

MATH0005 Algebra 1

Matthew Towers

July 9, 2024

Contents

Preface	5
1 Logic	7
1.1 Propositional calculus	7
1.2 Well-formed formulas	8
1.3 Truth tables	9
1.4 Truth values for WFFs	11
1.5 Logical equivalence	14
1.6 Useful logical equivalences	16
1.7 The contrapositive	18
1.8 Adequacy	19
1.9 First order logic	20
2 Sets and functions	27
2.1 Introduction to set theory	27
2.2 Set operators	29
2.3 Set algebra	32
2.4 De Morgan's laws	33
2.5 Functions	35
2.6 Function composition	36
2.7 Function properties	36
2.8 Invertibility	39
2.9 Conditions for invertibility	40
2.10 Permutations	42
2.11 Inverses and composition	44
2.12 Cycles	45
2.13 Products of disjoint cycles	48
2.14 Powers and orders	50
2.15 Transpositions	51
2.16 Sign	55
3 Matrices	59
3.1 Matrix definitions	59
3.2 Matrix multiplication	61
3.3 Transpose	68
3.4 Multiplication properties	69
3.5 Invertible matrices	71
3.6 Systems of linear equations	72

3.7	Row operations	75
3.8	Elementary matrices	76
3.9	Row reduced echelon form	77
3.10	RREF existence and uniqueness	79
3.11	Solving RREF systems	81
3.12	Invertibility and RREF	84
3.13	Finding inverses	85
4	Linear algebra	89
4.1	Fields	89
4.2	Vector spaces	93
4.3	Using the vector space axioms	94
4.4	Subspaces	95
4.5	Sums and intersections	97
4.6	Linear independence	98
4.7	Spanning sequences	99
4.8	Bases	102
4.9	Dimension	104
4.10	Basis and dimension examples	106
4.11	Extending to a basis	106
4.12	Finding dimensions	108
4.13	Linear maps	110
4.14	Kernel and image	112
4.15	The rank-nullity theorem	115
4.16	Matrix nullspace basis	116
4.17	Column space basis	118
4.18	Matrix of a linear map	120
4.19	Matrix of a composition	122
4.20	Change of basis	123
4.21	Coordinate isomorphisms	126

About these notes

These are the lecture notes for the 2024 edition of the first year UCL module MATH0005 Algebra 1. If you are reading the web version of these notes and you want a pdf copy, you can find one at [this link](#).

In 2021 this course was taught fully online. The lecture videos are available as a YouTube playlist on YouTube, but be warned that the treatment of some topics (especially first order logic) is slightly different this year.

I hope you enjoy the module. If you have questions or comments on the material, or if you find errors in the text, please email me at m.towers@ucl.ac.uk

You will find suggestions for further reading at the end of each chapter, and at the top of the Moodle page for MATH0005.

Chapter 1

Logic

This part of MATH0005 is about *logic*: the study of how we can reason and make deductions, of methods of argument, and of methods of proof.

1.1 Propositional calculus

We begin with propositional calculus, the study of propositions. A **proposition** is a mathematical statement which is either true or false.¹

Here are some example propositions.

- 34043 is the sum of two square numbers.
- The function $f(x) = \sin(x)$ is continuous.
- The square root of 2 is not a rational number.
- 11111111111111111111 is a prime number.
- $1 + 1 = 3$.
- $1 + 1 = 2$.
- The Riemann hypothesis is false.
- 25 is a square and 26 is a square.

Some of these are true and some are false, but each has a well-defined truth value, even if we don't know what it is. On the other hand, something like “ n is even” is not a proposition, because it doesn't have a truth value until we know what n is.

1.1.1 Logical connectives

Connectives combine simpler logical statements into more more complex ones. We use them to build complex propositions out of simpler ones. The standard connectives are ‘and’, ‘or’, ‘not’, ‘implies’, ‘if and only if’ (iff, for short).

Here are some examples of propositions which contain connectives.

¹Sometimes people use the word *proposition* for something that's a bit like a theorem but not quite as important. That's not what we're talking about here.

- and: “34043 is a sum of two squares and 34043 is divisible by 17”
- or: “34043 is a sum of two squares or 34043 is divisible by 17”
- not: “it is not true that 34043 is a sum of two squares”
- implies: “34043 is odd implies 34043 is divisible by 3”
- if and only if: “an odd prime number is a sum of two squares if and only if it leaves remainder 1 when you divide it by 4”

Implies is often expressed as “if... then”. The sentence “if 34043 is odd then 34043 is divisible by 3” means the same thing as “34043 is odd implies 34043 is divisible by 3.”

You might wonder if there are any more interesting “exotic connectives” that would allow us to create new statements not expressible using the connectives above. There are other connectives — common examples are exclusive or (written XOR), NAND (sometimes called Sheffer stroke), and NOR — but it’s a theorem that any connective you invent can be expressed in an equivalent way using just the connectives above (in fact, you don’t even need all of them).

1.2 Well-formed formulas

We’re now going to develop a formal language for expressing logical propositions and how they are combined using connectives.

1.2.1 Variables and connective symbols

Because we want to talk abstractly about how to reason, we don’t want to confine ourselves to particular propositions but to explore what can be said about all propositions. For that reason we introduce **propositional variables**: symbols that represent a proposition. Traditionally lower case English letters p , q , r , ... are used for propositional variables, or letters with subscripts p_1, p_2, \dots

In addition to propositional variables, the language we use will have symbols for some of the logical connectives we discussed before.

- \wedge represents *and*.
- \vee represents *or*.
- \rightarrow or \implies represents *implies*.
- \neg represents *not*.

Finally, we will also use brackets: (and).

We’ve now got the “letters” of our language: propositional variables, connective symbols, and brackets. Just like the letters a, b, c...z can be used to make English sentences, we can now build what we will call formulas, like $(p \vee q)$, or $(p \implies (q \wedge (\neg r)))$. But just like *eifaefeaioj* is a legitimate string of letters that isn’t a meaningful word, $\wedge \implies pq\neg$ doesn’t seem like something we can give a useful logical interpretation to. Collections of propositional variables, connectives, and brackets to which we can give a sensible meaning will be called well-formed formulas, and we are going to see next what the rules are for a formula to be well-formed.

1.2.2 Definition of a well-formed formula

We need rules to say which strings of connectives, brackets, and variables are **well-formed formulas**, or **WFFs** for short. We do this by specifying rules for constructing WFFs. By definition, something is a WFF if and only if it can be constructed using these rules.

1. A propositional variable is a WFF.
2. If ϕ and ψ are any two WFFs then
 - 2.1 $(\phi \wedge \psi)$ is a WFF,
 - 2.2 $(\phi \vee \psi)$ is a WFF,
 - 2.3 $(\phi \implies \psi)$ is a WFF, and
 - 2.4 $\neg\phi$ is a WFF.

1.2.3 WFF examples

Suppose p and q are propositional variables. Then the following are WFFs:

- p is a WFF because of rule 1.
- $(p \implies q)$ is a WFF by using rule 1 twice then rule 2.3.
- $\neg r$ by using rule 1 then rule 2.4.
- $((p \implies q) \vee \neg r)$ is a WFF as rule 1 says p, q, r are WFFs, rule 2.3 and rule 2.4 say that $(p \implies q)$ and $\neg r$ are WFFs, and finally rule 2.2 says the whole thing is a WFF.
- $\neg\neg(p \implies q)$ by rule 1, then rule 2.3, then rule 2.4 twice.

Only things that can be built using the rules are WFFs. You can't build

$$r \wedge \implies pq) \neg$$

using the rules above (can you prove it?), so it's not a WFF. You can't even build $p \vee q$ or $(p \wedge q \wedge r)$, so these aren't WFFs either.

1.3 Truth tables

We've seen what a WFF is. It's important to remember that a WFF like $(p \wedge q)$ isn't true or false on its own: that will depend on the truth or falsity of the statements represented by the propositional variables p and q . The aim of the next couple of sections is to see how, once we decide whether the propositional variables in a WFF are true or false, we can give a truth value to the whole WFF.

The way we do this is by making a *truth-table* definition for each connective of how the truth value of a WFF using that connective depends on the truth values of the WFFs it connects. We do this in such a way that the connective behaves like the informal logical idea it is supposed to represent: for example, \wedge is supposed to represent *and* so we will define $(\phi \wedge \psi)$ to be true if and only

if ϕ **and** ψ are both true. Once we've done this for every connective, we can determine the truth value of any WFF by looking at the simplest formulas contained in it, determining their truth values using our tables, and working our way upwards until we have the truth value of the whole formula.

1.3.1 Truth assignments for propositional variables

Let's start with giving truth values to propositional variables. Here and elsewhere T means true and F means false.

Definition 1.3.1. A **truth assignment** for a set V of propositional variables is a function $v : V \rightarrow \{T, F\}$.

(A better name for this concept would be 'truth-value assignment' since a truth assignment can make variables false as well as true, but this is the conventional name.)

Example 1.3.1. If p and q are propositional variables and $V = \{p, q\}$ then there is a truth assignment v for V such that $v(p) = T$ and $v(q) = F$.

This is one of the four different truth assignments for a set of two propositional variables. In general, if you have n propositional variables then there are 2^n different truth assignments for those variables, since each variable must be given one of two different truth values.

1.3.2 Extending a truth assignment to WFFs

Given a truth assignment for some propositional variables, we would like to extend it to get a truth value for all the WFFs using those variables in a way that takes into account the intended meaning of the logical connectives. This is a difficult problem for complex WFFs. For example, if you have a truth assignment which makes p and r true and q false, what should the truth value of the following WFF be?

$$((p \implies (q \vee r)) \implies (\neg p \vee q))$$

In order to approach the problem of extending a truth assignment so that it gives a sensible truth value to any WFF, suppose that we somehow already knew what truth values we were going to assign to the WFFs ϕ and ψ . What truth value should we give to the WFF $(\phi \wedge \psi)$? We are free to choose this of course, but since \wedge is supposed to represent the ordinary usage of the word "and" it would be sensible to assign $(\phi \wedge \psi)$ the value true if both ϕ and ψ were assigned true, and false otherwise.

This idea is summed up in the following **truth table** for \wedge :

ϕ	ψ	$(\phi \wedge \psi)$
T	T	T
T	F	F
F	T	F
F	F	F

Table 1.1: Truth table for \wedge

The meaning of the table is that given a truth assignment $v : V \rightarrow \{T, F\}$, our method of assigning a truth value to a WFF $(\phi \wedge \psi)$ using the variables V will be as follows. Row 1 means that if $v(\phi) = T$ and $v(\psi) = T$ then $v((\phi \wedge \psi))$ will be T . Row 2 means that if $v(\phi) = T$ and $v(\psi) = F$ then $v((\phi \wedge \psi))$ will be F , and so on.

Another way to think about this truth table is to use it to define \wedge as a way to combine two truth values into another truth value, just like $+$ combines two numbers into another number. We let $T \wedge T = T$, $T \wedge F = F$, $F \wedge T = F$, and $F \wedge F = F$. The advantage of this is that it lets us rewrite the last paragraph in a single sentence: we will define $v((\phi \wedge \psi))$ to be $v(\phi) \wedge v(\psi)$.

Here are the truth tables for the other connectives in our language.

ϕ	$\neg\phi$
T	F
F	T

Table 1.2: Truth table for \neg

ϕ	ψ	$(\phi \vee \psi)$
T	T	T
T	F	T
F	T	T
F	F	F

Table 1.3: Truth table for \vee

ϕ	ψ	$(\phi \implies \psi)$
T	T	T
T	F	F
F	T	T
F	F	T

Table 1.4: Truth table for \implies

Similarly to what we did for \wedge , we regard all of our connectives not just as symbols to be used in WFFs but as ways of combining truth values. For example, we define $\neg T = F$, $T \vee F = T$, and $F \implies T = T$.

People often find the truth table for implies confusing, especially the final two rows where ϕ is false. These last two rows tell us that $(\phi \implies \psi)$ is true whenever ϕ is false, regardless of the truth value given to ψ . If you'd like to read more about why this truth table is a sensible way to define truth values for statements containing implies, this short piece of writing by (Fields medallist) Tim Gowers, or this longer version is good.

1.4 Truth values for WFFs

Suppose we have a truth assignment $v : V \rightarrow \{T, F\}$. There is then a unique way to extend v so that it gives a truth value to any WFF using the propositional

variables V such that for any WFFs ϕ and ψ ,

$$\begin{aligned} v((\phi \wedge \psi)) &= v(\phi) \wedge v(\psi), \\ v((\phi \vee \psi)) &= v(\phi) \vee v(\psi), \\ v((\phi \implies \psi)) &= v(\phi) \implies v(\psi), \quad \text{and} \\ v(\neg\phi) &= \neg v(\phi). \end{aligned}$$

Recall that we use the connective symbols not just as parts of WFFs but as ways of combining truth values, for example $T \wedge F = F$, $T \implies T = T$, and $\neg F = T$. For example, if $V = \{p, q\}$ and $v(p) = T, v(q) = F$ we would have

$$\begin{aligned} v((p \wedge q)) &= v(p) \wedge v(q) \\ &= F \wedge T \\ &= F \end{aligned}$$

and

$$\begin{aligned} v((\neg p) \implies (p \vee q)) &= v(\neg p) \implies v(p \vee q) \\ &= (\neg v(p)) \implies ((v(p) \vee v(q))) \\ &= F \implies (T \vee F) \\ &= F \implies T \\ &= T. \end{aligned}$$

It's not completely obvious this really works, but you can read a proof in section 2.3 of the book by Goldrei mentioned in the further reading section at the end of this chapter.

This method of assigning truth values to WFFs can be thought of in a slightly different way: we just substitute in truth values in place of the propositional variables, and combine them using the truth tables for the connectives — exactly like how if you wanted to find the value of $x^2 + y + 3$ when $x = 1$ and $y = 2$, you would substitute the values in to get $1^2 + 2 + 3$ and combine them using the usual arithmetic operations to get 6.

Example 1.4.1. Let

$$\phi = ((p \wedge q) \vee (\neg p \wedge \neg q)).$$

Let $v(p) = T, v(q) = F$. We are going to find $v(\phi)$.

The method of assigning truth values to WFFs above tells us

$$\begin{aligned} v(\phi) &= v((p \wedge q) \vee (\neg p \wedge \neg q)) \\ &= v(p \wedge q) \vee v(\neg p \wedge \neg q) \end{aligned} \tag{1.1}$$

so we need to work out $v(p \wedge q)$ and $v(\neg p \wedge \neg q)$.

We have $v(p \wedge q) = v(p) \wedge v(q) = T \wedge F$. Looking at the T, F row of Table 1.1, the truth table for \wedge , we see that $T \wedge F$ is F .

Next,

$$\begin{aligned} v(\neg p \wedge \neg q) &= v(\neg p) \wedge v(\neg q) \\ &= \neg v(p) \wedge \neg v(q) \\ &= \neg T \wedge \neg F \\ &= F \wedge T. \end{aligned}$$

The F, T row of the same truth table tells us that this is F .

Finally, substituting the values we have just worked out for $v(p \wedge q)$ and $v(\neg p \wedge \neg q)$ into (1.1)

$$v(\phi) = F \vee F.$$

Looking at the F, F row of Table 1.3, the truth table for \vee , we see that $v(\phi) = F$.

Example 1.4.2. Consider the WFF $\phi = (p \implies (p \implies p))$ and the truth assignment $v(p) = T$. What is $v(\phi)$?

By definition,

$$\begin{aligned} v(\phi) &= v(p) \implies (v(p) \implies v(p)) \\ &= T \implies (T \implies T). \end{aligned}$$

Looking at the T, T row of the truth table for implies, Table 1.4, we see that $T \implies T$ is T . So $v(\phi) = T \implies T$. For the same reason, $v(\phi) = T$.

If you work out the truth value of ϕ when $v(p) = F$, you should find that the result is also T .

p	$(p \implies (p \implies p))$
T	T
F	T

Notice that the WFF ϕ from the previous example is true for *every* truth assignment of its variables. A WFF with this property is called a **tautology**, and a WFF which is false under every truth assignment, for example $(p \wedge \neg p)$, is a **contradiction**.

Example 1.4.3. Let

$$\phi = ((p \vee q) \wedge (p \vee \neg q)).$$

Given the truth assignment v such that $v(p) = T, v(q) = F$, let's work out the truth value $v(\phi)$. Since

$$v(\phi) = v(p \vee q) \wedge v(p \vee \neg q)$$

we can start by working out $v(p \vee q)$ and $v(p \vee \neg q)$ separately. From the truth table for \vee , Table 1.1, we see that

$$v(p \vee q) = T \vee F = T.$$

We have $v(\neg q) = \neg v(q) = \neg F = T$, so

$$v(p \vee \neg q) = T \vee T = T.$$

Finally $v(\phi) = T \wedge T = T$.

The **truth table for a WFF** lists its truth value under all possible truth assignments for its propositional variables. The truth table for the formula ϕ from the previous example (and for some of the formulas that make up ϕ) is given below. You should check that the following table is correct.

p	q	$(p \vee q)$	$\neg q$	$(p \vee \neg q)$	ϕ
T	T	T	F	T	T
T	F	T	T	T	T
F	T	T	F	F	F
F	F	F	T	T	F

1.5 Logical equivalence

To motivate the idea of logical equivalence, consider the two WFFs

$$\begin{aligned}\phi &= (p \wedge q) \\ \psi &= (q \wedge p).\end{aligned}$$

These are *different* WFFs because a WFF is purely a sequence of symbols and these are two different sequences of symbols. However, given any truth assignment, no matter what it is, ϕ and ψ always get equal truth values. You can see this by looking at the truth table for \wedge , Table 1.1 which is symmetrical in p and q , in the sense that if you swap the truth values for p and q , the truth value of $(p \wedge q)$ stays the same.

Definition 1.5.1. Two WFFs ϕ and ψ are called **logically equivalent**, and we write $\phi \equiv \psi$, if and only if they have the same truth value under every possible truth assignment.

Since the truth table for a WFF displays its truth values under every possible truth assignment, two WFFs are logically equivalent if and only if they have the same truth table.

When two WFFs are logically equivalent they may look different but they always have the same truth value, no matter what the truth values of their variables. This concept is useful in practise because if you want to prove something is true, you can prove some logically equivalent formula instead.

Theorem 1.5.1. *Let ϕ , ψ , and θ be WFFs. Then*

1. $(\phi \wedge \psi) \equiv (\psi \wedge \phi)$,
2. $(\phi \vee \psi) \equiv (\psi \vee \phi)$,
3. $(\phi \wedge (\psi \wedge \theta)) \equiv ((\phi \wedge \psi) \wedge \theta)$, and
4. $(\phi \vee (\psi \vee \theta)) \equiv ((\phi \vee \psi) \vee \theta)$.

The first two parts of this theorem are referred to as the commutativity properties for \wedge and \vee , and the second two parts as the associativity properties.

Proof. Parts 1 and 2 are very easy to check as they follow straight from the truth tables for \wedge , Table 1.1 and \vee , Table 1.3.

Parts 3 and 4 are tedious to check, but very easy. We will work out the truth values for one truth assignment and leave the others to you. Let v be a truth assignment such that $v(\phi) = v(\psi) = T$ and $v(\theta) = F$. For the left hand side of part 3 we have

$$\begin{aligned}v(\phi \wedge (\psi \wedge \theta)) &= v(\phi) \wedge v(\psi \wedge \theta) \\ &= T \wedge (v(\psi) \wedge v(\theta)) \\ &= T \wedge (T \wedge F) \\ &= T \wedge F \\ &= F\end{aligned}$$

and for the right hand side

$$\begin{aligned}
 v((\phi \wedge \psi) \wedge \theta) &= v(\phi \wedge \psi) \wedge v(\theta) \\
 &= (v(\phi) \wedge v(\psi)) \wedge F \\
 &= (T \wedge T) \wedge F \\
 &= T \wedge F \\
 &= F.
 \end{aligned}$$

Continuing like this you can show that the truth tables for both $(\phi \wedge (\psi \wedge \theta))$ and $((\phi \wedge \psi) \wedge \theta)$ are as follows.

ϕ	ψ	θ	
T	T	T	T
T	T	F	F
T	F	T	F
T	F	F	F
F	T	T	F
F	T	F	F
F	F	T	F
F	F	F	F

Part 4 can be done similarly. □

What the associativity laws, parts 3 and 4, do, is to allow us to drop *some* brackets while remaining logically unambiguous. Something like $p \wedge q \wedge r$ isn't a WFF — because it has \wedge symbols but no brackets — but part 3 guarantees us that any two ways we choose to bracket it give logically equivalent WFFs. Similarly

$$\begin{aligned}
 p_1 \wedge p_2 \wedge \cdots p_n \\
 p_1 \vee p_2 \vee \cdots p_n
 \end{aligned}$$

may not be WFFs, but any bracketings that do turn them into WFFs give logically equivalent formulas. For this reason, we often omit bracketings when they don't cause ambiguity, even though when we miss out the brackets we don't strictly speaking have a WFF.

Example 1.5.1. Sometimes brackets *are* essential. The WFFs

$$\begin{aligned}
 \phi &= (p \wedge (q \vee r)) \\
 \psi &= ((p \wedge q) \vee r)
 \end{aligned}$$

are *not* logically equivalent. Before you look at the truth tables below you should prove this by finding a truth assignment for the variables p, q, r which makes one of these WFFs true and the other false.

Here are the truth tables:

p	q	r	$(p \wedge (q \vee r))$	$((p \wedge q) \vee r)$
T	T	T	T	T
T	T	F	T	T
T	F	T	T	T
T	F	F	F	F
F	T	T	F	T
F	T	F	F	F
F	F	T	F	T
F	F	F	F	F

so they differ under the truth assignment making p false, q true, and r true, and also under the truth assignment making p false, q false, and r true.

1.6 Useful logical equivalences

1.6.1 Distributivity

The property that for all numbers a, b, c we have $a \times (b + c) = a \times b + a \times c$ is called *distributivity* of \times over $+$. Similar rules hold for \wedge and \vee .

Theorem 1.6.1. *Let ϕ , ψ , and θ be WFFs. Then*

- $(\phi \wedge (\psi \vee \theta)) \equiv ((\phi \wedge \psi) \vee (\phi \wedge \theta))$, and
- $(\phi \vee (\psi \wedge \theta)) \equiv ((\phi \vee \psi) \wedge (\phi \vee \theta))$.

Proof. Here are the truth tables for the four WFFs:

ϕ	ψ	θ	$(\phi \wedge (\psi \vee \theta))$	$((\phi \wedge \psi) \vee (\phi \wedge \theta))$
T	T	T	T	T
T	T	F	T	T
T	F	T	T	T
T	F	F	F	F
F	T	T	F	F
F	T	F	F	F
F	F	T	F	F
F	F	F	F	F

ϕ	ψ	θ	$(\phi \vee (\psi \wedge \theta))$	$((\phi \vee \psi) \wedge (\phi \vee \theta))$
T	T	T	T	T
T	T	F	T	T
T	F	T	T	T
T	F	F	T	T
F	T	T	T	T
F	T	F	F	F
F	F	T	F	F
F	F	F	F	F

The last two columns are the same in both tables, so the formulas are logically equivalent. \square

1.6.2 Double negation

Theorem 1.6.2. *Let ϕ be a WFF. Then $\neg\neg\phi \equiv \phi$.*

Proof. Let v be a truth assignment for the propositional variables involved in ϕ . If $v(\phi) = T$ then $v(\neg\phi) = \neg v(\phi) = F$ and so $v(\neg\neg\phi) = \neg v(\neg\phi) = \neg F = T$. Similarly if $v(\phi)$ is false so is $v(\neg\neg\phi)$. Therefore under any truth assignment v we have $v(\phi) = v(\neg\neg\phi)$. \square

1.6.3 De Morgan's laws

Theorem 1.6.3. *Let ϕ and ψ be WFFs. Then*

1. $\neg(\phi \vee \psi) \equiv (\neg\phi \wedge \neg\psi)$, and
2. $\neg(\phi \wedge \psi) \equiv (\neg\phi \vee \neg\psi)$.

You might find it clearer to write the right hand sides of these equivalences as $(\neg\phi) \wedge (\neg\psi)$ and $(\neg\phi) \vee (\neg\psi)$, even though these are not well-formed formulas. From now on we will add or remove brackets from formulas where it helps to make them clearer or more readable even if it means that they are not strictly WFFs.

Proof. Again, proving this is simply a matter of checking the possibilities for the truth values of ϕ and ψ under any assignment. In a table:

ϕ	ψ	$\neg(\phi \vee \psi)$	$(\neg\phi \wedge \neg\psi)$
T	T	F	F
T	F	F	F
F	T	F	F
F	F	T	T

The final columns are the same, so the two formulas have the same truth value no matter what truth assignment is used and are therefore logically equivalent. \square

De Morgan's laws can be generalized to more than two WFFs.

Theorem 1.6.4. *For any n and any WFFs ϕ_1, \dots, ϕ_n we have*

1. $\neg(\phi_1 \wedge \dots \wedge \phi_n) \equiv \neg\phi_1 \vee \dots \vee \neg\phi_n$, and
2. $\neg(\phi_1 \vee \dots \vee \phi_n) \equiv \neg\phi_1 \wedge \dots \wedge \neg\phi_n$.

While $\phi_1 \wedge \phi_2 \wedge \phi_3$, for example, isn't a WFF, every way of adding brackets to make it into one produces a logically equivalent WFF because of the associativity of \wedge , Theorem 1.5.1 part 3. Therefore it's OK for us to omit brackets here for the sake of making the formula easier to read.

1.7 The contrapositive

The following logical equivalence shows us that every WFF that uses \implies can be written with \neg and \vee instead.

Theorem 1.7.1. *Let ϕ and ψ be WFFs. Then*

$$(\phi \implies \psi) \equiv (\neg\phi \vee \psi).$$

Here is the truth table that proves this result:

ϕ	ψ	$(\phi \implies \psi)$	$(\neg\phi \vee \psi)$
T	T	T	T
T	F	F	F
F	T	T	T
F	F	T	T

This equivalence is commonly used when proving a statement like “ a implies b .” Proofs of statements in this form are often carried out by assuming that a is true and then deducing that b is also true. Why is that sufficient to prove ($a \implies b$)?

Suppose that if a is true, so is b . If a is false then $\neg a$ is true, so $(\neg a \vee b)$ is true *no matter what the statements a and b were*. On the other hand if a is true we know b is true as well, so $(\neg a \vee b)$ is true in that case too. So regardless of the truth value of a , the formula $(\neg a \vee b)$ is true. Because this is logically equivalent to $(a \implies b)$, we’re done.

1.7.1 The contrapositive

The **contrapositive** of an implication ($a \implies b$) is by definition $(\neg b \implies \neg a)$. For example, the contrapositive of “if it’s Monday, then it’s raining” is “if it’s not raining, then it’s not Monday.” We are going to use the logical equivalence of the previous section to show that an implication is logically equivalent to its contrapositive.

Theorem 1.7.2. *Let ϕ and ψ be WFFs. Then*

$$(\phi \implies \psi) \equiv (\neg\psi \implies \neg\phi).$$

Proof. You can check the truth tables for these two statements, or you can do this:

$$\begin{aligned} (\phi \implies \psi) &\equiv (\neg\phi \vee \psi) && \text{Theorem 1.7.1} \\ &\equiv (\neg\phi \vee \neg\neg\psi) && \text{Theorem 1.6.2} \\ &\equiv (\neg\neg\psi \vee \neg\phi) && \text{Theorem 1.5.1} \\ &\equiv (\neg\psi \implies \neg\phi) && \text{Theorem 1.7.1} \end{aligned}$$

□

Again this is very useful as a proof technique. If you want to prove ($a \implies b$), it is logically equivalent to prove the contrapositive $(\neg b \implies \neg a)$, and this is sometimes easier. An example is

x^2 is an irrational number implies x is an irrational number.

This statement is true, but the contrapositive “ x is rational implies x^2 is rational” is easier to prove because x being rational actually tells you something specific (that $x = p/q$ for some whole numbers p and q) which you can use to make the proof work. There are further examples given in this blog post by Timothy Gowers.

1.7.2 The converse

Don’t confuse the contrapositive of an implies statement with its **converse**. The converse of $(\phi \implies \psi)$ is defined to be $(\psi \implies \phi)$, and these two are **not** in general logically equivalent. (You should think of a truth assignment to show that $(p \implies q)$ and $(q \implies p)$ are not logically equivalent.)

1.8 Adequacy

One of the logical equivalences we proved earlier 1.7.1 was

$$p \implies q \equiv (\neg p) \vee q$$

which you could interpret as saying that we don’t really *need* the \implies connective, in the sense that if you are given any WFF using \vee , \wedge , \implies , and \neg you can convert it into a logically equivalent one that does not use \implies by replacing every occurrence of $\phi \implies \psi$ with $(\neg\phi) \vee \psi$.

Definition 1.8.1. A set of connectives is **adequate** if every WFF is logically equivalent to one using only the connectives from that set.

The argument above shows that the set $\{\wedge, \vee, \neg\}$ is adequate, but there are even smaller adequate sets.

Theorem 1.8.1. $\{\vee, \neg\}$ is adequate.

Proof. Every WFF is equivalent to one using only \wedge , \vee , and \neg . By the second of De Morgan’s laws, Theorem 1.6.3 part 2,

$$\neg(\phi \wedge \psi) \equiv (\neg\phi) \vee (\neg\psi)$$

so by double negation, Theorem 1.6.2,

$$\phi \wedge \psi \equiv \neg((\neg\phi) \vee (\neg\psi)). \tag{1.2}$$

This means every occurrence of \wedge in a formula can be replaced with the logically equivalent formula on the right hand side of (1.2) which only uses \vee and \neg . We’ve shown every WFF is equivalent to one only using \vee and \neg . \square

Example 1.8.1. Consider the formula ϕ given by

$$p \implies (q \wedge r).$$

Because $\{\vee, \neg\}$ is adequate there must exist a formula logically equivalent to ϕ using only \neg and \vee . Let’s find one.

$$\begin{aligned} p \implies (q \wedge r) &\equiv (\neg p) \vee (q \wedge r) && \text{Theorem 1.7.1} \\ &\equiv (\neg p) \vee \neg((\neg q) \vee (\neg r)) && (1.2) \end{aligned}$$

Theorem 1.8.2. $\{\wedge, \neg\}$ is adequate.

Proof. We already know that every WFF is logically equivalent to one only using \neg , \wedge , and \vee . By the first of De Morgan's laws, Theorem 1.6.3 part 1,

$$\neg(\phi \vee \psi) \equiv (\neg\phi) \wedge (\neg\psi)$$

and so using double negation (Theorem 1.6.2)

$$\phi \vee \psi \equiv \neg((\neg\phi) \wedge (\neg\psi)) \tag{1.3}$$

which means we can replace every occurrence of $\phi \vee \psi$ in a WFF with the right hand side of (1.3), which only involves \neg and \wedge . \square

1.8.1 Which sets of connectives are not adequate?

It's clear that we can't go any further: it isn't true that every WFF is equivalent to one using \vee only (any such formula is true when all its variables are true, so we can't find one equivalent to $\neg p$) or using \wedge only (same argument) or \implies only (same argument).

There *are* single connectives which are adequate on their own. For example, if we define $p \uparrow q$ to have the same truth table as $\neg(p \wedge q)$ (the *Sheffer stroke* or NAND), and $p \downarrow q$ (the *Pierce arrow* or NOR) to have the truth table of $\neg(p \vee q)$, it can be shown that both $\{\uparrow\}$ and $\{\downarrow\}$ are adequate.

ϕ	ψ	$(\phi \downarrow \psi)$
T	T	F
T	F	F
F	T	F
F	F	T

Table 1.5: Truth table for \downarrow

1.8.2 Why should you care about adequacy?

Firstly it can be useful for proving theorems to be able to find logical equivalents to a WFF in simple standard forms, e.g. disjunctive normal form and its and-analogue conjunctive normal form. Second, *logic gates* (electronic devices implementing logical connectives) are a fundamental part of digital circuit design. A computer chip is largely made up of many logic gates connected together. In the early days of digital electronics using only one type of logic gate helped make the design much easier. The Apollo guidance computer, used on the first ever moon landing, was built using only NOR gates (about 6000 of them).

1.9 First order logic

The WFFs we have studied so far only capture logical statements of a very simple form. Very commonly we want to work with more complex statements, especially those that depend on some kind of parameter or variable. Here are some examples.

Example 1.9.1. • There exists a rational number x with $x^2 = 2$.

- For every natural number² n there exists a prime number p with $p > n$.
- For all real numbers m there exists a real number n such that for all real numbers x greater than n it holds that e^x is greater than m .

These kinds of statement are especially common in analysis, but they arise everywhere in mathematics. Propositional calculus doesn't have a way of talking about statements that depend on a variable, and might be true for some values, or all values, or no values that variable could take. It also has no way to talk about functions or relations. The logical theory we're going to learn about that *can* deal with statements like this is called **first order logic** or **predicate calculus**.

In propositional calculus we had WFFs. The corresponding thing in first order logic is called a **first order formula**. We will treat these very informally — if you want to learn about these in detail, read chapter 4 of the book by Goldrei mentioned in the further reading section at the end of this chapter, or take MATH0037 in year 3.

Here is a simple example of a first order formula:

$$\forall x \in X \exists y \in Y : R(x, y)$$

The meaning of this is “for all x in the set X , there exists a y in the set Y such that the property $R(x, y)$ is true for x and y .” The new things here are the symbol \in which means “in,” the *quantifiers* \forall and \exists , and the *predicate* $R(x, y)$. In the next two subsections we'll explain what these last two are.

1.9.1 Predicates

A **predicate** $P(x)$ on a set X is a statement that becomes true or false when an element of X is substituted for the variable x . For example, “ x is even” is a predicate on the integers, since for every integer n the statement “ n is even” is either true or false.

We need predicates like $P(x, y)$ that require more than one variable. For example, $P(x, y)$ could be the statement “ $|x - y| < 1$ ” which is a true or false statement about any two real numbers x and y .

1.9.2 Quantifiers

Lots of statements in mathematics involve the idea of a predicate being true for some, or for all, values of a particular variable. For example:

- There exists some rational number q such that $q^2 = 2$.
- For all real numbers x , we have $x^2 \geq 0$.
- For all natural numbers n , there exists a prime number p such that $p > n$.

²A natural number is a non-negative whole number.

To write these formally we use quantifiers. There are two types of quantifier: \forall , the **universal quantifier**, and \exists , the **existential quantifier**.

The universal quantifier \forall is used to say that something is true for every element of a particular set, called the **domain** of the quantifier. If $P(x)$ is a predicate,

$$\forall x \in X : P(x)$$

is read as “for all x belonging to X , the statement $P(x)$ is true.” So

$$\forall x \in \mathbb{R} : x^2 > 0$$

is false, because not every real number x satisfies $x^2 > 0$, but

$$\forall x \in \mathbb{R} : x^2 \geq 0$$

is true, since every real number x satisfies $x^2 \geq 0$.

The existential quantifier \exists is used to say that something is true for at least one element of a particular set, again called the **domain**. If $P(x)$ is a predicate

$$\exists x \in X : P(x)$$

says “there exists x in X such that $P(x)$ is true.” Such a statement is true when there is *at least one* x belonging to the set X making $P(x)$ true. For example,

$$\exists x \in \mathbb{R} : x^2 > 0$$

is true, because there is at least one real number x whose square is positive (in fact, there’s infinitely many!), but

$$\exists x \in \mathbb{R} : x^2 < 0$$

is false since there is no real number x whose square is negative.

Let’s write the three examples at the start of this subsection using quantifier notation. We use \mathbb{Q} to represent the set of rational numbers, \mathbb{R} for the set of real numbers, \mathbb{N} for the set of natural numbers, and \mathbb{P} for the set of prime numbers. The statements are then:

- $\exists q \in \mathbb{Q} : q^2 = 2$
- $\forall x \in \mathbb{R} : x^2 > 0$
- $\forall n \in \mathbb{N} \exists p \in \mathbb{P} : p > n$.

