


NO BULLSHIT

guide to
LINEAR ALGEBRA



by Ivan Savov

NO BULLSHIT GUIDE TO LINEAR ALGEBRA

Ivan Savov

June 4, 2021

No bullshit guide to linear algebra

by Ivan Savov

Copyright © Ivan Savov, 2014, 2017. All rights reserved.

Published by Minireference Co.

Montréal, Québec, Canada

minireference.com | [@minireference](https://twitter.com/minireference) | [fb.me/noBSguide](https://facebook.com/noBSguide)

For inquiries, contact the author at ivan@minireference.com

Library and Archives Canada Cataloguing in Publication

Savov, Ivan, 1982-, author

No bullshit guide to linear algebra / Ivan Savov. — Second edition.

ISBN 978-0-9920010-2-5 (softcover)

1. Algebras, Linear--Textbooks. 2. Mathematics--Textbooks.
3. Textbooks. I. Title.

QA184.2.S29 2017

512'.5

C2017-902491-4

Cover image: Jim Hefferon and Nicholas Longo, CC-BY-SA-2.5

http://commons.wikimedia.org/wiki/File:Linalg_parallelepiped.png

Second edition

v2.2

git commits `master@1067:d370768 + vers/LA@1271:11193c2`

ISBN 978-0-9920010-2-5

10 9 8 7 6 5 4 3 2 1

Contents

Preface	ix
Introduction	1
1 Math fundamentals	9
1.1 Solving equations	10
1.2 Numbers	12
1.3 Variables	18
1.4 Functions and their inverses	19
1.5 Basic rules of algebra	23
1.6 Solving quadratic equations	30
1.7 The Cartesian plane	35
1.8 Functions	37
1.9 Functions reference	50
1.10 Geometry	66
1.11 Trigonometry	71
1.12 Trigonometric identities	78
1.13 Vectors	80
1.14 Complex numbers	96
1.15 Solving systems of linear equations	103
1.16 Set notation	109
1.17 Math problems	124
2 Intro to linear algebra	131
2.1 Definitions	131
2.2 Vector operations	138
2.3 Matrix operations	145
2.4 Linearity	151
2.5 Overview of linear algebra	156
2.6 Introductory problems	160

3	Computational linear algebra	163
3.1	Reduced row echelon form	164
3.2	Matrix equations	177
3.3	Matrix multiplication	181
3.4	Determinants	185
3.5	Matrix inverse	196
3.6	Computational problems	204
4	Geometric aspects of linear algebra	209
4.1	Lines and planes	209
4.2	Projections	219
4.3	Coordinate projections	225
4.4	Vector spaces	231
4.5	Vector space techniques	243
4.6	Geometric problems	253
5	Linear transformations	257
5.1	Linear transformations	257
5.2	Finding matrix representations	269
5.3	Change of basis for matrices	282
5.4	Invertible matrix theorem	287
5.5	Linear transformations problems	294
6	Theoretical linear algebra	297
6.1	Eigenvalues and eigenvectors	298
6.2	Special types of matrices	312
6.3	Abstract vector spaces	318
6.4	Abstract inner product spaces	322
6.5	Gram–Schmidt orthogonalization	328
6.6	Matrix decompositions	332
6.7	Linear algebra with complex numbers	339
6.8	Theory problems	356
7	Applications	363
7.1	Balancing chemical equations	364
7.2	Input–output models in economics	366
7.3	Electric circuits	368
7.4	Graphs	373
7.5	Fibonacci sequence	375
7.6	Linear programming	378
7.7	Least squares approximate solutions	379
7.8	Computer graphics	388
7.9	Cryptography	402
7.10	Error-correcting codes	412
7.11	Fourier analysis	421
7.12	Applications problems	435

8	Probability theory	439
8.1	Probability distributions	439
8.2	Markov chains	450
8.3	Google's PageRank algorithm	456
8.4	Probability problems	463
9	Quantum mechanics	465
9.1	Introduction	466
9.2	Polarizing lenses experiment	472
9.3	Dirac notation for vectors	479
9.4	Quantum information processing	486
9.5	Postulates of quantum mechanics	489
9.6	Polarizing lenses experiment revisited	506
9.7	Quantum mechanics is not that weird	510
9.8	Quantum mechanics applications	515
9.9	Quantum mechanics problems	531
	End matter	535
	Conclusion	535
	Social stuff	537
	Acknowledgements	537
	General linear algebra links	538
A	Answers and solutions	539
B	Notation	571
	Math notation	571
	Set notation	572
	Vectors notation	572
	Complex numbers notation	573
	Vector space notation	573
	Abstract vector spaces notation	574
	Notation for matrices and matrix operations	574
	Notation for linear transformations	575
	Matrix decompositions	575
	Index	577

Concept maps




Figure 1: This diagram shows the concepts and topics covered in this book.

Consult the index on page 577 to find the exact location in the book where each concept is defined.




Figure 2: Chapter 5 is about linear transformations and their properties.




Figure 3: Chapter 6 covers theoretical aspects of linear algebra.




Figure 4: Matrix computations play an important role throughout this book. Matrices are used to represent linear transformations, systems of linear equations, and various geometric operations.




Figure 5: The book concludes with three chapters on linear algebra applications. In Chapter 7 we'll discuss applications to science, economics, business, computing, and signal processing. In Chapter 8 we'll explain probability theory, and finally in Chapter 9 we'll introduce quantum mechanics.

You can annotate the concept maps with your current knowledge of each concept to keep track of your progress. Add a single dot (•) next to all concepts you've heard of, two dots (••) next to concepts you think you know, and three dots (•••) next to concepts you've used in exercises and problems. By collecting some dots every week, you'll be able to move through the material in no time at all.

If you don't want to mark up your book, you can download a printable version of the concept maps here: bit.ly/LAcmaps.

Preface

This is a book about linear algebra and its applications. The material is presented at the level of a first-year university course, in an approachable style that cuts to the point. It covers both practical and theoretical aspects of linear algebra, with extra emphasis on explaining the connections between concepts and building a solid understanding of the material.

This book is designed to **give readers access to advanced math modelling tools regardless of their academic background**. Since the book includes all the prerequisites needed to learn linear algebra, it's suitable for readers of any skill level—including those who don't feel comfortable with fundamental math concepts.

Why learn linear algebra?

Linear algebra is one of the most fundamental and all-around useful subjects in mathematics. The practical skills learned by studying linear algebra—such as manipulating vectors and matrices—form an essential foundation for applications in physics, computer science, statistics, machine learning, and many other fields of scientific study. Learning linear algebra can also be a lot of fun. Readers will experience numerous *knowledge buzz* moments when they feel math ideas “click” into place and understand the connections between concepts.

The tools of linear algebra form a bridge to more advanced areas of mathematics. For example, learning about *abstract vector spaces* will help readers recognize the common “vector space structure” in seemingly unrelated mathematical objects like matrices, polynomials, and functions. Linear algebra techniques apply not only to standard vectors, but to *all* mathematical objects that are vector-like!

What's in this book?

Each section is a self-contained tutorial that covers the definitions, formulas, and explanations associated with a single topic. Check out the concept maps on the preceding pages to see the book's many topics and the connections between them.

The book begins with a review chapter on numbers, algebra, sets, equations, functions, geometry, and trigonometry (Chapter 1). If you haven't previously studied these concepts, or if you feel your math skills are a little "rusty," read this chapter and work through the exercises and problems provided. If you feel confident in your high school math abilities, jump straight to Chapter 2, where the linear algebra begins.

Chapters 2–6 cover the core topics of linear algebra: vectors, bases, analytical geometry, matrices, linear transformations, matrix representations, vector spaces, inner product spaces, eigenvectors, and matrix decompositions. These chapters contain the material required for every university-level linear algebra course. Each section contains plenty of exercises so you can test your understanding as you read; and each chapter concludes with an extensive list of problems for further practice.

Chapters 7, 8, and 9 discuss various applications of linear algebra. Though this material isn't likely to appear on any final exam, these chapters serve to demonstrate the power of linear algebra techniques and their relevance to many areas of science. The mini-course on quantum mechanics (Chapter 9) is unique to this book. Read this chapter to understand the fascinating laws of physics that govern the behaviour of atoms and photons.

Is this book for you?

The quick pace and lively explanations in this book provide interesting reading for students and non-students alike. Whether you're learning linear algebra for a course, reviewing material as a prerequisite for more advanced topics, or generally curious about the subject, this book will help you find your way in the land of linear algebra.

Students and educators can use this book as the main textbook for any university-level linear algebra course. It contains everything students need to know to prepare for a linear algebra final exam. Don't be fooled by the book's small size compared to other textbooks: it's all in here. The text is compact because we've distilled the essentials and removed the unnecessary crud.

Publisher

The starting point of the **No Bullshit Guide** textbook series dates back to my student days when I was required to purchase expensive course textbooks, which were long and tedious to read. I said to myself, “Something must be done,” and started a publishing company to produce textbooks that explain math and physics concepts clearly, concisely, and affordably.

Minireference Co. is a publisher specializing in math and science. Our goal is to make advanced math modelling tools accessible to everyone by producing affordable textbooks that explain math concepts clearly and concisely. It’s time for a break from traditional textbooks that are expensive, heavy, and tedious to read. The books in the **No Bullshit Guide** series are an efficient and enjoyable alternative for learning mathematical subjects.

About the author

I have been teaching math and physics for more than 17 years as a private tutor. Through this experience, I learned to break complicated ideas into smaller, interconnected chunks that are easy to understand. An interesting feedback loop occurs when students learn concepts in small, manageable chunks: they experience *knowledge buzz* whenever concepts “click” into place, and this excitement motivates them to continue learning more. I know this from first-hand experience, both as a teacher and as a student. I completed my undergraduate studies in electrical engineering, then stayed on to earn a M.Sc. in physics, and a Ph.D. in computer science from McGill University. Linear algebra played a central role throughout my studies. With this book, I want to share with you some of what I’ve learned about this expansive subject.

Ivan Savov
Montreal, 2020

Introduction

In recent years we've seen countless advances in science and technology. Modern science and engineering fields have developed advanced models for understanding the real world, predicting the outcomes of experiments, and building useful technology. Although we're still far from obtaining a "theory of everything" that can fully explain reality and predict the future, we do have a significant understanding of the natural world on many levels: physical, chemical, biological, ecological, psychological, and social. And, since mathematical models are leveraged throughout these fields of study, anyone interested in contributing to scientific and technological advances must also understand mathematics.

The linear algebra techniques you'll learn in this book are some of the most powerful mathematical modelling tools that exist. At the core of linear algebra lies a very simple idea: *linearity*. A function f is *linear* if it obeys the equation

$$f(ax_1 + bx_2) = af(x_1) + bf(x_2),$$

where x_1 and x_2 are any two inputs of the function. We use the term *linear combination* to describe any expression constructed from a set of variables by multiplying each variable by a constant and adding the results. In the above equation, the linear combination $ax_1 + bx_2$ of the inputs x_1 and x_2 is transformed into the linear combination $af(x_1) + bf(x_2)$ of the outputs of the function $f(x_1)$ and $f(x_2)$. **Essentially, linear functions transform a linear combination of inputs into the same linear combination of outputs.** If the input to the linear function f consists of five parts x_1 and three parts x_2 , then the output of the function will consist of five parts $f(x_1)$ and three parts $f(x_2)$. That's it, that's all! Now you know everything there is to know about linear algebra. The rest of the book is just details.

Linear models are super useful

A significant proportion of the math models used in science describe *linear relationships* between quantities. Mathematicians, scientists, engineers, and business analysts develop and use linear models to make sense of the systems they study. Linear models are popular because they are **easy to describe mathematically**. We can obtain the parameters of a linear model for a real-world system by analyzing the system's behaviour for relatively few inputs. Let's illustrate this important point with an example.

Example You're visiting an art gallery. Inside, the screen of a tablet computer is being projected onto a giant wall. Anything you draw on the tablet instantly appears projected onto the wall. However, the tablet's user interface doesn't give any indication about how to hold the tablet "right side up." How can you find the correct orientation of the tablet so your drawing won't appear rotated or upside-down?




Figure 6: An unknown linear transformation T maps “tablet coordinates” to “wall coordinates.” How can we characterize T ?

The tablet's screen is a two-dimensional *input space* described by coordinates (x, y) and the wall projection is a two-dimensional *output space* described by wall coordinates (x', y') . You're looking for the unknown transformation T that maps the pixels of the tablet screen (the input space) to the projection on the wall (the output space):

$$(x, y) \xrightarrow{T} (x', y').$$

This task is directly analogous to the tasks scientists and engineers face every day when trying to model real-world systems by observing how systems transform inputs to outputs. If the unknown transformation T is linear, you can learn what it is very quickly, using only two swipes on the tablet screen.

To understand how T transforms screen coordinates (x, y) to wall coordinates (x', y') , you can use this two-step “probing” procedure:

1. Draw a horizontal line on the tablet to represent the x -direction in the input space $\rightarrow = (1, 0)$. You observe the output $\nearrow$ projected on the wall. This tells you horizontal lines are transformed to northeast diagonal lines in the wall-projection space.
2. Draw a vertical line in the y -direction $\uparrow = (0, 1)$ on the tablet. You observe the output $\nwarrow$ appears on the wall. This means vertical lines on the tablet screen turn into northwest diagonal lines when projected on the wall.




Figure 7: Drawing a short horizontal arrow $\rightarrow$ on the tablet screen results in a northeast diagonal projection on the wall $\nearrow$. Drawing a vertical arrow $\uparrow$ on the tablet results in a northwest diagonal line projected on the wall $\nwarrow$.

Here comes the interesting part: now that you know the outputs $\nearrow$ and $\nwarrow$ produced for the two input directions, you can **predict the linear transformation's output for any other input**. Let's look at the math equations that show why this is true.

Suppose you want to predict what will appear on the wall if you draw a line on the tablet in the direction $(3, 2)$. The coordinates $(3, 2)$ describe a swipe with length 3 units in the x -direction and 2 units in the y -direction. The input coordinates $(3, 2)$ can be written as $3(1, 0) + 2(0, 1) = 3\rightarrow + 2\uparrow$. Because you know T is linear, the wall projection of this input will have a length equal to 3 times the x -direction output $\nearrow$ plus 2 times the y -direction output $\nwarrow$:

$$T(3\rightarrow + 2\uparrow) = 3T(\rightarrow) + 2T(\uparrow) = 3\nearrow + 2\nwarrow.$$




Figure 8: The linear transformation T maps the input $3\rightarrow + 2\uparrow$ to the output $3T(\rightarrow) + 2T(\uparrow) = 3\nearrow + 2\nwarrow$.

Knowing that the input $\rightarrow$ produces the output $\nearrow$ and the input $\uparrow$ produces the output $\nwarrow$ allows you to determine the linear transformation's output for all other inputs. Every input (a, b) can be written as a linear combination: $(a, b) = a(1, 0) + b(0, 1) = a \rightarrow + b \uparrow$. Since you know T is linear, you know the corresponding output will be

$$T(a \rightarrow + b \uparrow) = aT(\rightarrow) + bT(\uparrow) = a \nearrow + b \nwarrow.$$

Since you can predict the output of T for all possible inputs, you have obtained a complete characterization of the linear transformation T .

The probing procedure we used to characterize the two-dimensional tablet-to-wall linear transformation (denoted $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$) can be used to study arbitrary linear transformations with n -dimensional inputs and m -dimensional outputs (denoted $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$). **Knowing the outputs of a linear transformation T for all “directions” in its input space gives us a complete characterization of T .**

TL;DR The linear property allows us to analyze multidimensional systems and processes by studying their effects on a small set of inputs. This is the essential reason linear models are used so widely in science. Without this linear structure, characterizing the behaviour of unknown input-output systems would be a much harder task.

Linear transformations

Linear transformations will be a central topic throughout this book. You can think of linear transformations as “vector functions” and understand their properties as analogous to the properties of the regular functions you’re familiar with. The action of a function on a number is similar to the action of a linear transformation on a vector:

$$\begin{aligned} \text{function } f : \mathbb{R} &\rightarrow \mathbb{R} &\Leftrightarrow &\text{linear transformation } T : \mathbb{R}^n \rightarrow \mathbb{R}^m \\ \text{input } x \in \mathbb{R} &&\Leftrightarrow &\text{input } \vec{x} \in \mathbb{R}^n \\ \text{output } f(x) \in \mathbb{R} &&\Leftrightarrow &\text{output } T(\vec{x}) \in \mathbb{R}^m \\ \text{inverse function } f^{-1} &&\Leftrightarrow &\text{inverse transformation } T^{-1} \\ \text{roots of } f &&\Leftrightarrow &\text{kernel of } T \end{aligned}$$

Studying linear algebra will expose you to many new topics associated with linear transformations. You’ll learn about concepts like vector spaces, projections, rotations, and orthogonalization procedures. Indeed, a first linear algebra course introduces many advanced, abstract ideas; yet all the new ideas you’ll encounter can be seen as extensions of ideas you’re already familiar with. Linear algebra is the vector-upgrade to your high school knowledge of functions.

Prerequisites

To understand linear algebra, you must have some preliminary knowledge of fundamental math concepts like numbers, equations, and functions. For example, you should be able to tell me the meaning of the parameters m and b in the equation $f(x) = mx + b$. If you do not feel confident about your basic math skills, don't worry. Chapter 1 is a prerequisites chapter specially designed to help bring you quickly up to speed on the high school math material.

Executive summary

The book is organized into nine chapters. Chapters 2 through 6 are the core of linear algebra. Chapters 7 through 9 contain optional reading about linear algebra applications. The concept maps on pages v, vi, and vii illustrate the connections between the topics we'll cover. I know the maps may seem informationally intimidating at first sight, but don't worry—the book is split into tiny chunks, and we'll navigate the material step by step. It will be like Mario World, but in n dimensions and with a lot of bonus levels.

Chapter 2 is a general introduction to linear algebra. Linear algebra is the math of vectors and matrices, so we'll start by defining the mathematical operations we can perform on vectors and matrices.

In Chapter 3, we'll tackle the computational aspects of linear algebra. By the end of this chapter you'll know how to solve systems of equations, transform a matrix into its *reduced row echelon form*, compute the product of two matrices, and find the *determinant* and the *inverse* of a square matrix. Each of these computational tasks can be tedious to carry out by hand and can require lots of steps. There is no way around this; we must do the grunt work before we get to the cool stuff.

In Chapter 4, we'll review the properties and the equations that describe basic geometric objects like points, lines, and planes. We'll learn how to compute projections onto vectors, projections onto planes, and distances between objects. We'll also review the meaning of vector coordinates, which are lengths measured with respect to a basis. We'll learn about linear combinations of vectors, the *span* of a set of vectors, and formally define what a *vector space* is.

Chapter 5 is about linear transformations. Armed with the computational tools from Chapter 3 and the geometric intuition from Chapter 4, we can tackle the core subject of linear algebra: linear transformations. We'll explore in detail the correspondence between linear transformations ($T : \mathbb{R}^n \rightarrow \mathbb{R}^m$) and their representation as $m \times n$ matrices. We'll also learn how the entries in a matrix represen-

tation depend on the choice of basis for the input and output spaces of the transformation. Section 5.4 on the invertible matrix theorem serves as a midway checkpoint for your understanding of linear algebra. This theorem connects several seemingly disparate concepts: reduced row echelon forms, matrix inverses, and determinants. The invertible matrix theorem links all these concepts and highlights the properties of invertible linear transformations that distinguish them from non-invertible transformations.

Chapter 6 covers more advanced theoretical topics of linear algebra. We'll define the *eigenvalues* and the *eigenvectors* of a square matrix. We'll see how the eigenvalues of a matrix tell us important information about the properties of the matrix, and learn about the special names given to different types of matrices based on the properties of their eigenvalues. In Section 6.3 we'll discuss *abstract vector spaces*. Abstract vectors are mathematical objects that—like vectors—have components and can be scaled, added, and subtracted by manipulating their components. Section 6.7 will discuss linear algebra with complex numbers.

In Chapter 7, we'll discuss the applications of linear algebra. If you've done your job learning the material in the first six chapters, you'll get to learn all the cool things you can do with linear algebra. Chapter 8 will introduce the basic concepts of probability theory. Chapter 9 contains an introduction to quantum mechanics.




Figure 9: The dependency structure of the chapters in the book.

Figure 9 shows the prerequisite structure between the chapters. If you're new to linear algebra, it would be best to read the chapters in order. If you already have some experience with the subject, you can skip ahead to the parts you're interested in. Use the table of contents, the concept maps on pages v–vii, and the index on page 577 to navigate the book. The chapters and sections in the book are self-contained so they will make sense even if you read them out of order.

Difficulty level

In terms of difficulty, I must prepare you to get ready for some serious uphill pushes. As your personal “trail guide” up the mountain of linear algebra, it’s my obligation to warn you about the difficulties that lie ahead.

Linear algebra is a difficult subject because it requires developing your computational skills, your geometric intuition, and your abstract thinking—all at the same time! The computational aspects of linear algebra are not particularly difficult, but they can be boring and repetitive. You’ll have to carry out hundreds of steps of basic arithmetic. The geometric problems you’ll encounter in Chapter 4 can be tough at first, but they’ll get easier once you learn to draw diagrams and develop your geometric reasoning. The theoretical aspects of linear algebra are difficult because they require a new way of thinking, which resembles what doing “real math” is like. You must not only understand and use the material; you must also know how to *prove* mathematical statements using the definitions and properties of math objects.

In summary, much toil awaits you as you learn the concepts of linear algebra, but the effort is totally worth it. All the brain sweat you put into understanding vectors and matrices will lead to mind-expanding insights. You will reap the benefits of your efforts for the rest of your life as your knowledge of linear algebra opens many doors for you.

Chapter 1

Math fundamentals

In this chapter we'll review the fundamental ideas of mathematics—the prerequisites for learning linear algebra. We'll define the different types of numbers and the concept of a function, which is a transformation that takes numbers as inputs and produces numbers as outputs. Linear algebra is the extension of these ideas to many dimensions: instead of doing math with numbers and functions, in linear algebra we'll be doing math with vectors and linear transformations.




Figure 1.1: A concept map showing the mathematical topics covered in this chapter. We'll learn how to solve equations using algebra, how to model the world using functions, and some important facts about geometry. The material in this chapter is required for your understanding of the more advanced topics in this book.

1.1 Solving equations

Most math skills boil down to being able to manipulate and solve equations. Solving an equation means finding the value of the unknown in the equation.

Check this shit out:

$$x^2 - 4 = 45.$$

To solve the above equation is to answer the question “What is x ?” More precisely, we want to find the number that can take the place of x in the equation so that the equality holds. In other words, we’re asking,

“Which number times itself minus four gives 45?”

That is quite a mouthful, don’t you think? To remedy this verbosity, mathematicians often use specialized symbols to describe math operations. The problem is that these specialized symbols can be very confusing. Sometimes even the simplest math concepts are inaccessible if you don’t know what the symbols mean.

What are your feelings about math, dear reader? Are you afraid of it? Do you have anxiety attacks because you think it will be too difficult for you? Chill! Relax, my brothers and sisters. There’s nothing to it. Nobody can magically guess the solution to an equation immediately. To find the solution, you must break the problem into simpler steps. Let’s walk through this one together.

To find x , we can manipulate the original equation, transforming it into a different equation (as true as the first) that looks like this:

$$x = \text{only numbers.}$$

That’s what it means to *solve* an equation: the equation is solved because the unknown is isolated on one side, while the constants are grouped on the other side. You can type the numbers on the right-hand side into a calculator and obtain the numerical value of x .

By the way, before we continue our discussion, let it be noted: the equality symbol ($=$) means that all that is to the left of $=$ is equal to all that is to the right of $=$. To keep this equality statement true, **for every change you apply to the left side of the equation, you must apply the same change to the right side of the equation.**

To find x , we need to manipulate the original equation into its final form, simplifying it step by step until it can’t be simplified any further. The only requirement is that the manipulations we make transform one true equation into another true equation. In this example, the first simplifying step is to add the number four to both sides of the equation:

$$x^2 - 4 + 4 = 45 + 4,$$

which simplifies to

$$x^2 = 49.$$

Now the expression looks simpler, yes? How did I know to perform this operation? I wanted to “undo” the effects of the operation -4 . We undo an operation by applying its *inverse*. In the case where the operation is the subtraction of some amount, the inverse operation is the addition of the same amount. We’ll learn more about function inverses in Section 1.4.

We’re getting closer to our goal of *isolating* x on one side of the equation, leaving only numbers on the other side. The next step is to undo the square x^2 operation. The inverse operation of squaring a number x^2 is to take its square root $\sqrt{\quad}$, so that’s what we’ll do next. We obtain

$$\sqrt{x^2} = \sqrt{49}.$$

Notice how we applied the square root to both sides of the equation? If we don’t apply the same operation to both sides, we’ll break the equality!

The equation $\sqrt{x^2} = \sqrt{49}$ simplifies to

$$|x| = 7.$$

What’s up with the vertical bars around x ? The notation $|x|$ stands for the *absolute value* of x , which is the same as x except we ignore the sign that indicates whether x is positive or negative. For example $|5| = 5$ and $|-5| = 5$, too. The equation $|x| = 7$ indicates that both $x = 7$ and $x = -7$ satisfy the equation $x^2 = 49$. Seven squared is 49, $7^2 = 49$, and negative seven squared is also 49, $(-7)^2 = 49$, because the two negative signs cancel each other out.

The final solutions to the equation $x^2 - 4 = 45$ are

$$x = 7 \quad \text{and} \quad x = -7.$$

Yes, there are *two* possible answers. You can check that both of the above values of x satisfy the initial equation $x^2 - 4 = 45$.

If you are comfortable with all the notions of high school math and you feel you could have solved the equation $x^2 - 4 = 45$ on your own, then you can skim through this chapter quickly. If on the other hand you are wondering how the squiggle killed the power two, then this chapter is for you! In the following sections we will review all the essential concepts from high school math that you will need to power through the rest of this book. First, let me tell you about the different kinds of numbers.

1.2 Numbers

In the beginning, we must define the main players in the world of math: numbers.

Definitions

Numbers are the basic objects we use to count, measure, quantify, and calculate things. Mathematicians like to classify the different kinds of number-like objects into categories called *sets*:

- The natural numbers: $\mathbb{N} = \{0, 1, 2, 3, 4, 5, 6, 7, \dots\}$
- The integers: $\mathbb{Z} = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$
- The rational numbers: $\mathbb{Q} = \{\frac{5}{3}, \frac{22}{7}, 1.5, 0.125, -7, \dots\}$
- The real numbers: $\mathbb{R} = \{-1, 0, 1, \sqrt{2}, e, \pi, 4.94\dots, \dots\}$
- The complex numbers: $\mathbb{C} = \{-1, 0, 1, i, 1 + i, 2 + 3i, \dots\}$

These categories of numbers should be somewhat familiar to you. Think of them as neat classification labels for everything that you would normally call a number. Each group in the above list is a *set*. A set is a collection of items of the same kind. Each collection has a name and a precise definition for which items belong in that collection. Note also that each of the sets in the list contains all the sets above it, as illustrated in Figure 1.2. For now, we don't need to go into the details of sets and set notation, but we do need to be aware of the different sets of numbers.




Figure 1.2: An illustration of the nested containment structure of the different number sets. The set of natural numbers is contained in the set of integers, which in turn is contained in the set of rational numbers. The set of rational numbers is contained in the set of real numbers, which is contained in the set of complex numbers.

Why do we need so many different sets of numbers? Each set of numbers is associated with more and more advanced mathematical problems.

The simplest numbers are the natural numbers $\mathbb{N}$, which are sufficient for all your math needs if all you're going to do is *count* things. How many goats? Five goats here and six goats there so the total is

11 goats. The sum of any two natural numbers is also a natural number.

As soon as you start using *subtraction* (the inverse operation of addition), you start running into negative numbers, which are numbers outside the set of natural numbers. If the only mathematical operations you will ever use are *addition* and *subtraction*, then the set of integers $\mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$ will be sufficient. Think about it. Any integer plus or minus any other integer is still an integer.

You can do a lot of interesting math with integers. There is an entire field in math called *number theory* that deals with integers. However, to restrict yourself solely to integers is somewhat limiting—a rotisserie menu that offers $\frac{1}{2}$ of a chicken would be totally confusing.

If you want to use division in your mathematical calculations, you'll need the rationals $\mathbb{Q}$. The set of rational numbers corresponds to all numbers that can be expressed as *fractions* of the form $\frac{m}{n}$ where m and n are integers, and $n \neq 0$. You can add, subtract, multiply, and divide rational numbers, and the result will always be a rational number. However, even the rationals are not enough for all of math!

In geometry, we can obtain *irrational* quantities like $\sqrt{2}$ (the diagonal of a square with side 1) and π (the ratio between a circle's circumference and its diameter). There are no integers x and y such that $\sqrt{2} = \frac{x}{y}$, therefore we say that $\sqrt{2}$ is *irrational* (not in the set $\mathbb{Q}$). An irrational number has an infinitely long decimal expansion that doesn't repeat. For example, $\pi = 3.141592653589793\dots$ where the dots indicate that the decimal expansion of π continues all the way to infinity.

Combining the irrational numbers with the rationals gives us all the useful numbers, which we call the set of real numbers $\mathbb{R}$. The set $\mathbb{R}$ contains the integers, the rational numbers $\mathbb{Q}$, as well as irrational numbers like $\sqrt{2} = 1.4142135\dots$. By using the reals you can compute pretty much anything you want. From here on in the text, when I say *number*, I mean an element of the set of real numbers $\mathbb{R}$.

The only thing you can't do with the reals is to take the square root of a negative number—you need the complex numbers $\mathbb{C}$ for that. We defer the discussion on complex numbers until Section 1.14.

Operations on numbers

Addition

You can add numbers. I'll assume you're familiar with this stuff:

$$2 + 3 = 5, \quad 45 + 56 = 101, \quad 9999 + 1 = 10\,000.$$

You can visualize numbers as sticks of different length. Adding numbers is like adding sticks together: the resulting stick has a length

equal to the sum of the lengths of the constituent sticks, as illustrated in Figure 1.3.



Figure 1.3: The addition of numbers corresponds to adding lengths.

Addition is *commutative*, which means that $a + b = b + a$. In other words, the order of the numbers in a summation doesn't matter. It is also *associative*, which means that if you have a long summation like $a + b + c$ you can compute it in any order $(a + b) + c$ or $a + (b + c)$, and you'll get the same answer.

Subtraction

Subtraction is the inverse operation of addition:

$$2 - 3 = -1, \quad 45 - 56 = -11, \quad 999 - 1 = 998.$$

Unlike addition, subtraction is not a commutative operation. The expression $a - b$ is not equal to the expression $b - a$, or written mathematically:

$$a - b \neq b - a.$$

Instead we have $b - a = -(a - b)$, which shows that changing the order of a and b in the expression changes its sign.

Subtraction is not associative either:

$$(a - b) - c \neq a - (b - c).$$

For example $(7 - 2) - 3 = 2$ while $7 - (2 - 3) = 8$.

Multiplication

You can also multiply numbers together:

$$ab = \underbrace{a + a + \cdots + a}_{b \text{ times}} = \underbrace{b + b + \cdots + b}_{a \text{ times}}.$$

Note that multiplication can be defined in terms of repeated addition.

The visual way to think about multiplication is as an area calculation. The area of a rectangle of width a and height b is equal to ab . A rectangle with a height equal to its width is a square, and this is why we call $aa = a^2$ “ a squared.”




Figure 1.4: The area of a rectangle with width 3 m and height 2 m is equal to 6 m^2 , which is equivalent to six squares with area 1 m^2 each.

Multiplication of numbers is also commutative, $ab = ba$, and associative, $abc = (ab)c = a(bc)$. In modern math notation, no special symbol is required to denote multiplication; we simply put the two factors next to each other and say the multiplication is *implicit*. Some other ways to denote multiplication are $a \cdot b$, $a \times b$, and, on computer systems, $a * b$.

Division

Division is the inverse operation of multiplication.

$$a/b = \frac{a}{b} = a \div b = \text{one } b^{\text{th}} \text{ of } a.$$

Whatever a is, you need to divide it into b equal parts and take one such part.

Division is not a commutative operation since a/b is not equal to b/a . Division is not associative either: $(a \div b) \div c \neq a \div (b \div c)$. For example, when $a = 6$, $b = 3$, and $c = 2$, we get $(6/3)/2 = 1$ while $6/(3/2) = 4$.

Note that you cannot divide by 0. Try it on your calculator or computer. It will say “error divide by zero” because this action simply doesn’t make sense. After all, what would it mean to divide something into zero equal parts?

Exponentiation

The act of multiplying a number by itself many times is called *exponentiation*. We denote “ a exponent n ” using a superscript, where n is the number of times the base a is multiplied by itself:

$$a^n = \underbrace{aaa \cdots a}_{n \text{ times}}.$$

In words, we say “ a raised to the power of n .”

To visualize how exponents work, we can draw a connection between the value of exponents and the dimensions of geometric objects. Figure 1.5 illustrates how the same length 2 corresponds to different geometric objects when raised to different exponents. The

number 2 corresponds to a line segment of length two, which is a geometric object in a one-dimensional space. If we add a line segment of length two in a second dimension, we obtain a square with area 2^2 in a two-dimensional space. Adding a third dimension, we obtain a cube with volume 2^3 in a three-dimensional space. Indeed, raising a base a to the exponent 2 is commonly called “ a squared,” and raising a to the power of 3 is called “ a cubed.”

The geometrical analogy about one-dimensional quantities as lengths, two-dimensional quantities as areas, and three-dimensional quantities as volumes is good to keep in mind.




Figure 1.5: Geometric interpretation for exponents 1, 2, and 3. A length raised to exponent 2 corresponds to the area of a square. The same length raised to exponent 3 corresponds to the volume of a cube.

Our visual intuition works very well up to three dimensions, but we can use other means of visualizing higher exponents, as demonstrated in Figure 1.6.

Operator precedence

There is a standard convention for the order in which mathematical operations must be performed. The basic algebra operations have the following precedence:

1. Parentheses
2. Exponents
3. Multiplication and Division
4. Addition and Subtraction

If you’re seeing this list for the first time, the acronym PEMDAS and the associated mnemonic “Please Excuse My Dear Aunt Sally,” might help you remember the order of operations.

For instance, the expression $5 \cdot 3^2 + 13$ is interpreted as “First find the square of 3, then multiply it by 5, and then add 13.” Parentheses are needed to carry out the operations in a different order: to multiply 5 times 3 first and *then* take the square, the equation should read




Figure 1.6: Visualization of numbers raised to different exponents. Each box in this grid contains a^n dots, where the base a varies from one through five, and the exponent n varies from one through five. In the first row we see that the number $a = 1$ raised to any exponent is equal to itself. The second row corresponds to the base $a = 2$ so the number of dots doubles each time we increase the exponent by one. Starting from $2^1 = 2$ in the first column, we end up with $2^5 = 32$ in the last column. The rest of the rows show how exponentiation works for different bases.

$(5 \cdot 3)^2 + 13$, where parentheses indicate that the square acts on $(5 \cdot 3)$ as a whole and not on 3 alone.

Exercises

E1.1 Solve for the unknown x in the following equations:

a) $3x + 2 - 5 = 4 + 2$

b) $\frac{1}{2}x - 3 = \sqrt{3} + 12 - \sqrt{3}$

c) $\frac{7x-4}{2} + 1 = 8 - 2$

d) $5x - 2 + 3 = 3x - 5$

E1.2 Indicate all the number sets the following numbers belong to.

a) -2

b) $\sqrt{-3}$

c) $8 \div 4$

d) $\frac{5}{3}$

e) $\frac{\pi}{2}$

E1.3 Calculate the values of the following expressions:

a) $2^3 3 - 3$

b) $2^3(3 - 3)$

c) $\frac{4-2}{3^3}(6 \cdot 7 - 41)$

1.3 Variables

In math we use a lot of *variables* and *constants*, which are placeholder names for *any* number or unknown. Variables allow us to perform calculations without knowing all the details.

Example You're having tacos for lunch today and wondering how many you can eat without going over your caloric budget. Your goal is to eat 800 calories for lunch and you want to do the calculation before getting to the restaurant because you fear your math abilities might be affected in the presence of tacos. You're not sure how many calories each taco contains, so you invent the variable c to denote this unknown. You also define the variable x to represent the number of tacos you will eat, and come up with the equation $800 = cx$ to represent the total number of calories of your lunch. Solving for x , you find the total number of tacos you should order is $x = \frac{800}{c}$. If the restaurant serves tacos that contain $c = 200$ calories each, then you should order $x = \frac{800}{200} = 4$ of them. If the restaurant serves only giant tacos worth $c = 400$ calories each, then you can only eat $x = \frac{800}{400} = 2$ of them. Observe we were able to solve for x even before knowing the value of c .

