the other hand, over \mathbb{F}_3 we have $\gcd(p, q) \neq 1$, since $3 = 0 \pmod{3}$ (see Example 7.16).

Notice that $\text{res}(p, q)$ is itself a (multivariate) *polynomial* in terms of the coefficients of p and q . For instance, when $n = 2$ and $m = 1$, we have

$$\text{res}(p, q) = \det \begin{bmatrix} a_0 & b_0 & 0 \\ a_1 & b_1 & b_0 \\ a_2 & 0 & b_1 \end{bmatrix} = a_0 b_1^2 + a_2 b_0^2 - a_1 b_0 b_1,$$

which is a polynomial in a_0, a_1, a_2, b_0, b_1 . This fact will turn out to be important later on.

7.D. Multiple roots and derivatives

Let F be a field and let $p \in F[x]$ and $c \in F$. We already know (from Exercise 7.9) that c is a root of p if and only if $x - c$ divides p .

Definition 7.25. Let F be a field and let $p \in F[x]$ be a nonzero polynomial. Suppose that $c \in F$ is a root of p . The **multiplicity** of c is the largest integer k such that the polynomial $(x - c)^k$ divides p . If the multiplicity of c is at least 2, then c is called a **multiple root** of p ; otherwise (i.e., if the multiplicity of c is precisely 1), c is called a **simple root**.

You should know from calculus that there is a simple criterion for when $c \in \mathbb{R}$ is a multiple root of a polynomial $p \in \mathbb{R}[x]$: c is a multiple root of p if and only if $p(c) = p'(c) = 0$, where p' is the *derivative* of p . It turns out that this criterion, correctly interpreted, can be extended to polynomials

over an *arbitrary* field. Of course, the usual “ ε - δ ” definition of the derivative doesn’t make much sense in an arbitrary field (what would it mean in a finite field, for example?), but we can define derivatives of polynomials without invoking limits, by simply declaring the derivative of p to be what we know it “ought to be”.¹⁹

Definition 7.26. Let F be a field and let $p = \sum_k a_k x^k$ be a polynomial over F in a variable x . The **derivative** of p is the polynomial p' given by

$$p' := \sum_k (ka_k)x^{k-1},$$

where ka_k is, as usual, a shortcut for $a_k + \dots + a_k$ (k summands).

Example 7.27. Let p be a prime number. Viewing x^p as a polynomial over \mathbb{F}_p , we get

$$(x^p)' = px^{p-1} = 0 \pmod{p}.$$

Hence, x^p is a polynomial over \mathbb{F}_p of degree $p > 0$ whose derivative is zero.

Exercise 7.28. Let F be a field of characteristic zero and let $p \in F[x]$ be polynomial of positive degree. Show that $p' \neq 0$ and, in fact, $\deg p' = \deg p - 1$. (In other words, the situation described in Example 7.27 cannot occur in a field of characteristic zero.)

Exercise 7.29. Let F be a field and let $p, q \in F[x]$. Show that $(p + q)' = p' + q'$.

Lemma 7.30 (Product rule). *Let F be a field and let $p, q \in F[x]$. Then $(pq)' = p'q + pq'$.*

PROOF. It is enough to consider the case when $p = x^k$ and $q = x^\ell$ for some $k, \ell \in \mathbb{N}$ (why?). We simply compute and compare both sides:

$$\begin{aligned} (pq)' &= (x^k \cdot x^\ell)' = (x^{k+\ell})' = (k + \ell)x^{k+\ell-1}; \\ p'q + pq' &= (x^k)'x^\ell + x^k(x^\ell)' = (kx^{k-1})x^\ell + x^k(\ell x^{\ell-1}) = (k + \ell)x^{k+\ell-1}. \quad \blacksquare \end{aligned}$$

Theorem 7.31. *Let F be a field and let $p \in F[x]$ be a nonzero polynomial. Then $c \in F$ is a multiple root of p if and only if $p(c) = p'(c) = 0$.*

PROOF. If $p(c) = 0$, then we can write $p = (x - c)q$ for some polynomial $q \in F[x]$, and c is a multiple root of p if and only if $q(c) = 0$. Using the product rule, we get

$$p' = ((x - c)q)' = q + (x - c)q'.$$

Hence, $p'(c) = q(c)$. In particular, $q(c) = 0$ if and only if $p'(c) = 0$, as desired. ■

For a polynomial $p \in F[x]$, let $p^{(k)}$ denote the k -th **derivative** of p , i.e., let $p^{(k)} := p''\dots'$ (k primes).

Exercise 7.32. Let F be a field of characteristic 0 and let $p \in F[x]$ be a nonzero polynomial. Show that $c \in F$ is a root of p of multiplicity k if and only if

$$p(c) = p'(c) = \dots = p^{(k-1)}(c) = 0 \quad \text{and} \quad p^{(k)}(c) \neq 0.$$

In Exercise 7.32, it is important to assume that the characteristic of F is 0 (since in a field of positive characteristic, the derivative of a polynomial of positive degree can be zero—see Example 7.27).

Exercise 7.33. Let F be a field of characteristic $p > 0$ and let $f \in F[x]$. Show that $f^{(p)} = 0$.

¹⁹This philosophy plays a significant role in the area of **algebraic geometry**, which is often concerned with extending “geometric” or “analytic” concepts to more general “algebraic” settings.

7.E. Algebraically closed fields

A field F is **algebraically closed** if every polynomial $p \in F[x]$ of degree at least 1 has at least one root in F .

Example 7.34. The field \mathbb{R} is not algebraically closed, because the polynomial $x^2 + 1$ has no real roots. The field \mathbb{F}_2 is also not algebraically closed, because $x^2 + x + 1$ has no roots in \mathbb{F}_2 .

Exercise 7.35. Show that every algebraically closed field is infinite. *Hint:* Let F be a finite field and consider the polynomial $1 + \prod_{a \in F} (x - a)$.

The following theorem, called the *fundamental theorem of algebra*, gives the main example of an algebraically closed field:

Theorem 7.36 (Fundamental theorem of algebra). *The field \mathbb{C} of complex numbers is algebraically closed.*

Exercise 7.37. Use Theorem 7.36 and Exercise 4.27 to show that the field $\overline{\mathbb{Q}}$ of algebraic numbers is algebraically closed.

There are many proofs of Theorem 7.36, most of which rely on facts from complex analysis. We will not prove Theorem 7.36 in these notes.

Even though \mathbb{C} is the most familiar algebraically closed field, there are many others. In particular, the characteristic of an algebraically closed field can be positive (recall that if the characteristic of a field F is $p > 0$, then $p \cdot 1 = 0$ in F).

Theorem 7.38. *Let F be a field. Then there is an algebraically closed field extension $K \supseteq F$ of F .*

We won't prove Theorem 7.38 here either.

Lemma 7.39. *Let F be an algebraically closed field and let $p \in F[x]$ be a polynomial of degree $n \geq 1$. Then p **splits** in F , i.e., there exist elements $a, a_1, \dots, a_n \in F$ such that $p = a(x - a_1) \cdots (x - a_n)$.*

Exercise 7.40. Prove Lemma 7.39.

Exercise 7.41. Let F be an algebraically closed field and let $p, q \in F[x]$ be nonzero polynomials. Show that $\gcd(p, q) \neq 1$ if and only if p and q have a common root.

Corollary 7.42. *Let F be an algebraically closed field and let $p \in F[x]$ be a nonzero polynomial. Suppose that $p' \neq 0$. Then p has a multiple root if and only if $\text{res}(p, p') = 0$. ■*

Example 7.43. Let F be an algebraically closed field and let $p \in F[x]$ be a polynomial of degree 2:

$$p = a_0 + a_1x + a_2x^2,$$

where $a_0, a_1, a_2 \in F$, $a_2 \neq 0$. If $\text{char}(F) \neq 2$, then $p' = a_1 + 2a_2x$ is a polynomial of degree 1, and

$$\text{res}(p, p') = \det \begin{bmatrix} a_0 & a_1 & 0 \\ a_1 & 2a_2 & a_1 \\ a_2 & 0 & 2a_2 \end{bmatrix} = a_2(4a_0a_2 - a_1^2).$$

Since $a_2 \neq 0$, we recover the familiar fact that p has a multiple root if and only if $4a_0a_2 - a_1^2 = 0$.

7.F. The Schwartz–Zippel lemma

Let F be a field and let $p \in F[x_1, \dots, x_n]$ be a polynomial over F in n variables; that is, p is a combination of finitely many **monomials** (i.e., expressions of the form $x_1^{t_1} \cdots x_n^{t_n}$ with $t_1, \dots, t_n \in \mathbb{N}$) with coefficients in F . The **degree** of a monomial $x_1^{t_1} \cdots x_n^{t_n}$ is

$$\deg(x_1^{t_1} \cdots x_n^{t_n}) := t_1 + \cdots + t_n,$$

and the **(total) degree** of p , denoted by $\deg p$, is the largest degree of a monomial appearing in p with a nonzero coefficient. For instance,

$$\deg(x_1x_2^2 + x_1^2x_2 - 2x_1 - 2x_2 + 3) = 3.$$

As usual, the degree of the zero polynomial is, by convention, equal to $-\infty$.

Exercise 7.44. Let $P_{d,n}(F)$ denote the F -vector space of all polynomials over a field F in n variables of degree at most d . Show that

$$\dim P_{d,n}(F) = \binom{n+d}{n}.$$

Hint: https://youtu.be/w0i_ZF1GTVY.

Example 7.45. The monomials in n variables of degree at most d form a basis for $P_{d,n}(F)$. Hence, for example, a basis for $P_{2,2}(F)$ is

$$\{1, x_1, x_2, x_1^2, x_1x_2, x_2^2\},$$

and thus the dimension of the space $P_{2,2}(F)$ is 6, which is indeed equal to $\binom{2+2}{2}$.