1.9.3 Note on how domains for quantifiers are specified

In all the examples above, we specified the domain of the quantifier explicitly. Sometimes, especially in analysis, the domain is specified using a condition. Analysts will write

$$\forall \epsilon > 0 \exists \delta > 0 : \dots$$

to mean “for all positive real numbers ϵ , there exists a positive real number δ , ...”. You can think of this as meaning

$$\forall \epsilon \in (0, \infty) \exists \delta \in (0, \infty) : \dots$$

where $(0, \infty)$ means the set of all positive real numbers.

There are even times when the domain of one variable seems to depend on another — you might see an expression like

$$\forall n \in \mathbb{N} \exists p \geq n : P(p, n)$$

If p is intended to be a natural number, for example, we could rewrite this as

$$\forall n \in \mathbb{N} \exists p \in \mathbb{N} : (p \geq n) \wedge P(p, n).$$

1.9.4 Multiple quantifiers

We need more than one quantifier to express statements whose truth depends on more than one variable. For example, in analysis, the definition of a function f being bounded is that

$$\exists M \in \mathbb{R} \forall x \in \mathbb{R} : |f(x)| \leq M,$$

that is, there exists a real number M such that for all real numbers x we have $|f(x)| \leq M$. The definition of a function being continuous at a point is even more complicated, requiring three quantifiers.

When we have more than one quantifier, it's important to get the quantifiers in the correct order because if you change the order of the quantifiers you may change whether or not the statement is true.

Example 1.9.2.

$$\forall x \in \mathbb{R} \exists y \in \mathbb{R} : x + y = 0$$

is **true**, because no matter what real number x you give me, I can find a real number y (equal to $-x$) such that $x + y = 0$. On the other hand

$$\exists y \in \mathbb{R} \forall x \in \mathbb{R} : x + y = 0$$

which is the same statement except that the quantifiers are the other way round, is **false** because there is no real number y such that for all real numbers x we have $x + y = 0$.

However, quantifiers *of the same type* can be swapped without changing the truth or falsity of the statement, for example both

$$\forall x \in X \forall y \in Y : P(x, y)$$

and

$$\forall y \in Y \forall x \in X : P(x, y)$$

are true if and only if $P(x, y)$ is true for every x in X and every y in Y .

1.9.5 Negation

It's often useful to ask what it would mean for a quantified statement *not* to be true. For example, you might want to prove that a certain sequence (a_n) *does not* tend to 0. The definition of a sequence tending to 0 is that

$$\forall \epsilon \in (0, \infty) \exists N \in \mathbb{N} \forall n \in \mathbb{N} : n \geq N \implies |a_n| < \epsilon.$$

We could negate this just by adding a \neg to the front: to say that a sequence does not tend to 0 is to say that

$$\neg \forall \epsilon \in (0, \infty) \exists N \in \mathbb{N} \forall n \in \mathbb{N} : n \geq N \implies |a_n| < \epsilon$$

but it would be helpful if we could simplify this in some way, to make it easier to prove. We can do this with the following quantifier rules.

Theorem 1.9.1. *For any predicate $P(x)$ on a set X ,*

1. $\neg(\forall x \in X : P(x))$ is true if and only if $\exists x \in X : \neg P(x)$ is true, and
2. $\neg(\exists x \in X : P(x))$ is true if and only if $\forall x \in X : \neg P(x)$ is true.

The first of these holds because to say $\forall x \in X : P(x)$ is false means that not every $P(x)$ for $x \in X$ is true, that is, there exists $x \in X$ such that $P(x)$ is false. The second can be seen similarly.

By using these rules repeatedly, we can find equivalent forms of negated quantified statements even when they contain multiple quantifiers. Returning to our example of a sequence not tending to zero,

$$\begin{aligned} \neg \forall \epsilon \in (0, \infty) \exists N \in \mathbb{N} \forall n \in \mathbb{N} : n \geq N \implies |a_n| < \epsilon \\ \text{is true if and only if} \\ \exists \epsilon \in (0, \infty) \neg \exists N \in \mathbb{N} \forall n \in \mathbb{N} : n \geq N \implies |a_n| < \epsilon \\ \text{is true, if and only if} \\ \exists \epsilon \in (0, \infty) \forall N \in \mathbb{N} \neg \forall n \in \mathbb{N} : n \geq N \implies |a_n| < \epsilon \\ \text{is true, if and only if} \\ \exists \epsilon \in (0, \infty) \forall N \in \mathbb{N} \exists n \in \mathbb{N} : \neg(n \geq N \implies |a_n| < \epsilon). \end{aligned}$$

At this point we can use what we know about logical equivalences for WFFs to simplify further: this is equivalent to

$$\exists \epsilon \in (0, \infty) \forall N \in \mathbb{N} \exists n \in \mathbb{N} : (n \geq N) \wedge (|a_n| \geq \epsilon).$$

We now know what we need to do to prove (a_n) doesn't tend to 0: we have to find some $\epsilon > 0$ such that for every N , there is an $n \geq N$ such that $|a_n| \geq \epsilon$.

We will now look at some more examples of producing useful equivalent forms for negations of quantified statements. We're going to use real-life examples from bits of mathematics you may not have met yet, but this won't be a problem as our negation procedure doesn't require understanding anything about the meaning of the statement!

Example 1.9.3. The statement "every value the function $f : \mathbb{R} \rightarrow \mathbb{R}$ takes is less than 10" can be written

$$\forall x \in \mathbb{R} : f(x) < 10.$$

What does it mean for this to be false? Let's negate it, using the negation of quantifiers theorem 1.9.1. That tells us

$$\neg \forall x \in \mathbb{R} : P(x)$$

is equivalent to

$$\exists x \in \mathbb{R} : \neg P(x)$$

which in our case says $\exists x \in \mathbb{R} : \neg(f(x) < 10)$, or $\exists x \in \mathbb{R} : f(x) \geq 10$.

Example 1.9.4. Consider the statement “the function $f : \mathbb{R} \rightarrow \mathbb{R}$ is bounded”, which we could write as

$$\exists M \in \mathbb{R} \forall x \in \mathbb{R} |f(x)| \leq M.$$

Let’s negate it. To keep things short, we’ll write $P(x, M)$ for the predicate $|f(x)| \leq M$. We know that

$$\neg \exists M \in \mathbb{R} \forall x \in \mathbb{R} : P(x, M)$$

is true if and only if

$$\forall M \in \mathbb{R} \neg \forall x \in \mathbb{R} : P(x, M)$$

is true, if and only if

$$\forall M \in \mathbb{R} \exists x \in \mathbb{R} : \neg P(x, M).$$

so “the function f is not bounded” can be written as $\forall M \in \mathbb{R} \exists x \in \mathbb{R} : \neg(|f(x)| \leq M)$, or equivalently, $\forall M \in \mathbb{R} \exists x \in \mathbb{R} : |f(x)| > M$.

Further reading

Chapter 6 of *The Foundations of Mathematics* by Ian Stewart and David Tall covers a lot of this material at a similar level. The book *Propositional and Predicate Calculus: A Model of Argument* by Derek Goldrei goes far beyond what we cover in MATH0005, but I recommend it if you want to know about logic in much more depth. You can also take the 3rd year course MATH0037 Logic.

Chapter 3 of the free online book *Discrete Mathematics: An Open Introduction* by Oscar Levin has material on propositional calculus and first order logic, though it doesn’t use the same framework of well-formed formulas that we do and the vocabulary they use is slightly different.

Chapter 2

Sets and functions

2.1 Introduction to set theory

2.1.1 Definition of a set

A **set** is a collection of (mathematical) objects. There is an entire field of mathematics called set theory dedicated to the study of sets and to their use as a foundation for mathematics, but in MATH0005 we are going to give only an informal introduction to sets and their properties. If you want to know more, see the further reading section at the end of this chapter.

2.1.2 Set builder notation

We specify sets using logical notation: if $P(x)$ is a property that is true or false for each thing x , then

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

is the collection of all x such that $P(x)$ is true.¹ For example, if $P(x)$ is the property that x is an even integer then $\{x : P(x)\}$ is the set consisting of all the even integers $\dots, -4, -2, 0, 2, 4, \dots$

Sometimes we already have a set X and a property $P(x)$ that is true or false for each x in X . For example, X might be the set \mathbb{R} of all real numbers, and $P(x)$ the property $x > 0$. We write

$$\{x \in X : P(x)\} \tag{2.1}$$

for the set of all things x in X such that $P(x)$ is true. In the example, this set would be the set of all strictly positive real numbers, sometimes written $(0, \infty)$ in interval notation. The symbol \in means ‘is an element of.’

These two ways of specifying a set are called **set builder notation**.

2.1.3 Listing the elements of a set

Sometimes we want to specify a set simply by listing its elements one by one. For example, the set whose elements are 1, 2, and 3 is written $\{1, 2, 3\}$, which

¹You have to be slightly careful with this kind of *unrestricted comprehension* because it can lead to contradictions. You can ignore this for the purposes of MATH0005, but if you want to know more then check out the Further Reading section at the end.

is a short way of writing for the set builder notation

$$\{x : (x = 1) \vee (x = 2) \vee (x = 3)\}.$$

2.1.4 Elements of a set

The things in a set are called its **elements** or **members**. We write $a \in X$ to mean that a is an element or member of the set X , and $a \notin X$ to mean that a is not an element or member of X .

There is a unique set with no elements, called the **empty set** and written \emptyset . No matter what a is, $a \notin \emptyset$.

We allow any kind of mathematical object, including sets themselves, as elements of sets. Sets can contain functions, matrices, vectors, numbers, and sets themselves.

Example 2.1.1.

$$\{\emptyset, 1, \{2\}, \{\{3\}\}$$

is a set whose four elements are the empty set, the number 1, the set containing the number 2, and the set containing the set containing 3.

Example 2.1.2. Let $X = \{1, 2, \{3\}\}$. Then $1 \in X$, $0 \notin X$, $3 \notin X$, $\{3\} \in X$.

2.1.5 Set equality

Two sets X and Y are defined to be equal if for all things a we have $a \in X$ if and only if $a \in Y$ — that is, if and only if they have the same elements. It's helpful to write this in terms of logical equivalence: the two sets

$$X = \{x : P(x)\}, Y = \{x : Q(x)\}$$

are equal if and only if $P(x)$ and $Q(x)$ have the same truth value for every x . This definition of set equality has some important consequences. First, if

$$X = \{1, 2\}, Y = \{2, 1\}$$

then $X = Y$ because these are shorthand for

$$X = \{x : (x = 1) \vee (x = 2)\}, Y = \{x : (x = 2) \vee (x = 1)\}$$

and for any propositional variables p and q , the WFFs $p \vee q$ and $q \vee p$ are logically equivalent. **The order in which we list the elements of a set doesn't matter.** Similarly, if

$$X = \{1, 1\}, Y = \{1\}$$

then $X = Y$, because the two definitions above are shorthand for

$$X = \{x : (x = 1) \vee (x = 1)\}, Y = \{x : x = 1\}$$

and the two formulas $p \vee p$ and p are logically equivalent. **Repetition doesn't matter.**

2.1.6 Subsets

We need vocabulary for talking about one set being contained in another.

Definition 2.1.1. • X is a **subset** of Y , written $X \subseteq Y$, if and only if every element of X is also an element of Y .

- If X is not a subset of Y we write $X \not\subseteq Y$.
- X is a **proper subset** of Y , written $X \subsetneq Y$, if and only if $X \subseteq Y$ but $X \neq Y$.

Thus X being a proper subset of Y means that X is a subset of Y and Y contains something that X does not contain.

There is an important way to rephrase the definition of two sets being equal: $X = Y$ if and only if $X \subseteq Y$ and $Y \subseteq X$. This is sometimes useful as a proof technique, as you can split a proof of $X = Y$ into first checking $X \subseteq Y$ and then checking $Y \subseteq X$.

Example 2.1.3. • $\{0\} \subseteq \{0, 1\}$

- $\{0\} \subsetneq \{0, 1\}$
- $\{0, 1\} \not\subseteq \{1, 2\}$
- $\{1, 2, 1\} = \{2, 1\}$

Why is the last equality true? The only things which are elements of $\{1, 2, 1\}$ are 1 and 2. The only things which are elements of $\{1, 2\}$ are 1 and 2. So the two sets are equal according to our definition. There's no concept of something being an element of a set "more than once."

This is the way in which our definition of set equality captures the idea of sets being unordered collections of objects which disregard repetition.

The definition of subset means that the empty set is a subset of any set. $\emptyset \subseteq X$ for any set X , because $\forall x : x \in \emptyset \implies x \in X$ is *vacuously true*: there's nothing in \emptyset which could fail to be in the set X in order to make $\emptyset \subseteq X$ false.

2.2 Set operators

2.2.1 Union, intersection, difference, complement

Definition 2.2.1. Let A and B be sets.

- $A \cup B$, the **union** of A and B , is $\{x : x \in A \vee x \in B\}$.
- $A \cap B$, the **intersection** of A and B , is $\{x : x \in A \wedge x \in B\}$.
- $A \setminus B$, the **set difference** of A and B , is $\{x \in A : x \notin B\}$.
- If A is a subset of a set Ω then A^c , the **complement** of A in Ω , is $\{x \in \Omega : x \notin A\}$.

We can express complements using set differences. If A is a subset of Ω then its complement A^c in Ω is equal to $\Omega \setminus A$.

Example 2.2.1. Suppose $A = \{0, 1, 2\}$, $B = \{1, 2, 3\}$, $C = \{4\}$.

- $A \cup B = \{0, 1, 2, 3\}$
- $A \cap B = \{1, 2\}$
- $A \cap C = \emptyset$
- $A \setminus B = \{0\}$
- $A \setminus C = A$
- If $\Omega = \{0, 1, 2, \dots\}$ then A^c would be $\{3, 4, \dots\}$.

The set $\mathbb{N} = \{0, 1, 2, \dots\}$ is called the **natural numbers**. Some people exclude 0 from \mathbb{N} but in MATH0005 the natural numbers include 0.

It's typical to draw Venn diagrams to represent set operations. We draw a circle, or a blob, for each set. The elements of the set A are represented by the area inside the circle labelled A . Here are some examples:




Figure 2.1: Venn diagram for the set difference of A and B







Figure 2.2: Venn diagram for the union of A and B

Figure 2.3: Venn diagram for the intersection of A and B Figure 2.4: Venn diagram for the complement of the union of A and B Figure 2.5: Venn diagram for the complement of the intersection of A and B

2.2.2 Size of a set

Definition 2.2.2. The **size** or **cardinality** of a set X , written $|X|$, is the number of *distinct* elements it has.

Example 2.2.2. • $|\{1, 2\}| = 2$

- $|\emptyset| = 0$
- $|\{1, 2, 1, 3\}| = 3$

Definition 2.2.3. A set is **finite** if it has 0, or 1, or 2, or any other natural number of elements. A set that is not finite is called **infinite**.

\mathbb{N} and \mathbb{Z} are infinite sets while the sets in Example 2.2.2 are all finite.

2.3 Set algebra

2.3.1 Commutativity and associativity

This section is about the laws that union, intersection, difference, and complement obey.

Theorem 2.3.1. *For all sets A and B ,*

- $A \cap B = B \cap A$, and
- $A \cup B = B \cup A$.

These are called the **commutativity** properties for intersection and union. These results might seem obvious, but we will write out the proofs carefully because the method of using logical equivalences will be applied to more complex set identities later.

Proof. By definition,

$$\begin{aligned} A \cap B &= \{x : x \in A \wedge x \in B\} \\ B \cap A &= \{x : x \in B \wedge x \in A\}. \end{aligned}$$

Theorem 1.5.1 tells us that for any two WFFs ϕ and ψ , the formulas $(\phi \wedge \psi)$ and $(\psi \wedge \phi)$ are logically equivalent: one is true if and only if the other is true. So

$$x \in A \wedge x \in B$$

is true if and only if

$$x \in B \wedge x \in A$$

is true. This shows that for any x we have $x \in A \cap B$ if and only if $x \in B \cap A$, so by definition of set equality, $A \cap B = B \cap A$.

The argument for \cup is the same, except that we use the logical equivalence $(\phi \vee \psi) \equiv (\psi \vee \phi)$. \square

What this proof shows is that if you have a set X defined in set builder notation using a logical formula

$$X = \{x : P(x)\}$$

then it is equal to any other set defined using a logically equivalent formula.

Theorem 2.3.2. *For all sets A , B , C we have*

- $A \cap (B \cap C) = (A \cap B) \cap C$ and
- $A \cup (B \cup C) = (A \cup B) \cup C$.

This is the **associativity** property for \cap and \cup .

Proof. Like the proof of the last theorem, these equalities follow from the associativity properties for \wedge and \vee we saw in Theorem 1.5.1. \square

Associativity tells us means there's no ambiguity in writing

$$A \cup B \cup C \text{ or } A \cap B \cap C$$

without any brackets to indicate which union or intersection should be done first. Compare this with

$$1 + 2 + 3.$$

There's no need for brackets because it doesn't matter whether you do $1+2$ first then add 3, or whether you add 1 to the result of $2 + 3$. On the other hand $1 + 2 \times 3$ or $1 - 2 - 3$ require either brackets or a convention on which operation to do first. Similarly $A \cup (B \cap C)$ is different to $(A \cup B) \cap C$ in general, so the brackets here are obligatory, and $A \setminus (B \setminus C)$ is different to $(A \setminus B) \setminus C$.

2.3.2 The distributive laws

Because we defined unions and intersections using logical conditions on set elements, they should obey laws that come from the results we proved about \wedge and \vee .

Theorem 2.3.3. *For any sets A, B, C*

1. $A \cup (B \cap C) = (A \cup B) \cap (A \cup C)$ and
2. $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$.

Proof. Consider the first of these identities. The left hand side consists of all things x such that

$$x \in A \vee (x \in B \wedge x \in C). \quad (2.2)$$

The right hand side consists of all things x such that

$$(x \in A \vee x \in B) \wedge (x \in A \vee x \in C) \quad (2.3)$$

By Theorem 1.6.1, for any WFFs ϕ, ψ, θ we have $\phi \vee (\psi \wedge \theta) \equiv (\phi \vee \psi) \wedge (\phi \vee \theta)$. Thus any x makes (2.2) true if and only if it makes (2.3) true. So x belongs to the first set if and only if it belongs to the second, therefore the two sets are equal.

The second identity can be proved similarly. □

2.4 De Morgan's laws

Take a look at this Venn diagram:




Figure 2.6: Venn diagram for the complement of the union of A and B

You can see that the shaded area is exactly the area not in $A \cup B$, so this is the Venn diagram for $(A \cup B)^c$. Now consider the Venn diagrams for A^c and B^c :




Figure 2.7: Venn diagram for the complement of A




Figure 2.8: Venn diagram for the complement of B

You can see from the diagrams that $A^c \cap B^c = (A \cup B)^c$. This is a general and useful fact, one of *De Morgan's laws*.

Theorem 2.4.1. (*De Morgan's laws for sets*). Let $A, B \subseteq \Omega$ and let A^c and B^c denote the complement with respect to Ω . Then

1. $(A \cup B)^c = A^c \cap B^c$, and
2. $(A \cap B)^c = A^c \cup B^c$.

Proof. These follow from De Morgan's laws in logic. The left hand side of the first of these is the set of all $x \in \Omega$ such that

$$\neg(x \in A \vee x \in B)$$

and the right hand side is the set of all $x \in \Omega$ such that

$$\neg(x \in A) \wedge \neg(x \in B).$$

Since $\neg(p \vee q)$ is logically equivalent to $(\neg p \wedge \neg q)$ (Theorem 1.6.3), the two sets have the same elements and so are equal. The second equality follows from the other logical De Morgan law. \square

De Morgan's laws also work for unions and intersections of more than two sets.

Theorem 2.4.2. For any sets A_1, A_2, \dots

1. $(A_1 \cup A_2 \cup \dots)^c = A_1^c \cap A_2^c \cap \dots$, and
2. $(A_1 \cap A_2 \cap \dots)^c = A_1^c \cup A_2^c \cup \dots$

2.5 Functions

2.5.1 Definition of function, domain, codomain

Informally, given two sets X and Y a **function** or **map** f from X to Y is a definite rule which associates to each $x \in X$ an element $f(x) \in Y$. A ‘rule’ here could be some kind of mathematical formula, but that isn’t necessary — it could simply be a list specifying the output $f(x)$ for each input $x \in X$.

Definition 2.5.1. We write $f : X \rightarrow Y$ to mean that f is a function from X to Y . X is called the **domain** of f and Y is called the **codomain** of f .

We refer to the element $f(x)$ of Y as being the “output” or “value” of f when it is given the “input” or “argument” x .

This might seem vague: what is a definite rule? What does associates mean? Should we say that two functions with the same domain and codomain are equal if and only if they have the same rule, or should it be if and only if they have the same output for every input?²

2.5.2 When are two functions equal?

Definition 2.5.2. Two functions f and g are said to be equal, and we write $f = g$, if and only if

- they have the same domain, say X , and
- they have the same codomain, and
- for all $x \in X$ we have $f(x) = g(x)$.

Sometimes the definition has slightly strange-looking consequences.

Example 2.5.1. Let $f, g : \{0, 1\} \rightarrow \{0, 1\}$. $f(x) = x^2$. $g(x) = x$. Are they equal?

(the answer is yes — they have the same domain, same codomain, and the same output for every input in their common domain).

Definition 2.5.3. For any set X , the **identity function** $\text{id}_X : X \rightarrow X$ is defined by $\text{id}_X(x) = x$ for all $x \in X$.

Sometimes we just write id instead of id_X if it is clear which set we are talking about.

²These concepts are called *intensional* and *extensional* equality.

2.6 Function composition

2.6.1 Definition of function composition

Suppose you have two functions $f : X \rightarrow Y$ and $g : Y \rightarrow Z$:

$$X \xrightarrow{f} Y \xrightarrow{g} Z$$

Then you can make a new function $X \rightarrow Z$ whose rule is “do f , then do g ”.

Definition 2.6.1. Let $f : X \rightarrow Y$ and $g : Y \rightarrow Z$. The **composition** of g and f , written $g \circ f$ or gf , is the function $X \rightarrow Z$ with rule $(g \circ f)(x) = g(f(x))$.

This makes sense because $f(x)$ is an element of Y and g has domain Y so we can use any element of Y as an input to g .

It’s important to remember that $g \circ f$ is the function whose rule is “do f , then do g ”.

Proposition 2.6.1. If $f : X \rightarrow Y$ then $f \circ \text{id}_X = \text{id}_Y \circ f = f$.

Proof. For any $x \in X$ we have $(f \circ \text{id}_X)(x) = f(\text{id}_X(x)) = f(x)$ and $(\text{id}_Y \circ f)(x) = \text{id}_Y(f(x)) = f(x)$. \square

2.6.2 Associativity

Functions f and g such that the codomain of f equals the domain of g , in other words, functions such that $g \circ f$ makes sense, are called **composable**. Suppose that f and g are composable and g and h are also composable, so that we can draw a diagram

$$X \xrightarrow{f} Y \xrightarrow{g} Z \xrightarrow{h} W.$$

It seems there are two different ways to compose these three functions: you could first compose f and g , then compose the result with h , or you could compose g with h and then compose the result with f . But they both give the same result, because function composition is **associative**.

Lemma 2.6.2. Let $f : X \rightarrow Y, g : Y \rightarrow Z, h : Z \rightarrow W$. Then $h \circ (g \circ f) = (h \circ g) \circ f$.

Proof. Both $h \circ (g \circ f)$ and $(h \circ g) \circ f$ have the same domain X , same codomain W , and same rule that sends x to $h(g(f(x)))$. \square

The associativity property says that a composition like $h \circ g \circ f$ doesn’t need any brackets to make it unambiguous: however you bracket it, the result is the same. In fact we can omit brackets from a composition of any length without ambiguity.

2.7 Function properties

2.7.1 Image of a function

We need to extend our set builder notation a little. If $f : X \rightarrow Y$ is a function, we write $\{f(x) : x \in X\}$ for the subset of Y consisting of all elements $f(x)$ for $x \in X$. This set has a special name.

Definition 2.7.1. Let $f : X \rightarrow Y$. Then the **image** of f , written $\text{im}(f)$, is defined to be $\{f(x) : x \in X\}$.

Don't confuse codomain and image. Y is the codomain of f and the image $\text{im}(f)$ is a *subset* of Y , but it need not equal Y .

Some people use the word *range* to refer to one of these two concepts, but since different people use it for different things we will only use the terms image and codomain in MATH0005.

Example 2.7.1. • Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be the function $f(x) = x^2$. Every element $f(x)$ of the image of f is a nonnegative number, and every nonnegative number is the square of some real number, so $\text{im}(f) = [0, \infty)$.

- Let $g : \mathbb{Z} \rightarrow \mathbb{Z}$ be defined by $g(z) = 3z$. Then $\text{im}(g) = \{g(z) : z \in \mathbb{Z}\} = \{3z : z \in \mathbb{Z}\}$.

2.7.2 Injection, surjection, bijection

Definition 2.7.2. Let $f : X \rightarrow Y$ be a function.

- We say f is **injective** or **one-to-one** if and only if for all $a, b \in X$, if $f(a) = f(b)$ then $a = b$.
- We say f is **surjective** or **onto** if and only if for all $y \in Y$ there is at least one $x \in X$ such that $f(x) = y$.
- We say f is a **bijection** if and only if it is injective and surjective.

Another way to write the definition of surjective would be that a function is surjective if and only if its image equals its codomain.

As an example, here's a picture of a function $f : \{1, 2, 3, 4, 5, 6\} \rightarrow \{1, 2, 3, 4, 5\}$. There is an arrow from x to $f(x)$ for each x in the domain of f .




Figure 2.9: Drawing of a function from $\{1, 2, 3, 4, 5, 6\}$ to $\{1, 2, 3, 4, 5\}$ such that $f(1) = f(2) = 1$, $f(3) = 5$, $f(4) = 3$, $f(5) = 1$, $f(6) = 2$.

The function f shown in Figure 2.9 is not onto because $\text{im}(f)$ is a proper subset of the codomain, specifically, the codomain contains 4 but $\text{im}(f)$ does not. f is not one-to-one because $f(1) = f(2)$ but $1 \neq 2$.

Example 2.7.2. Here are some more examples to illustrate the injective, surjective, and bijective properties.

- $f : \mathbb{R} \rightarrow \mathbb{R}, f(x) = x^2$. It isn't injective as $f(-1) = f(1)$ and it isn't surjective as -1 is in the codomain, but there's no element x in the domain such that $f(x) = -1$.
- $g : \mathbb{R} \rightarrow [0, \infty), g(x) = x^2$. This is not injective for the same reason as before, but this time it *is* surjective: for each $y \geq 0$ we can find an element of the domain which g sends to y : for example. $g(\sqrt{y}) = y$.
- $h : [0, \infty) \rightarrow \mathbb{R}, h(x) = x^2$. Not surjective, for the same reason f isn't surjective (the codomain contains negative numbers, but the image doesn't contain any negative numbers, so the image doesn't equal the codomain). But h is injective: if x and y are in the domain of h and $h(x) = h(y)$ then $x^2 = y^2$, so $x = \pm y$. Since elements of the domain of h are nonnegative, it must be that $x = y$.
- $j : (-\infty, 0] \rightarrow [0, \infty), j(x) = x^2$. This is injective (for a similar reason to h) and surjective (for a similar reason to g), so it is a bijection.

All of these functions had their rules described in the same way, but their properties differed. This shows how important it is to specify the domain and codomain when you talk about a function. A question like "is the function $f(x) = x^2$ injective?" doesn't make any sense unless you do this.

2.7.3 Bijections and sizes of sets

How do we know when two sets have the same size? If you see an alien creature with an apple in each of its hundreds of hands you know it has the same number of apples as it does hands, even if you haven't counted either the apples or the hands.

You know that because you can pair each apple with the hand holding it. Every apple is held by one hand, and every hand holds one apple.

Suppose there is a bijection f between two sets X and Y . This gives us a way to pair up elements of X and elements of Y such that every element of X is paired with exactly one element of Y .

Consider the pairs $(x, f(x))$ for x in X . Every element of Y appears in exactly one of these pairs (at least one pair because f is onto, at most one pair because f is one-to-one). So a bijection pairs up each element of X with a unique element $f(x)$ of Y .




Figure 2.10: Picture of a bijection f from $\{1, 2, 3\}$ to $\{a, b, c\}$ such that $f(1) = c, f(2) = b, f(3) = a$

The picture shows a bijection $f : \{1, 2, 3\} \rightarrow \{a, b, c\}$. If we pair each element x of the domain with its image $f(x)$ we get the pairs $\langle 1, c \rangle, \langle 2, b \rangle, \langle 3, a \rangle$. Because f is a bijection, every element of the domain is paired with exactly one element of the codomain and every element of the codomain is paired with exactly one element of the domain. This leads us to make the definition that two sets have the **same size** (or the **same cardinality**) if and only if there is a bijection between them.

This definition works even for infinite sets — though it sometimes provides some counter-intuitive results. The set of integers \mathbb{Z} and the set of even integers $2\mathbb{Z} = \{\dots - 4, -2, 0, 2, 4, 6, \dots\}$ have the same size since there is a bijection

$$\begin{aligned} f : \mathbb{Z} &\rightarrow 2\mathbb{Z} \\ f(z) &= 2z \end{aligned}$$

even though one is a proper subset of the other.

2.8 Invertibility

Definition 2.8.1. Let $f : X \rightarrow Y$.

- A **left inverse** to f is a function $g : Y \rightarrow X$ such that $g \circ f = \text{id}_X$.
- A **right inverse** to f is a function $h : Y \rightarrow X$ such that $f \circ h = \text{id}_Y$.
- An **inverse** (or a **two sided inverse**) to f is a function $k : Y \rightarrow X$ which is a left and a right inverse to f .

We say f is **invertible** if it has a two sided inverse.

Notice that if g is left inverse to f then f is right inverse to g . A function can have more than one left inverse, or more than one right inverse: you will investigate this further in the problem sets.

The idea is that a left inverse “undoes” its right inverse, in the sense that if you have a function f with a left inverse g , and you start with $x \in X$ and apply f to get to $f(x) \in Y$, then doing g gets you back to where you started because $g(f(x)) = x$.

Example 2.8.1. • $f : \mathbb{R} \rightarrow [0, \infty), f(x) = x^2$ has a right inverse $g : [0, \infty) \rightarrow \mathbb{R}, g(x) = \sqrt{x}$. $f(g(x)) = x$ for all $x \in [0, \infty)$. It is not the case that g is a left inverse to f because $g(f(-1)) \neq -1$.

- This function f **does not have** a left inverse. Suppose h is left inverse to f , so that $hf = id_{\mathbb{R}}$. Then $h(f(-1)) = -1$, so $h(1) = -1$. Similarly $h(f(1)) = 1$, so $h(1) = 1$. Impossible! (The problem, as we will see in the next section, is that f isn't one-to-one.)
- The function g has a left inverse, f . But it **does not have** a right inverse. If $g \circ h = id_{\mathbb{R}}$ then $g(h(-1)) = -1$ so $g(h(-1)) = -1$. But there's no element of $[0, \infty)$ that g takes to -1 . (This time the problem is that g isn't onto.)

2.9 Conditions for invertibility

Here is the connexion between function properties and invertibility.

Theorem 2.9.1. *Let $f : X \rightarrow Y$ be a function between nonempty sets.*

1. f has a left inverse if and only if it is injective.
2. f has a right inverse if and only if it is surjective.
3. f has a two sided inverse if and only if it is bijective.

- Proof.*
1.
 - ONLY IF. Let g be a left inverse to f , so $g \circ f = id_X$. Suppose $f(a) = f(b)$. Then applying g to both sides, $g(f(a)) = g(f(b))$, so $a = b$.
 - IF. Let f be injective. Choose any x_0 in the domain of f . Define $g : Y \rightarrow X$ as follows. Each y in Y is either in the image of f or not. If y is in the image of f , it equals $f(x)$ for a *unique* x in X (uniqueness is because of the injectivity of f), so define $g(y) = x$. If y is not in the image of f , define $g(y) = x_0$. Clearly $g \circ f = id_X$.
 2.
 - IF. Suppose f is surjective. Let $y \in Y$. Then y is in the image of f , so we can choose an element $g(y) \in X$ such that $f(g(y)) = y$. This defines a function $g : Y \rightarrow X$ which is evidently a right inverse to f .
 - ONLY IF. Suppose f has a right inverse g , so $f \circ g = id_Y$. If $y \in Y$ then $f(g(y)) = id_Y(y) = y$, so $y \in \text{im}(f)$. Every element of Y is therefore in the image of f , so f is onto.
 3. If f has a left inverse and a right inverse, it is injective (by part 1 of this theorem) and surjective (by part 2), so is a bijection. Conversely if f is a bijection it has a left inverse $g : Y \rightarrow X$ and a right inverse $h : Y \rightarrow X$ by part 1 and part 2 again. We will now show $g = h$, so that g is a two sided inverse to f .

$$\begin{array}{ll}
 g = g \circ id_Y & \text{Proposition 2.6.1} \\
 = g \circ (f \circ h) & \text{as } f \circ h = id_Y \\
 = (g \circ f) \circ h & \text{associativity} \\
 = id_X \circ h & \text{as } g \circ f = id_X \\
 = h &
 \end{array}$$

so $g = h$ is a two sided inverse of f . □




Figure 2.11: A diagram illustrating the construction, in part 1 of the theorem, of the left inverse to an injective function $f : X \rightarrow Y$ where $X = \{x, y, z\}$, $Y = \{a, b, c, d, e\}$, and $f(x) = c$, $f(y) = b$, $f(z) = a$. Left-to-right arrows show where f sends elements of X and right-to-left arrows show where g sends elements of Y . The elements d and e of Y which are not in the image of f are all sent to the element x of X .

Figure 2.11 illustrates the construction in part 1 of the theorem. Arrows from left to right show where f sends each element of X . Arrows from right to left show where the left inverse g we have constructed sends each element of Y .

Definition 2.9.1. If $f : X \rightarrow Y$ is invertible, we write f^{-1} for the two sided inverse of f .

It makes sense to talk about *the* two sided inverse to f because there really is only one: if g and h are two sided inverses of f then certainly g is a left inverse and h is a right inverse, so the argument in the proof of part 3 of the theorem above shows $g = h$.

2.9.1 Inverse of a composition

Theorem 2.9.2. If $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ are invertible then so is $g \circ f$, and $(g \circ f)^{-1} = f^{-1} \circ g^{-1}$.

Proof. $f^{-1} \circ g^{-1}$ is a left inverse to $g \circ f$, because

$$\begin{aligned} (f^{-1} \circ g^{-1}) \circ (g \circ f) &= f^{-1} \circ (g^{-1} \circ g) \circ f && \text{associativity} \\ &= f^{-1} \circ \text{id}_Y \circ f \\ &= f^{-1} \circ f \\ &= \text{id}_X . \end{aligned}$$

A similar calculation shows that it is a right inverse as well. \square

It is important to get this the right way round. The inverse of $g \circ f$ is **not** normally $g^{-1} \circ f^{-1}$, indeed this composition may not even make sense. The correct result is easy to remember when you think about getting dressed. Each morning you put on your socks, then you put on your shoes: if k is the put-on-socks function and h is the put-on-shoes function then you apply the function $h \circ k$ to your feet. The inverse of this is taking off your shoes, then taking off your socks: $k^{-1} \circ h^{-1}$. Not the other way round — it's not even (normally) possible to take off your socks, then take off your shoes, just as it is not normally possible to form the composition $g^{-1} \circ f^{-1}$ in the context of the theorem above.³

A similar result applies when you compose more than two invertible functions: if f_1, f_2, \dots, f_n are invertible and if the composition

$$f_1 \circ \dots \circ f_n$$

makes sense, it is also invertible and its inverse is

$$f_n^{-1} \circ \dots \circ f_1^{-1}.$$

2.10 Permutations

Definition 2.10.1. • A **permutation** of a set X is a bijection $X \rightarrow X$.

- S_n , the **symmetric group** on n letters, is the set of all permutations of $\{1, 2, \dots, n\}$.

Example 2.10.1. • For any set X , the identity function $\text{id}_X : X \rightarrow X$ is a permutation.