Variable names

There are common naming patterns for variables:

- x : name used for the unknown in equations. We also use x to denote function inputs and the position of objects in physics.
- i, j, k, m, n : common names for integer variables
- a, b, c, d : letters near the beginning of the alphabet are often used to denote constants (fixed quantities that do not change).
- θ, ϕ : the Greek letters *theta* and *phi* are used to denote angles
- C : costs in business, along with P for profit, and R for revenue
- X : capital letters are used to denote random variables in probability theory

Variable substitution

We can often *change variables* and replace one unknown variable with another to simplify an equation. For example, say you don't feel comfortable around square roots. Every time you see a square root, you freak out until one day you find yourself taking an exam trying

to solve for x in the following equation:

$$\frac{6}{5 - \sqrt{x}} = \sqrt{x}.$$

Don't freak out! In crucial moments like this, substitution can help with your root phobia. Just write, "Let $u = \sqrt{x}$ " on your exam, and voila, you can rewrite the equation in terms of the variable u :

$$\frac{6}{5 - u} = u,$$

which contains no square roots.

The next step to solve for u is to undo the division operation. Multiply both sides of the equation by $(5 - u)$ to obtain

$$\frac{6}{5 - u}(5 - u) = u(5 - u),$$

which simplifies to

$$6 = 5u - u^2.$$

This can be rewritten as the equation $u^2 - 5u + 6 = 0$, which in turn can be rewritten as $(u - 2)(u - 3) = 0$ using the techniques we'll learn in Section 1.6.

We now see that the solutions are $u_1 = 2$ and $u_2 = 3$. The last step is to convert our u -answers into x -answers by using $u = \sqrt{x}$, which is equivalent to $x = u^2$. The final answers are $x_1 = 2^2 = 4$ and $x_2 = 3^2 = 9$. Try plugging these x values into the original square root equation to verify that they satisfy it.

Compact notation

Symbolic manipulation is a powerful tool because it allows us to manage complexity. Say you're solving a physics problem in which you're told the mass of an object is $m = 140$ kg. If there are many steps in the calculation, would you rather use the number 140 kg in each step, or the shorter symbol m ? It's much easier to use m throughout your calculation, and wait until the last step to substitute the value 140 kg when computing the final numerical answer.

1.4 Functions and their inverses

As we saw in the section on solving equations, the ability to "undo" functions is a key skill for solving equations.

Example Suppose we're solving for x in the equation

$$f(x) = c,$$

where f is some function and c is some constant. We're looking for the unknown x such that $f(x)$ equals c . Our goal is to isolate x on one side of the equation, but the function f stands in our way.

By using the *inverse function* (denoted f^{-1}) we “undo” the effects of f . We apply the inverse function f^{-1} to both sides of the equation to obtain

$$f^{-1}(f(x)) = f^{-1}(c).$$

By definition, the inverse function f^{-1} performs the opposite action of the function f , so together the two functions cancel each other out. We have $f^{-1}(f(x)) = x$ for any number x .

Provided everything is kosher (the function f^{-1} must be defined for the input c), the manipulation we made above is valid and we have obtained the answer $x = f^{-1}(c)$.

The above example introduces the notation f^{-1} for denoting the inverse function. This notation is inspired by the notation for reciprocals. Recall that multiplication by the reciprocal number a^{-1} is the inverse operation of multiplication by the number a : $a^{-1}ax = 1x = x$. In the case of functions, however, the negative-one exponent does not refer to “one over- $f(x)$ ” as in $\frac{1}{f(x)} = (f(x))^{-1}$; rather, it refers to the inverse function. In other words, the number $f^{-1}(y)$ is equal to the number x such that $f(x) = y$.

Be careful: sometimes an equation can have multiple solutions. For example, the function $f(x) = x^2$ maps two input values (x and $-x$) to the same output value $x^2 = f(x) = f(-x)$. The inverse function of $f(x) = x^2$ is $f^{-1}(y) = \sqrt{y}$, but both $x = +\sqrt{c}$ and $x = -\sqrt{c}$ are solutions to the equation $x^2 = c$. In this case, this equation's solutions can be indicated in shorthand notation as $x = \pm\sqrt{c}$.

Formulas

Here is a list of common functions and their inverses:

$$\text{function } f(x) \Leftrightarrow \text{inverse } f^{-1}(x)$$

$$x + 2 \Leftrightarrow x - 2$$

$$2x \Leftrightarrow \frac{1}{2}x$$

$$-1x \Leftrightarrow -1x$$

$$x^2 \Leftrightarrow \pm\sqrt{x}$$

$$2^x \Leftrightarrow \log_2(x)$$

$$\begin{aligned}
3x + 5 &\Leftrightarrow \frac{1}{3}(x - 5) \\
a^x &\Leftrightarrow \log_a(x) \\
\exp(x) = e^x &\Leftrightarrow \ln(x) = \log_e(x) \\
\sin(x) &\Leftrightarrow \sin^{-1}(x) = \arcsin(x) \\
\cos(x) &\Leftrightarrow \cos^{-1}(x) = \arccos(x)
\end{aligned}$$

The function-inverse relationship is *symmetric*—if you see a function on one side of the above table (pick a side, any side), you’ll find its inverse on the opposite side.

Don’t be surprised to see $-1x \Leftrightarrow -1x$ in the list of function inverses. Indeed, the opposite operation of multiplying by -1 is to multiply by -1 once more: $(-(-x) = x)$.

Example 1

If you want to solve the equation $x - 4 = 5$, you can apply the inverse function of $x - 4$, which is $x + 4$. After adding four to both sides of the equation, $x - 4 + 4 = 5 + 4$, we obtain the answer $x = 9$.

Example 2

Let’s say your teacher doesn’t like you and right away, on the first day of class, he gives you a serious equation and tells you to find x :

$$\log_5 \left(3 + \sqrt{6\sqrt{x} - 7} \right) = 34 + \sin(8) - \Psi(1).$$

See what I mean when I say the teacher doesn’t like you?

First, note that it doesn’t matter what Ψ (the Greek letter *psi*) is, since x is on the other side of the equation. You can keep copying $\Psi(1)$ from line to line, until the end, when you throw the ball back to the teacher. “My answer is in terms of *your* variables, dude. *You* go figure out what the hell Ψ is since you brought it up in the first place!” By the way, it’s not actually recommended to quote me verbatim should a situation like this arise. The same goes with $\sin(8)$. If you don’t have a calculator handy, don’t worry about it. Keep the expression $\sin(8)$ instead of trying to find its numerical value. In general, try to work with variables as much as possible and leave the numerical computations for the last step.

Okay, enough beating about the bush. Let’s just find x and get it over with! On the right-hand side of the equation, we have the sum of a bunch of terms with no x in them, so we’ll leave them as they are. On the left-hand side, the outermost function is a logarithm

base 5. Cool. Looking at the table of inverse functions, we find the exponential function is the inverse of the logarithm: $a^x \Leftrightarrow \log_a(x)$. To get rid of $\log_5$, we must apply the exponential function base 5 to both sides:

$$5^{\log_5(3 + \sqrt{6\sqrt{x} - 7})} = 5^{34 + \sin(8) - \Psi(1)},$$

which simplifies to

$$3 + \sqrt{6\sqrt{x} - 7} = 5^{34 + \sin(8) - \Psi(1)},$$

since 5^x cancels $\log_5 x$.

From here on, it is going to be as if Bruce Lee walked into a place with lots of bad guys. Addition of 3 is undone by subtracting 3 on both sides:

$$\sqrt{6\sqrt{x} - 7} = 5^{34 + \sin(8) - \Psi(1)} - 3.$$

To undo a square root we take the square:

$$6\sqrt{x} - 7 = \left(5^{34 + \sin(8) - \Psi(1)} - 3\right)^2.$$

Add 7 to both sides,

$$6\sqrt{x} = \left(5^{34 + \sin(8) - \Psi(1)} - 3\right)^2 + 7,$$

divide by 6

$$\sqrt{x} = \frac{1}{6} \left(\left(5^{34 + \sin(8) - \Psi(1)} - 3\right)^2 + 7 \right),$$

and square again to find the final answer:

$$x = \left[\frac{1}{6} \left(\left(5^{34 + \sin(8) - \Psi(1)} - 3\right)^2 + 7 \right) \right]^2.$$

Did you see what I was doing in each step? Next time a function stands in your way, hit it with its inverse so it knows not to challenge you ever again.

Discussion

The recipe I have outlined above is not universally applicable. Sometimes x isn't alone on one side. Sometimes x appears in several places in the same equation. In these cases, you can't effortlessly work your way, Bruce Lee-style, clearing bad guys and digging toward x —you need other techniques.

The bad news is there's no general formula for solving complicated equations. The good news is the above technique of "digging toward the x " is sufficient for 80% of what you are going to be doing. You can get another 15% if you learn how to solve the quadratic equation:

$$ax^2 + bx + c = 0.$$

We'll show a formula for solving quadratic equations in Section 1.6. Solving cubic equations like $ax^3 + bx^2 + cx + d = 0$ using a formula is also possible, but at this point you might as well start using a computer to solve for the unknowns.

There are all kinds of other equations you can learn how to solve: equations with multiple variables, equations with logarithms, equations with exponentials, and equations with trigonometric functions. The principle of "digging" toward the unknown by applying inverse functions is the key for solving all these types of equations, so be sure to practice using it.

Exercises

E1.4 Solve for x in the following equations:

a) $3x = 6$

b) $\log_5(x) = 2$

c) $\log_{10}(\sqrt{x}) = 1$

E1.5 Find the function inverse and use it to solve the problems.

a) Solve the equation $f(x) = 4$, where $f(x) = \sqrt{x}$.

b) Solve for x in the equation $g(x) = 1$, given $g(x) = e^{-2x}$.

1.5 Basic rules of algebra

It's important that you know the general rules for manipulating numbers and variables, a process otherwise known as—you guessed it—*algebra*. This little refresher will cover these concepts to make sure you're comfortable on the algebra front. We'll also review some important algebraic tricks, like *factoring* and *completing the square*, which are useful when solving equations.

Let's define some terminology for referring to different parts of math expressions. When an expression contains multiple things added together, we call those things *terms*. Furthermore, terms are usually composed of many things multiplied together. When a number x is obtained as the product of other numbers like $x = abc$, we say " x factors into a , b , and c ." We call a , b , and c the *factors* of x .

Given any three numbers a , b , and c , we can apply the following algebraic properties:




Figure 1.7: Diagram showing the names used to describe the different parts of the equation $abc + de = 0$.

1. Associative property: $a + b + c = (a + b) + c = a + (b + c)$ and $abc = (ab)c = a(bc)$
2. Commutative property: $a + b = b + a$ and $ab = ba$
3. Distributive property: $a(b + c) = ab + ac$

We use the distributive property every time we *expand* brackets. For example $a(b + c + d) = ab + ac + ad$. The brackets, also known as parentheses, indicate the expression $(b + c + d)$ must be treated as a whole; as a factor consisting of three terms. Multiplying this expression by a is the same as multiplying each term by a .

The opposite operation of expanding is called *factoring*, which consists of rewriting the expression with the common parts taken out in front of a bracket: $ab + ac = a(b + c)$. In this section, we'll discuss all algebra operations and illustrate what they're capable of.

Example Suppose we are asked to solve for t in the equation

$$7(3 + 4t) = 11(6t - 4).$$

Since the unknown t appears on both sides of the equation, it is not immediately obvious how to proceed.

To solve for t , we can bring all t terms to one side and all constant terms to the other side. First, expand the two brackets to obtain

$$21 + 28t = 66t - 44.$$

Then move things around to relocate all t s to the equation's right-hand side and all constants to the left-hand side:

$$21 + 44 = 66t - 28t.$$

We see t is contained in both terms on the right-hand side, so we can "factor it out" by rewriting the equation as

$$21 + 44 = t(66 - 28).$$

The answer is within close reach: $t = \frac{21+44}{66-28} = \frac{65}{38}$.

Expanding brackets

To *expand* a bracket is to multiply each term inside the bracket by the factor outside the bracket. The key thing to remember when expanding brackets is to apply the *distributive* property: $a(x + y) = ax + ay$. For longer expressions, we may need to apply the distributive property several times, until there are no more brackets left:

$$\begin{aligned}(a + b)(x + y + z) &= a(x + y + z) + b(x + y + z) \\ &= ax + ay + az + bx + by + bz.\end{aligned}$$

After expanding the brackets in this expression, we end up with six terms—one term for each of the six possible combinations of products between the terms in $(a + b)$ and the terms in $(x + y + z)$.

The distributive property is often used to manipulate expressions containing different powers of the variable x . For instance,

$$(x + 3)(x + 2) = x(x + 2) + 3(x + 2) = x^2 + x2 + 3x + 6.$$

We can use the commutative property on the second term $x2 = 2x$, then combine the two x terms into a single term to obtain

$$(x + 3)(x + 2) = x^2 + 5x + 6.$$

The bracket-expanding and simplification techniques demonstrated above are very common in math, and I recommend you solve some algebra practice problems to get the hang of them. Most math textbooks skip simplification steps and jump straight to the answer, since they assume readers are capable of doing simplifications on their own. It would be too long (and annoying) to show the simplifications for each expression. For example, the sentence “We can rewrite $(x + 3)(x + 2)$ as $x^2 + 5x + 6$,” is the short version of the longer sentence, “We can apply the distributive property twice on $(x + 3)(x + 2)$ then combine the terms with the same power of x to get $x^2 + 5x + 6$.”

It’s not unusual for people to make math mistakes when expanding brackets and manipulating long algebra expressions. To avoid mistakes, use a step-by-step approach and apply only one operation in each step. Write legibly and keep the equations “organized” so it’s easy to check the calculations performed in each step. Consider this slightly-more-complicated algebraic expression and its expansion:

$$\begin{aligned}(x + a)(bx^2 + cx + d) &= x(bx^2 + cx + d) + a(bx^2 + cx + d) \\ &= bx^3 + cx^2 + dx + abx^2 + acx + ad \\ &= bx^3 + (c + ab)x^2 + (d + ac)x + ad.\end{aligned}$$

Note how we sorted the terms in the final expression according to the different powers of x , with the terms containing x^2 grouped together, and the terms containing x grouped together. This approach helps keep things organized when dealing with expressions containing many terms.

Factoring

Factoring involves “taking out” the common parts of a complicated expression in order to make the expression more compact. Suppose we’re given the expression $6x^2y + 15x$. We can simplify this expression by taking out the common factors and moving them in front of a bracket. Let’s see how to do this, step by step.

The expression $6x^2y + 15x$ has two terms. Let’s split each term into its constituent factors:

$$6x^2y + 15x = (3)(2)(x)(x)y + (5)(3)x.$$

Since factors x and 3 appear in both terms, we can *factor them out* like this:

$$6x^2y + 15x = 3x(2xy + 5).$$

The expression on the right shows $3x$ is common to both terms.

Here’s another example of factoring—notice the common factors are taken out and moved in front of the bracket:

$$2x^2y + 2x + 4x = 2x(xy + 1 + 2) = 2x(xy + 3).$$

Factoring quadratic expressions

A *quadratic expression* is an expression of the form $ax^2 + bx + c$. The expression contains a *quadratic term* ax^2 , a *linear term* bx , and a constant term c . The numbers a , b , and c are called *coefficients*: the quadratic coefficient is a , the linear coefficient is b , and the constant coefficient is c .

To *factor* the quadratic expression $ax^2 + bx + c$ is to rewrite it as the product of a constant and two factors like $(x + p)$ and $(x + q)$:

$$ax^2 + bx + c = a(x + p)(x + q).$$

Rewriting quadratic expressions in factored form helps us better understand and describe their properties.

Example Suppose we're asked to describe the properties of the function $f(x) = x^2 - 5x + 6$. Specifically, we're asked to find the function's *roots*, which are the values of x for which the function equals zero.

Factoring the expression $x^2 - 5x + 6$ helps us see its properties more clearly, and makes our job of finding the roots of $f(x)$ easier. The factored form of this quadratic expression is

$$f(x) = x^2 - 5x + 6 = (x - 2)(x - 3).$$

Now we can see at a glance that the values of x for which $f(x) = 0$ are $x = 2$ and $x = 3$. When $x = 2$, the factor $(x - 2)$ is zero and hence $f(x) = 0$. Similarly, when $x = 3$, the factor $(x - 3)$ is zero so $f(x) = 0$.

How did we know that the factors of $x^2 - 5x + 6$ are $(x - 2)$ and $(x - 3)$ in the above example? For simple quadratics like the one above, we can simply *guess* the values of p and q in the equation $x^2 - 5x + 6 = (x + p)(x + q)$. Before we start guessing, let's look at the expanded version of the product between $(x + p)$ and $(x + q)$:

$$(x + p)(x + q) = x^2 + (p + q)x + pq.$$

Note the linear term on the right-hand side contains the sum of the unknowns $(p + q)$, while the constant term contains their product pq . If we want the equation $x^2 - 5x + 6 = x^2 + (p + q)x + pq$ to hold, we must find two numbers p and q whose sum equals -5 and whose product equals 6 . After a couple of attempts we find $p = -2$ and $q = -3$. This guessing approach is an effective strategy for many of the factoring problems we will likely be asked to solve, since math teachers often choose simple numbers like ± 1 , ± 2 , ± 3 , or ± 4 for the constants p and q . For more complicated quadratic expressions, we'll need to use the quadratic formula, which we'll talk about in Section 1.6.

Common quadratic forms

Let's look at some common variations of quadratic expressions you might encounter when doing algebra calculations.

The quadratic expression $x^2 - p^2$ is called a *difference of squares*, and it can be obtained by multiplying the factors $(x + p)$ and $(x - p)$:

$$(x + p)(x - p) = x^2 \cancel{-xp} \cancel{+px} - p^2 = x^2 - p^2.$$

There's no linear term because the $-xp$ term cancels the $+px$ term. Any time you see an expression like $x^2 - p^2$, you can know it comes from a product of the form $(x + p)(x - p)$.

A *perfect square* is a quadratic expression that can be written as the product of repeated factors $(x + p)$:

$$x^2 + 2px + p^2 = (x + p)(x + p) = (x + p)^2.$$

Note $x^2 - 2qx + q^2 = (x - q)^2$ is also a perfect square.

Completing the square

In this section we'll learn about an ancient algebra technique called *completing the square*, which allows us to rewrite *any* quadratic expression of the form $x^2 + Bx + C$ as a perfect square plus some constant correction factor $(x + p)^2 + k$. This algebra technique was described in one of the first books on *al-jabr* (algebra), written by Al-Khwarizmi around the year 800 CE. The name "completing the square" comes from the ingenious geometric construction used by this procedure. Yes, we can use geometry to solve algebra problems!

We assume the starting point for the procedure is a quadratic expression whose quadratic coefficient is one, $x^2 + Bx + C$, and use capital letters B and C to denote the linear and constant coefficients. The capital letters are to avoid any confusion with the quadratic expression $ax^2 + bx + c$, for which $a \neq 1$. Note we can always write $ax^2 + bx + c$ as $a(x^2 + \frac{b}{a}x + \frac{c}{a})$ and apply the procedure to the expression inside the brackets, identifying $\frac{b}{a}$ with B and $\frac{c}{a}$ with C .

First let's rewrite the quadratic expression $x^2 + Bx + C$ by splitting the linear term into two equal parts:

$$x^2 + \frac{B}{2}x + \frac{B}{2}x + C.$$

We can interpret the first three terms geometrically as follows: the x^2 term corresponds to a square with side length x , while the two $\frac{B}{2}x$ terms correspond to rectangles with sides $\frac{B}{2}$ and x . See the left side of Figure 1.8 for an illustration.




Figure 1.8: To complete the square in the expression $x^2 + Bx + C$, we need to add the quantity $(\frac{B}{2})^2$, which corresponds to a square (shown in darker colour) with sides equal to half the coefficient of the linear term. We also subtract $(\frac{B}{2})^2$ so the overall value of the expression remains unchanged.

The square with area x^2 and the two rectangles can be positioned to form a larger square with side length $(x + \frac{B}{2})$. Note there's a small

piece of sides $\frac{B}{2}$ by $\frac{B}{2}$ missing from the corner. To *complete the square*, we can add a term $(\frac{B}{2})^2$ to this expression. To preserve the equality, we also subtract $(\frac{B}{2})^2$ from the expression to obtain:

$$\begin{aligned} x^2 + \frac{B}{2}x + \frac{B}{2}x + C &= \underbrace{x^2 + \frac{B}{2}x + \frac{B}{2}x + (\frac{B}{2})^2}_{(x + \frac{B}{2})^2} - (\frac{B}{2})^2 + C \\ &= (x + \frac{B}{2})^2 - (\frac{B}{2})^2 + C. \end{aligned}$$

The right-hand side of this equation describes the area of the square with side length $(x + \frac{B}{2})$, minus the area of the small square $(\frac{B}{2})^2$, plus the constant C , as illustrated on the right side of Figure 1.8.

We can summarize the entire procedure in one equation:

$$x^2 + Bx + C = \underbrace{(x + \frac{B}{2})^2}_{(1)} + \underbrace{C - (\frac{B}{2})^2}_{(2)}.$$

There are two things to remember when you want to apply the complete-the-square trick: (1) choose the constant inside the bracket to be $\frac{B}{2}$ (half of the linear coefficient), and (2) subtract $(\frac{B}{2})^2$ outside the bracket in order to keep the equation balanced.

Solving quadratic equations

Suppose we want to solve the quadratic equation $x^2 + Bx + C = 0$. It's not possible to solve this equation with the digging-toward-the- x approach from Section 1.1 (since x appears in both the quadratic term x^2 and the linear term Bx). Enter the completing-the-square trick!

Example Let's find the solutions of the equation $x^2 + 5x + 6 = 0$. The coefficient of the linear term is $B = 5$, so we choose $\frac{B}{2} = \frac{5}{2}$ for the constant inside the bracket, and subtract $(\frac{B}{2})^2 = (\frac{5}{2})^2$ outside the bracket to keep the equation balanced. Completing the square gives

$$x^2 + 5x + 6 = (x + \frac{5}{2})^2 + 6 - (\frac{5}{2})^2 = 0.$$

Next we use fraction arithmetic to simplify the constant terms in the expression: $6 - (\frac{5}{2})^2 = 6 \cdot \frac{4}{4} - \frac{25}{4} = \frac{24-25}{4} = \frac{-1}{4} = -0.25$.

We're left with the equation

$$(x + 2.5)^2 - 0.25 = 0,$$

which we can now solve by digging toward x . First move 0.25 to the right-hand side to get $(x + 2.5)^2 = 0.25$. Then take the square root on both sides to obtain $(x + 2.5) = \pm 0.5$, which simplifies to

$x = -2.5 \pm 0.5$. The two solutions are $x = -2.5 + 0.5 = -2$ and $x = -2.5 - 0.5 = -3$. You can verify these solutions by substituting the values in the original equation $(-2)^2 + 5(-2) + 6 = 0$ and similarly $(-3)^2 + 5(-3) + 6 = 0$. Congratulations, you just solved a quadratic equation using a 1200-year-old algebra technique!

In the next section, we'll learn how to leverage the complete-the-square trick to obtain a general-purpose formula for quickly solving quadratic equations.

Exercises

E1.6 Factor the following quadratic expressions:

a) $x^2 - 8x + 7$

b) $x^2 + 4x + 4$

c) $x^2 - 9$

Hint: Guess the values p and q in the expression $(x + p)(x + q)$.

E1.7 Solve the equations by completing the square.

a) $x^2 + 2x - 15 = 0$

b) $x^2 + 4x + 1 = 0$

1.6 Solving quadratic equations

What would you do if asked to solve for x in the quadratic equation $2x^2 = 4x + 6$? This is called a *quadratic equation* since it contains the unknown variable x squared. The name comes from the Latin *quadratus*, which means square. Quadratic equations appear often, so mathematicians created a general formula for solving them. In this section, we'll learn about this formula and use it to put some quadratic equations in their place.

Before we can apply the formula, we need to rewrite the equation we are trying to solve in the following form:

$$ax^2 + bx + c = 0.$$

This is called the *standard form* of the quadratic equation. We obtain this form by moving all the numbers and x s to one side and leaving only 0 on the other side. For example, to transform the quadratic equation $2x^2 = 4x + 6$ into standard form, we subtract $4x + 6$ from both sides of the equation to obtain $2x^2 - 4x - 6 = 0$. What are the values of x that satisfy this equation?

Quadratic formula

The solutions to the equation $ax^2 + bx + c = 0$ for $a \neq 0$ are

$$x_1 = \frac{-b + \sqrt{b^2 - 4ac}}{2a} \quad \text{and} \quad x_2 = \frac{-b - \sqrt{b^2 - 4ac}}{2a}.$$

The quadratic formula is usually abbreviated $x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$, where the sign “ $\pm$ ” stands for both “ $+$ ” and “ $-$.” The notation “ $\pm$ ” allows us to express both solutions x_1 and x_2 in one equation, but you should keep in mind there are really two solutions.

Let’s see how the quadratic formula is used to solve the equation $2x^2 - 4x - 6 = 0$. Finding the two solutions requires the simple mechanical task of identifying $a = 2$, $b = -4$, and $c = -6$, then plugging these values into the two parts of the formula:

$$\begin{aligned} x_1 &= \frac{4 + \sqrt{4^2 - 4(2)(-6)}}{4} = \frac{4 + \sqrt{16 + 48}}{4} = \frac{4 + \sqrt{64}}{4} = 3, \\ x_2 &= \frac{4 - \sqrt{4^2 - 4(2)(-6)}}{4} = \frac{4 - \sqrt{16 + 48}}{4} = \frac{4 - \sqrt{64}}{4} = -1. \end{aligned}$$

We can easily verify that value $x_1 = 3$ and $x_2 = -1$ both satisfy the original equation $2x^2 = 4x + 6$.

Proof of the quadratic formula

Every claim made by a mathematician comes with a *proof*, which is a step-by-step argument that shows why the claim is true. It’s easy to see where a proof starts and where a proof ends in mathematical texts. Each proof begins with the heading *Proof* (usually in italics) and has the symbol “ $\square$ ” at its end. The purpose of these demarcations is to give readers the option to skip the proof. It’s not necessary to read and understand the proofs of all math statements, but reading proofs can often lead you to a more solid understanding of the material.

I want you to see the proof of the quadratic formula because it’s an important result that you’ll use very often to solve math problems. Reading the proof will help you understand where the quadratic formula comes from. The proof relies on the completing-the-square technique from the previous section, and some basic algebra operations. You can totally handle this!

Proof. We're starting from the quadratic equation $ax^2 + bx + c = 0$, and we're making the additional assumption that $a \neq 0$. We want to find the value or values of x that satisfy this equation.

The first thing we want to do is divide by a to obtain the equivalent equation

$$x^2 + \frac{b}{a}x + \frac{c}{a} = 0.$$

We are allowed to divide by a since we assumed that $a \neq 0$.

Next we apply the *complete the square* trick to the quadratic expression, to obtain an equivalent expression of the form $(x+?)^2 + ?$. Recall that the trick for completing the square is to choose the number inside the bracket to be half the coefficient of the linear term of the quadratic expression, which is $\frac{b}{2a}$ in this case. We must also subtract the square of this term outside the bracket in order to maintain the equality. After completing the square, we're left with the following equation:

$$\left(x + \frac{b}{2a}\right)^2 + \frac{c}{a} - \frac{b^2}{4a^2} = 0.$$

From here, we use the standard digging-toward-the- x procedure. Move all constants to the right-hand side,

$$\left(x + \frac{b}{2a}\right)^2 = \frac{b^2}{4a^2} - \frac{c}{a},$$

and take the square root of both sides to undo the square function:

$$x + \frac{b}{2a} = \pm \sqrt{\frac{b^2}{4a^2} - \frac{c}{a}}.$$

Since any number and its opposite have the same square, taking the square root gives us two possible solutions, which we denote using the " $\pm$ " symbol.

Next we subtract $\frac{b}{2a}$ from both sides of the equation to isolate x and obtain $x = -\frac{b}{2a} \pm \sqrt{\frac{b^2}{4a^2} - \frac{c}{a}}$. We tidy up the mess under the square root, $\sqrt{\frac{b^2}{4a^2} - \frac{c}{a}} = \sqrt{\frac{b^2}{4a^2} - \frac{4a \cdot c}{4a \cdot a}} = \sqrt{\frac{b^2 - 4ac}{4a^2}} = \frac{\sqrt{b^2 - 4ac}}{2a}$, and add the fractions on the right-hand side to obtain $x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$. The solutions to the quadratic equation $ax^2 + bx + c = 0$ are

$$x_1 = \frac{-b + \sqrt{b^2 - 4ac}}{2a} \quad \text{and} \quad x_2 = \frac{-b - \sqrt{b^2 - 4ac}}{2a}.$$

This completes the proof of the quadratic formula. □

The expression $b^2 - 4ac$ is called the *discriminant* of the equation. The discriminant tells us important information about the solutions of the equation $ax^2 + bx + c = 0$. The solutions x_1 and x_2 correspond to real numbers if the discriminant is positive or zero: $b^2 - 4ac \geq 0$. When the discriminant is zero ($b^2 - 4ac = 0$), the equation has only one solution since $x_1 = x_2 = \frac{-b}{2a}$. If the discriminant is negative, $b^2 - 4ac < 0$, the quadratic formula requires computing the square root of a negative number, which is not allowed for real numbers.

Alternative proof

To prove the quadratic formula, we don't necessarily need to show the algebra steps we followed to obtain the formula as outlined above. The quadratic formula states that x_1 and x_2 are solutions. To prove the formula is correct we can simply plug x_1 and x_2 into the equation $ax^2 + bx + c = 0$ to verify that x_1 and x_2 are solutions. Verify this on your own.

Applications

The golden ratio

The *golden ratio* is an essential proportion in geometry, art, aesthetics, biology, and mysticism, and is usually denoted as $\varphi = \frac{1+\sqrt{5}}{2} = 1.6180339\dots$. This ratio is determined as the positive solution to the quadratic equation

$$x^2 - x - 1 = 0.$$

Applying the quadratic formula to this equation yields two solutions,

$$x_1 = \frac{1 + \sqrt{5}}{2} = \varphi \quad \text{and} \quad x_2 = \frac{1 - \sqrt{5}}{2} = -\frac{1}{\varphi}.$$

You can learn more about the various contexts in which the golden ratio appears from the Wikipedia article on the subject.

Explanations

Multiple solutions

Often, we are interested in only one of the two solutions to the quadratic equation. It will usually be obvious from the context of the problem which of the two solutions should be kept and which should be discarded. For example, the *time of flight* of a ball

thrown in the air from a height of 3 metres with an initial velocity of 12 metres per second is obtained by solving the equation $(-4.9)t^2 + 12t + 3 = 0$. The two solutions of the quadratic equation are $t_1 = -0.229$ and $t_2 = 2.678$. The first answer t_1 corresponds to a time in the past so we reject it as invalid. The correct answer is t_2 . The ball will hit the ground after $t = 2.678$ seconds.

Relation to factoring

In the previous section we discussed the *quadratic factoring* operation by which we could rewrite a quadratic function as the product of a constant and two factors:

$$f(x) = ax^2 + bx + c = a(x - x_1)(x - x_2).$$

The two numbers x_1 and x_2 are called the *roots* of the function: these points are where the function $f(x)$ touches the x -axis.

You now have the ability to factor any quadratic equation: use the quadratic formula to find the two solutions, x_1 and x_2 , then rewrite the expression as $a(x - x_1)(x - x_2)$.

Some quadratic expressions cannot be factored, however. These “unfactorable” expressions correspond to quadratic functions whose graphs do not touch the x -axis. They have no real solutions (no roots). There is a quick test you can use to check if a quadratic function $f(x) = ax^2 + bx + c$ has roots (touches or crosses the x -axis) or doesn’t have roots (never touches the x -axis). If $b^2 - 4ac > 0$ then the function f has two roots. If $b^2 - 4ac = 0$, the function has only one root, indicating the special case when the function touches the x -axis at only one point. If $b^2 - 4ac < 0$, the function has no roots. In this case, the quadratic formula fails because it requires taking the square root of a negative number, which is not allowed (for now). We’ll come back to the idea of taking square roots of negative numbers in Section 1.14 (see page 96).

Links

[Algebra explanation of the quadratic formula]

<https://www.youtube.com/watch?v=r3SEkdtpobo>

[Visual explanation of the quadratic formula derivation]

<https://www.youtube.com/watch?v=EBbtoFMJvFc>

Exercises

E1.8 Solve for x in the quadratic equation $2x^2 - x = 3$.

E1.9 Solve for x in the equation $x^4 - 4x^2 + 4 = 0$.

Hint: Use the substitution $y = x^2$.

1.7 The Cartesian plane

The Cartesian plane, named after famous philosopher and mathematician René Descartes, is used to visualize pairs of numbers (x, y) .

Consider first the *number line* representation for numbers.




Figure 1.9: Every real number x corresponds to a point on the number line. The number line extends indefinitely to the left (toward negative infinity) and to the right (toward positive infinity).

The Cartesian plane is the two-dimensional generalization of the number line. Generally, we call the plane's horizontal axis "the x -axis" and its vertical axis "the y -axis." We put notches at regular intervals on each axis so we can measure distances.




Figure 1.10: Every point in the Cartesian plane corresponds to a pair of real numbers (x, y) . Points $P = (P_x, P_y)$, vectors $\vec{v} = (v_x, v_y)$, and graphs of functions $(x, f(x))$ live here.

Figure 1.10 is an example of an empty Cartesian coordinate system. Think of the coordinate system as an empty canvas. What can you draw on this canvas?

Vectors and points

A point $P = (P_x, P_y)$ in the Cartesian plane has an x -coordinate and a y -coordinate. To find this point, start from the origin—the point $(0,0)$ —and move a distance P_x on the x -axis, then move a distance P_y on the y -axis.




Figure 1.11: A Cartesian plane which shows the point $P = (-3, 2)$ and the vectors $\vec{v}_1 = (3, 1)$ and $\vec{v}_2 = \vec{v}_3 = (-1, -2)$.

Similar to a point, a vector $\vec{v} = (v_x, v_y)$ is a pair of coordinates. Unlike points, we don't necessarily start from the plane's origin when mapping vectors. We draw vectors as arrows that explicitly mark where the vector starts and where it ends. Note that vectors $\vec{v}_2$ and $\vec{v}_3$ illustrated in Figure 1.11 are actually the *same* vector—the “displace left by 1 and down by 2” vector. It doesn't matter where you draw this vector, it will always be the same whether it begins at the plane's origin or elsewhere.

Graphs of functions

The Cartesian plane is great for visualizing functions. You can think of a function as a set of input-output pairs $(x, f(x))$. You can draw the *graph* of a function by letting the y -coordinate represent the function's output value:

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

For example, with the function $f(x) = x^2$, we can pass a line through the set of points

$$(x, y) = (x, x^2),$$

and obtain the graph shown in Figure 1.12.




Figure 1.12: The graph of the function $f(x) = x^2$ consists of all pairs of points (x, y) in the Cartesian plane that satisfy $y = x^2$.

When plotting functions by setting $y = f(x)$, we use a special terminology for the two axes. The x -axis represents the *independent* variable (the one that varies freely), and the y -axis represents the *dependent* variable $f(x)$, since $f(x)$ depends on x .

To draw the graph of any function $f(x)$, use the following procedure. Imagine making a sweep over all of the possible input values for the function. For each input x , put a point at the coordinates $(x, y) = (x, f(x))$ in the Cartesian plane. Using the graph of a function, you can literally *see* what the function does: the “height” y of the graph at a given x -coordinate tells you the value of the function $f(x)$.

Dimensions

The number line is one-dimensional. Every number x can be visualized as a point on the number line. The Cartesian plane has two dimensions: the x dimension and the y dimension. If we need to visualize math concepts in 3D, we can use a three-dimensional coordinate system with x , y , and z axes (see Figure 1.55 on page 92).

1.8 Functions

We need to have a relationship talk. We need to talk about functions. We use functions to describe the relationships between variables. In particular, functions describe how one variable *depends* on another.