Let $p \in F[x_1, \dots, x_n]$. The **zero locus** of p is the set

$$\mathcal{Z}_F(p) := \{(c_1, \dots, c_n) \in F^n : p(c_1, \dots, c_n) = 0\}.$$

Theorem 7.46 (Schwartz–Zippel lemma). *Let F be a field and let $p \in F[x_1, \dots, x_n]$ be a nonzero polynomial in n variables of degree at most d . Let $S \subseteq F$ be a nonempty finite set. Then*

$$|\mathcal{Z}_F(p) \cap S^n| \leq d|S|^{n-1}.$$

Equivalently, if we choose n elements $c_1, \dots, c_n \in S$ independently and uniformly at random, then

$$\mathbb{P}[p(c_1, \dots, c_n) = 0] \leq \frac{d}{|S|}.$$

Corollary 7.47. *Let F be an infinite field and let $p \in F[x_1, \dots, x_n]$. If $\mathcal{Z}_F(p) = F^n$ (in other words, if $p(c) = 0$ for all $c \in F^n$), then $p = 0$ as a polynomial.*

PROOF. Suppose that $p \neq 0$ and let $d := \deg p$. Let S be any finite subset of F of size greater than d . Then, with probability at least $1 - d/|S| > 0$, a random tuple $(c_1, \dots, c_n) \in S^n$ satisfies $p(c_1, \dots, c_n) \neq 0$, meaning that the zero locus of p cannot be all of F^n . ■

Corollary 7.48. *Let F be an infinite field and let $p_1, \dots, p_k \in F[x_1, \dots, x_n]$ be a finite collection of polynomials. If $\mathcal{Z}_F(p_1) \cup \dots \cup \mathcal{Z}_F(p_k) = F^n$, then $p_i = 0$ for some $1 \leq i \leq k$.*

PROOF. Follows from the observation that $\mathcal{Z}_F(p_1) \cup \dots \cup \mathcal{Z}_F(p_k) = \mathcal{Z}_F(p_1 \cdots p_k)$. ■

What makes Corollary 7.47 particularly useful is that it allows one to show that two polynomials are equal (i.e., their difference is the zero polynomial) without explicitly computing and comparing their coefficients. Consider, for example, the product rule for derivatives:

$$(pq)' = p'q + pq'. \tag{7.49}$$

If we write $p = a_0 + a_1x + \dots + a_nx^n$ and $q = b_0 + b_1x + \dots + b_mx^m$, then both $(pq)'$ and $p'q + pq'$ can be viewed as polynomials in the variables $x, a_0, \dots, a_n, b_0, \dots, b_m$ with integer coefficients. For instance, when $n = m = 2$, both these polynomials are equal to

$$a_0b_1 + a_1b_0 + 2a_0b_2x + 2a_1b_1x + 2a_2b_0x + 3a_1b_2x^2 + 3a_2b_1x^2 + 4a_2b_2x^3.$$

Now, we know from calculus that the product rule holds for polynomials *with real coefficients*. This means that the two multivariate polynomials representing the two sides of (7.49) take the same value for every choice of $x, a_0, \dots, a_n, b_0, \dots, b_m \in \mathbb{R}$. But then, by Corollary 7.47, this shows that they must be equal *as polynomials*. In particular, if we now plug in values for the coefficients $a_0,$

$\dots, a_n, b_0, \dots, b_m$ from *any field* F whatsoever, we would get two equal polynomials in x over F . In other words, the product rule (7.49) must hold over *every field*²⁰. Of course, in this particular example, it is easy enough to compute the two sides of (7.49) explicitly (as is done in the proof of Lemma 7.30). However, we will soon encounter situations where explicit computation is too difficult, and the above proof strategy becomes indispensable.

PROOF OF THEOREM 7.46. The proof is by induction on n . For $n = 1$, Theorem 7.46 says that if $p \in F[x]$ is a nonzero univariate polynomial of degree at most d , then it has at most d roots—but this we already know (see Exercise 7.10). Now assume that $n > 1$ and that the Schwartz–Zippel lemma holds for all polynomials in fewer than n variables. Let $p \in F[x_1, \dots, x_n]$ be a nonzero polynomial of degree at most d . Let k be the largest degree in which the variable x_n appears in p (it is possible that $k = 0$). We can write (uniquely)

$$p = p_0 + p_1x_n + p_2x_n^2 + \cdots + p_kx_n^k,$$

where p_0, \dots, p_k are polynomials in x_1, \dots, x_{n-1} . Note that, by the choice of k , p_k is a nonzero polynomial. Furthermore, $\deg p_k \leq d - k$. If we plug in any specific values, say c_1, \dots, c_{n-1} , for x_1, \dots, x_{n-1} into p , then $p(c_1, \dots, c_{n-1}, x_n)$ becomes a polynomial in a single variable x_n of degree at most k . Such a polynomial can have at most k roots—unless it's the zero polynomial, which can only happen when $p_k(c_1, \dots, c_{n-1}) = 0$. This motivates splitting the set $\mathcal{Z}_F(p) \cap S^n$ as follows:

$$\begin{aligned} |\mathcal{Z}_F(p) \cap S^n| = & |\{(c_1, \dots, c_n) \in S : p(c_1, \dots, c_n) = 0 \text{ and } p_k(c_1, \dots, c_{n-1}) \neq 0\}| \\ & + |\{(c_1, \dots, c_n) \in S : p(c_1, \dots, c_n) = 0 \text{ and } p_k(c_1, \dots, c_{n-1}) = 0\}|. \end{aligned}$$

Set

$$\begin{aligned} N_1 &:= |\{(c_1, \dots, c_n) \in S : p(c_1, \dots, c_n) = 0 \text{ and } p_k(c_1, \dots, c_{n-1}) \neq 0\}|; \\ N_2 &:= |\{(c_1, \dots, c_n) \in S : p(c_1, \dots, c_n) = 0 \text{ and } p_k(c_1, \dots, c_{n-1}) = 0\}|. \end{aligned}$$

To upper bound N_1 , observe that there are (trivially) at most $|S|^{n-1}$ choices for (c_1, \dots, c_{n-1}) , and for each such choice, assuming that $p_k(c_1, \dots, c_{n-1}) \neq 0$, there are at most k choices for c_n (since the polynomial $p(c_1, \dots, c_{n-1}, x_n)$ can have at most k roots). Thus, we have

$$N_1 \leq |S|^{n-1} \cdot k.$$

On the other hand, to upper bound N_2 , notice that, by the inductive hypothesis applied to p_k , there are at most $(d - k)|S|^{n-2}$ choices for (c_1, \dots, c_{n-1}) with $p_k(c_1, \dots, c_{n-1}) = 0$, and for each such choice there can be (trivially) at most $|S|$ choices for c_n . Therefore,

$$N_2 \leq (d - k)|S|^{n-2} \cdot |S| = (d - k)|S|^{n-1}.$$

Hence,

$$|\mathcal{Z}_F(p) \cap S^n| = N_1 + N_2 \leq |S|^{n-1}k + (d - k)|S|^{n-1} = d|S|^{n-1},$$

and we are done. ■

7.G. Application: identity testing

At this point, it would be amiss not to mention the numerous applications of the Schwartz–Zippel lemma in computer science. It often happens that a computational problem can be reduced to the question of whether or not two multivariate polynomials are equal. Calculating and comparing the coefficients of the polynomials in question can sometimes be too time-consuming; on the other hand, evaluating the given polynomials at a particular input may be more feasible. The Schwartz–Zippel lemma then gives an upper bound on the probability that two *distinct* polynomials will take the *same* value at a random input, which can then be used to show that the two given polynomials are equal—if not with certainty, then at least with overwhelming probability.

²⁰And even over every commutative ring.

The approach is best explained with the help of an example. Consider the following problem:

Let \star be a binary operation on an n -element set X . Is this operation associative?

In other words, we are given an n -by- n multiplication table for a binary operation \star , and our goal is to decide, as efficiently as we can, whether \star is associative, i.e., whether we have

$$(x \star y) \star z = x \star (y \star z) \quad \text{for all } x, y, z \in X.$$

We say that $(x, y, z) \in X^3$ is a **nonassociative triple** if

$$(x \star y) \star z \neq x \star (y \star z),$$

so our question can be restated as, is there a nonassociative triple for \star ?

Example 7.50. Let $X := \{1, \dots, n\}$, where $n \geq 3$, and define a binary operation \star on X by

$$i \star j := \begin{cases} 3 & \text{if } (i, j) \neq (1, 2); \\ 2 & \text{if } (i, j) = (1, 2). \end{cases}$$

Then the operation \star is nonassociative, and it has *only one* nonassociative triple, namely $(1, 1, 2)$.

An obvious algorithm for this problem is to test every single triple $(x, y, z) \in X^3$ for nonassociativity. Since there are n^3 triples to check, this algorithm requires roughly n^3 steps. It turns out that if we can handle a small probability of making a mistake, then there is a much faster approach:

Theorem 7.51 (Rajagopalan–Schulman). *Suppose that we are given a multiplication table for a binary operation \star on an n -element set X . There is a randomized algorithm with running time $O(n^2)$ that outputs an answer **Yes** or **No** so that:*

- if the operation \star is associative, the answer is always **Yes**;
- if the operation \star is nonassociative, then the answer is **No** with probability at least $1/2$.

Remark 7.52. It may seem like $1/2$ is a rather high probability of making a mistake. However, this probability can be made arbitrarily small by simply running the same algorithm several times. For instance, suppose that we run the algorithm twenty times in a row. This produces a sequence of twenty **Yes/No** answers. If at least one of the answers was **No**, we would know for sure that the operation \star is nonassociative. On the other hand, if all twenty answers were **Yes**, then we should feel fairly confident that \star is actually associative. Indeed, if \star weren't associative, then each one of the answers would be **Yes** with probability at most $1/2$, so the probability of answering **Yes** twenty times in a row is at most $(1/2)^{20}$, which is less than one in a million.