- The function $f : \{1, 2, 3\} \rightarrow \{1, 2, 3\}$ given by $f(1) = 3, f(2) = 2, f(3) = 1$ is a permutation. f is an element of S_3 .
- The function $g : \{1, 2, 3, 4\} \rightarrow \{1, 2, 3, 4\}$ given by $g(1) = 2, g(2) = 3, g(3) = 4, g(4) = 1$ is a permutation. g is an element of S_4 .

Here are some diagrams illustrating the permutations f and g :





Figure 2.12: Diagram of the permutation f from the example above.

³The shoes and socks illustration comes from Gilbert Strang's famous 18.06 linear algebra course.

Figure 2.13: Diagram of the permutation g from the example above.

2.10.1 Two row notation

We need a way of writing down elements of S_n . The simplest is called **two row notation**. To represent $f \in S_n$, you write two rows of numbers. The top row is $1, 2, \dots, n$. Then underneath each number i on the top row you write $f(i)$:

$$\begin{pmatrix} 1 & 2 & \cdots & n \\ f(1) & f(2) & \cdots & f(n) \end{pmatrix}$$

As an example, here are the two row notations for the two permutations of the previous example.

$$f : \begin{pmatrix} 1 & 2 & 3 \\ 3 & 2 & 1 \end{pmatrix}$$

$$g : \begin{pmatrix} 1 & 2 & 3 & 4 \\ 2 & 3 & 4 & 1 \end{pmatrix}$$

The two row notation for the identity in S_n is particularly simple:

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

This is a simple and not-that-efficient notation (it is not feasible to write down an element of S_{100} this way, even if it is a very simple permutation e.g. swaps 1 and 2, leaves 3-100 alone), but it is at least concrete and simple.

2.10.2 How many permutations?

$n!$, pronounced *n factorial*, means $n \times (n-1) \times \cdots \times 2 \times 1$.

Theorem 2.10.1. $|S_n| = n!$

Proof. We will count how many possible bottom rows of two row notations for elements of S_n there are. Because a permutation is a bijection — one-to-one and onto — this bottom row consists of the numbers $1, 2, \dots, n$ in some order. We just need to show that there are exactly $n!$ different ways to order the numbers $1, 2, \dots, n$.

We prove this by induction on n . For the base case $n = 1$ we have $1! = 1$ and it is clear that there is only one way to order a single 1.

For the inductive step, suppose $|S_{n-1}| = (n-1)!$. An ordering of $1, 2, \dots, n$ arises in exactly one way as an ordering of $1, 2, \dots, (n-1)$ with the number n inserted into one of n places (the first, or second, or \dots , or n th position). So the number of such orderings is $|S_{n-1}|$ (the number of ways to choose an ordering of $1, 2, \dots, (n-1)$) times the number of ways to insert an n , giving $|S_n| = |S_{n-1}| \times n = (n-1)! \times n = n!$. This completes the inductive step. \square

For example, there are $2! = 2$ elements of S_2 : they are

$$\begin{pmatrix} 1 & 2 \\ 1 & 2 \end{pmatrix}, \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$$

The first one is the identity function on $\{1, 2\}$.

2.11 Inverses and composition

2.11.1 Inverse of a permutation

Permutations are bijections, so by Theorem 2.9.1 they have inverse functions. The inverse function to a permutation σ undoes what σ did, in the sense that if $\sigma(x) = y$ then $\sigma^{-1}(y) = x$. In two row notation you write $\sigma(x)$ beneath x , so you can get the two row notation for σ^{-1} by swapping the rows (and reordering).

Example 2.11.1.

$$\begin{aligned} \sigma &= \begin{pmatrix} 1 & 2 & 3 & 4 \\ 2 & 3 & 4 & 1 \end{pmatrix} \\ \sigma^{-1} &= \begin{pmatrix} 2 & 3 & 4 & 1 \\ 1 & 2 & 3 & 4 \end{pmatrix} = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 4 & 1 & 2 & 3 \end{pmatrix} \end{aligned}$$

2.11.2 Composition of permutations

We know by Theorem 2.9.2 that the composition of two bijections is a bijection, so the composition of two permutations of a set X is again a permutation of X .

Example 2.11.2. Let

$$\begin{aligned} \sigma &= \begin{pmatrix} 1 & 2 & 3 \\ 2 & 1 & 3 \end{pmatrix} \\ \tau &= \begin{pmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{pmatrix} \end{aligned}$$

Then $\sigma \circ \tau$ is the function $\{1, 2, 3\} \rightarrow \{1, 2, 3\}$ whose rule is “do τ , then do σ .” Thus

$$\begin{aligned} (\sigma \circ \tau)(1) &= \sigma(\tau(1)) = \sigma(1) = 2 \\ (\sigma \circ \tau)(2) &= \sigma(\tau(2)) = \sigma(3) = 3 \\ (\sigma \circ \tau)(3) &= \sigma(\tau(3)) = \sigma(2) = 1 \end{aligned}$$

In two row notation,

$$\sigma \circ \tau = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{pmatrix}.$$

There are several similarities between composing permutations and multiplying nonzero numbers. For example, if a , b , and c are nonzero real number then $a(bc) = (ab)c$. Furthermore the identity permutation behaves for composition just like the number 1 behaves for multiplication. For each nonzero real number a we have $a \times 1 = 1 \times a = a$, and for each permutation s we have $s \circ \text{id} = \text{id} \circ s = s$. Equally, for each nonzero real number a there is another nonzero real number a^{-1} such that $a \times a^{-1} = 1 = a^{-1} \times a$, and for each permutation s there is an inverse permutation s^{-1} such that $s \circ s^{-1} = \text{id} = s^{-1} \circ s$. Because of these similarities we often talk about multiplying two permutations when we mean composing them, and given two permutations s and t we usually write st for their composition instead of $s \circ t$.

2.11.3 Composition isn't commutative

Composition has one big difference with real number multiplication: the order matters.

Example 2.11.3. With σ and τ as before,

$$\sigma\tau = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{pmatrix}$$

$$\tau\sigma = \begin{pmatrix} 1 & 2 & 3 \\ 3 & 1 & 2 \end{pmatrix}$$

Comparing this to the example in the previous section, $\sigma\tau$ and $\tau\sigma$ are different. Composition of permutations is not commutative in general.

Definition 2.11.1. Two permutations s and t are said to **commute** if $st = ts$.

2.12 Cycles

2.12.1 Cycle definition and notation

We're going to introduce a more efficient way of writing permutations. This involves thinking about a special kind of permutation called a cycle.

Let $m > 0$, let a_1, \dots, a_m be distinct positive integers. Then

$$a = (a_1, \dots, a_m)$$

is defined to be the permutation such that

- $a(a_i) = a_{i+1}$ for $i < m$,
- $a(a_m) = a_1$, and
- $a(x) = x$ for any number x which isn't equal to one of the a_i .

Remember that a permutation is a special kind of function, so the notation $a(a_i)$ means 'apply the function a to the number a_i .' If we let a_{m+1} be a_1 then we could just say $a(a_i) = a_{i+1}$ for all $1 \leq i \leq m$.

Definition 2.12.1. A permutation of the form (a_1, \dots, a_m) is called an m -**cycle**. A permutation which is an m -cycle for some m is called a **cycle**.

There are two important things to note:

- Any 1-cycle, e.g. (1) or (2), is equal to the identity permutation.
- If we just write down the cycle (1, 2, 3), say, it could be an element of S_3 , or S_4 , or S_5 , or any other S_n with $n \geq 3$. When it matters, we will make this clear.

Example 2.12.1. • In S_3 , the 2-cycle (1, 2) is the permutation that sends 1 to 2, 2 to 1, and 3 to 3. In two row notation $(1, 2) = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 1 & 3 \end{pmatrix}$.

- In S_4 , the 3-cycle (2, 4, 3) is the permutation that sends 1 to 1, 2 to 4, 4 to 3, and 3 to 2. In two row notation, $(2, 4, 3) = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 1 & 4 & 2 & 3 \end{pmatrix}$.

The picture below is of the 5-cycle (1, 2, 3, 4, 5), illustrating why these permutations are called “cycles”.




Figure 2.14: Picture of the 5-cycle (1,2,3,4,5). The numbers 1, 2, 3, 4, 5 are arranged in a circle with an arrow pointing from 1 to 2, 2 to 3, 3 to 4, 4 to 5, 5 back to 1

2.12.2 Composing cycles

Let's compose two cycles. Let $s = (1, 2, 3, 4, 5)$, $t = (4, 3, 5, 1)$ be elements of S_5 . We'll work out the two row notation for $s \circ t$. Remember that this is the permutation whose rule is to do t then do s .

$t(1) = 4$, so $s(t(1)) = s(4) = 5$. Therefore the two row notation for $s \circ t$ looks like

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 5 & ? & ? & ? & ? \end{pmatrix}.$$

Next, $t(2) = 2$ (as 2 doesn't appear in the cycle defining t), so $s(t(2)) = s(2) = 3$. Now we know the next bit of the two row notation:

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 5 & 3 & ? & ? & ? \end{pmatrix}$$

You should continue this procedure and check that what you end up with is

$$s \circ t = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 5 & 3 & 1 & 4 & 2 \end{pmatrix}.$$

2.12.3 Multiple ways to write the same cycle

Example 2.12.2. Consider the two cycles $a = (1, 2, 3, 4)$ and $b = (2, 3, 4, 1)$. The permutation a sends 1 to 2, 2 to 3, 3 to 4, 4 to 1, and any other number to itself. So does b . So $a = b$. Similarly if $c = (3, 4, 1, 2)$ and $d = (4, 1, 2, 3)$ then $a = b = c = d$.

In general, every m -cycle can be written m different ways since you can put any one of the m things in the cycle first.

Example 2.12.3. In S_5 ,

$$(5, 3, 2) = (3, 2, 5) = (2, 5, 3).$$

2.12.4 Disjoint cycles

Definition 2.12.2. Two cycles (a_0, \dots, a_{m-1}) and (b_0, \dots, b_{k-1}) are **disjoint** if no a_i equals any b_j .

Example 2.12.4. • $(1, 2, 7)$ is disjoint from $(5, 4)$

- $(1, 2, 3)$ and $(3, 5)$ are **not** disjoint.

One reason disjoint cycles are important is that disjoint cycles commute, that is, if a and b are disjoint cycles then $a \circ b = b \circ a$. This is special as you have seen that in general, for two permutations s and t , $s \circ t \neq t \circ s$. You will prove this in the problem sets for MATH0005, but we'll record it here for future use.

Theorem 2.12.1. Let a and b be disjoint cycles. Then $ab = ba$.

2.12.5 Non-uniqueness

There can be many different ways to write a given permutation as a product of disjoint cycles. For example, letting $s = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 7 & 3 & 2 & 1 & 7 & 9 & 4 & 6 & 8 \end{pmatrix}$,

$$\begin{aligned} s &= (1, 7, 4)(2, 3)(5)(6, 9, 8) \\ &= (7, 4, 1)(2, 3)(6, 9, 8) \\ &= (2, 3)(6, 9, 8)(7, 4, 1) \\ &= \dots \end{aligned}$$

It is important to remember are that an m -cycle can be written in m different ways, for example $(1, 2, 3) = (2, 3, 1) = (3, 1, 2)$, and that disjoint cycles commute, for example $(1, 2)(3, 4) = (3, 4)(1, 2)$.

2.12.6 Inverse of a cycle

For every permutation s there is an inverse permutation s^{-1} such that $s \circ s^{-1} = s^{-1} \circ s = \text{id}$. How do we find the inverse of a cycle? Let

$$a = (a_1, \dots, a_m)$$

Then a sends a_i to a_{i+1} for all $1 \leq i \leq m$ (and every number not equal to an a_i to itself), so a^{-1} should send a_{i+1} to a_i for all $1 \leq i \leq m$ (and every number not equal to an a_i to itself). In other words, a^{-1} is the cycle $(a_m, a_{m-1}, \dots, a_2, a_1)$:

$$(a_1, \dots, a_m)^{-1} = (a_m, a_{m-1}, \dots, a_2, a_1)$$

As a special case, the inverse of the 2-cycle (i, j) is (j, i) . But $(i, j) = (j, i)$, so in fact every 2-cycle is its own inverse.

If we draw cycles as we did in Figure 2.14, their inverses are obtained by “reversing the arrows.”



Figure 2.15: On the left is a diagram showing the numbers 1, 2, and 3 in a cycle with an arrow from 1 to 2, 2 to 3, and 3 to 1 illustrating the 3-cycle $(1, 2, 3)$. On the right is its inverse $(3, 2, 1)$, the same picture with the arrows reversed

2.12.7 Not all permutations are cycles

Not every permutation is a cycle.

Example 2.12.5. The permutation $\sigma = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 2 & 1 & 4 & 3 \end{pmatrix}$ is not a cycle. Suppose for a contradiction that it was. σ sends 1 to 2, and 2 to 1, so if it were a cycle it would have to be $(1, 2)$. But $(1, 2)$ sends 3 to 3, whereas $\sigma(3) = 4$, so $\sigma \neq (1, 2)$.

Here is a diagram of the permutation σ from the previous example.




Figure 2.16: Picture of the permutation sending 1 to 2, 2 to 1, 3 to 4, and 4 to 3. Arrows indicate where each number is sent.

While σ is not a cycle, it *is* the composition of two cycles: $\sigma = (1, 2) \circ (3, 4)$. In fact, every permutation can be written this way, which we’ll prove in the next section.

2.13 Products of disjoint cycles

2.13.1 Every permutation is a product of disjoint cycles

Theorem 2.13.1. *Every $s \in S_n$ equals a product of disjoint cycles.*

Proof. Choose a number i and form the sequence $i, s(i), s^2(i), \dots$. Since these numbers belong to the set $\{1, 2, \dots, n\}$ not all of them can repeat, that is, there exist $p < q$ such that $s^p(i) = s^q(i)$, and so $s^{q-p}(i) = i$. Choose the smallest positive number r such that $s^r(i) = i$. Then all of the numbers

$$i, s(i), s^2(i), \dots, s^{r-1}(i) \quad (2.4)$$

are different, since if there were numbers $0 \leq p < q < r$ such that $s^p(i) = s^q(i)$ then we'd have $s^{q-p}(i) = i$ and $0 < q-p < r$ contradicting r being the smallest such power. Observe that, on the numbers (2.4), s behaves exactly like the r -cycle $(i, s(i), \dots, s^{r-1}(i))$.

If there is a number $j \in \{1, 2, \dots, n\}$ that doesn't belong to this cycle, we can repeat the procedure above to get another cycle containing j and such that s does the same thing as the cycle to its elements. This will be disjoint from the previous one since if $s^p(j) = s^q(i)$ for some p, q then $j = s^{q-p}(i)$ which belongs to the previous cycle, contradicting the choice of j .

By repeating this procedure until there are no longer any numbers between 1 and n not contained in one of our disjoint cycles; the product of these is s . \square

Definition 2.13.1. An expression for a permutation s as a product of disjoint cycles c_1, c_2, \dots, c_r

$$s = c_1 c_2 \cdots c_r$$

is called a **disjoint cycle decomposition** of s .

We've just proved that every permutation has at least one disjoint cycle decomposition. In fact a permutation can have lots of disjoint cycle decompositions, e.g.

$$(1, 2)(3, 4) = (3, 4)(1, 2) = (4, 3)(1, 2) = \dots$$

2.13.2 How to find a disjoint cycle decomposition

To find a disjoint cycle decomposition for an element of S_n :

1. Pick a number that doesn't yet appear in a cycle.
2. Compute its image, and the image of that, and so on, until you have a cycle. Write down that cycle.
3. If all elements of $1, \dots, n$ are in one of your cycles, stop, else go back to step 1.

Example 2.13.1. Let $s = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 7 & 6 & 3 & 1 & 2 & 5 & 4 \end{pmatrix}$. We pick a number not yet in a cycle, say 1. 1 goes to 7, 7 goes to 4, 4 goes to 1. We are back to the number we started with, so our first cycle is $(1, 7, 4)$. Now we pick another number not in a cycle, say 2. s sends 2 to 6, 6 to 5, and 5 to 2. That's another cycle, so we have $(1, 7, 4)(2, 6, 5)$. Now we pick another number not yet in a cycle — the only one left is 3. s sends 3 to 3, so this is immediately a cycle. We have $s = (1, 7, 4)(2, 6, 5)(3)$.

As we saw when we defined cycles in Definition 2.12.1, any 1-cycle is equal to the identity function. For that reason (and because 1-cycles look confusingly like what we write when we evaluate a function) we usually omit 1-cycles like (3) from disjoint cycle decompositions, so we'd write the permutation s of the previous example as $(1, 7, 4)(2, 6, 5)$.

2.13.3 Composing permutations given as products of disjoint cycles

Example 2.13.2. Let's work out a disjoint cycle decomposition for $\sigma\tau$ where

$$\sigma = (4, 2, 6, 1, 5)$$

$$\tau = (5, 4, 7, 3, 8)$$

are elements of S_8 .

Remember that $\sigma\tau$ means do τ , then do σ . Keeping that in mind, all we have to do is follow the instructions from before. Start with 1:

$$\sigma(\tau(1)) = \sigma(1) = 5$$

$$\sigma(\tau(5)) = \sigma(4) = 2$$

$$\sigma(\tau(2)) = \sigma(2) = 6$$

$$\sigma(\tau(6)) = \sigma(6) = 1$$

... and we have our first cycle, $(1, 5, 2, 6)$. Continuing with a number not yet in a cycle, say 3, we get

$$\sigma(\tau(3)) = \sigma(8) = 8$$

$$\sigma(\tau(8)) = \sigma(5) = 4$$

$$\sigma(\tau(4)) = \sigma(7) = 7$$

$$\sigma(\tau(7)) = \sigma(3) = 3$$

... and we have our next cycle, $(3, 8, 4, 7)$. There are no numbers left, so

$$\sigma\tau = (1, 5, 2, 6)(3, 8, 4, 7).$$

You should now find a disjoint cycle decomposition for $\tau\sigma$. You'll find that it has two 4-cycles again, but isn't the same as the decomposition we got for $\sigma\tau$.

Example 2.13.3. Let's find a disjoint cycle decomposition for $(1, 2, 3, 4)(2, 3, 4, 5)$. Write $a = (1, 2, 3, 4)$, $b = (2, 3, 4, 5)$. Starting with 1 as usual,

$$a(b(1)) = a(1) = 2$$

$$a(b(2)) = a(3) = 4$$

$$a(b(4)) = a(5) = 5$$

$$a(b(5)) = a(2) = 3$$

$$a(b(3)) = a(4) = 1$$

and so $ab = (1, 2, 4, 5, 3)$.

2.14 Powers and orders

2.14.1 Powers of a permutation

Since the composition of two permutations is another permutation, we can form powers of a permutation by composing it with itself some number of times.

Definition 2.14.1. Let s be a permutation and let m be an integer. Then

$$s^m = \begin{cases} s \cdots s \text{ (} m \text{ times)} & m > 0 \\ \text{id} & m = 0 \\ s^{-1} \cdots s^{-1} \text{ (} -m \text{ times)} & m < 0 \end{cases}$$

It's tedious but straightforward to check that for any integers a, b ,

- $s^a \circ s^b = s^{a+b}$, and
- $(s^a)^b = s^{ab}$

so that some of the usual exponent laws for real numbers hold for composing permutations. The two facts above are called the **exponent laws** for permutations.

2.14.2 Order of a permutation

Definition 2.14.2. The **order** of a permutation σ , written $o(\sigma)$, is the smallest strictly positive number n such that $\sigma^n = \text{id}$.

For example, let

$$s = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{pmatrix}$$

$$t = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 1 & 3 \end{pmatrix}$$

You should check that $s^2 \neq \text{id}$ but $s^3 = \text{id}$, so the order of s is 3, and that $t \neq \text{id}$ but $t^2 = \text{id}$ so the order of t is 2.

2.14.3 Order of an m -cycle

Lemma 2.14.1. *The order of an m -cycle is m .*

Proof. Let the m -cycle be $a = (a_1, \dots, a_m)$. If $r < m$ then $a^r(a_1) = a_{r+1} \neq a_1$, so $a^r \neq \text{id}$. On the other hand $a^m(a_1) = a(a_m) = a_1$ and in general $a^m(a_i) = a^i(a^{m-i}(a_i)) = a^i(a_1) = a_i$, so $a^m = \text{id}$. \square

2.15 Transpositions

2.15.1 Definition of a transposition

Definition 2.15.1. A **transposition** is a 2-cycle.

For example, the only transpositions in S_3 are $(1, 2)$, $(2, 3)$, and $(1, 3)$.

2.15.2 Every permutation is a product of transpositions

In this section we're going to prove that every permutation can be written as a product of transpositions. Before we do so, here is some motivation for why we should expect this to be true.

One way we've already seen of illustrating a permutation s in S_n is to draw the numbers $1, 2, \dots, n$ in a column, and then in another column to its right, and draw a line from each number i in the first column to $s(i)$ in the second. You get what looks like a lot of pieces of string tangled together.

Here is a diagram of the permutation $s = (1, 2, 3, 4, 5)$. Think of the lines as pieces of string connecting i on the left with $s(i)$ on the right.




Figure 2.17: A string diagram of the permutation $(1, 2, 3, 4, 5)$. Two columns contain the numbers $1, 2, 3, 4, 5$. Strings connect the numbers $1, 2, 3, 4, 5$ in the left-hand column to $2, 3, 4, 5, 1$ respectively in the right-hand column.

Composing two permutations drawn in this way corresponds to placing their diagrams side-by-side:




Figure 2.18: String diagrams for $(1, 2)$ and $(2, 3)$ are shown side by side, and then the right-hand column of the diagram for $(1, 2)$ is joined to the left-hand column for $(2, 3)$. The resulting diagram represents $(2, 3)(1, 2) = (1, 3, 2)$

Imagine taking such a diagram and stretching it out. You could divide it up into smaller diagrams, each of which contains only one crossing of strings.




Figure 2.19: String diagram for $(1, 2, 3, 4)$, with dotted vertical lines to divide the strings into sections with only one crossing. From left to right, the first crossing is between strings 3 and 4, then 2 and 3, then 1 and 2

A diagram with a single string crossing is a transposition, since only two

numbers change place. The diagram above illustrates the fact that

$$(1, 2, 3, 4) = (1, 2)(2, 3)(3, 4).$$

Now we're ready for a formal proof of the result that every permutation equals a product of transpositions. We first do it for cycles, for which we need a lemma.

Lemma 2.15.1. *Let $a_1, a_2, \dots, a_m, a_{m+1}$ be distinct numbers. Then*

$$(a_1, a_2)(a_2, a_3, \dots, a_{m+1}) = (a_1, a_2, \dots, a_{m+1}).$$

For example, you should check by calculating the two row notation for both sides that

$$\begin{aligned} (1, 2)(2, 3) &= (1, 2, 3) \\ (2, 3)(3, 1, 4, 5) &= (2, 3, 1, 4, 5) \\ (7, 2)(2, 8, 9, 6, 4) &= (7, 2, 8, 9, 6, 4). \end{aligned}$$

Proof. Let

$$\begin{aligned} r &= (a_1, a_2, \dots, a_{m+1}) \\ t &= (a_1, a_2) \\ s &= (a_2, a_3, \dots, a_{m+1}) \end{aligned}$$

so that we have to show $r(s) = t(s(x))$ for all integers x . If x is not one of the a_i then $r(x) = x$, $s(x) = x$, and $t(x) = x$ so $r(x) = t(s(x)) = x$. We only need to do the case when x is equal to a_i for some $1 \leq i \leq m+1$.

- Let $i = 1$. Then $r(a_1) = a_2$, and $s(a_1) = a_1$ and $t(a_1) = a_2$ so $t(s(a_1)) = a_2 = r(a_1)$.
- Let $i = 2$, so $r(a_2) = a_3$. We have $s(a_2) = a_3$ and $t(a_3) = a_3$, so $t(s(a_2)) = a_3 = r(a_2)$.
- Let $3 \leq i < m+1$, so, $r(a_i) = a_{i+1}$, $s(a_i) = a_{i+1}$, $t(a_{i+1}) = a_{i+1}$, so $t(s(a_i)) = t(a_{i+1}) = a_{i+1} = r(a_i)$.
- Finally $r(a_{m+1}) = a_1$, $s(a_{m+1}) = a_2$, $t(a_2) = a_1$, so $t(s(a_{m+1})) = a_1 = r(a_{m+1})$. \square

Lemma 2.15.2. *Every cycle equals a product of transpositions.*

Proof. Let $a = (a_1, \dots, a_m)$ be a cycle. Lemma 2.15.1 tells us that

$$a = (a_1, a_2)(a_2, a_3, \dots, a_m).$$

Using that lemma again on the cycle (a_2, \dots, a_m) we get

$$a = (a_1, a_2)(a_2, a_3)(a_3, a_4, \dots, a_m).$$

Repeating this gives

$$a = (a_1, a_2)(a_2, a_3) \cdots (a_{m-1}, a_m)$$

which shows that a can be written as a product of transpositions. \square

To illustrate this result you should check by computing both sides that

$$(1, 2, 3, 4) = (1, 2)(2, 3)(3, 4).$$

Theorem 2.15.3. *Every permutation in S_n is equal to a product of transpositions.*

Proof. Let p be a permutation. We have seen that every permutation can be written as a product of cycles, so there are cycles c_1, \dots, c_k such that $p = c_1 \cdots c_k$. The lemma above shows how to write each c_i as a product of transpositions, which expresses p as a product of transpositions too. \square

Example 2.15.1. Suppose we want to express

$$s = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 2 & 3 & 1 & 5 & 6 & 4 \end{pmatrix}$$

as a product of transpositions. A disjoint cycle decomposition for s is

$$s = (1, 2, 3)(4, 5, 6)$$

and applying the lemma above, we get

$$\begin{aligned} (1, 2, 3) &= (1, 2)(2, 3) \\ (4, 5, 6) &= (4, 5)(5, 6) \end{aligned}$$

So

$$s = (1, 2, 3)(4, 5, 6) = (1, 2)(2, 3)(4, 5)(5, 6).$$

2.16 Sign

2.16.1 Definition of odd and even permutations

Theorem 2.15.3 says that every permutation can be expressed as a product of transpositions.

Definition 2.16.1. A permutation is **odd** if it can be expressed as a product of an odd number of transpositions and **even** if it can be expressed as a product of an even number of transpositions.

(Sometimes people refer to the *parity* of a permutation to mean whether it is odd or even.)

Example 2.16.1. • $(1, 2)$ is odd.

- $\text{id} = (1, 2)(1, 2)$ is even.
- $(1, 2, 3) = (1, 2)(2, 3)$ is even.
- The expression for an m -cycle $a = (a_1, \dots, a_m)$ as a product of $m - 1$ transpositions

$$(a_1, \dots, a_m) = (a_1, a_2)(a_2, a_3) \cdots (a_{m-1}, a_m)$$

in Lemma 2.15.2 shows that an m cycle is even if m is odd and odd if m is even.

2.16.2 The odd xor even theorem

It seems possible that a permutation could be odd AND even at the same time, but this isn't the case. The proof is a little strange because it involves a seemingly-unrelated polynomial — if you prefer a proof that refers only to permutations, take a look at the unassessed exercises.

Theorem 2.16.1. *No permutation is both odd and even.*

Proof. Let $s \in S_n$. If s could be written both as a product of an even number and an odd number of transpositions, say

$$s = t_1 t_2 \cdots t_{2a} = t'_1 t'_2 \cdots t'_{2b+1}$$

then we could rearrange to get

$$\text{id} = t_{2a} t_{2a-1} \cdots t_2 t_1 t'_1 t'_2 \cdots t'_{2b+1}$$

which expresses the identity permutation as a product of an odd number $2(a+b)+1$ of transpositions. So it's enough to show that this is impossible.

Introduce variables x_1, x_2, \dots, x_n . Given any polynomial f in these variables and any permutation $s \in S_n$, we define a new polynomial $s(f)$ which is obtained by replacing every variable x_i in f by $x_{s(i)}$. For example, if $s = (1, 2, 3)$ and $f = x_1 x_2 + x_3^4 + x_5$ then

$$\begin{aligned} s(f) &= x_{s(1)} x_{s(2)} + x_{s(3)}^4 + x_{s(5)} \\ &= x_2 x_3 + x_1^4 + x_5. \end{aligned}$$

Now define

$$\Delta_n = \prod_{1 \leq i < j \leq n} (x_i - x_j)$$

so that, for example,

$$\Delta_3 = (x_1 - x_2)(x_1 - x_3)(x_2 - x_3).$$

Let $t = (a, b)$ be any transposition, where $a < b$. We have $t(\Delta_n) = -\Delta_n$, because t only affects factors of Δ_n containing x_a or x_b , and the only factors to change sign are $(x_a - x_i)$ with $a < i \leq b$ and $(x_i - x_b)$ with $a \leq i < b$, a total of $2(b-a) - 1$ sign changes (as $(x_a - x_b)$ appears in both cases). It follows that any product of an odd number of transpositions sends Δ_n to $-\Delta_n$. Since the identity sends Δ_n to Δ_n , it cannot be written as the product of an odd number of transpositions. \square

2.16.3 Sign of a permutation

Definition 2.16.2. • The **sign** of a permutation is 1 if it is even and -1 if it is odd.

- We write $\text{sign}(s)$ for the sign of the permutation s .

So if s can be written as a product of m transpositions, $\text{sign}(s) = (-1)^m$.

Lemma 2.16.2. *For any two permutations s and t , $\text{sign}(st) = \text{sign}(s) \text{sign}(t)$*

Proof. If s can be written as a product of m transpositions and t can be written as a product of n transpositions, then st can be written as a product of $m + n$ transpositions. So

$$\begin{aligned}\text{sign}(st) &= (-1)^{n+m} \\ &= (-1)^m (-1)^n \\ &= \text{sign}(s) \text{sign}(t)\end{aligned}$$

□

This rule about the sign of a product means that

- an even permutation times an even permutation is even,
- an even permutation times an odd permutation is odd, and
- an odd permutation times an odd permutation is even

just like when we multiply odd and even integers.

2.16.4 Two results on the sign of a permutation

Lemma 2.16.3. 1. *Even length cycles are odd and odd length cycles are even.*

2. *If s is any permutation, $\text{sign}(s) = \text{sign}(s^{-1})$.*

Proof. 1. We saw in the proof of Lemma 2.15.2 that if $a = (a_1, \dots, a_m)$ is any m -cycle,

$$(a_1 \dots, a_m) = (a_1, a_2)(a_2, a_3) \cdots (a_{m-1}, a_m)$$

so an m -cycle can be written as a product of $m - 1$ transpositions. The number of transpositions in this expression therefore has the opposite parity to m , as required.

2. If $s = t_1 \cdots t_m$ is a product of m transpositions, $s^{-1} = t_m^{-1} \cdots t_1^{-1}$. But the inverse of a transposition is a transposition, so s^{-1} is also the product of m transpositions.

□

Another way to express the first part of this lemma would be to say that $\text{sign}(a_1, \dots, a_m) = (-1)^{m-1}$.

Further reading

You don't need to read any of these for the purposes MATH0005, but if you want to learn more about the topics covered here are my recommendations.

Set theory

The third year course MATH0037 Logic contains some material on set theory. If you want to learn about formal (ZFC) set theory and can't wait for MATH0037, *Classic Set Theory* by Derek Goldrei is a great introduction. It was written as an Open University textbook so is designed for self-study. *Naive Set Theory* by Paul Halmos gives an idea of what formal set theory is all about without getting into all of the axiomatic details.

The problems with unrestricted set comprehension mentioned briefly in the text are explained nicely in the Stanford Encyclopedia of Philosophy entry for Russell's Paradox, but you can find hundreds of other examples with an internet search. This short pdf by the philosopher Richard Pettigrew gives a short sketch of what goes wrong and how it is fixed formally.

Permutations

Most basic algebra textbooks go into more detail on permutations than we do in 0005. I like *A Concise Introduction to Pure Mathematics* by Martin Liebeck a lot, and it has a nice application of the sign of a permutation to (not) solving the 15 puzzle. *Topics in Algebra* by I. Herstein is not always the easiest text but contains loads of interesting material if algebra is your thing, some of which is covered in MATH0006 Algebra 2. C. Pinter's *Book of Abstract Algebra* is published by Dover so is cheap even if you want a hard copy, and covers permutations in chapters 7 and 8. It's especially worthwhile if you want to learn more abstract algebra.

Chapter 3

Matrices

3.1 Matrix definitions

We begin with a lot of definitions.

Definition 3.1.1. • A $m \times n$ **matrix** is a rectangular grid of numbers with m rows and n columns.

- A **square** matrix is one which is $n \times n$ for some n .
- A (height m) **column vector** is an $m \times 1$ matrix.
- A (width n) **row vector** is a $1 \times n$ matrix.
- \mathbb{R}^n is the set of all column vectors with height n and real numbers as entries, \mathbb{C}^n is the set of all height n column vectors with complex numbers as entries.
- $M_{m \times n}(\mathbb{R})$ is the set of all $m \times n$ matrices with real number entries.
- The $m \times n$ **zero matrix**, written $0_{m \times n}$, is the $m \times n$ matrix all of whose entries are zero.

Example 3.1.1. • $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is a 2×1 column vector, an element of \mathbb{R}^2 .

- $\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix}$ is a 2×3 matrix
- $(-1 \quad -2)$ is a 1×2 row vector
- $\begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$ is a 2×2 square matrix.
- $0_{2 \times 2} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$.

3.1.1 Matrix entries

The i, j entry of a matrix means the number in row i and column j . It is important to get these the correct way round. Usually when you give (x, y) coordinates, x refers to the horizontal direction and y refers to the vertical direction. When we talk about the i, j entry of a matrix, however, the first number i refers to the row number (i.e. the vertical direction) and the second number j refers to the column number (i.e. the horizontal direction).

We often write $A = (a_{ij})$ to mean that A is the matrix whose i, j entry is called a_{ij} . For example, if A is 2×2 then saying $A = (a_{ij})$ means that

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}.$$

If you're using this notation you must also specify the size of the matrix, of course.

We often talk about the columns and rows of a matrix. If A is an $m \times n$ matrix

$$A = \begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{pmatrix}$$

then the i th row of A means the $1 \times n$ row vector

$$(a_{i1} \quad a_{i2} \quad \cdots \quad a_{in})$$

and the j th column of A is the $m \times 1$ column vector

$$\begin{pmatrix} a_{1j} \\ a_{2j} \\ \vdots \\ a_{mj} \end{pmatrix}.$$

For example, if

$$A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$$

then the first row is $(1 \quad 2)$ and the second column is $\begin{pmatrix} 2 \\ 4 \end{pmatrix}$.

3.1.2 Matrix addition and scalar multiplication

We can add matrices of the same size. If $A = (a_{ij})$ and $B = (b_{ij})$ are the same size, then $A + B$ is defined to be the matrix whose i, j entry is $a_{ij} + b_{ij}$.

Example 3.1.2.

$$\begin{pmatrix} 1 & 2 \\ 4 & 5 \end{pmatrix} + \begin{pmatrix} 0 & 1 \\ 2 & 3 \end{pmatrix} = \begin{pmatrix} 1+0 & 2+1 \\ 4+2 & 5+3 \end{pmatrix} = \begin{pmatrix} 1 & 3 \\ 6 & 8 \end{pmatrix}.$$

In other words, we add matrices by adding corresponding entries. We never add matrices of different sizes.

We also multiply matrices by numbers. This is called **scalar multiplication**. If $A = (a_{ij})$ is a matrix and λ a number then λA means the matrix obtained by multiplying every entry in A by λ , so the i, j entry of λA is λa_{ij} .

Example 3.1.3.

$$2 \begin{pmatrix} 1 & -3 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 2 & -6 \\ 0 & 2 \end{pmatrix}.$$

3.1.3 Laws for addition and scalar multiplication

These operations have some familiar properties.