For example, the revenue R from a music concert depends on the number of tickets sold n . If each ticket costs \$25, the revenue from the concert can be written *as a function of n* as follows: $R(n) = 25n$. Solving for n in the equation $R(n) = 7000$ tells us the number of ticket sales needed to generate \$7000 in revenue. This is a simple model of a function; as your knowledge of functions builds, you'll learn how to build more detailed models of reality. For instance, if you need to include a 5% processing charge for issuing the tickets, you can update the revenue model to $R(n) = 0.95 \cdot 25 \cdot n$. If the estimated cost of hosting the concert is $C = \$2000$, then the profit from the concert P can be modelled as

$$\begin{aligned} P(n) &= R(n) - C \\ &= 0.95 \cdot \$25 \cdot n - \$2000 \end{aligned}$$

The function $P(n) = 23.75n - 2000$ models the profit from the concert as a function of the number of tickets sold. This is a pretty good model already, and you can always update it later as you learn more information.

The more functions you know, the more tools you have for modelling reality. To “know” a function, you must be able to understand and connect several of its aspects. First you need to know the function's mathematical **definition**, which describes exactly what the function does. Starting from the function's definition, you can use your existing math skills to find the function's **properties**. You must also know the **graph** of the function; what the function looks like if you plot x versus $f(x)$ in the Cartesian plane. It's also a good idea to remember the **values** of the function for some important inputs. Finally—and this is the part that takes time—you must learn about the function's **relations** to other functions.

Definitions

A *function* is a mathematical object that takes numbers as inputs and produces numbers as outputs. We use the notation

$$f: A \rightarrow B$$

to denote a function from the input set A to the output set B . In this book, we mostly study functions that take real numbers as inputs and give real numbers as outputs: $f: \mathbb{R} \rightarrow \mathbb{R}$.




Figure 1.13: An abstract representation of a function f from the set A to the set B . The function f is the arrow which *maps* each input x in A to an output $f(x)$ in B . The output of the function $f(x)$ is also denoted y .

A function is not a number; rather, it is a *mapping* from numbers to numbers. We say “ f maps x to $f(x)$.” For any input x , the output value of f for that input is denoted $f(x)$, which is read as “ f of x .”

We’ll now define some fancy technical terms used to describe the input and output sets of functions.

- A : the *source set* of the function describes the types of numbers that the function takes as inputs.
- $\text{Dom}(f)$: the *domain* of a function is the set of allowed input values for the function.
- B : the *target set* of a function describes the type of outputs the function has. The target set is sometimes called the *codomain*.
- $\text{Im}(f)$: the *image* of the function is the set of all possible output values of the function. The image is sometimes called the *range*.

See Figure 1.14 for an illustration of these concepts. The purpose of introducing all this math terminology is so we’ll have words to distinguish the general types of inputs and outputs of the function (real numbers, complex numbers, vectors) from the specific properties of the function like its domain and image.

Let’s look at an example to illustrate the difference between the source set and the domain of a function. Consider the square root function $f: \mathbb{R} \rightarrow \mathbb{R}$ defined as $f(x) = \sqrt{x}$, which is shown in Figure 1.15. The source set of f is the set of real numbers—yet only nonnegative real numbers are allowed as inputs, since $\sqrt{x}$ is not defined for negative numbers. Therefore, the domain of the square root function is only the nonnegative real numbers: $\text{Dom}(f) = \mathbb{R}_+ = \{x \in \mathbb{R} \mid x \geq 0\}$. Knowing the domain of a function is essential to using the function correctly. In this case, whenever you use the square root function, you need to make sure that the inputs to the function are nonnegative numbers.




Figure 1.14: Illustration of the input and output sets of a function $f: A \rightarrow B$. The *source set* is denoted A and the *domain* is denoted $\text{Dom}(f)$. Note that the function's domain is a subset of its source set. The *target set* is denoted B and the *image* is denoted $\text{Im}(f)$. The image is a subset of the target set.

The complicated-looking expression between the curly brackets uses *set notation* to define the set of nonnegative numbers $\mathbb{R}_+$. In words, the expression $\mathbb{R}_+ = \{x \in \mathbb{R} \mid x \geq 0\}$ states that “ $\mathbb{R}_+$ is defined as the set of all real numbers x such that x is greater than or equal to zero.” We'll discuss set notation in more detail in Section 1.16. For now, you can just remember that $\mathbb{R}_+$ represents the set of nonnegative real numbers.




Figure 1.15: The input and output sets of the function $f(x) = \sqrt{x}$. The domain of f is the set of nonnegative real numbers $\mathbb{R}_+$ and its image is $\mathbb{R}_+$.

To illustrate the difference between the image of a function and its target set, let's look at the function $f(x) = x^2$ shown in Figure 1.16. The quadratic function is of the form $f: \mathbb{R} \rightarrow \mathbb{R}$. The function's source set is $\mathbb{R}$ (it takes real numbers as inputs) and its target set is $\mathbb{R}$ (the outputs are real numbers too); however, not all real numbers are possible outputs. The *image* of the function $f(x) = x^2$ consists only of the nonnegative real numbers $\mathbb{R}_+ = \{y \in \mathbb{R} \mid y \geq 0\}$, since $f(x) \geq 0$ for all x .




Figure 1.16: The function $f(x) = x^2$ is defined for all reals: $\text{Dom}(f) = \mathbb{R}$. The image of the function is the set of nonnegative real numbers: $\text{Im}(f) = \mathbb{R}_+$.

Function properties

We'll now introduce some additional terminology for describing three important function properties. Every function is a mapping from a source set to a target set, but what kind of mapping is it?

- A function is *injective* if it maps two different inputs to two different outputs. If x_1 and x_2 are two input values that are not equal $x_1 \neq x_2$, then the output values of an injective function will also not be equal $f(x_1) \neq f(x_2)$.
- A function is *surjective* if its image is equal to its target set. For every output y in the target set of a surjective function, there is at least one input x in its domain such that $f(x) = y$.
- A function is *bijective* if it is both injective and surjective.

I know this seems like a lot of terminology to get acquainted with, but it's important to have names for these function properties. We'll need this terminology to give a precise definition of the *inverse function* in the next section.

Injective property We can think of *injective* functions as pipes that transport fluids between containers. Since fluids cannot be compressed, the “output container” must be at least as large as the “input container.” If there are two distinct points x_1 and x_2 in the input container of an injective function, then there will be two distinct points $f(x_1)$ and $f(x_2)$ in the output container of the function as well. In other words, injective functions don't smoosh things together.

In contrast, a function that doesn't have the injective property can map several different inputs to the same output value. The function $f(x) = x^2$ is not injective since it sends inputs x and $-x$ to the same output value $f(x) = f(-x) = x^2$, as illustrated in Figure 1.16.

The maps-distinct-inputs-to-distinct-outputs property of injective functions has an important consequence: given the output of

an injective function y , there is only one input x such that $f(x) = y$. If a second input x' existed that also leads to the same output $f(x) = f(x') = y$, then the function f wouldn't be injective. For each of the outputs y of an injective function f , there is a *unique* input x such that $f(x) = y$. In other words, injective functions have a unique-input-for-each-output property.

Surjective property A function is *surjective* if its outputs cover the entire target set: every number in the target set is a possible output of the function for some input. For example, the function $f: \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x) = x^3$ is surjective: for every number y in the target set $\mathbb{R}$, there is an input x , namely $x = \sqrt[3]{y}$, such that $f(x) = y$. The function $f(x) = x^3$ is surjective since its image is equal to its target set, $\text{Im}(f) = \mathbb{R}$, as shown in Figure 1.17.

On the other hand, the function $f: \mathbb{R} \rightarrow \mathbb{R}$ defined by the equation $f(x) = x^2$ is not surjective since its image is only the nonnegative numbers $\mathbb{R}_+$ and not the whole set of real numbers (see Figure 1.16). The outputs of this function do not include the negative numbers of the target set, because there is no real number x that can be used as an input to obtain a negative output value.




Figure 1.17: For the function $f(x) = x^3$ the image is equal to the target set of the function, $\text{Im}(f) = \mathbb{R}$, therefore the function f is surjective. The function f maps two different inputs $x_1 \neq x_2$ to two different outputs $f(x_1) \neq f(x_2)$, so f is injective. Since f is both injective and surjective, it is a *bijective* function.

Bijective property A function is *bijective* if it is both injective and surjective. When a function $f: A \rightarrow B$ has both the injective and surjective properties, it defines a *one-to-one correspondence* between the numbers of the source set A and the numbers of the target set B . This means for every input value x , there is exactly one corresponding output value y , and for every output value y , there is exactly one input value x such that $f(x) = y$. An example of a bijective function is the function $f: \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x) = x^3$ (see Figure 1.17). For every input x in the source set $\mathbb{R}$, the corresponding output y is given

by $y = f(x) = x^3$. For every output value y in the target set $\mathbb{R}$, the corresponding input value x is given by $x = \sqrt[3]{y}$.

A function is not bijective if it lacks one of the required properties. Examples of non-bijective functions are $f(x) = \sqrt{x}$, which is not surjective and $f(x) = x^2$, which is neither injective nor surjective.

Counting solutions Another way to understand the injective, surjective, and bijective properties of functions is to think about the solutions to the equation $f(x) = b$, where b is a number in the target set B . The function f is injective if the equation $f(x) = b$ has *at most one* solution for every number b . The function f is surjective if the equation $f(x) = b$ has *at least one* solution for every number b . If the function f is bijective then it is both injective and surjective, which means the equation $f(x) = b$ has *exactly one* solution.

Inverse function

We used inverse functions repeatedly in previous chapters, each time describing the inverse function informally as an “undo” operation. Now that we have learned about bijective functions, we can give a the precise definition of the inverse function and explain some of the details we glossed over previously.

Recall that a *bijective* function $f : A \rightarrow B$ is a *one-to-one correspondence* between the numbers in the source set A and numbers in the target set B : for every output y , there is exactly one corresponding input value x such that $f(x) = y$. The *inverse function*, denoted f^{-1} , is the function that takes any output value y in the set B and finds the corresponding input value x that produced it $f^{-1}(y) = x$.




Figure 1.18: The inverse f^{-1} undoes the operation of the function f .

For every bijective function $f : A \rightarrow B$, there exists an inverse function $f^{-1} : B \rightarrow A$ that performs the *inverse mapping* of f . If we start from some x , apply f , and then apply f^{-1} , we’ll arrive—full circle—back to the original input x :

$$f^{-1}(f(x)) = x.$$

In Figure 1.18 the function f is represented as a forward arrow, and the inverse function f^{-1} is represented as a backward arrow that puts the value $f(x)$ back to the x it came from.

Similarly, we can start from any y in the set B and apply f^{-1} followed by f to get back to the original y we started from:

$$f(f^{-1}(y)) = y.$$

In words, this equation tells us that f is the “undo” operation for the function f^{-1} , the same way f^{-1} is the “undo” operation for f .

If a function is missing the injective property or the surjective property then it isn’t bijective and it doesn’t have an inverse. Without the injective property, there could be two inputs x and x' that both produce the same output $f(x) = f(x') = y$. In this case, computing $f^{-1}(y)$ would be impossible since we don’t know which of the two possible inputs x or x' was used to produce the output y . Without the surjective property, there could be some output y' in B for which the inverse function f^{-1} is not defined, so the equation $f(f^{-1}(y)) = y$ would not hold for all y in B . The inverse function f^{-1} exists only when the function f is bijective.

Wait a minute! We know the function $f(x) = x^2$ is not bijective and therefore doesn’t have an inverse, but we’ve repeatedly used the square root function as an inverse function for $f(x) = x^2$. What’s going on here? Are we using a double standard like a politician that espouses one set of rules publicly, but follows a different set of rules in their private dealings? Is mathematics corrupt?

Don’t worry, mathematics is not corrupt—it’s all legit. We can use inverses for non-bijective functions by imposing *restrictions* on the source and target sets. The function $f(x) = x^2$ is not bijective when defined as a function $f : \mathbb{R} \rightarrow \mathbb{R}$, but it is bijective if we define it as a function from the set of nonnegative numbers to the set of nonnegative numbers, $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. Restricting the source set to $\mathbb{R}_+ = \{x \in \mathbb{R} \mid x \geq 0\}$ makes the function injective, and restricting the target set to $\mathbb{R}_+$ also makes the function surjective. The function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ defined by the equation $f(x) = x^2$ is bijective and its inverse is $f^{-1}(y) = \sqrt{y}$.

It’s important to keep track the of restrictions on the source set we applied when solving equations. For example, solving the equation $x^2 = c$ by restricting the solution space to nonnegative numbers will give us only the positive solution $x = \sqrt{c}$. We have to manually add the negative solution $x = -\sqrt{c}$ in order to obtain the complete solutions: $x = \sqrt{c}$ or $x = -\sqrt{c}$, which is usually written $x = \pm\sqrt{c}$. The possibility of multiple solutions is present whenever we solve equations involving non-injective functions.

Function composition

We can combine two simple functions by chaining them together to build a more complicated function. This act of applying one function after another is called *function composition*. Consider for example the composition:

$$f \circ g(x) = f(g(x)) = z.$$




Figure 1.19: The function composition $f \circ g$ describes the combination of first applying the function g , followed by the function f : $f \circ g(x) = f(g(x))$.

Figure 1.19 illustrates the concept of function composition. First, the function $g : A \rightarrow B$ acts on some input x to produce an intermediary value $y = g(x)$ in the set B . The intermediary value y is then passed through the function $f : B \rightarrow C$ to produce the final output value $z = f(y) = f(g(x))$ in the set C . We can think of the *composite function* $f \circ g$ as a function in its own right. The function $f \circ g : A \rightarrow C$ is defined through the formula $f \circ g(x) = f(g(x))$.

Don't worry too much about the “ $\circ$ ” symbol—it's just a convenient math notation I wanted you to know about. Writing $f \circ g$ is the same as writing $f(g(x))$. The important takeaway from Figure 1.19 is that functions can be combined by using the outputs of one function as the inputs to the next. This is a very useful idea for building math models. You can understand many complicated input-output transformations by describing them as compositions of simple functions.

Example 1 Consider the function $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ given by $g(x) = \sqrt{x}$, and the function $f : \mathbb{R} \rightarrow \mathbb{R}_+$ defined by $f(x) = x^2$. The composite function $f \circ g(x) = (\sqrt{x})^2 = x$ is defined for all nonnegative reals. The composite function $g \circ f$ is defined for all real numbers, and we have $g \circ f(x) = \sqrt{x^2} = |x|$.

Example 2 The composite functions $f \circ g$ and $g \circ f$ describe different operations. If $g(x) = \ln(x)$ and $f(x) = x^2$, the functions

$g \circ f(x) = \ln(x^2)$ and $f \circ g(x) = (\ln x)^2$ have different domains and produce different outputs, as you can verify using a calculator.

Using the notation “ $\circ$ ” for function composition, we can give a concise description of the properties of a bijective function $f : A \rightarrow B$ and its inverse function $f^{-1} : B \rightarrow A$:

$$(f^{-1} \circ f)(x) = x \quad \text{and} \quad (f \circ f^{-1})(y) = y,$$

for all x in A and all y in B .

Function names

We use short symbols like $+$, $-$, $\times$, and $\div$ to denote most of the important functions used in everyday life. We also use the squiggle notation $\sqrt{\quad}$ for square roots and superscripts to denote exponents. All other functions are identified and denoted by their *name*. If I want to compute the *cosine* of the angle 60° (a function describing the ratio between the length of one side of a right-angle triangle and the hypotenuse), I write $\cos(60^\circ)$, which means I want the value of the $\cos$ function for the input 60° .

Incidentally, the function $\cos$ has a nice output value for that specific angle: $\cos(60^\circ) = \frac{1}{2}$. Therefore, seeing $\cos(60^\circ)$ somewhere in an equation is the same as seeing $\frac{1}{2}$. To find other values of the function, say $\cos(33.13^\circ)$, you’ll need a calculator. All scientific calculators have a convenient little cos button for this very purpose.

Handles on functions

When you learn about functions you learn about the different “handles” by which you can “grab” these mathematical objects. The main handle for a function is its **definition**: it tells you the precise way to calculate the output when you know the input. The function definition is an important handle, but it is also important to “feel” what the function does intuitively. How does one get a feel for a function?

Table of values

One simple way to represent a function is to look at a list of input-output pairs: $\{\{\text{in} = x_1, \text{out} = f(x_1)\}, \{\text{in} = x_2, \text{out} = f(x_2)\}, \{\text{in} = x_3, \text{out} = f(x_3)\}, \dots\}$. A more compact notation for the input-output pairs is $\{(x_1, f(x_1)), (x_2, f(x_2)), (x_3, f(x_3)), \dots\}$, where the first number of each pair represents an input value and the second represents the output value given by the function.

We can also build a **table of values** by writing the input values in one column and recording the corresponding output values in a second column. You can choose inputs at random or focus on the important-looking x values in the function's domain.

<u>input = x</u>	$\rightarrow$	<u>$f(x)$ = output</u>
0	$\rightarrow$	$f(0)$
1	$\rightarrow$	$f(1)$
55	$\rightarrow$	$f(55)$
x_4	$\rightarrow$	$f(x_4)$

Table 1.1: Table of input-output values of the function $f(x)$. The input values $x = 0$, $x = 1$ and $x = 55$ are chosen to “test” what the function does.

You can create a table of values for any function you want to study. Follow the example shown in Table 1.1. Use the input values that interest you and fill out the right side of the table by calculating the value of $f(x)$ for each input x .

Function graph

One of the best ways to feel a function is to look at its graph. A graph is a line on a piece of paper that passes through all input-output pairs of a function. Imagine you have a piece of paper, and on it you draw a blank *coordinate system* as in Figure 1.20.

The horizontal axis is used to measure x . The vertical axis is used to measure $f(x)$. Because writing out $f(x)$ every time is long and tedious, we use a short, single-letter alias to denote the output value of f as follows:

$$y = f(x) = \text{output.}$$

Think of each input-output pair of the function f as a point (x, y) in the coordinate system. The graph of a function is a representational drawing of everything the function does. If you understand how to interpret this drawing, you can infer everything there is to know about the function.

Facts and properties

Another way to feel a function is by knowing the function's properties. This approach boils down to learning facts about the function




Figure 1.20: An empty (x, y) -coordinate system that you can use to draw function graphs. The graph of $f(x)$ consists of all the points for which $(x, y) = (x, f(x))$. See Figure 1.12 on page 37 for the graph of $f(x) = x^2$.

and its connections to other functions. An example of a mathematical connection is the equation $\log_B(x) = \frac{\log_b(x)}{\log_b(B)}$, which describes a link between the logarithmic function base B and the logarithmic function base b .

The more you know about a function, the more “paths” your brain builds to connect to that function. Real math knowledge is not about memorization; it is about establishing a network of associations between different areas of information in your brain. See the concept maps on page v for an illustration of the paths that link math concepts. Mathematical thought is the usage of these associations to carry out calculations and produce mathematical arguments. For example, knowing about the connection between logarithmic functions will allow you compute the value of $\log_7(e^3)$, even though calculators don’t have a button for logarithms base 7. We find $\log_7(e^3) = \frac{\ln e^3}{\ln 7} = \frac{3}{\ln 7}$, which can be computed using the ln button.

To develop mathematical skills, it is vital to practice path-building between concepts by solving exercises. With this book, I will introduce you to some of the many paths linking math concepts, but it’s on you to reinforce these paths through practice.

Example 3 Consider the function f from the real numbers to the real numbers ($f: \mathbb{R} \rightarrow \mathbb{R}$) defined as $f(x) = x^2 + 2x - 3$. The value of f when $x = 1$ is $f(1) = 1^2 + 2(1) - 3 = 0$. When $x = 2$, the output is $f(2) = 2^2 + 2(2) - 3 = 5$. What is the value of f when $x = 0$? You can

use algebra to rewrite this function as $f(x) = (x + 3)(x - 1)$, which tells you the graph of this function crosses the x -axis at $x = -3$ and at $x = 1$. The values above will help you plot the graph of $f(x)$.

Example 4 Consider the exponential function with base 2 defined by $f(x) = 2^x$. This function is crucial to computer systems. For instance, RAM memory chips come in powers of two because the memory space is exponential in the number of “address lines” used on the chip. When $x = 1$, $f(1) = 2^1 = 2$. When x is 2 we have $f(2) = 2^2 = 4$. The function is therefore described by the following input-output pairs: $(0, 1)$, $(1, 2)$, $(2, 4)$, $(3, 8)$, $(4, 16)$, $(5, 32)$, $(6, 64)$, $(7, 128)$, $(8, 256)$, $(9, 512)$, $(10, 1024)$, $(11, 2048)$, $(12, 4096)$, etc. Recall that any number raised to exponent 0 gives 1. Thus, the exponential function passes through the point $(0, 1)$. Recall also that negative exponents lead to fractions, so we have the points $(-1, \frac{1}{2})$, $(-2, \frac{1}{4})$, $(-3, \frac{1}{8})$, etc. You can plot these $(x, f(x))$ coordinates in the Cartesian plane to obtain the graph of the function.

Discussion

To describe a function we specify its source and target sets $f: A \rightarrow B$, then give an equation of the form $f(x) = \text{“expression involving } x\text{”}$ that defines the function. Since functions are defined using equations, does this mean that functions and equations are the same thing? Let’s take a closer look.

In general, any equation containing two variables describes a *relation* between these variables. For example, the equation $x - 3 = y - 4$ describes a relation between the variables x and y . We can isolate the variable y in this equation to obtain $y = x + 1$ and thus find the value of y when the value of x is given. We can also isolate x to obtain $x = y - 1$ and use this equation to find x when the value of y is given. In the context of an equation, the relationship between the variables x and y is symmetrical and no special significance is attached to either of the two variables.

We also can describe the same relationship between x and y as a function $f: \mathbb{R} \rightarrow \mathbb{R}$. We choose to identify x as the input variable and y as the output variable of the function f . Having identified y with the output variable, we can interpret the equation $y = x + 1$ as the definition of the function $f(x) = x + 1$.

Note that the equation $x - 3 = y - 4$ and the function $f(x) = x + 1$ describe the same relationship between the variables x and y . For example, if we set the value $x = 5$ we can find the value of y by solving the equation $5 - 3 = y - 4$ to obtain $y = 6$, or by computing the output of the function $f(x)$ for the input $x = 5$, which gives us the

same answer $f(5) = 6$. In both cases we arrive at the same answer, but modelling the relationship between x and y as a function allows us to use the whole functions toolbox, like function composition and function inverses.

* * *










In this section we talked a lot about functions in general but we haven't said much about any function specifically. There are many useful functions out there, and we can't discuss them all here. In the next section, we'll introduce 10 functions of strategic importance for all of science. If you get a grip on these functions, you'll be able to understand all of physics and calculus and handle *any* problem your teacher may throw at you.

1.9 Functions reference

Your *function vocabulary* determines how well you can express yourself mathematically in the same way your English vocabulary determines how well you can express yourself in English. The following pages aim to embiggen your function vocabulary, so you'll know how to handle the situation when a teacher tries to pull some trick on you at the final.

If you're seeing these functions for the first time, don't worry about remembering all the facts and properties on the first reading. We'll use these functions throughout the rest of the book, so you'll have plenty of time to become familiar with them. Remember to return to this section if you ever get stuck on a function.

To build mathematical intuition, it's essential you understand functions' graphs. Memorizing the definitions and properties of functions gets a lot easier with visual accompaniment. Indeed, remembering what the function "looks like" is a great way to train yourself to recognize various types of functions. Figure 1.21 shows the graphs of some of the most important functions we'll use in this book.

(a) $f(x) = 2x - 3$ (b) $f(x) = x^2$ (c) $f(x) = \sqrt{x}$ (d) $f(x) = |x|$ (e) $f(x) = e^x$ (f) $f(x) = \ln(x)$ (g) $f(x) = \sin(x)$ (h) $f(x) = \cos(x)$ (i) $f(x) = \tan(x)$ **Figure 1.21:** We'll see many types of function graphs in the next pages.

Line

The equation of a line describes an input-output relationship where the change in the output is *proportional* to the change in the input. The equation of a line is

$$f(x) = mx + b.$$

The constant m describes the slope of the line. The constant b is called the y -intercept and it is the value of the function when $x = 0$.

Consider what relationship the equation of $f(x)$ describes for different values of m and b . What happens when m is positive? What happens when m is negative?

Graph




Figure 1.22: The graph of the function $f(x) = 2x - 3$. The slope is $m = 2$. The y -intercept of this line is $b = -3$. The x -intercept is at $x = \frac{3}{2}$.

Properties

- Domain: $\mathbb{R}$. The function $f(x) = mx + b$ is defined for all reals.
- Image: $\mathbb{R}$ if $m \neq 0$. If $m = 0$ the function is constant $f(x) = b$, so the image set contains only a single number $\{b\}$.
- $x = -b/m$: the x -intercept of $f(x) = mx + b$. The x -intercept is obtained by solving $f(x) = 0$.
- The inverse to the line $f(x) = mx + b$ is $f^{-1}(x) = \frac{1}{m}(x - b)$, which is also a line.

General equation

A line can also be described in a more symmetric form as a relation:

$$Ax + By = C.$$

This is known as the *general* equation of a line. The general equation for the line shown in Figure 1.22 is $2x - 1y = 3$.

Given the general equation of a line $Ax + By = C$ with $B \neq 0$, you can convert to the function form $y = f(x) = mx + b$ by computing the slope $m = \frac{-A}{B}$ and the y -intercept $b = \frac{C}{B}$.

Square

The function x squared, is also called the *quadratic* function, or *parabola*. The formula for the quadratic function is

$$f(x) = x^2.$$

The name “quadratic” comes from the Latin *quadratus* for square, since the expression for the area of a square with side length x is x^2 .




Figure 1.23: Plot of the quadratic function $f(x) = x^2$. The graph of the function passes through the following (x,y) coordinates: $(-2,4)$, $(-1,1)$, $(0,0)$, $(1,1)$, $(2,4)$, $(3,9)$, etc.

Properties

- Domain: $\mathbb{R}$. The function $f(x) = x^2$ is defined for all numbers.
- Image: $\mathbb{R}_+ = \{y \in \mathbb{R} \mid y \geq 0\}$. The outputs are nonnegative numbers since $x^2 \geq 0$, for all real numbers x .
- The function x^2 is the inverse of the square root function $\sqrt{x}$.
- $f(x) = x^2$ is *two-to-one*: it sends both x and $-x$ to the same output value $x^2 = (-x)^2$.
- The quadratic function is *convex*, meaning it curves upward.

The set expression $\{y \in \mathbb{R} \mid y \geq 0\}$ that we use to define the non-negative real numbers ($\mathbb{R}_+$) is read “the set of real numbers that are greater than or equal to zero.”

Square root

The square root function is denoted

$$f(x) = \sqrt{x} = x^{\frac{1}{2}}.$$

The square root $\sqrt{x}$ is the inverse function of the square function x^2 when the two functions are defined as $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. The symbol $\sqrt{c}$ refers to the *positive* solution of $x^2 = c$. Note that $-\sqrt{c}$ is also a solution of $x^2 = c$.

Graph




Figure 1.24: The graph of the function $f(x) = \sqrt{x}$. The domain of the function is $\mathbb{R}_+$ because we can't take the square root of a negative number.

Properties

- **Domain:** $\mathbb{R}_+ = \{x \in \mathbb{R} \mid x \geq 0\}$. The function $f(x) = \sqrt{x}$ is only defined for nonnegative inputs. There is no real number y such that y^2 is negative, hence the function $f(x) = \sqrt{x}$ is not defined for negative inputs x .
- **Image:** $\mathbb{R}_+ = \{y \in \mathbb{R} \mid y \geq 0\}$. The outputs of the function $f(x) = \sqrt{x}$ are nonnegative numbers since $\sqrt{x} \geq 0$.

In addition to *square* root, there is also *cube* root $f(x) = \sqrt[3]{x} = x^{\frac{1}{3}}$, which is the inverse function for the cubic function $f(x) = x^3$. We have $\sqrt[3]{8} = 2$ since $2 \times 2 \times 2 = 8$. More generally, we can define the n^{th} -root function $\sqrt[n]{x}$ as the inverse function of x^n .

Absolute value

The *absolute value* function tells us the size of numbers without paying attention to whether the number is positive or negative. We can compute a number's absolute value by *ignoring the sign* of the number. A number's absolute value corresponds to its distance from the origin of the number line.

Another way of thinking about the absolute value function is to say it multiplies negative numbers by -1 to “cancel” their negative sign:

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

Graph




Figure 1.25: The graph of the absolute value function $f(x) = |x|$.

Properties

- Domain: $\mathbb{R}$. The function $f(x) = |x|$ is defined for all inputs.
- Image: $\mathbb{R}_+ = \{y \in \mathbb{R} \mid y \geq 0\}$
- The combination of squaring followed by square-root is equivalent to the absolute value function:

$$\sqrt{x^2} = |x|,$$

since squaring destroys the sign.

Polynomials

The polynomials are a very useful family of functions. For example, quadratic polynomials of the form $f(x) = ax^2 + bx + c$ often arise when describing physics phenomena.

The general equation for a polynomial function of degree n is

$$f(x) = a_0 + a_1x + a_2x^2 + a_3x^3 + \cdots + a_nx^n.$$

The constants a_i are known as the *coefficients* of the polynomial.

Parameters

- x : the variable
- a_0 : the constant term
- a_1 : the *linear* coefficient, or *first-order* coefficient
- a_2 : the *quadratic* coefficient
- a_3 : the *cubic* coefficient
- a_n : the n^{th} order coefficient
- n : the *degree* of the polynomial. The degree of $f(x)$ is the largest power of x that appears in the polynomial.

A polynomial of degree n has $n + 1$ coefficients: $a_0, a_1, a_2, \dots, a_n$.

Properties

- Domain: $\mathbb{R}$. Polynomials are defined for all inputs.
- The roots of $f(x)$ are the values of x for which $f(x) = 0$.
- The image of a polynomial function depends on the coefficients.
- The sum of two polynomials is also a polynomial.

The most general first-degree polynomial is a line $f(x) = mx + b$, where m and b are arbitrary constants. The most general second-degree polynomial is $f(x) = a_2x^2 + a_1x + a_0$, where again a_0, a_1 , and a_2 are arbitrary constants. We call a_k the *coefficient* of x^k , since this is the number that appears in front of x^k . Following the pattern, a third-degree polynomial will look like $f(x) = a_3x^3 + a_2x^2 + a_1x + a_0$.

In general, a polynomial of degree n has the equation

$$f(x) = a_nx^n + a_{n-1}x^{n-1} + \cdots + a_2x^2 + a_1x + a_0.$$

You can add two polynomials by adding together their coefficients:

$$\begin{aligned} f(x) + g(x) &= (a_nx^n + \cdots + a_1x + a_0) + (b_nx^n + \cdots + b_1x + b_0) \\ &= (a_n + b_n)x^n + \cdots + (a_1 + b_1)x + (a_0 + b_0). \end{aligned}$$

The subtraction of two polynomials works similarly. We can also multiply polynomials together using the general algebra rules for expanding brackets.

Solving polynomial equations

Very often in math, you will have to *solve* polynomial equations of the form

$$A(x) = B(x),$$

where $A(x)$ and $B(x)$ are both polynomials. Recall from earlier that to *solve*, we must find the values of x that make the equality true.

Say the revenue of your company is a function of the number of products sold x , and can be expressed as $R(x) = 2x^2 + 2x$. Say also the cost you incur to produce x objects is $C(x) = x^2 + 5x + 10$. You want to determine the amount of product you need to produce to break even, that is, so that revenue equals cost: $R(x) = C(x)$. To find the break-even value x , solve the equation

$$2x^2 + 2x = x^2 + 5x + 10.$$

This may seem complicated since there are x s all over the place. No worries! We can turn the equation into its “standard form,” and then use the quadratic formula. First, move all the terms to one side until only zero remains on the other side:

$$\begin{aligned} 2x^2 + 2x - x^2 &= \cancel{x^2} + 5x + 10 - \cancel{x^2} \\ x^2 + 2x - 5x &= \cancel{5x} + 10 - \cancel{5x} \\ x^2 - 3x - 10 &= \cancel{10} - \cancel{10} \\ x^2 - 3x - 10 &= 0. \end{aligned}$$

Remember, if we perform the same operations on both sides of the equation, the resulting equation has the same solutions. Therefore, the values of x that satisfy $x^2 - 3x - 10 = 0$, namely $x = -2$ and $x = 5$, also satisfy $2x^2 + 2x = x^2 + 5x + 10$, which is the original problem we’re trying to solve.

This “shuffling of terms” approach will work for any polynomial equation $A(x) = B(x)$. We can always rewrite it as $C(x) = 0$, where $C(x)$ is a new polynomial with coefficients equal to the difference of the coefficients of A and B . Don’t worry about which side you move all the coefficients to because $C(x) = 0$ and $0 = -C(x)$ have exactly the same solutions. Furthermore, the degree of the polynomial C can be no greater than that of A or B .

The form $C(x) = 0$ is the *standard form* of a polynomial, and we’ll explore several formulas you can use to find its solution(s).

Formulas

The formula for solving the polynomial equation $P(x) = 0$ depends on the *degree* of the polynomial in question.

For a first-degree polynomial equation, $P_1(x) = mx + b = 0$, the solution is $x = \frac{-b}{m}$: just move b to the other side and divide by m .

For a second-degree polynomial,

$$P_2(x) = ax^2 + bx + c = 0,$$

the solutions are $x_1 = \frac{-b + \sqrt{b^2 - 4ac}}{2a}$ and $x_2 = \frac{-b - \sqrt{b^2 - 4ac}}{2a}$.

If $b^2 - 4ac < 0$, the solutions will involve taking the square root of a negative number. In those cases, we say no real solutions exist.

There is also a formula for polynomials of degree 3 and 4, but they are complicated. For polynomials with order ≥ 5 , there does not exist a general analytical solution.

Using a computer

When solving real-world problems, you'll often run into much more complicated equations. To find the solutions of anything more complicated than the quadratic equation, I recommend using a computer algebra system like SymPy: <http://live.sympy.org>.

To make SymPy solve the standard-form equation $C(x) = 0$, call the function `solve(expr, var)`, where the expression `expr` corresponds to $C(x)$, and `var` is the variable you want to solve for. For example, to solve $x^2 - 3x + 2 = 0$, type in the following:

```
>>> solve(x**2 - 3*x + 2, x)          # usage: solve(expr, var)
[1, 2]
```

The function `solve` will find the solutions to any equation of the form `expr = 0`. In this case, we see the solutions are $x = 1$ and $x = 2$.

Another way to solve the equation is to factor the polynomial $C(x)$ using the function `factor` like this:

```
>>> factor(x**2 - 3*x + 2)           # usage: factor(expr)
(x - 1)*(x - 2)
```

We see that $x^2 - 3x + 2 = (x - 1)(x - 2)$, which confirms the two roots are indeed $x = 1$ and $x = 2$.

Substitution trick

Sometimes you can solve fourth-degree polynomials by using the quadratic formula. Say you're asked to solve for x in

$$x^4 - 7x^2 + 10 = 0.$$

Imagine this problem is on your exam, where you are not allowed to use a computer. How does the teacher expect you to solve for x ? The trick is to substitute $y = x^2$ and rewrite the same equation as

$$y^2 - 7y + 10 = 0,$$

which you can solve by applying the quadratic formula. If you obtain the solutions $y = \alpha$ and $y = \beta$, then the solutions to the original fourth-degree polynomial are $x = \pm\sqrt{\alpha}$ and $x = \pm\sqrt{\beta}$, since $y = x^2$.

Since we're not taking an exam right now, we are allowed to use the computer to find the roots:

```
>>> solve(y**2 - 7*y + 10, y)
[2, 5]
>>> solve(x**4 - 7*x**2 + 10, x)
[sqrt(2), -sqrt(2), sqrt(5), -sqrt(5)]
```

Note how the second-degree polynomial has two roots, while the fourth-degree polynomial has four roots.

Even and odd functions

The polynomials form an entire family of functions. Depending on the choice of degree n and coefficients $a_0, a_1, \dots, a_n$, a polynomial function can take on many different shapes. Consider the following observations about the symmetries of polynomials:

- If a polynomial contains only even powers of x , like $f(x) = 1 + x^2 - x^4$ for example, we call this polynomial *even*. Even polynomials have the property $f(x) = f(-x)$. The sign of the input doesn't matter.
- If a polynomial contains only odd powers of x , for example $g(x) = x + x^3 - x^9$, we call this polynomial *odd*. Odd polynomials have the property $g(x) = -g(-x)$.
- If a polynomial has both even and odd terms then it is neither even nor odd.

The terminology of *odd* and *even* applies to functions in general and not just to polynomials. All functions that satisfy $f(x) = f(-x)$ are called *even functions*, and all functions that satisfy $f(x) = -f(-x)$ are called *odd functions*.

Sine

The sine function represents a fundamental unit of vibration. The graph of $\sin(x)$ *oscillates* up and down and crosses the x -axis multiple times. The shape of the graph of $\sin(x)$ corresponds to the shape of a vibrating string. See Figure 1.26.