PROOF. The difficulty of the problem is that even if \star is nonassociative, it can still have only very few nonassociative triples (see Example 7.50). The idea is to use a bit of linear algebra to construct a *new* binary operation based on \star in such a way that if \star is associative, then so is the new operation, while if \star is nonassociative, then *at least half* of the possible inputs for the new operation form nonassociative triples.

Let F be a field of size at least 7 (in practice, it is convenient to make F a finite field, say \mathbb{F}_7). We may then view the n -element set $X = \{e_1, \dots, e_n\}$ as a basis for an n -dimensional vector space V over F . We extend \star to a bilinear operation on V in the usual way:

$$\left(\sum_{i=1}^n a_i e_i \right) \star \left(\sum_{j=1}^n b_j e_j \right) := \sum_{i=1}^n \sum_{j=1}^n (a_i b_j) (e_i \star e_j). \tag{7.53}$$

Exercise 7.54. Show that the extended operation \star on V is associative if and only if so is the original operation \star on X .

Note that, given the coordinates of two vectors $x, y \in V$, we can use formula (7.53) to compute the coordinates of the vector $x \star y$ in $O(n^2)$ steps (*why?*). And here's the randomized algorithm for testing associativity:

- Fix a set $S \subseteq F$ of size 6.
- Choose elements $a_1, \dots, a_n, b_1, \dots, b_n, c_1, \dots, c_n \in S$ independently and uniformly at random and set

$$x := \sum_{i=1}^n a_i e_i, \quad y := \sum_{j=1}^n b_j e_j, \quad \text{and} \quad z := \sum_{k=1}^n c_k e_k.$$

- Compute the vectors $(x \star y) \star z$ and $x \star (y \star z)$. If the results are equal, output **Yes**; otherwise, output **No**.

The most time-consuming part of this procedure is computing the coordinates of the vectors $(x \star y) \star z$ and $x \star (y \star z)$, as this requires performing the operation \star on elements of V four times—but this still amounts to $O(n^2)$ steps.

Obviously, if the operation \star is associative, the above algorithm always outputs **Yes**. Now suppose that \star is nonassociative. We have to prove that the probability that the algorithm's answer is **Yes** in this case is at most $1/2$.

Since we are assuming that \star is nonassociative, there exist indices $1 \leq \alpha, \beta, \gamma \leq n$ such that

$$e_\ell := (e_\alpha \star e_\beta) \star e_\gamma \neq e_\alpha \star (e_\beta \star e_\gamma) =: e_m.$$

Let $p(x, y, z)$ denote the coefficient of e_ℓ in $(x \star y) \star z$. Since we have

$$(x \star y) \star z = \left(\left(\sum_{i=1}^n a_i e_i \right) \star \left(\sum_{j=1}^n b_j e_j \right) \right) \star \left(\sum_{k=1}^n c_k e_k \right) = \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n (a_i b_j c_k) ((e_i \star e_j) \star e_k),$$

we conclude that

$$p(x, y, z) = \sum_{(e_i \star e_j) \star e_k = e_\ell} a_i b_j c_k,$$

where the sum is over all triples of indices (i, j, k) such that $(e_i \star e_j) \star e_k = e_\ell$. In particular, p is a polynomial in the $3n$ variables $a_1, \dots, a_n, b_1, \dots, b_n, c_1, \dots, c_n$ of degree 3. (Note that $p \neq 0$ since the monomial $e_\alpha e_\beta e_\gamma$ appears in p with coefficient 1.) Similarly, if we let $q(x, y, z)$ be the coefficient of e_ℓ in $x \star (y \star z)$, then

$$q(x, y, z) = \sum_{e_i \star (e_j \star e_k) = e_\ell} a_i b_j c_k,$$

which is also a polynomial in of degree at most 3. (We say “at most 3” because q may be zero.) By the choice of α, β , and γ , we have

$$p(e_\alpha, e_\beta, e_\gamma) = 1 \neq 0 = q(e_\alpha, e_\beta, e_\gamma),$$

so the polynomials p and q are distinct. This means that $p - q$ is a nonzero polynomial of degree at most 3, and thus, by the Schwartz–Zippel lemma, the probability that $p(x, y, z) = q(x, y, z)$ for randomly chosen $x, y, z \in S^n$ is at most $3/|S| = 3/6 = 1/2$, as desired. ■

Extra exercises for Section 7

Exercise 7.55. For this exercise, we need to define resultants of multivariate polynomials. Let F be a field and let $p, q \in F[x_1, \dots, x_n]$ be a pair of nonzero polynomials over F in n variables. For each $1 \leq i \leq n$, we may consider p and q as polynomials in x_i whose coefficients are, in turn, polynomials in the remaining $n - 1$ variables. This allows us to compute the **resultant** of p and q with respect to the variable x_i , denoted by $\text{res}_{x_i}(p, q)$. For instance,

$$\text{res}_x(x^2 + y^2 + z^2, xyz) = \det \begin{bmatrix} y^2 + z^2 & 0 & 0 \\ 0 & yz & 0 \\ 1 & 0 & yz \end{bmatrix} = y^4 z^2 + y^2 z^4.$$

Note that $\text{res}_{x_i}(p, q)$ is itself a polynomial over F in the remaining $n - 1$ variables.

In this exercise we work over \mathbb{C} . Given a pair of polynomials $f, g \in \mathbb{C}[t]$ of degree at least 1, we consider the set $\mathcal{S}_{f,g} \subseteq \mathbb{C}^2$ **parameterized** by f and g :

$$\mathcal{S}_{f,g} := \{(x, y) \in \mathbb{C}^2 : x = f(t), y = g(t) \text{ for some } t \in \mathbb{C}\}.$$

- (a) Consider the polynomials $x - f(t)$ and $y - g(t)$ in the three variables t, x, y . Let $p \in \mathbb{C}[x, y]$ be the polynomial given by the formula

$$p := \text{res}_t(x - f(t), y - g(t)).$$

Show that $\mathcal{S}_{f,g}$ is the zero locus of p .

- (b) Find a polynomial $p \in \mathbb{C}[x, y]$ whose zero locus is parameterized by the polynomials

$$f(t) = t^2 + t \quad \text{and} \quad g(t) = t^3 + t.$$

Exercise 7.56 (Kakeya problem over a finite field). For this exercise, F is a finite field of size q . A subset $E \subseteq F^n$ is called a **Kakeya set**²¹ if it “contains a line in every direction,” i.e., if for each nonzero vector $v \in F^n$, there is some $a \in E$ such that

$$\{a + tv : t \in F\} \subseteq E.$$

Our goal is to establish the following lower bound on the size of a Kakeya set:

Theorem 7.57 (Dvir). *If $E \subseteq F^n$ is a Kakeya set, then*

$$|E| \geq \binom{q+n-1}{n}.$$

Notice that $\binom{q+n-1}{n} \geq q^n/n! = (1/n!)|F^n|$, so it follows from Theorem 7.57 that a Kakeya set must occupy at least a $(1/n!)$ proportion of the entire space F^n , regardless of the size of the finite field F .

Suppose, towards a contradiction, that $E \subseteq F^n$ is a Kakeya set such that $|E| < \binom{q+n-1}{n}$.

- (a) Show that there is a nonzero polynomial $p \in F[x_1, \dots, x_n]$ of degree $d < q$ such that $E \subseteq \mathcal{Z}_F(p)$. *Hint:* Compare the dimension of the space of all polynomials in n variables of degree less than q with that of the space F^E of all functions from E to F .

Let p be the polynomial obtained in part (a) and let $d := \deg p$. Write

$$p = p_0 + p_1 + \dots + p_d,$$

where p_i is the i -th **homogeneous component** of p , i.e., the polynomial obtained from p by only retaining the monomials of degree i . By definition, $p_d \neq 0$ and $d \geq 1$.

Take any nonzero vector $v = (v_1, \dots, v_n) \in F^n$ and let $a = (a_1, \dots, a_n) \in E$ be such that

$$\{a + tv : t \in F\} \subseteq E.$$

(Such a exists because E is a Kakeya set.) Define $f_{v,a} \in F[t]$ to be the polynomial given by

$$f_{v,a}(t) := p(a + tv) = p(a_1 + tv_1, \dots, a_n + tv_n).$$

- (b) Show that $f_{v,a}$ is the zero polynomial.
 (c) Show that $[t^d]f_{v,a} = p_d(v)$, and hence $p_d(v) = 0$ for all $v \in F^n$.
 (d) Finish the proof of Theorem 7.57. *Hint:* Use Schwartz–Zippel.

²¹Named after the Japanese mathematician Sōichi Kakeya.

