

Linear Algebra

Joseph R. Mileti

August 28, 2024

Contents

1	Introduction	5
1.1	What is Linear Algebra?	5
1.2	Mathematical Statements and Mathematical Truth	10
1.3	Quantifiers and Proofs	14
1.4	Evens and Odds	24
1.5	Sets, Set Construction, and Subsets	31
1.6	Functions	37
1.7	Injective, Surjective, and Bijective Functions	44
1.8	Solving Equations	48
2	Spans and Linear Transformations in Two Dimensions	51
2.1	Intersections of Lines in \mathbb{R}^2	51
2.2	Vectors in \mathbb{R}^2	54
2.3	Spans	56
2.4	Linear Transformations of \mathbb{R}^2	68
2.5	The Standard Matrix of a Linear Transformation	82
2.6	Matrix Algebra	91
2.7	Range, Null Space, and Inverses	101
3	Coordinates and Eigenvectors in Two Dimensions	115
3.1	Coordinates and Change of Basis	115
3.2	Matrices with Respect to Other Coordinates	122
3.3	Eigenvalues and Eigenvectors	132
3.4	Determinants	150
4	Beyond Two Dimensions	157
4.1	Vector Spaces and Subspaces	157
4.2	Solving Linear Systems	169
4.3	Spanning Sequences	185
4.4	Linear Independence	189
4.5	Bases and Coordinates	196
4.6	Dimension	205
5	Linear Transformations and Matrices	213
5.1	Linear Transformations	213
5.2	The Range and Null Space of a Linear Transformation	230
5.3	Determinants	244
5.4	Eigenvalues and Eigenvectors	253

Chapter 1

Introduction

1.1 What is Linear Algebra?

The most elementary, yet honest, way to describe linear algebra is that it is the basic mathematics of high dimensions. By “basic”, we do not mean that the theory is easy, but only that it is essential to a more nuanced understanding of the mathematics of high dimensions. For example, the simplest curves in two dimensions are lines, and we use lines to approximate more complicated curves in Calculus. Similarly, the most basic surfaces in three dimensions are planes, and we use planes to approximate more complicated surfaces. One of our major goals will be to generalize the concepts of lines and planes to the “flat” objects in high dimensions. Another major goal will be to understand the simplest kinds of functions (so-called *linear transformations*) that arise in this setting. If this sounds abstract and irrelevant to real world problems, we will explain in the rest of this section, and throughout the course, why these concepts are incredibly important to mathematics and its applications.

Generalizing Lines and Planes

Let’s begin by recalling how we can describe a line in two dimensions. One often thinks about a line as the solution set to an equation of the form $y = mx + b$ (where m and b are fixed numbers). Although most lines in two dimensions can arise from such equations, these descriptions omit vertical lines. A better, more symmetric, and universal way to describe a line is as the set of solutions to an equation of the form $ax + by = c$, where at least one of a or b is nonzero. For example, we can now describe the vertical line $x = 5$ using the numbers $a = 1$, $b = 0$, and $c = 5$. Notice that if b is nonzero, then we can equivalently write this equation as $y = (-\frac{a}{b})x + \frac{c}{b}$, which fits into the above model. Similarly, if a is nonzero, then we can solve for x in terms of y .


With this approach, we are thinking about a line as described algebraically by an equation. However, there is another, more geometric, way to describe a line in two dimensions. Start with a nonzero vector \vec{v} in the plane \mathbb{R}^2 (with its tail at the origin), and think about taking the collection of all scalar multiples of \vec{v} . In other words, stretch the vector \vec{v} in all possible ways, including switching its direction around using negative scalars, and consider all possible outcomes. In set-theoretic notation, we are forming the set $\{t \cdot \vec{v} : t \in \mathbb{R}\}$ (if this symbolism is unfamiliar to you, we will discuss set-theoretic notation in detail later). When viewed as a collection of points, this set consists of a line through the origin in the direction of \vec{v} .

For example, consider the vector \vec{v} in the plane whose first coordinate is 1 and whose second coordinate is 2. In past courses, you may have written \vec{v} as the vector $\langle 1, 2 \rangle$, and thought about it as the position vector

of the point $(1, 2)$. In linear algebra, we will typically write such vectors vertically as

$$\vec{v} = \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

for reasons that we will explain later. Moreover, we will typically blur the distinction the point $(1, 2)$ with the position vector of this, i.e. the vector whose tail is at the origin and whose head is at $(1, 2)$ (see Section 2.2). If we visualize this vector with its tail at the origin, we obtain the picture on the left:



Now consider taking all multiples of the vector \vec{v} . That is, we consider the vector $2\vec{v}$, the vector $(-1)\vec{v} = -\vec{v}$, the vector $.24\vec{v}$, etc. If we collect *all* of these vectors into a set, we obtain


$$\left\{ t \cdot \begin{pmatrix} 1 \\ 2 \end{pmatrix} : t \in \mathbb{R} \right\} = \left\{ \begin{pmatrix} t \\ 2t \end{pmatrix} : t \in \mathbb{R} \right\}.$$

If we think about the corresponding set of points, we obtain a line through the origin, which is depicted in the above picture on the right. Notice that every point on this line has second coordinate equal to 2 times the first coordinate, and so this line equals the solution set to $y = 2x$. Written in our more symmetric notation, our line is the solution set of $2x - y = 0$.

What if we want a line that is not through the origin? We then add an “offset” vector to serve as a common shift to all the points on a line through the origin. If we call our offset vector \vec{p} , then we need to add \vec{p} to each of the points on the original line $\{t \cdot \vec{v} : t \in \mathbb{R}\}$ to obtain $\{\vec{p} + t\vec{v} : t \in \mathbb{R}\}$. For example, if $\vec{p} = \begin{pmatrix} 4 \\ 1 \end{pmatrix}$, then we can consider

$$\left\{ \begin{pmatrix} 4 \\ 1 \end{pmatrix} + t \cdot \begin{pmatrix} 1 \\ 2 \end{pmatrix} : t \in \mathbb{R} \right\} = \left\{ \begin{pmatrix} 4+t \\ 1+2t \end{pmatrix} : t \in \mathbb{R} \right\}.$$

Visually, we are taking our original line, and adding shifting the entire line over through the addition of \vec{p} , which is represented as a dotted vector below:



The above set-theoretic notation might remind you of the parametric equations

$$\begin{aligned}x &= 4 + t \\ y &= 1 + 2t,\end{aligned}$$

which is an equivalent way to describe the points on the line in terms of a parameter t .

If we try to generalize equations of the form $ax + by = c$ to three dimensions, we obtain equations of the form $ax + by + cz = d$. In this case, the solution set is a plane (as long as at least one of a, b, c is nonzero) that can be geometrically described as having $\langle a, b, c \rangle$ as a normal vector. Thus, instead of describing a 1-dimensional line in 2-dimensional space, this generalization describes a 2-dimensional plane in 3-dimensional space.

In contrast, our parametric description of lines in 2-dimensions carry over nicely to lines in 3-dimensions. If we fix a vector \vec{v} in 3-dimensional space \mathbb{R}^3 , and consider the set of all scalar multiples of \vec{v} , then we obtain a line in 3-dimensions through the origin. As above, we can add an offset to describe lines more generally. For example, you may be used to describing the line going through the point $(7, -2, 1)$ with direction vector $\langle 5, 0, 3 \rangle$ by the parametric equations:

$$\begin{aligned}x &= 7 + 5t \\ y &= -2 \\ z &= 1 + 3t.\end{aligned}$$

Alternatively, we can describe the points on the line by the following set:

$$\left\{ \begin{pmatrix} 7 \\ -2 \\ 1 \end{pmatrix} + t \cdot \begin{pmatrix} 5 \\ 0 \\ 3 \end{pmatrix} : t \in \mathbb{R} \right\} = \left\{ \begin{pmatrix} 7 + 5t \\ -2 \\ 1 + 3t \end{pmatrix} : t \in \mathbb{R} \right\}.$$

Thus, our two different ways to describe a line in 2-dimensions generalize to different objects in 3-dimensions. Can we describe a line in 3-dimensions in a way other than the parametric one we just established? One way is to recognize a line as the intersection of two (nonparallel) planes. Algebraically, we are then defining a line as the common solution set to two equations:

$$\begin{aligned}a_1x + b_1y + c_1z &= d_1 \\ a_2x + b_2y + c_2z &= d_2,\end{aligned}$$

with appropriate conditions that enforce that the planes are not parallel.

Can we also describe a plane parametrically? Consider an arbitrary plane in \mathbb{R}^3 through the origin. Think about taking two nonzero and nonparallel vectors \vec{u} and \vec{w} parallel to the plane, and then using them to “sweep out” the rest of the plane. In other words, we want to focus on \vec{u} and \vec{w} , and then stretch and add them in all possible ways to “fill in” the remaining points, similar to how we filled in points in 2-dimensions by just stretching one vector. In set-theoretic notation, we are considering the set $\{t\vec{u} + s\vec{w} : t, s \in \mathbb{R}\}$. For example, if

$$\vec{u} = \begin{pmatrix} 2 \\ -3 \\ 1 \end{pmatrix} \quad \text{and} \quad \vec{w} = \begin{pmatrix} -7 \\ 1 \\ 4 \end{pmatrix},$$

then we have the set

$$\left\{ t \cdot \begin{pmatrix} 2 \\ -3 \\ 1 \end{pmatrix} + s \cdot \begin{pmatrix} -7 \\ 1 \\ 4 \end{pmatrix} : t, s \in \mathbb{R} \right\} = \left\{ \begin{pmatrix} 2t - 7s \\ -3t + s \\ t + 4s \end{pmatrix} : t, s \in \mathbb{R} \right\}.$$

This idea of combining vectors in all possible ways through scaling and adding will be a fundamental tool for us. As above, if we want to think about planes in general (not just through the origin), then we should add an offset.

With all of this in mind, can we generalize the ideas of lines and planes to higher dimensions? What would be the analogue of a plane in 4-dimensions? How can we describe a 3-dimensional “flat” object (like a line or a plane) in 7-dimensions? Although these are fascinating pure mathematics questions, you may wonder why we would care.

In calculus, you learned how to compute projections in \mathbb{R}^2 and \mathbb{R}^3 based on the dot product (such as how to project one vector onto another, or how to project a point onto a plane). One of our goals will be to generalize the ideas behind dot products and projections to higher dimensions. An important situation where this arises in practice is how to fit the “best” possible curve to some data points. For instance, suppose that we have n data points and want to fit the “best” line, parabola, etc. to these points. We will define “best” by saying that it minimizes a certain “distance”, defined in terms of our generalized dot product (this is analogous to the fact that the projection of a point onto a plane minimizes the distance between the given point and the projected point). By viewing the collection of all possible lines as a certain 2-dimensional object inside of \mathbb{R}^n , we will see that we can take a similar approach in \mathbb{R}^n . Hence, fitting a line to n points can be thought of as projecting onto a certain “generalized plane” in \mathbb{R}^n . Similarly, fitting a parabola amounts to projecting a point in \mathbb{R}^n onto a 3-dimensional object. These examples demonstrate the power that arises from understanding high dimensions. Moreover, these ideas lie at the heart of many other applications, such as filling in missing data in order to reconstruct parts of an image that have been lost, or predicting which movies you might like on Netflix in order to provide recommendations.

Transformations of Space

In Calculus, you studied functions $f: \mathbb{R} \rightarrow \mathbb{R}$ where both the input and output are elements of \mathbb{R} . In Multivariable Calculus, you looked at functions $f: \mathbb{R} \rightarrow \mathbb{R}^n$ (where the inputs are numbers and the outputs are elements of n -dimensional space) which can be thought of as parametric descriptions of curves in \mathbb{R}^n . For example, the function $f: \mathbb{R} \rightarrow \mathbb{R}^2$ given by $f(t) = (\cos t, \sin t)$ can be viewed as a parametric description of the unit circle. After that, you investigated functions $f: \mathbb{R}^n \rightarrow \mathbb{R}$ where, for instance, the graph of a function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ like $f(x, y) = x^2 + y^2$ can be visualized as a surface in \mathbb{R}^3 . Perhaps at the end of Multivariable Calculus you started to think about functions $f: \mathbb{R}^m \rightarrow \mathbb{R}^n$. For example, consider the function $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by

$$f(x, y) = (x - y, x + y).$$


In our new notation where we write vectors vertically, we could write this function as follows:

$$f\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} x - y \\ x + y \end{pmatrix}.$$

Notice that we have double parentheses on the left because we are thinking of the vector $\begin{pmatrix} x \\ y \end{pmatrix}$ as one input to the function f . For example, we have

$$f\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad f\left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} 0 \\ 2 \end{pmatrix}, \quad \text{and} \quad f\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

In Calculus, you might have viewed these as vector fields where we put the vector $\langle 1, 1 \rangle$ at the point $(1, 0)$, put the vector $\langle 0, 2 \rangle$ at the point $(1, 1)$, etc. However, we want to think about the function a bit differently now by interpreting it as “moving” the plane via operations like rotations, projections, and scalings. Visually, think about having two copies of the plane, and imagine all vectors as having their tails at the origin. On the left we have the input vectors, and on the right we have the output vectors.



In the above picture, we have used the same three input/output pairs calculated above. From these three sample inputs, it appears that the above function rotates the plane by 45° counterclockwise, and simultaneously scales the plane by a factor of $\sqrt{2}$. We will eventually develop techniques to understand and analyze these so-called “linear” ways to transform space, which will allow us to easily verify that *every* vector is rotated and scaled as just described. More interestingly, we will be able to easily determine formulas for functions that have rotation and scaling factors. As we will see, matrices will play a key role in representing, or coding, these kinds of functions.

How do these types of functions arise in practice? One simple example is in computer graphics. Many 3-dimensional images in a computer are formed by fitting polygons across points that serve as a skeleton for the underlying model. Now if we want to change the position of the camera or move objects around, then we need to alter the position of the skeletal points. In order to do this, we need to apply a certain transformation $f: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ to each of them. If we rotate the camera, then we are applying a rotation in \mathbb{R}^3 by a given angle along a certain axis. How can we compute f so that we can implement it quickly?

Another example arises when changing coordinates. For example, when setting up a problem in physics, one needs to define axes to orient the objects under consideration. Two different people might set up the

axes differently, and hence might use different coordinates for the position of a given object. I might label an object by the coordinates $(3, 1, -8)$ in my system, while you label it as $(-2, 5, 0)$ in your system. We would like to have a way to translate between our different ways of measuring, and it turns out that we can do this by a “linear” transformation (sometimes with an offset similar to that used in our lines and planes) from \mathbb{R}^3 to \mathbb{R}^3 . Working out how to do this is not only important for calculations, but is essential in understanding the invariance of physical laws under certain changes of coordinates.

In an even more applied setting, these kinds of transformation arise in many modeling situations. Consider a city made up of 4 districts. Suppose that some people move between the districts in each given year but (for simplicity here) that the total population in the city is constant. If we know the rate of movement between the districts, we can record all of the information in a 4×4 table:

	1	2	3	4
1	86%	3%	7%	4%
2	2%	75%	11%	12%
3	3%	1%	93%	3%
4	2%	2%	5%	91%

We interpret this table by using the entry in row i and column j to mean the percentage of people in district i that move to district j in a given year. Notice that if we know the current population distribution in the current year, then we can compute the population distribution in the next year. As a result, this table defines a certain function from \mathbb{R}^4 to \mathbb{R}^4 . Also, if our city has 12 districts, then we can view the situation as a transformation from \mathbb{R}^{12} to \mathbb{R}^{12} . One natural question that arises in this modeling situation is to ask what happens to the population over time. What will be the distribution in 10 years under these assumptions? What about 100 years? Is there some kind of stability that happens in the long run?

A similar situation arises when analyzing games like Monopoly. Suppose that we label the squares with the numbers from 1 to 40 around the board, starting with “Go” as 1. After one roll of the pair of dice, there is a certain probability that we are on a given square. For example, there is a $\frac{1}{18}$ chance that we will be on square 4, and a $\frac{5}{36}$ chance that we will be on square 7. In general, given a 40-dimensional vector representing the probability that we are on each square at a certain stage of the game, we can determine another 40-dimensional vector representing the probability that we are on each square at the next stage of the game. In this way we have defined a function from \mathbb{R}^{40} to \mathbb{R}^{40} . Understanding the long term behavior arising from iterating this function tells us which squares of the game receive the most traffic.

1.2 Mathematical Statements and Mathematical Truth

Unfortunately, many people view mathematics only as a collection of complicated equations and elaborate computational techniques (or algorithms) that lead to the correct answers to a narrow class of problems. Although these are indeed aspects of mathematics, they do not reflect the fundamental nature of the subject. Mathematics, at its core, is about determining *truth*, at least for certain precise mathematical statements. Before we consider some examples, let’s recall some notation and terminology for the standard “universes” of numbers:

- $\mathbb{N} = \{0, 1, 2, 3, \dots\}$ is the set of *natural numbers*.
- $\mathbb{N}^+ = \{1, 2, 3, \dots\}$ is the set of positive natural numbers.
- $\mathbb{Z} = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$ is the set of *integers*.
- \mathbb{Q} is the set of *rational* numbers, i.e. those numbers that can be written as a fraction (i.e. quotient) of integers. For example, $\frac{1}{2}$, $\frac{-3}{17}$, etc. are all rational numbers. Notice that all integers are rational numbers because we can view 5 as $\frac{5}{1}$, for example.

- \mathbb{R} is the set of *real* numbers, i.e. those numbers that we can express as a (possibly infinite) decimal. Every rational number is a real number, but π , e , $\sqrt{2}$, etc. are all real numbers that are not rational.

There are other important universes of numbers, such as the complex numbers \mathbb{C} , and many more that you will encounter in Abstract Algebra. However, we will focus on the above examples in our study. To denote that a given number n belongs to one of the above collections, we will use the \in symbol. For example, we will write $n \in \mathbb{Z}$ as shorthand for “ n is an integer”. We will elaborate on how to use the symbol \in more broadly when we discuss general set theory notation.

Returning to our discussion of truth, a mathematical statement is either objectively true or objectively false, without reference to the outside world and without any additional conditions or information. For some examples, consider the following (we’ve highlighted some key words that we will discuss in the next few sections):

1. $35 + 81$ is equal to 116.
2. The sum of two odd integers is **always** an even integer.
3. The difference of two prime numbers is **always** an even integer.
4. **There exists** a simultaneous solution to the three equations

$$\begin{array}{rcccccc} 2x & & & + & 8z & = & 6 \\ 7x & - & 3y & + & 18z & = & 15 \\ -3x & + & 3y & - & 2z & = & -1 \end{array}$$

in \mathbb{R}^3 , i.e. **there exists** a choice of real numbers for x , y , and z making all three equations true.

5. The remainder when dividing 333^{2856} by 2857 is 1.
6. **Every** continuous function is differentiable.
7. **Every** differentiable function is continuous.
8. **There exist** positive natural numbers a, b, c with $a^3 + b^3 = c^3$.
9. The digits of π eventually form a repeating sequence.
10. The values of $0, 1, 2, \dots, 9$ occur with equal frequency (i.e. each about $\frac{1}{10}$ of the time) in the infinite decimal expansion of π .

Which of these 10 assertions are true and which are false? In many cases, the answer is not obvious. Here are the results:

1. True. This statement can be verified by a simple calculation.
2. True. However, it’s not immediately obvious how we could ever verify it. After all, there are infinitely many odd numbers, so we can’t simply try them all.
3. False. To show that it is false, it suffices to give just one counterexample. Notice that 7 and 2 are prime, but $7 - 2 = 5$ and 5 is not even.
4. False. Again, it may not be obvious how to show that *no* possible choice of x , y , and z exist. We will develop systemic ways to solve such problems later.
5. True. It is possible to verify this by calculation (by using a suitably programmed computer). However, there are better ways to understand why this is true, as you will see in Elementary Number Theory or Abstract Algebra.

6. False. The function $f(x) = |x|$ is continuous everywhere but is not differentiable at 0.
7. True. See Calculus or Analysis.
8. False. This is a special case of something called Fermat's Last Theorem, and it is quite difficult to show (see Algebraic Number Theory).
9. False. This follows from the fact that π is an irrational number, i.e. not an element of \mathbb{Q} , but this is not easy to show.
10. We still don't know whether this is true or false! Numerical evidence (checking the first billion digits directly, for example) suggests that it may be true. Mathematicians have thought about this problem for a century, but we still do not know how to answer it definitively.

Recall that a mathematical statement must be either true or false. In contrast, an equation is typically neither true nor false when viewed in isolation, and hence is not a mathematical statement. For example, it makes no sense to ask whether $y = 2x + 3$ is true or false, because it depends on which numbers we plug in for x and y . When $x = 6$ and $y = 15$, then the statement becomes true, but when $x = 3$ and $y = 7$, the statement is false. For a more interesting example, the equation

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

is not a mathematical statement as given, because we have not been told how to interpret the x and the y . Is the statement true when x is my cat Cayley and y is my cat Maschke? (Adding them together is scary enough, and I don't even want to think about what it would mean to *square* them.) In order to assign a truth value, we need to provide context for where x and y can come from. To fix this, we can write

$$\text{“For all real numbers } x \text{ and } y, \text{ we have } (x + y)^2 = x^2 + 2xy + y^2\text{”},$$

which is now a true mathematical statement. As we will eventually see, if we replace *real numbers* with 2×2 *matrices*, the corresponding statement is false.

For a related example, it is natural to think that the statement $(x + y)^2 = x^2 + y^2$ is false, but again it is not a valid mathematical statement as written. We can instead say that the statement

$$\text{“For all real numbers } x \text{ and } y, \text{ we have } (x + y)^2 = x^2 + y^2\text{”}$$

is false, because $(1 + 1)^2 = 4$ while $1^2 + 1^2 = 2$. However, the mathematical statement

$$\text{“There exist real numbers } x \text{ and } y \text{ such that } (x + y)^2 = x^2 + y^2\text{”}$$

is true, because $(1 + 0)^2$ does equal $1^2 + 0^2$. Surprisingly, there are contexts (i.e. replacing *real numbers* with more exotic number systems) where the corresponding “for all” statement is true (see Abstract Algebra).

Here are a few other examples of statements that are *not* mathematical statements:

- $F = ma$ and $E = mc^2$: Our current theories of physics say that these equations are true in the real world whenever the symbols are interpreted properly, but mathematics on its own is a different beast. As written, these equations are neither true nor false from a mathematical perspective. For example, if $F = 4$, $m = 1$, and $a = 1$, then $F = ma$ is certainly false.
- $a^2 + b^2 = c^2$: Unfortunately, most people “remember” this as the Pythagorean Theorem. However, it is not even a mathematical statement as written. We could fix it by writing “For all right triangles with side lengths a , b , and c , where c is the length the hypotenuse, we have that $a^2 + b^2 = c^2$ ”, in which case we have a true mathematical statement.

- Talking Heads is the greatest band of all time: Of course, different people can have different opinions about this. I may believe that the statement is true, but the notion of “truth” here is very different from the objective notion of truth necessary for a mathematical statement.
- Shakespeare wrote *Hamlet*: This is almost certainly true, but it’s not a mathematical statement. First, it references the outside world. Also, it’s at least conceivable that with new evidence, we might change our minds. For example, perhaps we’ll learn that Shakespeare stole the work of somebody else.

In many subjects, a primary goal is to determine whether certain statements are true or false. However, the methods for determining truth vary between disciplines. In the natural sciences, truth is often gauged by appealing to observations and experiments, and then building a logical structure (perhaps using some mathematics) to convince others of a claim. Economics arguments are built through a combination of current and historical data, mathematical modeling, and rhetoric. In both of these examples, truth is always subject to revision based on new evidence.

In contrast, mathematics has a unique way of determining the truth or falsity of a given statement: we provide an airtight, logical *proof* that verifies its truth with certainty. Once we’ve succeeded in finding a correct proof of a mathematical statement, we know that it must be true for all eternity. Unlike the natural sciences, we do not have tentative theories that are extremely well-supported but may be overthrown with new evidence. Thus, mathematics does not have the same types of revolutions like plate tectonics, evolution by natural selection, the oxygen theory of combustion (in place of phlogiston), relativity, quantum mechanics, etc. which overthrow the core structure of a subject and cause a fundamental shift in what statements are understood to be true.

To many, the fact that mathematicians require a complete logical proof with absolute certainty seems strange. Doesn’t it suffice to simply check the truth of statement in many instances, and then generalize it to a universal law? Consider the following example. One of the true statements mentioned above is that there are no positive natural numbers a, b, c with $a^3 + b^3 = c^3$, i.e. we can not obtain a cube by adding two cubes. The mathematician Leonhard Euler conjectured that a similar statement held for fourth powers, i.e. that we can not obtain a fourth power by adding three fourth powers. More formally, he conjectured that there are no positive natural numbers a, b, c, d with $a^4 + b^4 + c^4 = d^4$. For over 200 years it seemed reasonable to believe this might be true, as it held for all small examples and was a natural generalization of a true statement. However, it was eventually shown that there indeed are examples where the sum of 3 fourth powers equals a fourth power, such as the following:

$$95800^4 + 217519^4 + 414560^4 = 422481^4.$$

In fact, this example is the smallest one possible. Thus, even though $a^4 + b^4 + c^4 \neq d^4$ for all values positive natural numbers a, b, c , and d having at most 5 digits, the statement does not hold generally.

In spite of this example, you may question the necessity of proofs for mathematics relevant to the sciences and applications, where approximations and occasional errors or exceptions may not matter so much. There are many historical reasons why mathematicians have embraced complete, careful, and logical proofs as **the** way to determine truth in mathematics independently from applications. In later math classes, you may explore some of these internal historical aspects, but here are three direct reasons for this approach:

- Mathematics should exist independently from the sciences because sometimes the same mathematics applies to different subjects. It is possible that edge cases which do not matter in one subject (say economics or physics) might matter in another (like computer science). The math needs to be consistent and coherent on its own without reference to the application.
- In contrast to the sciences where two generally accepted theories that contradict each other in some instance can coexist for long periods of time (such as relativity and quantum mechanics), mathematics can not sustain such inconsistencies. As we’ll see, one reason for this is that mathematics allows a certain type of argument called proof by contradiction. Any inconsistency at all would allow us to draw all sorts of erroneous conclusions, and the logical structure of mathematics would unravel.

- Unlike the sciences, many areas of math are not subject to direct validation through a physical test. An idea in physics or chemistry, arising from either a theoretical predication or a hunch, can be verified by running an experiment. However, in mathematics we often have no way to reliably verify our guesses through such means. As a result, proofs in mathematics can be viewed as the analogues of experiments in the sciences. In other words, since mathematics exists independently from the sciences, we need an internal check for our intuitions and hunches, and proofs play this role.

1.3 Quantifiers and Proofs

In the examples of mathematical statements in the previous section, we highlighted two key phrases that appear incredibly often in mathematical statements: **for all** and **there exists**. These two phrases are called *quantifiers* in mathematics, and they form the building blocks of more complicated expressions. Occasionally, these quantifiers appear disguised by a different word choice. Here are a few phrases that mean precisely the same thing in mathematics:

- **For all:** For every, For any, Every, Always,
- **There exists:** There is, For some, We can find,

These phrases mean what you might expect. For example, saying that a statement of the form “For all a , . . .” is true means that whenever we plug in any particular value for a into the . . . part, the resulting statement is true. Similarly, saying that a statement of the form “There exists a , . . .” is true means that there is at least one (but possibly more) choice of a value to plug in for a so that the resulting statement is true. Notice that we are *not* saying that there is exactly one choice. Also, be careful in that the phrase “for some” used in everyday conversation could be construed to mean that there need to be several (i.e. more than one) values to plug in for a to make the result true, but in math it is completely synonymous with “there exists”.

So how do we prove that a statement that begins with a “there exists” quantifier is true? For example, consider the following statement:

“There exists $a \in \mathbb{Z}$ such that $2a^2 - 1 = 71$ ”.

From your training in mathematics up to this point, you may see the equation at the end and immediately rush to manipulate it using the procedures that you’ve been taught for years. Before jumping into that, let’s examine the logical structure here. As mentioned above, our “there exists” statement is true exactly when there is a concrete integer that we can plug in for a so that “ $2a^2 - 1 = 71$ ” is a true statement. Thus, in order to convince somebody that the statement is true, we need only find (at least) one particular value to plug in for a so that when we compute $2a^2 - 1$ we obtain 71. Right? In other words, if all that we care about is knowing for sure that the statement is true, we just need to verify that some $a \in \mathbb{Z}$ has this property. Suppose that we happen to stumble across the number 6 and notice that

$$\begin{aligned} 2 \cdot 6^2 - 1 &= 2 \cdot 36 - 1 \\ &= 72 - 1 \\ &= 71. \end{aligned}$$

At this point, we can assert with confidence that the statement is true, and in fact we’ve just carried out a complete proof. Now you may ask yourself “How did we know to plug in 6 there?”, and that is a good question. However, there is a difference between the creative leap (or leap of faith) we took in choosing 6, and the routine verification that it worked. Perhaps we arrived at 6 by plugging in numbers until we got lucky. Perhaps we sacrificed a chicken to get the answer. Perhaps we had a vision. Maybe you copied the answer from a friend or from online (note: don’t do this). Now we do care very much about the underlying

methods to find a , both for ethical reasons and because sacrificing a chicken may not work if we change the equation slightly. However, for the logical purposes of this argument, the way that we arrived at our value for a does not matter.

We're (hopefully) all convinced that we have verified that the statement "There exists $a \in \mathbb{Z}$ such that $2a^2 - 1 = 71$ " is true, but as mentioned we would like to have routine methods to solve similar problems in the future so that we do not have to stumble around in the dark nor invest in chicken farms. Of course, the tools to do this are precisely the material that you learned years ago in elementary algebra. One approach is to perform operations on both sides of the equality with the goal of isolating the a . If we add 1 to both sides, we arrive at $2a^2 = 72$, and after dividing both sides by 2 we conclude that $a^2 = 36$. At this point, we realize that there are two solutions, namely 6 and -6 . Alternatively, we can try bringing the 71 over and factoring. By the way, this method found two solutions, and indeed -6 would have worked above. However, remember that proving a "there exists" statement means just finding at least one value that works, so it didn't matter that there was more than one solution.

Let's consider the following more interesting example of a mathematical statement:

$$\text{"There exists } a \in \mathbb{R} \text{ such that } 2a^5 + 2a^3 - 6a^2 + 1 = 0\text{"}$$

It's certainly possible that we might get lucky and find a real number to plug in that verifies the truth of this statement. But if the chicken sacrificing doesn't work, you may be stymied about how to proceed. However, if you remember Calculus, then there is a nice way to argue that this statement is true without actually finding a particular value of a . The key fact is the Intermediate Value Theorem from Calculus, which says that if $f: \mathbb{R} \rightarrow \mathbb{R}$ is a continuous function that is positive at some point and negative at another, then it must be 0 at some point as well. Letting $f(x) = 2x^5 + 2x^3 - 6x^2 + 1$, we know from Calculus that $f(x)$ is continuous. Since $f(0) = 1$ and $f(1) = -1$, it follows from the Intermediate Value Theorem that there is an $a \in \mathbb{R}$ (in fact between 0 and 1) such that $f(a) = 0$. Thus, we've proven that the above statement is true, so long as you accept the Intermediate Value Theorem. Notice again that we've established the statement without actually exhibiting an a that works.

We can make the above question harder by performing the following small change to the statement:

$$\text{"There exists } a \in \mathbb{Q} \text{ such that } 2a^5 + 2a^3 - 6a^2 + 1 = 0\text{"}$$

Since we do not know what the value of a that worked above was, we are not sure whether it is an element of \mathbb{Q} . In fact, questions like this are a bit harder. There is indeed a method to determine the truth of a statement like this, but that's for another course (see Abstract Algebra). The takeaway lesson here is that mathematical statements that look quite similar might require very different methods to solve.

Summing up, a statement of the form "There exists a such that \dots " is true exactly when there is some concrete value of a that we can plug into " \dots " so the resulting statement is true. Thus, saying that the statement

$$\text{"There exists } a \in \mathbb{Z} \text{ such that } 2a^2 - 1 = 71\text{"}$$

is true is the same as saying that at least one of the following (infinitely many) mathematical statements is true:

- \dots
- $2 \cdot (-2)^2 - 1 = 71.$
- $2 \cdot (-1)^2 - 1 = 71.$
- $2 \cdot 0^2 - 1 = 71.$
- $2 \cdot 1^2 - 1 = 71.$

- $2 \cdot 2^2 - 1 = 71$.
- ...

Now almost all of these statements are false, but the fact that at least one of them is true (in fact, exactly two are true) means that the original “there exists” statement is true.

Let’s move on to statements involving our other quantifier. Consider the following “for all” statement:

$$\text{“For all } a, b \in \mathbb{R}, \text{ we have } (a + b)^2 = a^2 + 2ab + b^2\text{”}.$$

In the previous section, we briefly mentioned this statement, but wrote it slightly differently as:

$$\text{“For all real numbers } x \text{ and } y, \text{ we have } (x + y)^2 = x^2 + 2xy + y^2\text{”}.$$

Notice that these are both expressing the exact same thing. We only replaced the phrase “real numbers” by the symbol \mathbb{R} and changed our choice of letters. Since the letters are just placeholders for the “for all” quantifier, these two mean precisely the same thing. What does it mean to say that the first statement above is true? As mentioned above, our “for all” statement is true exactly when *whenever* we plug in concrete real numbers for a and b into “ $(a + b)^2 = a^2 + 2ab + b^2$ ”, the result is a true statement. In other words, this is the same as saying that every single one of the following (infinitely many) mathematical statements are true:

- ...
- $(3 + 7)^2 = 3^2 + 2 \cdot 3 \cdot 7 + 7^2$.
- $(1 + 0)^2 = 1^2 + 2 \cdot 1 \cdot 0 + 0^2$.
- $((-5) + \pi)^2 = (-5)^2 + 2 \cdot (-5) \cdot \pi + \pi^2$.
- $((-11) + (-11))^2 = (-11)^2 + 2 \cdot (-11) \cdot (-11) + (-11)^2$.
- $(e + \sqrt{2})^2 = e^2 + 2 \cdot e \cdot \sqrt{2} + (\sqrt{2})^2$.
- ...

Notice that we are allowing the possibility of plugging in the same value for both a and b . We use different letters because they *could* correspond to different values, not because they *must* correspond to different values.

Ok, so how do we prove that the statement

$$\text{“For all } a, b \in \mathbb{R}, \text{ we have } (a + b)^2 = a^2 + 2ab + b^2\text{”}.$$

is true? The problem is that there are infinitely many elements of \mathbb{R} (so infinitely many choices for *each* of a and b), and hence there is no way to examine each possible pair in turn and ever hope to finish. Moreover, there are lots of real numbers that you’ve never thought about before, so it’s hard to even conceive of being able to work with them all.

The way around this obstacle is write a general argument that works regardless of the values for a and b . In other words, we’re going to take two completely *arbitrary* elements of \mathbb{R} that we will name as a and b (so that we can refer to them), and then argue that the result of computing $(a + b)^2$ is the same as the result of computing $a^2 + 2ab + b^2$. When we say “arbitrary”, think about sticking your hand into a bag containing all of the real numbers, and pulling out values for a and b . You are not allowed to pick “nice” values, and you have to work with anything that comes from the bag. In fact, it might help to think of the values of a and b as being picked out of the bag and handed to you by an evil adversary. By taking these arbitrary elements of \mathbb{R} that we call a and b , and then arguing that the value $(a + b)^2$ equals the value $a^2 + 2ab + b^2$, our argument

will work no matter which particular numbers are actually chosen for a and b . In other words, if we are able to write a general argument that works using only the assumption that a and b are real numbers (i.e. using only the fact that they came from the bag), then we can conclude with confidence that each of the infinitely many statements written above are true, and hence the “for all” statement is true.

Now in order to do this, we have to start somewhere. After all, with no assumptions at all about how $+$ and \cdot work, or what squaring means, we have no way to proceed. Ultimately, mathematics starts with basic axioms explaining how certain fundamental mathematical objects and operations work, and builds up everything from there. We won’t go into all of those axioms here, but for the purposes of this discussion we will make use of the following fundamental facts about the real numbers:

- Commutative Law (for multiplication): For all $x, y \in \mathbb{R}$, we have $x \cdot y = y \cdot x$.
- Distributive Law: For all $x, y, z \in \mathbb{R}$, we have $x \cdot (y + z) = x \cdot y + x \cdot z$.

These facts are often taken as two (of about 12) of the axioms for the real numbers. It is also possible to prove them from a construction of the real numbers (see Analysis) using more fundamental axioms. In any event, we can use them to prove the above statement as follows. Let $a, b \in \mathbb{R}$ be arbitrary. We then have that $a + b \in \mathbb{R}$, and

$$\begin{aligned}
 (a + b)^2 &= (a + b) \cdot (a + b) && \text{(by definition)} \\
 &= (a + b) \cdot a + (a + b) \cdot b && \text{(by the Distributive Law)} \\
 &= a \cdot (a + b) + b \cdot (a + b) && \text{(by the Commutative Law)} \\
 &= a \cdot a + a \cdot b + b \cdot a + b \cdot b && \text{(by the Distributive Law)} \\
 &= a^2 + a \cdot b + a \cdot b + b^2 && \text{(by the Commutative Law)} \\
 &= a^2 + 2ab + b^2 && \text{(by definition).}
 \end{aligned}$$

Focus on the logic, and not the algebraic manipulations. We begin by taking arbitrary $a, b \in \mathbb{R}$. Once that sentence is complete, a and b each now represent a specific concrete real number. That is, they are no longer “varying” or serving as placeholders for all real numbers. The act of taking arbitrary values fixes them as concrete numbers, and hence we are now faced with one of mathematical statements in the above infinite list of mathematical statements.

Now let’s turn our attention to the chain of equalities. Read them in consecutive order. We are claiming that $(a + b)^2$ equals $(a + b) \cdot (a + b)$ in the first line. Then the second line says that $(a + b) \cdot (a + b)$ equals $(a + b) \cdot a + (a + b) \cdot b$ by the Distributive Law. Here is the underlying logic behind that equality. Notice that $a + b$ is some particular real number because a and b are particular real numbers (recall that they were instantiated as such when we took arbitrary values at the beginning). Now since $a + b$ is a particular real number (think of it as x), and we can view $a + b$ as the sum of two real numbers (playing the role of y and z , respectively), we can apply the Distributive Law to conclude that $(a + b) \cdot (a + b) = (a + b) \cdot a + (a + b) \cdot b$. The next line is an assertion that the third and fourth expressions are equal by the Commutative Law, etc. If you believe all of the steps, then we have shown that for our completely arbitrary choice of a and b in \mathbb{R} , the first and second expressions are equal, the second and third expressions are equal, the third and fourth expressions are equal, etc. Since equality is transitive (i.e. if $x = y$ and $y = z$, then $x = z$), we conclude that $(a + b)^2 = a^2 + 2ab + b^2$. We have taken completely arbitrary $a, b \in \mathbb{R}$, and verified the statement in question, so we can now assert that the “For all” statement is true.

As a quick aside, now that we know that $(a + b)^2 = a^2 + 2ab + b^2$ for all $a, b \in \mathbb{R}$, we can use this fact whenever we have two real numbers. We can even conclude that the statement

$$\text{“For all } a, b \in \mathbb{R}, \text{ we have } (2a + 3b)^2 = (2a)^2 + 2(2a)(3b) + (3b)^2\text{”}$$

is true. How does this follow? Consider completely arbitrary $a, b \in \mathbb{R}$. We then have that $2a \in \mathbb{R}$ and $3b \in \mathbb{R}$, and thus we can *apply* our previous result to the two numbers $2a$ and $3b$. We are *not* setting “ $a = 2a$ ” or

“ $b = 3b$ ” because it does not make sense to say that $a = 2a$ if a is anything other than 0. We are simply using the fact that if a and b are real numbers, then $2a$ and $3b$ are also real numbers, so we can insert them in for the placeholder values of a and b in our result. Always think of the (arbitrary) choice of letters used in “there exists” and “for all” statements as empty vessels that could be filled with any appropriate value.

We’ve discussed the basic idea behind how to prove that a “there exists” statement or a “for all” statement is true. How do we prove that statements of these forms are false? For example, suppose that we want to show that the statement

$$\text{“There exists } a \in \mathbb{R} \text{ such that } a^2 + 2a = -5\text{”}$$

is false. That is, we want to argue that there is *no* concrete real number that can be plugged into a so that the resulting statement “ $a^2 + 2a = -5$ ” is a true statement. In other words, we want to show that none of the following (infinitely many) mathematical statements are true:

- ...
- $3^2 + 2 \cdot 3 = -5$.
- $0^2 + 2 \cdot 0 = -5$.
- $(-11)^2 + 2 \cdot (-11) = -5$.
- $e^2 + 2 \cdot e = -5$.
- $(\sqrt{2})^2 + 2 \cdot \sqrt{2} = -5$.
- ...

Turning this on its head, we want to show that all of the following (infinitely many) mathematical statements are true:

- ...
- $3^2 + 2 \cdot 3 \neq -5$.
- $0^2 + 2 \cdot 0 \neq -5$.
- $(-11)^2 + 2 \cdot (-11) \neq -5$.
- $e^2 + 2 \cdot e \neq -5$.
- $(\sqrt{2})^2 + 2 \cdot \sqrt{2} \neq -5$.
- ...

Notice how the word *all* found its way into the above reasoning. Showing that *none* of a collection of statements is true is the same as showing that *all* of them are false, which is the same as showing that *all* of the corresponding negations are true! Thinking through the logic, the statement

$$\text{“There exists } a \in \mathbb{R} \text{ such that } a^2 + 2a = -5\text{”}$$

is false exactly when the statement

$$\text{“For all } a \in \mathbb{R}, \text{ we have } a^2 + 2a \neq -5\text{”}$$

is true.

We made use of negations above by noticing that the negation of the statement “ $a^2 + 2a = -5$ ” is the statement “ $a^2 + 2a \neq -5$ ”. But we can use the above ideas to understand the negation of more complicated statements involving quantifiers. The negation of the statement

$$\text{“There exists } a \in \mathbb{R} \text{ such that } a^2 + 2a = -5\text{”}$$

is simply

$$\text{“Not(There exists } a \in \mathbb{R} \text{ such that } a^2 + 2a = -5\text{).”}$$

However, this statement, with the giant **Not** in front, is not easily amenable to analysis because it is neither a “there exists” statement nor a “for all” statement. But as we saw above, this negation is equivalent to the statement

$$\text{“For all } a \in \mathbb{R}, \text{ we have } a^2 + 2a \neq -5\text{”}.$$

In other words, we can move the negation past the “there exists” as long as we change it to a “for all” when doing so. Thus, in order to prove that

$$\text{“There exists } a \in \mathbb{R} \text{ such that } a^2 + 2a = -5\text{”}$$

is false, we can instead prove that

$$\text{“For all } a \in \mathbb{R}, \text{ we have } a^2 + 2a \neq -5\text{”}$$

is true, because it is equivalent to the negation. And we have a strategy for proving that a “for all” statement is true by working with an arbitrary element! Let’s carry out the argument here. Consider an arbitrary $a \in \mathbb{R}$. Notice that

$$\begin{aligned} a^2 + 2a &= (a^2 + 2a + 1) - 1 \\ &= (a + 1)^2 - 1 \\ &\geq 0 - 1 && \text{(because squares of reals are nonnegative)} \\ &= -1. \end{aligned}$$

We have shown that given any arbitrary $a \in \mathbb{R}$, we have $a^2 + 2a \geq -1$, and hence $a^2 + 2a \neq -5$. We conclude that the statement

$$\text{“For all } a \in \mathbb{R}, \text{ we have } a^2 + 2a \neq -5\text{”}$$

is true, and hence the statement

$$\text{“There exists } a \in \mathbb{R} \text{ such that } a^2 + 2a = -5\text{”}$$

is false. Can you see a way to solve this problem using Calculus?

Everything that we said in this example works generally. That is, suppose that we have a statement of the form

$$\text{“There exists } a \text{ such that } \dots\text{”},$$

and we want to argue that the statement is false. We instead prove that the negation

$$\text{“Not(There exists } a \text{ such that } \dots\text{)”}$$

is true. To show that there does not exist an a with a certain property, we need to show that every a *fails* to have that property. Thus, we can instead show that the statement

$$\text{“For all } a, \text{ we have Not}(\dots)\text{”}$$

is true.

Similarly, suppose that we have a statement of the form

“For all a , we have ...”,

and we want to argue that the statement is false. We instead prove that the negation

“**Not**(For all a , we have ...)”

is true. To show that not every a has a certain property, we need to show that there exists some a that *fails* to have that property. Thus, we can instead show that the statement

“There exists a such that **Not**(...)”

is true. In general, we can move a **Not** past one of our two quantifiers at the expense of *flipping* the quantifier to the other type. Although this provides a useful mechanical rule to apply when thinking about negations, it is better to think through the underlying logic each time until the reasoning becomes completely natural.

Life becomes more complicated when a mathematical statement involves both types of quantifiers in an alternating fashion. For example, consider the following two statements:

1. “For all $x \in \mathbb{Z}$, there exists $y \in \mathbb{Z}$ such that $3x + y = 5$ ”.
2. “There exists $y \in \mathbb{Z}$ such that for all $x \in \mathbb{Z}$, we have $3x + y = 5$ ”.

At first glance, these two statements appear to be essentially the same. After all, they both have “for all $x \in \mathbb{Z}$ ”, both have “there exists $y \in \mathbb{Z}$ ”, and both end with the expression “ $3x + y = 5$ ”. Does the fact that these quantifiers appear in different order matter?

Let’s examine the first statement more closely. Notice that it has the form “For all $x \in \mathbb{Z}$...”. In order for this “for all” statement to be true, we want to know whether we obtain a true statement *whenever* we plug in a particular integer x in the “...” part. In other words, we’re asking if *all* of the following (infinitely many) mathematical statements are true:

- ...
- “There exists $y \in \mathbb{Z}$ such that $3 \cdot (-2) + y = 5$ ”.
- “There exists $y \in \mathbb{Z}$ such that $3 \cdot (-1) + y = 5$ ”.
- “There exists $y \in \mathbb{Z}$ such that $3 \cdot 0 + y = 5$ ”.
- “There exists $y \in \mathbb{Z}$ such that $3 \cdot 1 + y = 5$ ”.
- “There exists $y \in \mathbb{Z}$ such that $3 \cdot 2 + y = 5$ ”.
- ...

Looking through each of these, it does indeed appear that they are all true: We can use $y = 11$ in the first displayed one (i.e. when $x = -2$), then $y = 8$ in the next, then $y = 5$ in the next one, etc. However, there are infinitely many statements, so we can’t actually check each one in turn and hope to finish. We need a general argument that works no matter which value x takes. Now given any *arbitrary* $x \in \mathbb{Z}$, we can verify if consider the value of y to be $5 - 3x$, then we obtain a true statement. Here is how we would write this argument up formally.

Proposition 1.3.1. *For all $x \in \mathbb{Z}$, there exists $y \in \mathbb{Z}$ such that $3x + y = 5$.*

Proof. Let $x \in \mathbb{Z}$ be arbitrary. Since \mathbb{Z} is closed under multiplication and subtraction, we know that $5 - 3x \in \mathbb{Z}$. Now

$$\begin{aligned} 3x + (5 - 3x) &= (3x - 3x) + 5 \\ &= 0 + 5 \\ &= 5. \end{aligned}$$

Thus, we have shown the existence of a $y \in \mathbb{Z}$ with $3x + y = 5$ (namely $y = 5 - 3x$). \square

Let's pause to note a few things about this argument. First, we've labeled the statement as a proposition. By doing so, we are making a claim that the statement to follow is a true statement, and that we will be providing a proof. Alternatively, we sometimes will label a statement as a *theorem* instead of a proposition if we want to elevate it to a position of prominence (typically theorems say something powerful, surprising, or incredibly useful). In the proof, we are trying to argue that a "for all" statement is true, so we start by taking an *arbitrary* element of \mathbb{Z} . Although this x is arbitrary, it is *not* varying. Instead, once we take an arbitrary x , it is now one fixed concrete integer that we can use in the rest of the argument. For this particular but arbitrary $x \in \mathbb{Z}$, we now want to argue that a certain "there exists" statement is true. In order to do this, we need to exhibit an example of a y that works, and verify it for the reader. Since we have a particular $x \in \mathbb{Z}$ in hand, the y that we pick can depend on that x . In this case, we simply verify that $5 - 3x$ works as a value for y . As in the examples given above, we do not need to explain why we chose to use $5 - 3x$, only that the resulting statement is true. That is, although you might have done some algebra to figure out what y might work, the process you used to find such a potential y is irrelevant to the logical demonstration that such a y exists. Finally, the square box at the end of the argument indicates that the proof is over, and so the next paragraph (i.e. this one) is outside the scope of the argument.

Let's move on to the second of our two statements above. Notice that it has the form "There exists $y \in \mathbb{Z} \dots$ ". In order for this statement to be true, we want to know whether we can find *one* value for y such that we obtain a true statement in the " \dots " part after plugging it in. In other words, we're asking if *any* of the following (infinitely many) mathematical statements are true:

- \dots
- "For all $x \in \mathbb{Z}$, we have $3x + (-2) = 5$ ".
- "For all $x \in \mathbb{Z}$, we have $3x + (-1) = 5$ ".
- "For all $x \in \mathbb{Z}$, we have $3x + 0 = 5$ ".
- "For all $x \in \mathbb{Z}$, we have $3x + 1 = 5$ ".
- "For all $x \in \mathbb{Z}$, we have $3x + 2 = 5$ ".
- \dots

Looking through each of these, it appears that every single one of them is false, i.e. *none* of them are true. Thus, it appears that the second statement is false. We can formally prove that it is false by proving that its negation is true. Applying our established rules for how to negate across quantifiers, to show that

"Not(There exists $y \in \mathbb{Z}$ such that for all $x \in \mathbb{Z}$, we have $3x + y = 5$)"

is true, we can instead show that

"For all $y \in \mathbb{Z}$, Not(for all $x \in \mathbb{Z}$, we have $3x + y = 5$)"

is true, which is same as showing that

“For all $y \in \mathbb{Z}$, there exists $x \in \mathbb{Z}$ such that **Not**($3x + y = 5$)”

is true, which is the same as showing that

“For all $y \in \mathbb{Z}$, there exists $x \in \mathbb{Z}$ such that $3x + y \neq 5$ ”.

is true. We now prove that this final statement is true, which is the same as showing that our original second statement is false.

Proposition 1.3.2. *For all $y \in \mathbb{Z}$, there exists $x \in \mathbb{Z}$ such that $3x + y \neq 5$.*

Proof. Let $y \in \mathbb{Z}$ be arbitrary. We have two cases:

- *Case 1:* Suppose that $y \neq 5$. We then have that $0 \in \mathbb{Z}$ and

$$\begin{aligned} 3 \cdot 0 + y &= 0 + y \\ &= y. \end{aligned}$$

Since $y \neq 5$, we have shown the existence of an $x \in \mathbb{Z}$ with $3x + y \neq 5$ (namely $x = 0$).

- *Case 2:* Suppose that $y = 5$. We then have that $1 \in \mathbb{Z}$ and

$$\begin{aligned} 3 \cdot 1 + y &= 3 + 5 \\ &= 8. \end{aligned}$$

Since $8 \neq 5$, we have shown the existence of an $x \in \mathbb{Z}$ with $3x + y \neq 5$ (namely $x = 1$).

As these two cases exhaust all possibilities for y , we have shown that such an x exists unconditionally. \square

Again, let’s pause to examine the structure of this proof. Since we are trying to prove a “for all” statement, we start by taking an arbitrary element of \mathbb{Z} . Once we have this y in hand, our task is to prove a “there exists” statement. Intuitively, once we have a specific y in hand, i.e. once y is given a concrete value, it appears that almost any value of x will work. In fact, it seems that at most one value of x might cause a problem. But our task is to prove a “there exists” statement, so we need to provide a value of x and prove that it works. In the above argument, we chose to respond to the given y with the value of $x = 0$, so long as the given value of y is not 5. Why did we make this case distinction and use this particular value of x ? From the logical perspective of the argument, our motivations do not matter. Of course, it is a good exercise to consider *why* the case distinction was made, and *what* motivated the choice of x . That’s not to say that these are the only choices we could make! It is certainly possible to make a different case distinction and/or choose different values of x in response to specific y ’s. Try to write a different argument yourself!

In general, consider statements of the following two types:

1. “For all a , there exists b such that ...”.
2. “There exists b such that for all a , we have ...”.

Let’s examine the difference between them in a more informal way. Think about a game with two players where Player I goes first. For the first statement to be true, it needs to be the case that no matter how Player I moves, Player II can respond in such a way so that ... happens. Notice in this scenario Player II’s move can depend on Player I’s move, i.e. the value of b can depend on the value of a . For the second statement to be true, it needs to be the case that Player I can make a move so brilliant that no matter how Player II responds, we have that ... happens. In this scenario, b needs to be chosen *first* without knowing

a , so b can not depend on a in any way.

Finally, let's discuss one last construct in mathematical statements, which is an "if...then..." clause. We call such statements *implications*, and they naturally arise when we want to quantify only over part of a set. For example, the statement

$$\text{"For all } a \in \mathbb{R}, \text{ we have } a^2 - 4 \geq 0\text{"}$$

is false because $0 \in \mathbb{R}$ and $0^2 - 4 < 0$. However, the statement

$$\text{"For all } a \in \mathbb{R} \text{ with } a \geq 2, \text{ we have } a^2 - 4 \geq 0\text{"}$$

is true. Instead of coupling the condition " $a \geq 2$ " with the "for all" statement, we can instead write this statement as

$$\text{"For all } a \in \mathbb{R}, (\text{If } a \geq 2, \text{ then } a^2 - 4 \geq 0)\text{"}.$$

We often write this statement in shorthand by dropping the "for all" as:

$$\text{"If } a \in \mathbb{R} \text{ and } a \geq 2, \text{ then } a^2 - 4 \geq 0\text{"}.$$

One convention, that initially seems quite strange, arises from this. Since we want to allow "if...then..." statements, we need to assign truth values to them because every mathematical statement should either be true or false. If we plug the value 3 for a into this last statement (or really past the "for all" in the penultimate statement), we arrive at the statement

$$\text{"If } 3 \geq 2, \text{ then } 3^2 - 4 \geq 0\text{"},$$

which we naturally say is true because both the "if" part and the "then" part are true. However, it's less clear how we should assign a truth value to

$$\text{"If } 1 \geq 2, \text{ then } 1^2 - 4 \geq 0\text{"}$$

because both the "if" part and the "then" part are false. We also have an example like

$$\text{"If } -5 \geq 2, \text{ then } (-5)^2 - 4 \geq 0\text{"},$$

where the "if" part is false and the "then" part is true. In mathematics, we make the convention that an "if...then..." statement is false only when the "if" part is true and the "then" part is false. Thus, these last two examples we declare to be true. The reason why we do this is to be consistent with the intent of the "for all" quantifier. In the example

$$\text{"For all } a \in \mathbb{R}, (\text{If } a \geq 2, \text{ then } a^2 - 4 \geq 0)\text{"},$$

we do not want values of a with $a < 2$ to have any effect at all on the truth value of the "for all" statement. Thus, we want the parenthetical statement to be true whenever the "if" part is false. In general, given two mathematical statements P and Q , we *define* the following:

- If P is true and Q is true, we say that "If P , then Q " is true.
- If P is true and Q is false, we say that "If P , then Q " is false.
- If P is false and Q is true, we say that "If P , then Q " is true.
- If P is false and Q is false, we say that "If P , then Q " is true.

We can compactly illustrate these conventions with the following simple table, known as a *truth table*, where we use T for true and F for false:

P	Q	If P , then Q
T	T	T
T	F	F
F	T	T
F	F	T

Compare these with the simple truth tables that arise from the word *and* and the word *or* (remembering that *or* is always the inclusive or in mathematics, unless stated otherwise):

P	Q	P and Q
T	T	T
T	F	F
F	T	F
F	F	F

P	Q	P or Q
T	T	T
T	F	T
F	T	T
F	F	F

1.4 Evens and Odds

We will spend this section discussing even and odd integers, and culminate with a proof that $\sqrt{2}$ is irrational. As we've discussed, all mathematics ultimately relies upon a few core concepts and axioms. Thus, whenever we introduce a new word like *even* or *odd*, we need to define it in terms of more basic concepts. We accomplish this using our “there exists” quantifier.

Definition 1.4.1. Let $a \in \mathbb{Z}$.

- We say that a is even if there exists $m \in \mathbb{Z}$ with $a = 2m$.
- We say that a is odd if there exists $m \in \mathbb{Z}$ with $a = 2m + 1$.

Since this is our first formal definition, let's pause for a moment to understand the role of definitions in mathematics. First, in contrast to our “if...then...” statements, the word “if” when used alone in a definition is really shorthand for “to mean that”. Now a mathematical definition tells us *exactly* what we mean by the words or notation that we introduce. There is no more subtlety to add. Every time we use the word “even”, we are really just using it so that we do not have to say “there exists $m \in \mathbb{Z}$ with $a = 2m$ ”. In other words, everything about an integer being “even” should *always* eventually go back to the definition.

We can use this definition to now assert that certain integers are even or odd. For example, we can assert that 10 is even because $10 = 2 \cdot 5$ and $5 \in \mathbb{Z}$. We can also see that 71 is odd because we can write $71 = 2 \cdot 35 + 1$ and $35 \in \mathbb{Z}$. Also notice that 0 is even by our definition because $0 = 2 \cdot 0$ and $0 \in \mathbb{Z}$.

Now you might have thought to define the word *even* in a different way. For example, you could consider defining a to be even if the remainder when dividing a by 2 is 0. This is certainly a natural approach, and for many people that is how it was explained to them when they were young. However, since mathematical terms should be precisely defined down to our ultimately basic concepts, such a definition would require us to work through what we mean by “division” and “remainder” for integers. Although it is certainly possible to do this, our official definition introduces no new concepts and is easier to work with. Eventually, if we were to formally define “division” and “remainder” (like you might do in Elementary Number Theory or Abstract Algebra), then you can *prove* that our official definition means the same thing as the one obtained by such an approach. In general, however, there is no strict rule for choosing which definition to use when several competing alternatives are available. Ideally, we settle in on a definition that is simple, useful, and elegant. In mathematical subjects that have been developed over the course of several centuries, mathematicians have settled on the “right” core definitions over time, but in newer areas finding the “right” definitions is often an important step.

We now prove our first result. We'll write it formally, and then discuss its structure after the proof.

Proposition 1.4.2. *If $a \in \mathbb{Z}$ is even, then a^2 is even.*

Proof. Let $a \in \mathbb{Z}$ be an arbitrary even integer. Since a is even, we can fix $n \in \mathbb{Z}$ with $a = 2n$. Notice that

$$\begin{aligned} a^2 &= (2n)^2 \\ &= 4n^2 \\ &= 2 \cdot (2n^2). \end{aligned}$$

Since $2n^2 \in \mathbb{Z}$, we conclude that a^2 is even.

When starting this proof, we have to remember that there is a hidden “for all” in the “if...then...”. That is, the statement that we are trying to prove is:

“For all $a \in \mathbb{Z}$, if a is even, then a^2 is even”.

Thus, we should start the argument by taking an arbitrary $a \in \mathbb{Z}$. However, we're trying to prove an "if...then..." statement about such an a . Whenever the "if..." part is false, we do not care about it (or alternatively we assign it true by the discussion at the end of the previous section), so instead of taking an arbitrary $a \in \mathbb{Z}$, we should take an arbitrary $a \in \mathbb{Z}$ *that is even*. With this even a in hand, our goal is to prove that a^2 is even.

Recall that whenever we think about even numbers now, we should always eventually go back to our definition. Thus, we next unwrap what it means to say that “ a is even”. By definition of even, we know that there exists $m \in \mathbb{Z}$ with $a = 2m$. In other words, there is at least one choice of $m \in \mathbb{Z}$ so that the statement “ $a = 2m$ ” is true.

Now it's conceivable that there are many m that work (the definition does not rule that out), but there is at least one that works. We invoke this true statement by *picking* some value of $m \in \mathbb{Z}$ that works, and we do this by giving it a name n . This was an arbitrary choice of name, and we could have chosen almost any other name for it. We could have called it k , b , ℓ , x , δ , Maschke, \heartsuit , or $\$$. The only really awful choice would be to call it a , because we have already given the letter a a meaning (namely as our arbitrary element). We could even have called it m , and in the future we will likely do this. However, to avoid confusion in our first arguments, we've chosen to use a different letter than the one in the definition to make it clear that we are now fixing one value that works. We encapsulate this entire paragraph in the key phrase “*we can fix*”. In general, when we want to invoke a true “there exists” statement in our argument, we use the phrase *we can fix* to pick a corresponding witness.

Ok, we've successfully taken our assumption and unwrapped it, so that we now have a fixed $n \in \mathbb{Z}$ with $a = 2n$. Before jumping into the algebra of the middle part of the argument, let's think about our goal. We want to show that a^2 is even. In other words, we want to argue that there exists $m \in \mathbb{Z}$ with $a^2 = 2m$. Don't think about the letter. We want to end by writing $a^2 = 2__$ where whatever we fill in for $__$ is an integer.

With this in mind, we start with what we know is true, i.e. that $a = 2n$, and hope to drive forward with true statements every step of the way until we arrive at our goal. Since $a = 2n$ is true, we know that $a^2 = (2n)^2$ is true. We also know that $(2n)^2 = 4n^2$ is true and that $4n^2 = 2 \cdot (2n^2)$ is true. Putting it all together, we conclude that $a^2 = 2 \cdot (2n^2)$ is true. Have we arrived at our goal? We've written a^2 as 2 times something, namely it is 2 times $2n^2$. Finally, we notice that $2n^2 \in \mathbb{Z}$ because $n \in \mathbb{Z}$. Thus, starting with the true statement $a = 2n$, we have derived a sequence of true statements culminating with the true statement that a^2 equals 2 times some integer. Therefore, by definition, we are able to conclude that a^2 is even. Since a was arbitrary, we are done.

Pause to make sure that you understand all of the logic in the above argument. Mathematical proofs are typically written in very concise ways where each word matters. Furthermore, these words often pack in complex thoughts, such as with the “we can fix” phrase above. Eventually, we will just write our arguments succinctly without all of this commentary, and it’s important to make sure that you understand how to unpack both the language and the logic used in proofs.

In fact, we can prove a stronger result than what is stated in the proposition. It turns out that if $a \in \mathbb{Z}$ is even and $b \in \mathbb{Z}$ is arbitrary, then ab is even (i.e. the product of an even integer and *any* integer is an even integer). Try to give a proof! From this fact, we can immediately conclude that the previous proposition is true, because given any $a \in \mathbb{Z}$ that is even, we can apply this stronger result when using a for both of the values (i.e. for both a and b). Remember that the letters are placeholders, so we can fill them both with the same value if we want. Different letters do not necessarily mean different values!

Let's move on to another argument that uses the definitions of both even and odd.

Proposition 1.4.3. *If $a \in \mathbb{Z}$ is even and $b \in \mathbb{Z}$ is odd, then $a + b$ is odd.*

Proof. Let $a, b \in \mathbb{Z}$ be arbitrary with a even and b odd. Since a is even, we can fix $n \in \mathbb{Z}$ with $a = 2n$. Since b is odd, we can fix $k \in \mathbb{Z}$ with $b = 2k + 1$. Notice that

$$\begin{aligned} a + b &= 2n + (2k + 1) \\ &= (2n + 2k) + 1 \\ &= 2 \cdot (n + k) + 1. \end{aligned}$$

Now $n + k \in \mathbb{Z}$ because both $n \in \mathbb{Z}$ and $k \in \mathbb{Z}$, so we can conclude that $a + b$ is odd. \square

This argument is similar to the last one, but now we have two arbitrary elements $a, b \in \mathbb{Z}$, with the additional assumption that a is even and b is odd. As in the previous proof, we unwrapped the definitions involving “there exists” quantifiers to fix witnessing elements n and k . Notice that we had to give these witnessing elements different names because the n that we pick to satisfy $a = 2n$ might be a completely different number from the k that we pick to satisfy $b = 2k + 1$. Once we've unwrapped those definitions, our goal is to prove that $a + b$ is odd, which means that we want to show that $a + b = 2__ + 1$, where we fill in $__$ with an integer. Now using algebra we proceed forward from our given information to conclude that $a + b = 2 \cdot (n + k) + 1$, so since $n + k \in \mathbb{Z}$ (because both $n \in \mathbb{Z}$ and $k \in \mathbb{Z}$), we have reached our goal.

We now ask a seemingly simple question: Is 1 even? We might notice that 1 is odd because $1 = 2 \cdot 0 + 1$ and $0 \in \mathbb{Z}$, but how does that help us? At the moment, we only have our definitions, and it is not immediately obvious from the definitions that a number can not be both even and odd. To prove that 1 is not even, we have to argue that

$$\text{“There exists } m \in \mathbb{Z} \text{ with } 1 = 2m\text{”}$$

is false, which is the same as showing that

$$\text{“Not(There exists } m \in \mathbb{Z} \text{ with } 1 = 2m\text{)”}$$

is true, which is the same as showing that

$$\text{“For all } m \in \mathbb{Z}, \text{ we have } 1 \neq 2m\text{”}$$

is true. Thus, we need to prove a “for all” statement, which we do by considering two cases.

Proposition 1.4.4. *The integer 1 is not even.*

Proof. We show that $2m \neq 1$ for all $m \in \mathbb{Z}$. Let $m \in \mathbb{Z}$ be arbitrary. We then have that either $m \leq 0$ or $m \geq 1$, giving us two cases:

- *Case 1:* Suppose that $m \leq 0$. Multiplying both sides by 2, we see that $2m \leq 0$, so $2m \neq 1$.
- *Case 2:* Suppose that $m \geq 1$. Multiplying both sides by 2, we see that $2m \geq 2$, so $2m \neq 1$.

Since these two cases exhaust all possibilities for m , we have shown that $2m \neq 1$ unconditionally. \square

This is a perfectly valid argument, but it's very specific to the number 1. We would like to prove the far more general result that no integer is both even and odd. We'll introduce a new method of proof to accomplish this task. Up until this point, if we have a statement P that we want to prove is true, we tackle the problem directly by working through the quantifiers one by one. Similarly, if we want to prove that P is false, we instead prove that $\text{Not}(P)$ is true, and use our rules for moving the **Not** inside so that we can prove a statement involving quantifiers on the outside directly.

However, there is another method to prove that a statement P is true that is beautifully sneaky. The idea is as follows. We *assume* that $\text{Not}(P)$ is true, and show that under this assumption we can logically derive another statement, say Q , that we *know* to be false. Thus, *if* $\text{Not}(P)$ was true, *then* Q would have to be both true and false at the same time. Madness would ensue. **Human sacrifice, dogs and cats living together, mass hysteria.** This is inconceivable, so the only possible explanation is that $\text{Not}(P)$ must be false, which is the same as saying that P must be true. A proof of this type is called a *proof by contradiction*, because under the assumption that $\text{Not}(P)$ was true, we derived a contradiction, and hence we can conclude that P must be true.

Proposition 1.4.5. *No integer is both even and odd.*

Before jumping into the proof, let's examine what the proposition is saying formally. If we write it out carefully, the claim is that

“ $\text{Not}(\text{There exists } a \in \mathbb{Z} \text{ such that } a \text{ is even and } a \text{ is odd})$ ”

is true. If we were trying to prove this statement directly, we would move the **Not** inside and instead try to prove that

“For all $a \in \mathbb{Z}$, we have $\text{Not}(a \text{ is even and } a \text{ is odd})$ ”

is true, which is the same as showing that

“For all $a \in \mathbb{Z}$, either a is not even or a is not odd”

is true (recalling that “or” is always the inclusive or in mathematics). To prove this directly, we would then need to take an arbitrary $a \in \mathbb{Z}$, and argue that (at least one) of “ a is not even” or “ a is not odd” is true. Since this looks a bit difficult, let's think about how we would structure a proof by contradiction. Recall that we are trying to prove that

“ $\text{Not}(\text{There exists } a \in \mathbb{Z} \text{ such that } a \text{ is even and } a \text{ is odd})$ ”

is true. Now instead of moving the **Not** inside and proving the corresponding “for all” statement directly, we are going to do a proof by contradiction. Thus, we *assume* that

“ $\text{Not}(\text{Not}(\text{There exists } a \in \mathbb{Z} \text{ such that } a \text{ is even and } a \text{ is odd}))$ ”,

is true, which is the same as assuming that

“There exists $a \in \mathbb{Z}$ such that a is even and a is odd”,

is true, and then derive a contradiction. Let's do it.

Proof of Proposition 1.4.5. Assume, for the sake of obtaining a contradiction, that there exists an integer that is both even and odd. We can then fix an $a \in \mathbb{Z}$ that is both even and odd. Since a is even, we can fix $m \in \mathbb{Z}$ with $a = 2m$. Since a is odd, we can fix $n \in \mathbb{Z}$ with $a = 2n + 1$. We then have $2m = 2n + 1$, so $2(m - n) = 1$. Notice that $m - n \in \mathbb{Z}$ because both $m \in \mathbb{Z}$ and $n \in \mathbb{Z}$. Thus, we can conclude that 1 is even, which contradicts Proposition 1.4.4. Therefore, our assumption must be false, and hence no integer can be both even and odd. \square

Ok, so no integer can be both even and odd. Is it true that every integer is either even or odd? Intuitively, the answer is clearly yes, but it's not obvious how to prove it without developing a theory of division with remainder. It is possible to accomplish this task by carrying out the steps in the following outline:

1. Start with 0. We know from above that 0 is even, so clearly 0 is either even or odd.
2. Show that if $a \in \mathbb{N}$ is either even or odd, then $a + 1$ is either even or odd.
3. Intuitively, every natural number can be obtained by starting with 0, and then iteratively adding 1 some finite number of times. Thus, using (1) and (2) repeatedly, it seems that we can conclude that every $a \in \mathbb{N}$ is either even or odd. The way to formalize this argument is to use a technique called “mathematical induction”. However, since a careful discussion of this technique will take us too far afield, we will leave the details of the argument to a later course (see Bridges to Advanced Mathematics).
4. Show that if $a \in \mathbb{N}$ is either even or odd, then $-a$ is either even or odd. This then allows us to extend the statement in (3) from \mathbb{N} to \mathbb{Z} .

Since we are leaving the details of part (3) to a later course, we will simply assert that the following is true, and you'll have to suffer through the anticipation for a semester.

Fact 1.4.6. *Every integer is either even or odd.*

We turn now to another result, which will require a new technique to prove.

Proposition 1.4.7. *If $a \in \mathbb{Z}$ and a^2 is even, then a is even.*

Before jumping into a proof of this fact, we pause to notice that a direct approach looks infeasible. Why? Suppose that we try to prove it directly by starting with the assumption that a^2 is even. Then, by definition, we can fix $n \in \mathbb{Z}$ with $a^2 = 2n$. Since $a^2 \geq 0$, we conclude that we must have that $n \geq 0$. It is now natural to take the square root of both sides and write $a = \sqrt{2n}$. Recall that our goal is to write a as 2 times some integer, but this looks bad. We have $a = \sqrt{2} \cdot \sqrt{n}$, but $\sqrt{2}$ is not 2, and \sqrt{n} is probably not an integer. We can force a 2 by noticing that $\sqrt{2n} = 2 \cdot \sqrt{\frac{n}{2}}$, but $\sqrt{\frac{n}{2}}$ seems even less likely to be an integer.

Let's take a step back. Notice that Proposition 1.4.7 looks an awful lot like Proposition 1.4.2. In fact, one is of the form “If P, then Q” while the other is of the form “If Q, then P”. We say that “If Q, then P” is the *converse* of “If P, then Q”. Unfortunately, if an “If...then...” statement is true, its converse might be false. For example,

“If $f(x)$ is differentiable, then $f(x)$ is continuous”

is true, but

“If $f(x)$ is continuous, then $f(x)$ is differentiable”.

is false. For an even more basic example, the statement

“If $a \in \mathbb{Z}$ and $a \geq 7$, then $a \geq 4$ ”

is true, but the converse statement

“If $a \in \mathbb{Z}$ and $a \geq 4$, then $a \geq 7$ ”

is false.

The reason why Proposition 1.4.2 was easier to prove was that we started with the assumption that a was even, and by squaring both sides of $a = 2n$ we were able to write a^2 as 2 times an integer by using the fact that the square of an integer was an integer. However, starting with an assumption about a^2 , it seems difficult to conclude much about a without taking square roots. Here's where a truly clever idea comes in.

Instead of looking at the converse of our statement, which says “If $a \in \mathbb{Z}$ and a is even, then a^2 is even”, consider the following statement:

“If $a \in \mathbb{Z}$ and a is not even, then a^2 is not even”

Now this statement is a strange twist on the first. We’ve switched the hypothesis and conclusion around and included negations that were not there before. At first sight, it may appear that this statement has nothing to do with the one in Proposition 1.4.7. However, suppose that we are somehow able to prove it. I claim that Proposition 1.4.7 follows. How? Suppose that $a \in \mathbb{Z}$ is such that a^2 is even. We want to argue that a must be even. Well, suppose not. Then a is not even, so by this new statement (which we are assuming we know is true), we could conclude that a^2 is not even. However, this contradicts our assumption. Therefore, it must be the case that a is even!

We want to give this general technique a name. The *contrapositive* of a statement of the form “If P , then Q ” is the statement “If **Not**(Q), then **Not**(P)”. In other words, we flip the two parts of the “If...then...” statement and put a **Not** on both of them. In general, suppose that we are successful in proving that the contrapositive statement

“If **Not**(Q), then **Not**(P)”

is true. From this, it turns out that we can conclude that

“If P , then Q ”

is true. Let’s walk through the steps. Remember, we are assuming that we know that “If **Not**(Q), then **Not**(P)” is true. To prove that “If P , then Q ” is true, we assume that P is true, and have as our goal to show that Q is true. Now under the assumption that Q is false, we would be able to conclude that **Not**(Q) is true, but this would imply that **Not**(P) is true, contradicting the fact that we are assuming that P is true! The only logical possibility is that the truth of P must imply the truth of Q .

We are now ready to prove Proposition 1.4.7.

Proof of 1.4.7. We prove the contrapositive. That is, we show that whenever a is not even, then a^2 is not even. Let $a \in \mathbb{Z}$ be an arbitrary integer that is not even. Using Fact 1.4.6, it follows that a is odd. Thus, we can fix $n \in \mathbb{Z}$ with $a = 2n + 1$. We then have

$$\begin{aligned} a^2 &= (2n + 1)^2 \\ &= 4n^2 + 4n + 1 \\ &= 2 \cdot (2n^2 + 2n) + 1. \end{aligned}$$

Notice that $2n^2 + 2n \in \mathbb{Z}$ because $n \in \mathbb{Z}$, so we can conclude that a^2 is odd. Using Proposition 1.4.5, it follows that a^2 is not even. We have shown that if a is not even, then a^2 is not even. Since we’ve proven the contrapositive, it follows that if a^2 is even, then a is even. \square

We can now prove the following fundamental theorem.

Theorem 1.4.8. *There does not exist $q \in \mathbb{Q}$ with $q^2 = 2$. In other words, $\sqrt{2}$ is irrational.*

Proof. Suppose for the sake of obtaining a contradiction that there does exist $q \in \mathbb{Q}$ with $q^2 = 2$. Fix $a, b \in \mathbb{Z}$ with $q = \frac{a}{b}$ and such that $\frac{a}{b}$ is in lowest terms, i.e. where a and b have no common factors greater than 1. We have

$$\left(\frac{a}{b}\right)^2 = 2,$$

hence

$$\frac{a^2}{b^2} = 2,$$

and so

$$a^2 = 2 \cdot b^2.$$

Since $b^2 \in \mathbb{Z}$, we conclude that a^2 is even. Using Proposition 1.4.7, it follows that a is even, so we can fix $c \in \mathbb{Z}$ with $a = 2c$. We then have

$$\begin{aligned} 2b^2 &= a^2 \\ &= (2c)^2 \\ &= 4c^2. \end{aligned}$$

Dividing each side by 2, we conclude that

$$b^2 = 2c^2.$$

Since $c^2 \in \mathbb{Z}$, it follows that b^2 is even. Using Proposition 1.4.7 again, we conclude that b is even. Thus, we can fix $d \in \mathbb{Z}$ with $b = 2d$. We then have

$$q = \frac{a}{b} = \frac{2c}{2d} = \frac{c}{d}.$$

This is a contradiction because $\frac{a}{b}$ was assumed to be in lowest terms, but we have reduced it further. Therefore, there does not exist $q \in \mathbb{Q}$ with $q^2 = 2$. \square

We end this section with an interesting fact, which gives another example of proving a “there exists” statement.

Proposition 1.4.9. *If $a \in \mathbb{Z}$ is odd, then there exist $b, c \in \mathbb{Z}$ with $a = b^2 - c^2$. In other words, every odd integer is the difference of two perfect squares.*

Before jumping into the proof, we first try out some examples:

$$\begin{aligned} 1 &= 1^2 - 0^2 \\ 3 &= 2^2 - 1^2 \\ 5 &= 3^2 - 2^2 \\ 7 &= 4^2 - 3^2 \\ 9 &= 5^2 - 4^2. \end{aligned}$$

There seems to be a clear pattern here, but notice that $9 = 3^2 - 0^2$ as well, so there are sometimes other ways to write numbers as the difference of squares. Nonetheless, to prove a “there exists” statement, we just need to give an example. The above pattern suggests that we can just use adjacent numbers for b and c , and we now give a proof of this fact generally.

Proof. Let $a \in \mathbb{Z}$ be an arbitrary odd integer. By definition, we can fix $n \in \mathbb{Z}$ with $a = 2n + 1$. Notice that $n + 1 \in \mathbb{Z}$ and that

$$\begin{aligned} (n + 1)^2 - n^2 &= n^2 + 2n + 1 - n^2 \\ &= 2n + 1 \\ &= a. \end{aligned}$$

Therefore, we have shown the existence of b and c (namely $b = n + 1$ and $c = n$) for which $a = b^2 - c^2$. Since $a \in \mathbb{Z}$ was an arbitrary odd integer, we conclude that every odd integer is the difference of two perfect squares. \square

1.5 Sets, Set Construction, and Subsets

Sets and Set Construction

We now discuss the fundamental structure that mathematicians use to package objects together.

Definition 1.5.1. *A set is a collection of elements without regard to repetition or order.*

Intuitively, a set is a box where the only thing that matters are the objects that are inside it, and furthermore the box does not have more than one of any given object. We use $\{$ and $\}$ as delimiters for sets. For example, $\{3, 5\}$ is a set with two elements. Since all that matters are the elements, we define two sets to be equal if they have the same elements, regardless of how the sets themselves are defined or described.

Definition 1.5.2. *Given two sets A and B , we say that $A = B$ if A and B have exactly the same elements.*

Since only the actual elements matter, not their order, we have $\{3, 7\} = \{7, 3\}$ and $\{1, 2, 3\} = \{3, 1, 2\}$. Also, although we typically would not even write something like $\{2, 5, 5\}$, if we choose to do so then we would have $\{2, 5, 5\} = \{2, 5\}$ because both have the same elements, namely 2 and 5.

Notation 1.5.3. *Given an object x and a set A , we write $x \in A$ to mean that x is an element of A , and we write $x \notin A$ to mean that x is not an element of A .*

For example, we have $2 \in \{2, 5\}$ and $3 \notin \{2, 5\}$. Since sets are mathematical objects, they may be elements of other sets. For example, we can form the set $S = \{1, \{2, 3\}\}$. Notice that we have $1 \in S$ and $\{2, 3\} \in S$, but $2 \notin S$ and $3 \notin S$. As a result, S has only 2 elements, namely 1 and $\{2, 3\}$. Thinking of a set as a box, one element of S is the number 1, and the other is a different box (which happens to have the two elements 2 and 3 inside it).

The empty set is the unique set with no elements. We can write it as $\{\}$, but instead we typically denote it by \emptyset . There is only *one* empty set, because if both A and B have no elements, then they have exactly the same elements for vacuous reasons, and hence $A = B$. Notice that $\{\emptyset\}$ does not equal \emptyset . After all, $\{\emptyset\}$ has one element! You can think of $\{\emptyset\}$ as a box that has one empty box inside it.

Notice that sets can be either finite or infinite. At this point, our standard examples of infinite sets are the various universes of numbers:

- $\mathbb{N} = \{0, 1, 2, 3, \dots\}$.
- $\mathbb{N}^+ = \{1, 2, 3, \dots\}$.
- $\mathbb{Z} = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$.
- \mathbb{Q} is the set of rational numbers.
- \mathbb{R} is the set of real numbers.

Beyond these fundamental sets, there are various ways to define new sets. In some cases, we can simply list the elements as we did above. Although this often works for small finite sets, it is almost never a good idea to list the elements of a set with 20 or more elements, and it rarely works for infinite sets (unless there is an obvious pattern like $\{5, 10, 15, 20, \dots\}$). One of the standard ways to define a set S is to carve it out of some bigger set A by describing a certain property that may or may not be satisfied by an element of A . For example, we could define

$$S = \{n \in \mathbb{N} : 5 < n < 13\}.$$

We read this line by saying that S is defined to be the set of all $n \in \mathbb{N}$ such that $5 < n < 13$. Thus, in this case, we are taking $A = \mathbb{N}$, and forming a set S by carving out those elements of A that satisfy the condition that $5 < n < 13$. In other words, think about going through each of element n , checking if $5 < n < 13$ is a

true statement, and collecting those $n \in \mathbb{N}$ that make it true into a set that we call S . In more simple terms, we can also describe S as follows:

$$S = \{6, 7, 8, 9, 10, 11, 12\}.$$

It is important that we put the “ \mathbb{N} ” in the above description, because if we wrote $\{n : 5 < n < 13\}$ then it would be unclear what n we should consider. For example, should $\frac{11}{2}$ be in this set? How about $\sqrt{17}$? Sometimes the “universe” of numbers (or other mathematical objects) that we are working within is clear, but typically it is best to write the global set that we are picking elements from in order to avoid such ambiguity. Notice that when we define a set, there is no guarantee that it has any elements. For example, $\{q \in \mathbb{N} : q^2 = 2\} = \emptyset$ by Theorem 1.4.8. Keep in mind that we can also use words in our description of sets, such as $\{n \in \mathbb{N} : n \text{ is an even prime}\}$. As mentioned above, two sets that have quite different descriptions can be equal. For example, we have

$$\{n \in \mathbb{N} : n \text{ is an even prime}\} = \{n \in \mathbb{N} : 3 < n^2 < 8\}$$

because both sets equal $\{2\}$. Always remember the structure of sets formed in this way. We write

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

where A is a known set and $P(x)$ is a “property” such that given a particular $y \in A$, the statement $P(y)$ is either true or false.

Another way to describe a set is through a “parametric” description. Rather than carving out a certain subset of a given set by describing a property that the elements must satisfy, we can instead form all the elements one obtains by varying a value through a particular set. For example, consider the following description of a set:

$$S = \{3x^2 + 1 : x \in \mathbb{R}\}.$$

Although the notation looks quite similar to the above (in both case we have curly braces, with a $:$ in the middle), this set is described differently. Notice that instead of having a set that elements are coming from on the left of the colon, we now have a set that elements are coming from on the right. Furthermore, we now have a formula on the left rather than a property on the right. The difference is that for a property, when we plug in an element from the given set, we either obtain a true or false value, but that isn’t the case for a formula like $3x^2 + 1$. The idea here is that instead of carving out a subset of \mathbb{R} by using a property (i.e. taking those elements that make the property *true*), we let x vary through all real numbers, plug each of these real numbers x into $3x^2 + 1$, and form the set of all possible outputs. For example, we have $4 \in S$ because $4 = 3 \cdot 1^2 + 1$. In other words, when $x = 1$, the left hand side gives the value 4, so we should put $4 \in S$. Notice also that $4 = 3 \cdot (-1)^2 + 1$, so we can also see that $4 \in S$ because of the “witness” -1 . Of course, we are forming a set, so we do not repeat the number 4. We also have $1 \in S$ because $1 = 3 \cdot 0^2 + 1$, and we have $76 \in S$ because $76 = 3 \cdot 5^2 + 1$. Notice also that $7 \in S$ because $7 = 3 \cdot (\sqrt{2})^2 + 1$.

In a general parametric set description, we will have a set A and a *function* $f(x)$ that allows inputs from A , and we write

$$\{f(x) : x \in A\}$$

for the set of all possible outputs of the function as we vary the inputs through the set A . We will discuss the general definition of a function in the next section, but for the moment you can think of them as given by formulas.

Now it is possible and indeed straightforward to turn any parametric description of a set into one where we carve out a subset by a property. In our case of $S = \{3x^2 + 1 : x \in \mathbb{R}\}$ above, we can alternatively write it as

$$S = \{y \in \mathbb{R} : \text{There exists } x \in \mathbb{R} \text{ with } y = 3x^2 + 1\}.$$

Notice how we flipped the way we described the set by introducing a “there exists” quantifier in order to form a property. This is always possible for a parametric description. For example, we have

$$\{5n + 4 : n \in \mathbb{N}\} = \{m \in \mathbb{N} : \text{There exists } n \in \mathbb{N} \text{ with } m = 5n + 4\}.$$

Thus, these parametric descriptions are not essentially new ways to describe sets, but they can often be more concise and clear.

By the way, we can use multiple parameters in our description. For example, consider the set

$$S = \{18m + 33n : m, n \in \mathbb{Z}\}.$$

Now we are simply letting m and n vary through all possible values in \mathbb{Z} and collecting all of the values $18m + 33n$ that result. For example, we have $15 \in S$ because $15 = 18 \cdot (-1) + 33 \cdot 1$. We also have $102 \in S$ because $102 = 18 \cdot 2 + 33 \cdot 2$. Notice that we are varying m and n independently, so they might take different values, or the same value (as in the case of $m = n = 2$). Don’t be fooled by the fact that we used different letters! As above, we can flip this description around by writing

$$S = \{k \in \mathbb{Z} : \text{There exists } m, n \in \mathbb{Z} \text{ with } k = 18m + 33n\}.$$

Subsets and Set Equality

Definition 1.5.4. *Given two sets A and B , we write $A \subseteq B$ to mean that every element of A is an element of B . More formally, $A \subseteq B$ means that for all $a \in A$, we have $a \in B$.*

To prove that $A \subseteq B$, one takes an arbitrary $a \in A$, and argues that $a \in B$. For example, let $A = \{6n : n \in \mathbb{Z}\}$ and let $B = \{2n : n \in \mathbb{Z}\}$. Since both of these sets are infinite, we can’t show that $A \subseteq B$ by taking each element of A in turn and showing that it is an element of B . Instead, we take an *arbitrary* $a \in A$, and show that $a \in B$. Here’s the proof.

Proposition 1.5.5. *Let $A = \{6n : n \in \mathbb{Z}\}$ and $B = \{2n : n \in \mathbb{Z}\}$. We have $A \subseteq B$.*

Proof. Let $a \in A$ be arbitrary. By definition of A , this means that we can fix an $m \in \mathbb{Z}$ with $a = 6m$. Notice then that $a = 2 \cdot (3m)$. Since $3m \in \mathbb{Z}$, it follows that $a \in B$. Since $a \in A$ was arbitrary, we conclude that $A \subseteq B$. \square

As usual, pause to make sure that you understand the logic of the argument above. First, we took an arbitrary element a from the set A . Now since $A = \{6n : n \in \mathbb{Z}\}$ and this is a parametric description with an implicit “there exists” quantifier, there must be one fixed integer value of n that puts a into the set A . In our proof, we chose to call that one fixed integer m . Now in order to show that $a \in B$, we need to exhibit a $k \in \mathbb{Z}$ with $a = 2k$. In order to do this, we hope to manipulate $a = 6m$ to introduce a 2, and ensure that the element we are multiplying by 2 is an integer.

What would go wrong if we tried to prove that $B \subseteq A$? Let’s try it. Let $b \in B$ be arbitrary. Since $b \in B$, we can fix $m \in \mathbb{Z}$ with $b = 2m$. Now our goal is to try to prove that we can find a $n \in \mathbb{Z}$ with $b = 6n$. It’s not obvious how to obtain a 6 from that 2, but we can try to force a 6 in the following way. Since $b = 2m$ and $2 = \frac{6}{3}$, we can write $b = 6 \cdot \frac{m}{3}$. We have indeed found a number n such that $b = 6n$, but we have not checked that this n is an integer. In general, dividing an integer by 3 does not result in an integer, so this argument currently has a hole in it.

Although that argument has a problem, we can not immediately conclude that $B \not\subseteq A$. Our failure to find an argument does not mean that an argument does not exist. So how can we show that $B \not\subseteq A$? All that we need to do is find just *one example* of an element of B that is not an element of A (because the negation of the “for all” statement $A \subseteq B$ is a “there exists” statement). We choose 2 as our example. However,

we need to convince everybody that this choice works. So let's do it! First, notice that $2 = 2 \cdot 1$, so $2 \in B$ because $1 \in \mathbb{Z}$. We now need to show that $2 \notin A$, and we will do this using a proof by contradiction. Suppose instead that $2 \in A$. Then, by definition, we can fix an $m \in \mathbb{Z}$ with $2 = 6m$. We then have that $m = \frac{2}{6} = \frac{1}{3}$. However, this is a contradiction because $\frac{1}{3} \notin \mathbb{Z}$. Since our assumption that $2 \in A$ led to a contradiction, we conclude that $2 \notin A$. We found an example of an element that is in B but not in A , so we conclude that $B \not\subseteq A$.

Recall that two sets A and B are defined to be equal if they have the same elements. Therefore, we have $A = B$ exactly when both $A \subseteq B$ and $B \subseteq A$ are true. Thus, given two sets A and B , we can prove that $A = B$ by performing two proofs like the one above. Such a strategy is called a *double containment* proof. We give an example of such an argument now.

Proposition 1.5.6. *Let $A = \{7n - 3 : n \in \mathbb{Z}\}$ and $B = \{7n + 11 : n \in \mathbb{Z}\}$. We have $A = B$.*

Proof. We prove that $A = B$ by showing that both $A \subseteq B$ and also that $B \subseteq A$.

- We first show that $A \subseteq B$. Let $a \in A$ be arbitrary. By definition of A , we can fix an $m \in \mathbb{Z}$ with $a = 7m - 3$. Notice that

$$\begin{aligned} a &= 7m - 3 \\ &= 7m - 14 + 11 \\ &= 7(m - 2) + 11. \end{aligned}$$

Now $m - 2 \in \mathbb{Z}$ because $m \in \mathbb{Z}$, so it follows that $a \in B$. Since $a \in A$ was arbitrary, we conclude that $A \subseteq B$.

- We now show that $B \subseteq A$. Let $b \in B$ be arbitrary. By definition of B , we can fix an $m \in \mathbb{Z}$ with $b = 7m + 11$. Notice that

$$\begin{aligned} b &= 7m + 11 \\ &= 7m + 14 - 3 \\ &= 7(m + 2) - 3. \end{aligned}$$

Now $m + 2 \in \mathbb{Z}$ because $m \in \mathbb{Z}$, so it follows that $b \in A$. Since $b \in B$ was arbitrary, we conclude that $B \subseteq A$.

We have shown that both $A \subseteq B$ and $B \subseteq A$ are true, so it follows that $A = B$. □

Here is a more interesting example. Consider the set

$$S = \{9m + 15n : m, n \in \mathbb{Z}\}.$$

For example, we have $9 \in S$ because $9 = 9 \cdot 1 + 15 \cdot 0$. We also have $3 \in S$ because $3 = 9 \cdot 2 + 15 \cdot (-1)$ (or alternatively because $3 = 9 \cdot (-3) + 15 \cdot 2$). We can always generate new values of S by simply plugging in values for m and n , but is there another way to describe the elements of S in an easier way? We now show that an integer is in S exactly when it is a multiple of 3.

Proposition 1.5.7. *We have $\{9m + 15n : m, n \in \mathbb{Z}\} = \{3m : m \in \mathbb{Z}\}$.*

Proof. We give a double containment proof.

- We first show that $\{9m + 15n : m, n \in \mathbb{Z}\} \subseteq \{3m : m \in \mathbb{Z}\}$. Let $a \in \{9m + 15n : m, n \in \mathbb{Z}\}$ be arbitrary. By definition, we can fix $k, \ell \in \mathbb{Z}$ with $a = 9k + 15\ell$. Notice that

$$\begin{aligned} a &= 9k + 15\ell \\ &= 3 \cdot (3k + 5\ell). \end{aligned}$$

Now $3k + 5\ell \in \mathbb{Z}$ because $k, \ell \in \mathbb{Z}$, so it follows that $a \in \{3m : m \in \mathbb{Z}\}$. Since $a \in \{9m + 15n : m, n \in \mathbb{Z}\}$ was arbitrary, we conclude that $\{9m + 15n : m, n \in \mathbb{Z}\} \subseteq \{3m : m \in \mathbb{Z}\}$.

- We now show that $\{3m : m \in \mathbb{Z}\} \subseteq \{9m + 15n : m, n \in \mathbb{Z}\}$. Let $a \in \{3m : m \in \mathbb{Z}\}$ be arbitrary. By definition, we can fix $k \in \mathbb{Z}$ with $a = 3k$. Notice that

$$\begin{aligned} a &= 3k \\ &= (9 \cdot (-3) + 15 \cdot 2) \cdot k \\ &= 9 \cdot (-3k) + 15 \cdot 2k. \end{aligned}$$

Now $-3k, 2k \in \mathbb{Z}$ because $k \in \mathbb{Z}$, so it follows that $a \in \{9m + 15n : m, n \in \mathbb{Z}\}$. Since $a \in \{3m : m \in \mathbb{Z}\}$ was arbitrary, we conclude that $\{3m : m \in \mathbb{Z}\} \subseteq \{9m + 15n : m, n \in \mathbb{Z}\}$.

We have shown that both $\{9m + 15n : m, n \in \mathbb{Z}\} \subseteq \{3m : m \in \mathbb{Z}\}$ and $\{3m : m \in \mathbb{Z}\} \subseteq \{9m + 15n : m, n \in \mathbb{Z}\}$ are true, so it follows that $\{9m + 15n : m, n \in \mathbb{Z}\} = \{3m : m \in \mathbb{Z}\}$. \square

Ordered Pairs and Sequences

Definition 1.5.8. An ordered pair is a collection of two (not necessarily distinct) objects, where order and repetition do matter.

We denote ordered pairs by using normal parentheses rather than curly braces. For example, we let $(2, 5)$ be the ordered pair whose first element is 2 and whose second element is 5. Notice that we have $(2, 5) \neq (5, 2)$ despite the fact that $\{2, 5\} = \{5, 2\}$. Make sure to keep a clear distinction between the ordered pair $(2, 5)$ and the set $\{2, 5\}$. We *do* allow the possibility of an ordered pair such as $(2, 2)$, and here the repetition of 2's is meaningful. Furthermore, we do not use \in in ordered pairs, so we would **not** write $2 \in (2, 5)$. We'll talk about ways to refer to the two elements of an ordered pair later.

We can generalize ordered pairs to the possibility of having more than two elements. In this case, we have an ordered list of n elements, like $(5, 4, 5, -2)$. We call such an object an *n-tuple*, a *list* with n elements, or a finite *sequence* of length n . Thus, for example, we could call $(5, 4, 5, -2)$ a 4-tuple. It is also possible to have infinite sequences (i.e. infinite lists), but we will wait to discuss these until the time comes.

Operations on Sets

Aside from listing elements, carving out subsets of a known set using a certain property, and giving a parametric description (which as mentioned above is just a special case of the previous type), there are other ways to build sets by using certain set-theoretic operations.

Definition 1.5.9. Given two sets A and B , we define $A \cup B$ to be the set consisting of those elements that are in A or B (or both). In other words, we define

$$A \cup B = \{x : x \in A \text{ or } x \in B\}.$$

We call this set the union of A and B .

Here, as in mathematics generally, we use *or* to mean “inclusive or”. In other words, if x is an element of both A and B , then we still put x into $A \cup B$. We give a few examples without proof:

- $\{1, 2, 7\} \cup \{4, 9\} = \{1, 2, 4, 7, 9\}$.
- $\{1, 2, 3\} \cup \{2, 3, 5\} = \{1, 2, 3, 5\}$.
- $\{2n : n \in \mathbb{Z}\} \cup \{2n + 1 : n \in \mathbb{Z}\} = \mathbb{Z}$. This is a restatement of Fact 1.4.6.

- $\{2n : n \in \mathbb{N}^+\} \cup \{2n + 1 : n \in \mathbb{N}^+\} = \{2, 3, 4, \dots\}$.
- $\{2n : n \in \mathbb{N}^+\} \cup \{2n - 1 : n \in \mathbb{N}^+\} = \{1, 2, 3, 4, \dots\} = \mathbb{N}^+$.
- $A \cup \emptyset = A$ for every set A .

Definition 1.5.10. Given two sets A and B , we define $A \cap B$ to be the set consisting of those elements that are in both A and B . In other words, we define

$$A \cap B = \{x : x \in A \text{ and } x \in B\}.$$

We call this set the intersection of A and B .

Here are a few examples, again without proof:

- $\{1, 2, 7\} \cap \{4, 9\} = \emptyset$.
- $\{1, 2, 3\} \cap \{2, 3, 5\} = \{2, 3\}$.
- $\{1, \{2, 3\}\} \cap \{1, 2, 3\} = \{1\}$.
- $\{2n : n \in \mathbb{Z}\} \cap \{3n : n \in \mathbb{Z}\} = \{6n : n \in \mathbb{Z}\}$.
- $\{3n + 1 : n \in \mathbb{N}^+\} \cap \{3n + 2 : n \in \mathbb{N}^+\} = \emptyset$.
- $A \cap \emptyset = \emptyset$ for every set A .

Definition 1.5.11. We say that two sets A and B are disjoint if $A \cap B = \emptyset$.

Definition 1.5.12. Given two sets A and B , we define $A \setminus B$ to be the set consisting of those elements that are in A , but not in B . In other words, we define

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

We call this set the (relative) complement of B (in A).