In the remainder of this book, we'll meet the function $\sin(x)$ many times. We'll define the function $\sin(x)$ more formally as a trigonometric ratio in Section 1.11. In Section 1.13 we'll use $\sin(x)$ and $\cos(x)$ (another trigonometric ratio) to work out the *components* of vectors.

At this point in the book, however, we don't want to go into too much detail about all these applications. Let's hold off on the discussion about vectors, triangles, angles, and ratios of lengths of sides and instead just focus on the graph of the function $f(x) = \sin(x)$.

Graph




Figure 1.26: The graph of the function $y = \sin(x)$ passes through the following (x, y) coordinates: $(0, 0)$, $(\frac{\pi}{6}, \frac{1}{2})$, $(\frac{\pi}{4}, \frac{\sqrt{2}}{2})$, $(\frac{\pi}{3}, \frac{\sqrt{3}}{2})$, $(\frac{\pi}{2}, 1)$, $(\frac{2\pi}{3}, \frac{\sqrt{3}}{2})$, $(\frac{3\pi}{4}, \frac{\sqrt{2}}{2})$, $(\frac{5\pi}{6}, \frac{1}{2})$, and $(\pi, 0)$. For x between π and 2π , the function's graph has the same shape it has for x between 0 and π , but with negative values.




Figure 1.27: The function $f(x) = \sin(x)$ crosses the x -axis at $x = \pi$.

Let's start at $x = 0$ and follow the graph of the function $\sin(x)$ as it goes up and down. The graph starts from $(0, 0)$ and smoothly

increases until it reaches the maximum value at $x = \frac{\pi}{2}$. Afterward, the function comes back down to cross the x -axis at $x = \pi$. After π , the function drops below the x -axis and reaches its minimum value of -1 at $x = \frac{3\pi}{2}$. It then travels up again to cross the x -axis at $x = 2\pi$. This 2π -long cycle repeats after $x = 2\pi$. This is why we call the function *periodic*—the shape of the graph repeats.




Figure 1.28: The graph of $\sin(x)$ from $x = 0$ to $x = 2\pi$ repeats periodically everywhere else on the number line.

Properties

- Domain: $\mathbb{R}$. The function $f(x) = \sin(x)$ is defined for all input values.
- Image: $\{y \in \mathbb{R} \mid -1 \leq y \leq 1\}$. The outputs of the sine function are always between -1 and 1 .
- Roots: $\{\dots, -3\pi, -2\pi, -\pi, 0, \pi, 2\pi, 3\pi, \dots\}$.
The function $\sin(x)$ has roots at all multiples of π .
- The function is periodic, with period 2π : $\sin(x) = \sin(x + 2\pi)$.
- The sin function is *odd*: $\sin(x) = -\sin(-x)$
- Relation to cos: $\sin^2 x + \cos^2 x = 1$
- Relation to csc: $\csc(x) = \frac{1}{\sin x}$ (csc is read *cosecant*)
- The inverse function of $\sin(x)$ is denoted as $\sin^{-1}(x)$ or $\arcsin(x)$, not to be confused with $(\sin(x))^{-1} = \frac{1}{\sin(x)} = \csc(x)$.
- The number $\sin(\theta)$ is the length-ratio of the vertical side and the hypotenuse in a right-angle triangle with angle θ at the base.

Links

[See the Wikipedia page for nice illustrations]
<http://en.wikipedia.org/wiki/Sine>

Cosine

The cosine function is the same as the sine function *shifted* by $\frac{\pi}{2}$ to the left: $\cos(x) = \sin(x + \frac{\pi}{2})$. Thus everything you know about the sine function also applies to the cosine function.

Graph




Figure 1.29: The graph of the function $y = \cos(x)$ passes through the following (x, y) coordinates: $(0, 1)$, $(\frac{\pi}{6}, \frac{\sqrt{3}}{2})$, $(\frac{\pi}{4}, \frac{\sqrt{2}}{2})$, $(\frac{\pi}{3}, \frac{1}{2})$, $(\frac{\pi}{2}, 0)$, $(\frac{2\pi}{3}, -\frac{1}{2})$, $(\frac{3\pi}{4}, -\frac{\sqrt{2}}{2})$, $(\frac{5\pi}{6}, -\frac{\sqrt{3}}{2})$, and $(\pi, -1)$.

The $\cos$ function starts at $\cos(0) = 1$, then drops down to cross the x -axis at $x = \frac{\pi}{2}$. $\cos$ continues until it reaches its minimum value at $x = \pi$. The function then moves upward, crossing the x -axis again at $x = \frac{3\pi}{2}$, and reaching its maximum value again at $x = 2\pi$.

Properties

- Domain: $\mathbb{R}$
- Image: $\{y \in \mathbb{R} \mid -1 \leq y \leq 1\}$
- Roots: $\{\dots, -\frac{3\pi}{2}, -\frac{\pi}{2}, \frac{\pi}{2}, \frac{3\pi}{2}, \frac{5\pi}{2}, \dots\}$
- Relation to $\sin$: $\sin^2 x + \cos^2 x = 1$
- Relation to $\sec$: $\sec(x) = \frac{1}{\cos x}$ ($\sec$ is read *secant*)
- The inverse function of $\cos(x)$ is denoted $\cos^{-1}(x)$ or $\arccos(x)$.
- The $\cos$ function is *even*: $\cos(x) = \cos(-x)$
- The number $\cos(\theta)$ is the length-ratio of the horizontal side and the hypotenuse in a right-angle triangle with angle θ at the base

Tangent

The tangent function is the ratio of the sine and cosine functions:

$$f(x) = \tan(x) = \frac{\sin(x)}{\cos(x)}.$$

Graph




Figure 1.30: The graph of the function $f(x) = \tan(x)$.

Properties

- Domain: $\{x \in \mathbb{R} \mid x \neq \frac{(2n+1)\pi}{2} \text{ for any } n \in \mathbb{Z}\}$
- Image: $\mathbb{R}$
- The function $\tan$ is periodic with period π .
- The $\tan$ function “blows up” at values of x where $\cos x = 0$. These are called *asymptotes* of the function and their locations are $x = \dots, -\frac{3\pi}{2}, -\frac{\pi}{2}, \frac{\pi}{2}, \frac{3\pi}{2}, \dots$
- Value at $x = 0$: $\tan(0) = \frac{0}{1} = 0$, because $\sin(0) = 0$.
- Value at $x = \frac{\pi}{4}$: $\tan\left(\frac{\pi}{4}\right) = \frac{\sin(\frac{\pi}{4})}{\cos(\frac{\pi}{4})} = \frac{\frac{\sqrt{2}}{2}}{\frac{\sqrt{2}}{2}} = 1$.
- The number $\tan(\theta)$ is the length-ratio of the vertical and the horizontal sides in a right-angle triangle with angle θ .
- The inverse function of $\tan(x)$ is denoted $\tan^{-1}(x)$ or $\arctan(x)$.
- The inverse tangent function is used to compute the angle at the base in a right-angle triangle with horizontal side length ℓ_h and vertical side length ℓ_v : $\theta = \tan^{-1}\left(\frac{\ell_v}{\ell_h}\right)$.

Exponential

The exponential function base $e = 2.7182818\dots$ is denoted

$$f(x) = e^x = \exp(x).$$

Graph




Figure 1.31: The graph of the exponential function $f(x) = e^x$ passes through the following points: $(-2, \frac{1}{e^2})$, $(-1, \frac{1}{e})$, $(0, 1)$, $(1, e)$, $(2, e^2)$, $(3, e^3)$, $(4, e^4)$, etc.

Properties

- Domain: $\mathbb{R}$
- Image: $\{y \in \mathbb{R} \mid y > 0\}$
- $f(a)f(b) = f(a+b)$ since $e^a e^b = e^{a+b}$

A more general exponential function would be $f(x) = Ae^{\gamma x}$, where A is the initial value, and γ (the Greek letter *gamma*) is the *rate* of the exponential. For $\gamma > 0$, the function $f(x)$ is increasing, as in Figure 1.31. For $\gamma < 0$, the function is decreasing and tends to zero for large values of x . The case $\gamma = 0$ is special since $e^0 = 1$, so $f(x)$ is a constant of $f(x) = A1^x = A$.

Links

[The exponential function 2^x evaluated]

<http://www.youtube.com/watch?v=e4MSN6IImpI>

Natural logarithm

The natural logarithm function is denoted

$$f(x) = \ln(x) = \log_e(x).$$

The function $\ln(x)$ is the inverse function of the exponential e^x .

Graph




Figure 1.32: The graph of the function $\ln(x)$ passes through the following coordinates: $(\frac{1}{e^2}, -2)$, $(\frac{1}{e}, -1)$, $(1, 0)$, $(e, 1)$, $(e^2, 2)$, $(e^3, 3)$, $(e^4, 4)$, etc.

Properties

- Domain: $\{x \in \mathbb{R} \mid x > 0\}$
- Image: $\mathbb{R}$

Exercises

E1.10 Find the domain, the image, and the roots of $f(x) = 2 \cos(x)$.

E1.11 What are the degrees of the following polynomials? Are they even, odd, or neither?

a) $p(x) = x^2 - 5x^4 + 1$

b) $q(x) = x - x^3 + x^5 - x^7$

E1.12 Solve for x in the following polynomial equations.

a) $3x + x^2 = x - 15 + 2x^2$

b) $3x^2 - 4x - 4 + x^3 = x^3 + 2x + 2$

1.10 Geometry

The word “geometry” comes from the Greek roots *geo*, which means “earth,” and *metron*, which means “measurement.” This name is linked to one of the early applications of geometry, which was to measure the total amount of land contained within a certain boundary region. Over the years, the study of geometry evolved to be more abstract. Instead of developing formulas for calculating the area of specific regions of land, mathematicians developed general area formulas that apply to *all* regions that have a particular shape.

In this section we’ll present formulas for calculating the perimeters, areas, and volumes for various shapes (also called “figures”) commonly encountered in the real world. For two-dimensional figures, the main quantities of interest are the figures’ areas and the figures’ perimeters (the length of the walk around the figure). For three-dimensional figures, the quantities of interest are the surface area (how much paint it would take to cover all sides of the figure), and volume (how much water it would take to fill a container of this shape). The formulas presented are by no means an exhaustive list of everything there is to know about geometry, but they represent a core set of facts that you want to add to your toolbox.

Triangles

The area of a triangle is equal to $\frac{1}{2}$ times the length of its base times its height:

$$A = \frac{1}{2}ah_a.$$

Note that h_a is the height of the triangle *relative to* the side a .




Figure 1.33: A triangle with side lengths a , b , and c . The height of the triangle with respect to the side a is denoted h_a .

The perimeter of a triangle is given by the sum of its side lengths:

$$P = a + b + c.$$

Interior angles of a triangle rule The sum of the inner angles in any triangle is equal to 180° . Consider a triangle with internal angles α , β and γ as shown in Figure 1.34. We may not know the values of

the individual angles α , β , and γ , but we know their sum is $\alpha + \beta + \gamma = 180^\circ$.




Figure 1.34: A triangle with inner angles α , β , and γ and sides a , b , and c .

Sine rule The sine rule states the following equation is true:

$$\frac{a}{\sin(\alpha)} = \frac{b}{\sin(\beta)} = \frac{c}{\sin(\gamma)},$$

where α is the angle opposite to side a , β is the angle opposite to side b , and γ is the angle opposite to side c , as shown in Figure 1.34.

Cosine rule The cosine rule states the following equations are true:

$$a^2 = b^2 + c^2 - 2bc \cos(\alpha),$$

$$b^2 = a^2 + c^2 - 2ac \cos(\beta),$$

$$c^2 = a^2 + b^2 - 2ab \cos(\gamma).$$

These equations are useful when you know two sides of a triangle and the angle between them, and you want to find the third side.

Circle

The circle is a beautiful shape. If we take the centre of the circle at the origin $(0,0)$, the circle of radius r corresponds to the equation

$$x^2 + y^2 = r^2.$$

This formula describes the set of points (x,y) with a distance from the centre equal to r .

Area

The area enclosed by a circle of radius r is given by $A = \pi r^2$. A circle of radius $r = 1$ has area π .

Circumference and arc length

The circumference of a circle of radius r is

$$C = 2\pi r.$$

A circle of radius $r = 1$ has circumference 2π . This is the total length you can measure by following the curve all the way around to trace the outline of the entire circle. For example, the circumference of a circle of radius 3 m is $C = 2\pi(3) = 18.85$ m. This is how far you'll need to walk to complete a full turn around a circle of radius $r = 3$ m.

What is the length of a part of the circle? Say you have a piece of the circle, called an *arc*, and that piece corresponds to the angle $\theta = 57^\circ$ as shown in Figure 1.35. What is the arc's length ℓ ? If the circle's total length $C = 2\pi r$ represents a full 360° turn around the circle, then the arc length ℓ for a portion of the circle corresponding to the angle θ is

$$\ell = 2\pi r \frac{\theta}{360}.$$

The arc length ℓ depends on r , the angle θ , and a factor of $\frac{2\pi}{360}$.




Figure 1.35: The arc length ℓ equals $\frac{57}{360}$ of the circle's circumference $2\pi r$.

Radians

While scientists and engineers commonly use degrees as a measurement unit for angles, mathematicians prefer to measure angles in *radians*, denoted rad.

Measuring an angle in radians is equivalent to measuring the arc length ℓ on a circle with radius $r = 1$, as illustrated in Figure 1.36.




Figure 1.36: The angle θ measured in radians corresponds to the arc length ℓ on a circle with radius 1. The full circle corresponds to the angle 2π rad.

The conversion ratio between degrees and radians is

$$2\pi \text{ rad} = 360^\circ.$$

When the angle θ is measured in radians, the arc length is given by:

$$\ell = r\theta.$$

To find the arc length ℓ , we simply multiply the circle radius r times the angle θ measured in radians.

Note the arc-length formula with θ measured in radians is simpler than the arc-length formula with θ measured in degrees, since we don't need the conversion factor of 360° .

Sphere

A sphere of radius r is described by the equation $x^2 + y^2 + z^2 = r^2$. The surface area of the sphere is $A = 4\pi r^2$, and its volume is given by $V = \frac{4}{3}\pi r^3$.




Figure 1.37: A sphere of radius r has surface area $4\pi r^2$ and volume $\frac{4}{3}\pi r^3$.

Cylinder

The surface area of a cylinder consists of the top and bottom circular surfaces, plus the area of the side of the cylinder:

$$A = 2(\pi r^2) + (2\pi r)h.$$

The volume of a cylinder is the product of the area of the cylinder's base times its height:

$$V = (\pi r^2) h.$$




Figure 1.38: A cylinder with radius r and height h has volume $\pi r^2 h$.

Example You open the hood of your car and see 2.0 L written on top of the engine. The 2.0 L refers to the combined volume of the four pistons, which are cylindrical in shape. The owner's manual tells you the radius of each piston is 43.75 mm, and the height of each piston is 83.1 mm. Verify the total engine volume is $1998789 \text{ mm}^3 \approx 2 \text{ L}$.

Cones and pyramids

The volume of a square pyramid with side length a and height h is given by the formula $V = \frac{1}{3}a^2h$. The volume of a cone of radius r and height h is given by the formula $V = \frac{1}{3}\pi r^2h$. Note the factor $\frac{1}{3}$ appears in both formulas. These two formulas are particular cases of the general volume formula that applies to all pyramids:

$$V = \frac{1}{3}Ah,$$

where A is the area of the pyramid's base and h is its height. This formula applies for pyramids with a base that is a triangle (triangular pyramids), a square (square pyramids), a rectangle (rectangular pyramids), a circle (cones), or any other shape.





Figure 1.39: The volumes of pyramids and cones are described by the formula $V = \frac{1}{3}Ah$, where A is the area of the base and h is the height.

Exercises

E1.13 Find the length of side x in the triangle below.



Hint: Use the cosine rule.

E1.14 Find the volume and the surface area of a sphere with radius 2.

E1.15 On a rainy day, Laura brings her bike indoors, and the wet bicycle tires leave a track of water on the floor. What is the length of the water track left by the bike's rear tire (diameter 73 cm) if the wheel makes five full turns along the floor?

1.11 Trigonometry

If one of the angles in a triangle is equal to 90° , we call this triangle a *right-angle triangle*. In this section we'll discuss right-angle triangles in great detail and get to know their properties. We'll learn some fancy new terms like *hypotenuse*, *opposite*, and *adjacent*, which are used to refer to the different sides of a triangle. We'll also use the functions *sine*, *cosine*, and *tangent* to compute the *ratios of lengths* in right triangles.

Understanding triangles and their associated trigonometric functions is of fundamental importance: you'll need this knowledge for your future understanding of mathematical concepts like vectors and complex numbers.




Figure 1.40: A right-angle triangle. The angle at the base is denoted θ and the names of the sides of the triangle are indicated.

Concepts

- A, B, C : the three *vertices* of the triangle
- θ : the angle at the vertex C . Angles can be measured in degrees or radians.
- $\text{opp} = AB$: the length of the *opposite* side to θ
- $\text{adj} = BC$: the length of side *adjacent* to θ
- $\text{hyp} = AC$: the *hypotenuse*. This is the triangle's longest side.
- h : the "height" of the triangle (in this case $h = \text{opp} = AB$)
- $\sin \theta = \frac{\text{opp}}{\text{hyp}}$: the *sine* of theta is the ratio of the length of the opposite side and the length of the hypotenuse
- $\cos \theta = \frac{\text{adj}}{\text{hyp}}$: the *cosine* of theta is the ratio of the adjacent length and the hypotenuse length
- $\tan \theta = \frac{\sin \theta}{\cos \theta} = \frac{\text{opp}}{\text{adj}}$: the *tangent* is the ratio of the opposite length divided by the adjacent length

Pythagoras' theorem

In a right-angle triangle, the length of the hypotenuse squared is equal to the sum of the squares of the lengths of the other sides:

$$\text{adj}^2 + \text{opp}^2 = \text{hyp}^2.$$

If we divide both sides of the above equation by hyp^2 , we obtain

$$\frac{\text{adj}^2}{\text{hyp}^2} + \frac{\text{opp}^2}{\text{hyp}^2} = 1.$$

Since $\frac{\text{adj}}{\text{hyp}} = \cos \theta$ and $\frac{\text{opp}}{\text{hyp}} = \sin \theta$, this equation can be rewritten as

$$\cos^2 \theta + \sin^2 \theta = 1.$$

This is a powerful *trigonometric identity* that describes an important relation between sine and cosine functions. In case you've never seen this notation before, the expression $\cos^2 \theta$ is used to denote $(\cos(\theta))^2$.

Sin and cos

Meet the trigonometric functions, or trigs for short. These are your new friends. Don't be shy now, say hello to them.

"Hello."

"Hi."

"Sooooooo, you are like functions right?"

"Yep," sin and cos reply in chorus.

"Okay, so what do you do?"

"Who me?" asks cos. "Well I tell the ratio...hmm...Wait, are you asking what I do as a *function* or specifically what *I* do?"

"Both I guess?"

"Well, as a function, I take angles as inputs and I give ratios as answers. More specifically, I tell you how 'wide' a triangle with that angle will be," says cos all in one breath.

"What do you mean wide?" you ask.

"Oh yeah, I forgot to say, the triangle must have a hypotenuse of length 1. What happens is there is a point P that moves around on a circle of radius 1, and we *imagine* a triangle formed by the point P , the origin, and the point on the x -axis located directly below the point P ."

"I am not sure I get it," you confess.

"Let me try explaining," says sin. "Look at Figure 1.41 and you'll see a circle. This is the unit circle because it has a radius of 1. You see it, yes?"

"Yes."

"Now imagine a point P that moves along the circle of radius 1, starting from the point $P(0) = (1, 0)$. The x and y coordinates of the point $P(\theta) = (P_x(\theta), P_y(\theta))$ as a function of θ are

$$P(\theta) = (P_x(\theta), P_y(\theta)) = (\cos \theta, \sin \theta).$$

So, *either* you can think of us in the context of triangles, or in the context of the unit circle."

"Cool. I kind of get it. Thanks so much," you say, but in reality you are weirded out. Talking functions? "Well guys. It was nice to meet you, but I have to get going, to finish the rest of the book."

"See you later," says cos.

"Peace out," says sin.

The unit circle

The *unit circle* is a circle of radius one centred at the origin. The unit circle consists of all points (x, y) that satisfy the equation $x^2 + y^2 = 1$. A point P on the unit circle has coordinates $(P_x, P_y) = (\cos \theta, \sin \theta)$, where θ is the angle P makes with the x -axis.




Figure 1.41: The unit circle corresponds to the equation $x^2 + y^2 = 1$. The coordinates of the point P on the unit circle are $P_x = \cos \theta$ and $P_y = \sin \theta$.




Figure 1.42: The function $f(\theta) = \sin \theta$ describes the vertical position of a point P that travels along the unit circle. The graph shows the values of the function $f(\theta) = \sin \theta$ for angles between $\theta = 0$ and $\theta = \pi$.

Figure 1.42 shows the graph of the function $f(\theta) = \sin \theta$. The values $\sin \theta$ for the angles 0 , $\frac{\pi}{6}$ (30°), $\frac{\pi}{3}$ (60°), and $\frac{\pi}{2}$ (90°) are marked. There are three values to remember: $\sin \theta = 0$ when $\theta = 0$, $\sin \theta = \frac{1}{2}$ when $\theta = \frac{\pi}{6}$ (30°), and $\sin \theta = 1$ when $\theta = \frac{\pi}{2}$ (90°). See Figure 1.26 (page 60) for a graph of $\sin \theta$ that shows a complete cycle around the circle. Also see Figure 1.29 (page 62) for the graph of $\cos \theta$.

Instead of trying to memorize the values of the functions $\cos \theta$ and $\sin \theta$ separately, it's easier to remember them as a combined “package” $(\cos \theta, \sin \theta)$, which describes the x - and y -coordinates of the point P for the angle θ . Figure 1.43 shows the values of $\cos \theta$ and $\sin \theta$ for the angles 0 , $\frac{\pi}{6}$ (30°), $\frac{\pi}{4}$ (45°), $\frac{\pi}{3}$ (60°), and $\frac{\pi}{2}$ (90°). These are the most common angles that often show up on homework and exam questions. For each angle, the x -coordinate (the first number in the bracket) is $\cos \theta$, and the y -coordinate (the second number in the bracket) is $\sin \theta$.




Figure 1.43: The combined $(\cos \theta, \sin \theta)$ coordinates for the points on the unit circle at the most common angles: 0 , $\frac{\pi}{6}$ (30°), $\frac{\pi}{4}$ (45°), $\frac{\pi}{3}$ (60°), and $\frac{\pi}{2}$ (90°).

Note the values of $\cos \theta$ and $\sin \theta$ for the angles shown in Figure 1.43 are all combinations of the fractions $\frac{1}{2}$, $\frac{\sqrt{2}}{2}$, and $\frac{\sqrt{3}}{2}$. The square roots appear as a consequence of the trigonometric identity $\cos^2 \theta + \sin^2 \theta = 1$. This identity tells us that the sum of the squared coordinates of each point on the unit circle is equal to one. Let's look at what this equation tells us for the angle $\theta = \frac{\pi}{6}$ (30°). Remember that $\sin(30^\circ) = \frac{1}{2}$ (the length of the dashed line in Figure 1.43). We can plug this value into the equation $\cos^2(30^\circ) + \sin^2(30^\circ) = 1$ to find the value: $\cos(30^\circ) = \sqrt{1 - \sin^2(30^\circ)} = \sqrt{1 - \frac{1}{4}} = \sqrt{\frac{3}{4}} = \frac{\sqrt{3}}{2}$.

The coordinates $(\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2})$ for the angle $\theta = \frac{\pi}{4}$ (45°) are obtained from a similar calculation. We know the values of $\sin \theta$ and $\cos \theta$ must be equal for that angle, so we're looking for the number a that satisfies the equation $a^2 + a^2 = 1$, which is $a = \frac{1}{\sqrt{2}} = \frac{\sqrt{2}}{2}$. The values of $\cos(60^\circ)$ and $\sin(60^\circ)$ can be obtained from a symmetry argument. Measuring 60° from the x -axis is the same as measuring 30° from the y -axis, so $\cos(60^\circ) = \sin(30^\circ) = \frac{1}{2}$ and $\sin(60^\circ) = \cos(30^\circ) = \frac{\sqrt{3}}{2}$.

We can extend the calculations described above for all other angles that are multiples of $\frac{\pi}{6}$ (30°) and $\frac{\pi}{4}$ (45°) to obtain the $\cos \theta$ and $\sin \theta$ values for the whole unit circle, as shown in Figure 1.44.

Don't be intimidated by all the information shown in Figure 1.44! You're not expected to memorize all these values. The primary reason for including this figure is so you can appreciate the symmetries of the sine and cosine values that we find as we go around the circle. The values of $\sin \theta$ and $\cos \theta$ for all angles are the same as the values for the angles between 0° and 90° , but one or more of their coordinates has a negative sign. For example, 150° is just like 30° , except its x -coordinate is negative since the point lies to the left of the y -axis. Another use for Figure 1.44 is to convert between angles measured in degrees and radians, since both units for angles are indicated.




Figure 1.44: The coordinates of the point on the unit circle $(\cos \theta, \sin \theta)$ are indicated for all multiples of $\frac{\pi}{6}$ (30°) and $\frac{\pi}{4}$ (45°). Note the symmetries.

Non-unit circles

Consider a point $Q(\theta)$ at an angle of θ on a circle with radius $r \neq 1$. How can we find the x - and y -coordinates of the point $Q(\theta)$?

We saw that the coefficients $\cos \theta$ and $\sin \theta$ correspond to the x - and y -coordinates of a point on the *unit* circle ($r = 1$). To obtain the coordinates for a point on a circle of radius r , we must *scale* the coordinates by a factor of r :

$$Q(\theta) = (Q_x(\theta), Q_y(\theta)) = (r \cos \theta, r \sin \theta).$$

The take-away message is that you can use the functions $\cos \theta$ and $\sin \theta$ to find the “horizontal” and “vertical” components of any length r . From this point on in the book, we’ll always talk about the length of the *adjacent* side as $x = r \cos \theta$, and the length of the *opposite* side as $y = r \sin \theta$. It is extremely important you get comfortable with this notation.

The reasoning behind the above calculations is as follows:

$$\cos \theta = \frac{\text{adj}}{\text{hyp}} = \frac{x}{r} \quad \Rightarrow \quad x = r \cos \theta,$$




Figure 1.45: The x - and y -coordinates of a point at the angle θ and distance of r from the origin are given by $x = r \cos \theta$ and $y = r \sin \theta$.

and

$$\sin \theta = \frac{\text{opp}}{\text{hyp}} = \frac{y}{r} \quad \Rightarrow \quad y = r \sin \theta.$$

Calculators

Watch out for the units of angle measures when using calculators and computers. Make sure you know what kind of angle units the functions $\sin$, $\cos$, and $\tan$ expect as inputs, and what kind of outputs the functions $\sin^{-1}$, $\cos^{-1}$, and $\tan^{-1}$ return.

For example, let's see what we should type into the calculator to compute the sine of 30 degrees. If the calculator is set to degrees, we simply type: $\boxed{3}$, $\boxed{0}$, $\boxed{\sin}$, $\boxed{=}$, and obtain the answer 0.5.

If the calculator is set to radians, we have two options:

1. Change the mode of the calculator so it works in degrees.
2. Convert 30° to radians

$$30^\circ \times \frac{2\pi \text{ rad}}{360^\circ} = \frac{\pi}{6} \text{ rad},$$

and type: $\boxed{\pi}$, $\boxed{/}$, $\boxed{6}$, $\boxed{\sin}$, $\boxed{=}$ on the calculator.

Try computing $\cos(60^\circ)$, $\cos(\frac{\pi}{3} \text{ rad})$, and $\cos^{-1}(\frac{1}{2})$ using your calculator to make sure you know how it works.

Exercises

E1.16 Given a circle with radius $r = 5$, find the x - and y -coordinates of the point at $\theta = 45^\circ$. What is the circumference of the circle?

E1.17 Convert the following angles from degrees to radians.

a) 30°

b) 45°

c) 60°

d) 270°

Links

[Unit-circle walkthrough and tricks by patrickJMT on YouTube]
<http://bit.ly/1mQg9Cj> and <http://bit.ly/1hvA702>

1.12 Trigonometric identities

There are a number of important relationships between the values of the functions $\sin$ and $\cos$. Here are three of these relationships, known as *trigonometric identities*. There about a dozen other identities that are less important, but you should memorize these three.

The three identities to remember are:

1. Unit hypotenuse

$$\sin^2 \theta + \cos^2 \theta = 1.$$

The unit hypotenuse identity is true by the Pythagoras theorem and the definitions of $\sin$ and $\cos$. The sum of the squares of the sides of a triangle is equal to the square of the hypotenuse.

2. Sine angle sum

$$\sin(a + b) = \sin(a) \cos(b) + \sin(b) \cos(a).$$

The mnemonic for this identity is “sico + sico.”

3. Cosine angle sum

$$\cos(a + b) = \cos(a) \cos(b) - \sin(a) \sin(b).$$

The mnemonic for this identity is “coco – sisi.” The negative sign is there because it’s not good to be a sissy.

Derived formulas

If you remember the above three formulas, you can derive pretty much all the other trigonometric identities.

Double angle formulas

Starting from the sico + sico identity and setting $a = b = x$, we can derive the following identity:

$$\sin(2x) = 2 \sin(x) \cos(x).$$

Starting from the coco-sisi identity, we obtain

$$\begin{aligned}\cos(2x) &= \cos^2(x) - \sin^2(x) \\ &= 2\cos^2(x) - 1 = 2(1 - \sin^2(x)) - 1 = 1 - 2\sin^2(x).\end{aligned}$$

The formulas for expressing $\sin(2x)$ and $\cos(2x)$ in terms of $\sin(x)$ and $\cos(x)$ are called *double angle formulas*.

If we rewrite the double-angle formula for $\cos(2x)$ to isolate the $\sin^2$ or the $\cos^2$ term, we obtain the *power-reduction formulas*:

$$\cos^2(x) = \frac{1}{2}(1 + \cos(2x)), \quad \sin^2(x) = \frac{1}{2}(1 - \cos(2x)).$$

Self-similarity

Sin and cos are periodic functions with period 2π . Adding a multiple of 2π to the function's input does not change the function:

$$\sin(x + 2\pi) = \sin(x), \quad \cos(x + 2\pi) = \cos(x).$$

This follows because adding a multiple of 2π brings us back to the same point on the unit circle.

Furthermore, sin and cos have symmetries with respect to zero,

$$\sin(-x) = -\sin(x), \quad \cos(-x) = \cos(x),$$

within each π half-cycle,

$$\sin(\pi - x) = \sin(x), \quad \cos(\pi - x) = -\cos(x),$$

and within each full 2π cycle,

$$\sin(2\pi - x) = -\sin(x), \quad \cos(2\pi - x) = \cos(x).$$

Take the time to revisit Figure 1.26 (page 60), Figure 1.29 (page 62), and Figure 1.44 (page 76) to visually confirm that all the equations shown above are true. Knowing the points where the functions take on the same values (symmetries) or take on opposite values (anti-symmetries) is very useful in calculations.

Sin is cos, cos is sin

It shouldn't be surprising if I tell you that sin and cos are actually $\frac{\pi}{2}$ -shifted versions of each other:

$$\cos(x) = \sin\left(x + \frac{\pi}{2}\right), \quad \sin(x) = \cos\left(x - \frac{\pi}{2}\right).$$

Formulas for sums and products

Here are some formulas for transforming sums into products:

$$\sin(a) + \sin(b) = 2 \sin\left(\frac{1}{2}(a+b)\right) \cos\left(\frac{1}{2}(a-b)\right),$$

$$\sin(a) - \sin(b) = 2 \sin\left(\frac{1}{2}(a-b)\right) \cos\left(\frac{1}{2}(a+b)\right),$$

$$\cos(a) + \cos(b) = 2 \cos\left(\frac{1}{2}(a+b)\right) \cos\left(\frac{1}{2}(a-b)\right),$$

$$\cos(a) - \cos(b) = -2 \sin\left(\frac{1}{2}(a+b)\right) \sin\left(\frac{1}{2}(a-b)\right).$$

And here are some formulas for transforming products into sums:

$$\sin(a) \cos(b) = \frac{1}{2} \left(\sin(a+b) + \sin(a-b) \right),$$

$$\sin(a) \sin(b) = \frac{1}{2} \left(\cos(a-b) - \cos(a+b) \right),$$

$$\cos(a) \cos(b) = \frac{1}{2} \left(\cos(a-b) + \cos(a+b) \right).$$

Discussion

The above formulas will come in handy when you need to find some unknown in an equation, or when you are trying to simplify

Exercises

E1.18 Given $a = \pi$ and $b = \frac{\pi}{2}$, find

a) $\sin(a+b)$

b) $\cos(2a)$

c) $\cos(a+b)$

E1.19 Simplify the following expressions and compute their value without using a calculator.

a) $\cos(x) + \cos(\pi - x)$

b) $2 \sin^2(x) + \cos(2x)$

c) $\sin\left(\frac{5\pi}{4}\right) \sin\left(-\frac{\pi}{4}\right)$

d) $2 \cos\left(\frac{5\pi}{4}\right) \cos\left(-\frac{\pi}{4}\right) \cos(\pi)$

1.13 Vectors

In this section we'll learn how to manipulate multi-dimensional objects called vectors. Vectors are the precise way to describe directions in space. We need vectors in order to describe physical quantities like the velocity of an object, its acceleration, and the net force acting on the object. Vectors are used more broadly in the study of computer graphics, probability theory, machine learning, and other fields of

science and mathematics. It's all about vectors these days, so you better get to know them.




Figure 1.46: The vector $\vec{v} = (3,2)$ is an arrow in the Cartesian plane. The horizontal component of $\vec{v}$ is $v_x = 3$ and the vertical component is $v_y = 2$.

Vectors are built from *components*, which are ordinary numbers. You can think of a vector as a list of numbers, and *vector algebra* as operations performed on the numbers in the list. Vectors can also be manipulated as geometric objects, represented by arrows in space. The arrow that corresponds to the vector $\vec{v} = (v_x, v_y)$ starts at the origin $(0,0)$ and ends at the point (v_x, v_y) . The word vector comes from the Latin *vehere*, which means *to carry*. Indeed, the vector $\vec{v}$ takes the point $(0,0)$ and carries it to the point (v_x, v_y) .




Figure 1.47: This figure illustrates the new concepts related to vectors. As you can see, there is quite a bit of new vocabulary to learn, but don't be fazed—all these terms are just fancy ways of talking about arrows.

Vectors are extremely useful in all areas of life. In physics, for example, we use a vector to describe the velocity of an object. It is not sufficient to say that the speed of a tennis ball is 200 kilometres per

hour: we must also specify the direction in which the ball is moving. Both of the two velocities

$$\vec{v}_1 = (200, 0) \quad \text{and} \quad \vec{v}_2 = (0, 200)$$

describe motion at the speed of 200 kilometres per hour; but since one velocity points along the x -axis, and the other points along the y -axis, they are *completely* different velocities. The velocity vector contains information about the object's speed *and* its direction. The direction makes a big difference. If it turns out the tennis ball is hurtling toward you, you'd better get out of the way!

Definitions

A two-dimensional vector $\vec{v}$ corresponds to a *pair of numbers*:

$$\vec{v} = (v_x, v_y),$$

where v_x is the x -*component* of the vector and v_y is its y -*component*. We denote the set of two-dimensional vectors as $\mathbb{R}^2$, since the components of a two-dimensional vector are specified by two real numbers. We'll use the mathematical shorthand $\vec{v} \in \mathbb{R}^2$ to define a two-dimensional vector $\vec{v}$. Vectors in $\mathbb{R}^2$ can be represented as arrows in the Cartesian plane. See the vector $\vec{v} = (3, 2)$ illustrated in Figure 1.46.

We can also define three-dimensional vectors like the vector $\vec{v} = (v_x, v_y, v_z) \in \mathbb{R}^3$, which has three components. Three-dimensional vectors can be represented as arrows in a coordinate system that has three axes, like the one shown in Figure 1.55 on page 92. A three-dimensional coordinate system is similar to the Cartesian coordinate system you're familiar with, and includes the additional z -axis that measures the height above the plane. In fact, there's no limit to the number of dimensions for vectors. We can define vectors in an n -dimensional space: $\vec{v} = (v_1, v_2, \dots, v_n) \in \mathbb{R}^n$. For the sake of simplicity, we'll define all the vector operation formulas using two-dimensional vectors. Unless otherwise indicated in the text, all the formulas we give for two-dimensional vectors $\vec{v} \in \mathbb{R}^2$ also apply to n -dimensional vectors $\vec{v} \in \mathbb{R}^n$.