8. CLASSIFICATION OF LINEAR TRANSFORMATIONS

8.A. Conjugacy

Let V be a vector space over a field F . A **(linear) transformation** of V is a map $\varphi \in \text{Lin}(V, V)$. We can view the pair (V, φ) as an algebraic structure in its own right, and there is a natural notion of isomorphism for such structures. Namely, if W is another vector space over F and $\psi \in \text{Lin}(W, W)$, then the structures (V, φ) and (W, ψ) are **isomorphic** if there is a linear bijection $\pi: V \rightarrow W$ (an **isomorphism**) such that for all $x \in V$,

$$\pi(\varphi(x)) = \psi(\pi(x));$$

or, to put it more concisely,

$$\pi \circ \varphi = \psi \circ \pi.$$

The diagram below illustrates this situation:

$$\begin{array}{ccc} V & \xrightarrow{\varphi} & V \\ \pi \downarrow & \searrow & \downarrow \pi \\ W & \xrightarrow{\psi} & W \end{array} \quad \begin{array}{c} \pi \circ \varphi = \psi \circ \pi \end{array}$$

For the structures (V, φ) and (W, ψ) to be isomorphic in the above sense, the vector spaces V and W must themselves be isomorphic. Therefore, it is convenient to assume that $V = W$ and just compare pairs of transformations $\varphi, \psi \in \text{Lin}(V, V)$. This motivates the following definition:

Definition 8.1. Let V be a vector space over a field F . We say that linear transformations $\varphi, \psi \in \text{Lin}(V, V)$ are **conjugate**, in symbols $\varphi \cong \psi$, if there is a linear bijection $\pi: V \rightarrow V$ such that

$$\pi \circ \varphi = \psi \circ \pi.$$

There are a number of equivalent ways to define conjugacy of transformations $\varphi, \psi \in \text{Lin}(V, V)$. For instance, $\varphi \cong \psi$ if and only if there is a linear bijection $\pi: V \rightarrow V$ such that $\psi = \pi \circ \varphi \circ \pi^{-1}$. Also, we have $\varphi \cong \psi$ if and only if there exist bases X, Y for V such that $[\varphi]_{X,X} = [\psi]_{Y,Y}$ (**why?**).

Our goal in this section is to classify transformations $\varphi \in \text{Lin}(V, V)$ up to conjugacy, *provided* that the space V is finite-dimensional and the field F is algebraically closed. When $\dim V = n$, transformations of V can be identified with n -by- n matrices over F , and so we will freely switch between working with transformations and with matrices. Thus, for example, we say that two matrices $A, B \in M_{n \times n}(F)$ are **conjugate** if there is a matrix $C \in M_{n \times n}(F)$ of rank n such that $CA = BC$. Also, describing a transformation $\varphi \in \text{Lin}(V, V)$ of an n -dimensional space V up to conjugacy is tantamount to finding a basis X for V in which the matrix $[\varphi]_{X,X}$ has a particularly simple form.

Example 8.2. If $\dim V = 1$, then every linear transformation of V has the form $x \mapsto ax$ for a fixed scalar $a \in F$, and it is clear that any such transformation is only conjugate to itself.

Example 8.3. Already in the case when $\dim V = 2$, the situation becomes somewhat complicated. We will show that if F is algebraically closed, then for each transformation $\varphi \in \text{Lin}(V, V)$ of a 2-dimensional F -vector space V , there is a basis $X = (x_1, x_2)$ for V in which the matrix $[\varphi]_{X,X}$ looks

$$\text{either like this: } \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}, \quad \lambda_1, \lambda_2 \in F, \quad \text{or like this: } \begin{bmatrix} \lambda & 0 \\ 1 & \lambda \end{bmatrix}, \quad \lambda \in F.$$

In other words, we either have

$$\varphi(x_1) = \lambda_1 x_1 \quad \text{and} \quad \varphi(x_2) = \lambda_2 x_2,$$

or else,

$$\varphi(x_1) = \lambda x_1 + x_2 \quad \text{and} \quad \varphi(x_2) = \lambda x_2.$$

8.B. Eigenvectors and eigenvalues

Definition 8.4. Let V be a vector space over a field F and let $\varphi \in \text{Lin}(V, V)$. A subspace $W \subseteq V$ is φ -invariant if for all $x \in W$, we have $\varphi(x) \in W$ as well.

If $\varphi \in \text{Lin}(V, V)$ and $W \subseteq V$ is a φ -invariant subspace, then the restriction $\varphi|_W$ of φ to W can be thought of as a linear transformation of W . Thus, if we find a φ -invariant subspace W , then we can first analyze the behavior φ on W and then hope to deal with the action of φ on the rest of V separately.

Example 8.5. The entire space V itself and the zero space $\{0\}$ are φ -invariant for every linear transformation $\varphi \in \text{Lin}(V, V)$.

Example 8.6. Suppose that $\varphi \in \text{Lin}(\mathbb{R}^3, \mathbb{R}^3)$ is the \mathbb{R} -linear transformation of \mathbb{R}^3 given by rotation around the vertical axis by some fixed angle α (see Fig. 5(a)). Then the vertical axis is φ -invariant, and φ acts on it as the identity transformation; the horizontal plane is also φ -invariant, and φ acts on it as the rotation around the origin by the angle α (see Fig. 5(b)(c)).




Figure 5. Rotation in \mathbb{R}^3 around the vertical axis.

Of particular interest to us are the simplest nontrivial invariant subspaces, i.e., invariant subspaces of dimension 1. Let V be a vector space over a field F and let $\varphi \in \text{Lin}(V, V)$. A nonzero element $x \in V$ is called an **eigenvector**²² of φ if the 1-dimensional subspace $\text{Span}(\{x\})$ is φ -invariant, i.e., if $\varphi(x) = \lambda x$ for some $\lambda \in F$. This λ is called the **eigenvalue** of φ corresponding to x . The set of all eigenvalues of φ is called the **spectrum** of φ and is denoted $\text{Spec}(\varphi)$. For $\lambda \in \text{Spec}(\varphi)$, the set

$$E(\lambda) := \{x \in V : \varphi(x) = \lambda x\}$$

is called the **eigenspace** of φ corresponding to λ . The term “*eigenspace*” is justified, because

$$E(\lambda) = \ker(\varphi - \lambda \text{id}_V),$$

and thus it is indeed a subspace of V . Note that $E(\lambda)$ is the set of all eigenvectors corresponding to λ , together with the zero vector. By definition, if $\lambda \in \text{Spec}(\varphi)$, then $\dim E(\lambda) \geq 1$.

Now suppose that V is a finite-dimensional F -vector space and let $n := \dim V$. Suppose that $\varphi \in \text{Lin}(V, V)$ is a linear transformation of V . What are the eigenvalues of φ ? By definition, $t \in F$ is an eigenvalue of φ if and only if

$$\ker(\varphi - t \text{id}_V) \neq \{0\},$$

²²Not named after the famous German mathematician Eugen Eigen.

which in turn is equivalent to

$$\det(\varphi - t \operatorname{id}_V) = 0.$$

(This is where we use that V is finite-dimensional.) To compute the determinant, pick any ordered basis X for V and let $A := [\varphi]_{X,X}$ be the matrix representing φ in this basis. Then

$$\det(\varphi - t \operatorname{id}_V) = \det(A - tI_n) = \det \begin{bmatrix} A(1,1) - t & A(1,2) & \cdots & A(1,n) \\ A(2,1) & A(2,2) - t & \cdots & A(2,n) \\ \vdots & \vdots & \ddots & \vdots \\ A(n,1) & A(n,2) & \cdots & A(n,n) - t \end{bmatrix}$$

After expanding this determinant using the Leibniz formula, we obtain a polynomial in t over F of degree n . This polynomial is called the **characteristic polynomial** of the transformation φ (or of the matrix A) and we denote it by $\operatorname{Char}_\varphi$ (or Char_A).

Exercise 8.7. Show that $[t^n]\operatorname{Char}_\varphi = (-1)^n$.

Example 8.8. In the 2-by-2 case, the characteristic polynomial of a matrix

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$$

is

$$\operatorname{Char}_A(t) = \det \begin{bmatrix} a_{11} - t & a_{12} \\ a_{21} & a_{22} - t \end{bmatrix} = t^2 - (a_{11} + a_{22})t + a_{11}a_{22} - a_{12}a_{21}.$$

From the above discussion, we obtain the following conclusion:

Theorem 8.9. Let V be a finite-dimensional F -vector space and let $\varphi \in \operatorname{Lin}(V, V)$. Then $\lambda \in F$ is an eigenvalue of F if and only if λ is a root of $\operatorname{Char}_\varphi$. ■

Corollary 8.10. Let V be a finite-dimensional vector space over an algebraically closed field F and let $\varphi \in \operatorname{Lin}(V, V)$. If $\dim V \geq 1$, then φ has an eigenvalue. ■

Example 8.11. Consider the 2-by-2 matrix

$$A := \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Its characteristic polynomial is $\operatorname{Char}_A(t) = t^2 - 1$, so, viewed as a matrix over \mathbb{R} , it has two eigenvalues: 1 and -1 . On the other hand, the matrix

$$B := \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$$

has characteristic polynomial $\operatorname{Char}_B(t) = t^2 + 1$, so B has no eigenvalues over \mathbb{R} . Nevertheless, it has two eigenvalues over \mathbb{C} , namely i and $-i$.

Example 8.12. Even over an algebraically closed field, a linear transformation of an *infinite-dimensional* space may have no eigenvalues. For example, let F be any field and consider the F -vector space $F^{\mathbb{N}}$ of all infinite sequences (x_0, x_1, \dots) of elements of F . Then the map $\varphi \in \operatorname{Lin}(F^{\mathbb{N}}, F^{\mathbb{N}})$ given by the formula

$$\varphi(x_0, x_1, x_2, \dots) := (0, x_0, x_1, \dots)$$

has no eigenvalues ([exercise!](#)).

Exercise 8.13. Let V be an n -dimensional vector space over a field F and let $\varphi \in \operatorname{Lin}(V, V)$. Show that $|\operatorname{Spec}(\varphi)| \leq n$.

Exercise 8.14. Let F be a field and let $A \in M_{n \times n}(F)$. Show that $\operatorname{Char}_A = \operatorname{Char}_{A^\top}$ and conclude that $\operatorname{Spec}(A) = \operatorname{Spec}(A^\top)$.