Theorem 3.1.1. *If a and b are numbers and $A, B,$ and C are matrices of the same size,*

1. $A + B = B + A$ (*commutativity*)
2. $A + (B + C) = (A + B) + C$ (*associativity*)
3. $(a + b)A = aA + bA$ (*distributivity*),
4. $a(A + B) = aA + aB$ (*distributivity*), and
5. $a(bA) = (ab)A$. □

These can be proved using the usual laws for addition and multiplication of numbers.

3.2 Matrix multiplication

We are going to define a way to multiply certain matrices together. After that we will see several different ways to understand this definition, and we will see how the definition arises as a kind of function composition.

Definition 3.2.1. Let $A = (a_{ij})$ be a $m \times n$ matrix and $B = (b_{ij})$ be an $n \times p$ matrix. Then the matrix product AB is defined to be the $m \times p$ matrix whose i, j entry is

$$\sum_{k=1}^n a_{ik}b_{kj}. \quad (3.1)$$

Before we even start thinking about this definition we record one key point about it. There are two n s in the definition above: one is the number of columns of A and the other is the number of rows of B . These really must be the same. We **only** define the matrix product AB when the number of columns of A equals the number of rows of B . The reason for this will become clear when we interpret matrix multiplication in terms of function composition later.

Example 3.2.1. The 1, 2 entry of a matrix product AB is obtained by putting $i = 1$ and $j = 2$ in the formula (3.1). If $A = (a_{ij})$ is $m \times n$ and $B = (b_{ij})$ is $n \times p$ then this is

$$a_{11}b_{12} + a_{12}b_{22} + a_{13}b_{32} + \cdots + a_{1n}b_{n2}$$

You can see that we are multiplying each entry in the first row of A by the corresponding entry in the second column of B and adding up the results. In general, the i, j entry of AB is obtained by multiplying the entries of row i of A with the entries of column j of B and adding them up.

Example 3.2.2. Let's look at a generic example first. Let

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}, B = \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix}.$$

The number of columns of A equals the number of rows of B , so the matrix product AB is defined, and since (in the notation of the definition) $m = n = p = 2$, the size of AB is $m \times p$ which is 2×2 . From the formula, we get

$$AB = \begin{pmatrix} a_{11}b_{11} + a_{12}b_{21} & a_{11}b_{12} + a_{12}b_{22} \\ a_{21}b_{11} + a_{22}b_{21} & a_{21}b_{12} + a_{22}b_{22} \end{pmatrix}.$$

Example 3.2.3. Making the previous example concrete, if

$$A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}, B = \begin{pmatrix} 5 & 6 \\ 7 & 8 \end{pmatrix}.$$

then A is 2×2 , B is 2×2 , so the matrix product AB is defined and will be another 2×2 matrix:

$$\begin{aligned} AB &= \begin{pmatrix} 1 \times 5 + 2 \times 7 & 1 \times 6 + 2 \times 8 \\ 3 \times 5 + 4 \times 7 & 3 \times 6 + 4 \times 8 \end{pmatrix} \\ &= \begin{pmatrix} 19 & 22 \\ 43 & 50 \end{pmatrix}. \end{aligned}$$

Matrix multiplication is so important that it is helpful to have several different ways of looking at it. The formula above is useful when we want to prove general properties of matrix multiplication, but we can get further insight when we examine the definition carefully from different points of view.

3.2.1 Matrix multiplication happens columnwise

A very important special case of matrix multiplication is when we multiply a $m \times n$ matrix by an $n \times 1$ column vector. Let

$$A = \begin{pmatrix} a & b & c \\ d & e & f \end{pmatrix}, \mathbf{x} = \begin{pmatrix} x \\ y \\ z \end{pmatrix}.$$

Then we have

$$A\mathbf{x} = \begin{pmatrix} ax + by + cz \\ dx + ey + fz \end{pmatrix}$$

Another way to write the result of this matrix multiplication is

$$x \begin{pmatrix} a \\ d \end{pmatrix} + y \begin{pmatrix} b \\ e \end{pmatrix} + z \begin{pmatrix} c \\ f \end{pmatrix}$$

showing that the result is obtained by adding up scalar multiples of the columns of A . If we write \mathbf{c}_j for the j th column of A then the expression

$$x\mathbf{c}_1 + y\mathbf{c}_2 + z\mathbf{c}_3,$$

where we add up scalar multiples of the \mathbf{c}_j s, is called a **linear combination** of \mathbf{c}_1 , \mathbf{c}_2 , and \mathbf{c}_3 . Linear combinations are a fundamental idea and we will return to them again and again in the rest of MATH0005.

Definition 3.2.2. Let $\mathbf{v}_1, \mathbf{v}_2, \dots$ be matrices all of the same shape. A **linear combination** of $\mathbf{v}_1, \mathbf{v}_2, \dots$ is a matrix of the form

$$a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + \dots$$

where the a_i are numbers.

This result is true whenever we multiply an $m \times n$ matrix and an $n \times 1$ column vector, not just in the example above.

Proposition 3.2.1. Let $A = (a_{ij})$ be an $m \times n$ matrix and \mathbf{x} an $n \times 1$ column vector with entries x_1, \dots, x_n . If $\mathbf{c}_1, \dots, \mathbf{c}_n$ are the columns of A then

$$A\mathbf{x} = \sum_{k=1}^n x_k \mathbf{c}_k.$$

Proof. From the matrix multiplication formula (3.1) we get

$$A\mathbf{x} = \begin{pmatrix} \sum_{k=1}^n a_{1k}x_k \\ \sum_{k=1}^n a_{2k}x_k \\ \vdots \\ \sum_{k=1}^n a_{mk}x_k \end{pmatrix} = \sum_{k=1}^n x_k \begin{pmatrix} a_{1k} \\ a_{2k} \\ \vdots \\ a_{mk} \end{pmatrix}$$

The column vector whose entries are $a_{1k}, a_{2k}, \dots, a_{mk}$ is exactly the k th column of A , so this completes the proof. \square

Definition 3.2.3. For a fixed n , the **standard basis vectors** $\mathbf{e}_1, \dots, \mathbf{e}_n$ are the vectors

$$\begin{pmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \dots, \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix}.$$

The vector \mathbf{e}_i with a 1 in position i and zeroes elsewhere is called the i th standard basis vector.

For example, if $n = 3$ then there are three standard basis vectors

$$\mathbf{e}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \mathbf{e}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \mathbf{e}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

The special case of the proposition above when we multiply a matrix by a standard basis vector is often useful, so we'll record it here.

Corollary 3.2.2. Let A be a $m \times n$ matrix and \mathbf{e}_j the j th standard basis vector of height n . Then $A\mathbf{e}_j$ is equal to the j th column of A .

Proof. According to Proposition 3.2.1 we have $A\mathbf{e}_j = \sum_{k=1}^n x_k \mathbf{c}_k$ where x_k is the k th entry of \mathbf{e}_j and \mathbf{c}_k is the k th column of A . The entries of \mathbf{e}_j are all zero except for the j th which is 1, so

$$A\mathbf{e}_j = 0 \times \mathbf{c}_1 + \cdots + 1 \times \mathbf{c}_j + \cdots + 0 \times \mathbf{c}_n = \mathbf{c}_j. \quad \square$$

Example 3.2.4. Let $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$. You should verify that $A \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ equals the first column of A and $A \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ equals the second column of A .

The next theorem tells us that we can do any matrix multiplication AB column-by-column, multiplying A into each of the columns of B in turn.

Theorem 3.2.3. Let A be an $m \times n$ matrix and B an $n \times p$ matrix with columns $\mathbf{d}_1, \dots, \mathbf{d}_p$. Then

$$AB = \left(\begin{array}{c|ccc} | & \cdots & | \\ A\mathbf{d}_1 & \cdots & A\mathbf{d}_p \\ | & \cdots & | \end{array} \right).$$

The notation means that the first column of AB is equal to what you get by multiplying A into the first column of B , the second column of AB is what you get by multiplying A into the second column of B , and so on. That's what it means to say that matrix multiplication works *columnwise*.

Proof. From the matrix multiplication formula (3.1) the j th column of AB has entries

$$\begin{pmatrix} \sum_{k=1}^n a_{1k} b_{kj} \\ \sum_{k=1}^n a_{2k} b_{kj} \\ \vdots \\ \sum_{k=1}^n a_{mk} b_{kj} \end{pmatrix} \quad (3.2)$$

The entries b_{kj} for $k = 1, 2, \dots, n$ are exactly the entries in column j of B , so (3.2) is $A\mathbf{d}_j$ as claimed. \square

Corollary 3.2.4. Every column of AB is a linear combination of the columns of A .

Proof. Theorem 3.2.3 tells us that each column of AB equals $A\mathbf{d}$ for certain vectors \mathbf{d} , and Proposition 3.2.1 tells us that any such vector $A\mathbf{d}$ is a linear combination of the columns of A . \square

Example 3.2.5. Let's look at how Proposition 3.2.1 and Theorem 3.2.3 apply to Example 3.2.3, when A was $\begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$ and the columns of B are $\mathbf{d}_1 = \begin{pmatrix} 5 \\ 7 \end{pmatrix}$ and $\mathbf{d}_2 = \begin{pmatrix} 6 \\ 8 \end{pmatrix}$.

You can check that

$$\begin{aligned} A\mathbf{d}_1 &= \begin{pmatrix} 19 \\ 43 \end{pmatrix} \\ &= 5 \begin{pmatrix} 1 \\ 3 \end{pmatrix} + 7 \begin{pmatrix} 2 \\ 4 \end{pmatrix} \\ A\mathbf{d}_2 &= \begin{pmatrix} 22 \\ 50 \end{pmatrix} \\ &= 6 \begin{pmatrix} 1 \\ 3 \end{pmatrix} + 8 \begin{pmatrix} 2 \\ 4 \end{pmatrix} \end{aligned}$$

and that these are the columns of AB we computed before.

3.2.2 Matrix multiplication happens rowwise

There are analogous results when we multiply an $1 \times n$ row vector and an $n \times p$ matrix.

Proposition 3.2.5. *Let \mathbf{a} be a $1 \times n$ row vector with entries a_1, \dots, a_n and let B be an $n \times p$ matrix with rows $\mathbf{s}_1, \dots, \mathbf{s}_n$. Then $\mathbf{a}B = \sum_{k=1}^n a_k \mathbf{s}_k$.*

Proof. From the matrix multiplication formula (3.1) we get

$$\begin{aligned} \mathbf{a}B &= \left(\sum_{k=1}^n a_k b_{k1} \quad \cdots \quad \sum_{k=1}^n a_k b_{kp} \right) \\ &= \sum_{k=1}^n a_k (b_{k1} \quad \cdots \quad b_{kp}) \\ &= \sum_{k=1}^n a_k \mathbf{s}_k. \quad \square \end{aligned}$$

In particular, $\mathbf{a}B$ is a linear combination of the rows of B .

Theorem 3.2.6. *Let A be a $m \times n$ matrix with rows $\mathbf{r}_1, \dots, \mathbf{r}_m$ and let B be an $n \times p$ matrix. Then*

$$AB = \begin{pmatrix} - & \mathbf{r}_1 B & - \\ \cdots & \cdots & \cdots \\ - & \mathbf{r}_m B & - \end{pmatrix}$$

The notation means that the first row of AB is equal to $\mathbf{r}_1 B$, the second row is equal to $\mathbf{r}_2 B$, and so on.

Proof. From the matrix multiplication formula (3.1), the i th row of AB has entries

$$\begin{aligned} &\left(\sum_{k=1}^n a_{ik} b_{k1} \quad \cdots \quad \sum_{k=1}^n a_{ik} b_{kp} \right) \\ &= \sum_{k=1}^n a_{ik} (b_{k1} \quad \cdots \quad b_{kp}). \end{aligned} \quad (3.3)$$

Row i of A is $\mathbf{r}_i = (a_{i1} \quad a_{i2} \quad \cdots \quad a_{in})$, so $\mathbf{r}_i B$ agrees with (3.3) by Proposition 3.2.5. \square

The theorem combined with Proposition 3.2.5 show that the rows of AB are linear combinations of the rows of B .

Example 3.2.6. Returning to the example where

$$A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}, B = \begin{pmatrix} 5 & 6 \\ 7 & 8 \end{pmatrix}$$

the rows of A are $\mathbf{r}_1 = (1 \ 2)$ and $\mathbf{r}_2 = (3 \ 4)$ and the rows of B are $\mathbf{s}_1 = (5 \ 6)$ and $\mathbf{s}_2 = (7 \ 8)$. We have

$$\begin{aligned} \mathbf{r}_1 B &= (1 \ 2) \begin{pmatrix} 5 & 6 \\ 7 & 8 \end{pmatrix} \\ &= \mathbf{s}_1 + 2\mathbf{s}_2 \\ &= (19 \ 22) \\ \mathbf{r}_2 B &= (3 \ 4) \begin{pmatrix} 5 & 6 \\ 7 & 8 \end{pmatrix} \\ &= 3\mathbf{s}_1 + 4\mathbf{s}_2 \\ &= (43 \ 50). \end{aligned}$$

and these are the rows of the matrix product AB .

3.2.3 Row times column

The matrix multiplication formula says that the i, j entry of AB is

$$\sum_{k=1}^n a_{ik} b_{kj}.$$

Let's think about where the entries in this sum come from. The entries from A involved are $a_{i,1}, a_{i,2}, \dots, a_{i,n}$. These are exactly the entries in the i th row of A .

The entries from B in this sum are $b_{1,j}, b_{2,j}, \dots, b_{n,j}$. These are the entries from the j th column of B . So to get the i, j entry of AB , we matrix multiply the i th row of A by the j th column of B . That is, if the i th row of A is \mathbf{r}_i and the j th column of B is \mathbf{c}_j , then

$$AB = (\mathbf{r}_i \mathbf{c}_j). \quad (3.4)$$

(If you're thinking 'wait, isn't $\mathbf{r}_i \mathbf{c}_j$ a 1×1 matrix, not a number?' then you are correct. We will identify the 1×1 matrix (x) with the number x .)

As an example, consider

$$A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}, B = \begin{pmatrix} 2 & 1 & 2 \\ 3 & 2 & 3 \\ 4 & 3 & 4 \end{pmatrix}$$

Multiplying the second row \mathbf{r}_2 of A into the third column \mathbf{c}_3 of B gives

$$(4 \ 5 \ 6) \begin{pmatrix} 2 \\ 3 \\ 4 \end{pmatrix} = 8 + 15 + 24 = 47$$

which is the 2,3 entry of

$$AB = \begin{pmatrix} 20 & 14 & 20 \\ 47 & 32 & 47 \\ 74 & 50 & 74 \end{pmatrix}.$$

Example 3.2.7. When the result of a matrix multiplication is a 1×1 matrix we will think of it as a number. This is like a dot product, if you've seen those before.

$$(1 \ 2 \ 3) \begin{pmatrix} 4 \\ 5 \\ 6 \end{pmatrix} = 1 \times 4 + 2 \times 5 + 3 \times 6 = 32.$$

Example 3.2.8. Let $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{pmatrix}$, a 3×2 matrix, and $\mathbf{c} = \begin{pmatrix} 7 \\ 8 \end{pmatrix}$, a 2×1 column vector. The number of columns of A and the number of rows of \mathbf{c} are equal, so we can compute $A\mathbf{c}$.

$$A\mathbf{c} = \begin{pmatrix} 1 \times 7 + 2 \times 8 \\ 3 \times 7 + 4 \times 8 \\ 5 \times 7 + 6 \times 8 \end{pmatrix}.$$

Example 3.2.9. Let

$$A = (1 \ 2), B = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}.$$

A is 1×2 , B is 2×3 , so the matrix product AB is defined, and is a 1×3 matrix. The columns of B are $\mathbf{c}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, $\mathbf{c}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, and $\mathbf{c}_3 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$. The product AB is therefore

$$\begin{aligned} (A\mathbf{c}_1 \ A\mathbf{c}_2 \ A\mathbf{c}_3) &= (1 \times 1 + 2 \times 0 \quad 1 \times 0 + 2 \times 1 \quad 1 \times 1 + 2 \times 0) \\ &= (1 \ 2 \ 1) \end{aligned}$$

Example 3.2.10. Let

$$A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}, B = \begin{pmatrix} 5 & 6 \\ 7 & 8 \end{pmatrix}.$$

Then A is 2×2 , B is 2×2 , so the matrix product AB is defined and will be another 2×2 matrix:

$$AB = \begin{pmatrix} 1 \times 5 + 2 \times 7 & 1 \times 6 + 2 \times 8 \\ 3 \times 5 + 4 \times 7 & 3 \times 6 + 4 \times 8 \end{pmatrix}.$$

3.2.4 Matrix multiplication motivation

In this section we'll try to answer two questions: where does this strange-looking notion of matrix multiplication come from? Why can we only multiply A and B if the number of columns of A equals the number of rows of B ?

Definition 3.2.4. Let A be a $m \times n$ matrix. Then $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is the function defined by

$$T_A(\mathbf{x}) = A\mathbf{x}.$$

Notice that this definition really does make sense. If $\mathbf{x} \in \mathbb{R}^n$ then it is an $n \times 1$ column vector, so the matrix product $A\mathbf{x}$ exists and has size $m \times 1$, so it is an element of \mathbb{R}^m .

Now suppose we have an $m \times n$ matrix A and a $q \times p$ matrix B , so that $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $T_B : \mathbb{R}^p \rightarrow \mathbb{R}^q$. Can we form the composition $T_A \circ T_B$? The answer is no, unless $q = n$, that is, unless the number of columns of A equals the number of rows of B . So let's assume that $q = n$ so that B is $n \times p$ and the composition

$$T_A \circ T_B : \mathbb{R}^p \rightarrow \mathbb{R}^m$$

makes sense. What can we say about it?

Theorem 3.2.7. *If A is $m \times n$ and B is $n \times p$ then $T_A \circ T_B = T_{AB}$.*

You will prove this on a problem sheet.

The theorem shows that matrix multiplication is related to composition of functions. That's useful because it suggests something: we know that function composition is always associative, so can we use that to show matrix multiplication is associative too? That is, if the products AB and BC make sense, is $A(BC)$ equal to $(AB)C$? This is not exactly obvious if you just write down the horrible formulas for the i, j entries of both matrices. If we believe the theorem though it's easy: we know

$$T_A \circ (T_B \circ T_C) = (T_A \circ T_B) \circ T_C$$

because function composition is associative, and so

$$\begin{aligned} T_A \circ T_{BC} &= T_{AB} \circ T_C \\ T_{A(BC)} &= T_{(AB)C}. \end{aligned}$$

If $T_X = T_Y$ then $X = Y$ (for example, you could evaluate at the standard basis vector \mathbf{e}_j to see that the j th column of X equals the j th column of Y for any j), so we get $A(BC) = (AB)C$.

Since we didn't prove the theorem here, we'll prove the associativity result in a more pedestrian way in the next section.

3.3 Transpose

Definition 3.3.1. Let $A = (a_{ij})$ be a $m \times n$ matrix. The **transpose** of A , written A^T , is the $n \times m$ matrix whose i, j entry is a_{ji} .

You can think of the transpose as being obtained by reflecting A in the south east diagonal starting in the top left hand corner, or as the matrix whose columns are the transposes of the rows of A , or the matrix whose rows are the transposes of the columns of A .

Example 3.3.1. • If $A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix}$ then $A^T = \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$.

- If $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$ then $A^T = \begin{pmatrix} 1 & 3 \\ 2 & 4 \end{pmatrix}$.
- If $A = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$ then $A^T = (1 \ 2 \ 3)$.

It's common to use transposes when we want to think geometrically, because if $\mathbf{x} \in \mathbb{R}^n$ then $\mathbf{x}^T \mathbf{x}$ is equal to

$$x_1^2 + x_2^2 + \cdots + x_n^2$$

which is the square of the length of \mathbf{x} . (As usual, we have identified the 1×1 matrix $\mathbf{x}^T \mathbf{x}$ with a number here).

When \mathbf{z} is a *complex* column vector, that is, an element of \mathbb{C}^n for some n , this doesn't quite work. If $\mathbf{z} = \begin{pmatrix} 1 \\ i \end{pmatrix}$ for example, then $\mathbf{z}^T \mathbf{z} = 0$, which is not a good measure of the length of \mathbf{z} . For this reason, when people work with complex vectors they often use the *conjugate transpose* A^H defined to be the matrix whose entries are the complex conjugates of the entries of A^T . With this

definition, for a complex vector $\mathbf{z} = \begin{pmatrix} z_1 \\ \vdots \\ z_n \end{pmatrix}$ we get

$$\mathbf{z}^H \mathbf{z} = |z_1|^2 + \cdots + |z_n|^2.$$

3.4 Multiplication properties

Proposition 3.4.1. *Let A and A' be $m \times n$ matrices, let B and B' be $n \times p$ matrices, let C be a $p \times q$ matrix, and let λ be a number. Then*

1. $A(BC) = (AB)C$ (*associativity*),
2. $(A + A')B = AB + A'B$, and $A(B + B') = AB + AB'$ (*distributivity*),
3. $(\lambda A)B = \lambda(AB) = A(\lambda B)$, and
4. $(AB)^T = B^T A^T$.

Proof. Let $A = (a_{ij})$, $A' = (a'_{ij})$, $B = (b_{ij})$, $B' = (b'_{ij})$, $C = (c_{ij})$. During this proof we also write X_{ij} to mean the i, j entry of a matrix X .

1. AB has i, j entry $\sum_{k=1}^n a_{ik} b_{kj}$, so the i, j entry of $(AB)C$ is

$$\sum_{l=1}^p (AB)_{il} c_{lj} = \sum_{l=1}^p \sum_{k=1}^n a_{ik} b_{kl} c_{lj}. \quad (3.5)$$

On the other hand, the i, j entry of BC is $\sum_{l=1}^p b_{il} c_{lj}$ so the i, j entry of $A(BC)$ is

$$\begin{aligned} \sum_{k=1}^n a_{ik} (BC)_{kj} &= \sum_{k=1}^n a_{ik} \sum_{l=1}^p b_{kl} c_{lj} \\ &= \sum_{k=1}^n \sum_{l=1}^p a_{ik} b_{kl} c_{lj}. \end{aligned} \quad (3.6)$$

(3.6) and (3.5) are the same because it doesn't matter if we do the k or l summation first: we just get the same terms in a different order.

2. The i, j entry of $(A+A')B$ is $\sum_{k=1}^n (a_{ik}+a'_{ik})b_{kj}$ which equals $\sum_{k=1}^n a_{ik}b_{kj} + \sum_{k=1}^n a'_{ik}b_{kj}$, but this is the sum of the i, j entry of AB and the i, j entry of $A'B$, proving the first equality. The second is similar.
3. The i, j entry of λA is λa_{ij} , so the i, j entry of $(\lambda A)B$ is

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

so $(\lambda A)B$ and $\lambda(AB)$ have the same i, j entry for any i, j , and are therefore equal. The second equality can be proved similarly.

4. One special case of this result is very easy: for any row vector $\mathbf{a} =$

$$(a_1 \quad \cdots \quad a_n) \text{ and column vector } \mathbf{b} = \begin{pmatrix} b_1 \\ \vdots \\ b_n \end{pmatrix} \text{ we have}$$

$$\mathbf{ab} = \sum_{k=1}^n a_k b_k = \sum_{k=1}^n b_k a_k = \mathbf{b}^T \mathbf{a}^T.$$

For the general case, we start off by checking that $(AB)^T$ and $B^T A^T$ have the same size. AB is $m \times p$ so $(AB)^T$ is $p \times m$, while B^T and A^T are $p \times n$ and $n \times m$ respectively so $B^T A^T$ is also $p \times m$. Now we only need to show that for any i and j , they have the same i, j entry. Let \mathbf{r}_i be the i th row of A and \mathbf{c}_j the j th column of B . The i, j entry of AB is $\mathbf{r}_i \mathbf{c}_j$, so the i, j entry of $(AB)^T$ is $\mathbf{r}_j \mathbf{c}_i$. On the other hand, the i th row of B^T is \mathbf{c}_i^T and the j th column of A^T is \mathbf{r}_j^T , so the i, j entry of $B^T A^T$ is $\mathbf{c}_i^T \mathbf{r}_j^T$. These two are equal by the special case mentioned at the start of this proof. \square

These results tell you that you can use some of the normal rules of algebra when you work with matrices, like what happened for permutations. Again, like permutations, what you can't do is use the commutative property.

3.4.1 Matrix multiplication isn't commutative

Definition 3.4.1. Two matrices A and B are said to **commute** if AB and BA are both defined and $AB = BA$.

For some pairs of matrices, the product AB is defined but BA is not. For example, if A is 2×3 and B is 3×4 then AB is defined but BA isn't. Even when both AB and BA are defined and have the same size they won't in general be equal.

Example 3.4.1. let $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$ and $B = \begin{pmatrix} 5 & 6 \\ 7 & 8 \end{pmatrix}$. Then

$$AB = \begin{pmatrix} 19 & 22 \\ 43 & 50 \end{pmatrix}$$

$$BA = \begin{pmatrix} 23 & 34 \\ 31 & 46 \end{pmatrix}.$$

3.4.2 The identity matrix

Definition 3.4.2. The $n \times n$ **identity matrix** I_n is the matrix with i, j entry 1 if $i = j$ and 0 otherwise.

For example,

$$I_2 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, I_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

The most important property of identity matrices is that they behave like the number 1 does when you multiply by them.

Theorem 3.4.2. If A is an $m \times n$ matrix then $I_m A = A I_n = A$.

Proof. Let $A = (a_{ij})$, $I_n = (\delta_{ij})$, so δ_{ij} is 1 if $i = j$ and 0 otherwise. The formula for matrix multiplication tells us that for any i and j , the i, j entry of $I_m A$ is $\sum_{k=1}^m \delta_{ik} a_{kj}$. The only term in this sum that can be nonzero is the one when $k = i$, so the sum equals $1 \times a_{ij} = a_{ij}$. Thus the i, j entry of $I_m A$ equals a_{ij} , the i, j entry of A .

The other equality can be proved similarly. \square

3.5 Invertible matrices

Definition 3.5.1. An $n \times n$ matrix A is called **invertible** if and only if there exists an $n \times n$ matrix B such that $AB = BA = I_n$.

If there is such a matrix B , we can prove that there is only one such matrix B :

Proposition 3.5.1. If $AB = BA = I_n$ and $AC = CA = I_n$ then $B = C$.

Proof.

$$\begin{aligned} B &= B I_n && \text{Theorem 3.4.2} \\ &= B(AC) \\ &= (BA)C && \text{associativity} \\ &= I_n C \\ &= C && \text{Theorem 3.4.2} \end{aligned}$$

\square

This means that when a matrix is invertible we can talk about *the* inverse of A . We write A^{-1} for the inverse of A when it exists.

3.5.1 Matrices with rows or columns of zeroes are not invertible

Theorem 3.5.2. If an $n \times n$ matrix A has a row of zeroes, or a column of zeroes, then it is not invertible.

Proof. Suppose A has a column of zeroes and that B is any other $n \times n$ matrix. By Theorem 3.2.3, the columns of BA are B times the columns of A . In particular, one of these columns is B times the zero vector, which is the zero vector. Since one of the columns of BA is all zeroes, BA is not the identity.

If A has a row of zeroes, we can make a similar argument using Theorem 3.2.6. \square

3.5.2 Inverse of a product of matrices

If you multiply any number of invertible matrices together, the result is invertible. Recall the shoes-and-socks result about the inverse of a composition of two functions: exactly the same thing is true.

Theorem 3.5.3. *If A_1, \dots, A_k are invertible $n \times n$ matrices then $A_1 \cdots A_k$ is invertible with inverse $A_k^{-1} \cdots A_1^{-1}$.*

The proof is the same as for functions: you can simply check that $A_k^{-1} \cdots A_1^{-1}$ is a two sided inverse to $A_1 \cdots A_k$ using the associativity property for matrix multiplication.

This theorem has a useful corollary about when matrix products are invertible.

Corollary 3.5.4. *Let A and E be $n \times n$ matrices with E invertible. Then EA is invertible if and only if A is invertible, and AE is invertible if and only if A is invertible.*

Proof. If A is invertible then the theorem tells us that so are EA and AE .

Suppose EA is invertible. Certainly E^{-1} is invertible (its inverse is E), so by the theorem $E^{-1}EA$ is invertible, that is, A is invertible. The argument for AE is similar. \square

3.5.3 Inverse of a transpose

Theorem 3.5.5. *Let A be an $n \times n$ matrix. Then A is invertible if and only if A^T is invertible.*

Proof. Suppose A is invertible. We claim that A^T is invertible with inverse $(A^{-1})^T$. This is true because

$$\begin{aligned} A^T(A^{-1})^T &= (A^{-1}A)^T && \text{Proposition 3.4.1 part 4} \\ &= I_n^T \\ &= I_n \end{aligned}$$

and similarly $(A^{-1})^T A^T = I_n$.

Conversely, suppose A^T is invertible. The above argument shows that $(A^T)^T$ is invertible, that is, A is invertible. \square

3.6 Systems of linear equations

3.6.1 Definition of a linear system

Definition 3.6.1. A system of m linear equations in n unknowns x_1, \dots, x_n with coefficients $a_{ij}, 1 \leq i \leq m, 1 \leq j \leq n$ and b_1, \dots, b_m is a list of simultaneous

equations

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n &= b_2 \\ &\vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= b_m \end{aligned}$$

As the notation suggests, we can turn a system of linear equations into a matrix equation and study it using matrix methods.

3.6.2 Matrix form of a linear system

Every system of linear equations can be written in **matrix form**: the above system is equivalent to saying that $A\mathbf{x} = \mathbf{b}$, where $A = (a_{ij})$, $\mathbf{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$, and $\mathbf{b} = \begin{pmatrix} b_1 \\ \vdots \\ b_m \end{pmatrix}$.

Example 3.6.1. The system of linear equations

$$\begin{aligned} 2x + 3y + 4z &= 5 \\ x + 5z &= 0 \end{aligned} \tag{3.7}$$

has matrix form

$$\begin{pmatrix} 2 & 3 & 4 \\ 1 & 0 & 5 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 5 \\ 0 \end{pmatrix}.$$

This connection means that we can use systems of linear equations to learn about matrices, and use matrices to learn about systems of linear equations. For example, if A is invertible then we can solve the matrix equation

$$A\mathbf{x} = \mathbf{b}$$

by multiplying both sides by A^{-1} to get that there is a *unique* solution $\mathbf{x} = A^{-1}\mathbf{b}$.

We are going to make two more observations about solving linear systems based on what we know about matrix multiplication. The first is that by Proposition 3.2.1, the vectors which can be written as $A\mathbf{u}$ for some \mathbf{u} are exactly the ones which are linear combinations of the columns of A , that is, vectors of the form

$$u_1\mathbf{c}_1 + \cdots + u_n\mathbf{c}_n$$

where \mathbf{c}_j is the j th column of A . So the matrix equation $A\mathbf{x} = \mathbf{b}$ has a solution if and only if \mathbf{b} can be written as a linear combination of the columns of A . This set of linear combinations is therefore important enough to have a name.

Definition 3.6.2. The **column space** of a matrix A , written $C(A)$, is the set of all linear combinations of the columns of A .

We can now restate the above by saying that $A\mathbf{x} = \mathbf{b}$ has a solution if and only if \mathbf{b} belongs to the column space of A .

Definition 3.6.3. A **homogeneous** matrix equation is one of the form $A\mathbf{x} = \mathbf{0}$.

What does it mean for $A\mathbf{k} = \mathbf{0}$ to be true? Using Proposition 3.2.1 again, it says that

$$k_1\mathbf{c}_1 + \cdots + k_n\mathbf{c}_n = \mathbf{0} \quad (3.8)$$

where the k_j are the entries of \mathbf{k} and the \mathbf{c}_j are the columns of A . An equation of the form (3.8) is called a **linear dependence relation**, or just a linear dependence, on $\mathbf{c}_1, \dots, \mathbf{c}_n$.

Definition 3.6.4. Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ be matrices of the same size. A **linear dependence relation** on $\mathbf{v}_1, \dots, \mathbf{v}_n$ is an equation

$$a_1\mathbf{v}_1 + \cdots + a_n\mathbf{v}_n = \mathbf{0}$$

where the a_i are numbers. A linear dependence relation is called **nontrivial** or **nonzero** if not all of the a_i are zero.

We've seen that solutions of the matrix equation $A\mathbf{x} = \mathbf{0}$ correspond to linear dependence relations on the columns of A .

The solutions of the matrix equation $A\mathbf{x} = \mathbf{0}_m$ are so important that they get their own name.

Definition 3.6.5. The **nullspace** of an $m \times n$ matrix A , written $N(A)$, is $\{\mathbf{v} \in \mathbb{R}^n : A\mathbf{v} = \mathbf{0}_m\}$.

The homogeneous equation $A\mathbf{x} = \mathbf{0}_m$ has the property that the zero vector is a solution, if \mathbf{u} and \mathbf{v} are solutions then so is $\mathbf{u} + \mathbf{v}$, and if λ is a number then $\lambda\mathbf{u}$ is also a solution. This is what it means to say that $N(A)$ is a *subspace* of \mathbb{R}^n , something we will cover in the final chapter of MATH0005.

3.6.3 Augmented matrix

The **augmented matrix** of a system of linear equations whose matrix form is $A\mathbf{x} = \mathbf{b}$ is the matrix which you get by adding \mathbf{b} as an extra column on the right of A . We write this as $(A \mid \mathbf{b})$ or just $(A \ \mathbf{b})$.

For example, the augmented matrix for the system of linear equations (3.7) above would be

$$\begin{pmatrix} 2 & 3 & 4 & 5 \\ 1 & 0 & 5 & 0 \end{pmatrix}.$$

Definition 3.6.6. • A **solution** to a matrix equation $A\mathbf{x} = \mathbf{b}$ is a vector \mathbf{y} (of numbers this time, not unknowns) such that $A\mathbf{y} = \mathbf{b}$.

- A matrix equation $A\mathbf{x} = \mathbf{b}$ is called **consistent** if it has a solution and **inconsistent** otherwise.

A system of linear equations may have a unique solution, many different solutions, or no solutions at all. In future lectures we will see how to find out how many solutions, if any, a system has.

3.7 Row operations

3.7.1 How we solve linear systems

If you are given a system of linear equations in variables x, y, z and asked to solve them, what you probably do is to manipulate the equations by adding multiples of one equation to another until you have “eliminated” some of the variables and you can read off the solutions. We are going to try to formalise this method of solving linear equations.

Because we want to use matrix methods, let’s solve an example system and keep track of what our equation manipulation does to the corresponding augmented matrix.

Consider the linear system

$$\begin{aligned} 3x + 4y &= 6 \\ x + 2y &= 5. \end{aligned}$$

The corresponding matrix equation is $A\mathbf{x} = \mathbf{b}$ where

$$A = \begin{pmatrix} 3 & 4 \\ 1 & 2 \end{pmatrix}, \mathbf{b} = \begin{pmatrix} 6 \\ 5 \end{pmatrix}, \mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

The augmented matrix is $\begin{pmatrix} 3 & 4 & 6 \\ 1 & 2 & 5 \end{pmatrix}$.

To solve the system, we first eliminate x from first equation by subtracting 3 times the second equation from the first. The equations become

$$\begin{aligned} -2y &= -9 \\ x + 2y &= 5 \end{aligned}$$

The augmented matrix $\begin{pmatrix} -2 & 0 & -9 \\ 1 & 2 & 5 \end{pmatrix}$ of this new system is obtained by adding -3 times the second row of the old augmented matrix to the first row.

Next we get the coefficient of y in the first equation to 1 by multiplying the first equation by $-1/2$. The equations become

$$\begin{aligned} y &= 9/2 \\ x + 2y &= 5 \end{aligned}$$

The augmented matrix $\begin{pmatrix} 0 & 1 & 9/2 \\ 1 & 2 & 5 \end{pmatrix}$ of this new system is obtained by multiplying every entry in the first row of the old augmented matrix by $-1/2$.