In many cases where we consider $A \setminus B$, we will have that $B \subseteq A$, but we will occasionally use it even when $B \not\subseteq A$. Here are a few examples:

- $\{5, 6, 7, 8, 9\} \setminus \{5, 6, 8\} = \{7, 9\}$.
- $\{1, 2, 7\} \setminus \{4, 9\} = \{1, 2, 7\}$.
- $\{1, 2, 3\} \setminus \{2, 3, 5\} = \{1\}$.
- $\{2n : n \in \mathbb{Z}\} \setminus \{4n : n \in \mathbb{Z}\} = \{4n + 2 : n \in \mathbb{Z}\}$.
- $A \setminus \emptyset = A$ for every set A .
- $A \setminus A = \emptyset$ for every set A .

Definition 1.5.13. Given two sets A and B , we let $A \times B$ be the set of all ordered pairs (a, b) such that $a \in A$ and $b \in B$, and we call this set the Cartesian product of A and B .

For example, we have

$$\{1, 2, 3\} \times \{6, 8\} = \{(1, 6), (1, 8), (2, 6), (2, 8), (3, 6), (3, 8)\}$$

and

$$\mathbb{N} \times \mathbb{N} = \{(0, 0), (0, 1), (1, 0), (2, 0), \dots, (4, 7), \dots\}.$$

Notice that elements of $\mathbb{R} \times \mathbb{R}$ correspond to points in the plane.

We can also generalize the concept of a Cartesian product to more than 2 sets. If we are given n sets A_1, A_2, \dots, A_n , we let $A_1 \times A_2 \times \dots \times A_n$ be the set of all n -tuples (a_1, a_2, \dots, a_n) such that $a_i \in A_i$ for each i . For example, we have

$$\{1, 2\} \times \{3\} \times \{4, 5\} = \{(1, 3, 4), (1, 3, 5), (2, 3, 4), (2, 3, 5)\}.$$

In the special case when A_1, A_2, \dots, A_n are all the same set A , we use the notation A^n to denote the set $A \times A \times \dots \times A$ (where we have n copies of A). Thus, A^n is the set of all finite sequences of elements of A of length n . For example, $\{0, 1\}^n$ is the set of all finite sequences of 0's and 1's of length n . Notice that this notation fits in with the notation \mathbb{R}^n that we are used to from Calculus, and we will continue to use it throughout Linear Algebra.

1.6 Functions

We're all familiar with functions from Calculus. In that context, a function is often given by a “formula”, such as $f(x) = x^4 - 4x^3 + 2x - 1$. However, we also encounter piecewise-defined functions, such as

$$f(x) = \begin{cases} x^2 + 1 & \text{if } x \geq 2, \\ x - 1 & \text{if } x < 2, \end{cases}$$

and the function $g(x) = |x|$, which is really piecewise defined as

$$g(x) = \begin{cases} x & \text{if } x \geq 0, \\ -x & \text{if } x < 0. \end{cases}$$

For a more exotic example of a piecewise defined function, consider

$$f(x) = \begin{cases} 1 & \text{if } x \in \mathbb{Q}, \\ 0 & \text{if } x \notin \mathbb{Q}. \end{cases}$$

Despite these examples, only the most basic functions in mathematics are defined through formulas on pieces. For instance, the function $f(x) = \sin x$ is *not* given by a formula, and it is difficult to compute values of this function with any accuracy using only basic operations like $+$ and \cdot . In fact, we give this function the strange new name of “sine” because we can not express it easily using more basic operations. The function $f(x) = 2^x$ is easy to compute for $x \in \mathbb{Z}$, but it's actually nontrivial to compute, and even define, in general (after all, do you remember the definition of 2^π ?). Even more fundamentally, the function $f(x) = \sqrt{x}$ is also not given by a formula, because the definition, i.e. $f(x)$ is the unique $y \geq 0$ with the property that $y^2 = x$, does not give us an easy way to compute it.

Beyond these fundamental functions that you encounter before Calculus, you learn more exotic ways to define functions in Calculus. Given a function f , you learn how to define a new function f' , called the derivative of f , using a certain limit operation. Now in many cases, you can compute f' more easily using

facts like the Product Rule and the Chain Rule, but these rules do not always apply. Moreover, given any continuous function g , we can define a new function f by letting

$$f(x) = \int_0^x g(t) \, dt.$$

In other words, f is defined as the “(signed) area of g so far” function, in that $f(x)$ is defined to be the (signed) area between the graph of g and the x -axis over the interval from 0 to x . Formally, f is defined as a limit of Riemann sums. Again, in Calculus you learn ways to compute $f(x)$ more easily in many special cases using the Fundamental Theorem of Calculus. For example, if

$$f(x) = \int_0^x (3t^2 + t) \, dt,$$

then we can also compute f as

$$f(x) = x^3 + \frac{x^2}{2},$$

while if

$$f(x) = \int_0^x \sin t \, dt,$$

then we can also compute f as

$$f(x) = 1 - \cos x.$$

However, not all integrals can be evaluated so easily. In fact, it turns out that the perfectly well-defined function

$$f(x) = \int_0^x e^{-t^2} \, dt$$

can not be expressed through polynomials, exponentials, logs, and trigonometric functions using only operations like $+$, \cdot , and function composition. Of course, we can still approximate it using Riemann sums (or Simpson’s Rule), and this is important for us to be able to do since this function represents the area under a normal curve, which is essential in statistics.

If we move away from functions whose inputs and outputs are real numbers, we can think about other interesting ways to define functions. For example, suppose we define a function whose inputs and outputs are elements of \mathbb{R}^2 by letting $f(\vec{u})$ be the result of rotating \vec{u} by 27° clockwise around the origin. This seems to be a well-defined function despite the fact that it is not clear how to compute it (though we will learn how to compute it in time).

Alternatively, consider a function whose inputs and outputs are natural numbers defined by letting $f(n)$ be the number of primes less than or equal to n . For example, we have $f(3) = 2$, $f(4) = 2$, $f(9) = 4$, and $f(30) = 10$. Although it is possible to compute this function, it’s not clear whether we can compute it quickly. In other words, it’s not obvious if we can compute something like $f(2^{50})$ without a huge amount of work.

Perhaps you have some exposure to the concept of a function as it is used in computer programming. From this perspective, a function is determined by a sequence of imperative statements or function compositions as defined by a precise programming language. Since a computer is doing the interpreting, of course all such functions can be computed in principle (or if your computations involve real numbers, then at least up to good approximations). However, if you take this perspective, an interesting question arises. If we write two different functions f and g that do not follow the same steps, and perhaps even act qualitatively differently in structure, but they always produce the same output on the same input, should we consider them to be the same function? We can even ask this question outside of the computer science paradigm. For example, if we define $f(x) = \sin^2 x + \cos^2 x$ and $g(x) = 1$, then should we consider f and g to be the same function?

We need to make a choice about how to define a function in general. Intuitively, given two sets A and B , a function $f: A \rightarrow B$ is an input-output “mechanism” that produces a *unique* output $b \in B$ for any given input $a \in A$. As we’ve seen, the vast majority of functions that we have encountered so far can be computed in principle, so up until this point, we could interpret “mechanism” in an algorithmic and computational sense. However, we want to allow as much freedom as possible in this definition so that we can consider new ways to define functions in time. In fact, as you might see in later courses (like Automata, Formal Languages, and Computational Complexity), there are some natural functions that are not computable even in theory. As a result, we choose to abandon the notion of computation in our definition. By making this choice, we will be able to sidestep some of the issues in the previous paragraph, but we still need to make a choice about whether to consider the functions $f(x) = \sin^2 x + \cos^2 x$ and $g(x) = 1$ to be equal.

With all of this background, we are now in a position to define functions as certain special types of sets. Thinking about functions from this more abstract point of view eliminates the vague “mechanism” concept because they will simply be sets. With this perspective, we’ll see that functions can be defined in any way that a set can be defined. Our approach both clarifies the concept of a function and also provides us with some much needed flexibility in defining functions in more interesting ways. Here is the formal definition.

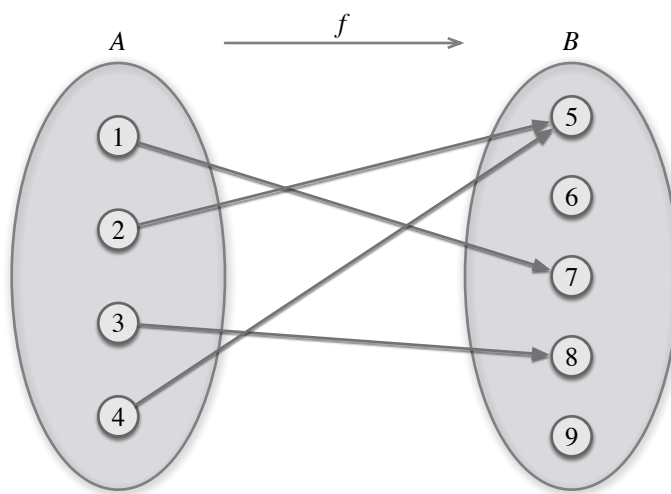
Definition 1.6.1. *Let A and B be sets. A function from A to B is a subset f of $A \times B$ with the property that for all $a \in A$, there exists a unique $b \in B$ with $(a, b) \in f$. Also, instead of writing “ f is a function from A to B ”, we typically use the shorthand notation “ $f: A \rightarrow B$ ”.*

For example, let $A = \{1, 2, 3, 4\}$ and let $B = \{5, 6, 7, 8, 9\}$. An example of a function $f: A \rightarrow B$ is the set

$$f = \{(1, 7), (2, 5), (3, 8), (4, 5)\}.$$

Intuitively, we think of f as an input-output mechanism that sends the input 1 to the output 7, the input 2 to the output 5, the input 3 to the output 8, and the input 4 to the output 5. Notice that each input is sent to exactly one output, which is the content of the phrase “there exists a unique” in the above definition.

You are certainly used to visualizing functions $f: \mathbb{R} \rightarrow \mathbb{R}$ by looking at their “graphs”, i.e. the curve in the plane that contains the corresponding input-output pairs. These visualizations are incredibly useful, but they do not generalize well to situations where the inputs and outputs are not real numbers. A more abstract, but significantly more flexible, way to visualize a function $f: A \rightarrow B$ is to draw the set A on the left, the set B on the right, and then include an arrow from each input on the left to the corresponding output on the right. Here is the corresponding picture for our function above:



Notice that the fact that every input produces exactly one output corresponds to the fact that every element of A has exactly one arrow coming out of it. In contrast, for the set B on the right, notice that some elements have no arrows coming in to them, some have one arrow coming in, and one has multiple arrows coming into it. This represents a fundamental (and incredibly important) asymmetry in the definition of a function. From the definition of a function $f: A \rightarrow B$, we know that for every $a \in A$, there is a unique $b \in B$ such that $(a, b) \in f$. However, as this example shows, it might *not* be the case that for every $b \in B$, there is a unique $a \in A$ with $(a, b) \in f$. In fact, for a given $b \in B$, there might be *no* $a \in A$ with $f(a) = b$, or there might be *many* $a \in A$ with $f(a) = b$. Be careful with the order of quantifiers!

Thinking of functions as special types of sets is helpful to clarify the definition, but is often awkward to work with in practice. For example, writing $(2, 5) \in f$ to mean that f sends 2 to 5 quickly becomes annoying. Thus, we introduce some new notation matching up with our old experience with functions.

Notation 1.6.2. *Let A and B be sets. If $f: A \rightarrow B$ and $a \in A$, we write $f(a)$ to mean the unique $b \in B$ such that $(a, b) \in f$.*

For instance, in the above example of f , we can instead write

$$f(1) = 7, \quad f(2) = 5, \quad f(3) = 8, \quad \text{and} \quad f(4) = 5.$$

We can also convert the typical way of defining a function into this formal set theoretic way. For example, consider the function $f: \mathbb{R} \rightarrow \mathbb{R}$ by letting $f(x) = x^2$. We can instead define f by the set

$$\{(x, y) \in \mathbb{R} \times \mathbb{R} : y = x^2\},$$

or parametrically as

$$\{(x, x^2) : x \in \mathbb{R}\}.$$

One side effect of our definition of a function is that we immediately obtain a nice definition for when two functions $f: A \rightarrow B$ and $g: A \rightarrow B$ are equal because we have defined when two sets are equal. Given two functions $f: A \rightarrow B$ and $g: A \rightarrow B$, if we unwrap our definition of set equality, we see that $f = g$ exactly when f and g have the same elements, which is precisely the same thing as saying that $f(a) = g(a)$ for all $a \in A$. In particular, the *manner* in which we describe functions does not matter so long as the functions behave the same on all inputs. For example, if we define $f: \mathbb{R} \rightarrow \mathbb{R}$ and $g: \mathbb{R} \rightarrow \mathbb{R}$ by letting $f(x) = \sin^2 x + \cos^2 x$ and $g(x) = 1$, then we have that $f = g$ because $f(x) = g(x)$ for all $x \in \mathbb{R}$.

Definition 1.6.3. *Let $f: A \rightarrow B$ be a function.*

- *We call A the domain of f .*
- *We call B the codomain of f .*
- *We define $\text{range}(f) = \{b \in B : \text{There exists } a \in A \text{ with } f(a) = b\}$.*

Notice that given a function $f: A \rightarrow B$, we have $\text{range}(f) \subseteq B$, but it is possible that $\text{range}(f) \neq B$. For example, in the above case, we have that the codomain of f is $\{5, 6, 7, 8, 9\}$, but $\text{range}(f) = \{5, 7, 8\}$. In general, given a function $f: A \rightarrow B$, it may be very difficult to determine $\text{range}(f)$ because we may need to search through all $a \in A$.

For an interesting example of a function with a mysterious looking range, fix $n \in \mathbb{N}^+$ and define $f: \{0, 1, 2, \dots, n-1\} \rightarrow \{0, 1, 2, \dots, n-1\}$ by letting $f(a)$ be the remainder when dividing a^2 by n . For example, if $n = 10$, then we have the following table of values of f :

$$\begin{array}{ccccc} f(0) = 0 & f(1) = 1 & f(2) = 4 & f(3) = 9 & f(4) = 6 \\ f(5) = 5 & f(6) = 6 & f(7) = 9 & f(8) = 4 & f(9) = 1. \end{array}$$

Thus, for $n = 10$, we have $\text{range}(f) = \{0, 1, 4, 5, 6, 9\}$. This simple but strange looking function has many interesting properties. Given a reasonably large number $n \in \mathbb{N}$, it looks potentially difficult to determine whether an element is in $\text{range}(f)$ because we might need to search through a huge number of inputs to see if a given output actually occurs. If n is prime, then it turns out that there are much faster ways to determine if a given element is in $\text{range}(f)$ (see Number Theory). However, it is widely believed (although we do not currently have a proof!) that there is no efficient method to do this when n is the product of two large primes, and this is the basis for some cryptosystems (Goldwasser-Micali) and pseudo-random number generators (Blum-Blum-Shub).

We now turn our attention to a fundamental operation on functions. You are familiar with the idea of composition from Calculus (and earlier mathematics courses), as the Chain Rule is about how to differentiate the composition of two differentiable functions. Let's recall the basic idea of function composition through an example. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be the function $f(x) = x + 1$, and let $g: \mathbb{R} \rightarrow \mathbb{R}$ be the function $g(x) = x^2$. We can then form a new function $f \circ g$ by first performing g and then performing f . That is, given an input x , we first compute $g(x)$, and then use that output as an input to f in order to obtain the value $f(g(x))$. For example, on input 2, we notice that $g(2) = 2^2 = 4$, and that $f(4) = 4 + 1 = 5$, so $(f \circ g)(2) = 5$. In terms of symbols and formulas, the function $f \circ g: \mathbb{R} \rightarrow \mathbb{R}$ is given by

$$\begin{aligned}(f \circ g)(x) &= f(g(x)) \\ &= f(x^2) \\ &= x^2 + 1.\end{aligned}$$

We can also form the function $g \circ f: \mathbb{R} \rightarrow \mathbb{R}$. Now given an input x , we first compute $f(x)$, and then use that output as an input to g in order to obtain the value $g(f(x))$. For example, on input 2, we notice that $f(2) = 2 + 1 = 3$, and that $g(3) = 3^2 = 9$, so $(g \circ f)(2) = 9$. In terms of symbols and formulas, the function $g \circ f: \mathbb{R} \rightarrow \mathbb{R}$ is given by

$$\begin{aligned}(g \circ f)(x) &= g(f(x)) \\ &= g(x + 1) \\ &= (x + 1)^2 \\ &= x^2 + 2x + 1.\end{aligned}$$

In particular, notice that order that we composed the functions matters! We have $(f \circ g)(2) = 5$ while $(g \circ f)(2) = 9$. Since we have found one example of an $x \in \mathbb{R}$ with $(f \circ g)(x) \neq (g \circ f)(x)$, we conclude that $f \circ g \neq g \circ f$. It does not matter that there do exist some values of x with $(f \circ g)(x) = (g \circ f)(x)$ (for example, this is true when $x = 0$). Remember that two functions are equal precisely when they agree on *all* inputs, so to show that the two functions are not equal it suffices to find just one value where they disagree (again remember that the negation of a “for all” statement is a “there exists” statement).

We now define composition of functions in general. Suppose that $f: A \rightarrow B$ and $g: C \rightarrow D$ are functions. Can we always form $g \circ f$? Given an $a \in A$, we can use it as an input to f to obtain $f(a) \in B$. But it might not make sense to use $f(a)$ as an input to g . After all, the domain of g is C , and there is no reason to believe that $f(a) \in C$. As a result, when trying to form $g \circ f$, we restrict attention to the case where the codomain of f is equal to the domain of g .

Definition 1.6.4. Suppose that $f: A \rightarrow B$ and $g: B \rightarrow C$ are functions. The composition of g and f , denoted $g \circ f$, is the function $g \circ f: A \rightarrow C$ defined by letting $(g \circ f)(a) = g(f(a))$ for all $a \in A$.


Let's see an example. Let $A = \{1, 2, 3, 4\}$, let $B = \{5, 6, 7, 8, 9\}$, and let $C = \{3, 6, 9\}$. Let $f: A \rightarrow B$ be our function from above defined by letting

$$f(1) = 7, \quad f(2) = 5, \quad f(3) = 8, \quad \text{and} \quad f(4) = 5.$$

Let $g: B \rightarrow C$ be the function defined by

$$g(5) = 9, \quad g(6) = 6, \quad g(7) = 9, \quad g(8) = 3, \quad \text{and} \quad g(9) = 6.$$


Using our more abstract visualization of functions, we can picture this situation in the following way:



Since the codomain of f equals the domain of g (both are equal to B), we can form the composition $g \circ f$, which is a function with domain A and codomain C . Let's see how $g \circ f$ behaves. By definition, we have

$$\begin{aligned} (g \circ f)(1) &= g(f(1)) \\ &= g(5) \\ &= 9, \end{aligned}$$

so $g \circ f$ sends the input $1 \in A$ to the output $9 \in C$. If we carry out the corresponding computation for the other three elements of A , we see that $g \circ f$ can be visualized as follows:



Notice that in these more general situations, we should *not* think of the composition $g \circ f$ as plugging in the “formula” for f in for the input variable of g . After all, there is no “formula” for a function in general! Instead, the composition $g \circ f$ is the result of chaining together the input-output mechanisms of the functions f and g . That is, given $a \in A$, feed it as input to f to obtain the output $f(a) \in B$, and then take this value and use it as input to g in order to obtain $g(f(a)) \in C$. In the above picture, we can visualize this chaining process as simply following the arrows. That is, given an input $a \in A$, follow the arrow with tail at a to obtain $f(a) \in B$, and then follow the resulting arrow with tail at $f(a)$ to obtain $g(f(a)) \in C$.

We saw above that even for functions $f, g: \mathbb{R} \rightarrow \mathbb{R}$, we might have $f \circ g \neq g \circ f$. In other words, function composition is not commutative in general. However, the situation is even worse! It is possible for one of these compositions to make sense while the other one does not. Consider our example functions $f: A \rightarrow B$ and $g: B \rightarrow C$ given above. We noticed that $g \circ f$ makes sense because the codomain of f equaled the domain of g . However, it does *not* make sense to form $f \circ g$. After all, the codomain of g is C , while the domain of f is A . Thus, we can not chain together the two functions on all inputs. That is, even though $g(8) = 3 \in A$ so we can form $f(g(8))$, we have $g(5) = 9 \notin A$, so $f(g(5))$ is meaningless.

Despite the fact that function composition is not commutative, it is at least associative whenever the corresponding compositions make sense.

Proposition 1.6.5. *Let A, B, C, D be sets. Suppose that $f: A \rightarrow B$, that $g: B \rightarrow C$, and that $h: C \rightarrow D$ are functions. We then have that $(h \circ g) \circ f = h \circ (g \circ f)$. Stated more simply, function composition is associative whenever it is defined.*

Proof. Let $a \in A$ be arbitrary. We then have

$$\begin{aligned} ((h \circ g) \circ f)(a) &= (h \circ g)(f(a)) \\ &= h(g(f(a))) \\ &= h((g \circ f)(a)) \\ &= (h \circ (g \circ f))(a), \end{aligned}$$

where each step follows by definition of composition. Therefore $((h \circ g) \circ f)(a) = (h \circ (g \circ f))(a)$ for all $a \in A$. It follows that $(h \circ g) \circ f = h \circ (g \circ f)$. \square

Although the above argument is a complete and rigorous proof of the associativity of function composition, it is also useful to visualize the situation in pictures. Imagine that $f: A \rightarrow B$, that $g: B \rightarrow C$, and that $h: C \rightarrow D$ are all functions. Think about adding another oval D to the above picture that we used for function composition, and including arrows from elements of C to elements of D according to the function h . Now convince yourself that each of $(h \circ g) \circ f$ and $h \circ (g \circ f)$ represent “follow the arrows through the picture”. That is, for $(h \circ g) \circ f$, we first follow the arrows from B to D to form the function $h \circ g$, and then chain these resulting arrows with the arrows from A to B to form $(h \circ g) \circ f$. In contrast, for $h \circ (g \circ f)$, we first follow the arrows from A to C to form the function $g \circ f$, and then chain these resulting arrows with the arrows from C to D to form $h \circ (g \circ f)$. In both cases, we are just following the arrows through the picture, so we should expect $(h \circ g) \circ f$ to equal $h \circ (g \circ f)$.

Before moving on, we pause to define and examine what is probably the most boring example of a function on a set.

Definition 1.6.6. *Let A be a set. The function $\text{id}_A: A \rightarrow A$ defined by $\text{id}_A(a) = a$ for all $a \in A$ is called the identity function on A .*

We call this function the identity function because it leaves other functions alone when we compose with it. However, we have to be careful that we compose with the identity function on the correct set and the correct side.

Proposition 1.6.7. *For any function $f: A \rightarrow B$, we have $f \circ \text{id}_A = f$ and $\text{id}_B \circ f = f$.*

Proof. Let $f: A \rightarrow B$ be an arbitrary function.

- We first show that $f \circ id_A = f$. Let $a \in A$ be arbitrary. We have

$$\begin{aligned} (f \circ id_A)(a) &= f(id_A(a)) && \text{(by definition of composition)} \\ &= f(a) \end{aligned}$$

Since $a \in A$ was arbitrary, it follows that $f \circ id_A = f$.

- We now show that $id_B \circ f = f$. Let $a \in A$ be arbitrary. We have

$$\begin{aligned} (id_B \circ f)(a) &= id_B(f(a)) && \text{(by definition of composition)} \\ &= f(a) && \text{(because } f(a) \in B \text{)} \end{aligned}$$

Since $a \in A$ was arbitrary, it follows that $id_B \circ f = f$.

□

1.7 Injective, Surjective, and Bijective Functions

Recall that the key property in the definition of a function $f: A \rightarrow B$ is that every input element from A produces a unique output element from B . However, this does not work in reverse. Given $b \in B$, it may be the case that b is the output of zero, one, or many elements from A . We give special names to the types of functions that have limitations on how often each element of B actually occurs as an output.

Definition 1.7.1. Let $f: A \rightarrow B$ be a function.

- We say that f is *injective* (or *one-to-one*) if whenever $a_1, a_2 \in A$ satisfy $f(a_1) = f(a_2)$, we have $a_1 = a_2$.
- We say that f is *surjective* (or *onto*) if for all $b \in B$, there exists $a \in A$ such that $f(a) = b$.
- We say that f is *bijective* if f is both injective and surjective.

Let's take a moment to unpack these definitions. First, saying that a function $f: A \rightarrow B$ is surjective is simply saying that every $b \in B$ is hit at least once by an element $a \in A$. In terms of our above visualizations, a function $f: A \rightarrow B$ is surjective if every element of B has at least one arrow coming in to it. We can rephrase this idea using Definition 1.6.3 by saying that $f: A \rightarrow B$ is surjective exactly when $\text{range}(f) = B$.

The definition of injective is slightly more mysterious at first. Intuitively, a function $f: A \rightarrow B$ is injective when every $b \in B$ is hit by at most one $a \in A$, i.e. if every element of B has at most one arrow coming into it. Now saying this precisely takes a little bit of thought. After all, it is difficult to say “there exists at most one” because our “there exists” quantifier is used to mean that there is at least one! The idea is to turn this around and not directly talk about $b \in B$ at all. Instead, we want to say that we never have a situation where we have two distinct elements $a_1, a_2 \in A$ that go to the same place under f . Thus, we want to say

$$\text{“Not(There exists } a_1, a_2 \in A \text{ with } a_1 \neq a_2 \text{ and } f(a_1) = f(a_2)\text{)”}.$$

We can rewrite this statement as

$$\text{“For all } a_1, a_2 \in A, \text{ we have Not}(a_1 \neq a_2 \text{ and } f(a_1) = f(a_2)\text{)”},$$

which is equivalent to

$$\text{“For all } a_1, a_2 \in A, \text{ we have either } a_1 = a_2 \text{ or } f(a_1) \neq f(a_2)\text{”}$$

(notice that the negation of the “and” statement became an “or” statement). Finally, we can rewrite this as the following “if...then...” statement:

“For all $a_1, a_2 \in A$, if $a_1 \neq a_2$, then $f(a_1) \neq f(a_2)$ ”.

Looking at our statement here, it captures what we want to express perfectly because it says that distinct inputs always go to distinct outputs, which exactly says no element of B is hit by 2 or more elements, and hence that every element of B is hit by at most 1 element. Thus, we could indeed take this as our definition of injective. The problem is that this definition is difficult to use in practice. To see why, think about how we would argue that a given function $f: A \rightarrow B$ is injective. It appears that we would want to take arbitrary $a_1, a_2 \in A$ with $a_1 \neq a_2$, and argue that under this assumption we must have that $f(a_1) \neq f(a_2)$. Now the problem with this is that is very difficult to work with an expression involving \neq in ways that preserve truth. For example, we have that $-1 \neq 1$, but $(-1)^2 = 1^2$, so we can not square both sides and preserve non-equality. To get around this problem, we instead take the contrapositive of the statement in question, which turns into our formal definition of injective:

“For all $a_1, a_2 \in A$, if $f(a_1) = f(a_2)$, then $a_1 = a_2$ ”.

Notice that in our definition above, we simply replace the “for all... if... then...” construct with a “when-ever...we have...” for clarity, but these are saying precisely the same thing, i.e. that whenever we have two elements of A that happen to be sent to the same element of B , then in fact those two elements of A must be the same. Although our official definition is slightly harder to wrap one’s mind around, it is *much* easier to work with in practice. To prove that a given $f: A \rightarrow B$ is injective, we take arbitrary $a_1, a_2 \in A$ with $f(a_1) = f(a_2)$, and use this equality to derive the conclusion that $a_1 = a_2$.

To recap the colloquial ways to understand these concepts, a function $f: A \rightarrow B$ is injective if every $b \in B$ is hit by at most one $a \in A$, and is surjective if every $b \in B$ is hit by at least one $a \in A$. It follows that a function $f: A \rightarrow B$ is bijective if every $b \in B$ is hit by exactly one $a \in A$. These ways of thinking about injective and surjective are great, but we need to be careful when proving that a function is injective or surjective. Given a function $f: A \rightarrow B$, here is the general process for proving that it has one or both of these properties:

- Suppose we want to prove that a function $f: A \rightarrow B$ is injective. Notice that the definition involves a “for all” quantifier. Thus, we should start by taking arbitrary $a_1, a_2 \in A$ that satisfy $f(a_1) = f(a_2)$, and then work forward to derive that $a_1 = a_2$. In this way, we show that whenever two elements of A happen to produce the same output, then they must have been the same element all along.
- Suppose we want to prove that a function $f: A \rightarrow B$ is surjective. Notice that the definition is a “for all...there exists...” statement. Thus, we should start by taking an arbitrary $b \in B$, and then show how to build an $a \in A$ with $f(a) = b$. In other words, we need to take an arbitrary $b \in B$ and fill in the blank in $f(\text{---}) = b$ with an element of A .

Here is an example.

Proposition 1.7.2. *The function $f: \mathbb{R} \rightarrow \mathbb{R}$ given by $f(x) = 5x - 7$ is bijective.*

Proof. We need to show that f is both injective and surjective.

- We first show that f is injective. Let $x_1, x_2 \in \mathbb{R}$ be arbitrary with $f(x_1) = f(x_2)$. By definition of f , we then have that $5x_1 - 7 = 5x_2 - 7$. Adding 7 to both sides, we see that $5x_1 = 5x_2$. Dividing both sides by 5, we conclude that $x_1 = x_2$. Therefore, f is injective.

- We next show that f is surjective. Let $y \in \mathbb{R}$ be arbitrary. Notice that $\frac{y+7}{5} \in \mathbb{R}$ and that

$$\begin{aligned} f\left(\frac{y+7}{5}\right) &= 5 \cdot \frac{y+7}{5} - 7 \\ &= (y+7) - 7 \\ &= y. \end{aligned}$$

Thus, we have shown the existence of an $x \in \mathbb{R}$ with $f(x) = y$. Therefore, f is surjective.

Since f is both injective and surjective, it follows that f is bijective. \square

Here's another example. In this case, we show that a function is not surjective. Let's think the logic here. Given a function $f: A \rightarrow B$, the statement that f is surjective is:

“For all $b \in B$, there exists $a \in A$ such that $f(a) = b$ ”.

Saying that a function $f: A \rightarrow B$ is not surjective is saying that the negation of this statement is true. Now we can express the negation of this statement as follows:

“There exists $b \in B$ such that for all $a \in A$, we have $f(a) \neq b$ ”.

Thus, to argue that a function $f: A \rightarrow B$ is not surjective, we have to give an example of an element $b \in B$ that is not hit by any $a \in A$.

Proposition 1.7.3. *The function $f: \mathbb{Z} \rightarrow \mathbb{Z}$ given by $f(n) = 2n$ is injective but not surjective.*

Proof. We show each part:

- We first show that f is injective. Let $n_1, n_2 \in \mathbb{Z}$ be arbitrary with $f(n_1) = f(n_2)$. By definition of f , we then have that $2n_1 = 2n_2$. Dividing both sides by 2, we conclude that $n_1 = n_2$. Therefore, f is injective.
- We next show that f is not surjective. We claim that $1 \in \mathbb{Z}$ has the property that $f(n) \neq 1$ for all $n \in \mathbb{Z}$. To see this, suppose instead that there does exist an $n \in \mathbb{Z}$ with $f(n) = 1$, and fix such an n . We would then have that $2n = 1$, so 1 would be even, which contradicts Proposition 1.4.4. It follows that $f(n) \neq 1$ for all $n \in \mathbb{Z}$, and hence f is not surjective. \square

Notice that if we define $g: \mathbb{R} \rightarrow \mathbb{R}$ by letting $g(x) = 2x$, then g is injective by the same proof, but g is surjective (think about why). Thus, changing the domain or codomain of a function can change the properties of that function. We will have much more to say about injective and surjective functions in time.

Given a general function $f: A \rightarrow B$, we've seen that every element of A produces a unique element of B (by the definition of a function), but it might not be the case that every element of B is hit by a unique element of A . Thus, there is a serious asymmetry between the domain and the codomain. In particular, we can not typically “undo” a function. That is, if we have a function $f: A \rightarrow B$ together with a particular $b \in B$, we can not generally go backwards and produce the unique element of A that hits it. After all, if a given $b \in B$ is not in $\text{range}(f)$, then there is no natural candidate element of A , and if $b \in B$ is in $\text{range}(f)$, then there might be several candidates.

Now if $f: A \rightarrow B$ is a bijection, then all of these problems go away, because then every $b \in B$ is hit by a unique $a \in A$. In this setting, we can define a “backwards” function $g: B \rightarrow A$ by just reversing the arrows. The resulting function is naturally called the *inverse* of f . Here is the formal definition.

Definition 1.7.4. *Let $f: A \rightarrow B$ be a bijective function. We define a new function $f^{-1}: B \rightarrow A$ as follows. Given $b \in B$, let $f^{-1}(b)$ be the unique $a \in A$ that satisfies $f(a) = b$.*

At first, this definition might look strange, because f^{-1} is defined in a wordy way without information about how to “compute” it. However, remember that a function need not be given by a formula! For any $b \in B$, we have identified a unique element of A to assign to the value $f^{-1}(b)$, and that is all we need to do to define a function.

For example, consider $A = \{1, 2, 3\}$ and $B = \{4, 5, 6\}$. Define $f: A \rightarrow B$ as follows:

$$f(1) = 6, \quad f(2) = 5, \quad \text{and} \quad f(3) = 4.$$

Notice that f is bijective, so f^{-1} is defined as a function with domain B and codomain A . What is $f^{-1}(4)$? Looking at the definition, we need to find the unique $a \in A$ with $f(a) = 4$. In this case, the unique such value is 3, so $f^{-1}(4) = 3$. Working through the other values, we determine the following:

$$f^{-1}(4) = 3, \quad f^{-1}(5) = 2, \quad \text{and} \quad f^{-1}(6) = 1.$$

Here is another way to think about the inverse function in terms of sets. As a set, we have

$$f = \{(1, 6), (2, 5), (3, 4)\}.$$

Notice that

$$f^{-1} = \{(4, 3), (5, 2), (6, 1)\},$$

which is just the result of flipping each of the ordered pairs in f . In general, if you choose to view a function $f: A \rightarrow B$ as the corresponding set of ordered pairs, then

$$f^{-1} = \{(b, a) \in B \times A : (a, b) \in f\}.$$

Although this is a clever set-theoretic way to think about f^{-1} , and it feels closer to the “flipping the arrows” intuition, it is a bit abstract and typically not as useful when the domain or codomain are infinite.

Let’s see how this works in a context where the domain and codomain are infinite. As mentioned in Section 1.6, given $x \geq 0$ we typically think of \sqrt{x} as the unique $y \geq 0$ with the property that $y^2 = x$. That’s a wordy definition that looks like an inverse! To explore this idea, we start by looking at the function $g: \mathbb{R} \rightarrow \mathbb{R}$ defined by $g(x) = x^2$. Unfortunately, g is not bijective. It is not injective because $g(1) = 1 = g(-1)$, but $1 \neq -1$. Also, it is not surjective because $-1 \notin \text{range}(g)$.

Let’s fix this issue by changing the domain and codomain. Let $A = \{x \in \mathbb{R} : x \geq 0\}$ be the set of nonnegative real numbers. Now define $f: A \rightarrow A$ by letting $f(x) = x^2$. It is possible to show that f is bijective (one uses the fact that f is increasing to show that f is injective, and the Intermediate Value Theorem from Analysis to argue that f is surjective). Thus, f^{-1} is defined as a function with domain A and codomain A , i.e. we have a function $f^{-1}: A \rightarrow A$. By definition, given $x \in A$, we have defined $f^{-1}(x)$ to be the unique $y \in A$ that satisfies $f(y) = x$. That is, for any $x \geq 0$, we have that $f^{-1}(x)$ is the unique $y \geq 0$ that satisfies $y^2 = x$. In other words, the square root function is really defined as an inverse! Similarly, if we define $h: \mathbb{R} \rightarrow \{x \in \mathbb{R} : x > 0\}$ by letting $h(x) = e^x$, then h is bijective, and so has an inverse. We call this inverse function the *natural log*.

Before moving on, we discuss one other abstract but powerful way to think about inverse functions. Let $f: A \rightarrow B$ be a bijective function. We then have an inverse function $f^{-1}: B \rightarrow A$. Can we compose these two functions? Notice that the codomain of f equals the domain of f^{-1} (both are B), so we can form $f^{-1} \circ f$, and this function has domain A and codomain A . That is, we have $f^{-1} \circ f: A \rightarrow A$. What is this composition? Intuitively, given any $a \in A$, we first feed the input a into f and follow the arrow to get $f(a) \in B$, and then feed this result into f^{-1} to go back across the arrow to arrive back at a . In other words, it seems that we should have $f^{-1}(f(a)) = a$ for all $a \in A$. We can also compose the functions in the other order, but notice now that $f \circ f^{-1}: B \rightarrow B$. A similar argument suggests that we should have $f(f^{-1}(b)) = b$ for all $b \in B$. We now verify these suspicions carefully.

Proposition 1.7.5. *Let $f: A \rightarrow B$ be a bijective function.*

1. For all $a \in A$, we have $f^{-1}(f(a)) = a$. Thus, $f^{-1} \circ f = id_A$.
2. For all $b \in B$, we have $f(f^{-1}(b)) = b$. Thus, $f \circ f^{-1} = id_B$.

Proof.

1. Let $a \in A$ be arbitrary. By definition of the inverse function $f^{-1}(f(a))$ is defined to be the unique $c \in A$ such that $f(c) = f(a)$. But of course a is an example of such an element of A , as we trivially have $f(a) = f(a)$. Thus, we must have $f^{-1}(f(a)) = a$.

We now show that $f^{-1} \circ f = id_A$. Let $a \in A$ be arbitrary. We have

$$\begin{aligned}
 (f^{-1} \circ f)(a) &= f^{-1}(f(a)) && \text{(by definition of composition)} \\
 &= a && \text{(from above)} \\
 &= id_A(a) && \text{(by definition of } id_A \text{)}.
 \end{aligned}$$

Since $a \in A$ was arbitrary, we conclude that $f^{-1} \circ f = id_A$.

2. Let $b \in B$ be arbitrary. By definition of the inverse function, we know that $f^{-1}(b)$ is the unique $a \in A$ that satisfies $f(a) = b$. Thus, we must have $f(f^{-1}(b)) = b$.

We now show that $f \circ f^{-1} = id_B$. Let $b \in B$ be arbitrary. We have

$$\begin{aligned}
 (f \circ f^{-1})(b) &= f(f^{-1}(b)) && \text{(by definition of composition)} \\
 &= b && \text{(from above)} \\
 &= id_B(b) && \text{(by definition of } id_B \text{)}.
 \end{aligned}$$

Since $b \in B$ was arbitrary, we conclude that $f \circ f^{-1} = id_B$.

□

1.8 Solving Equations

Suppose that we want to solve the equation $3x - 7 = 4$ in \mathbb{R} . In other words, we want to find all real numbers such that when we plug them in for x , the result is a true statement. Of course, we can naturally follow the procedures from algebra. Thus, we add 7 to both sides to conclude that $3x = 11$, and then divide both sides by 3 to conclude that $x = \frac{11}{3}$. The idea of this process is to perform the same operation to both sides in an effort to isolate x .

That all sounds good, but now suppose that we try to solve the equation $x - 2 = 1$ in \mathbb{R} . Maybe we did not think to add 2 to both sides, and instead decide to square both sides because we like quadratics. We then arrive at $(x - 2)^2 = 1^2$, so $x^2 - 4x + 4 = 1$. Now we remember to subtract 1 from both sides to conclude that $x^2 - 4x + 3 = 0$, and hence $(x - 3)(x - 1) = 0$. It follows that $x = 3$ and $x = 1$ are both solutions, right? Right? Of course, 1 is not actually a solution to $x - 2 = 1$, but it may not be obvious what went wrong. In order to understand what is happening here, we need to think about the logic of what we are doing when we mindlessly apply this “do the same thing to both sides” procedure.

Let’s start to think through the logic of solving equations carefully. We want to find all $x \in \mathbb{R}$ such that $x - 2 = 1$ is true. The idea then is to *assume* that we have a real number that gives a true statement when we plug it in, and then see what we can say about that real number. Let’s formalize this a bit by giving the real number a name. *Assume* that we have an $a \in \mathbb{R}$ such that $a - 2 = 1$ is true. Now we have an actual real number a in hand and it has the property that $a - 2$ equals the same thing as 1. Since these two real numbers $a - 2$ and 1 are the same real number, we can then conclude that $(a - 2)^2 = 1^2$ because we’ve simply squared the same real number. We now do algebra to conclude that $a^2 - 4a + 4 = 1$. Since these two real numbers

$a^2 - 4a + 4$ and 1 are the same real number, if we subtract one from both of them, then the resulting two real numbers are equal. Thus, we can conclude that $a^2 - 4a + 3 = 0$. Since $(a - 3)(a - 1) = a^2 - 4a + 3$ (by using the Distributive Law twice), we can conclude that $(a - 3)(a - 1) = 0$. Now we use the fundamental fact that if the product of two real numbers is 0, then one of the two numbers must be 0, so we can conclude that either $a - 3 = 0$ or $a - 1 = 0$. By adding the same thing to both sides of each of these, it follows that $a = 3$ or $a = 1$.

Did you follow all of that logic? If so, the *only* conclusion that we can draw is that *if* $a \in \mathbb{R}$ is such that $a - 2 = 1$, *then* it must be the case that either $a = 3$ or $a = 1$. Notice that the logic flows forward, which is why we can only conclude that an “if...then...” statement is true. The logic of the argument says nothing about the *converse* of this “if...then...”. In other words, there is no reason to believe that the statement “if either $a = 3$ or $a = 1$, then $a - 2 = 1$ ” is true, and in fact it is false! In general, all that we can say is that if we perform an operation on both sides of an equality, then elements that made the original equation true will make the resulting equation true.

Let’s formalize all of this in a definition. To talk about an “equation”, we are really talking about the two functions on each side of the equation. For example, asking for those $x \in \mathbb{R}$ such that $3x - 7 = 4$ is asking the following question: If $f: \mathbb{R} \rightarrow \mathbb{R}$ is $f(x) = 3x - 7$, and $g: \mathbb{R} \rightarrow \mathbb{R}$ is $g(x) = 4$, then determine the values $a \in \mathbb{R}$ such that $f(a) = g(a)$ is true. We give this collection of values a special name.

Definition 1.8.1. Suppose that we have two functions $f: A \rightarrow B$ and $g: A \rightarrow B$. We define the solution set of the equation $f(x) = g(x)$ to be $\{a \in A : f(a) = g(a)\}$.

Of course, the question we have to ask ourselves is *how* we determine the solution set of a given equation. In general, this is a hard problem. However, the techniques of algebra encourage us to perform the same operations on both sides of the equality with the hopes of isolating x . Suppose that we have two functions $f: A \rightarrow B$ and $g: A \rightarrow B$, and we let S be the solution set of the equation $f(x) = g(x)$. Suppose that we perform an operation on both sides of this equality, and we let T be the solution set of the resulting equation. By the arguments and examples given at the beginning of this section, we see that $S \subseteq T$, but it is possible that $T \not\subseteq S$ in general! However, if we are able to determine T , we can always go back and plug all of these values of T into the equation to determine which of them are in S .

Consider the following examples:

- If $f: \mathbb{R} \rightarrow \mathbb{R}$ is $f(x) = 3x - 7$, and $g: \mathbb{R} \rightarrow \mathbb{R}$ is $g(x) = 4$, then the solution set of the equation $f(x) = g(x)$ is $\{\frac{11}{3}\}$.
- If $f: \mathbb{R} \rightarrow \mathbb{R}$ is $f(x) = x^2$, and $g: \mathbb{R} \rightarrow \mathbb{R}$ is $g(x) = 7x - 12$, then the solution set of the equation $f(x) = g(x)$ is $\{3, 4\}$.
- If $f: \mathbb{R} \rightarrow \mathbb{R}$ is $f(x) = x$, and $g: \mathbb{R} \rightarrow \mathbb{R}$ is $g(x) = x + 1$, then the solution set of the equation $f(x) = g(x)$ is \emptyset .
- If $f: \mathbb{R} \rightarrow \mathbb{R}$ is $f(x) = \sin x$, and $g: \mathbb{R} \rightarrow \mathbb{R}$ is $g(x) = \cos x$, then the solution set of the equation $f(x) = g(x)$ is $\{\frac{\pi}{4} + \pi n : n \in \mathbb{Z}\}$.

As these examples illustrate, the solution set of an equation can have 1 element, no elements, several elements, or infinitely many elements.

We can also consider equations with many variables, such as $x + 2y = 1$ where the variables range over \mathbb{R} . We can try to generalize the above definition, but in fact this is not necessary. We can instead consider the two sides of the equation as functions of two variables. Thus, we are considering the functions $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ given by $f((x, y)) = x + 2y$ and $g: \mathbb{R}^2 \rightarrow \mathbb{R}$ given by $g((x, y)) = 1$. To determine the solution set, we just need to find those elements in \mathbb{R}^2 that make the equation true, i.e. we need to find the solution set S . Of course, one way to describe S is simply as

$$S = \{(x, y) \in \mathbb{R}^2 : x + 2y = 1\}.$$

Can we give a parametric description of this set? One way is to say that the equation $x + 2y = 1$ has the same solution set as $x = 1 - 2y$. Hence, no matter what value we plug in for y , there is a unique value of x that we can pair with it so that we obtain a solution. Thus, it appears that we can describe the solution set as

$$A = \{(1 - 2t, t) : t \in \mathbb{R}\}.$$

How can we formally argue that this is another way to describe the solution set? To show that $A = S$, we give a double containment proof:

- We first show that $A \subseteq S$ (i.e. that every element of A is a solution). Let $(a, b) \in A$ be arbitrary. By definition of A , we can fix $t \in \mathbb{R}$ with $(a, b) = (1 - 2t, t)$, so $a = 1 - 2t$ and $b = t$. Notice that

$$\begin{aligned} a + 2b &= (1 - 2t) + 2 \cdot t \\ &= 1 - 2t + 2t \\ &= 1, \end{aligned}$$

so $(a, b) \in S$. Since $(a, b) \in A$ was arbitrary, it follows that $A \subseteq S$.

- We now show that $S \subseteq A$ (i.e. that every solution is an element of A). Let $(a, b) \in S$ be arbitrary. We then have that $a + 2b = 1$. Subtracting $2b$ from both sides, we conclude that $a = 1 - 2b$. Thus, we have that $(a, b) = (1 - 2b, b)$. Since $b \in \mathbb{R}$, it follows that $(a, b) \in A$. Since $(a, b) \in S$ was arbitrary, we conclude that $S \subseteq A$.

Since we have shown that both $A \subseteq S$ and $S \subseteq A$ are true, it follows that $A = S$.

Using our new method of writing vectors vertically, notice that we can write the solution set of $x + 2y = 1$ as

$$\left\{ \begin{pmatrix} 1 - 2t \\ t \end{pmatrix} : t \in \mathbb{R} \right\},$$

which we can rewrite as

$$\left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix} + t \cdot \begin{pmatrix} -2 \\ 1 \end{pmatrix} : t \in \mathbb{R} \right\}.$$

Geometrically, we can think of this as taking the vector $(-2, 1)$ and scaling it in all possible ways, and then adding the offset $(1, 0)$ to all the vectors that result. Also, if we instead solved for y in terms of x as $y = \frac{1-x}{2}$, then a similar argument shows that we can describe the solution set as

$$\left\{ \begin{pmatrix} t \\ \frac{1-t}{2} \end{pmatrix} : t \in \mathbb{R} \right\},$$

or equivalently as

$$\left\{ \begin{pmatrix} 0 \\ 1/2 \end{pmatrix} + t \cdot \begin{pmatrix} 1 \\ -1/2 \end{pmatrix} : t \in \mathbb{R} \right\}.$$

From this perspective, we can describe the solution set as taking the vector $(1, -\frac{1}{2})$ and scaling it in all possible ways, and then adding the offset $(0, \frac{1}{2})$ to all the vectors that result. Thus, it is possible to describe one solution set in many seemingly different ways.

Notice that we also have to know the domain of the functions in question. For example, consider the equation $x - y = 0$. If we are thinking of this as two functions from \mathbb{R}^2 to \mathbb{R} , say $f(x, y) = x - y$ and $g(x, y) = 0$, then the solution set is $\{(t, t) : t \in \mathbb{R}\}$ which graphically is a line through the origin in \mathbb{R} . However, if we are thinking of this equation as two functions from \mathbb{R}^3 to \mathbb{R} , say $f(x, y, z) = x - y$ and $g(x, y, z) = 0$, then the solution set is $\{(t, t, s) : t, s \in \mathbb{R}\}$, which graphically is a plane in \mathbb{R}^3 through the origin.

Chapter 2

Spans and Linear Transformations in Two Dimensions

2.1 Intersections of Lines in \mathbb{R}^2

Recall that we can describe a line in \mathbb{R}^2 as the solution set of an equation of the form

$$ax + by = c,$$

where at least one of a or b is nonzero. Notice that if $a \neq 0$, then for every value that we plug in for y , there exists a unique value that can be plugged in for x such that the resulting pair is a solution to the equation. Geometrically, this statement says that if $a \neq 0$, then every horizontal line will pass through the graph of the given line exactly once. Similarly, if $b \neq 0$, then for every value that we plug in for x , there exists a unique value that can be plugged in for y such that the resulting pair is a solution to the equation. In other words, if $b \neq 0$, then every vertical line will pass through the graph of the given line exactly once.

Consider the following system of equations:

$$\begin{array}{rcl} x & - & 4y = -2 \\ 2x & - & 5y = 8. \end{array}$$

The solution set to a system of equations is just the set of elements (or in this case pairs of elements) that satisfy *all* of the given equations. In other words, the solution set to a system of equations is the *intersection* of the solution sets of the individual equations. Geometrically, we are asking for the points that are on both of the above lines.

Suppose that (x, y) is a solution to both of these equations. Taking the second equation and subtracting two times the first, we conclude that (x, y) must also satisfy $3y = 12$. Dividing both sides by 3, it follows that (x, y) must also satisfy $y = 4$. Plugging this into the first equation, we conclude that (x, y) must satisfy $x - 16 = -2$, and hence $x = 14$. It follows that (x, y) must equal $(14, 4)$. In other words, we've shown that $(14, 4)$ is the only possible solution. Now we can plug this the two equations to notice that

$$14 - 4 \cdot 4 = 14 - 16 = -2$$

and that

$$2 \cdot 14 - 5 \cdot 4 = 28 - 20 = 8.$$

Thus, the pair $(14, 4)$ really is a solution. It follows that the solution set is $\{(14, 4)\}$, and we verified through algebraic means that the two given lines intersect in a unique point.

Suppose instead that we had the following system:

$$\begin{array}{rcl} 5x & - & 3y = -1 \\ -2x & + & y = 0 \\ -x & - & 4y = 6. \end{array}$$

Now we are looking for the points that are on all three lines. Suppose that (x, y) is a solution to all of these equations. Taking the first equation and adding three times the second, we conclude that (x, y) must also satisfy $-x = -1$, and hence $x = 1$. Plugging this into the second equation, we conclude that (x, y) must satisfy $-2 + y = 0$, and hence $y = 2$. It follows that (x, y) must equal $(1, 2)$, so $(1, 2)$ is the only possible solution. If we plug this into the first two equations, we can verify that it does indeed satisfy them. However, notice that

$$-1 - 4 \cdot 2 = -9,$$

so $(1, 2)$ does not satisfy the third equation. Thus, we can conclude that the solution set of this system is \emptyset .

Suppose more abstractly that we have fixed numbers $a, b, c, d, j, k \in \mathbb{R}$, and we consider a system of two equations in two unknowns like the following:

$$\begin{array}{rcl} ax & + & by = j \\ cx & + & dy = k. \end{array}$$

Can we say something about the solution set in general? Notice that the solution set of each equation is a line (assuming that at least one of a or b is nonzero, and at least one of c or d is nonzero). As long as the lines are not parallel, it seems geometrically clear that they must intersect in a unique point. How can we express, in an algebraically succinct manner, that the lines are not parallel? If $b \neq 0$, then the slope of the first line equals $-\frac{a}{b}$. Similarly, if $d \neq 0$, then the slope of the second line is $-\frac{c}{d}$. Thus, if both $b \neq 0$ and $d \neq 0$, then saying that the lines are not parallel is the same as saying that $-\frac{a}{b} \neq -\frac{c}{d}$. Can we express this in a way that does not involve division, so that we have a hope of interpreting it if some values are 0? By cross-multiplying and adding ad to both sides, we can rephrase this as saying that $ad - bc \neq 0$, which now has the added benefit of making sense even when some of the values are 0. We now argue through purely algebraic means that in this case, there is always a unique solution.

Proposition 2.1.1. *Suppose that $a, b, c, d, j, k \in \mathbb{R}$ and consider the system of two equations in two unknowns x and y :*

$$\begin{array}{rcl} ax & + & by = j \\ cx & + & dy = k. \end{array}$$

Assume that $ad - bc \neq 0$. If $S \subseteq \mathbb{R}^2$ is the solution set of this system, then S has a unique element, and in fact

$$S = \left\{ \left(\frac{dj - bk}{ad - bc}, \frac{ak - cj}{ad - bc} \right) \right\}.$$

Proof. We first show that the given pair is the only possible solution. We are assuming that $a, b, c, d \in \mathbb{R}$ satisfy $ad - bc \neq 0$. Suppose now that (x, y) is an arbitrary solution to the system. Since we know that $ax + by = j$ is true, we can multiply both sides by d to conclude that

$$adx + bdy = dj.$$

Similarly, since we know that $cx + dy = k$ is true, we can multiply both sides by b to conclude that

$$bcx + bdy = bk.$$

Taking the first of these equations and subtracting the second, it follows that $(ad - bc)x = dj - bk$. Now since $ad - bc \neq 0$, we can divide by it to deduce that

$$x = \frac{dj - bk}{ad - bc}.$$

We can perform a similar process to determine what y must equal. Multiplying the first equation by c , we conclude that

$$acx + bcy = cj.$$

Multiplying the second equation by a , we also must have

$$acx + ady = ak.$$

Taking the second of these equations and subtracting the first, it follows that $(ad - bc)y = ak - cj$. Now since $ad - bc \neq 0$, we can divide by it to deduce that

$$y = \frac{ak - cj}{ad - bc}.$$

Thus, the only possible solution is the ordered pair

$$\left(\frac{dj - bk}{ad - bc}, \frac{ak - cj}{ad - bc} \right).$$

We now check that this is indeed a solution. We have

$$\begin{aligned} a \cdot \frac{dj - bk}{ad - bc} + b \cdot \frac{ak - cj}{ad - bc} &= \frac{adj - abk + bak - bcj}{ad - bc} \\ &= \frac{adj - bcj}{ad - bc} \\ &= j \cdot \frac{ad - bc}{ad - bc} \\ &= j \end{aligned}$$

and also

$$\begin{aligned} c \cdot \frac{dj - bk}{ad - bc} + d \cdot \frac{ak - cj}{ad - bc} &= \frac{cdj - cbk + dak - dcj}{ad - bc} \\ &= \frac{adk - bck}{ad - bc} \\ &= k \cdot \frac{ad - bc}{ad - bc} \\ &= k. \end{aligned}$$

Therefore, the given pair really is a solution. □

Notice that if we have a system

$$\begin{aligned} ax + by &= j \\ cx + dy &= k \end{aligned}$$

where $ad - bc = 0$, then we do not have enough information to determine the solution set, or even the number of solutions. For example, the system

$$\begin{aligned} 2x + 4y &= 1 \\ 3x + 6y &= 5 \end{aligned}$$

has no solutions, because if (x, y) is a solution, then subtracting $\frac{3}{2}$ times the first equation from the second, we would be able conclude that (x, y) is a solution to $0 = \frac{7}{2}$, which is impossible. In contrast, the system

$$\begin{array}{rcl} 2x & + & 4y = 6 \\ 3x & + & 6y = 9 \end{array}$$

has infinitely many solutions, because any pair of the form $(3 - 2t, t)$ where $t \in \mathbb{R}$ is a solution.

2.2 Vectors in \mathbb{R}^2

You have already encountered vectors in Calculus (and possibly physics) courses. In that context, vectors are typically described as objects with both magnitude and direction, and are visualized as arrows. We often think of the arrows as “free-floating” in space, so that if we pick up and move the arrow to another location, it is still the same vector. Although the intuition that comes with such a view is often very helpful, it is difficult to perform calculations with such a description. As a result, we also learn to think about vectors in terms of ordered pairs (or triples) of numbers, and you might have used notation like $\langle 3, 4 \rangle$ to represent a vector whose head is 3 units to the right and 4 units above its tail. From a psychological perspective, it is useful to distinguish between the vector $\langle 3, 4 \rangle$ and the point $(3, 4)$, but formally both are representing an object by using two numbers (where order matters and repetition is allowed). In fact, if you think about taking the vector $\langle 3, 4 \rangle$ and placing its tail at the origin, then the head of the vector will be sitting at the point $(3, 4)$. In calculus, you might have enforced a technical distinction between the point $(3, 4)$ and the so-called “position vector” $\langle 3, 4 \rangle$ anchored at the origin.

In Linear Algebra, we will avoid all of these small pedantic distinctions, and will typically use different notation. Thus, instead of writing $\langle 3, 4 \rangle$ or $(3, 4)$, we will usually write

$$\begin{pmatrix} 3 \\ 4 \end{pmatrix}.$$

Also, we typically call this object a *vector* in all circumstances. You are free to visualize the above object as either a point, a position vector, or just a packaged collection of two numbers. From our perspective, although geometric visualization or numeric interpretation can be very helpful in certain contexts, the fundamental “meaning” of a vector is not important. Formally, we are just working with an ordered pair, but writing it vertically. As a result, the collection of all vectors in the plane (i.e. vectors with 2 numbers) is just the set $\mathbb{R}^2 = \mathbb{R} \times \mathbb{R}$, and so we will view \mathbb{R}^2 as the set of all such vectors.

Although we do not care about how we interpret and visualize the elements of \mathbb{R}^2 , we do care about two fundamental operations: vector addition and scalar multiplication. Given two vectors $\vec{v}, \vec{w} \in \mathbb{R}^2$, say

$$\vec{v} = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{w} = \begin{pmatrix} b \\ d \end{pmatrix},$$

we define the sum of these vectors to be

$$\vec{v} + \vec{w} = \begin{pmatrix} a + b \\ c + d \end{pmatrix}.$$

In other words, we add vectors componentwise, by treated the first coordinates and second coordinates in isolation. Recall that this definition also matches the geometric interpretation of vector addition, where we place the tail of \vec{w} at the head of \vec{v} , and look at the vector whose tail is the tail of \vec{v} , and whose head is the head of \vec{w} .

The other core operation is scalar multiplication. In this context, we are using the word *scalar* as a synonym for “real number”. Given a vector $\vec{v} \in \mathbb{R}^2$, say

$$\vec{v} = \begin{pmatrix} a \\ b \end{pmatrix},$$

together with a scalar $r \in \mathbb{R}$, we define

$$r \cdot \vec{v} = \begin{pmatrix} ra \\ rb \end{pmatrix}.$$

That is, to multiply a vector by a number, we multiply *each* entry of the vector by the number. When viewing vectors as geometric arrows, this operation corresponds to stretching (and flipping the orientation when $r < 0$) the arrow.

The core algebraic properties of these operations are compiled in the following proposition, where we define the zero vector to be

$$\vec{0} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

Proposition 2.2.1. *Vector addition and scalar multiplication in \mathbb{R}^2 have the following properties:*

1. For all $\vec{v}, \vec{w} \in \mathbb{R}^2$, we have $\vec{v} + \vec{w} \in \mathbb{R}^2$ (closure under addition).
2. For all $\vec{v} \in \mathbb{R}^2$ and all $r \in \mathbb{R}$, we have $r \cdot \vec{v} \in \mathbb{R}^2$ (closure under scalar multiplication).
3. For all $\vec{v}, \vec{w} \in \mathbb{R}^2$, we have $\vec{v} + \vec{w} = \vec{w} + \vec{v}$ (commutativity of addition).
4. For all $\vec{u}, \vec{v}, \vec{w} \in \mathbb{R}^2$, we have $(\vec{u} + \vec{v}) + \vec{w} = \vec{u} + (\vec{v} + \vec{w})$ (associativity of addition).
5. For all $\vec{v} \in \mathbb{R}^2$, we have $\vec{v} + \vec{0} = \vec{v}$ ($\vec{0}$ is an additive identity).
6. For all $\vec{v} \in \mathbb{R}^2$, there exists $\vec{w} \in \mathbb{R}^2$ such that $\vec{v} + \vec{w} = \vec{0}$ (existence of additive inverses).
7. For all $\vec{v}, \vec{w} \in \mathbb{R}^2$ and all $r \in \mathbb{R}$, we have $r \cdot (\vec{v} + \vec{w}) = r \cdot \vec{v} + r \cdot \vec{w}$.
8. For all $\vec{v} \in \mathbb{R}^2$ and all $r, s \in \mathbb{R}$, we have $(r + s) \cdot \vec{v} = r \cdot \vec{v} + s \cdot \vec{v}$.
9. For all $\vec{v} \in \mathbb{R}^2$ and all $r, s \in \mathbb{R}$, we have $r \cdot (s \cdot \vec{v}) = (rs) \cdot \vec{v}$.
10. For all $\vec{v} \in \mathbb{R}^2$, we have $1 \cdot \vec{v} = \vec{v}$.

The proof of this proposition consists of opening up each of the vectors to reveal the two components, and then appealing to the corresponding algebraic properties of the real numbers themselves. For example, here is a proof of the third statement. Let $\vec{v}, \vec{w} \in \mathbb{R}^2$ be arbitrary. By definition of \mathbb{R}^2 , we can fix $a, b, c, d \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{w} = \begin{pmatrix} b \\ d \end{pmatrix}.$$

We have

$$\begin{aligned} \vec{v} + \vec{w} &= \begin{pmatrix} a \\ c \end{pmatrix} + \begin{pmatrix} b \\ d \end{pmatrix} \\ &= \begin{pmatrix} a + b \\ c + d \end{pmatrix} \\ &= \begin{pmatrix} b + a \\ d + c \end{pmatrix} && \text{(since } + \text{ is commutative on } \mathbb{R}) \\ &= \begin{pmatrix} b \\ d \end{pmatrix} + \begin{pmatrix} a \\ c \end{pmatrix} \\ &= \vec{w} + \vec{v}. \end{aligned}$$

All of the other proofs are completely analogous, and worthwhile exercises.

As mentioned above, Linear Algebra is built upon the two core operations of vector addition and scalar multiplication. In fact, when we will eventually generalize to settings beyond \mathbb{R}^2 in Section 4.1, all that we will care about is the the 10 algebraic properties described in the above proposition are true.

2.3 Spans

In Section 1.8, we explored how to write the solution set of $ax + by = c$ in a parametric way. Now if $c = 0$, then we are looking at a line through the origin, and we can write the solution set as the set of multiples of one vector in \mathbb{R}^2 . For example, we can describe the solution set of $4x - y = 0$ as

$$\left\{ t \cdot \begin{pmatrix} 1 \\ 4 \end{pmatrix} : t \in \mathbb{R} \right\}.$$

We now give a name to sets of this type.

Definition 2.3.1. Let $\vec{u} \in \mathbb{R}^2$. We define a subset of \mathbb{R}^2 as follows:

$$\text{Span}(\vec{u}) = \{c \cdot \vec{u} : c \in \mathbb{R}\}.$$

We call this set the span of the vector \vec{u} .

Notice that for every $\vec{u} \in \mathbb{R}^2$, we have that $\text{Span}(\vec{u}) \subseteq \mathbb{R}^2$, i.e. $\text{Span}(\vec{u})$ is a *subset* of \mathbb{R}^2 . Also, we have that $\text{Span}((0,0)) = \{(0,0)\}$, and if $\vec{u} \in \mathbb{R}^2$ is nonzero, then $\text{Span}(\vec{u})$ is an infinite set that geometrically consists of the points on the line through the origin and the point described by \vec{u} . One reason that we are interested in these sets is that they have nice *closure* properties. For example, we will show in the next result that if we take two vectors in $\text{Span}(\vec{u})$, then the sum of those vectors will also be in $\text{Span}(\vec{u})$. Think of this statement as analogous to the fact that the sum of two integers is another integer, or the product of real numbers is real number.

Proposition 2.3.2. Let $\vec{u} \in \mathbb{R}^2$ be arbitrary, and let $S = \text{Span}(\vec{u})$. We have the following:

1. $\vec{0} \in S$.
2. For all $\vec{v}_1, \vec{v}_2 \in S$, we have $\vec{v}_1 + \vec{v}_2 \in S$ (i.e. S is closed under addition).
3. For all $\vec{v} \in S$ and all $d \in \mathbb{R}$, we have $d\vec{v} \in S$ (i.e. S is closed under scalar multiplication).

Proof. We prove each of the three properties individually.

1. Notice that $\vec{0} = (0,0) = 0 \cdot \vec{u}$, so $\vec{0} \in S$.
2. Since the statement in question is a “for all” statement, we start by taking two arbitrary elements of S . Let $\vec{v}_1, \vec{v}_2 \in S$ be arbitrary. Since $\vec{v}_1 \in S$, we can fix $c_1 \in \mathbb{R}$ with $\vec{v}_1 = c_1 \cdot \vec{u}$. Similarly, since $\vec{v}_2 \in S$, we can fix $c_2 \in \mathbb{R}$ with $\vec{v}_2 = c_2 \cdot \vec{u}$. Notice that

$$\begin{aligned} \vec{v}_1 + \vec{v}_2 &= c_1 \cdot \vec{u} + c_2 \cdot \vec{u} \\ &= (c_1 + c_2) \cdot \vec{u}. \end{aligned}$$

Since $c_1 + c_2 \in \mathbb{R}$, it follows that $\vec{v}_1 + \vec{v}_2 \in S$.

3. Let $\vec{v} \in S$ and $d \in \mathbb{R}$ be arbitrary. Since $\vec{v} \in S$, we can fix $c \in \mathbb{R}$ with $\vec{v} = c\vec{u}$. Notice that

$$\begin{aligned} d\vec{v} &= d \cdot (c\vec{u}) \\ &= (dc) \cdot \vec{u}. \end{aligned}$$

Since $dc \in \mathbb{R}$, it follows that $d\vec{v} \in S$.

□

The previous result is very special, and does not work for arbitrary subsets of \mathbb{R}^2 . In fact, if we consider a set consisting of points on a line that is *not* through the origin, then it does not hold. For example, consider the set

$$A = \left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix} + t \cdot \begin{pmatrix} -2 \\ 1 \end{pmatrix} : t \in \mathbb{R} \right\}.$$

In Section 1.8, we showed that A was the solution set of the equation $x + 2y = 1$. Notice that we have the following:

- $\begin{pmatrix} 1 \\ 0 \end{pmatrix} \in A$ because $\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} + 0 \cdot \begin{pmatrix} -2 \\ 1 \end{pmatrix}$.
- $\begin{pmatrix} -1 \\ 1 \end{pmatrix} \in A$ because $\begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} + 1 \cdot \begin{pmatrix} -2 \\ 1 \end{pmatrix}$.
- $\begin{pmatrix} 0 \\ 1 \end{pmatrix} \notin A$. To see this, notice that $0 + 2 \cdot 1 = 2 \neq 1$, so the given point is not a solution to the equation $x + 2y = 1$. Alternatively, you can show that there is no $t \in \mathbb{R}$ such that

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} + t \cdot \begin{pmatrix} -2 \\ 1 \end{pmatrix}$$

by assuming that such a t exists and deriving a contradiction.

Thus, we have that

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} \in A \quad \text{and} \quad \begin{pmatrix} -1 \\ 1 \end{pmatrix} \in A,$$

but

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} + \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \notin A.$$

It follows that A is *not* closed under addition, i.e. that A does not satisfy property (2) in Proposition 2.3.2.

Recall that $\text{Span}(\vec{u})$ consists of the set of points on a line through the origin and \vec{u} (as long as $\vec{u} \neq \vec{0}$). Thus, the following result is geometrically clear. However, we work through a careful algebraic proof here, because when we generalize to higher dimensions we will not be able to rely on the geometry.

Proposition 2.3.3. *For all $\vec{u} \in \mathbb{R}^2$, we have that $\text{Span}(\vec{u}) \neq \mathbb{R}^2$.*

Proof. Let $\vec{u} \in \mathbb{R}^2$ be arbitrary. Fix $a, b \in \mathbb{R}$ with $\vec{u} = \begin{pmatrix} a \\ b \end{pmatrix}$. We have the following cases:

- *Case 1:* Suppose that $a = 0$. We claim that

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} \notin \text{Span}(\vec{u}).$$

To see this, suppose instead that $\begin{pmatrix} 1 \\ 0 \end{pmatrix} \in \text{Span}(\vec{u})$. By definition of $\text{Span}(\vec{u})$, we can then fix $c \in \mathbb{R}$ with

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} = c \cdot \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} ca \\ cb \end{pmatrix}.$$

We then have $1 = ca$ and $0 = cb$. Since $a = 0$, this first equation implies that $1 = 0$, which is a contradiction. It follows that our assumption is false, hence

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} \notin \text{Span}(\vec{u}).$$

Therefore, $\text{Span}(\vec{u}) \neq \mathbb{R}^2$.

- *Case 2:* Suppose that $a \neq 0$. We claim that

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} \notin \text{Span}(\vec{u}).$$

To see this, suppose instead that $\begin{pmatrix} 0 \\ 1 \end{pmatrix} \in \text{Span}(\vec{u})$. We can then fix $c \in \mathbb{R}$ with

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} = c \cdot \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} ca \\ cb \end{pmatrix}.$$

We then have $0 = ca$ and $1 = cb$. Since $a \neq 0$, we can divide both sides by it in the first equation to conclude that $c = 0$. Plugging this into the second equation implies that $1 = 0$, which is a contradiction. It follows that our assumption is false, hence

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} \notin \text{Span}(\vec{u}).$$

Therefore, $\text{Span}(\vec{u}) \neq \mathbb{R}^2$.

Now the two cases exhaust all possibilities, so we conclude that $\text{Span}(\vec{u}) \neq \mathbb{R}^2$ unconditionally. \square

How should we generalize the definition of $\text{Span}(\vec{u})$ to two vectors, say $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$? We can certainly consider each of the sets

$$\{c_1 \vec{u}_1 : c_1 \in \mathbb{R}\}$$

and

$$\{c_2 \vec{u}_2 : c_2 \in \mathbb{R}\}$$

separately. As we've seen in Proposition 2.3.2, each of these sets individually have nice closure properties. To combine them, we could consider taking the union of these two sets

$$U = \{c_1 \vec{u}_1 : c_1 \in \mathbb{R}\} \cup \{c_2 \vec{u}_2 : c_2 \in \mathbb{R}\}.$$

If each of \vec{u}_1 and \vec{u}_2 are nonzero, then this set consists of the points that are on at least one of the two corresponding lines. Sadly, such a set loses some of the nice properties of the sets described in the Proposition 2.3.2, because it is not closed under addition. For example, consider the case where

$$\vec{u}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

so we are looking at the set

$$U = \left\{ c_1 \begin{pmatrix} 1 \\ 0 \end{pmatrix} : c_1 \in \mathbb{R} \right\} \cup \left\{ c_2 \begin{pmatrix} 0 \\ 1 \end{pmatrix} : c_2 \in \mathbb{R} \right\} = \left\{ \begin{pmatrix} c_1 \\ 0 \end{pmatrix} : c_1 \in \mathbb{R} \right\} \cup \left\{ \begin{pmatrix} 0 \\ c_2 \end{pmatrix} : c_2 \in \mathbb{R} \right\}.$$

Notice that we have

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} \in U \quad \text{and} \quad \begin{pmatrix} 0 \\ 1 \end{pmatrix} \in U, \quad \text{but} \quad \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \notin U.$$

Hence, U is not closed under addition.

Instead of just taking the union, we want to instead think about all elements of \mathbb{R}^2 that we can “reach” from our two vectors \vec{u}_1 and \vec{u}_2 using scalings and additions. For example, starting with the vectors

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} -1 \\ 1 \end{pmatrix},$$

we can scale them and add them in the following way:

$$5 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + (-3) \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 8 \\ 2 \end{pmatrix}.$$

In other words, we can “reach” $\begin{pmatrix} 8 \\ 2 \end{pmatrix}$ from the other two vectors. Thinking in this way leads to the following fundamental definition.

Definition 2.3.4. Let $\vec{u}_1, \vec{u}_2, \vec{v} \in \mathbb{R}^2$. We say that \vec{v} is a linear combination of \vec{u}_1 and \vec{u}_2 if there exists $c_1, c_2 \in \mathbb{R}$ with $\vec{v} = c_1\vec{u}_1 + c_2\vec{u}_2$.

In other words, \vec{v} is a linear combination of \vec{u}_1 and \vec{u}_2 if we can “reach it” through some combination of scaling and adding. For example, we have that

$$\begin{pmatrix} 8 \\ 2 \end{pmatrix} \text{ is a linear combination of } \begin{pmatrix} 1 \\ 1 \end{pmatrix} \text{ and } \begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

by the computation just before the definition. We also have that

$$\begin{pmatrix} 5 \\ -4 \end{pmatrix} \text{ is a linear combination of } \begin{pmatrix} 1 \\ 2 \end{pmatrix} \text{ and } \begin{pmatrix} 4 \\ 1 \end{pmatrix}$$

because

$$\begin{pmatrix} 5 \\ -4 \end{pmatrix} = (-3) \cdot \begin{pmatrix} 1 \\ 2 \end{pmatrix} + 2 \cdot \begin{pmatrix} 4 \\ 1 \end{pmatrix},$$

and

$$\begin{pmatrix} 20 \\ -76/3 \end{pmatrix} \text{ is a linear combination of } \begin{pmatrix} 7 \\ -5 \end{pmatrix} \text{ and } \begin{pmatrix} 3 \\ 2 \end{pmatrix}$$

because

$$\begin{pmatrix} 20 \\ -76/3 \end{pmatrix} = 4 \cdot \begin{pmatrix} 7 \\ -5 \end{pmatrix} + (-8/3) \cdot \begin{pmatrix} 3 \\ 2 \end{pmatrix}.$$

Definition 2.3.5. Let $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$. We define a subset of \mathbb{R}^2 as follows:

$$\text{Span}(\vec{u}_1, \vec{u}_2) = \{c_1\vec{u}_1 + c_2\vec{u}_2 : c_1, c_2 \in \mathbb{R}\}.$$

In other words, $\text{Span}(\vec{u}_1, \vec{u}_2)$ is the set of all linear combinations of \vec{u}_1 and \vec{u}_2 . We call this set the span of the vectors \vec{u}_1 and \vec{u}_2 .

Always keep in mind that $\text{Span}(\vec{u}_1, \vec{u}_2)$ is a set of vectors. In fact, it is often an infinite set of vectors. We’re starting with just 1 or 2 vectors, looking at all of the (often infinitely) many vectors that can be reached using \vec{u}_1 and \vec{u}_2 , and then collecting them all together into one set.

Just like the case for one vector in Proposition 2.3.2, the span of two vectors is a set with nice closure properties.

Proposition 2.3.6. Let $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$ be arbitrary, and let $S = \text{Span}(\vec{u}_1, \vec{u}_2)$. We have the following.

1. $\vec{0} \in S$.
2. For all $\vec{v}_1, \vec{v}_2 \in S$, we have $\vec{v}_1 + \vec{v}_2 \in S$ (i.e. S is closed under addition).
3. For all $\vec{v} \in S$ and all $d \in \mathbb{R}$, we have $d\vec{v} \in S$ (i.e. S is closed under scalar multiplication).

Proof. We prove each of the three properties individually:

1. Notice that $\vec{0} = (0, 0) = 0 \cdot \vec{u}_1 + 0 \cdot \vec{u}_2$, so $\vec{0} \in S$.
2. Let $\vec{v}_1, \vec{v}_2 \in S$ be arbitrary. Since $\vec{v}_1 \in S$, we can fix $c_1, c_2 \in \mathbb{R}$ with $\vec{v}_1 = c_1 \vec{u}_1 + c_2 \vec{u}_2$. Similarly, since $\vec{v}_2 \in S$, we can fix $d_1, d_2 \in \mathbb{R}$ with $\vec{v}_2 = d_1 \vec{u}_1 + d_2 \vec{u}_2$. Notice that

$$\begin{aligned}\vec{v}_1 + \vec{v}_2 &= c_1 \vec{u}_1 + c_2 \vec{u}_2 + d_1 \vec{u}_1 + d_2 \vec{u}_2 \\ &= c_1 \vec{u}_1 + d_1 \vec{u}_1 + c_2 \vec{u}_2 + d_2 \vec{u}_2 \\ &= (c_1 + d_1) \vec{u}_1 + (c_2 + d_2) \vec{u}_2.\end{aligned}$$

Since $c_1 + d_1 \in \mathbb{R}$ and $c_2 + d_2 \in \mathbb{R}$, it follows that $\vec{v}_1 + \vec{v}_2 \in S$.

3. Let $\vec{v} \in S$ and $d \in \mathbb{R}$ be arbitrary. Since $\vec{v} \in S$, we can fix $c_1, c_2 \in \mathbb{R}$ with $\vec{v} = c_1 \vec{u}_1 + c_2 \vec{u}_2$. Notice that

$$\begin{aligned}d\vec{v} &= d \cdot (c_1 \vec{u}_1 + c_2 \vec{u}_2) \\ &= d \cdot (c_1 \vec{u}_1) + d \cdot (c_2 \vec{u}_2) \\ &= (dc_1) \cdot \vec{u}_1 + (dc_2) \cdot \vec{u}_2.\end{aligned}$$

Since $dc_1, dc_2 \in \mathbb{R}$, it follows that $d\vec{v} \in S$.

□

As we've seen above, we have

$$\begin{pmatrix} 8 \\ 2 \end{pmatrix} \in \text{Span} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right)$$

because

$$\begin{pmatrix} 8 \\ 2 \end{pmatrix} = 5 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + (-3) \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

We can always generate more elements of

$$\text{Span} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right)$$

by simply plugging in values for the constant scaling factors. For example, we have

$$7 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + 3 \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 4 \\ 10 \end{pmatrix}$$

so

$$\begin{pmatrix} 4 \\ 10 \end{pmatrix} \in \text{Span} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right).$$


Although it is easy to produce elements of the $\text{Span}(\vec{u}_1, \vec{u}_2)$ by simply choosing scalars and performing the computation, it appears more difficult to determine if a given element of \mathbb{R}^2 is an element of $\text{Span}(\vec{u}_1, \vec{u}_2)$. For example, is

$$\begin{pmatrix} 2 \\ 5 \end{pmatrix} \in \text{Span} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right)?$$

To answer this question, we want to know whether there exists $c_1, c_2 \in \mathbb{R}$ with

$$c_1 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + c_2 \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 2 \\ 5 \end{pmatrix}.$$

Geometrically, we are asking whether we can reach the point $\begin{pmatrix} 2 \\ 5 \end{pmatrix}$ by adding suitably scaled versions of $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} -1 \\ 1 \end{pmatrix}$. We can visualize this situation as follows:



On the left, we see the point we are trying to reach from the two given vectors. On the right, we see scaled versions of these vectors that add up to the given point. Thus, it is geometrically clear that we can indeed find such $c_1, c_2 \in \mathbb{R}$. Now let's work through the algebra to find the precise values of c_1 and c_2 . Again, we want to find $c_1, c_2 \in \mathbb{R}$ with

$$c_1 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + c_2 \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 2 \\ 5 \end{pmatrix}.$$

Working out the straightforward algebra, this is the same as finding $c_1, c_2 \in \mathbb{R}$ such that

$$\begin{pmatrix} c_1 - c_2 \\ c_1 + c_2 \end{pmatrix} = \begin{pmatrix} 2 \\ 5 \end{pmatrix},$$

which is the same as finding $c_1, c_2 \in \mathbb{R}$ such both of the following equations are true:

$$\begin{aligned} c_1 - c_2 &= 2 \\ c_1 + c_2 &= 5. \end{aligned}$$

In other words, we are trying to solve the following system of equations:

$$\begin{aligned} x - y &= 2 \\ x + y &= 5. \end{aligned}$$


By either working through the algebra, or applying Proposition 2.1.1, there is a unique solution, and it is $(\frac{7}{2}, \frac{3}{2})$. Thus, we have

$$\begin{pmatrix} 2 \\ 5 \end{pmatrix} = (7/2) \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + (3/2) \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix},$$


and hence

$$\begin{pmatrix} 2 \\ 5 \end{pmatrix} \in \text{Span} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right).$$

Looking back at the geometric picture above, we see that $\frac{7}{2}$ and $\frac{3}{2}$ are the correct scaling factors for the two vectors needed to reach the given point. To see this connection more clearly, it might help to draw a grid system based on the two vectors:



With this overlaid grid, we can see where the scalar values $\frac{7}{2}$ and $\frac{3}{2}$ arise from a geometric perspective. Moreover, as we study the grid on the left, it starts to seem like the point $\begin{pmatrix} 2 \\ 5 \end{pmatrix}$ is not important. That is, you might start to believe that we can reach *any* point in the plane using a suitable linear combination of $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} -1 \\ 1 \end{pmatrix}$. Let's explore this idea in pictures using the point $(-4, 2)$ instead of $(2, 5)$:




Using these pictures, we expect that the linear combination

$$(-1) \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + 3 \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

equals $(-4, 2)$, and this can be easily verified algebraically.

Now consider the point $(2.15, -1.35)$, illustrated on the left:



In this case, the particular values of the scalars to form the correct linear combination are not obvious, but the picture on the right surely suggests that such scalars exist. As a rough approximation, the scalar multiple of the first vector looks to be slightly less than $\frac{1}{2}$, while the scalar multiple of the second seems to be somewhere between $-\frac{3}{2}$ and -2 . In fact, it can be easily verified that

$$\begin{pmatrix} 2.15 \\ -1.35 \end{pmatrix} = (0.4) \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + (-1.75) \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

Ok, we now have some good reason to believe that every point in the plane can be written as linear combination of $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} -1 \\ 1 \end{pmatrix}$. Let's try to prove this carefully. In set-theoretic notation, we are claiming that

$$\text{Span} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right) = \mathbb{R}^2$$

is true. In order to try to prove this, we should give a double containment proof as described in Section 1.5. Now we certainly have

$$\text{Span} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right) \subseteq \mathbb{R}^2,$$

because any scalar multiple of an element of \mathbb{R}^2 is an element of \mathbb{R}^2 , and the sum of two elements of \mathbb{R}^2 is an element of the \mathbb{R}^2 .

The reverse containment is the interesting one. To prove it, let $\vec{v} \in \mathbb{R}^2$ be arbitrary. By definition of \mathbb{R}^2 , we can fix $j, k \in \mathbb{R}$ with $\vec{v} = \begin{pmatrix} j \\ k \end{pmatrix}$. We want to determine whether there exists $c_1, c_2 \in \mathbb{R}$ with

$$c_1 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + c_2 \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} j \\ k \end{pmatrix}.$$

Working through the algebra as above, we are really asking whether the system

$$\begin{array}{rclcl} x & - & y & = & j \\ x & + & y & = & k \end{array}$$

has a solution. Notice that the existence of a solution follows immediately from Proposition 2.1.1 because $1 \cdot 1 - 1 \cdot (-1) = 2 \neq 0$. Thus, there do exist $c_1, c_2 \in \mathbb{R}$ with

$$c_1 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + c_2 \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} j \\ k \end{pmatrix} = \vec{v}.$$

Since $\vec{v} \in \mathbb{R}^2$ was arbitrary, it follows that

$$\mathbb{R}^2 \subseteq \text{Span} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right).$$

Combining this with the other containment, we conclude that

$$\text{Span} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right) = \mathbb{R}^2.$$

In contrast, consider the set

$$\text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ 4 \end{pmatrix} \right).$$

Notice that the two vectors that we are looking at here are multiples of each other, and hence lie on the same line through the origin. As a result, it appears from the geometry that we can not move off of this line using only scaling and addition. Let's think through this algebraically using our set theoretic notation. We have

$$\text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ 4 \end{pmatrix} \right) = \left\{ c_1 \begin{pmatrix} 1 \\ 2 \end{pmatrix} + c_2 \begin{pmatrix} 2 \\ 4 \end{pmatrix} : c_1, c_2 \in \mathbb{R} \right\}.$$

Now for any $c_1, c_2 \in \mathbb{R}$, we have

$$\begin{aligned} c_1 \begin{pmatrix} 1 \\ 2 \end{pmatrix} + c_2 \begin{pmatrix} 2 \\ 4 \end{pmatrix} &= \begin{pmatrix} c_1 + 2c_2 \\ 2c_1 + 4c_2 \end{pmatrix} \\ &= \begin{pmatrix} c_1 + 2c_2 \\ 2 \cdot (c_1 + 2c_2) \end{pmatrix} \\ &= (c_1 + 2c_2) \begin{pmatrix} 1 \\ 2 \end{pmatrix}, \end{aligned}$$

so since $c_1 + 2c_2 \in \mathbb{R}$, we conclude that

$$c_1 \begin{pmatrix} 1 \\ 2 \end{pmatrix} + c_2 \begin{pmatrix} 2 \\ 4 \end{pmatrix} \in \text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix} \right).$$

Since this was true for any $c_1, c_2 \in \mathbb{R}$, it follows that

$$\text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ 4 \end{pmatrix} \right) \subseteq \text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix} \right).$$

Now we also have

$$\text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix} \right) \subseteq \text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ 4 \end{pmatrix} \right)$$

(think about why), so we can conclude that

$$\text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ 4 \end{pmatrix} \right) = \text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix} \right).$$

The next two propositions generalize these ideas.

Proposition 2.3.7. *For all $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$, we have*

$$\text{Span}(\vec{u}_1) \subseteq \text{Span}(\vec{u}_1, \vec{u}_2).$$

Proof. Let $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$ be arbitrary. Consider an arbitrary $\vec{v} \in \text{Span}(\vec{u}_1)$. By definition, we can fix $c \in \mathbb{R}$ with $\vec{v} = c\vec{u}_1$. We then have

$$\begin{aligned} \vec{v} &= c\vec{u}_1 \\ &= c\vec{u}_1 + \vec{0} \\ &= c\vec{u}_1 + 0\vec{u}_2. \end{aligned}$$

Since $c \in \mathbb{R}$ and $0 \in \mathbb{R}$, we can conclude that $\vec{v} \in \text{Span}(\vec{u}_1, \vec{u}_2)$. As $\vec{v} \in \text{Span}(\vec{u}_1)$ was arbitrary, it follows that $\text{Span}(\vec{u}_1) \subseteq \text{Span}(\vec{u}_1, \vec{u}_2)$. \square

Proposition 2.3.8. *Let $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$. The following are equivalent, i.e. if one is true, then so is the other (and hence if one is false, then so is the other):*

1. $\text{Span}(\vec{u}_1, \vec{u}_2) = \text{Span}(\vec{u}_1)$.
2. $\vec{u}_2 \in \text{Span}(\vec{u}_1)$.

Proof. Exercise. \square

We also have the following simple (but useful) fact because vector addition is commutative. See if you can write up a careful proof.

Proposition 2.3.9. *For all $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$, we have $\text{Span}(\vec{u}_1, \vec{u}_2) = \text{Span}(\vec{u}_2, \vec{u}_1)$.*

Proof. Exercise. \square

With all of this work in hand, we now characterize the pairs of vectors whose span is all of \mathbb{R}^2 in several different ways.

Theorem 2.3.10. *Let $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$, say*

$$\vec{u}_1 = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} b \\ d \end{pmatrix}.$$

The following are equivalent, i.e. if any one of the statements is true, then so are all of the others (and hence if any one of the statements is false, then so are all of the others).

1. For all $\vec{v} \in \mathbb{R}^2$, there exist $r_1, r_2 \in \mathbb{R}$ with $\vec{v} = r_1\vec{u}_1 + r_2\vec{u}_2$, i.e. $\text{Span}(\vec{u}_1, \vec{u}_2) = \mathbb{R}^2$.
2. $\vec{u}_1 \neq \vec{0}$ and $\vec{u}_2 \notin \text{Span}(\vec{u}_1)$.
3. $\vec{u}_2 \neq \vec{0}$ and $\vec{u}_1 \notin \text{Span}(\vec{u}_2)$.
4. $ad - bc \neq 0$.
5. For all $\vec{v} \in \mathbb{R}^2$, there exists a unique pair of numbers $r_1, r_2 \in \mathbb{R}$ with $\vec{v} = r_1\vec{u}_1 + r_2\vec{u}_2$.

Proof. In order to prove that all five of these statements are equivalent, we go around in a circle. In other words, we prove that (1) implies (2) (i.e. if (1) is true, then (2) is true), then prove that (2) implies (3), then (3) implies (4), then (4) implies (5), and finally that (5) implies (1).

- (1) implies (2): We instead prove the contrapositive that **Not**(2) implies **Not**(1). Suppose then that **Not**(2) is true, i.e. that (2) is false. We then have two possibilities:

- *Case 1:* Suppose that $\vec{u}_1 = \vec{0}$. In this case, we have

$$\text{Span}(\vec{u}_1, \vec{u}_2) = \text{Span}(\vec{0}, \vec{u}_2) = \text{Span}(\vec{u}_2, \vec{0})$$

by Proposition 2.3.9. Now $\vec{0} \in \text{Span}(\vec{u}_2)$ by Proposition 2.3.2, so

$$\text{Span}(\vec{u}_2, \vec{0}) = \text{Span}(\vec{u}_2)$$

by Proposition 2.3.8. Thus, we have

$$\text{Span}(\vec{u}_1, \vec{u}_2) = \text{Span}(\vec{u}_2).$$

Since $\text{Span}(\vec{u}_2) \neq \mathbb{R}^2$ by Proposition 2.3.3, we conclude that $\text{Span}(\vec{u}_1, \vec{u}_2) \neq \mathbb{R}^2$.

- *Case 2:* Suppose that $\vec{u}_2 \in \text{Span}(\vec{u}_1)$. We then have

$$\text{Span}(\vec{u}_1, \vec{u}_2) = \text{Span}(\vec{u}_1)$$

by Proposition 2.3.8. Since $\text{Span}(\vec{u}_1) \neq \mathbb{R}^2$ by Proposition 2.3.3, we conclude that $\text{Span}(\vec{u}_1, \vec{u}_2) \neq \mathbb{R}^2$.

Thus, in either case, we see that **Not**(1) is true.

- (2) implies (3): We instead prove the contrapositive that **Not**(3) implies **Not**(2). Suppose then that **Not**(3) is true, i.e. that (3) is false. We then have two possibilities:

- *Case 1:* Suppose that $\vec{u}_2 = \vec{0}$. In this case, we have $\vec{u}_2 \in \text{Span}(\vec{u}_1)$ by Proposition 2.3.2, so (2) is false.
- *Case 2:* Suppose that $\vec{u}_1 \in \text{Span}(\vec{u}_2)$. By definition, we can then fix $r \in \mathbb{R}$ with $\vec{u}_1 = r\vec{u}_2$. If $\vec{u}_1 = \vec{0}$, then (2) is false, and we are done. Suppose then that $\vec{u}_1 \neq \vec{0}$. Since $\vec{u}_1 = r\vec{u}_2$, it must be the case that $r \neq 0$. Thus, we can multiply both sides of $\vec{u}_1 = r\vec{u}_2$ by $\frac{1}{r}$ to conclude that $\vec{u}_2 = \frac{1}{r} \cdot \vec{u}_1$. Since $\frac{1}{r} \in \mathbb{R}$, it follows that $\vec{u}_2 \in \text{Span}(\vec{u}_1)$, and hence (2) is false.

Thus, in either case, we see that **Not**(2) is true.

- (3) implies (4): We instead prove the contrapositive that **Not**(4) implies **Not**(3). Suppose then that **Not**(4) is true, so then $ad - bc = 0$. Now if $\vec{u}_2 = \vec{0}$, then **Not**(3) is true and we are done. Assume then that $\vec{u}_2 \neq \vec{0}$. We now have two cases:

- *Case 1:* Suppose that $b = 0$. Notice that we must have $d \neq 0$ because $\vec{u}_2 \neq \vec{0}$. Now since $ad - bc = 0$, we conclude that $ad = 0$, and dividing both sides by the nonzero d , we conclude that $a = 0$. Now notice that

$$\begin{aligned} \frac{c}{d} \cdot \vec{u}_2 &= \frac{c}{d} \cdot \begin{pmatrix} b \\ d \end{pmatrix} \\ &= \frac{c}{d} \cdot \begin{pmatrix} 0 \\ d \end{pmatrix} \\ &= \begin{pmatrix} 0 \\ c \end{pmatrix} \\ &= \begin{pmatrix} a \\ c \end{pmatrix} \\ &= \vec{u}_1. \end{aligned}$$

Since $\frac{c}{d} \in \mathbb{R}$, we conclude that $\vec{u}_1 \in \text{Span}(\vec{u}_2)$, and hence **Not**(3) is true.

- *Case 2:* Suppose that $b \neq 0$. Since $ad - bc = 0$, we must have that $ad = bc$. Now $b \neq 0$, so we can divide both sides by it to conclude that $\frac{ad}{b} = c$. Now notice that

$$\begin{aligned} \frac{a}{b} \cdot \vec{u}_2 &= \frac{a}{b} \cdot \begin{pmatrix} b \\ d \end{pmatrix} \\ &= \begin{pmatrix} a \\ \frac{ad}{b} \end{pmatrix} \\ &= \begin{pmatrix} a \\ c \end{pmatrix} && \text{(from above)} \\ &= \vec{u}_1. \end{aligned}$$

Since $\frac{a}{b} \in \mathbb{R}$, we conclude that $\vec{u}_1 \in \text{Span}(\vec{u}_2)$, and hence **Not**(3) is true.

Since the two cases exhaust all possibilities, we conclude that **Not**(3) is true unconditionally, which completes the proof.

- (4) implies (5): Suppose that (4) is true, so we are assuming that $ad - bc \neq 0$. Let $\vec{v} \in \mathbb{R}^2$ be arbitrary, and fix $j, k \in \mathbb{R}$ with $\vec{v} = (j, k)$. Since $ad - bc \neq 0$, we can use Proposition 2.1.1 to conclude that the system

$$\begin{aligned} ax + by &= j \\ cx + dy &= k \end{aligned}$$

has a unique solution. From this, it follows that there is a unique solution to

$$x \begin{pmatrix} a \\ c \end{pmatrix} + y \begin{pmatrix} b \\ d \end{pmatrix} = \begin{pmatrix} j \\ k \end{pmatrix},$$

which is the same as saying that there exists a unique pair of numbers $r_1, r_2 \in \mathbb{R}$ with $\vec{v} = r_1 \vec{u}_1 + r_2 \vec{u}_2$. Since $\vec{v} \in \mathbb{R}^2$ was arbitrary, we conclude that (5) is true.

- (5) implies (1): This is immediate because (5) is just a strengthening of (1) (if there exists a unique pair of numbers, then there certainly exists a pair of numbers).

□

Suppose then that we have $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$ with $\text{Span}(\vec{u}_1, \vec{u}_2) = \mathbb{R}^2$. By definition, we can express every $\vec{v} \in \mathbb{R}^2$ as a linear combination of \vec{u}_1 and \vec{u}_2 . However, by the theorem we just proved, we now know the stronger fact that no matter which $\vec{v} \in \mathbb{R}^2$ we have, we can always find *unique* scaling factors c_1 and c_2 in \mathbb{R} to apply to \vec{u}_1 and \vec{u}_2 so that the corresponding linear combination equals \vec{v} . Since this stronger property of uniqueness is fundamental, we introduce new, more concise, terminology to describe this situation.

Definition 2.3.11. Let $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$. We say that the ordered pair (\vec{u}_1, \vec{u}_2) is a basis for \mathbb{R}^2 if $\text{Span}(\vec{u}_1, \vec{u}_2) = \mathbb{R}^2$.

We will often use lowercase Greek letters, like α or β , to denote a basis. For example, suppose that

$$\vec{u}_1 = \begin{pmatrix} 5 \\ 1 \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} -6 \\ -3 \end{pmatrix}.$$

Notice that $5 \cdot (-3) - 1 \cdot (-6) = -9$, so $\text{Span}(\vec{u}_1, \vec{u}_2) = \mathbb{R}^2$ by Theorem 2.3.10. If we let α denote the ordered pair (\vec{u}_1, \vec{u}_2) , i.e. let

$$\alpha = \left(\begin{pmatrix} 5 \\ 1 \end{pmatrix}, \begin{pmatrix} -6 \\ -3 \end{pmatrix} \right),$$

then α is a basis of \mathbb{R}^2 .

As we will see, bases will play an essential role throughout our study. What makes them so special? Intuitively, our two fundamental operations on vectors are addition and scalar multiplication. If we have a basis α , then we can reach all elements of \mathbb{R}^2 from the vectors in α by using only these two operations. Thus, if we have a function $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, and we know (for some reason) how f interacts with addition and scalar multiplication, then we can use knowledge of how f behaves on a basis α to understand how f behaves on other inputs. With this idea in mind, we turn now to special functions that behave well with respect to addition and scalar multiplication.

2.4 Linear Transformations of \mathbb{R}^2

Back in Section 1.1, we discussed a function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by letting

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} x - y \\ x + y \end{pmatrix}.$$

Computing the values on some of the points, we obtain the following:

- $T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$
- $T\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$
- $T\left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} 0 \\ 2 \end{pmatrix}.$

As mentioned in Section 1.1, by computing the values at enough points, it starts to look like T might be rotating the plane by 45° counterclockwise around the origin, and also scaling the plane by a factor of $\sqrt{2}$.