The following (almost trivial) observation is often useful. Suppose that V is a finite-dimensional vector space over an algebraically closed field F and let $\varphi \in \text{Lin}(V, V)$. If a subspace $W \subseteq V$ is φ -invariant, then the restriction $\varphi|_W$ is a linear transformation of W , and hence, assuming that $W \neq \{0\}$, it has an eigenvalue as well as a corresponding eigenvector. But every eigenvector of $\varphi|_W$ is, of course, also an eigenvector of φ itself, and thus W contains an eigenvector of φ .

8.C. Diagonalizable transformations

There isn't a question that one can't immediately answer about a diagonal matrix.

Kevin O'Meara, John Clark, and Charles Vinsonhaler

Definition 8.15. Let V be a vector space over a field F and let $\varphi \in \text{Lin}(V, V)$. We say that φ is **diagonalizable** if V has a basis consisting of eigenvectors of φ .

The word “diagonalizable” is explained by the following observation. Suppose that $\varphi \in \text{Lin}(V, V)$ is a diagonalizable linear transformation of an n -dimensional vector space V and let $X = (x_1, \dots, x_n)$ be an ordered basis for V consisting of eigenvectors of φ . Then the matrix $[\varphi]_{X,X}$ looks like this:

$$[\varphi]_{X,X} = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_n \end{bmatrix},$$

where $\lambda_1, \dots, \lambda_n$ are the eigenvalues of φ corresponding to x_1, \dots, x_n , respectively. A matrix of this form is called **diagonal**.

Diagonalizable transformations are particularly easy to understand. Unfortunately, not every transformation is diagonalizable, even if we are working in a finite-dimensional vector space over an algebraically closed field.

Example 8.16. Consider the matrix

$$A := \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \in M_{2 \times 2}(\mathbb{C}).$$

Its characteristic polynomial is $\text{Char}_A(t) = (t - 1)^2$, and thus its only eigenvalue is 1. This means that if A were diagonalizable, then there would exist a basis for $M_{2 \times 1}(\mathbb{C})$ consisting of two elements x_1 and x_2 such that $Ax_1 = x_1$ and $Ax_2 = x_2$. But since $\{x_1, x_2\}$ is a basis, that would imply $Ax = x$ for all $x \in M_{2 \times 1}(\mathbb{C})$, which is absurd.

Nevertheless, over an algebraically closed field, *most* linear transformations are diagonalizable, in a certain precise sense (explained below), and so the bulk of the work in classifying linear transformations up to conjugacy really consists in “taming” the nasty “exceptional” cases.

Lemma 8.17. *Let V be an F -vector space and let $\varphi \in \text{Lin}(V, V)$. Suppose that $S \subseteq V$ is a set whose elements are eigenvectors of φ corresponding to distinct eigenvalues. Then S is independent.*

PROOF. Suppose, towards a contradiction, that S is not independent, and let

$$a_1x_1 + \cdots + a_nx_n = 0$$

be an equality, where: a_1, \dots, a_n are nonzero elements of F ; x_1, \dots, x_n are pairwise distinct elements of S ; and $n \geq 1$ is the smallest possible. Note that we must have $n \geq 2$, since an equality

of the form $a_1x_1 = 0$ is impossible. For each $1 \leq i \leq n$, let λ_i be the eigenvalue of φ corresponding to x_i . By the assumption on S , the elements $\lambda_1, \dots, \lambda_n$ are pairwise distinct. Now we write

$$\begin{aligned} 0 &= (\varphi - \lambda_n \text{id}_V)(0) = (\varphi - \lambda_n \text{id}_V)(a_1x_1 + \dots + a_nx_n) \\ &= a_1(\lambda_1 - \lambda_n)x_1 + \dots + \cancel{a_n(\lambda_n - \lambda_n)x_n}. \end{aligned}$$

Let $b_i := a_i(\lambda_i - \lambda_n)$. Then the coefficients b_1, \dots, b_{n-1} are nonzero, while

$$b_1x_1 + \dots + b_{n-1}x_{n-1} = 0,$$

contradicting the choice of n . ■

Exercise 8.18. Lemma 8.17 can be used to give an alternative solution to Exercise 3.43. Consider the \mathbb{R} -vector space $\mathbb{R}^{\mathbb{N}}$ of all infinite sequences of reals. For each $\alpha \in \mathbb{R}$, let

$$e_\alpha := (1, \alpha, \alpha^2, \alpha^3, \dots).$$

Let $\varphi: \mathbb{R}^{\mathbb{N}} \rightarrow \mathbb{R}^{\mathbb{N}}$ be the linear map given by

$$\varphi(x_0, x_1, x_2, \dots) := (x_1, x_2, x_3, \dots).$$

Show that for each $\alpha \in \mathbb{R}$, the sequence e_α is an eigenvector of φ with eigenvalue α , and conclude that the set $\{e_\alpha : \alpha \in \mathbb{R}\}$ is independent.

Corollary 8.19. *Let V be an n -dimensional vector space over a field F and let $\varphi \in \text{Lin}(V, V)$. If $|\text{Spec}(\varphi)| = n$, then φ is diagonalizable.*

PROOF. Suppose the eigenvalues of φ are $\lambda_1, \dots, \lambda_n$. For each $1 \leq i \leq n$, pick any eigenvector x_i corresponding to λ_i . By Lemma 8.17, the set $\{x_1, \dots, x_n\}$ is independent. But $\dim V = n$, so this set must be a basis. ■

In the setting of Corollary 8.19, saying that $|\text{Spec}(\varphi)| = n$ is the same as to say that the polynomial Char_φ has n distinct roots. Over an algebraically closed field, a “typical” polynomial of degree n does have n distinct roots—meaning that a “typical” transformation $\varphi \in \text{Lin}(V, V)$ is diagonalizable. (We will use a precise version of this idea in the next subsection.)

8.D. The Cayley–Hamilton theorem

Let V be an F -vector space and let $\varphi \in \text{Lin}(V, V)$. If we wish to understand the structure of φ , there are two kinds of things we can try: we could investigate what happens when φ is applied repeatedly, and we could use the vector space structure of V to take linear combinations of φ with other linear transformations. This leads us to the following definitions. For each $n \in \mathbb{N}$, let

$$\varphi^n := \underbrace{\varphi \circ \varphi \circ \dots \circ \varphi}_{n \text{ terms}}.$$

By definition, $\varphi^0 = \text{id}_V$, $\varphi^1 = \varphi$, $\varphi^2 = \varphi \circ \varphi$, and so on. Given a polynomial

$$p = a_0 + a_1t + \dots + a_nt^n \in F[t],$$

define

$$p(\varphi) := a_0 \text{id}_V + a_1\varphi + \dots + a_n\varphi^n \in \text{Lin}(V, V).$$

Similarly, for a matrix $A \in M_{n \times n}(F)$, we let

$$p(A) := a_0I_n + a_1A + \dots + a_nA^n \in M_{n \times n}(F).$$

Lemma 8.20. *Let V be an F -vector space and let $\varphi \in \text{Lin}(V, V)$. If $p, q \in F[t]$, then*

$$p(\varphi) \circ q(\varphi) = q(\varphi) \circ p(\varphi) = (pq)(\varphi).$$

PROOF. It is enough to consider the case when $p = t^k$ and $q = t^\ell$ (why?). Then

$$p(\varphi) = \varphi^k, \quad q(\varphi) = \varphi^\ell, \quad (pq)(\varphi) = \varphi^{k+\ell},$$

and, of course, $\varphi^k \circ \varphi^\ell = \varphi^\ell \circ \varphi^k = \varphi^{k+\ell}$. ■

Exercise 8.21. Let V be an F -vector space and let $\varphi \in \text{Lin}(V, V)$. Let $x \in V$ be an eigenvector of φ with eigenvalue λ . Show that for each polynomial $p \in F[t]$, x is an eigenvector of $p(\varphi)$ with eigenvalue $p(\lambda)$. In other words, if $\varphi(x) = \lambda \cdot x$, then $p(\varphi)(x) = p(\lambda) \cdot x$.

Lemma 8.22. Let V be a finite-dimensional vector space over a field F and let $\varphi \in \text{Lin}(V, V)$. Then there is a nonzero polynomial $p \in F[t]$ such that $p(\varphi) = 0$.

PROOF. We have already seen a very similar argument in the proof of Lemma 4.20. Let $n := \dim V$. Then $\dim \text{Lin}(V, V) = n^2$, and hence the tuple

$$(\text{id}_V, \varphi, \varphi^2, \dots, \varphi^{n^2})$$

is not independent. Thus, there exist coefficients $a_0, \dots, a_{n^2} \in F$, not all zero, such that

$$a_0 \text{id}_V + a_1 \varphi + \dots + a_{n^2} \varphi^{n^2} = 0.$$

This means that the polynomial $p := a_0 + a_1 t + \dots + a_{n^2} t^{n^2}$ is as desired. ■

The above proof of Lemma 8.22 produces a polynomial p of degree at most n^2 , where n is the dimension of V . It turns out that one can find a suitable polynomial of degree n ; in fact, the *characteristic polynomial* of φ does the trick:

Theorem 8.23 (Frobenius, Cayley–Hamilton theorem²³). Let V be a finite-dimensional vector space over a field F and let $\varphi \in \text{Lin}(V, V)$. Then $\text{Char}_\varphi(\varphi) = 0$. Equivalently, if $A \in M_{n \times n}(F)$ is an n -by- n matrix, then $\text{Char}_A(A) = 0$.

Remark 8.24. It is tempting to give the following “proof”: “Let $A \in M_{n \times n}(F)$. By definition, $\text{Char}_A(t) = \det(A - tI_n)$. Hence,

$$\text{Char}_A(A) = \det(A - AI_n) = \det(A - A) = \det(0) = 0,$$

as desired.” This is, of course, absurd. For instance, this so-called “argument” shows that $\text{Char}_A(A)$, an n -by- n matrix, is equal to $0 \in F$, a scalar.