Next we eliminate y from the second equation by subtracting 2 times the first equation from the second. The equations become

$$\begin{aligned} y &= 9/2 \\ x &= -4 \end{aligned}$$

The augmented matrix $\begin{pmatrix} 0 & 1 & 9/2 \\ 1 & 0 & -4 \end{pmatrix}$ of this new system is obtained by adding -2 times the first row to the second row.

row operation	inverse
$\mathbf{r}_j \mapsto \mathbf{r}_j + \lambda \mathbf{r}_i$ (add λ times \mathbf{r}_i to \mathbf{r}_j)	$\mathbf{r}_j \mapsto \mathbf{r}_j - \lambda \mathbf{r}_i$ (add $-\lambda$ times \mathbf{r}_i to \mathbf{r}_j)
$\mathbf{r}_i \mapsto \lambda \mathbf{r}_i$ (multiply \mathbf{r}_i by $\lambda \neq 0$)	$\mathbf{r}_i \mapsto \lambda^{-1} \mathbf{r}_i$ (multiply \mathbf{r}_i by λ^{-1})
$\mathbf{r}_i \leftrightarrow \mathbf{r}_j$ (swap \mathbf{r}_i and \mathbf{r}_j)	$\mathbf{r}_i \leftrightarrow \mathbf{r}_j$ (swap \mathbf{r}_i and \mathbf{r}_j)

Lastly, if we wanted the first equation to tell us the value of the first variable and the second equation to tell us about the second variable, we could swap the order of the two equations, corresponding to swapping the rows of the augmented matrix so that it becomes $\begin{pmatrix} 1 & 0 & -4 \\ 0 & 1 & 9/2 \end{pmatrix}$.

3.7.2 Row operations

The manipulations we do to systems of linear equations correspond to doing *row operations* to the augmented matrices.

Definition 3.7.1. A **row operation** is one of the following things we can do to a matrix.

1. Add λ times row i to row j (for $j \neq i$, λ any number), written $\mathbf{r}_j \mapsto \mathbf{r}_j + \lambda \mathbf{r}_i$.
2. Multiply row i by λ , where $\lambda \neq 0$, written $\mathbf{r}_i \mapsto \lambda \mathbf{r}_i$.
3. Swap rows i and j , written $\mathbf{r}_i \leftrightarrow \mathbf{r}_j$.

3.7.3 Row operations are invertible

For each row operation r there is another row operation s such that doing r then s , or doing s then r , gets you back to the matrix you started with. Here is a table of the three types of row operations and their inverses.

3.8 Elementary matrices

Definition 3.8.1. An **elementary matrix** is one obtained by doing a single row operation to an identity matrix.

Example 3.8.1. • The elementary matrix $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ results from doing the row operation $\mathbf{r}_1 \leftrightarrow \mathbf{r}_2$ to I_2 .

• The elementary matrix $\begin{pmatrix} 1 & 2 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ results from doing the row operation $\mathbf{r}_1 \mapsto \mathbf{r}_1 + 2\mathbf{r}_2$ to I_3 .

• The elementary matrix $\begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}$ results from doing the row operation $\mathbf{r}_1 \mapsto (-1)\mathbf{r}_1$ to I_2 .

Doing a row operation r to a matrix has the same effect as multiplying that matrix on the left by the elementary matrix corresponding to r .

Theorem 3.8.1. *Let r be a row operation and B an $m \times n$ matrix. Then $r(B) = r(I_m)B$.*

Proof. Any row operation r changes B to a matrix $r(B)$ whose rows are linear combinations (Definition 3.2.2) of the rows of B . Proposition 3.2.5 and Theorem 3.2.6 show that we can choose an $m \times m$ matrix A such that $r(B) = AB$. By putting $B = I_m$ we see that $A = r(I_m)$. \square

Example 3.8.2. The theorem tells you that doing the row operation $\mathbf{r}_1 \mapsto \mathbf{r}_1 + 2\mathbf{r}_2$ to a matrix with three rows is the same as left-multiplying by the 3×3

elementary matrix $\begin{pmatrix} 1 & 2 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$. For example, if we do this row operation to $\begin{pmatrix} a \\ b \\ c \end{pmatrix}$ we get $\begin{pmatrix} a + 2b \\ b \\ c \end{pmatrix}$ which is the same as

$$\begin{pmatrix} 1 & 2 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} a + 2b \\ b \\ c \end{pmatrix}.$$

Corollary 3.8.2. *Elementary matrices are invertible.*

Proof. Let r be a row operation, s be the inverse row operation to r , and let I_n be an identity matrix. By Theorem 3.8.1, $r(I_n)s(I_n) = r(s(I_n))$. Because s is inverse to r , this equals I_n . Similarly, $s(I_n)r(I_n) = s(r(I_n)) = I_n$. It follows that $r(I_n)$ is invertible with inverse $s(I_n)$. \square

Theorem 3.8.1 combined with Corollary 3.8.2 shows that if A results from doing a row operation to B , then $A = EB$ for some invertible matrix E . What about if A results from doing a sequence of row operations?

Theorem 3.8.3. *Suppose that A is a matrix obtained by doing a sequence of row operations to another matrix B . Then $A = EB$ for some invertible matrix E .*

Proof. If the row operations are r_1, \dots, r_k and if E_i is the elementary matrix corresponding to r_i then using Theorem 3.8.1 repeatedly gives

$$A = E_k E_{k-1} \cdots E_2 E_1 B.$$

Let $E = E_k E_{k-1} \cdots E_2 E_1$, so $A = EB$. The matrix E is a product of invertible matrices, by Corollary 3.8.2, so it is invertible by Theorem 3.5.3. \square

3.9 Row reduced echelon form

3.9.1 Row operations don't change the solutions to a matrix equation

Our informal method of solving linear systems is to do certain manipulations to the equations until they are in a form where the solutions are easy to read off. This method only works if the manipulations we do don't change the set of solutions.

When we introduced row operations, it was because their effect on the augmented matrix of a linear system corresponded to the kind of manipulations we perform when solving such a linear system. We're now going to prove that these row operations don't change the set of solutions.

Theorem 3.9.1. *Suppose that $(A' \mid \mathbf{b}')$ results from $(A \mid \mathbf{b})$ by doing a sequence of row operations. Then the matrix equations $A\mathbf{x} = \mathbf{b}$ and $A'\mathbf{x} = \mathbf{b}'$ have the same solutions.*

Proof. By Theorem 3.8.3, there is an invertible matrix E such that

$$E(A \mid \mathbf{b}) = (A' \mid \mathbf{b}')$$

and so (because matrix multiplication works columnwise, Theorem 3.2.3) $EA = A'$ and $E\mathbf{b} = \mathbf{b}'$.

We must show that a column vector \mathbf{v} is a solution of $A\mathbf{x} = \mathbf{b}$ if and only if it is a solution of $A'\mathbf{x} = \mathbf{b}'$. If $A\mathbf{v} = \mathbf{b}$ then multiplying on the left by E gives $E A \mathbf{v} = E \mathbf{b}$, that is, $A' \mathbf{v} = \mathbf{b}'$. Conversely if $A' \mathbf{v} = \mathbf{b}'$ then $E A \mathbf{v} = E \mathbf{b}$, so multiplying on the left by E^{-1} gives $A \mathbf{v} = \mathbf{b}$. \square

3.9.2 Row reduced echelon form

We talked about manipulating equations into a simple form where the solutions could be easily read off. The simple form we will use is called row reduced echelon form. To define it we need the notion of a leading entry.

Definition 3.9.1. The **leading entry** in a nonzero row of a matrix is the first non-zero entry in that row, starting from the left.

Of course, if a row is all zeroes then it doesn't have a leading entry. In the matrix $\begin{pmatrix} 0 & 2 \\ 0 & 0 \end{pmatrix}$ the leading entry in the first row is the 2 in position 1, 2 while the second row has no leading entry.

Definition 3.9.2. A matrix is in **row reduced echelon form** (RREF) if

1. all leading entries are 1,
2. any rows which are all zero are below any rows which are not all zero,
3. all entries in the same column as a leading entry are zero, and
4. for every i , if row i and $i + 1$ have a leading entry then the leading entry in row $i + 1$ is to the right of that in row i .

Example 3.9.1. • $\begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$ isn't in RREF: the zero row is at the top, violating condition 2

• $\begin{pmatrix} 2 & 0 \\ 0 & 0 \end{pmatrix}$ isn't in RREF: there is a row in which the left-most non-zero entry is not 1 which breaks condition 1.

• $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ isn't in RREF because of condition 4: the left-most 1 in row 2 is not to the right of the left-most 1 in the row above it.

- $\begin{pmatrix} 1 & \alpha & \beta & 3 \\ 0 & 0 & 1 & -2 \end{pmatrix}$ is in RREF if and only if $\beta = 0$: the left-most 1 in row 2 is in column 3, but it is not the only non-zero entry in column 3 (as required by condition 3) unless $\beta = 0$.
- $\begin{pmatrix} 1 & 0 & 0 & 3 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$ is in RREF.

3.10 RREF existence and uniqueness

3.10.1 Existence and uniqueness

Here are two facts about row reduced echelon form.

Theorem 3.10.1. *For every matrix A , there is a sequence of row operations taking A to a matrix in row reduced echelon form.*

Proof. We prove this by induction on the number of columns of A . When A has one column, either A is the zero vector (in which case it is in RREF) or it has a nonzero entry a . Swap a to the top row, multiply the top row by $1/a$, and subtract multiples of the top row from the other rows to make all other rows 0.

The result is the vector $\begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$ which is in RREF.

For the inductive step, suppose that A is $m \times n$ and that the result is true for all matrices with $n - 1$ columns. We then know that there is a series of row operations we can do to A that result in a matrix X whose first $n - 1$ columns form a RREF matrix. Suppose the matrix formed by these $n - 1$ columns has k rows of zeroes at the bottom. If the final column has zeroes in its bottom k entries, the matrix is in RREF. If not, swap a nonzero entry to the top of these k rows, subtract multiples of this row from the others to eliminate all other nonzero entries in the final column, and multiply by a scalar so that its entry is 1. The result is in RREF. \square

Theorem 3.10.2. *Let A be a matrix. If R and S are RREF matrices that can be obtained by doing row operations to A , then $R = S$.*

This theorem says that there is only one RREF matrix which can be obtained by doing row operations to A , so we are justified in calling the unique RREF matrix reachable from A the row reduced echelon form of A .

Proof. Again, the proof is by induction on the number n of columns of A . The base case is when A is an $m \times 1$ column vector. There are only two RREF height m column vectors: the zero vector $\mathbf{0}_m$ and a vector \mathbf{v} with a 1 at the top and 0s everywhere else. If there are sequences of row operations taking A to \mathbf{v} and also to $\mathbf{0}_m$ then there is a sequence of row operations taking $\mathbf{0}_m$ to \mathbf{v} . This is impossible because any row operation sends $\mathbf{0}_m$ to $\mathbf{0}_m$.

For the inductive step, suppose that A is $m \times n$ and that R and S are RREF matrices reachable from A . Let A' , R' , and S' be the matrices formed by the first $n - 1$ columns of A , R , and S respectively. The matrices R' and S' are

RREF matrices formed by doing row operations to A' , so by induction they are equal. Suppose for a contradiction that $R \neq S$, so that there is some j such that the j th entry r_{jn} in the last column of R which differs from the corresponding entry s_{jn} of S .

Theorem 3.9.1 tells us that the equations $A\mathbf{x} = \mathbf{0}$, $R\mathbf{x} = \mathbf{0}$, and $S\mathbf{x} = \mathbf{0}$ all have exactly the same solutions.

Let \mathbf{u} be any solution of $A\mathbf{x} = \mathbf{0}$. We have $R\mathbf{u} = S\mathbf{u} = \mathbf{0}$, so $(R - S)\mathbf{u} = \mathbf{0}$. Since the first $n - 1$ columns of $R - S$ are all zeroes, we get $(r_{jn} - s_{jn})u_n = 0$, so $u_n = 0$. In other words, every solution of $A\mathbf{x} = \mathbf{0}$ has last entry zero.

The RREF matrix R' has some nonzero rows and then some zero rows. Suppose there are k zero rows. We can then write R' like this

$$\begin{pmatrix} X \\ \mathbf{0}_{k \times (n-1)} \end{pmatrix}$$

where X has no zero rows. Then R and S have the form

$$R = \begin{pmatrix} X & \mathbf{r} \\ \mathbf{0}_{k \times (n-1)} & \mathbf{t} \end{pmatrix}, S = \begin{pmatrix} X & \mathbf{s} \\ \mathbf{0}_{k \times (n-1)} & \mathbf{u} \end{pmatrix}$$

I claim that $\mathbf{t} \neq \mathbf{0}$. Suppose for a contradiction that $\mathbf{t} = \mathbf{0}$. The first $m - k$ rows of R all have leading entries. For $1 \leq i \leq m - k$, let the leading entry in row i of R occur in column number c_i . Let \mathbf{r} have entries r_1, \dots, r_{m-k} , where m is the number of rows of A . Notice that column c_i of R has a 1 in position i and zeroes everywhere else, so if we add up r_1 times column c_1 , r_2 times column c_2 , and so on, up to r_{m-k} times column c_{m-k} we get the vector

$$\begin{pmatrix} r_1 \\ \vdots \\ r_{m-k} \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

which is the last column of R . It follows that the vector with -1 in position n , with r_i in position c_i for $1 \leq i \leq m - k$, and with zeroes elsewhere is a solution to $R\mathbf{x} = \mathbf{0}$. This contradicts every solution to $A\mathbf{x} = \mathbf{0}$ having last entry zero.

Since R is in RREF, \mathbf{t} must have a 1 at the top and all other entries zero, and $\mathbf{r} = \mathbf{0}$. The same argument applies to S , so $\mathbf{u} = \mathbf{t}$ and $\mathbf{s} = \mathbf{0}$. This shows $R = S$. \square

3.10.2 Example of putting a matrix into RREF

Let $A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$. We want to do a sequence of row operations to A which ends up with a matrix in RREF. Row 1 has a leading entry at position 1, 1, but the other entries in column 1 aren't 0. We use the 1, 1 entry as a **pivot** to eliminate the other entries in column 1. That is, we apply row operations of the

form $\mathbf{r}_j \mapsto \mathbf{r}_j + \lambda \mathbf{r}_1$ to make the other entries in column 1 equal to 0.

$$A \xrightarrow{\mathbf{r}_2 \mapsto \mathbf{r}_2 - 4\mathbf{r}_1} \begin{pmatrix} 1 & 2 & 3 \\ 0 & -3 & -6 \\ 7 & 8 & 9 \end{pmatrix} \xrightarrow{\mathbf{r}_3 \mapsto \mathbf{r}_3 - 7\mathbf{r}_1} \begin{pmatrix} 1 & 2 & 3 \\ 0 & -3 & -6 \\ 0 & -6 & -12 \end{pmatrix}$$

This matrix isn't in RREF. One reason is that the leading entry in row 2, in position 2, 2, isn't equal to 1. To make that leading entry 1 we can use the row operation that multiplies row 2 by $-1/3$:

$$\begin{pmatrix} 1 & 2 & 3 \\ 0 & -3 & -6 \\ 0 & -6 & -12 \end{pmatrix} \xrightarrow{\mathbf{r}_2 \mapsto (-1/3)\mathbf{r}_2} \begin{pmatrix} 1 & 2 & 3 \\ 0 & 1 & 2 \\ 0 & -6 & -12 \end{pmatrix}$$

Now we have a leading entry in row 2, column 2 which is equal to 1, but there are other nonzero entries in that column. We use the 2, 2 entry as the next pivot to eliminate entries in column 2.

$$\begin{pmatrix} 1 & 2 & 3 \\ 0 & 1 & 2 \\ 0 & -6 & -12 \end{pmatrix} \xrightarrow{\mathbf{r}_1 \mapsto \mathbf{r}_1 - 2\mathbf{r}_2} \begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & 2 \\ 0 & -6 & -12 \end{pmatrix} \\ \xrightarrow{\mathbf{r}_3 \mapsto \mathbf{r}_3 + 6\mathbf{r}_2} \begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & 2 \\ 0 & 0 & 0 \end{pmatrix}$$

This matrix is in RREF, so we are done.

3.11 Solving RREF systems

Suppose we start with a linear system with matrix form $A\mathbf{x} = \mathbf{b}$ then put the augmented matrix $(A \mid \mathbf{b})$ into RREF. Suppose the resulting matrix in RREF is $(A' \mid \mathbf{b}')$. The whole point of RREF was that the solutions of $A\mathbf{x} = \mathbf{b}$ are the same as those of $A'\mathbf{x} = \mathbf{b}'$ but it should be "easy" to find the solutions of $A'\mathbf{x} = \mathbf{b}'$. How do we actually find those solutions?

Example 3.11.1. Here is an augmented matrix in RREF

$$\begin{pmatrix} 1 & 0 & 2 & 0 & 0 \\ 0 & 1 & 4 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

If the variables are called x, y, z, w then the corresponding equations are

$$\begin{aligned} x + 2z &= 0 \\ y + 4z &= 0 \\ w &= 0 \\ 0 &= 1 \end{aligned}$$

The last equation is false, so there are no solutions to this linear system.

Example 3.11.2. Here is the same augmented matrix with a different final column.

$$\begin{pmatrix} 1 & 0 & 2 & 0 & 2 \\ 0 & 1 & 4 & 0 & 3 \\ 0 & 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

In this case, if the variables are x, y, z, w , the equations are

$$\begin{aligned} x + 2z &= 2 \\ y + 4z &= 3 \\ w &= 4 \\ 0 &= 0 \end{aligned}$$

Clearly we can choose any value for z and this will uniquely determine the values of the other three variables through the equations $x = 2 - 2z, y = 3 - 4z, w = 4$. The **general solution** is

$$\begin{pmatrix} x \\ y \\ z \\ w \end{pmatrix} = \begin{pmatrix} 2 - 2z \\ 3 - 4z \\ z \\ 4 \end{pmatrix}.$$

In general, for an augmented matrix in row reduced echelon form:

- If the last column of the augmented matrix has a leading entry (like in the first example above), there are no solutions. Otherwise,
- variables corresponding to a column with no leading entry (like z in the second example) can be chosen freely, and
- the other variables are uniquely determined in terms of these free variables.

Variables whose column has no leading entry are called **free variables** and variables whose column does contain a leading entry are called **pivot variables**.

It's easy to see why the first bullet point is true: in that case the row with a leading entry in its final column would correspond to an equation saying

$$0 = 1$$

which has no solutions, so the system is inconsistent. Let's prove the second part of the claim.

Theorem 3.11.1. *Suppose $R\mathbf{x} = \mathbf{b}$ is a system of linear equations whose augmented matrix is in RREF and does not have a leading entry in the final column. Let R be $m \times n$, let x_{j_1}, \dots, x_{j_k} be the free variables, and let x_{i_1}, \dots, x_{i_l} be the pivot variables. Then for any list of numbers a_1, \dots, a_k there is exactly one solution to $R\mathbf{x} = \mathbf{b}$ in which $x_{j_p} = a_p$ for all $1 \leq p \leq k$.*

The point of this theorem is to make precise the idea that in a RREF linear system, the values of the free variables can be chosen arbitrarily, and these choices completely determine the values of the other (pivot) variables.

Proof. The l nonzero equations corresponding to the first l rows of the augmented matrix take the form

$$x_{i_q} + \sum_{p=1}^k r_{q,j_p} x_{j_p} = b_q \quad (3.9)$$

for $1 \leq q \leq l$. (The coefficient of x_{i_q} is 1 because leading entries in a RREF matrix are 1, and no pivot variables except x_{i_q} appear in the q th equation with nonzero coefficient because pivot variables correspond to columns with a leading entry, and by the definition of RREF every other entry in those columns must be zero).

Now let a_1, \dots, a_k be any numbers. If $x_{j_p} = a_p$ for $1 \leq p \leq k$ then the equations (3.9) force

$$x_{i_q} = b_q - \sum_{p=1}^k r_{q,j_p} a_p.$$

All values of the free and pivot variables are determined, so there is exactly one such solution. \square

To give the **general solution** of a system of linear equations means to identify the free variables by putting the system into RREF and then give an expression for the pivot variables in terms of the free variables. Here is another example of finding a general solution of a system of linear equations.

Example 3.11.3. Consider the system of linear equations

$$\begin{aligned} 2x_1 + x_2 + x_3 - 2x_4 &= 3 \\ x_1 + x_3 - 2x_4 &= 2 \\ x_1 + 2x_3 - 4x_4 &= 1 \\ 3x_1 + x_2 + 3x_3 - 6x_4 &= 4. \end{aligned}$$

By doing row operations to the corresponding augmented matrix, you can reach the RREF matrix

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

with corresponding system of linear equations

$$\begin{aligned} x_1 &= 3 \\ x_2 &= -2 \\ x_3 - 2x_4 &= -1. \end{aligned}$$

There is one free variable, x_4 , which can take any value. The general solution is

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = \begin{pmatrix} 3 \\ -2 \\ 2x_4 - 1 \\ x_4 \end{pmatrix}.$$

We can now show that homogeneous linear systems with more variables than equations have nonzero solutions (that is, solutions other than the zero vector $\mathbf{0}$).

Proposition 3.11.2. *If A is $m \times n$ and $n > m$ then the matrix equation $A\mathbf{x} = \mathbf{0}$ has a nonzero solution.*

Proof. When we do row operations to A to get a RREF matrix, that RREF matrix has at most one leading entry in each of its m rows. Since there are more columns than rows, there must be a column with no leading entry. The corresponding variable is a free variable, so there is a solution in which it is nonzero. \square

3.12 Invertibility and RREF

We are going to prove a theorem connecting invertibility, RREF, and solutions of matrix equations. To do it we need a lemma on the RREF of square matrices.

Lemma 3.12.1. *Let R be an $n \times n$ matrix in RREF which is not I_n . Then R has a column with no leading entry.*

Proof. We'll prove the contrapositive: if every column of R has a leading entry, then $R = I_n$. Suppose every column has a leading entry, so there are n leading entries. There's at most one leading entry per row and there are n rows, so every row must have a leading entry.

The leading entries go from left to right as we move down the rows of the matrix, so the leading entries in row 1, 2, \dots , n must be in columns 1, 2, \dots , n otherwise there would be no room to fit them in.

Because R is in RREF, columns with leading entries have zeroes in all other

positions. So the first column is $\begin{pmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$, the second column is $\begin{pmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$, and so on.

These are the columns of the identity matrix, so $R = I_n$. \square

Now we can prove the theorem.

Theorem 3.12.2. *For an $n \times n$ matrix A , the following statements are equivalent.*

- (1) A is invertible.
- (2) The only solution to $A\mathbf{x} = \mathbf{0}$ is $\mathbf{x} = \mathbf{0}$.
- (3) There is a sequence of row operations taking A to I_n .

Proof. We will show that (1) implies (2), (2) implies (3), and (3) implies (1). When we've done this, it's easy to see that the truth of any one of these three statements implies the truth of any other.

- (1) \Rightarrow (2). If $A\mathbf{x} = \mathbf{0}$ then multiplying on the left by A^{-1} gives $\mathbf{x} = A^{-1}\mathbf{0} = \mathbf{0}$.

- (2) \Rightarrow (3). Let R be a matrix in RREF obtained by doing a sequence of row operations to A ; we know such a matrix exists by Theorem 3.10.1. The solutions to $A\mathbf{x} = \mathbf{0}$ are the same as the solutions to $R\mathbf{x} = \mathbf{0}$ by Theorem 3.9.1, so $R\mathbf{x} = \mathbf{0}$ has only the solution $\mathbf{x} = \mathbf{0}$. R must be I_n , otherwise by Lemma 3.12.1 it would have a column with no leading entry so there would be a free variable in the solutions to $R\mathbf{x} = \mathbf{0}$ and in particular there would be a nonzero solution.
- (3) \Rightarrow (1). Since there is a sequence of row operations taking A to I_n , there is an invertible matrix E such that $EA = I_n$ by Theorem 3.8.3. Then $A = E^{-1}I_n = E^{-1}$, so A is invertible with inverse E . \square

3.12.1 Left and right inverses

We saw in the previous chapter that an arbitrary function could have a left inverse but not a right inverse, or a right inverse but not a left inverse. The same is *not* true of square matrices.

Theorem 3.12.3. *Let A be an $n \times n$ matrix. The following are equivalent:*

- (i) A is invertible.
- (ii) There is a matrix L such that $LA = I_n$.
- (iii) There is a matrix R such that $AR = I_n$.

Proof. We will show that (i) \iff (ii) and then that (i) \iff (iii).

Certainly (i) implies (ii). Conversely, suppose that $LA = I_n$. Let \mathbf{v} be any solution of $A\mathbf{v} = \mathbf{0}$. Multiplying on the left by L we get $\mathbf{v} = \mathbf{0}$, so A is invertible by Theorem 3.12.2.

Again, (i) obviously implies (iii). If $AR = I_n$ then transposing both sides, $R^T A^T = I_n^T = I_n$. By the same argument as above, A^T is invertible, and so A is invertible by Theorem 3.5.5. \square

3.13 Finding inverses

Let A be a square matrix. We now have a method of determining whether or not A is invertible: do row operations to A until you reach a matrix in RREF. Then by Theorem 3.12.2 A is invertible if and only if the RREF matrix is the identity.

What if we actually want to know what the inverse matrix is? You probably already know that a 2×2 matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ is invertible if and only if $ad - bc \neq 0$, and in this case

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

This formula does generalise to larger matrices, but not in a way which is easy to use: for example, the general formula for the inverse of a 3×3 invertible

matrix $A = (a_{ij})$ is

$$A^{-1} = \frac{1}{\Delta} \begin{pmatrix} \begin{vmatrix} a_{22} & a_{23} \\ a_{32} & a_{33} \end{vmatrix} & -\begin{vmatrix} a_{12} & a_{13} \\ a_{32} & a_{33} \end{vmatrix} & \begin{vmatrix} a_{12} & a_{13} \\ a_{22} & a_{23} \end{vmatrix} \\ -\begin{vmatrix} a_{21} & a_{23} \\ a_{31} & a_{33} \end{vmatrix} & \begin{vmatrix} a_{11} & a_{13} \\ a_{31} & a_{33} \end{vmatrix} & -\begin{vmatrix} a_{11} & a_{13} \\ a_{21} & a_{23} \end{vmatrix} \\ \begin{vmatrix} a_{21} & a_{22} \\ a_{31} & a_{32} \end{vmatrix} & -\begin{vmatrix} a_{11} & a_{12} \\ a_{31} & a_{32} \end{vmatrix} & \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix} \end{pmatrix}$$

where $\begin{vmatrix} a & b \\ c & d \end{vmatrix}$ means $ad - bc$ and

$$\Delta = a_{11}a_{22}a_{33} + a_{12}a_{23}a_{31} + a_{13}a_{21}a_{32} - a_{11}a_{23}a_{32} - a_{12}a_{21}a_{33} - a_{13}a_{22}a_{31}.$$

This isn't a formula that you want to use. Luckily we can use RREF techniques to determine invertibility and find inverses.

3.13.1 How to determine invertibility and find inverses

Let A be an $n \times n$ matrix, and suppose we want to find out whether A is invertible and if so what its inverse is. Let I_n be the $n \times n$ identity matrix. Here is a method:

1. Form the super-augmented matrix $(A \mid I_n)$.
2. Do row operations to put this into RREF.
3. If you get $(I_n \mid B)$ then A is invertible with inverse B .
4. If the first part of the matrix isn't I_n then A isn't invertible.

It works because the first part of the matrix is a RREF matrix resulting from doing row operations to A , so if it is I_n then by Theorem 3.12.2 A is invertible, and if it is not I_n then A is not invertible. It just remains to explain why, in the case A is invertible, you end up with $(I_n \mid A^{-1})$.

Think about the columns $\mathbf{c}_1, \dots, \mathbf{c}_n$ of the inverse of A . We have $A(\mathbf{c}_1 \cdots \mathbf{c}_n) = I_n$, so $A\mathbf{c}_1 = \mathbf{e}_1$, $A\mathbf{c}_2 = \mathbf{e}_2$, etc, where \mathbf{e}_i is the i th column of I_n . So \mathbf{c}_1 is the unique solution of the matrix equation $A\mathbf{x} = \mathbf{e}_1$. You find that by putting $(A \mid \mathbf{e}_1)$ into RREF, and you must get $(I_n \mid \mathbf{c}_1)$ since \mathbf{c}_1 is the only solution.

Repeating that argument for every column, when we put $(A \mid \mathbf{e}_1 \cdots \mathbf{e}_n)$ into RREF we get $(I_n \mid \mathbf{c}_1 \cdots \mathbf{c}_n)$, that is, $(I_n \mid A^{-1})$.

Example 3.13.1. Let $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$. To find whether A is invertible, and if so what its inverse is, we put $(A \mid I_2)$ into RREF:

$$\begin{aligned} \begin{pmatrix} 1 & 2 & 1 & 0 \\ 3 & 4 & 0 & 1 \end{pmatrix} &\xrightarrow{\mathbf{r}_2 \mapsto \mathbf{r}_2 - 3\mathbf{r}_1} \begin{pmatrix} 1 & 2 & 1 & 0 \\ 0 & -2 & -3 & 1 \end{pmatrix} \\ &\xrightarrow{\mathbf{r}_2 \mapsto (-1/2)\mathbf{r}_2} \begin{pmatrix} 1 & 2 & 1 & 0 \\ 0 & 1 & 3/2 & -1/2 \end{pmatrix} \\ &\xrightarrow{\mathbf{r}_1 \mapsto \mathbf{r}_1 - 2\mathbf{r}_2} \begin{pmatrix} 1 & 0 & -2 & 1 \\ 0 & 1 & 3/2 & -1/2 \end{pmatrix} \end{aligned}$$

This is in RREF, so the inverse of A is

$$\begin{pmatrix} -2 & 1 \\ 3/2 & -1/2 \end{pmatrix}$$

as you can check by multiplying them together.

Further reading

There are literally hundreds of textbooks about matrices and linear algebra, so it is worth browsing the library and finding one that you like. If you get a good one, let me know.

As an undergraduate I found *Linear Algebra* by A.O. Morris very clear and easy to read. *Advanced Engineering Mathematics* by E. Kreyszig (any edition, there are loads) is also well written and despite the title has a great deal of material relevant to a math degree, not just linear algebra. I haven't read *Vectors, Pure and Applied: A General Introduction to Linear Algebra* by T. Körner, but the author is an excellent writer. The previous MATH0005 lecturer recommended *Guide to Linear Algebra* by D. Towers (no relation to me), *Elementary Linear Algebra* by H. Anton, and the more sophisticated *Linear Algebra* by S. Lang.

The MIT class *18.06 Linear Algebra* lectured by Gilbert Strang is really interesting, and Strang is a famously good lecturer. Lecture videos and assignments are available online. The course has almost no proofs which means it has time to cover a really wide range of material, far beyond what goes in a normal first year linear algebra course. Don't watch the lectures without also doing the assignments!

Chapter 4

Linear algebra

4.1 Fields

This part of the module is about generalizing what we know about matrices and vectors.

When we talk about vectors, or matrices, there's an important thing we have to decide: where do the entries come from? For example, we might work with matrices with real numbers as entries, or with complex numbers as entries. We never really discussed this in the first part of the module, because it doesn't make any difference to the theory we developed. So it's natural to ask which other kinds of numbers we could use as entries in our vectors and matrices and still have everything work OK.

The answer is that the entries must come from what is called a **field**. Roughly speaking, a field is a set with multiplication and addition operations that obey the usual rules of algebra, and where you can divide by any non-zero element. Examples are \mathbb{R} , the set of all real numbers, \mathbb{C} , the set of all complex numbers, \mathbb{Q} , the set of all rational numbers. Non-examples are \mathbb{Z} : you can't divide by 2 in \mathbb{Z} , and $M_{2 \times 2}(\mathbb{R})$, the set of all 2×2 real matrices, again because we know there are non-zero 2×2 matrices which aren't invertible.

The usual way to define a field is to write down a list of axioms. Everything that satisfies the axioms is a field. If you do want to know the field axioms, here they are: you can read them here.

The next section is about an important family of fields we have not seen yet.

4.1.1 Finite fields of prime order

Let p be a prime number. Then \mathbb{F}_p , the **finite field of order p** , is the set $\{0, 1, 2, \dots, p-1\}$ with addition and multiplication exactly the same as for ordinary whole numbers, except that we regard all multiples of p as equal to 0.

This is easier to understand in an example. Let's work in \mathbb{F}_5 . Then

- $3 + 4 = 7 = 2$ because $5 = 0$ in \mathbb{F}_5 . Similarly,
- $3 + 2 = 0$,
- $4 \times 3 = 12 = 2$.

We can deal with subtractions and negative numbers too:

- $3 - 4 = -1 = -1 + 5 = 4$
- $4 \times (-3) = -12 = -12 + 15 = 3$.

4.1.2 Addition and multiplication tables

There's an easy way to summarize all the possible additions and multiplications in \mathbb{F}_5 : draw tables with rows and columns labelled by 0, 1, 2, 3, 4, and in row i and column j put the result of $i + j$ or $i \times j$.

+	0	1	2	3	4
0	0	1	2	3	4
1	1	2	3	4	0
2	2	3	4	0	1
3	3	4	0	1	2
4	4	0	1	2	3

Table 4.1: Addition table for \mathbb{F}_5 .

×	0	1	2	3	4
0	0	0	0	0	0
1	0	1	2	3	4
2	0	2	4	1	3
3	0	3	1	4	2
4	0	4	3	2	1

Table 4.2: Multiplication table for \mathbb{F}_5 .

Here are the addition and multiplication tables for \mathbb{F}_3 .

+	0	1	2
0	0	1	2
1	1	2	0
2	2	0	1

Table 4.3: Addition table for \mathbb{F}_3 .

×	0	1	2
0	0	0	0
1	0	1	2
2	0	2	1

Table 4.4: Multiplication table for \mathbb{F}_3 .

4.1.3 Multiplicative inverses

In a field, you have to be able to divide by non-zero elements. Being able to divide by x is equivalent to the existence of an element $1/x$ or x^{-1} , called a

multiplicative inverse to x (because y/x is just $y \times 1/x$). So to check that we can divide, we need to check that for each nonzero x there is another element which you can multiply x by to get 1. We will do this for the examples of \mathbb{F}_3 and \mathbb{F}_5 .

In \mathbb{F}_3 , you can see from the table that

- $1 \times 1 = 1$, so $1^{-1} = 1$
- $2 \times 2 = 1$, so $2^{-1} = 2$

and so every non-zero element of \mathbb{F}_3 has a multiplicative inverse.

In \mathbb{F}_5 , you can see from the table that

- $1 \times 1 = 1$, so $1^{-1} = 1$
- $2 \times 3 = 1$, so $2^{-1} = 3$ and $3^{-1} = 2$
- $4 \times 4 = 1$, so $4^{-1} = 4$.

and so every non-zero element of \mathbb{F}_5 has a multiplicative inverse.

We're going to give a proof that every element of \mathbb{F}_p has a multiplicative inverse and give a method for finding these multiplicative inverses.

First, here is a reminder about division with remainder. There is a fundamental fact about the natural numbers: if you have any two positive whole numbers x and y , you can try to divide x by y . When you do that you get a quotient q and a remainder r , which is a nonnegative integer less than y :

$$x = qy + r \qquad 0 \leq r < y$$

It's easy to convince yourself of the truth of this by imagining starting at zero on the number line and taking steps of size y to the right, stopping when the next step would take you past x . Call the number of steps taken at the point you stop q . You've stopped at the number qy , and the distance from this point to x , which is the remainder $r = x - qy$, must be less than y since otherwise you could take another full step without passing x .