Vector operations

Consider two vectors, $\vec{u} = (u_x, u_y)$ and $\vec{v} = (v_x, v_y)$, and assume that $\alpha \in \mathbb{R}$ is an arbitrary constant. The following operations are defined for these vectors:

- **Addition:** $\vec{u} + \vec{v} = (u_x + v_x, u_y + v_y)$

- **Subtraction:** $\vec{u} - \vec{v} = (u_x - v_x, u_y - v_y)$
- **Scaling:** $\alpha\vec{u} = (\alpha u_x, \alpha u_y)$
- **Dot product:** $\vec{u} \cdot \vec{v} = u_x v_x + u_y v_y$
- **Length:** $\|\vec{u}\| = \sqrt{\vec{u} \cdot \vec{u}} = \sqrt{u_x^2 + u_y^2}$. The vector's length is also called the *norm* of the vector. We sometimes use the letter u to denote the length of the vector $\vec{u}$.

Note there is no vector division operation.

For vectors in a three-dimensional space $\vec{u} = (u_x, u_y, u_z) \in \mathbb{R}^3$ and $\vec{v} = (v_x, v_y, v_z) \in \mathbb{R}^3$, we can also define the **cross product** operation $\vec{u} \times \vec{v} = (u_y v_z - u_z v_y, u_z v_x - u_x v_z, u_x v_y - u_y v_x)$. The dot product and the cross product are new operations that you probably haven't seen before. We'll talk more about dot products and the cross products later on. For now let's start with the basics.

Vector representations

We'll use three equivalent ways to denote vectors in two dimensions:

- $\vec{v} = (v_x, v_y)$: component notation. The vector is written as a pair of numbers called the *components* or *coordinates* of the vector.
- $\vec{v} = v_x \hat{i} + v_y \hat{j}$: unit vector notation. The vector is expressed as a combination of the unit vectors $\hat{i} = (1, 0)$ and $\hat{j} = (0, 1)$.
- $\vec{v} = \|\vec{v}\| \angle \theta$: length-and-direction notation (also known as *polar coordinates*). The vector is expressed in terms of its *length* $\|\vec{v}\|$ and the angle θ that the vector makes with the x -axis.




Figure 1.48: The vector $\vec{v} = (v_x, v_y) = v_x \hat{i} + v_y \hat{j} = \|\vec{v}\| \angle \theta$.

We use the component notation for doing vector algebra calculations since it is most compact. The unit vector notation shows explicitly that the vector $\vec{v}$ corresponds to the sum of $v_x \hat{i}$ (a displacement of v_x

steps in the direction of the x -axis) and $v_y \hat{j}$ (a displacement of v_y steps in the direction of the y -axis). The length-and-direction notation describes the vector $\vec{v}$ as a displacement of $\|\vec{v}\|$ steps in the direction of the angle θ . We'll use all three ways of denoting vectors throughout the rest of the book, and we'll learn how to convert between them.

Vector algebra

Addition and subtraction Just like numbers, you can add vectors

$$\vec{v} + \vec{w} = (v_x, v_y) + (w_x, w_y) = (v_x + w_x, v_y + w_y),$$

subtract them

$$\vec{v} - \vec{w} = (v_x, v_y) - (w_x, w_y) = (v_x - w_x, v_y - w_y),$$

and solve all kinds of equations where the unknown variable is a vector. This is not a formidably complicated new development in mathematics. Performing arithmetic calculations on vectors simply requires **carrying out arithmetic operations on their components**. Given two vectors, $\vec{v} = (4, 2)$ and $\vec{w} = (3, 7)$, their difference is computed as $\vec{v} - \vec{w} = (4, 2) - (3, 7) = (1, -5)$.

Scaling We can also *scale* a vector by any number $\alpha \in \mathbb{R}$:

$$\alpha \vec{v} = (\alpha v_x, \alpha v_y),$$

where each component is multiplied by the scale factor α . Scaling changes the length of a vector. If $\alpha > 1$ the vector will get longer, and if $0 \leq \alpha < 1$ then the vector will become shorter. If α is a negative number, the scaled vector will point in the opposite direction.

Length A vector's length is obtained from Pythagoras' theorem. Imagine a right-angle triangle with one side of length v_x and the other side of length v_y ; the length of the vector is equal to the length of the triangle's hypotenuse:

$$\|\vec{v}\|^2 = v_x^2 + v_y^2 \quad \Rightarrow \quad \|\vec{v}\| = \sqrt{v_x^2 + v_y^2}.$$

A common technique is to scale a vector $\vec{v}$ by one over its length $\frac{1}{\|\vec{v}\|}$ to obtain a unit vector that points in the same direction as $\vec{v}$:

$$\hat{v} = \frac{\vec{v}}{\|\vec{v}\|} = \left(\frac{v_x}{\|\vec{v}\|}, \frac{v_y}{\|\vec{v}\|} \right).$$

Unit vectors (denoted with a hat instead of an arrow) are useful when you want to describe only a direction in space without any specific length in mind. Verify that $\|\hat{v}\| = 1$.

Vector as arrows

So far, we described how to perform algebraic operations on vectors in terms of their components. Vector operations can also be interpreted geometrically, as operations on arrows in the Cartesian plane.

Vector addition The sum of two vectors corresponds to the combined displacement of the two vectors. Figure 1.49 illustrates the addition of two vectors, $\vec{v}_1 = (3, 0)$ and $\vec{v}_2 = (2, 2)$. The sum of the two vectors is the vector $\vec{v}_1 + \vec{v}_2 = (3, 0) + (2, 2) = (5, 2)$.




Figure 1.49: The addition of the vectors $\vec{v}_1$ and $\vec{v}_2$ produces the vector $(5, 2)$.

Vector subtraction Before we describe vector subtraction, note that multiplying a vector by a scale factor $\alpha = -1$ gives a vector of the same length as the original, but pointing in the opposite direction.

This fact is useful if you want to subtract two vectors using the graphical approach. Subtracting a vector is the same as adding the negative of the vector:

$$\vec{w} - \vec{v}_1 = \vec{w} + (-\vec{v}_1) = \vec{v}_2.$$




Figure 1.50: The vector subtraction $\vec{w} - \vec{v}_1$ is equivalent to the vector addition $\vec{w} + (-\vec{v}_1)$, where $(-\vec{v}_1)$ is like $\vec{v}_1$ but points in the opposite direction.

Figure 1.50 illustrates the graphical procedure for subtracting the vector $\vec{v}_1 = (3, 0)$ from the vector $\vec{w} = (5, 2)$. Subtraction of $\vec{v}_1 = (3, 0)$ is the same as addition of $-\vec{v}_1 = (-3, 0)$.

Scaling The scaling operation acts to change the length of a vector. Suppose we want to obtain a vector in the same direction as the vector $\vec{v} = (3, 2)$, but half as long. “Half as long” corresponds to a scale factor of $\alpha = 0.5$. The scaled-down vector is $\vec{w} = 0.5\vec{v} = (1.5, 1)$. Conversely, we can think of the vector $\vec{v}$ as being twice as long as the vector $\vec{w}$.




Figure 1.51: Vectors $\vec{v}$ and $\vec{w}$ are related by the equation $\vec{v} = 2\vec{w}$.

Multiplying a vector by a negative number reverses its direction.

Length-and-direction representation

So far, we’ve seen how to represent a vector in terms of its components. There is another way of representing two-dimensional vectors: we can describe the vector $\vec{v} \in \mathbb{R}^2$ in terms of its length $\|\vec{v}\|$ and its direction θ —the angle it makes with the x -axis. For example, the vector $(1, 1)$ can also be written as $\sqrt{2}/45^\circ$ using length-and-direction notation. This length-and-direction notation is useful because it makes it easy to see the “size” of vectors. On the other hand, vector arithmetic operations are much easier to carry out in the component notation. It’s therefore good to know the formulas for converting between the two vector representations.




Figure 1.52: The x - and y -components of a vector with length $\|\vec{v}\|$ in the direction θ are given by $v_x = \|\vec{v}\| \cos \theta$ and $v_y = \|\vec{v}\| \sin \theta$.

To convert the length-and-direction vector $\vec{v} = \|\vec{v}\| \angle \theta$ into an x -component and a y -component (v_x, v_y) , use the formulas

$$v_x = \|\vec{v}\| \cos \theta \quad \text{and} \quad v_y = \|\vec{v}\| \sin \theta.$$

To convert from component notation $\vec{v} = (v_x, v_y)$ to length-and-direction $\|\vec{v}\| \angle \theta$, use

$$\|\vec{v}\| = \sqrt{v_x^2 + v_y^2}, \quad \theta = \begin{cases} \tan^{-1}\left(\frac{v_y}{v_x}\right) & \text{if } v_x > 0, \\ 180^\circ + \tan^{-1}\left(\frac{v_y}{v_x}\right) & \text{if } v_x < 0, \\ 90^\circ & \text{if } v_x = 0 \text{ and } v_y > 0, \\ -90^\circ & \text{if } v_x = 0 \text{ and } v_y < 0. \end{cases}$$

Finding the angle θ is a little tricky. We must use a different formula for computing θ depending on the value of v_x and there are four different cases to consider. The basic idea is to use the inverse tangent function $\tan^{-1}$, which is also called arctan, or atan on computer systems. By convention, the function $\tan^{-1}$ returns values between -90° ($-\frac{\pi}{2}$ [rad]) and 90° ($\frac{\pi}{2}$ [rad]), which correspond to vectors with a positive v_x component. If the v_x component is negative, we must add 180° (π [rad]) to the output of the inverse-tangent calculation to obtain the correct angle. When $v_x = 0$ we can't compute the fraction $\frac{v_y}{v_x}$ because we cannot divide by zero, so must handle the cases with $v_x = 0$ separately as described above. Computer algebra systems provide the two-input arctangent function atan2(y, x) that will correctly compute the angle θ for any coordinate pair (x, y) .

Unit vector notation

In two dimensions, we can think of a vector $\vec{v} = (v_x, v_y)$ as a command to “Go a distance v_x in the x -direction and a distance v_y in the y -direction.” To write this set of commands more explicitly, we can use multiples of the vectors $\hat{i}$ and $\hat{j}$. These are the unit vectors pointing in the x and y directions:

$$\hat{i} = (1, 0) \quad \text{and} \quad \hat{j} = (0, 1).$$

Any number multiplied by $\hat{i}$ corresponds to a vector with that number in the first coordinate. For example, $3\hat{i} = (3, 0)$ and $4\hat{j} = (0, 4)$.

In physics, we tend to perform a lot of numerical calculations with vectors; to make things easier, we often use unit vector notation:

$$v_x \hat{i} + v_y \hat{j} \quad \Leftrightarrow \quad (v_x, v_y).$$

The addition rule remains the same for the new notation:

$$\underbrace{2\hat{i} + 3\hat{j}}_{\vec{v}} + \underbrace{5\hat{i} - 2\hat{j}}_{\vec{w}} = \underbrace{7\hat{i} + 1\hat{j}}_{\vec{v} + \vec{w}}.$$

It's the same story repeating all over again: we need to add $\hat{i}$ s with $\hat{i}$ s, and $\hat{j}$ s with $\hat{j}$ s.

Examples

Simple example

Compute the sum $\vec{s} = 4\hat{i} + 5\angle 30^\circ$. Express your answer in the length-and-direction notation.

Since we want to carry out an addition, and since addition is performed in terms of components, our first step is to convert $5\angle 30^\circ$ into component notation: $5\angle 30^\circ = 5 \cos 30^\circ \hat{i} + 5 \sin 30^\circ \hat{j} = \frac{5\sqrt{3}}{2}\hat{i} + \frac{5}{2}\hat{j}$. We can now compute the sum:

$$\vec{s} = 4\hat{i} + \frac{5\sqrt{3}}{2}\hat{i} + \frac{5}{2}\hat{j} = (4 + \frac{5\sqrt{3}}{2})\hat{i} + (\frac{5}{2})\hat{j}.$$

The x -component of the sum is $s_x = (4 + \frac{5\sqrt{3}}{2})$ and the y -component of the sum is $s_y = (\frac{5}{2})$. To express the answer as a length and a direction, we compute the length $\|\vec{s}\| = \sqrt{s_x^2 + s_y^2} = 8.697$ and the direction $\tan^{-1}(s_y/s_x) = 16.7^\circ$. The answer is $\vec{s} = 8.697\angle 16.7^\circ$.

Relative motion example

A boat can reach a top speed of 12 knots in calm seas. Instead of cruising through a calm sea, however, the boat's crew is trying to sail up the St-Laurence river. The speed of the current is 5 knots.

If the boat travels directly upstream at full throttle $12\hat{i}$, then the speed of the boat relative to the shore will be

$$12\hat{i} - 5\hat{i} = 7\hat{i},$$

since we must “deduct” the speed of the current from the speed of the boat relative to the water. See the vector diagram in Figure 1.53.



Figure 1.53: A boat travels with speed 12 knots against a current of 5 knots.

If the crew wants to cross the river perpendicular to the current flow, they can use some of the boat's thrust to counterbalance the current, and the remaining thrust to push across. The situation is

illustrated in Figure 1.54. In what direction should the boat sail to cross the river? We are looking for the direction of $\vec{v}$ the boat should take such that, after adding in the velocity of the current, the boat moves in a straight line between the two banks (in the $\hat{j}$ direction).




Figure 1.54: Part of the boat's thrust cancels the current.

Let's analyze the vector diagram. The opposite side of the triangle is parallel to the current flow and has length 5. We take the up-the-river component of the velocity $\vec{v}$ to be equal to $5\hat{i}$, so that it cancels exactly the $-5\hat{i}$ flow of the river. The hypotenuse has length 12 since this is the speed of the boat relative to the surface of the water.

From all of this we can answer the question like professionals. You want the angle? Well, we have that $\frac{\text{opp}}{\text{hyp}} = \frac{5}{12} = \sin(\theta)$, where θ is the angle of the boat's course relative to the straight line between the two banks. We can use the inverse-sin function to solve for the angle:

$$\theta = \sin^{-1}\left(\frac{5}{12}\right) = 24.62^\circ.$$

The across-the-river component of the velocity can be calculated using $v_y = 12 \cos(\theta) = 10.91$, or from Pythagoras' theorem if you prefer $v_y = \sqrt{\|\vec{v}\|^2 - v_x^2} = \sqrt{12^2 - 5^2} = 10.91$.

Discussion

We did a lot of hands-on activities with vectors in this section and skipped over some of the theoretical details. Now that you've been exposed to the practical side of vector calculations, it's worth clarifying certain points that we glossed over.

Vectors vs. points

We used the notation $\mathbb{R}^2$ to describe two kinds of math objects: the set of points in the Cartesian plane and the set of vectors in a two-

dimensional space. The point $P = (P_x, P_y)$ and the vector $\vec{v} = (v_x, v_y)$ are both represented by pairs of real numbers, so we use the notation $P \in \mathbb{R}^2$ and $\vec{v} \in \mathbb{R}^2$ to describe them. This means that a pair of numbers $(3, 2) \in \mathbb{R}^2$ could represent the *coordinates* of a point, or the *components* of a vector, depending on the context.

Let's take a moment to review the definitions of points and vectors and clarify the types of operations we can perform on them:

- **Space of points** $\mathbb{R}^2$: the set of points $P = (P_x, P_y)$ corresponds to locations in the Cartesian plane. The point $P = (P_x, P_y)$ corresponds to the geometric instructions: "Starting at the origin $(0, 0)$, move P_x units along the x -axis and P_y units along the y -axis." The distance between points P and Q is denoted $d(P, Q)$.
- **Vector space** $\mathbb{R}^2$: the set of vectors $\vec{v} = (v_x, v_y)$ describes displacements in the Cartesian plane. The vector $\vec{v} = (v_x, v_y)$ corresponds to the instructions: "Starting anywhere, move v_x units along the x -axis and v_y units along the y -axis." Vectors can be combined and manipulated using the vector algebra operations $\vec{u} + \vec{v}$, $\vec{u} - \vec{v}$, $\alpha\vec{u}$, $\vec{u} \cdot \vec{v}$, and $\|\vec{v}\|$.

Note the geometric instructions for points and vectors are very similar; the only difference is the starting point. The coordinates of a point (P_x, P_y) specify a *fixed position* relative to the origin $(0, 0)$, while the components of a vector (v_x, v_y) describe a *relative displacement* that can have any starting point.

Let's look at some examples of calculations that combine points and vectors. Consider the points P and Q in the Cartesian plane, and the displacement vector $\vec{v}_{PQ}$ between them. The displacement vector $\vec{v}_{PQ}$ gives the "move instructions" for getting from point P to point Q and is defined by the equation:

$$\vec{v}_{PQ} = Q - P.$$

This equation says that subtracting two points produces a vector, which make sense if you think about it—the "difference" between two points is a displacement vector.

We can use the displacement vector $\vec{v}_{PQ}$ in calculations like this:

$$P + \vec{v}_{PQ} = P + (Q - P) = Q.$$

In words, this calculation shows that "Starting at the point P and moving by $\vec{v}_{PQ}$ brings us to the point Q ."

The above equations use addition and subtraction operations between a mix of points and vectors. This is rather unusual: normally we only use operations like "+" and "-" between math objects of the same kind. In this case, we're allowed to mix points and vectors because they both describe "move instructions" of the same kind.

Let's keep going. What other useful calculations can we do by combining points and vectors? Suppose we wanted to find the midpoint M that lies exactly in the middle between points P and Q . We can find the midpoint M using the displacement vector $\vec{v}_{PQ}$ and some basic vector algebra. If starting from P and moving by $\vec{v}_{PQ}$ brings us all the way to the point Q , then starting from P and moving by $\frac{1}{2}\vec{v}_{PQ}$ will bring us to the midpoint: $M = P + \frac{1}{2}\vec{v}_{PQ}$.

The mathematical bridge between points and vectors allows us to use vector techniques to solve geometry problems. By learning to describe geometric objects like points, lines, and circles using vectors, we can do complicated geometry calculations using simple algebraic manipulations like vector operations. This exemplifies a general pattern in mathematics: applying techniques developed in one domain to solve problems in another domain.

Example You come to class one day and there's a surprise quiz that asks you to write the formula for the distance $d(P, Q)$ between two points $P = (P_x, P_y)$ and $Q = (Q_x, Q_y)$. You don't remember ever learning about such a formula and feel caught off guard. How can the teacher ask for a formula they haven't covered in class yet? This seems totally unfair!

After a minute of stressing out, you take a deep breath, come back to your senses, and resolve to give this problem a shot. You start by sketching a coordinate system, placing points P and Q in it, and drawing the line that connects the two points. What is the formula that describes the length of this line?

The line from P to Q looks like the hypotenuse of a triangle, which makes you think that trigonometry could somehow be used to find the answer. Unfortunately, trying to remember the trigonometry formulas has only the effect of increasing your math anxiety. You take this as a sign that you should look for other options. In math, it's important to trust your gut instincts.

By a fortunate coincidence, you were recently reading about the connection between points and vectors, and specifically about the displacement vector $\vec{v}_{PQ} = Q - P$. The line in your sketch represents the vector $\vec{v}_{PQ}$. You realize that the distance between the points P and Q is the same as the length of the vector $\vec{v}_{PQ}$. You remember the formula for the length of a vector $\vec{v}$ is $\|\vec{v}\| = \sqrt{v_x^2 + v_y^2}$ and you know the formula for the displacement vector is $\vec{v}_{PQ} = (Q_x - P_x, Q_y - P_y)$, so you combine these formulas to obtain the answer: $d(P, Q) = \|\vec{v}_{PQ}\| = \sqrt{(Q_x - P_x)^2 + (Q_y - P_y)^2}$. One more win for the "don't worry and try it" strategy for solving math problems!

Vectors in three dimensions

A three-dimensional coordinate system consists of three axes: the x -axis, the y -axis, and the z -axis. The three axes point in perpendicular directions to each other, as illustrated in Figure 1.55. Look around you and find a corner of the room you're in where two walls and the floor meet. The x -axis and the y -axis are the edges where the floor meets the walls. The vertical edge where the two walls meet represents the z -axis.




Figure 1.55: A three-dimensional coordinate system with x , y , and z axes.

The vector $\vec{v} = (v_x, v_y, v_z) \in \mathbb{R}^3$ describes the following displacement instructions: “Move v_x units in the direction of the x -axis, then move v_y along the y -axis, and finally move v_z in the direction of the z -axis.” In three dimensions, there are three unit vectors that describe unit steps in the direction of each of the axes:

$$\hat{i} = (1, 0, 0), \quad \hat{j} = (0, 1, 0) \quad \text{and} \quad \hat{k} = (0, 0, 1).$$

We can therefore describe the vector $\vec{v} = (v_x, v_y, v_z)$ in terms of unit vectors as $\vec{v} = v_x\hat{i} + v_y\hat{j} + v_z\hat{k}$.

High-dimensional vectors

The most common types of vectors you'll encounter in math and physics are two-dimensional and three-dimensional vectors. In other fields of science like genetics and machine learning, it's common to see vectors with many more dimensions. For example, in machine learning we often represent “rich data” like images, videos, and text as vectors with thousands of dimensions.

An example of an n -dimensional vector is

$$\vec{v} = (v_1, v_2, \dots, v_n) \in \mathbb{R}^n.$$

The vector algebra operations you learned in this section also apply to these high-dimensional vectors.

Vectors and vector coordinates

One final point we need to clarify is the difference between real-world vector quantities like the velocity of a tennis ball $\vec{v}$ and its mathematical representation as a coordinate vector (v_x, v_y, v_z) . If you know the coordinate vector (v_x, v_y, v_z) then you know what the real-world velocity is, right? Not quite.

Let's say you're doing a physics research project on tennis serves. You define an xyz -coordinate system for the tennis court, which allows you to represent the ball's velocity $\vec{v}$ as a triple of components (v_x, v_y, v_z) interpreted as: "The ball is moving with velocity v_x units in the x -direction, v_y units in the y -direction, and v_z units in the z -direction."

Suppose you want to describe the velocity vector $\vec{v}$ to a fellow physicist via text message. Referring to your sheet of calculations, you find the values $\vec{v} = (60, 3, -2)$, which you know were measured in metres per second. You send this message:

The velocity is $(60, 3, -2)$ measured in metres per second.

A few minutes later the following reply comes back:

Wait whaaat? What coordinate system are you using?

Indeed the information you sent is incomplete. Vector components depend on the coordinate system in which the vectors are represented. The triple of numbers $(60, 3, -2)$ only makes sense once you know the directions of the axes in the xyz -coordinate system. Realizing your mistake, you send a text with all the required information:

Using the coordinate system centred at the south post of the net, with the x -axis pointing east along the court, the y -axis pointing north along the net, and the z -axis pointing up, the velocity is $(60, 3, -2)$ in metres per second.

A few seconds later, you get the reply:

OK got it now. Thx!

This hypothetical situation illustrates the importance of the coordinate systems for describing vectors. If you don't know what the coordinate system is, knowing the coordinates (v_x, v_y, v_z) doesn't tell you much. Only when you know the directions of the unit vectors $\hat{i}$, $\hat{j}$, and $\hat{k}$ can you interpret the instructions $\vec{v} = v_x\hat{i} + v_y\hat{j} + v_z\hat{k}$.

It turns out, using the xyz -coordinate system with the three vectors $\{\hat{i}, \hat{j}, \hat{k}\}$ is just one of many possible ways we can represent vectors. We can represent a vector $\vec{v}$ as coordinates (v_1, v_2, v_3) with respect to any *basis* $\{\hat{e}_1, \hat{e}_2, \hat{e}_3\}$ using the expression $\vec{v} = v_1\hat{e}_1 + v_2\hat{e}_2 +$

$v_3\hat{e}_3$, which corresponds to the instructions: “Move v_1 units in the direction of $\hat{e}_1$, move v_2 units in the direction of $\hat{e}_2$, and move v_3 units in the direction of $\hat{e}_3$.”

What’s a basis, you ask? I’m glad you asked, because this is the subject of the next section.

Basis

One of the most important concepts in the study of vectors is the concept of a *basis*. Consider the three-dimensional vector space $\mathbb{R}^3$. A *basis* for $\mathbb{R}^3$ is a set of vectors $\{\hat{e}_1, \hat{e}_2, \hat{e}_3\}$ that can be used as a coordinate system for $\mathbb{R}^3$. If the set of vectors $\{\hat{e}_1, \hat{e}_2, \hat{e}_3\}$ is a basis, then you can *represent* any vector $\vec{v} \in \mathbb{R}^3$ as coordinates (v_1, v_2, v_3) with respect to that basis:

$$\vec{v} = v_1\hat{e}_1 + v_2\hat{e}_2 + v_3\hat{e}_3.$$

The vector $\vec{v}$ is obtained by measuring out a distance v_1 in the $\hat{e}_1$ direction, a distance v_2 in the $\hat{e}_2$ direction, and a distance v_3 in the $\hat{e}_3$ direction.

You are already familiar with the *standard* basis $\{\hat{i}, \hat{j}, \hat{k}\}$, which is associated with the *xyz*-coordinate system. You know that any vector $\vec{v} \in \mathbb{R}^3$ can be expressed as a triple (v_x, v_y, v_z) with respect to the basis $\{\hat{i}, \hat{j}, \hat{k}\}$ through the formula $\vec{v} = v_x\hat{i} + v_y\hat{j} + v_z\hat{k}$. The whole point of this section is to let you know that other bases (coordinate systems) exist, and to get you into the habit of asking, “With respect to which coordinate system?” every time you see a coordinate vector (a, b, c) .

An analogy

Let’s start with a simple example of a basis. If you look at the HTML source code behind any web page, you’re sure to find at least one mention of the colour stylesheet directive such as `color:#336699;`. The numbers should be interpreted as a triple of values (33,66,99), each value describing the amount of red, green, and blue needed to create a given colour. Let us call the colour described by the triple (33,66,99) CoolBlue. This convention for colour representation is called the RGB colour model and we can think of it as the *RGB basis*. A basis is a set of elements that can be combined together to express something more complicated. In our case, the **R**, **G**, and **B** elements are pure colours that can create any colour when mixed appropriately. Schematically, we can write this mixing idea as

$$\text{CoolBlue} = (33,66,99)_{RGB} = 33\mathbf{R} + 66\mathbf{G} + 99\mathbf{B},$$

where the *components* determine the strength of each colour. To create the colour, we combine its components as symbolized by the + operation.

The cyan, magenta, and yellow (CMY) colour model is another basis for representing colours. To express the “cool blue” colour in the CMY basis, you will need the following components:

$$(33, 66, 99)_{RGB} = \text{CoolBlue} = (222, 189, 156)_{CMY} = 222\mathbf{C} + 189\mathbf{M} + 156\mathbf{Y}.$$

The *same* colour CoolBlue is represented by a *different* set of components when the CMY colour basis is used.

Note that a triple of components by itself doesn’t mean anything unless we know the basis being used. For example, if we were to interpret the triple of components (33, 66, 99) with respect to the CMY basis, we would obtain a completely different colour, which would not be cool at all.

A basis is required to convert mathematical objects like the triple (a, b, c) into real-world ideas like colours. As exemplified above, to avoid any ambiguity we can use a subscript after the bracket to indicate the basis associated with each triple of components. Writing $(222, 189, 156)_{CMY}$ and $(33, 66, 99)_{RGB}$ clarifies which basis to use for each triple of components.

Discussion

It would be hard to over-emphasize the importance of the basis—the coordinate system you use to describe vectors. The choice of coordinate system is the bridge between real-world vector quantities and their mathematical representation in terms of components. Every time you start a new problem that involves vector calculations, the first thing you should do is choose the coordinate system you want to use, and indicate it clearly in the diagram.

Using a non-standard coordinate system can sometimes simplify the equations you have to solve. For example, let’s say we want to study the motion of a block sliding down an incline with velocity $\vec{v}$, as illustrated in Figure 1.56. Using the standard xy -basis, the velocity vector is represented as $(v \cos \theta, -v \sin \theta)_{xy}$, which has components in both the x - and y -directions and requires using trigonometric functions. If instead you use the non-standard $x'y'$ -basis, the components of the velocity will be $(v, 0)_{x'y'}$. Note the velocity only has a component along the x' -direction, which will simplify all subsequent calculations.




Figure 1.56: The vector $\vec{v}$ is described by the coordinates $(v \cos \theta, -v \sin \theta)_{xy}$ with respect to the standard basis xy . The same vector $\vec{v}$ is described by the coordinates $(v, 0)_{x'y'}$ with respect to the “tilted” basis $x'y'$.

Links

[Vectors and vector operations explained by 3Blue1Brown]

https://www.youtube.com/watch?v=fNk_zzaMoSs

[More vector illustrations and definitions from Wikipedia]

https://en.wikipedia.org/wiki/Euclidean_vector

Exercises

E1.20 Given the vectors $\vec{v}_1 = (2, 1)$, $\vec{v}_2 = (2, -1)$, and $\vec{v}_3 = (3, 3)$, calculate the following expressions:

a) $\vec{v}_1 + \vec{v}_2$

b) $\vec{v}_2 - 2\vec{v}_1$

c) $\vec{v}_1 + \vec{v}_2 + \vec{v}_3$

E1.21 Express the following vectors as components:

a) $\vec{v}_1 = 10\angle 30^\circ$

b) $\vec{v}_2 = 12\angle -90^\circ$

c) $\vec{v}_3 = 3\angle 170^\circ$

E1.22 Express the following vectors in length-and-direction notation:

a) $\vec{u}_1 = (4, 0)$

b) $\vec{u}_2 = (1, 1)$

c) $\vec{u}_3 = (-1, 3)$

1.14 Complex numbers

By now, you’ve heard about complex numbers $\mathbb{C}$. The word “complex” is an intimidating word. Surely it must be a complex task to learn about the complex numbers. That may be true in general, but it helps if you know about vectors. Complex numbers are similar to two-dimensional vectors $\vec{v} \in \mathbb{R}^2$. We add and subtract complex numbers like vectors. Complex numbers also have components, length, and “direction.” If you understand vectors, you will understand complex numbers at almost no additional mental cost.

We’ll begin with a practical problem.

Example

Suppose you're asked to solve the following quadratic equation:

$$x^2 + 1 = 0.$$

You're looking for a number x , such that $x^2 = -1$. If you are only allowed to give real answers (the set of real numbers is denoted $\mathbb{R}$), then there is no answer to this question. In other words, this equation has no solutions. Graphically speaking, this is because the quadratic function $f(x) = x^2 + 1$ does not cross the x -axis.

However, we're not taking no for an answer! If we insist on solving for x in the equation $x^2 + 1 = 0$, we can imagine a new number i that satisfies $i^2 = -1$. We call i the unit imaginary number. The solutions to the equation are therefore $x_1 = i$ and $x_2 = -i$. There are two solutions because the equation is quadratic. We can check that $i^2 + 1 = -1 + 1 = 0$ and also $(-i)^2 + 1 = (-1)^2 i^2 + 1 = i^2 + 1 = 0$.

Thus, while the equation $x^2 + 1 = 0$ has no real solutions, it *does* have solutions if we allow the answers to be imaginary numbers.

Definitions

Complex numbers have a real part and an imaginary part:

- i : the unit imaginary number $i = \sqrt{-1}$ or $i^2 = -1$
- bi : an imaginary number that is equal to b times i
- $\mathbb{R}$: the set of real numbers
- $\mathbb{C}$: the set of complex numbers $\mathbb{C} = \{a + bi \mid a, b \in \mathbb{R}\}$
- $z = a + bi$: a complex number
 - ▷ $\operatorname{Re}\{z\} = a$: the real part of z
 - ▷ $\operatorname{Im}\{z\} = b$: the imaginary part of z
- $\bar{z}$: the *complex conjugate* of z . If $z = a + bi$, then $\bar{z} = a - bi$.

The polar representation of complex numbers:

- $z = |z| \angle \varphi_z = |z| \cos \varphi_z + i|z| \sin \varphi_z$
- $|z| = \sqrt{z\bar{z}} = \sqrt{a^2 + b^2}$: the *magnitude* of $z = a + bi$
- $\varphi_z = \tan^{-1}(b/a)$: the *phase* or *argument* of $z = a + bi$
- $\operatorname{Re}\{z\} = |z| \cos \varphi_z$
- $\operatorname{Im}\{z\} = |z| \sin \varphi_z$

Formulas

Addition and subtraction

Just as we performed the addition of vectors component by component, we perform addition on complex numbers by adding the real parts together and adding the imaginary parts together:

$$(a + bi) + (c + di) = (a + c) + (b + d)i.$$

Polar representation

We can give a geometric interpretation of the complex numbers by extending the real number line into a two-dimensional plane called the *complex plane*. The horizontal axis in the complex plane measures the *real* part of the number. The vertical axis measures the *imaginary* part. Complex numbers are points in the complex plane.




Figure 1.57: The complex number $z = a + bi$ corresponds to the point with coordinates (a, b) in the complex plane.

It is possible to represent any complex number $z = a + bi$ in terms of its *magnitude* and its *phase*:

$$z = |z| \angle \varphi_z = \underbrace{|z| \cos \varphi_z}_a + \underbrace{|z| \sin \varphi_z}_b i.$$

The *magnitude* (or *absolute value*) of a complex number $z = a + bi$ is

$$|z| = \sqrt{a^2 + b^2}.$$

This corresponds to the *length* of the vector that represents the complex number in the complex plane. The formula is obtained by using Pythagoras' theorem.

The *phase*, also known as the *argument* of the complex number $z = a + bi$ is given by the formula

$$\varphi_z = \arg z = \text{atan2}(b, a) = \begin{cases} \tan^{-1}\left(\frac{b}{a}\right) & \text{if } a > 0, \\ \pi + \tan^{-1}\left(\frac{b}{a}\right) & \text{if } a < 0, \\ \frac{\pi}{2} & \text{if } a = 0 \text{ and } b > 0, \\ -\frac{\pi}{2} & \text{if } a = 0 \text{ and } b < 0. \end{cases}$$

The phase corresponds to the angle that z forms with the real axis.

We previously saw this complicated-looking formula with four cases when we talked about vectors (Section 1.13). When a certain formula comes up two times in a math book, this should tell you the author *really* wants you to know it. Seriously, do me a favour and revisit the exercise E1.22 (page 96).

Computer algebra systems provide the two-input inverse tangent function `atan2`, which is the easiest way to calculate the phase φ_z for the complex number $z = a + bi$. The function `atan2` handles all four cases automatically and always computes the correct phase φ_z . Try doing some calculations with `atan2` using the computer algebra systems at <https://live.sympy.org>.

* * *

In addition to the vector-like operations we can perform on complex numbers, like computing their magnitude and phase, we can also perform other operations on complex numbers that are not defined for vectors. The set of complex numbers $\mathbb{C}$ is a *field*. This means, in addition to the addition and subtraction operations, we can also perform multiplication and division with complex numbers.

Multiplication

The product of two complex numbers is computed using the usual rules of algebra:

$$\begin{aligned} (a + bi)(c + di) &= a(c + di) + bi(c + di) \\ &= ac + adi + bci + bdi^2 \\ &= (ac - bd) + (ad + bc)i. \end{aligned}$$

In the polar representation, the product formula is

$$(p\angle\phi)(q\angle\psi) = pq\angle(\phi + \psi).$$

To multiply two complex numbers, multiply their magnitudes and add their phases.

Example Verify that $z\bar{z} = a^2 + b^2 = |z|^2$.

Division

Let's look at the procedure for dividing complex numbers:

$$\frac{(a + bi)}{(c + di)} = \frac{(a + bi)}{(c + di)} \frac{(c - di)}{(c - di)} = (a + bi) \frac{(c - di)}{(c^2 + d^2)} = (a + bi) \frac{\overline{c + di}}{|c + di|^2}.$$

In other words, to divide the number z by the complex number s , compute $\bar{s}$ and $|s|^2 = s\bar{s}$ and then use

$$z/s = z \frac{\bar{s}}{|s|^2}.$$

You can think of $\frac{\bar{s}}{|s|^2}$ as being equivalent to s^{-1} .

Cardano's example One of the earliest examples of reasoning involving complex numbers was given by Gerolamo Cardano in his 1545 book *Ars Magna*. Cardano wrote, "If someone says to you, divide 10 into two parts, one of which multiplied into the other shall produce 40, it is evident that this case or question is impossible." We want to find numbers x_1 and x_2 such that $x_1 + x_2 = 10$ and $x_1 x_2 = 40$. This sounds kind of impossible. Or is it?