Example 8.25. Suppose that A is a 2-by-2 matrix and write

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}.$$

Then $\text{Char}_A(t) = t^2 - (a_{11} + a_{22})t + a_{11}a_{22} - a_{12}a_{21}$, and thus

$$\begin{aligned} \text{Char}_A(A) &= A^2 - (a_{11} + a_{22})A + (a_{11}a_{22} - a_{12}a_{21})I_2 \\ &= \begin{bmatrix} a_{11}^2 + a_{12}a_{21} & a_{11}a_{12} + a_{12}a_{22} \\ a_{11}a_{21} + a_{21}a_{22} & a_{12}a_{21} + a_{22}^2 \end{bmatrix} - \begin{bmatrix} a_{11}^2 + a_{11}a_{22} & a_{11}a_{12} + a_{12}a_{22} \\ a_{11}a_{21} + a_{21}a_{22} & a_{11}a_{22} + a_{22}^2 \end{bmatrix} \\ &\quad + \begin{bmatrix} a_{11}a_{22} - a_{12}a_{21} & 0 \\ 0 & a_{11}a_{22} - a_{12}a_{21} \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \end{aligned}$$

as claimed. Notice that, regardless of the size of the matrix A , the entries of $\text{Char}_A(A)$ are going to be polynomials in the entries of A with integer coefficients, and the Cayley–Hamilton theorem asserts that all these polynomials are zero. This observation plays a crucial role in the proof of the Cayley–Hamilton theorem given below.

²³The first complete proof of this theorem was given by Ferdinand Georg Frobenius in 1878. However, it is usually called the *Cayley–Hamilton theorem*, after Arthur Cayley and William Rowan Hamilton, who considered some of its special cases in the 1850s.

PROOF. We proceed in three steps.

First, assume that φ is diagonalizable. Let $\{x_1, \dots, x_n\}$ be a basis for V consisting of eigenvectors of φ and for each $1 \leq i \leq n$, let λ_i be the eigenvalue of φ corresponding to x_i . By Exercise 8.21,

$$\text{Char}_\varphi(\varphi)(x_i) = \text{Char}_\varphi(\lambda_i) \cdot x_i = 0 \cdot x_i = 0,$$

where we use the fact that eigenvalues of φ are roots of the characteristic polynomial Char_φ . Since $\text{Char}_\varphi(\varphi)(x_i) = 0$ for every basis vector x_i , $\text{Char}_\varphi(\varphi)$ must be the zero function.

Second, we shall consider arbitrary φ but assume that F is an algebraically closed field of characteristic 0, for example, the field \mathbb{C} of complex numbers. It will be more convenient to work with matrices rather than with linear transformations. Let A be an n -by- n matrix and write

$$A = \begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{bmatrix} \quad (8.26)$$

We already know that $\text{Char}_A(A) = 0$ if A is diagonalizable. On the other hand, if A is *not* diagonalizable, then, by Corollary 8.19, A has fewer than n eigenvalues, which, since F is algebraically closed, means that the polynomial Char_A has a multiple root. Since $\deg \text{Char}_A = n$ and $\text{char}(F) = 0$, Char'_A is a nonzero polynomial of degree $n - 1$, and hence, by Corollary 7.42, Char_A has a multiple root if and only if $\text{res}(\text{Char}_A, \text{Char}'_A) = 0$. To summarize, every n -by- n matrix A over F has at least one of the following properties:

$$\text{Char}_A(A) = 0 \quad \text{or} \quad \text{res}(\text{Char}_A, \text{Char}'_A) = 0. \quad (8.27)$$

Take any $1 \leq i, j \leq n$. As observed in Example 8.25, the (i, j) -th entry of the matrix $\text{Char}_A(A)$ is a polynomial in the n^2 variables a_{11}, \dots, a_{nn} with integer coefficients; denote this polynomial by $p_{ij}(a_{11}, \dots, a_{nn})$. Similarly, the expression $\text{res}(\text{Char}_A, \text{Char}'_A)$ is *also* a polynomial in a_{11}, \dots, a_{nn} with integer coefficients, which we denote by $q(a_{11}, \dots, a_{nn})$. For example, if $n = 2$, then

$$\begin{aligned} q(a_{11}, a_{12}, a_{21}, a_{22}) &= \text{res}(t^2 - (a_{11} + a_{22})t + a_{11}a_{22} - a_{12}a_{21}, 2t - a_{11} - a_{22}) \\ &= \det \begin{bmatrix} a_{11}a_{22} - a_{12}a_{21} & -a_{11} - a_{22} & 0 \\ -a_{11} - a_{22} & 2 & -a_{11} - a_{22} \\ 1 & 0 & 2 \end{bmatrix} \\ &= -a_{11}^2 + 2a_{11}a_{22} - 4a_{12}a_{21} - a_{22}^2. \end{aligned}$$

From (8.27), we know that

$$p_{ij}(a_{11}, \dots, a_{nn}) = 0 \quad \text{or} \quad q(a_{11}, \dots, a_{nn}) = 0 \quad \text{for all } a_{11}, \dots, a_{nn} \in F.$$

Since F is infinite, Corollary 7.48 implies that we have $p_{ij} = 0$ or $q = 0$ *as polynomials*. Since $q \neq 0$ (because there exist matrices with n distinct eigenvalues), we conclude that $p_{ij} = 0$. But then $p_{ij}(a_{11}, \dots, a_{nn}) = 0$ for all $a_{11}, \dots, a_{nn} \in F$, regardless of whether the matrix A is diagonalizable. Since this is true for all i, j , we conclude that $\text{Char}_A(A) = 0$ for all $A \in M_{n \times n}(F)$.

Third, we consider the general case. The above argument carried out over any algebraically closed field of characteristic 0, e.g., over \mathbb{C} , shows that the integer polynomial p_{ij} representing the (i, j) -th entry of $\text{Char}_A(A)$ for A as in (8.26) is, in fact, the zero polynomial. But then $p_{ij}(a_{11}, \dots, a_{nn}) = 0$ for all $a_{11}, \dots, a_{nn} \in F$ *regardless of the choice of the field F* . Thus, $\text{Char}_A(A) = 0$ for all n -by- n matrices A over *any* field F , as desired. \blacksquare ²⁴

²⁴In fact, this argument shows that the Cayley–Hamilton theorem holds not only over every field, but over every commutative ring as well.

8.E. Generalized eigenspaces

Definition 8.28. Let V be a finite-dimensional vector space over a field F and let $\varphi \in \text{Lin}(V, V)$. Let $\lambda \in \text{Spec}(\varphi)$ and suppose that the multiplicity of λ as a root of Char_φ is m . Define

$$G(\lambda) := \ker((\varphi - \lambda \text{id}_V)^m).$$

We call $G(\lambda)$ the **generalized eigenspace** of φ corresponding to the eigenvalue λ .

Exercise 8.29. Let V be a finite-dimensional F -vector space and let $\varphi \in \text{Lin}(V, V)$. Show that for each $\lambda \in \text{Spec}(\varphi)$, we have $E(\lambda) \subseteq G(\lambda)$.

Exercise 8.30. Let V be a finite-dimensional F -vector space and let $\varphi \in \text{Lin}(V, V)$. Show that for each $\lambda \in \text{Spec}(\varphi)$, the space $G(\lambda)$ is φ -invariant.

Theorem 8.31 (Generalized eigenspace decomposition). *Let V be a finite-dimensional vector space over an algebraically closed field F and let $\varphi \in \text{Lin}(V, V)$. Let $\lambda_1, \dots, \lambda_k$ be the distinct eigenvalues of φ and let G_1, \dots, G_k be the corresponding generalized eigenspaces. Then every vector $x \in V$ can be expressed uniquely as $x = x_1 + \dots + x_k$, where $x_1 \in G_1, \dots, x_k \in G_k$.*

PROOF. Let $n := \dim V$. For brevity, let $p := \text{Char}_\varphi$. For each $1 \leq i \leq k$, let m_i be the multiplicity of λ_i as a root of p and define

$$p_i := (t - \lambda_i)^{m_i}.$$

Thus, we can write

$$p = (-1)^n (t - \lambda_1)^{m_1} \dots (t - \lambda_k)^{m_k} = (-1)^n p_1 \dots p_k.$$

Also, let $q_i := p/p_i$; in other words,

$$q_i = (-1)^n p_1 \dots p_{i-1} p_{i+1} \dots p_k.$$

Note that, by definition, $G_i = \ker(p_i(\varphi))$.

Claim. For each $1 \leq i \leq k$, $\text{im}(q_i(\varphi)) \subseteq G_i$.