Theorem 4.1.1. *Let p be prime and $0 < a < p$ be a whole number. Then there are whole numbers s and t such that $as + pt = 1$.*

Proof. Let g be the smallest positive number that can be written as an integer multiple of p plus an integer multiple of a , so $g = as + pt$ for some integers s and t . Note that $g < p$ since $p - a$ is a positive number that can be written this way. Divide p by g to get quotient q and remainder r , so $p = qg + r$ where $0 \leq r < g$. Then

$$\begin{aligned} r &= p - qg \\ &= p - q(as + pt) \\ &= (1 - qt)p - qsa. \end{aligned}$$

Since $0 \leq r < g$ and g is the smallest *positive* number that can be written as a multiple of p plus a multiple of a we must have $r = 0$, that is, g divides the prime number p . Since $g < p$ we have $g = 1$. \square

We can even calculate numbers s and t satisfying $as + pt = 1$ quickly. To do this divide p by a , getting a quotient q_2 and a remainder r_2

$$p = q_2a + r_2 \quad 0 \leq r_2 < a$$

(we start with a subscript 2 here because we will think of p as r_0 and a as r_1). If $r_2 = 0$, stop. Otherwise divide a by r_2 getting a quotient q_3 and remainder r_3

$$a = q_3r_2 + r_3 \quad 0 \leq r_3 < r_2$$

If $r_3 = 0$, stop. Otherwise divide r_2 by r_3 getting a quotient q_4 and remainder r_4 :

$$r_2 = q_4r_3 + r_4 \quad 0 \leq r_4 < r_3$$

If $r_4 = 0$, stop. Otherwise divide r_3 by r_4 getting quotient q_5 and remainder r_5 :

$$r_3 = q_5r_4 + r_5 \quad 0 \leq r_5 < r_4$$

and so on. The sequence r_2, r_3, \dots satisfies

$$a > r_2 > r_3 > r_4 > r_5 > \dots \geq 0$$

You can't have a decreasing sequence of positive integers that goes on forever, so there is some m such that $r_m = 0$. The last few divisions were

$$r_{m-5} = q_{m-3}r_{m-4} + r_{m-3} \tag{4.1}$$

$$r_{m-4} = q_{m-2}r_{m-3} + r_{m-2} \tag{4.2}$$

$$r_{m-3} = q_{m-1}r_{m-2} + r_{m-1} \tag{4.3}$$

$$r_{m-2} = q_m r_{m-1} + 0 \tag{4.4}$$

We claim that $r_{m-1} = 1$. For r_{m-1} divides r_{m-2} , because of the last equation. In (4.3) the terms on the right hand side are multiples of r_{m-1} , so the left hand side r_{m-3} is a multiple of r_{m-1} as well. Repeating the same argument we end up with r_{m-1} dividing all the left-hand-sides, in particular, r_{m-1} divides the prime number p . Since $r_{m-1} < a < p$ we have $r_{m-1} = 1$. So (4.3) is really

$$r_{m-3} = q_{m-1}r_{m-2} + 1$$

or equivalently

$$r_{m-3} - q_{m-1}r_{m-2} = 1. \tag{4.5}$$

Now we can express r_{m-3} in terms of rs with a smaller subscript with equation (4.1), and we can express r_{m-2} in terms of rs with a smaller subscript using (4.2). When we substitute this in (4.5), we get $1 =$ some multiple of r_{m-3} plus some multiple of r_{m-4} . And we can keep doing that over and over again until we eventually get $1 =$ a multiple of r_1 plus a multiple of r_0 , that is, $as + pt$ for some whole numbers s and t .

It helps to see an example. Take $p = 13$ and $a = 5$. We have

$$13 = 2 \times 5 + 3$$

$$5 = 1 \times 3 + 2$$

$$3 = 1 \times 2 + 1$$

$$2 = 2 \times 1 + 0$$

and so

$$\begin{aligned} 1 &= 3 - 1 \times 2 \\ &= (13 - 2 \times 5) - 1 \times (5 - 1 \times 3) \\ &= (13 - 2 \times 5) - 1 \times (5 - 1 \times (13 - 2 \times 5)) \\ &= 2 \times 13 - 5 \times 5 \end{aligned}$$

This helps us find multiplicative inverses in \mathbb{F}_p because if $as + pt = 1$ then in \mathbb{F}_p we have, since multiples of p are equal to zero, $as = 1$ and s is the multiplicative inverse of a . In our example, $p = 13, a = 5, s = -5, t = 2$, so the multiplicative inverse of 5 in \mathbb{F}_{13} is -5 which is the same as 8.

4.2 Vector spaces

We are now ready to define vector spaces. The idea is to observe that sets of column vectors, or row vectors, or more generally matrices of a given size, all come equipped with a notion of addition and scalar multiplication and all obey the same collection of simple algebraic rules, for example, that addition is commutative, that scalar multiplication distributes over vector addition, and so on. We will define a vector space as any set with operations of addition and scalar multiplication obeying similar rules to those satisfied by column vectors. The power of doing this is that it lets us apply our theory in seemingly entirely different contexts.

4.2.1 The vector space axioms

Definition 4.2.1. Let \mathbb{F} be a field. An \mathbb{F} -vector space is a set V with

- a special element $\mathbf{0}_V$ called the **zero vector**
- an operation $+$ called addition
- a way to multiply elements of V by elements of \mathbb{F} , called **scalar multiplication**

such that for all $\mathbf{u}, \mathbf{v}, \mathbf{w}$ in V and all λ, μ in \mathbb{F} ,

1. $\mathbf{v} + \mathbf{w} = \mathbf{w} + \mathbf{v}$
2. $\mathbf{u} + (\mathbf{v} + \mathbf{w}) = (\mathbf{u} + \mathbf{v}) + \mathbf{w}$
3. $\mathbf{0}_V + \mathbf{v} = \mathbf{v}$
4. there exists $\mathbf{x} \in V$ such that $\mathbf{x} + \mathbf{v} = \mathbf{0}_V$
5. $\lambda(\mu\mathbf{v}) = (\lambda\mu)\mathbf{v}$
6. $1\mathbf{v} = \mathbf{v}$
7. $\lambda(\mathbf{v} + \mathbf{w}) = \lambda\mathbf{v} + \lambda\mathbf{w}$
8. $(\lambda + \mu)\mathbf{v} = \lambda\mathbf{v} + \mu\mathbf{v}$

You sometimes see the phrase “vector space over \mathbb{F} ”, which means the same thing as \mathbb{F} -vector space.

4.2.2 Examples of vector spaces

The elements of vector spaces can be anything at all. They don't have to look like column or row vectors. Here are some examples of vector spaces.

- \mathbb{R}^n is a real vector space, \mathbb{C}^n is a complex vector space, and if \mathbb{F} is any field then \mathbb{F}^n , the set of all height n column vectors with entries from \mathbb{F} is an \mathbb{F} -vector space.
- $M_{m \times n}(\mathbb{R})$, the set of all $m \times n$ matrices with real entries, is a real vector space with the zero vector being the all-zeroes matrix. Similarly for any other field.
- $\{0\}$ with the only possible operations is an \mathbb{F} -vector space, for any field \mathbb{F} , the zero vector space.
- Let \mathcal{F} be the set of all functions $\mathbb{R} \rightarrow \mathbb{R}$. Define $f + g$ to be the function $\mathbb{R} \rightarrow \mathbb{R}$ given by $(f + g)(x) = f(x) + g(x)$ and, for a real number λ and a function f , define λf by $(\lambda f)(x) = \lambda f(x)$. Then \mathcal{F} is a real vector space with the zero vector being the constant function taking the value 0.
- If A is a $m \times n$ matrix, the set of all solutions of $A\mathbf{x} = 0$ is a vector space. This is the nullspace $N(A)$ we met in Definition 3.6.5.
- The set of all real solutions to the differential equation $y'' + ay' + by = 0$ is a vector space, with the definitions of addition and scalar multiplication as in \mathcal{F} above.
- The set $\mathbb{F}[x]$ of all polynomials in one variable x is a \mathbb{F} -vector space, as is the set $\mathbb{F}_{\leq n}[x]$ of all polynomials in x of degree at most n .
- The set of magic matrices, those whose row sums and column sums are all equal, is a vector space with the usual matrix addition and scalar multiplication.

4.3 Using the vector space axioms

There are some familiar properties of vector addition and scalar multiplication — like the fact that if you multiply a vector by the scalar zero, you get the zero vector — which aren't listed in the axioms. Are they special to column vectors, or do they hold in every vector space?

To answer questions like this we can give a proof that uses only the vector space axioms, not the specific form of a particular vector space's elements.

Lemma 4.3.1. *Let V be a vector space and $\mathbf{v} \in V$. Then $0\mathbf{v} = \mathbf{0}_V$.*

Be careful that you understand the notation here. $\mathbf{0}_V$ means the special zero vector given in the definition of the vector space V , and $0\mathbf{v}$ means the vector \mathbf{v} scalar multiplied by the scalar 0. They're not *obviously* the same thing.

Proof. $0\mathbf{v} = (0 + 0)\mathbf{v} = 0\mathbf{v} + 0\mathbf{v}$ (by axiom 8 of Definition 4.2.1). Axiom 4 says there's an element \mathbf{u} of V such that $\mathbf{u} + 0\mathbf{v} = \mathbf{0}_V$, so add it to both sides:

$$\begin{aligned} \mathbf{u} + 0\mathbf{v} &= \mathbf{u} + (0\mathbf{v} + 0\mathbf{v}) \\ \mathbf{0}_V &= (\mathbf{u} + 0\mathbf{v}) + 0\mathbf{v} && \text{axiom 2, definition of } \mathbf{u} \\ \mathbf{0}_V &= \mathbf{0}_V + 0\mathbf{v} && \text{definition of } \mathbf{u} \\ \mathbf{0}_V &= 0\mathbf{v} && \text{axiom 3.} \end{aligned}$$

□

Lemma 4.3.2. *Let V be a vector space and let $\mathbf{x} \in V$. Then $\mathbf{x} + (-1)\mathbf{x} = \mathbf{0}_V$.*

Proof.

$$\begin{aligned} \mathbf{x} + (-1)\mathbf{x} &= 1\mathbf{x} + (-1)\mathbf{x} && \text{axiom 6} \\ &= (1 + (-1))\mathbf{x} && \text{axiom 8} \\ &= 0\mathbf{x} \\ &= \mathbf{0}_V && \text{Lemma 4.3.1.} \end{aligned}$$

□

We write $-\mathbf{x}$ for the additive inverse of \mathbf{x} which axiom 4 provides, and $\mathbf{y} - \mathbf{x}$ as shorthand for $\mathbf{y} + -\mathbf{x}$. Here are two more proofs using the axioms.

Lemma 4.3.3. 1. *Let λ be a scalar. Then $\lambda\mathbf{0}_V = \mathbf{0}_V$.*

2. *Suppose $\lambda \neq 0$ is a scalar and $\lambda\mathbf{x} = \mathbf{0}_V$. Then $\mathbf{x} = \mathbf{0}_V$.*

Proof. 1.

$$\begin{aligned} \lambda\mathbf{0}_V &= \lambda(\mathbf{0}_V + \mathbf{0}_V) && \text{axiom 3} \\ &= \lambda\mathbf{0}_V + \lambda\mathbf{0}_V && \text{axiom 7} \end{aligned}$$

Axiom 2 tells there's an additive inverse to $\lambda\mathbf{0}_V$. Adding it to both sides and using axiom 2, we get $\mathbf{0}_V = \lambda\mathbf{0}_V$.

2.

$$\begin{aligned} \lambda\mathbf{x} &= \mathbf{0}_V \\ \lambda^{-1}(\lambda\mathbf{x}) &= \lambda^{-1}\mathbf{0}_V \\ (\lambda^{-1}\lambda)\mathbf{x} &= \mathbf{0}_V && \text{axiom 5 and part 1} \\ 1\mathbf{x} &= \mathbf{0}_V \\ \mathbf{x} &= \mathbf{0}_V && \text{axiom 6.} \end{aligned}$$

□

4.4 Subspaces

When we talk about a vector space over a field \mathbb{F} , the word *scalar* refers to an element of \mathbb{F} .

Definition 4.4.1. A **subspace** of a vector space V is a subset U of V which

1. contains the zero vector $\mathbf{0}_V$,
2. is *closed under addition*, meaning that for all $\mathbf{v}, \mathbf{w} \in U$ we have $\mathbf{v} + \mathbf{w} \in U$, and
3. is *closed under scalar multiplication*, meaning that for all scalars λ and all $\mathbf{u} \in U$ we have $\lambda\mathbf{u} \in U$.

We write $U \leq V$ to mean that U is a subspace of V .

The idea this definition captures is that a subspace of V is a nonempty subset which is itself a vector space under the same addition and scalar multiplication as V .

If $U \leq V$ and $\mathbf{u}_1, \dots, \mathbf{u}_n \in U$ and $\lambda_1, \dots, \lambda_n$ are scalars then $\sum_{i=1}^n \lambda_i \mathbf{u}_i \in U$. This follows by using closure under scalar multiplication and closure under addition lots of times.

4.4.1 Subspace examples

Example 4.4.1. If V is any vector space, $V \leq V$. This is because, as a vector space, V contains the zero vector, is closed under addition, and is closed under scalar multiplication.

A subspace of V other than V is called a **proper subspace**.

Example 4.4.2. For any vector space V we have $\{\mathbf{0}_V\} \leq V$. Certainly this set contains the zero vector. It is closed under addition because $\mathbf{0}_V + \mathbf{0}_V = \mathbf{0}_V$, and it is closed under scalar multiplication by Lemma 4.3.3. This is called the zero subspace.

Example 4.4.3. Let U be the set of vectors in \mathbb{R}^2 whose first entry is zero. Then $U \leq \mathbb{R}^2$. We check the three conditions in the definition of subspace.

1. The zero vector in \mathbb{R}^2 is $\begin{pmatrix} 0 \\ 0 \end{pmatrix}$. This has first coordinate 0, so it is an element of U .
2. Let $\mathbf{v}, \mathbf{w} \in U$, so that $\mathbf{v} = \begin{pmatrix} 0 \\ x \end{pmatrix}$ and $\mathbf{w} = \begin{pmatrix} 0 \\ y \end{pmatrix}$ for some real numbers x and y . Then $\mathbf{v} + \mathbf{w} = \begin{pmatrix} 0 \\ x + y \end{pmatrix}$ has first coordinate 0, so it is an element of U .
3. Let \mathbf{v} be as above and $\lambda \in \mathbb{R}$. Then $\lambda\mathbf{v} = \begin{pmatrix} 0 \\ \lambda x \end{pmatrix}$ which has first coordinate 0, so $\lambda\mathbf{v} \in U$.

All three conditions hold, so $U \leq \mathbb{R}^2$. Of course, a similar argument shows the vectors in \mathbb{F}^n with first entry 0 are a subspace of \mathbb{F}^n for any field \mathbb{F} and any n .

To every matrix A we associate two important subspaces. The nullspace $N(A)$ (Definition 3.6.5) is the set of all vectors \mathbf{x} such that $A\mathbf{x} = \mathbf{0}$, and the column space $C(A)$ is the set of all linear combinations of the columns of A .

Example 4.4.4. Let A be an $m \times n$ matrix with entries from the field \mathbb{F} . The nullspace $N(A)$ contains the zero vector as $A\mathbf{0}_n = \mathbf{0}_m$. It is closed under addition as if $\mathbf{u}, \mathbf{v} \in N(A)$ then $A\mathbf{v} = \mathbf{0}_m$ and $A\mathbf{u} = \mathbf{0}_m$ so

$$\begin{aligned} A(\mathbf{u} + \mathbf{v}) &= A\mathbf{u} + A\mathbf{v} \\ &= \mathbf{0}_m + \mathbf{0}_m \\ &= \mathbf{0}_m \end{aligned}$$

and therefore $\mathbf{u} + \mathbf{v} \in N(A)$. It is closed under scalar multiplication because if λ is any scalar then $A(\lambda\mathbf{u}) = \lambda A\mathbf{u} = \lambda\mathbf{0}_m = \mathbf{0}_m$ so $\lambda\mathbf{u} \in N(A)$. Therefore $N(A) \leq \mathbb{F}^n$.

The column space $C(A)$, defined to be the set of all linear combinations of the columns of A , is a subspace of \mathbb{F}^m . We won't prove that here, because it is a special case of Proposition 4.7.1 which we prove later.

Example 4.4.5. The set U of all vectors in \mathbb{R}^3 with first entry 1 is not a subspace of \mathbb{R}^3 . It doesn't contain the zero vector (and it doesn't meet the other two conditions either).

Example 4.4.6. \mathbb{Z} is not a subspace of \mathbb{R} . It contains the zero vector 0, it is closed under addition because if you add two integers you get another integer. But it is not closed under scalar multiplication: $\sqrt{2}$ is a scalar, $1 \in \mathbb{Z}$, but $\sqrt{2} \times 1$ is not in \mathbb{Z} .

Example 4.4.7. Let U be the set of all functions $f : \mathbb{R} \rightarrow \mathbb{R}$ with $f(1) = 0$. This is a subspace of the vector space \mathcal{F} of all functions $\mathbb{R} \rightarrow \mathbb{R}$. The zero vector in \mathcal{F} is the constant function that always takes the value zero, so certainly it belongs to U . If $f, g \in U$ then $(f+g)(1) = f(1) + g(1) = 0 + 0 = 0$, so $f+g \in U$. If $\lambda \in \mathbb{R}$ and $f \in U$ then $(\lambda f)(1) = \lambda f(1) = \lambda \times 0 = 0$ so $\lambda f \in U$.

Example 4.4.8. $\{A \in M_{n \times n}(\mathbb{R}) : A^T = A\} \leq M_{n \times n}(\mathbb{R})$. The transpose operation satisfies $(A+B)^T = A^T + B^T$ and $(\lambda A)^T = \lambda A^T$, which you should check. This makes the three conditions straightforward to check.

Example 4.4.9. $U = \{A \in M_{n \times n}(\mathbb{R}) : A^2 = \mathbf{0}_{m \times n}\}$ is not a subspace of $M_{n \times n}(\mathbb{R})$. For example, U contains $E_{12} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ and $E_{21} = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$ but you can check that $E_{12} + E_{21} \notin U$.

4.5 Sums and intersections

Proposition 4.5.1. Let X and Y be subspaces of a vector space V .

1. $X \cap Y \leq V$.
2. $X + Y = \{\mathbf{x} + \mathbf{y} : \mathbf{x} \in X, \mathbf{y} \in Y\} \leq V$.

Proof. To show something is a subspace we have to check the three properties: containing the zero vector, closure under addition, and closure under scalar multiplication.

1. • $\mathbf{0}_V \in X \cap Y$ as X and Y are subspaces so contain $\mathbf{0}_V$.

- Let $\mathbf{x}, \mathbf{y} \in X \cap Y$. X is a subspace, so closed under addition, so $\mathbf{x} + \mathbf{y} \in X$. For the same reason $\mathbf{x} + \mathbf{y} \in Y$. Therefore $\mathbf{x} + \mathbf{y} \in X \cap Y$.
 - Let λ be a scalar and $\mathbf{x} \in X \cap Y$. X is a subspace, so closed under scalar multiplication, so $\lambda\mathbf{x} \in X$. For the same reason $\lambda\mathbf{x} \in Y$. Therefore $\lambda\mathbf{x} \in X \cap Y$.
2. • $\mathbf{0}_V$ is in X and Y as they are subspaces, so $\mathbf{0}_V + \mathbf{0}_V = \mathbf{0}_V$ is in $X + Y$.
- Any two elements of $X + Y$ have the form $\mathbf{x}_1 + \mathbf{y}_1$ and $\mathbf{x}_2 + \mathbf{y}_2$, where $\mathbf{x}_i \in X$ and $\mathbf{y}_i \in Y$.

$$(\mathbf{x}_1 + \mathbf{y}_1) + (\mathbf{x}_2 + \mathbf{y}_2) = (\mathbf{x}_1 + \mathbf{x}_2) + (\mathbf{y}_1 + \mathbf{y}_2)$$

by associativity and commutativity. But $\mathbf{x}_1 + \mathbf{x}_2 \in X$ as X is a subspace and $\mathbf{y}_1 + \mathbf{y}_2 \in Y$ as Y is a subspace, so this is in $X + Y$ which is therefore closed under addition.

- Let λ be a scalar.

$$\lambda(\mathbf{x}_1 + \mathbf{y}_1) = \lambda\mathbf{x}_1 + \lambda\mathbf{y}_1$$

$\lambda\mathbf{x}_1 \in X$ as X is a subspace so closed under scalar multiplication, $\lambda\mathbf{y}_1 \in Y$ for the same reason, so their sum is in $X + Y$ which is therefore closed under scalar multiplication. \square

4.6 Linear independence

4.6.1 Linear combinations

We met the idea of a linear combination of matrices in chapter 3. We will now expand the definition so that we can talk about linear combinations of elements of any vector space.

Definition 4.6.1. Let V be a vector space and $\mathbf{v}_1, \dots, \mathbf{v}_n \in V$. A **linear combination** of $\mathbf{v}_1, \dots, \mathbf{v}_n$ is an element of V of the form

$$\lambda_1\mathbf{v}_1 + \lambda_2\mathbf{v}_2 + \dots + \lambda_n\mathbf{v}_n$$

where the λ_i are scalars.

4.6.2 Linear independence

Definition 4.6.2. Let V be a vector space.

- A sequence $\mathbf{v}_1, \dots, \mathbf{v}_n$ of elements of V is **linearly independent** if and only if the only scalars $\lambda_1, \dots, \lambda_n$ such that $\sum_{i=1}^n \lambda_i\mathbf{v}_i = \mathbf{0}_V$ are $\lambda_1 = \dots = \lambda_n = 0$.
- A sequence which is not linearly independent is called **linearly dependent**.

It is important that linear independence is a property of *sequences* (not sets) of vectors. Sequences have a particular order, and they can contain the same element multiple times.

Checking whether elements of a vector space are linearly independent is simple. You just have to try and find a linear combination that gives the zero vector where not all the scalars are zero. If you can do it, the sequence is linearly dependent, if you can't it is linearly independent. When we're talking about vectors in \mathbb{F}^n , or matrices, this is just solving linear equations.

4.6.3 Examples of linear (in)dependence

Example 4.6.1. $\mathbf{u} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, $\mathbf{v} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, $\mathbf{w} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ are not linearly independent in \mathbb{R}^2 , because $1 \times \mathbf{u} + 1 \times \mathbf{v} + (-1) \times \mathbf{w} = \mathbf{0}$.

Example 4.6.2. $\mathbf{u} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$, $\mathbf{v} = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ are linearly independent in \mathbb{R}^2 . For if $\alpha\mathbf{u} + \beta\mathbf{v} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$ then $\begin{pmatrix} \alpha + \beta \\ \alpha - \beta \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$. This is a system of linear equations:

$$\begin{aligned}\alpha + \beta &= 0 \\ \alpha - \beta &= 0\end{aligned}$$

For such a simple system it's easy to see that the only solution is $\alpha = \beta = 0$. This tells you that the only solution to $\alpha\mathbf{u} + \beta\mathbf{v} = \mathbf{0}$ is $\alpha = \beta = 0$, which is the definition of linear independence for \mathbf{u}, \mathbf{v} .

Example 4.6.3. $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ are linearly independent in \mathbb{R}^2 . You can prove this in a similar (but easier) way to the previous example.

More generally if \mathbf{e}_i is the height n column vector with 0 everywhere except 1 at position i , then the sequence $\mathbf{e}_1, \dots, \mathbf{e}_n$ is linearly independent.

Example 4.6.4. In \mathcal{F} , the vector space of all functions $\mathbb{R} \rightarrow \mathbb{R}$, let's show that the functions $f(x) = \cos(x)$ and $g(x) = \sin(x)$ are linearly independent. Suppose that $\alpha f + \beta g = 0_{\mathcal{F}}$, that is, suppose $\alpha \cos(x) + \beta \sin(x) = 0$ for all x .

Take $x = 0$. Since $\alpha \cos(0) + \beta \sin(0) = 0$ we get $\alpha = 0$. Now take $x = \pi/2$ to get $\beta \sin(\pi/2) = 0$, that is $\beta = 0$. We have shown $\alpha = \beta = 0$ and so these functions are linearly independent.

Often it turns out that deciding whether a sequence of vectors is linearly independent is equivalent to seeing whether a system of linear equations has only the solution where every variable is zero — so you can apply the methods we learned in chapter 3.

4.7 Spanning sequences

4.7.1 Definition of span

Definition 4.7.1. Let V be an \mathbb{F} -vector space and $\mathbf{v}_1, \dots, \mathbf{v}_n \in V$. The **span** of $\mathbf{v}_1, \dots, \mathbf{v}_n$, written $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_n)$ is the set of all linear combinations of

$\mathbf{v}_1, \dots, \mathbf{v}_n$, so

$$\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_n) = \{\lambda_1 \mathbf{v}_1 + \dots + \lambda_n \mathbf{v}_n : \lambda_1, \dots, \lambda_n \in \mathbb{F}\}.$$

For technical reasons we define the span of the empty sequence of vectors to be $\{\mathbf{0}_V\}$.

To understand the definition a bit better, let's look at two simple special cases. The span of a single element \mathbf{s} of an \mathbb{F} -vector space V is

$$\{\lambda \mathbf{s} : \lambda \in \mathbb{F}\},$$

since any linear combination of \mathbf{s} is just a scalar multiple of \mathbf{s} . The span of two elements \mathbf{u}, \mathbf{v} of V is

$$\{a\mathbf{u} + b\mathbf{v} : a, b \in \mathbb{F}\}.$$

4.7.2 Spans are subspaces

Proposition 4.7.1. *If $\mathbf{s}_1, \dots, \mathbf{s}_n$ are elements of a vector space V then $\text{span}(\mathbf{s}_1, \dots, \mathbf{s}_n)$ is a subspace of V .*

Proof. Write S for $\text{span}\{\mathbf{s}_1, \dots, \mathbf{s}_n\}$. Recall that S consists of every linear combination $\sum_{i=1}^n \lambda_i \mathbf{s}_i$, where the λ_i are scalars.

1. S contains the zero vector because it contains $\sum_{i=1}^n 0\mathbf{s}_i$, and each $0\mathbf{s}_i$ is the zero vector.
2. S is closed under addition because if $\sum_{i=1}^n \lambda_i \mathbf{s}_i$ and $\sum_{i=1}^n \mu_i \mathbf{s}_i$ are any two elements of S then

$$\sum_{i=1}^n \lambda_i \mathbf{s}_i + \sum_{i=1}^n \mu_i \mathbf{s}_i = \sum_{i=1}^n (\lambda_i + \mu_i) \mathbf{s}_i$$

is in S .

3. S is closed under scalar multiplication because if $\sum_{i=1}^n \lambda_i \mathbf{s}_i$ is in S and λ is a scalar then

$$\lambda \sum_{i=1}^n \lambda_i \mathbf{s}_i = \sum_{i=1}^n (\lambda \lambda_i) \mathbf{s}_i$$

is also in S .

S fulfils all three conditions in the Definition 4.4.1 of a subspace, so $S \leq V$. \square

4.7.3 Spanning sequences

Definition 4.7.2. Elements $\mathbf{v}_1, \dots, \mathbf{v}_n$ of a vector space V are a **spanning sequence** for V if and only if $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_n) = V$.

The term *spanning set* is also used.

We also say $\mathbf{v}_1, \dots, \mathbf{v}_n$ spans V to mean that it is a spanning sequence.

Often deciding whether or not a sequence of vectors is a spanning sequence is equivalent to solving some linear equations.

Example 4.7.1. If you want to check whether $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} -1 \\ 1 \end{pmatrix}$ are a spanning sequence for \mathbb{R}^2 , what you need to do is to verify that for every $\begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{R}^2$ there are real numbers α and β such that

$$\alpha \begin{pmatrix} 1 \\ 1 \end{pmatrix} + \beta \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

In other words, you have to prove that for **every** $x, y \in \mathbb{R}$ the system of linear equations

$$\begin{aligned} \alpha + \beta &= x \\ \alpha - \beta &= y \end{aligned}$$

has a solution. That's easy in this case, because you can just notice that $\alpha = (x+y)/2, \beta = (x-y)/2$ is a solution, but for bigger and more complicated systems you can use the method of RREF.

Example 4.7.2. $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} -1 \\ 1 \end{pmatrix}$ are a spanning sequence for \mathbb{R}^2 , as we have just seen.

Example 4.7.3. Let's try to determine whether $\mathbf{v}_1 = \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}, \mathbf{v}_2 = \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix}, \mathbf{v}_3 = \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}$ are a spanning sequence for \mathbb{R}^3 . We need to find out whether it's true that for all $\begin{pmatrix} x \\ y \\ z \end{pmatrix} \in \mathbb{R}^3$ there exist $\alpha, \beta, \gamma \in \mathbb{R}^3$ such that

$$\alpha \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix} + \beta \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix} + \gamma \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} = \begin{pmatrix} x \\ y \\ z \end{pmatrix}.$$

This is equivalent to asking whether for every x, y, z the simultaneous equations

$$\begin{aligned} \alpha + \gamma &= x \\ -\alpha + \beta &= y \\ -\beta - \gamma &= z \end{aligned}$$

have a solution. Again, in this special case you might just notice that (adding the three equations) there is no solution unless $x+y+z=0$, so this collection of vectors is **not** a spanning sequence. In general, to find out if a system of linear equations has a solution you can put the augmented matrix into row reduced echelon form. In this case the augmented matrix is

$$\begin{pmatrix} 1 & 0 & 1 & x \\ -1 & 1 & 0 & y \\ 0 & -1 & -1 & z \end{pmatrix}$$

Doing the row operations $r_2 \mapsto r_2 + r_1$ followed by $r_3 \mapsto r_3 + r_2$ leads to

$$\begin{pmatrix} 1 & 0 & 1 & x \\ 0 & 1 & 1 & y+x \\ 0 & 0 & 0 & z+y+x \end{pmatrix}$$

These equations have no solutions if $x + y + z \neq 0$, so for example $\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$ is not in the span of $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ because $1 + 0 + 0 \neq 0$.

4.8 Bases

4.8.1 Basis definition

Definition 4.8.1. A sequence $\mathbf{v}_1, \dots, \mathbf{v}_n$ of elements of a vector space V is a **basis** for V if and only if

1. it is linearly independent, and
2. it is a spanning sequence for V .

Importantly, bases are sequences not sets. This is because the order of a basis matters to some of the definitions we will make later, like the matrix of a linear map.

4.8.2 The standard basis for \mathbb{F}^n

The most important example is the **standard basis** of \mathbb{F}^n (no matter which field \mathbb{F} is). Let \mathbf{e}_i be the column vector in \mathbb{F}^n with a 1 in position i and 0s elsewhere. When $n = 3$, for example, we have

$$\mathbf{e}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \mathbf{e}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \mathbf{e}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

Then $\mathbf{e}_1, \dots, \mathbf{e}_n$ is a basis of \mathbb{F}^n , called the standard basis. To check this, we must verify the two parts of the definition of basis.

1. (Linear independence). Suppose $\sum_{i=1}^n \lambda_i \mathbf{e}_i = \mathbf{0}$. To verify linear independence we have to prove all the λ_i are zero. Using the definition of the \mathbf{e}_i

we get $\begin{pmatrix} \lambda_1 \\ \vdots \\ \lambda_n \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}$. So $\lambda_i = 0$ for all i as required.

2. (Spanning) We have to show that any element of \mathbb{F}^n is a linear combination of $\mathbf{e}_1, \dots, \mathbf{e}_n$. Let $\mathbf{v} = \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix} \in \mathbb{F}^n$. Then $\mathbf{v} = \sum_{i=1}^n v_i \mathbf{e}_i$, so \mathbf{v} is a linear combination of the \mathbf{e}_i as required.

4.8.3 More basis examples

Example 4.8.1. $\mathbb{R}_{\leq 3}[x]$ consists of all polynomials of degree at most 3 in the variable x . It has a basis $1, x, x^2, x^3$, because

- (linear independence) if $a + bx + cx^2 + dx^3$ is the zero polynomial, that is if it is zero for every value of x , then $a = b = c = d = 0$. This is because a polynomial of degree m has at most m roots.
- (spanning) every polynomial of degree at most 3 has the form $a + bx + cx^2 + dx^3$ for some a, b, c, d , and so is a linear combination of $1, x, x^2, x^3$.

Example 4.8.2. Let $V = M_{m \times n}(\mathbb{F})$ be the \mathbb{F} -vector space of all $m \times n$ matrices. Let E_{ij} be the matrix which has a 1 in position i, j and zeroes elsewhere. Then

$$E_{11}, E_{21}, \dots, E_{m1}, E_{12}, E_{22}, \dots, E_{mn}$$

is a basis for V . This can be proved in exactly the same way as we proved that the standard basis of \mathbb{F}^n really was a basis.

4.8.4 What is a basis good for?

Lemma 4.8.1. *If $\mathbf{v}_1, \dots, \mathbf{v}_n$ is a basis of V , every $\mathbf{v} \in V$ can be written uniquely as $\sum_{i=1}^n \lambda_i \mathbf{v}_i$ for some scalars λ_i .*

Proof. Every $\mathbf{v} \in V$ can be written this way because the \mathbf{v}_i are a basis and hence a spanning sequence for V . The problem is to prove that every $\mathbf{v} \in V$ can be written like this in only one way.

Suppose that

$$\sum_{i=1}^n \lambda_i \mathbf{v}_i = \sum_{i=1}^n \mu_i \mathbf{v}_i.$$

Then subtracting one side from the other,

$$\sum_{i=1}^n (\lambda_i - \mu_i) \mathbf{v}_i = \mathbf{0}_V.$$

Linear independence of the \mathbf{v}_i tells us $\lambda_i - \mu_i = 0$ for all i , so $\lambda_i = \mu_i$ for all i . We have proved that there is only one expression for \mathbf{v} as a linear combination of the elements of the basis $\mathbf{v}_1, \dots, \mathbf{v}_n$. \square

This means that a basis gives a way of giving coordinates to an arbitrary vector space, no matter what the elements look like. Once we fix a basis of V , there is a one-one correspondence between the elements of V and the coefficients needed to express them in terms of that basis — you could call these the coordinates of the vector in terms of this basis.

A basis also allows us to compare coefficients. Suppose $\mathbf{v}_1, \dots, \mathbf{v}_n$ is a basis of a vector space V and that

$$\lambda_1 \mathbf{v}_1 + \dots + \lambda_n \mathbf{v}_n = \mu_1 \mathbf{v}_1 + \dots + \mu_n \mathbf{v}_n.$$

Then the uniqueness result Lemma 4.8.1 tells us we can compare coefficients to get that $\lambda_1 = \mu_1$, $\lambda_2 = \mu_2$, and so on.

4.8.5 Multiple bases for the same vector space

A given vector space can have many different bases. This is true in a trivial sense: as we saw before, basis are sequences, the order matters, so $\begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is different to $\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ but clearly still a basis of \mathbb{R}^2 . But it is also true in a more interesting way. Take \mathbb{R}^2 , for example: we know $\mathbf{e}_1, \mathbf{e}_2$ is a basis, but so also is

$$\mathbf{u} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \mathbf{v} = \begin{pmatrix} 1 \\ -1 \end{pmatrix}.$$

Let's check this. Suppose $a\mathbf{u} + b\mathbf{v} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$. Then $\begin{pmatrix} a+b \\ a-b \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$, so $a+b = 0 = a-b$ from which it follows $a = b = 0$ and \mathbf{u}, \mathbf{v} is linearly independent. To show \mathbf{u}, \mathbf{v} spans \mathbb{R}^2 , let $\begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{R}^2$. We must show there exist $a, b \in \mathbb{R}$ such that $a\mathbf{u} + b\mathbf{v} = \begin{pmatrix} x \\ y \end{pmatrix}$. The condition a and b must satisfy is $a+b = x, a-b = y$. It is always possible to find such a and b : solving the equations you get $a = (x+y)/2, b = (x-y)/2$, so \mathbf{u}, \mathbf{v} spans \mathbb{R}^2 .

Here's why a vector space having several different bases is useful. The expression of an element \mathbf{v} in terms of different bases can tell us different things about \mathbf{v} . In other words, different bases give different ways of looking at the elements of the vector space.

Say for example you are representing an image as an element of \mathbb{R}^n . The smallest possible example is a 2-pixel image which we could represent as an element $\begin{pmatrix} a \\ b \end{pmatrix} = a\mathbf{e}_1 + b\mathbf{e}_2$ in \mathbb{R}^2 , where the first coordinate tells me how bright the first pixel is and the second tells me how bright the second is.