"Nevertheless," Cardano said, "we shall solve it in this fashion:

$$x_1 = 5 + \sqrt{15}i \text{ and } x_2 = 5 - \sqrt{15}i."$$

When you add $x_1 + x_2$ you obtain 10. When you multiply the two numbers the answer is

$$\begin{aligned} x_1 x_2 &= (5 + \sqrt{15}i)(5 - \sqrt{15}i) \\ &= 25 - 5\sqrt{15}i + 5\sqrt{15}i - \sqrt{15}^2 i^2 = 25 + 15 = 40. \end{aligned}$$

Hence $5 + \sqrt{15}i$ and $5 - \sqrt{15}i$ are two numbers whose sum is 10 and whose product is 40.

Example 2 Let's compute the product of -1 and i . The answer is obviously $-i$, but let's look at this simple calculation geometrically. The polar representation of the number i is $1\angle\frac{\pi}{2}$. Multiplication of any complex number $z = |z|\angle\varphi_z$ by i corresponds to adding $\frac{\pi}{2}$ to the phase of the number:

$$zi = (|z|\angle\varphi_z)(1\angle\frac{\pi}{2}) = (|z|\cdot 1)\angle(\varphi_z + \frac{\pi}{2}) = |z|\angle(\varphi_z + \frac{\pi}{2}).$$

In other words, multiplication by i is equivalent to applying a $\frac{\pi}{2}$ (90°) counterclockwise rotation in the complex plane. We can therefore interpret the answer $(-1)(i) = -i$ as the complex number $-1 = 1\angle\pi$ experiencing a $\frac{\pi}{2}$ rotation to arrive at $1\angle(\pi + \frac{\pi}{2}) = 1\angle\frac{3\pi}{2} = -i$.

Example 3 Find the polar representation of $z = -3 - i$ and compute z^6 . Let's denote the polar representation of z by $z = r\angle\varphi$ as shown in Figure 1.58. We find $r = \sqrt{3^2 + 1^2} = \sqrt{10}$ and $\varphi = \tan^{-1}(\frac{1}{3}) + \pi = 0.322 + \pi$. Using the polar representation, we can easily compute z^6 :

$$z^6 = r^6\angle(6\varphi) = (\sqrt{10})^6\angle 6(0.322 + \pi) = 10^3\angle 1.932 + 6\pi = 10^3\angle 1.932.$$

Note we can ignore multiples of 2π in the phase. We thus find the value of z^6 is $1000 \cos(1.932) + 1000 \sin(1.932)i = -353.4 + 935.5i$.




Figure 1.58: The complex number $z = 3 - i$ has magnitude $r = \sqrt{10}$ and phase $\varphi = 0.322 + \pi = 3.463$.

Fundamental theorem of algebra

The fundamental theorem of algebra states that any polynomial of degree n , $P(x) = a_n x^n + \cdots + a_2 x^2 + a_1 x + a_0$, can be written as

$$P(x) = a_n(x - z_1)(x - z_2) \cdots (x - z_n),$$

where $z_i \in \mathbb{C}$ are the polynomial's roots. In other words, the equation $P(x) = 0$ has n solutions: the complex numbers $z_1, z_2, \dots, z_n$. Before today, you might have said the equation $x^2 + 1 = 0$ has no solutions. Now you know its solutions are the complex numbers $z_1 = i$ and $z_2 = -i$.

The theorem is “fundamental” because it tells us we’ll never need to invent numbers “fancier” than the complex numbers to solve polynomial equations. To understand why this is important, recall that each set of numbers is associated with a different class of equations. Figure 1.2 on page 12 shows the nested containment structure of the number sets $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{C}$. The natural numbers $\mathbb{N}$ appear as solutions of the equation $m + n = x$, where m and n are natural numbers (denoted $m, n \in \mathbb{N}$). The integers $\mathbb{Z}$ are

the solutions to equations of the form $x + m = n$, where $m, n \in \mathbb{N}$. The rational numbers $\mathbb{Q}$ are necessary to solve for x in $mx = n$, with $m, n \in \mathbb{Z}$. To find the solutions of $x^2 = 2$, we need the real numbers $\mathbb{R}$. And in this section, we learned that the solutions to the equation $x^2 = -1$ are complex numbers $\mathbb{C}$. At this point you might be wondering if you're attending some sort of math party, where mathematicians write down complicated equations and—just for the fun of it—invent new sets of numbers to describe the solutions to these equations. Can this process continue indefinitely?

Nope. The party ends with $\mathbb{C}$. The fundamental theorem of algebra guarantees that any polynomial equation you could come up with—no matter how complicated it is—has solutions that are complex numbers $\mathbb{C}$.

Euler's formula

It turns out the exponential function is related to the functions sine and cosine. Lo and behold, we have *Euler's formula*:

$$e^{i\theta} = \cos \theta + i \sin \theta.$$

Inputting an imaginary number to the exponential function outputs a complex number that contains both $\cos$ and $\sin$. Euler's formula gives us an alternate notation for the polar representation of complex numbers: $z = |z|\angle\varphi_z = |z|e^{i\varphi_z}$.

If you want to impress your friends with your math knowledge, plug $\theta = \pi$ into the above equation to find

$$e^{i\pi} = \cos(\pi) + i \sin(\pi) = -1,$$

which can be rearranged to obtain the equation $e^{i\pi} + 1 = 0$. The equation $e^{i\pi} + 1 = 0$ is called *Euler's identity*, and it shows a relationship between the five most important numbers in all of mathematics: Euler's number $e = 2.71828\dots$, $\pi = 3.14159\dots$, the imaginary number i , 1, and zero. It's kind of cool to see all these important numbers reunited in one equation, don't you agree?

One way to understand the equation $e^{i\pi} + 1 = 0$, is to think of $e^{i\pi}$ as the polar representation of the complex number $z = 1e^{i\pi} = 1\angle\pi$, which is the same as 1 rotated counterclockwise by π radians (180°) in the complex plane. We know $e^{i\pi} = 1\angle\pi = -1$ and so $e^{i\pi} + 1 = 0$.

De Moivre's formula

By replacing θ in Euler's formula with $n\theta$, we obtain de Moivre's formula:

$$(\cos \theta + i \sin \theta)^n = \cos n\theta + i \sin n\theta.$$

De Moivre's formula makes sense if you think of the complex number $z = e^{i\theta} = \cos \theta + i \sin \theta$, raised to the n^{th} power:

$$(\cos \theta + i \sin \theta)^n = z^n = (e^{i\theta})^n = e^{in\theta} = \cos n\theta + i \sin n\theta.$$

Setting $n = 2$ in de Moivre's formula, we can derive the double angle formulas (page 78) as the real and imaginary parts of the following equation:

$$(\cos^2 \theta - \sin^2 \theta) + (2 \sin \theta \cos \theta)i = \cos(2\theta) + \sin(2\theta)i.$$

Links

[Intuitive proof of the fundamental theorem of algebra]

<https://www.youtube.com/watch?v=shEk8sz1o0w>

1.15 Solving systems of linear equations

Solving equations with one unknown—like $2x + 4 = 7x$, for instance—requires manipulating both sides of the equation until the unknown variable is *isolated* on one side. For this instance, we can subtract $2x$ from both sides of the equation to obtain $4 = 5x$, which simplifies to $x = \frac{4}{5}$.

What about the case when you are given *two* equations and must solve for *two* unknowns? For example,

$$\begin{aligned}x + 2y &= 5, \\ 3x + 9y &= 21.\end{aligned}$$

Can you find values of x and y that satisfy both equations?

Concepts

- x, y : the two unknowns in the equations
- $eq1, eq2$: a system of two equations that must be solved *simultaneously*. These equations will look like

$$\begin{aligned}a_1x + b_1y &= c_1, \\ a_2x + b_2y &= c_2,\end{aligned}$$

where as , bs , and cs are given constants.

Principles

If you have n equations and n unknowns, you can solve the equations *simultaneously* and find the values of the unknowns. There are several different approaches for solving equations simultaneously. We'll show three of these approaches for the case $n = 2$.

Solution techniques

When solving for two unknowns in two equations, the best approach is to *eliminate* one of the variables from the equations. By combining the two equations appropriately, we can simplify the problem to the problem of finding one unknown in one equation.

Solving by substitution

We want to solve the following system of equations:

$$\begin{aligned}x + 2y &= 5, \\ 3x + 9y &= 21.\end{aligned}$$

We can isolate x in the first equation to obtain

$$\begin{aligned}x &= 5 - 2y, \\ 3x + 9y &= 21.\end{aligned}$$

Now *substitute* the expression for x from the top equation into the bottom equation:

$$3(5 - 2y) + 9y = 21.$$

We just eliminated one of the unknowns by substitution. Continuing, we expand the bracket to find

$$15 - 6y + 9y = 21,$$

or

$$3y = 6.$$

We find $y = 2$, but what is x ? Easy. To solve for x , plug the value $y = 2$ into any of the equations we started from. Using the equation $x = 5 - 2y$, we find $x = 5 - 2(2) = 1$.

Solving by subtraction

Let's now look at another way to solve the same system of equations:

$$\begin{aligned}x + 2y &= 5, \\ 3x + 9y &= 21.\end{aligned}$$

Observe that any equation will remain true if we multiply the whole equation by some constant. For example, we can multiply the first equation by 3 to obtain an equivalent set of equations:

$$\begin{aligned}3x + 6y &= 15, \\ 3x + 9y &= 21.\end{aligned}$$

Why did I pick 3 as the multiplier? By choosing this constant, the x terms in both equations now have the same coefficient.

Subtracting two true equations yields another true equation. Let's subtract the top equation from the bottom one:

$$\cancel{3x} - \cancel{3x} + 9y - 6y = 21 - 15 \Rightarrow 3y = 6.$$

The $3x$ terms cancel. This subtraction eliminates the variable x because we multiplied the first equation by 3. We find $y = 2$. To find x , substitute $y = 2$ into one of the original equations:

$$x + 2(2) = 5,$$

from which we deduce that $x = 1$.

Solving by equating

There is a third way to solve the system of equations

$$\begin{aligned}x + 2y &= 5, \\ 3x + 9y &= 21.\end{aligned}$$

We can isolate x in both equations by moving all other variables and constants to the right-hand sides of the equations:

$$\begin{aligned}x &= 5 - 2y, \\ x &= \frac{1}{3}(21 - 9y) = 7 - 3y.\end{aligned}$$

Though the variable x is unknown to us, we know two facts about it: x is equal to $5 - 2y$ and x is equal to $7 - 3y$. Therefore, we can eliminate x by equating the right-hand sides of the equations:

$$5 - 2y = 7 - 3y.$$

We solve for y by adding $3y$ to both sides and subtracting 5 from both sides. We find $y = 2$ then plug this value into the equation $x = 5 - 2y$ to find x . The solutions are $x = 1$ and $y = 2$.

Discussion

The repeated use of the three algebraic techniques presented in this section will allow you to solve any system of n linear equations in n unknowns. Each time you eliminate one variable using a substitution, a subtraction, or an elimination by equating, you're simplifying the problem to a problem of finding $(n - 1)$ unknowns in a system of $(n - 1)$ equations. In Chapter 3 we'll develop a more advanced, systematic approach for solving systems of linear equations.

Geometric solution

Solving a system of two linear equations in two unknowns can be understood geometrically as finding the point of intersection between two lines in the Cartesian plane. In this section we'll explore this correspondence between algebra and geometry to develop yet another way of solving systems of linear equations.

The algebraic equation $ax + by = c$ containing the unknowns x and y can be interpreted as a *constraint* equation on the set of possible values for the variables x and y . We can visualize this constraint geometrically by considering the coordinate pairs (x, y) that lie in the Cartesian plane. Recall that every point in the Cartesian plane can be represented as a coordinate pair (x, y) , where x and y are the coordinates of the point.

Figure 1.59 shows the geometrical representation of three equations. The line ℓ_a corresponds to the set of points (x, y) that satisfy the equation $x = 1$, the line ℓ_b is the set of points (x, y) that satisfy the equation $y = 2$, and the line ℓ_c corresponds to the set of points that satisfy $x + 2y = 2$.




Figure 1.59: Graphical representations of three linear equations.

You can convince yourself that the geometric lines shown in Figure 1.59 are equivalent to the algebraic equations by considering individual points (x, y) in the plane. For example, the points $(1, 0)$, $(1, 1)$, and $(1, 2)$ are all part of the line ℓ_a since they satisfy the equation $x = 1$. For the line ℓ_c , you can verify that the line's x -intercept $(2, 0)$ and its y -intercept $(0, 1)$ both satisfy the equation $x + 2y = 2$.

The Cartesian plane as a whole corresponds to the set $\mathbb{R}^2$, which describes all possible pairs of coordinates. To understand the equivalence between the algebraic equation $ax + by = c$ and the line ℓ in the Cartesian plane, we can use the following precise math notation:

$$\ell : \{(x, y) \in \mathbb{R}^2 \mid ax + by = c\}.$$

In words, this means that the line ℓ is defined as the subset of the pairs of real numbers (x, y) that satisfy the equation $ax + by = c$.

Figure 1.60 shows the graphical representation of the line ℓ .

You don't have to take my word for it, though! Think about it and convince yourself that all points on the line ℓ shown in Figure 1.60 satisfy the equation $ax + by = c$. For example, you can check that the x -intercept $(\frac{c}{a}, 0)$ and the y -intercept $(0, \frac{c}{b})$ satisfy the equation $ax + by = c$.




Figure 1.60: Graphical representation of the equation $ax + by = c$.

Solving the system of two equations

$$a_1x + b_1y = c_1,$$

$$a_2x + b_2y = c_2,$$

corresponds to finding the intersection of the lines ℓ_1 and ℓ_2 that represent each equation. The pair (x, y) that satisfies both algebraic equations simultaneously is equivalent to the point (x, y) that is the intersection of lines ℓ_1 and ℓ_2 , as illustrated in Figure 1.61.




Figure 1.61: The point (x, y) that lies at the intersection of lines ℓ_1 and ℓ_2 .

Example Let's see how we can use the geometric interpretation to solve the system of equations

$$\begin{aligned}x + 2y &= 5, \\ 3x + 9y &= 21.\end{aligned}$$

We've already seen three different *algebraic* techniques for finding the solution to this system of equations; now let's see a *geometric* approach for finding the solution. I'm not kidding you, we're going to solve the exact same system of equations a fourth time!

The first step is to draw the lines that correspond to each of the equations using pen and paper or a graphing calculator. The second step is to find the coordinates of the point where the two lines intersect as shown in Figure 1.62. The point $(1, 2)$ that lies on both lines ℓ_1 and ℓ_2 corresponds to the x and y values that satisfy both equations simultaneously.




Figure 1.62: The line ℓ_1 with equations $x + 2y = 5$ intersects the line ℓ_2 with equation $3x + 9y = 21$ at the point $(1, 2)$.

Visit the webpage at www.desmos.com/calculator/exikik615f to play with an interactive version of the graphs shown in Figure 1.62. Try changing the equations and see how the graphs change.

Exercises

E1.23 Plot the lines ℓ_a , ℓ_b , and ℓ_c shown in Figure 1.59 (page 106) using the Desmos graphing calculator. Use the graphical representation of these lines to find: **a)** the intersection of lines ℓ_c and ℓ_a , **b)** the intersection of ℓ_a and ℓ_b , and **c)** the intersection of lines ℓ_b and ℓ_c .

E1.24 Solve the system of equations simultaneously for x and y :

$$\begin{aligned}2x + 4y &= 16, \\ 5x - y &= 7.\end{aligned}$$

E1.25 Solve the system of equations for the unknowns x , y , and z :

$$2x + y - 4z = 28,$$

$$x + y + z = 8,$$

$$2x - y - 6z = 22.$$

E1.26 Solve for p and q given the equations $p + q = 10$ and $p - q = 4$.

1.16 Set notation

A *set* is the mathematically precise notion for describing a group of objects. You don't need to know about sets to perform simple math; but more advanced topics require an understanding of what sets are and how to denote set membership, set operations, and set containment relations. This section introduces all the relevant concepts.

Definitions

- *set*: a collection of mathematical objects
- S, T : the usual variable names for sets
- $s \in S$: this statement is read " s is an element of S " or " s is in S "
- $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}$: some important number sets: the naturals, the integers, the rationals, and the real numbers, respectively.
- $\emptyset$: the *empty set* is a set that contains no elements
- $\{ \dots \}$: the curly brackets are used to define sets, and the expression inside the curly brackets describes the set contents.

Set operations:

- $S \cup T$: the *union* of two sets. The union of S and T corresponds to the elements in either S or T .
- $S \cap T$: the *intersection* of the two sets. The intersection of S and T corresponds to the elements that are in both S and T .
- $S \setminus T$: *set difference* or *set minus*. The set difference $S \setminus T$ corresponds to the elements of S that are not in T .

Set relations:

- $\subset$: is a strict subset of
- $\subseteq$: is a subset of or equal to

Here is a list of special mathematical shorthand symbols and their corresponding meanings:

- $\in$: element of
- $\notin$: not an element of
- $\forall$: for all
- $\exists$: there exists
- $\nexists$: there doesn't exist
- $|$: such that

These symbols are used in math proofs because they allow us to express complex mathematical arguments succinctly and precisely.

An *interval* is a subset of the real line. We denote an interval by specifying its endpoints and surrounding them with either square brackets “[” or round brackets “(” to indicate whether or not the corresponding endpoint is included in the interval.

- $[a, b]$: the *closed* interval from a to b . This corresponds to the set of numbers between a and b on the real line, including the endpoints a and b . $[a, b] = \{x \in \mathbb{R} \mid a \leq x \leq b\}$.
- (a, b) : the *open* interval from a to b . This corresponds to the set of numbers between a and b on the real line, *not* including the endpoints a and b . $(a, b) = \{x \in \mathbb{R} \mid a < x < b\}$.
- $[a, b)$: the half-open interval that includes the left endpoint a but not the right endpoint b . $[a, b) = \{x \in \mathbb{R} \mid a \leq x < b\}$.

Sometimes we encounter intervals that consist of two disjointed parts. We use the notation $[a, b] \cup [c, d]$ to denote the union of the two intervals, which is the set of numbers *either* between a and b (inclusive) *or* between c and d (inclusive).

Sets

Much of math's power comes from *abstraction*: the ability to see the bigger picture and think *meta* thoughts about the common relationships between math objects. We can think of individual numbers like 3, -5 , and π , or we can talk about the set of *all* numbers.

It is often useful to restrict our attention to a specific *subset* of the numbers as in the following examples.

Example 1: The nonnegative real numbers

Define $\mathbb{R}_+ \subset \mathbb{R}$ (read “ $\mathbb{R}_+$ is a subset of $\mathbb{R}$ ”) to be the set of nonnegative real numbers:

$$\mathbb{R}_+ \stackrel{\text{def}}{=} \{\text{all } x \text{ in } \mathbb{R} \text{ such that } x \geq 0\},$$

or expressed more compactly,

$$\mathbb{R}_+ \stackrel{\text{def}}{=} \{x \in \mathbb{R} \mid x \geq 0\}.$$

If we were to translate the above expression into plain English, it would read “The set $\mathbb{R}_+$ is defined as the set of all real numbers x such that x is greater or equal to zero.”

Note we used the “is defined as” symbol “ $\stackrel{\text{def}}{=}$ ” instead of the basic “ $=$ ” to give an extra hint that we’re defining a new variable $\mathbb{R}_+$ that is equal to the set expression on the right. In this book, we’ll sometimes use the symbol “ $\stackrel{\text{def}}{=}$ ” when defining new variables and math quantities. Some other books use the notation “ $:=$ ” or “ $\equiv$ ” for this purpose. The meaning of “ $\stackrel{\text{def}}{=}$ ” is identical to “ $=$ ” but it tells us the variable on the left of the equality is new.

Example 2: Even and odd integers

Define the set of even integers as

$$E \stackrel{\text{def}}{=} \{m \in \mathbb{Z} \mid m = 2n, n \in \mathbb{Z}\} = \{\dots, -4, -2, 0, 2, 4, \dots\}$$

and the set of odd integers as

$$O \stackrel{\text{def}}{=} \{m \in \mathbb{Z} \mid m = 2n + 1, n \in \mathbb{Z}\} = \{\dots, -3, -1, 1, 3, 5, \dots\}.$$

Indeed, every even number is divisible by two, so it can be written in the form $2n$ for some integer n . Odd numbers can be obtained from the “template” $2n + 1$, with n varying over all integers.

In both of the above examples, we use the *set-builder* notation $\{\dots \mid \dots\}$ to define the sets. Inside the curly braces we first describe the general kind of mathematical objects we are talking about, followed by the symbol “ $\mid$ ” (read “such that”), followed by the conditions that must be satisfied by all elements of the set.

Number sets

Recall the fundamental number sets we defined in Section 1.2 in the beginning of the book. It is worthwhile to review them briefly.

The *natural* numbers form the set derived when you start from 0 and add 1 any number of times:

$$\mathbb{N} \stackrel{\text{def}}{=} \{0, 1, 2, 3, 4, 5, 6, \dots\}.$$

We use the notation $\mathbb{N}^*$ to denote the set of *positive natural numbers*. The set $\mathbb{N}^*$ is the same as $\mathbb{N}$ but excludes zero.

The integers are the numbers derived by adding or subtracting 1 some number of times:

$$\mathbb{Z} \stackrel{\text{def}}{=} \{x \mid x = \pm n, n \in \mathbb{N}\}.$$

If we allow for divisions between integers, we require the set of rational numbers to represent the results:

$$\mathbb{Q} \stackrel{\text{def}}{=} \left\{ \frac{m}{n} \mid m \in \mathbb{Z}, n \in \mathbb{N}^* \right\},$$

In words, this expression is telling us that every rational number can be written as a fraction $\frac{m}{n}$, where m is an integer ($m \in \mathbb{Z}$), and n is a positive natural number ($n \in \mathbb{N}^*$).

The broader class of real numbers also includes all rationals as well as irrational numbers like $\sqrt{2}$ and π :

$$\mathbb{R} \stackrel{\text{def}}{=} \{\pi, e, -1.53929411\dots, 4.99401940129401\dots, \dots\}.$$

Finally, we have the set of complex numbers:

$$\mathbb{C} \stackrel{\text{def}}{=} \{1, i, 1 + i, 2 + 3i, \dots\},$$

where $i \stackrel{\text{def}}{=} \sqrt{-1}$ is the unit imaginary number.

Note that the definitions of $\mathbb{R}$ and $\mathbb{C}$ are not very precise. Rather than give a precise definition of each set inside the curly braces as we did for $\mathbb{Z}$ and $\mathbb{Q}$, we instead stated some examples of the elements in the set. Mathematicians sometimes do this and expect you to guess the general pattern for all the elements in the set.

The following inclusion relationship holds for the fundamental sets of numbers:

$$\mathbb{N} \subset \mathbb{Z} \subset \mathbb{Q} \subset \mathbb{R} \subset \mathbb{C}.$$

This relationship means every natural number is also an integer. Every integer is a rational number. Every rational number is a real. And every real number is also a complex number. See Figure 1.2 (page 12) for an illustration of the subset relationship between the number sets.

Rational numbers and fractions

So far in this book, we've used the notions of "fraction" and "rational number" somewhat interchangeably. Now that we've learned about sets, we can clarify the differences and equivalencies between these related concepts.

The same rational number $\frac{2}{3}$ can be written as a fraction in multiple, equivalent ways. The fractions $\frac{2}{3}$, $\frac{4}{6}$, $\frac{6}{9}$, $\frac{8}{12}$, and $\frac{2k}{3k}$ all correspond

to the same rational number. Keep in mind the existence of these *equivalent fractions* whenever checking whether two rational numbers are equal. For example, one person could obtain the answer $\frac{2}{3}$ to a given problem, while another person obtains the answer $\frac{4}{6}$. Since the two fractions look different, we might think these are different answers, when in fact both answers correspond to the same rational number.

A *reduced fraction* is a fraction of the form $\frac{m}{n}$ such that the numbers m and n are the smallest possible. We can obtain the reduced fraction by getting rid of any common factors that appear both in the numerator and denominator. For example,

$$\frac{4}{6} = \frac{2 \cdot 2}{3 \cdot 2} = \frac{2 \cdot \cancel{2}}{3 \cdot \cancel{2}} = \frac{2}{3},$$

where we cancelled the common factor 2 to obtain the equivalent reduced fraction. Reduced fractions are a useful representation for the set of rational numbers, because each rational number corresponds to a unique reduced fraction. Two rational numbers are equal if and only if they correspond to the same reduced fraction.

Subsets of the real line

Recall that the real numbers $\mathbb{R}$ have a graphical representation as points on the number line. See Figure 1.9 on page 35 for a reminder. The number line is also useful for representing various subsets of the real numbers, which we call *intervals*. We can graphically represent an interval by setting a section of the number line in **bold**. For example, the set of numbers that are strictly greater than 2 and strictly smaller than 4 is represented mathematically either as “ $(2, 4)$,” or more explicitly as

$$\{x \in \mathbb{R} \mid 2 < x < 4\},$$

or graphically as in Figure 1.63.




Figure 1.63: The open interval $(2, 4) = \{x \in \mathbb{R} \mid 2 < x < 4\}$.

Let's read the mathematical definition of this set carefully, and try to connect it with the graphical representation. Recall that the symbol $\in$ denotes set membership and the vertical bar stands for “such that,” so the whole expression “ $\{x \in \mathbb{R} \mid 2 < x < 4\}$ ” is read “the set of

real numbers x , such that $2 < x < 4$." Indeed this is also the region shown in bold in Figure 1.63.

Note that this interval is described by *strict* inequalities, which means the subset contains 2.000000001 and 3.99999999, but doesn't contain the endpoints 2 and 4. These *open* endpoints 2 and 4 are denoted on the number line as empty dots. An empty dot indicates that the endpoint is not included in the set.

We use the *union* symbol ($\cup$) to denote subsets of the number line that consist of several parts. For example, the set of numbers that lies *either* between -3 and 0 *or* between 1 and 2 is written as

$$\{x \in \mathbb{R} \mid -3 \leq x \leq 0\} \cup \{x \in \mathbb{R} \mid 1 \leq x \leq 2\}.$$




Figure 1.64: The graphical representation of the set $[-3, 0] \cup [1, 2]$.

This set is defined by less-than-or-equal inequalities, so the intervals contain their endpoints. These *closed* endpoints are denoted on the number line with filled-in dots.

Set relations

We'll now introduce a useful graphical representation for set relations and set operations. Although sets are purely mathematical constructs and they have no "shape," we can draw *Venn diagrams* to visualize relationships between sets and different subsets.

Consider the notion of a set B that is strictly contained in another set A . We write $B \subset A$ if $\forall b \in B, b \in A$ as well. Written in words, $B \subset A$ tells us every element of B is also an element of A .




Figure 1.65: Venn diagram showing an example of the set relation $B \subset A$. The set B is strictly contained in the set A .

Figure 1.65 shows the picture that mathematicians have in mind when they say, "The set B is contained in the set A ." The picture helps us visualize this abstract mathematical notion.

Mathematicians use two different symbols to describe set containment, in order to specify either a *strict* containment relation or a *subset-or-equal* relation. The two types of containment relations between sets are similar to the *less-than* ($<$) and *less-than-or-equal* ($\leq$) relations between numbers. A strict containment relation is denoted by the symbol $\subset$. We write $B \subset A$ if and only if every element of B is also an element of A , and there exists at least one element of A that is not an element of B . Using set notation, the previous sentence is expressed as

$$B \subset A \quad \Leftrightarrow \quad \forall b \in B, b \in A \text{ and } \exists a \in A \text{ such that } a \notin B.$$

For example, the expression $E \subset \mathbb{Z}$ shows that the even numbers are a strict subset of the integers. Every even number is an integer, but there exist integers that are not even (the odd numbers). Some mathematicians prefer the more descriptive symbol $\subsetneq$ to describe strict containment relations.

A subset-or-equal relation is denoted $B \subseteq A$. In writing $B \subseteq A$, a mathematician claims, “Every element of B is also an element of A ,” but makes no claim about the existence of elements that are contained in A but not in B . The statement $B \subset A$ implies $B \subseteq A$; however, $B \subseteq A$ does not imply $B \subset A$. This is analogous to how $b < a$ implies $b \leq a$, but $b \leq a$ doesn’t imply $b < a$, since a and b could be equal.

Set operations

Venn diagrams also help us visualize the subsets obtained from set operations. Figure 1.66 illustrates the set union $A \cup B$, the set intersection $A \cap B$, and the set difference $A \setminus B$, for two sets A and B .

The union $A \cup B$ describes all elements that are in either set A or set B , or both. If $e \in A \cup B$, then $e \in A$ or $e \in B$.

Recall the set of even numbers $E \subset \mathbb{Z}$ and the set of odd numbers $O \subset \mathbb{Z}$ defined above. Since every integer is either an even number or an odd number, we know $\mathbb{Z} \subseteq E \cup O$. The union of two subsets is always contained within the parent set, so we also know $E \cup O \subseteq \mathbb{Z}$. Combining these facts, we can establish the equality $E \cup O = \mathbb{Z}$, which states the fact, “The combination of all even and odd numbers is the same as all integers.”

The set intersection $A \cap B$ and set difference $A \setminus B$ are also illustrated in Figure 1.66. The intersection of two sets contains the elements that are part of both sets. The set difference $A \setminus B$ contains all the elements that are in A but not in B .

Note the meaning of the conjunction “or” in English is ambiguous. The expression “in A or B ” could be interpreted as either an




Figure 1.66: Venn diagrams showing different subsets obtained using the set operations: set union $A \cup B$, set intersection $A \cap B$, and set difference $A \setminus B$.

“inclusive or,” meaning “in A or B , or in both”—or as an “exclusive or,” meaning “in A or B , but not both.” Mathematicians always use “or” in the inclusive sense, so $A \cup B$ denotes elements that are in A or B , or in both sets. We can obtain an expression that corresponds to the “exclusive or” of two sets by taking the union of the sets and subtracting their intersection: $(A \cup B) \setminus (A \cap B)$.

Example 3: Set operations

Consider the three sets $A = \{a, b, c\}$, $B = \{b, c, d\}$, and $C = \{c, d, e\}$. Using set operations, we can define new sets, such as

$$A \cup B = \{a, b, c, d\}, \quad A \cap B = \{b, c\}, \quad \text{and} \quad A \setminus B = \{a\},$$

which correspond to elements in either A or B , the set of elements in A and B , and the set of elements in A but not in B , respectively.

We can also construct expressions involving three sets:

$$A \cup B \cup C = \{a, b, c, d, e\}, \quad A \cap B \cap C = \{c\}.$$

And we can write more elaborate set expressions, like

$$(A \cup B) \setminus C = \{a, b\},$$

which is the set of elements that are in A or B but not in C .

Another example of a complicated set expression is

$$(A \cap B) \cup (B \cap C) = \{b, c, d\},$$

which describes the set of elements in both A and B or in both B and C . As you can see, set notation is a compact, precise language for writing complicated set expressions.

Example 4: Word problem

A startup is looking to hire student interns for the summer. Define C to be the subset of students who are good with computers, M

the subset of students who know math, D the students with design skills, and L the students with good language skills.

Using set notation, we can specify different subsets of the students the startup might hire. Let's say the startup is a math textbook publisher; they want to hire students from the set $M \cap L$ —the students who are good at math and who also have good language skills. A startup that builds websites needs both designers and coders, and therefore would choose students from the set $D \cup C$.

New vocabulary

The specialized notation used by mathematicians can be difficult to get used to. You must learn how to read symbols like $\exists$, $\subset$, $|$, and $\in$ and translate their meaning in the sentence. Indeed, learning advanced mathematics notation is akin to learning a new language.

To help you practice the new vocabulary, we'll look at a simple mathematical proof that makes use of the new symbols.

Simple proof example

Claim: Given $J(n) = 3n + 2 - n$, $J(n) \in E$ for all $n \in \mathbb{Z}$.

The claim is that the function $J(n)$ outputs an even number, whenever the input n is an integer. To prove this claim, we have to show that the expression $3n + 2 - n$ is even for all numbers $n \in \mathbb{Z}$.

Proof. We want to show $J(n) \in E$ for all $n \in \mathbb{Z}$. Let's first review the definition of the set of even numbers $E \stackrel{\text{def}}{=} \{m \in \mathbb{Z} \mid m = 2n, n \in \mathbb{Z}\}$. A number is even if it is equal to $2n$ for some integer n . Next let's simplify the expression for $J(n)$ as follows:

$$J(n) = 3n + 2 - n = 2n + 2 = 2(n + 1).$$

Observe that the number $(n + 1)$ is always an integer whenever n is an integer. Since the output of $J(n) = 2(n + 1)$ is equal to $2m$ for some integer m , we've proven that $J(n) \in E$, for all $n \in \mathbb{Z}$. $\square$

Sets as solutions to equations

Another context where sets come up is when describing solutions to equations and inequalities. In Section 1.1 we learned how to solve for the unknown x in equations. To solve the equation $f(x) = c$ is to find all the values of x that satisfy this equation. For simple equations like $x - 3 = 6$, the solution is a single number $x = 9$, but more complex equations can have multiple solutions. For example, the solution to

the equation $x^2 = 4$ is the set $\{-2, 2\}$, since both $x = -2$ and $x = 2$ satisfy the equation.

Please update your definition of the math verb “to solve” (an equation) to include the new notion of a *solution set*—the set of values that satisfy the equation. A solution set is the mathematically precise way to describe an equation’s solutions:

- The solution set to the equation $x - 3 = 6$ is the set $\{9\}$.
- The solution set for the equation $x^2 = 4$ is the set $\{-2, 2\}$.
- The solution set of $\sin(x) = 0$ is the set $\{x \mid x = \pi n, \forall n \in \mathbb{Z}\}$.
- The solution set for the equation $\sin(x) = 2$ is $\emptyset$ (the empty set), since there is no number x that satisfies the equation.

The SymPy function `solve` returns the solutions of equations as a list. To solve the equation $f(x) = c$ using SymPy, we first rewrite it as expression that equals zero $f(x) - c = 0$, then call the function `solve`:

```
>>> solve(x-3 -6, x)          # usage: solve(expr, var)
[9]

>>> solve(x**2 -4, x)
[-2, 2]

>>> solve(sin(x), x)
[0, pi]                        # found only solutions in [0,2*pi]

>>> solve(sin(x) -2, x)
[]                              # empty list = empty set
```

In the next section we’ll learn how the notion of a solution set is used for describing the solutions to systems of equations.

Solution sets to systems of equations

Let’s revisit what we learned in Section 1.15 about the solutions to systems of linear equations, and define their solution sets more precisely. The solution set for the system of equations

$$\begin{aligned}a_1x + b_1y &= c_1, \\a_2x + b_2y &= c_2,\end{aligned}$$

corresponds to the intersection of two sets:

$$\underbrace{\{(x, y) \in \mathbb{R}^2 \mid a_1x + b_1y = c_1\}}_{\ell_1} \cap \underbrace{\{(x, y) \in \mathbb{R}^2 \mid a_2x + b_2y = c_2\}}_{\ell_2}.$$

Recall that the lines ℓ_1 and ℓ_2 are the geometric interpretation of these sets. Each line corresponds to a set of coordinate pairs (x, y) that satisfy the equation of the line. The solution to the system of equations

is the set of points at the intersection of the two lines $\ell_1 \cap \ell_2$. Note the word *intersection* is used in two different mathematical contexts: the solution is the *set intersection* of two sets, and also the *geometric intersection* of two lines.

Let's take advantage of this correspondence between set intersections and geometric line intersections to understand the solutions to systems of equations in a little more detail. In the next three sections, we'll look at three possible cases that can occur when trying to solve a system of two linear equations in two unknowns. So far we've only discussed Case A, which occurs when the two lines intersect at a point, as in the example shown in Figure 1.67. To fully understand the possible solutions to a system of equations, we need to think about all other cases; like Case B when $\ell_1 \cap \ell_2 = \emptyset$ as in Figure 1.68, and Case C when $\ell_1 \cap \ell_2 = \ell_1 = \ell_2$ as in Figure 1.69.

Case A: One solution. When the lines ℓ_1 and ℓ_2 are non-parallel, they will intersect at a point as shown in Figure 1.67. In this case, the solution set to the system of equations contains a single point:

$$\{(x, y) \in \mathbb{R}^2 \mid x + 2y = 2\} \cap \{(x, y) \in \mathbb{R}^2 \mid x = 1\} = \{(1, \tfrac{1}{2})\}.$$




Figure 1.67: Case A: The intersection of the lines with equations $x + 2y = 2$ and $x = 1$ is the point $(1, \frac{1}{2}) \in \mathbb{R}^2$.