Before delving into this idea further, let's take a step back and think about a somewhat philosophical question. In this example, we computed a few points and used those to make a guess about what would happen on other points. Is it reasonable to believe that we can make global predictions based on a few points? Given a function $f: \mathbb{R} \rightarrow \mathbb{R}$, if we know the values of f at several points, can we use those values to predict what f will be at other points? Without some sort of additional assumption, it seems impossible to predict the values of f on other points, because f could be piecewise-defined or worse. Now in the natural sciences and in statistics, you often extrapolate from certain data points to a general function that fits the points well, at least up to small errors. The general philosophical assumptions underlying such an extrapolation is that the function that we are searching for is *not* random or erratic, but somehow will be constrained to behave in reasonable ways. In these areas, we tend to follow this approach either because we believe in the fundamentally regularity and simplicity of nature, or because we need to make some simplifying assumptions to make any sort of reasonable predictions.

However, in mathematics, we want to be precise and codify what we mean by saying something like “the function f is not random or erratic” or “the function f has nice structural properties” so that we can understand the assumptions behind these models. We could make the assumption that “ f is a polynomial”, which will restrict the possibilities. In calculus, you might make the assumption that “ f is continuous” to avoid large jumps in the function, or even assume the more restrictive assumption that “ f is differentiable everywhere and $-3 \leq f'(x) \leq 3$ for all $x \in \mathbb{R}$ ” so that the function will be smooth and neither increase nor decrease too rapidly.

Returning to our function T above, it turns out that it has a couple of really interesting structural properties that restrict its behavior. Let $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ be arbitrary, and fix $x_1, y_1, x_2, y_2 \in \mathbb{R}$ with

$$\vec{v}_1 = \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} \quad \text{and} \quad \vec{v}_2 = \begin{pmatrix} x_2 \\ y_2 \end{pmatrix}.$$

Notice that

$$\begin{aligned}
 T(\vec{v}_1 + \vec{v}_2) &= T\left(\begin{pmatrix} x_1 \\ y_1 \end{pmatrix} + \begin{pmatrix} x_2 \\ y_2 \end{pmatrix}\right) \\
 &= T\left(\begin{pmatrix} x_1 + x_2 \\ y_1 + y_2 \end{pmatrix}\right) \\
 &= \begin{pmatrix} x_1 + x_2 - (y_1 + y_2) \\ x_1 + x_2 + y_1 + y_2 \end{pmatrix} \\
 &= \begin{pmatrix} x_1 - y_1 + x_2 - y_2 \\ x_1 + x_2 + y_1 + y_2 \end{pmatrix} \\
 &= \begin{pmatrix} x_1 - y_1 \\ x_1 + y_1 \end{pmatrix} + \begin{pmatrix} x_2 - y_2 \\ x_2 + y_2 \end{pmatrix} \\
 &= T\left(\begin{pmatrix} x_1 \\ y_1 \end{pmatrix}\right) + T\left(\begin{pmatrix} x_2 \\ y_2 \end{pmatrix}\right) \\
 &= T(\vec{v}_1) + T(\vec{v}_2).
 \end{aligned}$$

In other words, we can interchange the order between adding vectors and applying T . That is, we can either add two vectors first, and then apply the function T , or we can apply the function T to each vector, and then add the results. This is a remarkable property! Very (very) few functions $f: \mathbb{R} \rightarrow \mathbb{R}$ have the property that $f(a + b) = f(a) + f(b)$ for all $a, b \in \mathbb{R}$. For example, in general, $(a + b)^2$ does not equal $a^2 + b^2$, and $\sin(a + b)$ does not equal $\sin a + \sin b$.

However, if we think geometrically about a rotation around the origin by a fixed angle θ , it appears that such a rotation also has this property. Given two vectors \vec{v}_1 and \vec{v}_2 in \mathbb{R}^2 , it looks like the following two operations result in the same vector:

- First add the vectors to form $\vec{v}_1 + \vec{v}_2$, and then rotate the result by θ .
- First rotate \vec{v}_1 and \vec{v}_2 each by θ , and then add the results.

Beyond addition, the other fundamental operation on vectors is scalar multiplication. Let $\vec{v} \in \mathbb{R}^2$ and $c \in \mathbb{R}$ be arbitrary, and fix $x, y \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

Notice that

$$\begin{aligned}
 T(c \cdot \vec{v}) &= T\left(c \cdot \begin{pmatrix} x \\ y \end{pmatrix}\right) \\
 &= T\left(\begin{pmatrix} cx \\ cy \end{pmatrix}\right) \\
 &= \begin{pmatrix} cx - cy \\ cx + cy \end{pmatrix} \\
 &= \begin{pmatrix} c \cdot (x - y) \\ c \cdot (x + y) \end{pmatrix} \\
 &= c \cdot \begin{pmatrix} x - y \\ x + y \end{pmatrix} \\
 &= c \cdot T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) \\
 &= c \cdot T(\vec{v}).
 \end{aligned}$$

Thus, we can interchange the order between scalar multiplication and applying T . Just like for addition, rotation behaves well with respect to scalar multiplication. More precisely, given a vector $\vec{v} \in \mathbb{R}^2$ and a number $c \in \mathbb{R}$, it looks like the following two operations results in the same vector:

- First scale \vec{v} by c to form $c \cdot \vec{v}$, and then rotate the result by θ .
- First rotate \vec{v} by θ , and then scale the result by c .

Functions from \mathbb{R}^2 to \mathbb{R}^2 that respect the two basic operations of vector addition and scalar multiplication are given a special name, and will be a fundamental object of our study.

Definition 2.4.1. A linear transformation from \mathbb{R}^2 to \mathbb{R}^2 is a function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with the following two properties:

1. $T(\vec{v}_1 + \vec{v}_2) = T(\vec{v}_1) + T(\vec{v}_2)$ for all $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ (i.e. T preserves addition).
2. $T(c \cdot \vec{v}) = c \cdot T(\vec{v})$ for all $\vec{v} \in \mathbb{R}^2$ and $c \in \mathbb{R}$ (i.e. T preserves scalar multiplication).

Consider the function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} x^2 + y \\ 5x \end{pmatrix}.$$

Is T a linear transformation? If we try to compute whether T preserves addition, we might notice that

$$\begin{aligned} T\left(\begin{pmatrix} x_1 \\ y_1 \end{pmatrix} + \begin{pmatrix} x_2 \\ y_2 \end{pmatrix}\right) &= T\left(\begin{pmatrix} x_1 + x_2 \\ y_1 + y_2 \end{pmatrix}\right) \\ &= \begin{pmatrix} (x_1 + x_2)^2 + (y_1 + y_2) \\ 5 \cdot (x_1 + x_2) \end{pmatrix} \\ &= \begin{pmatrix} x_1^2 + 2x_1x_2 + x_2^2 + y_1 + y_2 \\ 5x_1 + 5x_2 \end{pmatrix}, \end{aligned}$$

while

$$\begin{aligned} T\left(\begin{pmatrix} x_1 \\ y_1 \end{pmatrix}\right) + T\left(\begin{pmatrix} x_2 \\ y_2 \end{pmatrix}\right) &= \begin{pmatrix} x_1^2 + y_1 \\ 5x_1 \end{pmatrix} + \begin{pmatrix} x_2^2 + y_2 \\ 5x_2 \end{pmatrix} \\ &= \begin{pmatrix} x_1^2 + x_2^2 + y_1 + y_2 \\ 5x_1 + 5x_2 \end{pmatrix}. \end{aligned}$$

Since one of these has an extra $2x_1x_2$ in the first coordinate, it appears that these two outputs are not the same. However, looks can be deceiving, and occasionally different looking formulas can produce the same output on all inputs (like the fact that $\sin(2x)$ equals $2\sin x \cos x$ for all $x \in \mathbb{R}$). In order to argue that T is not a linear transformation, we should instead give a specific counterexample. Looking at the above computation, we should probably choose two vectors whose first coordinate is nonzero. Notice that

$$\begin{aligned} T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} + \begin{pmatrix} 2 \\ 0 \end{pmatrix}\right) &= T\left(\begin{pmatrix} 3 \\ 0 \end{pmatrix}\right) \\ &= \begin{pmatrix} 3^2 + 0 \\ 5 \cdot (3 + 0) \end{pmatrix} \\ &= \begin{pmatrix} 9 \\ 15 \end{pmatrix}, \end{aligned}$$

while

$$\begin{aligned} T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) + T\left(\begin{pmatrix} 2 \\ 0 \end{pmatrix}\right) &= \begin{pmatrix} 1^2 + 0 \\ 5 \cdot 1 \end{pmatrix} + \begin{pmatrix} 2^2 + 0 \\ 5 \cdot 2 \end{pmatrix} \\ &= \begin{pmatrix} 1 \\ 5 \end{pmatrix} + \begin{pmatrix} 4 \\ 10 \end{pmatrix} \\ &= \begin{pmatrix} 5 \\ 15 \end{pmatrix}. \end{aligned}$$

Therefore, we have

$$T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} + \begin{pmatrix} 2 \\ 0 \end{pmatrix}\right) \neq T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) + T\left(\begin{pmatrix} 2 \\ 0 \end{pmatrix}\right),$$

so T is not a linear transformation.

In contrast, let's take a look at the function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} x + 3y \\ y \end{pmatrix}.$$

It turns out that T is a linear transformation (this fact will follow from Proposition 2.4.3 below, but it's worth working through the argument carefully on your own). Let's examine how this linear transformation behaves. Trying out a few points on the x -axis, we see that

$$T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad \text{and} \quad T\left(\begin{pmatrix} 2 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 2 \\ 0 \end{pmatrix}$$

In fact, for any $x \in \mathbb{R}$, we have

$$\begin{aligned} T\left(\begin{pmatrix} x \\ 0 \end{pmatrix}\right) &= \begin{pmatrix} x + 3 \cdot 0 \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} x \\ 0 \end{pmatrix}, \end{aligned}$$

so every point on the x -axis is fixed. Looking at the line $y = 1$, we see that for all $x \in \mathbb{R}$, we have

$$\begin{aligned} T\left(\begin{pmatrix} x \\ 1 \end{pmatrix}\right) &= \begin{pmatrix} x + 3 \cdot 1 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} x + 3 \\ 1 \end{pmatrix}, \end{aligned}$$

so the line $y = 1$ is shifted to the right by 3. A similar computation shows that the line $y = 2$ is shifted to the right by 6, the line $y = \frac{1}{3}$ is shifted to the right by 1, and the line $y = -1$ is shifted to the left by 3. Examining the formula again, it becomes clear geometrically that for a fixed c , the line $y = c$ is shifted to the right by $3c$ when c is positive, and is shifted to the left by $3c$ when c is negative. Thus, T takes the plane and “shifts” it along horizontal lines, with the shift becoming larger as we move away from the x -axis. A linear transformation T of this type is called a *shear transformation*.

Proposition 2.4.2. *Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. We have the following:*

1. $T(\vec{0}) = \vec{0}$.
2. $T(-\vec{v}) = -T(\vec{v})$ for all $\vec{v} \in \mathbb{R}^2$.

3. $T(c_1\vec{v}_1 + c_2\vec{v}_2) = c_1 \cdot T(\vec{v}_1) + c_2 \cdot T(\vec{v}_2)$ for all $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ and all $c_1, c_2 \in \mathbb{R}$.

Proof.

1. Notice that

$$\begin{aligned} T(\vec{0}) &= T(\vec{0} + \vec{0}) \\ &= T(\vec{0}) + T(\vec{0}) \end{aligned} \quad (\text{since } T \text{ preserves addition}),$$

so $T(\vec{0}) = T(\vec{0}) + T(\vec{0})$. Subtracting the vector $T(\vec{0})$ from both sides, we conclude that

$$\vec{0} = T(\vec{0}),$$

which completes the proof of 1.

2. Let $\vec{v} \in \mathbb{R}^2$ be arbitrary. Notice that

$$\begin{aligned} T(-\vec{v}) &= T((-1) \cdot \vec{v}) \\ &= (-1) \cdot T(\vec{v}) \quad (\text{since } T \text{ preserves scalar multiplication}) \\ &= -T(\vec{v}). \end{aligned}$$

Since $\vec{v} \in \mathbb{R}^2$ was arbitrary, this completes the proof of (2).

3. Let $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ and $c_1, c_2 \in \mathbb{R}$ be arbitrary. Since $c_1\vec{v}_1$ and $c_2\vec{v}_2$ are both elements of \mathbb{R}^2 , we have

$$\begin{aligned} T(c_1\vec{v}_1 + c_2\vec{v}_2) &= T(c_1\vec{v}_1) + T(c_2\vec{v}_2) \quad (\text{since } T \text{ preserves addition}) \\ &= c_1 \cdot T(\vec{v}_1) + c_2 \cdot T(\vec{v}_2) \quad (\text{since } T \text{ preserves scalar multiplication}). \end{aligned}$$

Since $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ and $c_1, c_2 \in \mathbb{R}$ were arbitrary, this completes the proof of (3).

□

Let's return to our above example of a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} x - y \\ x + y \end{pmatrix}.$$

Notice that we can also write T as

$$\begin{aligned} T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) &= \begin{pmatrix} x \\ x \end{pmatrix} + \begin{pmatrix} -y \\ y \end{pmatrix} \\ &= x \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + y \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix}. \end{aligned}$$

In other words, we can view T as follows. For any $x, y \in \mathbb{R}$:

T sends $\begin{pmatrix} x \\ y \end{pmatrix}$ to the linear combination of $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} -1 \\ 1 \end{pmatrix}$ given by using the scalars x and y .

Hence, in this case, we see that our linear transformation T can be viewed through the lens of linear combinations. Moreover, T sends (x, y) to the point obtained by making a grid system using the vectors $(1, 1)$ and $(-1, 1)$, and then finding the point that is x units along $(1, 1)$ and y units along $(-1, 1)$.

Similarly, if we look at the linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} x + 3y \\ y \end{pmatrix},$$

then we can also write T as

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = x \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} + y \cdot \begin{pmatrix} 3 \\ 1 \end{pmatrix}.$$

Thus, we can view T as taking a vector (x, y) and using its values as coordinates within the grid system formed by $(1, 0)$ and $(3, 1)$.

Proposition 2.4.3. *Let $a, b, c, d \in \mathbb{R}$. Define a function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by*

$$\begin{aligned} T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) &= \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix} \\ &= x \cdot \begin{pmatrix} a \\ c \end{pmatrix} + y \cdot \begin{pmatrix} b \\ d \end{pmatrix}. \end{aligned}$$

We then have that T is a linear transformation.

Proof. To prove that T is a linear transformation, we need to check that T preserves both addition and scalar multiplication:

- We first check that T preserves addition. Let $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ be arbitrary. Fix $x_1, y_1, x_2, y_2 \in \mathbb{R}$ with

$$\vec{v}_1 = \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} \quad \text{and} \quad \vec{v}_2 = \begin{pmatrix} x_2 \\ y_2 \end{pmatrix}.$$

We have

$$\begin{aligned} T(\vec{v}_1 + \vec{v}_2) &= T\left(\begin{pmatrix} x_1 \\ y_1 \end{pmatrix} + \begin{pmatrix} x_2 \\ y_2 \end{pmatrix}\right) \\ &= T\left(\begin{pmatrix} x_1 + x_2 \\ y_1 + y_2 \end{pmatrix}\right) \\ &= \begin{pmatrix} a(x_1 + x_2) + b(y_1 + y_2) \\ c(x_1 + x_2) + d(y_1 + y_2) \end{pmatrix} \\ &= \begin{pmatrix} ax_1 + ax_2 + by_1 + by_2 \\ cx_1 + cx_2 + dy_1 + dy_2 \end{pmatrix} \\ &= \begin{pmatrix} ax_1 + by_1 + ax_2 + by_2 \\ cx_1 + dy_1 + cx_2 + dy_2 \end{pmatrix} \\ &= \begin{pmatrix} ax_1 + by_1 \\ cx_1 + dy_1 \end{pmatrix} + \begin{pmatrix} ax_2 + by_2 \\ cx_2 + dy_2 \end{pmatrix} \\ &= T\left(\begin{pmatrix} x_1 \\ y_1 \end{pmatrix}\right) + T\left(\begin{pmatrix} x_2 \\ y_2 \end{pmatrix}\right) \\ &= T(\vec{v}_1) + T(\vec{v}_2). \end{aligned}$$

Since $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ were arbitrary, it follows that T preserves addition.

- We now check that T preserves scalar multiplication. Let $\vec{v} \in \mathbb{R}^2$ and $r \in \mathbb{R}$ be arbitrary. Fix $x, y \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

We have

$$\begin{aligned}
 T(r \cdot \vec{v}) &= T\left(r \cdot \begin{pmatrix} x \\ y \end{pmatrix}\right) \\
 &= T\left(\begin{pmatrix} rx \\ ry \end{pmatrix}\right) \\
 &= \begin{pmatrix} a(rx) + b(ry) \\ c(rx) + d(ry) \end{pmatrix} \\
 &= \begin{pmatrix} rax + rby \\ rcx + rdy \end{pmatrix} \\
 &= r \cdot \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix} \\
 &= r \cdot T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) \\
 &= r \cdot T(\vec{v}).
 \end{aligned}$$

Since $\vec{v} \in \mathbb{R}^2$ and $r \in \mathbb{R}$ were arbitrary, it follows that T preserves scalar multiplication.

We've shown that T preserves both addition and scalar multiplication, so T is a linear transformation. \square

With this result in hand, we now have an infinite stock of examples of linear transformations from \mathbb{R}^2 to \mathbb{R}^2 .

Suppose that we have a function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, and assume that we know that T is a linear transformation with both

$$T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 3 \\ 2 \end{pmatrix} \quad \text{and} \quad T\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} -1 \\ 7 \end{pmatrix}.$$

Can we compute other values of T ? It turns out that this is possible for *all* other inputs of T . For example, to determine

$$T\left(\begin{pmatrix} 8 \\ 3 \end{pmatrix}\right)$$

we can first notice that

$$\begin{pmatrix} 8 \\ 3 \end{pmatrix} = 8 \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} + 3 \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

so we have

$$\begin{aligned}
 T\left(\begin{pmatrix} 8 \\ 3 \end{pmatrix}\right) &= T\left(8 \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} + 3 \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) \\
 &= 8 \cdot T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) + 3 \cdot T\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) && \text{(by Proposition 2.4.2)} \\
 &= 8 \cdot \begin{pmatrix} 3 \\ 2 \end{pmatrix} + 3 \cdot \begin{pmatrix} -1 \\ 7 \end{pmatrix} \\
 &= \begin{pmatrix} 21 \\ 37 \end{pmatrix}.
 \end{aligned}$$

There was nothing special about the vector $\begin{pmatrix} 8 \\ 3 \end{pmatrix}$ here. Using the same technique, we can determine the value of $T\left(\begin{pmatrix} a \\ b \end{pmatrix}\right)$ for any $a, b \in \mathbb{R}$ by simply noticing that

$$\begin{pmatrix} a \\ b \end{pmatrix} = a \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} + b \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

and using the fact that T is a linear transformation.

In the previous example, we assumed that we knew the values of T on the two basic vectors $(1, 0)$ and $(0, 1)$. Suppose instead that we know that values of T on different inputs. For example, suppose that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a linear transformation, and that we know that

$$T\left(\begin{pmatrix} 2 \\ 8 \end{pmatrix}\right) = \begin{pmatrix} 3 \\ 2 \end{pmatrix} \quad \text{and} \quad T\left(\begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) = \begin{pmatrix} -1 \\ 7 \end{pmatrix}.$$

Can we compute other values of T ? Although the process is slightly more involved, we can indeed determine the value of T on all other inputs. The key fact here is that

$$\text{Span}\left(\begin{pmatrix} 2 \\ 8 \end{pmatrix}, \begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) = \mathbb{R}^2$$

because $2 \cdot (-3) - 8 \cdot 1 = -14$, which is nonzero. Thus, given any vector $\vec{v} \in \mathbb{R}^2$, we can express \vec{v} as a linear combination of

$$\begin{pmatrix} 2 \\ 8 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 \\ -3 \end{pmatrix},$$

and then use the fact that T is a linear transformation to determine $T(\vec{v})$. For example, if we want to determine

$$T\left(\begin{pmatrix} 9 \\ 1 \end{pmatrix}\right),$$

we first find the unique pair of numbers $c_1, c_2 \in \mathbb{R}$ with

$$\begin{pmatrix} 9 \\ 1 \end{pmatrix} = c_1 \cdot \begin{pmatrix} 2 \\ 8 \end{pmatrix} + c_2 \cdot \begin{pmatrix} 1 \\ -3 \end{pmatrix}.$$

Working through the algebra, it turns out that

$$\begin{pmatrix} 9 \\ 1 \end{pmatrix} = 2 \cdot \begin{pmatrix} 2 \\ 8 \end{pmatrix} + 5 \cdot \begin{pmatrix} 1 \\ -3 \end{pmatrix}.$$

From this, we conclude that

$$\begin{aligned} T\left(\begin{pmatrix} 9 \\ 1 \end{pmatrix}\right) &= T\left(2 \cdot \begin{pmatrix} 2 \\ 8 \end{pmatrix} + 5 \cdot \begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) \\ &= 2 \cdot T\left(\begin{pmatrix} 2 \\ 8 \end{pmatrix}\right) + 5 \cdot T\left(\begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) && \text{(by Proposition 2.4.2)} \\ &= 2 \cdot \begin{pmatrix} 3 \\ 2 \end{pmatrix} + 5 \cdot \begin{pmatrix} -1 \\ 7 \end{pmatrix} \\ &= \begin{pmatrix} 1 \\ 39 \end{pmatrix}. \end{aligned}$$

Generalizing these ideas leads to the conclusion that if $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a linear transformation and $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$ satisfy $\text{Span}(\vec{u}_1, \vec{u}_2) = \mathbb{R}^2$, then T is completely determined by the two values $T(\vec{u}_1)$ and $T(\vec{u}_2)$. We state this “completely determined” property formally in the following way.

Theorem 2.4.4. *Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 . Suppose that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $S: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ are linear transformations with the property that $T(\vec{u}_1) = S(\vec{u}_1)$ and $T(\vec{u}_2) = S(\vec{u}_2)$. We then have that $T = S$, i.e. $T(\vec{v}) = S(\vec{v})$ for all $\vec{v} \in \mathbb{R}^2$.*

Proof. Let $\vec{v} \in \mathbb{R}^2$ be arbitrary. Since α is a basis of \mathbb{R}^2 , we can fix $c_1, c_2 \in \mathbb{R}$ with

$$\vec{v} = c_1 \vec{u}_1 + c_2 \vec{u}_2.$$

We then have

$$\begin{aligned} T(\vec{v}) &= T(c_1 \vec{u}_1 + c_2 \vec{u}_2) \\ &= c_1 \cdot T(\vec{u}_1) + c_2 \cdot T(\vec{u}_2) && \text{(by Proposition 2.4.2)} \\ &= c_1 \cdot S(\vec{u}_1) + c_2 \cdot S(\vec{u}_2) && \text{(by assumption)} \\ &= S(c_1 \vec{u}_1 + c_2 \vec{u}_2) && \text{(by Proposition 2.4.2)} \\ &= S(\vec{v}). \end{aligned}$$

Since $\vec{v} \in \mathbb{R}^2$ was arbitrary, it follows that $T = S$. □

Notice that the theorem we just proved says nothing about what happens if $\alpha = (\vec{u}_1, \vec{u}_2)$ is not a basis of \mathbb{R}^2 . For example, if $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a linear transformation and we know that

$$T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 4 \\ 1 \end{pmatrix} \quad \text{and} \quad T\left(\begin{pmatrix} 2 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 8 \\ 2 \end{pmatrix},$$

then it seems impossible to determine

$$T\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right)$$

because

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} \notin \text{Span}\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ 0 \end{pmatrix}\right).$$

Up until now, we have determined other values of T under the *assumption* that we had a linear transformation T that produced certain outputs on a couple of given values. We have not yet dealt with the question of whether such a linear transformation T exists at all. In other words, we've discussed a uniqueness question (there is at most one such linear transformation) without addressing the corresponding existence question (whether there is at least one). For instance, if we change the values above and think about whether there is a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with

$$T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 4 \\ 1 \end{pmatrix} \quad \text{and} \quad T\left(\begin{pmatrix} 2 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 5 \\ 3 \end{pmatrix},$$

then we can quickly see that no such T exists. Why? If $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a linear transformation with

$$T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 4 \\ 1 \end{pmatrix},$$

then we must have

$$\begin{aligned} T\left(\begin{pmatrix} 2 \\ 0 \end{pmatrix}\right) &= T\left(2 \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) \\ &= 2 \cdot T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) \\ &= 2 \cdot \begin{pmatrix} 4 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} 8 \\ 2 \end{pmatrix}. \end{aligned}$$

In this instance, we run into a problem because our two vectors are multiples of each other. What happens if this is not the case?

Returning to our example before the proposition, we may ask whether there exists a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that

$$T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 3 \\ 2 \end{pmatrix} \quad \text{and} \quad T\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} -1 \\ 7 \end{pmatrix}.$$

Now in this situation, there is no clear conflict like the one above, but it is also not obvious whether such a linear transformation exists. To prove the existence of such a T , it looks like we need to build it directly. Clearly, there is a *function* $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with

$$f\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 3 \\ 2 \end{pmatrix} \quad \text{and} \quad f\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} -1 \\ 7 \end{pmatrix}$$

because we can just define f like this on these two points and define it to be $\vec{0}$ everywhere else. However, such a piecewise defined function is unlikely to be a linear transformation (and the given f certainly is not). In order to try to obtain a linear transformation T with these properties, we can look to Proposition 2.4.3, which provides an infinite supply of linear transformations. From this perspective, we want to choose $a, b, c, d \in \mathbb{R}$ so that if we let

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix},$$

then we will have

$$T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 3 \\ 2 \end{pmatrix} \quad \text{and} \quad T\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} -1 \\ 7 \end{pmatrix}.$$

Working through the simple algebra, we see that it works if we let $a = 3$, $b = -1$, $c = 2$, and $d = 7$. In other words, the function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} 3x - y \\ 2x + 7y \end{pmatrix}$$

is a linear transformation (by Proposition 2.4.3) that takes the correct values on our two points. There was nothing special about the vectors $(3, 2)$ and $(-1, 7)$. In fact, for any choice of $a, b, c, d \in \mathbb{R}$, there exists a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that

$$T\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad T\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} b \\ d \end{pmatrix},$$

namely

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix}.$$

However, this simple construction does not work as easily if we change the input vectors away from $(1, 0)$ and $(0, 1)$.

Returning to the other example above, we may ask whether there exists a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that

$$T\left(\begin{pmatrix} 2 \\ 8 \end{pmatrix}\right) = \begin{pmatrix} 3 \\ 2 \end{pmatrix} \quad \text{and} \quad T\left(\begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) = \begin{pmatrix} -1 \\ 7 \end{pmatrix}.$$

In this case, it's much harder to try to directly build $a, b, c, d \in \mathbb{R}$ to make this happen. Although it is certainly possible to work through the algebra directly, we instead take a different approach here (we will eventually see a more direct computational approach that avoids the ugly algebra just alluded to). Rather than build

T through a formula, we instead build T more abstractly. Before diving into the general description of T , let's consider how T would have to behave on some specific inputs. For example, we must have

$$\begin{aligned} T\left(\begin{pmatrix} 3 \\ -9 \end{pmatrix}\right) &= T\left(3 \cdot \begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) \\ &= 3 \cdot T\left(\begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) \\ &= 3 \cdot \begin{pmatrix} -1 \\ 7 \end{pmatrix} \\ &= \begin{pmatrix} -3 \\ 21 \end{pmatrix}. \end{aligned}$$

We would also need

$$\begin{aligned} T\left(\begin{pmatrix} 3 \\ 5 \end{pmatrix}\right) &= T\left(\begin{pmatrix} 2 \\ 8 \end{pmatrix} + \begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) \\ &= T\left(\begin{pmatrix} 2 \\ 8 \end{pmatrix}\right) + T\left(\begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) \\ &= \begin{pmatrix} 3 \\ 2 \end{pmatrix} + \begin{pmatrix} -1 \\ 7 \end{pmatrix} \\ &= \begin{pmatrix} 2 \\ 9 \end{pmatrix}. \end{aligned}$$

Following in a similar manner, we can determine what $T(\vec{v})$ would have to be for any $\vec{v} \in \mathbb{R}^2$ because

$$\text{Span}\left(\begin{pmatrix} 2 \\ 8 \end{pmatrix}, \begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) = \mathbb{R}^2.$$

These computations provide us with a candidate function for T . Now we turn this around by simply *defining* a function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ in this way, and check that the resulting function is in fact a linear transformation. Here is the general argument.

Theorem 2.4.5. *Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 . For all $\vec{w}_1, \vec{w}_2 \in \mathbb{R}^2$, there exists a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with $T(\vec{u}_1) = \vec{w}_1$ and $T(\vec{u}_2) = \vec{w}_2$.*

Before jumping into the proof, notice that we are *not* assuming that $\text{Span}(\vec{w}_1, \vec{w}_2) = \mathbb{R}^2$, as such an assumption is unnecessary.

Proof. Define a function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ as follows. Given $\vec{v} \in \mathbb{R}^2$, fix the unique pair $c_1, c_2 \in \mathbb{R}$ with $\vec{v} = c_1\vec{u}_1 + c_2\vec{u}_2$, and define

$$T(\vec{v}) = c_1\vec{w}_1 + c_2\vec{w}_2.$$

Notice that we have described T completely as a function, so we need only check that it is a linear transformation with the required properties.

- We first check that T preserves addition. Let $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ be arbitrary. Since (\vec{u}_1, \vec{u}_2) is a basis of \mathbb{R}^2 by assumption, we can fix the unique pair of numbers $c_1, c_2 \in \mathbb{R}$ with

$$\vec{v}_1 = c_1\vec{u}_1 + c_2\vec{u}_2,$$

and also to fix the unique pair of numbers $d_1, d_2 \in \mathbb{R}^2$ with

$$\vec{v}_2 = d_1\vec{u}_1 + d_2\vec{u}_2.$$

We then have

$$\begin{aligned}\vec{v}_1 + \vec{v}_2 &= (c_1\vec{u}_1 + c_2\vec{u}_2) + (d_1\vec{u}_1 + d_2\vec{u}_2) \\ &= (c_1 + d_1)\vec{u}_1 + (c_2 + d_2)\vec{u}_2,\end{aligned}$$

hence

$$\begin{aligned}T(\vec{v}_1 + \vec{v}_2) &= (c_1 + d_1)\vec{w}_1 + (c_2 + d_2)\vec{w}_2 && \text{(by definition of } T) \\ &= c_1\vec{w}_1 + d_1\vec{w}_1 + c_2\vec{w}_2 + d_2\vec{w}_2 \\ &= c_1\vec{w}_1 + c_2\vec{w}_2 + d_1\vec{w}_1 + d_2\vec{w}_2 \\ &= T(\vec{v}_1) + T(\vec{v}_2) && \text{(by definition of } T).\end{aligned}$$

Since $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ were arbitrary, it follows that T preserves addition.

- We next check that T preserves scalar multiplication. Let $\vec{v} \in \mathbb{R}^2$ and $c \in \mathbb{R}$ be arbitrary. Since (\vec{u}_1, \vec{u}_2) is a basis of \mathbb{R}^2 by assumption, we can fix the unique pair of numbers $d_1, d_2 \in \mathbb{R}$ with

$$\vec{v} = d_1\vec{u}_1 + d_2\vec{u}_2.$$

We then have

$$\begin{aligned}c \cdot \vec{v} &= c \cdot (d_1\vec{u}_1 + d_2\vec{u}_2) \\ &= (cd_1)\vec{u}_1 + (cd_2)\vec{u}_2,\end{aligned}$$

hence

$$\begin{aligned}T(c \cdot \vec{v}) &= (cd_1)\vec{w}_1 + (cd_2)\vec{w}_2 && \text{(by definition of } T) \\ &= c \cdot (d_1\vec{w}_1 + d_2\vec{w}_2) \\ &= c \cdot T(\vec{v}) && \text{(by definition of } T).\end{aligned}$$

Since $\vec{v} \in \mathbb{R}^2$ and $c \in \mathbb{R}$ were arbitrary, it follows that T preserves scalar multiplication.

- We finally check that $T(\vec{u}_1) = \vec{w}_1$ and $T(\vec{u}_2) = \vec{w}_2$. Notice that

$$\vec{u}_1 = 1 \cdot \vec{u}_1 + 0 \cdot \vec{u}_2$$

so by definition of T , we have

$$\begin{aligned}T(\vec{u}_1) &= 1 \cdot \vec{w}_1 + 0 \cdot \vec{w}_2 \\ &= \vec{w}_1.\end{aligned}$$

Similarly, we have

$$\vec{u}_2 = 0 \cdot \vec{u}_1 + 1 \cdot \vec{u}_2$$

so by definition of T , we have

$$\begin{aligned}T(\vec{u}_2) &= 0 \cdot \vec{w}_1 + 1 \cdot \vec{w}_2 \\ &= \vec{w}_2.\end{aligned}$$

To recap, we have built a function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ that is a linear transformation and satisfies both $T(\vec{u}_1) = \vec{w}_1$ and $T(\vec{u}_2) = \vec{w}_2$. This completes the proof. \square

As we've mentioned, linear transformations are the functions from \mathbb{R}^2 to \mathbb{R}^2 that are “nice” from the point of view of linear algebra. In Calculus, you think of a function $f: \mathbb{R} \rightarrow \mathbb{R}$ as being “nice” if it is differentiable (or maybe continuous). Now back in Calculus, you talk about results that say that if you combine two “nice” functions in a basic way, then the resulting function is “nice” as well. For example, we have the following.

- If $f: \mathbb{R} \rightarrow \mathbb{R}$ and $g: \mathbb{R} \rightarrow \mathbb{R}$ are both differentiable functions, then the function $f + g: \mathbb{R} \rightarrow \mathbb{R}$ defined by $(f + g)(x) = f(x) + g(x)$ for all $x \in \mathbb{R}$ is also a differentiable function. Moreover, you learn how to compute $(f + g)'$ in terms of f' and g' . This is the *Sum Rule*.
- If $f: \mathbb{R} \rightarrow \mathbb{R}$ and $g: \mathbb{R} \rightarrow \mathbb{R}$ are both differentiable functions, then the function $f \cdot g: \mathbb{R} \rightarrow \mathbb{R}$ defined by $(f \cdot g)(x) = f(x) \cdot g(x)$ for all $x \in \mathbb{R}$ is also a differentiable function. Moreover, you learn how to compute $(f \cdot g)'$ in terms of f' and g' . This is the *Product Rule*.
- If $f: \mathbb{R} \rightarrow \mathbb{R}$ and $g: \mathbb{R} \rightarrow \mathbb{R}$ are both differentiable functions, then the function $f \circ g: \mathbb{R} \rightarrow \mathbb{R}$ defined by $(f \circ g)(x) = f(g(x))$ for all $x \in \mathbb{R}$ is also a differentiable function. Moreover, you learn how to compute $(f \circ g)'$ in terms of f' and g' . This is the *Chain Rule*.

We could also think about other ways to build new functions from old, such as taking the quotient of two functions (as long as the denominator is nonzero), or multiplying a function by a constant (although this can be seen as a special case of the Product Rule).

We want to ask similar questions about linear transformations. Suppose that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $S: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ are both linear transformations. What operations can we perform on them? Recall that linear transformations are certain special types of functions. We can certainly add the two functions T and S by adding them pointwise like we do in Calculus, except now we are adding the output vectors rather than the output real numbers. Although this provides a well-defined function, it is not completely obvious that the resulting function is a linear transformation. It seems reasonable to believe this might be true because linear transformations “play nice” with addition. We will verify that the resulting function is indeed a linear transformation in Proposition 2.4.7 below.

Can we also multiply T and S together? If we try to do it in the naive manner, then we need to understand how we will define the multiplication of two output vectors in \mathbb{R}^2 . We could define multiplication componentwise, but such a definition lacks a nice geometric interpretation, and it feels unlikely that the resulting function is a linear transformation anyway (since the span of a vector is not closed under such an operation). We might consider the dot product, but then remember that the dot product of two elements of \mathbb{R}^2 is an element of \mathbb{R} rather than an element of \mathbb{R}^2 . Although we can try to think of other potential ways to define a multiplication of vectors, it seems improbable that any of these will result in a linear transformation anyway. In place of a general notion of multiplication of linear transformations, we can instead think of multiplying a linear transformation by a scalar, just as we do for Calculus functions (i.e. we can multiply the function $f(x) = x^2$ by 5 to obtain the function $g(x) = 5x^2$). Since linear transformations “play nice” with scalar multiplication, there is a reasonable hope that the result is a linear transformation.

Definition 2.4.6. *We define the following:*

- Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $S: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be functions. We define a new function $T + S: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by letting $(T + S)(\vec{v}) = T(\vec{v}) + S(\vec{v})$ for all $\vec{v} \in \mathbb{R}^2$.
- Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a function and let $r \in \mathbb{R}$. We define a new function $r \cdot T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by letting $(r \cdot T)(\vec{v}) = r \cdot T(\vec{v})$ for all $\vec{v} \in \mathbb{R}^2$.

Now that we have defined these operations, we argue that if T and S are both linear transformations, then their sum is also a linear transformation (and similarly for scalar multiplication). Moreover, we also show that the *composition* of two linear transformations is a linear transformation. Recall that we defined the composition of two functions f and g generally in Definition 1.6.4, as long as the codomain of the

function on the right equals the domain of the function on the left. Since both of these equal \mathbb{R}^2 for linear transformations, we can indeed also perform composition.

Proposition 2.4.7. *Let $T, S: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be linear transformations.*

1. *The function $T + S$ is a linear transformation.*
2. *For all $r \in \mathbb{R}$, then function $r \cdot T$ is a linear transformation.*
3. *The function $T \circ S$ is a linear transformation.*

Proof. We prove each statement individually:

1. We first check that $T + S$ preserves addition. Let $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ be arbitrary. We have

$$\begin{aligned}
 (T + S)(\vec{v}_1 + \vec{v}_2) &= T(\vec{v}_1 + \vec{v}_2) + S(\vec{v}_1 + \vec{v}_2) && \text{(by definition)} \\
 &= T(\vec{v}_1) + T(\vec{v}_2) + S(\vec{v}_1) + S(\vec{v}_2) && \text{(since } T \text{ and } S \text{ are linear transformations)} \\
 &= T(\vec{v}_1) + S(\vec{v}_1) + T(\vec{v}_2) + S(\vec{v}_2) \\
 &= (T + S)(\vec{v}_1) + (T + S)(\vec{v}_2) && \text{(by definition).}
 \end{aligned}$$

Therefore, the function $T + S$ preserves addition.

We now check that $T + S$ preserves scalar multiplication. Let $\vec{v} \in \mathbb{R}^2$ and $c \in \mathbb{R}$ be arbitrary. We have

$$\begin{aligned}
 (T + S)(c \cdot \vec{v}) &= T(c \cdot \vec{v}) + S(c \cdot \vec{v}) && \text{(by definition)} \\
 &= c \cdot T(\vec{v}) + c \cdot S(\vec{v}) && \text{(since } T \text{ and } S \text{ are linear transformations)} \\
 &= c \cdot (T(\vec{v}) + S(\vec{v})) \\
 &= c \cdot (T + S)(\vec{v}) && \text{(by definition).}
 \end{aligned}$$

Therefore, the function $T + S$ preserves scalar multiplication as well. It follows that $T + S$ is a linear transformation.

2. Let $r \in \mathbb{R}$ be arbitrary. We first check that $r \cdot T$ preserves addition. Let $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ be arbitrary. We have

$$\begin{aligned}
 (r \cdot T)(\vec{v}_1 + \vec{v}_2) &= r \cdot T(\vec{v}_1 + \vec{v}_2) && \text{(by definition)} \\
 &= r \cdot (T(\vec{v}_1) + T(\vec{v}_2)) && \text{(since } T \text{ is a linear transformation)} \\
 &= r \cdot T(\vec{v}_1) + r \cdot T(\vec{v}_2) \\
 &= (r \cdot T)(\vec{v}_1) + (r \cdot T)(\vec{v}_2) && \text{(by definition).}
 \end{aligned}$$

Therefore, the function $r \cdot T$ preserves addition.

We now check that $r \cdot T$ preserves scalar multiplication. Let $\vec{v} \in \mathbb{R}^2$ and $c \in \mathbb{R}$ be arbitrary. We have

$$\begin{aligned}
 (r \cdot T)(c \cdot \vec{v}) &= r \cdot T(c \cdot \vec{v}) && \text{(by definition)} \\
 &= r \cdot (c \cdot T(\vec{v})) && \text{(since } T \text{ is a linear transformation)} \\
 &= (rc) \cdot T(\vec{v}) \\
 &= (cr) \cdot T(\vec{v}) \\
 &= c \cdot (r \cdot T(\vec{v})) \\
 &= c \cdot (r \cdot T)(\vec{v}) && \text{(by definition).}
 \end{aligned}$$

Therefore, the function $r \cdot T$ preserves scalar multiplication as well. It follows that $r \cdot T$ is a linear transformation.

3. Exercise.

□

2.5 The Standard Matrix of a Linear Transformation

In this section, we develop some shorthand notation for referring to a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and its action on vectors in \mathbb{R}^2 (i.e. the result of plugging \vec{v} into T to form $T(\vec{v})$). In Calculus, the notation $\frac{dy}{dx}$ for derivatives aids both computation and theoretical discussions because it greatly simplifies what one needs to think about, naturally reminds us of the limit operation that defines a derivative, and suggests derivative and integral rules that turn out to be true. We want to develop some concise and efficient notation for linear algebra that provide similar benefits.

Throughout this section, we will denote the two standard unit vectors along the axes with the following notation:

$$\vec{e}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad \text{and} \quad \vec{e}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

Notice that $\text{Span}(\vec{e}_1, \vec{e}_2) = \mathbb{R}^2$, either by directly working with the definition of $\text{Span}(\vec{e}_1, \vec{e}_2)$, or by simply using Theorem 2.3.10 together with the fact that $1 \cdot 1 - 0 \cdot 0 = 1$.

Suppose now that $a, b, c, d \in \mathbb{R}$. We know from Theorem 2.4.5 and Proposition 2.4.4 that there exists a unique linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that

$$T(\vec{e}_1) = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad T(\vec{e}_2) = \begin{pmatrix} b \\ d \end{pmatrix}.$$

In other words, the two vectors $T(\vec{e}_1)$ and $T(\vec{e}_2)$ encode everything there is to know about T . Therefore, a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is completely determined by 4 numbers (i.e. the 4 numbers comprising the two vectors $T(\vec{e}_1)$ and $T(\vec{e}_2)$). Instead of writing down the two vectors

$$T(\vec{e}_1) = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad T(\vec{e}_2) = \begin{pmatrix} b \\ d \end{pmatrix}$$

in order, we can arrange them in a certain 2×2 table:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

We call such a table a *matrix*, or more specifically a 2×2 matrix. In other words, a 2×2 matrix is just an arrangement of 4 real numbers into 2 rows and 2 columns, where the position of a number matters (like a 4-tuple). We will eventually see many uses for matrices and interpret their table structure in different ways depending on the context. However, for the moment, we want to think of a matrix as “coding” a linear transformation. To enable this, we will use the following terminology.

Definition 2.5.1. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. Fix $a, b, c, d \in \mathbb{R}$ with

$$T(\vec{e}_1) = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad T(\vec{e}_2) = \begin{pmatrix} b \\ d \end{pmatrix}.$$

We define the standard matrix of T to be the following 2×2 matrix:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

In other words, the standard matrix of T has the entries of $T(\vec{e}_1)$ in the first column, and the entries of $T(\vec{e}_2)$ in the second column.

Notation 2.5.2. If $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a linear transformation, then we use the notation $[T]$ to denote the standard matrix of T .

As mentioned, we want to think of the standard matrix as a shorthand “code” for the corresponding linear transformation. That is, instead of saying:

Consider the unique linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that $T(\vec{e}_1) = \begin{pmatrix} a \\ c \end{pmatrix}$ and $T(\vec{e}_2) = \begin{pmatrix} b \\ d \end{pmatrix}$,

we can instead simply say:

Consider the linear transformation with standard matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$.

Now we want to push this shorthand notation further by developing a more succinct notation for plugging a vector as input into the function T . Recall that if we know the values of T on the vectors \vec{e}_1 and \vec{e}_2 , then we can in fact determine a formula for T on an arbitrary vector as follows. Assuming that

$$T(\vec{e}_1) = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad T(\vec{e}_2) = \begin{pmatrix} b \\ d \end{pmatrix},$$

then for any $x, y \in \mathbb{R}$, we have

$$\begin{aligned} T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) &= T\left(x \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} + y \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) \\ &= T(x \cdot \vec{e}_1 + y \cdot \vec{e}_2) \\ &= x \cdot T(\vec{e}_1) + y \cdot T(\vec{e}_2) \\ &= x \cdot \begin{pmatrix} a \\ c \end{pmatrix} + y \cdot \begin{pmatrix} b \\ d \end{pmatrix} \\ &= \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix}. \end{aligned}$$

Since we are using this matrix to “code” the linear transformation T , can we create natural notation that dispenses with T entirely and instead just uses the standard matrix? A function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ takes a vector $\vec{v} \in \mathbb{R}^2$ and feeds it in as input in order to form the output $T(\vec{v})$. We can replace the direct reference to the function T by instead using the standard matrix $[T]$, but now rather than feeding \vec{v} as an input to T (resulting in the notation $T(\vec{v})$), we instead think of “hitting” \vec{v} with the matrix $[T]$. Thus, if

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix}$$

then we will write

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

in place of the usual function notation $T(\vec{v})$. If we want to make this notation match up with and equal the output of the function T , then we should *define* the above matrix-vector “product” to be

$$x \cdot \begin{pmatrix} a \\ c \end{pmatrix} + y \cdot \begin{pmatrix} b \\ d \end{pmatrix} = \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix}.$$

In other words, we are guided to the following definition.

Definition 2.5.3. Let $a, b, c, d, x, y \in \mathbb{R}$. Let A be the 2×2 matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

and let $\vec{v} \in \mathbb{R}^2$ be

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

We define the matrix-vector product, written as $A\vec{v}$, to be the vector

$$\begin{pmatrix} ax + by \\ cx + dy \end{pmatrix}.$$

In other words, we define

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix}.$$

For example, we have

$$\begin{aligned} \begin{pmatrix} 5 & -1 \\ 2 & 4 \end{pmatrix} \begin{pmatrix} -3 \\ 8 \end{pmatrix} &= \begin{pmatrix} 5 \cdot (-3) + (-1) \cdot 8 \\ 2 \cdot (-3) + 4 \cdot 8 \end{pmatrix} \\ &= \begin{pmatrix} -23 \\ 26 \end{pmatrix}. \end{aligned}$$

Notice that if A is a 2×2 matrix, and $\vec{v} \in \mathbb{R}^2$, then $A\vec{v} \in \mathbb{R}^2$. Furthermore, the first entry of $A\vec{v}$ is the dot product of the first row of A with \vec{v} , and the second entry of $A\vec{v}$ is the dot product of the second row of A with \vec{v} .

Proposition 2.5.4. *Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. If A is the standard matrix of T , then $T(\vec{v}) = A\vec{v}$ for all $\vec{v} \in \mathbb{R}^2$. In other words, $T(\vec{v}) = [T]\vec{v}$ for all $\vec{v} \in \mathbb{R}^2$.*

Proof. The proof appears in the discussion before the definition of the matrix-vector product, and in fact we defined the matrix-vector product so that this result would be true. \square

Putting everything together, we can view the above computation

$$\begin{pmatrix} 5 & -1 \\ 2 & 4 \end{pmatrix} \begin{pmatrix} -3 \\ 8 \end{pmatrix} = \begin{pmatrix} -23 \\ 26 \end{pmatrix}$$

as shorthand for:

$$\begin{aligned} \text{If } T: \mathbb{R}^2 \rightarrow \mathbb{R}^2 \text{ is the unique linear transformation with } T(\vec{e}_1) = \begin{pmatrix} 5 \\ 2 \end{pmatrix} \text{ and } T(\vec{e}_2) = \begin{pmatrix} -1 \\ 4 \end{pmatrix}, \\ \text{then we have } T\left(\begin{pmatrix} -3 \\ 8 \end{pmatrix}\right) = \begin{pmatrix} -23 \\ 26 \end{pmatrix}. \end{aligned}$$

We can also view the computation

$$\begin{pmatrix} 3 & -1 \\ 2 & 7 \end{pmatrix} \begin{pmatrix} 8 \\ 3 \end{pmatrix} = \begin{pmatrix} 3 \cdot 8 + (-1) \cdot 3 \\ 2 \cdot 8 + 7 \cdot 3 \end{pmatrix} = \begin{pmatrix} 21 \\ 37 \end{pmatrix}$$

as shorthand for the long-winded work we carried out in Section 2.4 after the proof of Proposition 2.4.3.

Now which of these two approaches would you rather take when performing computations? Although our new notation allows us the luxury of performing computations without working through the detailed thought process every time, always remember that we are thinking of a matrix as encoding a certain type of *function* from \mathbb{R}^2 to \mathbb{R}^2 . In other words, we take in an element of \mathbb{R}^2 , and by “hitting it” with a matrix we are turning it into another element of \mathbb{R}^2 . This input-output behavior is just like that of a function, because after all a matrix is *encoding* a function (in fact, it is encoding a linear transformation).

Before working out the standard matrices for various linear transformations, let's pause for a moment to interpret the matrix-vector product in a different light. Given $a, b, c, d, x, y \in \mathbb{R}$, we defined

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix}.$$

Notice then that

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = x \cdot \begin{pmatrix} a \\ c \end{pmatrix} + y \cdot \begin{pmatrix} b \\ d \end{pmatrix},$$

so we can think of the matrix-vector product as follows: Given a vector, take its entries as the “weights” used to form a linear combination of the columns of the matrix. In other words, we can view the computation above as computing the following linear combination:

$$\begin{pmatrix} 5 & -1 \\ 2 & 4 \end{pmatrix} \begin{pmatrix} -3 \\ 8 \end{pmatrix} = (-3) \cdot \begin{pmatrix} 5 \\ 2 \end{pmatrix} + 8 \cdot \begin{pmatrix} -1 \\ 4 \end{pmatrix}.$$

As we will see, viewing a matrix-vector product from this perspective can be very useful.

Now since we are viewing a matrix as a table where the order of the entries matter, is it natural to define equality of matrices as follows.

Definition 2.5.5. Given two 2×2 matrices, say

$$A = \begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} a_2 & b_2 \\ c_2 & d_2 \end{pmatrix},$$

we define $A = B$ to mean that $a_1 = a_2$, $b_1 = b_2$, $c_1 = c_2$, and $d_1 = d_2$.

Fortunately, this definition matches up perfectly with equality of the corresponding linear transformations.

Proposition 2.5.6. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $S: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ both be linear transformations. We then have that $T = S$ if and only if $[T] = [S]$.

Proof. Suppose first that $T = S$, i.e. that $T(\vec{v}) = S(\vec{v})$ for all $\vec{v} \in \mathbb{R}^2$. In particular, we then have that both $T(\vec{e}_1) = S(\vec{e}_1)$ and $T(\vec{e}_2) = S(\vec{e}_2)$. Therefore, by the definition of the standard matrix, it follows that $[T] = [S]$.

Suppose instead that $[T] = [S]$. By the definition of the standard matrix, it follows that both $T(\vec{e}_1) = S(\vec{e}_1)$ and $T(\vec{e}_2) = S(\vec{e}_2)$ are true. Since $\text{Span}(\vec{e}_1, \vec{e}_2) = \mathbb{R}^2$, we can use Proposition 2.4.4 to conclude that $T = S$. \square

Over the next few results, we work to determine the standard matrices of several fundamental linear transformations.

Proposition 2.5.7. Let $id: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the function defined by $id(\vec{v}) = \vec{v}$ for all $\vec{v} \in \mathbb{R}^2$. In other words, id is the identity function. We then have that id is a linear transformation and

$$[id] = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Proof. We first check that id is a linear transformation.

- Let $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ be arbitrary. We have

$$\begin{aligned} id(\vec{v}_1 + \vec{v}_2) &= \vec{v}_1 + \vec{v}_2 \\ &= id(\vec{v}_1) + id(\vec{v}_2). \end{aligned}$$

Therefore, id preserves addition.

- Let $\vec{v} \in \mathbb{R}^2$ and $c \in \mathbb{R}$ be arbitrary. We have

$$\begin{aligned} id(c \cdot \vec{v}) &= c \cdot \vec{v} \\ &= c \cdot id(\vec{v}). \end{aligned}$$

Therefore, id preserves scalar multiplication.

It follows that id is a linear transformation.

We now notice that

$$id(\vec{e}_1) = \vec{e}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

and

$$id(\vec{e}_2) = \vec{e}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

so

$$[id] = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

This completes the proof. □

In general, if we have a simple formula for T , then determining $[T]$ is easy.

Proposition 2.5.8. *Let $a, b, c, d \in \mathbb{R}$ and let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be defined by*

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix}.$$

We then have that T is a linear transformation and

$$[T] = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Proof. By Proposition 2.4.3, we know that T is a linear transformation. Now

$$T(\vec{e}_1) = \begin{pmatrix} a \cdot 1 + b \cdot 0 \\ c \cdot 1 + d \cdot 0 \end{pmatrix} = \begin{pmatrix} a \\ c \end{pmatrix}$$

and

$$T(\vec{e}_2) = \begin{pmatrix} a \cdot 0 + b \cdot 1 \\ c \cdot 0 + d \cdot 1 \end{pmatrix} = \begin{pmatrix} b \\ d \end{pmatrix}.$$

Therefore,

$$[T] = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

which completes the proof. □

We now move on to determine the standard matrix of a linear transformation for which we have not yet developed a formula: rotations.

Proposition 2.5.9. *Let $\theta \in \mathbb{R}$. Define a function $R_\theta: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by letting $R_\theta(\vec{v})$ be the result of rotating \vec{v} by θ radians counterclockwise around the origin. We then have that R_θ is a linear transformation and*

$$[R_\theta] = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}.$$

Proof. When we first introduced the idea of a linear transformation at the beginning of Section 2.4, we argued through geometric means that a rotation through an angle of θ preserves both addition and multiplication. It follows that R_θ is a linear transformation. To determine the standard matrix $[R_\theta]$, notice first that

$$R_\theta(\vec{e}_1) = \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix}$$

by definition of these trigonometric functions (recall that the definition of $\cos \theta$ and $\sin \theta$ is that if we walk a distance θ along the unit circle counterclockwise from the point $(1, 0)$, then $\cos \theta$ is the x -coordinate of this point and $\sin \theta$ is the y -coordinate of this point). Now \vec{e}_2 has already been rotated from \vec{e}_1 by $\frac{\pi}{2}$ radians, so

$$R_\theta(\vec{e}_2) = \begin{pmatrix} \cos(\theta + \frac{\pi}{2}) \\ \sin(\theta + \frac{\pi}{2}) \end{pmatrix}.$$

Putting these together, we see that

$$[R_\theta] = \begin{pmatrix} \cos \theta & \cos(\theta + \frac{\pi}{2}) \\ \sin \theta & \sin(\theta + \frac{\pi}{2}) \end{pmatrix}.$$

Using some basic trigonometric identities, we have $\cos(\theta + \frac{\pi}{2}) = -\sin \theta$ and $\sin(\theta + \frac{\pi}{2}) = \cos \theta$. Therefore,

$$[R_\theta] = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}.$$

□

For example, we have

$$[R_{\frac{\pi}{4}}] = \begin{pmatrix} \cos \frac{\pi}{4} & -\sin \frac{\pi}{4} \\ \sin \frac{\pi}{4} & \cos \frac{\pi}{4} \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}.$$

Thus, if we want to rotate $(2, 7)$ counterclockwise by $\frac{\pi}{4}$ radians, we can just compute

$$\begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} 2 \\ 7 \end{pmatrix} = \begin{pmatrix} \frac{-5}{\sqrt{2}} \\ \frac{8}{\sqrt{2}} \end{pmatrix}.$$

Notice that once we have computed the standard matrix for a linear transformation (such as a rotation), it becomes completely straightforward to compute how the linear transformation behaves on *any* input by just following the formula for the matrix-vector product.

In the previous proof, we asserted that R_θ was a linear transformation through geometric reasoning. If you don't find this argument completely convincing, we can approach the problem from a different perspective. Following the above proof, we can derive a formula for R_θ under the *assumption* that R_θ is a linear transformation. With such a formula in hand, we can then check after the fact that the output really has the desired properties. In our case, if we *assume* that R_θ is a linear transformation, then the above argument tell us that we must have

$$\begin{aligned} R_\theta \left(\begin{pmatrix} x \\ y \end{pmatrix} \right) &= \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \\ &= \begin{pmatrix} x \cos \theta - y \sin \theta \\ x \sin \theta + y \cos \theta \end{pmatrix} \end{aligned}$$

for all $x, y \in \mathbb{R}$. Now that we have this as a potential formula for rotation through an angle of θ , let's go ahead and verify that the output really is the rotation of (x, y) by an angle of θ . In order to do that, we

have to show that the angle between \vec{v} and $R_\theta(\vec{v})$ equals θ , and also that the length of \vec{v} equals the length of $R_\theta(\vec{v})$, for all $\vec{v} \in \mathbb{R}^2$. We first check the latter, and will use the notation $\|\vec{v}\|$ for the length of the vector \vec{v} . Given an arbitrary $\vec{v} \in \mathbb{R}^2$, say with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix},$$

we have

$$\begin{aligned} \|R_\theta(\vec{v})\| &= \sqrt{(x \cos \theta - y \sin \theta)^2 + (x \sin \theta + y \cos \theta)^2} \\ &= \sqrt{x^2 \cos^2 \theta - 2xy \cos \theta \sin \theta + y^2 \sin^2 \theta + x^2 \sin^2 \theta + 2xy \sin \theta \cos \theta + y^2 \cos^2 \theta} \\ &= \sqrt{x^2(\cos^2 \theta + \sin^2 \theta) + y^2(\cos^2 \theta + \sin^2 \theta)} \\ &= \sqrt{x^2 + y^2} \\ &= \|\vec{v}\|. \end{aligned}$$

We now check that the angle between $R_\theta(\vec{v})$ and \vec{v} really is θ . Recall from Calculus that if $\vec{u}, \vec{w} \in \mathbb{R}^2$, and if α is the angle between them, then

$$\vec{u} \cdot \vec{w} = \|\vec{u}\| \|\vec{w}\| \cos \alpha,$$

(where the \cdot on the left is the dot product), so we have

$$\cos \alpha = \frac{\vec{u} \cdot \vec{w}}{\|\vec{u}\| \|\vec{w}\|}.$$

Let's use this to compute the angle between $R_\theta(\vec{v})$ and \vec{v} in our case. For any $\vec{v} \in \mathbb{R}^2$, say with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix},$$

we have

$$\begin{aligned} \frac{R_\theta(\vec{v}) \cdot \vec{v}}{\|R_\theta(\vec{v})\| \|\vec{v}\|} &= \frac{R_\theta(\vec{v}) \cdot \vec{v}}{\|\vec{v}\| \|\vec{v}\|} && \text{(from above)} \\ &= \frac{x(x \cos \theta - y \sin \theta) + y(x \sin \theta + y \cos \theta)}{x^2 + y^2} \\ &= \frac{x^2 \cos \theta - xy \sin \theta + xy \sin \theta + y^2 \cos \theta}{x^2 + y^2} \\ &= \frac{(x^2 + y^2) \cos \theta}{x^2 + y^2} \\ &= \cos \theta. \end{aligned}$$

Thus, the angle between $R_\theta(\vec{v})$ and \vec{v} really is θ . Therefore, we have verified that the above formula for R_θ really does rotate an arbitrary vector \vec{v} by an angle of θ . Now without the geometric reasoning, we have not checked whether this angle of θ is clockwise or counterclockwise (i.e. we have not checked which “side” of \vec{v} the vector $R_\theta(\vec{v})$ is on), but we will discuss such issues when we get to determinants.

More generally, if we have a function $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ either described geometrically or through a complicated procedure, and we guess (or believe) that T might be a linear transformation, then we can use that guess to derive a potential formula for T based on the values $T(\vec{e}_1)$ and $T(\vec{e}_2)$. Once we have that formula, we can verify after the fact that the formula works (since we will have a potential formula in hand to play with).

We can also work backwards. In other words, we can take a matrix, and interpret it as coding a linear transformation from \mathbb{R}^2 to \mathbb{R}^2 . For example, consider the following matrix:

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

Now we know that there exists a unique linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ having this matrix as its standard matrix. Let's try to interpret the action of this unique T geometrically. For any $x, y \in \mathbb{R}$, we must have

$$\begin{aligned} T \left(\begin{pmatrix} x \\ y \end{pmatrix} \right) &= \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \\ &= \begin{pmatrix} x + y \\ y \end{pmatrix}. \end{aligned}$$

In particular, since

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ 0 \end{pmatrix} = \begin{pmatrix} x \\ 0 \end{pmatrix}$$

for all $x \in \mathbb{R}$, we have that T fixes every point on the x -axis. Also, we have

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ 1 \end{pmatrix} = \begin{pmatrix} x + 1 \\ 1 \end{pmatrix}$$

for all $x \in \mathbb{R}$, so T shifts every point on the line $y = 1$ to the right by 1. Similarly, for all $b \in \mathbb{R}$ with $b > 0$, we have that T shifts every point on the line $y = b$ to the right by b , and for all $b \in \mathbb{R}$ with $b < 0$, we have that T shifts every point on the line $y = b$ to the left by b . We call such a function a “shear transformation” because it shifts the plane in opposite directions relative to one axis (in this case, the x -axis). In other words, we say that the matrix

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

“codes” a horizontal shear transformation. Generalizing this, given $k \in \mathbb{R}$ with $k > 0$, the matrix

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

is the standard matrix of a horizontal shear transformation that shifts the line $y = b$ to the right by kb if $b > 0$, and shifts the line $y = b$ to the left by kb if $b < 0$. If $k < 0$, then the right/left shifts are switched above and below the x -axis. Similarly, the matrix

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

is the standard matrix of a vertical shear transformation that shifts the line $x = a$ up or down based on the signs of k and a .

We now show that another important class of functions are linear transformations: *projections* of points onto lines through the origin. We also derive the standard matrices for these linear transformations.

Proposition 2.5.10. *Let $\vec{w} \in \mathbb{R}^2$ be a nonzero vector and let $W = \text{Span}(\vec{w})$. Define $P_{\vec{w}}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by letting $P_{\vec{w}}(\vec{v})$ be the vector in W that is closest to \vec{v} . We then have that $P_{\vec{w}}$ is a linear transformation. Moreover, if*

$$\vec{w} = \begin{pmatrix} a \\ b \end{pmatrix}$$

where $a, b \in \mathbb{R}$, then

$$[P_{\vec{w}}] = \begin{pmatrix} \frac{a^2}{a^2+b^2} & \frac{ab}{a^2+b^2} \\ \frac{ab}{a^2+b^2} & \frac{b^2}{a^2+b^2} \end{pmatrix}.$$

Proof. We have defined $P_{\vec{w}}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by letting $P_{\vec{w}}(\vec{v})$ be the vector in W that is closest to \vec{v} . Given $\vec{v} \in \mathbb{R}^2$, we know from Calculus that the projection of \vec{v} onto the vector \vec{w} equals

$$P_{\vec{w}}(\vec{v}) = \frac{\vec{v} \cdot \vec{w}}{\|\vec{w}\|^2} \cdot \vec{w},$$

where the \cdot in the numerator is the dot product of the two vectors. Thus, if we fix $x, y \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix},$$

then

$$\begin{aligned} P_{\vec{w}}(\vec{v}) &= \frac{xa + yb}{(\sqrt{a^2 + b^2})^2} \cdot \begin{pmatrix} a \\ b \end{pmatrix} \\ &= \frac{ax + by}{a^2 + b^2} \cdot \begin{pmatrix} a \\ b \end{pmatrix} \\ &= \begin{pmatrix} \frac{a^2x + aby}{a^2 + b^2} \\ \frac{abx + b^2y}{a^2 + b^2} \end{pmatrix}. \end{aligned}$$

Therefore, the function $P_{\vec{w}}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ can be written as

$$\begin{aligned} P_{\vec{w}}\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) &= \begin{pmatrix} \frac{a^2x + aby}{a^2 + b^2} \\ \frac{abx + b^2y}{a^2 + b^2} \end{pmatrix} \\ &= \begin{pmatrix} \frac{a^2}{a^2 + b^2} \cdot x + \frac{ab}{a^2 + b^2} \cdot y \\ \frac{ab}{a^2 + b^2} \cdot x + \frac{b^2}{a^2 + b^2} \cdot y \end{pmatrix}. \end{aligned}$$

Notice that $P_{\vec{w}}$ is a linear transformation by Proposition 2.4.3. Furthermore, we can use Proposition 2.5.8 to conclude that

$$[P_{\vec{w}}] = \begin{pmatrix} \frac{a^2}{a^2 + b^2} & \frac{ab}{a^2 + b^2} \\ \frac{ab}{a^2 + b^2} & \frac{b^2}{a^2 + b^2} \end{pmatrix}.$$

□

For example, suppose that we want to understand the linear transformation that projects a point onto the line $y = 2x$. Notice that if we let

$$\vec{w} = \begin{pmatrix} 1 \\ 2 \end{pmatrix},$$

then the solution set to the equation $y = 2x$ equals $\text{Span}(\vec{w})$. Thus, we have

$$[P_{\vec{w}}] = \begin{pmatrix} \frac{1}{5} & \frac{2}{5} \\ \frac{2}{5} & \frac{4}{5} \end{pmatrix}.$$

For example, the projection of the point

$$\begin{pmatrix} 13 \\ -3 \end{pmatrix}$$

onto the line $y = 2x$ is given by

$$\begin{pmatrix} \frac{1}{5} & \frac{2}{5} \\ \frac{2}{5} & \frac{4}{5} \end{pmatrix} \begin{pmatrix} 13 \\ -3 \end{pmatrix} = \begin{pmatrix} \frac{7}{5} \\ \frac{14}{5} \end{pmatrix}.$$

We can use our work on projections to show that the reflection across a line through the origin is always a linear transformation, and also to compute the corresponding standard matrix.

Proposition 2.5.11. *Let $\vec{w} \in \mathbb{R}^2$ be a nonzero vector and let $W = \text{Span}(\vec{w})$. Define $F_{\vec{w}}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by letting $F_{\vec{w}}(\vec{v})$ be the result of reflecting \vec{v} across the line W . We then have that $F_{\vec{w}}$ is a linear transformation. Moreover, if*

$$\vec{w} = \begin{pmatrix} a \\ b \end{pmatrix}$$

where $a, b \in \mathbb{R}$, then

$$[F_{\vec{w}}] = \begin{pmatrix} \frac{a^2-b^2}{a^2+b^2} & \frac{2ab}{a^2+b^2} \\ \frac{2ab}{a^2+b^2} & \frac{b^2-a^2}{a^2+b^2} \end{pmatrix}.$$

Proof. Consider an arbitrary $\vec{v} \in \mathbb{R}^2$. Notice the vector with tail at the head of \vec{v} and tip at the point $P_{\vec{w}}(\vec{v})$ can be written as $P_{\vec{w}}(\vec{v}) - \vec{v}$. If we add this vector to \vec{v} , then of course we land at $P_{\vec{w}}(\vec{v})$. Now if we want to reflect *across* W , then we want to add this vector again. In other words, we want to add 2 times $P_{\vec{w}}(\vec{v}) - \vec{v}$ to \vec{v} . Therefore, we have

$$F_{\vec{w}}(\vec{v}) = \vec{v} + 2 \cdot (P_{\vec{w}}(\vec{v}) - \vec{v})$$

for all $\vec{v} \in \mathbb{R}^2$, so

$$F_{\vec{w}}(\vec{v}) = 2 \cdot P_{\vec{w}}(\vec{v}) - \vec{v}$$

for all $\vec{v} \in \mathbb{R}^2$. Now given arbitrary $x, y \in \mathbb{R}$, we can use Proposition 2.5.10 to compute

$$\begin{aligned} F_{\vec{w}} \left(\begin{pmatrix} x \\ y \end{pmatrix} \right) &= 2 \cdot P_{\vec{w}} \left(\begin{pmatrix} x \\ y \end{pmatrix} \right) - \begin{pmatrix} x \\ y \end{pmatrix} \\ &= 2 \cdot \left(\begin{pmatrix} \frac{a^2}{a^2+b^2} \cdot x + \frac{ab}{a^2+b^2} \cdot y \\ \frac{ab}{a^2+b^2} \cdot x + \frac{b^2}{a^2+b^2} \cdot y \end{pmatrix} \right) - \begin{pmatrix} x \\ y \end{pmatrix} \\ &= \begin{pmatrix} \left(\frac{2a^2}{a^2+b^2} - 1 \right) \cdot x + \frac{2ab}{a^2+b^2} \cdot y \\ \frac{2ab}{a^2+b^2} \cdot x + \left(\frac{2b^2}{a^2+b^2} - 1 \right) \cdot y \end{pmatrix} \\ &= \begin{pmatrix} \frac{a^2-b^2}{a^2+b^2} \cdot x + \frac{2ab}{a^2+b^2} \cdot y \\ \frac{2ab}{a^2+b^2} \cdot x + \frac{b^2-a^2}{a^2+b^2} \cdot y \end{pmatrix}. \end{aligned}$$

Therefore, using Proposition 2.4.3 and 2.5.8, we conclude that $F_{\vec{w}}$ is a linear transformation and that

$$[F_{\vec{w}}] = \begin{pmatrix} \frac{a^2-b^2}{a^2+b^2} & \frac{2ab}{a^2+b^2} \\ \frac{2ab}{a^2+b^2} & \frac{b^2-a^2}{a^2+b^2} \end{pmatrix}.$$

□

2.6 Matrix Algebra

Now that we have developed formulas for the standard matrix of several important classes of linear transformations, we can explore what happens when we build new linear transformations from old ones (just like how we develop the formulas for the derivatives of some simple functions, and then learn the Sum Rule, Product Rule, and Chain Rule to handle more complicated functions in Calculus). For example, suppose that $T_1: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $T_2: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ are both linear transformations. We know that $T_1 + T_2$ is a linear transformation from Proposition 2.4.7, so we can consider its standard matrix $[T_1 + T_2]$. Now we would like to know how to compute this based on knowledge of the two standard matrices $[T_1]$ and $[T_2]$. If we find a

nice way to do this, then it will be natural to *define* the sum of two matrices in such a way so that this works. Suppose then that

$$[T_1] = \begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} \quad \text{and} \quad [T_2] = \begin{pmatrix} a_2 & b_2 \\ c_2 & d_2 \end{pmatrix}.$$

To compute the standard matrix for $[T_1 + T_2]$, we need to determine what this linear transformation does to the vectors \vec{e}_1 and \vec{e}_2 . We have

$$\begin{aligned} (T_1 + T_2)(\vec{e}_1) &= T_1(\vec{e}_1) + T_2(\vec{e}_1) \\ &= \begin{pmatrix} a_1 \\ c_1 \end{pmatrix} + \begin{pmatrix} a_2 \\ c_2 \end{pmatrix} \\ &= \begin{pmatrix} a_1 + a_2 \\ c_1 + c_2 \end{pmatrix} \end{aligned}$$

and

$$\begin{aligned} (T_1 + T_2)(\vec{e}_2) &= T_1(\vec{e}_2) + T_2(\vec{e}_2) \\ &= \begin{pmatrix} b_1 \\ d_1 \end{pmatrix} + \begin{pmatrix} b_2 \\ d_2 \end{pmatrix} \\ &= \begin{pmatrix} b_1 + b_2 \\ d_1 + d_2 \end{pmatrix}. \end{aligned}$$

Thus,

$$[T_1 + T_2] = \begin{pmatrix} a_1 + a_2 & b_1 + b_2 \\ c_1 + c_2 & d_1 + d_2 \end{pmatrix}.$$

In other words, if we want to define addition of matrices in such a way that it corresponds to addition of linear transformations, then we should define

$$\begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} + \begin{pmatrix} a_2 & b_2 \\ c_2 & d_2 \end{pmatrix} = \begin{pmatrix} a_1 + a_2 & b_1 + b_2 \\ c_1 + c_2 & d_1 + d_2 \end{pmatrix}.$$

Now that we are writing matrix-vector products in place of function notation, we are interpreting the matrix-vector product

$$\begin{pmatrix} 1 & 7 \\ -2 & 1 \end{pmatrix} \begin{pmatrix} -5 \\ 0 \end{pmatrix}$$

as the element of \mathbb{R}^2 that results from feeding the vector in as input to the linear transformation coded by the matrix. In other words, the matrix takes in the vector and transforms it into an output. We can then take this output and feed it into another linear transformation. If we want to put this output into the linear transformation given by

$$\begin{pmatrix} 3 & 1 \\ 4 & 9 \end{pmatrix},$$

then it is natural to chain all of the pieces together and write

$$\begin{pmatrix} 3 & 1 \\ 4 & 9 \end{pmatrix} \begin{pmatrix} 1 & 7 \\ -2 & 1 \end{pmatrix} \begin{pmatrix} -5 \\ 0 \end{pmatrix}.$$

In this case, the action goes from right to left, as the input vector is fed into the matrix next to it and transformed into a new vector, whose result is then fed into the matrix on the left. In other words, we want

this chaining to represent the composition of the corresponding linear transformations. Since we are writing this as what looks like a product, we should *define* matrix multiplication so that it corresponds to function composition. In other words, we want to define

$$\begin{pmatrix} 3 & 1 \\ 4 & 9 \end{pmatrix} \begin{pmatrix} 1 & 7 \\ -2 & 1 \end{pmatrix}$$

to be the standard matrix of the composition of the linear transformations. If we are successful in doing this, then instead of multiplying right to left above, we can multiply the matrices to correspond to the composition first, and then feed the vector into the resulting matrix.

Let's figure out how to make this work. Let $T_1: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $T_2: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be linear transformations. Using Proposition 2.4.7, we know that $T_1 \circ T_2$ is a linear transformation. Suppose then that

$$[T_1] = \begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} \quad \text{and} \quad [T_2] = \begin{pmatrix} a_2 & b_2 \\ c_2 & d_2 \end{pmatrix}.$$

We now compute $[T_1 \circ T_2]$. In order to do this, we want to compute what $T_1 \circ T_2$ does to both \vec{e}_1 and \vec{e}_2 . We have

$$\begin{aligned} (T_1 \circ T_2)(\vec{e}_1) &= T_1(T_2(\vec{e}_1)) \\ &= T_1 \left(\begin{pmatrix} a_2 \\ c_2 \end{pmatrix} \right) \\ &= \begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} \begin{pmatrix} a_2 \\ c_2 \end{pmatrix} \\ &= \begin{pmatrix} a_1 a_2 + b_1 c_2 \\ c_1 a_2 + d_1 c_2 \end{pmatrix} \end{aligned}$$

and

$$\begin{aligned} (T_1 \circ T_2)(\vec{e}_2) &= T_1(T_2(\vec{e}_2)) \\ &= T_1 \left(\begin{pmatrix} b_2 \\ d_2 \end{pmatrix} \right) \\ &= \begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} \begin{pmatrix} b_2 \\ d_2 \end{pmatrix} \\ &= \begin{pmatrix} a_1 b_2 + b_1 d_2 \\ c_1 b_2 + d_1 d_2 \end{pmatrix}. \end{aligned}$$

Therefore, if we want matrix multiplication to correspond to composition of linear transformations, then we should define

$$\begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} \begin{pmatrix} a_2 & b_2 \\ c_2 & d_2 \end{pmatrix} = \begin{pmatrix} a_1 a_2 + b_1 c_2 & a_1 b_2 + b_1 d_2 \\ c_1 a_2 + d_1 c_2 & c_1 b_2 + d_1 d_2 \end{pmatrix}.$$

We now turn these into formal definitions.

Definition 2.6.1. *Given two matrices*

$$\begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} a_2 & b_2 \\ c_2 & d_2 \end{pmatrix},$$

we define

$$\begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} + \begin{pmatrix} a_2 & b_2 \\ c_2 & d_2 \end{pmatrix} = \begin{pmatrix} a_1 + a_2 & b_1 + b_2 \\ c_1 + c_2 & d_1 + d_2 \end{pmatrix}$$

and

$$\begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} \begin{pmatrix} a_2 & b_2 \\ c_2 & d_2 \end{pmatrix} = \begin{pmatrix} a_1a_2 + b_1c_2 & a_1b_2 + b_1d_2 \\ c_1a_2 + d_1c_2 & c_1b_2 + d_1d_2 \end{pmatrix}.$$

As written, the formula for matrix multiplication is extremely difficult to remember. It becomes easier to understand and remember if we label the entries differently. Suppose then that

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} b_{1,1} & b_{1,2} \\ b_{2,1} & b_{2,2} \end{pmatrix}.$$

In other words, we let $a_{i,j}$ be the entry of A that is in row i and column j . Although the double subscripts may look confusing at first, always remember that the first number is the row (so how far down we go) while the second is the column (so how far across we go). In this notation, we then have

$$AB = \begin{pmatrix} a_{1,1}b_{1,1} + a_{1,2}b_{2,1} & a_{1,1}b_{1,2} + a_{1,2}b_{2,2} \\ a_{2,1}b_{1,1} + a_{2,2}b_{2,1} & a_{2,1}b_{1,2} + a_{2,2}b_{2,2} \end{pmatrix}.$$

Looking through this table, we see that the entry in row i and column j is

$$a_{i,1}b_{1,j} + a_{i,2}b_{2,j},$$

which is just the dot product of row i of A with column j of B . For example, to determine the entry in row 2 and column 1 of AB , just take the dot product of row 2 of A with column 1 of B . To see this in action, consider the following matrix product:

$$\begin{pmatrix} 5 & 1 \\ -1 & 3 \end{pmatrix} \begin{pmatrix} 2 & 7 \\ 3 & -6 \end{pmatrix}.$$

To compute the result in the upper left-hand corner, we take the dot product of the first row of the matrix on the left and the first column of the matrix on the right. Thus, the entry in the upper left-hand corner is $5 \cdot 2 + 1 \cdot 3 = 13$. In general, we have

$$\begin{aligned} \begin{pmatrix} 5 & 1 \\ -1 & 3 \end{pmatrix} \begin{pmatrix} 2 & 7 \\ 3 & -6 \end{pmatrix} &= \begin{pmatrix} 5 \cdot 2 + 1 \cdot 3 & 5 \cdot 7 + 1 \cdot (-6) \\ (-1) \cdot 2 + 3 \cdot 3 & (-1) \cdot 7 + 3 \cdot (-6) \end{pmatrix} \\ &= \begin{pmatrix} 13 & 29 \\ 7 & -25 \end{pmatrix}. \end{aligned}$$

Proposition 2.6.2. *Let $T_1: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $T_2: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be linear transformations. We have the following:*

1. $[T_1 + T_2] = [T_1] + [T_2]$.
2. $[T_1 \circ T_2] = [T_1] \cdot [T_2]$.

Proof. The proof appears in the discussion before the definition of matrix addition and multiplication, and in fact we *defined* these operations so that this result would be true. \square

We also can define the product of a number and a matrix. The idea is that if $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a linear transformation, and $r \in \mathbb{R}$, then $r \cdot T$ is a linear transformation by Proposition 2.4.7.

Definition 2.6.3. *Given a matrix*

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

and an $r \in \mathbb{R}$, we define

$$r \cdot \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} ra & rb \\ rc & rd \end{pmatrix}.$$

Proposition 2.6.4. *Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation and let $r \in \mathbb{R}$. We have $[r \cdot T] = r \cdot [T]$.*

Proof. Exercise. □

Let's see matrix multiplication in action. Suppose that we consider the following function from \mathbb{R}^2 to \mathbb{R}^2 . Given a point, we first rotate the point counterclockwise by 45° , and then we project the result onto the line $y = 2x$. This procedure describes a function, which is a linear transformation because it is a composition of two linear transformations (the projection and the rotation). To determine the standard matrix of this linear transformation, we can simply multiply the two individual matrices, but we have to be careful about the order. Since function composition happens from right to left (i.e. in $f \circ g$ we first apply g and then apply f), we should put the rotation matrix on the right because we apply it first. Thus, the matrix of the linear transformation that comes from this composition is:

$$\begin{pmatrix} \frac{1}{5} & \frac{2}{5} \\ \frac{2}{5} & \frac{4}{5} \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} = \begin{pmatrix} \frac{3}{5\sqrt{2}} & \frac{1}{5\sqrt{2}} \\ \frac{6}{5\sqrt{2}} & \frac{2}{5\sqrt{2}} \end{pmatrix}.$$

Suppose we perform the operations in the other order. In other words, suppose that we take a point in \mathbb{R}^2 , project it onto the line $y = 2x$, and then rotate the result counterclockwise by 45° . The result is a linear transformation whose standard matrix is

$$\begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \frac{1}{5} & \frac{2}{5} \\ \frac{2}{5} & \frac{4}{5} \end{pmatrix} = \begin{pmatrix} \frac{-1}{5\sqrt{2}} & \frac{-2}{5\sqrt{2}} \\ \frac{3}{5\sqrt{2}} & \frac{6}{5\sqrt{2}} \end{pmatrix}.$$

Notice that we arrived at different answers depending on the order in which we multiplied the matrices! This may seem jarring at first, but it makes sense if you think about the product as the composition of linear transformations. Geometrically, it's clear that if we first rotate and then project onto $y = 2x$, then we obtain a different result than if we first project onto $y = 2x$ and then rotate (in particular, in the latter case it is unlikely that we end up on $y = 2x$).

To see another example of how to interpret matrix multiplication, let's go back and look at the linear transformation from the introduction. In that case, we looked at the transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by

$$T \left(\begin{pmatrix} x \\ y \end{pmatrix} \right) = \begin{pmatrix} x - y \\ x + y \end{pmatrix}.$$

The standard matrix of T is

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

Notice that

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

Now the second matrix on the left is just $[R_{\frac{\pi}{4}}]$, while the first represents scaling by a factor of $\sqrt{2}$. Thus, we can now confidently assert what seemed geometrically reasonable, i.e. that this transformation rotates the plane counterclockwise around the origin by $\frac{\pi}{4}$, and then scales by a factor of $\sqrt{2}$. Notice that we also have

$$\begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \sqrt{2} & 0 \\ 0 & \sqrt{2} \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$$

which also seems geometrically reasonable (if we first rotate and then scale, then we obtain the same result as if we first scale and then rotate).

Let's explore matrix multiplication in the context of two rotations. Notice that for any $\alpha, \beta \in \mathbb{R}$, if we rotate counterclockwise by an angle of β , and then rotate counterclockwise by an angle of α , this is the same as rotating counterclockwise once by an angle of $\alpha + \beta$. Thus, we must have

$$R_\alpha \circ R_\beta = R_{\alpha+\beta}$$

for all $\alpha, \beta \in \mathbb{R}$. It follows that we must have

$$[R_\alpha] \cdot [R_\beta] = [R_{\alpha+\beta}]$$

for all $\alpha, \beta \in \mathbb{R}$. Computing the left-hand side, we obtain

$$\begin{aligned} [R_\alpha] \cdot [R_\beta] &= \begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} \cos \beta & -\sin \beta \\ \sin \beta & \cos \beta \end{pmatrix} \\ &= \begin{pmatrix} \cos \alpha \cos \beta - \sin \alpha \sin \beta & -\cos \alpha \sin \beta - \sin \alpha \cos \beta \\ \sin \alpha \cos \beta + \cos \alpha \sin \beta & -\sin \alpha \sin \beta + \cos \alpha \cos \beta \end{pmatrix}. \end{aligned}$$

Since the right-hand side is

$$[R_{\alpha+\beta}] = \begin{pmatrix} \cos(\alpha + \beta) & -\sin(\alpha + \beta) \\ \sin(\alpha + \beta) & \cos(\alpha + \beta) \end{pmatrix},$$

and we must have $[R_\alpha] \cdot [R_\beta] = [R_{\alpha+\beta}]$, we conclude that

$$\begin{aligned} \cos(\alpha + \beta) &= \cos \alpha \cos \beta - \sin \alpha \sin \beta \\ \sin(\alpha + \beta) &= \sin \alpha \cos \beta + \cos \alpha \sin \beta \end{aligned}$$

for all $\alpha, \beta \in \mathbb{R}$. Thus, we have derived these trigonometric identities using Linear Algebra!

We end this section by cataloging the algebraic properties of matrices and vectors. In all these cases, it is possible to see why these results are true by simply playing around with the formulas. However, in many cases, we can understand why they are true by interpreting matrices as coding linear transformations. We first work with matrix-vector products.

Proposition 2.6.5. *Let A and B be 2×2 matrices.*

1. *For all $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$, we have $A(\vec{v}_1 + \vec{v}_2) = A\vec{v}_1 + A\vec{v}_2$.*
2. *For all $\vec{v} \in \mathbb{R}^2$ and $c \in \mathbb{R}$, we have $A(c\vec{v}) = c \cdot A\vec{v}$.*
3. *For all $\vec{v} \in \mathbb{R}^2$, we have $(A + B)\vec{v} = A\vec{v} + B\vec{v}$.*
4. *For all $\vec{v} \in \mathbb{R}^2$, we have $A(B\vec{v}) = (AB)\vec{v}$.*

Proof.

1. Let T be the unique linear transformation with $[T] = A$. We have

$$\begin{aligned} A(\vec{v}_1 + \vec{v}_2) &= T(\vec{v}_1 + \vec{v}_2) \\ &= T(\vec{v}_1) + T(\vec{v}_2) && \text{(since } T \text{ is a linear transformation)} \\ &= A\vec{v}_1 + A\vec{v}_2. \end{aligned}$$

2. Let T be the unique linear transformation with $[T] = A$. We have

$$\begin{aligned} A(c\vec{v}) &= T(c\vec{v}) \\ &= c \cdot T(\vec{v}) && \text{(since } T \text{ is a linear transformation)} \\ &= c \cdot A\vec{v}. \end{aligned}$$

3. Let S and T be the unique linear transformations with $[S] = A$ and $[T] = B$. We then have $[S + T] = A + B$, so

$$\begin{aligned}(A + B)\vec{v} &= (S + T)(\vec{v}) \\ &= S(\vec{v}) + T(\vec{v}) && \text{(by definition)} \\ &= A\vec{v} + B\vec{v}.\end{aligned}$$

4. Let S and T be the unique linear transformations with $[S] = A$ and $[T] = B$. We then have $[S \circ T] = AB$, so

$$\begin{aligned}A(B\vec{v}) &= S(T(\vec{v})) \\ &= (S \circ T)(\vec{v}) && \text{(by definition)} \\ &= (AB)\vec{v}.\end{aligned}$$

□

We now develop the algebraic properties of matrix-matrix products. In some cases, it is easy to do this through simple algebra. However, in other cases, the algebra becomes extremely messy. For example, we will argue that if A , B , and C are 2×2 matrices, then $(AB)C = A(BC)$. If we try to do this by “opening up” the matrices and following the formulas, it gets hairy pretty fast (although it is possible with patience). However, we can interpret this problem differently in a way that eliminates the calculation. Remember that matrix multiplication corresponds to function composition, and we know that function composition is associative (i.e. that $f \circ (g \circ h) = (f \circ g) \circ h$) by Proposition 1.6.5. Thus, $(AB)C = A(BC)$ should follow. We work out the details in part 3 of the next proposition.

Proposition 2.6.6. *Let A , B , and C be 2×2 matrices. We have the following:*

1. $A + B = B + A$.
2. $A + (B + C) = (A + B) + C$.
3. $(AB)C = A(BC)$.
4. $A(B + C) = AB + AC$.
5. $(A + B)C = AC + BC$.
6. For all $r \in \mathbb{R}$, we have $A(r \cdot B) = r \cdot (AB) = (r \cdot A)B$.

Proof.

1. Although it is possible to prove this by interpreting the matrices as coding linear transformations, in this case we can just compute easily. Suppose that

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} b_{1,1} & b_{1,2} \\ b_{2,1} & b_{2,2} \end{pmatrix}.$$

We then have

$$\begin{aligned}
 A + B &= \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix} + \begin{pmatrix} b_{1,1} & b_{1,2} \\ b_{2,1} & b_{2,2} \end{pmatrix} \\
 &= \begin{pmatrix} a_{1,1} + b_{1,1} & a_{1,2} + b_{1,2} \\ a_{2,1} + b_{2,1} & a_{2,2} + b_{2,2} \end{pmatrix} \\
 &= \begin{pmatrix} b_{1,1} + a_{1,1} & b_{1,2} + a_{1,2} \\ b_{2,1} + a_{2,1} & b_{2,2} + a_{2,2} \end{pmatrix} \\
 &= \begin{pmatrix} b_{1,1} & b_{1,2} \\ b_{2,1} & b_{2,2} \end{pmatrix} + \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix} \\
 &= B + A.
 \end{aligned}$$

2. This is similar to (1), except we use the fact that $+$ is associative on \mathbb{R} rather than commutative.
3. Let R , S and T be the unique linear transformations with $[R] = A$, $[S] = B$, and $[T] = C$. We know that $(R \circ S) \circ T = R \circ (S \circ T)$ by Proposition 1.6.5. Since these linear transformations are equal as functions, we know that they have the same standard matrices, so $[(R \circ S) \circ T] = [R \circ (S \circ T)]$. Now

$$\begin{aligned}
 [(R \circ S) \circ T] &= [R \circ S] \cdot [T] \\
 &= ([R] \cdot [S]) \cdot [T] \\
 &= (AB)C.
 \end{aligned}$$

Similarly, we have

$$\begin{aligned}
 [R \circ (S \circ T)] &= [R] \cdot [S \circ T] \\
 &= [R] \cdot ([S] \cdot [T]) \\
 &= A(BC).
 \end{aligned}$$

Therefore, we must have $(AB)C = A(BC)$.

4. Let R , S and T be the unique linear transformations with $[R] = A$, $[S] = B$, and $[T] = C$. We claim that $R \circ (S + T) = (R \circ S) + (R \circ T)$. To see this, let $\vec{v} \in \mathbb{R}^2$ be arbitrary. We have

$$\begin{aligned}
 (R \circ (S + T))(\vec{v}) &= R((S + T)(\vec{v})) \\
 &= R(S(\vec{v}) + T(\vec{v})) && \text{(by definition)} \\
 &= R(S(\vec{v})) + R(T(\vec{v})) && \text{(since } R \text{ is a linear transformation)} \\
 &= (R \circ S)(\vec{v}) + (R \circ T)(\vec{v}) \\
 &= ((R \circ S) + (R \circ T))(\vec{v}) && \text{(by definition).}
 \end{aligned}$$

Since $\vec{v} \in \mathbb{R}^2$ was arbitrary, we conclude that $R \circ (S + T) = (R \circ S) + (R \circ T)$. It follows that

$$[R] \cdot ([S] + [T]) = [R] \cdot [S] + [R] \cdot [T],$$

and hence $A(B + C) = AB + AC$.

5. The proof is similar to (4).

6. Let S and T be the unique linear transformations with $[S] = A$, $[T] = B$. We claim that $S \circ (r \cdot T) = r \cdot (S \circ T)$. To see this, let $\vec{v} \in \mathbb{R}^2$ be arbitrary. We have

$$\begin{aligned} (S \circ (r \cdot T))(\vec{v}) &= S((r \cdot T)(\vec{v})) \\ &= S(r \cdot T(\vec{v})) && \text{(by definition)} \\ &= r \cdot (S(T(\vec{v}))) && \text{(since } S \text{ is a linear transformation)} \\ &= r \cdot (S \circ T)(\vec{v}) && \text{(by definition).} \end{aligned}$$

Since $\vec{v} \in \mathbb{R}^2$ was arbitrary, we conclude that $S \circ (r \cdot T) = r \cdot (S \circ T)$. It follows that

$$[S] \cdot [r \cdot T] = r \cdot ([S] \cdot [T]),$$

and hence $A(r \cdot B) = r \cdot (AB)$. The proof of the other equality is similar.

□

Notice one key algebraic property that is missing: the commutative law of multiplication. In general, we can have $AB \neq BA$ for 2×2 matrices A and B ! We saw an example of this above with projections and rotations. For an easier example, notice that

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

while

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

We can interpret this geometrically as follows. The matrix

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

gives a horizontal shear transformation of the plane, while the matrix

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

gives a vertical shear transformation of the plane. Now when we compute

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

we are taking the composition that comes from first shearing the plane horizontally, and then shearing the plane vertically. Notice that if we do this, then \vec{e}_1 is sent to \vec{e}_1 by the horizontal shear, and then this is fed into the vertical shear to produce

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

which is the first column of the product above. In contrast, when we compute

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

we are taking the composition that comes from first shearing the plane vertically, and then shearing the plane horizontally. If we do this, then \vec{e}_1 is sent to

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

by the vertical shear, then this is fed into the horizontal shear to produce

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

which is the first column of the product above. In other words, we get a different result if we first horizontally shear the plane and then vertically shear it, versus if we first vertically shear the plane and then horizontally shear it.

Definition 2.6.7. *We define two special matrices.*

- We let 0 denote the 2×2 matrix of all zeros, i.e.

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

We have to distinguish between the number 0 and the matrix 0 from context.

- We let $I = [id]$ where $id: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the function given by $id(\vec{v}) = \vec{v}$ for all $\vec{v} \in \mathbb{R}^2$. Thus,

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

by Proposition 2.5.7.

We end with a few properties demonstrating that these two matrices act like the numbers 0 and 1. All of the proofs can be carried out using simple calculations (although it is also possible to do them theoretically), so will be left as exercises.

Proposition 2.6.8. *Let A be a 2×2 matrix.*

1. $A + 0 = A = 0 + A$.
2. $A + (-1) \cdot A = 0 = (-1) \cdot A + A$.
3. $A \cdot 0 = 0 = 0 \cdot A$.
4. $A \cdot I = A = I \cdot A$.

Taken together, we can interpret Proposition 2.6.6 and Proposition 2.6.8 as saying that addition and multiplication of matrices behave quite similarly to addition and multiplication of numbers. Thinking of matrices as “abstract numbers” is often a fruitful approach as we will see, but keep in mind that not all of the properties of numbers carry over to matrices. For example, we have seen that it is possible that AB does not equal BA . We will also encounter a few other instances where matrices differ from numbers. For example, 1 plays the role of I for multiplication in the real numbers, and every nonzero number has a multiplicative inverse. In other words, for every $a \in \mathbb{R}$ with $a \neq 0$, there exists $b \in \mathbb{R}$ with $ab = 1$. Is the same true for matrices? In other words, if A is a 2×2 matrix with $A \neq 0$, can we find a 2×2 matrix B with $AB = I$? Since multiplication of matrices is still a little mysterious, the answer is not obvious. We take up this question and related issues in the next section.

2.7 Range, Null Space, and Inverses

Recall that if $f: A \rightarrow B$ is a function, then

$$\text{range}(f) = \{b \in B : \text{There exists } a \in A \text{ with } f(a) = b\}.$$

We begin by studying the ranges of linear transformations $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$. It turns out that we can get a nice description of the range of such a function by looking at its standard matrix. Remember that one way to think about the matrix-vector product is that

$$\begin{aligned} \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} &= \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix} \\ &= x \cdot \begin{pmatrix} a \\ c \end{pmatrix} + y \cdot \begin{pmatrix} b \\ d \end{pmatrix}. \end{aligned}$$

Thus, it appears that if we have one fixed matrix, and use it to hit all possible vectors in \mathbb{R}^2 , then we will obtain all possible linear combinations of (the vectors represented by) columns of the matrix, i.e. we will obtain the span of the columns. We now formalize this into the following result.

Proposition 2.7.1. *Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. Let \vec{u}_1 be the first column of $[T]$, and let \vec{u}_2 be the second column of $[T]$. We then have that $\text{range}(T) = \text{Span}(\vec{u}_1, \vec{u}_2)$.*

Proof. Fix $a, b, c, d \in \mathbb{R}$ with

$$[T] = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

We then have

$$\vec{u}_1 = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} b \\ d \end{pmatrix}.$$

We argue that $\text{range}(T) = \text{Span}(\vec{u}_1, \vec{u}_2)$ is true by doing a double containment proof.

We first show that $\text{range}(T) \subseteq \text{Span}(\vec{u}_1, \vec{u}_2)$. Let $\vec{w} \in \text{range}(T)$ be arbitrary. By definition, we can fix $\vec{v} \in \mathbb{R}^2$ with $T(\vec{v}) = \vec{w}$. Fix $x, y \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

We then have

$$\begin{aligned} \vec{w} &= T(\vec{v}) \\ &= \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \\ &= x \cdot \begin{pmatrix} a \\ c \end{pmatrix} + y \cdot \begin{pmatrix} b \\ d \end{pmatrix} \\ &= x \cdot \vec{u}_1 + y \cdot \vec{u}_2. \end{aligned}$$

Since $x, y \in \mathbb{R}$, it follows that $\vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2)$. Since $\vec{w} \in \text{range}(T)$ was arbitrary, we conclude that $\text{range}(T) \subseteq \text{Span}(\vec{u}_1, \vec{u}_2)$.

We now show that $\text{Span}(\vec{u}_1, \vec{u}_2) \subseteq \text{range}(T)$. Let $\vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2)$ be arbitrary. By definition, we can fix

$r_1, r_2 \in \mathbb{R}$ with $\vec{w} = r_1 \cdot \vec{u}_1 + r_2 \cdot \vec{u}_2$. Notice that

$$\begin{aligned}\vec{w} &= r_1 \cdot \begin{pmatrix} a \\ c \end{pmatrix} + r_2 \cdot \begin{pmatrix} b \\ d \end{pmatrix} \\ &= \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} \\ &= T \left(\begin{pmatrix} r_1 \\ r_2 \end{pmatrix} \right).\end{aligned}$$

Since $r_1, r_2 \in \mathbb{R}$, it follows that $\vec{w} \in \text{range}(T)$. Since $\vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2)$ was arbitrary, we conclude that $\text{Span}(\vec{u}_1, \vec{u}_2) \subseteq \text{range}(T)$. \square

For example, if $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the linear transformation with

$$[T] = \begin{pmatrix} 5 & 6 \\ 1 & 2 \end{pmatrix},$$

then we have

$$\text{range}(T) = \text{Span} \left(\begin{pmatrix} 5 \\ 1 \end{pmatrix}, \begin{pmatrix} 6 \\ 2 \end{pmatrix} \right) = \mathbb{R}^2,$$

where the latter inequality follows from the fact that $5 \cdot 2 - 6 \cdot 1 = 4$ is nonzero. Thus, T is surjective.