Proof. Take any $y \in V$. We need to show that $q_i(\varphi)(y) \in G_i$, i.e., $p_i(\varphi)q_i(\varphi)(y) = 0$. But $p_i q_i = p$, and $p(\varphi) = 0$ by the Cayley–Hamilton theorem, so $p_i(\varphi)q_i(\varphi)(y) = p(\varphi)(y) = 0$, as claimed. \dashv

Since $\lambda_1, \dots, \lambda_k$ are pairwise distinct, we have $\gcd(q_1, \dots, q_k) = 1$. Therefore, by Exercise 7.19, there exist polynomials $u_1, \dots, u_k \in F[t]$ such that $q_1 u_1 + \dots + q_k u_k = 1$. Take any $x \in V$. We have

$$x = \text{id}_V(x) = (q_1 u_1 + \dots + q_k u_k)(\varphi)(x) = \underbrace{q_1(\varphi)u_1(\varphi)(x)}_{\in \text{im}(q_1(\varphi)) \subseteq G_1} + \dots + \underbrace{q_k(\varphi)u_k(\varphi)(x)}_{\in \text{im}(q_k(\varphi)) \subseteq G_k},$$

as desired. It remains to prove uniqueness. To that end, it is enough to show that if $x_1 + \dots + x_k = 0$ and $x_1 \in G_1, \dots, x_k \in G_k$, then $x_1 = \dots = x_k = 0$ (**why?**). Consider any $1 \leq i \leq k$. Since $x_i \in G_i$, we have $p_i(\varphi)(x_i) = 0$. On the other hand, for each $j \neq i$, the polynomial p_j divides q_i , and hence $q_i(\varphi)(x_j) = 0$. Since $\gcd(p_i, q_i) = 1$, there are polynomials $u, v \in F[t]$ such that $u p_i + v q_i = 1$. Then

$$\begin{aligned} x_i = \text{id}_V(x_i) &= (u p_i + v q_i)(\varphi)(x_i) = \underbrace{u(\varphi)p_i(\varphi)(x_i)}_{= 0} + v(\varphi)q_i(\varphi)(x_i) \\ &= v(\varphi)q_i(\varphi)(x_i) \end{aligned}$$

$$[\text{since } q_i(\varphi)(x_j) = 0 \text{ for } j \neq i] = v(\varphi)q_i(\varphi)(x_1 + \dots + x_k) = 0,$$

and we are done. ■

Exercise 8.32. In the setting of Theorem 8.31, let m_i be the multiplicity of λ_i as a root of Char_φ . Show that $\dim G_i = m_i$. *Hint:* What is the characteristic polynomial of $\varphi|_{G_i}$?

Exercise 8.33. Let V be a finite-dimensional vector space over an algebraically closed field F and let $\varphi \in \text{Lin}(V, V)$. Show that for each $\lambda \in \text{Spec}(\varphi)$, we have

$$G(\lambda) = \{x \in V : (\varphi - \lambda \text{id}_V)^m(x) = 0 \text{ for some } m \in \mathbb{N}\}.$$

Suppose that we are in the setting of Theorem 8.31; i.e., let $\varphi \in \text{Lin}(V, V)$ be a linear transformation of a finite-dimensional vector space V over an algebraically closed field F , let $\lambda_1, \dots, \lambda_k$ be the distinct eigenvalues of φ , and let G_1, \dots, G_k be the corresponding generalized eigenspaces. For each $1 \leq i \leq k$, let m_i be the multiplicity of λ_i as a root of Char_φ (by Exercise 8.32, $m_i = \dim G_i$). Set $\varphi_i := \varphi|_{G_i}$, so φ_i is a linear transformation of G_i . Thanks to Theorem 8.31, to understand the structure of φ , we just need to understand the transformations $\varphi_1, \dots, \varphi_k$ individually. Indeed, if we know $\varphi_1, \dots, \varphi_k$, then, for each $x \in V$, the value $\varphi(x)$ is determined uniquely as follows: If we write $x = x_1 + \dots + x_k$ with $x_1 \in G_1, \dots, x_k \in G_k$, then $\varphi(x) = \varphi_1(x_1) + \dots + \varphi_k(x_k)$.

Another way to phrase this is in terms of matrices. For each $1 \leq i \leq k$, pick an arbitrary ordered basis X_i for G_i . Set $X := X_1 \hat{\ } \dots \hat{\ } X_k$, where $\hat{\ }$ indicates concatenation of finite tuples.²⁵ Then, according to Theorem 8.31, X is an ordered basis for V . If we set $A_i := [\varphi_i]_{X_i, X_i}$ to be the matrix representing φ_i with respect to the basis X_i , then the matrix $A := [\varphi]_{X, X}$ representing φ in the basis X has the following “block-diagonal” form:

$$A = \begin{bmatrix} A_1 & & & \\ & A_2 & & \\ & & \ddots & \\ & & & A_k \end{bmatrix} \quad (\text{the entries outside of the diagonal “blocks” are zero}).$$

Thus, if we manage to choose the bases X_1, \dots, X_k so that the corresponding matrices A_1, \dots, A_k have a particularly “simple” structure, then we would obtain a “simple” matrix A representing the transformation φ .

It remains to investigate the structure of the transformations φ_i , $1 \leq i \leq k$. It is actually more natural to look at the transformation $\psi_i := \varphi_i - \lambda_i \text{id}$ instead. Of course, if we know ψ_i , then we know φ_i as well, since for all $x \in G_i$, $\varphi_i(x) = \psi_i(x) + \lambda_i x$. By the definition of G_i , the transformation ψ_i has the property that $\psi_i^{m_i} = 0$. Such transformations are called *nilpotent*, and in the next subsection we shall see that their behavior can be analyzed very precisely.

8.F. Structure of nilpotent transformations

Definition 8.34. Let V be a vector space over a field F . A linear transformation $\varphi \in \text{Lin}(V, V)$ is called **nilpotent** if $\varphi^m = 0$ for some $m \in \mathbb{N}$. The least such m is called the **nilpotency degree** of φ and is denoted by $\text{ndeg}(\varphi)$.

Let V be a finite-dimensional F -vector space and let $\varphi \in \text{Lin}(V, V)$ be nilpotent. The structure of φ can be understood particularly well using a special type of basis for V , called a chain basis. A **chain basis** for φ is a basis $X \subseteq V$ for V such that:

- for all $x \in X$, either $\varphi(x) \in X$ or $\varphi(x) = 0$; and
- for each $x \in X$, there is at most one element $y \in X$ with $\varphi(y) = x$.

The name “chain basis” is motivated by the following considerations. Let X be a chain basis for φ . If we represent each element of X by a dot and put an arrow pointing from the dot corresponding to each element $x \in X$ to the dot corresponding to $\varphi(x)$ (whenever $\varphi(x) \neq 0$), then the resulting directed graph will look like a collection of disjoint “chains”:

²⁵The **concatenation** of two sequences $X = (x_1, \dots, x_s)$ and $Y = (y_1, \dots, y_t)$ is the sequence

$$X \hat{\ } Y := (x_1, \dots, x_s, y_1, \dots, y_t).$$




Figure 6. A chain basis represented by a directed graph (the rightmost elements are mapped to 0 by φ).

Indeed, since φ is nilpotent, if we follow the arrows starting from any dot, we must eventually reach a “dead end” (corresponding to an element of X that is sent to 0 by φ). On the other hand, starting from any of the “dead ends,” we can follow the arrows backwards and build the corresponding “chain” (when we try to follow the arrows backwards, there are no “forks in the road,” because for each $x \in X$, there is at most one arrow pointing to x).

If X is a chain basis for φ , then the elements of X can be put in a sequence (x_1, \dots, x_n) so that for each $1 \leq i \leq n$, we have either $\varphi(x_i) = x_{i+1}$ or $\varphi(x_i) = 0$ (see Fig. 6). With respect to such an ordered chain basis, the matrix A that represents φ is “block-diagonal,” where the blocks B_1, \dots, B_k correspond to the “chains,” and each individual block B_i looks like this:

$$B_i = \begin{bmatrix} 0 & & & & & \\ 1 & 0 & & & & \\ & 1 & 0 & & & \\ & & \ddots & \ddots & & \\ & & & 1 & 0 & \end{bmatrix},$$

i.e., the only nonzero entries of B_i are located immediately below the main diagonal and are all equal to 1.²⁶

Theorem 8.35. *Let V be a finite-dimensional vector space over a field F and let $\varphi \in \text{Lin}(V, V)$ be nilpotent. Then there is a chain basis $X \subseteq V$ for φ .*

PROOF. Before proceeding with the proof, we need some terminology. Let $W \subseteq V$ be a subspace. A tuple $(x_1, \dots, x_k) \in V^k$ is **independent over W** if for all $a_1, \dots, a_k \in F$, we have

$$a_1x_1 + \dots + a_kx_k \in W \iff a_1 = \dots = a_k = 0.$$

A tuple $X = (x_1, \dots, x_k)$ is a **basis over W** if it is independent over W and

$$\text{Span}(\{x_1, \dots, x_k\} \cup W) = V.$$

Thus, “independent” means the same as “independent over $\{0\}$ ” and a basis is the same as a basis over $\{0\}$. Notice that a tuple X is a basis over W if and only if for every ordered basis Y for W , the concatenation $X \frown Y$ is an ordered basis for V .

Exercise 8.36. Show that every tuple $X = (x_1, \dots, x_k)$ that is independent over a subspace $W \subseteq V$ can be extended to a basis over W .

²⁶It is common to order the elements of a chain basis differently, resulting in blocks with 1s immediately *above* the main diagonal. The ordering used here seems more natural to me personally; of course, there is no substantial difference between the two conventions.

To establish Theorem 8.35, we shall prove the following stronger claim using induction on $\dim V$:

Let $m := \text{ndeg}(\varphi)$ and suppose that $X = (x_1, \dots, x_k) \in V^k$ is a tuple independent over $\ker(\varphi^{m-1})$. Then X can be extended to a chain basis for φ .

The base case $\dim V = 0$ is trivial, so assume that $\dim V \geq 1$. In this case $\text{id}_V \neq 0$, and hence $m \geq 1$. Saying that m is the nilpotency degree of φ means that $\ker(\varphi^m) = V$ and $\ker(\varphi^{m-1}) \subsetneq V$. The restriction of φ to $\ker(\varphi^{m-1})$ is a nilpotent transformation of $\ker(\varphi^{m-1})$ of nilpotency degree $m - 1$. Our plan is to apply the inductive hypothesis to this transformation.

To begin with, we use Exercise 8.36 to extend X to a basis X' over $\ker(\varphi^{m-1})$. Since any chain basis that extends X' also extends X , we will, without loss of generality, assume that $X = X'$.