Case B: No solution. If the lines ℓ_1 and ℓ_2 are parallel then they will never intersect. The intersection of these lines is the empty set:

$$\{(x, y) \in \mathbb{R}^2 \mid x + 2y = 2\} \cap \{(x, y) \in \mathbb{R}^2 \mid x + 2y = 4\} = \emptyset.$$

Think about it—there is no point (x, y) that lies on both ℓ_1 and ℓ_2 . Using algebra terminology, we say this system of equations has no solution, since there are no numbers x and y that satisfy both equations.




Figure 1.68: Case B: The lines with equations $x + 2y = 2$ and $x + 2y = 4$ are parallel and do not intersect. Using set notation, we can describe the solution set as $\emptyset$ (the empty set).

Case C: Infinitely many solutions. If the lines l_1 and l_2 are parallel and overlapping then they intersect everywhere. This case occurs when one of the equations in a system of equations is a multiple of the other equation, as in the case of equations $x + 2y = 2$ and $3x + 6y = 6$. The lines l_1 and l_2 that correspond to these equations are shown in Figure 1.69. Any point (x, y) that satisfies $x + 2y = 2$ also satisfies $3x + 6y = 6$. Since both equations describe the same geometric line, the intersection of the two lines is equal to the lines: $l_1 \cap l_2 = l_1 = l_2$. In this case, the solution to the system of equations is described by the set $\{(x, y) \in \mathbb{R}^2 \mid x + 2y = 2\}$.




Figure 1.69: Case C: the line l_1 described by equation $x + 2y = 2$ and the line l_2 described by equation $3x + 6y = 6$ correspond to the same line in the Cartesian plane. The intersection of these lines is the set $\{(x, y) \in \mathbb{R}^2 \mid x + 2y = 2\} = l_1 = l_2$.

We need to consider all three cases when thinking about the solutions to systems of linear equations: the solution set can be a point (Case A), the empty set (Case B), or a line (Case C). Observe that

the same mathematical notion (a set) is able to describe the solutions in all three cases even though the solutions correspond to very different geometric objects. In Case A the solution is a set that contains a single point $\{(x, y)\}$. In Case B the solution is the empty set $\emptyset$. And in Case C the solution set is described by the infinite set $\{(x, y) \in \mathbb{R}^2 \mid ax + by = c\}$, which corresponds to a line ℓ in the Cartesian plane. I hope you'll agree with me that set notation is useful for describing mathematical concepts precisely and for handling solutions to linear equations.

Sets are also useful for describing the solutions to inequalities, which is what we'll learn about next.

Inequalities

In this section, we'll learn how to solve inequalities. The solution set to an inequality is an *interval*—a subset of the number line. Consider the inequality $x^2 \leq 4$, which is equivalent to asking the question, “For which values of x is x^2 less than or equal to 4?” The answer to this question is the interval $[-2, 2] = \{x \in \mathbb{R} \mid -2 \leq x \leq 2\}$.

Working with inequalities is essentially the same as working with their endpoints. To solve the inequality $x^2 \leq 4$, we first solve $x^2 = 4$ to find the endpoints and then use trial and error to figure out which part of the space to the left and right of the endpoints satisfies the inequality.

It's important to distinguish the different types of inequality conditions. The four different types of inequalities are

- $f(x) < g(x)$: a strict inequality. The function $f(x)$ is always *strictly less than* the function $g(x)$.
- $f(x) \leq g(x)$: the function $f(x)$ is *less than or equal to* $g(x)$.
- $f(x) > g(x)$: $f(x)$ is *strictly greater than* $g(x)$.
- $f(x) \geq g(x)$: $f(x)$ is *greater than or equal to* $g(x)$.

Depending on the type of inequality, the answer will be either a *open* or *closed* interval.

To solve inequalities we use the techniques we learned for solving equations: we perform simplifying steps **on both sides of the inequality** until we obtain the answer. The only new aspect when dealing with inequalities is the following. When multiplying an inequality by a negative number on both sides, we must flip the direction of the inequality:

$$f(x) \leq g(x) \quad \Rightarrow \quad -f(x) \geq -g(x).$$

Example 5 To solve the inequality $7 - x \leq 5$ we must *dig* toward the x and *undo* all the operations that stand in our way:

$$\begin{aligned} 7 - x &\leq 5, \\ (-x) + 7 &\leq 5, \\ (-x) + 7 - 7 &\leq 5 - 7, \\ -x &\leq -2, \\ x &\geq 2. \end{aligned}$$

To obtain the second line we simply rewrote the order of operations on the left side of the inequality. In the third line we subtracted 7 from both sides of the inequality to undo the $+7$ operation. In the last step we multiplied both sides of the inequality by -1 , which had the effect of changing the inequality from $\leq$ to $\geq$. The solution set to the inequality $7 - x \leq 5$ is the interval $[2, \infty)$.

Example 6 To solve the inequality $x^2 \leq 4$, we must undo the quadratic function by taking the square root of both sides of the inequality. Note the equation $x^2 = 4$ has two solutions: $x = -2$ and $x = 2$. Similarly, we'll need to consider two separate cases for the inequality conditions. Simplifying the inequality $x^2 \leq 4$ by taking the square root on both sides results in two inequality conditions

$$x \geq -2 \quad \text{and} \quad x \leq 2,$$

which we can express more concisely as $-2 \leq x \leq 2$. If x is a negative number, it must be greater than -2 ; and if x is a positive number, it must be less than 2 in order for $x^2 \leq 4$. The solution set for the inequality $x^2 \leq 4$ is the interval $[-2, 2] = \{x \in \mathbb{R} \mid -2 \leq x \leq 2\}$. Note the solution is a closed interval (square brackets), which means the endpoints are included.

The best way to convince yourself that the above algebraic reasoning is correct is to think about the graph of the function $f(x) = x^2$. The inequality $x^2 \leq 4$ corresponds to the condition $f(x) \leq 4$. For what values of x is the graph of the function $f(x)$ below the line with equation $y = 4$?

As you can see, solving inequalities is no more complicated than solving equations. You can think about an inequality in terms of its endpoints, which correspond to the equality conditions. Whenever things get complicated (as in Example 6), you can sketch the function graphs for the different terms in the inequality and visually determine the appropriate directions for the inequality signs.

Sets related to functions

A function that takes real variables as inputs and produces real numbers as outputs is denoted $f : \mathbb{R} \rightarrow \mathbb{R}$. The *domain* of a function is the set of all possible inputs to the function that produce an output:

$$\text{Dom}(f) \stackrel{\text{def}}{=} \{x \in \mathbb{R} \mid f(x) \in \mathbb{R}\}.$$

Inputs for which the function is undefined are not part of the domain. For instance the function $f(x) = \sqrt{x}$ is not defined for negative inputs, so we have $\text{Dom}(f) = \mathbb{R}_+$.

The *image* of a function is the set of all possible outputs of the function:

$$\text{Im}(f) \stackrel{\text{def}}{=} \{y \in \mathbb{R} \mid \exists x \in \mathbb{R}, y = f(x)\}.$$

For example, the function $f(x) = x^2$ has the image set $\text{Im}(f) = \mathbb{R}_+$ since the outputs it produces are always nonnegative.

Discussion

Knowledge of the precise mathematical jargon introduced in this section is not crucial to understanding basic mathematics. That said, I wanted to expose you to some technical math notation here because this is the language in which mathematicians think and communicate. Most advanced math textbooks will assume you understand technical math notation, so it's good to be prepared.

Exercises

E1.27 Given the three sets $A = \{1, 2, 3, 4, 5, 6, 7\}$, $B = \{1, 3, 5\}$, and $C = \{2, 4, 6\}$, compute the following set expressions.

a) $A \setminus B$	b) $B \cup C$	c) $A \cap B$	d) $B \cap C$
e) $A \cup B \cup C$	f) $A \setminus (B \cup C)$	g) $(A \setminus B) \cup C$	h) $A \cap B \cap C$

E1.28 Find the values of x that satisfy the following inequalities.

a) $2x < 3$	b) $-4x \geq 20$	c) $ 2x - 3 < 5$
d) $3x + 3 < 5x - 5$	e) $\frac{1}{2}x - 2 \geq \frac{1}{3}$	f) $(x + 1)^2 \geq 9$

Express your answer as an interval with appropriate endpoints.

1.17 Math problems

We've now reached the first section of problems in this book. The purpose of these problems is to give you a way to comprehensively practice your math fundamentals. Knowing how to solve math problems is a very useful skill to develop. At times, honing your math chops might seem like tough mental work, but at the end of each problem, you'll gain a stronger foothold on all the topics you've been learning about. You'll also experience a small *achievement buzz* after each problem you vanquish.

Sit down and take a crack at these practice problems today, or another time when you're properly caffeinated. If you make time for some math practice, you'll develop long-lasting comprehension and true math fluency.

Without solving any problems, you're likely to forget most of what you've learned in the next few months. You might still remember the big ideas, but the details will be fuzzy and faded. By solving some of the practice problems, you'll remember a lot more stuff. Don't break the pace now: with math, it's very much *use it or lose it!*

Make sure you put your phone away while you're working on the problems. You don't need fancy technology to do math; grab a pen and some paper from the printer and you'll be fine. The great mathematicians like Descartes, Hilbert, Leibniz, and Noether did most of their work with pen and paper and they did well. Spend some time with math the way they did.

P1.1 Solve for x in the equation $x^2 - 9 = 7$.

P1.2 Solve for x in the equation $\cos^{-1}\left(\frac{x}{\lambda}\right) - \phi = \omega t$.

P1.3 Solve for x in the equation $\frac{1}{x} = \frac{1}{a} + \frac{1}{b}$.

P1.4 Use a calculator to find the values of the following expressions:

a) $\sqrt[4]{3^3}$

b) 2^{10}

c) $7^{\frac{1}{4}} - 10$

d) $\frac{1}{2} \ln(e^{22})$

P1.5 Compute the following expressions involving fractions:

a) $\frac{1}{2} + \frac{1}{4}$

b) $\frac{4}{7} - \frac{23}{5}$

c) $1\frac{3}{4} + 1\frac{31}{32}$

P1.6 Use the basic rules of algebra to simplify the following expressions:

a) $ab \frac{1}{a} b^2 cb^{-3}$

b) $\frac{abc}{bca}$

c) $\frac{27a^2}{\sqrt{9abba}}$

d) $\frac{a(b+c) - ca}{b}$

e) $\frac{a}{c\sqrt[3]{b}} \frac{b^{\frac{4}{3}}}{a^2}$

f) $(x+a)(x+b) - x(a+b)$

P1.7 Expand the brackets in the following expressions:

a) $(x + a)(x - b)$

b) $(2x + 3)(x - 5)$

c) $(5x - 2)(2x + 7)$

P1.8 Factor the following expressions as a product of linear terms:

a) $x^2 - 2x - 8$

b) $3x^3 - 27x$

c) $6x^2 + 11x - 21$

P1.9 Complete the square in the following quadratic expressions to obtain expressions of the form $A(x - h)^2 + k$.

a) $x^2 - 4x + 7$

b) $2x^2 + 12x + 22$

c) $6x^2 + 11x - 21$

P1.10 A golf club and a golf ball cost \$1.10 together. The golf club costs one dollar more than the ball. How much does the ball cost?

P1.11 A father is 35 years old and his son is 5 years old. In how many years will the father's age be four times the son's age?

P1.12 A boy and a girl collected 120 nuts. The girl collected twice as many nuts as the boy. How many nuts did each collect?

P1.13 Alice is 5 years older than Bob. The sum of their ages is 25 years. How old is Alice?

P1.14 A publisher needs to bind 4500 books. One print shop can bind these books in 30 days, another shop can do it in 45 days. How many days are necessary to bind all the books if both shops work in parallel?

Hint: Find the books-per-day rate of each shop.

P1.15 A plane leaves Vancouver travelling at 600 km/h toward Montreal. One hour later, a second plane leaves Vancouver heading for Montreal at 900 km/h. How long will it take for the second plane to overtake the first?

Hint: Distance travelled is equal to velocity multiplied by time: $d = vt$.

P1.16 The golden ratio, denoted φ , is the positive solution to the quadratic equation $x^2 - x - 1 = 0$. Find the golden ratio.

P1.17 Solve for x in the equation $\frac{1}{x} + \frac{2}{1-x} = \frac{4}{x^2}$.

Hint: Multiply both sides of the equation by $x^2(1-x)$.

P1.18 Use substitution to solve for x in the following equations:

a) $x^6 - 4x^3 + 4 = 0$

b) $\frac{1}{2 - \sin x} = \sin x$

P1.19 Find the range of values of the parameter m for which the equation $2x^2 - mx + m = 0$ has no real solutions.

Hint: Use the quadratic formula.

P1.20 Use the properties of exponents and logarithms to simplify

a) $e^x e^{-x} e^z$

b) $\left(\frac{xy^{-2}z^{-3}}{x^2y^3z^{-4}}\right)^{-3}$

c) $(8x^6)^{-\frac{2}{3}}$

P1.21 Two algorithms, P and Q, can be used to solve a certain problem. The running time of Algorithm P as a function of the size of the problem n is described by the function $P(n) = 0.002n^2$. The running time of Algorithm Q is described by $Q(n) = 0.5n$. For small problems, Algorithm P runs faster. Starting from what n will Algorithm Q be faster?

P1.22 Consider a right-angle triangle in which the shorter sides are 8 cm and 6 cm. What is the length of the triangle's longest side?

P1.23 A television screen measures 26 inches on the diagonal. The screen height is 13 inches. How wide is the screen?

P1.24 Given the angle and distance measurements labelled in Figure 1.70, calculate the distance d and the height of the mountain peak h .





Figure 1.70: Measuring the height of a mountain using angles.

Hint: Use the definition of $\tan \theta$ to obtain two equations in two unknowns.


P1.25 Find x . Express your answer in terms of a , b , c and θ .



Hint: Use Pythagoras' theorem twice; then use the function $\tan$.

P1.26 Use the power-reduction trigonometric identities (page 78) to express $\sin^2 \theta \cos^2 \theta$ in terms of $\cos 4\theta$.

P1.27 Find the length of side c in the triangle:



Hint: Use the sine rule.

P1.28 Consider the obtuse triangle shown in Figure 1.71.

- Express h in terms of a and θ .
- What is the area of this triangle?
- Express c in terms of the variables a , b , and θ .

Hint: You can use the cosine rule for part c).




Figure 1.71: A triangle with base b and height h .

P1.29 Satoshi likes warm saké. He places 1 litre of water in a sauce pan with diameter 17 cm. How much will the height of the water level rise when Satoshi immerses a saké bottle with diameter 7.5 cm?

Hint: You'll need the volume conversion ratio $1 \text{ litre} = 1000 \text{ cm}^3$.

P1.30 The length of a rectangle is $c + 2$ and its height is 5. What is the area of the rectangle?

P1.31 A box of facial tissues has dimensions 10.5 cm by 7 cm by 22.3 cm. What is the volume of the box in litres?

Hint: $1 \text{ L} = 1000 \text{ cm}^3$.

P1.32 A swimming pool has length $\ell = 20 \text{ m}$, width $w = 10 \text{ m}$, and depth $d = 1.5 \text{ m}$. Calculate the volume of water in the swimming pool in litres?

Hint: $1 \text{ m}^3 = 1000 \text{ L}$.

P1.33 How many litres of water remain in a tank that is 15 m long, 6 m wide, and 5 m high, if 30% of its capacity is spent?

P1.34 A building has two water tanks, each with capacity 4000 L. One of them is $\frac{1}{4}$ full and the other contains three times more water. How many litres of water does the building have in total?

P1.35 A man sells firewood. To make standard portions, he uses a standard length of rope ℓ to surround a pack of logs. One day, a customer asks him for a double portion of firewood. What length of rope should he use to measure this order? Assume the packs of logs are circular in shape.

P1.36 How much pure water should be added to 10 litres of a solution that is 60% acid to make a solution that is 20% acid?

P1.37 A tablet screen has a resolution of 768 pixels by 1024 pixels, and the physical dimensions of the screen are 6 inches by 8 inches. One might conclude that the best size of a PDF document for such a screen would be 6 inches by 8 inches. At first I thought so too, but I forgot to account for the status bar, which is 20 pixels tall. The actual usable screen area is only 768 pixels by 1004 pixels. Assuming the width of the PDF is chosen to be 6 inches, what height should the PDF be so that it fits perfectly in the content area of the tablet screen?

P1.38 Find the sum of the natural numbers 1 through 100.

Hint: Imagine pairing the biggest number with the smallest number in the sum, the second biggest number with the second smallest number, etc.

P1.39 Solve for x and y simultaneously in the following system of equations: $-x - 2y = -2$ and $3x + 3y = 0$.

P1.40 Solve the following system of equations for the three unknowns:

$$\begin{aligned}1x + 2y + 3z &= 14, \\2x + 5y + 6z &= 30, \\-1x + 2y + 3z &= 12.\end{aligned}$$

P1.41 Express the following vectors in length-and-direction notation:

$$\text{a) } \vec{u}_1 = (0, 5) \qquad \text{b) } \vec{u}_2 = (1, 2) \qquad \text{c) } \vec{u}_3 = (-1, -2)$$

P1.42 Express the following vectors as components:

$$\text{a) } \vec{v}_1 = 20\angle 30^\circ \qquad \text{b) } \vec{v}_2 = 10\angle -90^\circ \qquad \text{c) } \vec{v}_3 = 5\angle 150^\circ$$

P1.43 Express the following vectors in terms of unit vectors $\hat{i}$, $\hat{j}$, and $\hat{k}$:

$$\text{a) } \vec{w}_1 = 10\angle 25^\circ \qquad \text{b) } \vec{w}_2 = 7\angle -90^\circ \qquad \text{c) } \vec{w}_3 = (3, -2, 3)$$

P1.44 Given the vectors $\vec{v}_1 = (1, 1)$, $\vec{v}_2 = (2, 3)$, and $\vec{v}_3 = 5\angle 30^\circ$, calculate the following expressions:

$$\text{a) } \vec{v}_1 + \vec{v}_2 \qquad \text{b) } \vec{v}_2 - 2\vec{v}_1 \qquad \text{c) } \vec{v}_1 + \vec{v}_2 + \vec{v}_3$$

P1.45 Starting from the point $P = (2, 6)$, the three displacement vectors shown in Figure 1.72 are applied to obtain the point Q . What are the coordinates of the point Q ?

P1.46 Compute the following expressions:

$$\text{a) } \sqrt{-4} \qquad \text{b) } \frac{2+3i}{2+2i} \qquad \text{c) } e^{3i}(2+i)e^{-3i}$$




Figure 1.72: A point P is displaced by three vectors to obtain point Q .

P1.47 Solve for $x \in \mathbb{C}$ in the following equations:

a) $x^2 = -4$

b) $\sqrt{x} = 4i$

c) $x^2 + 2x + 2 = 0$

d) $x^4 + 4x^2 + 3 = 0$

Hint: To solve **d**), use the substitution $u = x^2$.

P1.48 Given the numbers $z_1 = 2 + i$, $z_2 = 2 - i$, and $z_3 = -1 - i$, compute

a) $|z_1|$

b) $\frac{z_1}{z_3}$

c) $z_1 z_2 z_3$

P1.49 Let B be the set of people who are bankers and C be the set of crooks. Rewrite the math statement $\exists b \in B \mid b \notin C$ in plain English.

P1.50 Let M denote the set of people who run Monsanto, and H denote the people who ought to burn in hell for all eternity. Write the math statement $\forall p \in M, p \in H$ in plain English.

P1.51 When starting a business, one sometimes needs to find investors. Define M to be the set of investors with money, and C to be the set of investors with connections. Describe the following sets in words: a) $M \setminus C$, b) $C \setminus M$, and the most desirable set c) $M \cap C$.

Chapter 2

Intro to linear algebra

The first chapter reviewed core ideas of mathematics. Now that we're done with the prerequisites, we can begin the main discussion of linear algebra: the study of vectors and matrices.

2.1 Definitions

Vectors and matrices are the objects of study in linear algebra, and in this chapter we'll define them and learn the basic operations we can perform on them.

We denote the set of n -dimensional vectors with real coefficients as $\mathbb{R}^n$. A vector $\vec{v} \in \mathbb{R}^n$ is an n -tuple of real numbers.¹ For example, a three-dimensional vector is defined as a triple of numbers:

$$\vec{v} \stackrel{\text{def}}{=} (v_1, v_2, v_3).$$

To specify the vector $\vec{v}$, we must specify the values for its three components, v_1 , v_2 , and v_3 . We'll use the terms *components* and *coordinates* interchangeably throughout the book.

A matrix $A \in \mathbb{R}^{m \times n}$ is a rectangular array of real numbers with m rows and n columns. For example, a 3×2 matrix is defined like this:

$$A \stackrel{\text{def}}{=} \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix}.$$

To specify the matrix A , we need to specify the values of its six *entries*: a_{11} , a_{12} , a_{21} , a_{22} , a_{31} , and a_{32} .

¹The notation " $s \in S$ " is read " s is an element of S " or " s in S ."

In the remainder of this chapter we'll learn about the mathematical operations we can perform on vectors and matrices. Many problems in science, business, and technology can be described in terms of vectors and matrices, so it's important you understand how to work with these math objects.

Context

To illustrate what's new about vectors and matrices, let's begin by reviewing the properties of something more familiar: the set of real numbers $\mathbb{R}$. The basic operations for real numbers are:

- Addition (denoted $+$)
- Subtraction, the inverse of addition (denoted $-$)
- Multiplication (denoted implicitly)
- Division, the inverse of multiplication (denoted by fractions)

You're familiar with these operations and know how to use them to evaluate math expressions and solve equations.

You should also be familiar with *functions* that take real numbers as inputs and give real numbers as outputs, denoted $f : \mathbb{R} \rightarrow \mathbb{R}$. Recall that, by definition, the *inverse function* f^{-1} *undoes* the effect of f . If you are given $f(x)$ and want to find x , you can use the inverse function as follows: $f^{-1}(f(x)) = x$. For example, the function $f(x) = \ln(x)$ has the inverse $f^{-1}(x) = e^x$, and the inverse of $g(x) = \sqrt{x}$ is $g^{-1}(x) = x^2$.

Having reviewed the basic operations for real numbers $\mathbb{R}$, let's now introduce the basic operations for vectors $\mathbb{R}^n$ and matrices $\mathbb{R}^{m \times n}$.

Vector operations

The operations we can perform on vectors are:

- Addition (denoted $+$)
- Subtraction, the inverse of addition (denoted $-$)
- Scaling (denoted implicitly)
- Dot product (denoted $\cdot$)
- Cross product (denoted $\times$)

We'll discuss each of these vector operations in Section 2.2. Although you should already be familiar with vectors and vector operations from Section 1.13, it's worth revisiting these concepts in greater depth, because vectors are the foundation of linear algebra.

Matrix operations

The mathematical operations defined for matrices A and B are:

- Addition (denoted $A + B$)
- Subtraction, the inverse of addition (denoted $A - B$)
- Scaling by a constant α (denoted αA)
- Matrix product (denoted AB)
- Matrix-vector product (denoted $A\vec{v}$)
- Matrix inverse (denoted A^{-1})
- Trace (denoted $\text{Tr}(A)$)
- Determinant (denoted $\det(A)$ or $|A|$)

We'll define each of these operations in Section 2.3, and we'll learn about the various computational, geometric, and theoretical considerations associated with these matrix operations throughout the remainder of the book.

Let's now examine one important matrix operation in closer detail: the matrix-vector product $A\vec{x}$.

Matrix-vector product

Consider the matrix $A \in \mathbb{R}^{m \times n}$ and the vector $\vec{v} \in \mathbb{R}^n$. The matrix-vector product $A\vec{x}$ produces a linear combination of the columns of the matrix A with coefficients $\vec{x}$. For example, the product of a 3×2 matrix A and a 2×1 vector $\vec{x}$ results in a 3×1 vector, which we'll denote $\vec{y}$:

$$\vec{y} = A\vec{x},$$

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \underbrace{\begin{bmatrix} x_1 a_{11} + x_2 a_{12} \\ x_1 a_{21} + x_2 a_{22} \\ x_1 a_{31} + x_2 a_{32} \end{bmatrix}}_{\text{row picture}} = \underbrace{x_1 \begin{bmatrix} a_{11} \\ a_{21} \\ a_{31} \end{bmatrix} + x_2 \begin{bmatrix} a_{12} \\ a_{22} \\ a_{32} \end{bmatrix}}_{\text{column picture}}.$$

The key thing to observe in the above formula is the dual interpretation of the matrix-vector product $A\vec{x}$ in the “row picture” and in the “column picture.” In the row picture, we obtain the vector $\vec{y}$ by computing the dot product of the vector $\vec{x}$ with each of the rows of the matrix A . In the column picture, we interpret the vector $\vec{y}$ as x_1 times the first column of A plus x_2 times the second column of A . In other words, $\vec{y}$ is a linear combination of the columns of A . For example, if you want to obtain the linear combination consisting of three times the first column of A and four times the second column of A , you can multiply A by the vector $\vec{x} = \begin{bmatrix} 3 \\ 4 \end{bmatrix}$.

Linear combinations as matrix products

Consider some set of vectors $\{\vec{e}_1, \vec{e}_2\}$, and a third vector $\vec{y}$ that is a *linear combination* of the vectors $\vec{e}_1$ and $\vec{e}_2$:

$$\vec{y} = \alpha \vec{e}_1 + \beta \vec{e}_2.$$

The numbers $\alpha, \beta \in \mathbb{R}$ are the coefficients in this linear combination.

The matrix-vector product is defined expressly for the purpose of studying linear combinations. We can describe the linear combination $\vec{y} = \alpha \vec{e}_1 + \beta \vec{e}_2$ as the following matrix-vector product:

$$\vec{y} = \begin{bmatrix} | & | \\ \vec{e}_1 & \vec{e}_2 \\ | & | \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}.$$

The matrix E has $\vec{e}_1$ and $\vec{e}_2$ as columns. The dimensions of the matrix E will be $n \times 2$, where n is the dimension of the vectors $\vec{e}_1$, $\vec{e}_2$, and $\vec{y}$.

Linear transformations

Dear readers, we've reached the key notion in the study of linear algebra. This is the crux. The essential fibre. The main idea. I know you're ready to handle it because you're familiar with functions of a real variable $f : \mathbb{R} \rightarrow \mathbb{R}$, and you just learned the definition of the matrix-vector product (in which the variables were chosen to subliminally remind you of the standard conventions for the function input x and the function output $y = f(x)$). Without further ado, I present to you the concept of a *linear transformation*.

The matrix-vector product corresponds to the abstract notion of a *linear transformation*, which is one of the key notions in the study of linear algebra. Multiplication by a matrix $A \in \mathbb{R}^{m \times n}$ can be thought of as computing a linear transformation T_A that takes n -vectors as inputs and produces m -vectors as outputs:

$$T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m.$$

Instead of writing $\vec{y} = T_A(\vec{x})$ to denote the linear transformation T_A applied to the vector $\vec{x}$, we can write $\vec{y} = A\vec{x}$. Since the matrix A has m rows, the result of the matrix-vector product is an m -vector. Applying the linear transformation T_A to the vector $\vec{x}$ corresponds to the product of the matrix A and the column vector $\vec{x}$. We say T_A is *represented by* the matrix A .

Inverse When a matrix A is square and invertible, there exists an inverse matrix A^{-1} which *undoes* the effect of A to restore the original

input vector:

$$A^{-1}(A\vec{x}) = A^{-1}A\vec{x} = \vec{x}.$$

Using the matrix inverse A^{-1} to undo the effects of the matrix A is analogous to using the inverse function f^{-1} to undo the effects of the function f .

Example 1 Consider the linear transformation that multiplies the first components of input vectors by 3 and multiplies the second components by 5, as described by the matrix

$$A = \begin{bmatrix} 3 & 0 \\ 0 & 5 \end{bmatrix}, \quad A\vec{x} = \begin{bmatrix} 3 & 0 \\ 0 & 5 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 3x_1 \\ 5x_2 \end{bmatrix}.$$

The inverse of the matrix A is

$$A^{-1} = \begin{bmatrix} \frac{1}{3} & 0 \\ 0 & \frac{1}{5} \end{bmatrix}, \quad A^{-1}(A\vec{x}) = \begin{bmatrix} \frac{1}{3} & 0 \\ 0 & \frac{1}{5} \end{bmatrix} \begin{bmatrix} 3x_1 \\ 5x_2 \end{bmatrix} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \vec{x}.$$

The inverse matrix multiplies the first component by $\frac{1}{3}$ and the second component by $\frac{1}{5}$, which effectively undoes what A did.

Example 2 Things get a little more complicated when matrices *mix* the different components of the input vector, as in this example:

$$B = \begin{bmatrix} 1 & 2 \\ 0 & 3 \end{bmatrix}, \quad \text{which acts as } B\vec{x} = \begin{bmatrix} 1 & 2 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} x_1 + 2x_2 \\ 3x_2 \end{bmatrix}.$$

Make sure you understand how to compute $B\vec{x}$ using both the *row picture* and the *column picture* of the matrix-vector product.

The inverse of the matrix B is the matrix

$$B^{-1} = \begin{bmatrix} 1 & -\frac{2}{3} \\ 0 & \frac{1}{3} \end{bmatrix}.$$

Multiplication by the matrix B^{-1} is the “undo action” for multiplication by B :

$$B^{-1}(B\vec{x}) = \begin{bmatrix} 1 & -\frac{2}{3} \\ 0 & \frac{1}{3} \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 & -\frac{2}{3} \\ 0 & \frac{1}{3} \end{bmatrix} \begin{bmatrix} x_1 + 2x_2 \\ 3x_2 \end{bmatrix} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \vec{x}.$$

By definition, the inverse A^{-1} *undoes* the effects of the matrix A . The cumulative effect of applying A^{-1} after A is the *identity matrix* $\mathbb{1}$, which has 1s on the diagonal and 0s everywhere else:

$$A^{-1}A\vec{x} = \mathbb{1}\vec{x} = \vec{x} \quad \Rightarrow \quad A^{-1}A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = \mathbb{1}.$$

Note that $\mathbb{1}\vec{x} = \vec{x}$ for any vector $\vec{x}$.

We'll discuss matrix inverses and how to compute them in more detail later (Section 3.5). For now, it's important you know they exist.

An overview of linear algebra

In the remainder of the book, we'll learn all about the properties of vectors and matrices. Matrix-vector products play an important role in linear algebra because of their relation to *linear transformations*.

Functions are transformations from an input space (the domain) to an output space (the image). A linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a function that takes n -vectors as inputs and produces m -vectors as outputs. If the function T is linear, the output $\vec{y} = T(\vec{x})$ of T applied to $\vec{x}$ can be computed as the matrix-vector product $A_T\vec{x}$, for some matrix $A_T \in \mathbb{R}^{m \times n}$. We say T is *represented by* the matrix A_T . Equivalently, every matrix $A \in \mathbb{R}^{m \times n}$ corresponds to some linear transformation $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$. Given the equivalence between matrices and linear transformations, we can reinterpret the statement “linear algebra is about vectors and matrices” by saying “linear algebra is about vectors and linear transformations.”

You can adapt your existing knowledge about functions to the world of linear transformations. The action of a function on a number is similar to the action of a linear transformation on a vector. The table below summarizes several useful correspondences between functions and linear transformations.

function $f : \mathbb{R} \rightarrow \mathbb{R}$	$\Leftrightarrow$	linear transformation $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ represented by the matrix $A \in \mathbb{R}^{m \times n}$
input $x \in \mathbb{R}$	$\Leftrightarrow$	input $\vec{x} \in \mathbb{R}^n$
output $f(x) \in \mathbb{R}$	$\Leftrightarrow$	output $T_A(\vec{x}) = A\vec{x} \in \mathbb{R}^m$
$g \circ f(x) = g(f(x))$	$\Leftrightarrow$	$T_B(T_A(\vec{x})) = BA\vec{x}$
function inverse f^{-1}	$\Leftrightarrow$	matrix inverse A^{-1}
roots of f	$\Leftrightarrow$	kernel of T_A = null space of $A = \mathcal{N}(A)$
image of f	$\Leftrightarrow$	image of T_A = column space of $A = \mathcal{C}(A)$

Table 2.1: Correspondences between functions and linear transformations.

This table of correspondences serves as a roadmap for the rest of the material in this book. You'll notice the table introduces several new linear algebra concepts like *kernel*, *null space*, and *column space*, but not too many. You can totally do this!

Remember to always connect the new concepts of linear algebra to concepts you're already familiar with. For example, the roots of a function $f(x)$ are the set of inputs for which the function's output is zero. Similarly, the *kernel* of a linear transformation T is the set of inputs that T sends to the zero vector. The roots of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ and the kernel of a linear transformation $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are essentially the same concept; we're just upgrading functions to vector inputs.

In Chapter 1, I explained why functions are useful tools for modelling the real world. Well, linear algebra is the “vector upgrade” to your real-world modelling skills. With linear algebra you'll be able to model complex relationships between multivariable inputs and multivariable outputs. To build modelling skills, you must first develop your geometric intuition about lines, planes, vectors, bases, linear transformations, vector spaces, vector subspaces, etc. It's a lot of work, but the effort you invest will pay dividends.

Links

[Linear algebra lecture series by Prof. Strang from MIT]

<http://bit.ly/1ayRcrj> (row and column picture example)

[A system of equations in the row picture and column picture]

https://www.youtube.com/watch?v=uNKDw46_Ev4

Exercises

E2.1 Find the inverse matrix A^{-1} for the matrix $A = \begin{bmatrix} 7 & 0 \\ 0 & 2 \end{bmatrix}$. Verify that $A^{-1}(A\vec{v}) = \vec{v}$ for any vector $\vec{v} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}$.

E2.2 Given the matrices $A = \begin{bmatrix} 1 & 3 \\ 4 & 5 \end{bmatrix}$ and $B = \begin{bmatrix} -1 & 0 \\ 3 & 3 \end{bmatrix}$, and the vectors $\vec{v} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$ and $\vec{w} = \begin{bmatrix} -3 \\ -4 \end{bmatrix}$, compute the following expressions.

a) $A\vec{v}$ b) $B\vec{v}$ c) $A(B\vec{v})$ d) $B(A\vec{v})$ e) $A\vec{w}$ f) $B\vec{w}$

E2.3 Find the components v_1 and v_2 of the vector $\vec{v} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}$ so that $E\vec{v} = 3\vec{e}_2 - 2\vec{e}_1$, where E is the following matrix:

$$E = \begin{bmatrix} | & | \\ \vec{e}_1 & \vec{e}_2 \\ | & | \end{bmatrix}.$$

What next?

We won't bring geometry, vector spaces, algorithms, and the applications of linear algebra into the mix all at once. Instead, let's start

with the basics. Since linear algebra is about vectors and matrices, let's define vectors and matrices precisely, and describe the math operations we can perform on them.

2.2 Vector operations

Section 1.13 introduced some basic notions about vectors. Understanding vectors is so important for linear algebra that it's worth going beyond the rudimentary understanding of vectors as "directional quantities," and so we took the time to describe vectors more abstractly—as math objects. With vectors defined, our next step is to specify their properties and the operations we can perform on vectors. This is what this section is all about.

Definitions

Consider the vectors $\vec{u} = (u_1, u_2, u_3)$ and $\vec{v} = (v_1, v_2, v_3)$, and an arbitrary constant $\alpha \in \mathbb{R}$. Vector algebra can be summarized as the following operations:

- Addition: $\vec{u} + \vec{v} \stackrel{\text{def}}{=} (u_1 + v_1, u_2 + v_2, u_3 + v_3)$
- Subtraction: $\vec{u} - \vec{v} \stackrel{\text{def}}{=} (u_1 - v_1, u_2 - v_2, u_3 - v_3)$
- Scaling: $\alpha \vec{u} \stackrel{\text{def}}{=} (\alpha u_1, \alpha u_2, \alpha u_3)$
- Dot product: $\vec{u} \cdot \vec{v} \stackrel{\text{def}}{=} u_1 v_1 + u_2 v_2 + u_3 v_3$
- Cross product: $\vec{u} \times \vec{v} \stackrel{\text{def}}{=} (u_2 v_3 - u_3 v_2, u_3 v_1 - u_1 v_3, u_1 v_2 - u_2 v_1)$
- Length: $\|\vec{u}\| \stackrel{\text{def}}{=} \sqrt{u_1^2 + u_2^2 + u_3^2}$

In the next few pages we'll see what these operations can do for us.

Notation

The set of real numbers is denoted $\mathbb{R}$. An n -dimensional real vector consists of n real numbers slapped together in a bracket. We denote the set of three-dimensional vectors as $\mathbb{R}^3 \stackrel{\text{def}}{=} (\mathbb{R}, \mathbb{R}, \mathbb{R})$. Similarly, the set of n -dimensional real vectors is denoted $\mathbb{R}^n$.