In contrast, consider the linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with

$$[T] = \begin{pmatrix} 1 & 3 \\ 4 & 12 \end{pmatrix}.$$

We have

$$\text{range}(T) = \text{Span} \left(\begin{pmatrix} 1 \\ 4 \end{pmatrix}, \begin{pmatrix} 3 \\ 12 \end{pmatrix} \right) = \text{Span} \left(\begin{pmatrix} 1 \\ 4 \end{pmatrix} \right),$$

where the latter equality follows from Proposition 2.3.8 and the fact that

$$\begin{pmatrix} 3 \\ 12 \end{pmatrix} = 3 \cdot \begin{pmatrix} 1 \\ 4 \end{pmatrix},$$

so

$$\begin{pmatrix} 3 \\ 12 \end{pmatrix} \in \text{Span} \left(\begin{pmatrix} 1 \\ 4 \end{pmatrix} \right).$$

Therefore, we have that $\text{range}(T) \neq \mathbb{R}^2$ by Proposition 2.3.3, and so T is not surjective.

We now define another important subset associated to a linear transformation.

Definition 2.7.2. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. We define

$$\text{Null}(T) = \{\vec{v} \in \mathbb{R}^2 : T(\vec{v}) = \vec{0}\}.$$

We call $\text{Null}(T)$ the null space of T (or the kernel of T).

In contrast to $\text{range}(T)$, which consists of the possible *outputs* that are actually hit, we have that $\text{Null}(T)$ is the set of *inputs* of T that get sent to $\vec{0}$. In other words, $\text{Null}(T)$ collects into one set all of the inputs elements from \mathbb{R}^2 that are “trivialized” or “killed” by T . Notice that we always have that $\vec{0} \in \text{Null}(T)$ because $T(\vec{0}) = \vec{0}$ by Proposition 2.4.2.

For a specific example, suppose that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the linear transformation with

$$[T] = \begin{pmatrix} 5 & 6 \\ 1 & 2 \end{pmatrix}.$$

We claim that $\text{Null}(T) = \{\vec{0}\}$. To see this, notice first that $\vec{0} \in \text{Null}(T)$ from above. Suppose now that $\vec{v} \in \text{Null}(T)$ is arbitrary. Fix $x, y \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

Notice that

$$\begin{aligned} T(\vec{v}) &= \begin{pmatrix} 5 & 6 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \\ &= \begin{pmatrix} 5x + 6y \\ x + 2y \end{pmatrix}. \end{aligned}$$

Since we are assuming that $\vec{v} \in \text{Null}(T)$, we know that $T(\vec{v}) = \vec{0} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$. Thus, both of the following are true:

$$\begin{aligned} 5x + 6y &= 0 \\ x + 2y &= 0. \end{aligned}$$

From here, it is possible to work through the algebra to conclude that both $x = 0$ and $y = 0$. Alternatively, we can just apply Proposition 2.1.1 (since $5 \cdot 2 - 6 \cdot 1 = 4$ is nonzero) to conclude that there is a unique pair (x, y) that satisfies both of these equations, and since $(0, 0)$ is a solution, we must have $x = 0$ and $y = 0$. In either case, we conclude that $\vec{v} = \vec{0}$. Since $\vec{v} \in \text{Null}(T)$ was arbitrary, we conclude that $\text{Null}(T) \subseteq \{\vec{0}\}$. Combining this with the reverse containment, we conclude that $\text{Null}(T) = \{\vec{0}\}$.

In general, however, $\text{Null}(T)$ can contain vectors other than $\vec{0}$. For example, consider the linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with

$$[T] = \begin{pmatrix} 1 & 3 \\ 4 & 12 \end{pmatrix}.$$

In this case, we claim that

$$\text{Null}(T) = \text{Span} \left(\begin{pmatrix} -3 \\ 1 \end{pmatrix} \right).$$

Although it is not obvious how to come up with such a guess at this point, we will see techniques to help us later. Regardless, we can prove this equality directly now. We first show that

$$\text{Span} \left(\begin{pmatrix} -3 \\ 1 \end{pmatrix} \right) \subseteq \text{Null}(T).$$

Let

$$\vec{v} \in \text{Span} \left(\begin{pmatrix} -3 \\ 1 \end{pmatrix} \right)$$

be arbitrary. By definition, we can fix $c \in \mathbb{R}$ with

$$\vec{v} = c \cdot \begin{pmatrix} -3 \\ 1 \end{pmatrix} = \begin{pmatrix} -3c \\ c \end{pmatrix}.$$

We then have

$$\begin{aligned} T(\vec{v}) &= T\left(\begin{pmatrix} -3c \\ c \end{pmatrix}\right) \\ &= \begin{pmatrix} 1 & 3 \\ 4 & 12 \end{pmatrix} \begin{pmatrix} -3c \\ c \end{pmatrix} \\ &= \begin{pmatrix} -3c + 3c \\ -12c + 12c \end{pmatrix} \\ &= \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \end{aligned}$$

so $\vec{v} \in \text{Null}(T)$. Therefore,

$$\text{Span}\left(\begin{pmatrix} -3 \\ 1 \end{pmatrix}\right) \subseteq \text{Null}(T).$$

For the reverse containment, let $\vec{v} \in \text{Null}(T)$ be arbitrary, and fix $x, y \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

Since $\vec{v} \in \text{Null}(T)$, we have $T(\vec{v}) = \vec{0}$, hence

$$\begin{pmatrix} 1 & 3 \\ 4 & 12 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

which implies that

$$\begin{pmatrix} x + 3y \\ 4x + 12y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

Therefore, we must have $x + 3y = 0$ (and also $4x + 12y = 0$, but this is not important to us here). From this, we conclude that $x = -3y$, so

$$\vec{v} = \begin{pmatrix} -3y \\ y \end{pmatrix} = y \cdot \begin{pmatrix} -3 \\ 1 \end{pmatrix}.$$

Since $y \in \mathbb{R}$, we conclude that

$$\vec{v} \in \text{Span}\left(\begin{pmatrix} -3 \\ 1 \end{pmatrix}\right).$$

As $\vec{v} \in \text{Null}(T)$ was arbitrary, we conclude that

$$\text{Null}(T) \subseteq \text{Span}\left(\begin{pmatrix} -3 \\ 1 \end{pmatrix}\right).$$

We have shown both containments, so we can conclude that

$$\text{Null}(T) = \text{Span}\left(\begin{pmatrix} -3 \\ 1 \end{pmatrix}\right).$$

The previous two examples illustrate a phenomenon that we now prove in general.

Theorem 2.7.3. *Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation, and fix $a, b, c, d \in \mathbb{R}$ with*

$$[T] = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

We have the following cases:

1. If $ad - bc \neq 0$, then $\text{Null}(T) = \{\vec{0}\}$ and $\text{range}(T) = \mathbb{R}^2$.
2. If all of a, b, c, d equal 0, then $\text{Null}(T) = \mathbb{R}^2$ and $\text{range}(T) = \{\vec{0}\}$.
3. If $ad - bc = 0$ and at least one of a, b, c, d is nonzero, then there exist nonzero $\vec{u}, \vec{w} \in \mathbb{R}^2$ with $\text{Null}(T) = \text{Span}(\vec{u})$ and $\text{range}(T) = \text{Span}(\vec{w})$.

Proof.

1. Suppose first that $ad - bc \neq 0$. We first show that $\text{Null}(T) = \{\vec{0}\}$. Notice that $T(\vec{0}) = \vec{0}$ by Proposition 2.4.2, so $\vec{0} \in \text{Null}(T)$. Now let $\vec{v} \in \text{Null}(T)$ be arbitrary. Fix $x, y \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

Since $T(\vec{v}) = \vec{0}$, it follows that

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

Thus, both of the following are true:

$$\begin{aligned} ax + by &= 0 \\ cx + dy &= 0. \end{aligned}$$

Since $ad - bc \neq 0$, we can apply Proposition 2.1.1 to conclude that there is a unique pair (x, y) that satisfies both of these equations, and since $(0, 0)$ is a solution, we must have $x = 0$ and $y = 0$. In either case, we conclude that $\vec{v} = \vec{0}$. Therefore, $\text{Null}(T) = \{\vec{0}\}$.

We now show that $\text{range}(T) = \mathbb{R}^2$. Using Proposition 2.7.1, we know that

$$\text{range}(T) = \text{Span} \left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix} \right).$$

Since $ad - bc \neq 0$ by assumption, we know from Theorem 2.3.10 that

$$\text{Span} \left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix} \right) = \mathbb{R}^2.$$

Combining these two equalities, it follows that $\text{range}(T) = \mathbb{R}^2$.

2. Suppose that all of a, b, c, d are 0. In this case, we have that $T(\vec{v}) = \vec{0}$ for all $\vec{v} \in \mathbb{R}^2$, so $\text{Null}(T) = \mathbb{R}^2$ and $\text{range}(T) = \{\vec{0}\}$.
3. Finally, suppose that $ad - bc = 0$, but at least one of a, b, c, d is nonzero. For notation in the rest of the argument, let \vec{u}_1 be the first column of $[T]$ and let \vec{u}_2 be the second column of $[T]$, so

$$\vec{u}_1 = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} b \\ d \end{pmatrix}.$$

We now have a couple of cases, depending on whether one of these two vectors equals $\vec{0}$:

- *Case 1:* Suppose that $\vec{u}_1 = \vec{0}$, i.e. that both $a = 0$ and $c = 0$. In this case, we must have that at least one of b or d is nonzero, and hence $\vec{u}_2 \neq \vec{0}$. We first claim that $\text{Null}(T) = \text{Span}(\vec{e}_1)$. Notice

that if $x \in \mathbb{R}$ is arbitrary, then

$$\begin{aligned} T\left(\begin{pmatrix} x \\ 0 \end{pmatrix}\right) &= \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} ax \\ cx \end{pmatrix} \\ &= \begin{pmatrix} 0 \\ 0 \end{pmatrix} \end{aligned}$$

because $a = 0$ and $c = 0$, so

$$\begin{pmatrix} x \\ 0 \end{pmatrix} \in \text{Null}(T).$$

Since $x \in \mathbb{R}$ was arbitrary, we conclude that $\text{Span}(\vec{e}_1) \subseteq \text{Null}(T)$. For the reverse containment, let $\vec{v} \in \text{Null}(T)$ be arbitrary. Fix $x, y \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

Since $T(\vec{v}) = \vec{0}$, it follows that

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

Since $a = 0$ and $c = 0$, this implies that

$$\begin{pmatrix} by \\ dy \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

so $by = 0$ and $dy = 0$. Now since at least one of b or d is nonzero, we can divide the corresponding equation (i.e. either $by = 0$ or $dy = 0$), by it to conclude that $y = 0$. Therefore,

$$\vec{v} = \begin{pmatrix} x \\ 0 \end{pmatrix} = x \cdot \vec{e}_1,$$

and hence $\vec{v} \in \text{Span}(\vec{e}_1)$. Since $\vec{v} \in \text{Null}(T)$ was arbitrary, it follows that $\text{Null}(T) \subseteq \text{Span}(\vec{e}_1)$. Combining this with the above containment, we conclude that $\text{Null}(T) = \text{Span}(\vec{e}_1)$.

We now claim that $\text{range}(T) = \text{Span}(\vec{u}_2)$. To see this, notice that $\text{range}(T) = \text{Span}(\vec{u}_1, \vec{u}_2)$ by Proposition 2.7.1. Now we know that $\text{Span}(\vec{u}_1, \vec{u}_2) = \text{Span}(\vec{u}_2, \vec{u}_1)$ by Proposition 2.3.9. Now $\vec{u}_1 = \vec{0}$ by assumption, and $\vec{0} \in \text{Span}(\vec{u}_2)$ trivially, so $\text{Span}(\vec{u}_2, \vec{u}_1) = \text{Span}(\vec{u}_2)$ by Proposition 2.3.8. Putting it all together, we conclude that $\text{range}(T) = \text{Span}(\vec{u}_2)$.

- *Case 2:* Suppose that $\vec{u}_2 = \vec{0}$. By a similar argument to Case 1, we then have that $\text{Null}(T) = \text{Span}(\vec{e}_2)$ and $\text{range}(T) = \text{Span}(\vec{u}_1)$.
- *Case 3:* Suppose then that both \vec{u}_1 and \vec{u}_2 are nonzero. Since $ad - bc = 0$, we may use Theorem 2.3.10 to conclude that there exists $r \in \mathbb{R}$ with $\vec{u}_2 = r \cdot \vec{u}_1$. We then have $b = ra$ and $d = rc$. We claim that

$$\text{Null}(T) = \text{Span}\left(\begin{pmatrix} -r \\ 1 \end{pmatrix}\right).$$

For any $x \in \mathbb{R}$, we have

$$\begin{aligned} T\left(\begin{pmatrix} -rx \\ x \end{pmatrix}\right) &= \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} -rx \\ x \end{pmatrix} \\ &= \begin{pmatrix} a & ra \\ c & rc \end{pmatrix} \begin{pmatrix} -rx \\ x \end{pmatrix} \\ &= \begin{pmatrix} 0 \\ 0 \end{pmatrix} \end{aligned}$$

so

$$\begin{pmatrix} -rx \\ x \end{pmatrix} \in \text{Null}(T).$$

Since $x \in \mathbb{R}$ was arbitrary, it follows that

$$\text{Span}\left(\begin{pmatrix} -r \\ 1 \end{pmatrix}\right) \subseteq \text{Null}(T).$$

For the reverse containment, let $\vec{v} \in \text{Null}(T)$ be arbitrary. Fix $x, y \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

Since $T(\vec{v}) = \vec{0}$, it follows that

$$\begin{pmatrix} a & ra \\ c & rc \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

so

$$\begin{pmatrix} ax + ray \\ cx + rcy \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

Thus, we have both $ax + ray = 0$ and $cx + rcy = 0$, so $a(x + ry) = 0$ and $c(x + ry) = 0$. Now $\vec{u}_1 \neq 0$, so at least one of a or c is nonzero, and hence we can divide the corresponding equation by this nonzero value to conclude that $x + ry = 0$. From here, it follows that $x = -ry$, so

$$\vec{v} = \begin{pmatrix} -ry \\ y \end{pmatrix} = y \cdot \begin{pmatrix} -r \\ 1 \end{pmatrix},$$

and hence

$$\vec{v} \in \text{Span}\left(\begin{pmatrix} -r \\ 1 \end{pmatrix}\right).$$

Since $\vec{v} \in \text{Null}(T)$ was arbitrary, it follows that

$$\text{Null}(T) \subseteq \text{Span}\left(\begin{pmatrix} -r \\ 1 \end{pmatrix}\right).$$

Combining this with the above containment, we conclude that

$$\text{Null}(T) = \text{Span}\left(\begin{pmatrix} -r \\ 1 \end{pmatrix}\right).$$

We now show that $\text{range}(T) = \text{Span}(\vec{u}_1)$. We know that $\text{range}(T) = \text{Span}(\vec{u}_1, \vec{u}_2)$ by Proposition 2.7.1. Now since $u_2 = r \cdot \vec{u}_1$, we know that $\vec{u}_2 \in \text{Span}(\vec{u}_1)$, and hence $\text{Span}(\vec{u}_1, \vec{u}_2) = \text{Span}(\vec{u}_1)$ by Proposition 2.3.8. Putting it all together, we conclude that $\text{range}(T) = \text{Span}(\vec{u}_1)$.

□

Notice that our theorem establishes a complementary relationship between $\text{Null}(T)$ and $\text{range}(T)$, despite the fact that one of these sets concerns inputs and the other is about outputs. In the case where $\text{Null}(T)$ collapses to a point, we have that $\text{range}(T)$ is the 2-dimensional plane \mathbb{R}^2 . In the case where $\text{Null}(T)$ is the 2-dimensional plane \mathbb{R}^2 , we have that $\text{range}(T)$ collapses to a point. Finally, in all other cases, both $\text{Null}(T)$ and $\text{range}(T)$ consist of 1-dimensional lines through the origin. We will drastically generalize this phenomenon later in the Rank-Nullity Theorem.

Recall that $\text{Null}(T)$ collects together all of the vectors that T sends to $\vec{0}$. Now if $\text{Null}(T)$ has at least 2 elements, then there are at least two elements that go to $\vec{0}$, so T will not be injective. Perhaps surprisingly, the converse of this works as well. In other words, simply by knowing that only one element goes to $\vec{0}$, we can conclude that every element of \mathbb{R}^2 is hit by at most one element. Of course, for general functions $f: A \rightarrow B$, knowledge of the number of elements that hit a given $b \in B$ provides little to no information about how many elements hit a given $c \in B$ with $c \neq b$. However, linear transformations are very special.

Proposition 2.7.4. *Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. We then have that T is injective if and only if $\text{Null}(T) = \{\vec{0}\}$.*

Proof. Suppose first that T is injective. We show that $\text{Null}(T) = \{\vec{0}\}$ by giving a double containment proof. First recall that $T(\vec{0}) = \vec{0}$ by Proposition 2.4.2, so $\vec{0} \in \text{Null}(T)$, and hence $\{\vec{0}\} \subseteq \text{Null}(T)$. Now let $\vec{v} \in \text{Null}(T)$ be arbitrary. By definition, we then have that $T(\vec{v}) = \vec{0}$. Since we also have $T(\vec{0}) = \vec{0}$ by Proposition 2.4.2, it follows that $T(\vec{v}) = T(\vec{0})$. Using the fact that T is injective, we can conclude that $\vec{v} = \vec{0}$. Therefore, $\text{Null}(T) \subseteq \{\vec{0}\}$.

Suppose now that $\text{Null}(T) = \{\vec{0}\}$. Let $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ be arbitrary with $T(\vec{v}_1) = T(\vec{v}_2)$. We then have

$$\begin{aligned} T(\vec{v}_1 - \vec{v}_2) &= T(\vec{v}_1) - T(\vec{v}_2) && \text{(since } T \text{ is a linear transformation)} \\ &= \vec{0}, \end{aligned}$$

so $\vec{v}_1 - \vec{v}_2 \in \text{Null}(T)$. Since $\text{Null}(T) = \{\vec{0}\}$, we conclude that $\vec{v}_1 - \vec{v}_2 = \vec{0}$, and hence $\vec{v}_1 = \vec{v}_2$. Since $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ were arbitrary with $T(\vec{v}_1) = T(\vec{v}_2)$, it follows that T is injective. □

Corollary 2.7.5. *Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation, and fix $a, b, c, d \in \mathbb{R}$ with*

$$[T] = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

The following are equivalent:

1. $ad - bc \neq 0$.
2. T is bijective.
3. T is injective.
4. T is surjective.

Proof. We again prove these equivalences in a circle (as we did in Theorem 2.3.10).

- (1) implies (2): Suppose that $ad - bc \neq 0$. Using Theorem 2.7.3, we then have that $\text{Null}(T) = \{\vec{0}\}$ and $\text{range}(T) = \mathbb{R}^2$. Since $\text{Null}(T) = \{\vec{0}\}$, we may use Proposition 2.7.4 to conclude that T is injective. Since $\text{range}(T) = \mathbb{R}^2$, we also know that T is surjective. Thus, T is bijective.
- (2) implies (3): Immediate from the definition of bijective.

- (3) implies (4): Suppose that T is injective. By Proposition 2.7.4, we know that $\text{Null}(T) = \{\vec{0}\}$. Looking at the three cases in Theorem 2.7.3, it follows that we must have $\text{range}(T) = \mathbb{R}^2$, so T is surjective.
- (4) implies (1): Suppose that T is surjective, so that $\text{range}(T) = \mathbb{R}^2$. Looking at the three cases in Theorem 2.7.3 and using Proposition 2.3.3 (that the span of one vector can not be all of \mathbb{R}^2), it follows that we must have $ad - bc \neq 0$.

□

Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the unique linear transformation with

$$[T] = \begin{pmatrix} 3 & 4 \\ 1 & 2 \end{pmatrix}.$$

Since $3 \cdot 2 - 4 \cdot 1 = 2$ is nonzero, we may apply Corollary 2.7.5 to conclude that T is bijective. In particular, we can consider the function $T^{-1}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, although it might not be clear how to compute it. For example, what does

$$T^{-1} \left(\begin{pmatrix} 1 \\ 5 \end{pmatrix} \right)$$

equal? In order to determine this, we want to find $x, y \in \mathbb{R}$ with

$$T \left(\begin{pmatrix} x \\ y \end{pmatrix} \right) = \begin{pmatrix} 1 \\ 5 \end{pmatrix}.$$

Thus, we want to find $x, y \in \mathbb{R}$ with

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

which is the same as solving the following system:

$$\begin{array}{rcrcrcrcl} 3x & + & 4y & = & 1 \\ x & + & 2y & = & 5. \end{array}$$

Working through algebra, or applying Proposition 2.1.1, we determine that the unique solution is $(-9, 7)$. Thus, we have

$$T \left(\begin{pmatrix} -9 \\ 7 \end{pmatrix} \right) = \begin{pmatrix} 1 \\ 5 \end{pmatrix}$$

and hence

$$T^{-1} \left(\begin{pmatrix} 1 \\ 5 \end{pmatrix} \right) = \begin{pmatrix} -9 \\ 7 \end{pmatrix}.$$

At this point, two natural questions arise. First, is there a better way to compute the values of $T^{-1}(\vec{w})$ for a given $\vec{w} \in \mathbb{R}^2$, perhaps by making use of the standard matrix $[T]$? Second, although we've defined a function $T^{-1}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, is it even a linear transformation so that we can code it as a matrix? We begin by handling the latter question.

Proposition 2.7.6. *Suppose that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a bijective linear transformation. We then have that the function $T^{-1}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a linear transformation.*

Proof. We check that T^{-1} preserves the two operations:

- Let $\vec{w}_1, \vec{w}_2 \in \mathbb{R}^2$ be arbitrary. We need to show that $T^{-1}(\vec{w}_1 + \vec{w}_2) = T^{-1}(\vec{w}_1) + T^{-1}(\vec{w}_2)$. Notice that

$$\begin{aligned} T(T^{-1}(\vec{w}_1) + T^{-1}(\vec{w}_2)) &= T(T^{-1}(\vec{w}_1)) + T(T^{-1}(\vec{w}_2)) \\ &= \vec{w}_1 + \vec{w}_2. \end{aligned}$$

Thus, we have found the (necessarily unique because T is bijective) element of \mathbb{R}^2 that T maps to $\vec{w}_1 + \vec{w}_2$, so we can conclude that

$$T^{-1}(\vec{w}_1 + \vec{w}_2) = T^{-1}(\vec{w}_1) + T^{-1}(\vec{w}_2).$$

Therefore, the function T^{-1} preserves addition.

- Let $\vec{w} \in \mathbb{R}^2$ and $c \in \mathbb{R}$ be arbitrary. We need to show that $T^{-1}(c \cdot \vec{w}) = c \cdot T^{-1}(\vec{w})$. Notice that

$$\begin{aligned} T(c \cdot T^{-1}(\vec{w})) &= c \cdot T(T^{-1}(\vec{w})) \\ &= c \cdot \vec{w}. \end{aligned}$$

Thus, we have found the (necessarily unique) element of \mathbb{R}^2 that T maps to $c \cdot \vec{w}$, so we can conclude that

$$T^{-1}(c \cdot \vec{w}) = c \cdot T^{-1}(\vec{w}).$$

Therefore, the function T^{-1} preserves scalar multiplication.

Since T^{-1} preserves both addition and scalar multiplication, it follows that T^{-1} is a linear transformation. \square

Suppose then that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a bijective linear transformation. We now know that T^{-1} is also a linear transformation, and we have $T \circ T^{-1} = id$ and $T^{-1} \circ T = id$ by Proposition 1.7.5 (where we are just writing id in place of $id_{\mathbb{R}^2}$ to keep the notation simple). Can we compute $[T^{-1}]$ from knowledge of $[T]$? Suppose that

$$[T] = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Since T is bijective, we know that $ad - bc \neq 0$ by Corollary 2.7.5. Now to determine $[T^{-1}]$, we need to calculate both $T^{-1}(\vec{e}_1)$ and $T^{-1}(\vec{e}_2)$. To determine $T^{-1}(\vec{e}_1)$, we need to find the unique vector $\vec{v} \in \mathbb{R}^2$ with $T(\vec{v}) = \vec{e}_1$, which amounts to finding $x, y \in \mathbb{R}$ with

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix},$$

which amounts to solving the following system:

$$\begin{aligned} ax + by &= 1 \\ cx + dy &= 0. \end{aligned}$$

Since $ad - bc \neq 0$, we can use Proposition 2.1.1 to conclude that this system has a unique solution, and it is

$$\left(\frac{d}{ad - bc}, \frac{-c}{ad - bc} \right).$$

Therefore, we have

$$T^{-1}(\vec{e}_1) = \begin{pmatrix} \frac{d}{ad - bc} \\ \frac{-c}{ad - bc} \end{pmatrix}.$$

Similarly, to determine $T^{-1}(\vec{e}_2)$, we need to solve the system

$$\begin{aligned} ax + by &= 0 \\ cx + dy &= 1. \end{aligned}$$

Since $ad - bc \neq 0$, we can use Proposition 2.1.1 to conclude that this system has a unique solution, and it is

$$\left(\frac{-b}{ad - bc}, \frac{a}{ad - bc} \right).$$

Therefore, we have

$$T^{-1}(\vec{e}_2) = \begin{pmatrix} \frac{-b}{ad - bc} \\ \frac{a}{ad - bc} \end{pmatrix}.$$

It follows that

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

These calculations give the following result.

Proposition 2.7.7. *Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a bijective linear transformation, and fix $a, b, c, d \in \mathbb{R}$ with*

$$[T] = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

We then have that $ad - bc \neq 0$, and

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

Proof. We carried out this calculation before the proposition. □

Suppose that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a bijective linear transformation. Let $A = [T]$ and let $B = [T^{-1}]$. Since $T \circ T^{-1} = id$, we know from Proposition 2.6.2 that we must have $AB = I$, where

$$I = [id] = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Similarly, since $T^{-1} \circ T = id$, we must have $BA = I$. Intuitively, if we think about matrices as certain abstract numbers, with the matrix I playing the role of the number 1 (the identity for multiplication), then A and B are multiplicative inverses of each other. We can codify this idea for matrices with the following definition.

Definition 2.7.8. *Let A be a 2×2 matrix. We say that A is invertible if there exists a 2×2 matrix B with both $AB = I$ and $BA = I$.*

Since a matrix A is invertible if and only if the corresponding linear transformation T has an inverse, which is the same as saying that T is bijective, we obtain the following result.

Proposition 2.7.9. *Let A be a 2×2 matrix, say*

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

We then have that A is invertible if and only if $ad - bc \neq 0$. In this case, A has a unique inverse, given by

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

Notice the pattern here. We are swapping the position of a and d along the main diagonal, negating the two values off of the main diagonal, and then dividing everything by the term $ad - bc$.

Notation 2.7.10. *If A is an invertible 2×2 matrix, we denote its unique inverse by A^{-1} .*

Let's take a moment to reflect on our work connecting matrix inverses with function inverses. As we just mentioned, a matrix has an inverse exactly when the corresponding linear transformation has an inverse (as a function). Using our theory about when a function has an inverse, and combining this with our hard work to determine when a linear transformation is bijective, we were able to obtain the simple condition that a matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

is invertible exactly when $ad - bc \neq 0$. Without all of our background work on functions, if we were handed a random 2×2 matrix A , and we wanted to know whether we could find a matrix B with $AB = I$, then we would have had to work directly with the strange formula for matrix multiplication. However, by shifting our perspective from the formula for matrix multiplication to composition of functions, we were able to attack this problem from a more powerful and enlightening perspective.

Now that we have done the hard work, we can greatly simplify some computations. For example, if $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the unique linear transformation with

$$[T] = \begin{pmatrix} 3 & 4 \\ 1 & 2 \end{pmatrix},$$

then we computed

$$T^{-1} \left(\begin{pmatrix} 1 \\ 5 \end{pmatrix} \right)$$

above by solving a system. We can attack this problem differently now. Since $3 \cdot 2 - 4 \cdot 1 = 2$ is nonzero, we know that T is bijective. Furthermore, we know that T^{-1} is a linear transformation by Proposition 2.7.6 and that

$$\begin{aligned} [T^{-1}] &= \frac{1}{2} \begin{pmatrix} 2 & -4 \\ -1 & 3 \end{pmatrix} \\ &= \begin{pmatrix} 1 & -2 \\ -\frac{1}{2} & \frac{3}{2} \end{pmatrix} \end{aligned}$$

by Proposition 2.7.7. Thus, we can just compute

$$\begin{aligned} T^{-1} \left(\begin{pmatrix} 1 \\ 5 \end{pmatrix} \right) &= \begin{pmatrix} 1 & -2 \\ -\frac{1}{2} & \frac{3}{2} \end{pmatrix} \begin{pmatrix} 1 \\ 5 \end{pmatrix} \\ &= \begin{pmatrix} -9 \\ 7 \end{pmatrix}. \end{aligned}$$

Furthermore, now that we know $[T^{-1}]$, we can compute T^{-1} on other inputs easily. For example, we have

$$\begin{aligned} T^{-1} \left(\begin{pmatrix} 6 \\ 10 \end{pmatrix} \right) &= \begin{pmatrix} 1 & -2 \\ -\frac{1}{2} & \frac{3}{2} \end{pmatrix} \begin{pmatrix} 6 \\ 10 \end{pmatrix} \\ &= \begin{pmatrix} -14 \\ 12 \end{pmatrix}. \end{aligned}$$

We can also turn this whole idea on its head to use inverses of matrices to solve linear systems. For example, suppose that we want to solve the system

$$\begin{aligned} 4x + 6y &= 7 \\ 3x + 5y &= 1 \end{aligned}$$

Let

$$A = \begin{pmatrix} 4 & 6 \\ 3 & 5 \end{pmatrix}$$

and let

$$\vec{b} = \begin{pmatrix} 7 \\ 1 \end{pmatrix}.$$

Thus, we are trying to find $x, y \in \mathbb{R}$ with

$$\begin{pmatrix} 4 & 6 \\ 3 & 5 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 7 \\ 1 \end{pmatrix}$$

i.e. we are trying to find those $\vec{v} \in \mathbb{R}^2$ such that $A\vec{v} = \vec{b}$. Hopefully, this reminds you of an equation of the form $ax = b$ in \mathbb{R} . In that case, we divide both sides by a to solve the equation, as long as a is not zero. Now we can't divide by a matrix, but dividing by a is the same as multiplying by $\frac{1}{a} = a^{-1}$. If \vec{v} satisfies $A\vec{v} = \vec{b}$, then we can multiply both sides on the left by A^{-1} to conclude that $A^{-1}(A\vec{v}) = A^{-1}\vec{b}$. Using the algebraic properties from the last section, we have $(A^{-1}A)\vec{v} = A^{-1}\vec{b}$, so $I\vec{v} = A^{-1}\vec{b}$, and hence $\vec{v} = A^{-1}\vec{b}$. We can also plug back in to conclude that $A^{-1}\vec{b}$ really is a solution.

Now in our case, we have that A is invertible because $4 \cdot 5 - 6 \cdot 3 = 2$ is nonzero, and we have

$$\begin{aligned} A^{-1} &= \frac{1}{4 \cdot 5 - 6 \cdot 3} \begin{pmatrix} 5 & -6 \\ -3 & 4 \end{pmatrix} \\ &= \begin{pmatrix} \frac{5}{2} & -3 \\ -\frac{3}{2} & 2 \end{pmatrix}. \end{aligned}$$

Therefore, a solution to our system is given by

$$\begin{aligned} \vec{v} &= A^{-1}\vec{b} \\ &= \begin{pmatrix} \frac{5}{2} & -3 \\ -\frac{3}{2} & 2 \end{pmatrix} \begin{pmatrix} 7 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} \frac{29}{2} \\ -\frac{17}{2} \end{pmatrix}. \end{aligned}$$

We end with two important properties of inverses. Both are reasonably natural, but be careful with the second! You might have guessed that $(AB)^{-1}$ would equal $A^{-1}B^{-1}$. However, remember that the order in which we multiply matrices matters! Think about this in terms of socks and shoes. If you first put your socks on, and then put your shoes on, then to undo this you should first take your shoes off, and then take your socks off.

Proposition 2.7.11. *We have the following.*

1. *If A is an invertible 2×2 matrix, then A^{-1} is invertible and $(A^{-1})^{-1} = A$.*
2. *If A and B are both invertible 2×2 matrices, then AB is invertible and $(AB)^{-1} = B^{-1}A^{-1}$.*

Proof.

1. Let A be an invertible 2×2 matrix. We then have that

$$A \cdot A^{-1} = I \quad \text{and} \quad A^{-1} \cdot A = I.$$

Looking at these two equalities from a different perspective, they say that A works as an inverse for the matrix A^{-1} on both sides. In other words, we have shown the existence of an inverse for A^{-1} , namely A . It follows that A^{-1} is invertible and $(A^{-1})^{-1} = A$.

2. Let A and B both be invertible 2×2 matrices. We claim that $B^{-1}A^{-1}$ works as an inverse for AB . To show this, it suffices to multiply these in both directions and check that we get I . We have

$$\begin{aligned} (AB) \cdot (B^{-1}A^{-1}) &= A(BB^{-1})A^{-1} \\ &= AIA^{-1} \\ &= AA^{-1} \\ &= I \end{aligned}$$

and also

$$\begin{aligned} (B^{-1}A^{-1}) \cdot (AB) &= B^{-1}(A^{-1}A)B \\ &= B^{-1}IB \\ &= B^{-1}B \\ &= I. \end{aligned}$$

Thus, we have shown the existence of an inverse for AB , namely $B^{-1}A^{-1}$. It follows that AB is invertible and that $(AB)^{-1} = B^{-1}A^{-1}$.

□

Chapter 3


Coordinates and Eigenvectors in Two Dimensions

3.1 Coordinates and Change of Basis

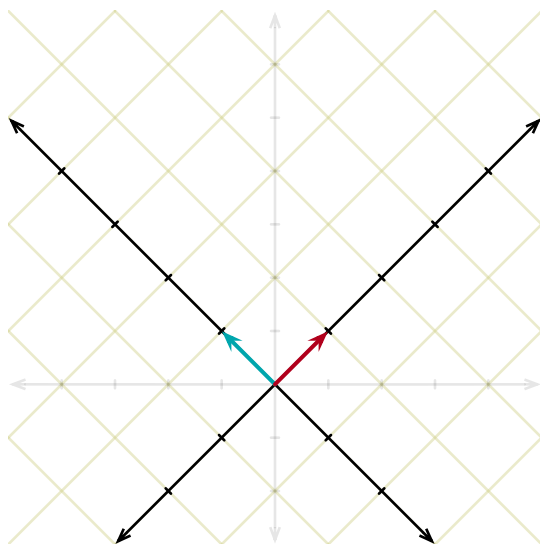
Suppose that $\alpha = (\vec{u}_1, \vec{u}_2)$ is a basis for \mathbb{R}^2 . By Theorem 2.3.10, we know that for every $\vec{v} \in \mathbb{R}^2$, there exists a unique pair of numbers $c_1, c_2 \in \mathbb{R}$ with $\vec{v} = c_1\vec{u}_1 + c_2\vec{u}_2$. Geometrically, we are saying that if we form axes in \mathbb{R}^2 based on \vec{u}_1 and \vec{u}_2 , and form new “graph paper” based on the grid system created in this way, then every point in \mathbb{R}^2 can be described uniquely by these scaling factors along the axes. In more detail, think about the line through the origin and \vec{u}_1 as a new “ x -axis” and the line through the origin and \vec{u}_2 as a new “ y -axis”. For example, if

$$\alpha = \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right),$$

then α is a basis of \mathbb{R}^2 (because $1 \cdot 1 - 1 \cdot (-1) \neq 0$), and the new grid system we drew in Section 2.3 gave us the following picture:



To really emphasize this new way of visualizing the plane, think of replacing the standard axes with axes along each of the basis vectors (we dim the old axes to help focus):



We can then measure any point in \mathbb{R}^2 through using these new axes and the corresponding new grid system. We will often want to calculate and understand the scaling factors (i.e. how far along each of the new axes we need) to describe a given vector. To this end, we introduce the following definition.

Definition 3.1.1. Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 . We define a function $\text{Coord}_\alpha: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ as follows. Given $\vec{v} \in \mathbb{R}^2$, let $c_1, c_2 \in \mathbb{R}$ be the unique scalars such that $\vec{v} = c_1\vec{u}_1 + c_2\vec{u}_2$, and define

$$\text{Coord}_\alpha(\vec{v}) = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}.$$

For example, if we continue working with the basis

$$\alpha = \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right),$$

then we saw in Section 2.3 that

$$\begin{pmatrix} 2 \\ 5 \end{pmatrix} = (7/2) \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + (3/2) \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix},$$

so

$$\text{Coord}_\alpha \left(\begin{pmatrix} 2 \\ 5 \end{pmatrix} \right) = \begin{pmatrix} 7/2 \\ 3/2 \end{pmatrix},$$

and we also saw that

$$\begin{pmatrix} -4 \\ 2 \end{pmatrix} = (-1) \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + 3 \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix},$$

so

$$\text{Coord}_\alpha \left(\begin{pmatrix} -4 \\ 2 \end{pmatrix} \right) = \begin{pmatrix} -1 \\ 3 \end{pmatrix},$$

Let's try another example. Let

$$\alpha = \left(\begin{pmatrix} 5 \\ 1 \end{pmatrix}, \begin{pmatrix} -6 \\ -3 \end{pmatrix} \right).$$

Notice that α is a basis of \mathbb{R}^2 because $5 \cdot (-3) - 1 \cdot (-6) = -9 \neq 0$. We then have Coord_α is a function from \mathbb{R}^2 to \mathbb{R}^2 , which takes as input a vector $\vec{v} \in \mathbb{R}^2$, and outputs the element of \mathbb{R}^2 coding how \vec{v} is measured

in the coordinate system based on α , i.e. how much along the \vec{u}_1 -axis we need to go (measured in steps of length equal to the length of \vec{u}_1) and how much along the \vec{u}_2 -axis we need to go. Suppose that we want to compute

$$\text{Coord}_\alpha \left(\begin{pmatrix} 12 \\ 0 \end{pmatrix} \right).$$

In order to do this, we need to find the unique $c_1, c_2 \in \mathbb{R}$ with

$$\begin{pmatrix} 12 \\ 0 \end{pmatrix} = c_1 \cdot \begin{pmatrix} 5 \\ 1 \end{pmatrix} + c_2 \cdot \begin{pmatrix} -6 \\ -3 \end{pmatrix}.$$

We find these values by solving the following linear system:

$$\begin{array}{rcl} 5x & - & 6y = 12 \\ x & - & 3y = 0. \end{array}$$

Solving this system directly, or just applying Proposition 2.1.1, we see that the unique solution is $(4, \frac{4}{3})$. Therefore, we have

$$\begin{pmatrix} 12 \\ 0 \end{pmatrix} = 4 \cdot \begin{pmatrix} 5 \\ 1 \end{pmatrix} + (4/3) \cdot \begin{pmatrix} -6 \\ -3 \end{pmatrix}.$$

and hence

$$\text{Coord}_\alpha \left(\begin{pmatrix} 12 \\ 0 \end{pmatrix} \right) = \begin{pmatrix} 4 \\ 4/3 \end{pmatrix}.$$

Using this idea, we can obtain a formula for the function Coord_α by applying Proposition 2.1.1.

Proposition 3.1.2. *Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 . Fix $a, b, c, d \in \mathbb{R}$ with*

$$\vec{u}_1 = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} b \\ d \end{pmatrix}.$$

For any $j, k \in \mathbb{R}$, we have

$$\begin{aligned} \text{Coord}_\alpha \left(\begin{pmatrix} j \\ k \end{pmatrix} \right) &= \begin{pmatrix} \frac{dj-bk}{ad-bc} \\ \frac{ak-cj}{ad-bc} \end{pmatrix} \\ &= \frac{1}{ad-bc} \cdot \begin{pmatrix} dj-bk \\ ak-cj \end{pmatrix} \end{aligned}$$

Proof. Let $j, k \in \mathbb{R}$ be arbitrary. To calculate

$$\text{Coord}_\alpha \left(\begin{pmatrix} j \\ k \end{pmatrix} \right),$$

we need to find the unique $r_1, r_2 \in \mathbb{R}$ with

$$\begin{pmatrix} j \\ k \end{pmatrix} = r_1 \cdot \begin{pmatrix} a \\ c \end{pmatrix} + r_2 \cdot \begin{pmatrix} b \\ d \end{pmatrix}.$$

Finding these values is the same thing as solving the following linear system:

$$\begin{array}{rcl} ax & + & by = j \\ cx & + & dy = k \end{array}$$

Since α is a basis of \mathbb{R}^2 , we know that $ad - bc \neq 0$ by Theorem 2.3.10. Using Proposition 2.1.1, we know that this system has a unique solution, namely

$$\left(\frac{dj - bk}{ad - bc}, \frac{ak - cj}{ad - bc} \right).$$

Therefore, we have

$$\begin{aligned} \text{Coord}_\alpha \left(\begin{pmatrix} j \\ k \end{pmatrix} \right) &= \begin{pmatrix} \frac{dj - bk}{ad - bc} \\ \frac{ak - cj}{ad - bc} \end{pmatrix} \\ &= \frac{1}{ad - bc} \cdot \begin{pmatrix} dj - bk \\ ak - cj \end{pmatrix}. \end{aligned}$$

□

Although this formula certainly works, and is often convenient for short calculations, it is typically best *not* to think of the function Coord_α as given by some ugly formula. We will discuss the advantages and disadvantages of viewing these coordinate changes geometrically, algebraically, and computationally throughout the course.

Proposition 3.1.3. *Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 . The function $\text{Coord}_\alpha: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a linear transformation.*

One way to see this is to use the formula we developed for Coord_α in Proposition 3.1.2.

Proof 1 of Proposition 3.1.3. Fix $a, b, c, d \in \mathbb{R}$ with

$$\vec{u}_1 = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} b \\ d \end{pmatrix}.$$

Since α is a basis of \mathbb{R}^2 , we know from Theorem 2.3.10 that $ad - bc \neq 0$. Now we know that Coord_α can be described by the following formula:

$$\text{Coord}_\alpha \left(\begin{pmatrix} x \\ y \end{pmatrix} \right) = \begin{pmatrix} \frac{dx - by}{ad - bc} \\ \frac{ay - cx}{ad - bc} \end{pmatrix}.$$

Rewriting this formula as

$$\text{Coord}_\alpha \left(\begin{pmatrix} x \\ y \end{pmatrix} \right) = \begin{pmatrix} \frac{d}{ad - bc} \cdot x + \frac{-b}{ad - bc} \cdot y \\ \frac{-c}{ad - bc} \cdot x + \frac{a}{ad - bc} \cdot y \end{pmatrix} = x \cdot \begin{pmatrix} \frac{d}{ad - bc} \\ \frac{-c}{ad - bc} \end{pmatrix} + y \cdot \begin{pmatrix} \frac{-b}{ad - bc} \\ \frac{a}{ad - bc} \end{pmatrix},$$

we can apply Proposition 2.4.3 to immediately conclude that Coord_α is a linear transformation. □

Although this works, there is a more elegant way to prove this result which avoids ugly formulas and “opening up” the vectors \vec{u}_1 and \vec{u}_2 .

Proof 2 of Proposition 3.1.3. We first check that Coord_α preserves addition. Let $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ be arbitrary. Since α is a basis of \mathbb{R}^2 , we can apply Theorem 2.3.10 and fix the unique pair of numbers $c_1, c_2 \in \mathbb{R}$ and the unique pair of numbers $d_1, d_2 \in \mathbb{R}$ with

$$\vec{v}_1 = c_1 \vec{u}_1 + c_2 \vec{u}_2$$

and

$$\vec{v}_2 = d_1 \vec{u}_1 + d_2 \vec{u}_2.$$

By definition of Coord_α , we have

$$\text{Coord}_\alpha(\vec{v}_1) = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} \quad \text{and} \quad \text{Coord}_\alpha(\vec{v}_2) = \begin{pmatrix} d_1 \\ d_2 \end{pmatrix}.$$

Notice that

$$\begin{aligned} \vec{v}_1 + \vec{v}_2 &= c_1\vec{u}_1 + c_2\vec{u}_2 + d_1\vec{u}_1 + d_2\vec{u}_2 \\ &= c_1\vec{u}_1 + d_1\vec{u}_1 + c_2\vec{u}_2 + d_2\vec{u}_2 \\ &= (c_1 + d_1)\vec{u}_1 + (c_2 + d_2)\vec{u}_2. \end{aligned}$$

Since $c_1 + d_1 \in \mathbb{R}$ and $c_2 + d_2 \in \mathbb{R}$, we have found the (necessarily unique by Theorem 2.3.10) pair of numbers that express $\vec{v}_1 + \vec{v}_2$ as a linear combination of \vec{u}_1 and \vec{u}_2 . Therefore, we have

$$\begin{aligned} \text{Coord}_\alpha(\vec{v}_1 + \vec{v}_2) &= \begin{pmatrix} c_1 + d_1 \\ c_2 + d_2 \end{pmatrix} \\ &= \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} + \begin{pmatrix} d_1 \\ d_2 \end{pmatrix} \\ &= \text{Coord}_\alpha(\vec{v}_1) + \text{Coord}_\alpha(\vec{v}_2). \end{aligned}$$

Since $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^2$ were arbitrary, it follows that Coord_α preserves addition.

We now check that Coord_α preserves scalar multiplication. Let $\vec{v} \in \mathbb{R}^2$ and $c \in \mathbb{R}$ be arbitrary. Since α is a basis of \mathbb{R}^2 , we can apply Theorem 2.3.10 and fix the unique pair of numbers $d_1, d_2 \in \mathbb{R}^2$ with

$$\vec{v} = d_1\vec{u}_1 + d_2\vec{u}_2$$

By definition of Coord_α , we have

$$\text{Coord}_\alpha(\vec{v}) = \begin{pmatrix} d_1 \\ d_2 \end{pmatrix}.$$

Notice that

$$\begin{aligned} c\vec{v} &= c \cdot (d_1\vec{u}_1 + d_2\vec{u}_2) \\ &= c \cdot (d_1\vec{u}_1) + c \cdot (d_2\vec{u}_2) \\ &= (cd_1)\vec{u}_1 + (cd_2)\vec{u}_2. \end{aligned}$$

Since $cd_1 \in \mathbb{R}$ and $cd_2 \in \mathbb{R}$, we have found the (necessarily unique by Theorem 2.3.10) pair of numbers that express $c\vec{v}$ as a linear combination of \vec{u}_1 and \vec{u}_2 . Therefore, we have

$$\begin{aligned} \text{Coord}_\alpha(c\vec{v}) &= \begin{pmatrix} cd_1 \\ cd_2 \end{pmatrix} \\ &= c \cdot \begin{pmatrix} d_1 \\ d_2 \end{pmatrix} \\ &= c \cdot \text{Coord}_\alpha(\vec{v}). \end{aligned}$$

Since $\vec{v} \in \mathbb{R}^2$ and $c \in \mathbb{R}$ were arbitrary, it follows that Coord_α preserves scalar multiplication.

We've shown that Coord_α preserves both addition and scalar multiplication, so Coord_α is a linear transformation. \square

For example, since we computed a formula for the functions Coord_α , we can immediately determine the standard matrix of these linear transformations.

Proposition 3.1.4. Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 . If

$$\vec{u}_1 = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} b \\ d \end{pmatrix},$$

then

$$[\text{Coord}_\alpha] = \begin{pmatrix} \frac{d}{ad-bc} & \frac{-b}{ad-bc} \\ \frac{-c}{ad-bc} & \frac{a}{ad-bc} \end{pmatrix}.$$

Thus, if we let

$$P = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

then $[\text{Coord}_\alpha] = P^{-1}$.

Proof. We proved in Proposition 3.1.3 that Coord_α is a linear transformation. We also developed a formula for Coord_α in Proposition 3.1.2:

$$\begin{aligned} \text{Coord}_\alpha \left(\begin{pmatrix} x \\ y \end{pmatrix} \right) &= \begin{pmatrix} \frac{dx-by}{ad-bc} \\ \frac{ay-cx}{ad-bc} \end{pmatrix} \\ &= \begin{pmatrix} \frac{d}{ad-bc} \cdot x + \frac{-b}{ad-bc} \cdot y \\ \frac{-c}{ad-bc} \cdot x + \frac{a}{ad-bc} \cdot y \end{pmatrix}. \end{aligned}$$

Therefore, using Proposition 2.5.8, we can immediately conclude that

$$[\text{Coord}_\alpha] = \begin{pmatrix} \frac{d}{ad-bc} & \frac{-b}{ad-bc} \\ \frac{-c}{ad-bc} & \frac{a}{ad-bc} \end{pmatrix}.$$

The last statement is now immediate from Proposition 2.7.9. □

For example, suppose that

$$\vec{u}_1 = \begin{pmatrix} 1 \\ 2 \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} 3 \\ 7 \end{pmatrix}.$$

We then have that

$$\begin{aligned} [\text{Coord}_\alpha] &= \begin{pmatrix} 1 & 3 \\ 2 & 7 \end{pmatrix}^{-1} \\ &= \frac{1}{7 \cdot 1 - (-3) \cdot (-2)} \begin{pmatrix} 7 & -3 \\ -2 & 1 \end{pmatrix} \\ &= \begin{pmatrix} 7 & -3 \\ -2 & 1 \end{pmatrix}, \end{aligned}$$

so we can compute values of Coord_α by simply multiplying by this matrix. In other words, we have

$$\begin{aligned} \text{Coord}_\alpha \left(\begin{pmatrix} 4 \\ 9 \end{pmatrix} \right) &= \begin{pmatrix} 7 & -3 \\ -2 & 1 \end{pmatrix} \begin{pmatrix} 4 \\ 9 \end{pmatrix} \\ &= \begin{pmatrix} 7 \cdot 4 + (-3) \cdot 9 \\ (-2) \cdot 4 + 1 \cdot 9 \end{pmatrix} \\ &= \begin{pmatrix} 1 \\ 1 \end{pmatrix}. \end{aligned}$$

Since we will be changing coordinates a lot in the coming sections, and since the notation $\text{Coord}_\alpha(\vec{v})$ will become cumbersome, we introduce some simplified notation.

Notation 3.1.5. Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 . Given $\vec{v} \in \mathbb{R}^2$, we use the notation $[\vec{v}]_\alpha$ as shorthand for $\text{Coord}_\alpha(\vec{v})$.

Thus, working with the basis

$$\alpha = \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right),$$

we can write

$$\left[\begin{pmatrix} 2 \\ 5 \end{pmatrix} \right]_\alpha = \begin{pmatrix} 7/2 \\ 3/2 \end{pmatrix}$$

and

$$\left[\begin{pmatrix} -4 \\ 2 \end{pmatrix} \right]_\alpha = \begin{pmatrix} -1 \\ 3 \end{pmatrix},$$

With this new notation, we can summarize Proposition 3.1.4 as follows. Given any basis $\alpha = (\vec{u}_1, \vec{u}_2)$ of \mathbb{R}^2 , say

$$\vec{u}_1 = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} b \\ d \end{pmatrix},$$

if we let

$$P = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

then we have

$$[\vec{v}]_\alpha = P^{-1}\vec{v}$$

for all $\vec{v} \in \mathbb{R}^2$. In hindsight, we can understand the appearance of the inverse of P both algebraically and geometrically:

- *Algebraic Approach:* Suppose that we have a basis α described above, are given some $\vec{v} \in \mathbb{R}^2$, and are trying to compute $[\vec{v}]_\alpha$. If $\vec{v} = \begin{pmatrix} j \\ k \end{pmatrix}$, then we are trying to find the unique $x, y \in \mathbb{R}$ with

$$x \cdot \begin{pmatrix} a \\ c \end{pmatrix} + y \cdot \begin{pmatrix} b \\ d \end{pmatrix} = \begin{pmatrix} j \\ k \end{pmatrix},$$

which means that we are trying to solve the following system of equations:

$$\begin{array}{rcrcrcrcrcrcrc} ax & + & by & = & j \\ cx & + & dy & = & k. \end{array}$$

Using our notation P above for the corresponding matrix, we are trying to find the unique vector $\vec{x} \in \mathbb{R}^2$ with $P\vec{x} = \vec{v}$. As described at the end of Section 2.7, the unique such vector is $\vec{x} = P^{-1}\vec{v}$. Thus, we have $[\vec{v}]_\alpha = P^{-1}\vec{v}$ for all $\vec{v} \in \mathbb{R}^2$.

- *Geometric Approach:* Consider the situation from the point of view of the new coordinate system. Suppose we have a vector $\vec{v} \in \mathbb{R}^2$, and we happen to know for some reason that

$$[\vec{v}]_\alpha = \begin{pmatrix} r_1 \\ r_2 \end{pmatrix}.$$

Geometrically, we are saying that you get \vec{v} by using r_1 and r_2 as the scaling factors on each of the axes in the new coordinate system. More formally, we are saying that

$$\begin{aligned}\vec{v} &= r_1 \cdot \begin{pmatrix} a \\ c \end{pmatrix} + r_2 \cdot \begin{pmatrix} b \\ d \end{pmatrix} \\ &= \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} r_1 \\ r_2 \end{pmatrix}.\end{aligned}$$

Thus, for any $\vec{v} \in \mathbb{R}^2$, we have $\vec{v} = P[\vec{v}]_\alpha$. Multiplying both sides by P^{-1} on the left, it follows that $[\vec{v}]_\alpha = P^{-1}\vec{v}$ for all $\vec{v} \in \mathbb{R}^2$.

Before moving on, let's take a moment to reflect on this situation and compare it to similar situations in the past. In Calculus, you saw a different way to describe points in \mathbb{R}^2 by using polar coordinates. For example, the point $(1, 1)$ in the Cartesian plane can be described by the polar point $(\sqrt{2}, \frac{\pi}{4})$. In other words, you already have some experience labeling points in the plane using a pair of numbers different from the usual Cartesian pair of numbers. We are doing the same thing now, except in this setting we now have infinitely many different coordinate changes: one for each choice of basis $\alpha = (\vec{u}_1, \vec{u}_2)$. However, in many ways our coordinate changes are much better behaved than the polar coordinate changes:

- Geometrically, the grid system formed by our new coordinates consists entirely of parallelograms. In contrast, in the polar system, the “regions” that we obtain are cut out between circles (for constant value of r) and rays from the origin (for constant values of θ). Thus, some of the boundaries in the polar system are curves rather than straight lines.
- Not only are the regions of the grid system in our new coordinates parallelograms, but in fact they consist of repeating parallelograms of precisely the same size and shape throughout the plane. In the polar system, the regions become larger (for fixed change in r and θ) as we move further from the origin. In other words, there is a distortion in the polar system that does not arise in our new systems. This “area distortion” is at the root of the r that must be inserted when performing double integrals in polar coordinates, because the regions become scaled by a factor of r depending on the distance from the origin. We will begin to discuss this general idea of *area distortion* when we talk about determinants.
- In our new coordinate systems, every point can be described by a *unique* pair of numbers. In contrast, every point has infinitely many polar descriptions. For example, in addition to representing $(1, 1)$ by the polar point $(\sqrt{2}, \frac{\pi}{4})$, we can also represent it by the polar point $(\sqrt{2}, \frac{9\pi}{4})$, or even by $(-\sqrt{2}, \frac{5\pi}{4})$. Now if we restrict to $r \geq 0$ and $0 \leq \theta < 2\pi$, then we fix most of these issues, but the origin can still be described in the infinitely many ways $(0, \theta)$ for any value of θ .

3.2 Matrices with Respect to Other Coordinates

Consider the unique linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with standard matrix

$$[T] = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}.$$

Unlike some of the simpler examples that we have seen so far, it is not immediately clear how this transformation acts geometrically on the plane. We can try to plug in points, and perhaps with enough such values we might see a pattern. However, it turns out that this linear transformation T can be best understood and visualized by *changing coordinates*. That is, instead of looking at the plane as demarcated using the usual Cartesian grid system based on \vec{e}_1 and \vec{e}_2 , we can get a much better sense of what is happening by looking through the eyes of somebody using a different grid system.

Recall that in our definition of the standard matrix of a linear transformation T , we coded the action of T based on the values on \vec{e}_1 and \vec{e}_2 . However, if (\vec{u}_1, \vec{u}_2) is a basis of \mathbb{R}^2 , then by Proposition 2.4.4, we know that T is completely determined by its action on \vec{u}_1 and \vec{u}_2 as well. Furthermore, in conjunction with Theorem 2.4.5, we can pick $T(\vec{u}_1)$ and $T(\vec{u}_2)$ arbitrarily and there will then be exactly one linear transformation sending \vec{u}_1 and \vec{u}_2 to these values. Imagine living a world that has a grid system based on \vec{u}_1 and \vec{u}_2 as the foundation. How would the people living in this world code T as a matrix?

For example, suppose that in this world they are using the two vectors

$$\vec{u}_1 = \begin{pmatrix} 3 \\ -1 \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} 2 \\ 1 \end{pmatrix}$$

as the basis of their coordinates. Notice that $3 \cdot 1 - (-1) \cdot 2 = 5$, which is nonzero, so (\vec{u}_1, \vec{u}_2) is indeed a basis of \mathbb{R}^2 . To us, these two vectors may look like a strange choice, but to the person living in this world they are simply their versions of \vec{e}_1 and \vec{e}_2 , and they view *our* choice of two vectors as bizarre. Now if the people in this world knew the values $T(\vec{u}_1)$ and $T(\vec{u}_2)$, then they could determine the value of T on everything else. More specifically, given $\vec{v} \in \mathbb{R}^2$, in order to determine $T(\vec{v})$, they would start by finding $c_1, c_2 \in \mathbb{R}$ with

$$\vec{v} = c_1 \vec{u}_1 + c_2 \vec{u}_2.$$

In other words, they would first determine $[\vec{v}]_\alpha$, the coordinates of \vec{v} in their system. Moreover, they would literally view \vec{v} as the vector

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix}$$

even though we would look at its entries differently (we would open up \vec{u}_1 and \vec{u}_2 , multiply by the constants, and add to obtain the standard coordinates of \vec{v}). From here, they would compute

$$\begin{aligned} T(\vec{v}) &= T(c_1 \vec{u}_1 + c_2 \vec{u}_2) \\ &= c_1 \cdot T(\vec{u}_1) + c_2 \cdot T(\vec{u}_2) \end{aligned}$$

in order to determine $T(\vec{v})$. Now in our above example with the given T , \vec{u}_1 , and \vec{u}_2 , then you might imagine that such a person would compute as follows:

$$\begin{aligned} T(\vec{u}_1) &= \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} 3 \\ -1 \end{pmatrix} = \begin{pmatrix} 5 \\ 1 \end{pmatrix} \\ T(\vec{u}_2) &= \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} 2 \\ 1 \end{pmatrix} = \begin{pmatrix} 5 \\ 4 \end{pmatrix} \end{aligned}$$

However, be careful here! We computed these values using the standard matrix of T , which is the matrix of T relative to *our* coordinates. Just like we do not yet know the matrix of the people working in a different grid system, they do not know our (standard) matrix, so they certainly would not compute in this way. Moreover, when we write the vectors on the right, we are really thinking about them as

$$\begin{pmatrix} 5 \\ 1 \end{pmatrix} = 5 \cdot \vec{e}_1 + 1 \cdot \vec{e}_2,$$

so by “opening up” these outputs we are looking at them through the tainted eyes of our coordinates!

Although the people in this world would not follow our computations, we can still figure out their answer. How? Well, they are looking to find $d_1, d_2 \in \mathbb{R}$ with

$$T(\vec{u}_1) = d_1 \vec{u}_1 + d_2 \vec{u}_2,$$

i.e. from our perspective they are trying to calculate $[T(\vec{u}_1)]_\alpha$. Notice that

$$\begin{pmatrix} 5 \\ 1 \end{pmatrix} = \frac{3}{5} \cdot \begin{pmatrix} 3 \\ -1 \end{pmatrix} + \frac{8}{5} \cdot \begin{pmatrix} 2 \\ 1 \end{pmatrix},$$

or in other words

$$T(\vec{u}_1) = \frac{3}{5} \cdot \vec{u}_1 + \frac{8}{5} \cdot \vec{u}_2.$$

Thus, in their coordinates, they would view $T(\vec{u}_1)$ as the vector

$$\begin{pmatrix} \frac{3}{5} \\ \frac{8}{5} \end{pmatrix}.$$

Now we also have

$$\begin{pmatrix} 5 \\ 4 \end{pmatrix} = -\frac{3}{5} \cdot \begin{pmatrix} 3 \\ -1 \end{pmatrix} + \frac{17}{5} \cdot \begin{pmatrix} 2 \\ 1 \end{pmatrix},$$

or in other words

$$T(\vec{u}_2) = -\frac{3}{5} \cdot \vec{u}_1 + \frac{17}{5} \cdot \vec{u}_2.$$

Thus, in their coordinates, they would view $T(\vec{u}_2)$ as the vector

$$\begin{pmatrix} -\frac{3}{5} \\ \frac{17}{5} \end{pmatrix}$$

Putting it all together, if the people in this world used their coordinates throughout, it looks like they would code T with the matrix

$$\begin{pmatrix} \frac{3}{5} & -\frac{3}{5} \\ \frac{8}{5} & \frac{17}{5} \end{pmatrix}.$$

Here is the general definition.

Definition 3.2.1. Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 , and let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. Fix $a, b, c, d \in \mathbb{R}$ with

$$[T(\vec{u}_1)]_\alpha = \begin{pmatrix} a \\ c \end{pmatrix}$$

and

$$[T(\vec{u}_2)]_\alpha = \begin{pmatrix} b \\ d \end{pmatrix}.$$

We define the matrix of T relative to α to be

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

We denote this matrix by $[T]_\alpha$. In other words, we let the first column of $[T]_\alpha$ be the coordinates of $T(\vec{u}_1)$ relative to α , and we let the second column of $[T]_\alpha$ be the coordinates of $T(\vec{u}_2)$ relative to α .

We now prove the following analogue of Proposition 2.5.4 which says that the matrix $[T]_\alpha$ really does code the transformation correctly from the point of view of people working with α -coordinates throughout. To see how these relate, notice that if we let $\varepsilon = (\vec{e}_1, \vec{e}_2)$, then we have $[\vec{v}]_\varepsilon = \vec{v}$ for all $\vec{v} \in \mathbb{R}^2$ and $[T]_\varepsilon = [T]$ for all linear transformations $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, so we can rewrite Proposition 2.5.4 as saying that $[T(\vec{v})]_\varepsilon = [T]_\varepsilon \cdot [\vec{v}]_\varepsilon$ for all linear transformations $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and all $\vec{v} \in \mathbb{R}^2$. We now generalize this to any coordinate system.

Proposition 3.2.2. *Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 , and let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. We have*

$$[T(\vec{v})]_\alpha = [T]_\alpha \cdot [\vec{v}]_\alpha$$

for all $\vec{v} \in \mathbb{R}^2$.

Proof. Fix $a, b, c, d \in \mathbb{R}$ with

$$[T]_\alpha = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Now let $\vec{v} \in \mathbb{R}^2$ be arbitrary. Fix $r_1, r_2 \in \mathbb{R}$ with

$$[\vec{v}]_\alpha = \begin{pmatrix} r_1 \\ r_2 \end{pmatrix}.$$

By definition, we then have $\vec{v} = r_1 \cdot \vec{u}_1 + r_2 \cdot \vec{u}_2$. Since T is a linear transformation, notice that

$$\begin{aligned} T(\vec{v}) &= T(r_1 \cdot \vec{u}_1 + r_2 \cdot \vec{u}_2) \\ &= r_1 \cdot T(\vec{u}_1) + r_2 \cdot T(\vec{u}_2). \end{aligned}$$

Now using the fact that Coord_α is a linear transformation from Proposition 3.1.3, we have

$$\begin{aligned} [T(\vec{v})]_\alpha &= [r_1 \cdot T(\vec{u}_1) + r_2 \cdot T(\vec{u}_2)]_\alpha \\ &= r_1 \cdot [T(\vec{u}_1)]_\alpha + r_2 \cdot [T(\vec{u}_2)]_\alpha \\ &= r_1 \cdot \begin{pmatrix} a \\ c \end{pmatrix} + r_2 \cdot \begin{pmatrix} b \\ d \end{pmatrix} \\ &= \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} \\ &= [T]_\alpha \cdot [\vec{v}]_\alpha. \end{aligned}$$

□

Let's see this result in action with the T and $\alpha = (\vec{u}_1, \vec{u}_2)$ that we have been using, i.e. with

$$[T] = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}, \quad \vec{u}_1 = \begin{pmatrix} 3 \\ -1 \end{pmatrix}, \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} 2 \\ 1 \end{pmatrix}.$$

Let

$$\vec{v} = \begin{pmatrix} 10 \\ 5 \end{pmatrix}.$$

In the standard coordinate system, we have

$$T(\vec{v}) = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} 10 \\ 5 \end{pmatrix} = \begin{pmatrix} 25 \\ 20 \end{pmatrix}.$$

Now notice that

$$\vec{v} = 0 \cdot \vec{u}_1 + 5 \cdot \vec{u}_2,$$

so in α coordinate system, we have

$$[\vec{v}]_\alpha = \begin{pmatrix} 0 \\ 5 \end{pmatrix}.$$

Thus, to determine $[T(\vec{v})]_\alpha$, people living in the α -coordinate world would compute

$$[T(\vec{v})]_\alpha = \begin{pmatrix} \frac{3}{5} & -\frac{3}{5} \\ \frac{8}{5} & \frac{17}{5} \end{pmatrix} \begin{pmatrix} 0 \\ 5 \end{pmatrix} = \begin{pmatrix} -3 \\ 17 \end{pmatrix}.$$

Don't be alarmed that these results look different! After all, we have

$$-3 \cdot \begin{pmatrix} 3 \\ -1 \end{pmatrix} + 17 \cdot \begin{pmatrix} 2 \\ 1 \end{pmatrix} = \begin{pmatrix} 25 \\ 20 \end{pmatrix} = 25 \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} + 20 \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

so we of course get the same answer, just viewed through the lens of different coordinate systems.

Let's recap the process for computing $[T]_\alpha$ given a basis $\alpha = (\vec{u}_1, \vec{u}_2)$. Compute $T(\vec{u}_1)$, but rather than use the two numbers inside this vector (which are really the coordinates relative to \vec{e}_1 and \vec{e}_2), instead find the unique $c_1, c_2 \in \mathbb{R}$ with

$$T(\vec{u}_1) = c_1 \vec{u}_1 + c_2 \vec{u}_2.$$

Similarly, find the unique $d_1, d_2 \in \mathbb{R}$ with

$$T(\vec{u}_2) = d_1 \vec{u}_1 + d_2 \vec{u}_2.$$

We then have that

$$[T]_\alpha = \begin{pmatrix} c_1 & d_1 \\ c_2 & d_2 \end{pmatrix}.$$

At this point, you may ask yourself a very reasonable question. We started with a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with

$$[T] = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}.$$

We looked at a new coordinate system α and found that

$$[T]_\alpha = \begin{pmatrix} \frac{3}{5} & -\frac{3}{5} \\ \frac{8}{5} & \frac{17}{5} \end{pmatrix}.$$

So we replaced a matrix that we did not understand geometrically by a much uglier matrix that we do not understand geometrically. What progress! In this case, it is correct to say that this coordinate system made things worse. However, let's try a different coordinate system. Suppose that we let

$$\vec{w}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad \text{and} \quad \vec{w}_2 = \begin{pmatrix} -1 \\ 1 \end{pmatrix},$$

and let $\beta = (\vec{w}_1, \vec{w}_2)$. We have that β is a basis of \mathbb{R}^2 because $1 \cdot 1 - 1 \cdot (-1) = 2$ is nonzero. Notice that

$$\begin{aligned} T(\vec{w}_1) &= \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} 3 \\ 3 \end{pmatrix} \\ &= 3 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + 0 \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix} \\ &= 3 \cdot \vec{w}_1 + 0 \cdot \vec{w}_2 \end{aligned}$$

so

$$[T(\vec{w}_1)]_\beta = \begin{pmatrix} 3 \\ 0 \end{pmatrix}.$$

Also, we have

$$\begin{aligned} T(\vec{w}_2) &= \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} -1 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} -1 \\ 1 \end{pmatrix} \\ &= 0 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + 1 \cdot \begin{pmatrix} -1 \\ 1 \end{pmatrix} \\ &= 0 \cdot \vec{w}_1 + 1 \cdot \vec{w}_2, \end{aligned}$$

so

$$[T(\vec{w}_2)]_\beta = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

It follows that

$$[T]_\beta = \begin{pmatrix} 3 & 0 \\ 0 & 1 \end{pmatrix}.$$

Now this is a much simpler matrix to think about! In the standard coordinates, this matrix expands the plane horizontally by a factor of 3 because

$$\begin{pmatrix} 3 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 3c_1 \\ c_2 \end{pmatrix}$$

for all $c_1, c_2 \in \mathbb{R}$. In other words, if a point is a distance of c_1 away from the y -axis, then it stays on the same horizontal line but is now 3 times as far from the y -axis. The same matrix-vector calculation applies here, but we just need to interpret it differently. Instead of thinking about c_1 and c_2 as the distances along the x -axis and y -axis that we travel (i.e. the amount we scale \vec{e}_1 and \vec{e}_2 by), we instead think about them as the amount we scale \vec{w}_1 and \vec{w}_2 by. In other words, we have

$$T(c_1 \cdot \vec{w}_1 + c_2 \cdot \vec{w}_2) = 3c_1 \cdot \vec{w}_1 + c_2 \cdot \vec{w}_2$$

for all $c_1, c_2 \in \mathbb{R}$. It follows that our transformation is an expansion by a factor of 3, but parallel to a different line than the x -axis. Notice that $\text{Span}(\vec{w}_1)$ is the solution set to $y = x$, and $\text{Span}(\vec{w}_2)$ is the solution set to $y = -x$. Thus, our grid system consists of lines parallel to $y = x$ or $y = -x$. From this perspective, our transformation takes a point, and triples the amount it has along the line $y = x$, while keeping the same amount it has along the line $y = -x$. Since $y = x$ and $y = -x$ are perpendicular, we can view T as simply expanding the plane by a factor of 3 in a perpendicular direction away from the line $y = -x$.

As we've just witnessed, by using a change in coordinates, we can sometime understand the action of a linear transformation more deeply. However, there are several questions that remain. How do we find a good coordinate system so that the matrix is "simple"? How are the matrices between different coordinate systems related? Can we compute these matrices more quickly? We will answer the latter two questions in this section, and begin to answer the first (and probably most important) question in the next section. It turns out that answering the first question definitively is quite difficult, and the situation in more than two dimensions is beyond the scope of a first course.

Suppose then that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a linear transformation. Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 . For some simplified notation without square brackets, let $A = [T]$, let $B = [T]_\alpha$, and let $C = [\text{Coord}_\alpha]$. How do

A and B relate to each other? The matrix C takes a vector in standard coordinates and converts it into α -coordinates, and we can then feed that output into B to get the result of $T(\vec{v})$ in α -coordinates (using Proposition 3.2.2). In other words, BC takes a vector \vec{v} and outputs $[T(\vec{v})]_\alpha$. For the other product, A takes a vector in standard coordinates and outputs $T(\vec{v})$ in standard coordinates, which we can feed into C to get the α coordinates of $T(\vec{v})$, i.e. to get $[T(\vec{v})]_\alpha$. Since these are the same, it looks like given any $\vec{v} \in \mathbb{R}^2$, we can follow the diagram below in either direction to arrive at the same result:

$$\begin{array}{ccccc} \vec{v} & \rightarrow & C & \rightarrow & [\vec{v}]_\alpha \\ \downarrow & & & & \downarrow \\ A & & & & B \\ \downarrow & & & & \downarrow \\ T(\vec{v}) & \rightarrow & C & \rightarrow & [T(\vec{v})]_\alpha \end{array}$$

In other words, it appears that $(CA)\vec{v} = (BC)\vec{v}$ for all $\vec{v} \in \mathbb{R}^2$. We now prove this.

Proposition 3.2.3. *Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 and let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. Let $A = [T]$, let $B = [T]_\alpha$, and let $C = [\text{Coord}_\alpha]$. We then have that $(CA)\vec{v} = (BC)\vec{v}$ for all $\vec{v} \in \mathbb{R}^2$.*

Proof. For any $\vec{v} \in \mathbb{R}^2$, we have

$$\begin{aligned} (CA)\vec{v} &= C(A\vec{v}) \\ &= C \cdot T(\vec{v}) \\ &= [T(\vec{v})]_\alpha \\ &= [T]_\alpha \cdot [\vec{v}]_\alpha && \text{(by Proposition 3.2.2)} \\ &= B \cdot [\vec{v}]_\alpha \\ &= B(C\vec{v}) \\ &= (BC)\vec{v}. \end{aligned}$$

□

In general, if two matrices act the same on all vectors, then they must be the same matrix.

Proposition 3.2.4. *Let A and B be 2×2 matrices. If $A\vec{v} = B\vec{v}$ for all $\vec{v} \in \mathbb{R}^2$, then $A = B$.*

Proof. Exercise. □

Corollary 3.2.5. *Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 and let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. Let $A = [T]$, let $B = [T]_\alpha$, and let $C = [\text{Coord}_\alpha]$. We then have that $CA = BC$.*

Proof. This follows immediately from Proposition 3.2.3 and Proposition 3.2.4. □

We can now use the formula for $[\text{Coord}_\alpha]$ that we developed in Proposition 3.1.4 in order to establish a straightforward way to compute $[T]_\alpha$ from $[T]$.

Proposition 3.2.6. *Let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 and let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. Fix $a, b, c, d \in \mathbb{R}$ with*

$$\vec{u}_1 = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} b \\ d \end{pmatrix}$$

and notice that $ad - bc \neq 0$ (because α is a basis of \mathbb{R}^2). Let

$$P = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

We then have that P is invertible and

$$[T]_\alpha = P^{-1}[T]P.$$

Proof. Let $A = [T]$, let $B = [T]_\alpha$, and let $C = [\text{Coord}_\alpha]$. We know Corollary 3.2.5 that $CA = BC$. Now Proposition 3.1.4 tells us that $C = P^{-1}$, so $P^{-1}A = BP^{-1}$. Multiplying both sides of this equation on the right by P , we conclude that $P^{-1}AP = B$, i.e. that $[T]_\alpha = P^{-1}[T]P$. \square

Let's take a step back and interpret what this is saying. Suppose that $\alpha = (\vec{u}_1, \vec{u}_2)$ and P are as in the setup of the proposition. We are claiming that

$$[T]_\alpha = \begin{pmatrix} a & b \\ c & d \end{pmatrix}^{-1} \cdot [T] \cdot \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

How does this work? Suppose that we are given a vector \vec{v} and its α -coordinates are

$$[\vec{v}]_\alpha = \begin{pmatrix} x \\ y \end{pmatrix}.$$

When we feed this vector $[\vec{v}]_\alpha$ as input to the matrix P on the right, we get as output

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = x \cdot \begin{pmatrix} a \\ c \end{pmatrix} + y \cdot \begin{pmatrix} b \\ d \end{pmatrix},$$

which when expanded out will give us the standard coordinates for \vec{v} . When we feed this result into the matrix $[T]$, we see that it gets transformed into $T(\vec{v})$ in standard coordinates. We then go ahead and feed this result into the matrix $P^{-1} = C$ on the left, which takes $T(\vec{v})$ in standard coordinates and turns into $[T(\vec{v})]_\alpha$, i.e. the α -coordinates of $T(\vec{v})$. In other words, the chaining of three matrices on the right converts α -coordinates to standard coordinates, then transforms this result in standard coordinates, then converts from standard coordinates to α -coordinates, and this three-step process can also be accomplished by just using the matrix $[T]_\alpha$ alone.

With this result in hand, we now have a significantly easier way to compute the matrix $[T]_\alpha$ using $[T]$ and the two vectors \vec{u}_1 and \vec{u}_2 . To see this in action, let's go back to our original example in this section. So let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the unique linear transformation with

$$[T] = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$$

and let $\alpha = (\vec{u}_1, \vec{u}_2)$ where

$$\vec{u}_1 = \begin{pmatrix} 3 \\ -1 \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} 2 \\ 1 \end{pmatrix}.$$

Letting

$$P = \begin{pmatrix} 3 & 2 \\ -1 & 1 \end{pmatrix},$$

we have

$$P^{-1} = \frac{1}{5} \begin{pmatrix} 1 & -2 \\ 1 & 3 \end{pmatrix},$$

hence

$$\begin{aligned}
 [T]_\alpha &= P^{-1}AP \\
 &= \frac{1}{5} \begin{pmatrix} 1 & -2 \\ 1 & 3 \end{pmatrix} \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} 3 & 2 \\ -1 & 1 \end{pmatrix} \\
 &= \frac{1}{5} \begin{pmatrix} 1 & -2 \\ 1 & 3 \end{pmatrix} \begin{pmatrix} 5 & 5 \\ 1 & 4 \end{pmatrix} \\
 &= \frac{1}{5} \begin{pmatrix} 3 & -3 \\ 8 & 17 \end{pmatrix} \\
 &= \begin{pmatrix} \frac{3}{5} & -\frac{3}{5} \\ \frac{8}{5} & \frac{17}{5} \end{pmatrix},
 \end{aligned}$$

as we calculated above. Again, the advantage of this approach over our previous methods is that we can simply invert a matrix and perform a straightforward computation without having to think through the concepts every single time.

For the other coordinate system, consider $\beta = (\vec{w}_1, \vec{w}_2)$ where

$$\vec{w}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad \text{and} \quad \vec{w}_2 = \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

Letting

$$P = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix},$$

we have

$$P^{-1} = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix},$$

hence

$$\begin{aligned}
 [T]_\beta &= P^{-1}[T]P \\
 &= \frac{1}{2} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \\
 &= \frac{1}{2} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 3 & -1 \\ 3 & 1 \end{pmatrix} \\
 &= \frac{1}{2} \begin{pmatrix} 6 & 0 \\ 0 & 2 \end{pmatrix} \\
 &= \begin{pmatrix} 3 & 0 \\ 0 & 1 \end{pmatrix},
 \end{aligned}$$

as we calculated above.

We end with a new example. Consider the unique linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with

$$[T] = \begin{pmatrix} \frac{4}{3} & \frac{1}{6} \\ -\frac{2}{3} & \frac{2}{3} \end{pmatrix}.$$

As in our previous example, this matrix is sufficiently complicated that it seems extremely difficult to view the corresponding action geometrically. However, consider the following coordinate system. Let $\alpha = (\vec{u}_1, \vec{u}_2)$ where

$$\vec{u}_1 = \begin{pmatrix} 1 \\ -2 \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} 3 \\ 0 \end{pmatrix}.$$

Letting

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

we have

$$P^{-1} = \frac{1}{6} \begin{pmatrix} 0 & -3 \\ 2 & 1 \end{pmatrix},$$

hence

$$\begin{aligned} [T]_{\alpha} &= P^{-1}[T]P \\ &= \frac{1}{6} \begin{pmatrix} 0 & -3 \\ 2 & 1 \end{pmatrix} \begin{pmatrix} \frac{4}{3} & \frac{1}{6} \\ -\frac{2}{3} & \frac{2}{3} \end{pmatrix} \begin{pmatrix} 1 & 3 \\ -2 & 0 \end{pmatrix} \\ &= \frac{1}{6} \begin{pmatrix} 0 & -3 \\ 2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 4 \\ -2 & -2 \end{pmatrix} \\ &= \frac{1}{6} \begin{pmatrix} 6 & 6 \\ 0 & 6 \end{pmatrix} \\ &= \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}. \end{aligned}$$

Aha, this is a matrix that we understand! In standard coordinates, it represents a horizontal shear relative to the x -axis, i.e. it shifts the half-plane above the x -axis to the right, and shifts the half-plane below the x -axis to the left. The underlying reason for this is that

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x + y \\ y \end{pmatrix}$$

for all $x, y \in \mathbb{R}$. Thus, every point on the line $y = 0$ (which is $\text{Span}(\vec{e}_1)$) is fixed, the line $y = 1$ is shifted to the right by 1, the line $y = 3$ is shifted to the right by 3, the line $y = -2$ is shifted to the left by 2, etc.

However, we have to geometrically interpret our matrix differently because it is in α -coordinates. Although the vectors \vec{u}_1 and \vec{u}_2 are not perpendicular, we still obtain a certain kind of shear transformation. Notice that the above matrix tells us that if \vec{v} has α -coordinates

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix}$$

then $T(\vec{v})$ will have α -coordinates

$$\begin{pmatrix} c_1 + c_2 \\ c_2 \end{pmatrix}.$$

In other words, we have

$$T(c_1 \cdot \vec{u}_1 + c_2 \cdot \vec{u}_2) = (c_1 + c_2) \cdot \vec{u}_1 + c_2 \cdot \vec{u}_2$$

for all $c_1, c_2 \in \mathbb{R}$. Thus, instead of the shear being relative to the line $y = 0$, which is $\text{Span}(\vec{e}_1)$, it will be relative to $\text{Span}(\vec{u}_1)$, i.e. relative to the line $y = -2x$. All points on this line are fixed because the second α -coordinate of such a point is 0. Every point with second α -coordinate equal to 1, i.e. every point that is 3 to the right of the line $y = -2x$, will have its first α -coordinate shifted by 1. In other words, points on the line $y = -2(x - 3) = -2x + 6$ will stay on this line, but will move down along it by adding \vec{u}_1 . Continuing in this way, anything to the right of the line $y = -2x$ will be sheared in the direction of \vec{u}_1 by an amount proportional to its horizontal distance from $y = -2x$, and anything to the left of the line $y = -2x$ will be sheared in the direction of $-\vec{u}_1$ by an amount proportional to its horizontal distance from $y = -2x$.

3.3 Eigenvalues and Eigenvectors

In the previous section, we learned how a well-chosen change in coordinates can produce a simpler matrix for a linear transformation T , which can lead to both geometric and algebraic understanding of how T acts. However, we gave no indication of how to choose $\alpha = (\vec{u}_1, \vec{u}_2)$ in a way so that $[T]_\alpha$ is “nice”. In this section, we first define what we consider “nice” behavior, and then explore how to find bases α such that $[T]_\alpha$ will be as simple as possible.

In our first example in the previous section, where

$$[T] = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix},$$

we worked with $\beta = (\vec{w}_1, \vec{w}_2)$ where

$$\vec{w}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad \text{and} \quad \vec{w}_2 = \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

For this choice of coordinates, we have

$$[T]_\beta = \begin{pmatrix} 3 & 0 \\ 0 & 1 \end{pmatrix}.$$

In other words, we have $T(\vec{w}_1) = 3\vec{w}_1$ and $T(\vec{w}_2) = \vec{w}_2$. This latter equation where we found a fixed point of T looks ideal. What could be better than a point that is not moved by T ? The former equation is almost as good. We have that $T(\vec{w}_1)$ is simply stretched, and so the result is that $T(\vec{w}_1)$ is still an element of $\text{Span}(\vec{w}_1)$. In other words, although \vec{w}_1 was moved, it was not rotated in any way, and hence was sent to another point in $\text{Span}(\vec{w}_1)$. We now give a name to the points and scaling factors that arise in this way.

Definition 3.3.1. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation.

- An eigenvector of T is a nonzero vector $\vec{v} \in \mathbb{R}^2$ such that there exists $\lambda \in \mathbb{R}$ with $T(\vec{v}) = \lambda\vec{v}$.
- An eigenvalue of T is a scalar $\lambda \in \mathbb{R}$ such that there exists a nonzero $\vec{v} \in \mathbb{R}^2$ with $T(\vec{v}) = \lambda\vec{v}$.

When $\vec{v} \in \mathbb{R}^2$ is nonzero and $\lambda \in \mathbb{R}$ are such that $T(\vec{v}) = \lambda\vec{v}$, we say that \vec{v} is an eigenvector of T corresponding to the eigenvalue λ .

For example, suppose that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the unique linear transformation with

$$[T] = \begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix}.$$

Notice that

$$\begin{aligned} T\left(\begin{pmatrix} 2 \\ 1 \end{pmatrix}\right) &= \begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 2 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} 6 \\ 3 \end{pmatrix} \\ &= 3 \cdot \begin{pmatrix} 2 \\ 1 \end{pmatrix}, \end{aligned}$$

so $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ is an eigenvector of T , 3 is an eigenvalue of T , and hence $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ is an eigenvector of T corresponding to 3. In contrast, we have

$$\begin{aligned} T\left(\begin{pmatrix} 1 \\ -3 \end{pmatrix}\right) &= \begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ -3 \end{pmatrix} \\ &= \begin{pmatrix} 10 \\ -2 \end{pmatrix}. \end{aligned}$$

Since there does not exist $\lambda \in \mathbb{R}$ with

$$\begin{pmatrix} 10 \\ -2 \end{pmatrix} = \lambda \cdot \begin{pmatrix} 1 \\ -3 \end{pmatrix}$$

(by contradiction, since such a λ would have to satisfy both $\lambda = 10$ and $\lambda = \frac{2}{3}$), it follows that $\begin{pmatrix} 1 \\ -3 \end{pmatrix}$ is not an eigenvector of A .

Notice that, in the definition, we insist that an eigenvector \vec{v} is nonzero. One reason for this is that $T(\vec{0}) = \vec{0}$ for every linear transformation T , so always having a “trivial” eigenvector cheapens the name. More importantly, since $T(\vec{0}) = \vec{0}$, we have that $T(\vec{0}) = \lambda \cdot \vec{0}$ for every $\lambda \in \mathbb{R}$. Thus, if we do not rule out the zero vector, then every real number would become an eigenvalue, which we certainly do not want to be the case.

However, it is possible that 0 could be an eigenvalue for a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$. For this to happen, we just need there to exist a nonzero vector $\vec{v} \in \mathbb{R}^2$ with $T(\vec{v}) = \vec{0}$, which is the same thing as saying that $\text{Null}(T) \neq \{\vec{0}\}$. For example, if $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ consider the unique linear transformation with

$$[T] = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}.$$

Notice that $\text{Null}(T) \neq \{\vec{0}\}$ by Theorem 2.7.3 because $1 \cdot 1 - 1 \cdot 1 = 0$. In fact, we have

$$\begin{aligned} T\left(\begin{pmatrix} 1 \\ -1 \end{pmatrix}\right) &= \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ -1 \end{pmatrix} \\ &= \begin{pmatrix} 0 \\ 0 \end{pmatrix} \\ &= 0 \cdot \begin{pmatrix} 1 \\ -1 \end{pmatrix}. \end{aligned}$$

so $\begin{pmatrix} 1 \\ -1 \end{pmatrix}$ is an eigenvector corresponding to 0. We will see a close relationship between eigenvalues, eigenvectors, and null spaces more generally below.

Before moving on, we first show that every eigenvector corresponds to exactly one eigenvalue. As mentioned above, for this to be true, it is essential that we not allow the zero vector in our definition of eigenvector.

Proposition 3.3.2. *Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. If \vec{v} is an eigenvector of T , then there exists a unique $\lambda \in \mathbb{R}$ such that \vec{v} is an eigenvector of T corresponding to λ . In other words, if \vec{v} is an eigenvector of T corresponding to both $\lambda \in \mathbb{R}$ and $\mu \in \mathbb{R}$, then $\lambda = \mu$.*

Proof. Let \vec{v} be an eigenvector of T . Suppose that \vec{v} is an eigenvector of T corresponding to both $\lambda \in \mathbb{R}$ and $\mu \in \mathbb{R}$. By definition of an eigenvector, we then have that $\vec{v} \neq \vec{0}$, that $T(\vec{v}) = \lambda\vec{v}$, and that $T(\vec{v}) = \mu\vec{v}$. Combining these latter two equalities, we have $\lambda\vec{v} = \mu\vec{v}$, and subtracting $\mu\vec{v}$ from both sides, we conclude that $\lambda\vec{v} - \mu\vec{v} = \vec{0}$. It follows that $(\lambda - \mu)\vec{v} = \vec{0}$. Now if $\lambda - \mu \neq 0$, then we can divide both sides by $\lambda - \mu$ to conclude that $\vec{v} = \vec{0}$, which would be a contradiction. Therefore, we must have $\lambda - \mu = 0$, and hence $\lambda = \mu$. \square

We have defined eigenvalues and eigenvectors for a linear transformation T , but we can also define them for matrices. Since matrices code linear transformations, you might think that there is no distinction. However, we do want to keep a sharp distinction between linear transformations and the matrices that code them because we need *both* a linear transformation and an α to build a matrix. In particular, one linear transformation can be represented by distinct matrices when we change the coordinate system. Nonetheless, we can also define these concepts for matrices in the completely analogous way.

Definition 3.3.3. Let A be a 2×2 matrix.

- An eigenvector of A is a nonzero vector $\vec{v} \in \mathbb{R}^2$ such that there exists $\lambda \in \mathbb{R}$ with $A\vec{v} = \lambda\vec{v}$.
- An eigenvalue of A is a scalar $\lambda \in \mathbb{R}$ such that there exists a nonzero $\vec{v} \in \mathbb{R}^2$ with $A\vec{v} = \lambda\vec{v}$.

When $\vec{v} \in \mathbb{R}^2$ is nonzero and $\lambda \in \mathbb{R}$ are such that $A\vec{v} = \lambda\vec{v}$, we say that \vec{v} is an eigenvector of A corresponding to the eigenvalue λ .

Now if we work with the standard matrix of a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, then it is immediate from the definitions that \vec{v} is an eigenvector of T corresponding to λ if and only if \vec{v} is an eigenvector of $[T]$ corresponding to λ . However, if we use a different matrix representation, then the eigenvectors might change. We will not dwell on these issues now, but they will play a role later.

Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation, and let $A = [T]$. Given $\vec{v} \in \mathbb{R}^2$ and $\lambda \in \mathbb{R}$, it is easy to determine if \vec{v} is an eigenvector of T corresponding to λ : just calculate $A\vec{v}$ and check if it equals $\lambda\vec{v}$. Now if we are only given $\vec{v} \in \mathbb{R}^2$ (so no λ), it is still straightforward to determine if \vec{v} is an eigenvector of T . We just need to compute $A\vec{v}$ and then check if the result is *some* scalar multiple of \vec{v} . However, given $\lambda \in \mathbb{R}$, it is not obvious how we can check whether λ is an eigenvalue and, if so, find eigenvectors corresponding to it. In other words, we have the following two questions:

- Given a real number λ , how do we determine if λ is an eigenvalue of A ?
- Given an eigenvalue λ of A , how do we determine the eigenvectors corresponding to λ ?

Here is an idea for how to answer the first question. Given $\lambda \in \mathbb{R}$, we want to know if there are any nonzero vectors $\vec{v} \in \mathbb{R}^2$ such that $A\vec{v} = \lambda\vec{v}$. Intuitively, we want to “solve” for \vec{v} here. To make this happen, we try to bring the \vec{v} ’s together to the other side and instead work with $A\vec{v} - \lambda\vec{v} = \vec{0}$. Now the natural idea is to try to factor out the common \vec{v} , but $A - \lambda$ does not make sense because we can not subtract a number from a matrix. However, there is a hack to get around this. We can view the scalar multiplication $\lambda\vec{v}$ as a certain matrix-vector product. For example, if we want to multiply \vec{v} by 5, we can instead hit \vec{v} by the matrix

$$\begin{pmatrix} 5 & 0 \\ 0 & 5 \end{pmatrix}.$$

In other words, we can simply write $\lambda\vec{v}$ as $\lambda I\vec{v}$. Since I is the identity matrix, this is just like multiplying by 1 in an algebraic equation. With this perspective, we can rewrite $A\vec{v} - \lambda\vec{v} = \vec{0}$ as $A\vec{v} - \lambda I\vec{v} = \vec{0}$. From here, we can factor out the \vec{v} and rewrite this as $(A - \lambda I)\vec{v} = \vec{0}$. Thus, to determine if there are any nonzero vectors \vec{v} with $A\vec{v} = \lambda\vec{v}$, we can instead determine if the matrix $A - \lambda I$ kills off any nonzero vectors. In other words, we are asking if the null space of the matrix $A - \lambda I$ is nontrivial. We now formalize this argument in the following result.

Proposition 3.3.4. Let A be a 2×2 matrix, let $\vec{v} \in \mathbb{R}^2$, and let $\lambda \in \mathbb{R}$. We have that $A\vec{v} = \lambda\vec{v}$ if and only if $\vec{v} \in \text{Null}(A - \lambda I)$. Therefore, \vec{v} is an eigenvector of A corresponding to λ if and only if $\vec{v} \neq \vec{0}$ and $\vec{v} \in \text{Null}(A - \lambda I)$.

Proof. Suppose first that $A\vec{v} = \lambda\vec{v}$. Subtracting $\lambda\vec{v}$ from both sides, we then have $A\vec{v} - \lambda\vec{v} = \vec{0}$, and hence $A\vec{v} - \lambda I\vec{v} = \vec{0}$. We then have $(A - \lambda I)\vec{v} = \vec{0}$, and thus $\vec{v} \in \text{Null}(A - \lambda I)$.

Conversely, suppose that $\vec{v} \in \text{Null}(A - \lambda I)$. We then have $(A - \lambda I)\vec{v} = \vec{0}$, so $A\vec{v} - \lambda I\vec{v} = \vec{0}$, and hence $A\vec{v} - \lambda\vec{v} = \vec{0}$. Adding $\lambda\vec{v}$ to both sides, we conclude that $A\vec{v} = \lambda\vec{v}$. \square

From this result, we immediately obtain the following corollary.

Corollary 3.3.5. Let A be a 2×2 matrix and let $\lambda \in \mathbb{R}$. We have that λ is an eigenvalue of A if and only if $\text{Null}(A - \lambda I) \neq \{\vec{0}\}$.

Combining this fact with Theorem 2.7.3, we obtain a method to determine if a given number $\lambda \in \mathbb{R}$ is an eigenvalue of a matrix A . For example, suppose that we are working with the matrix

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

as above. Is 5 an eigenvalue of A ? We calculate

$$A - 5I = \begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix} - \begin{pmatrix} 5 & 0 \\ 0 & 5 \end{pmatrix} = \begin{pmatrix} -1 & -2 \\ 1 & -4 \end{pmatrix}.$$

Notice that $\text{Null}(A - 5I) = \{\vec{0}\}$ by Theorem 2.7.3 because $(-1) \cdot (-4) - (-2) \cdot 1 = 6$ is nonzero. Using Corollary 3.3.5, it follows that 5 is not an eigenvalue of A . Now we noticed that 3 was an eigenvalue of A above by stumbling across a particular eigenvector corresponding to 3. We can see that 3 is an eigenvalue in another way now. Notice that

$$A - 3I = \begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix} - \begin{pmatrix} 3 & 0 \\ 0 & 3 \end{pmatrix} = \begin{pmatrix} 1 & -2 \\ 1 & -2 \end{pmatrix}.$$

Notice that $\text{Null}(A - 3I) \neq \{\vec{0}\}$ by Theorem 2.7.3 because $(-1) \cdot (-2) - (-2) \cdot 1 = 0$. Therefore, 3 is an eigenvalue of A by Corollary 3.3.5. Now if we want to find all eigenvectors corresponding to 3, we just want to find the null space of the matrix

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

Following the proof of Theorem 2.7.3 (in Case 3 at the end), we conclude that

$$\text{Null}(A - 3I) = \text{Span} \left(\begin{pmatrix} 2 \\ 1 \end{pmatrix} \right).$$

Thus, we rediscover the eigenvector

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

that we stumbled across above, and in fact any constant multiple of this eigenvector is also an eigenvector corresponding to 3. In other words, every element of

$$\text{Span} \left(\begin{pmatrix} 2 \\ 1 \end{pmatrix} \right)$$

is an eigenvector of A corresponding to 3.

Given a matrix A and a scalar $\lambda \in \mathbb{R}$, we now have a procedure to determine if λ is an eigenvalue of A . Using this, we can go through numbers one at time, and if we're lucky, we might find eigenvalues. Is there a better way to find *all* of the eigenvalues simultaneously? The key fact is that we perform the same procedure no matter what λ is, and so we can simply look at the matrix $A - \lambda I$ in general, and work with it (i.e. we do not plug a specific value in for λ). Suppose then that

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

and we want to find all eigenvalues of A . In other words, we want to find those $\lambda \in \mathbb{R}$ such that $\text{Null}(A - \lambda I) \neq \{\vec{0}\}$. We have

$$A - \lambda I = \begin{pmatrix} a & b \\ c & d \end{pmatrix} - \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix} = \begin{pmatrix} a - \lambda & b \\ c & d - \lambda \end{pmatrix}.$$

Therefore, using Theorem 2.7.3, we have that $\text{Null}(A - \lambda I) \neq \{\vec{0}\}$ if and only if $(a - \lambda)(d - \lambda) - bc = 0$. When we expand this out, we obtain a degree two polynomial in the variable λ . We give this polynomial a special name.

Definition 3.3.6. Given a 2×2 matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

we define the characteristic polynomial of A to be the following polynomial in variable λ :

$$(a - \lambda)(d - \lambda) - bc = \lambda^2 - (a + d)\lambda + (ad - bc).$$

Let's see a few examples.

Example 3.3.7. Let

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

Find all eigenvalues of A , and then find at least one eigenvector of A corresponding to each eigenvalue.