Claim. *The tuple $Y := (\varphi(x_1), \dots, \varphi(x_k))$ is independent over $\ker(\varphi^{m-2})$.*

Proof. Take any $a_1, \dots, a_k \in F$ such that

$$a_1\varphi(x_1) + \dots + a_k\varphi(x_k) \in \ker(\varphi^{m-2}).$$

Since

$$a_1\varphi(x_1) + \dots + a_k\varphi(x_k) = \varphi(a_1x_1 + \dots + a_kx_k),$$

we conclude that

$$a_1x_1 + \dots + a_kx_k \in \ker(\varphi^{m-1}).$$

But the tuple X is independent over $\ker(\varphi^{m-1})$, and hence $a_1 = \dots = a_k = 0$, as desired. \dashv

Since $\ker(\varphi^m) = V$, we have $\text{im}(\varphi) \subseteq \ker(\varphi^{m-1})$. Thus, the elements of Y belong to $\ker(\varphi^{m-1})$, and we may apply the inductive hypothesis with $\ker(\varphi^{m-1})$ in place of V , $\varphi|_{\ker(\varphi^{m-1})}$ in place of φ , and Y in place of X . This shows that we can extend Y to a chain basis Z for the restriction of φ to $\ker(\varphi^{m-1})$. We now claim that $X \cup Z$ is a desired chain basis for φ .

Claim. *The tuple $X \cup Z$ is independent.*

Proof. Write $Z = (z_1, \dots, z_\ell)$ and suppose that

$$a_1x_1 + \dots + a_kx_k + b_1z_1 + \dots + b_\ell z_\ell = 0.$$

Then

$$a_1x_1 + \dots + a_kx_k = -b_1z_1 - \dots - b_\ell z_\ell \in \ker(\varphi^{m-1}).$$

Since X is independent over $\ker(\varphi^{m-1})$, we conclude that $a_1 = \dots = a_k = 0$. But then

$$b_1z_1 + \dots + b_\ell z_\ell = 0,$$

and since Z is independent, $b_1 = \dots = b_\ell = 0$ as well. \dashv

Since Z is a basis for $\ker(\varphi^{m-1})$, while X is a basis *over* $\ker(\varphi^{m-1})$, the tuple $X \cup Z$ is spanning. Hence, $X \cup Z$ is a basis. It remains to verify that it is a *chain* basis, which is left as an [exercise](#). ■

8.G. The Jordan normal form

Now it's time to put the results of this section together. Let V be a finite-dimensional vector space over an algebraically closed field F and let $\varphi \in \text{Lin}(V, V)$. Let $\lambda_1, \dots, \lambda_k$ be the distinct eigenvalues of φ and let G_1, \dots, G_k be the corresponding generalized eigenspaces. For each $1 \leq i \leq k$, set $\varphi_i := \varphi|_{G_i}$ and $\psi_i := \varphi_i - \lambda_i \text{id}$. By construction, ψ_i is a nilpotent transformation of G_i . Hence, by




Figure 7. A cartoon of the proof of Theorem 8.35.

Theorem 8.35, we can find a chain basis X_i for ψ_i . We can view X_i as an ordered basis such that the matrix $[\psi_i]_{X_i, X_i}$ is “block-diagonal,” with blocks of the form

$$\begin{bmatrix} 0 & & & & \\ 1 & 0 & & & \\ & 1 & 0 & & \\ & & \ddots & \ddots & \\ & & & 1 & 0 \end{bmatrix}.$$

Since $\varphi_i = \psi_i + \lambda_i \text{id}$, the matrix $A_i := [\varphi_i]_{X_i, X_i}$ is also “block-diagonal,” with blocks of the form

$$\begin{bmatrix} \lambda_i & & & & \\ 1 & \lambda_i & & & \\ & 1 & \lambda_i & & \\ & & \ddots & \ddots & \\ & & & 1 & \lambda_i \end{bmatrix}.$$

Let $X := X_1 \cap \dots \cap X_k$. From Theorem 8.31, we know that the matrix $A := [\varphi]_{X, X}$ is simply the “block-diagonal” matrix assembled from A_1, \dots, A_k . Combining all of this information, we conclude that the matrix A is “block-diagonal,” with blocks of the form

$$\begin{bmatrix} \lambda & & & & \\ 1 & \lambda & & & \\ & 1 & \lambda & & \\ & & \ddots & \ddots & \\ & & & 1 & \lambda \end{bmatrix}, \tag{8.37}$$

where $\lambda \in \text{Spec}(\varphi)$. Such a matrix is called a **Jordan normal form** of φ . Thus, we have established the following fundamental result:

Theorem 8.38 (Jordan). *Let V be a finite-dimensional vector space over an algebraically closed field F and let $\varphi \in \text{Lin}(V, V)$. Then φ has a Jordan normal form; i.e., there is a basis X for V such that $[\varphi]_{X, X}$ is a “block-diagonal” matrix with blocks of the form (8.37), where $\lambda \in \text{Spec}(\varphi)$. ■*

Exercise 8.39. Let V be a finite-dimensional vector space over an algebraically closed field F and let $\varphi \in \text{Lin}(V, V)$. Show that the Jordan normal form of φ is unique up to reordering of the blocks. *Hint:* Analyze our proof of Theorem 8.38.

Theorem 8.38, combined with the result of Exercise 8.39, gives a very satisfying answer to the classification problem: *Two linear transformations $\varphi, \psi \in \text{Lin}(V, V)$ of a finite-dimensional vector*

space over an algebraically closed field are conjugate if and only if their Jordan normal forms are the same up to reordering of the blocks.

Example 8.40. Let F be an algebraically closed field and suppose that V is a 2-dimensional F -vector space. What can the Jordan normal form of a transformation $\varphi \in \text{Lin}(V, V)$ look like? It must be “block-diagonal,” so there are two possibilities:

- one block of size 2; or
- two blocks of size 1.

In the first case, the matrix has the form

$$\begin{bmatrix} \lambda & 0 \\ 1 & \lambda \end{bmatrix},$$

where λ is the unique eigenvalue of φ . In the second case, the matrix is

$$\begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix},$$

and $\text{Spec}(\varphi) = \{\lambda_1, \lambda_2\}$ (it is still possible that $\lambda_1 = \lambda_2$). When $\dim V = 3$, there are three options:

$$\begin{bmatrix} \lambda & 0 & 0 \\ 1 & \lambda & 0 \\ 0 & 1 & \lambda \end{bmatrix}, \quad \begin{bmatrix} \lambda_1 & 0 & 0 \\ 1 & \lambda_1 & 0 \\ 0 & 0 & \lambda_2 \end{bmatrix}, \quad \text{and} \quad \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix}.$$

Extra exercises for Section 8

Exercise 8.41. Let F be a field and let V be an F -vector space. Let $\varphi, \psi: V \rightarrow V$ be linear functions. Suppose that φ and ψ commute, i.e., that $\varphi \circ \psi = \psi \circ \varphi$. Let $\lambda \in \text{Spec}(\varphi)$ and let $W \subseteq V$ be the corresponding eigenspace.

- (a) Show that the space W is ψ -invariant.
- (b) Conclude that if the field F is algebraically closed and the space V is finite-dimensional, then W contains an eigenvector of ψ ; in particular, φ and ψ have a common eigenvector.

Exercise 8.42. In this exercise we prove the following result of Sylvester:

Theorem 8.43 (Sylvester). *Let F be an algebraically closed field and let $A \in M_{n \times n}(F)$, $B \in M_{m \times m}(F)$. If A and B have no common eigenvalues, then for each $C \in M_{n \times m}(F)$, the equation*

$$AX - XB = C$$

has a unique solution $X \in M_{n \times m}(F)$.

Let $V := M_{n \times m}(F)$ and define linear functions $\varphi_A, \varphi_B: V \rightarrow V$ by

$$\varphi_A(X) := AX \quad \text{and} \quad \varphi_B(X) := XB.$$

- (a) Show that the functions φ_A and φ_B commute with each other.
- (b) Let $\lambda \in \text{Spec}(\varphi_A - \varphi_B)$. Show that $\lambda = \mu - \nu$ for some $\mu \in \text{Spec}(\varphi_A)$ and $\nu \in \text{Spec}(\varphi_B)$.
Hint: Apply the result of Exercise 8.41 with $\varphi_A - \varphi_B$ in place of φ and φ_A in place of ψ .
- (c) Show that $\text{Spec}(\varphi_A) = \text{Spec}(A)$ and $\text{Spec}(\varphi_B) = \text{Spec}(B)$. *Hint:* To show that $\text{Spec}(\varphi_B) = \text{Spec}(B)$, use the fact that $\text{Spec}(B) = \text{Spec}(B^\top)$ (see Exercise 8.14).
- (d) Conclude that all eigenvalues of $\varphi_A - \varphi_B$ are nonzero and finish the proof of Theorem 8.43.

Exercise 8.44. Let F be an algebraically closed field and let $A \in M_{n \times n}(F)$. Use Theorem 8.38 to show that the matrices A and A^\top are conjugate.

REFERENCES

- [BL88] G. BIRKHOFF and S. MAC LANE. *Algebra*. New York: Chelsea Pub. Co., 1988
- [Gol12] J.S. GOLAN. *The Linear Algebra a Beginning Graduate Student Ought to Know*. Dordrecht: Springer, 2012
- [Loe14] N. LOEHR. *Advanced Linear Algebra*. Boca Raton: CRC Press, Taylor & Francis Group, 2014
- [Mat10] J. MATOUŠEK. *Thirty-Three Miniatures: Mathematical and Algorithmic Applications of Linear Algebra*. Providence: American Mathematical Society, 2010
- [OCV11] K. O'MEARA, J. CLARK, and C. VINSONHALER. *Advanced Topics in Linear Algebra: Weaving Matrix Problems through the Weyr Form*. Cary: Oxford University Press, 2011