Now consider the alternative basis $\mathbf{e}_1 + \mathbf{e}_2, \mathbf{e}_1 - \mathbf{e}_2$. Any image $a\mathbf{e}_1 + b\mathbf{e}_2$ can be re-written in terms of the new basis:

$$a\mathbf{e}_1 + b\mathbf{e}_2 = \frac{a+b}{2}(\mathbf{e}_1 + \mathbf{e}_2) + \frac{a-b}{2}(\mathbf{e}_1 - \mathbf{e}_2).$$

So the new basis is giving us a different description of the image. It tells us how bright the image is overall (the coefficient $(a+b)/2$ of $\mathbf{e}_1 + \mathbf{e}_2$ is the average brightness of the two pixels, so it measures the overall image brightness) and how different in brightness the two pixels are (the coefficient $(a-b)/2$ of $\mathbf{e}_1 - \mathbf{e}_2$ is a measure of how different the brightnesses a and b of the two pixels are).

4.9 Dimension

4.9.1 Basis size

We are going to define the **dimension** of a finite-dimensional vector space V as the size of a basis of V . But as we've seen, a vector space can have many different bases. So we have some proving to do before this definition makes sense. We need to know that any two bases have the same size.

4.9.2 Spanning sequences are at least as big as linearly independent sequences (Steinitz exchange)

Theorem 4.9.1. *Let V be a vector space and suppose $\mathbf{s}_1, \dots, \mathbf{s}_m$ spans V and $\mathbf{l}_1, \dots, \mathbf{l}_n$ is a linearly independent sequence of elements of V . Then $m \geq n$.*

Proof. Assume for a contradiction that $m < n$. Since the \mathbf{s}_i span V we can write each \mathbf{l}_j as a linear combination of the \mathbf{s}_i so there are scalars a_{ij} such that $\mathbf{l}_j = \sum_{i=1}^m a_{ij} \mathbf{s}_i$. Let A be the $m \times n$ matrix (a_{ij}) , which has more columns than rows.

By Proposition 3.11.2, the matrix equation $A\mathbf{x} = \mathbf{0}$ has at least one nonzero solution $\mathbf{v} = \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix}$. Because $A\mathbf{v} = \mathbf{0}$, for any i we have $\sum_{j=1}^n a_{ij} v_j = 0$. Now

$$\begin{aligned} \sum_{j=1}^n v_j \mathbf{l}_j &= \sum_{j=1}^n v_j \sum_{i=1}^m a_{ij} \mathbf{s}_i \\ &= \sum_{i=1}^m \left(\sum_{j=1}^n a_{ij} v_j \right) \mathbf{s}_i \\ &= \sum_{i=1}^m 0 \mathbf{s}_i \\ &= \mathbf{0}_V \end{aligned}$$

and since the v_j are not all zero, this contradicts the linear independence of $\mathbf{l}_1, \dots, \mathbf{l}_n$. \square

4.9.3 All bases of a finite-dimensional vector space have the same size

To make life slightly easier, we are going to work only with finite-dimensional vector spaces. A vector space is called **finite-dimensional** if it contains a finite spanning sequence.

Theorem 4.9.2. *Any two bases of a finite-dimensional vector space V have the same size.*

Proof. V has a finite spanning sequence $\mathbf{s}_1, \dots, \mathbf{s}_m$ because it is finite-dimensional. Therefore every linearly independent sequence has size at most m , so is finite, so every basis is finite. (We haven't actually shown that a basis exists, but this will follow from something we prove later).

Let $\mathbf{b}_1, \dots, \mathbf{b}_k$ and $\mathbf{c}_1, \dots, \mathbf{c}_l$ be bases of V . Then $k \leq l$ (as the \mathbf{b}_i s are linearly independent and the \mathbf{c}_i s span). By the same argument with the two bases swapped, $l \leq k$. Therefore $k = l$. \square

Now that we know any two bases have the same size, we can make our definition of dimension:

Definition 4.9.1. The **dimension** of a vector space V , written $\dim V$, is the size of any basis of V .

There's a special case: the dimension of the zero vector space $\{0\}$ is defined to be 0. If you want you can talk yourself into believing that the empty set is a basis of the zero vector space, so that this is covered by the definition above, but it's easier just to think of this as a special case.

4.10 Basis and dimension examples

We've already seen a couple of examples, the most important being the **standard basis** of \mathbb{F}^n , the space of height n column vectors with entries in \mathbb{F} . This standard basis was $\mathbf{e}_1, \dots, \mathbf{e}_n$ where \mathbf{e}_i is the height n column vector with a 1 in position i and 0s elsewhere. The basis has size n , so $\dim \mathbb{F}^n = n$.

We can do a similar thing for the vector space of all $m \times n$ matrices over a field \mathbb{F} . Let E_{ij} be the $m \times n$ matrix with a 1 in position i, j and 0s elsewhere. Then the E_{ij} , for $1 \leq i \leq m$, $1 \leq j \leq n$ are a basis of $M_{m \times n}(\mathbb{F})$, which therefore has dimension mn .

Example 4.10.1. The trace of a matrix is the sum of the elements of its leading diagonal. We will find a basis of the set S of 2×2 matrices with trace zero.

First note that this really is a vector space (a subspace of $M_{2 \times 2}(\mathbb{F})$), so its dimension is at most 4.

A good start is to write down an expression for a general matrix with trace zero. It must have the form $\begin{pmatrix} a & b \\ c & -a \end{pmatrix}$. This matrix can be written

$$a \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} + b \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} + c \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$$

Call the three matrices above H, E, F so that our expression was $aH + bE + cF$. Since H, E, F are in S , they are a spanning sequence for S . You can check that they're linearly independent, so they are a basis and $\dim S = 3$.

Example 4.10.2. $\dim \mathbb{R}_{\leq n}[x] = n + 1$, because $1, x, \dots, x^n$ is a basis.

Example 4.10.3. Let $S = \text{span}(\sin, \cos)$, a subspace of the \mathbb{R} -vector space of all functions $\mathbb{R} \rightarrow \mathbb{R}$. We will find $\dim S$.

The functions \cos and \sin are linearly independent by Example 4.6.4, and they span S by definition. Therefore they form a basis of S and $\dim S = 2$.

4.11 Extending to a basis

Our goal in this section is to show that every linearly independent sequence in a finite-dimensional vector space can be extended, by adding some more vectors to the sequence, to a basis.

4.11.1 The extension lemma

Lemma 4.11.1. *Suppose $\mathbf{v}_1, \dots, \mathbf{v}_n$ is a linearly independent sequence in a vector space V , and $\mathbf{u} \in V$. Then $\mathbf{u} \notin \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_n)$ implies $\mathbf{v}_1, \dots, \mathbf{v}_n, \mathbf{u}$ is linearly independent.*

Proof. We prove the contrapositive, which is that if $\mathbf{v}_1, \dots, \mathbf{v}_n, \mathbf{u}$ is linearly dependent then $\mathbf{u} \in \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_n)$.

Suppose $\mathbf{v}_1, \dots, \mathbf{v}_n, \mathbf{u}$ is linearly dependent. There are scalars $\lambda, \lambda_1, \dots, \lambda_n$, not all of which are zero, such that

$$\lambda \mathbf{u} + \sum_{i=1}^n \lambda_i \mathbf{v}_i = \mathbf{0}_V.$$

λ can't be zero, for then this equation would say that $\mathbf{v}_1, \dots, \mathbf{v}_n$ was linearly dependent. Therefore we can rearrange to get

$$\mathbf{u} = -\lambda^{-1} \sum_{i=1}^n \lambda_i \mathbf{v}_i = \sum_{i=1}^n -\lambda^{-1} \lambda_i \mathbf{v}_i \in \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_n)$$

as required. \square

4.11.2 Every linearly independent sequence can be extended to a basis

Proposition 4.11.2. *Let $\mathbf{l}_1, \dots, \mathbf{l}_n$ be a linearly independent sequence of elements of a finite-dimensional vector space V . Then there is a basis of V containing $\mathbf{l}_1, \dots, \mathbf{l}_n$.*

Proof. : Let $\mathcal{L} = \mathbf{l}_1, \dots, \mathbf{l}_n$. Since V is finite-dimensional there are elements $\mathbf{v}_1, \dots, \mathbf{v}_m$ of V that span V .

Define a sequence of sequences of elements of V as follows: $\mathcal{S}_0 = \mathcal{L}$, and for $i \geq 0$,

$$\mathcal{S}_{i+1} = \begin{cases} \mathcal{S}_i & \text{if } \mathbf{v}_{i+1} \in \text{span } \mathcal{S}_i \\ \mathcal{S}_i, \mathbf{v}_{i+1} & \text{otherwise.} \end{cases}$$

Here $\mathcal{S}_i, \mathbf{v}_{i+1}$ just means take the sequence \mathcal{S}_i and add \mathbf{v}_{i+1} on to the end.

Note that in either case $\mathbf{v}_{i+1} \in \text{span } \mathcal{S}_{i+1}$, and also that $\mathcal{S}_0 \subseteq \mathcal{S}_1 \subseteq \dots \subseteq \mathcal{S}_m$.

Each sequence \mathcal{S}_i is linearly independent by the extension lemma, Lemma 4.11.1 and in particular \mathcal{S}_m is linearly independent. Furthermore $\text{span } \mathcal{S}_m$ contains the spanning sequence $\{\mathbf{v}_1, \dots, \mathbf{v}_m\}$ because for each i we have $\mathbf{v}_i \in \text{span } \mathcal{S}_i \subseteq \text{span } \mathcal{S}_m$, so since subspaces are closed under taking linear combinations, $\text{span } \mathcal{S}_m = V$. Therefore \mathcal{S}_m is a basis containing \mathcal{L} . This completes the proof. \square

As a corollary, we can prove that every finite-dimensional vector space has a basis. Start with any nonzero vector you like — this forms a linearly independent sequence of length 1. The above result lets us extend that to a basis, and in particular, a basis exists.

Example 4.11.1. Consider the sequence of elements $\mathcal{L} = \mathbf{l}_1, \mathbf{l}_2$ where $\mathbf{l}_1 = (0, 1, 1, 0)$, $\mathbf{l}_2 = (1, 0, 1, 0)$ of the vector space V of all width 4 row vectors with real number entries. It's easy to check that they are linearly independent. We are going to use the procedure above, together with the spanning sequence

$$\begin{aligned} \mathbf{v}_1 &= (1, 0, 0, 0), \mathbf{v}_2 = (0, 1, 0, 0) \\ \mathbf{v}_3 &= (0, 0, 1, 0), \mathbf{v}_4 = (0, 0, 0, 1) \end{aligned}$$

of V to produce a basis of V containing \mathcal{L} .

We begin with the sequence $\mathcal{S}_0 = \mathcal{L}$. To find \mathcal{S}_1 we have to determine if $\mathbf{v}_1 \in \text{span } \mathcal{S}_0$. It isn't (to see this, show that the system of linear equations

$$\mathbf{v}_1 = a\mathbf{l}_1 + b\mathbf{l}_2$$

has no solutions), so \mathcal{S}_1 is \mathcal{S}_0 with \mathbf{v}_1 added, which is $\mathbf{l}_1, \mathbf{l}_2, \mathbf{v}_1$.

To find \mathcal{S}_2 we have to determine if $\mathbf{v}_2 \in \text{span } \mathcal{S}_2$. It is, because

$$\mathbf{v}_2 = (0, 1, 0, 0) = \mathbf{l}_1 - \mathbf{l}_2 + \mathbf{v}_1$$

so \mathcal{S}_2 is the same as \mathcal{S}_1 .

To find \mathcal{S}_3 we have to determine if $\mathbf{v}_3 \in \text{span } \mathcal{S}_3$. It is, because

$$\mathbf{v}_3 = \mathbf{l}_2 - \mathbf{v}_1$$

so \mathcal{S}_3 is the same as \mathcal{S}_2 .

Finally to find \mathcal{S}_4 we have to determine if $\mathbf{v}_4 \in \text{span } \mathcal{S}_3$. It is not (no linear combination of \mathcal{S}_3 can have a nonzero entry in the last position), so \mathcal{S}_4 is \mathcal{S}_3 with \mathbf{v}_4 added. We have run out of \mathbf{v}_i s, so \mathcal{S}_4 is the required basis containing \mathcal{L} .

4.12 Finding dimensions

The extension lemma has all sorts of consequences that are very useful for making arguments about the dimension of a vector space. In this section we'll write down the most common ones.

4.12.1 Lower bound for the dimension of a vector space

As soon as you see k linearly independent elements in a vector space, you know its dimension is at least k .

Corollary 4.12.1. *Let V be a vector space and let $\mathbf{v}_1, \dots, \mathbf{v}_k$ be linearly independent elements of V . Then $\dim V \geq k$.*

Proof. You can extend these elements to a basis of V having size at least k , and the size of that basis is the dimension of V . \square

4.12.2 Any $\dim V + 1$ elements must be linearly dependent

Theorem 4.12.2. *Any sequence of at least $n + 1$ elements in a vector space of dimension n is linearly dependent.*

Proof. The vector space has a basis of size n , which is in particular a spanning sequence of size n . By Theorem 4.9.1 (Steinitz Exchange), any linearly independent sequence has size at most n . \square

For example, if you have 4 vectors in \mathbb{R}^3 you know they must be linearly dependent, no matter what they are.

4.12.3 Dimensions of subspaces

Proposition 4.12.3. *If $U \leq V$ then*

1. $\dim U \leq \dim V$, and
2. if $\dim U = \dim V$ then $U = V$.

Proof. 1. A basis of U is a linearly independent sequence in V and a basis of V is (in particular) a spanning sequence for V , so by Theorem 4.9.1 the size of a basis of U is less than or equal to the size of a basis of V .

2. Let $\dim V = n$ and let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be a basis of U , so $U = \text{span}(\mathbf{u}_1, \dots, \mathbf{u}_n)$. Suppose for a contradiction that $U \neq V$, and let \mathbf{v} be an element of V not in U . Then $\mathbf{u}_1, \dots, \mathbf{u}_n, \mathbf{v}$ is linearly independent (by the extension lemma, Lemma 4.11.1), which contradicts Theorem 4.12.2. \square

As soon as you have n linearly independent elements in a vector space of dimension n , they must be a basis.

Corollary 4.12.4. *Let V be a vector space of dimension n . Any sequence of n linearly independent elements of V are a basis of V .*

Proof. Let U be the span of this sequence. This length n sequence spans U by definition, and it is linearly independent, so it is a basis of U and $\dim U = n$. The previous proposition tells us $U = V$, so in fact the sequence is a basis of V . \square

4.12.4 Dimension of a sum of subspaces

Consider two sets X and Y . What's the size of $X \cup Y$ in terms of the size of X and the size of Y ? It isn't $|X| + |Y|$, in general, because elements belonging to X and Y get counted twice when you add the sizes like this. The correct answer is $|X| + |Y| - |X \cap Y|$. We would like a similar result for sums of subspaces.

Theorem 4.12.5. *Let V be a vector space and $X, Y \leq V$. Then*

$$\dim(X + Y) = \dim X + \dim Y - \dim X \cap Y.$$

Proof. Take a basis $\mathcal{I} = \mathbf{i}_1, \dots, \mathbf{i}_k$ of $X \cap Y$. Extend \mathcal{I} to a basis $\mathcal{X} = \mathbf{i}_1, \dots, \mathbf{i}_k, \mathbf{x}_1, \dots, \mathbf{x}_n$ of X , using Proposition 4.11.2. Extend \mathcal{I} to a basis $\mathcal{Y} = \mathbf{i}_1, \dots, \mathbf{i}_k, \mathbf{y}_1, \dots, \mathbf{y}_m$ of Y . It's now enough to prove that $\mathcal{J} = \mathbf{i}_1, \dots, \mathbf{i}_k, \mathbf{x}_1, \dots, \mathbf{x}_n, \mathbf{y}_1, \dots, \mathbf{y}_m$ is a basis of $X + Y$, because if we do that then we will know the size of \mathcal{J} , which is $k + n + m$, equals the size of a basis of \mathcal{I} (which is $k + n$) plus the size of a basis of Y (which is $k + m$) minus the size of a basis of $X \cap Y$ (which is k).

To check something is a basis for $X + Y$, as always, we must check that it is a spanning sequence for $X + Y$ and that it is linearly independent.

Spanning: let $\mathbf{x} + \mathbf{y} \in X + Y$, where $\mathbf{x} \in X, \mathbf{y} \in Y$. Then there are scalars such that

$$\begin{aligned} \mathbf{x} &= \sum_{j=1}^k a_j \mathbf{i}_j + \sum_{j=1}^n c_j \mathbf{x}_j \\ \mathbf{y} &= \sum_{j=1}^k b_j \mathbf{i}_j + \sum_{j=1}^m d_j \mathbf{y}_j \end{aligned}$$

and so

$$\mathbf{x} + \mathbf{y} = \sum_{j=1}^k (a_j + b_j) \mathbf{i}_j + \sum_{j=1}^n c_j \mathbf{x}_j + \sum_{j=1}^m d_j \mathbf{y}_j$$

Linear independence: suppose

$$\sum_{j=1}^k a_j \mathbf{i}_j + \sum_{j=1}^n c_j \mathbf{x}_j + \sum_{j=1}^m d_j \mathbf{y}_j = \mathbf{0}.$$

Rearrange it:

$$\sum_{j=1}^k a_j \mathbf{i}_j + \sum_{j=1}^n c_j \mathbf{x}_j = - \sum_{j=1}^m d_j \mathbf{y}_j.$$

The left hand side is in X and the right hand side is in Y . So both sides are in $X \cap Y$, in particular, the right hand side is in $X \cap Y$. Since \mathcal{I} is a basis of $X \cap Y$, there are scalars e_j such that

$$\sum_{j=1}^k e_j \mathbf{i}_j = - \sum_{j=1}^m d_j \mathbf{y}_j$$

This is a linear dependence on \mathcal{Y} which is linearly independent, so all the d_j are 0. Similarly all the c_j are 0. So the a_j are 0 too, and we have linear independence. \square

4.13 Linear maps

4.13.1 Motivation

Suppose that you have two finite sets X and Y and a function $f : X \rightarrow Y$. If you know that f is onto then you get some information about X and Y : you know that X must be at least as large as Y .

But an arbitrary function between two vector spaces doesn't necessarily give you any information about their relationship as vector spaces. To get such information, we need to restrict to functions that respect the vector space structure — that is, the scalar multiplication and the vector addition.

Functions with this property, which we're going to define shortly, are called linear maps. They allow us to do something similar to the finite set example above: for example, if you have a surjective linear map from a vector space X to another vector space Y , it is true that $\dim X \geq \dim Y$.

4.13.2 Definition of a linear map

Definition 4.13.1. Let V and W be vector spaces over the same field \mathbb{F} . A function $T : V \rightarrow W$ is called a **linear map** or a **linear transformation** if

1. for all $\lambda \in \mathbb{F}$ and all $\mathbf{v} \in V$ we have $T(\lambda \mathbf{v}) = \lambda T(\mathbf{v})$, and
2. for all $\mathbf{v}, \mathbf{w} \in V$ we have $T(\mathbf{v} + \mathbf{w}) = T(\mathbf{v}) + T(\mathbf{w})$.

Point 1 is what it means to say that T respects scalar multiplication and point 2 is what it means to say that T respects vector addition.

This concept is so common that it has many names. For us,

- T is a linear map
- T is a linear function
- T is a linear transformation
- T is linear

all mean exactly the same thing, namely that T satisfies Definition 4.13.1.

4.13.3 Examples of linear maps

1. For any vector space V , the identity map $\text{id} : V \rightarrow V$ and the zero map $z : V \rightarrow V$ given by $z(v) = \mathbf{0}_V$ for all $v \in V$ are linear.
2. Let A be a $m \times n$ matrix with entries in a field \mathbb{F} . Then $T_A : \mathbb{F}^n \rightarrow \mathbb{F}^m$ defined by $T_A(\mathbf{x}) = A\mathbf{x}$ is linear.
3. $T : M_{n \times n}(\mathbb{R}) \rightarrow M_{n \times n}(\mathbb{R}), T(A) = A^2$ is **not** linear.
4. $T : \mathbb{R}^n \rightarrow \mathbb{R}, T \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = \sum_{i=1}^n x_i$ is linear.
5. $D : \mathbb{R}_{\leq n}[x] \rightarrow \mathbb{R}_{\leq n}[x]$ given by $D(f) = \frac{df}{dx}$ is linear.

Let's look at why some of these are true, starting with example 2. To show that T_A is a linear map we have to check the two parts of the definition of being a linear map. Both of these are going to follow from properties of matrix multiplication and addition that you learned in the previous section.

1. Let $\mathbf{x} \in \mathbb{F}^n$ and $\lambda \in \mathbb{F}$. Then

$$\begin{aligned} T_A(\lambda\mathbf{x}) &= A(\lambda\mathbf{x}) && \text{definition of } T_A \\ &= \lambda A\mathbf{x} && \text{matrix multiplication properties} \\ &= \lambda T_A(\mathbf{x}) && \text{definition of } T_A \end{aligned}$$

2. Let $\mathbf{x}, \mathbf{y} \in \mathbb{F}^n$. Then

$$\begin{aligned} T_A(\mathbf{x} + \mathbf{y}) &= A(\mathbf{x} + \mathbf{y}) && \text{definition of } T_A \\ &= A\mathbf{x} + A\mathbf{y} && \text{matrix multiplication properties} \\ &= T_A(\mathbf{x}) + T_A(\mathbf{y}) && \text{definition of } T_A \end{aligned}$$

The properties of matrix multiplication used were proved in Proposition 3.4.1.

Similarly, the fact that the differentiation map D of example 5 is linear follows from standard properties of derivatives: you know, for example, that for any two functions (not just polynomials) f and g we have $\frac{d}{dx}(f + g) = \frac{df}{dx} + \frac{dg}{dx}$, which shows that D satisfies the second part of the linearity definition.

As an example where the linearity of a map doesn't just come from standard facts you already know, consider

$$T : \mathbb{R}^2 \rightarrow \mathbb{R} \quad T \begin{pmatrix} x \\ y \end{pmatrix} = 2x - y$$

To show T is linear we have to show that it has properties 1 and 2 from the definition.

1.

$$\begin{aligned} T \left(\lambda \begin{pmatrix} x \\ y \end{pmatrix} \right) &= T \begin{pmatrix} \lambda x \\ \lambda y \end{pmatrix} \\ &= 2\lambda x - \lambda y \\ &= \lambda(2x - y) \\ &= \lambda T \begin{pmatrix} x \\ y \end{pmatrix} \end{aligned}$$

2.

$$\begin{aligned} T \left(\begin{pmatrix} x_1 \\ y_1 \end{pmatrix} + \begin{pmatrix} x_2 \\ y_2 \end{pmatrix} \right) &= T \begin{pmatrix} x_1 + x_2 \\ y_1 + y_2 \end{pmatrix} \\ &= 2(x_1 + x_2) - (y_1 + y_2) \\ &= (2x_1 - y_1) + (2x_2 - y_2) \\ &= T \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} + T \begin{pmatrix} x_2 \\ y_2 \end{pmatrix}. \end{aligned}$$

Here are some examples of things which are **not** linear maps:

Example 4.13.1. • $T : \mathbb{R} \rightarrow \mathbb{R}, T(x) = |x|$ isn't linear. It doesn't satisfy either linearity property. $T(-2 \cdot 3) \neq -2 \cdot T(3)$, and $T(-1 + 1) \neq T(-1) + T(1)$.

• $T : \mathbb{R}^3 \rightarrow \mathbb{R}^3, T(\mathbf{x}) = \mathbf{x} + \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$. Again it doesn't satisfy either part of the definition - you should check that.

4.14 Kernel and image

4.14.1 Definition of kernel and image

To every linear transformation we associate two important subspaces.

Definition 4.14.1. let $T : V \rightarrow W$ be linear.

1. the **kernel** of T , written $\ker T$, is $\{\mathbf{v} \in V : T(\mathbf{v}) = \mathbf{0}_W\}$
2. the **image** of T , written $\text{im } T$, is $\{T(\mathbf{v}) : \mathbf{v} \in V\}$

In other words, the image is what we normally mean by the image of a function.

An important family of examples are the linear maps $T_A : \mathbb{F}^n \rightarrow \mathbb{F}^m$ defined by left-multiplication by an $m \times n$ matrix A with entries from the field \mathbb{F} . In that case the image $\text{im } T_A$ is equal to the column space $C(A)$ by Proposition 3.2.1, and the kernel $\ker T_A$ is the nullspace $N(A)$.

4.14.2 A property of all linear maps

Lemma 4.14.1. *Let $T : V \rightarrow W$ be a linear map. Then $T(\mathbf{0}_V) = \mathbf{0}_W$.*

Proof.

$$\begin{aligned} T(\mathbf{0}_V) &= T(\mathbf{0}_V + \mathbf{0}_V) \\ &= T(\mathbf{0}_V) + T(\mathbf{0}_V) \end{aligned}$$

by the second part of the definition of linearity. Now add $-T(\mathbf{0}_V)$ to both sides:

$$\begin{aligned} T(\mathbf{0}_V) - T(\mathbf{0}_V) &= T(\mathbf{0}_V) + T(\mathbf{0}_V) - T(\mathbf{0}_V) \\ \mathbf{0}_W &= T(\mathbf{0}_V) \end{aligned} \quad \square$$

4.14.3 Kernels and images are subspaces

Lemma 4.14.2. *Let $T : V \rightarrow W$ be linear. Then $\ker T \leq V$ and $\text{im } T \leq W$.*

Proof. To show something is a subspace you must check the three conditions: it contains the zero vector, it is closed under addition, it is closed under scalar multiplication.

First, the kernel.

1. To show that the kernel contains $\mathbf{0}_V$, we must show that $T(\mathbf{0}_V) = \mathbf{0}_W$. That's exactly Lemma 4.14.1.
2. If $\mathbf{v}, \mathbf{w} \in \ker T$ then $T(\mathbf{v} + \mathbf{w}) = T(\mathbf{v}) + T(\mathbf{w}) = \mathbf{0}_W + \mathbf{0}_W = \mathbf{0}_W$, so $\mathbf{v} + \mathbf{w} \in \ker T$.
3. If $\mathbf{v} \in \ker T$ and $\lambda \in \mathbb{F}$ then $T(\lambda\mathbf{v}) = \lambda T(\mathbf{v})$ by the second part of the definition of linearity, and this is $\lambda\mathbf{0}_W$ which equals $\mathbf{0}_W$. Since $T(\lambda\mathbf{v}) = \mathbf{0}_W$, we have $\lambda\mathbf{v} \in \ker T$.

Next, the image.

1. We know from Lemma 4.14.1 that $T(\mathbf{0}_V) = \mathbf{0}_W$, so $\mathbf{0}_W \in \text{im } T$.
2. Any two elements of $\text{im } T$ have the form $T(\mathbf{u}), T(\mathbf{v})$ some $\mathbf{u}, \mathbf{v} \in V$. Then $T(\mathbf{u}) + T(\mathbf{v}) = T(\mathbf{u} + \mathbf{v})$ (second part of the linearity definition), which is an element of $\text{im } T$, so $\text{im } T$ is closed under addition.
3. If $T(\mathbf{u}) \in \text{im } T$ and $\lambda \in \mathbb{F}$ then $\lambda T(\mathbf{u}) = T(\lambda\mathbf{u})$ by the second part of the definition of linearity, and this is an element of $\text{im } T$ as it is T applied to something, so $\text{im } T$ is closed under scalar multiplication.

□

Example 4.14.1. Let $A = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ so that we have a linear map $T_A : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by $T_A(\mathbf{x}) = A\mathbf{x}$. We will find $\text{im } T_A$ and $\ker T_A$.

$$\begin{aligned} \text{im } T_A &= \left\{ T_A \begin{pmatrix} x \\ y \end{pmatrix} : x, y \in \mathbb{R} \right\} \\ &= \left\{ \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} : x, y \in \mathbb{R} \right\} \\ &= \left\{ \begin{pmatrix} y \\ 0 \end{pmatrix} : x, y \in \mathbb{R} \right\} \end{aligned}$$

Another way to write this is that $\text{im } T_A = \text{span} \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, and so $\dim \text{im } T_A = 1$.

Now we'll do the kernel.

$$\begin{aligned} \ker T_A &= \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{R}^2 : T_A \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \right\} \\ &= \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{R}^2 : \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \right\} \\ &= \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{R}^2 : \begin{pmatrix} y \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \right\} \\ &= \left\{ \begin{pmatrix} x \\ 0 \end{pmatrix} : x \in \mathbb{R} \right\} \end{aligned}$$

Again we could write this as $\ker T_A = \text{span} \begin{pmatrix} 1 \\ 0 \end{pmatrix}$. The kernel and image are equal in this case.

Example 4.14.2. Let $D : \mathbb{R}_{\leq n}[x] \rightarrow \mathbb{R}_{\leq n}[x]$ be $D(f) = \frac{df}{dx}$. We will describe $\ker D$ and $\text{im } D$.

A polynomial has derivative zero if and only if it is constant, so $\ker D$ is the set of all constant polynomials. This is spanned by any (nonzero) constant polynomial, so it has dimension one.

Next consider $\text{im } D$. Let $S \leq \mathbb{R}_{\leq n}[x]$ be the subspace spanned by $1, x, \dots, x^{n-1}$, that is, the subspace consisting of all polynomials of degree at most $n-1$. Certainly $\text{im } D \leq S$, since when you differentiate a polynomial of degree at most n you get a polynomial of degree at most $n-1$. But if $s(x) \in S$ then $s(x)$ has an indefinite integral $t(x)$ in $\mathbb{R}_{\leq n}[x]$ and $D(t) = s$, so every $s \in S$ is in $\text{im } D$, so $\text{im } D = S$.

A useful property of the kernel of a linear map is that it tells you whether or not the map is injective.

Proposition 4.14.3. Let T be a linear map $T : U \rightarrow V$.

1. T is injective if and only if $\ker T = \{\mathbf{0}_U\}$.
2. T is surjective if and only if $\dim \text{im } T = \dim V$.

Proof. 1. Suppose T is injective. By Lemma 4.14.1, $T(\mathbf{0}_U) = \mathbf{0}_V$, so $\mathbf{0}_U \in \ker T$. Now if $\mathbf{u} \in \ker T$ then $T(\mathbf{u}) = \mathbf{0}_V$ so $T(\mathbf{u}) = T(\mathbf{0}_U)$, so by injectivity $\mathbf{u} = \mathbf{0}_U$. It follows $\ker T = \{\mathbf{0}_U\}$.

Conversely suppose $\ker T = \{\mathbf{0}_U\}$ and that $T(\mathbf{x}) = T(\mathbf{y})$. By linearity $T(\mathbf{x} - \mathbf{y}) = \mathbf{0}_V$, so $\mathbf{x} - \mathbf{y} \in \ker T$, so $\mathbf{x} - \mathbf{y} = \mathbf{0}_U$, so $\mathbf{x} = \mathbf{y}$ and T is injective.

2. If T is surjective then $\text{im } T = V$ so certainly $\dim \text{im } T = \dim V$.

Conversely, suppose $\dim \text{im } T = \dim V$. The image $\text{im } T$ is a subspace of V by Lemma 4.14.2. It has the same dimension as V so it equals V by Proposition 4.12.3. We have shown $\text{im } T = V$, so T is surjective. \square

4.15 The rank-nullity theorem

4.15.1 Definition of rank and nullity

Definition 4.15.1. Let $T : V \rightarrow W$ be a linear map.

- The **nullity** of T , written $\text{null } T$, is $\dim \ker T$.
- The **rank** of T , written $\text{rank } T$ is $\dim \text{im } T$.

Example 4.15.1. Returning to the differentiation example from the start of section 4.13.3, $D : \mathbb{R}_{\leq n}[x] \rightarrow \mathbb{R}_{\leq n}[x]$ has nullity 1 (since its kernel was one-dimensional, spanned by the constant polynomial 1) and rank n , since its image had a basis $1, x, \dots, x^{n-1}$ of size n . Notice that $\text{rank}(D) + \text{null}(D) = \dim \mathbb{R}_{\leq n}[x]$, this isn't a coincidence.

4.15.2 Statement of the rank-nullity theorem

Theorem 4.15.1. Let $T : V \rightarrow W$ be a linear map. Then

$$\text{rank } T + \text{null } T = \dim V.$$

This is called the **rank-nullity theorem**.

Proof. We'll assume V and W are finite-dimensional, not that it matters. Here is an outline of how the proof is going to work.

1. Choose a basis $\mathcal{K} = \mathbf{k}_1, \dots, \mathbf{k}_m$ of $\ker T$
2. Extend it to a basis $\mathcal{B} = \mathbf{k}_1, \dots, \mathbf{k}_m, \mathbf{v}_1, \dots, \mathbf{v}_n$ of V using Proposition 4.11.2. When we've done this, $\dim V = m + n$ and we need only show $\dim \text{im } T = n$
3. Show that $T(\mathbf{v}_1), \dots, T(\mathbf{v}_n)$ is a basis of $\text{im } T$.

The only part needing elaboration is the last part. First, let's show that the sequence $T(\mathbf{v}_1), \dots, T(\mathbf{v}_n)$ spans $\text{im } T$. Any element of the image is equal to $T(\mathbf{v})$ for some $\mathbf{v} \in V$. We have to show that any such $T(\mathbf{v})$ lies in the span of the $T(\mathbf{v}_i)$ s.

Since \mathcal{B} is a basis of V we may write \mathbf{v} as $\sum_{i=1}^m \lambda_i \mathbf{k}_i + \sum_{i=1}^n \mu_i \mathbf{v}_i$ for some scalars λ_i, μ_i . Then

$$\begin{aligned} T(\mathbf{v}) &= T\left(\sum_{i=1}^m \lambda_i \mathbf{k}_i + \sum_{i=1}^n \mu_i \mathbf{v}_i\right) \\ &= \sum_{i=1}^m \lambda_i T(\mathbf{k}_i) + \sum_{i=1}^n \mu_i T(\mathbf{v}_i) && \text{by linearity} \\ &= \sum_{i=1}^n \mu_i T(\mathbf{v}_i) && \text{as } \mathbf{k}_i \in \ker T \\ &\in \text{span}(T(\mathbf{v}_1), \dots, T(\mathbf{v}_n)) \end{aligned}$$

as required.

Now let's show $T(\mathbf{v}_1), \dots, T(\mathbf{v}_n)$ is linearly independent. Suppose

$$\sum_{i=1}^n \mu_i T(\mathbf{v}_i) = 0,$$

so that we need to show the μ_i are all 0. Using linearity,

$$T\left(\sum_{i=1}^n \mu_i \mathbf{v}_i\right) = 0$$

which means $\sum_{i=1}^n \mu_i \mathbf{v}_i \in \ker T$. As \mathcal{K} is a basis for $\ker T$, we can write

$$\sum_{i=1}^n \mu_i \mathbf{v}_i = \sum_{i=1}^m \lambda_i \mathbf{k}_i$$

for some scalars λ_i . But \mathcal{B} , being a basis, is linearly independent and so all the scalars are 0. In particular all the μ_i are 0, which completes the proof. \square

We can use the

4.16 Matrix nullspace basis

We are ready to prove that the fundamental solutions of $A\mathbf{x} = \mathbf{0}$ are a basis for $N(A)$. To set up the notation, suppose A is a $m \times n$ matrix, R is a RREF matrix obtained by doing row operations to A , the number of columns of R with a leading entry is r and the number of columns with no leading entry is k , so $r + k = n$. Let the numbers of the columns with no leading entry be

$$j_1 < j_2 < \dots < j_k$$

and the numbers of the columns that do have a leading entry be

$$i_1 < i_2 < \dots < i_r$$

so that the numbers $j_1, \dots, j_k, i_1, \dots, i_r$ contain all the numbers $1, \dots, n$ exactly once. The variable corresponding to a column with no leading entry is called a free parameter because its value can be chosen freely. There are k fundamental solutions to $A\mathbf{x} = \mathbf{0}$, with the j th fundamental solution defined to be the one where the j th free parameter is 1 and all the other free parameters are 0.