Solution. For any $\lambda \in \mathbb{R}$, we have

$$A - \lambda I = \begin{pmatrix} 4 - \lambda & -2 \\ 1 & 1 - \lambda \end{pmatrix},$$

so the characteristic polynomial of A is

$$\begin{aligned} (4 - \lambda)(1 - \lambda) - (-2) &= 4 - 5\lambda + \lambda^2 + 2 \\ &= \lambda^2 - 5\lambda + 6 \\ &= (\lambda - 3)(\lambda - 2). \end{aligned}$$

Therefore, $\text{Null}(A - \lambda I) \neq \{\vec{0}\}$ if and only if either $\lambda = 3$ or $\lambda = 2$. It follows that the eigenvalues of A are 3 and 2. We have already seen that

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

is an eigenvector corresponding to 3, and that the set of all eigenvectors corresponding to 3 is

$$\text{Null}(A - 3I) = \text{Span} \left(\begin{pmatrix} 2 \\ 1 \end{pmatrix} \right).$$

Now consider $\lambda = 2$. We have

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

To find an eigenvector corresponding to 2, we need only find a nonzero element of the null space of this matrix. One example is

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

Following the proof of Theorem 2.7.3, we conclude that the set of all eigenvectors corresponding to 3 is

$$\text{Null}(A - 2I) = \text{Span} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix} \right).$$

□

Example 3.3.8. *Let*

$$A = \begin{pmatrix} 5 & 4 \\ -7 & -6 \end{pmatrix}.$$

Find all eigenvalues of A , and then find at least one eigenvector of A corresponding to each eigenvalue.

Solution. For any $\lambda \in \mathbb{R}$, we have

$$A - \lambda I = \begin{pmatrix} 5 - \lambda & 4 \\ -7 & -6 - \lambda \end{pmatrix},$$

so the characteristic polynomial of A is

$$\begin{aligned} (5 - \lambda)(-6 - \lambda) - (-28) &= -30 + \lambda + \lambda^2 + 28 \\ &= \lambda^2 + \lambda - 2 \\ &= (\lambda - 1)(\lambda + 2) \end{aligned}$$

Therefore, $\text{Null}(A - \lambda I) \neq \{\vec{0}\}$ if and only if either $\lambda = 1$ or $\lambda = -2$. It follows that the eigenvalues of A are 1 and -2 . For $\lambda = 1$, we have

$$A - 1I = \begin{pmatrix} 4 & 4 \\ -7 & -7 \end{pmatrix}.$$

To find an eigenvector corresponding to 1, we need only find a nonzero element of the null space of this matrix. One example is

$$\begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

Following the proof of Theorem 2.7.3, we conclude that the set of all eigenvectors corresponding to 1 is

$$\text{Null}(A - 1I) = \text{Span} \left(\begin{pmatrix} -1 \\ 1 \end{pmatrix} \right).$$

Now consider $\lambda = -2$. We have

$$A + 2I = \begin{pmatrix} 7 & 4 \\ -7 & -4 \end{pmatrix}.$$

To find an eigenvector corresponding to -2 , we need only find a nonzero element of the null space of this matrix. One example is

$$\begin{pmatrix} -4 \\ 7 \end{pmatrix}.$$

Again, following the proof of Theorem 2.7.3, we conclude that the set of all eigenvectors corresponding to -2 is

$$\text{Null}(A + 2I) = \text{Span} \left(\begin{pmatrix} -4 \\ 7 \end{pmatrix} \right).$$

□

Example 3.3.9. *Let*

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

Find all eigenvalues of A , and then find at least one eigenvector of A corresponding to each eigenvalue.

Solution. For any $\lambda \in \mathbb{R}$, we have

$$A - \lambda I = \begin{pmatrix} -\lambda & -1 \\ 1 & -\lambda \end{pmatrix},$$

so the characteristic polynomial of A is

$$(-\lambda)(-\lambda) - (-1) = \lambda^2 + 1.$$

Therefore, A has no eigenvalues because $\lambda^2 + 1$ has no roots in \mathbb{R} . □

Example 3.3.10. *Let*

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

Find all eigenvalues of A , and then find at least one eigenvector of A corresponding to each eigenvalue.

Solution. For any $\lambda \in \mathbb{R}$, we have

$$A - \lambda I = \begin{pmatrix} 1 - \lambda & 1 \\ 0 & 1 - \lambda \end{pmatrix}$$

so the characteristic polynomial of A is

$$(1 - \lambda)^2 - 0 = (\lambda - 1)^2,$$

Therefore, 1 is the only eigenvalue of A . We have

$$A - 1I = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix},$$

Notice that

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

is an element of $\text{Null}(A - 1I)$, so is an eigenvector of A corresponding to 1. Following the proof of Theorem 2.7.3, we conclude that the set of all eigenvectors corresponding to 1 is

$$\text{Null}(A - 1I) = \text{Span} \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} \right).$$

□

Let's go back to our matrix

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

We know that there exists a unique linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that $[T] = A$. Our original motivation for eigenvectors was that they appeared to be smart choices to use as vectors for a new coordinate system. Let's go ahead and make use of that now. We know from Example 3.3.7 that the eigenvalues of T are 3 and 2. We also know that

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

is an eigenvector corresponding to 3, and that

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

is an eigenvector corresponding to 2. With this in mind, let $\alpha = (\vec{u}_1, \vec{u}_2)$ where

$$\vec{u}_1 = \begin{pmatrix} 2 \\ 1 \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

Since $2 \cdot 1 - 1 \cdot 1 = 1$ is nonzero, we know that (\vec{u}_1, \vec{u}_2) is a basis of \mathbb{R}^2 . Now using the fact that \vec{u}_1 and \vec{u}_2 are eigenvectors of T , we have

$$\begin{aligned} T(\vec{u}_1) &= 3\vec{u}_1 \\ &= 3 \cdot \vec{u}_1 + 0 \cdot \vec{u}_2, \end{aligned}$$

so

$$[T(\vec{u}_1)]_\alpha = \begin{pmatrix} 3 \\ 0 \end{pmatrix}.$$

We also have

$$\begin{aligned} T(\vec{u}_2) &= 2\vec{u}_2 \\ &= 0 \cdot \vec{u}_1 + 2 \cdot \vec{u}_2. \end{aligned}$$

so

$$[T(\vec{u}_2)]_\alpha = \begin{pmatrix} 0 \\ 2 \end{pmatrix}.$$

It follows that

$$[T]_\alpha = \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix}.$$

Thus, we have indeed found a coordinate system α where $[T]_\alpha$ is particularly nice. We can also compute $[T]_\alpha$ using Proposition 3.2.6. Letting

$$P = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$$

we then have that P is invertible and that

$$P^{-1} = \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix}.$$

Therefore, we have

$$\begin{aligned} [T]_\alpha &= P^{-1} \cdot [T] \cdot P \\ &= \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix} \begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \\ &= \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix} \begin{pmatrix} 6 & 2 \\ 3 & 2 \end{pmatrix} \\ &= \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix}. \end{aligned}$$

The reason why we use eigenvectors of T for our choice of \vec{u}_1 and \vec{u}_2 in our basis $\alpha = (\vec{u}_1, \vec{u}_2)$ is that we always obtain a matrix that looks like this.

Definition 3.3.11. A diagonal 2×2 matrix is a 2×2 matrix D such that there exists $a, d \in \mathbb{R}$ with

$$D = \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix}$$

In other words, a diagonal matrix is a matrix that has 0's off of the main diagonal, but may have nonzero entries on the main diagonal.

Definition 3.3.12. A linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is diagonalizable if there exists a basis $\alpha = (\vec{u}_1, \vec{u}_2)$ of \mathbb{R}^2 such that $[T]_\alpha$ is a diagonal matrix.

For example, if $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the unique linear transformation with

$$[T] = \begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix}$$

then we've just seen that T is diagonalizable. We now argue generally that a linear transformation is diagonalizable if and only if we can find a basis of two eigenvectors.

Proposition 3.3.13. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation and let $\alpha = (\vec{u}_1, \vec{u}_2)$ be a basis of \mathbb{R}^2 . The following are equivalent:

1. $[T]_\alpha$ is a diagonal matrix.
2. \vec{u}_1 and \vec{u}_2 are eigenvectors of T .

Furthermore, in this case, the diagonal entries of $[T]_\alpha$ are the eigenvalues corresponding to \vec{u}_1 and \vec{u}_2 , i.e. the upper left entry of $[T]_\alpha$ is the eigenvalue of $[T]$ corresponding to \vec{u}_1 , and the lower right entry of $[T]_\alpha$ is the eigenvalue of T corresponding to \vec{u}_2 .

Proof. Suppose first that $[T]_\alpha$ is a diagonal matrix. Fix $a, d \in \mathbb{R}$ with

$$[T]_\alpha = \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix}.$$

By the definition of $[T]_\alpha$, we then have that

$$\begin{aligned} T(\vec{u}_1) &= a \cdot \vec{u}_1 + 0 \cdot \vec{u}_2 \\ &= a \cdot \vec{u}_1, \end{aligned}$$

so \vec{u}_1 is an eigenvector corresponding to a . Similarly, we have

$$\begin{aligned} T(\vec{u}_2) &= 0 \cdot \vec{u}_1 + d \cdot \vec{u}_2 \\ &= d \cdot \vec{u}_2, \end{aligned}$$

so \vec{u}_2 is an eigenvector corresponding to d . Thus, we have shown that (2) is true, and also the additional statement.

Suppose now that \vec{u}_1 and \vec{u}_2 are eigenvectors of T . Fix $\lambda_1, \lambda_2 \in \mathbb{R}$ with $T(\vec{u}_1) = \lambda_1 \cdot \vec{u}_1$ and $T(\vec{u}_2) = \lambda_2 \cdot \vec{u}_2$. We then have

$$\begin{aligned} T(\vec{u}_1) &= \lambda_1 \cdot \vec{u}_1 \\ &= \lambda_1 \cdot \vec{u}_1 + 0 \cdot \vec{u}_2, \end{aligned}$$

so

$$[T(\vec{u}_1)]_\alpha = \begin{pmatrix} \lambda_1 \\ 0 \end{pmatrix}.$$

We also have

$$\begin{aligned} T(\vec{u}_2) &= \lambda_2 \cdot \vec{u}_2 \\ &= 0 \cdot \vec{u}_1 + \lambda_2 \cdot \vec{u}_2, \end{aligned}$$

so

$$[T(\vec{u}_2)]_\alpha = \begin{pmatrix} 0 \\ \lambda_2 \end{pmatrix}.$$

Therefore, by the definition of $[T]_\alpha$, we have

$$[T]_\alpha = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}.$$

Thus, we have shown that (1) is true, and also the additional statement. \square

Corollary 3.3.14. *A linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is diagonalizable if and only if there exists a basis (\vec{u}_1, \vec{u}_2) of \mathbb{R}^2 such that both \vec{u}_1 and \vec{u}_2 are eigenvectors of T .*

Proof. This follows immediately from Proposition 3.3.13. \square

Example 3.3.15. *Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the unique linear transformation with*

$$[T] = \begin{pmatrix} 2 & -3 \\ 2 & -5 \end{pmatrix}.$$

Determine if T is diagonalizable, and if so, find α such that $[T]_\alpha$ is a diagonal matrix.

Solution. Let $A = [T]$. For any $\lambda \in \mathbb{R}$, we have

$$A - \lambda I = \begin{pmatrix} 2 - \lambda & -3 \\ 2 & -5 - \lambda \end{pmatrix},$$

so the characteristic polynomial of A is

$$\begin{aligned} (2 - \lambda)(-5 - \lambda) - (-6) &= -10 + 5\lambda - 2\lambda + \lambda^2 + 6 \\ &= \lambda^2 + 3\lambda - 4 \\ &= (\lambda - 1)(\lambda + 4). \end{aligned}$$

Therefore, $\text{Null}(A - \lambda I) \neq \{\vec{0}\}$ if and only if either $\lambda = 1$ or $\lambda = -4$. It follows that the eigenvalues of A are 1 and -4 . Since

$$A - 1I = \begin{pmatrix} 1 & -3 \\ 2 & -6 \end{pmatrix},$$

we have that one eigenvector of A corresponding to 1 is

$$\begin{pmatrix} 3 \\ 1 \end{pmatrix}$$

Since

$$A + 4I = \begin{pmatrix} 6 & -3 \\ 2 & -1 \end{pmatrix},$$

we have that one eigenvector of A corresponding to -4 is

$$\begin{pmatrix} 1 \\ 2 \end{pmatrix}.$$

Letting

$$\vec{u}_1 = \begin{pmatrix} 3 \\ 1 \end{pmatrix} \quad \text{and} \quad \vec{u}_2 = \begin{pmatrix} 1 \\ 2 \end{pmatrix},$$

we have that \vec{u}_1 and \vec{u}_2 are eigenvectors of A and that (\vec{u}_1, \vec{u}_2) is a basis of \mathbb{R}^2 (because $3 \cdot 2 - 1 \cdot 1 = 5$ is nonzero). Therefore, if we let $\alpha = (\vec{u}_1, \vec{u}_2)$, then using Proposition 3.3.13, we know that $[T]_\alpha$ is a diagonal matrix and in fact

$$[T]_\alpha = \begin{pmatrix} 1 & 0 \\ 0 & -4 \end{pmatrix}.$$

□

Example 3.3.16. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the unique linear transformation with

$$[T] = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

Determine if T is diagonalizable, and if so, find α such that $[T]_\alpha$ is a diagonal matrix.

Solution. In Example 3.3.9, we showed that the characteristic polynomial of $[T]$ is $\lambda^2 + 1$, so there are no eigenvalues of T . Thus, there are no eigenvectors of T , so T is not diagonalizable by Corollary 3.3.14. □

Example 3.3.17. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the unique linear transformation with

$$[T] = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

Determine if T is diagonalizable, and if so, find α such that $[T]_\alpha$ is a diagonal matrix.

Solution. In Example 3.3.10, we showed that the characteristic polynomial of $[T]$ is $(\lambda - 1)^2$, so the only eigenvalue of T is 1. We also alluded to the fact that the eigenvectors of T corresponding to 1 are exactly the elements of $\text{Span}(\vec{e}_1)$ (without the zero vector), i.e. that the set of eigenvectors corresponding to 1 is

$$\left\{ \begin{pmatrix} x \\ 0 \end{pmatrix} : x \in \mathbb{R} \setminus \{0\} \right\}.$$

First notice that for any $x \in \mathbb{R} \setminus \{0\}$, we have

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ 0 \end{pmatrix} = \begin{pmatrix} x \\ 0 \end{pmatrix} = 1 \cdot \begin{pmatrix} x \\ 0 \end{pmatrix},$$

so every element of the given set is indeed an eigenvector corresponding to 1. Conversely, suppose that

$$\begin{pmatrix} x \\ y \end{pmatrix}$$

is an eigenvector corresponding to 1. We then have that

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x \\ y \end{pmatrix},$$

so

$$\begin{pmatrix} x + y \\ y \end{pmatrix} = \begin{pmatrix} x \\ y \end{pmatrix}.$$

It follows that we must have $x + y = x$, so we must have $y = 0$, and thus our eigenvector is in the given set.

Notice then that T does have infinitely many eigenvectors. However, given any two eigenvectors \vec{u}_1 and \vec{u}_2 of T , we have that each is a multiple of the other (geometrically they are on the same line), so $\text{Span}(\vec{u}_1, \vec{u}_2) \neq \mathbb{R}^2$. Therefore, T is not diagonalizable by Corollary 3.3.14. □

We've just seen an example where the characteristic polynomial has a double root, and the corresponding linear transformation is not diagonalizable. However, there also exist diagonalizable linear transformations T where the characteristic polynomial of $[T]$ has a double root.

Example 3.3.18. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the unique linear transformation with

$$[T] = \begin{pmatrix} 3 & 0 \\ 0 & 3 \end{pmatrix}.$$

Determine if T is diagonalizable, and if so, find α such that $[T]_\alpha$ is a diagonal matrix.

Solution. We can quickly see that T is diagonalizable because $[T]$ itself is diagonal, i.e. if we let $\alpha = (\vec{e}_1, \vec{e}_2)$, then $[T]_\alpha = [T]$ is diagonal. However, we can also attack this problem the usual way. The characteristic polynomial is $(3 - \lambda)^2 = (\lambda - 3)^2$, so the only eigenvalue is 3. Now if we let $A = [T]$, then

$$A - 3I = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

Thus, $\text{Null}(A - 3I) = \mathbb{R}^2$, and hence every nonzero vector in \mathbb{R}^2 is an eigenvector of T . Therefore, given any basis $\alpha = (\vec{u}_1, \vec{u}_2)$ of \mathbb{R}^2 , we have that

$$[T]_\alpha = \begin{pmatrix} 3 & 0 \\ 0 & 3 \end{pmatrix}.$$

□

We can use these examples to state some general phenomena. Suppose that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a linear transformation, and let $A = [T]$. If the characteristic polynomial of A has no real roots, then T has no eigenvalues, so T has no eigenvectors, and hence T is not diagonalizable by Corollary 3.3.14. Now if the characteristic polynomial of A has a repeated root, then we've just seen that sometimes T is diagonalizable and sometimes it is not, and so we have to handle those on a case-by-case basis. What happens if the characteristic polynomial of A has two distinct real roots? In this case, we know that T has 2 eigenvalues, and hence we can find an eigenvector for each of these eigenvalues. The only question is whether these two eigenvectors will span \mathbb{R}^2 . We now argue that this is always the case.

Proposition 3.3.19. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. Suppose that \vec{u}_1, \vec{u}_2 are eigenvectors of T corresponding to eigenvalues λ_1 and λ_2 respectively. If $\lambda_1 \neq \lambda_2$, then $\alpha = (\vec{u}_1, \vec{u}_2)$ is a basis of \mathbb{R}^2 .

Proof. We prove the contrapositive. Suppose that $\text{Span}(\vec{u}_1, \vec{u}_2) \neq \mathbb{R}^2$. Since \vec{u}_1 and \vec{u}_2 are both eigenvectors, we know that \vec{u}_1 and \vec{u}_2 are both nonzero. Using Theorem 2.3.10, we then have that $\vec{u}_2 \in \text{Span}(\vec{u}_1)$, so we can fix $c \in \mathbb{R}$ with $\vec{u}_2 = c \cdot \vec{u}_1$. We now have

$$\begin{aligned} T(\vec{u}_2) &= T(c \cdot \vec{u}_1) \\ &= c \cdot T(\vec{u}_1) \\ &= c \cdot \lambda_1 \vec{u}_1 \\ &= \lambda_1 \cdot (c \cdot \vec{u}_1) \\ &= \lambda_1 \vec{u}_2. \end{aligned}$$

Thus, \vec{u}_2 is an eigenvector of T corresponding to λ_1 as well. Using Proposition 3.3.2, it follows that $\lambda_1 = \lambda_2$. □

Corollary 3.3.20. If a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ has 2 distinct eigenvalues, then it is diagonalizable.

Proof. Suppose that T has 2 distinct eigenvalues λ_1 and λ_2 . Fix eigenvectors \vec{u}_1 and \vec{u}_2 of T corresponding to λ_1 and λ_2 respectively. Using Proposition 3.3.19, we know that (\vec{u}_1, \vec{u}_2) is a basis of \mathbb{R}^2 . Therefore, T is diagonalizable by Corollary 3.3.14. \square

Suppose that $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a diagonalizable linear transformation. Fix a basis $\alpha = (\vec{u}_1, \vec{u}_2)$ of \mathbb{R}^2 where both \vec{u}_1 and \vec{u}_2 are eigenvectors of T . We then have that $[T]_\alpha$ is a diagonal matrix by Corollary 3.3.14. Furthermore, if we let P be the matrix whose first column is \vec{u}_1 and whose second column is \vec{u}_2 , then using Proposition 3.2.6, we have that P is invertible and that

$$[T]_\alpha = P^{-1} \cdot [T] \cdot P.$$

Multiplying on the left by P and on the right by P^{-1} , we conclude that

$$[T] = P \cdot [T]_\alpha \cdot P^{-1}.$$

Thus, we can write $[T] = PDP^{-1}$ for some invertible matrix P and diagonal matrix D .

This setup works in the reverse direction as well. Suppose that there exists a invertible matrix P and a diagonal matrix D with

$$[T] = PDP^{-1}.$$

Multiplying on the left by P^{-1} and on the right by P , we then have that

$$D = P^{-1} \cdot [T] \cdot P.$$

Letting \vec{u}_1 be the first column of P and \vec{u}_2 be the second column of P , we know that (\vec{u}_1, \vec{u}_2) is a basis of \mathbb{R}^2 because P is invertible. Thus, if we let $\alpha = (\vec{u}_1, \vec{u}_2)$, then we have that $[T]_\alpha = D$, and hence T is diagonalizable.

In other words, instead of thinking about a diagonalizable linear transformation as one that can be represented by a diagonal matrix in some coordinate system, we can think about a diagonalizable linear transformation as one whose standard matrix can be written as $[T] = PDP^{-1}$ where P is invertible and D is diagonal.

Given a diagonalizable linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, we've seen how we use a coordinate system α with $[T]_\alpha$ diagonal to help understand the action of T geometrically on the plane. However, there are direct computational benefits as well. Here's one example. Suppose that we understand what happens when we repeatedly apply T to one point. That is, suppose that we have $\vec{v} \in \mathbb{R}^2$, and we want to look at $T(\vec{v})$, then $T(T(\vec{v}))$, then $T(T(T(\vec{v})))$, etc. Situations like this arose in the introduction when we were studying population models and Monopoly (although those were happening in \mathbb{R}^4 and \mathbb{R}^{40} rather than \mathbb{R}^2). Now if we let $A = [T]$, then we are trying to understand powers of the matrix A .

For a simple example, let

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

Now we could simply compute powers of A directly. We have

$$\begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix}^2 = \begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 14 & -10 \\ 5 & -1 \end{pmatrix}$$

and so

$$\begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix}^3 = \begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 14 & -10 \\ 5 & -1 \end{pmatrix} = \begin{pmatrix} 46 & -38 \\ 19 & -11 \end{pmatrix}.$$

Although we can keep performing repeated multiplication directly, this process is tedious and it's not clear whether a pattern is emerging. However, we can use the fact that if $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the unique linear transformation with $[T] = A$, then T is diagonalizable. Recall that we showed that if

$$P = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \quad \text{and} \quad D = \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix}$$

then we have

$$D = P^{-1}AP.$$

Multiplying on the left by P and on the right by P^{-1} , we conclude that

$$A = PDP^{-1}.$$

Written out with numbers, we have

$$\begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix}.$$

It may appear that writing A as the product of three matrices will only make finding powers of A worse. However, let's examine what happens. We have

$$\begin{aligned} A^2 &= PDP^{-1}PDP^{-1} \\ &= PD^2P^{-1}, \end{aligned}$$

and then

$$\begin{aligned} A^3 &= A^2A \\ &= PD^2P^{-1}PDP^{-1} \\ &= PD^3P^{-1}. \end{aligned}$$

In general, we have

$$A^n = PD^nP^{-1}$$

for all $n \in \mathbb{N}^+$. Of course, we still have to raise D to a power, but this is easy because D is diagonal! If

$$D = \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix}$$

then we have

$$\begin{aligned} D^2 &= \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix} \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix} \\ &= \begin{pmatrix} a^2 & 0 \\ 0 & d^2 \end{pmatrix} \end{aligned}$$

and then

$$\begin{aligned} D^3 &= D^2D \\ &= \begin{pmatrix} a^2 & 0 \\ 0 & d^2 \end{pmatrix} \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix} \\ &= \begin{pmatrix} a^3 & 0 \\ 0 & d^3 \end{pmatrix}. \end{aligned}$$

In general, we have

$$D^n = \begin{pmatrix} a^n & 0 \\ 0 & d^n \end{pmatrix}$$

for all $n \in \mathbb{N}^+$.

Let's go back to our example where we had $A = PDP^{-1}$ as follows:

$$\begin{pmatrix} 4 & -2 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix}.$$

For any $n \in \mathbb{N}^+$, we have

$$\begin{aligned} A^n &= PD^nP^{-1} \\ &= \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 3^n & 0 \\ 0 & 2^n \end{pmatrix} \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix} \\ &= \begin{pmatrix} 2 \cdot 3^n & 2^n \\ 3^n & 2^n \end{pmatrix} \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix} \\ &= \begin{pmatrix} 2 \cdot 3^n - 2^n & 2^{n+1} - 2 \cdot 3^n \\ 3^n - 2^n & 2^{n+1} - 3^n \end{pmatrix}. \end{aligned}$$

Thus, we have a general formula that works for any n . It's worthwhile to check that we get the same answers when $n = 2$ and $n = 3$ as the direct calculations above show.

Although most of the really interesting applications of this technique will have to wait until we extend our work into higher dimensions, we look at a toy example here. Suppose that we have two *states* that we simply call state 1 and state 2. We think of time as happening in discrete increments, and as each tick of time goes by, we either stay in the current state, or switch to another. Suppose that if we are currently in state 1, then there is a 60% chance that we will stay in state 1 and a 40% chance that we will switch over into state 2 after the next tick. Suppose also that if we are currently in state 2, then there is a 10% chance that we will switch to state 1 and a 90% chance that we will stay in state 2 after the next tick. We can code this using the following matrix:

$$A = \begin{pmatrix} 0.6 & 0.1 \\ 0.4 & 0.9 \end{pmatrix}.$$

If we label the entries of A with

$$A = \begin{pmatrix} p_{1,1} & p_{1,2} \\ p_{2,1} & p_{2,2} \end{pmatrix}$$

then we are interpreting $p_{i,j}$ as the probability that, starting in state j , we end in state i after 1 step. Notice that the columns of A sum to 1, which they must because we need to move somewhere after each step.

For each i and j , let $p_{i,j}^{(2)}$ be the probability that starting in state j , we end in state i after 2 steps. How do we calculate $p_{i,j}^{(2)}$? Now there are two ways to get to state i from state j in 2 steps, because each of the two states can serve as the intermediate state. One way is to move from state j to state 1 after the first step, and from there move to state i after the second step. The probability that this happens is $p_{i,1}p_{1,j}$. The other possibility is to move from state j to state 2 after the first step, and from there move to state i after the second step. The probability that this happens is $p_{i,2}p_{2,j}$. Therefore, we have

$$p_{i,j}^{(2)} = p_{i,1} \cdot p_{1,j} + p_{i,2} \cdot p_{2,j}.$$

In other words, the number $p_{i,j}^{(2)}$ that we are looking for is the dot product of row i of A with column j of A , so it is the (i, j) -entry of A^2 ! Thus, we could compute

$$\begin{aligned} A^2 &= \begin{pmatrix} 0.6 & 0.1 \\ 0.4 & 0.9 \end{pmatrix} \begin{pmatrix} 0.6 & 0.1 \\ 0.4 & 0.9 \end{pmatrix} \\ &= \begin{pmatrix} 0.6 \cdot 0.6 + 0.1 \cdot 0.4 & 0.6 \cdot 0.1 + 0.1 \cdot 0.9 \\ 0.4 \cdot 0.6 + 0.9 \cdot 0.4 & 0.4 \cdot 0.1 + 0.9 \cdot 0.9 \end{pmatrix} \\ &= \begin{pmatrix} 0.4 & 0.15 \\ 0.6 & 0.85 \end{pmatrix}. \end{aligned}$$

Thus, $p_{1,2}^{(2)} = 0.15$, which means that there is 15% chance that, starting in state 2, we end in state 1 after exactly 2 steps. Again, notice that the columns of A^2 sum to 1, because we need to move somewhere after 2 steps.

Suppose now that we generalize this by defining $p_{i,j}^{(n)}$ to be the probability that, starting in state j , we end in state i after exactly n steps. How can we compute these numbers? Let's start by thinking about $p_{i,j}^{(3)}$. How can we get from state j to state i in 3 steps? We could think about all the possible 3 step routes that we can take, but there is a better way. We can either get from state j to state 1 in 2 steps, and then hop to state i in 1 step from there. Or we can get from j to state 2 in 2 steps, and then hop to state i in 1 step from there. In other words, we have

$$p_{i,j}^{(3)} = p_{i,1} \cdot p_{1,j}^{(2)} + p_{i,2} \cdot p_{2,j}^{(2)}.$$

Thus, $p_{i,j}^{(3)}$ is the dot product of row i of A with column j of A^2 , so it is the (i, j) -entry of A^3 ! Generalizing this idea, we have that

$$p_{i,j}^{(n+1)} = p_{i,1} \cdot p_{1,j}^{(n)} + p_{i,2} \cdot p_{2,j}^{(n)}$$

for all i, j , and n , so if we want to know $p_{i,j}^{(n)}$, then we just need to compute A^n and read off the values. Thus, if we want to know the long term behavior, we need to compute large powers of A . If we can diagonalize the corresponding linear transformation, then we can do this quickly and easily!

Without working through all of the details here, it turns out that the characteristic polynomial of A is

$$\lambda^2 - 1.5 \cdot \lambda + 0.5 = (\lambda - 1)(\lambda - 1/2)$$

Thus, the eigenvalues are 1 and $\frac{1}{2}$. An eigenvector corresponding to 1 is

$$\begin{pmatrix} 1 \\ 4 \end{pmatrix}$$

and an eigenvector corresponding to $\frac{1}{2}$ is

$$\begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

Letting

$$P = \begin{pmatrix} 1 & -1 \\ 4 & 1 \end{pmatrix} \quad \text{and} \quad D = \begin{pmatrix} 1 & 0 \\ 0 & \frac{1}{2} \end{pmatrix}$$

we have

$$P^{-1} = \begin{pmatrix} \frac{1}{5} & \frac{1}{5} \\ -\frac{4}{5} & \frac{1}{5} \end{pmatrix}$$

and

$$D = P^{-1}AP.$$

It follows that

$$A = PDP^{-1}.$$

Therefore, for any $n \in \mathbb{N}^+$, we have

$$\begin{aligned} A^n &= PD^nP^{-1} \\ &= \begin{pmatrix} 1 & -1 \\ 4 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & (\frac{1}{2})^n \end{pmatrix} \begin{pmatrix} \frac{1}{5} & \frac{1}{5} \\ -\frac{4}{5} & \frac{1}{5} \end{pmatrix} \\ &= \begin{pmatrix} 1 & -1 \\ 4 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{5} & \frac{1}{5} \\ -\frac{4}{5 \cdot 2^n} & \frac{1}{5 \cdot 2^n} \end{pmatrix} \\ &= \begin{pmatrix} \frac{1}{5} + \frac{4}{5 \cdot 2^n} & \frac{1}{5} - \frac{1}{5 \cdot 2^n} \\ \frac{4}{5} - \frac{4}{5 \cdot 2^n} & \frac{4}{5} + \frac{1}{5 \cdot 2^n} \end{pmatrix}. \end{aligned}$$

Reading off the various entries, we conclude the following for every $n \in \mathbb{N}^+$:

$$\begin{aligned} p_{1,1}^{(n)} &= \frac{1}{5} + \frac{4}{5 \cdot 2^n} \\ p_{1,2}^{(n)} &= \frac{1}{5} - \frac{1}{5 \cdot 2^n} \\ p_{2,1}^{(n)} &= \frac{4}{5} - \frac{4}{5 \cdot 2^n} \\ p_{2,2}^{(n)} &= \frac{4}{5} + \frac{1}{5 \cdot 2^n} \end{aligned}$$

We have now calculated the exact values for each $n \in \mathbb{N}^+$, but we obtain a clearer picture qualitatively if we think about what happens in long run. If $n \in \mathbb{N}^+$ is reasonably large, then regardless of which state we start in, there is an approximately 20% chance that we are in state 1 after n steps, and there is an approximately 80% chance that we are in state 2 after n steps.

We end this section with an application of diagonalization to solving a fundamental recurrence.

Example 3.3.21. Consider the Fibonacci numbers, defined by $f_0 = 0$, $f_1 = 1$, and $f_n = f_{n-1} + f_{n-2}$ for all $n \in \mathbb{N}$ with $n \geq 2$. Find a general formula for f_n .

Solution. Let

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \quad \text{and} \quad \vec{x}_0 = \begin{pmatrix} f_1 \\ f_0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

For any $x, y \in \mathbb{R}$, we have

$$A \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x+y \\ x \end{pmatrix}.$$

Thus, for all $n \in \mathbb{N}$, we have

$$A \begin{pmatrix} f_{n+1} \\ f_n \end{pmatrix} = \begin{pmatrix} f_{n+1} + f_n \\ f_{n+1} \end{pmatrix} = \begin{pmatrix} f_{n+2} \\ f_{n+1} \end{pmatrix}.$$

It follows that the second coordinate of

$$A^n \vec{x}_0$$

equals f_n . Thus, to find a general formula for f_n , we will first find a formula for A^n .

For any $\lambda \in \mathbb{R}$, we have

$$A - \lambda I = \begin{pmatrix} 1 - \lambda & 1 \\ 1 & -\lambda \end{pmatrix}$$

so the characteristic polynomial of A is

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

Using the quadratic formula, the eigenvalues of A are

$$\frac{1 \pm \sqrt{5}}{2}.$$

Let

$$\lambda_1 = \frac{1 + \sqrt{5}}{2} \quad \text{and} \quad \lambda_2 = \frac{1 - \sqrt{5}}{2}.$$

Working through the computations, we see that

$$\begin{pmatrix} \lambda_1 \\ 1 \end{pmatrix}$$

is an eigenvector of A corresponding to λ_1 , and

$$\begin{pmatrix} \lambda_2 \\ 1 \end{pmatrix}$$

is an eigenvector of A corresponding to λ_2 . Letting

$$P = \begin{pmatrix} \lambda_1 & \lambda_2 \\ 1 & 1 \end{pmatrix} \quad \text{and} \quad D = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix},$$

we have

$$P^{-1} = \frac{1}{\sqrt{5}} \cdot \begin{pmatrix} 1 & -\lambda_2 \\ -1 & \lambda_1 \end{pmatrix}.$$

Thus, we have

$$\begin{aligned} A &= PDP^{-1} \\ &= \begin{pmatrix} \lambda_1 & \lambda_2 \\ 1 & 1 \end{pmatrix} \cdot \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \cdot \frac{1}{\sqrt{5}} \cdot \begin{pmatrix} 1 & -\lambda_2 \\ -1 & \lambda_1 \end{pmatrix}. \end{aligned}$$

Therefore, for any $n \in \mathbb{N}$, we have

$$\begin{aligned} A^n \vec{x}_0 &= PD^n P^{-1} \vec{x}_0 \\ &= \begin{pmatrix} \lambda_1 & \lambda_2 \\ 1 & 1 \end{pmatrix} \cdot \begin{pmatrix} \lambda_1^n & 0 \\ 0 & \lambda_2^n \end{pmatrix} \cdot \frac{1}{\sqrt{5}} \cdot \begin{pmatrix} 1 & -\lambda_2 \\ -1 & \lambda_1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} \lambda_1 & \lambda_2 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \lambda_1^n & 0 \\ 0 & \lambda_2^n \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{5}} \\ \frac{-1}{\sqrt{5}} \end{pmatrix} \\ &= \begin{pmatrix} \lambda_1 & \lambda_2 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{5}} \cdot \lambda_1^n \\ \frac{-1}{\sqrt{5}} \cdot \lambda_2^n \end{pmatrix} \\ &= \begin{pmatrix} \frac{1}{\sqrt{5}} \cdot (\lambda_1^{n+1} - \lambda_2^{n+1}) \\ \frac{1}{\sqrt{5}} \cdot (\lambda_1^n - \lambda_2^n) \end{pmatrix}. \end{aligned}$$

Now f_n is the second coordinate of $A^n \vec{x}_0$, so

$$\begin{aligned} f_n &= \frac{1}{\sqrt{5}} \cdot (\lambda_1^n - \lambda_2^n) \\ &= \frac{1}{\sqrt{5}} \cdot \left(\left(\frac{1 + \sqrt{5}}{2} \right)^n - \left(\frac{1 - \sqrt{5}}{2} \right)^n \right) \end{aligned}$$

for all $n \in \mathbb{N}$. □

3.4 Determinants

Given two vectors in \mathbb{R}^2 , how do we determine the area of the parallelogram that they form? We know that the area of a parallelogram equals the product of the base and the height, but it may be difficult to find these, and may involve some trigonometry. We approach this question from a different perspective, by thinking about the properties of the corresponding function.

Consider then the function $g: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ defined by letting $g(\vec{v}, \vec{w})$ be the area of the parallelogram determined by \vec{v} and \vec{w} . We consider a few properties of g :

1. $g(\vec{e}_1, \vec{e}_2) = 1$ where $\vec{e}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\vec{e}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$.
2. $g(\vec{v}, \vec{v}) = 0$ for all $\vec{v} \in \mathbb{R}^2$.
3. $g(c \cdot \vec{v}, \vec{w}) = c \cdot g(\vec{v}, \vec{w})$ and $g(\vec{v}, c \cdot \vec{w}) = c \cdot g(\vec{v}, \vec{w})$ for all $\vec{v}, \vec{w} \in \mathbb{R}^2$ and all $c \in \mathbb{R}$ with $c \geq 0$.
4. $g(\vec{v}, \vec{u} + \vec{w}) = g(\vec{v}, \vec{u}) + g(\vec{v}, \vec{w})$ and $g(\vec{u} + \vec{w}, \vec{v}) = g(\vec{u}, \vec{v}) + g(\vec{w}, \vec{v})$ for all $\vec{u}, \vec{v}, \vec{w} \in \mathbb{R}^2$ such that \vec{u} and \vec{w} are on the same “side” of \vec{v} .

Property 1 simply asserts that the area of the unit square is 1. For property 2, notice that if we consider the same vector twice, then the “parallelogram” they form is a flattened 1-dimensional figure with no interior, and hence has area equal to 0. For property 3, if we scale one of the vectors by a number $c \geq 0$, then thinking of that vector as forming the base of the parallelogram, the resulting parallelogram will have a base c times as large but still have the same height. Notice that we need $c \geq 0$ because we certainly want the area to be positive. If $c < 0$, then $c \cdot \vec{v}$ points in the opposite direction of \vec{v} , and we would want to put $|c|$ in place of the c on the right of each equation. We’ll come back to this issue below.

Property 4 is the most interesting and subtle one. To see why it is true, consider first the special case when \vec{v} is on the positive x -axis, so $\vec{v} = \begin{pmatrix} r \\ 0 \end{pmatrix}$ with $r \geq 0$. Suppose that

$$\vec{u} = \begin{pmatrix} a \\ b \end{pmatrix} \quad \vec{w} = \begin{pmatrix} c \\ d \end{pmatrix}$$

with both $b \geq 0$ and $d \geq 0$, so that both \vec{u} and \vec{w} are “above” the line spanned by vector \vec{v} . Notice that we have $g(\vec{v}, \vec{u}) = rb$ because we can view the parallelogram as having base of length r (the vector \vec{v}) and height b (the y -component of \vec{u}). Similarly, we have $g(\vec{v}, \vec{w}) = rd$. Now

$$\vec{u} + \vec{w} = \begin{pmatrix} a + c \\ b + d \end{pmatrix}$$

with $b + d \geq 0$, so by the same argument we have $g(\vec{v}, \vec{u} + \vec{w}) = r(b + d)$. Since $g(\vec{v}, \vec{u}) + g(\vec{v}, \vec{w}) = rb + rd$, we conclude that

$$g(\vec{v}, \vec{u} + \vec{w}) = g(\vec{v}, \vec{u}) + g(\vec{v}, \vec{w}).$$

Now in this example, we assumed that \vec{u} and \vec{w} both had positive y -components. If instead they both had negative y -components, then we can carry out a similar argument except that $g(\vec{v}, \vec{u}) = r \cdot |b|$, $g(\vec{v}, \vec{w}) = r \cdot |d|$, and

$$g(\vec{v}, \vec{u} + \vec{w}) = r \cdot |b + d| = r \cdot (|b| + |d|)$$

where the last equality follows because b and d are both negative.

However, things get interesting in the above example (still thinking of \vec{v} as being on the positive x -axis) when exactly one of b or d is positive, while the other is negative. For example, we have

$$g(\vec{e}_1, \vec{e}_2) = 1$$

and

$$g(\vec{e}_1, -\vec{e}_2) = 1,$$

but

$$g(\vec{e}_1, \vec{e}_2 + (-\vec{e}_2)) = g(\vec{e}_1, \vec{0}) = 0,$$

so

$$g(\vec{e}_1, \vec{e}_2 + (-\vec{e}_2)) \neq g(\vec{e}_1, \vec{e}_2) + g(\vec{e}_1, -\vec{e}_2).$$

The essential problem is that when \vec{u} and \vec{w} are on different sides of the x -axis, which is the line spanned by \vec{v} , their “heights” work in opposite directions and hence do not add as we would hope.

Now the above arguments assume that \vec{v} was on the positive x -axis, but we certainly change axes by setting up a different coordinate system. If we take an arbitrary \vec{v} , then we could think of forming one axis as the line spanned by \vec{v} . Now given \vec{u} and \vec{w} , the area of the corresponding parallelogram will equal the length of \vec{v} times the distance of \vec{u} (respectively \vec{w}) from this line. If \vec{u} and \vec{w} are on the same side of the line, then the distance from $\vec{u} + \vec{w}$ to this line will be the sum of the two distances. However, if they are on different sides of this line, then this will not be true.

The restrictions on properties 3 and 4 are incredibly annoying from the perspective of computing g . For example, it is often not immediately obvious whether two vectors are on the same “side” of a given vector. Furthermore, the restrictions of 3 and 4 are not present in our much loved concept of linear transformations. So what do we do? We abandon the concept of area and move to a concept of *signed* area. That is, we want a function that works without restrictions, and so we just interpret negative signs appropriately. Although this looks like a hack, it’s actually brilliant because now the sign of the result will also give information about the *orientation* of the two vectors \vec{v} and \vec{w} with respect to each other. In other words, we will also be coding which “side” of the line spanned by \vec{v} the vector \vec{w} is on. More precisely, if \vec{w} is on the side of \vec{v} described by rotating counterclockwise by less than 180° , then we will have a positive area, while if \vec{w} is on the side of \vec{v} described by rotating clockwise by less than 180° , then we will have a negative area.

Therefore, we consider a function $f: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ with the following properties (notice that we have abandoned the restrictions on 3 and 4):

1. $f(\vec{e}_1, \vec{e}_2) = 1$ where $\vec{e}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\vec{e}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$.
2. $f(\vec{v}, \vec{v}) = 0$ for all $\vec{v} \in \mathbb{R}^2$.
3. $f(c \cdot \vec{v}, \vec{w}) = c \cdot f(\vec{v}, \vec{w})$ and $f(\vec{v}, c \cdot \vec{w}) = c \cdot f(\vec{v}, \vec{w})$ for all $\vec{v}, \vec{w} \in \mathbb{R}^2$ and all $c \in \mathbb{R}$.
4. $f(\vec{v}, \vec{u} + \vec{w}) = f(\vec{v}, \vec{u}) + f(\vec{v}, \vec{w})$ and $f(\vec{u} + \vec{w}, \vec{v}) = f(\vec{u}, \vec{v}) + f(\vec{w}, \vec{v})$ for all $\vec{u}, \vec{v}, \vec{w} \in \mathbb{R}^2$.

Before we try to compute f in general, we first examine the special case of $f(\vec{e}_2, \vec{e}_1)$. Of course, the two vectors \vec{e}_1 and \vec{e}_2 form a square with area 1, but let’s examine the difference in orientation when considering the ordered pair (\vec{e}_1, \vec{e}_2) versus the ordered pair (\vec{e}_2, \vec{e}_1) . To think about this, consider starting at the first vector, and rotating it counterclockwise around the origin until it becomes parallel to the second. When

working with the ordered pair (\vec{e}_1, \vec{e}_2) , we only have to rotate \vec{e}_1 counterclockwise by 90° (which is less than 180°) to make this happen. In contrast, when we perform the same operation with the ordered pair (\vec{e}_2, \vec{e}_1) , we have to rotate \vec{e}_2 counterclockwise by 270° (which is more than 180°) in order to line it up with \vec{e}_1 . Thus, while $f(\vec{e}_1, \vec{e}_2) = 1$, we might expect that we can derive that $f(\vec{e}_2, \vec{e}_1) = -1$ from the above properties. Instead of doing this one special case, we show that if we flip the arguments to f , then that affects the result by introducing a negative sign.

Proposition 3.4.1. *If $f: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ is a function with the above 4 properties, then $f(\vec{w}, \vec{v}) = -f(\vec{v}, \vec{w})$ for all $\vec{v}, \vec{w} \in \mathbb{R}^2$.*

Proof. Let $\vec{v}, \vec{w} \in \mathbb{R}^2$. We compute the value of $f(\vec{v} + \vec{w}, \vec{v} + \vec{w})$ in two different ways. On the one hand, we know that

$$f(\vec{v} + \vec{w}, \vec{v} + \vec{w}) = 0$$

by Property 2. On the other hand, we have

$$\begin{aligned} f(\vec{v} + \vec{w}, \vec{v} + \vec{w}) &= f(\vec{v} + \vec{w}, \vec{v}) + f(\vec{v} + \vec{w}, \vec{w}) && \text{(by Property 4)} \\ &= f(\vec{v}, \vec{v}) + f(\vec{w}, \vec{v}) + f(\vec{v}, \vec{w}) + f(\vec{w}, \vec{w}) && \text{(by Property 4)} \\ &= 0 + f(\vec{w}, \vec{v}) + f(\vec{v}, \vec{w}) + 0 && \text{(by Property 2)} \\ &= f(\vec{w}, \vec{v}) + f(\vec{v}, \vec{w}). \end{aligned}$$

Therefore, we have

$$f(\vec{w}, \vec{v}) + f(\vec{v}, \vec{w}) = 0.$$

We now obtain the result by subtracting $f(\vec{v}, \vec{w})$ from both sides. □

As a consequence, if f satisfies the above 4 properties, then we have

$$\begin{aligned} f(\vec{e}_2, \vec{e}_1) &= -f(\vec{e}_1, \vec{e}_2) \\ &= -1. \end{aligned}$$

We can use these ideas to develop a formula for f in general. The idea is to use the linearity in each coordinate to reduce the problem to situations involving \vec{e}_1 and \vec{e}_2 , just as we do for linear transformations.

Proposition 3.4.2. *If $f: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ is a function with the above 4 properties, then*

$$f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) = ad - bc$$

for all $a, b, c, d \in \mathbb{R}$.

Proof. Let $a, b, c, d \in \mathbb{R}$ be arbitrary. We have

$$\begin{aligned} f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) &= f(a \cdot \vec{e}_1 + c \cdot \vec{e}_2, b \cdot \vec{e}_1 + d \cdot \vec{e}_2) \\ &= f(a \cdot \vec{e}_1 + c \cdot \vec{e}_2, b \cdot \vec{e}_1) + f(a \cdot \vec{e}_1 + c \cdot \vec{e}_2, d \cdot \vec{e}_2) && \text{(by Property 4)} \\ &= f(a \cdot \vec{e}_1, b \cdot \vec{e}_1) + f(c \cdot \vec{e}_2, b \cdot \vec{e}_1) + f(a \cdot \vec{e}_1, d \cdot \vec{e}_2) + f(c \cdot \vec{e}_2, d \cdot \vec{e}_2) && \text{(by Property 4)} \\ &= ab \cdot f(\vec{e}_1, \vec{e}_1) + cb \cdot f(\vec{e}_2, \vec{e}_1) + ad \cdot f(\vec{e}_1, \vec{e}_2) + cd \cdot f(\vec{e}_2, \vec{e}_2) && \text{(by Property 2)} \\ &= ab \cdot 0 + cb \cdot (-1) + ad \cdot 1 + cd \cdot 0 \\ &= ad - bc. \end{aligned}$$

□

We've shown that *if* f is a function with the above properties, *then* it must satisfy the given formula. However, we have not actually shown that this function does indeed have the required properties. We do that now.

Proposition 3.4.3. *If we define $f: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ by*

$$f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) = ad - bc,$$

then f satisfies the above 4 properties.

Proof. We work through the 4 properties in turn.

1. We have

$$\begin{aligned} f(\vec{e}_1, \vec{e}_2) &= f\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) \\ &= 1 \cdot 1 - 0 \cdot 0 \\ &= 1. \end{aligned}$$

2. Let $\vec{v} \in \mathbb{R}^2$ be arbitrary. Fix $a, c \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} a \\ c \end{pmatrix}.$$

We then have

$$\begin{aligned} f(\vec{v}, \vec{v}) &= f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} a \\ c \end{pmatrix}\right) \\ &= a \cdot c - c \cdot a \\ &= 0. \end{aligned}$$

3. Let $\vec{v}, \vec{w} \in \mathbb{R}^2$ and $r \in \mathbb{R}$ be arbitrary. Fix $a, b, c, d \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} a \\ c \end{pmatrix} \quad \text{and} \quad \vec{w} = \begin{pmatrix} b \\ d \end{pmatrix}.$$

We have

$$\begin{aligned} f(r \cdot \vec{v}, \vec{w}) &= f\left(r \cdot \begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) \\ &= f\left(\begin{pmatrix} ra \\ rc \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) \\ &= rad - rbc \\ &= r \cdot (ad - bc) \\ &= r \cdot f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) \\ &= r \cdot f(\vec{v}, \vec{w}). \end{aligned}$$

A similar argument shows that $f(\vec{v}, r \cdot \vec{w}) = r \cdot f(\vec{v}, \vec{w})$.

4. Let $\vec{v}, \vec{u}, \vec{w} \in \mathbb{R}^2$ be arbitrary. Fix $a, b, c, d, r, s \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} a \\ c \end{pmatrix}, \quad \vec{u} = \begin{pmatrix} b \\ d \end{pmatrix} \quad \text{and} \quad \vec{w} = \begin{pmatrix} r \\ s \end{pmatrix}$$

We have

$$\begin{aligned} f(\vec{v}, \vec{u} + \vec{w}) &= f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix} + \begin{pmatrix} r \\ s \end{pmatrix}\right) \\ &= f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b+r \\ d+s \end{pmatrix}\right) \\ &= a(d+s) - (b+r)c \\ &= ad - bc + as - rc \\ &= f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) + f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} r \\ s \end{pmatrix}\right) \\ &= f(\vec{v}, \vec{u}) + f(\vec{v}, \vec{w}). \end{aligned}$$

A similar argument shows that $f(\vec{u} + \vec{w}, \vec{v}) = f(\vec{u}, \vec{v}) + f(\vec{w}, \vec{v})$.

□

How do parallelograms relate to linear transformations? The key fact is that linear transformations transform parallelograms into parallelograms. In other words, although a linear transformation can rotate, shear, etc., it will send straight lines to straight lines and hence not distort the plane through complicated means. To see why this is true, let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. Consider a parallelogram P , and assume for the moment that P is based at the origin and determined by the vectors \vec{v} and \vec{w} . We then have that the vertices of P are $\vec{0}$, \vec{v} , \vec{w} , and $\vec{v} + \vec{w}$. Letting $[0, 1] = \{x \in \mathbb{R} : 0 \leq x \leq 1\}$, we can then describe the set of points in P set-theoretically as

$$P = \{r \cdot \vec{v} + s \cdot \vec{w} : r, s \in [0, 1]\}.$$

If we apply T to every point in P , we obtain the set

$$\{T(r \cdot \vec{v} + s \cdot \vec{w}) : r, s \in [0, 1]\}.$$

Now T is a linear transformation, so $T(r\vec{v} + s\vec{w}) = r \cdot T(\vec{v}) + s \cdot T(\vec{w})$ for all $r, s \in [0, 1]$, and hence we can also describe this set as

$$\{r \cdot T(\vec{v}) + s \cdot T(\vec{w}) : r, s \in [0, 1]\}.$$

Thus, T sends the parallelogram P with vertices $\vec{0}$, \vec{v} , \vec{w} , and $\vec{v} + \vec{w}$ to the parallelogram with vertices $\vec{0}$, $T(\vec{v})$, $T(\vec{w})$, and $T(\vec{v}) + T(\vec{w}) = T(\vec{v} + \vec{w})$. In other words, T sends a parallelogram at the origin to a (possibly different) parallelogram at the origin. From here, we can obtain general parallelograms by using an offset. Suppose then that $\vec{p}, \vec{v}, \vec{w} \in \mathbb{R}^2$ and consider the parallelogram

$$P = \{\vec{p} + r \cdot \vec{v} + s \cdot \vec{w} : r, s \in [0, 1]\}.$$

As above, if we apply T to all elements of P and use the fact that T is a linear transformation, then we see that P gets sent to the set

$$\{T(\vec{p}) + r \cdot T(\vec{v}) + s \cdot T(\vec{w}) : r, s \in [0, 1]\},$$

which is also a parallelogram (determined by the vectors $T(\vec{v})$ and $T(\vec{w})$, but offset by the vector $T(\vec{p})$).

Now although linear transformations send parallelograms to parallelograms, they may change the area of these parallelograms. For example, consider the unique linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with

$$[T] = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}.$$

We've seen before that this linear transformation rotates the plane by 45° counterclockwise and scales the plane by a factor of $\sqrt{2}$. Consider the unit square, i.e. the parallelogram formed by \vec{e}_1 and \vec{e}_2 , which clearly has area 1. Now since

$$T(\vec{e}_1) = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad \text{and} \quad T(\vec{e}_2) = \begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

the above argument shows that T maps the unit square to the parallelogram determined by

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

Since

$$f\left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix}\right) = 1 \cdot 1 - (-1) \cdot 1 = 2,$$

it follows that T turns the unit square of area 1 into a parallelogram (in fact a square) with area 2. We also could have realized this geometrically, by noticing that the side lengths of the resulting square are $\sqrt{2}$. Now it turns out that T will double the area of *every* parallelogram. One can argue this geometrically using the fact that T scales by $\sqrt{2}$, but in fact this is not special to this particular T at all, as we now show.

Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a linear transformation. Let $\vec{p}, \vec{v}, \vec{w} \in \mathbb{R}^2$, and consider the parallelogram

$$P = \{\vec{p} + r\vec{v} + s\vec{w} : r, s \in [0, 1]\}.$$

Since \vec{p} is just serving as an offset, we know that the signed area of this parallelogram is $f(\vec{v}, \vec{w})$. Now we know from above that T sends P to the parallelogram

$$\{T(\vec{p}) + r \cdot T(\vec{v}) + s \cdot T(\vec{w}) : r, s \in [0, 1]\},$$

and we know that the signed area of this parallelogram is $f(T(\vec{v}), T(\vec{w}))$. Can we determine how this number relates to the number $f(\vec{v}, \vec{w})$? Fixing $v_1, v_2, w_1, w_2 \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \quad \text{and} \quad \vec{w} = \begin{pmatrix} w_1 \\ w_2 \end{pmatrix},$$

and using the properties of f , we have

$$\begin{aligned} f(T(\vec{v}), T(\vec{w})) &= f(T(v_1 \cdot \vec{e}_1 + v_2 \cdot \vec{e}_2), T(w_1 \cdot \vec{e}_1 + w_2 \cdot \vec{e}_2)) \\ &= f(v_1 \cdot T(\vec{e}_1) + v_2 \cdot T(\vec{e}_2), w_1 \cdot T(\vec{e}_1) + w_2 \cdot T(\vec{e}_2)) \\ &= v_1 w_1 \cdot f(T(\vec{e}_1), T(\vec{e}_1)) + v_1 w_2 \cdot f(T(\vec{e}_1), T(\vec{e}_2)) \\ &\quad + v_2 w_1 \cdot f(T(\vec{e}_2), T(\vec{e}_1)) + v_2 w_2 \cdot f(T(\vec{e}_2), T(\vec{e}_2)) \\ &= v_1 w_1 \cdot 0 + v_1 w_2 \cdot f(T(\vec{e}_1), T(\vec{e}_2)) - v_2 w_1 \cdot f(T(\vec{e}_1), T(\vec{e}_2)) + v_2 w_2 \cdot 0 \\ &= (v_1 w_2 - v_2 w_1) \cdot f(T(\vec{e}_1), T(\vec{e}_2)) \\ &= f(\vec{v}, \vec{w}) \cdot f(T(\vec{e}_1), T(\vec{e}_2)). \end{aligned}$$

Now if we fix $a, b, c, d \in \mathbb{R}$ with

$$[T] = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

then we have

$$\begin{aligned} f(T(\vec{e}_1), T(\vec{e}_2)) &= f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) \\ &= ad - bc, \end{aligned}$$

and hence

$$f(T(\vec{v}), T(\vec{w})) = f(\vec{v}, \vec{w}) \cdot (ad - bc).$$

for all $\vec{v}, \vec{w} \in \mathbb{R}^2$. Therefore, T sends any parallelogram with signed area r to a parallelogram with signed area $r \cdot (ad - bc)$. In other words, this number $ad - bc$ gives the (*signed*) *area distortion factor* of the linear transformation T . The sign here represents orientation. If $ad - bc$ is positive, then T will preserve the orientation of two input vectors. If $ad - bc$ is negative, then T will reverse the orientation of two input vectors. We give this numerical distortion factor a special name.

Definition 3.4.4. *Given a matrix*

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

we define $\det(A) = ad - bc$ and call this number the determinant of A . We also write

$$\begin{vmatrix} a & b \\ c & d \end{vmatrix}$$

for the determinant of A .

Always remember that the determinant of a matrix is not just some random number, but gives the (signed) area distortion factor of the corresponding linear transformation. With this perspective, we can now geometrically interpret why a matrix A is not invertible when $\det(A) = 0$, because in this case the corresponding linear transformation will collapse parallelograms with positive area onto flattened 1-dimensional “parallelograms”, and hence the corresponding linear transformation will not be injective.

Chapter 4

Beyond Two Dimensions

Now that we have a solid understanding of the linear algebra of \mathbb{R}^2 , it's time to move on to higher dimensions. A natural approach would be to generalize everything to the setting of \mathbb{R}^n for an arbitrary $n \in \mathbb{N}^+$. Although several of the concepts, and much of the work, that we developed in \mathbb{R}^2 carry over naturally to this setting, it turns out that some wrinkles arise. For example, in two dimensions, if we have two nonzero vectors $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$ that are not multiples of each other, then $\text{Span}(\vec{u}_1, \vec{u}_2) = \mathbb{R}^2$, i.e. we can reach everything in \mathbb{R}^2 by using only scalar multiplication and addition on \vec{u}_1 and \vec{u}_2 . In contrast, given two vectors in \mathbb{R}^3 that are not multiples of each other, their span geometrically appears to be a plane through the origin, so it looks as though we cannot reach everything in \mathbb{R}^3 . If we have three nonzero vectors $\vec{u}_1, \vec{u}_2, \vec{u}_3 \in \mathbb{R}^3$ such that none is a multiple of another, do we necessarily have $\text{Span}(\vec{u}_1, \vec{u}_2, \vec{u}_3) = \mathbb{R}^3$? This generalization may seem natural, but unfortunately it is false, because the three vectors may still all lie in the same plane through the origin. For example, if

$$\vec{u}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \quad \vec{u}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad \text{and} \quad \vec{u}_3 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix},$$

then $\text{Span}(\vec{u}_1, \vec{u}_2, \vec{u}_3) \neq \mathbb{R}^3$ (because every element of the span will lie in the xy -plane). Thus, the added dimension adds a new level of complexity.

When we move beyond three dimensions, we lose much of our geometric intuition, and certain questions that seem visibly clear in two or three dimensions become opaque. For example, we just gave an intuitive geometric argument that we cannot reach everything in \mathbb{R}^3 by just scaling and adding 2 vectors. Is it possible to find 6 vectors in \mathbb{R}^7 such that we can reach everything in \mathbb{R}^7 by just scaling and adding these 6 vectors? It is natural to hope that since we are in \mathbb{R}^7 , we might need 7 vectors to reach everything. However, this is far from obvious, and our inability to visualize \mathbb{R}^7 makes this question particularly tricky. We will soon work to develop algebraic techniques to answer these questions.

Before jumping into a general study of \mathbb{R}^n , let's take a step back. All of the linear algebra of \mathbb{R}^2 relied on two fundamental operations: vector addition and scalar multiplication. We can certainly define these operations in \mathbb{R}^n , but there are many other types of mathematical objects that we can add together and multiply by scalars. In fact, whenever we were able to avoid “opening up” an element of \mathbb{R}^2 into its components, all that we were using was the fact that these operations satisfied certain basic rules, so those arguments should carry over to *any* situation where we can add and scalar multiply coherently.

4.1 Vector Spaces and Subspaces

We begin by defining the abstract concept of a *vector space*, which is a world where we can both add elements and multiply elements by scalars. In other words, a vector space is any place where we can take

linear combinations of its elements in a natural way. Although vector space has the word *vector* in its name, the elements of it need not be “vectors” in any traditional algebraic or geometric sense. In other words, they need not have qualities like magnitude or direction. All that matters is that we can add the elements together and multiply the elements by scalars. For every positive natural number n , the set \mathbb{R}^n will be a vector space, but we will see many other examples as well.

One of the most important examples is *functions*. For example, if $f: \mathbb{R} \rightarrow \mathbb{R}$ is the function $f(x) = x^2 + 2x - 5$, and $g: \mathbb{R} \rightarrow \mathbb{R}$ is the function $g(x) = x + 9$, then we can *add* f and g to form the function $f + g: \mathbb{R} \rightarrow \mathbb{R}$ given by $(f + g)(x) = f(x) + g(x) = x^2 + 3x + 4$. We can also multiply functions by scalars. If $f: \mathbb{R} \rightarrow \mathbb{R}$ is the function $f(x) = x^2 + 2x - 5$, then we can multiply f by 7 to form the function $7 \cdot f: \mathbb{R} \rightarrow \mathbb{R}$ given by $(7 \cdot f)(x) = 7 \cdot f(x) = 7x^2 + 14x - 35$. Now functions may appear to be fundamentally different objects than elements of \mathbb{R}^n , but it turns out that addition and scalar multiplication of functions behaves in many ways just like addition and scalar multiplication of elements of \mathbb{R}^n .

Is this level of generality useful? Why would we want to apply the ideas of linear algebra to *functions*? One natural place where linear algebra on functions arises is in the study of differential equations. Unlike algebraic equations where we want to find which *numbers* $x \in \mathbb{R}$ making the equation true, in differential equations we want to find which *functions* we can plug in to get a true statement. For example, suppose that we want to determine all functions $f: \mathbb{R} \rightarrow \mathbb{R}$ such that $f'' + f = 0$, where f'' is the second derivative of f . In other words, we want to determine all functions $f: \mathbb{R} \rightarrow \mathbb{R}$ such that $f''(x) + f(x) = 0$ for all $x \in \mathbb{R}$. With some thought, you might stumble on the solution $f(x) = \sin x$ and the solution $f(x) = \cos x$. Are there any others? It turns out that we can build new solutions from old solutions. Suppose that g_1 and g_2 are arbitrary solutions, and that $c_1, c_2 \in \mathbb{R}$ are arbitrary. For any $x \in \mathbb{R}$, we have

$$\begin{aligned} (c_1 g_1 + c_2 g_2)''(x) + (c_1 g_1 + c_2 g_2)(x) &= c_1 g_1''(x) + c_2 g_2''(x) + c_1 g_1(x) + c_2 g_2(x) \\ &= c_1 g_1''(x) + c_1 g_1(x) + c_2 g_2''(x) + c_2 g_2(x) \\ &= c_1 \cdot (g_1''(x) + g_1(x)) + c_2 \cdot (g_2''(x) + g_2(x)) \\ &= c_1 \cdot 0 + c_2 \cdot 0 \\ &= 0. \end{aligned}$$

In other words, if g_1 and g_2 are both solutions to $f'' + f = 0$, then any linear combination of g_1 and g_2 is also a solution to $f'' + f = 0$. In particular, the function $h(x) = 2 \sin x - 5 \cos x$ is also a solution, as is *any* linear combination of $\sin x$ and $\cos x$. Thus, we see how linear combinations of functions can arise naturally when trying to find solutions to these differential equations.

As we progress, we will see many other places where linear algebra ideas can be applied to functions in order to solve fundamental problems. In order to include these examples (and many others to come in future math courses), we define the following very general concept.

Definition 4.1.1. A vector space is a nonempty set V of objects, called *vectors*, equipped with operations of addition and scalar multiplication, along with an element $\vec{0}$, such that the following properties hold:

1. For all $\vec{v}, \vec{w} \in V$, we have $\vec{v} + \vec{w} \in V$ (closure under addition).
2. For all $\vec{v} \in V$ and all $c \in \mathbb{R}$, we have $c \cdot \vec{v} \in V$ (closure under scalar multiplication).
3. For all $\vec{v}, \vec{w} \in V$, we have $\vec{v} + \vec{w} = \vec{w} + \vec{v}$ (commutativity of addition).
4. For all $\vec{u}, \vec{v}, \vec{w} \in V$, we have $(\vec{u} + \vec{v}) + \vec{w} = \vec{u} + (\vec{v} + \vec{w})$ (associativity of addition).
5. For all $\vec{v} \in V$, we have $\vec{v} + \vec{0} = \vec{v}$ ($\vec{0}$ is an additive identity).
6. For all $\vec{v} \in V$, there exists $\vec{w} \in V$ such that $\vec{v} + \vec{w} = \vec{0}$ (existence of additive inverses).
7. For all $\vec{v}, \vec{w} \in V$ and all $c \in \mathbb{R}$, we have $c \cdot (\vec{v} + \vec{w}) = c \cdot \vec{v} + c \cdot \vec{w}$.

8. For all $\vec{v} \in V$ and all $c, d \in \mathbb{R}$, we have $(c + d) \cdot \vec{v} = c \cdot \vec{v} + d \cdot \vec{v}$.
9. For all $\vec{v} \in V$ and all $c, d \in \mathbb{R}$, we have $c \cdot (d \cdot \vec{v}) = (cd) \cdot \vec{v}$.
10. For all $\vec{v} \in V$, we have $1 \cdot \vec{v} = \vec{v}$.

Read the above definition carefully. A vector space is a world where we can add elements and multiply by scalars. However, we don't necessarily know *how* these operations are calculated, only that they satisfy certain fundamental properties. Also, even though we call the elements of V *vectors*, they may not be vectors in any traditional sense. In particular, they may not have anything resembling “magnitude” or “direction” that we are used to. Furthermore, we may not be able to break them into “components” like we do in \mathbb{R}^2 , and moreover we may not be able to visualize the elements of V at all. A vector space is *any* collection of objects where we have operations of addition and scalar multiplication that satisfy the above properties.

We now give a few important examples of vector spaces. Notice that Proposition 2.2.1 exactly says that \mathbb{R}^2 , with the normal operations of addition and scalar multiplication, is a vector space. We can easily generalize this to any \mathbb{R}^n .

Example 4.1.2. Let $n \in \mathbb{N}^+$. Consider the set \mathbb{R}^n of all n -tuples of real numbers. Define addition and scalar multiplication on \mathbb{R}^n by letting

$$\begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = \begin{pmatrix} a_1 + b_1 \\ a_2 + b_2 \\ \vdots \\ a_n + b_n \end{pmatrix} \quad \text{and} \quad c \cdot \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} = \begin{pmatrix} ca_1 \\ ca_2 \\ \vdots \\ ca_n \end{pmatrix}.$$

Finally, let

$$\vec{0} = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}.$$

With these operations, \mathbb{R}^n is a vector space.

To verify that \mathbb{R}^n with these operations is indeed a vector space, we need to check that all 10 of the above axioms are true. The first two are trivial. For the third, notice that for any $a_1, a_2, \dots, a_n, b_1, b_2, \dots, b_n \in \mathbb{R}$, we have

$$\begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = \begin{pmatrix} a_1 + b_1 \\ a_2 + b_2 \\ \vdots \\ a_n + b_n \end{pmatrix}$$

and

$$\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} + \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} = \begin{pmatrix} b_1 + a_1 \\ b_2 + a_2 \\ \vdots \\ b_n + a_n \end{pmatrix},$$

so

$$\begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} + \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix}$$

because $+$ is commutative on \mathbb{R} . Working through all of the other properties, we see that since addition and scalar multiplication happens component-wise, each of these properties follow from the fact that similar properties hold for the real numbers. The other properties can all be verified in a similar way, and fundamentally all boil down to the fact that \mathbb{R} with the usual operations have the necessary properties.

Example 4.1.3. Consider the set \mathcal{F} of all functions $f: \mathbb{R} \rightarrow \mathbb{R}$ (e.g. the function $f(x) = \sin x$ is an element of \mathcal{F} , the function $g(x) = |x|$ is an element of \mathcal{F} , etc.). Define addition and scalar multiplication on \mathcal{F} in the usual way:

- Given $f, g: \mathbb{R} \rightarrow \mathbb{R}$, let $f + g: \mathbb{R} \rightarrow \mathbb{R}$ be the function defined by $(f + g)(x) = f(x) + g(x)$ for all $x \in \mathbb{R}$.
- Given $f: \mathbb{R} \rightarrow \mathbb{R}$ and $c \in \mathbb{R}$, let $c \cdot f: \mathbb{R} \rightarrow \mathbb{R}$ be the function defined by $(c \cdot f)(x) = c \cdot f(x)$ for all $x \in \mathbb{R}$.

Finally, let $\vec{0}$ be the zero function, i.e. $\vec{0}$ is the function $z: \mathbb{R} \rightarrow \mathbb{R}$ given by $z(x) = 0$ for all $x \in \mathbb{R}$. With these operations, \mathcal{F} is a vector space.

Since function addition and scalar operation happens “pointwise”, we can check each of the 10 axioms by referencing the fact that \mathbb{R} satisfies the necessary properties. For example, consider Property 4. Let $f, g, h: \mathbb{R} \rightarrow \mathbb{R}$ be arbitrary functions. For any $x \in \mathbb{R}$, we have

$$\begin{aligned} (f + (g + h))(x) &= f(x) + (g + h)(x) \\ &= f(x) + (g(x) + h(x)) \\ &= (f(x) + g(x)) + h(x) && \text{(since } + \text{ is associative on } \mathbb{R}) \\ &= (f + g)(x) + h(x) \\ &= ((f + g) + h)(x). \end{aligned}$$

Since $x \in \mathbb{R}$ was arbitrary, we conclude that $f + (g + h) = (f + g) + h$. The other axioms are similar.

Example 4.1.4. Let $\mathcal{M}_{2 \times 2}$ be the set of all 2×2 matrices. Define addition and scalar multiplication on \mathcal{M} as we defined them earlier:

$$\begin{pmatrix} a_1 & b_1 \\ c_1 & d_1 \end{pmatrix} + \begin{pmatrix} a_2 & b_2 \\ c_2 & d_2 \end{pmatrix} = \begin{pmatrix} a_1 + a_2 & b_1 + b_2 \\ c_1 + c_2 & d_1 + d_2 \end{pmatrix} \quad \text{and} \quad r \cdot \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} ra & rb \\ rc & rd \end{pmatrix}.$$

Finally, let $\vec{0}$ be the zero matrix, i.e. $\vec{0}$ is the matrix

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

With these operations, $\mathcal{M}_{2 \times 2}$ is a vector space.

In this case, most of the vector space axioms follow from Proposition 2.6.6 and Proposition 2.6.8.

Example 4.1.5. Let \mathcal{S} be the set of all infinite sequences of real numbers. For example, $(1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \dots)$ is an element of \mathcal{S} , and a general element of \mathcal{S} can be written as $(a_1, a_2, a_3, a_4, \dots)$ where each $a_i \in \mathbb{R}$. Define addition and scalar multiplication on \mathcal{S} as follows:

- $(a_1, a_2, a_3, a_4, \dots) + (b_1, b_2, b_3, b_4, \dots) = (a_1 + b_1, a_2 + b_2, a_3 + b_3, a_4 + b_4, \dots)$.
- $c \cdot (a_1, a_2, a_3, a_4, \dots) = (ca_1, ca_2, ca_3, ca_4, \dots)$.

Finally, let $\vec{0}$ be the sequence of all zeros, i.e. $\vec{0} = (0, 0, 0, 0, \dots)$. With these operations, \mathcal{S} is a vector space.

The verification of the axioms here is very similar to the verification of the axioms for \mathbb{R}^n . In fact, an element of \mathcal{S} is just like an element of \mathbb{R}^n , except the entries go on forever. We end with a more bizarre example.

Example 4.1.6. Let $V = \{x \in \mathbb{R} : x > 0\}$ be the set of all positive real numbers. We will define a strange “addition” and “scalar multiplication” on V , and to avoid confusion, we use \oplus for our definition of addition, and \odot for our definition of scalar multiplication.

- Given $a, b \in V$, let $a \oplus b = a \cdot b$, i.e. $a \oplus b$ is usual product of a and b .
- Given $a \in V$ and $c \in \mathbb{R}$, let $c \odot a = a^c$ (notice that $a^c > 0$ because $a > 0$).

Finally, let $\vec{0} = 1$. With these operations, V is a vector space.

Let’s verify Property 7. Let $a, b \in V$ and $c \in \mathbb{R}$ be arbitrary. We have

$$\begin{aligned} c \odot (a \oplus b) &= c \odot (ab) \\ &= (ab)^c \end{aligned}$$

and

$$\begin{aligned} (c \odot a) \oplus (c \odot b) &= a^c \oplus b^c \\ &= a^c b^c, \end{aligned}$$

so since $(ab)^c = a^c b^c$, we conclude that $c \odot (a \oplus b) = a^c b^c$. The others follow from similar properties of multiplication and exponentiation.

We have now seen several examples of vector spaces, but we will see that there are many others. Moreover, you will see more exotic examples of vector spaces in later courses. For example, in quantum mechanics, one can view the possible quantum states as elements of a certain vector space. In fact, in that setting, you will use complex numbers as scalars in place of real numbers. Linear combinations of these quantum states using complex scalars are at the heart of many of the strange and counterintuitive features of the quantum world, such as quantum entanglement and quantum computation.

The beauty of encapsulating many very different concrete realizations into the one abstract concept of a vector space is that we can prove theorems about general vector spaces, and then use such theorems in each particular manifestation. We start this process now by proving several additional algebraic properties of vector spaces that one might have considered including in the above 10, but were omitted because they can be derived from the others.

Proposition 4.1.7. Let V be a vector space.

1. For all $\vec{v} \in V$, we have $\vec{0} + \vec{v} = \vec{v}$.
2. For all $\vec{v}, \vec{w} \in V$, if $\vec{v} + \vec{w} = \vec{0}$, then $\vec{w} + \vec{v} = \vec{0}$.

Proof. 1. Let $\vec{v} \in V$ be arbitrary. We have

$$\begin{aligned} \vec{0} + \vec{v} &= \vec{v} + \vec{0} && \text{(by Property 3)} \\ &= \vec{v} && \text{(by Property 5).} \end{aligned}$$

2. Let $\vec{v}, \vec{w} \in V$ be arbitrary with $\vec{v} + \vec{w} = \vec{0}$. We then have

$$\begin{aligned} \vec{w} + \vec{v} &= \vec{v} + \vec{w} && \text{(by Property 3)} \\ &= \vec{0}. \end{aligned}$$

□

Proposition 4.1.8. *Suppose that V is a vector space, and let $\vec{u}, \vec{v}, \vec{w} \in V$.*

1. *If $\vec{v} + \vec{u} = \vec{w} + \vec{u}$, then $\vec{v} = \vec{w}$.*
2. *If $\vec{u} + \vec{v} = \vec{u} + \vec{w}$, then $\vec{v} = \vec{w}$.*

Proof. 1. Let $\vec{u}, \vec{v}, \vec{w} \in V$ be arbitrary with $\vec{v} + \vec{u} = \vec{w} + \vec{u}$. By Property 6, we can fix $\vec{z} \in V$ with $\vec{u} + \vec{z} = \vec{0}$. We then have

$$\begin{aligned}
 \vec{v} &= \vec{v} + \vec{0} && \text{(by Property 5)} \\
 &= \vec{v} + (\vec{u} + \vec{z}) \\
 &= (\vec{v} + \vec{u}) + \vec{z} && \text{(by Property 4)} \\
 &= (\vec{w} + \vec{u}) + \vec{z} && \text{(by assumption)} \\
 &= \vec{w} + (\vec{u} + \vec{z}) && \text{(by Property 4)} \\
 &= \vec{w} + \vec{0} \\
 &= \vec{w} && \text{(by Property 5),}
 \end{aligned}$$

so $\vec{v} = \vec{w}$.

2. Let $\vec{u}, \vec{v}, \vec{w} \in V$ be arbitrary with $\vec{u} + \vec{v} = \vec{u} + \vec{w}$. Using Property 3 on both sides, we then have $\vec{v} + \vec{u} = \vec{w} + \vec{u}$. Using the first part of this proposition, we conclude that $\vec{v} = \vec{w}$. □

Proposition 4.1.9. *Let V be a vector space.*

1. *If $\vec{z} \in V$ and $\vec{v} + \vec{z} = \vec{v}$ for all $\vec{v} \in V$, then $\vec{z} = \vec{0}$, i.e. $\vec{0}$ is the only additive identity.*
2. *For all $\vec{v} \in V$, there is a unique $\vec{w} \in V$ with $\vec{v} + \vec{w} = \vec{0}$, i.e. additive inverses are unique.*

Proof. 1. Let $\vec{z} \in V$, and suppose that $\vec{v} + \vec{z} = \vec{v}$ for all $\vec{v} \in V$. We have

$$\begin{aligned}
 \vec{z} &= \vec{z} + \vec{0} && \text{(by Property 5)} \\
 &= \vec{0} + \vec{z} && \text{(by Property 3)} \\
 &= \vec{0} && \text{(by assumption with } \vec{v} = \vec{0}\text{).}
 \end{aligned}$$

2. Let $\vec{v} \in V$ be arbitrary, and suppose that \vec{u} and \vec{w} are both inverses. We then have both $\vec{v} + \vec{u} = \vec{0}$ and also $\vec{v} + \vec{w} = \vec{0}$. It follows that $\vec{v} + \vec{u} = \vec{v} + \vec{w}$, so $\vec{u} = \vec{w}$. □

Definition 4.1.10. *Let V be a vector space.*

1. *Given $\vec{v} \in V$, we denote the unique additive inverse of \vec{v} by $-\vec{v}$.*
2. *Given $\vec{v}, \vec{w} \in V$, we define $\vec{v} - \vec{w}$ to be $\vec{v} + (-\vec{w})$, i.e. $\vec{v} - \vec{w}$ is the sum of \vec{v} and the unique additive inverse of \vec{w} .*

Proposition 4.1.11. *Let V be a vector space.*

1. $0 \cdot \vec{v} = \vec{0}$ for all $\vec{v} \in V$.
2. $c \cdot \vec{0} = \vec{0}$ for all $c \in \mathbb{R}$.
3. $(-1) \cdot \vec{v} = -\vec{v}$ for all $\vec{v} \in V$.

Proof. 1. Let $\vec{v} \in V$ be arbitrary. We have

$$\begin{aligned} 0 \cdot \vec{v} + \vec{0} &= 0 \cdot \vec{v} && \text{(by Property 5)} \\ &= (0 + 0) \cdot \vec{v} \\ &= 0 \cdot \vec{v} + 0 \cdot \vec{v} && \text{(by Property 8),} \end{aligned}$$

so $0 \cdot \vec{v} + \vec{0} = 0 \cdot \vec{v} + 0 \cdot \vec{v}$. Using Proposition 4.1.8, it follows that $0 \cdot \vec{v} = \vec{0}$.

2. Let $c \in \mathbb{R}$ be arbitrary. We have

$$\begin{aligned} c \cdot \vec{0} + \vec{0} &= c \cdot \vec{0} && \text{(by Property 5)} \\ &= c \cdot (\vec{0} + \vec{0}) && \text{(by Property 5)} \\ &= c \cdot \vec{0} + c \cdot \vec{0} && \text{(by Property 7),} \end{aligned}$$

so $c \cdot \vec{0} + \vec{0} = c \cdot \vec{0} + c \cdot \vec{0}$. Using Proposition 4.1.8, it follows that $c \cdot \vec{0} = \vec{0}$.

3. Let $\vec{v} \in V$ be arbitrary. We have

$$\begin{aligned} \vec{v} + (-1) \cdot \vec{v} &= 1 \cdot \vec{v} + (-1) \cdot \vec{v} && \text{(by Property 10)} \\ &= (1 + (-1)) \cdot \vec{v} && \text{(by Property 8)} \\ &= 0 \cdot \vec{v} \\ &= \vec{0} && \text{(by part 1),} \end{aligned}$$

so $(-1) \cdot \vec{v}$ is the additive inverse for \vec{v} . Therefore, $(-1) \cdot \vec{v} = -\vec{v}$ by definition of $-\vec{v}$. □

Suppose that V is a vector space, and that W is an arbitrary *subset* of V , i.e. that $W \subseteq V$. Can we view W as a vector space in its own right by “inheriting” the operations of addition and scalar multiplication from V , i.e. by using the addition and scalar multiplication that V provides? Given any $\vec{w}_1, \vec{w}_2 \in W$, we automatically know that $\vec{w}_1 + \vec{w}_2 = \vec{w}_2 + \vec{w}_1$ because this is true in V , and we are using the same addition that we are using in V (this is what we mean by “inheriting” the operation). We also have associativity of addition and Properties 7 through 10 for similar reasons. However, some of the other properties are less clear. It certainly seems possible that given $\vec{w}_1, \vec{w}_2 \in W$, we might have that $\vec{w}_1 + \vec{w}_2 \notin W$, i.e. although $\vec{w}_1 + \vec{w}_2$ must be an element of V (because V is a vector space), there is no reason to believe that it falls into our subset W . For example, if $V = \mathbb{R}$ and $W = \{2, 3\}$, then $2 + 3 \notin W$ even though $2 + 3 \in V$. Similarly, there is no reason to believe that Property 2 and Property 5 hold in W either. Subsets of W where these three properties do hold are given a special name.

Definition 4.1.12. Let V be a vector space. A *subspace* of V is a subset $W \subseteq V$ with the following properties:

- $\vec{0} \in W$.
- For all $\vec{w}_1, \vec{w}_2 \in W$, we have $\vec{w}_1 + \vec{w}_2 \in W$.
- For all $\vec{w} \in W$ and all $c \in \mathbb{R}$, we have $c \cdot \vec{w} \in W$.

You may rightly point out that Property 6 does not look like it will hold automatically either. However, if W is a subspace of V , then given any $\vec{w} \in W$, we have that $(-1) \cdot \vec{w} \in W$ because $-1 \in \mathbb{R}$, and since $-\vec{w} = (-1)\vec{w}$ by Proposition 4.1.11, we conclude that the additive inverse of \vec{w} is in W as well. Therefore, if W is a subspace of V , then we can view W as a vector space in its own right under the inherited operations from V .

For an example of a subspace, we claim that the set

$$W = \left\{ \begin{pmatrix} a \\ 2a \end{pmatrix} : a \in \mathbb{R} \right\}$$

is a subspace of \mathbb{R}^2 . To see this, we check the three properties.

- First notice that $\vec{0} \in W$ by taking $a = 0$.
- Let $\vec{w}_1, \vec{w}_2 \in W$ be arbitrary. By definition, we can fix $a_1, a_2 \in \mathbb{R}$ with

$$\vec{w}_1 = \begin{pmatrix} a_1 \\ 2a_1 \end{pmatrix} \quad \text{and} \quad \vec{w}_2 = \begin{pmatrix} a_2 \\ 2a_2 \end{pmatrix}.$$

We have

$$\begin{aligned} \vec{w}_1 + \vec{w}_2 &= \begin{pmatrix} a_1 \\ 2a_1 \end{pmatrix} + \begin{pmatrix} a_2 \\ 2a_2 \end{pmatrix} \\ &= \begin{pmatrix} a_1 + a_2 \\ 2a_1 + 2a_2 \end{pmatrix} \\ &= \begin{pmatrix} a_1 + a_2 \\ 2(a_1 + a_2) \end{pmatrix}, \end{aligned}$$

so since $a_1 + a_2 \in \mathbb{R}$, it follows that $\vec{w}_1 + \vec{w}_2 \in W$.

- Let $\vec{w} \in W$ and $c \in \mathbb{R}$ be arbitrary. By definition, we can fix $a \in \mathbb{R}$ with

$$\vec{w} = \begin{pmatrix} a \\ 2a \end{pmatrix}.$$

We have

$$\begin{aligned} c \cdot \vec{w} &= c \cdot \begin{pmatrix} a \\ 2a \end{pmatrix} \\ &= \begin{pmatrix} ac \\ 2 \cdot (ac) \end{pmatrix}, \end{aligned}$$

so since $ac \in \mathbb{R}$, it follows that $c \cdot \vec{w} \in W$.

Therefore, W is a subspace of \mathbb{R}^2 .

For another example, given any vector space V , we have that $\{\vec{0}\}$ is a subspace of V . To see this, first notice that $\vec{0} \in \{\vec{0}\}$ trivially. We also have $\vec{0} + \vec{0} \in \{\vec{0}\}$ by definition of $\vec{0}$ and $c \cdot \vec{0} \in \{\vec{0}\}$ for all $c \in \mathbb{R}$ by Proposition 4.1.11.

For an interesting example of a subset of a vector space that is *not* a subspace, consider the vector space $V = \mathbb{R}^2$, and let

$$W = \left\{ \begin{pmatrix} a \\ b \end{pmatrix} \in \mathbb{R}^2 : a \in \mathbb{Z} \text{ and } b \in \mathbb{Z} \right\}.$$

To see that W is not a subspace of \mathbb{R}^2 , notice that

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix} \in W,$$

but

$$\frac{1}{2} \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix} \notin W,$$

so W is not closed under scalar multiplication. Since we have found a counterexample to the third property, it follows that W is not a subspace of \mathbb{R}^2 . However, it is straightforward to check that W satisfies the first two properties.

Let \mathcal{F} be the vector space of all functions $f: \mathbb{R} \rightarrow \mathbb{R}$. There are many interesting examples of subspaces of \mathcal{F} . For example, we have that following.

- Let \mathcal{C} be the subset of \mathcal{F} consisting of all continuous functions, i.e.

$$\mathcal{C} = \{f \in \mathcal{F} : f \text{ is continuous at every } x \in \mathbb{R}\}.$$

Now the function $f(x) = 0$ (which is the zero vector of \mathcal{F}) is continuous, the sum of two continuous functions is continuous, and a scalar multiple of a continuous function is continuous (the latter two statements are established in Calculus). Therefore, \mathcal{C} is a subspace of \mathcal{F} .

- Let \mathcal{D} be the subset of \mathcal{F} consisting of all differentiable functions, i.e.

$$\mathcal{D} = \{f \in \mathcal{F} : f \text{ is differentiable at every } x \in \mathbb{R}\}.$$

Now the function $f(x) = 0$ (which is the zero vector of \mathcal{F}) is differentiable, the sum of two differentiable functions is differentiable, and a scalar multiple of a differentiable function is differentiable (the latter two statements are established in Calculus). Therefore, \mathcal{D} is a subspace of \mathcal{F} . In fact, since every differentiable function is continuous, we can view \mathcal{D} as a subspace of \mathcal{C} .

We can obtain other examples of subspaces of \mathcal{F} , indeed of \mathcal{D} , by considering certain special functions.

Definition 4.1.13. A function $f: \mathbb{R} \rightarrow \mathbb{R}$ is called a polynomial function if there exists an $n \in \mathbb{N}$ and $a_0, a_1, \dots, a_{n-1}, a_n \in \mathbb{R}$ such that $f(x) = a_n x^n + a_{n-1} x^{n-1} + \dots + a_1 x + a_0$ for all $x \in \mathbb{R}$.

Recall that the degree of a polynomial is the largest power of x having a nonzero coefficient (and we do not define the degree of the zero polynomial because it has no nonzero coefficients). Using these polynomials, we have the following additional subspaces of \mathcal{F} .

- Let \mathcal{P} be the subset of \mathcal{F} consisting of all polynomial functions. Since the zero function is a polynomial function, the sum of two polynomial functions is a polynomial function, and a scalar multiple of a polynomial function is a polynomial function, it follows that \mathcal{P} is a subspace of \mathcal{F} . In fact, since all polynomials are differentiable, we have that \mathcal{P} is a subspace of \mathcal{D} .
- Fix $n \in \mathbb{N}$, and let \mathcal{P}_n be the subset of \mathcal{F} consisting of all polynomial functions of degree at most n (together with the zero function). Thus, for this fixed $n \in \mathbb{N}$, we consider those functions f such that there exists $a_0, a_1, \dots, a_{n-1}, a_n \in \mathbb{R}$ with $f(x) = a_n x^n + a_{n-1} x^{n-1} + \dots + a_1 x + a_0$ for all $x \in \mathbb{R}$. As above, it is straightforward to check that \mathcal{P}_n is a subspace of \mathcal{F} . In fact, each \mathcal{P}_n is a subspace of \mathcal{P} .

For example, the function $f: \mathbb{R} \rightarrow \mathbb{R}$ given by $f(x) = x^2 + 3x - 7$ is an element of \mathcal{P}_2 , as is the function $g: \mathbb{R} \rightarrow \mathbb{R}$ given by $g(x) = 5x - 2$ (because we can write it as $g(x) = 0x^2 + 5x - 2$).

We now expand our definition of linear combinations and spans to general vector spaces and an arbitrary finite number of vectors.

Definition 4.1.14. Let V be a vector space and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{v} \in V$. We say that \vec{v} is a linear combination of $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$ if there exists $c_1, c_2, \dots, c_n \in \mathbb{R}$ with $\vec{v} = c_1 \vec{u}_1 + c_2 \vec{u}_2 + \dots + c_n \vec{u}_n$.

Definition 4.1.15. Let V be a vector space, and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in V$. We define

$$\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = \{c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n : c_1, c_2, \dots, c_n \in \mathbb{R}\}.$$

In other words, $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is the set of all linear combinations of $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$.

One reason why we like taking the span of a finite collection of vectors in a vector space is that the resulting set is always a *subspace* of the vector space (not just a subset), as we now show.

Proposition 4.1.16. Let V be a vector space and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in V$. The set $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is a subspace of V .

Proof. We check the three properties.

- Notice that

$$\begin{aligned} 0 \cdot \vec{u}_1 + 0 \cdot \vec{u}_2 + \dots + 0 \cdot \vec{u}_n &= \vec{0} + \vec{0} + \dots + \vec{0} && \text{(by Proposition 4.1.11)} \\ &= \vec{0}, \end{aligned}$$

so $\vec{0} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$.

- Let $\vec{w}_1, \vec{w}_2 \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ be arbitrary. By definition, we can fix $c_1, c_2, \dots, c_n \in \mathbb{R}$ with

$$\vec{w}_1 = c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n,$$

and we can fix $d_1, d_2, \dots, d_n \in \mathbb{R}$ with

$$\vec{w}_2 = d_1\vec{u}_1 + d_2\vec{u}_2 + \dots + d_n\vec{u}_n.$$

Using properties 3, 4, and 8 of vector spaces, we then have

$$\begin{aligned} \vec{w}_1 + \vec{w}_2 &= (c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n) + (d_1\vec{u}_1 + d_2\vec{u}_2 + \dots + d_n\vec{u}_n) \\ &= (c_1\vec{u}_1 + d_1\vec{u}_1) + (c_2\vec{u}_2 + d_2\vec{u}_2) + \dots + (c_n\vec{u}_n + d_n\vec{u}_n) \\ &= (c_1 + d_1)\vec{u}_1 + (c_2 + d_2)\vec{u}_2 + \dots + (c_n + d_n)\vec{u}_n. \end{aligned}$$

Since $c_i + d_i \in \mathbb{R}$ for all i , we conclude that $\vec{w}_1 + \vec{w}_2 \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$.

- Let $\vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ and $r \in \mathbb{R}$ be arbitrary. By definition, we can fix $c_1, c_2, \dots, c_n \in \mathbb{R}$ with

$$\vec{w} = c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n.$$

Using properties 7 and 9 of vector spaces, we have

$$\begin{aligned} r \cdot \vec{w} &= r \cdot (c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n) \\ &= r \cdot (c_1\vec{u}_1) + r \cdot (c_2\vec{u}_2) + \dots + r \cdot (c_n\vec{u}_n) \\ &= (rc_1) \cdot \vec{u}_1 + (rc_2) \cdot \vec{u}_2 + \dots + (rc_n) \cdot \vec{u}_n. \end{aligned}$$

Since $rc_i \in \mathbb{R}$ for all i , we conclude that $r \cdot \vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$.

We have shown that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ contains the zero vector, is closed under addition, and is closed under scalar multiplication. Therefore, $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is a subspace of V . \square

We have now defined the span of any finite sequence of elements from a vector space. Well, *almost* any finite sequence. What about the empty sequence $()$, i.e. the sequence of length 0 that has no elements? How should we define $\text{Span}()$? Although this question might seem analogous to the question “What is the sound of one hand clapping?”, it turns out to be useful to assign a meaning here. Perhaps the first idea would be to define $\text{Span}()$ to be the empty set $\emptyset = \{\}$. But if we went there, then Proposition 4.1.16 would be false for $\text{Span}()$, because \emptyset is not a subspace of V . So if we want Proposition 4.1.16 to be true for the span of the empty sequence, we must define it to be a subspace. Given a vector space V , the smallest subspace of V is $\{\vec{0}\}$, so we choose to adopt the following convention.

Definition 4.1.17. *Given a vector space V , we define the span of the empty sequence to be $\text{Span}() = \{\vec{0}\}$.*

We will eventually see other good reasons for adopting this convention. For now, here is some additional motivation. Suppose that we are working in \mathbb{Z} , and want to define the sum of a finite sequence of numbers. For example, we want to assign the value 21 to the sequence $(9, 2, 7, 3)$ because $9 + 2 + 7 + 3 = 21$. Think about writing a computer program that takes as input a finite sequence of integers and outputs the sum of the elements of that sequence. There are two natural ways to write such a program:

- We might write the program by looping through the elements of the sequence and adding value to an accumulator variable. If we go this route, then we should probably start by setting the accumulator variable to 0 at the beginning, and then add each element in turn to the current accumulator. Once the list has been traversed, we return the value in the accumulator variable.
- Alternatively, we could write a recursive program. Given a list $(a_1, a_2, \dots, a_{n-1}, a_n)$, we should recursively get the sum of $(a_1, a_2, \dots, a_{n-1})$, and then add a_n to the result. If we write such a recursive program, we need a base case. In this situation, the natural base case is the empty sequence $()$. And to make the recursive process work, we should output the neutral element 0 when given the input $()$, because 0 is the number that leaves everything else alone under $+$.

In either case, notice that the program produces 0 on the sequence $()$. That is, the most natural way to define the sum of the elements of the empty sequence $()$ is to assign it the number 0.

Now let's go back to working in a vector space V and thinking about $\text{Span}()$. Instead of working with $+$ on \mathbb{Z} , we are now working with $+$ on V . As above, we can naturally define the sum of any finite nonempty sequence of vectors. And analogous to how we define the integer sum of $()$ to be 0, we should define the the vector sum of $()$ to be $\vec{0}$, since $\vec{0}$ leaves everything else alone under $+$. Now $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is the set of all sums of scalar multiples of the \vec{u}_i . So $\text{Span}()$ should be the set of all sums we get by taking scalar multiples of the elements of $()$. This results in just the empty sum, which is $\vec{0}$. Therefore, we should define $\text{Span}()$ to be $\{\vec{0}\}$.

With that convention taken care of, let's return to a discussion of the value of Proposition 4.1.16. Notice that this result generalizes Proposition 2.3.2 and Proposition 2.3.6 by extending them to arbitrary vector spaces and arbitrary finite collections of vectors. In fact, using this proposition, we can easily see that

$$\left\{ \begin{pmatrix} a \\ 2a \end{pmatrix} : a \in \mathbb{R} \right\}$$

is a subspace of \mathbb{R}^2 , which we argued directly above, by simply noticing that it equals

$$\text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix} \right).$$

Using this proposition, we can immediately form subspaces of some vector spaces. For example, we have that

$$W = \text{Span} \left(\begin{pmatrix} 1 \\ 2 \\ -5 \end{pmatrix}, \begin{pmatrix} -2 \\ -3 \\ 8 \end{pmatrix} \right)$$

is a subspace of \mathbb{R}^3 . Geometrically, by taking the span of two vectors in \mathbb{R}^3 that are not scalar multiples of each other, we are sweeping out a plane in \mathbb{R}^3 through the origin. Is

$$\begin{pmatrix} -3 \\ -2 \\ 7 \end{pmatrix} \in W?$$

To answer this question, we want to know whether there exists $c_1, c_2 \in \mathbb{R}$ with

$$c_1 \cdot \begin{pmatrix} 1 \\ 2 \\ -5 \end{pmatrix} + c_2 \cdot \begin{pmatrix} -2 \\ -3 \\ 8 \end{pmatrix} = \begin{pmatrix} -3 \\ -2 \\ 7 \end{pmatrix}.$$

In other words, we want to know if the system of equations

$$\begin{array}{rclcl} x & - & 2y & = & -3 \\ 2x & - & 3y & = & -2 \\ -5x & + & 8y & = & 7 \end{array}$$

has a solution. Subtracting twice the first equation from the second, we conclude that any solution must satisfy $y = 4$. Plugging this into the first equation tells us that any solution must satisfy $x - 8 = -3$, and hence $x = 5$. It follows that $(5, 4)$ is the only possible solution, and by plugging this in we can indeed verify that it satisfies all three equations. Therefore,

$$\begin{pmatrix} -3 \\ -2 \\ 7 \end{pmatrix} \in W.$$

However, notice that

$$\begin{pmatrix} -3 \\ -2 \\ 8 \end{pmatrix} \notin W$$

because again the only possible solution to the new system obtained by replacing 7 by 8 is $(5, 4)$ (because the logic above only used the first two equations), but $(5, 4)$ would not be a solution to the new last equation. Generalizing this idea, if we are working in \mathbb{R}^n , then determining if a given vector is in a span of m vectors, we will have to determine if a certain system of n equations in m variables has a solution. We will develop general techniques for this problem in the next section.

Although not obvious, it turns out that systems of equations can naturally arise even when working in vector spaces other than \mathbb{R}^n . For example, let \mathcal{D} be the vector space of all differentiable functions, let $f_1: \mathbb{R} \rightarrow \mathbb{R}$ be the function $f_1(x) = e^x$, and let $f_2: \mathbb{R} \rightarrow \mathbb{R}$ be the function $f_2(x) = e^{-x}$. Notice that $f_1, f_2 \in \mathcal{D}$ and that

$$W = \text{Span}(f_1, f_2)$$

is a subspace of \mathcal{D} by Proposition 4.1.16. Let $g: \mathbb{R} \rightarrow \mathbb{R}$ be the function $g(x) = x^2$. Is $g \in W$? We want to know whether there exists $c_1, c_2 \in \mathbb{R}$ with $g = c_1 f_1 + c_2 f_2$. In other words, we want to know if the following statement is true:

$$\text{“There exists } c_1, c_2 \in \mathbb{R} \text{ such that for all } x \in \mathbb{R}, \text{ we have } g(x) = c_1 f_1(x) + c_2 f_2(x)\text{”}.$$

In other words, we want to know whether the following statement is true:

$$\text{“There exists } c_1, c_2 \in \mathbb{R} \text{ such that for all } x \in \mathbb{R}, \text{ we have } x^2 = c_1 e^x + c_2 e^{-x}\text{”}.$$

We show this is false by giving a proof by contradiction. Suppose instead that the statement is true, and fix $c_1, c_2 \in \mathbb{R}$ such that $x^2 = c_1 e^x + c_2 e^{-x}$ for all $x \in \mathbb{R}$. The key idea is to avoid thinking about all x , and plug in particular values of x to obtain information about c_1 and c_2 (since, after all, we are assuming that this holds *for all* $x \in \mathbb{R}$). Plugging in $x = 0$, $x = 1$, and $x = -1$, we then must have

$$\begin{aligned} 0 &= c_1 + c_2 \\ 1 &= c_1 \cdot e + c_2 \cdot \frac{1}{e} \\ 1 &= c_1 \cdot \frac{1}{e} + c_2 \cdot e. \end{aligned}$$

Taking the second of these equations and subtracting e times the first, we conclude that c_1 and c_2 must satisfy

$$1 = c_2 \cdot \left(\frac{1}{e} - e \right).$$

Similarly, taking the third of these equations and subtracting $\frac{1}{e}$ times the first, we conclude that c_1 and c_2 must satisfy

$$1 = c_2 \cdot \left(e - \frac{1}{e} \right).$$

Finally, adding these latter two equations, we conclude that we must have $2 = 0$. This is a contradiction, so our original assumption that such c_1 and c_2 exist must be false. It follows that $g \notin \text{Span}(f_1, f_2)$, i.e. that $g \notin W$.

Notice that there was no guiding reason why we plugged in $0, 1, -1$ into the above equations. We could certainly have plugged in other values. However, the idea was to try to plug in enough values to obtain a system of equations that does not have a solution. If we instead noticed that there was a solution to these three equations, then we could try plugging in more points. If we still had a solution for more and more points, we might expect that g was really an element of $\text{Span}(f_1, f_2)$, and we might end up with a candidate choice for c_1 and c_2 . At this point, we would still need to show that $x^2 = c_1 e^x + c_2 e^{-x}$ *for all* $x \in \mathbb{R}$, which cannot be accomplished by merely plugging in values. However, with a candidate choice of c_1 and c_2 in hand, we can try to give an algebraic argument that works for all $x \in \mathbb{R}$.

As we are seeing, solving systems of equations seems to come up in many of these types of problems. We now turn to a general method to solve these systems.

4.2 Solving Linear Systems

Consider the following system of equations:

$$\begin{aligned} x + 2y + z &= 3 \\ 3x - y - 3z &= -1 \\ 2x + 3y + z &= 4. \end{aligned}$$

Suppose that we want to determine if this system has a solution, i.e. whether there exists an $(x, y, z) \in \mathbb{R}^3$ that makes all three equations simultaneously true. Geometrically, each equation represents a plane in \mathbb{R}^3 , so asking if the system has a solution is the same as asking if the three planes intersect. Instead of looking at each equation as describing a plane, we can write the system by looking vertically:

$$x \cdot \begin{pmatrix} 1 \\ 3 \\ 2 \end{pmatrix} + y \cdot \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix} + z \cdot \begin{pmatrix} 1 \\ -3 \\ 1 \end{pmatrix} = \begin{pmatrix} 3 \\ -1 \\ 4 \end{pmatrix}.$$

In other words, asking if this system has a solution is the same as asking whether

$$\begin{pmatrix} 3 \\ -1 \\ 4 \end{pmatrix} \in \text{Span} \left(\begin{pmatrix} 1 \\ 3 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix}, \begin{pmatrix} 1 \\ -3 \\ 1 \end{pmatrix} \right).$$

Therefore, we can interpret our one system of such equations in many ways.

How then could we determine if the system has a solution? The idea is to derive new equations from these three that any solution must also satisfy. Ideally, we seek to find “simpler” equations, such as ones that do not have all three variables simultaneously. Although there are many different possibilities for how to eliminate a variable, let’s consider starting at the top left. The coefficient of x in the first equation is 1, and we can use this value to eliminate the x in the second equation. If we take the second equation and subtract off 3 times the first, then we arrive at the equation

$$-7y - 6z = -10,$$

which is really shorthand for

$$0x - 7y - 6z = -10.$$

What does this operation mean in terms of solutions? If we have a solution (x, y, z) of the original system, then our (x, y, z) must also satisfy this equation. We can continue on in this way to derive other new equations that our (x, y, z) must satisfy, but in the process we will likely build many new equations. If instead of 3 equations in 3 variables, we had 8 equations in 7 variables, then the number of equations would balloon, and it would be difficult to keep track of them all.

Is it possible to *ignore* some of our old equations now that we have a new one? It’s certainly not true that we can ignore them all, because a solution to this one little equation might not be a solution to the original three. For example, $(0, 0, \frac{5}{3})$ is a solution to our new equation, but is certainly not a solution to the system. However, is it possible to *replace* one of our original equations with this new one, so that we can prevent this explosive growth of equations? Since we adjusted the second equation, let’s think about replacing it with the new one. Consider then the two systems

$$\begin{array}{rrcr} x & + & 2y & + & z & = & 3 \\ 3x & - & y & - & 3z & = & -1 \\ 2x & + & 3y & + & z & = & 4 \end{array}$$

and

$$\begin{array}{rrcr} x & + & 2y & + & z & = & 3 \\ & & -7y & - & 6z & = & -10 \\ 2x & + & 3y & + & z & = & 4. \end{array}$$

Remember that the logic goes forward, so any solution of the first system will indeed be a solution to the second (for the first and third equations this is trivial, and for the second this follows from the fact that it was derived from the original system by subtracting 3 times the first equation from the second). However, can we reverse the logic to argue that any solution to the second system will also be a solution to the first system? Since the first and third equations are the same, the fundamental question is whether we can derive the old second equation from the new system. Since the process of going from the old system to the new system was that we subtracted 3 times the first equation from the old second, we can *go backwards* by adding three times the first equation to the new second in order to recover the old second. Check that it works in this instance! This logic shows that any solution to the new system will also be a solution to the old system. In other words, the solution sets of the two systems are the same, and so we have *transformed* our original system into a slightly easier new system.

We want to generalize this situation and develop some simple rules that we can apply to similar situations without affecting the solution set. We first define the types of systems of equations that we will focus on.

Definition 4.2.1. A linear equation in the variables x_1, x_2, \dots, x_n is an equation of the form

$$a_1x_1 + a_2x_2 + \dots + a_nx_n = b,$$

where a_1, a_2, \dots, a_n, b are fixed real numbers.

For example

$$4x_1 - 9x_2 + 0x_3 + 7x_4 = 0$$

is a linear equation in x_1, x_2, x_3, x_4 . However, $x^2 + y - 7z = 12$ is not a linear equation in x, y, z because of the presence of x^2 . Also, $3x + xy + 2y = -4$ is not a linear equation in x, y, z because of the multiplication of x and y . In other words, a linear equation is about the most basic type of equation involving x_1, x_2, \dots, x_n because all that we can do is multiply by constants and add.

Definition 4.2.2. A system of linear equations, or a linear system, is a finite collection of linear equations in the same variables:

$$\begin{array}{ccccccc} a_{1,1}x_1 & + & a_{1,2}x_2 & + & \dots & + & a_{1,n}x_n & = & b_1 \\ a_{2,1}x_1 & + & a_{2,2}x_2 & + & \dots & + & a_{2,n}x_n & = & b_2 \\ \vdots & & \vdots & & & & \vdots & & \vdots \\ a_{m,1}x_1 & + & a_{m,2}x_2 & + & \dots & + & a_{m,n}x_n & = & b_m. \end{array}$$

A solution to this system of linear equations is an n -tuple of numbers $(s_1, s_2, \dots, s_n) \in \mathbb{R}^n$ which makes each of the above m equations true. The solution set of the system is the subset of \mathbb{R}^n consisting of all solutions to the system.

We now codify the basic moves that we will use to eliminate variables like we did in the example above.

Definition 4.2.3. An elementary row operation on a system of equations is one of following:

- Swapping the position of two equations in the list (swapping).
- Multiplying an equation by a nonzero constant (rescaling).
- Replacing one equation by the sum of itself and a multiple of another equation (row combination).

As we saw in our example, if we replaced an equation by a suitable combination, we were able to go backwards and recover the old system. We now argue that this is always possible when we apply an elementary row operation to a linear system.

Proposition 4.2.4 (Elementary Row Operations are Reversible). Suppose that we have a linear system, and we apply one elementary row operation to obtain a new system. We can then apply one elementary row operation to the new system in order to recover the old system.

Proof. Let A_1 be the original system, and let A_2 be the system that we obtain after applying one elementary row operation.

- If the elementary row operation was swapping row i and row j , then we can obtain A_1 from A_2 by swapping row i and row j .
- If the elementary row operation was multiplying row i by the nonzero number $c \in \mathbb{R}$, then we can obtain A_1 from A_2 by multiplying row i by the nonzero number $\frac{1}{c}$.
- If the elementary row operation was replacing row i with the sum of row i and c times row j (where $c \in \mathbb{R}$ and $j \neq i$), then we can obtain A_1 from A_2 by replacing row i with the sum of row i and $-c$ times row j .

Therefore, in all cases, we can obtain the old system from the new system using one elementary row operation. \square

Corollary 4.2.5. *Suppose that we have a linear system A_1 , and we apply one elementary row operation to obtain a new system A_2 . If S_1 is the solution set to the original system, and S_2 is the solution set to the new system, then $S_1 = S_2$.*

Proof. Since A_2 is derived from A_1 , any solution to A_1 will be a solution to A_2 , hence $S_1 \subseteq S_2$. Since elementary row operations are reversible, we can derive A_1 from A_2 , so any solution to A_2 will be a solution to A_1 , and hence $S_2 \subseteq S_1$. It follows that $S_1 = S_2$. \square

This corollary generalizes our earlier example where we argued that the system

$$\begin{array}{rrcr} x & + & 2y & + & z & = & 3 \\ 3x & - & y & - & 3z & = & -1 \\ 2x & + & 3y & + & z & = & 4 \end{array}$$

has the same solution set as the system

$$\begin{array}{rrcr} x & + & 2y & + & z & = & 3 \\ & & -7y & - & 6z & = & -10 \\ 2x & + & 3y & + & z & = & 4. \end{array}$$

Now we can continue to apply elementary row operations to further simplify this system without changing the solution set. The idea is to continue by eliminating the x in the third equation, and then moving on to eliminate a y as follows:

$$\begin{array}{l} \begin{array}{rrcr} x & + & 2y & + & z & = & 3 \\ 3x & - & y & - & 3z & = & -1 \\ 2x & + & 3y & + & z & = & 4 \end{array} & \rightarrow & \begin{array}{rrcr} x & + & 2y & + & z & = & 3 \\ & & -7y & - & 6z & = & -10 \\ 2x & + & 3y & + & z & = & 4 \end{array} & (-3R_1 + R_2) \\ \\ & \rightarrow & \begin{array}{rrcr} x & + & 2y & + & z & = & 3 \\ & & -7y & - & 6z & = & -10 \\ & & -y & - & z & = & -2 \end{array} & (-2R_1 + R_3) \\ \\ & \rightarrow & \begin{array}{rrcr} x & + & 2y & + & z & = & 3 \\ & & -y & - & z & = & -2 \\ & & -7y & - & 6z & = & -10 \end{array} & \begin{array}{l} (R_2 \leftrightarrow R_3) \\ (R_2 \leftrightarrow R_3) \end{array} \\ \\ & \rightarrow & \begin{array}{rrcr} x & + & 2y & + & z & = & 3 \\ & & -y & - & z & = & -2 \\ & & & & z & = & 4 \end{array} & (-7R_2 + R_3). \end{array}$$

Notice that at each step, we notate to the right of each system which elementary row operation we are applying. Furthermore, once we have finished using the first equation to eliminate all x 's below it, we moved over to y 's. We did not use the y in the first equation to eliminate below, because adding a multiple of the first equation to an equation below it would reintroduce an x there, and hence undo our progress. Instead, we want to use a y in an equation that no longer contains an x to help do the elimination. Although we could have kept the equation with $-7y$ in place and used that to eliminate the y below, it is easier to have a simple coefficient on a y to use for row combination operations. We also could have multiplied the $-7y$ equation by $\pm\frac{1}{7}$, but it seemed cleaner to swap the equation in order to avoid fractions.

In this case, the final system has a nice “staircase” structure and is much easier to understand. Moreover, by repeatedly applying Corollary 4.2.5, we know that the solution set of the original system is the same as the solution set of the final system. Looking at the final system, we realize that any solution (x, y, z) must have the property that $z = 4$. Back-substituting into the second equation gives $-y - z = -2$, so we conclude that any solution must have $-y - 4 = -2$, hence must have $y = -2$. Back-substituting these values of z and y into the first equation, we see that we must have $x - 4 + 4 = 3$, so $x = 3$. Thus, the only possible solution is $(3, -2, 4)$.

Is $(3, -2, 4)$ actually a solution of the system? We can indeed argue that it will be a solution (without just plugging it in) as follows. Notice that anything of the form $(*, *, 4)$ is a solution to the last equation. Following the back-substitution process, anything of the form $(*, -2, 4)$ is a solution to the second equation. By back-substitution again, anything of the form $(3, -2, 4)$ is a solution to the first equation. Therefore, $(3, -2, 4)$ is a solution to each equation, and hence the whole system. Therefore, the solution set of the final system is $\{(3, -2, 4)\}$, and hence the solution set of the original system is $\{(3, -2, 4)\}$.

Although this method works well, carrying around the variables is cumbersome and annoying. Notice that the variable names do not matter, and that they line up throughout. In other words, we can code the coefficients of the variables and the constants on the right by placing them into a matrix. For now, the matrices will simply serve as bookkeeping devices to keep everything lined up. Also, to speed up the process a bit more, we can do all of the elimination steps for one variable at once because the corresponding rows do not interact. Thus, we can write the above steps in the following more compact way:

$$\begin{aligned}
\begin{pmatrix} 1 & 2 & 1 & 3 \\ 3 & -1 & -3 & -1 \\ 2 & 3 & 1 & 4 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 2 & 1 & 3 \\ 0 & -7 & -6 & -10 \\ 0 & -1 & -1 & -2 \end{pmatrix} &\begin{matrix} (-3R_1 + R_2) \\ (-2R_1 + R_3) \end{matrix} \\
&\rightarrow \begin{pmatrix} 1 & 2 & 1 & 3 \\ 0 & -1 & -1 & -2 \\ 0 & -7 & -6 & -10 \end{pmatrix} &\begin{matrix} (R_2 \leftrightarrow R_3) \\ (R_2 \leftrightarrow R_3) \end{matrix} \\
&\rightarrow \begin{pmatrix} 1 & 2 & 1 & 3 \\ 0 & -1 & -1 & -2 \\ 0 & 0 & 1 & 4 \end{pmatrix} &(-7R_2 + R_3).
\end{aligned}$$

Now that we have reached this point, we can reinterpret the last matrix as corresponding to the system

$$\begin{array}{rclcl} x & + & 2y & + & z & = & 3 \\ & & -y & - & z & = & -2 \\ & & & & z & = & 4. \end{array}$$

and proceed to solve it as above.

Let's take a look at another example. Consider the following linear system:

$$\begin{array}{rcccccl} 2x & & & + & 8z & = & 6 \\ 7x & - & 3y & + & 18z & = & 15 \\ -3x & + & 3y & - & 2z & = & -1. \end{array}$$

Again, we have to interpret this situation properly. Our system actually is

$$\begin{array}{rrrrrr} x & + & 0y & + & 4z & = & 3 \\ 0x & - & 3y & - & 10z & = & -6 \\ 0x & + & 0y & + & 0z & = & 0. \end{array}$$

Using Corollary 4.2.5, we know that the solution set of our original system is the same as the solution set of this final system. No matter what $(x, y, z) \in \mathbb{R}^3$ we have, the final equation will always be true because the left-hand side will be 0 and the right-hand will be 0. Hence, we obtain no information at all from the last equation (i.e. it provides no restrictions), and hence our original system has the same solution set as the system

$$\begin{array}{rrrr} x & & + & 4z & = & 3 \\ & -3y & - & 10z & = & -6 \end{array}$$

that is obtained by omitting the last (trivial) equation. Now an ordered triple (x, y, z) of numbers satisfies this system if and only if

$$\begin{aligned} x &= 3 - 4z \\ -3y &= -6 + 10z, \end{aligned}$$

which is if and only if

$$\begin{aligned} x &= 3 - 4z \\ y &= 2 - \frac{10}{3}z. \end{aligned}$$

From this perspective, we can see that there are many solutions. For example, $(3, 2, 0)$ is a solution, as is $(-1, -\frac{4}{3}, 1)$ and $(-9, -8, 3)$. In general, we can let z be anything, say $z = t$, and then we obtain a unique solution with this value as follows:

$$\begin{aligned} x &= 3 - 4t \\ y &= 2 - \frac{10}{3} \cdot t \\ z &= t. \end{aligned}$$

This should remind you of a parametric equation of a line in 3-dimensions, and indeed it is the line through the point $(3, 2, 0)$ with direction vector $\langle -4, -\frac{10}{3}, 1 \rangle$. In other words, we have shown that the three planes

$$\begin{array}{rrrrrr} 2x & & + & 8z & = & 6 \\ 7x & - & 3y & + & 18z & = & 15 \\ -3x & + & 3y & - & 2z & = & -3 \end{array}$$

intersect in a line, and we have found a parametric description of the line using our new tools (no cross products or geometry needed!). We can also describe the solution set of the system using our parametric set notation as

$$\left\{ \begin{pmatrix} 3 \\ 2 \\ 0 \end{pmatrix} + t \cdot \begin{pmatrix} -4 \\ -10/3 \\ 1 \end{pmatrix} : t \in \mathbb{R} \right\}.$$

Notice that this set resembles the set

$$\text{Span} \left(\begin{pmatrix} -4 \\ -10/3 \\ 1 \end{pmatrix} \right) = \left\{ t \cdot \begin{pmatrix} -4 \\ -10/3 \\ 1 \end{pmatrix} : t \in \mathbb{R} \right\},$$

except that everything inside is being added to the “offset” vector

$$\begin{pmatrix} 3 \\ 2 \\ 0 \end{pmatrix}.$$

We now formalize some of the ideas we have developed above with some important definitions.

Definition 4.2.6. A linear system is consistent if it has at least one solution. Otherwise, it is inconsistent.

Definition 4.2.7. The augmented matrix of the linear system

$$\begin{array}{ccccccccc} a_{1,1}x_1 & + & a_{1,2}x_2 & + & \dots & + & a_{1,n}x_n & = & b_1 \\ a_{2,1}x_1 & + & a_{2,2}x_2 & + & \dots & + & a_{2,n}x_n & = & b_2 \\ \vdots & & \vdots & & & & \vdots & & \vdots \\ a_{m,1}x_1 & + & a_{m,2}x_2 & + & \dots & + & a_{m,n}x_n & = & b_m \end{array}$$

is the following $m \times (n+1)$ matrix:

$$\begin{pmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} & b_1 \\ a_{2,1} & a_{2,2} & \dots & a_{2,n} & b_2 \\ \vdots & \vdots & & \vdots & \vdots \\ a_{m,1} & a_{m,2} & \dots & a_{m,n} & b_m \end{pmatrix}.$$

Definition 4.2.8. A leading entry of a matrix is an entry that is the leftmost nonzero element in its row.

In our above examples, we used elementary row operations to arrive at a matrix with a “staircase”-like pattern. We now formalize this pattern with a definition.

Definition 4.2.9. A matrix is said to be in echelon form if the following two conditions are true:

- All zero rows appear below nonzero rows.
- In each nonzero row (aside from the first row), the leading entry is to the right of the leading entry in the row above it.

By applying elementary row operations, we can always arrive at a matrix in echelon form, as the following result states. The process of doing this is known as *Gaussian elimination*.

Definition 4.2.10. Let A and B be two $m \times n$ matrices.

- We say that A is row equivalent to B if there exists a finite sequence of elementary row operations that we can apply to A to obtain B .
- We say that B is an echelon form of A if A is row equivalent to B , and B is in echelon form.

Since row operations are reversible, notice that if A is row equivalent to B , then B is row equivalent to A .

Proposition 4.2.11. For any matrix A , there exists an echelon form of A .

Proof. If A is the zero matrix, then it is trivially in echelon form. Suppose then that A has at least one nonzero entry. Look at the first column that has a nonzero entry, and say that it is column j . By swapping rows, if necessary, we can obtain a matrix with a nonzero entry in the $(1, j)$ position. Then, by scaling the first row (if necessary), we can obtain a matrix with a 1 in the $(1, j)$ position. Now we can use row

combination on each of the rows below, and subtract off a suitable multiple of the first row from each so that every other entry in the column j is a 0. We will then have a matrix where columns $1, 2, \dots, j-1$ are all zero columns, while column j will have a 1 in the $(1, j)$ entry and 0's elsewhere. At this point, we ignore the first row and the first j columns, and continue this process on the rest of the matrix recursively. Eventually, we will arrive at a matrix in echelon form. \square

Notice that the above argument produces a matrix with an additional property beyond the requirements of echelon form: every leading entry of the final matrix is a 1. By performing these scalings, it can make it easier to understand how to eliminate using row combination. However, such scalings occasionally introduce unsightly fractions. Thus, when performing computations by hand, we will sometimes skip this step if we see other ways to eliminate variables.

For example, let's solve the following system of linear equations:

$$\begin{array}{rrrrrr} & & 2z & + & 10w & = & -6 \\ 3x & & & + & 9z & - & w & = & 5 \\ -x & + & 2y & - & 4z & + & w & = & 0. \end{array}$$

We first write down the augmented matrix of the linear system:

$$\left(\begin{array}{ccccc} 0 & 0 & 2 & 10 & -6 \\ 3 & 0 & 9 & -1 & 5 \\ -1 & 2 & -4 & 1 & 0 \end{array} \right).$$

We now apply elementary row operations until we obtain a matrix in echelon form:

$$\begin{aligned} \left(\begin{array}{ccccc} 0 & 0 & 2 & 10 & -6 \\ 3 & 0 & 9 & -1 & 5 \\ -1 & 2 & -4 & 1 & 0 \end{array} \right) &\rightarrow \left(\begin{array}{ccccc} -1 & 2 & -4 & 1 & 0 \\ 3 & 0 & 9 & -1 & 5 \\ 0 & 0 & 2 & 10 & -6 \end{array} \right) & \begin{array}{l} (R_1 \leftrightarrow R_3) \\ (R_1 \leftrightarrow R_3) \end{array} \\ &\rightarrow \left(\begin{array}{ccccc} -1 & 2 & -4 & 1 & 0 \\ 0 & 6 & -3 & 2 & 5 \\ 0 & 0 & 2 & 10 & -6 \end{array} \right) & (3R_1 + R_2). \end{aligned}$$

This final matrix here is in echelon form, and is the augmented matrix of the following linear system:

$$\begin{array}{rrrrrr} -x & + & 2y & - & 4z & + & w & = & 0 \\ & & 6y & - & 3z & + & 2w & = & 5 \\ & & & & 2z & + & 10w & = & -6. \end{array}$$

At this point, we think about the system as follows. Using Corollary 4.2.5, we know that our original system and this new system have the same solution set. Looking at the first equation, if we know the y, z, w values of a potential solution, then we can solve uniquely for x in a way that satisfies the first equation. Similarly, for the second equation, if we know z and w , then we can solve uniquely for y . Finally, for the third equation, if we know w , then we can solve uniquely for z . However, we don't have any obvious restriction on w . This is our cue for the fact that we can assign a value to w *arbitrarily*, and then work backwards through the equations to fill out the solutions uniquely for z, y, x in turn. As alluded to above, this method is called *back-substitution*. We begin by assigning w an arbitrary value that will serve as a parameter. Let's set $w = t$ for some arbitrary number t . The last equation now says

$$2z = -6 - 10w,$$

which has the same solution set as

$$z = -3 - 5w.$$

Thus, if $w = t$, then letting

$$z = -3 - 5t,$$

we see that (z, w) will satisfy the last equation. With this in mind, the second equation says

$$\begin{aligned} 6y &= 5 + 3z - 2w \\ &= 5 + 3 \cdot (-3 - 5t) - 2t \\ &= 5 - 9 - 15t - 2t \\ &= -4 - 17t, \end{aligned}$$

hence if we set

$$y = -\frac{2}{3} - \frac{17}{6}t$$

then (y, z, w) will satisfy the last two equations. Finally, the first equation gives us

$$\begin{aligned} x &= 2y - 4z + w \\ &= 2 \cdot \left(-\frac{2}{3} - \frac{17}{6}t \right) - 4 \cdot (-3 - 5t) + t \\ &= -\frac{4}{3} - \frac{17}{3}t + 12 + 20t + t \\ &= \frac{32}{3} + \frac{46}{3}t. \end{aligned}$$

Thus, the solution set is

$$\left\{ \begin{pmatrix} \frac{32}{3} + \frac{46}{3}t \\ -\frac{2}{3} - \frac{17}{6}t \\ -3 - 5t \\ t \end{pmatrix} : t \in \mathbb{R} \right\},$$

or equivalently

$$\left\{ \begin{pmatrix} 32/3 \\ -2/3 \\ -3 \\ 0 \end{pmatrix} + t \begin{pmatrix} 46/3 \\ -17/6 \\ -5 \\ 1 \end{pmatrix} : t \in \mathbb{R} \right\}.$$

Notice again that the solution set is the span of one vector, with an added offset. Although we lose some of the geometric intuition, this is what a 1-dimensional line in \mathbb{R}^4 should look like.

In general, we have the following result.

Proposition 4.2.12. *Suppose that we have a linear system with augmented matrix A , and that B is an echelon form of A .*

1. *If the last column of B contains a leading entry, then the system is inconsistent.*
2. *If the last column of B contains no leading entry, but every other column of B has a leading entry, then the system is consistent and has a unique solution.*
3. *If the last column of B contains no leading entry, and there is at least one other column of B without a leading entry, then the system is consistent and has infinitely many solutions. Moreover, for each choice of values for the variables that do not correspond to leading entries, there is a unique solution for the system taking these values.*

- Proof.* 1. Suppose that the last column of B contains a leading entry, say $c \neq 0$. We then have that B is the augmented matrix of a system including the equation $0 = c$. Since this equation has no solution, we can use Corollary 4.2.5 to conclude that our original system has no solutions, and hence is inconsistent.
2. Suppose that the last column of B does not contain a leading entry, but every other column does. We can work backwards through the equations in order to solve uniquely for each variable. In this way, we build a unique solution to the system corresponding to B , and hence to our original system by Corollary 4.2.5. It follows that our system is consistent and has a unique solution.
3. Suppose that the last column of B does not contain a leading entry, and at least one other column B does not have a leading entry. Now for each column of B that does not have a leading entry, we introduce a distinct parameter for the corresponding variable. We can then work backwards through the equation to solve uniquely for each variable that corresponds to a leading entry in terms of these parameters. In this way, we build unique solutions for each choice of parameter for the system corresponding to B , and hence (by Corollary 4.2.5) to our original system. It follows that our system is consistent and has infinitely many solutions (one for each choice of each of the parameters). \square

For another example, suppose that we want to solve the following system:

$$\begin{array}{rrrrrrr} x_1 & + & 5x_2 & - & x_3 & - & 2x_4 & - & 3x_5 & = & 16 \\ 3x_1 & + & 15x_2 & - & 2x_3 & - & 4x_4 & - & 2x_5 & = & 56 \\ -2x_1 & - & 10x_2 & & & + & x_4 & - & 10x_5 & = & -46 \\ 4x_1 & + & 20x_2 & - & x_3 & - & 3x_4 & + & 11x_5 & = & 86. \end{array}$$

We first write down the augmented matrix of the linear system:

$$\left(\begin{array}{cccccc} 1 & 5 & -1 & -2 & -3 & 16 \\ 3 & 15 & -2 & -4 & -2 & 56 \\ -2 & -10 & 0 & 1 & -10 & -46 \\ 4 & 20 & -1 & -3 & 11 & 86 \end{array} \right).$$

We now apply elementary row operations until we obtain a matrix in echelon form:

$$\begin{aligned} \left(\begin{array}{cccccc} 1 & 5 & -1 & -2 & -3 & 16 \\ 3 & 15 & -2 & -4 & -2 & 56 \\ -2 & -10 & 0 & 1 & -10 & -46 \\ 4 & 20 & -1 & -3 & 11 & 86 \end{array} \right) &\rightarrow \left(\begin{array}{cccccc} 1 & 5 & -1 & -2 & -3 & 16 \\ 0 & 0 & 1 & 2 & 7 & 8 \\ 0 & 0 & -2 & -3 & -16 & -14 \\ 0 & 0 & 3 & 5 & 23 & 22 \end{array} \right) &\begin{array}{l} (-3R_1 + R_2) \\ (2R_1 + R_3) \\ (-4R_1 + R_4) \end{array} \\ &\rightarrow \left(\begin{array}{cccccc} 1 & 5 & -1 & -2 & -3 & 16 \\ 0 & 0 & 1 & 2 & 7 & 8 \\ 0 & 0 & 0 & 1 & -2 & 2 \\ 0 & 0 & 0 & -1 & 2 & -2 \end{array} \right) &\begin{array}{l} (2R_2 + R_3) \\ (-3R_2 + R_4) \end{array} \\ &\rightarrow \left(\begin{array}{cccccc} 1 & 5 & -1 & -2 & -3 & 16 \\ 0 & 0 & 1 & 2 & 7 & 8 \\ 0 & 0 & 0 & 1 & -2 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right) &(R_3 + R_4). \end{aligned}$$

Notice that the last column does not have a leading entry, that three other columns do have leading entries, and that two other columns do not have leading entries. Using Proposition 4.2.12, we introduce

parameters for the variables corresponding to columns without leading entries, say by letting $x_2 = s$ and $x_5 = t$. We then back-substitute to conclude that

$$\begin{aligned}x_4 &= 2 + 2x_5 \\ &= 2 + 2t,\end{aligned}$$

and then

$$\begin{aligned}x_3 &= 8 - 2x_4 - 7x_5 \\ &= 8 - 2(2 + 2t) - 7t \\ &= 4 - 11t,\end{aligned}$$

and finally

$$\begin{aligned}x_1 &= 16 - 5x_2 + x_3 + 2x_4 + 3x_5 \\ &= 16 - 5s + (4 - 11t) + 2(2 + 2t) + 3t \\ &= 24 - 5s - 4t.\end{aligned}$$

Thus, the solution set can be described as

$$\left\{ \begin{pmatrix} 24 - 5s - 4t \\ s \\ 4 - 11t \\ 2 + 2t \\ t \end{pmatrix} : s, t \in \mathbb{R} \right\}$$

which we can equivalently write as:

$$\left\{ \begin{pmatrix} 24 \\ 0 \\ 4 \\ 2 \\ 0 \end{pmatrix} + s \begin{pmatrix} -5 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} -4 \\ 0 \\ -11 \\ 2 \\ 1 \end{pmatrix} : s, t \in \mathbb{R} \right\}.$$

Notice in this case that the solution set can be written as the span of 2 vectors, plus an offset. Geometrically, this feels like a 2-dimensional “plane” in \mathbb{R}^5 , where the offset moves the plane off of the origin. We will discuss these geometric ideas in more detail soon.

As we saw in the previous example, the algebra involved in back-substitution can be a bit messy. To avoid this, we can continue to perform elementary row operations after reaching echelon form. The idea is to use leading entries to also eliminate values *above*, rather than just below. To see this in action, we start with the above matrix in echelon form and continue:

$$\begin{aligned} \begin{pmatrix} 1 & 5 & -1 & -2 & -3 & 16 \\ 0 & 0 & 1 & 2 & 7 & 8 \\ 0 & 0 & 0 & 1 & -2 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 5 & -1 & 0 & -7 & 20 \\ 0 & 0 & 1 & 0 & 11 & 4 \\ 0 & 0 & 0 & 1 & -2 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} &\begin{array}{l} (2R_3 + R_1) \\ (-2R_3 + R_2) \end{array} \\ &\rightarrow \begin{pmatrix} 1 & 5 & 0 & 0 & 4 & 24 \\ 0 & 0 & 1 & 0 & 11 & 4 \\ 0 & 0 & 0 & 1 & -2 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} &(R_2 + R_1). \end{aligned}$$

From here, we still introduce parameters $x_2 = s$ and $x_5 = t$ and back-substitute to conclude that

$$\begin{aligned}x_4 &= 2 + 2x_5 \\ &= 2 + 2t.\end{aligned}$$

However, now notice that we immediately have

$$\begin{aligned}x_3 &= 4 - 11x_5 \\ &= 4 - 11t,\end{aligned}$$

and

$$\begin{aligned}x_1 &= 24 - 5x_2 - 4x_5 \\ &= 24 - 5s - 4t,\end{aligned}$$

avoiding some of the annoying algebra that arises in the back-substitution process. Matrices with these additional properties (that allow one to skip essentially all of the algebra when performing back-substitution) are given a special name.

Definition 4.2.13. *A matrix is said to be in reduced echelon form if the following conditions are true:*

- *It is in echelon form.*
- *The leading entry of every nonzero row is a 1.*
- *In each column containing a leading entry, all numbers aside from this entry are 0.*

One situation where proceeding to reduced echelon form is particularly valuable is when we want to solve several different linear systems that have the same coefficients on the variables. For example, suppose that we want to solve each of the following:

$$\begin{array}{rrcrcl}x & + & 2y & + & z & = & -4 \\ 3x & + & 5y & - & 2z & = & -5 \\ -2x & - & 3y & + & z & = & 7\end{array}\quad \text{and} \quad \begin{array}{rrcrcl}x & + & 2y & + & z & = & 1 \\ 3x & + & 5y & - & 2z & = & -9 \\ -2x & - & 3y & + & z & = & 0.\end{array}$$

Notice that the only difference here is the constants on the right. We could of course form the augmented matrix of the first system and go down to echelon form as follows:

$$\begin{aligned}\begin{pmatrix} 1 & 2 & 1 & -4 \\ 3 & 5 & -2 & -5 \\ -2 & -3 & 1 & 7 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 2 & 1 & -4 \\ 0 & -1 & -5 & 7 \\ 0 & 1 & 3 & -1 \end{pmatrix} && \begin{array}{l} (-3R_1 + R_2) \\ (2R_1 + R_3) \end{array} \\ &\rightarrow \begin{pmatrix} 1 & 2 & 1 & -4 \\ 0 & -1 & -5 & 7 \\ 0 & 0 & -2 & 6 \end{pmatrix} && (R_2 + R_3).\end{aligned}$$

We could now do back-substitution to solve the first system, and then move on to the second system. However, after a moment's thought, it becomes clear that we would perform the exact same row operations on the augmented matrix of the second system in order to get to echelon form. Instead of doing the same operations twice, why not add another column to the above matrix? In other words, let's form the following matrix:

$$\begin{pmatrix} 1 & 2 & 1 & -4 & 1 \\ 3 & 5 & -2 & -5 & -9 \\ -2 & -3 & 1 & 7 & 0 \end{pmatrix}.$$

We can now perform elementary row operations on this matrix, and interpret the final echelon form matrix accordingly. The key fact that allows this to work is that elementary row operations perform arithmetic on each column separately. That is, we never “cross the streams” of the columns when performing Gaussian Elimination, so there is no harm in treating several columns as “augmented columns”. If we proceed to echelon form on this matrix, we obtain:

$$\begin{aligned} \begin{pmatrix} 1 & 2 & 1 & -4 & 1 \\ 3 & 5 & -2 & -5 & -9 \\ -2 & -3 & 1 & 7 & 0 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 2 & 1 & -4 & 1 \\ 0 & -1 & -5 & 7 & -12 \\ 0 & 1 & 3 & -1 & 2 \end{pmatrix} &\begin{array}{l} (-3R_1 + R_2) \\ (2R_1 + R_3) \end{array} \\ &\rightarrow \begin{pmatrix} 1 & 2 & 1 & -4 & 1 \\ 0 & -1 & -5 & 7 & -12 \\ 0 & 0 & -2 & 6 & -10 \end{pmatrix} &(R_2 + R_3). \end{aligned}$$

Notice that if we just look at the first four columns, then we have captured the old computation on the first linear system. But if we instead ignore the fourth column (so focus on columns 1, 2, 3, and 5), then we have coded the way we would solve the second linear system. Of course, we could now solve each system separately using back-substitution. Instead, we continue the process of performing elementary row operations in order to get to a matrix in reduced echelon form:

$$\begin{aligned} \begin{pmatrix} 1 & 2 & 1 & -4 & 1 \\ 0 & -1 & -5 & 7 & -12 \\ 0 & 0 & -2 & 6 & -10 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 2 & 1 & -4 & 1 \\ 0 & 1 & 5 & -7 & 12 \\ 0 & 0 & 1 & -3 & 5 \end{pmatrix} &\begin{array}{l} ((-1) \cdot R_2) \\ ((-1/2) \cdot R_3). \end{array} \\ &\rightarrow \begin{pmatrix} 1 & 2 & 0 & -1 & -4 \\ 0 & 1 & 0 & 8 & -13 \\ 0 & 0 & 1 & -3 & 5 \end{pmatrix} &\begin{array}{l} (-R_3 + R_1) \\ (-5R_3 + R_2) \end{array} \\ &\rightarrow \begin{pmatrix} 1 & 0 & 0 & -17 & 22 \\ 0 & 1 & 0 & 8 & -13 \\ 0 & 0 & 1 & -3 & 5 \end{pmatrix} &(-2R_2 + R_1) \end{aligned}$$

Take some time to think about how this computation (and the previous one where we introduced reduced echelon form) was performed. The initial process of starting with a matrix and then going to echelon form focuses on column 1 first (by finding a nonzero entry there, flipping it to the top, and then using it to eliminate below), then moves on to column 2, etc. However, once we are in echelon form, we work with later columns and move backwards. In this backward process, we began by focusing on column 3, which is the rightmost column with a leading entry. We used the 1 in position (3, 3) to eliminate the nonzero entries *above* it. Since the entries to the left of this 1 were all 0 (because we were already in echelon form), these elementary row operations kept the matrix in echelon form. Now once we have made the (3, 3) entry the only nonzero element in its column, we moved backwards to another column with a leading entry, and used it to eliminate the nonzero entries above it.

Now that we are in reduced echelon form, we can immediately read off the solutions to the two linear systems. If we focus on the first four columns, then we see that the first linear system has the same solution set as the linear system

$$\begin{array}{rcl} x & = & -17 \\ y & = & 8 \\ z & = & -3. \end{array}$$

In other words, our first linear system has a unique solution, which is $(-17, 8, -3)$. Similarly, our second linear system has a unique solution, which is $(22, -13, 5)$. No back-substitution necessary!

Before moving on, we should pause to reflect on one important consequence of the last example. A matrix

of numbers is not just a code of one linear system. For example, consider the matrix:

$$\begin{pmatrix} 2 & 1 & -5 & 3 & 0 \\ 6 & -4 & 1 & 7 & 9 \end{pmatrix}$$

As given, this is just a grid of numbers. We might choose to think about it as the augmented matrix of the linear system

$$\begin{array}{rrrrrr} 2x & + & y & - & 5z & + & 3w & = & 0 \\ 6x & - & 4y & + & z & + & 7w & = & 9. \end{array}$$

We can also choose to imagine the last two columns as augmented columns, and so think about it as coding both of the following two systems simultaneously:

$$\begin{array}{rrrr} 2x & + & y & - & 5z & = & 3 \\ 6x & - & 4y & + & z & = & 7 \end{array} \quad \text{and} \quad \begin{array}{rrrr} 2x & + & y & - & 5z & = & 0 \\ 6x & - & 4y & + & z & = & 9. \end{array}$$

Or we could imagine the last three columns as augmented, and thus consider the matrix as coding three linear systems:

$$\begin{array}{rrrr} 2x & + & y & = & -5 \\ 6x & - & 4y & = & 1, \end{array} \quad \begin{array}{rrrr} 2x & + & y & = & 3 \\ 6x & - & 4y & = & 7, \end{array} \quad \text{and} \quad \begin{array}{rrrr} 2x & + & y & = & 0 \\ 6x & - & 4y & = & 9. \end{array}$$

Indeed, we can even think of the last four columns as augmented, and so consider our matrix as coding four (sad) linear systems in one variable. We will soon see many other ways to interpret a matrix that do not involve viewing it as coding a linear system (after all, in earlier chapters, we thought about a matrix as coding a linear transformation). For now, the key takeaway is that a matrix is just a grid of numbers, and you can choose to interpret it how you want.

Although the techniques that we have developed in this section are naturally suited to problems in \mathbb{R}^m , these linear systems arise when studying other vector spaces as well. For example, suppose that we are working in the vector space \mathcal{P}_2 of polynomials with degree at most 2. Let $f_1: \mathbb{R} \rightarrow \mathbb{R}$ be the function $f_1(x) = 5x^2 - 2x - 1$ and let $f_2: \mathbb{R} \rightarrow \mathbb{R}$ be the function $f_2(x) = -3x^2 + x - 4$. We then have that

$$\text{Span}(f_1, f_2) = \{c_1 f_1 + c_2 f_2 : c_1, c_2 \in \mathbb{R}\}$$

is the set of all linear combinations of f_1 and f_2 . Let $g: \mathbb{R} \rightarrow \mathbb{R}$ be the polynomial function given by $g(x) = -7x^2 + x - 40$. Notice that $g \in \mathcal{P}_2$. Is $g \in \text{Span}(f_1, f_2)$? In other words, can we obtain the parabola given by the graph of g by suitably scaling and adding the parabolas given by f_1 and f_2 ? We want to know whether there exists $c_1, c_2 \in \mathbb{R}$ with $g = c_1 f_1 + c_2 f_2$. Recall that two functions are equal if and only if they give the same output on every possible input. Thus, we want to know whether there exists $c_1, c_2 \in \mathbb{R}$ such that for all $x \in \mathbb{R}$, we have $g(x) = c_1 f_1(x) + c_2 f_2(x)$. In other words, does there exist $c_1, c_2 \in \mathbb{R}$ such that for all $x \in \mathbb{R}$, we have

$$-7x^2 + x - 40 = c_1 \cdot (5x^2 - 2x - 1) + c_2 \cdot (-3x^2 + x - 4).$$

On the face of it, this question looks challenging. However, we can expand the right hand side and then collect terms having the same powers of x . In other words, for all $c_1, c_2, x \in \mathbb{R}$, we have

$$c_1 \cdot (5x^2 - 2x - 1) + c_2 \cdot (-3x^2 + x - 4) = (5c_1 - 3c_2)x^2 + (-2c_1 + c_2)x + (-c_1 - 4c_2).$$

Thus, we want to know whether there exist $c_1, c_2 \in \mathbb{R}$ such that for all $x \in \mathbb{R}$, we have

$$-7x^2 + x - 40 = (5c_1 - 3c_2)x^2 + (-2c_1 + c_2)x + (-c_1 - 4c_2).$$

Now it is natural to rephrase this question by equating the coefficients on each side, but does this actually work? Certainly if we happen to find $c_1, c_2 \in \mathbb{R}$ with

$$\begin{array}{rclcl} 5c_1 & - & 3c_2 & = & -7 \\ -2c_1 & + & c_2 & = & 1 \\ -c_1 & - & 4c_2 & = & -40, \end{array}$$

then we will have

$$-7x^2 + x - 40 = (5c_1 - 3c_2)x^2 + (-2c_1 + c_2)x + (-c_1 - 4c_2)$$

for all $x \in \mathbb{R}$. However, is it possible that two polynomials of degree 2 behave the same on all inputs but actually look different? It turns out that the answer is no. We will build our argument based on the following fundamental fact saying that a polynomial of degree n has at most n roots. You likely believe it from earlier work, but we will defer a proof to Abstract Algebra.

Fact 4.2.14. *Let $a_0, a_1, \dots, a_{n-1}, a_n \in \mathbb{R}$ with $a_n \neq 0$. If $f: \mathbb{R} \rightarrow \mathbb{R}$ is the function $f(x) = a_n x^n + a_{n-1} x^{n-1} + \dots + a_1 x + a_0$, then there are at most n distinct $r \in \mathbb{R}$ with $f(r) = 0$.*

With this fact in hand, we can argue that if two polynomial functions have the same input/output behavior, then they must actually be the same polynomial.

Proposition 4.2.15. *Let $a_0, a_1, \dots, a_n \in \mathbb{R}$ and $b_0, b_1, \dots, b_n \in \mathbb{R}$. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ and $g: \mathbb{R} \rightarrow \mathbb{R}$ be the polynomial functions:*

$$\begin{aligned} f(x) &= a_n x^n + a_{n-1} x^{n-1} + \dots + a_1 x + a_0 \\ g(x) &= b_n x^n + a_{n-1} x^{n-1} + \dots + b_1 x + b_0. \end{aligned}$$

If $f(x) = g(x)$ for all $x \in \mathbb{R}$, then $a_i = b_i$ for all i .

Proof. Suppose that $f(x) = g(x)$ for all $x \in \mathbb{R}$. Define $h: \mathbb{R} \rightarrow \mathbb{R}$ by letting $h(x) = f(x) - g(x)$ for all $x \in \mathbb{R}$, so

$$h(x) = (a_n - b_n)x^n + (a_{n-1} - b_{n-1})x^{n-1} + \dots + (a_1 - b_1)x + (a_0 - b_0)$$

for all $x \in \mathbb{R}$. Since $f(x) = g(x)$ for all $x \in \mathbb{R}$, we have that $h(x) = 0$ for all $x \in \mathbb{R}$, so h has infinitely many roots. Since $h(x)$ is a polynomial, the only possibility is that every coefficient of h must equal 0. Thus, $a_i - b_i = 0$ for all i , so $a_i = b_i$ for all i . \square

We now return to our original question of whether $g \in \text{Span}(f_1, f_2)$. Using the result that two polynomial functions agree on all inputs if and only if the coefficients are equal, we simply want to know whether the system

$$\begin{array}{rclcl} 5c_1 & - & 3c_2 & = & -7 \\ -2c_1 & + & c_2 & = & 1 \\ -c_1 & - & 4c_2 & = & -40 \end{array}$$

has a solution. Applying elementary row operations, we obtain

$$\begin{aligned}
 \begin{pmatrix} 5 & -3 & -7 \\ -2 & 1 & 1 \\ -1 & -4 & -40 \end{pmatrix} &\rightarrow \begin{pmatrix} -1 & -4 & -40 \\ -2 & 1 & 1 \\ 5 & -3 & -7 \end{pmatrix} && (R_1 \leftrightarrow R_3) \\
 &\rightarrow \begin{pmatrix} -1 & -4 & -40 \\ 0 & 9 & 81 \\ 0 & -23 & -207 \end{pmatrix} && \begin{array}{l} (R_1 \leftrightarrow R_3) \\ (-2R_1 + R_2) \\ (5R_1 + R_3) \end{array} \\
 &\rightarrow \begin{pmatrix} -1 & -4 & -40 \\ 0 & 1 & 9 \\ 0 & 1 & 9 \end{pmatrix} && \begin{array}{l} (\frac{1}{9} \cdot R_2) \\ (\frac{-1}{23} \cdot R_3) \end{array} \\
 &\rightarrow \begin{pmatrix} -1 & -4 & -40 \\ 0 & 1 & 9 \\ 0 & 0 & 0 \end{pmatrix} && (-R_2 + R_3).
 \end{aligned}$$

Thus, there is a unique solution to this system, and it is $(4, 9)$. We can check this result by noticing that for all $x \in \mathbb{R}$, we have

$$\begin{aligned}
 4 \cdot f_1(x) + 9 \cdot f_2(x) &= 4 \cdot (5x^2 - 2x - 1) + 9 \cdot (-3x^2 + x - 4) \\
 &= 20x^2 - 8x - 4 + (-27x^2 + 9x - 36) \\
 &= -7x^2 + x - 40 \\
 &= g(x).
 \end{aligned}$$

In contrast, consider the function $h: \mathbb{R} \rightarrow \mathbb{R}$ given by $h(x) = -x^2 + 6$. We claim that $h \notin \text{Span}(f_1, f_2)$. We argue this as follows. Suppose that $h \in \text{Span}(f_1, f_2)$, and fix $c_1, c_2 \in \mathbb{R}$ with $h = c_1 f_1 + c_2 f_2$. We then have that $h(x) = c_1 f_1(x) + c_2 f_2(x)$ for all $x \in \mathbb{R}$, so

$$\begin{aligned}
 -x^2 + 6 &= c_1 \cdot f_1(x) + c_2 \cdot f_2(x) \\
 &= c_1 \cdot (5x^2 - 2x - 1) + c_2 \cdot (-3x^2 + x - 4) \\
 &= (5c_1 - 3c_2)x^2 + (-2c_1 + c_2)x + (-c_1 - 4c_2)
 \end{aligned}$$

for all $x \in \mathbb{R}$. It follows that c_1 and c_2 must satisfy the following system of equations:

$$\begin{array}{rclcl}
 5c_1 & - & 3c_2 & = & -1 \\
 -2c_1 & + & c_2 & = & 0 \\
 -c_1 & - & 4c_2 & = & 6.
 \end{array}$$

We showed at the beginning of Section 2.1 that this system has no solutions, so it follows that there does not exist $c_1, c_2 \in \mathbb{R}$ with $h = c_1 f_1 + c_2 f_2$, and hence $h \notin \text{Span}(f_1, f_2)$.

4.3 Spanning Sequences

Now that we have handled the basics of solving linear systems, we can move on to other places where linear systems arise. For a simple example of such a question, consider asking the following: Does

$$\text{Span} \left(\begin{pmatrix} 1 \\ 3 \\ -4 \end{pmatrix}, \begin{pmatrix} 3 \\ 1 \\ -8 \end{pmatrix}, \begin{pmatrix} -3 \\ 7 \\ 2 \end{pmatrix} \right) = \mathbb{R}^3?$$

In other words, we want to know if for all $b_1, b_2, b_3 \in \mathbb{R}$, there exists $c_1, c_2, c_3 \in \mathbb{R}$ with

$$c_1 \cdot \begin{pmatrix} 1 \\ 3 \\ -4 \end{pmatrix} + c_2 \cdot \begin{pmatrix} 3 \\ 1 \\ -8 \end{pmatrix} + c_3 \cdot \begin{pmatrix} -3 \\ 7 \\ 2 \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix}.$$

Now for each choice of $b_1, b_2, b_3 \in \mathbb{R}$, we are asking if a certain linear system is consistent. Thus, given arbitrary $b_1, b_2, b_3 \in \mathbb{R}$, we want to examine the augmented matrix

$$\begin{pmatrix} 1 & 3 & -3 & b_1 \\ 3 & 1 & 7 & b_2 \\ -4 & -8 & 2 & b_3 \end{pmatrix}.$$

Applying elementary row operations, we have

$$\begin{aligned} \begin{pmatrix} 1 & 3 & -3 & b_1 \\ 3 & 1 & 7 & b_2 \\ -4 & -8 & 2 & b_3 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 3 & -3 & b_1 \\ 0 & -8 & 16 & -3b_1 + b_2 \\ 0 & 4 & -10 & 4b_1 + b_3 \end{pmatrix} && \begin{matrix} (-3R_1 + R_2) \\ (4R_1 + R_3) \end{matrix} \\ &\rightarrow \begin{pmatrix} 1 & 3 & -3 & b_1 \\ 0 & 1 & -2 & \frac{3}{8}b_1 - \frac{1}{8}b_2 \\ 0 & 4 & -10 & 4b_1 + b_3 \end{pmatrix} && (-\frac{1}{8} \cdot R_2) \\ &\rightarrow \begin{pmatrix} 1 & 3 & -3 & b_1 \\ 0 & 1 & -2 & \frac{3}{8}b_1 - \frac{1}{8}b_2 \\ 0 & 0 & -2 & -\frac{3}{2}b_1 + \frac{1}{2}b_2 + b_3 \end{pmatrix} && (-4R_2 + R_3). \end{aligned}$$

Notice that no matter what values $b_1, b_2, b_3 \in \mathbb{R}$ we have, there is not a leading entry in the last column. Therefore, the corresponding linear system has a solution by Proposition 4.2.12 (in fact, it has a unique solution because every other column has a leading entry). It follows that

$$\text{Span} \left(\begin{pmatrix} 1 \\ 3 \\ -4 \end{pmatrix}, \begin{pmatrix} 3 \\ 1 \\ -8 \end{pmatrix}, \begin{pmatrix} -3 \\ 7 \\ 2 \end{pmatrix} \right) = \mathbb{R}^3$$

is true. Suppose that we instead ask the following question: Does

$$\text{Span} \left(\begin{pmatrix} 0 \\ 1 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ 3 \\ 7 \end{pmatrix}, \begin{pmatrix} 8 \\ 5 \\ 7 \end{pmatrix} \right) = \mathbb{R}^3?$$

As above, we take arbitrary $b_1, b_2, b_3 \in \mathbb{R}$, and apply elementary row operations:

$$\begin{aligned} \begin{pmatrix} 0 & 2 & 8 & b_1 \\ 1 & 3 & 5 & b_2 \\ 3 & 7 & 7 & b_3 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 3 & 5 & b_2 \\ 0 & 2 & 8 & b_1 \\ 3 & 7 & 7 & b_3 \end{pmatrix} && \begin{matrix} (R_1 \leftrightarrow R_2) \\ (R_1 \leftrightarrow R_2) \end{matrix} \\ &\rightarrow \begin{pmatrix} 1 & 3 & 5 & b_2 \\ 0 & 2 & 8 & b_1 \\ 0 & -2 & -8 & -3b_2 + b_3 \end{pmatrix} && (-3R_1 + R_3) \\ &\rightarrow \begin{pmatrix} 1 & 3 & 5 & b_2 \\ 0 & 1 & 4 & \frac{1}{2}b_1 \\ 0 & -2 & -8 & -3b_2 + b_3 \end{pmatrix} && (\frac{1}{2} \cdot R_2) \\ &\rightarrow \begin{pmatrix} 1 & 3 & 5 & b_2 \\ 0 & 1 & 4 & \frac{1}{2}b_1 \\ 0 & 0 & 0 & b_1 - 3b_2 + b_3 \end{pmatrix} && (2R_2 + R_3). \end{aligned}$$

Looking at this final matrix, we see an interesting phenomenon. If the numbers b_1, b_2, b_3 satisfy $b_1 - 3b_2 + b_3 \neq 0$, then the final column has a leading entry, and hence the system is inconsistent by Proposition 4.2.12. For example, if we look at the vector

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

and notice that $1 - 3 \cdot 0 + 0 = 1$ is nonzero, we conclude that

$$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \notin \text{Span} \left(\begin{pmatrix} 0 \\ 1 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ 3 \\ 7 \end{pmatrix}, \begin{pmatrix} 8 \\ 5 \\ 7 \end{pmatrix} \right).$$

Therefore, we certainly have

$$\text{Span} \left(\begin{pmatrix} 0 \\ 1 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ 3 \\ 7 \end{pmatrix}, \begin{pmatrix} 8 \\ 5 \\ 7 \end{pmatrix} \right) \neq \mathbb{R}^3.$$

However, if the numbers b_1, b_2, b_3 satisfy $b_1 - 3b_2 + b_3 = 0$, then the final column does not have a leading entry, and hence the system is consistent by Proposition 4.2.12. In fact, we know that the system has infinitely many solutions. In other words, we have shown that

$$\text{Span} \left(\begin{pmatrix} 0 \\ 1 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ 3 \\ 7 \end{pmatrix}, \begin{pmatrix} 8 \\ 5 \\ 7 \end{pmatrix} \right) = \left\{ \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} \in \mathbb{R}^3 : b_1 - 3b_2 + b_3 = 0 \right\}.$$

Looking at the set on the right, we see that it is the solution set to the equation $x - 3y + z = 0$. In other words, our three vectors span a plane in \mathbb{R}^3 through the origin, and we have found an equation for that plane.

With these two situations in hand, you may wonder whether we need to carry around that last column (b_1, b_2, b_3) after all. In the first example, it did not matter what appeared in the last column in the end. The key reason was that each row had a leading entry that was not in the last column, so we could solve backwards for the variables no matter what appeared there. In the second example, we did use the results of the last column to explicitly describe the set spanned by the three vectors, but we needed much less information to determine that the three vectors did not span \mathbb{R}^3 . In particular, all that we needed to know was that there was a choice of $b_1, b_2, b_3 \in \mathbb{R}$ so that $b_1 - 3b_2 + b_3$ would be nonzero, providing us with an example where there would be no solution. Could we do the above process without the last column, notice the zero row at the bottom, and then conclude that there must be a choice of last column at the beginning that produces a nonzero entry in the bottom right? It turns out that the answer is yes, and the fundamental reason is that elementary row operations are reversible. We isolate this result.

Proposition 4.3.1. *Let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$. Let A be the $m \times n$ matrix where the i^{th} column is \vec{u}_i , and let B be an echelon form of A . We then have that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = \mathbb{R}^m$ if and only if every row of B has a leading entry.*

Proof. Suppose first that every row of B has a leading entry. Let $\vec{v} \in \mathbb{R}^m$ be arbitrary, and fix $b_1, b_2, \dots, b_m \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_{m-1} \\ b_m \end{pmatrix}.$$

We want to know if there exists $c_1, c_2, \dots, c_n \in \mathbb{R}$ with $c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n = \vec{v}$, which means we want to know if a certain linear system with m equations and n variables has a solution. Now the augmented matrix of this linear system is obtained by appending \vec{v} as a new final column onto the end of A . If we apply the same elementary row operations to this new matrix that produced B , then since every row of B has a leading entry, we end up with an $m \times (n+1)$ matrix whose last column does not have a leading entry. Using Proposition 4.2.12, it follows that the system does indeed have a solution, so $\vec{v} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. Since $\vec{v} \in \mathbb{R}^m$ was arbitrary, we conclude that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = \mathbb{R}^m$.

We now prove the other direction by proving the contrapositive. Suppose that some row of B does not have a leading entry. Then B must have a row of all zeros, and since B is in echelon form, it follows that the last row of B consists entirely of zeros. Consider the vector

$$\vec{e}_m = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix} \in \mathbb{R}^m.$$

Suppose that we append \vec{e}_m as a new final column onto the end of B . If we apply the same elementary row operations we did to get from A to B in reverse by starting from B (which is possible by Proposition 4.2.4), then we will end with an $m \times (n+1)$ column whose first n columns are A , and whose last column will be a vector $\vec{v} \in \mathbb{R}^m$. Since B with the appended vector \vec{e}_m has a leading entry in the last row, we know from Proposition 4.2.12 that the corresponding linear system does not have a solution. Therefore, there does not exist $c_1, c_2, \dots, c_n \in \mathbb{R}$ with $c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n = \vec{v}$. It follows that $\vec{v} \notin \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$, and so $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) \neq \mathbb{R}^m$. \square

In Proposition 2.3.3, we showed that $\text{Span}(\vec{u}) \neq \mathbb{R}^2$ for all $\vec{u} \in \mathbb{R}^2$. Geometrically, the span of one vector in \mathbb{R}^2 is a line through the origin (unless that vector is $\vec{0}$), and a line does not cover the entire plane. If we go up to \mathbb{R}^3 , it is geometrically reasonable to believe that $\text{Span}(\vec{u}_1, \vec{u}_2) \neq \mathbb{R}^3$ for all $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^3$, because $\text{Span}(\vec{u}_1, \vec{u}_2)$ looks like a plane through the origin (unless one vector is a multiple of the other). Although the geometric picture looks convincing, it's actually a bit challenging to argue this purely algebraically from first principles. When we move beyond \mathbb{R}^3 , we lose our direct geometric visualization skills, and so we have to rely on the algebra. However, it seems reasonable to believe that we cannot find three elements of \mathbb{R}^4 whose span is \mathbb{R}^4 . We now go ahead and use the algebraic tools that we have developed to solve this problem generally.

Corollary 4.3.2. *If $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$ and $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = \mathbb{R}^m$, then $n \geq m$.*

Proof. Let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$ and suppose that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = \mathbb{R}^m$. Let A be the $m \times n$ matrix where the i^{th} column is \vec{u}_i . Fix an echelon form B of the matrix A . By Proposition 4.3.1, we know that every row of B has a leading entry. Since B is in echelon form, every column has at most one leading entry. Therefore, the number of rows of B , which equals the number of leading entry of B , must be less than or equal to the number of columns of B . In other words, we must have $m \leq n$. \square

By taking the contrapositive of this result, we have the following fundamental fact.

Corollary 4.3.3. *If $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$ and $n < m$, then $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) \neq \mathbb{R}^m$.*

We can also use some of these techniques to determine whether a sequence of elements of \mathcal{P}_n spans all of \mathcal{P}_n . Consider the following elements of \mathcal{P}_2 :

- $f_1(x) = x^2 - 2x$.
- $f_2(x) = x^2 - 3x + 4$.

- $f_3(x) = x + 5$.

Does

$$\text{Span}(f_1, f_2, f_3) = \mathcal{P}_2?$$

We want to know whether for all $g \in \mathcal{P}_2$, there exist $c_1, c_2, c_3 \in \mathbb{R}$ with $g = c_1 f_1 + c_2 f_2 + c_3 f_3$. Now given $g \in \mathcal{P}_2$, we can fix $a_0, a_1, a_2 \in \mathbb{R}$ such that $g(x) = a_2 x^2 + a_1 x + a_0$ for all $x \in \mathbb{R}$. Thus, we want to know whether for all $a_0, a_1, a_2 \in \mathbb{R}$, there exist $c_1, c_2, c_3 \in \mathbb{R}$ such that for all $x \in \mathbb{R}$, we have

$$c_1 \cdot (x^2 - 2x) + c_2 \cdot (x^2 - 3x + 4) + c_3 \cdot (x + 5) = a_2 x^2 + a_1 x + a_0.$$

By factoring out the left-hand side and combining terms of the same power, we want to know whether for all $a_0, a_1, a_2 \in \mathbb{R}$, there exists $c_1, c_2, c_3 \in \mathbb{R}$ such that for all $x \in \mathbb{R}$, we have

$$(c_1 + c_2) \cdot x^2 + (-2c_1 - 3c_2 + c_3) \cdot x + (4c_2 + 5c_3) = a_2 x^2 + a_1 x + a_0.$$

Since polynomials are equal precisely when the corresponding coefficients are equal, this is equivalent to asking whether for all $a_0, a_1, a_2 \in \mathbb{R}$, the system

$$\begin{array}{ccccccc} c_1 & + & c_2 & & & = & a_2 \\ -2c_1 & - & 3c_2 & + & c_3 & = & a_1 \\ & & 4c_2 & + & 5c_3 & = & a_0 \end{array}$$

has a solution. Applying elementary row operations to the augmented matrix, we obtain

$$\begin{aligned} \begin{pmatrix} 1 & 1 & 0 & a_2 \\ -2 & -3 & 1 & a_1 \\ 0 & 4 & 5 & a_0 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 1 & 0 & a_2 \\ 0 & -1 & 1 & a_1 + 2a_2 \\ 0 & 4 & 5 & a_0 \end{pmatrix} && (2R_1 + R_2) \\ &\rightarrow \begin{pmatrix} 1 & 1 & 0 & a_2 \\ 0 & -1 & 1 & a_1 + 2a_2 \\ 0 & 0 & 9 & a_0 + 4a_1 + 8a_2 \end{pmatrix} && (4R_2 + R_3). \end{aligned}$$

Since we can always solve this system no matter what $a_0, a_1, a_2 \in \mathbb{R}$ happen to be, we can always find $c_1, c_2, c_3 \in \mathbb{R}$ making the above true. Therefore, $\text{Span}(f_1, f_2, f_3) = \mathcal{P}_2$.

Notice alternatively that we can rephrase the question of whether for all $a_0, a_1, a_2 \in \mathbb{R}$, the system

$$\begin{array}{ccccccc} c_1 & + & c_2 & & & = & a_2 \\ -2c_1 & - & 3c_2 & + & c_3 & = & a_1 \\ & & 4c_2 & + & 5c_3 & = & a_0 \end{array}$$

has a solution as asking whether

$$\text{Span} \left(\begin{pmatrix} 1 \\ -2 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ -3 \\ 4 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 5 \end{pmatrix} \right) = \mathbb{R}^3.$$

Thus, we can apply elementary row operations to the 3×3 matrix with these vectors as columns (and without the extra column of a_2, a_1, a_0), and apply Proposition 4.3.1.

4.4 Linear Independence

Given a vector space V and elements $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in V$, we have a way to say that we can reach every element of V by saying that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$. However, if $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$, then we might have some redundancy. For example, we have that

$$\text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} -2 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \end{pmatrix} \right) = \mathbb{R}^2.$$

One way to see this is to notice that

$$\text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} -2 \\ 3 \end{pmatrix} \right) = \mathbb{R}^2$$

because $1 \cdot 3 - 2 \cdot (-2) = 7$ is nonzero, and then use the fact that

$$\text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} -2 \\ 3 \end{pmatrix} \right) \subseteq \text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} -2 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \end{pmatrix} \right).$$

Alternatively, we can perform an elementary row operation:

$$\begin{pmatrix} 1 & -2 & 2 \\ 2 & 3 & -1 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & -2 & 2 \\ 0 & 7 & -5 \end{pmatrix} \quad (-2R_1 + R_2).$$

Since the matrix on the right has a leading entry in each row, we can apply Proposition 4.3.1.

Although we do have

$$\text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} -2 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \end{pmatrix} \right) = \mathbb{R}^2,$$

it turns out we can reach elements of \mathbb{R}^2 in many ways, and so there is some redundancy in using these three vectors. For example, we have

$$\begin{pmatrix} 2 \\ -1 \end{pmatrix} \in \text{Span} \left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} -2 \\ 3 \end{pmatrix} \right)$$

because the span of the latter two vectors is all of \mathbb{R}^2 . Explicitly, we can find the constants that work in this case by interpreting the above matrix as an augmented matrix and solving the corresponding system. If we do this, then we determine that

$$\begin{pmatrix} 2 \\ -1 \end{pmatrix} = 4/7 \cdot \begin{pmatrix} 1 \\ 2 \end{pmatrix} + (-5/7) \cdot \begin{pmatrix} -2 \\ 3 \end{pmatrix}.$$

Since the third vector is a linear combination of the first two, it seems wasteful to include it in the list.

Suppose more generally that we have a vector space V and a list of vectors $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in V$. When thinking about what we can reach from these vectors through linear combinations, how can we formally express the idea that there is some redundancy in the list $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$? Certainly, if a vector appears twice in the list, then that seems wasteful. Another possibility is that some \vec{u}_j is a multiple of a different \vec{u}_i . However, as the above example shows, there seems to be some redundancy in the list

$$\left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} -2 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \end{pmatrix} \right)$$

even though no vector in the list is a multiple of another. The real issue is that we can reach one of the vectors using a linear combination of the others (in fact, in this case, we can reach each of the three vectors using a linear combination of the other two).

Although checking if one vector in our list $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is a linear combination of the others is a natural idea, it is rather painful in practice. For example, consider the list

$$\left(\begin{pmatrix} 1 \\ 2 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -2 \\ -3 \\ 4 \\ 5 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ -13 \\ 9 \end{pmatrix} \right)$$

of vectors in \mathbb{R}^4 . If we want to determine if any vector is a linear combination of the other two, then we have to determine whether any of three different linear systems has a solution, and so we would have to perform Gaussian Elimination on three different 4×3 matrices. Is there a better way?

Let's go back and look at our above example with

$$\begin{pmatrix} 2 \\ -1 \end{pmatrix} = 4/7 \cdot \begin{pmatrix} 1 \\ 2 \end{pmatrix} + (-5/7) \cdot \begin{pmatrix} -2 \\ 3 \end{pmatrix}.$$

Clearing denominators and moving everything to one side, we conclude that

$$4 \cdot \begin{pmatrix} 1 \\ 2 \end{pmatrix} + (-5) \cdot \begin{pmatrix} -2 \\ 3 \end{pmatrix} + (-7) \cdot \begin{pmatrix} 2 \\ -1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

Thus, we see that we have an interesting linear combination of our three vectors that results in $\vec{0}$. Moreover, notice the symmetry in this equation because all three of our vectors are on the left, and we are finding a nontrivial linear combination of them that produces $\vec{0}$, rather than a linear combination of two of them that produces the third. With this example in mind, we make the following definition.

Definition 4.4.1. Let V be a vector space and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in V$. We say that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is a linearly independent sequence if the following statement is true:

For all $c_1, c_2, \dots, c_n \in \mathbb{R}$, if $c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n = \vec{0}$, then $c_1 = c_2 = \dots = c_n = 0$.

In other words, $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent if the only linear combination of the \vec{u}_i that produces $\vec{0}$ is the trivial one. If $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is not linearly independent, then we say that it is linearly dependent.

Be very careful when negating this statement. The negation turns the “for all” into a “there exists”, but the negation does *not* say that there is a linear combination of the \vec{u}_i giving $\vec{0}$ in which all of the c_i are nonzero. Instead, a sequence $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly dependent if there exists $c_1, c_2, \dots, c_n \in \mathbb{R}$ with $c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n = \vec{0}$ and such that at least one c_i is nonzero. For example, the sequence of vectors

$$\left(\begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ 4 \end{pmatrix}, \begin{pmatrix} 3 \\ 1 \end{pmatrix} \right)$$

is linear dependent because

$$(-2) \cdot \begin{pmatrix} 1 \\ 2 \end{pmatrix} + 1 \cdot \begin{pmatrix} 2 \\ 4 \end{pmatrix} + 0 \cdot \begin{pmatrix} 3 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

and at least one of the coefficients is nonzero (i.e. it doesn't matter that the last one is 0).

Although this definition may feel a bit strange at first, we will grow to really appreciate its strengths. Before doing that, however, we first show that if we have at least two vectors, then our definition is equivalent to the statement that none of the \vec{u}_i is a linear combination of the others. The fundamental idea is embodied in the above example where we moved elements to one side and cleared/introduced denominators.

Proposition 4.4.2. Let V be a vector space, and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in V$ where $n \geq 2$. The following are equivalent:

1. $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent.
2. There does not exist an i such that $\vec{u}_i \in \text{Span}(\vec{u}_1, \dots, \vec{u}_{i-1}, \vec{u}_{i+1}, \dots, \vec{u}_n)$.

Proof. We prove this by proving the contrapositive of each direction. Suppose first that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly dependent. Fix $c_1, c_2, \dots, c_n \in \mathbb{R}$ with $c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n = \vec{0}$ and such that at least one c_j is nonzero. Fix i with $c_i \neq 0$. We then have

$$(-c_i) \cdot \vec{u}_i = c_1\vec{u}_1 + \dots + c_{i-1}\vec{u}_{i-1} + c_{i+1}\vec{u}_{i+1} + \dots + c_n\vec{u}_n.$$

Multiplying both sides by $-\frac{1}{c_i}$, we see that

$$\vec{u}_i = (-c_1/c_i) \cdot \vec{u}_1 + \cdots + (-c_{i-1}/c_i) \cdot \vec{u}_{i-1} + (-c_{i+1}/c_i) \cdot \vec{u}_{i+1} + \cdots + (-c_n/c_i) \cdot \vec{u}_n.$$

Therefore, $\vec{u}_i \in \text{Span}(\vec{u}_1, \dots, \vec{u}_{i-1}, \vec{u}_{i+1}, \dots, \vec{u}_n)$.

Suppose conversely that there exists an i such that $\vec{u}_i \in \text{Span}(\vec{u}_1, \dots, \vec{u}_{i-1}, \vec{u}_{i+1}, \dots, \vec{u}_n)$. Fix such an i , and then fix $c_1, \dots, c_{i-1}, c_{i+1}, c_n \in \mathbb{R}$ with

$$\vec{u}_i = c_1 \vec{u}_1 + \cdots + c_{i-1} \vec{u}_{i-1} + c_{i+1} \vec{u}_{i+1} + \cdots + c_n \vec{u}_n.$$

We then have that

$$\vec{0} = c_1 \vec{u}_1 + \cdots + c_{i-1} \vec{u}_{i-1} + (-1) \cdot \vec{u}_i + c_{i+1} \vec{u}_{i+1} + \cdots + c_n \vec{u}_n.$$

Since $-1 \neq 0$, we have shown that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly dependent. \square

What happens if we only have 1 vector? In this case, it's unclear what we mean if we try to say that one vector is a linear combination of the others, because there are no others. However, our official definition does not have this problem.

Proposition 4.4.3. *Let V be a vector space, and let $\vec{u} \in V$. We then have that (\vec{u}) is linearly independent if and only if $\vec{u} \neq \vec{0}$.*

Proof. We prove the contrapositive of each direction:

- Suppose first that $\vec{u} = \vec{0}$. We then have that $1 \cdot \vec{u} = 1 \cdot \vec{0} = \vec{0}$. Since $1 \neq 0$, it follows that (\vec{u}) is linearly dependent.
- Suppose that (\vec{u}) is linearly dependent. By definition, we can fix $c \neq 0$ with $c \cdot \vec{u} = \vec{0}$. Since $c \neq 0$, we can multiply both sides of $c\vec{u} = \vec{0}$ by $1/c$ to conclude that

$$\frac{1}{c} \cdot (c\vec{u}) = \frac{1}{c} \cdot \vec{0}.$$

Now

$$\begin{aligned} \frac{1}{c} \cdot (c\vec{u}) &= \left(\frac{1}{c} \cdot c \right) \cdot \vec{u} \\ &= 1 \cdot \vec{u} \\ &= \vec{u} \end{aligned}$$

and

$$\frac{1}{c} \cdot \vec{0} = \vec{0}$$

by Proposition 4.1.11. Therefore, we have $\vec{u} = \vec{0}$. \square

Finally, what about the empty sequence? Should we consider $()$ to be linearly independent? If we follow the definition carefully, then the answer is yes. Why? Is it the case that each time we have a linearly combination of the vectors in $()$ that gives $\vec{0}$, each of the coefficients is 0. Yes, for the vacuous reason that there is no such linear combination giving $\vec{0}$! If you don't find that bit of logic convincing, then we can simply adopt it as a convention like we did for Definition 4.1.17.

Definition 4.4.4. *In any vector space V , the empty sequence $()$ is considered to be linearly independent.*

Let's return to our question of asking if

$$\left(\begin{pmatrix} 1 \\ 2 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -2 \\ -3 \\ 4 \\ 5 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ -13 \\ 9 \end{pmatrix} \right)$$

is a linearly independent sequence in \mathbb{R}^4 . As mentioned above, if we were to use the Proposition 4.4.2, then we would have to work with 3 different linear systems. However, let's instead directly use the definition. Let $c_1, c_2, c_3 \in \mathbb{R}$ be arbitrary with

$$c_1 \cdot \begin{pmatrix} 1 \\ 2 \\ 0 \\ 1 \end{pmatrix} + c_2 \cdot \begin{pmatrix} -2 \\ -3 \\ 4 \\ 5 \end{pmatrix} + c_3 \cdot \begin{pmatrix} 2 \\ -1 \\ -13 \\ 9 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

We want to show that $c_1 = c_2 = c_3 = 0$. Performing the scalar multiplication and addition on the left-hand side, we see that c_1, c_2, c_3 must satisfy the following system of equations:

$$\begin{array}{rcccccl} c_1 & - & 2c_2 & + & 2c_3 & = & 0 \\ 2c_1 & - & 3c_2 & - & c_3 & = & 0 \\ & & 4c_2 & - & 13c_3 & = & 0 \\ c_1 & + & 5c_2 & + & 9c_3 & = & 0. \end{array}$$

In order to show that we must have $c_1 = c_2 = c_3 = 0$, we just need to show that the only solution to this system is the trivial solution $(0, 0, 0)$. Performing Gaussian Elimination on the augmented matrix, we obtain

$$\begin{aligned} \left(\begin{array}{cccc} 1 & -2 & 2 & 0 \\ 2 & -3 & -1 & 0 \\ 0 & 4 & -13 & 0 \\ 1 & 5 & 9 & 0 \end{array} \right) &\rightarrow \left(\begin{array}{cccc} 1 & -2 & 2 & 0 \\ 0 & 1 & -5 & 0 \\ 0 & 4 & -13 & 0 \\ 0 & 7 & 7 & 0 \end{array} \right) & \begin{array}{l} (-2R_1 + R_2) \\ (-R_1 + R_4) \end{array} \\ &\rightarrow \left(\begin{array}{cccc} 1 & -2 & 2 & 0 \\ 0 & 1 & -5 & 0 \\ 0 & 0 & 7 & 0 \\ 0 & 0 & 42 & 0 \end{array} \right) & \begin{array}{l} (-4R_2 + R_3) \\ (-7R_2 + R_4) \end{array} \\ &\rightarrow \left(\begin{array}{cccc} 1 & -2 & 2 & 0 \\ 0 & 1 & -5 & 0 \\ 0 & 0 & 7 & 0 \\ 0 & 0 & 0 & 0 \end{array} \right) & (-6R_3 + R_4). \end{aligned}$$

Since the last column does not have a leading entry, but every other column does, we conclude from Proposition 4.2.12 that the system has a unique solution. Since $(0, 0, 0)$ is trivially one solution, it follows that the solution set is $\{(0, 0, 0)\}$. Therefore, we must have $c_1 = c_2 = c_3 = 0$. We conclude that

$$\left(\begin{pmatrix} 1 \\ 2 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -2 \\ -3 \\ 4 \\ 5 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ -13 \\ 9 \end{pmatrix} \right)$$

is linearly independent. Notice that we only needed to perform Gaussian Elimination on *one* 4×4 matrix, rather than on three 4×3 matrices. Moreover, notice that the last column was all zeros, which will always be the case when setting up a problem like this in \mathbb{R}^m . Since elementary row operations do not affect a zero column, there is no need to carry it around, as we now argue.

Proposition 4.4.5. *Let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$. Let A be the $m \times n$ matrix where the i^{th} column is \vec{u}_i , and let B be an echelon form of A . We then have that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent if and only if every column of B has a leading entry.*

Proof. Let

$$S = \left\{ \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix} \in \mathbb{R}^n : c_1 \vec{u}_1 + c_2 \vec{u}_2 + \dots + c_n \vec{u}_n = \vec{0} \right\}.$$

Notice that $(0, 0, \dots, 0) \in S$, and that we have $S = \{(0, 0, \dots, 0)\}$ if and only if $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent.

Now S is the solution set of the linear system with m equations and n variables obtained by opening up the \vec{u}_i and viewing the c_i as the variables. The augmented matrix of this linear system is obtained by appending $\vec{0}$ as a new final column onto the end of A . If we apply the same elementary row operations that produced B from A to this augmented matrix, then the first n columns will be B , and the last will be $\vec{0}$. Thus, this matrix does not have a leading entry in the last column. Now the system is consistent because $(0, 0, \dots, 0)$ is a solution, so applying Proposition 4.2.12, we conclude that $S = \{(0, 0, \dots, 0)\}$ if and only if every column of B has a leading entry. Therefore, $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent if and only if every column of B has a leading entry. \square

For example, suppose that we want to determine if the vectors

$$\left(\begin{pmatrix} 1 \\ 0 \\ -2 \end{pmatrix}, \begin{pmatrix} 0 \\ 3 \\ 6 \end{pmatrix}, \begin{pmatrix} -4 \\ -2 \\ 3 \end{pmatrix} \right)$$

are linearly independent in \mathbb{R}^3 . We simply form the matrix with these vectors as columns, perform Gaussian Elimination, and check if each column has a leading entry. We have

$$\begin{aligned} \begin{pmatrix} 1 & 0 & -4 \\ 0 & 3 & -2 \\ -2 & 6 & 3 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 0 & -4 \\ 0 & 3 & -2 \\ 0 & 6 & -5 \end{pmatrix} && (2R_1 + R_3) \\ &\rightarrow \begin{pmatrix} 1 & 0 & -4 \\ 0 & 3 & -2 \\ 0 & 0 & -1 \end{pmatrix} && (-2R_2 + R_3). \end{aligned}$$

We have reached an echelon form of A , and we notice that every column has a leading entry. We conclude that the vectors are linearly independent.

We now obtain some important results like we did after the corresponding fact for spans. For example, they imply that we cannot find 5 linearly independent vectors in \mathbb{R}^4 . Intuitively, this feels very reasonable because it feels impossible to push in 5 completely independent direction with the “4-dimensional space” \mathbb{R}^4 . However, this would have been difficult to prove directly without all our theory.

Corollary 4.4.6. *If $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$ and $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent, then $n \leq m$.*

Proof. Let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$ and suppose that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent. Let A be the $m \times n$ matrix where the i^{th} column is \vec{u}_i . Fix an echelon form B of the matrix A . By Proposition 4.4.5, we know that every column of B has a leading entry. Furthermore, since B is in echelon form, every column of B has exactly one leading entry. Since every row of B has at most one leading entry by definition, the number of columns of B must be less than or equal to the number of rows of B . In other words, we must have $n \leq m$. \square

By taking the contrapositive of this result, we have the following fundamental fact.

Corollary 4.4.7. *If $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$ and $n > m$, then $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly dependent.*

Be sure to keep the results about spans and linear independence straight! Let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$, let A be the $m \times n$ matrix where the i^{th} column is \vec{u}_i , and let B be an echelon form of A . We have that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = \mathbb{R}^m$ if and only if every *row* of B has a leading entry. We have that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent if and only if every *column* of B has a leading entry. Also, if $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = \mathbb{R}^m$, then $n \geq m$. On the other hand, if $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent, then $n \leq m$.

We can also apply these ideas to vector spaces other than \mathbb{R}^m . Let \mathcal{P}_2 be the vector space of all polynomials of degree at most 2. Consider the following elements of \mathcal{P}_2 :

- $f_1(x) = x^2 - 2x$.
- $f_2(x) = x^2 - 3x + 4$.
- $f_3(x) = x + 5$.

We claim that (f_1, f_2, f_3) is linearly independent. Let $c_1, c_2, c_3 \in \mathbb{R}$ be arbitrary with $c_1 f_1 + c_2 f_2 + c_3 f_3 = 0$. Since functions are equal if and only if they agree on all inputs, we know that for all $x \in \mathbb{R}$, we have

$$c_1 \cdot (x^2 - 2x) + c_2 \cdot (x^2 - 3x + 4) + c_3 \cdot (x + 5) = 0.$$

We rephrase this as saying that for all $x \in \mathbb{R}$, we have

$$(c_1 + c_2) \cdot x^2 + (-2c_1 - 3c_2 + c_3) \cdot x + (4c_2 + 5c_3) = 0x^2 + 0x + 0.$$

Since polynomials give equal values on all inputs precisely when the corresponding coefficients are equal, this implies that

$$\begin{array}{rcccccl} c_1 & + & c_2 & & & = & 0 \\ -2c_1 & - & 3c_2 & + & c_3 & = & 0 \\ & & 4c_2 & + & 5c_3 & = & 0. \end{array}$$

Thus, to show that $c_1 = c_2 = c_3 = 0$, we need only show that this system has only the trivial solution $(0, 0, 0)$. Applying elementary row operations to the augmented matrix, we obtain

$$\begin{aligned} \begin{pmatrix} 1 & 1 & 0 & 0 \\ -2 & -3 & 1 & 0 \\ 0 & 4 & 5 & 0 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 4 & 5 & 0 \end{pmatrix} && (2R_1 + R_2) \\ &\rightarrow \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & 9 & 0 \end{pmatrix} && (4R_1 + R_3). \end{aligned}$$

Since there is no leading entry in the last column, but there is a leading entry in all other columns, we conclude that the system has a unique solution, namely $\{(0, 0, 0)\}$. Therefore, we must have that $c_1 = c_2 = c_3 = 0$. It follows that (f_1, f_2, f_3) is linearly independent.

Notice that we performed almost exactly the same calculations when we checked that $\text{Span}(f_1, f_2, f_3) = \mathcal{P}_2$ in the previous section. The one key difference is that we previously put arbitrary values in the last column whereas now we used zeros, but in both cases the last column did not really matter. Also, we checked for leading entries in all rows when we were looking at the span, and now we check for leading entries in all columns (excluding the last). Since there is a leading entry in all rows and all but the last column in this case, we had *both* $\text{Span}(f_1, f_2, f_3) = \mathcal{P}_2$ and also that (f_1, f_2, f_3) was linearly independent.

For a more exotic example, consider the vector space \mathcal{F} of all functions $f: \mathbb{R} \rightarrow \mathbb{R}$. Recall that in \mathcal{F} , the zero vector is the constant zero function. Let $f_1: \mathbb{R} \rightarrow \mathbb{R}$ be the function $f_1(x) = \sin x$, and let $f_2: \mathbb{R} \rightarrow \mathbb{R}$ be the function $f_2(x) = \cos x$. We claim that (f_1, f_2) is linearly independent. To see this, let $c_1, c_2 \in \mathbb{R}$ be arbitrary with $c_1 f_1 + c_2 f_2 = 0$. Since functions are equal exactly when they give the same output on every input, we then have that $c_1 f_1(x) + c_2 f_2(x) = 0$ for all $x \in \mathbb{R}$. In other words, we have

$$c_1 \sin x + c_2 \cos x = 0$$

for all $x \in \mathbb{R}$. Since this statement is true for all $x \in \mathbb{R}$, it is true whenever we plug in a specific value. Plugging in $x = 0$, we see that

$$c_1 \cdot 0 + c_2 \cdot 1 = 0$$

so $c_2 = 0$. Plugging in $x = \frac{\pi}{2}$, we see that

$$c_1 \cdot 1 + c_2 \cdot 0 = 0$$

so $c_1 = 0$. We took arbitrary $c_1, c_2 \in \mathbb{R}$ with $c_1 f_1 + c_2 f_2 = 0$, and showed that $c_1 = c_2 = 0$. Therefore, (f_1, f_2) is linearly independent.

In contrast, suppose that we let $g_1: \mathbb{R} \rightarrow \mathbb{R}$ be the function $g_1(x) = \sin 2x$, and let $g_2: \mathbb{R} \rightarrow \mathbb{R}$ be the function $g_2(x) = \sin x \cos x$. We claim that (g_1, g_2) is linearly dependent. To see this, recall the trigonometric identity that for all $x \in \mathbb{R}$, we have

$$\sin 2x = 2 \sin x \cos x.$$

It follows that for all $x \in \mathbb{R}$, we have

$$\sin 2x - 2 \sin x \cos x = 0,$$

and hence

$$1 \cdot g_1 + (-2) \cdot g_2 = 0.$$

Since we have found a nontrivial linear combination of g_1 and g_2 which gives the zero vector of \mathcal{F} , we conclude that (g_1, g_2) is linearly dependent.

4.5 Bases and Coordinates

We can combine the concept of “enough vectors” provided by spanning with the “no redundancy requirement” of linear independence to form the following fundamental concept.

Definition 4.5.1. *Let V be a vector space. A basis of V is a sequence of vectors $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ of V such that both of the following are true:*

1. $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$.
2. $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent.

At first, it might appear that this definition conflicts with Definition 2.3.11. Recall that given $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^2$, we said that (\vec{u}_1, \vec{u}_2) was a basis for \mathbb{R}^2 if $\text{Span}(\vec{u}_1, \vec{u}_2) = \mathbb{R}^2$. In other words, we did not have the assumption that (\vec{u}_1, \vec{u}_2) was linearly independent in Definition 2.3.11. The reason is that if $\text{Span}(\vec{u}_1, \vec{u}_2) = \mathbb{R}^2$, then (\vec{u}_1, \vec{u}_2) is automatically linear independent! To see this, we can apply Theorem 2.3.10, which says that if $\text{Span}(\vec{u}_1, \vec{u}_2) = \mathbb{R}^2$, then $\vec{u}_1 \notin \text{Span}(\vec{u}_2)$ and $\vec{u}_2 \notin \text{Span}(\vec{u}_1)$, so (\vec{u}_1, \vec{u}_2) is linearly independent by Proposition 4.4.2. However, be aware that this simplification does not work in general (after all, at the beginning of Section 4.4, we saw an example of three vectors that span \mathbb{R}^2 but are not linearly independent). The key distinction that makes things work in this situation is that we have two vectors in \mathbb{R}^2 . We will eventually generalize this situation in Proposition 4.6.7.

Let's instead work with the new definition on its own terms. Let

$$\alpha = \left(\begin{pmatrix} 1 \\ -2 \end{pmatrix}, \begin{pmatrix} 5 \\ 1 \end{pmatrix} \right).$$

We claim that α is a basis for \mathbb{R}^2 . To see this, we could simply notice that $1 \cdot 1 - (-2) \cdot 5 = 11$ is nonzero, so

$$\text{Span} \left(\begin{pmatrix} 1 \\ -2 \end{pmatrix}, \begin{pmatrix} 5 \\ 1 \end{pmatrix} \right) = \mathbb{R}^2$$

by Theorem 2.3.10, and then refer to the previous paragraph. But let's imagine that we forgot all of Chapter 2. Instead, we look to Proposition 4.3.1 and Proposition 4.4.5, and apply elementary row operations to the matrix having these vectors as columns:

$$\begin{pmatrix} 1 & 5 \\ -2 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 5 \\ 0 & 11 \end{pmatrix} \quad (2R_1 + R_2).$$

Since there is a leading entry in each row, we know that

$$\text{Span} \left(\begin{pmatrix} 1 \\ -2 \end{pmatrix}, \begin{pmatrix} 5 \\ 1 \end{pmatrix} \right) = \mathbb{R}^2$$

And since there is a leading entry in each column, we know that α is linearly independent. Therefore, α is a basis for \mathbb{R}^2 .

A much simpler basis for \mathbb{R}^2 is

$$\varepsilon = \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right).$$

It is straightforward to check that this sequence of two vectors both spans \mathbb{R}^2 and is linearly independent. Alternatively, the matrix having these two vectors as columns is I , which is in echelon form and has a leading entry in every row and every column.

More generally, we can give natural bases for some of our favorite vector spaces:

- Let $n \in \mathbb{N}^+$. For each i with $1 \leq i \leq n$, let \vec{e}_i be the vector with a 1 in the i^{th} position and 0's everywhere else. The sequence

$$\varepsilon_n = (\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n)$$

is a basis for \mathbb{R}^n (again, either by checking directly, or noticing that the corresponding matrix is in echelon form and has leading entries in every row and column). We call ε_n the *standard* or *natural* basis of \mathbb{R}^n .

- Let $n \in \mathbb{N}$ and consider the vector space \mathcal{P}_n of all polynomial functions of degree at most n . Is there a simple basis for \mathcal{P}_n ? For each $k \in \mathbb{N}$ with $0 \leq k \leq n$, let $f_k: \mathbb{R} \rightarrow \mathbb{R}$ be the function $f_k(x) = x^k$. That is, $f_0(x) = 1$, $f_1(x) = x$, $f_2(x) = x^2$, \dots , and $f_n(x) = x^n$. It is straightforward to check that $(f_0, f_1, f_2, \dots, f_n)$ is a basis for \mathcal{P}_n .
- There is a natural basis for the vector space V of all 2×2 matrices given by

$$\alpha = \left(\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \right).$$

- Finally, we handle a somewhat trivial but important situation. Let $V = \{\vec{0}\}$ be a vector space consisting of just the zero vector. We have that $\text{Span}() = V$ by Definition 4.1.17 and that $()$ is linearly independent by Definition 4.4.4. Thus, the empty sequence $()$ is a basis for V .

Although we have natural bases for the above vector spaces, always keep in mind that vector spaces typically have many other bases. We saw one example of a different basis for \mathbb{R}^2 above. Consider now the following sequence in \mathbb{R}^3 :

$$\alpha = \left(\begin{pmatrix} 1 \\ 3 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix}, \begin{pmatrix} 1 \\ -3 \\ 1 \end{pmatrix} \right).$$

We claim that α is a basis for \mathbb{R}^3 . To see this, we apply elementary row operations to the matrix having these three vectors as columns:

$$\begin{aligned} \begin{pmatrix} 1 & 2 & 1 \\ 3 & -1 & -3 \\ 2 & 3 & 1 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 2 & 1 \\ 0 & -7 & -6 \\ 0 & -1 & -1 \end{pmatrix} && \begin{array}{l} (-3R_1 + R_2) \\ (-2R_1 + R_3) \end{array} \\ &\rightarrow \begin{pmatrix} 1 & 2 & 1 \\ 0 & -1 & -1 \\ 0 & -7 & -6 \end{pmatrix} && \begin{array}{l} (R_2 \leftrightarrow R_3) \\ (R_2 \leftrightarrow R_3) \end{array} \\ &\rightarrow \begin{pmatrix} 1 & 2 & 1 \\ 0 & -1 & -1 \\ 0 & 0 & 1 \end{pmatrix} && (-7R_2 + R_3). \end{aligned}$$

Since this last matrix is in echelon form and has a leading entry in each row and each column, we may use Proposition 4.3.1 and Proposition 4.4.5 to conclude that α is a basis for \mathbb{R}^3 .

We now develop an equivalent characterization of bases that will be important for all of our later work.

Theorem 4.5.2. *Let V be a vector space and let $\alpha = (\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ be a sequence of elements of V . We then have that α is a basis for V if and only if every $\vec{v} \in V$ can be expressed as a linear combination of $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$ in a unique way.*

Proof. Suppose first that α is a basis for V . Let $\vec{v} \in V$ be arbitrary. Since $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$, we know that we can express \vec{v} as a linear combination of $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$ giving \vec{v} . This gives existence, and we now prove uniqueness. Let $c_1, c_2, \dots, c_n, d_1, d_2, \dots, d_n \in \mathbb{R}$ be arbitrary such that both of the following are true:

$$\begin{aligned} \vec{v} &= c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n \\ \vec{v} &= d_1\vec{u}_1 + d_2\vec{u}_2 + \dots + d_n\vec{u}_n. \end{aligned}$$

We prove that $c_i = d_i$ for all $i \in \{1, 2, \dots, n\}$. Since both of the right-hand sides equal \vec{v} , we have

$$c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n = d_1\vec{u}_1 + d_2\vec{u}_2 + \dots + d_n\vec{u}_n.$$

Subtracting the right-hand side from both sides, we conclude that

$$c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n - d_1\vec{u}_1 - d_2\vec{u}_2 - \dots - d_n\vec{u}_n = \vec{0}.$$

Using the vector space axioms to rearrange this expression, we conclude that

$$(c_1 - d_1)\vec{u}_1 + (c_2 - d_2)\vec{u}_2 + \dots + (c_n - d_n)\vec{u}_n = \vec{0}.$$

Since $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent, it follows that $c_i - d_i = 0$ for all $i \in \{1, 2, \dots, n\}$. Therefore, $c_i = d_i$ for all $i \in \{1, 2, \dots, n\}$. Since the c_i and d_i were arbitrary, it follows that \vec{v} can be expressed as a linear combination of $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$ in a unique way.

Conversely, suppose that every $\vec{v} \in V$ can be expressed as a linear combination of $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$ in a unique way. In particular, every $\vec{v} \in V$ is a linear combination of $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$, so $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$. We now need to prove linear independence. Let $c_1, c_2, \dots, c_n \in \mathbb{R}$ be arbitrary with

$$c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n = \vec{0}.$$

Now we also have

$$0\vec{u}_1 + 0\vec{u}_2 + \dots + 0\vec{u}_n = \vec{0}.$$

Since these are two ways to write $\vec{0}$ as a linear combination of $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$, our assumption about uniqueness tells us that we must have $c_i = 0$ for all $i \in \{1, 2, \dots, n\}$. Therefore, $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent. It follows that α is a basis of V . \square

With this result in hand, we can now generalize the idea of coordinates to other vector spaces.

Definition 4.5.3. Suppose that V is a vector space and that $\alpha = (\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is a basis of V . We define a function

$$\text{Coord}_\alpha: V \rightarrow \mathbb{R}^n$$

as follows. Given $\vec{v} \in V$, let $c_1, c_2, \dots, c_n \in \mathbb{R}$ be the unique values such that $\vec{v} = c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n$, and define

$$\text{Coord}_\alpha(\vec{v}) = \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix}.$$

We call this vector the coordinates of \vec{v} relative to α . We also use the notation $[\vec{v}]_\alpha$ for $\text{Coord}_\alpha(\vec{v})$.

For example, let

$$\alpha = \left(\begin{pmatrix} 1 \\ 3 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix}, \begin{pmatrix} 1 \\ -3 \\ 1 \end{pmatrix} \right).$$

We checked above that α is a basis for \mathbb{R}^3 . Let's calculate

$$\text{Coord}_\alpha \left(\begin{pmatrix} -1 \\ 13 \\ 0 \end{pmatrix} \right) = \left[\begin{pmatrix} -1 \\ 13 \\ 0 \end{pmatrix} \right]_\alpha.$$

In order to do this, we want to calculate the unique values $c_1, c_2, c_3 \in \mathbb{R}$ such that

$$c_1 \begin{pmatrix} 1 \\ 3 \\ 2 \end{pmatrix} + c_2 \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix} + c_3 \begin{pmatrix} 1 \\ -3 \\ 1 \end{pmatrix} = \begin{pmatrix} -1 \\ 13 \\ 0 \end{pmatrix}.$$

To do this, we can form the augmented matrix

$$\begin{pmatrix} 1 & 2 & 1 & -1 \\ 3 & -1 & -3 & 13 \\ 2 & 3 & 1 & 0 \end{pmatrix}$$

and perform elementary row operations to solve the system. Working through the calculations, we see that the unique solution is $c_1 = 5$, $c_2 = -4$, and $c_3 = 2$. Therefore, we have

$$\left[\begin{pmatrix} -1 \\ 13 \\ 0 \end{pmatrix} \right]_\alpha = \begin{pmatrix} 5 \\ -4 \\ 2 \end{pmatrix}.$$

Remember that in \mathbb{R}^2 , we took a basis α , and the used vectors within it as the foundation for a grid system formed by parallelograms, and we interpreted Coord_α as geometrically describing how much we needed to stretch each vector to reach the given point. In \mathbb{R}^3 , a basis provides us with a similar grid system, but now formed using three directions and made up of parallelepipeds (slanty 3-dimensional cube-like things). Given a basis α and $\vec{v} \in \mathbb{R}^3$, we then have that $\text{Coord}_\alpha(\vec{v}) = [\vec{v}]_\alpha$ gives us the position of \vec{v} within this grid system.