Theorem 4.16.1. *The fundamental solutions to $A\mathbf{x} = \mathbf{0}$ are a basis of the nullspace $N(A)$.*

Before we do the proof, let's work an example to illustrate how it will go. Free parameters and their coefficients will be coloured red. Take

$$R = \begin{pmatrix} 0 & 1 & 2 & 0 & 3 \\ 0 & 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

so that we have $m = 3$ and $n = 5$. The columns with no leading entry are 1, 3, and 5 so $k = 3$, $j_1 = 1$, $j_2 = 3$, $j_3 = 5$, and variables x_1 , x_3 , and x_5 are the free parameters. The columns with leading entries are 2 and 4, so $r = 2$ and $i_1 = 2$, $i_2 = 4$. The equations corresponding to $R\mathbf{x} = \mathbf{0}$ are

$$\begin{aligned} x_2 + 2x_3 + 3x_5 &= 0 \\ x_4 + 4x_5 &= 0 \end{aligned}$$

(missing off the final one, since it just says $0 = 0$). These equations show that once values for the free parameters are known, the other variables x_2 and x_4 are completely determined by those values — we have $x_2 = -2x_3 - 3x_5$ and $x_4 = -4x_5$.

The three fundamental solutions are

$$\mathbf{s}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \mathbf{s}_2 = \begin{pmatrix} 0 \\ -2 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \mathbf{s}_3 = \begin{pmatrix} 0 \\ -3 \\ 0 \\ -4 \\ 1 \end{pmatrix}.$$

These are linearly independent, because if

$$a_1\mathbf{s}_1 + a_2\mathbf{s}_2 + a_3\mathbf{s}_3 = \mathbf{0}$$

then row 1 of this equation shows that $a_1 = 0$, row 3 shows that $a_2 = 0$, and row 5 shows that $a_3 = 0$. To show that the fundamental solutions span,

suppose $\mathbf{s} = \begin{pmatrix} s_1 \\ \vdots \\ s_5 \end{pmatrix}$ is any vector such that $R\mathbf{s} = \mathbf{0}$. We claim that \mathbf{s} is equal

to $s_1\mathbf{s}_1 + s_3\mathbf{s}_2 + s_5\mathbf{s}_3$. Certainly these two vectors have the same entries in rows 1, 3, and 5, since the entries in these rows are s_1 , s_3 , and s_5 . What about the entries in rows 2 and 4, that is, the values of x_2 and x_4 ? As above, these entries are completely determined by the values of the free parameters. The free parameters are the same for both vectors, so the values of x_2 and x_4 are the same as well. We have shown that any solution \mathbf{s} is in the span of the fundamental solutions, so we are done.

Proof. Theorem 3.9.1 shows that $N(A) = N(R)$, so we will show that the fundamental solutions are a basis of $N(R)$.

First we show that the fundamental solutions $\mathbf{s}_1, \dots, \mathbf{s}_k$ are linearly independent. Suppose that

$$\sum_j a_j \mathbf{s}_j = \mathbf{0}. \quad (4.6)$$

The first fundamental solution corresponds to the variable for column j_1 , so \mathbf{s}_1 has a 1 in row j_1 and all the other fundamental solutions have a 0 there. Thus the entry in row j_1 on the left hand side of (4.6) is a_1 , so $a_1 = 0$. Similarly all the other coefficients are zero.

Now let \mathbf{s} be any solution of $A\mathbf{x} = \mathbf{0}$ and let the entry of \mathbf{s} in row i be s_i . We are going to show that

$$\mathbf{s} = \sum_{i=1}^k s_{j_i} \mathbf{s}_i.$$

The left hand side and right hand side of this claimed equation are solutions to $R\mathbf{x} = \mathbf{0}$ with the same free parameter values s_{j_1}, \dots, s_{j_k} . We just have to show that the values of the other variables are the same. But the values of the variables that aren't free parameters are uniquely determined by the values of the free parameters, because by the RREF property, for $1 \leq a \leq r$ the only equation containing x_{i_a} is the one coming from row a of R which has the form

$$x_{i_1} + \dots = 0$$

where the only other variables occurring with nonzero coefficient are free parameters. Since \mathbf{s} and $\sum_{i=1}^k s_{j_i} \mathbf{s}_i$ are solutions to $R\mathbf{x} = \mathbf{0}$ with the same free parameter values s_{j_1}, \dots, s_{j_k} , they are equal. □

4.17 Column space basis

We can now calculate a basis for the nullspace of a $m \times n$ matrix A by putting it into RREF and reading off the fundamental solutions to $A\mathbf{x} = \mathbf{0}$. In this section we consider the problem of finding a basis of the column space $C(A)$, which we defined in Definition 3.6.2 to be the span of the columns of A .

By Proposition 3.2.1, the set of vectors $A\mathbf{v}$ is exactly the set of linear combinations of the columns of A . Therefore the column space $C(A)$ is equal to the image of $T_A : \mathbb{F}^n \rightarrow \mathbb{F}^m$ given by $T_A(\mathbf{x}) = A\mathbf{x}$.

It would be nice to be able to solve this problem using RREF, because it is very easy to find a basis for the column space of a RREF matrix like

$$R = \begin{pmatrix} 0 & 1 & 2 & 0 & 3 \\ 0 & 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

The columns containing a leading entry, in this example columns 2 and 4, are easily seen to be a basis for the column space of R . Unfortunately doing row operations can change the column space of a matrix, so knowing the column space of R does not immediately give you the column space of A .

One solution for this would be to introduce column operations and column reduced echelon form, and re-prove all the things about row operations and row

reduced echelon form. Instead we are going to stick with the row operations we already know and use the transpose to convert columns into rows.

We defined the column space of a matrix as the span of its columns. The row space is defined similarly.

Definition 4.17.1. Let A be a $m \times n$ matrix. The **row space** of A is defined to be the span of the rows of A .

Theorem 4.17.1. Let A be $m \times n$, let E be a $m \times m$ invertible matrix, and let F be a $n \times n$ invertible matrix. Then the row space of EA equals the row space of A and the column space of AF equals the column space of A .

Proof. We will do the second part only as the first one can be proved similarly. By Corollary 3.2.4, the columns of AF are linear combinations of the columns of A , that is, elements of the subspace $C(A)$. The span $C(AF)$ of the columns of AF is therefore also contained in $C(A)$.

Applying the same argument again with AF in place of A and F^{-1} in place of F , the column space $C(AFF^{-1})$, that is, the column space of $C(A)$, is contained in $C(AF)$. \square

Since doing a row operation to a matrix is the same as left multiplication by an elementary matrix (Theorem 3.8.1), this shows that doing row operations to a matrix doesn't change its row space.

Theorem 4.17.2. Let R be a $m \times n$ RREF matrix. Then the nonzero rows of R are a basis for the row space of R .

Proof. Certainly the nonzero rows span the row space, so we only need show they are linearly independent. Let the nonzero rows be $\mathbf{r}_1, \dots, \mathbf{r}_l$, and let the leading entry in row i occur in column c_i . Suppose $\sum_{i=1}^l a_i \mathbf{r}_i = \mathbf{0}$. Pick any $1 \leq i \leq l$ and consider the entry in column c_i of this sum. On the right we have 0. On the left $a_i \mathbf{r}_i$ has a a_i in column c_i , and all the other \mathbf{r}_j s have zeros in column c_i because R is in RREF. Thus we have $a_i = 0$ for $1 \leq i \leq l$, so the rows are linearly independent. \square

The columns of A are the transposes of the rows of A^T , so we can get a basis for the column space of A by forming the matrix A^T , doing row operations until we reach a RREF matrix, then taking the transposes of the nonzero rows of this RREF matrix.

Example 4.17.1. Let $A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$. To find a basis of $C(A)$ we take the transpose of A to get

$$A^T = \begin{pmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 9 \end{pmatrix}$$

Doing row operations, we reach the RREF matrix

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

The nonzero rows $(1 \ 0 \ -1)$ and $(0 \ 1 \ 2)$ are a basis for the row space of R , which equals the row space of A^T , so their transposes

$$\begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 2 \end{pmatrix}$$

are a basis for the column space of A .

4.18 Matrix of a linear map

Linear maps are abstractly defined things. We'd like to make them concrete. We do this by making the following observation: once you know what a linear transformation does on a basis, you know what it does everywhere.

Here's what that means exactly. Let $T : V \rightarrow W$ be linear. Let $\mathbf{b}_1, \dots, \mathbf{b}_n$ be a basis of V . Then we can write any \mathbf{v} as $\sum_{i=1}^n \lambda_i \mathbf{b}_i$ for some scalars λ_i , and so by linearity $T(\mathbf{v}) = T(\sum_{i=1}^n \lambda_i \mathbf{b}_i) = \sum_{i=1}^n \lambda_i T(\mathbf{b}_i)$.

This means that if we want to describe a linear map, we can just say what it does on a basis $\mathbf{b}_1, \dots, \mathbf{b}_n$. You can then work out $T(\mathbf{v})$ for any $\mathbf{v} \in V$ just by knowing the $T(\mathbf{b}_i)$.

We can record what T does to each \mathbf{b}_i by giving the coefficients needed to write the $T(\mathbf{b}_i)$ in terms of some fixed basis of W . Putting all this together, to completely describe a linear map $T : V \rightarrow W$ we only need to give a basis $\mathbf{b}_1, \dots, \mathbf{b}_n$ of V , a basis $\mathbf{c}_1, \dots, \mathbf{c}_m$ of W , and the coefficients needed to express $T(\mathbf{b}_j)$ in terms of the \mathbf{c}_i . These coefficients naturally fit into an $m \times n$ matrix which will be called the matrix of T with respect to initial basis $\mathbf{b}_1, \dots, \mathbf{b}_n$ and final basis $\mathbf{c}_1, \dots, \mathbf{c}_m$.

4.18.1 Definition of the matrix of a linear map

Definition 4.18.1. Let $T : V \rightarrow W$ be linear, and let

- $\mathcal{B} = \mathbf{b}_1, \dots, \mathbf{b}_n$ be a basis of V
- $\mathcal{C} = \mathbf{c}_1, \dots, \mathbf{c}_m$ be a basis of W .

Define scalars a_{ij} by $T(\mathbf{b}_j) = \sum_{i=1}^m a_{ij} \mathbf{c}_i$. Then **the matrix of T with respect to initial basis \mathcal{B} and final basis \mathcal{C}** , written $[T]_{\mathcal{C}}^{\mathcal{B}}$, is the $m \times n$ matrix (a_{ij}) .

Another way to think of this definition is that the j th column of $[T]_{\mathcal{C}}^{\mathcal{B}}$ records the image of the j th basis element from \mathcal{B} under T , in the sense that the entries are the coefficients used in expressing $T(\mathbf{b}_j)$ as a linear combination of \mathcal{C} .

When we have a linear map from a vector space to itself, we sometimes use a slightly different terminology. If $T : V \rightarrow V$ and \mathcal{B} is a basis of V , the **matrix of T with respect to \mathcal{B}** means $[T]_{\mathcal{B}}^{\mathcal{B}}$.

Notice that the order of the basis matters in this definition. If you order the basis elements differently, you change the order of the columns or rows in the matrix. That's why our bases are sequences, not sets.

4.18.2 Examples of the matrix of a linear map

Example 4.18.1. Let $T : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ be defined by $T \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} a+b \\ a-b \\ 2a+b \end{pmatrix}$. This is linear. Let's find the matrix of T with respect to

- initial basis $\mathcal{E} = \mathbf{e}_1, \mathbf{e}_2$, the standard basis for \mathbb{R}^2 , and
- final basis $\mathcal{E}' = \mathbf{e}'_1, \mathbf{e}'_2, \mathbf{e}'_3$, the standard basis for \mathbb{R}^3 .

We have

$$T(\mathbf{e}_1) = \begin{pmatrix} 1 \\ 1 \\ 2 \end{pmatrix} = 1\mathbf{e}'_1 + 1\mathbf{e}'_2 + 2\mathbf{e}'_3$$

$$T(\mathbf{e}_2) = \begin{pmatrix} 1 \\ -1 \\ 1 \end{pmatrix} = 1\mathbf{e}'_1 - 1\mathbf{e}'_2 + 1\mathbf{e}'_3$$

so the matrix $[T]_{\mathcal{E}'}^{\mathcal{E}}$ is $\begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 2 & 1 \end{pmatrix}$.

Example 4.18.2. Let

- V be the vector space of all polynomials with real coefficients of degree ≤ 3 in one variable x
- $D : V \rightarrow V$ be the differentiation map
- \mathcal{B} be the basis $1, x, x^2, x^3$ of V .

D is a linear map, so let's find the matrix $[D]_{\mathcal{B}}^{\mathcal{B}}$ of D with respect to \mathcal{B} . We have

$$D(1) = 0 = 0 \times 1 + 0 \times x + 0 \times x^2 + 0 \times x^3$$

$$D(x) = 1 = 1 \times 1 + 0 \times x + 0 \times x^2 + 0 \times x^3$$

$$D(x^2) = 2x = 0 \times 1 + 2 \times x + 0 \times x^2 + 0 \times x^3$$

$$D(x^3) = 3x^2 = 0 \times 1 + 0 \times x + 3 \times x^2 + 0 \times x^3$$

and so

$$[D]_{\mathcal{B}}^{\mathcal{B}} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

Example 4.18.3. Let $\text{id} : V \rightarrow V$ be the identity map $\text{id}(\mathbf{v}) = \mathbf{v}$. Let $\mathcal{B} = \mathbf{b}_1, \dots, \mathbf{b}_n$ be any basis for V . We're going to work out $[\text{id}]_{\mathcal{B}}^{\mathcal{B}}$. For any j ,

$$\text{id}(\mathbf{b}_j) = \mathbf{b}_j = 0 \times \mathbf{b}_1 + \dots + 1 \times \mathbf{b}_j + \dots + 0 \times \mathbf{b}_n.$$

This means the j th column of $[\text{id}]_{\mathcal{B}}^{\mathcal{B}}$ is all 0s, except a 1 in position j . In other words, $[\text{id}]_{\mathcal{B}}^{\mathcal{B}} = I_n$, the $n \times n$ identity matrix.

This shows that the matrix of the identity map is the identity matrix, *so long as the initial basis and the final basis are the same.*

On the other hand, if $\mathcal{C} = \mathbf{c}_1, \dots, \mathbf{c}_n$ is a different basis of V then $[\text{id}]_{\mathcal{B}}^{\mathcal{C}}$ will **not** be the identity matrix. To figure out what goes in the j th column of this matrix we have to work out $\text{id}(\mathbf{b}_j)$, which is just \mathbf{b}_j of course, as a linear combination of the \mathbf{c}_i s. The coefficients we have to use, whatever they are, make up this column of the matrix.

Example 4.18.4. Consider two bases for \mathbb{R}^2

$$\begin{aligned}\mathcal{B} : \mathbf{e}_1 &= \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \mathbf{e}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ \mathcal{C} : \mathbf{c}_1 &= \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \mathbf{c}_2 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}\end{aligned}$$

Both $[\text{id}]_{\mathcal{B}}^{\mathcal{B}}$ and $[\text{id}]_{\mathcal{C}}^{\mathcal{C}}$ will be the identity matrix I_2 . Let's work out $[\text{id}]_{\mathcal{B}}^{\mathcal{C}}$. To do that, we have to express $\text{id}(\mathbf{c}_j)$ as a linear combination of the \mathbf{e}_i for $j = 1, 2$:

$$\begin{aligned}\text{id}(\mathbf{c}_1) &= \mathbf{c}_1 = \mathbf{e}_1 + \mathbf{e}_2 \\ \text{id}(\mathbf{c}_2) &= \mathbf{c}_2 = \mathbf{e}_1 - \mathbf{e}_2\end{aligned}$$

and so $[\text{id}]_{\mathcal{B}}^{\mathcal{C}} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$.

Example 4.18.5. Let A be an m by n matrix, and $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be the linear map $T_A(\mathbf{x}) = A\mathbf{x}$. Then the matrix of T_A with respect to the standard bases of \mathbb{R}^n and \mathbb{R}^m is A .

4.19 Matrix of a composition

Suppose we have two composable linear maps, S and T . The composition $T \circ S$ is still linear, as you can check. There should be a connection between the matrix of $T \circ S$ with respect to some bases and the matrices for T and S .

Theorem 4.19.1. *Let $S : U \rightarrow V$ and $T : V \rightarrow W$ be linear maps. Let*

- $\mathcal{B} = \mathbf{b}_1, \dots, \mathbf{b}_l$ be a basis of U ,
- $\mathcal{C} = \mathbf{c}_1, \dots, \mathbf{c}_m$ be a basis of V , and
- $\mathcal{D} = \mathbf{d}_1, \dots, \mathbf{d}_n$ be a basis of W .

Then $[T \circ S]_{\mathcal{D}}^{\mathcal{B}} = [T]_{\mathcal{D}}^{\mathcal{C}}[S]_{\mathcal{C}}^{\mathcal{B}}$

Here is picture of this situation:

$$\begin{array}{ccccc} \mathcal{B} & & \mathcal{C} & & \mathcal{D} \\ U & \xrightarrow{S} & V & \xrightarrow{T} & W \end{array}$$

This theorem provides some justification for our definition of matrix multiplication: composition of linear maps corresponds to multiplication of matrices.

Proof. Let $[T]_{\mathcal{D}}^{\mathcal{C}} = (t_{ij})$ and $[S]_{\mathcal{C}}^{\mathcal{B}} = (s_{ij})$. We will work out $[T \circ S]_{\mathcal{D}}^{\mathcal{B}}$ using the definition of the matrix of a linear map. For any $1 \leq c \leq l$,

$$\begin{aligned}
 (T \circ S)(\mathbf{b}_c) &= T(S(\mathbf{b}_c)) \\
 &= T\left(\sum_{k=1}^m s_{kc} \mathbf{c}_k\right) && \text{as } [S]_{\mathcal{C}}^{\mathcal{B}} = (s_{ij}) \\
 &= \sum_{k=1}^m s_{kc} T(\mathbf{c}_k) && \text{linearity of } T \\
 &= \sum_{k=1}^m s_{kc} \sum_{i=1}^n t_{ik} \mathbf{d}_i && \text{as } [T]_{\mathcal{D}}^{\mathcal{C}} = (t_{ij}) \\
 &= \sum_{i=1}^n \left(\sum_{k=1}^m t_{ik} s_{kc}\right) \mathbf{d}_i && \text{for finite sums, } \sum_k \sum_i = \sum_i \sum_k
 \end{aligned}$$

so the r, c entry of $[T \circ S]_{\mathcal{D}}^{\mathcal{B}}$ is $\sum_{k=1}^m t_{rk} s_{kc}$, which is the same as the r, c entry of $[T]_{\mathcal{D}}^{\mathcal{C}} [S]_{\mathcal{C}}^{\mathcal{B}}$ by the matrix multiplication formula. \square

4.20 Change of basis

Suppose we have a linear map T from V to W and two different bases for V and two different bases for W . We can form the matrix of T with respect to the first initial and final bases, and the second. These record the same information (the linear map T) in different ways, so they should be related in some way.

4.20.1 The change of basis formula

Let

- $T : V \rightarrow W$ be a linear map,
- \mathcal{B} and \mathcal{B}' be bases of V , and
- \mathcal{C} and \mathcal{C}' be bases of W .

Now make the following observation:

$$T = \text{id}_W \circ T \circ \text{id}_V$$

which holds purely because composing with an identity map doesn't change anything.

Now apply Theorem 4.19.1 from the previous section twice: you get the **change of basis formula**:

$$[T]_{\mathcal{C}'}^{\mathcal{B}'} = [\text{id}_W]_{\mathcal{C}'}^{\mathcal{C}} [T]_{\mathcal{C}}^{\mathcal{B}} [\text{id}_V]_{\mathcal{B}}^{\mathcal{B}'} \quad (4.7)$$

4.20.2 The matrix of the identity map with respect to different bases

In this subsection we're going to work an example of computing matrices of linear maps using the change of basis formula. On the way we'll see the significance of the matrix of the identity map with respect to different bases.

Let $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the linear map

$$T \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} -16x + 6y \\ -45x + 17y \end{pmatrix}$$

Let \mathcal{E} be the standard basis $\mathbf{e}_1, \mathbf{e}_2$ of \mathbb{R}^2 and let \mathcal{F} be the basis $\mathbf{f}_1 = \begin{pmatrix} 2 \\ 5 \end{pmatrix}, \mathbf{f}_2 = \begin{pmatrix} 1 \\ 3 \end{pmatrix}$. The matrix of T with respect to \mathcal{E} is easy to find:

$$\begin{aligned} T(\mathbf{e}_1) &= T \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} -16 \\ -45 \end{pmatrix} \\ &= -16\mathbf{e}_1 - 45\mathbf{e}_2 \\ T(\mathbf{e}_2) &= T \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} 6 \\ 17 \end{pmatrix} \\ &= 6\mathbf{e}_1 + 17\mathbf{e}_2 \end{aligned}$$

$$\text{so } [T]_{\mathcal{E}}^{\mathcal{E}} = \begin{pmatrix} -16 & 6 \\ -45 & 17 \end{pmatrix}.$$

$[\text{id}]_{\mathcal{E}}^{\mathcal{F}}$ is also easy: it's the matrix which tells us how to express the elements of \mathcal{F} in terms of the standard basis.

$$\begin{aligned} \text{id}(\mathbf{f}_1) &= \mathbf{f}_1 = 2\mathbf{e}_1 + 5\mathbf{e}_2 \\ \text{id}(\mathbf{f}_2) &= \mathbf{f}_2 = 1\mathbf{e}_1 + 3\mathbf{e}_2 \end{aligned}$$

$$\text{and so } [\text{id}]_{\mathcal{E}}^{\mathcal{F}} = \begin{pmatrix} 2 & 1 \\ 5 & 3 \end{pmatrix}.$$

How to express the \mathbf{e}_i in terms of the \mathbf{f}_i isn't so obvious, so on the face of it computing $[\text{id}]_{\mathcal{F}}^{\mathcal{E}}$ is harder. But we can avoid that, because we know about the matrix of a composition.

$$\begin{aligned} [\text{id}]_{\mathcal{F}}^{\mathcal{E}}[\text{id}]_{\mathcal{E}}^{\mathcal{F}} &= [\text{id} \circ \text{id}]_{\mathcal{F}}^{\mathcal{F}} && \text{Theorem 4.19.1} \\ &= [\text{id}]_{\mathcal{F}}^{\mathcal{F}} && \text{as } \text{id} \circ \text{id} = \text{id} \\ &= I_2 \end{aligned}$$

so

$$\begin{aligned} [\text{id}]_{\mathcal{F}}^{\mathcal{E}} &= ([\text{id}]_{\mathcal{E}}^{\mathcal{F}})^{-1} \\ &= \begin{pmatrix} 2 & 1 \\ 5 & 3 \end{pmatrix}^{-1} \\ &= \begin{pmatrix} 3 & -1 \\ -5 & 2 \end{pmatrix} \end{aligned}$$

We could work out $[T]_{\mathcal{F}}^{\mathcal{F}}$ directly using the definition, but instead we are going to practise using the change of basis formula (4.7). It says

$$\begin{aligned} [T]_{\mathcal{F}}^{\mathcal{F}} &= [\text{id}]_{\mathcal{F}}^{\mathcal{E}} [T]_{\mathcal{E}}^{\mathcal{E}} [\text{id}]_{\mathcal{E}}^{\mathcal{F}} \\ &= \begin{pmatrix} 3 & -1 \\ -5 & 2 \end{pmatrix} \begin{pmatrix} -16 & 6 \\ -45 & 17 \end{pmatrix} \begin{pmatrix} 2 & 1 \\ 5 & 3 \end{pmatrix} \\ &= \begin{pmatrix} -1 & 0 \\ 0 & 2 \end{pmatrix} \end{aligned}$$

Now consider $[T]_{\mathcal{F}}^{\mathcal{E}}$. Again we could find it directly from the definition by computing $T(\mathbf{e}_1)$ and $T(\mathbf{e}_2)$ and expressing them in terms of the \mathbf{f}_i s. But we already have the information we need: by Theorem 4.19.1,

$$\begin{aligned} [T]_{\mathcal{F}}^{\mathcal{E}} &= [T \circ \text{id}]_{\mathcal{F}}^{\mathcal{E}} \\ &= [T]_{\mathcal{F}}^{\mathcal{F}} [\text{id}]_{\mathcal{F}}^{\mathcal{E}} \\ &= \begin{pmatrix} -1 & 0 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} 3 & -1 \\ -5 & 2 \end{pmatrix} \\ &= \begin{pmatrix} -3 & 1 \\ -10 & 4 \end{pmatrix} \end{aligned}$$

To check our answer we compute $T(\mathbf{e}_1)$, which is $\begin{pmatrix} -16 \\ -45 \end{pmatrix}$. If the matrix is correct this should be the same as $-3\mathbf{f}_1 - 10\mathbf{f}_2$, and you can check that it really is.

4.20.3 Why would we use different bases to represent a linear map?

We already saw, when we first met bases of vector spaces, that different bases of vector spaces give us a different perspective on their elements — recall the example about two-pixel images. The same idea applies to linear maps.

A linear transformation which looks complex with respect to one basis can become much easier to understand when you choose the correct basis.

Example 4.20.1. $A = \begin{pmatrix} 4 & -3 & -3 \\ 3 & -2 & -3 \\ -1 & 1 & 2 \end{pmatrix}$. Consider $T_A : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, so the

matrix of T_A with respect to the standard basis of \mathbb{R}^3 is A . There is no obvious structure to T_A .

Now consider a new basis \mathcal{B} of \mathbb{R}^3 : $\mathbf{b}_1 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}$, $\mathbf{b}_2 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$, $\mathbf{b}_3 = \begin{pmatrix} -3 \\ -3 \\ 1 \end{pmatrix}$.

(You should check it really is a basis.) Let's find the matrix of T_A with respect to \mathcal{B} , that is, $[T_A]_{\mathcal{B}}^{\mathcal{B}}$.

You can check that $T_A(\mathbf{b}_1) = \mathbf{b}_1, T_A(\mathbf{b}_2) = \mathbf{b}_2, T_A(\mathbf{b}_3) = 2\mathbf{b}_3$. Thus $[T_A]_{\mathcal{B}}^{\mathcal{B}} =$

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

Suddenly the behaviour of T_A is clear. Vectors in the direction $\begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$ are unchanged by T_A , vectors in the direction \mathbf{b}_3 are scaled by a factor of 2.

This technique is called **diagonalisation**, you will learn more about it when you study eigenvectors and eigenvalues in Algebra 2.

4.21 Coordinate isomorphisms

4.21.1 Vector space isomorphisms

Consider the real vector space of all height 3 column vectors, and the real vector space of all height 3 row vectors. These are different vector spaces but they are essentially the same: there's no vector space property which one has but the other doesn't.

A way to make this precise is to observe that there is a bijective linear map between them, namely the transpose which sends $\begin{pmatrix} x \\ y \\ z \end{pmatrix}$ to $(x \ y \ z)$.

Definition 4.21.1. • Linear maps which are bijections are called **vector space isomorphisms**, or just isomorphisms.

- If there is an isomorphism $U \rightarrow V$, we say that U and V are **isomorphic** and write $U \cong V$.

Isomorphic vector spaces share the same vector space properties — for example, they always have the same dimension.

Lemma 4.21.1. *If $U \cong V$ then $\dim U = \dim V$.*

Proof. There's a linear bijection $T : U \rightarrow V$. $\ker T = \{\mathbf{0}_U\}$, so applying the rank-nullity theorem $\dim U = \dim \ker T + \dim \operatorname{im} T = \dim \operatorname{im} T$. Since $\ker T = \{\mathbf{0}_U\}$ we get $\dim U = \dim \operatorname{im} T$, but T is onto so $\operatorname{im} T = V$ and $\dim U = \dim V$. \square

4.21.2 Coordinate isomorphisms

If V is a finite-dimensional \mathbb{F} -vector space with a basis $\mathcal{B} = \mathbf{b}_1, \dots, \mathbf{b}_n$ then every $\mathbf{v} \in V$ can be written uniquely as $a_1\mathbf{b}_1 + \dots + a_n\mathbf{b}_n$ for some $a_i \in \mathbb{F}$. The column vector

$$[\mathbf{v}]_{\mathcal{B}} = \begin{pmatrix} a_1 \\ \vdots \\ a_n \end{pmatrix}$$

is called the **coordinate vector** of \mathbf{v} with respect to \mathcal{B} . We can use this idea to define an isomorphism between V and \mathbb{F}^n .

Proposition 4.21.2. *Let V be a \mathbb{F} -vector space with basis $\mathcal{B} = \mathbf{b}_1, \dots, \mathbf{b}_n$. Define $[-]_{\mathcal{B}} : V \rightarrow \mathbb{F}^n$ to be the function that sends $\mathbf{v} \in V$ to its coordinate vector $[\mathbf{v}]_{\mathcal{B}}$. Then $[-]_{\mathcal{B}}$ is a vector space isomorphism.*

$[-]_{\mathcal{B}}$ is called the coordinate isomorphism with respect to \mathcal{B} .

Proof. First we have to show it is linear.

- Let $\mathbf{x}, \mathbf{y} \in V$ and let $\mathbf{x} = \sum_i x_i \mathbf{b}_i$ and $\mathbf{y} = \sum_i y_i \mathbf{b}_i$, so that $[\mathbf{x}]_{\mathcal{B}} = (x_1, \dots, x_n)^T$ and $[\mathbf{y}]_{\mathcal{B}} = (y_1, \dots, y_n)^T$. Then $\mathbf{x} + \mathbf{y} = \sum_i (x_i + y_i) \mathbf{b}_i$, so

$$\begin{aligned} [\mathbf{x} + \mathbf{y}]_{\mathcal{B}} &= (x_1 + y_1, \dots, x_n + y_n)^T \\ &= (x_1, \dots, x_n)^T + (y_1, \dots, y_n)^T \\ &= [\mathbf{x}]_{\mathcal{B}} + [\mathbf{y}]_{\mathcal{B}}. \end{aligned}$$

- Let $\mathbf{x} \in V$ and $\lambda \in \mathbb{F}$. If $\mathbf{x} = \sum_i x_i \mathbf{b}_i$ so that $[\mathbf{x}]_{\mathcal{B}} = (x_1, \dots, x_n)^T$, then $\lambda \mathbf{x} = \sum_i \lambda x_i \mathbf{b}_i$ so $[\lambda \mathbf{x}]_{\mathcal{B}} = (\lambda x_1, \dots, \lambda x_n)^T = \lambda (x_1, \dots, x_n)^T = \lambda [\mathbf{x}]_{\mathcal{B}}$.

Now we have to show it is injective and surjective.

- If $[\mathbf{v}]_{\mathcal{B}} = \mathbf{0}_n$ then by definition $\mathbf{v} = 0\mathbf{b}_1 + \dots + 0\mathbf{b}_n = \mathbf{0}_V$. Therefore $\ker[-]_{\mathcal{B}} = \{\mathbf{0}_V\}$ and by Proposition 4.14.3, $[-]_{\mathcal{B}}$ is injective.
- Any vector $(x_1, \dots, x_n)^T$ in \mathbb{F}^n is the image of $\sum_i x_i \mathbf{b}_i \in V$ under $[-]_{\mathcal{B}}$.

□

In particular, every finite-dimensional vector space is isomorphic to a space of column vectors.

4.21.3 Coordinate and matrices of linear maps

Suppose we have \mathbb{F} -vector spaces U, V of dimensions n and m with bases $\mathcal{B} = \mathbf{b}_1, \dots, \mathbf{b}_n, \mathcal{C} = \mathbf{c}_1, \dots, \mathbf{c}_m$ and a linear map $T : U \rightarrow V$. We have coordinate isomorphisms $[-]_{\mathcal{B}} : U \rightarrow \mathbb{F}^n$ and $[-]_{\mathcal{C}} : V \rightarrow \mathbb{F}^m$ and a matrix $[T]_{\mathcal{C}}^{\mathcal{B}} \in M_{m \times n}(\mathbb{F})$. What is the relationship between $T : U \rightarrow V$ and the linear map $\mathbb{F}^n \rightarrow \mathbb{F}^m$ given by left-multiplication by $[T]_{\mathcal{C}}^{\mathcal{B}}$?

Theorem 4.21.3. *In the notation above, for any $\mathbf{u} \in U$ we have $[T]_{\mathcal{C}}^{\mathcal{B}}[\mathbf{u}]_{\mathcal{B}} = [T(\mathbf{u})]_{\mathcal{C}}$.*

Proof. Let $[T]_{\mathcal{C}}^{\mathcal{B}} = (a_{ij})$, so that $T(\mathbf{b}_j) = \sum_{i=1}^m a_{ij} \mathbf{c}_i$. Notice that this says $[T(\mathbf{b}_j)]_{\mathcal{C}}$ is $(a_{1j}, \dots, a_{mj})^T$, the j th column of $[T]_{\mathcal{C}}^{\mathcal{B}}$. The coordinate vector $[\mathbf{b}_j]_{\mathcal{B}}$ is the j th standard basis vector \mathbf{e}_j , so $[T]_{\mathcal{C}}^{\mathcal{B}}[\mathbf{b}_j]_{\mathcal{B}} = [T]_{\mathcal{C}}^{\mathcal{B}}\mathbf{e}_j$ which equals the j th column of $[T]_{\mathcal{C}}^{\mathcal{B}}$. Therefore

$$[T]_{\mathcal{C}}^{\mathcal{B}}[\mathbf{b}_j]_{\mathcal{B}} = [T(\mathbf{b}_j)]_{\mathcal{C}}. \quad (4.8)$$

Let $\mathbf{u} \in U$ and write $\mathbf{u} = \sum_j u_j \mathbf{b}_j$ so that $[\mathbf{u}]_{\mathcal{B}} = (u_1, \dots, u_n)^T$. Then

$$\begin{aligned}
 [T(\mathbf{u})]_{\mathcal{C}} &= \left[T \left(\sum_j u_j \mathbf{b}_j \right) \right] \\
 &= \left[\sum_j u_j T(\mathbf{b}_j) \right]_{\mathcal{C}} && \text{linearity of } T \\
 &= \sum_j u_j [T(\mathbf{b}_j)]_{\mathcal{C}} && \text{linearity of } [-]_{\mathcal{C}} \\
 &= \sum_j u_j [T]_{\mathcal{C}}^{\mathcal{B}} [\mathbf{b}_j]_{\mathcal{B}} && (4.8) \\
 &= [T]_{\mathcal{C}}^{\mathcal{B}} \sum_j u_j [\mathbf{b}_j] \\
 &= [T]_{\mathcal{C}}^{\mathcal{B}} \left[\sum_j u_j \mathbf{b}_j \right]_{\mathcal{B}} && \text{linearity of } [-]_{\mathcal{B}} \\
 &= [T]_{\mathcal{C}}^{\mathcal{B}} [\mathbf{u}]_{\mathcal{B}}.
 \end{aligned}$$

□

We can now answer the following natural question: if \mathcal{B} and \mathcal{C} are two bases of the finite-dimensional vector space V and $\mathbf{v} \in V$, what is the relationship between $[\mathbf{v}]_{\mathcal{B}}$ and $[\mathbf{v}]_{\mathcal{C}}$? Using the previous theorem applied to the identity map $\text{id} : V \rightarrow V$,

$$\begin{aligned}
 [\mathbf{v}]_{\mathcal{C}} &= [\text{id}(\mathbf{v})]_{\mathcal{C}} \\
 &= [\text{id}]_{\mathcal{C}}^{\mathcal{B}} [\mathbf{v}]_{\mathcal{B}}.
 \end{aligned}$$

Further reading

Most of the recommendations from the previous chapter are also relevant to the material in this one. If you want to take linear algebra further and you like the style, the books *Introduction to Linear Algebra*, *Linear Algebra and Learning From Data*, and *Linear Algebra and its Applications* by G. Strang might be good for you. The second is especially relevant if you are interested in AI/ML.

You will learn more about linear algebra and matrices in MATH0006 Algebra 2 next term, and there are more advanced linear algebra courses in subsequent years such as MATH0014 Further Linear Algebra and MATH0058 Computational Methods. Full details and syllabuses can be found on my pathways webpage.