When we move beyond \mathbb{R}^3 to a general \mathbb{R}^n , our theory allows us to carry over this intuition despite our loss of the ability to picture it geometrically. Moreover, our theory works even for more exotic vector spaces. Notice that if α is a basis of a vector space V , and α consists of n vectors, then Coord_α takes as inputs elements of V and produces as output an element of \mathbb{R}^n . Thus, even if V consists of functions (or even less tangible objects), Coord_α will take an element of V and output a list of n numbers. Thus, a basis provides a kind of grid system even in these settings, and can turn very complicated objects into lists of numbers telling us how to “locate” them within the vector space V .

For example, let \mathcal{P}_2 be the vector space of all polynomials of degree at most 2. Consider the following elements of \mathcal{P}_2 :

- $f_1(x) = x^2 - 2x$.
- $f_2(x) = x^2 - 3x + 4$.
- $f_3(x) = x + 5$.

We used this example in past discussions and showed both that $\text{Span}(f_1, f_2, f_3) = \mathcal{P}_2$ and that (f_1, f_2, f_3) is linearly independent. Therefore, $\alpha = (f_1, f_2, f_3)$ is a basis of \mathcal{P}_2 . Let $g: \mathbb{R} \rightarrow \mathbb{R}$ be the function $g(x) = x^2 + 6x + 13$. Let's calculate $[g]_\alpha$. Intuitively, we want to find the amount that we need to scale the two parabolas and the line in α to reach g . More formally, we want to find the unique $c_1, c_2, c_3 \in \mathbb{R}$ with $g = c_1 f_1 + c_2 f_2 + c_3 f_3$. Working through the algebra and solving the resulting system, it turns out that $c_1 = 4$, $c_2 = -3$, and $c_3 = 5$. Therefore, we have

$$[g]_\alpha = \begin{pmatrix} 4 \\ -3 \\ 5 \end{pmatrix}.$$

For another example, let \mathcal{F} be the vector space of all functions from \mathbb{R} to \mathbb{R} . Consider the following two elements of \mathcal{F} :

- $f_1(x) = \sin^2 x$.
- $f_2(x) = \cos^2 x$.

Let $W = \text{Span}(f_1, f_2)$, and notice that W is a subspace of \mathcal{F} by Proposition 4.1.16. Thus, we can view W as a vector space in its own right. Now we clearly have $W = \text{Span}(f_1, f_2)$ by definition. We next show that (f_1, f_2) is linearly independent. Let $c_1, c_2 \in \mathbb{R}$ be arbitrary with $c_1 f_1 + c_2 f_2 = 0$, i.e. with

$$c_1 \sin^2 x + c_2 \cos^2 x = 0$$

for all $x \in \mathbb{R}$. Plugging in $x = 0$, we see that $c_2 = 0$, and plugging in $x = \frac{\pi}{2}$, we see that $c_1 = 0$. Therefore, (f_1, f_2) is linearly independent as well, and hence $\alpha = (f_1, f_2)$ is a basis for W . Let $g: \mathbb{R} \rightarrow \mathbb{R}$ be the function given by $g(x) = \cos 2x$. Recall that for all $x \in \mathbb{R}$, we have

$$\cos 2x = \cos^2 x - \sin^2 x,$$

so

$$g = (-1) \cdot f_1 + 1 \cdot f_2.$$

It follows that $g \in W$, and that

$$[g]_{\alpha} = \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

Although most of the vector spaces we have talked about so far have “natural” bases, this need not always be the case. For example, consider the set

$$W = \left\{ \begin{pmatrix} a \\ b \\ c \end{pmatrix} \in \mathbb{R}^3 : a + b + c = 0 \right\}.$$

It is straightforward to check that W is a subspace, so we can consider W as a vector space in its own right. Moreover, it is a good exercise to show that

$$W = \text{Span} \left(\begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} \right).$$

Notice that these two vectors do *not* span all of \mathbb{R}^3 , but just the smaller subspace W . Don’t be confused by the fact that if we apply elementary row operations to the corresponding 3×2 matrix, we do not have a leading entry in each row (that condition only applies to determining if the vectors span \mathbb{R}^3).

In fact, this sequence of two vectors is also linearly independent. One can show this by looking at the matrix, but we can also do it directly. Let $c_1, c_2 \in \mathbb{R}$ be arbitrary with

$$c_1 \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix} + c_2 \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

We then have that

$$\begin{pmatrix} c_1 + c_2 \\ -c_1 \\ -c_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

Therefore, $-c_1 = 0$ and $-c_2 = 0$, which implies that $c_1 = 0$ and $c_2 = 0$. It follows that

$$\alpha = \left(\begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} \right)$$

is a basis for W . Notice that even though W lives inside \mathbb{R}^3 , this basis consists of only two vectors. As we will see shortly, this reflects the fact that W is a plane in \mathbb{R}^3 , and hence “2-dimensional”. A similar calculation shows that

$$\beta = \left(\begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} \right)$$

is also a basis for W , as is

$$\gamma = \left(\begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix}, \begin{pmatrix} -2 \\ 1 \\ 1 \end{pmatrix} \right).$$

Amongst all of the bases for W , there does not appear to be a “best” or most “natural” one. If we want a basis, we have to simply choose one, and different bases may be easier to use for different tasks.

Finally, consider the vector

$$\vec{v} = \begin{pmatrix} 2 \\ 6 \\ -8 \end{pmatrix},$$

and notice that $\vec{v} \in W$. If we calculate the coordinates of V relative to these bases, we see that

$$[\vec{v}]_\alpha = \begin{pmatrix} -6 \\ 8 \end{pmatrix}, \quad [\vec{v}]_\beta = \begin{pmatrix} 6 \\ -8 \end{pmatrix}, \quad \text{and} \quad [\vec{v}]_\gamma = \begin{pmatrix} 7 \\ -1 \end{pmatrix}.$$

Intuitively, each of these three bases of W are putting different 2-dimensional grid systems onto the plane W , and we are measuring the point \vec{v} (which lives in \mathbb{R}^3 , but is also on W) using these 2-dimensional grids. However, since there is no “best” basis for W , there is no “best” way to assign 2 numbers to locate this point \vec{v} on the plane W .

Suppose that we have a vector space V along with vectors $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in V$ with $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$. If $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent, then it is a basis. But what if it is not linearly independent? Can we use the fact that we have half of the properties of a basis to build a basis from it? In order to answer this question, we first establish a condition for when we can drop a vector from a sequence without affecting the span.

Proposition 4.5.4. *Let V be a vector space and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w} \in V$. The following are equivalent:*

1. $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w}) = \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$.
2. $\vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$.

Proof. We first prove that (1) \Rightarrow (2). Assume then that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w}) = \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. Notice that we clearly have $\vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w})$ because

$$\vec{w} = 0 \cdot \vec{u}_1 + 0 \cdot \vec{u}_2 + \dots + 0 \cdot \vec{u}_n + 1 \cdot \vec{w}.$$

Since $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w}) = \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$, it follows that $\vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$.

We now prove that (2) \Rightarrow (1). Assume then that $\vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. We can then fix $d_1, d_2, \dots, d_n \in \mathbb{R}$ with

$$\vec{w} = d_1 \vec{u}_1 + d_2 \vec{u}_2 + \dots + d_n \vec{u}_n.$$

We now show that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w}) = \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ by giving a double containment proof.

- $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) \subseteq \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w})$: Let $\vec{v} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ be arbitrary. By definition, we can fix $c_1, c_2, \dots, c_n \in \mathbb{R}$ with

$$\vec{v} = c_1 \vec{u}_1 + c_2 \vec{u}_2 + \dots + c_n \vec{u}_n.$$

We then have

$$\vec{v} = c_1 \vec{u}_1 + c_2 \vec{u}_2 + \dots + c_n \vec{u}_n + 0 \cdot \vec{w},$$

so $\vec{v} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w})$.

- $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w}) \subseteq \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$: Let $\vec{v} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w})$ be arbitrary. By definition, we can fix $c_1, c_2, \dots, c_n, c_{n+1} \in \mathbb{R}$ with

$$\vec{v} = c_1 \vec{u}_1 + c_2 \vec{u}_2 + \dots + c_n \vec{u}_n + c_{n+1} \vec{w}.$$

Plugging our expression for \vec{w} into this equation, we have

$$\begin{aligned}\vec{v} &= c_1\vec{u}_1 + c_2\vec{u}_2 + \cdots + c_n\vec{u}_n + c_{n+1} \cdot (d_1\vec{u}_1 + d_2\vec{u}_2 + \cdots + d_n\vec{u}_n) \\ &= c_1\vec{u}_1 + c_2\vec{u}_2 + \cdots + c_n\vec{u}_n + c_{n+1}d_1\vec{u}_1 + c_{n+1}d_2\vec{u}_2 + \cdots + c_{n+1}d_n\vec{u}_n \\ &= c_1\vec{u}_1 + c_{n+1}d_1\vec{u}_1 + c_2\vec{u}_2 + c_{n+1}d_2\vec{u}_2 + \cdots + c_n\vec{u}_n + c_{n+1}d_n\vec{u}_n \\ &= (c_1 + c_{n+1}d_1) \cdot \vec{u}_1 + (c_2 + c_{n+1}d_2) \cdot \vec{u}_2 + \cdots + (c_n + c_{n+1}d_n) \cdot \vec{u}_n,\end{aligned}$$

so $\vec{v} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$.

We have shown both containments, so we can conclude that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w}) = \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. \square

With this result in hand, we now establish a fundamental fact. Given a vector space V , we can trim any spanning sequence of vectors down to a basis of V by dropping some of the vectors.

Proposition 4.5.5. *Let V be a vector space, and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in V$ be such that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$. There exists a basis of V that can be obtained by omitting some of the \vec{u}_i from the sequence (this includes the possibility of omitting either none or all of the \vec{u}_i).*

Proof. Let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in V$ be such that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$. First notice that if $V = \{\vec{0}\}$, then V has the basis $()$, which is obtained by omitting all of the \vec{u}_i . Thus, we can assume that $V \neq \{\vec{0}\}$. Notice then that we must have $n \geq 1$ (i.e. we do not have the empty sequence of vectors since $V \neq \{\vec{0}\}$).

If $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent, then it is a basis for V , and we are done. Suppose then that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly dependent, and also assume that $n \geq 2$ (otherwise, go to the next paragraph). Using Proposition 4.4.2, we then know that one of the \vec{u}_i is a linear combination of the others. Applying Proposition 4.5.4, we can omit this vector from our list and then have a sequence of $n - 1$ vectors amongst the \vec{u}_i that span V . Now if this sequence is linearly independent, then it is a basis and we are done. If not, and if we still have at least 2 vectors, then by Proposition 4.4.2 again, we know one of the remaining \vec{u}_i is a linear combination of the others, and by applying Proposition 4.5.4 we can omit this vector from our list and then have a sequence of $n - 2$ vectors amongst the \vec{u}_i that span V . Repeat this process until we either end up with a basis from omitting some of the \vec{u}_i , or we end with just one of the \vec{u}_i .

Suppose then that we end up in this latter case so that we have one $k \in \{1, 2, \dots, n\}$ with $\text{Span}(\vec{u}_k) = V$. We must have $\vec{u}_k \neq \vec{0}$ because $V \neq \{\vec{0}\}$. By Proposition 4.4.3, we know that (\vec{u}_k) is linearly independent, and hence is a basis for V . \square

Let's see how to use these results in practice to trim down a sequence to a basis. Let W be the following subspace of \mathbb{R}^5 :

$$W = \text{Span} \left(\begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix}, \begin{pmatrix} 3 \\ 6 \\ -3 \end{pmatrix}, \begin{pmatrix} -2 \\ -5 \\ 4 \end{pmatrix}, \begin{pmatrix} 0 \\ -2 \\ 4 \end{pmatrix}, \begin{pmatrix} 7 \\ 15 \\ -9 \end{pmatrix} \right).$$

Suppose that we want to find a basis of W . We start by forming a matrix with the given vectors as columns. Let

$$A = \begin{pmatrix} 1 & 3 & -2 & 0 & 7 \\ 2 & 6 & -5 & -2 & 15 \\ -1 & -3 & 4 & 4 & -9 \end{pmatrix}.$$

If we want to check linear independence, we first apply elementary row operations to A :

$$\begin{aligned}
 \begin{pmatrix} 1 & 3 & -2 & 0 & 7 \\ 2 & 6 & -5 & -2 & 15 \\ -1 & -3 & 4 & 4 & -9 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 3 & -2 & 0 & 7 \\ 0 & 0 & -1 & -2 & 1 \\ 0 & 0 & 2 & 4 & -2 \end{pmatrix} && \begin{array}{l} (-2R_1 + R_2) \\ (R_1 + R_3) \end{array} \\
 &\rightarrow \begin{pmatrix} 1 & 3 & -2 & 0 & 7 \\ 0 & 0 & 1 & 2 & -1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} && (2R_2 + R_3) \\
 &\rightarrow \begin{pmatrix} 1 & 3 & 0 & 4 & 5 \\ 0 & 0 & 1 & 2 & -1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} && (2R_2 + R_1).
 \end{aligned}$$

Notice that we went down to reduced echelon form here (see Definition 4.2.13) for some illustrative purposes later. Since we now have an echelon form of A , we see that there are columns without leading entries, and hence the vectors are linearly dependent (of course, we could also have realized this by noticing that we have 5 vectors in \mathbb{R}^3). Now we could solve a linear system to find a nontrivial linear combination of the columns giving $\vec{0}$, and then use that to write one of the vectors as a linear combination of the others. However, we can immediately see the necessary information from the above matrix. The key idea is to view the matrix as coding many linear systems, as we did in the example after Definition 4.2.13. Look at just the first two columns of the original matrix A , and view it as the augmented matrix of the following 3 equations in 1 variable:

$$\begin{array}{rcrcrcrcl}
 x & = & 3 \\
 2x & = & 6 \\
 -x & = & -3.
 \end{array}$$

From the first two columns of the final matrix, we realize that this system has a solution, namely $x = 3$. Thus, the second column of our original matrix is a multiple of the first, and hence we do not need it. At this point, we could delete the second column and do Gaussian elimination on a new matrix. However, there is no need to do that! Look at columns 1, 3, and 4 of the original matrix, and view them as the augmented matrix of the following 3 equations in 2 variables:

$$\begin{array}{rcrcrcrcrcrcl}
 x & - & 2y & = & 0 \\
 2x & - & 5y & = & -2 \\
 -x & + & 4y & = & 4.
 \end{array}$$

From columns 1, 3, and 4 of the final matrix, we immediately see that $(4, 2)$ is a solution, so the fourth vector is a linear combination of the first and third columns, namely:

$$\begin{pmatrix} 0 \\ -2 \\ 4 \end{pmatrix} = 4 \cdot \begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix} + 2 \cdot \begin{pmatrix} -2 \\ -5 \\ 4 \end{pmatrix}.$$

Therefore, we can omit the fourth vector. By looking at columns 1, 3, and 5, we realize that

$$\begin{pmatrix} 7 \\ 15 \\ -9 \end{pmatrix} = 5 \cdot \begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix} + (-1) \cdot \begin{pmatrix} -2 \\ -5 \\ 4 \end{pmatrix},$$

so we can omit the fifth vector. In other words, viewing the above matrix computations from different perspectives, we have shown that

$$W = \text{Span} \left(\begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix}, \begin{pmatrix} -2 \\ -5 \\ 4 \end{pmatrix} \right).$$

Finally, we can also conclude that these two vectors are linearly independent because the two corresponding columns in the above echelon form each have a leading entry. It follows that

$$\left(\begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix}, \begin{pmatrix} -2 \\ -5 \\ 4 \end{pmatrix} \right)$$

is a basis for W . Generalizing this computation, we arrive at the following result.

Proposition 4.5.6. *Let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$, and let $W = \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. Let A be the $m \times n$ matrix whose i^{th} column is \vec{u}_i , and let B be an echelon form of A . If we build the sequence consisting only of those \vec{u}_i such that the i^{th} column of B has a leading entry, then we obtain a basis for W .*

In general, given an $m \times n$ matrix, we always obtain a subspace of \mathbb{R}^m by taking the span of the columns. We give this subspace a special name.

Definition 4.5.7. *Let A be an $m \times n$ matrix. We define $\text{Col}(A)$ to be the span of the columns of A , and call $\text{Col}(A)$ the column space of A . Notice that $\text{Col}(A)$ is a subspace of \mathbb{R}^m .*

Thus, by Proposition 4.5.6, to find a basis for $\text{Col}(A)$, we find an echelon form B of A , and then take the columns of A corresponding to those columns of B that have a leading entry.

4.6 Dimension

Despite the fact that a vector space may have many different bases and may not have one “natural” basis, is there anything that the bases of a vector space have in common? Although there are many bases of \mathbb{R}^m , we can combine Corollary 4.3.2 and Corollary 4.4.6 to conclude that every basis of \mathbb{R}^m has exactly m elements. It turns out that this fact holds more generally. In other words, any two bases of the same vector space have the same number of elements.

Proposition 4.5.5 gives a characterization for when we can drop a vector from a list without affecting the span. The next question we want to ask is when we can *exchange* a vector in a list with another without affecting the span. To get some intuition here, suppose that we have a vector space V along with $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w} \in V$, and suppose that we are considering the set $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. We want to know when we can swap out one of the vectors \vec{u}_k for the vector \vec{w} without affecting the span. Certainly we will want \vec{w} to be an element of $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ in order for this to be the case (since \vec{w} will definitely be in the span after we do an exchange). Thus, we will be able to fix $c_1, \dots, c_k, \dots, c_n \in \mathbb{R}$ with

$$\vec{w} = c_1 \vec{u}_1 + \dots + c_k \vec{u}_k + \dots + c_n \vec{u}_n.$$

In this case, will we always have

$$\text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{u}_k, \vec{u}_{k+1}, \dots, \vec{u}_n) = \text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{w}, \vec{u}_{k+1}, \dots, \vec{u}_n)?$$

To get a handle on this question, let's take a look at an example. Suppose that $V = \mathbb{R}^3$, and we have

$$\vec{u}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \quad \vec{u}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad \text{and} \quad \vec{u}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

Notice that $\text{Span}(\vec{u}_1, \vec{u}_2, \vec{u}_3) = \mathbb{R}^3$. Let

$$\vec{w} = \begin{pmatrix} 3 \\ 5 \\ 0 \end{pmatrix}.$$

Can we exchange \vec{w} for the various \vec{u}_k without affecting the span? Notice that

$$\text{Span}(\vec{u}_1, \vec{u}_2, \vec{w}) = \text{Span} \left(\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 3 \\ 5 \\ 0 \end{pmatrix} \right),$$

and so $\text{Span}(\vec{u}_1, \vec{u}_2, \vec{w}) \neq \mathbb{R}^3$ because we can never achieve a nonzero entry in the last coordinate. Thus, we cannot always exchange vectors without affecting the span. In contrast, it does turn out that

$$\text{Span}(\vec{w}, \vec{u}_2, \vec{u}_3) = \text{Span} \left(\begin{pmatrix} 3 \\ 5 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right)$$

and

$$\text{Span}(\vec{u}_1, \vec{w}, \vec{u}_3) = \text{Span} \left(\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 3 \\ 5 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right),$$

because the two spans on the right equal \mathbb{R}^3 , as one can check by performing elementary row operations on the corresponding matrices. To see what is going on here, let's express \vec{w} as a linear combination of $\vec{u}_1, \vec{u}_2, \vec{u}_3$ as follows:

$$\begin{pmatrix} 3 \\ 5 \\ 0 \end{pmatrix} = 3 \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + 5 \cdot \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} + 0 \cdot \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

Notice that the coefficients of \vec{u}_1 and \vec{u}_2 are both nonzero, while the coefficient of \vec{u}_3 is zero. We now seize on this pattern and prove the following general result saying that if the coefficient of \vec{u}_k is nonzero, then we can indeed do an exchange without affecting the span.

Proposition 4.6.1 (Steinitz Exchange). *Let V be a vector space and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w} \in V$. Suppose that we have $d_1, \dots, d_n \in \mathbb{R}$ with*

$$\vec{w} = d_1 \vec{u}_1 + \dots + d_k \vec{u}_k + \dots + d_n \vec{u}_n.$$

If $d_k \neq 0$, then

$$\text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{u}_k, \vec{u}_{k+1}, \dots, \vec{u}_n) = \text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{w}, \vec{u}_{k+1}, \dots, \vec{u}_n).$$

Proof. Assume that $d_k \neq 0$. Notice that we are given \vec{w} as a linear combination of $\vec{u}_1, \dots, \vec{u}_k, \dots, \vec{u}_n$, so we can use this to convert expressions involving \vec{w} into ones involving the \vec{u}_i . The key assumption that $d_k \neq 0$ allows us to reverse this and express \vec{u}_k in terms of \vec{w} and the other \vec{u}_i . More precisely, by subtracting all other terms from both sides of the given equality, we know that

$$\begin{aligned} d_k \vec{u}_k &= \vec{w} - d_1 \vec{u}_1 - \dots - d_{k-1} \vec{u}_{k-1} - d_{k+1} \vec{u}_{k+1} - \dots - d_n \vec{u}_n \\ &= (-d_1) \vec{u}_1 + \dots + (-d_{k-1}) \vec{u}_{k-1} + \vec{w} + (-d_{k+1}) \vec{u}_{k+1} + \dots + (-d_n) \vec{u}_n. \end{aligned}$$

Since $d_k \neq 0$, we can multiply both sides by $\frac{1}{d_k}$ and use the vector space axioms to conclude that

$$\vec{u}_k = \frac{-d_1}{d_k} \cdot \vec{u}_1 + \dots + \frac{-d_{k-1}}{d_k} \cdot \vec{u}_{k-1} + \frac{1}{d_k} \cdot \vec{w} + \frac{-d_{k+1}}{d_k} \cdot \vec{u}_{k+1} + \dots + \frac{-d_n}{d_k} \cdot \vec{u}_n.$$

With this work in hand, we now show that the two sets are equal by given a double containment proof.

- $\text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{w}, \vec{u}_{k+1}, \dots, \vec{u}_n) \subseteq \text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{u}_k, \vec{u}_{k+1}, \dots, \vec{u}_n)$: Let

$$\vec{v} \in \text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{w}, \vec{u}_{k+1}, \dots, \vec{u}_n)$$

be arbitrary. By definition, we can fix $c_1, \dots, c_n \in \mathbb{R}$ with

$$\vec{v} = c_1 \vec{u}_1 + \dots + c_{k-1} \vec{u}_{k-1} + c_k \vec{w} + c_{k+1} \vec{u}_{k+1} + \dots + c_n \vec{u}_n.$$

Plugging in our given expression for \vec{w} into this equation, we have

$$\begin{aligned} \vec{v} &= c_1 \vec{u}_1 + \dots + c_{k-1} \vec{u}_{k-1} + c_k \vec{w} + c_{k+1} \vec{u}_{k+1} + \dots + c_n \vec{u}_n \\ &= c_1 \vec{u}_1 + \dots + c_{k-1} \vec{u}_{k-1} + c_k \cdot (d_1 \vec{u}_1 + \dots + d_k \vec{u}_k + \dots + d_n \vec{u}_n) + c_{k+1} \vec{u}_{k+1} + \dots + c_n \vec{u}_n \\ &= (c_1 + c_k d_1) \cdot \vec{u}_1 + \dots + (c_{k-1} + c_k d_{k-1}) \cdot \vec{u}_{k-1} + c_k d_k \cdot \vec{u}_k \\ &\quad + (c_{k+1} + c_k d_{k+1}) \cdot \vec{u}_{k+1} + \dots + (c_n + c_k d_n) \cdot \vec{u}_n, \end{aligned}$$

so $\vec{v} \in \text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{u}_k, \vec{u}_{k+1}, \dots, \vec{u}_n)$.

- $\text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{u}_k, \vec{u}_{k+1}, \dots, \vec{u}_n) \subseteq \text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{w}, \vec{u}_{k+1}, \dots, \vec{u}_n)$: Let

$$\vec{v} \in \text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{u}_k, \vec{u}_{k+1}, \dots, \vec{u}_n)$$

be arbitrary. By definition, we can fix $c_1, \dots, c_n \in \mathbb{R}$ with

$$\vec{v} = c_1 \vec{u}_1 + \dots + c_{k-1} \vec{u}_{k-1} + c_k \vec{u}_k + c_{k+1} \vec{u}_{k+1} + \dots + c_n \vec{u}_n.$$

Plugging in our derived expression for \vec{u}_k from above into this equation, we have

$$\begin{aligned} \vec{v} &= c_1 \vec{u}_1 + \dots + c_{k-1} \vec{u}_{k-1} + c_k \vec{u}_k + c_{k+1} \vec{u}_{k+1} + \dots + c_n \vec{u}_n \\ &= c_1 \vec{u}_1 + \dots + c_{k-1} \vec{u}_{k-1} \\ &\quad + c_k \cdot \left(\frac{-d_1}{d_k} \cdot \vec{u}_1 + \dots + \frac{-d_{k-1}}{d_k} \cdot \vec{u}_{k-1} + \frac{1}{d_k} \cdot \vec{w} + \frac{-d_{k+1}}{d_k} \cdot \vec{u}_{k+1} + \dots + \frac{-d_n}{d_k} \cdot \vec{u}_n \right) \\ &\quad + c_{k+1} \vec{u}_{k+1} + \dots + c_n \vec{u}_n \\ &= \left(c_1 - \frac{c_k d_1}{d_k} \right) \cdot \vec{u}_1 + \dots + \left(c_{k-1} - \frac{c_k d_{k-1}}{d_k} \right) \cdot \vec{u}_{k-1} \\ &\quad + \frac{c_k}{d_k} \cdot \vec{w} + \left(c_{k+1} - \frac{c_k d_{k+1}}{d_k} \right) \cdot \vec{u}_{k+1} + \dots + \left(c_n - \frac{c_k d_n}{d_k} \right) \cdot \vec{u}_n, \end{aligned}$$

so $\vec{v} \in \text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{w}, \vec{u}_{k+1}, \dots, \vec{u}_n)$.

We have shown both containments, so we can conclude that

$$\text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{u}_k, \vec{u}_{k+1}, \dots, \vec{u}_n) = \text{Span}(\vec{u}_1, \dots, \vec{u}_{k-1}, \vec{w}, \vec{u}_{k+1}, \dots, \vec{u}_n).$$

□

We are now ready to prove one of the most foundational and important results in linear algebra.

Theorem 4.6.2. *Let V be a vector space and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w}_1, \vec{w}_2, \dots, \vec{w}_m \in V$. If $m > n$ and $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$, then $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ is linearly dependent.*

In other words, if we know that a vector space can be spanned by n vectors, then any sequence of strictly more than n vectors must be linearly dependent.

Proof. We start by looking at \vec{w}_1 . Since $\vec{w}_1 \in V$ and $V = \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$, we can fix $c_1, c_2, \dots, c_n \in \mathbb{R}$ with

$$\vec{w}_1 = c_1 \vec{u}_1 + c_2 \vec{u}_2 + \dots + c_n \vec{u}_n.$$

Now we have two cases:

- *Case 1:* Suppose that $c_i = 0$ for all $i \in \{1, 2, \dots, n\}$. We then have

$$\begin{aligned} \vec{w}_1 &= 0 \cdot \vec{u}_1 + 0 \cdot \vec{u}_2 + \dots + 0 \cdot \vec{u}_n \\ &= \vec{0} + \vec{0} + \dots + \vec{0} \\ &= \vec{0}. \end{aligned}$$

Therefore, we have

$$1 \cdot \vec{w}_1 + 0 \cdot \vec{w}_2 + \dots + 0 \cdot \vec{w}_m = \vec{0},$$

and hence $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ is linearly dependent. Thus, we have achieved our goal in this case.

- *Case 2:* Suppose then that at least one of c_1, c_2, \dots, c_n is nonzero. Since we can reorder the vectors $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$ without affecting the span, we can assume that $c_1 \neq 0$. Using Proposition 4.6.1, we can exchange \vec{w}_1 and \vec{u}_1 to conclude that

$$\text{Span}(\vec{w}_1, \vec{u}_2, \vec{u}_3, \dots, \vec{u}_n) = \text{Span}(\vec{u}_1, \vec{u}_2, \vec{u}_3, \dots, \vec{u}_n),$$

and hence that $\text{Span}(\vec{w}_1, \vec{u}_2, \vec{u}_3, \dots, \vec{u}_n) = V$.

We now continue on to examine \vec{w}_2 . Since $\vec{w}_2 \in V$ and $V = \text{Span}(\vec{w}_1, \vec{u}_2, \vec{u}_3, \dots, \vec{u}_n)$, we can fix $d_1, d_2, d_3, \dots, d_n \in \mathbb{R}$ with

$$\vec{w}_2 = d_1 \vec{w}_1 + d_2 \vec{u}_2 + d_3 \vec{u}_3 + \dots + d_n \vec{u}_n.$$

Now we again have two cases:

- Suppose that $d_i = 0$ for all $i \in \{2, 3, \dots, n\}$. We then have

$$\begin{aligned} \vec{w}_2 &= d_1 \cdot \vec{w}_1 + 0 \cdot \vec{u}_2 + 0 \cdot \vec{u}_3 + \dots + 0 \cdot \vec{u}_n \\ &= d_1 \cdot \vec{w}_1 + \vec{0} + \vec{0} + \dots + \vec{0} \\ &= d_1 \cdot \vec{w}_1. \end{aligned}$$

Therefore, we have that $\vec{w}_2 \in \text{Span}(\vec{w}_1)$, and hence $\vec{w}_2 \in \text{Span}(\vec{w}_1, \vec{w}_3, \vec{w}_4, \dots, \vec{w}_m)$. Since one of the \vec{w}_i is a linear combination of the others, we may use Proposition 4.4.2 to conclude $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ is linearly dependent. Thus, we have achieved our goal in this case.

- Suppose that at least one of d_2, d_3, \dots, d_n is nonzero. Since we can reorder the vectors $\vec{u}_2, \vec{u}_3, \dots, \vec{u}_n$ without affecting the span, we can assume that $d_2 \neq 0$. Using Proposition 4.6.1, we can exchange \vec{w}_2 and \vec{u}_2 to conclude that

$$\text{Span}(\vec{w}_1, \vec{w}_2, \vec{u}_3, \dots, \vec{u}_n) = \text{Span}(\vec{u}_1, \vec{u}_2, \vec{u}_3, \dots, \vec{u}_n),$$

and hence that $\text{Span}(\vec{w}_1, \vec{w}_2, \vec{u}_3, \dots, \vec{u}_n) = V$.

We now continue this process down the line of the \vec{w}_i . Along the way, we might end up in one of the first cases and conclude that $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ is linearly dependent, at which point we've reached our goal, and the argument stops. Suppose then that we always are in the second case, and so end up exchanging each of the first n vectors in the list $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_n, \dots, \vec{w}_m)$ for the n vectors in the list $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. At this point, we will have concluded that $\text{Span}(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_n) = V$. Now we are assuming that $m > n$, so $m \geq n + 1$ and hence we have the vector \vec{w}_{n+1} in our original list $\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m$. Since $\vec{w}_{n+1} \in V$, we then have that $\vec{w}_{n+1} \in \text{Span}(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_n)$. Using Proposition 4.4.2, it follows that $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ is linearly dependent, and hence we have reached our goal in this case as well. \square

Taking the contrapositive of the theorem, we arrive at the following result.

Corollary 4.6.3. *Let V be a vector space and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w}_1, \vec{w}_2, \dots, \vec{w}_m \in V$ be such that both $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$ and $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ is linearly independent. We then have $m \leq n$.*

In other words, given any vector space V , the number of vectors in any linearly independent sequence is less than or equal to the number of vectors in any spanning sequence. With this result in hand, we can easily establish the fact alluded to above that any two bases of a vector space must have the same number of elements.

Corollary 4.6.4. *Suppose that V is a vector space and both $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ and $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ are bases for V . We then have that $m = n$.*

Proof. Since $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$ and $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ is linearly independent, we may use Corollary 4.6.3 to conclude that $m \leq n$. Similarly, since $\text{Span}(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m) = V$ and $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent, we may use Corollary 4.6.3 to conclude that $n \leq m$. Combining these, it follows that $m = n$. \square

Notice that we have not yet examined the question of whether there always exists a basis, but we have shown that when bases exist, then they all must have the same number of elements. With that in mind, we can define the following fundamental concept.

Definition 4.6.5. *Let V be a vector space, and assume that V has at least one basis. We define $\dim(V)$ to be the number of elements in any basis of V . We call $\dim(V)$ the dimension of V .*

Let's take a moment to figure out the dimension of various natural vector spaces. Since all bases will have the same number of elements, to determine the dimension of a vector space V , we need only exhibit *one* basis of V and count the number of elements. The first three examples follow immediately from the examples we gave after defining a basis.

- For each $n \in \mathbb{N}^+$, we have $\dim(\mathbb{R}^n) = n$. To see this, we need only give one example of a basis of \mathbb{R}^n . Given any $n \in \mathbb{N}^+$, let $\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n$ be the vectors where \vec{e}_i has a 1 in position i and zeros elsewhere. As mentioned above, we have that $(\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n)$ is a basis for \mathbb{R}^n , so $\dim(\mathbb{R}^n) = n$. Notice in particular that $\dim(\mathbb{R}) = 1$ because (1) is a basis for \mathbb{R} .
- For each $n \in \mathbb{N}$, we have $\dim(\mathcal{P}_n) = n + 1$. To see this, we need only give one example of a basis of \mathcal{P}_n . For each $k \in \mathbb{N}$ with $0 \leq k \leq n$, let $f_k: \mathbb{R} \rightarrow \mathbb{R}$ be the function $f_k(x) = x^k$. Thus, $f_0(x) = 1$, $f_1(x) = x$, $f_2(x) = x^2$, \dots , and $f_n(x) = x^n$. As mentioned above, we have that $(f_0, f_1, f_2, \dots, f_n)$ is a basis for \mathcal{P}_n , so $\dim(\mathcal{P}_n) = n + 1$.
- If we let V be the vector space of all 2×2 matrices, then

$$\left(\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \right)$$

is a basis for V , so $\dim(V) = 4$.

- If we let W be the subspace of \mathbb{R}^3 given by

$$W = \left\{ \begin{pmatrix} a \\ b \\ c \end{pmatrix} : a + b + c = 0 \right\},$$

then we showed above that

$$\left(\begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} \right)$$

is a basis of W . Therefore, $\dim(W) = 2$. Notice that this is the case even though W lives “inside” of the 3-dimensional space \mathbb{R}^3 .

- If $V = \{\vec{0}\}$, then $()$ is a basis for V by Definition 4.1.17 and Definition 4.4.3, so $\dim(V) = 0$.

Let’s return to the question of whether bases always exist. It turns out that there are vector spaces that do not have bases, and in fact we have seen several already. For example, consider the vector space \mathcal{P} of all polynomial functions (without restriction on the degree). We claim that \mathcal{P} has no finite spanning sequence, and hence has no basis. To see this, let’s first consider an example. Suppose that we have the following elements of \mathcal{P} :

- $f_1(x) = x^5 + 3x^2 - 2x + 7$.
- $f_2(x) = x^7 - 12x^6 + 9x^4 - x^3 - 1$.
- $f_3(x) = x^{42} + 2x$.

Can we show that $\text{Span}(f_1, f_2, f_3) \neq \mathcal{P}$ by constructing an element of \mathcal{P} that is not in $\text{Span}(f_1, f_2, f_3)$? We can surely do this, because every element of $\text{Span}(f_1, f_2, f_3)$ is a linear combination of these three polynomials, so has degree at most 42. Thus, if $g(x) = x^{43}$, then $g \in \mathcal{P}$ but $g \notin \text{Span}(f_1, f_2, f_3)$. The idea of this argument can be applied generally. Given any $f_1, f_2, \dots, f_n \in \mathcal{P}$, there is a maximum degree of the polynomials f_1, f_2, \dots, f_n , and if we let k be 1 more than this maximum degree and define $g(x) = x^k$, then $g \in \mathcal{P}$ but $g \notin \text{Span}(f_1, f_2, \dots, f_n)$. Thus, for all $f_1, f_2, \dots, f_n \in \mathcal{P}$ (no matter what $n \in \mathbb{N}^+$ is), we always have $\text{Span}(f_1, f_2, \dots, f_n) \neq \mathcal{P}$. It follows that \mathcal{P} has no finite spanning sequence, and hence no basis.

We can use Proposition 4.5.5, which says that any spanning sequence can be trimmed down to a basis, to give a weaker condition for when a vector space V to have a basis.

Definition 4.6.6. Let V be a vector space. We say that V is finite-dimensional if there exists a basis of V , or equivalently if there exists a finite sequence that spans V . Otherwise, we say that V is infinite-dimensional and write $\dim(V) = \infty$.

Therefore, \mathcal{P} is infinite-dimensional, and we have $\dim(\mathcal{P}) = \infty$.

We now return to a discussion that we started at the beginning of Section 4.5, where we notice that if $\text{Span}(\vec{u}_1, \vec{u}_2) = \mathbb{R}^2$, then (\vec{u}_1, \vec{u}_2) is automatically a basis of \mathbb{R}^2 . In general, if we have half of the properties of being a basis (either spanning or linear independence), and we happen to have a number of vectors that matches the dimension of the vector space, then we necessarily have a basis.

Proposition 4.6.7. Let V be a finite-dimensional vector space and let $n = \dim(V)$. Let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in V$ (notice the same n).

1. If $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent, then $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is a basis of V .
2. If $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$, then $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is a basis of V .

Proof. Suppose that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent. To show that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is a basis of V , we need only show that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$. We have $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) \subseteq V$ trivially, so we just need to show that $V \subseteq \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. Let $\vec{v} \in V$ be arbitrary. Since $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{v})$ is a sequence of $n + 1$ vectors and $\dim(V) = n$, we know from Theorem 4.6.2 that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{v})$ is linearly dependent. Therefore, $\vec{v} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ by Proposition 4.6.9. Since $\vec{v} \in V$ was arbitrary, it follows that $V \subseteq \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$.

Suppose now that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$. By Proposition 4.5.5, we know that we can omit some of the \vec{u}_i to obtain a basis of V . Since every basis of V has n elements, we must not omit any of the \vec{u}_i , and hence $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is a basis of V . \square

Now that we have built up all of this theory, we can characterize subspaces of vector spaces. Let's first think about some simple examples. It is straightforward to show directly that the only subspaces of \mathbb{R} are $\{0\}$ and \mathbb{R} . In \mathbb{R}^2 , the subspaces that we know are $\{0\}$, \mathbb{R}^2 , and the lines through the origin, each of which can be written as $\text{Span}(\vec{u})$ for a nonzero $\vec{u} \in \mathbb{R}^2$. In general, using Proposition 4.1.16, we can always build subspaces of a vector space V by simply taking the span of a sequence of vectors. With this perspective, we can see that the following are subspaces of \mathbb{R}^3 :

- $\{\vec{0}\}$, which is $\text{Span}()$.
- $\text{Span}(\vec{u})$ for any $\vec{u} \in \mathbb{R}^3$ (if $\vec{u} = 0$, then this gives the previous case).
- $\text{Span}(\vec{u}_1, \vec{u}_2)$ for any $\vec{u}_1, \vec{u}_2 \in \mathbb{R}^3$ (in some cases, this might collapse to a previous case).
- \mathbb{R}^3 , which is $\text{Span}(\vec{e}_1, \vec{e}_2, \vec{e}_3)$.

Of course, we can take the span of more than 2 vectors, but in general it appears that we can always write these subspaces in one of the above forms. We also know of some other ways to define subspaces of \mathbb{R}^3 without directly taking a span. For example, we mentioned above that the set

$$W = \left\{ \begin{pmatrix} a \\ b \\ c \end{pmatrix} \in \mathbb{R}^3 : a + b + c = 0 \right\}$$

is a subspace of \mathbb{R}^3 . However, it is also possible to write it as a span as follows:

$$W = \left(\begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} \right).$$

With these ideas in hand, it is natural to ask whether *every* subspace of \mathbb{R}^3 can be described as a span of at most 3 vectors. In fact, we can prove a general version of this idea for any finite-dimensional vector space.

Proposition 4.6.8. *Let V be a finite-dimensional vector space with $\dim(V) = n$, and let W be a subspace of V . There exists $k \in \{0, 1, 2, \dots, n\}$ and $\vec{w}_1, \vec{w}_2, \dots, \vec{w}_k \in W$ with $W = \text{Span}(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_k)$ and such that $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_k)$ is linearly independent. In particular, W is a finite-dimensional vector space and $\dim(W) \leq \dim(V)$.*

Proof. First notice that if $W = \{\vec{0}\}$, then $W = \text{Span}()$, and we are done trivially. Suppose then that $W \neq \{\vec{0}\}$ and fix $\vec{w}_1 \in W$ with $\vec{w}_1 \neq \vec{0}$. Since $\vec{w}_1 \neq \vec{0}$, we know that (\vec{w}_1) is linearly independent by Proposition 4.4.3. Furthermore, as $\vec{w}_1 \in W$ and W is a subspace of V , we know that $\text{Span}(\vec{w}_1) \subseteq W$. Now if $\text{Span}(\vec{w}_1) = W$, then we are done (see the next paragraph). If not, then we can fix $\vec{w}_2 \in W$ with $\vec{w}_2 \notin \text{Span}(\vec{w}_1)$. Using Proposition 4.6.9, we then have that (\vec{w}_1, \vec{w}_2) is linearly independent. Since $\vec{w}_1, \vec{w}_2 \in W$ and W is a subspace of V , we know that $\text{Span}(\vec{w}_1, \vec{w}_2) \subseteq W$. Again, if $\text{Span}(\vec{w}_1, \vec{w}_2) = W$, then we are done (see the next paragraph). Otherwise, we continue this process.

Notice that we cannot continue this process to stage $n + 1$, because in that case we would have a linearly independent sequence of $n + 1$ vectors in V , which is impossible by Theorem 4.6.2 because $\dim(V) = n$. Therefore, we must stop by stage n , at which point we have found a $k \in \{0, 1, 2, \dots, n\}$ and $\vec{w}_1, \vec{w}_2, \dots, \vec{w}_k \in W$ with $W = \text{Span}(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_k)$ and such that $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_k)$ is linearly independent. This proves the first part of the result. Finally, notice in this case that $(\vec{w}_1, \vec{w}_2, \dots, \vec{w}_k)$ is a basis for W , so W is a finite-dimensional vector space and $\dim(W) \leq \dim(V)$. \square

In Proposition 4.5.5, we showed that every spanning sequence can be trimmed down to a basis. We now go in the other direction and show that every linearly independent sequence can be expanded to a basis, as long as the vector space that we are working in is finite-dimensional. We start with the following result.

Proposition 4.6.9. *Let V be a vector space, let $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ be a linearly independent sequence of vectors in V , and let $\vec{w} \in V$. The following are equivalent:*

1. $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w})$ is linearly independent.
2. $\vec{w} \notin \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$.

Proof. We prove each direction.

- (1) \Rightarrow (2): We prove the contrapositive **Not**(2) \Rightarrow **Not**(1). Suppose that $\vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. Applying Proposition 4.4.2, we immediately conclude that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w})$ is linearly dependent.
- (2) \Rightarrow (1): We prove the contrapositive **Not**(1) \Rightarrow **Not**(2). Suppose that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n, \vec{w})$ is linearly dependent. By definition, we can fix $c_1, c_2, \dots, c_n, c_{n+1} \in \mathbb{R}$ with

$$c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n + c_{n+1}\vec{w} = \vec{0},$$

and such that $c_i \neq 0$ for at least one $i \in \{1, 2, \dots, n, n+1\}$. Now if $c_{n+1} = 0$, then we would have $c_{n+1}\vec{w} = \vec{0}$, which would imply

$$c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n = \vec{0},$$

where $c_i \neq 0$ for at least one $i \in \{1, 2, \dots, n\}$, contradicting the fact that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is linearly independent. Therefore, we must have $c_{n+1} \neq 0$. Subtracting the first n terms of the equality

$$c_1\vec{u}_1 + c_2\vec{u}_2 + \dots + c_n\vec{u}_n + c_{n+1}\vec{w} = \vec{0}$$

from both sides, we obtain

$$c_{n+1}\vec{w} = (-c_1) \cdot \vec{u}_1 + (-c_2) \cdot \vec{u}_2 + \dots + (-c_n) \cdot \vec{u}_n.$$

Multiplying both sides by $\frac{1}{c_{n+1}}$ and using the vector space axioms, we see that

$$\vec{w} = \frac{-c_1}{c_{n+1}} \cdot \vec{u}_1 + \frac{-c_2}{c_{n+1}} \cdot \vec{u}_2 + \dots + \frac{-c_n}{c_{n+1}} \cdot \vec{u}_n.$$

Therefore, $\vec{w} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. □

With this proposition in hand, we can prove that we can indeed extend any linearly independent sequence to a basis (again assuming that the vector space in question has at least one basis).

Proposition 4.6.10. *Let V be a finite-dimensional vector space, and let $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k)$ be a linearly independent sequence of vectors in V . There exists $\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m \in V$ (possibly with $m = 0$) such that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k, \vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ is a basis of V .*

Proof. Start with the sequence $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k)$. Now if $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k) = V$, then $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k)$ is a basis of V and we are done. Suppose then that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k) \neq V$, and fix a vector $\vec{w}_1 \in V$ with $\vec{w}_1 \notin \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k)$. Using Proposition 4.6.9, we then have that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k, \vec{w}_1)$ is linearly independent. If $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k, \vec{w}_1) = V$, then we have a basis and we are done. Otherwise, fix a vector $\vec{w}_2 \in V$ with $\vec{w}_2 \notin \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k, \vec{w}_1)$. Using Proposition 4.6.9, we then have that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_k, \vec{w}_1, \vec{w}_2)$ is linearly independent. Continue in this way, stopping at any point where we have reached a basis. Notice that we must eventually stop, because otherwise we would eventually reach a point where we have a sequence of $\dim(V) + 1$ many vectors, which we know is not possible by Theorem 4.6.2. □

Chapter 5

Linear Transformations and Matrices

5.1 Linear Transformations

We have discussed linear transformations from \mathbb{R}^2 to \mathbb{R}^2 , but now that we have a much wider collection of worlds where we can add and scalar multiply, we can now generalize the idea of a linear transformation. In fact, we can even consider the situation where the domain and codomain are potentially different vector spaces.

Definition 5.1.1. *Let V and W be vector spaces. A linear transformation from V to W is a function $T: V \rightarrow W$ with the following two properties:*

1. $T(\vec{v}_1 + \vec{v}_2) = T(\vec{v}_1) + T(\vec{v}_2)$ for all $\vec{v}_1, \vec{v}_2 \in V$ (i.e. T preserves addition).
2. $T(c \cdot \vec{v}) = c \cdot T(\vec{v})$ for all $\vec{v} \in V$ and $c \in \mathbb{R}$ (i.e. T preserves scalar multiplication).

Although we have $+$ and \cdot on both sides of the above equations, these symbols are interpreted differently on each side. If $\vec{v}_1, \vec{v}_2 \in V$, then the $+$ in $T(\vec{v}_1 + \vec{v}_2)$ means the $+$ in V (because \vec{v}_1 and \vec{v}_2 are elements of V), while the $+$ in $T(\vec{v}_1) + T(\vec{v}_2)$ means the $+$ of W (because $T(\vec{v}_1)$ and $T(\vec{v}_2)$ will be elements of W). Similarly, the \cdot in the expression $T(c \cdot \vec{v})$ is the scalar multiplication in V , while the \cdot in $c \cdot T(\vec{v})$ is the scalar multiplication in W .

Let's see some examples of linear transformations. Back when we were talking about linear transformations from \mathbb{R}^2 to \mathbb{R}^2 , the simplest way to build a linear transformation was to take $a, b, c, d \in \mathbb{R}$ and define $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix}.$$

In each such case, the resulting function T is a linear transformation from \mathbb{R}^2 to \mathbb{R}^2 as we showed in Proposition 2.4.3. We can naturally generalize these ideas to \mathbb{R}^n for values of $n \neq 2$. For example, the function $T: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ given by

$$T\left(\begin{pmatrix} x \\ y \\ z \end{pmatrix}\right) = \begin{pmatrix} 4x + 2y - z \\ x + 8y - 2z \\ 2x + y - 17z \end{pmatrix}$$

is a linear transformation. We can also create linear transformations like this from \mathbb{R}^n to \mathbb{R}^m where $m \neq n$. For example, the function $T: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ given by

$$T\left(\begin{pmatrix} x \\ y \\ z \end{pmatrix}\right) = \begin{pmatrix} 3x - 8y + 3z \\ 2x + z \end{pmatrix}$$

is a linear transformation. We now prove that all functions defined using constant coefficients are indeed linear transformations.

Proposition 5.1.2. *Let $m, n \in \mathbb{N}^+$ and suppose that we have numbers $a_{i,j} \in \mathbb{R}$ for all $i, j \in \mathbb{N}$ with both $1 \leq i \leq m$ and $1 \leq j \leq n$. Define $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ by*

$$T \left(\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \right) = \begin{pmatrix} a_{1,1}x_1 + a_{1,2}x_2 + \cdots + a_{1,n}x_n \\ a_{2,1}x_1 + a_{2,2}x_2 + \cdots + a_{2,n}x_n \\ \vdots \\ a_{m,1}x_1 + a_{m,2}x_2 + \cdots + a_{m,n}x_n \end{pmatrix}.$$

We then have that T is linear transformation.

Proof. We first check that T preserves addition. Let $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^n$ be arbitrary. Fix $x_1, x_2, \dots, x_n \in \mathbb{R}$ and $y_1, y_2, \dots, y_n \in \mathbb{R}$ with

$$\vec{v}_1 = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \quad \text{and} \quad \vec{v}_2 = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix}.$$

We have

$$\begin{aligned} T(\vec{v}_1 + \vec{v}_2) &= T \left(\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} + \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} \right) \\ &= T \left(\begin{pmatrix} x_1 + y_1 \\ x_2 + y_2 \\ \vdots \\ x_n + y_n \end{pmatrix} \right) \\ &= \begin{pmatrix} a_{1,1}(x_1 + y_1) + a_{1,2}(x_2 + y_2) + \cdots + a_{1,n}(x_n + y_n) \\ a_{2,1}(x_1 + y_1) + a_{2,2}(x_2 + y_2) + \cdots + a_{2,n}(x_n + y_n) \\ \vdots \\ a_{m,1}(x_1 + y_1) + a_{m,2}(x_2 + y_2) + \cdots + a_{m,n}(x_n + y_n) \end{pmatrix} \\ &= \begin{pmatrix} a_{1,1}x_1 + a_{1,1}y_1 + a_{1,2}x_2 + a_{1,2}y_2 + \cdots + a_{1,n}x_n + a_{1,n}y_n \\ a_{2,1}x_1 + a_{2,1}y_1 + a_{2,2}x_2 + a_{2,2}y_2 + \cdots + a_{2,n}x_n + a_{2,n}y_n \\ \vdots \\ a_{m,1}x_1 + a_{m,1}y_1 + a_{m,2}x_2 + a_{m,2}y_2 + \cdots + a_{m,n}x_n + a_{m,n}y_n \end{pmatrix} \\ &= \begin{pmatrix} a_{1,1}x_1 + a_{1,2}x_2 + \cdots + a_{1,n}x_n \\ a_{2,1}x_1 + a_{2,2}x_2 + \cdots + a_{2,n}x_n \\ \vdots \\ a_{m,1}x_1 + a_{m,2}x_2 + \cdots + a_{m,n}x_n \end{pmatrix} + \begin{pmatrix} a_{1,1}y_1 + a_{1,2}y_2 + \cdots + a_{1,n}y_n \\ a_{2,1}y_1 + a_{2,2}y_2 + \cdots + a_{2,n}y_n \\ \vdots \\ a_{m,1}y_1 + a_{m,2}y_2 + \cdots + a_{m,n}y_n \end{pmatrix} \\ &= T \left(\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \right) + T \left(\begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} \right) \\ &= T(\vec{v}_1) + T(\vec{v}_2). \end{aligned}$$

Since $\vec{v}_1, \vec{v}_2 \in \mathbb{R}^n$ were arbitrary, it follows that T preserves addition.

We now check that T preserves scalar multiplication. Let $\vec{v} \in \mathbb{R}^n$ and $c \in \mathbb{R}$ be arbitrary. Fix $x_1, x_2, \dots, x_n \in \mathbb{R}$ with

$$\vec{v} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix}.$$

We have

$$\begin{aligned} T(c \cdot \vec{v}) &= T \left(c \cdot \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \right) \\ &= T \left(\begin{pmatrix} cx_1 \\ cx_2 \\ \vdots \\ cx_n \end{pmatrix} \right) \\ &= \begin{pmatrix} a_{1,1}cx_1 + a_{1,2}cx_2 + \cdots + a_{1,n}cx_n \\ a_{2,1}cx_1 + a_{2,2}cx_2 + \cdots + a_{2,n}cx_n \\ \vdots \\ a_{m,1}cx_1 + a_{m,2}cx_2 + \cdots + a_{m,n}cx_n \end{pmatrix} \\ &= c \cdot \begin{pmatrix} a_{1,1}x_1 + a_{1,2}x_2 + \cdots + a_{1,n}x_n \\ a_{2,1}x_1 + a_{2,2}x_2 + \cdots + a_{2,n}x_n \\ \vdots \\ a_{m,1}x_1 + a_{m,2}x_2 + \cdots + a_{m,n}x_n \end{pmatrix} \\ &= c \cdot T \left(\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \right) \\ &= c \cdot T(\vec{v}). \end{aligned}$$

Since $\vec{v} \in \mathbb{R}^n$ and $c \in \mathbb{R}$ were arbitrary, it follows that T preserves scalar multiplication.

We've shown that T preserves both addition and scalar multiplication, so T is a linear transformation. \square

Although this proposition provides us with a simple way to build linear transformations from \mathbb{R}^n to \mathbb{R}^m , there also exist much more interesting linear transformations between more exotic vector spaces. Let \mathcal{F} be the vector space of all functions $f: \mathbb{R} \rightarrow \mathbb{R}$, let \mathcal{D} be the subspace of \mathcal{F} consisting of all differentiable functions, and let \mathcal{C} be the subspace of \mathcal{F} consisting of all continuous functions. Now although elements of each of these vector spaces are *functions*, we can still use these elements as inputs to other functions. For example, consider the function $T: \mathcal{D} \rightarrow \mathcal{F}$ defined by letting $T(f) = f'$. That is, given an $f \in \mathcal{D}$, we output the function f' which is the derivative of f . Although inputs and outputs of this function are themselves functions, this description is a perfectly well-defined function from \mathcal{D} to \mathcal{F} . In fact, the function T is a linear

transformation. To see this, we use results from Calculus. For any $f, g \in \mathcal{D}$, we have

$$\begin{aligned} T(f + g) &= (f + g)' \\ &= f' + g' \\ &= T(f) + T(g), \end{aligned}$$

so T preserves addition. Also, for any $f \in \mathcal{D}$ and $c \in \mathbb{R}$, we have

$$\begin{aligned} T(c \cdot f) &= (c \cdot f)' \\ &= c \cdot f' \\ &= c \cdot T(f), \end{aligned}$$

so T preserves scalar multiplication. Therefore, T is a linear transformation.

We can also have a linear transformation from a vector space whose elements are functions to a vector space whose elements are numbers. For example, consider the function $T: \mathcal{C} \rightarrow \mathbb{R}$ given by

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

We now check that T is a linear transformation using results from Calculus. For any $f, g \in \mathcal{C}$, we have

$$\begin{aligned} T(f + g) &= \int_0^1 (f + g)(x) \, dx \\ &= \int_0^1 (f(x) + g(x)) \, dx \\ &= \int_0^1 f(x) \, dx + \int_0^1 g(x) \, dx \\ &= T(f) + T(g), \end{aligned}$$

so T preserves addition. Also, for any $f \in \mathcal{C}$ and $c \in \mathbb{R}$, we have

$$\begin{aligned} T(c \cdot f) &= \int_0^1 (c \cdot f)(x) \, dx \\ &= \int_0^1 c \cdot f(x) \, dx \\ &= c \cdot \int_0^1 f(x) \, dx \\ &= c \cdot T(f), \end{aligned}$$

so T preserves scalar multiplication. Therefore, T is a linear transformation.

Another important family of linear transformation consists of the functions Coord_α for different choices of bases α . Suppose that we have a vector space V whose objects might be mysterious. Suppose also that we have a basis $\alpha = (\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ of V . Notice that the function Coord_α takes as input an element of V and produces as output an element of \mathbb{R}^n . In other words, this function takes elements of V , which might be functions or even more exotic objects, and turns them into a tidy list of numbers. We now check that these functions are always linear transformations. The proof is a direct generalization of Proof 2 of Proposition [3.1.3](#)

Proposition 5.1.3. *Let V be a vector space and let α be a basis of V with n elements. The function $\text{Coord}_\alpha: V \rightarrow \mathbb{R}^n$ is a linear transformation.*

Proof. We first check that Coord_α preserves addition. Let $\vec{v}_1, \vec{v}_2 \in V$ be arbitrary. Since α is a basis of V , we can apply Theorem 4.5.2 to fix the unique $c_1, c_2, \dots, c_n \in \mathbb{R}$ and the unique $d_1, d_2, \dots, d_n \in \mathbb{R}$ with

$$\vec{v}_1 = c_1 \vec{u}_1 + c_2 \vec{u}_2 + \cdots + c_n \vec{u}_n$$

and

$$\vec{v}_2 = d_1 \vec{u}_1 + d_2 \vec{u}_2 + \cdots + d_n \vec{u}_n.$$

By definition of Coord_α , we have

$$\text{Coord}_\alpha(\vec{v}_1) = \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix} \quad \text{and} \quad \text{Coord}_\alpha(\vec{v}_2) = \begin{pmatrix} d_1 \\ d_2 \\ \vdots \\ d_n \end{pmatrix}.$$

Notice that

$$\begin{aligned} \vec{v}_1 + \vec{v}_2 &= (c_1 \vec{u}_1 + c_2 \vec{u}_2 + \cdots + c_n \vec{u}_n) + (d_1 \vec{u}_1 + d_2 \vec{u}_2 + \cdots + d_n \vec{u}_n) \\ &= c_1 \vec{u}_1 + d_1 \vec{u}_1 + c_2 \vec{u}_2 + d_2 \vec{u}_2 + \cdots + c_n \vec{u}_n + d_n \vec{u}_n \\ &= (c_1 + d_1) \vec{u}_1 + (c_2 + d_2) \vec{u}_2 + \cdots + (c_n + d_n) \vec{u}_n. \end{aligned}$$

Since $c_i + d_i \in \mathbb{R}$ for all $i \in \{1, 2, \dots, n\}$, we have found the (necessarily unique by Theorem 4.5.2) sequence of numbers that express $\vec{v}_1 + \vec{v}_2$ as a linear combination of $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$. Therefore, we have

$$\begin{aligned} \text{Coord}_\alpha(\vec{v}_1 + \vec{v}_2) &= \begin{pmatrix} c_1 + d_1 \\ c_2 + d_2 \\ \vdots \\ c_n + d_n \end{pmatrix} \\ &= \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix} + \begin{pmatrix} d_1 \\ d_2 \\ \vdots \\ d_n \end{pmatrix} \\ &= \text{Coord}_\alpha(\vec{v}_1) + \text{Coord}_\alpha(\vec{v}_2). \end{aligned}$$

Since $\vec{v}_1, \vec{v}_2 \in V$ were arbitrary, it follows that Coord_α preserves addition.

We now check that Coord_α preserves scalar multiplication. Let $\vec{v} \in V$ and $c \in \mathbb{R}$ be arbitrary. Since α is a basis of V , we can apply Theorem 4.5.2 to fix the unique $d_1, d_2, \dots, d_n \in \mathbb{R}$ with

$$\vec{v} = d_1 \vec{u}_1 + d_2 \vec{u}_2 + \cdots + d_n \vec{u}_n.$$

By definition of Coord_α , we have

$$\text{Coord}_\alpha(\vec{v}) = \begin{pmatrix} d_1 \\ d_2 \\ \vdots \\ d_n \end{pmatrix}.$$

Notice that

$$\begin{aligned} c\vec{v} &= c \cdot (d_1\vec{u}_1 + d_2\vec{u}_2 + \cdots + d_n\vec{u}_n) \\ &= c \cdot (d_1\vec{u}_1) + c \cdot (d_2\vec{u}_2) + \cdots + c \cdot (d_n\vec{u}_n) \\ &= (cd_1)\vec{u}_1 + (cd_2)\vec{u}_2 + \cdots + (cd_n)\vec{u}_n. \end{aligned}$$

Since $cd_i \in \mathbb{R}$ for all $i \in \{1, 2, \dots, n\}$, we have found the (necessarily unique by Theorem 4.5.2) sequence of numbers that express $c\vec{v}$ as a linear combination of \vec{u}_1 and \vec{u}_2 . Therefore, we have

$$\begin{aligned} \text{Coord}_\alpha(c\vec{v}) &= \begin{pmatrix} cd_1 \\ cd_2 \\ \vdots \\ cd_n \end{pmatrix} \\ &= c \cdot \begin{pmatrix} d_1 \\ d_2 \\ \vdots \\ d_n \end{pmatrix} \\ &= c \cdot \text{Coord}_\alpha(\vec{v}). \end{aligned}$$

Since $\vec{v} \in V$ and $c \in \mathbb{R}$ were arbitrary, it follows that Coord_α preserves scalar multiplication.

We've shown that Coord_α preserves both addition and scalar multiplication, so Coord_α is a linear transformation. \square

Several of the other proofs that we carried out in the context of \mathbb{R}^2 also carry over with almost no change in the argument other than replacing \mathbb{R}^2 with V or W .

Proposition 5.1.4. *Let $T: V \rightarrow W$ be a linear transformation. We have the following:*

1. $T(\vec{0}_V) = \vec{0}_W$ (where $\vec{0}_V$ is the zero vector of the vector space V , and $\vec{0}_W$ is the zero vector of W).
2. $T(-\vec{v}) = -T(\vec{v})$ for all $\vec{v} \in V$.
3. $T(c_1\vec{v}_1 + c_2\vec{v}_2) = c_1 \cdot T(\vec{v}_1) + c_2 \cdot T(\vec{v}_2)$ for all $\vec{v}_1, \vec{v}_2 \in V$ and all $c_1, c_2 \in \mathbb{R}$.

Proof. See the proof of Proposition 2.4.2. \square

Theorem 5.1.5. *Let V and W be vector spaces, and suppose that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$. Suppose that $T: V \rightarrow W$ and $S: V \rightarrow W$ are linear transformations with the property that $T(\vec{u}_i) = S(\vec{u}_i)$ for all $i \in \{1, 2, \dots, n\}$. We then have that $T = S$, i.e. $T(\vec{v}) = S(\vec{v})$ for all $\vec{v} \in V$.*

Proof. See the proof of Theorem 2.4.4. \square

Theorem 5.1.6. *Let V and W be vector spaces, and suppose that $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ is a basis of V . Let $\vec{w}_1, \vec{w}_2, \dots, \vec{w}_n \in W$. There exists a unique linear transformation $T: V \rightarrow W$ with $T(\vec{u}_i) = \vec{w}_i$ for all $i \in \{1, 2, \dots, n\}$.*

Proof. See the proof of Theorem 2.4.5 for existence, and then use Theorem 5.1.5 for uniqueness. \square

Let $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear transformation. Letting $(\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n)$ be the standard basis in \mathbb{R}^n , we know from Theorem 5.1.5 that T is completely determined by the values $T(\vec{e}_1), T(\vec{e}_2), \dots, T(\vec{e}_n)$. Furthermore, no matter how we choose the values $T(\vec{e}_1), T(\vec{e}_2), \dots, T(\vec{e}_n)$, Theorem 5.1.6 tells us that there is in fact a linear transformation that satisfies these requirements. It is natural to code this information using an $m \times n$

matrix where the i^{th} column is the vector $T(\vec{e}_i)$ (which will be an element of \mathbb{R}^m). For example, there is a unique linear transformation $T: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ such that

$$T\left(\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 5 \\ 2 \end{pmatrix}, \quad T\left(\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} -3 \\ 1 \end{pmatrix}, \quad \text{and} \quad T\left(\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} 7 \\ 0 \end{pmatrix}.$$

As in the case of linear transformations from \mathbb{R}^2 to \mathbb{R}^2 , we can code this linear transformation by putting the three output vectors into the columns of a matrix:

$$\begin{pmatrix} 5 & -3 & 7 \\ 2 & 1 & 0 \end{pmatrix}.$$

As in the past, we call this matrix the *standard matrix* of T , because it is chosen with respect to the standard bases of \mathbb{R}^3 and \mathbb{R}^2 . Notice this is a 2×3 matrix, and its rectangular shape represents the fact that the dimension of the domain of T differs from the dimension of the codomain of T . Nonetheless, we can use this matrix in much the same way that we did in the past. For example, suppose that we want to determine the value of

$$T\left(\begin{pmatrix} -2 \\ 6 \\ 3 \end{pmatrix}\right).$$

Since T is a linear transformation, we have

$$\begin{aligned} T\left(\begin{pmatrix} -2 \\ 6 \\ 3 \end{pmatrix}\right) &= T\left((-2) \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + 6 \cdot \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} + 3 \cdot \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}\right) \\ &= (-2) \cdot T\left(\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}\right) + 6 \cdot T\left(\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}\right) + 3 \cdot T\left(\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}\right) \\ &= (-2) \cdot \begin{pmatrix} 5 \\ 2 \end{pmatrix} + 6 \cdot \begin{pmatrix} -3 \\ 1 \end{pmatrix} + 3 \cdot \begin{pmatrix} 7 \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} (-2) \cdot 5 + 6 \cdot (-3) + 3 \cdot 7 \\ (-2) \cdot 2 + 6 \cdot 1 + 3 \cdot 0 \end{pmatrix} \\ &= \begin{pmatrix} -7 \\ 2 \end{pmatrix}. \end{aligned}$$

Notice from the second to last line that this is the same result as if we naively multiplied the 2×3 standard matrix by the input vector:

$$\begin{pmatrix} 5 & -3 & 7 \\ 2 & 1 & 0 \end{pmatrix} \begin{pmatrix} -2 \\ 6 \\ 3 \end{pmatrix}.$$

Rather than work out all of the details underlying how the standard matrix properly codes a linear transformation from $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$, we instead generalize it all first. Beyond \mathbb{R}^n , we have many other examples of vector spaces. Can we assign a matrix to more general linear transformations from these more exotic vector spaces? For example, how would we code a linear transformation $T: \mathcal{P}_2 \rightarrow W$, where W is the vector space of all 2×2 matrices? Can we actually code it with a table of numbers even though the elements of \mathcal{P}_2 and W are not actually lists of numbers? The key idea is to use *bases* of the two vector spaces. These bases provide coordinates, which turn elements of the vector spaces into lists of numbers that we can record.

Suppose then that we have two vector spaces V and W along with a linear transformation $T: V \rightarrow W$. Suppose that we have a basis $\alpha = (\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ of V and a basis $\beta = (\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ of W . We know

that T is completely determined by the values $T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n)$. Now each of these n vectors are elements of W , and we can code each of them using a list of m numbers obtained by inputting them into the function $\text{Coord}_\beta: W \rightarrow \mathbb{R}^m$. Thus, we can completely understand T by the values of these $m \cdot n$ many real numbers. With this in mind, we have the following definition.

Definition 5.1.7. Let $T: V \rightarrow W$ be a linear transformation, let $\alpha = (\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ be a basis for V , and let $\beta = (\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ be a basis for W . We define the matrix of T relative to α and β to be the $m \times n$ matrix where the i^{th} column is $[T(\vec{u}_i)]_\beta$. We denote this matrix by $[T]_\alpha^\beta$.

For example, consider the situation above where $V = \mathcal{P}_2$ and W is the vector space of all 2×2 matrices. Let $T: V \rightarrow W$ be the function

$$T(ax^2 + bx + c) = \begin{pmatrix} 2a + 5b - c & 3b \\ -7a + 4c & a + b + c \end{pmatrix}.$$

It is straightforward (but a bit tedious) to check that T is a linear transformation. Now in order to code T as a matrix, we need a choice of bases for each of V and W . Suppose that we use the basis $\alpha = (1, x, x^2)$ of V and the basis

$$\beta = \left(\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \right)$$

of W . Now since $\dim(W) = 4$ and $\dim(V) = 3$, we know that the matrix $[T]_\alpha^\beta$ will be a 4×3 matrix. To calculate it, we apply T to each of the elements of α in turn, and then express the output in the coordinates given by β . For example, we have

$$\begin{aligned} T(x^2) &= T(1 \cdot x^2 + 0 \cdot x + 0) \\ &= \begin{pmatrix} 2 & 0 \\ -7 & 1 \end{pmatrix} \\ &= 2 \cdot \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + 0 \cdot \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} + (-7) \cdot \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} + 1 \cdot \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}. \end{aligned}$$

Therefore, we have

$$[T(x^2)]_\beta = \begin{pmatrix} 2 \\ 0 \\ -7 \\ 1 \end{pmatrix}$$

and so this will be the third column of our matrix $[T]_\alpha^\beta$. Working through the other two calculations, we end up with

$$[T]_\alpha^\beta = \begin{pmatrix} -1 & 5 & 2 \\ 0 & 3 & 0 \\ 4 & 0 & -7 \\ 1 & 1 & 1 \end{pmatrix}.$$

Notation 5.1.8. Let $n \in \mathbb{N}^+$. We let ε_n be the standard basis of \mathbb{R}^n , i.e. $\varepsilon_n = (\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n)$ where \vec{e}_i is the vector in \mathbb{R}^n whose i^{th} entry is 1 and all other entries are 0.

Although we used natural bases above, we can also pick more interesting examples of bases to use for our matrices. For example, consider the linear transformation $T: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ given by

$$T \left(\begin{pmatrix} x \\ y \\ z \end{pmatrix} \right) = \begin{pmatrix} x + 2y - z \\ y + 3z \end{pmatrix}.$$

Notice that T is indeed a linear transformation by Proposition 5.1.2. We first look at the standard matrix, which can be written in our new notation as $[T]_{\varepsilon_3}^{\varepsilon_2}$. We have

$$T\left(\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = 1 \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + 0 \cdot \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$$

so

$$\left[T\left(\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}\right)\right]_{\varepsilon_2} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}.$$

Similarly, we have

$$T\left(\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 2 \\ 1 \\ 0 \end{pmatrix} = 2 \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + 1 \cdot \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$$

so

$$\left[T\left(\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}\right)\right]_{\varepsilon_2} = \begin{pmatrix} 2 \\ 1 \\ 0 \end{pmatrix}.$$

and also

$$T\left(\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} -1 \\ 3 \\ 1 \end{pmatrix} = -1 \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + 3 \cdot \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} + 1 \cdot \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

so

$$\left[T\left(\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}\right)\right]_{\varepsilon_2} = \begin{pmatrix} -1 \\ 3 \\ 1 \end{pmatrix}.$$

Therefore, we have

$$[T]_{\varepsilon_3}^{\varepsilon_2} = \begin{pmatrix} 1 & 2 & -1 \\ 0 & 1 & 3 \\ 0 & 0 & 1 \end{pmatrix}.$$

Notice that this matrix is obtained by simply writing down the coefficients of the x, y, z in order. In general, if we have a linear transformation $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ and we know that it is given by a nice formula as in Proposition 5.1.2, then the standard matrix will simply be the table of coefficients.

Suppose instead that we want to compute the matrix of T relative to a different basis for \mathbb{R}^3 and \mathbb{R}^2 . For example, let

$$\alpha = \left(\begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix}, \begin{pmatrix} 5 \\ 0 \\ 2 \end{pmatrix} \right)$$

and let

$$\beta = \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right).$$

By applying elementary row operations to the corresponding 3×3 matrix and checking that each row and column have a leading entry, we can verify that α is indeed a basis of \mathbb{R}^3 . We can perform a similar calculation for β , but in that case we can simply check that $1 \cdot 0 - 1 \cdot 1 = -1$ is nonzero. We now compute $[T]_{\alpha}^{\beta}$. We have

$$T\left(\begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix}\right) = \begin{pmatrix} 6 \\ -1 \\ 1 \end{pmatrix}.$$

Now to determine

$$\left[T \left(\begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix} \right) \right]_{\beta},$$

we want to find the unique $c_1, c_2 \in \mathbb{R}$ with

$$\begin{pmatrix} 6 \\ -1 \end{pmatrix} = c_1 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + c_2 \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

Working through the calculations, we have

$$\begin{pmatrix} 6 \\ -1 \end{pmatrix} = (-1) \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + 7 \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix},$$

so

$$\left[T \left(\begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix} \right) \right]_{\beta} = \begin{pmatrix} -1 \\ 7 \end{pmatrix}.$$

For the next two vectors, we have

$$T \left(\begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} \right) = \begin{pmatrix} -5 \\ 8 \end{pmatrix} = 8 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + (-13) \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix},$$

so

$$\left[T \left(\begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} \right) \right]_{\beta} = \begin{pmatrix} 8 \\ -13 \end{pmatrix},$$

and we have

$$T \left(\begin{pmatrix} 5 \\ 0 \\ 2 \end{pmatrix} \right) = \begin{pmatrix} 3 \\ 6 \end{pmatrix} = 6 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + (-3) \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix},$$

so

$$\left[T \left(\begin{pmatrix} 5 \\ 0 \\ 2 \end{pmatrix} \right) \right]_{\beta} = \begin{pmatrix} 6 \\ -3 \end{pmatrix}.$$

Therefore, we have the matrix

$$[T]_{\alpha}^{\beta} = \begin{pmatrix} -1 & 8 & 6 \\ 7 & -13 & -3 \end{pmatrix}.$$

Given two vector spaces V and W , where $n = \dim(V)$ and $m = \dim(W)$, along with a linear transformation T , we have figured out a way to code T using an $m \times n$ matrix. In fact, for different choices of bases for V and W , we obtain potentially different $m \times n$ matrices. Of course, these codings are only useful if we learn how to calculate with them. As in the case of \mathbb{R}^2 , to determine $T(\vec{v})$, we would like to be able to simply “hit” the corresponding matrix by (a representation of) \vec{v} . In other words, we want a simple mechanical way to define a matrix times a vector so that it will correspond to feeding the vector as input into the corresponding function. The following result tells us how to define the matrix-vector product so that this will work.

Proposition 5.1.9. *Let $T: V \rightarrow W$ be a linear transformation, let $\alpha = (\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ be a basis for V , and let $\beta = (\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ be a basis for W . Suppose that*

$$[T]_{\alpha}^{\beta} = \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{pmatrix}$$

and that

$$[\vec{v}]_{\alpha} = \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix}.$$

We then have that

$$[T(\vec{v})]_{\beta} = c_1 \cdot [T(\vec{u}_1)]_{\beta} + c_2 \cdot [T(\vec{u}_2)]_{\beta} + \cdots + c_n \cdot [T(\vec{u}_n)]_{\beta},$$

and hence

$$[T(\vec{v})]_{\beta} = \begin{pmatrix} a_{1,1}c_1 + a_{1,2}c_2 + \cdots + a_{1,n}c_n \\ a_{2,1}c_1 + a_{2,2}c_2 + \cdots + a_{2,n}c_n \\ \vdots \\ a_{m,1}c_1 + a_{m,2}c_2 + \cdots + a_{m,n}c_n \end{pmatrix}.$$

Proof. Suppose that

$$[\vec{v}]_{\alpha} = \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix}.$$

By definition of Coord_{α} , this implies that

$$\vec{v} = c_1\vec{u}_1 + c_2\vec{u}_2 + \cdots + c_n\vec{u}_n.$$

Since T is a linear transformation, we have

$$\begin{aligned} T(\vec{v}) &= T(c_1\vec{u}_1 + c_2\vec{u}_2 + \cdots + c_n\vec{u}_n) \\ &= c_1 \cdot T(\vec{u}_1) + c_2 \cdot T(\vec{u}_2) + \cdots + c_n \cdot T(\vec{u}_n). \end{aligned}$$

Since Coord_{β} is linear transformation by Proposition 5.1.3, it follows that

$$\begin{aligned} [T(\vec{v})]_{\beta} &= [c_1 \cdot T(\vec{u}_1) + c_2 \cdot T(\vec{u}_2) + \cdots + c_n \cdot T(\vec{u}_n)]_{\beta} \\ &= c_1 \cdot [T(\vec{u}_1)]_{\beta} + c_2 \cdot [T(\vec{u}_2)]_{\beta} + \cdots + c_n \cdot [T(\vec{u}_n)]_{\beta}. \end{aligned}$$

Finally, using the fact that

$$[T(\vec{u}_j)]_{\beta} = \begin{pmatrix} a_{1,j} \\ a_{2,j} \\ \vdots \\ a_{m,j} \end{pmatrix}$$

for all $j \in \{1, 2, \dots, n\}$ by definition of $[T]_\alpha^\beta$, we conclude that

$$\begin{aligned} [T(\vec{v})]_\beta &= c_1 \cdot [T(\vec{u}_1)]_\beta + c_2 \cdot [T(\vec{u}_2)]_\beta + \cdots + c_n \cdot [T(\vec{u}_n)]_\beta \\ &= c_1 \cdot \begin{pmatrix} a_{1,1} \\ a_{2,1} \\ \vdots \\ a_{m,1} \end{pmatrix} + c_2 \cdot \begin{pmatrix} a_{1,2} \\ a_{2,2} \\ \vdots \\ a_{m,2} \end{pmatrix} + \cdots + c_n \cdot \begin{pmatrix} a_{1,n} \\ a_{2,n} \\ \vdots \\ a_{m,n} \end{pmatrix} \\ &= \begin{pmatrix} a_{1,1}c_1 + a_{1,2}c_2 + \cdots + a_{1,n}c_n \\ a_{2,1}c_1 + a_{2,2}c_2 + \cdots + a_{2,n}c_n \\ \vdots \\ a_{m,1}c_1 + a_{m,2}c_2 + \cdots + a_{m,n}c_n \end{pmatrix}. \end{aligned}$$

□

With this in mind, we generalize our definition of a matrix-vector product as follows.

Definition 5.1.10. Given an $m \times n$ matrix A and a vector $\vec{v} \in \mathbb{R}^n$, say

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{pmatrix} \quad \text{and} \quad \vec{v} = \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix},$$

we define $A\vec{v}$ to be the following vector in \mathbb{R}^m :

$$A\vec{v} = \begin{pmatrix} a_{1,1}c_1 + a_{1,2}c_2 + \cdots + a_{1,n}c_n \\ a_{2,1}c_1 + a_{2,2}c_2 + \cdots + a_{2,n}c_n \\ \vdots \\ a_{m,1}c_1 + a_{m,2}c_2 + \cdots + a_{m,n}c_n \end{pmatrix}.$$

We call this the matrix-vector product.

We can rephrase all of this as follows. Let $T: V \rightarrow W$ be a linear transformation, let $\alpha = (\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ be a basis for V , and let $\beta = (\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ be a basis for W . For all $\vec{v} \in V$, we have $[T]_\alpha^\beta \cdot [\vec{v}]_\alpha = [T(\vec{v})]_\beta$. In other words, once we have the matrix $[T]_\alpha^\beta$, we can mimic the computation of plugging an input \vec{v} into T and computing the coordinates of the result by instead multiplying the matrix $[T]_\alpha^\beta$ by the vector $[\vec{v}]_\alpha$.

Recall our above example where $T: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ is given by

$$T\left(\begin{pmatrix} x \\ y \\ z \end{pmatrix}\right) = \begin{pmatrix} x + 2y - z \\ y + 3z \end{pmatrix}$$

and where we use the bases

$$\alpha = \left(\begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix}, \begin{pmatrix} 5 \\ 0 \\ 2 \end{pmatrix} \right) \quad \text{and} \quad \beta = \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right).$$

We calculated that

$$[T]_{\beta}^{\alpha} = \begin{pmatrix} 1 & 2 & -1 \\ 0 & 1 & 3 \end{pmatrix}$$

and that

$$[T]_{\alpha}^{\beta} = \begin{pmatrix} -1 & 8 & 6 \\ 7 & -13 & -3 \end{pmatrix}.$$

Let

$$\vec{v} = \begin{pmatrix} 13 \\ 5 \\ 4 \end{pmatrix}.$$

Now one way to compute $T(\vec{v})$ is to simply multiply the standard matrix by this vector:

$$\begin{aligned} T(\vec{v}) &= \begin{pmatrix} 1 & 2 & -1 \\ 0 & 1 & 3 \end{pmatrix} \begin{pmatrix} 13 \\ 5 \\ 4 \end{pmatrix} \\ &= \begin{pmatrix} 1 \cdot 13 + 2 \cdot 5 + (-1) \cdot 4 \\ 0 \cdot 13 + 1 \cdot 5 + 3 \cdot 4 \end{pmatrix} \\ &= \begin{pmatrix} 19 \\ 17 \end{pmatrix}. \end{aligned}$$

Suppose that we wanted to do a computation with the matrix $[T]_{\alpha}^{\beta}$. We know that this matrix will convert the α -coordinates of a vector \vec{v} into the β -coordinates of $T(\vec{v})$. Let's go ahead and compute $[\vec{v}]_{\alpha}$ in our case. To do this, we want to find the unique $c_1, c_2, c_3 \in \mathbb{R}$ with

$$\begin{pmatrix} 13 \\ 5 \\ 4 \end{pmatrix} = c_1 \begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix} + c_2 \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} + c_3 \begin{pmatrix} 5 \\ 0 \\ 2 \end{pmatrix}.$$

Solving the corresponding linear system, we conclude that

$$[\vec{v}]_{\alpha} = \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix}.$$

Now one way to determine $[T(\vec{v})]_{\beta}$ is to simply find the β -coordinates of $T(\vec{v})$ directly, i.e. to find the $d_1, d_2 \in \mathbb{R}$ such that

$$\begin{pmatrix} 19 \\ 17 \end{pmatrix} = d_1 \begin{pmatrix} 1 \\ 1 \end{pmatrix} + d_2 \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

Solving this system, we see that

$$[T(\vec{v})]_{\beta} = \begin{pmatrix} 17 \\ 2 \end{pmatrix}.$$

However, we can also determine $[T(\vec{v})]_{\beta}$ by simply multiplying the matrix $[T]_{\alpha}^{\beta}$ by the vector $[\vec{v}]_{\alpha}$ to obtain

$$\begin{aligned} [T(\vec{v})]_{\beta} &= \begin{pmatrix} -1 & 8 & 6 \\ 7 & -13 & -3 \end{pmatrix} \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix} \\ &= \begin{pmatrix} (-1) \cdot 3 + 8 \cdot 1 + 6 \cdot 2 \\ 7 \cdot 3 + (-13) \cdot 1 + (-3) \cdot 2 \end{pmatrix} \\ &= \begin{pmatrix} 17 \\ 2 \end{pmatrix}. \end{aligned}$$

Now that we can code *any* linear transformation between finite-dimensional vector spaces as a matrix, let's take a look at a more interesting example. Define $T: \mathcal{P}_3 \rightarrow \mathcal{P}_3$ by letting $T(f) = f'$. Let's compute $[T]_\alpha^\alpha$ in the case where $\alpha = (1, x, x^2, x^3)$. For the third column of $[T]_\alpha^\alpha$ we have

$$\begin{aligned} T(x^2) &= 2x \\ &= 0 \cdot 1 + 2 \cdot x + 0 \cdot x^2 + 0 \cdot x^3, \end{aligned}$$

so we have

$$[T(x^2)]_\alpha = \begin{pmatrix} 0 \\ 2 \\ 0 \\ 0 \end{pmatrix}.$$

Working through these computations for the other basis elements in α , we end up with

$$[T]_\alpha^\alpha = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

We can now use this matrix to take derivatives for us! For example, let $g \in \mathcal{P}$ be the function $g(x) = 2x^3 + 7x^2 - 8x + 3$. We then have

$$[g]_\alpha = \begin{pmatrix} 3 \\ -8 \\ 7 \\ 2 \end{pmatrix}.$$

It follows that

$$\begin{aligned} [T(g)]_\alpha &= [T]_\alpha^\alpha \cdot [g]_\alpha \\ &= \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 3 \\ -8 \\ 7 \\ 2 \end{pmatrix} \\ &= \begin{pmatrix} -8 \\ 14 \\ 6 \\ 0 \end{pmatrix}. \end{aligned}$$

From this, we can read off that $T(g) = 0x^3 + 6x^2 + 14x - 8$. Of course, in this case, computing the derivative directly was certainly faster, but this example illustrates how a more complicated linear transformation can be coded as a matrix and then used for calculations.

Just like for linear transformations from \mathbb{R}^2 to \mathbb{R}^2 , we can both add linear transformations and multiply them by scalars to obtain new linear transformations. However, notice that in order to add T and S , they must have the same domain and codomain.

Definition 5.1.11. *Let V and W be vector spaces.*

- *Let $T: V \rightarrow W$ and $S: V \rightarrow W$ be linear transformations. We define a new function $T + S: V \rightarrow W$ by letting $(T + S)(\vec{v}) = T(\vec{v}) + S(\vec{v})$ for all $\vec{v} \in V$.*

- Let $T: V \rightarrow W$ be a linear transformation and let $r \in \mathbb{R}$. We define a new function $r \cdot T: V \rightarrow W$ by letting $(r \cdot T)(\vec{v}) = r \cdot T(\vec{v})$ for all $\vec{v} \in V$.

Proposition 5.1.12. *Let V and W be vector spaces, and let $T, S: V \rightarrow W$ be linear transformations.*

1. *The function $T + S$ is a linear transformation.*
2. *For all $r \in \mathbb{R}$, then function $r \cdot T$ is a linear transformation.*

Proof. The proof is exactly the same as the proof of Proposition 2.4.7, except for replacing \mathbb{R}^2 by V or W , as appropriate. \square

With this result in hand, we want to define matrix addition and multiplication of a matrix by a scalar in a way so that it reflects the operations of addition and scalar multiplication of linear transformations. As mentioned above, we need the domain and codomain of T and S to match in order to define $T + S$. Thus, we should expect that the sizes of two matrices have to match if we want to add them.

Definition 5.1.13. *Given two $m \times n$ matrices*

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} b_{1,1} & b_{1,2} & \cdots & b_{1,n} \\ b_{2,1} & b_{2,2} & \cdots & b_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ b_{m,1} & b_{m,2} & \cdots & b_{m,n} \end{pmatrix},$$

we define

$$A + B = \begin{pmatrix} a_{1,1} + b_{1,1} & a_{1,2} + b_{1,2} & \cdots & a_{1,n} + b_{1,n} \\ a_{2,1} + b_{2,1} & a_{2,2} + b_{2,2} & \cdots & a_{2,n} + b_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} + b_{m,1} & a_{m,2} + b_{m,2} & \cdots & a_{m,n} + b_{m,n} \end{pmatrix}.$$

Definition 5.1.14. *Given an $m \times n$ matrix*

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{pmatrix}$$

and a $c \in \mathbb{R}$, we define

$$c \cdot A = \begin{pmatrix} ca_{1,1} & ca_{1,2} & \cdots & ca_{1,n} \\ ca_{2,1} & ca_{2,2} & \cdots & ca_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ ca_{m,1} & ca_{m,2} & \cdots & ca_{m,n} \end{pmatrix}.$$

Proposition 5.1.15. *Let V and W be vector spaces, let $T: V \rightarrow W$ and $S: V \rightarrow W$ be linear transformations, let $\alpha = (\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ be a basis for V , and let $\beta = (\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m)$ be a basis for W . We have the following:*

1. $[T + S]_{\alpha}^{\beta} = [T]_{\alpha}^{\beta} + [S]_{\alpha}^{\beta}$.
2. $[c \cdot T]_{\alpha}^{\beta} = c \cdot [T]_{\alpha}^{\beta}$ for all $c \in \mathbb{R}$.

Proof. Exercise. For (1), apply $T + S$ to each of the \vec{u}_i in turn to compute $[T + S]_{\alpha}^{\beta}$. \square

We can also compose linear transformations (since they are functions after all), as long as the codomain of the first function that we apply equals the domain of the second. By performing this operation, we always do obtain another linear transformation.

Proposition 5.1.16. *Let V , Z , and W be vector spaces, and let $T: Z \rightarrow W$ and $S: V \rightarrow Z$ be linear transformations. We then have that $T \circ S: V \rightarrow W$ is a linear transformation.*

Proof. We first check that $T \circ S$ preserves addition. Let $\vec{v}_1, \vec{v}_2 \in V$ be arbitrary. We have

$$\begin{aligned} (T \circ S)(\vec{v}_1 + \vec{v}_2) &= T(S(\vec{v}_1 + \vec{v}_2)) && \text{(by definition)} \\ &= T(S(\vec{v}_1) + S(\vec{v}_2)) && \text{(since } S \text{ is a linear transformation)} \\ &= T(S(\vec{v}_1)) + T(S(\vec{v}_2)) && \text{(since } T \text{ is a linear transformation)} \\ &= (T \circ S)(\vec{v}_1) + (T \circ S)(\vec{v}_2) && \text{(by definition).} \end{aligned}$$

Therefore, the function $T \circ S$ preserves addition.

We now check that $T \circ S$ preserves scalar multiplication. Let $\vec{v} \in V$ and $c \in \mathbb{R}$ be arbitrary. We have

$$\begin{aligned} (T \circ S)(c \cdot \vec{v}) &= T(S(c \cdot \vec{v})) && \text{(by definition)} \\ &= T(c \cdot S(\vec{v})) && \text{(since } S \text{ is a linear transformation)} \\ &= c \cdot T(S(\vec{v})) && \text{(since } T \text{ is a linear transformation)} \\ &= c \cdot (T \circ S)(\vec{v}) && \text{(by definition).} \end{aligned}$$

Therefore, the function $T \circ S$ preserves scalar multiplication as well. It follows that $T \circ S$ is a linear transformation. \square

As in the case for linear transformations from \mathbb{R}^2 to \mathbb{R}^2 , we want to define matrix multiplication in a way so that it corresponds to composition of linear transformations. Recall that an $m \times n$ matrix codes a linear transformation from an n -dimensional vector space to an m -dimensional vector space. Since composition only makes sense if the codomain of the first function we apply matches the domain of the second, we will need the numbers to line up. In other words, we will only be able to multiply an $m \times p$ matrix and a $p \times n$ matrix in that order (since the right function is applied first, it's codomain will be p -dimensional, while the left function has a p -dimensional range).

Suppose then that V , Z , and W are finite-dimensional vector spaces, and that $T: Z \rightarrow W$ and $S: V \rightarrow Z$ are linear transformations. We want to consider the linear transformation $T \circ S: V \rightarrow W$. Let's assume that $\dim(V) = n$, that $\dim(Z) = p$, and that $\dim(W) = m$. Fix bases for each of the vector spaces as follows:

- Let $\alpha = (\vec{u}_1, \dots, \vec{u}_n)$ be a basis for V .
- Let $\gamma = (\vec{z}_1, \dots, \vec{z}_p)$ be a basis Z .
- Let $\beta = (\vec{w}_1, \dots, \vec{w}_m)$ be a basis W .

Let $A = [T]_\gamma^\beta$ and notice that A is an $m \times p$ matrix. Let $B = [S]_\alpha^\gamma$ and notice that B is a $p \times n$ matrix. Write out A and B as

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,p} \\ a_{2,1} & a_{2,2} & \dots & a_{2,p} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \dots & a_{m,p} \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} b_{1,1} & b_{1,2} & \dots & b_{1,n} \\ b_{2,1} & b_{2,2} & \dots & b_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ b_{p,1} & b_{p,2} & \dots & b_{p,n} \end{pmatrix}.$$

Now we want to define the matrix product AB so that it will equal $[T \circ S]_\alpha^\beta$, and so the result should be an $m \times n$ matrix. The key question that we need to ask is what $T \circ S$ does to an arbitrary \vec{u}_j . For any $j \in \{1, 2, \dots, n\}$, we have

$$\begin{aligned} (T \circ S)(\vec{u}_j) &= T(S(\vec{u}_j)) \\ &= T(b_{1,j}\vec{z}_1 + \dots + b_{p,j}\vec{z}_p) \\ &= b_{1,j} \cdot T(\vec{z}_1) + \dots + b_{p,j} \cdot T(\vec{z}_p) \\ &= b_{1,j} \cdot (a_{1,1}\vec{w}_1 + a_{2,1}\vec{w}_2 + \dots + a_{m,1}\vec{w}_m) + \dots + b_{p,j} \cdot (a_{1,p}\vec{w}_1 + a_{2,p}\vec{w}_2 + \dots + a_{m,p}\vec{w}_m) \\ &= (b_{1,j}a_{1,1} + b_{2,j}a_{1,2} + \dots + b_{p,j}a_{1,p}) \cdot \vec{w}_1 + \dots + (b_{1,j}a_{m,1} + b_{2,j}a_{m,2} + \dots + b_{p,j}a_{m,p}) \cdot \vec{w}_m \\ &= (a_{1,1}b_{1,j} + a_{1,2}b_{2,j} + \dots + a_{1,p}b_{p,j}) \cdot \vec{w}_1 + \dots + (a_{m,1}b_{1,j} + a_{m,2}b_{2,j} + \dots + a_{m,p}b_{p,j}) \cdot \vec{w}_m. \end{aligned}$$

It follows that we have

$$[(T \circ S)(\vec{u}_j)]_\beta = \begin{pmatrix} a_{1,1}b_{1,j} + a_{1,2}b_{2,j} + \dots + a_{1,p}b_{p,j} \\ a_{2,1}b_{1,j} + a_{2,2}b_{2,j} + \dots + a_{2,p}b_{p,j} \\ \vdots \\ a_{m,1}b_{1,j} + a_{m,2}b_{2,j} + \dots + a_{m,p}b_{p,j} \end{pmatrix}.$$

Therefore, if we want to define AB so that it equals $[T \circ S]_\alpha^\beta$, then we should define the AB to be the $m \times n$ matrix whose (i, j) entry is

$$a_{i,1}b_{1,j} + a_{i,2}b_{2,j} + \dots + a_{i,p}b_{p,j},$$

i.e. that (i, j) entry should be the dot product of row i of A and column j of B .

Definition 5.1.17. Given an $m \times p$ matrix A and a $p \times n$ matrix B with

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,p} \\ a_{2,1} & a_{2,2} & \dots & a_{2,p} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \dots & a_{m,p} \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} b_{1,1} & b_{1,2} & \dots & b_{1,n} \\ b_{2,1} & b_{2,2} & \dots & b_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ b_{p,1} & b_{p,2} & \dots & b_{p,n} \end{pmatrix}$$

we define AB to be the $m \times n$ matrix

$$C = \begin{pmatrix} c_{1,1} & c_{1,2} & \dots & c_{1,n} \\ c_{2,1} & c_{2,2} & \dots & c_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ c_{m,1} & c_{m,2} & \dots & c_{m,n} \end{pmatrix}$$

where

$$c_{i,j} = a_{i,1}b_{1,j} + a_{i,2}b_{2,j} + \dots + a_{i,p}b_{p,j}$$

for all $i \in \{1, 2, \dots, m\}$ and all $j \in \{1, 2, \dots, n\}$, i.e. $c_{i,j}$ is the dot product of row i of A and column j of B .

We've now defined matrix multiplication so that the following result is true:

Proposition 5.1.18. Let V , Z , and W be finite-dimensional vector spaces, and let $T: Z \rightarrow W$ and $S: V \rightarrow Z$ be linear transformations. Let α be a basis for V , let γ be a basis for Z , and let β be a basis for W . We then have $[T \circ S]_\alpha^\beta = [T]_\gamma^\beta \cdot [S]_\alpha^\gamma$.

Suppose now that

$$A = \begin{pmatrix} 1 & 5 & -2 \\ -1 & 2 & 4 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 4 & 1 \\ -3 & 0 \\ 7 & 2 \end{pmatrix}.$$

Now A could arise as the matrix of many different linear transformations and choices of bases for the domain and codomain. However, since A is a 2×3 matrix, in all of these cases A will code a linear transformation from a 3-dimensional vector space to a 2-dimensional vector space. For one particular example, we could view A as the standard matrix (i.e. relative to ε_3 and ε_2) of the linear transformation $T: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ given by

$$T\left(\begin{pmatrix} x \\ y \\ z \end{pmatrix}\right) = \begin{pmatrix} x + 5y - 2z \\ -x + 2y + 4z \end{pmatrix}.$$

Similarly, B is a 3×2 matrix, and so it codes a linear transformation from a 2-dimensional vector space to a 3-dimensional vector space. There are many ways that B could arise like this, but for the simplest example, we could view B as the standard matrix (i.e. relative to ε_2 and ε_3) of the linear transformation $S: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ given by

$$S\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} 4x + y \\ -3x \\ 7x + 2y \end{pmatrix}.$$

Now when we compute AB , we can view this as the standard matrix of $T \circ S$. Notice that we can multiply A times B in this order because the codomain of S and the domain of T both equal \mathbb{R}^3 . Furthermore, although \mathbb{R}^3 appears in the middle, the function $T \circ S$ has domain \mathbb{R}^2 and codomain \mathbb{R}^2 , so AB will be a 2×2 matrix. Recall that in general, when we multiply an $m \times p$ matrix by a $p \times n$ matrix, we obtain an $m \times n$ matrix. In our case, we have

$$\begin{aligned} AB &= \begin{pmatrix} 1 & 5 & -2 \\ -1 & 2 & 4 \end{pmatrix} \begin{pmatrix} 4 & 1 \\ -3 & 0 \\ 7 & 2 \end{pmatrix} \\ &= \begin{pmatrix} -25 & -3 \\ 18 & 7 \end{pmatrix}. \end{aligned}$$

We can also compute the product BA , which we can view as the standard matrix of $S \circ T$. This computation makes sense because the codomain of T and the domain of S both equal \mathbb{R}^2 . However, notice now that $S \circ T$ has domain \mathbb{R}^3 and codomain \mathbb{R}^3 , so BA will be a 3×3 matrix:

$$\begin{aligned} BA &= \begin{pmatrix} 4 & 1 \\ -3 & 0 \\ 7 & 2 \end{pmatrix} \begin{pmatrix} 1 & 5 & -2 \\ -1 & 2 & 4 \end{pmatrix} \\ &= \begin{pmatrix} 3 & 22 & -4 \\ -3 & -15 & 6 \\ 5 & 39 & -6 \end{pmatrix}. \end{aligned}$$

Thus, not only does $AB \neq BA$, but in fact they have different sizes! In contrast, all of the properties of matrix addition and multiplication that we discussed in the context of \mathbb{R}^2 carry over now. For example, matrix multiplication is associative (at least whenever it is defined).

5.2 The Range and Null Space of a Linear Transformation

Suppose that we have a linear transformation $T: V \rightarrow W$. Since T is a function, we know that we can talk about its range. Recall that

$$\text{range}(T) = \{\vec{w} \in W : \text{There exists } \vec{v} \in V \text{ with } T(\vec{v}) = \vec{w}\},$$

and so $\text{range}(T)$ is a subset of W . We can also generalize our null space, and in this context, the null space of a linear transformation will be a subset of V .

Definition 5.2.1. Let $T: V \rightarrow W$ be a linear transformation. We define

$$\text{Null}(T) = \{\vec{v} \in V : T(\vec{v}) = \vec{0}_W\}.$$

The set $\text{Null}(T)$ is called the null space, or kernel, of T .

Beyond simply being subsets of V and W , the sets $\text{range}(T)$ and $\text{Null}(T)$ are in fact *subspaces*.

Proposition 5.2.2. Let V and W be vector spaces, and let $T: V \rightarrow W$ be a linear transformation.

1. $\text{Null}(T)$ is a subspace of V .
2. $\text{range}(T)$ is a subspace of W .

Proof.

1. We check the three properties:

- We know that $T(\vec{0}_V) = \vec{0}_W$ by Proposition 5.1.4, so $\vec{0}_V \in \text{Null}(T)$.
- Let $\vec{v}_1, \vec{v}_2 \in \text{Null}(T)$ be arbitrary. By definition of $\text{Null}(T)$, we have that $T(\vec{v}_1) = \vec{0}$ and $T(\vec{v}_2) = \vec{0}$. Since T is a linear transformation, we have

$$\begin{aligned} T(\vec{v}_1 + \vec{v}_2) &= T(\vec{v}_1) + T(\vec{v}_2) \\ &= \vec{0} + \vec{0} \\ &= \vec{0}. \end{aligned}$$

Therefore $\vec{v}_1 + \vec{v}_2 \in \text{Null}(T)$.

- Let $\vec{v} \in \text{Null}(T)$ and $c \in \mathbb{R}$ be arbitrary. By definition of $\text{Null}(T)$, we have that $T(\vec{v}) = \vec{0}$. Since T is a linear transformation, we have

$$\begin{aligned} T(c \cdot \vec{v}) &= c \cdot T(\vec{v}) \\ &= c \cdot \vec{0} \\ &= \vec{0}. \end{aligned}$$

Therefore $c \cdot \vec{v} \in \text{Null}(T)$.

2. We check the three properties:

- We know that $T(\vec{0}_V) = \vec{0}_W$ by Proposition 5.1.4, so $\vec{0}_W \in \text{range}(T)$.
- Let $\vec{w}_1, \vec{w}_2 \in \text{range}(T)$ be arbitrary. By definition of $\text{range}(T)$, we can fix $\vec{v}_1, \vec{v}_2 \in V$ with $\vec{w}_1 = T(\vec{v}_1)$ and $\vec{w}_2 = T(\vec{v}_2)$. We then have

$$\begin{aligned} \vec{w}_1 + \vec{w}_2 &= T(\vec{v}_1) + T(\vec{v}_2) \\ &= T(\vec{v}_1 + \vec{v}_2) \end{aligned} \quad (\text{since } T \text{ is a linear transformation}).$$

Since $\vec{v}_1 + \vec{v}_2 \in V$, it follows that $\vec{w}_1 + \vec{w}_2 \in \text{range}(T)$.

- Let $\vec{w} \in \text{range}(T)$ and $c \in \mathbb{R}$ be arbitrary. By definition of $\text{range}(T)$, we can fix $\vec{v} \in V$ with $\vec{w} = T(\vec{v})$. We then have

$$\begin{aligned} c\vec{w} &= c \cdot T(\vec{v}) \\ &= T(c\vec{v}) \end{aligned} \quad (\text{since } T \text{ is a linear transformation}).$$

Since $c\vec{v} \in V$, it follows that $c\vec{w} \in \text{range}(T)$.

□

Let $T: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be the unique linear transformation whose standard matrix (i.e. relative to ε_3 and ε_3) is

$$[T] = \begin{pmatrix} 1 & -1 & 4 \\ -1 & 2 & -3 \\ 3 & -3 & 14 \end{pmatrix}.$$

Notice that for any $x, y, z \in \mathbb{R}$, we have that

$$\begin{aligned} T \left(\begin{pmatrix} x \\ y \\ z \end{pmatrix} \right) &= \begin{pmatrix} 1 & -1 & 4 \\ -1 & 2 & -3 \\ 3 & -3 & 14 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} \\ &= \begin{pmatrix} x - y + 4z \\ -x + 2y - 3z \\ 3x - y + 14z \end{pmatrix}. \end{aligned}$$

How can we determine $\text{range}(T)$ and $\text{Null}(T)$? Given a particular vector $\vec{b} \in \mathbb{R}^3$, say

$$\vec{b} = \begin{pmatrix} -3 \\ 2 \\ -11 \end{pmatrix},$$

to determine if $\vec{b} \in \text{range}(T)$, we want to know if there exists $x, y, z \in \mathbb{R}$ with

$$T \left(\begin{pmatrix} x \\ y \\ z \end{pmatrix} \right) = \begin{pmatrix} -3 \\ 2 \\ -11 \end{pmatrix}.$$

In other words, we want to know if the following system has a solution:

$$\begin{array}{rrcrcl} x & - & y & + & 4z & = & -3 \\ -x & + & 2y & - & 3z & = & 2 \\ 3x & - & y & + & 14z & = & -11. \end{array}$$

To answer this question, we look at the augmented matrix and perform Gaussian Elimination:

$$\left(\begin{array}{cccc} 1 & -1 & 4 & -3 \\ -1 & 2 & -3 & 2 \\ 3 & -1 & 14 & -11 \end{array} \right) \rightarrow \cdots \rightarrow \left(\begin{array}{cccc} 1 & -1 & 4 & -3 \\ 0 & 1 & 1 & -1 \\ 0 & 0 & 0 & 0 \end{array} \right).$$

Since the last column of this augmented matrix does not have a leading entry, we conclude that this system does have a solution. Setting $z = 0$ (because we can build a solution no matter how we choose z , as there is

no leading entry in the third column), we then solve to determine $y = -1$ and $x = -4$. In other words, we have

$$T\left(\begin{pmatrix} -4 \\ -1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} -3 \\ 2 \\ -11 \end{pmatrix}.$$

We can also find a solution for other values of z . Setting $z = 1$ and solving, we obtain $y = -2$ and $x = -9$, so

$$T\left(\begin{pmatrix} -9 \\ -2 \\ 1 \end{pmatrix}\right) = \begin{pmatrix} -3 \\ 2 \\ -11 \end{pmatrix}$$

as well. In any case, we have shown that

$$\begin{pmatrix} -3 \\ 2 \\ -11 \end{pmatrix} \in \text{range}(T).$$

In general, suppose that $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear transformation and that $\vec{b} \in \mathbb{R}^m$. We can always determine whether $\vec{b} \in \text{range}(T)$ by applying Gaussian elimination to check whether the corresponding system has a solution. Notice also that we can view this differently by reinterpreting the system: we have that $\vec{b} \in \text{range}(T)$ if and only if \vec{b} is a linear combination of the columns of $[T]$. Therefore, we obtain the following result.

Proposition 5.2.3. *Let $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear transformation, and let $\vec{b} \in \mathbb{R}^m$. The following are equivalent:*

1. $\vec{b} \in \text{range}(T)$.
2. \vec{b} is a linear combination of the columns of $[T]$.

Thus, $\text{range}(T)$ is the span of the columns of $[T]$.

Since $T: V \rightarrow W$ is surjective if and only if $\text{range}(T) = W$, we obtain the following consequence.

Corollary 5.2.4. *Let $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear transformation, and let B be an echelon form of the matrix $[T]$. The following are equivalent:*

1. T is surjective.
2. Every row of B has a leading entry.

Proof. Immediate from Proposition 5.2.3 and Proposition 4.3.1. □

In addition, we can use our work at the end of Section 4.5 to efficiently find a basis for $\text{range}(T)$ in the case where $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$. The key insight is Proposition 5.2.3, which says that $\text{range}(T)$ is the column space of $[T]$ (see Definition 4.5.7).

Corollary 5.2.5. *Let $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear transformation and let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$ be the columns of $[T]$. Let B be an echelon form of the matrix $[T]$. If we build the sequence consisting only of those \vec{u}_i such that the i^{th} column of B has a leading entry, then we obtain a basis for $\text{range}(T)$.*

Proof. Immediate from Proposition 5.2.3 and Proposition 4.5.6. □

Let's now work to understand $\text{Null}(T)$ in our specific example of a linear transformation $T: \mathbb{R}^3 \rightarrow \mathbb{R}^3$. We want to know all the values of $x, y, z \in \mathbb{R}$ such that

$$T\left(\begin{pmatrix} x \\ y \\ z \end{pmatrix}\right) = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

Since

$$\begin{aligned} T\left(\begin{pmatrix} x \\ y \\ z \end{pmatrix}\right) &= \begin{pmatrix} 1 & -1 & 4 \\ -1 & 2 & -3 \\ 3 & -1 & 14 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} \\ &= \begin{pmatrix} x - y + 4z \\ -x + 2y - 3z \\ 3x - y + 14z \end{pmatrix}, \end{aligned}$$

we want to find all solutions to the following system:

$$\begin{array}{rrrrrr} x & - & y & + & 4z & = & 0 \\ -x & + & 2y & - & 3z & = & 0 \\ 3x & - & y & + & 14z & = & 0. \end{array}$$

Applying elementary row operations to the corresponding augmented matrix, we obtain

$$\begin{pmatrix} 1 & -1 & 4 & 0 \\ -1 & 2 & -3 & 0 \\ 3 & -1 & 14 & 0 \end{pmatrix} \rightarrow \cdots \rightarrow \begin{pmatrix} 1 & -1 & 4 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

Since there is no leading entry in the last column, we know that there is a solution. Of course, we knew this anyway because $(0, 0, 0)$ is a solution (recall that $\vec{0}$ is always an element of $\text{Null}(T)$). However, we can find all elements of $\text{Null}(T)$ by finding all solutions. Letting $z = t$, we can solve to determine $y = -t$ and $x = -5t$. Therefore,

$$\text{Null}(T) = \left\{ \begin{pmatrix} -5t \\ -t \\ t \end{pmatrix} : t \in \mathbb{R} \right\}.$$

In other words, we have

$$\text{Null}(T) = \text{Span} \left(\begin{pmatrix} -5 \\ -1 \\ 1 \end{pmatrix} \right).$$

Notice that we can rewrite the above in terms of matrix arithmetic. Let $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$, and let $A = [T]$. Asking if a given $\vec{b} \in \mathbb{R}^m$ is an element of $\text{range}(T)$ is the same as asking whether the matrix-vector equation $A\vec{x} = \vec{b}$ has a solution, i.e. does there exist an $\vec{x} \in \mathbb{R}^n$ such that the equation is true. Determining $\text{Null}(T)$ is the same as finding all solutions to $A\vec{x} = \vec{0}$.

If we turn all of this around, and start with a linear system, then we can interpret some of our computations in a new light. For example, back when we first started looking at linear systems, we worked with the following:

$$\begin{array}{rrrrrrrr} x_1 & + & 5x_2 & - & x_3 & - & 2x_4 & - & 3x_5 & = & 16 \\ 3x_1 & + & 15x_2 & - & 2x_3 & - & 4x_4 & - & 2x_5 & = & 56 \\ -2x_1 & - & 10x_2 & & & + & x_4 & - & 10x_5 & = & -46 \\ 4x_1 & + & 20x_2 & - & x_3 & - & 3x_4 & + & 11x_5 & = & 86. \end{array}$$

Furthermore, we showed that the solution set is

$$\left\{ \begin{pmatrix} 24 \\ 0 \\ 4 \\ 2 \\ 0 \end{pmatrix} + s \begin{pmatrix} -5 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} -4 \\ 0 \\ -11 \\ 2 \\ 1 \end{pmatrix} : s, t \in \mathbb{R} \right\}.$$

Notice then that

$$\begin{pmatrix} 24 \\ 0 \\ 4 \\ 2 \\ 0 \end{pmatrix}$$

is a solution to the system because we can obtain it when $s = 0$ and $t = 0$. However, the vector

$$\begin{pmatrix} -5 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

does not appear to be in the solution set, because it does not look like we can obtain it by choosing s and t appropriately. In fact, this vector is not a solution to the original system, but surprisingly it is a solution to:

$$\begin{array}{rrrrrrrrcl} x_1 & + & 5x_2 & - & x_3 & - & 2x_4 & - & 3x_5 & = & 0 \\ 3x_1 & + & 15x_2 & - & 2x_3 & - & 4x_4 & - & 2x_5 & = & 0 \\ -2x_1 & - & 10x_2 & & & + & x_4 & - & 10x_5 & = & 0 \\ 4x_1 & + & 20x_2 & - & x_3 & - & 3x_4 & + & 11x_5 & = & 0. \end{array}$$

In other words, the vector that appears multiplied by one of the parameters is in fact a solution to a different, but related, system obtained by turning all the constants on the right to 0. The same is true for the vector that is being scaled by t (check it!). To give a name to this situation, we call a linear system where the constants on the right are all 0 a *homogeneous* linear system.

In order to interpret this in our new language, let $T: \mathbb{R}^5 \rightarrow \mathbb{R}^4$ be the linear transformation given by

$$T \left(\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{pmatrix} \right) = \begin{pmatrix} x_1 + 5x_2 - x_3 - 2x_4 - 3x_5 \\ 3x_1 + 15x_2 - 2x_3 - 4x_4 - 2x_5 \\ -2x_1 - 10x_2 + x_4 - 10x_5 \\ 4x_1 + 20x_2 - x_3 - 3x_4 + 11x_5 \end{pmatrix}.$$

Since the two vectors multiplied by parameters are solutions to the corresponding homogeneous linear system, we have

$$\begin{pmatrix} -5 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} \in \text{Null}(T) \quad \text{and} \quad \begin{pmatrix} -4 \\ 0 \\ -11 \\ 2 \\ 1 \end{pmatrix} \in \text{Null}(T).$$

Now since $\text{Null}(T)$ is a subspace of \mathbb{R}^5 , it follows that

$$s \begin{pmatrix} -5 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} -4 \\ 0 \\ -11 \\ 2 \\ 1 \end{pmatrix} \in \text{Null}(T)$$

for all $s, t \in \mathbb{R}$. Now we also have

$$\begin{pmatrix} 24 \\ 0 \\ 4 \\ 2 \\ 0 \end{pmatrix} \notin \text{Null}(T)$$

because it is a solution to our original system, not the homogeneous system. However, whenever we add an element of $\text{Null}(T)$ to it, we obtain another solution. The fundamental reason for this is the following result.

Proposition 5.2.6. *Let V and W be vector spaces, and let $T: V \rightarrow W$ be a linear transformation.*

1. *If $T(\vec{v}) = \vec{b}$ and $\vec{z} \in \text{Null}(T)$, then $T(\vec{v} + \vec{z}) = \vec{b}$.*
2. *If $T(\vec{v}_1) = \vec{b}$ and $T(\vec{v}_2) = \vec{b}$, then $\vec{v}_1 - \vec{v}_2 \in \text{Null}(T)$.*

Proof. For the first, simply notice that since $\vec{z} \in \text{Null}(T)$, we have $T(\vec{z}) = \vec{0}$, and hence

$$\begin{aligned} T(\vec{v} + \vec{z}) &= T(\vec{v}) + T(\vec{z}) \\ &= \vec{b} + \vec{0} \\ &= \vec{b}. \end{aligned}$$

For the second, we have

$$\begin{aligned} T(\vec{v}_1 - \vec{v}_2) &= T(\vec{v}_1) - T(\vec{v}_2) \\ &= \vec{b} - \vec{b} \\ &= \vec{0}, \end{aligned}$$

so $\vec{v}_1 - \vec{v}_2 \in \text{Null}(T)$. □

Corollary 5.2.7. *Let $T: V \rightarrow W$ be a linear transformation, and let $\vec{b} \in W$. Suppose that $\vec{v} \in V$ is such that $T(\vec{v}) = \vec{b}$. We have*

$$\{\vec{x} \in V : T(\vec{x}) = \vec{b}\} = \{\vec{v} + \vec{z} : \vec{z} \in \text{Null}(T)\}.$$

Proof. Do a double containment proof using the previous result. □

In other words, if we have an example of an element $\vec{v} \in V$ that hits \vec{b} , then we can obtain the set of *all* elements of the domain that hit \vec{b} by simply adding all elements of the null space to \vec{v} . This is a powerful idea. In particular, to find all solutions to a linear system, it suffices to find just one particular solution, and then add to it all solutions to the corresponding homogeneous linear system (with 0's on the right). More generally, you will see consequences of these past two results in Differential Equations, where you will work in the vector space of functions rather than in \mathbb{R}^n .

Let's return to computing $\text{Null}(T)$ and $\text{range}(T)$, but now let's work on an example where one of the vector spaces is not some \mathbb{R}^n . Define $T: \mathcal{P}_2 \rightarrow \mathbb{R}$ by letting

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

Let first think about $\text{range}(T)$. We can obtain elements of $\text{range}(T)$ by simply plugging inputs into T and recording the outputs. For example, since

$$\int_0^1 x \, dx = \frac{1}{2},$$

we have $\frac{1}{2} \in \text{range}(T)$. In fact, we also have $\frac{1}{2} \in \text{range}(T)$ because

$$\int_0^1 \frac{1}{2} \, dx = \frac{1}{2}.$$

Using this idea, we can immediately show that $\text{range}(T) = \mathbb{R}$, because for any $a \in \mathbb{R}$, if we consider the constant polynomial $a \in \mathcal{P}_2$, we have

$$\int_0^1 a \, dx = (ax)|_0^1 = a.$$

Therefore, $\text{range}(T) = \mathbb{R}$. Notice that $\text{range}(T) = \mathbb{R}$ is a 1-dimensional vector space because (1) is a basis for \mathbb{R} .

For $\text{Null}(T)$, we start by finding some examples of elements that are sent to zero. Of course, the zero polynomial is in $\text{Null}(T)$. For more interesting examples, notice that

$$\begin{aligned} \int_0^1 \left(x - \frac{1}{2}\right) \, dx &= \left(\frac{x^2}{2} - \frac{x}{2}\right)_0^1 = \frac{1}{2} - \frac{1}{2} = 0 \\ \int_0^1 \left(x^2 - \frac{1}{3}\right) \, dx &= \left(\frac{x^3}{3} - \frac{x}{3}\right)_0^1 = \frac{1}{3} - \frac{1}{3} = 0 \end{aligned}$$

so $x - \frac{1}{2}, x^2 - \frac{1}{3} \in \text{Null}(T)$. We now show that $(x - \frac{1}{2}, x^2 - \frac{1}{3})$ is a basis for $\text{Null}(T)$. We first check that $(x - \frac{1}{2}, x^2 - \frac{1}{3})$ is linearly independent. Suppose that $c_1, c_2 \in \mathbb{R}$ satisfy

$$c_1 \left(x - \frac{1}{2}\right) + c_2 \left(x^2 - \frac{1}{3}\right) = 0$$

for all $x \in \mathbb{R}$. Multiplying out and collecting terms, we see that

$$c_2 x^2 + c_1 x + \left(-\frac{c_2}{3} - \frac{c_1}{2}\right) = 0$$

for all $x \in \mathbb{R}$. Equating coefficients on each side, we conclude that both $c_1 = 0$ and $c_2 = 0$. Therefore, $(x - \frac{1}{2}, x^2 - \frac{1}{3})$ is linearly independent.

We now show that $\text{Span}(x - \frac{1}{2}, x^2 - \frac{1}{3}) = \text{Null}(T)$. Since we know from above that both $x - \frac{1}{2} \in \text{Null}(T)$ and $x^2 - \frac{1}{3} \in \text{Null}(T)$, and we know that $\text{Null}(T)$ is a subspace of \mathcal{P}_2 , we conclude that $\text{Span}(x - \frac{1}{2}, x^2 - \frac{1}{3}) \subseteq \text{Null}(T)$.

Let $p \in \text{Null}(T)$ and fix $a_0, a_1, a_2 \in \mathbb{R}$ with $p(x) = a_0 + a_1x + a_2x^2$. We then have

$$\begin{aligned} 0 &= T(p(x)) \\ &= \int_0^1 p(x) \, dx \\ &= \int_0^1 (a_0 + a_1x + a_2x^2) \, dx \\ &= \left(a_0x + \frac{a_1}{2}x^2 + \frac{a_2}{3}x^3 \right)_0^1 \\ &= a_0 + \frac{a_1}{2} + \frac{a_2}{3}. \end{aligned}$$

It follows that $a_0 = -\frac{a_1}{2} - \frac{a_2}{3}$, so

$$\begin{aligned} p(x) &= a_0 + a_1x + a_2x^2 \\ &= -\frac{a_1}{2} - \frac{a_2}{3} + a_1x + a_2x^2 \\ &= a_1x - \frac{a_1}{2} + a_2x^2 - \frac{a_2}{3} \\ &= a_1 \cdot \left(x - \frac{1}{2} \right) + a_2 \cdot \left(x^2 - \frac{1}{3} \right). \end{aligned}$$

We have written $p(x)$ as a linear combination of $x - \frac{1}{2}$ and $x^2 - \frac{1}{3}$, so $p(x) \in \text{Span}(x - \frac{1}{2}, x^2 - \frac{1}{3})$. Since $p(x) \in \text{Null}(T)$ was arbitrary, it follows that $\text{Null}(T) \subseteq \text{Span}(x - \frac{1}{2}, x^2 - \frac{1}{3})$. Combining this with the above, it follows that $\text{Null}(T) = \text{Span}(x - \frac{1}{2}, x^2 - \frac{1}{3})$, and hence $(x - \frac{1}{2}, x^2 - \frac{1}{3})$ is a basis for $\text{Null}(T)$.

Now that we have determined $\text{Null}(T)$, we can use it in conjunction with Corollary 5.2.7 to find *all* polynomials in \mathcal{P}_2 that integrate (over $[0, 1]$) to a certain value. For example, one polynomial that integrates (over $[0, 1]$) to 5 is the constant polynomial 5. Using Corollary 5.2.7, the set of all polynomials in \mathbb{P}_2 that integrate to 5 is

$$\{5 + p : p \in \text{Null}(T)\} = \left\{ 5 + a \cdot \left(x - \frac{1}{2} \right) + b \cdot \left(x^2 - \frac{1}{3} \right) : a, b \in \mathbb{R} \right\}.$$

There is another curious connection happening here. We have shown that $\text{range}(T)$ has dimension 1 by finding the basis (1), and that $\text{Null}(T)$ has dimension 2 by finding the basis $(x - \frac{1}{2}, x^2 - \frac{1}{3})$. Now we also know that $\dim(\mathcal{P}_2) = 3$. Thus, our transformation “killed off” two of the three dimensions in \mathcal{P}_2 , and left a 1-dimensional range. Looking back at Theorem 2.7.3, we see a similar balancing between the “sizes” of $\text{Null}(T)$ and $\text{range}(T)$. We now prove a very general result that says that this balancing act is always true.

Definition 5.2.8. Let $T: V \rightarrow W$ be a linear transformation of finite-dimensional vector spaces.

- We define the *rank* of T , denoted $\text{rank}(T)$, to be $\dim(\text{range}(T))$.
- We define the *nullity* of T , denoted $\text{nullity}(T)$, to be $\dim(\text{Null}(T))$.

If $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$, then we can quickly determine $\text{rank}(T)$.

Corollary 5.2.9. Let $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear transformation, and let B be an echelon form of $[T]$. We then have that $\text{rank}(T)$ is the number of leading entries in B .

Proof. Let $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n \in \mathbb{R}^m$ be the columns of $[T]$. Corollary 5.2.5 tells us that we obtain a basis for $\text{rank}(T)$ by taking the sequence consisting only of those \vec{u}_i such that the i^{th} column of B has a leading entry. The number of vectors in this basis is the number of leading entries in B . \square

We now prove one of the most important theorems in Linear Algebra.

Theorem 5.2.10 (Rank-Nullity Theorem). *Let $T: V \rightarrow W$ be a linear transformation with V and W finite-dimensional vector spaces. We then have that $\text{rank}(T) + \text{nullity}(T) = \dim(V)$.*

Proof. Since $\text{Null}(T)$ is a subspace of V , we can fix a basis $(\vec{z}_1, \vec{z}_2, \dots, \vec{z}_k)$ of $\text{Null}(T)$. Now this sequence of vectors is linearly independent in V , so we can extend this to a basis $(\vec{z}_1, \vec{z}_2, \dots, \vec{z}_k, \vec{u}_1, \vec{u}_2, \dots, \vec{u}_m)$ of V by Proposition 4.6.10. We claim that $(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_m))$ is a basis of $\text{range}(T)$. To do this, we check three things:

- $\text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_m)) \subseteq \text{range}(T)$: Let $\vec{w} \in \text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_m))$ be arbitrary. Fix $c_1, c_2, \dots, c_m \in \mathbb{R}$ with

$$\vec{w} = c_1 \cdot T(\vec{u}_1) + c_2 \cdot T(\vec{u}_2) + \dots + c_m \cdot T(\vec{u}_m).$$

Since T is a linear transformation, we have

$$\vec{w} = T(c_1 \vec{u}_1 + c_2 \vec{u}_2 + \dots + c_m \vec{u}_m).$$

Since $c_1 \vec{u}_1 + c_2 \vec{u}_2 + \dots + c_m \vec{u}_m \in V$, it follows that $\vec{w} \in \text{range}(T)$.

- $\text{range}(T) \subseteq \text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_m))$: Let $\vec{w} \in \text{range}(T)$ be arbitrary. Fix $\vec{v} \in V$ with $T(\vec{v}) = \vec{w}$. Since $(\vec{z}_1, \vec{z}_2, \dots, \vec{z}_k, \vec{u}_1, \vec{u}_2, \dots, \vec{u}_m)$ is a basis of V , we can fix $c_1, c_2, \dots, c_k, d_1, d_2, \dots, d_m \in \mathbb{R}$ with

$$\vec{v} = c_1 \vec{z}_1 + c_2 \vec{z}_2 + \dots + c_k \vec{z}_k + d_1 \vec{u}_1 + d_2 \vec{u}_2 + \dots + d_m \vec{u}_m.$$

We then have

$$\begin{aligned} \vec{w} &= T(\vec{v}) \\ &= T(c_1 \vec{z}_1 + c_2 \vec{z}_2 + \dots + c_k \vec{z}_k + d_1 \vec{u}_1 + d_2 \vec{u}_2 + \dots + d_m \vec{u}_m) \\ &= c_1 \cdot T(\vec{z}_1) + c_2 \cdot T(\vec{z}_2) + \dots + c_k \cdot T(\vec{z}_k) + d_1 \cdot T(\vec{u}_1) + d_2 \cdot T(\vec{u}_2) + \dots + d_m \cdot T(\vec{u}_m) \\ &= c_1 \cdot \vec{0} + c_2 \cdot \vec{0} + \dots + c_k \cdot \vec{0} + d_1 \cdot T(\vec{u}_1) + d_2 \cdot T(\vec{u}_2) + \dots + d_m \cdot T(\vec{u}_m) \\ &= d_1 \cdot T(\vec{u}_1) + d_2 \cdot T(\vec{u}_2) + \dots + d_m \cdot T(\vec{u}_m) \end{aligned}$$

so

$$\vec{w} \in \text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_m)).$$

- $(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_m))$ is linearly independent: Let $d_1, d_2, \dots, d_m \in \mathbb{R}$ be arbitrary with

$$d_1 \cdot T(\vec{u}_1) + d_2 \cdot T(\vec{u}_2) + \dots + d_m \cdot T(\vec{u}_m) = \vec{0}.$$

We then have

$$T(d_1 \vec{u}_1 + d_2 \vec{u}_2 + \dots + d_m \vec{u}_m) = \vec{0},$$

so

$$d_1 \vec{u}_1 + d_2 \vec{u}_2 + \dots + d_m \vec{u}_m \in \text{Null}(T).$$

Since $(\vec{z}_1, \vec{z}_2, \dots, \vec{z}_k)$ is a basis of $\text{Null}(T)$, we can fix $c_1, c_2, \dots, c_k \in \mathbb{R}$ with

$$d_1 \vec{u}_1 + d_2 \vec{u}_2 + \dots + d_m \vec{u}_m = c_1 \vec{z}_1 + c_2 \vec{z}_2 + \dots + c_k \vec{z}_k.$$

We then have

$$(-c_1) \vec{z}_1 + (-c_2) \vec{z}_2 + \dots + (-c_k) \vec{z}_k + d_1 \vec{u}_1 + d_2 \vec{u}_2 + \dots + d_m \vec{u}_m = \vec{0}.$$

Since $(\vec{z}_1, \vec{z}_2, \dots, \vec{z}_k, \vec{u}_1, \vec{u}_2, \dots, \vec{u}_m)$ is linearly independent, it follows that $-c_i = 0$ for all i and $d_j = 0$ for all j . In particular, we have shown that $d_1 = d_2 = \dots = d_m = 0$.

Putting it all together, these three facts allow us to conclude that $(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_m))$ is a basis for $\text{range}(T)$, so $\text{rank}(T) = m$. Since $\text{nullity}(T) = k$ (because our basis for $\text{Null}(T)$ had k elements) and $\dim(V) = k + m$ (because our basis of V had $k + m$ elements), we conclude that $\text{rank}(T) + \text{nullity}(T) = \dim(V)$. \square

One reason that we care about $\text{Null}(T)$ is that we can use it to determine if a linear transformation is injective.

Proposition 5.2.11. *Let $T: V \rightarrow W$ be a linear transformation. We have that T is injective if and only if $\text{Null}(T) = \{\vec{0}_V\}$.*

Proof. Notice that the proof of Proposition 2.7.4 works in general, even if V and W are not \mathbb{R}^2 . \square

Given a linear transformation $T: V \rightarrow W$, we've introduced both $\text{rank}(T)$ and $\text{nullity}(T)$, which are measures of the “size” of $\text{range}(T)$ and $\text{Null}(T)$, respectively. Furthermore, these latter two sets are directly related to understanding whether T is surjective or injective (since T is surjective if and only if $\text{range}(T) = W$, and T is injective if and only if $\text{Null}(T) = \{\vec{0}\}$). Connected with this is the idea of whether T is invertible because, after all, we know that a function has an inverse exactly when it is bijective. With these thoughts in mind, we turn our attention to the related question of asking when a matrix has an inverse.

We begin with a simple, but at first unsettling, example. Notice that we have

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

Looking at the product on the left, we multiplied a 2×3 matrix by a 3×2 matrix, and ended with the 2×2 identity matrix. At this point, it might seem natural to call these two matrices inverses of each other. However, when we multiply in the other order, we obtain

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

First notice that when we multiply in this order, we are multiplying a 3×2 matrix by a 2×3 matrix, so the result is a 3×3 matrix, and hence is a different size than the matrix that results from multiplying in the original order. Furthermore, perhaps surprisingly, the latter product did *not* result in the 3×3 identity matrix. Thus, although the matrix product in one order suggested that the two matrices might be inverses of each other, the matrix product in the other order told us otherwise.

To understand what is going on beneath the computations, it is best to think in terms of linear transformations coded by the two matrices. Since

$$B = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}$$

is a 3×2 matrix, we can fix the unique linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ whose standard matrix is B . Notice then that

$$T\left(\begin{pmatrix} x \\ y \end{pmatrix}\right) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x \\ y \\ 0 \end{pmatrix}$$

for all $x, y \in \mathbb{R}$. In other words, T takes a point in \mathbb{R}^2 and turns it into the point in \mathbb{R}^3 by simply tacking on 0 as a last component. Similarly, since

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

is a 2×3 matrix, we can fix the unique linear transformation $S: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ whose standard matrix is A . Notice then that

$$S \left(\begin{pmatrix} x \\ y \\ z \end{pmatrix} \right) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} x \\ y \end{pmatrix}$$

for all $x, y, z \in \mathbb{R}$. In other words, S takes a point in \mathbb{R}^3 and turns it into a point in \mathbb{R}^2 by deleting the last coordinate, i.e. S acts by projecting 3-dimensional space onto the 2-dimensional Cartesian plane. Now since matrix multiplication corresponds to function composition, the matrix product

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

gives the standard matrix of $S \circ T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$. In terms of these functions, the composition acts as follows:

$$\begin{pmatrix} x \\ y \end{pmatrix} \xrightarrow{T} \begin{pmatrix} x \\ y \\ 0 \end{pmatrix} \xrightarrow{S} \begin{pmatrix} x \\ y \end{pmatrix}.$$

In other words, if we first throw a point from \mathbb{R}^2 into \mathbb{R}^3 by adding a 0 at the end, and then we remove the last coordinate, we arrive where we started.

In contrast, the matrix product

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

gives the standard matrix of $T \circ S: \mathbb{R}^3 \rightarrow \mathbb{R}^3$. In terms of these functions, the composition acts as follows:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \xrightarrow{S} \begin{pmatrix} x \\ y \end{pmatrix} \xrightarrow{T} \begin{pmatrix} x \\ y \\ 0 \end{pmatrix}.$$

In other words, if we take a point in \mathbb{R}^3 , then rip off the last coordinate, and follow that up with adding 0 on as a new last coordinate, we do not always end up where we started.

Now that we remember that a function has an inverse if and only if it is bijective, we can make some progress in understanding matrix inverses by thinking about restrictions on when a linear transformation can be bijective. We have the following two fundamental facts.

Proposition 5.2.12. *Let V and W be vector spaces. Let $T: V \rightarrow W$ be an injective linear transformation and let $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ be a linearly independent sequence in V . We then have that $(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n))$ is a linearly independent sequence in W .*

Proof. Exercise. □

Proposition 5.2.13. *Let V and W be vector spaces. Let $T: V \rightarrow W$ be a surjective linear transformation and assume that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$. We then have that $\text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n)) = W$.*

Proof. Notice first that $\text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n)) \subseteq W$ immediately from the definition because $T(\vec{u}_i) \in W$ for all i and W is a vector space. We now show that $W \subseteq \text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n))$. Let $\vec{w} \in W$ be arbitrary. Since T is surjective, we can fix $\vec{v} \in V$ with $T(\vec{v}) = \vec{w}$. Now $\vec{v} \in V$ and we know that $\text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n) = V$, so we conclude that $\vec{v} \in \text{Span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$. Therefore, we can fix $c_1, c_2 \in \mathbb{R}$ with

$$\vec{v} = c_1 \cdot \vec{u}_1 + c_2 \cdot \vec{u}_2 + \dots + c_n \cdot \vec{u}_n.$$

Applying T to both sides and using the fact that T is a linear transformation, we conclude that

$$\begin{aligned} T(\vec{v}) &= T(c_1 \cdot \vec{u}_1 + c_2 \cdot \vec{u}_2 + \cdots + c_n \cdot \vec{u}_n) \\ &= c_1 \cdot T(\vec{u}_1) + c_2 \cdot T(\vec{u}_2) + \cdots + c_n \cdot T(\vec{u}_n). \end{aligned}$$

Since $T(\vec{v}) = \vec{w}$, it follows that

$$\vec{w} = c_1 \cdot T(\vec{u}_1) + c_2 \cdot T(\vec{u}_2) + \cdots + c_n \cdot T(\vec{u}_n).$$

Since $c_1, c_2, \dots, c_n \in \mathbb{R}$, we conclude that $\vec{w} \in \text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n))$. Since $\vec{w} \in W$ was arbitrary, it follows that $W \subseteq \text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n))$.

We have shown that both $\text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n)) \subseteq W$ and $W \subseteq \text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n))$ are true, so it follows that $\text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n)) = W$. \square

Corollary 5.2.14. *Let V and W be finite-dimensional vector spaces, and let $n = \dim(V)$ and $m = \dim(W)$. Let $T: V \rightarrow W$ be a linear transformation.*

1. *If T is injective, then $n \leq m$.*
2. *If T is surjective, then $m \leq n$.*
3. *If T is bijective, then $m = n$.*

Proof.

1. Suppose that $T: V \rightarrow W$ is injective. Since $\dim(V) = n$, we can fix a basis $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ of V . By Proposition 5.2.12, we then have that $(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n))$ is a linearly independent sequence of W . Now $\dim(W) = m$, so W has a spanning sequence (in fact a basis) with m elements. Using Corollary 4.6.3, we conclude that $n \leq m$.
2. Suppose that $T: V \rightarrow W$ is surjective. Since $\dim(V) = n$, we can fix a basis $(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ of V . By Proposition 5.2.13, we then have that $\text{Span}(T(\vec{u}_1), T(\vec{u}_2), \dots, T(\vec{u}_n)) = W$. Now $\dim(W) = m$, so W has a linearly independent sequence (in fact a basis) with m elements. Using Corollary 4.6.3, we conclude that $m \leq n$.
3. This follows immediately from the first two parts.

\square

In the case where $T: V \rightarrow W$ is a bijective linear transformation, the inverse function exists, and happens to be a linear transformation as well.

Proposition 5.2.15. *Suppose that $T: V \rightarrow W$ is a bijective linear transformation. We then have that the function $T^{-1}: W \rightarrow V$ is a linear transformation.*

Proof. See the proof of Proposition 2.7.6, which generalizes here. \square

We now turn our attention to thinking about the inverse of a matrix A . Since a matrix can be viewed as coding a linear transformation, the matrix should have an inverse if and only if the corresponding linear transformation has an inverse, which is to say when the corresponding linear transformation is bijective. If A is an $m \times n$ matrix, then the corresponding linear transformation T would have domain equal to an n -dimensional vector space, and codomain equal to an m -dimensional vector space. But we know from above that the only case where such a T could be bijective is when $m = n$. Thus, we can restrict attention to the case of $m = n$, i.e. to the situation where A is a square matrix.

Definition 5.2.16. Suppose that A is an $n \times n$ matrix. An inverse of A is an $n \times n$ matrix B with both $BA = I_n$ and $AB = I_n$. We say that A is invertible if it has an inverse.

Proposition 5.2.17. Let A be an $n \times n$ matrix, and let B be an echelon form of A . We then have that A is invertible if and only if every row and every column of B has a leading entry.

Proof. Since A is an $n \times n$ matrix, we can fix the unique linear transformation $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$ with $[T] = A$. Notice that A has an inverse if and only if T is bijective. By Corollary 5.2.4, we know that T is surjective if and only if every row of B has a leading entry. By Proposition 5.2.11, we know that T is injective if and only if $\text{Null}(T) = \{\vec{0}\}$, which is if and only if every column of B has a leading entry. Therefore, A has an inverse if and only if every row and every column of B has a leading entry. \square

Okay, we now have a criterion to determine if an $n \times n$ matrix A has an inverse. We can compute an echelon form B of A , and determine if every row and column of B has a leading entry. Since A is a square matrix, this is the same as saying that the leading entries of B appear on the diagonal. For example, suppose that

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

Applying elementary row operations, we obtain

$$\begin{aligned} \begin{pmatrix} 1 & 0 & -2 \\ -3 & 1 & 4 \\ 2 & -3 & 4 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 0 & -2 \\ 0 & 1 & -2 \\ 0 & -3 & 8 \end{pmatrix} && \begin{matrix} (3R_1 + R_2) \\ (-2R_2 + R_3) \end{matrix} \\ &\rightarrow \begin{pmatrix} 1 & 0 & -2 \\ 0 & 1 & -2 \\ 0 & 0 & 2 \end{pmatrix} && (3R_2 + R_3). \end{aligned}$$

Since there is a leading entry in each row and column, the columns of A form a basis for \mathbb{R}^3 , and hence A is invertible.

Once we check that an $n \times n$ matrix A is invertible, how can we determine A^{-1} ? The idea is as follows. Fix the unique linear transformation $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$ with $[T] = A$. We then have that A^{-1} will be the standard matrix of T^{-1} , i.e. that $A^{-1} = [T^{-1}]$. Recall that the $[T^{-1}]$ is the matrix whose j^{th} column is $T^{-1}(\vec{e}_j)$. Thus, to determine the first column of A^{-1} , we would want to find $T^{-1}(\vec{e}_1)$, so we would want to find the unique $\vec{v} \in \mathbb{R}^n$ with $T(\vec{v}) = \vec{e}_1$. In other words, we want to solve $A\vec{v} = \vec{e}_1$. Now in order to do this, we can look at the augmented matrix obtained by adding the column \vec{e}_1 onto the end of A , and solve the corresponding linear system. The unique solution will be the first column of A^{-1} . We can do the whole process again by adding on \vec{e}_2 as an augmented column in order to find the second column of A^{-1} , etc. However, instead of doing many computations by repeating this process n times, we can do them all at once by tacking on n different augmented columns, one for each \vec{e}_j .

In our above example where

$$A = \begin{pmatrix} 1 & 0 & -2 \\ -3 & 1 & 4 \\ 2 & -3 & 4 \end{pmatrix}$$

we would want to solve three different linear systems arising from augmenting with each of \vec{e}_1 , \vec{e}_2 , and \vec{e}_3 . Thus, we form the matrix

$$\begin{pmatrix} 1 & 0 & -2 & 1 & 0 & 0 \\ -3 & 1 & 4 & 0 & 1 & 0 \\ 2 & -3 & 4 & 0 & 0 & 1 \end{pmatrix}.$$

Performing elementary row operations, we obtain

$$\begin{aligned} \begin{pmatrix} 1 & 0 & -2 & 1 & 0 & 0 \\ -3 & 1 & 4 & 0 & 1 & 0 \\ 2 & -3 & 4 & 0 & 0 & 1 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 0 & -2 & 1 & 0 & 0 \\ 0 & 1 & -2 & 3 & 1 & 0 \\ 0 & -3 & 8 & -2 & 0 & 1 \end{pmatrix} &\begin{matrix} (3R_1 + R_2) \\ (-2R_2 + R_3) \end{matrix} \\ &\rightarrow \begin{pmatrix} 1 & 0 & -2 & 1 & 0 & 0 \\ 0 & 1 & -2 & 3 & 1 & 0 \\ 0 & 0 & 2 & 7 & 3 & 1 \end{pmatrix}. \end{aligned}$$

At this point, we could stop and solve each of the three systems using back-substitution. However, there is a better way. Let's keep row reducing the matrix so the solutions will stare us in the face. In other words, let's keep going until we arrive at the identity matrix on the left.

$$\begin{aligned} &\rightarrow \begin{pmatrix} 1 & 0 & 0 & 8 & 3 & 1 \\ 0 & 1 & 0 & 10 & 4 & 1 \\ 0 & 0 & 2 & 7 & 3 & 1 \end{pmatrix} &\begin{matrix} (R_3 + R_1) \\ (R_3 + R_2) \end{matrix} \\ &\rightarrow \begin{pmatrix} 1 & 0 & 0 & 8 & 3 & 1 \\ 0 & 1 & 0 & 10 & 4 & 1 \\ 0 & 0 & 1 & \frac{7}{2} & \frac{3}{2} & \frac{1}{2} \end{pmatrix} &(\frac{1}{2} \cdot R_3). \end{aligned}$$

Notice that we can immediately read off the solutions to each of the three linear systems encoded in the matrix by simply reading off the last columns. It follows that A^{-1} is the matrix on the right-hand side here, i.e. that

$$A^{-1} = \begin{pmatrix} 8 & 3 & 1 \\ 10 & 4 & 1 \\ \frac{7}{2} & \frac{3}{2} & \frac{1}{2} \end{pmatrix}.$$

In general, here is an algorithm that, given an $n \times n$ matrix A , determines if A invertible and if so then calculates the inverse of A :

1. Form the $n \times 2n$ matrix obtained by augmenting the matrix A with the $n \times n$ identity matrix.
2. Perform elementary row operations until the $n \times n$ matrix on the left is in echelon form.
3. If the resulting $n \times n$ matrix on the left (which is in echelon form) does not have a leading entry in every row (equivalently every column), then the linear transformation coded by A is not bijective, so A does not have an inverse. Otherwise, continue to the next step.
4. Suppose then that the resulting $n \times n$ matrix on the left does have a leading entry in every row. Since this is a square matrix, there is a leading entry in each diagonal position. We can continue performing row operations to eliminate nonzero entries above the diagonal, and to make the diagonal entries equal to 1.
5. Once we have completed these row operations, we will have the $n \times n$ identity matrix on the left. The right n columns will contain the matrix A^{-1} .

5.3 Determinants

We looked at the theory of determinants in \mathbb{R}^2 back in Section 3.4. In that context, we worked to develop a theory of “signed area” of parallelograms by starting with basic properties that this concept should have, and then using these properties to derive a formula for the unique function satisfying them. We now try to generalize those ideas to higher dimensions. Given n vectors $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n \in \mathbb{R}^n$, the idea is to let $f(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n)$

be the signed n -dimensional volume of the “generalized parallelogram” formed by $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n$. For example, given 3 vectors in \mathbb{R}^3 , we obtain a “slanty-boxy-thingy” called a parallelepiped, and $f(\vec{v}_1, \vec{v}_2, \vec{v}_3)$ will be the signed volume of this object depending on the orientation of the three vectors. What does the sign mean in this context? Just as the span of one nonzero vector in \mathbb{R}^2 is a line through the origin which breaks \mathbb{R}^2 in two pieces, the span of two linearly independent vectors in \mathbb{R}^3 is a plane through the origin which breaks \mathbb{R}^3 into two pieces. Thus, the third vector in our list could be on one of the two “sides” of this plane, and the sign is there to signify which side. We’ll assign positive volume if it is on one side, and negative volume if it is on the other (this is the origin of the right-hand and left-hand rules that you may have seen in multivariable Calculus or in Physics). More generally, you might expect that the span of three linearly independent vectors in \mathbb{R}^4 will break \mathbb{R}^4 into two pieces, etc.

Definition 5.3.1. A function $f: \mathbb{R}^n \times \mathbb{R}^n \times \dots \times \mathbb{R}^n \rightarrow \mathbb{R}$ (where there are n copies of \mathbb{R}^n on the left of the arrow) is called a determinant function if it satisfies the following:

1. $f(\vec{e}_1, \dots, \vec{e}_n) = 1$.
2. If there exists $i \neq j$ with $\vec{v}_i = \vec{v}_j$, then $f(\vec{v}_1, \dots, \vec{v}_n) = 0$.
3. $f(\vec{v}_1, \dots, c \cdot \vec{v}_k, \dots, \vec{v}_n) = c \cdot f(\vec{v}_1, \dots, \vec{v}_k, \dots, \vec{v}_n)$ for all $\vec{v}_1, \dots, \vec{v}_n \in \mathbb{R}^n$, all $c \in \mathbb{R}$ and all k with $1 \leq k \leq n$.
4. We have

$$\begin{aligned} f(\vec{v}_1, \dots, \vec{v}_{k-1}, \vec{u} + \vec{w}, \vec{v}_{k+1}, \dots, \vec{v}_n) = \\ f(\vec{v}_1, \dots, \vec{v}_{k-1}, \vec{u}, \vec{v}_{k+1}, \dots, \vec{v}_n) + f(\vec{v}_1, \dots, \vec{v}_{k-1}, \vec{w}, \vec{v}_{k+1}, \dots, \vec{v}_n) \end{aligned}$$

for all $\vec{v}_1, \dots, \vec{v}_{k-1}, \vec{v}_{k+1}, \dots, \vec{v}_n, \vec{u}, \vec{w} \in \mathbb{R}^n$.

To capture properties 3 and 4 in words, we say that f is linear in each component.

In Section 3.4, we proved that the unique function satisfying these properties in the case where $n = 2$ is given by

$$f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) = ad - bc.$$

Furthermore, it is straightforward to check there is a unique function that works when $n = 1$, which is given by $f(a) = a$. In fact, we have the following fundamental result. Rather than prove it (which is within our grasp, but takes several days), we will simply assert that it is true.

Theorem 5.3.2. For each $n \in \mathbb{N}^+$, there is a unique function f satisfying the above properties. We call this function \det .

Using this result, we now work to develop more properties of \det , and eventually use them to provide us with methods to compute it. Alternatively, most of the results that we prove below can be interpreted as providing additional properties that any determinant function f must satisfy, and hence can be viewed as steps toward the proof of uniqueness.

Proposition 5.3.3. Let $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n \in \mathbb{R}^n$. If $i < j$, then

$$\det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n) = -\det(\vec{v}_1, \dots, \vec{v}_j, \dots, \vec{v}_i, \dots, \vec{v}_n).$$

In other words, swapping two vectors in the input of \det changes the sign.

Proof. Let $\vec{v}_1, \dots, \vec{v}_n$ be arbitrary and suppose that $i \neq j$. By property 2, we know that

$$\det(\vec{v}_1, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_n) = 0.$$

On the other hand, using property 4, we have

$$\begin{aligned} \det(\vec{v}_1, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_n) \\ = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_n) + \det(\vec{v}_1, \dots, \vec{v}_j, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_n). \end{aligned}$$

Now using property 4 and property 2, we have

$$\begin{aligned} \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_n) \\ = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_i, \dots, \vec{v}_n) + \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n) \\ = 0 + \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n) \\ = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n) \end{aligned}$$

and also

$$\begin{aligned} \det(\vec{v}_1, \dots, \vec{v}_j, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_n) \\ = \det(\vec{v}_1, \dots, \vec{v}_j, \dots, \vec{v}_i, \dots, \vec{v}_n) + \det(\vec{v}_1, \dots, \vec{v}_j, \dots, \vec{v}_j, \dots, \vec{v}_n) \\ = \det(\vec{v}_1, \dots, \vec{v}_j, \dots, \vec{v}_i, \dots, \vec{v}_n) + 0 \\ = \det(\vec{v}_1, \dots, \vec{v}_j, \dots, \vec{v}_i, \dots, \vec{v}_n). \end{aligned}$$

Plugging these into the above equation, we conclude that

$$\begin{aligned} \det(\vec{v}_1, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_n) \\ = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n) + \det(\vec{v}_1, \dots, \vec{v}_j, \dots, \vec{v}_i, \dots, \vec{v}_n), \end{aligned}$$

Now using the fact that

$$\det(\vec{v}_1, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_i + \vec{v}_j, \dots, \vec{v}_n) = 0$$

we conclude that

$$0 = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n) + \det(\vec{v}_1, \dots, \vec{v}_j, \dots, \vec{v}_i, \dots, \vec{v}_n).$$

□

Proposition 5.3.4. Let $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n \in \mathbb{R}^n$ and let $c \in \mathbb{R}$. If $i < j$ and $c \in \mathbb{R}$, then

$$\det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j + c\vec{v}_i, \dots, \vec{v}_n) = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n)$$

and

$$\det(\vec{v}_1, \dots, \vec{v}_i + c\vec{v}_j, \dots, \vec{v}_j, \dots, \vec{v}_n) = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n).$$

Proof. We have

$$\begin{aligned} \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j + c\vec{v}_i, \dots, \vec{v}_n) \\ = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n) + \det(\vec{v}_1, \dots, \vec{v}_i, \dots, c\vec{v}_i, \dots, \vec{v}_n) & \quad (\text{by Property 4}) \\ = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n) + c \cdot \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_i, \dots, \vec{v}_n) & \quad (\text{by Property 3}) \\ = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n) + c \cdot 0 & \quad (\text{by Property 2}) \\ = \det(\vec{v}_1, \dots, \vec{v}_i, \dots, \vec{v}_j, \dots, \vec{v}_n). \end{aligned}$$

The other case is similar.

□

Proposition 5.3.5. *Let $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n \in \mathbb{R}^n$ and assume that $(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n)$ is linearly dependent. We then have that $\det(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n) = 0$.*

Proof. We may assume that $n \geq 2$, because the result is trivial when $n = 1$. Since $(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n)$ is linearly dependent, we may apply Proposition 4.4.2 to fix a k such that \vec{v}_k is a linear combination of the other \vec{v}_i . Thus, we may fix $c_1, \dots, c_{k-1}, c_{k+1}, \dots, c_n \in \mathbb{R}$ with

$$\vec{v}_k = c_1 \vec{v}_1 + \dots + c_{k-1} \vec{v}_{k-1} + c_{k+1} \vec{v}_{k+1} + \dots + c_n \vec{v}_n.$$

Using the fact that \det is linear in each component (i.e. using Properties 3 and 4), we see that

$$\begin{aligned} \det(\vec{v}_1, \dots, \vec{v}_n) &= \det(\vec{v}_1, \dots, \vec{v}_{k-1}, c_1 \vec{v}_1 + \dots + c_{k-1} \vec{v}_{k-1} + c_{k+1} \vec{v}_{k+1} + \dots + c_n \vec{v}_n, \vec{v}_{k+1}, \dots, \vec{v}_n) \\ &= \sum_{i=1}^{k-1} c_i \cdot \det(\vec{v}_1, \dots, \vec{v}_{k-1}, \vec{v}_i, \vec{v}_{k+1}, \dots, \vec{v}_n) + \sum_{i=k+1}^n c_i \cdot \det(\vec{v}_1, \dots, \vec{v}_{k-1}, \vec{v}_i, \vec{v}_{k+1}, \dots, \vec{v}_n). \end{aligned}$$

Now every term in each of the above sum has one argument repeated, so every term in the above sums is 0 by Property 2. It follows that $\det(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n) = 0$. \square

Given an $n \times n$ matrix A , we can view A as consisting of n elements of \mathbb{R}^n in two ways: the set of n rows or the set of n columns. Since several of our above properties correspond to row operations, we choose the following definition.

Definition 5.3.6. *Given an $n \times n$ matrix A with rows $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n$, we define*

$$\det(A) = \det(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n).$$

We typically write

$$\begin{vmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ a_{2,1} & a_{2,2} & \dots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n,1} & a_{n,2} & \dots & a_{n,n} \end{vmatrix}$$

rather than

$$\det \begin{pmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ a_{2,1} & a_{2,2} & \dots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n,1} & a_{n,2} & \dots & a_{n,n} \end{pmatrix}.$$

Interpreting Property 3, Proposition 5.3.3, and Proposition 5.3.4 in the context of matrices, we obtain the following facts:

- If B is obtained from A by interchanging two rows, then $\det(B) = -\det(A)$.
- If B is obtained from A by multiplying a row by c , then $\det(B) = c \cdot \det(A)$.
- If B is obtained from A by adding a multiple of one row of A to another, then $\det(B) = \det(A)$.

We can visualize the second of these as saying that

$$\begin{vmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ \vdots & \vdots & \vdots & \vdots \\ ca_{i,1} & ca_{i,2} & \dots & ca_{i,n} \\ \vdots & \vdots & \vdots & \vdots \\ a_{n,1} & a_{n,2} & \dots & a_{n,n} \end{vmatrix} = c \cdot \begin{vmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ \vdots & \vdots & \vdots & \vdots \\ a_{i,1} & a_{i,2} & \dots & a_{i,n} \\ \vdots & \vdots & \vdots & \vdots \\ a_{n,1} & a_{n,2} & \dots & a_{n,n} \end{vmatrix}.$$

In other words, we can pull a c out of any one row. For example, we can perform the following computations:

$$\begin{aligned}
 \begin{vmatrix} 3 & 1 & 7 \\ -2 & -8 & 4 \\ 1 & 4 & 5 \end{vmatrix} &= (-2) \cdot \begin{vmatrix} 3 & 1 & 7 \\ 1 & 4 & -2 \\ 1 & 4 & 5 \end{vmatrix} \\
 &= (-2) \cdot (-1) \cdot \begin{vmatrix} 1 & 4 & 5 \\ 1 & 4 & -2 \\ 3 & 1 & 7 \end{vmatrix} && (R_1 \leftrightarrow R_3) \\
 &= (-2) \cdot (-1) \cdot \begin{vmatrix} 1 & 4 & 5 \\ 0 & 0 & -7 \\ 3 & 1 & 7 \end{vmatrix} && (R_1 \leftrightarrow R_3) \\
 &= (-2) \cdot (-1) \cdot \begin{vmatrix} 1 & 4 & 5 \\ 0 & 0 & -7 \\ 0 & -11 & -4 \end{vmatrix} && (-R_1 + R_2).
 \end{aligned}$$

Before moving on to complete this and other examples we first establish a couple more properties that will ease our calculations.

Definition 5.3.7. Let A and B be $n \times n$ matrices. We write $A \sim_R B$ to mean that A and B are row equivalent, i.e. that we can obtain B from A using a finite sequence of elementary row operations.

Proposition 5.3.8. Suppose that A and B are $n \times n$ matrices with $A \sim_R B$. We then have that $\det(A) = 0$ if and only if $\det(B) = 0$.

Proof. Row operations either leave the determinant alone, scale the determinant by a nonzero value, or multiply the determinant by -1 . If $\det(A) = 0$, then each elementary row operations will maintain this. Similarly, if $\det(A) \neq 0$, then each elementary row operations will maintain this. \square

Proposition 5.3.9. If A is an $n \times n$ diagonal matrix with

$$A = \begin{pmatrix} a_{1,1} & 0 & \dots & 0 \\ 0 & a_{2,2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & a_{n,n} \end{pmatrix}$$

then $\det(A) = a_{1,1}a_{2,2} \cdots a_{n,n}$.

Proof. Using Property 3 a total of n times together with Property 1 of \det , we have

$$\begin{aligned}
 \det(A) &= \det(a_{1,1} \cdot \vec{e}_1, a_{2,2} \cdot \vec{e}_2, \dots, a_{n,n} \cdot \vec{e}_n) \\
 &= a_{1,1}a_{2,2} \cdots a_{n,n} \det(\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n) \\
 &= a_{1,1}a_{2,2} \cdots a_{n,n} \cdot 1 \\
 &= a_{1,1}a_{2,2} \cdots a_{n,n}.
 \end{aligned}$$

\square

Proposition 5.3.10. If A is an $n \times n$ upper triangular matrix, i.e. if

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ 0 & a_{2,2} & \dots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & a_{n,n} \end{pmatrix},$$

then $\det(A) = a_{1,1}a_{2,2} \cdots a_{n,n}$.

Proof. Suppose first that some $a_{i,i}$ equals 0. We claim that the rows of A are linearly dependent. Let k be the smallest number with $a_{k,k} = 0$. Consider the matrix B that is obtained by deleting the first $k - 1$ rows and the first k columns of A . We then have that B is an $(n - k + 1) \times (n - k)$ matrix. The rows of B then form a sequence of $n - k + 1$ many vectors in \mathbb{R}^{n-k} . Since $n - k + 1 > n - k$, we can use Corollary 4.4.7 to conclude that the rows of B are linearly dependent. Thus, the last $n - k + 1$ many rows of A are linearly dependent (the same nontrivial linear combination of the rows of B that produces $\vec{0}$ will still work, since we are just adding zeros to the front). It follows that the rows of A are linearly dependent, so $\det(A) = 0$ by Proposition 5.3.5.

Suppose now that all $a_{i,i}$ are nonzero. In this case, we can use only the third elementary row operation (i.e. row combination) repeatedly to row reduce A to the diagonal matrix

$$B = \begin{pmatrix} a_{1,1} & 0 & \cdots & 0 \\ 0 & a_{2,2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_{n,n} \end{pmatrix}$$

(we first create the 0's above $a_{n,n}$ in the last column by using row combination with the last row, then proceed backwards across the columns). Since row combination does not change the determinant, we conclude that

$$\det(A) = \det(B) = a_{1,1}a_{2,2} \cdots a_{n,n}$$

by Proposition 5.3.9. □

With this in hand, we can complete our above computation by proceeding to echelon form:

$$\begin{aligned} \begin{vmatrix} 3 & 1 & 7 \\ -2 & -8 & 4 \\ 1 & 4 & 5 \end{vmatrix} &= (-2) \cdot (-1) \cdot \begin{vmatrix} 1 & 4 & 5 \\ 0 & 0 & -7 \\ 3 & 1 & 7 \end{vmatrix} && (-R_1 + R_2) \\ &= (-2) \cdot (-1) \cdot \begin{vmatrix} 1 & 4 & 5 \\ 0 & 0 & -7 \\ 0 & -11 & -8 \end{vmatrix} && (-3R_1 + R_3) \\ &= (-2) \cdot (-1) \cdot (-1) \cdot \begin{vmatrix} 1 & 4 & 5 \\ 0 & -11 & -8 \\ 0 & 0 & -7 \end{vmatrix} && \begin{matrix} (R_2 \leftrightarrow R_3) \\ (R_2 \leftrightarrow R_3) \end{matrix} \\ &= (-2) \cdot (-1) \cdot (-1) \cdot 1 \cdot (-11) \cdot (-7) \\ &= -154. \end{aligned}$$

For another example, we have

$$\begin{aligned}
 \begin{vmatrix} -1 & 3 & 0 \\ 5 & 2 & 1 \\ 8 & 4 & -2 \end{vmatrix} &= (-1) \cdot \begin{vmatrix} 1 & -3 & 0 \\ 5 & 2 & 1 \\ 8 & 4 & -2 \end{vmatrix} \\
 &= (-1) \cdot \begin{vmatrix} 1 & -3 & 0 \\ 0 & 17 & 1 \\ 0 & 28 & -2 \end{vmatrix} && \begin{matrix} (-5R_1 + R_2) \\ (-8R_1 + R_3) \end{matrix} \\
 &= (-17) \cdot \begin{vmatrix} 1 & -3 & 0 \\ 0 & 1 & \frac{1}{17} \\ 0 & 28 & -2 \end{vmatrix} \\
 &= (-17) \cdot \begin{vmatrix} 1 & -3 & 0 \\ 0 & 1 & \frac{1}{17} \\ 0 & 0 & -\frac{62}{17} \end{vmatrix} && (-28R_2 + R_3) \\
 &= (-17) \cdot \left(-\frac{62}{17}\right) \\
 &= 62,
 \end{aligned}$$

where we have using the above result once we reached an upper triangular matrix.

Corollary 5.3.11. *If A is a square matrix, then A is invertible if and only if $\det(A) \neq 0$.*

Proof. Suppose first that A is invertible. By the algorithm given at the end of the previous section, we then know that $A \sim_R I_n$. Since $\det(I_n) = 1 \neq 0$, Proposition 5.3.8 implies that $\det(A) \neq 0$.

Suppose conversely that A is not invertible. Let B be an echelon form of A and notice that $A \sim_R B$. We know that B does not contain a leading entry in every row, so B must contain a row of all zeros (since B is a square matrix), and hence the rows of B are linearly dependent. Therefore, $\det(B) = 0$ by Proposition 5.3.5. Using Proposition 5.3.8, we conclude that $\det(A) = 0$. \square

In Section 3.4, we showed that

$$\det \left(\begin{pmatrix} a \\ b \end{pmatrix}, \begin{pmatrix} c \\ d \end{pmatrix} \right) = ad - bc,$$

which in terms of matrices simply says that

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

Is there a simple formula for the determinant of a 3×3 matrix? Let's see what happens if we apply the same procedure that worked that in 2×2 case. Consider the following determinant:

$$\begin{vmatrix} a & b & c \\ d & e & f \\ g & h & i \end{vmatrix}.$$

As in the 2×2 case, we can write this as

$$\det \left(\begin{pmatrix} a \\ b \\ c \end{pmatrix}, \begin{pmatrix} d \\ e \\ f \end{pmatrix}, \begin{pmatrix} g \\ h \\ i \end{pmatrix} \right)$$

which is

$$\det(a\vec{e}_1 + b\vec{e}_2 + c\vec{e}_3, d\vec{e}_1 + e\vec{e}_2 + f\vec{e}_3, g\vec{e}_1 + h\vec{e}_2 + i\vec{e}_3).$$

Now in the 2×2 case, when we expand determinant using linearity in each component, we obtain $2 \cdot 2 = 4$ many summands. In our case, we obtain $3 \cdot 3 \cdot 3 = 27$ many summands. However, many of these terms are 0. For example, one such summand is

$$afi \cdot \det(\vec{e}_1, \vec{e}_3, \vec{e}_3) = afi \cdot 0 = 0.$$

How many terms do not have a repeated vector? We have 3 choices for which term to take from the first argument, and based on this we have 2 choices for a different term from the second, and then the third term is completely determined (since it must be different from the first two). Thus, there will only be $3 \cdot 2 = 6$ many nonzero terms. One such term is

$$aei \cdot \det(\vec{e}_1, \vec{e}_2, \vec{e}_3) = aei \cdot 1 = aei.$$

Another term is

$$\begin{aligned} bdi \cdot \det(\vec{e}_2, \vec{e}_1, \vec{e}_3) &= bdi \cdot (-1) \cdot \det(\vec{e}_1, \vec{e}_2, \vec{e}_3) \\ &= bdi \cdot (-1) \cdot 1 \\ &= -bdi. \end{aligned}$$

Notice that we had to switch two of the standard basis vectors, which explains the -1 . Another term is

$$\begin{aligned} cdh \cdot \det(\vec{e}_3, \vec{e}_1, \vec{e}_2) &= cdh \cdot (-1) \cdot \det(\vec{e}_1, \vec{e}_3, \vec{e}_2) \\ &= cdh \cdot (-1) \cdot (-1) \cdot \det(\vec{e}_1, \vec{e}_2, \vec{e}_3) \\ &= cdh \cdot (-1) \cdot (-1) \cdot 1 \\ &= cdh. \end{aligned}$$

In general, the sign of the result will depend on whether we need to make an even or odd number of swaps to get to the standard order of the natural basis. Working out all of the possibilities, it follows that

$$\begin{vmatrix} a & b & c \\ d & e & f \\ g & h & i \end{vmatrix} = aei - ahf - bdi + bfg + cdh - ceg$$

There are few things to notice from this. First, going back to the 2×2 case, notice that

$$\begin{vmatrix} a & c \\ b & d \end{vmatrix} = ad - cb = ad - bc = \begin{vmatrix} a & b \\ c & d \end{vmatrix}$$

In other words, for a 2×2 matrix A , we have $\det(A) = \det(A^T)$ where A^T (called the *transpose* of A) is the result of swapping the rows and columns of A . One can show that this is still true for 3×3 matrices using the above formula. In fact, with some work (which we omit), it is possible to show the following.

Theorem 5.3.12. *For any $n \times n$ matrix A , we have $\det(A) = \det(A^T)$.*

Let's take another look at our formula

$$\begin{vmatrix} a & b & c \\ d & e & f \\ g & h & i \end{vmatrix} = aei - ahf - bdi + bfg + cdh - ceg.$$

We can rewrite this in various ways, such as

$$\begin{aligned} aei - ahf - bdi + bfg + cdh - ceg &= a(ei - fh) - b(di - fg) + c(dh - eg) \\ &= -d(bi - ch) + e(ai - cg) - f(ah - bg) \\ &= g(bf - ce) - h(af - cd) + i(ae - bd). \end{aligned}$$

Notice that in each case, the part in parentheses is the determinant of the 2×2 matrix obtained by deleting the row and column corresponding to the element outside of the parentheses. Furthermore, these factors outside the parentheses correspond to the elements of a given row, with certain signs attached. Instead of expanding across a row, we can instead expand down a column:

$$\begin{aligned} aei - ahf - bdi + bfg + cdh - ceg &= a(ei - fh) - d(bi - ch) + g(bf - ce) \\ &= -b(di - fg) + e(ai - cg) - h(af - cd) \\ &= c(dh - eg) - f(ah - bg) + i(ae - bd). \end{aligned}$$

Each of these are called *cofactor expansions*.

Definition 5.3.13. Given an $n \times n$ matrix A , the (i, j) cofactor of A is $C_{ij} = (-1)^{i+j} \det(A_{ij})$, where A_{ij} is the result of deleting the i^{th} row and j^{th} column of the matrix A .

For example, the $(2, 3)$ cofactor of the matrix

$$\begin{pmatrix} 4 & 0 & 5 & -2 \\ 7 & -1 & -9 & 3 \\ 2 & 10 & -4 & 3 \\ 0 & 1 & 1 & 6 \end{pmatrix}$$

is

$$C_{2,3} = (-1)^{2+3} \cdot \begin{vmatrix} 4 & 0 & -2 \\ 2 & 10 & 3 \\ 0 & 1 & 6 \end{vmatrix} = (-1) \cdot \begin{vmatrix} 4 & 0 & -2 \\ 2 & 10 & 3 \\ 0 & 1 & 6 \end{vmatrix}.$$

The above calculations establish the following in the 3×3 case, but one can show that it holds more generally:

Theorem 5.3.14. Let A be an $n \times n$ matrix. For any i , we have

$$\det(A) = a_{i1}C_{i1} + a_{i2}C_{i2} + \cdots + a_{in}C_{in}$$

and for any j , we have

$$\det(A) = a_{1j}C_{1j} + a_{2j}C_{2j} + \cdots + a_{nj}C_{nj}.$$

Above, we computed

$$\begin{vmatrix} -1 & 3 & 0 \\ 5 & 2 & 1 \\ 8 & 4 & -2 \end{vmatrix} = 62$$

using row operations. We now compute it using a cofactor expansion across the first row:

$$\begin{aligned} \begin{vmatrix} -1 & 3 & 0 \\ 5 & 2 & 1 \\ 8 & 4 & -2 \end{vmatrix} &= (-1) \cdot \begin{vmatrix} 2 & 1 \\ 4 & -2 \end{vmatrix} - 3 \cdot \begin{vmatrix} 5 & 1 \\ 8 & -2 \end{vmatrix} + 0 \cdot \begin{vmatrix} 5 & 2 \\ 8 & 4 \end{vmatrix} \\ &= (-1) \cdot (-4 - 4) - 3 \cdot (-10 - 8) + 0 \cdot (20 - 26) \\ &= 8 + 54 \\ &= 62. \end{aligned}$$

We can also compute a cofactor expansion down the second column:

$$\begin{aligned}
 \begin{vmatrix} -1 & 3 & 0 \\ 5 & 2 & 1 \\ 8 & 4 & -2 \end{vmatrix} &= (-3) \cdot \begin{vmatrix} 5 & 1 \\ 8 & -2 \end{vmatrix} + 2 \cdot \begin{vmatrix} -1 & 0 \\ 8 & -2 \end{vmatrix} - 4 \cdot \begin{vmatrix} -1 & 5 \\ 0 & 1 \end{vmatrix} \\
 &= (-3) \cdot (-10 - 8) + 2 \cdot (2 - 0) - 4 \cdot (-1 - 0) \\
 &= 54 + 4 + 4 \\
 &= 62.
 \end{aligned}$$

Theorem 5.3.15. *If A and B are $n \times n$ matrices, then $\det(AB) = \det(A) \cdot \det(B)$.*

5.4 Eigenvalues and Eigenvectors

In Section 3.3, we explored eigenvalues and eigenvectors in the context of \mathbb{R}^2 . We now generalize these ideas, not only to general \mathbb{R}^n , but to any vector space.

Definition 5.4.1. *Let V be a vector space, and let $T: V \rightarrow V$ be a linear transformation.*

- An eigenvector of T is a nonzero vector $\vec{v} \in V$ such that there exists $\lambda \in \mathbb{R}$ with $T(\vec{v}) = \lambda\vec{v}$.
- An eigenvalue of T is a scalar $\lambda \in \mathbb{R}$ such that there exists a nonzero $\vec{v} \in V$ with $T(\vec{v}) = \lambda\vec{v}$.

When $\vec{v} \in V$ is nonzero and $\lambda \in \mathbb{R}$ are such that $T(\vec{v}) = \lambda\vec{v}$, we say that \vec{v} is an eigenvector of T corresponding to the eigenvalue λ .

Notice that in our definition, we require that the domain and codomain of the linear transformation are the same vector space. To see why, consider what would happen if we thought about a general $T: V \rightarrow W$. Given $\vec{v} \in V$, we have that $T(\vec{v}) \in W$ but $\lambda\vec{v} \in V$ for all $\lambda \in \mathbb{R}$. Thus, it wouldn't even make sense to say that $T(\vec{v})$ equals $\lambda\vec{v}$ in such a general situation.

Back when we talked about \mathbb{R}^2 , we noticed that if we had an eigenvector of a linear transformation, then any (nonzero) scalar multiple of it would also be an eigenvector. More generally, we have the following.

Proposition 5.4.2. *Let $T: V \rightarrow V$ be a linear transformation and let $\lambda \in \mathbb{R}$. The set*

$$W = \{\vec{v} \in V : T(\vec{v}) = \lambda\vec{v}\},$$

which is the set of all eigenvectors of T corresponding to λ together with $\vec{0}$, is a subspace of V . It is called the eigenspace of T corresponding to λ .

Proof. We check the three properties.

- We have $T(\vec{0}) = \vec{0}$ and $\lambda\vec{0} = \vec{0}$, so $T(\vec{0}) = \lambda\vec{0}$. Therefore, $\vec{0} \in W$.
- Let $\vec{v}_1, \vec{v}_2 \in W$ be arbitrary. We then have $T(\vec{v}_1) = \lambda\vec{v}_1$ and also $T(\vec{v}_2) = \lambda\vec{v}_2$, so

$$\begin{aligned}
 T(\vec{v}_1 + \vec{v}_2) &= T(\vec{v}_1) + T(\vec{v}_2) \\
 &= \lambda\vec{v}_1 + \lambda\vec{v}_2 \\
 &= \lambda \cdot (\vec{v}_1 + \vec{v}_2).
 \end{aligned}$$

It follows that $\vec{v}_1 + \vec{v}_2 \in W$. Therefore, W is closed under addition.

- Let $\vec{v} \in W$ and $c \in \mathbb{R}$ be arbitrary. We then have $T(\vec{v}) = \lambda\vec{v}$, so

$$\begin{aligned} T(c\vec{v}) &= c \cdot T(\vec{v}) \\ &= c \cdot (\lambda\vec{v}) \\ &= (c\lambda) \cdot \vec{v} \\ &= (\lambda c) \cdot \vec{v} \\ &= \lambda \cdot (c\vec{v}). \end{aligned}$$

It follows that $c\vec{v} \in W$. Therefore, W is closed under scalar multiplication.

Since W satisfies the three properties, we conclude that W is a subspace of V . \square

Let's consider a couple of examples of eigenvectors and eigenvalues for a more interesting vector space. Let V be the subspace of \mathcal{F} consisting of all of the infinitely differentiable functions $f: \mathbb{R} \rightarrow \mathbb{R}$. In other words, given a function $f: \mathbb{R} \rightarrow \mathbb{R}$, we have $f \in V$ if and only if $f'(x)$ exists for all $x \in \mathbb{R}$, $f''(x)$ exists for all $x \in \mathbb{R}$, and in general $f^{(n)}(x)$ exists for all $x \in \mathbb{R}$ and $n \in \mathbb{N}^+$.

First consider the linear transformation $T: V \rightarrow V$ given by differentiation, i.e. $T(f) = f'$ for all $f \in V$. Let's think about eigenvalues and eigenvectors of T .

- We have that 1 is an eigenvalue of T , because if $f(x) = e^x$, then $T(f) = 1 \cdot f$. Since we know that the eigenspace corresponding to 1 is a subspace of V , we conclude that the function ce^x is an eigenvector of T corresponding to 1 for any (nonzero) $c \in \mathbb{R}$.
- We have that 2 is an eigenvalue of T , because if $f(x) = e^{2x}$, then $T(f) = 2 \cdot f$. As in the previous case, we also know that the function ce^{2x} is a eigenvectors corresponding to 2 for each (nonzero) $c \in \mathbb{R}$.
- In fact, for *any* $\lambda \in \mathbb{R}$, the function $ce^{\lambda x}$ is an eigenvectors of T corresponding to λ for each $c \in \mathbb{R}$. Notice that if $\lambda = 0$, then the function $ce^{\lambda x}$ is just the constant function c .

Therefore, *every* $\lambda \in \mathbb{R}$ is an eigenvalue of T . Notice how much this differs from our previous experience where a linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ had at most 2 eigenvalues. A natural questions arises at this point. We found an infinite family of eigenvectors for each given $\lambda \in \mathbb{R}$. Have we missed any? In other words, if we have an $f \in V$ with $T(f) = \lambda \cdot f$, must it be the case that $f(x) = ce^{\lambda x}$ for some $c \in \mathbb{R}$? The answer to this question is not obvious, but it turns out that the answer is yes (see Differential Equations). In other words, the eigenspace corresponding to a given $\lambda \in \mathbb{R}$ equals $\text{Span}(e^{\lambda x})$, so each of the eigenspaces has dimension 1.

Now let's consider the linear transformation $T: V \rightarrow V$ given by taking the second derivative, i.e. $T(f) = f''$ for all $f \in V$.

- As above, we have that 1 is an eigenvalue of T , because if $f(x) = e^x$, then $T(f) = 1 \cdot f$. Furthermore, for all $c \in \mathbb{R}$, the function ce^x is an eigenvector of T corresponding to 1. However, there is another eigenvector corresponding to 1 that is not of this form. If $g(x) = e^{-x}$, then $T(g) = g$ because taking two derivatives introduces two negatives. Since we know that the eigenspace corresponding to 1 is a subspace of V , it follows that for all $c_1, c_2 \in \mathbb{R}$, the function $c_1e^x + c_2e^{-x}$ is a eigenvector of T corresponding to 1.
- We have that 2 is an eigenvalue of T , because if $f(x) = e^{\sqrt{2} \cdot x}$, then $T(f) = 2 \cdot f$. Similar to the previous case, the function $e^{-\sqrt{2} \cdot x}$ is also an eigenvector corresponding to 2. It follows that for all $c_1, c_2 \in \mathbb{R}$, the function $c_1e^{\sqrt{2} \cdot x} + c_2e^{-\sqrt{2} \cdot x}$ is a eigenvector of T corresponding to 2.
- In fact, for any *positive* $\lambda \in \mathbb{R}$, we have that λ is an eigenvalue of T and that the functions $c_1e^{\sqrt{\lambda} \cdot x} + c_2e^{-\sqrt{\lambda} \cdot x}$ are eigenvectors corresponding to λ . Moreover, it is possible to show that $(e^{\sqrt{\lambda} \cdot x}, e^{-\sqrt{\lambda} \cdot x})$ is linearly independent.

- We also have that -1 is an eigenvalue of T , but with a very different looking function. Notice that if $f(x) = \sin x$, then $T(f) = (-1) \cdot f$, so f is an eigenvector of T corresponding to -1 . We also have that $g(x) = \cos x$ is an eigenvector corresponding to -1 . It follows that for all $c_1, c_2 \in \mathbb{R}$, the function $c_1 \sin x + c_2 \cos x$ is a eigenvector of T corresponding to -1 .
- Generalizing as in the positive case, for any *negative* $\lambda \in \mathbb{R}$, we have that λ is an eigenvalue of T and that the functions $c_1 \sin(\sqrt{-\lambda} \cdot x) + c_2 \cos(\sqrt{-\lambda} \cdot x)$ are eigenvectors corresponding to λ . Moreover, it is possible to show that $(\sin(\sqrt{-\lambda} \cdot x), \cos(\sqrt{-\lambda} \cdot x))$ is linearly independent.
- We also have that 0 is an eigenvalue of T and that for each $c_1, c_2 \in \mathbb{R}$, the function $c_1 x + c_2$ is an eigenvector corresponding to 0 . Moreover, $(x, 1)$ is linearly independent.

In fact, as in the previous examples, we have found all of the eigenvectors corresponding to each λ . In other words, we have the following:

- If $\lambda > 0$, then the eigenspace corresponding to λ is $\text{Span}(e^{\sqrt{\lambda} \cdot x}, e^{-\sqrt{\lambda} \cdot x})$.
- If $\lambda < 0$, then the eigenspace corresponding to λ is $\text{Span}(\sin(\sqrt{-\lambda} \cdot x), \cos(\sqrt{-\lambda} \cdot x))$.
- If $\lambda = 0$, then the eigenspace corresponding to λ is $\text{Span}(x, 1)$.

Therefore, every λ is an eigenvalue of T , and for each $\lambda \in \mathbb{R}$, the eigenspace of T corresponding to λ has dimension 2. At this point, there is a natural feeling that the dimension of the eigenspaces is related to the number of derivatives that we took. Again, see Differential Equations.

In the last example, the transition from exponential to trigonometric functions might seem a little strange. However, there is a deep connection between the two. When thinking about eigenvectors corresponding to -1 , we might think to continue the exponential pattern and look at the function $f(x) = e^{ix}$ where $i = \sqrt{-1}$. If we mindlessly follow derivative rules (i.e. ignore the fact that we have a complex number in the exponent), then $f'(x) = ie^{ix}$ and $f''(x) = i^2 e^{ix} = (-1) \cdot e^{ix}$. In fact, in later courses you'll see that the proper way to define e^{ix} is as $\cos x + i \cdot \sin x$. Thus, our example here is illustrating a much deeper phenomenon.

We've just seen an example of finding eigenvalues and eigenvectors for a linear transformation on a complicated vector space consisting of functions. As we've seen, a complete understanding of that example involves some nontrivial background in differential equations. However, for linear transformations $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$, we can extend out techniques from \mathbb{R}^2 to this setting. First, instead of working with a linear transformation T , we'll often work with a matrix representing T , such as the standard matrix $[T]$. In this context, we have the following analogous definition.

Definition 5.4.3. Let A be a $n \times n$ matrix.

- An eigenvector of A is a nonzero vector $\vec{v} \in \mathbb{R}^n$ such that there exists $\lambda \in \mathbb{R}$ with $A\vec{v} = \lambda\vec{v}$.
- An eigenvalue of A is a scalar $\lambda \in \mathbb{R}$ such that there exists a nonzero $\vec{v} \in \mathbb{R}^n$ with $A\vec{v} = \lambda\vec{v}$.

When $\vec{v} \in \mathbb{R}^n$ is nonzero and $\lambda \in \mathbb{R}$ are such that $A\vec{v} = \lambda\vec{v}$, we say that \vec{v} is an eigenvector of A corresponding to the eigenvalue λ .

As in the case for \mathbb{R}^2 , we have the following result.

Proposition 5.4.4. Let A be an $n \times n$ matrix, let $\vec{v} \in \mathbb{R}^n$, and let $\lambda \in \mathbb{R}$. We have that $A\vec{v} = \lambda\vec{v}$ if and only if $\vec{v} \in \text{Null}(A - \lambda I)$. Therefore, \vec{v} is an eigenvector of A corresponding to λ if and only if $\vec{v} \neq \vec{0}$ and $\vec{v} \in \text{Null}(A - \lambda I)$.

Proof. Suppose first that $A\vec{v} = \lambda\vec{v}$. Subtracting $\lambda\vec{v}$ from both sides, we then have $A\vec{v} - \lambda\vec{v} = \vec{0}$, and hence $A\vec{v} - \lambda I\vec{v} = \vec{0}$. We then have $(A - \lambda I)\vec{v} = \vec{0}$, and thus $\vec{v} \in \text{Null}(A - \lambda I)$.

Conversely, suppose that $\vec{v} \in \text{Null}(A - \lambda I)$. We then have $(A - \lambda I)\vec{v} = \vec{0}$, so $A\vec{v} - \lambda I\vec{v} = \vec{0}$, and hence $A\vec{v} - \lambda\vec{v} = \vec{0}$. Adding $\lambda\vec{v}$ to both sides, we conclude that $A\vec{v} = \lambda\vec{v}$. \square

Corollary 5.4.5. *Let A be an $n \times n$ matrix and let $\lambda \in \mathbb{R}$. We have that λ is an eigenvalue of A if and only if $\text{Null}(A - \lambda I) \neq \{\vec{0}\}$.*

For example, let $T: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be the unique linear transformation with

$$[T] = \begin{pmatrix} 5 & -3 & 2 \\ 2 & -2 & 4 \\ -4 & 12 & -4 \end{pmatrix}.$$

Let's determine if 4 is an eigenvalue, and if so, find the eigenspace of T corresponding to $\lambda = 4$. Let $A = [T]$. We have

$$A - 4I = \begin{pmatrix} 5 & -3 & 2 \\ 2 & -2 & 4 \\ -4 & 12 & -4 \end{pmatrix} - \begin{pmatrix} 4 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 4 \end{pmatrix} = \begin{pmatrix} 1 & -3 & 2 \\ 2 & -6 & 4 \\ -4 & 12 & -8 \end{pmatrix}.$$

Performing row operations on this matrix, we obtain

$$\begin{pmatrix} 1 & -3 & 2 \\ 2 & -6 & 4 \\ -4 & 12 & -8 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & -3 & 2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad \begin{array}{l} (-2R_1 + R_2) \\ (4R_1 + R_3). \end{array}$$

Notice that this matrix does not have a leading entry in every column, so $\text{Null}(A - 4I) \neq \{\vec{0}\}$. To find the eigenspace of A corresponding to 4, we just need to solve the homogeneous linear system coded by augmenting the matrix $A - 4I$ with a column of 0's. To do this, we introduce parameters for the latter two columns, say $y = s$ and $z = t$. The first equation then says that $x - 3y + 2z = 0$, so $x = 3y - 2z$, and hence $x = 3s - 2t$. Therefore, we have

$$\begin{aligned} \text{Null}(A - 4I) &= \left\{ \begin{pmatrix} 3s - 2t \\ s \\ t \end{pmatrix} : s, t \in \mathbb{R} \right\} \\ &= \left\{ s \cdot \begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix} + t \cdot \begin{pmatrix} -2 \\ 0 \\ 1 \end{pmatrix} : s, t \in \mathbb{R} \right\}. \end{aligned}$$

In other words, we have

$$\text{Null}(A - 4I) = \text{Span} \left(\begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -2 \\ 0 \\ 1 \end{pmatrix} \right).$$

We know from Proposition 5.4.2 that the eigenspace of T corresponding to 4 is a subspace of \mathbb{R}^3 , and now we have written that subspace as a span of two vectors. In fact, we have that

$$\left(\begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -2 \\ 0 \\ 1 \end{pmatrix} \right)$$

is linearly independent, because if $c_1, c_2 \in \mathbb{R}$ are arbitrary with

$$c_1 \cdot \begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix} + c_2 \cdot \begin{pmatrix} -2 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

then we have

$$\begin{pmatrix} 3c_1 - 2c_2 \\ c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

and hence $c_1 = 0$ and $c_2 = 0$. Therefore,

$$\left(\begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -2 \\ 0 \\ 1 \end{pmatrix} \right)$$

is a basis for the eigenspace of T corresponding to 4.

Although we can use this process to find (a basis for) the eigenspace of T corresponding to any given $\lambda \in \mathbb{R}$, we first need a way to find all of the eigenvalues. Given a 2×2 matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

we know that λ is an eigenvalue of A if and only if $\text{Null}(A - \lambda I) \neq \{\vec{0}\}$, which is true if and only if $(a - \lambda)(d - \lambda) - bc = 0$. Thus, we can find the eigenvalues by finding the roots of a certain quadratic polynomial called the characteristic polynomial of A . How can we generalize this to larger matrices?

Suppose that we have an $n \times n$ matrix A and a $\lambda \in \mathbb{R}$. We still know that λ is an eigenvalue of A if and only if $\text{Null}(A - \lambda I) \neq \{\vec{0}\}$ by Corollary 5.4.5. Now $\text{Null}(A - \lambda I) \neq \{\vec{0}\}$ exactly when an echelon form of $A - \lambda I$ has a column without a leading entry (and hence has a row without a leading entry because $A - \lambda I$ is a square matrix). By our work on invertible matrices, we know that this is the same as saying that $A - \lambda I$ is *not* invertible. Finally, we know a computational way to say this is that $\det(A - \lambda I) = 0$ (see Corollary 5.3.11). To summarize, we have that

$$\begin{aligned} \lambda \text{ is an eigenvalue of } A &\iff \text{Null}(A - \lambda I) \neq \{\vec{0}\} \\ &\iff \text{An echelon form of } A - \lambda I \text{ has a column without a leading entry} \\ &\iff A - \lambda I \text{ is not invertible} \\ &\iff \det(A - \lambda I) = 0. \end{aligned}$$

Thus, to find the eigenvalues, we should calculate $\det(A - \lambda I)$ and determine the values of λ that make this equal to 0. Since this generalizes our work in the 2×2 case, we define the following.

Definition 5.4.6. Given an $n \times n$ matrix A , we define the characteristic polynomial of A to be the polynomial $\det(A - \lambda I)$.

Proposition 5.4.7. If A is an $n \times n$ matrix, then the characteristic polynomial of A is a polynomial of degree n .

Proof. When taking the determinant of $A - \lambda I$, we obtain a sum of terms, each of which appears by choosing n elements from $A - \lambda I$ in such a way that we pick one element from each row and each column. Since $A - \lambda I$ has n terms involving a λ (those along the diagonal), each of the terms in this sum will be a polynomial of degree at most n . Therefore, adding up all of these terms, we obtain a polynomial that has degree at most n . To see why the polynomial has degree exactly equal to n , notice that the only entries with a λ appear on the diagonal, so the only term that has a product of n factors involving λ occurs when we pick all of the diagonal entries:

$$(a_{1,1} - \lambda)(a_{2,2} - \lambda) \cdots (a_{n,n} - \lambda)$$

When we multiply this out we obtain one term which is $(-1)^n \cdot \lambda^n$, and all the rest of the summands inside this terms have degree less than n . Therefore, the leading term of our resulting polynomial will be $(-1)^n \cdot \lambda^n$, and hence $\det(A - \lambda I)$ will be a polynomial of degree n . \square

Consider the following example. Let

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

Suppose that we want to find all eigenvalues along with bases for the corresponding eigenspaces of A . The first step is to find the eigenvalues, so we want to calculate $\det(A - \lambda I)$ and determine when it equals 0. We have

$$A - \lambda I = \begin{pmatrix} 2 - \lambda & -3 & 1 \\ 1 & -2 - \lambda & 1 \\ 1 & -3 & 2 - \lambda \end{pmatrix}$$

so using a cofactor expansion along the first row, we have

$$\begin{aligned} \det(A - \lambda I) &= (2 - \lambda) \cdot \begin{vmatrix} -2 - \lambda & 1 \\ -3 & 2 - \lambda \end{vmatrix} - (-3) \cdot \begin{vmatrix} 1 & 1 \\ 1 & 2 - \lambda \end{vmatrix} + 1 \cdot \begin{vmatrix} 1 & -2 - \lambda \\ 1 & -3 \end{vmatrix} \\ &= (2 - \lambda) \cdot ((-2 - \lambda)(2 - \lambda) + 3) + 3 \cdot ((2 - \lambda) - 1) + 1 \cdot (-3 - (-2 - \lambda)) \\ &= (2 - \lambda) \cdot (-1 + \lambda^2) + 3 \cdot (1 - \lambda) + ((-1) + \lambda) \\ &= (-2 + \lambda + 2\lambda^2 - \lambda^3) + (3 - 3\lambda) + ((-1) + \lambda) \\ &= -\lambda^3 + 2\lambda^2 - \lambda \\ &= -\lambda(\lambda^2 - 2\lambda + 1) \\ &= -\lambda(\lambda - 1)^2. \end{aligned}$$

Therefore, the eigenvalues of A are 0 and 1. Applying elementary row operations, we see that

$$\begin{aligned} A - 0I &= \begin{pmatrix} 2 & -3 & 1 \\ 1 & -2 & 1 \\ 1 & -3 & 2 \end{pmatrix} \\ &\rightarrow \begin{pmatrix} 1 & -2 & 1 \\ 0 & 1 & -1 \\ 0 & -1 & 1 \end{pmatrix} && \begin{matrix} (-2R_1 + R_2) \\ (-R_1 + R_2) \end{matrix} \\ &\rightarrow \begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{pmatrix} && \begin{matrix} (2R_2 + R_1) \\ (R_2 + R_3). \end{matrix} \end{aligned}$$

Since we are finding $\text{Null}(A - 0I)$, we think of putting an augmented column of 0's on the right. Letting $z = t$, we then have that $y - z = 0$ and $x - z = 0$, so $y = t$ and $x = t$. It follows that the solution to the homogeneous linear system is

$$\left\{ \begin{pmatrix} t \\ t \\ t \end{pmatrix} : t \in \mathbb{R} \right\} = \left\{ t \cdot \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} : t \in \mathbb{R} \right\} = \text{Span} \left(\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right).$$

Since

$$\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$$

is trivially linearly independent, it follows that a basis for the eigenspace of A corresponding to 0 is

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

We also have

$$\begin{aligned} A - 1I &= \begin{pmatrix} 1 & -3 & 1 \\ 1 & -3 & 1 \\ 1 & -3 & 1 \end{pmatrix} \\ &\rightarrow \begin{pmatrix} 1 & -3 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{array}{l} (-R_1 + R_2) \\ (-R_1 + R_3). \end{array} \end{aligned}$$

Since we are finding $\text{Null}(A - 1I)$, we think of putting an augmented column of 0's on the right. Letting $y = s$ and $z = t$, we then have that $x - 3y + z = 0$, so $x = 3s - t$. It follows that the solution to the homogeneous linear system is

$$\left\{ \begin{pmatrix} 3s - t \\ s \\ t \end{pmatrix} : s, t \in \mathbb{R} \right\} = \left\{ s \cdot \begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix} + t \cdot \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} : s, t \in \mathbb{R} \right\} = \text{Span} \left(\begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} \right).$$

Now we claim that

$$\left(\begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} \right)$$

is linearly independent. To see this, let $c_1, c_2 \in \mathbb{R}$ be arbitrary with

$$c_1 \cdot \begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix} + c_2 \cdot \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

We then have

$$\begin{pmatrix} 3c_1 - c_2 \\ c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

so $c_1 = c_2 = 0$. Putting it all together, we conclude that

$$\left(\begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} \right)$$

is a basis for the eigenspace of A corresponding to 1.

We can also generalize the concepts of diagonalizable to this context. The arguments are completely analogous.

Definition 5.4.8. A diagonal $n \times n$ matrix is an $n \times n$ matrix D such that there exists $a_{1,1}, a_{2,2}, \dots, a_{n,n} \in \mathbb{R}$ with

$$D = \begin{pmatrix} a_{1,1} & 0 & 0 & \cdots & 0 & 0 \\ 0 & a_{2,2} & 0 & \cdots & 0 & 0 \\ 0 & 0 & a_{3,3} & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & a_{n-1,n-1} & 0 \\ 0 & 0 & 0 & \cdots & 0 & a_{n,n} \end{pmatrix}.$$

Definition 5.4.9. A linear transformation $T: V \rightarrow V$ is diagonalizable if there exists a basis α of V such that $[T]_\alpha^\alpha$ is a diagonal matrix.

Proposition 5.4.10. Let V be a finite-dimensional vector space, let $T: V \rightarrow V$ be a linear transformation and let $\alpha = (\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n)$ be a basis of V . The following are equivalent.

1. $[T]_\alpha^\alpha$ is a diagonal matrix.
2. $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$ are all eigenvectors of T .

Furthermore, in this case, the diagonal entries of $[T]_\alpha^\alpha$ are the eigenvalues corresponding to $\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n$.

Corollary 5.4.11. Let V be a finite-dimensional vector space and let $T: V \rightarrow V$ be a linear transformation. We then have that T is diagonalizable if and only if there exists a basis of V consisting entirely of eigenvectors of T .

Let's return to our example of

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

We found that A has a 1-dimensional eigenspace corresponding to 0, and a 2-dimensional eigenspace corresponding to 1. If we take the corresponding bases and put them together, we obtain:

$$\alpha = \left(\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 3 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} \right).$$

We claim that α is a basis for \mathbb{R}^3 . To see this, let

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

We will show that α is a basis for \mathbb{R}^3 while showing that P is invertible and simultaneously calculating P^{-1} . We have

$$\begin{aligned} \begin{pmatrix} 1 & 3 & -1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} &\rightarrow \begin{pmatrix} 1 & 0 & 1 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 3 & -1 & 1 & 0 & 0 \end{pmatrix} && (R_1 \leftrightarrow R_3) \\ &&& (R_1 \leftrightarrow R_3) \\ &\rightarrow \begin{pmatrix} 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & -1 & 0 & 1 & -1 \\ 0 & 3 & -2 & 1 & 0 & -1 \end{pmatrix} && \begin{matrix} (-R_1 + R_2) \\ (-R_1 + R_3) \end{matrix} \\ &\rightarrow \begin{pmatrix} 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & -1 & 0 & 1 & -1 \\ 0 & 0 & 1 & 1 & -3 & 2 \end{pmatrix} && (-R_2 + R_3) \\ &\rightarrow \begin{pmatrix} 1 & 0 & 0 & -1 & 3 & -1 \\ 0 & 1 & 0 & 1 & -2 & 1 \\ 0 & 0 & 1 & 1 & -3 & 2 \end{pmatrix} && \begin{matrix} (-R_3 + R_1) \\ (R_3 + R_2). \end{matrix} \end{aligned}$$

It follows that P is invertible, hence α is a basis for \mathbb{R}^3 . We can also use this to diagonalize A . Letting

$$D = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

we have $A = PDP^{-1}$, so

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