When learning about new math operations, it's important to keep track of the types of inputs and the types of outputs of each operation. In computer science, this information is called the *type signature*. For example, the type signature " $\text{op} : \mathbb{R} \rightarrow \mathbb{R}$ " tells us the math operation op takes real numbers as inputs and produces real numbers as outputs. Certain operations take pairs of numbers as inputs, like addition for example: $\text{add}(a, b)$ (usually denoted $a + b$). The type

signature for addition is “ $\text{add} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$,” which tells us addition takes two real numbers as inputs and produces real numbers as outputs. The symbol “ $\times$ ” denotes the *Cartesian product* of two sets. The Cartesian product $\mathbb{R} \times \mathbb{R}$ contains all possible pairs of real numbers (a, b) , where $a \in \mathbb{R}$ and $b \in \mathbb{R}$. Note the math symbol “ $\times$ ” has two different meanings depending on the objects that surround it. The expression $\vec{u} \times \vec{v}$ denotes the cross product of vectors $\vec{u}$ and $\vec{v}$, while the expression $A \times B$ denotes the Cartesian product of sets A and B .

The use of such formal math notation may seem complicated at first, but it’s very helpful when defining new math operations.

Addition and subtraction

Addition and subtraction take pairs of vectors as inputs and produce vectors as outputs:

$$+ : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n \quad \text{and} \quad - : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n.$$

Addition and subtraction are performed component-wise:

$$\vec{w} = \vec{u} \pm \vec{v} \quad \Leftrightarrow \quad w_i = u_i \pm v_i, \quad \forall i \in [1, \dots, n].$$

Scaling by a constant

Scaling is an operation that takes a number and a vector as inputs and produces a vector output:

$$\text{scalar-mult} : \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}^n.$$

There is no symbol to denote scalar multiplication—we just write the scale factor in front of the vector and the multiplication is implicit.

Multiplying the vector $\vec{u}$ by the scale factor α is equivalent to multiplying each component of the vector by α :

$$\vec{w} = \alpha \vec{u} \quad \Leftrightarrow \quad w_i = \alpha u_i.$$

For example, choosing $\alpha = 2$, we obtain the vector $\vec{w} = 2\vec{u}$, which is two times longer than the vector $\vec{u}$:

$$\vec{w} = (w_1, w_2, w_3) = (2u_1, 2u_2, 2u_3) = 2(u_1, u_2, u_3) = 2\vec{u}.$$

Vector products

We’ll now define the *dot product* and the *cross product*: two geometric operations useful for working with three-dimensional vectors.




Figure 2.1: Vectors $\vec{u}$ and $\vec{w}$ are related by the equation $\vec{w} = 2\vec{u}$.

Dot product

The *dot product* takes two vectors as inputs and produces a single, real number as an output:

$$\cdot : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}.$$

The dot product between the vector $\vec{v} = (v_x, v_y, v_z)$ and the vector $\vec{w} = (w_x, w_y, w_z)$ can be computed using either the algebraic formula,

$$\vec{v} \cdot \vec{w} = v_x w_x + v_y w_y + v_z w_z,$$

or the geometric formula,

$$\vec{v} \cdot \vec{w} = \|\vec{v}\| \|\vec{w}\| \cos(\varphi),$$

where φ is the angle between the two vectors. Note the value of the dot product depends on the vectors' lengths and the cosine of the angle between them.

The dot product is the key tool for calculating vector projections, vector decompositions, and determining orthogonality. The name *dot product* comes from the symbol used to denote it. It is also known as the *scalar product*, since the result of the dot product is a scalar number—a number that does not change when the basis changes. The dot product is also sometimes called the *inner product*.

We can combine the algebraic and the geometric formulas for the dot product to obtain the formula,

$$\cos(\varphi) = \frac{\vec{v} \cdot \vec{w}}{\|\vec{v}\| \|\vec{w}\|} = \frac{v_x w_x + v_y w_y + v_z w_z}{\|\vec{v}\| \|\vec{w}\|} \quad \text{and} \quad \varphi = \cos^{-1}(\cos(\varphi)).$$

This formula makes it possible to find the angle between two vectors if we know their components.

The geometric factor $\cos(\varphi)$ depends on the relative orientation of the two vectors as follows:

- If the vectors point in the same direction, then $\cos(\varphi) = \cos(0^\circ) = 1$, so $\vec{v} \cdot \vec{w} = \|\vec{v}\| \|\vec{w}\|$.
- If the vectors are perpendicular to each other, then $\cos(\varphi) = \cos(90^\circ) = 0$, so $\vec{v} \cdot \vec{w} = 0$.
- If the vectors point in exactly opposite directions, then $\cos(\varphi) = \cos(180^\circ) = -1$, so $\vec{v} \cdot \vec{w} = -\|\vec{v}\| \|\vec{w}\|$.

The dot product is defined for vectors of any dimension; as long as two vectors are defined with respect to the same basis, we can compute the dot product between them.

Cross product

The *cross product* takes two vectors as inputs and produces another vector as the output:

$$\times : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}^3.$$

The cross product of two vectors is perpendicular to both vectors:

$$\vec{v} \times \vec{w} = \{ \text{a vector perpendicular to both } \vec{v} \text{ and } \vec{w} \} \in \mathbb{R}^3.$$

If you take the cross product of one vector pointing in the x -direction with another vector pointing in the y -direction, the result will be a vector in the z -direction: $\hat{i} \times \hat{j} = \hat{k}$. The name *cross product* comes from the symbol used to denote it. It is also sometimes called the *vector product*, since the output of this operation is a vector.

The cross products of individual basis elements are defined as

$$\hat{i} \times \hat{j} = \hat{k}, \quad \hat{j} \times \hat{k} = \hat{i}, \quad \hat{k} \times \hat{i} = \hat{j}.$$

Look at Figure 1.55 on page 92 and imagine the vectors $\hat{i}$, $\hat{j}$, and $\hat{k}$ pointing along each axis. Try to visualize the three equations above.

The cross product is *anticommutative*, which means swapping the order of the inputs introduces a negative sign in the output:

$$\hat{j} \times \hat{i} = -\hat{k}, \quad \hat{k} \times \hat{j} = -\hat{i}, \quad \hat{i} \times \hat{k} = -\hat{j}.$$

It's likely that, until now, the products you've seen in math have been *commutative*, which means the order of the inputs doesn't matter. The product of two numbers is commutative $ab = ba$, and the dot product of two vectors is commutative $\vec{u} \cdot \vec{v} = \vec{v} \cdot \vec{u}$, but the cross product of two vectors is *anticommutative* $\vec{v} \times \vec{w} = -\vec{w} \times \vec{v}$.

Given two vectors $\vec{a} = a_x\hat{i} + a_y\hat{j} + a_z\hat{k}$ and $\vec{b} = b_x\hat{i} + b_y\hat{j} + b_z\hat{k}$, their cross product is calculated as

$$\vec{a} \times \vec{b} = (a_y b_z - a_z b_y)\hat{i} + (a_z b_x - a_x b_z)\hat{j} + (a_x b_y - a_y b_x)\hat{k}.$$

Computing the cross product requires a specific combination of multiplications and subtractions of the input vectors' components. The result of this combination is the vector $\vec{a} \times \vec{b}$ which is perpendicular to both $\vec{a}$ and $\vec{b}$.

The length of the cross product of two vectors is proportional to the sine of the angle between the two vectors:

$$\|\vec{a} \times \vec{b}\| = \|\vec{a}\|\|\vec{b}\| \sin(\varphi).$$

The right-hand rule

Consider the plane formed by the vectors $\vec{a}$ and $\vec{b}$. There are actually *two* vectors perpendicular to this plane: one above the plane and one below the plane. We use the *right-hand rule* to figure out which of these vectors corresponds to the cross product $\vec{a} \times \vec{b}$.

Make a fist with your right hand and then extend your thumb, first finger, and middle finger. When your index finger points in the same direction as the vector $\vec{a}$ and your middle finger points in the direction of $\vec{b}$, your thumb will point in the direction of $\vec{a} \times \vec{b}$. The relationship encoded in the right-hand rule matches the relationship between the standard basis vectors: $\hat{i} \times \hat{j} = \hat{k}$.




Figure 2.2: Using the right-hand rule to determine the direction of the cross product $\vec{a} \times \vec{b}$ based on directions of $\vec{a}$ and $\vec{b}$.

Links

[Nice illustrations of the cross product]

<http://lucasvb.tumblr.com/post/76812811092/>

<https://www.youtube.com/watch?v=eu6i7WJeinw>

Length of a vector

The length of a vector is a nonnegative number that describes the extent of the vector in space, and is sometimes referred to as the vector's *magnitude* or the *norm*. We express the notion of length as an n -dimensional extension of Pythagoras' formula: it is the length of the hypotenuse in a right-angle triangle, given the lengths of its two sides, just applied to n dimensions. The length of a vector $\vec{u}$ is denoted $\|\vec{u}\|$ or $\|\vec{u}\|_2$ or sometimes simply u .

The length of the vector $\vec{u} \in \mathbb{R}^n$ is computed as follows:

$$\|\vec{u}\| = \sqrt{u_1^2 + u_2^2 + \cdots + u_n^2} = \sqrt{\vec{u} \cdot \vec{u}}.$$

There are many mathematical concepts that correspond to the intuitive notion of length. The formula above computes the *Euclidian length* (or *Euclidian norm*) of the vector. Another name for the Euclidian length is the ℓ^2 -norm (pronounced *ell-two norm*²).

Note a vector's length can also be computed as the square root of the dot product of the vector with itself: $\|\vec{v}\| = \sqrt{\vec{v} \cdot \vec{v}}$. Indeed, there is a deep mathematical connection between norms and inner products.

Unit vectors

A *unit vector* is a vector that has length one. Given a vector $\vec{v}$ of any length, we can build a unit vector in the same direction by dividing $\vec{v}$ by its own length:

$$\hat{v} = \frac{\vec{v}}{\|\vec{v}\|}.$$

Unit vectors are useful in many contexts. For instance, when we want to specify a direction in space, we use a unit vector in that direction.

Projection

Pop-quiz time! Can you remember some essentials from Section 1.13? Suppose I give you a direction $\hat{d}$ and some vector $\vec{v}$, and ask you how much of $\vec{v}$ is in the direction $\hat{d}$? To find the answer, you must compute the dot product:

$$v_d = \hat{d} \cdot \vec{v} = \|\hat{d}\| \|\vec{v}\| \cos \theta = 1 \|\vec{v}\| \cos \theta,$$

²The name ℓ^2 -norm refers to the process of squaring each of the vector's components, and then taking the square root. Another norm is the ℓ^4 -norm, defined as the fourth root of the sum of the vector's components raised to the fourth power:

$$\|\vec{u}\|_4 \stackrel{\text{def}}{=} \sqrt[4]{u_1^4 + u_2^4 + u_3^4}.$$

where θ is the angle between $\vec{v}$ and $\hat{d}$. This formula is used in physics to compute x -components of forces: $F_x = \vec{F} \cdot \hat{i} = \|\vec{F}\| \cos \theta$.

Define the *projection* of a vector $\vec{v}$ in the $\hat{d}$ direction as follows:

$$\Pi_{\hat{d}}(\vec{v}) = v_d \hat{d} = (\hat{d} \cdot \vec{v}) \hat{d}.$$

If the direction is specified by a vector $\vec{d}$ that doesn't have length one, then the projection formula becomes

$$\Pi_{\vec{d}}(\vec{v}) = \left(\frac{\vec{d} \cdot \vec{v}}{\|\vec{d}\|^2} \right) \vec{d}.$$

Division by the length squared transforms the two appearances of the vector $\vec{d}$ into the unit vectors $\hat{d}$ needed for the projection formula:

$$\Pi_{\hat{d}}(\vec{v}) = \underbrace{(\vec{v} \cdot \hat{d})}_{\|\vec{v}\| \cos \theta} \hat{d} = \left(\vec{v} \cdot \frac{\vec{d}}{\|\vec{d}\|} \right) \frac{\vec{d}}{\|\vec{d}\|} = \left(\frac{\vec{v} \cdot \vec{d}}{\|\vec{d}\|^2} \right) \vec{d} = \Pi_{\vec{d}}(\vec{v}).$$

Remember these projection formulas well because we'll need them later: when computing projections onto planes (Section 4.2), when computing vector coordinates (Section 4.3), and when describing the change-of-basis operation (Section 5.3).

Discussion

This section elaborated on the properties of n -dimensional vectors, which are ordered tuples (lists) of n components. It's important to think of vectors as whole mathematical objects, rather than as components. Although vector operations boil down to manipulations of their components, vectors are most useful (and best understood) when you think of them as whole objects that have components, rather than focussing on their components.

Links

[Nice illustration of the cross product]

<http://lucasvb.tumblr.com/post/76812811092/>

[Vectors explained by 3Blue1Brown]

https://youtube.com/watch?v=fNk_zzaMoSs

[Cross products explained by 3Blue1Brown]

<https://youtube.com/watch?v=eu6i7WJeiw>

Exercises

E2.4 Given the vectors $\vec{u} = (1, 1, 0)$ and $\vec{v} = (0, 0, 3)$, compute the following vector expressions:

a) $\vec{u} + \vec{v}$ b) $\vec{u} - \vec{v}$ c) $3\vec{u} + \vec{v}$ d) $\|\vec{u}\|$

E2.5 Given $\vec{v} = (1, 2, 3)$ and $\vec{w} = (0, 1, 1)$, compute the following vector products: a) $\vec{v} \cdot \vec{w}$; b) $\vec{v} \times \vec{w}$; c) $\vec{w} \times \vec{v}$; d) $\vec{w} \times \vec{w}$.

E2.6 For each of the following vectors, $\vec{v}_1 = 10\angle 10^\circ$, $\vec{v}_2 = 10\angle 30^\circ$, $\vec{v}_3 = 10\angle 60^\circ$, $\vec{v}_4 = 10\angle 120^\circ$, complete the following tasks:

- Draw the vector in a Cartesian plane.
- Compute the vector's x - and y -coordinates.
- Compute the projection of the vector in the direction $\hat{i}$. Your answer should be a vector quantity.
- Compute the projection of the vector in the direction $\hat{j}$.
- Compute the projection of the vector in the direction $\vec{d} = (1, 1)$, and find the length of the projection.

Hint: Recall the formula for the projection of the vector $\vec{v}$ in the direction $\vec{d}$ is defined as $\Pi_{\vec{d}}(\vec{v}) = \left(\frac{\vec{d} \cdot \vec{v}}{\|\vec{d}\|^2} \right) \vec{d}$.

2.3 Matrix operations

A matrix is a two-dimensional array (a table) of numbers. Consider the m by n matrix $A \in \mathbb{R}^{m \times n}$. We denote the matrix as a whole A and refer to its individual entries as a_{ij} , where a_{ij} is the number in the i^{th} row and the j^{th} column of A . What are the mathematical operations we can perform on this matrix?

Addition and subtraction

The matrix addition and subtraction operations take pairs of matrices as inputs and produce matrices as outputs:

$$+ : \mathbb{R}^{m \times n} \times \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n} \quad \text{and} \quad - : \mathbb{R}^{m \times n} \times \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}.$$

Addition and subtraction are performed as follows:

$$C = A \pm B \quad \Leftrightarrow \quad c_{ij} = a_{ij} \pm b_{ij}, \quad \forall i \in [1, \dots, m], j \in [1, \dots, n].$$

For example, addition for two 3×2 matrices is expressed as

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} + \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ b_{31} & b_{32} \end{bmatrix} = \begin{bmatrix} a_{11} + b_{11} & a_{12} + b_{12} \\ a_{21} + b_{21} & a_{22} + b_{22} \\ a_{31} + b_{31} & a_{32} + b_{32} \end{bmatrix}.$$

Matrices must have the same dimensions to be added or subtracted.

Multiplication by a constant

Recall that *scaling* is another word for multiplication by a constant. Given a number α and a matrix A , we can *scale* A by α as follows:

$$\alpha A = \alpha \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} = \begin{bmatrix} \alpha a_{11} & \alpha a_{12} \\ \alpha a_{21} & \alpha a_{22} \\ \alpha a_{31} & \alpha a_{32} \end{bmatrix}.$$

Matrix-vector multiplication

The result of the matrix-vector product between a matrix $A \in \mathbb{R}^{m \times n}$ and a vector $\vec{v} \in \mathbb{R}^n$ is an m -dimensional vector:

$$\text{matrix-vector product} : \mathbb{R}^{m \times n} \times \mathbb{R}^n \rightarrow \mathbb{R}^m.$$

The formula for the matrix-vector product is

$$\vec{w} = A\vec{v} \quad \Leftrightarrow \quad w_i = \sum_{j=1}^n a_{ij}v_j, \quad \forall i \in [1, \dots, m].$$

For example, the product of a 3×2 matrix A and the 2×1 column vector $\vec{v}$ results in a 3×1 vector:

$$\begin{aligned} A\vec{v} &= \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = v_1 \underbrace{\begin{bmatrix} a_{11} \\ a_{21} \\ a_{31} \end{bmatrix}}_{\text{column picture}} + v_2 \underbrace{\begin{bmatrix} a_{12} \\ a_{22} \\ a_{32} \end{bmatrix}}_{\text{column picture}} \\ &= \left\{ \begin{bmatrix} (a_{11}, a_{12}) \cdot \vec{v} \\ (a_{21}, a_{22}) \cdot \vec{v} \\ (a_{31}, a_{32}) \cdot \vec{v} \end{bmatrix} \right\}_{\text{row picture}} \\ &= \begin{bmatrix} a_{11}v_1 + a_{12}v_2 \\ a_{21}v_1 + a_{22}v_2 \\ a_{31}v_1 + a_{32}v_2 \end{bmatrix} \in \mathbb{R}^{3 \times 1}. \end{aligned}$$

Note the two equivalent ways to understand the matrix-vector product: the *column picture* and the *row picture*. In the column picture, the multiplication of the matrix A by the vector $\vec{v}$ produces a **linear combination of the columns of the matrix**. In the row picture, multiplication of the matrix A by the vector $\vec{v}$ produces a column vector with components equal to the **dot products of the rows of the matrix A with the vector $\vec{v}$** .

Matrix-matrix multiplication

The matrix product AB of matrices $A \in \mathbb{R}^{m \times \ell}$ and $B \in \mathbb{R}^{\ell \times n}$ results in an $m \times n$ matrix:

$$\text{matrix-product} : \mathbb{R}^{m \times \ell} \times \mathbb{R}^{\ell \times n} \rightarrow \mathbb{R}^{m \times n}.$$

The formula for matrix multiplication computes the dot product between each row of A and each column of B :

$$C = AB \quad \Leftrightarrow \quad c_{ij} = \sum_{k=1}^{\ell} a_{ik}b_{kj}, \forall i \in [1, \dots, m], j \in [1, \dots, n].$$

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \begin{bmatrix} a_{11}b_{11} + a_{12}b_{21} & a_{11}b_{12} + a_{12}b_{22} \\ a_{21}b_{11} + a_{22}b_{21} & a_{21}b_{12} + a_{22}b_{22} \\ a_{31}b_{11} + a_{32}b_{21} & a_{31}b_{12} + a_{32}b_{22} \end{bmatrix} \in \mathbb{R}^{3 \times 2}.$$

Transpose

The transpose matrix A^T is defined by the formula $a_{ij}^T = a_{ji}$. We obtain the transpose by “flipping” the matrix through its diagonal:

$$^T : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{n \times m}$$

$$\begin{bmatrix} \alpha_1 & \alpha_2 & \alpha_3 \\ \beta_1 & \beta_2 & \beta_3 \end{bmatrix}^T = \begin{bmatrix} \alpha_1 & \beta_1 \\ \alpha_2 & \beta_2 \\ \alpha_3 & \beta_3 \end{bmatrix}.$$

Note that entries on the diagonal of the matrix do not change when we apply the transpose operation.

Properties of the transpose operation

- $(A + B)^T = A^T + B^T$
- $(AB)^T = B^T A^T$
- $(ABC)^T = C^T B^T A^T$
- $(A^T)^{-1} = (A^{-1})^T$

Vectors as matrices

A vector is a special type of matrix. You can treat a vector $\vec{v} \in \mathbb{R}^n$ either as a *column vector* ($n \times 1$ matrix) or as a *row vector* ($1 \times n$ matrix).

Inner product

Recall the definition of the *dot product* or *inner product* for vectors:

$$\cdot : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R} \quad \Leftrightarrow \quad \vec{u} \cdot \vec{v} = \sum_{i=1}^n u_i v_i.$$

If we think of vectors as *column* vectors, we can write the dot product in terms of the matrix transpose operation ^{T} and the standard rules of matrix multiplication:

$$\vec{u} \cdot \vec{v} = \vec{u}^T \vec{v} = \begin{bmatrix} u_1 & u_2 & u_3 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = u_1 v_1 + u_2 v_2 + u_3 v_3.$$

The dot product for vectors is thus a special case of matrix multiplication. Alternatively, we could say that matrix multiplication is defined in terms of the dot product.

Outer product

Consider again two *column* vectors $\vec{u}$ and $\vec{v}$ ($n \times 1$ matrices). We obtain the inner product by applying the transpose to the *first* vector in the product: $\vec{u}^T \vec{v} = \vec{u} \cdot \vec{v}$. Instead, if we apply the transpose to the *second* vector, we'll obtain the *outer product* of $\vec{u}$ and $\vec{v}$. The outer product operation takes pairs of vectors as inputs and produces matrices as outputs:

$$\text{outer-product} : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}.$$

For example, the outer product of two vectors in $\mathbb{R}^3$ is

$$\vec{u} \vec{v}^T = \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} \begin{bmatrix} v_1 & v_2 & v_3 \end{bmatrix} = \begin{bmatrix} u_1 v_1 & u_1 v_2 & u_1 v_3 \\ u_2 v_1 & u_2 v_2 & u_2 v_3 \\ u_3 v_1 & u_3 v_2 & u_3 v_3 \end{bmatrix} \in \mathbb{R}^{3 \times 3}.$$

Observe that the matrix-matrix product of a 3×1 matrix and a 1×3 matrix results in a 3×3 matrix.

In Section 4.2 we'll see how the outer product is used to build *projection matrices*. For example, the matrix that corresponds to the projection onto the x -axis is $M_x = \hat{i} \hat{i}^T \in \mathbb{R}^{3 \times 3}$. The x -projection of any vector $\vec{v}$ is computed as the matrix-vector product, $M_x \vec{v} = \hat{i} \hat{i}^T \vec{v} = \hat{i}(\hat{i} \cdot \vec{v}) = v_x \hat{i}$. More on that later.

Matrix inverses

Multiplying an invertible matrix A by its inverse A^{-1} produces the identity matrix: $AA^{-1} = \mathbb{1} = A^{-1}A$. The *identity matrix* obeys $\mathbb{1}\vec{v} = \vec{v}$ for all vectors $\vec{v}$. The inverse matrix A^{-1} *undoes* whatever A did. The cumulative effect of multiplying by A and A^{-1} is equivalent to the identity transformation,

$$A^{-1}(A(\vec{v})) = (A^{-1}A)\vec{v} = \mathbb{1}\vec{v} = \vec{v}.$$

We can think of “finding the inverse” ($\text{inv}(A) = A^{-1}$) as an operation of the form

$$\text{inv} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}.$$

Note that only *invertible* matrices have an inverse. Some matrices are not invertible—there is no “undo” operation for them. We’ll postpone the detailed discussion of invertibility until Section 5.4.

Properties of matrix inverse operation

- $(AB)^{-1} = B^{-1}A^{-1}$
- $(ABC)^{-1} = C^{-1}B^{-1}A^{-1}$
- $(A^T)^{-1} = (A^{-1})^T$
- $(A^{-1})^{-1} = A$

The matrix inverse plays the role of “division by the matrix A ” in matrix equations. We’ll discuss matrix equations in Section 3.2.

Trace

The *trace* of an $n \times n$ matrix is the sum of the n values on its diagonal:

$$\text{Tr} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}, \quad \text{Tr}[A] \stackrel{\text{def}}{=} \sum_{i=1}^n a_{ii}.$$

Properties of the trace operation

- $\text{Tr}[\alpha A + \beta B] = \alpha \text{Tr}[A] + \beta \text{Tr}[B]$ (linear property)
- $\text{Tr}[AB] = \text{Tr}[BA]$
- $\text{Tr}[ABC] = \text{Tr}[CAB] = \text{Tr}[BCA]$ (cyclic property)
- $\text{Tr}[A^T] = \text{Tr}[A]$
- $\text{Tr}[A] = \sum_{i=1}^n \lambda_i$, where $\{\lambda_i\}$ are the eigenvalues of A

Determinant

The *determinant* of a matrix is a calculation that involves all the entries of the matrix, and whose output is a single number:

$$\det : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}.$$

The determinant describes the relative geometry of the vectors that make up the rows of the matrix. More specifically, the determinant of a matrix A tells you the *volume* of a box with sides given by rows of A .

For example, the determinant of a 2×2 matrix is

$$\det(A) = \det\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) = \begin{vmatrix} a & b \\ c & d \end{vmatrix} = ad - bc.$$

The quantity $ad - bc$ corresponds to the area of the parallelogram formed by the vectors (a, b) and (c, d) . Observe that if the rows of A point in the same direction, $(a, b) = \alpha(c, d)$ for some $\alpha \in \mathbb{R}$, then the area of the parallelogram will be zero. If the determinant of a matrix is nonzero, then the rows of that matrix are linearly independent.

Properties of determinants

- $\det(AB) = \det(A) \det(B)$
- $\det(A) = \prod_{i=1}^n \lambda_i$, where $\{\lambda_i\}$ are the eigenvalues of A
- $\det(A^T) = \det(A)$
- $\det(A^{-1}) = \frac{1}{\det(A)}$

Discussion

Understanding vector and matrix operations is essential for understanding more advanced theoretical topics and the applications of linear algebra. In the remainder of this book, you'll learn about various algebraic and geometric interpretations of the matrix operations defined in this section. Seeing all these definitions at the same time can be overwhelming, I know; but the good news is that we've defined all the math operations and notation we'll need for the rest of the book. It could be worse, right?

So far, we've defined two of the main actors in linear algebra: vectors and matrices. But our introduction to linear algebra won't be complete until we introduce *linearity*, the main thread that runs through all the topics in this book.

Exercises

E2.7 Given the matrices $A = \begin{bmatrix} 3 & 4 \\ 2 & 1 \end{bmatrix}$, $B = \begin{bmatrix} -1 & 0 & 1 & 2 \\ 4 & 3 & 2 & 1 \end{bmatrix}$, and $C = \begin{bmatrix} -2 & 3 & 0 \\ 2 & -2 & 1 \end{bmatrix}$, compute the expressions.

- a) A^T b) C^T c) A^2 d) AB e) AC f) BA g) $C^T A$
 h) $\det(A)$ i) $\det(B)$ j) $\det(C)$ k) $\det(A^T)$ l) $\det(AA^{-1})$
 m) $\text{Tr}(A)$ n) $\text{Tr}(A^T)$

Hint: Some of these expressions may not exist.

E2.8 Given the 1×3 matrices (row vectors) $\vec{u} = (1, 2, 3)$ and $\vec{v} = (2, -1, 0)$, compute the following products:

- a) $\vec{u} \vec{u}^T$ b) $\vec{v} \vec{v}^T$ c) $\vec{u} \vec{v}^T$ d) $\vec{u}^T \vec{u}$ e) $\vec{v}^T \vec{v}$ f) $\vec{u}^T \vec{v}$

Hint: The transpose of a 1×3 row vector is a 3×1 column vector.

E2.9 Find the unknowns α and β in the equation $\begin{bmatrix} 2 & \alpha \\ \beta & -3 \end{bmatrix} \begin{bmatrix} 1 \\ 4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$.

2.4 Linearity

What is linearity and why is this entire book dedicated to learning about it? Consider an arbitrary function that contains terms with different powers of the input variable x :

$$f(x) = \underbrace{a/x}_{\text{one-over-}x} + \underbrace{b}_{\text{constant}} + \underbrace{mx}_{\text{linear term}} + \underbrace{qx^2}_{\text{quadratic}} + \underbrace{cx^3}_{\text{cubic}}.$$

The term mx is the *linear* term in this expression—it contains x raised to the first power. All the other terms are *nonlinear*. The linear term is special because changes in the value of the input x lead to proportional changes in the value of mx . If the input is $2x$, the linear term will have value $2mx$. If the input is $100x$, the linear term will have value $100mx$. This input-output proportionality does not hold for nonlinear terms.

In this section we'll discuss the special properties of expressions and equations containing only linear terms.

Introduction

A single-variable function takes as its input a real number x and outputs a real number $f(x)$. The type signature of this function is

$$f: \mathbb{R} \rightarrow \mathbb{R}.$$

The most general linear function from $\mathbb{R}$ to $\mathbb{R}$ looks like this:

$$f(x) = mx,$$

where $m \in \mathbb{R}$ is called the *coefficient* of x . The action of a linear function is to multiply the input by the constant m . So far, so good.

Example of composition of linear functions Given the linear functions $f(x) = 2x$ and $g(x) = 3x$, what is the equation of the function $h(x) \stackrel{\text{def}}{=} g \circ f(x) = g(f(x))$? The composition of $f(x) = 2x$ and $g(x) = 3x$ is the function $h(x) = g(f(x)) = 3(2x) = 6x$. Note the composition of two linear functions is also a linear function. The coefficient of h is equal to the product of the coefficients of f and g .

Definition

Linear functions map any linear combination of inputs to the same linear combination of outputs. A function f is *linear* if it satisfies the equation

$$f(\alpha x_1 + \beta x_2) = \alpha f(x_1) + \beta f(x_2),$$

for any two inputs x_1 and x_2 , and for all constants α and β .

Lines are not linear functions!

Consider the equation of a line:

$$l(x) = mx + b,$$

where the constant m corresponds to the slope of the line, and the constant $b = f(0)$ is its initial value. A *line* $l(x) = mx + b$ with $b \neq 0$ is *not* a linear function. This logic may seem a bit weird, but if you don't trust me, you can check for yourself:

$$l(\alpha x_1) = m(\alpha x_1) + b \quad \neq \quad \alpha(mx_1 + b) = \alpha l(x_1).$$

A function with a linear part added to a constant term is called an *affine transformation*. Affine transformations are cool but a bit off-topic, since this book focuses on *linear* transformations.

Multivariable functions

The study of linear algebra is the study of *all* things linear. In particular, we'll learn how to work with functions that take multiple

variables as inputs. Consider the set of functions that take pairs of real numbers as inputs and produce real numbers as outputs:

$$f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}.$$

The most general linear function of two variables is

$$f(x, y) = m_x x + m_y y.$$

You can think of m_x as the x -slope and m_y as the y -slope of the function. We say m_x is the x -coefficient and m_y the y -coefficient in the linear expression $m_x x + m_y y$.

Linear expressions

A *linear expression* in the variables x_1 , x_2 , and x_3 has the form,

$$a_1 x_1 + a_2 x_2 + a_3 x_3,$$

where a_1 , a_2 , and a_3 are arbitrary constants. Note the new terminology “linear in v ,” which refers to an expression in which the variable v is raised to the first power. The expression $\frac{1}{a}x_1 + b^6 x_2 + \sqrt{c} x_3$ contains nonlinear factors ($\frac{1}{a}$, b^6 , and $\sqrt{c}$) but is a linear expression in the variables x_1 , x_2 , and x_3 .

Linear equations

A *linear equation* in the variables x_1 , x_2 , and x_3 has the form

$$a_1 x_1 + a_2 x_2 + a_3 x_3 = c.$$

This equation is linear because it contains only linear terms in the variables x_1 , x_2 , and x_3 .

Example Linear equations are very versatile. Suppose you know that the following equation describes some real-world phenomenon:

$$4k - 2m + 8p = 10,$$

where k , m , and p correspond to three variables of interest. You can interpret this equation as describing the variable m as a function of the variables k and p , and rewrite the equation as

$$m(k, p) = 2k + 4p - 5.$$

Using this function, you can predict the value of m given the knowledge of the quantities k and p .

Another option would be to interpret k as a function of m and p : $k(m, p) = \frac{10}{4} + \frac{m}{2} - 2p$. This model would be useful if you know the quantities m and p and want to predict the value of the variable k .

Applications

Geometric interpretation of linear equations

The linear equation in x and y ,

$$ax + by = c, \quad b \neq 0,$$

corresponds to a line with the equation $y(x) = mx + y_0$ in the Cartesian plane. The slope of the line is $m = -a/b$ and its y -intercept is $y_0 = c/b$. The special case when $b = 0$ corresponds to a vertical line with equation $x = \frac{c}{a}$.

The most general linear equation in x , y , and z ,

$$ax + by + cz = d,$$

corresponds to the equation of a plane in a three-dimensional space. Assuming $c \neq 0$, we can rewrite this equation so z (the “height”) is a function of the coordinates x and y : $z(x, y) = z_0 + m_x x + m_y y$. The slope of the plane in the x -direction is $m_x = -\frac{a}{c}$ while the slope in the y -direction is $m_y = -\frac{b}{c}$. The z -intercept of this plane is $z_0 = \frac{d}{c}$.

First-order approximations

When we use a linear function as a mathematical model for a nonlinear, real-world input-output process, we say the function represents a *linear model* or a *first-order approximation* for the process. Let’s analyze what this means in more detail, and see why linear models are so popular in science.

In calculus, we learn that functions can be represented as infinite Taylor series:

$$f(x) = \text{taylor}(f(x)) = a_0 + a_1 t + a_2 t^2 + a_3 t^3 + \cdots = \sum_{n=0}^{\infty} a_n x^n,$$

where the coefficient a_n depends on the n^{th} derivative of $f(x)$. The Taylor series is only equal to the function $f(x)$ if infinitely many terms in the series are calculated. If we sum together only a finite number of terms in the series, we obtain a *Taylor series approximation*. The first-order Taylor series approximation to $f(x)$ is

$$f(x) \approx \text{taylor}_1(f(x)) = a_0 + a_1 x = f(0) + f'(0)x.$$

The above equation describes the best approximation to $f(x)$ near $x = 0$, by a line of the form $l(x) = mx + b$. To build a linear model $f(x)$ of a real-world process, it is sufficient to measure two parameters: the initial value $b = f(0)$ and the rate of change $m = f'(0)$.

Scientists routinely use linear models because this kind of model allows for easy *parametrization*. To build a linear model, the first step is to establish the initial value $f(0)$ by inputting $x = 0$ to the process and seeing what comes out. Next, we vary the input by some amount Δx , and observe the resulting change in the output Δf . The rate of change parameter is equal to the change in the output divided by the change in the input $m = \frac{\Delta f}{\Delta x}$. Thus, we can obtain the parameters of a linear model in two simple steps. In contrast, finding the parametrization of nonlinear models is a more complicated task.

For a function $F(x, y, z)$ that takes three variables as inputs, the first-order Taylor series approximation is

$$F(x, y, z) \approx b + m_x x + m_y y + m_z z.$$

Except for the constant term, the function has the form of a linear expression. The “first-order approximation” to a function of n variables $F(x_1, x_2, \dots, x_n)$ has the form $b + m_1 x_1 + m_2 x_2 + \dots + m_n x_n$.

As in the single-variable case, finding the parametrization of a multivariable linear model is a straightforward task. Suppose we want to model some complicated real-world phenomenon that has n input variables. First, we input only zeros to obtain the initial value $F(0, \dots, 0) = b$. Next, we go through each of the input variables one-by-one and measure how a small change in each input Δx_i affects the output Δf . The rate of change with respect to the input x_i is $m_i = \frac{\Delta f}{\Delta x_i}$. By combining the knowledge of the initial value b and the “slopes” with respect to each input parameter, we’ll obtain a complete linear model of the phenomenon.

Discussion

In the next three chapters, we’ll learn about more mathematical objects and mathematical operations. Linear algebra is the study of vectors, matrices, linear transformations, vector spaces, and other abstract, vector-like objects. The mathematical operations we’ll perform on these objects will be linear: $f(\alpha \mathbf{obj}_1 + \beta \mathbf{obj}_2) = \alpha f(\mathbf{obj}_1) + \beta f(\mathbf{obj}_2)$. Linearity is the core assumption of linear algebra.

Exercises

E2.10 Are these expressions linear in the variables x , y , and z ?

a) $2x + 5y + \sqrt{m}z$ **b)** $10\sqrt{x} + 2(y + z)$ **c)** $42x + \alpha^2 \sin(\frac{\pi}{3})y + z \cos(\frac{\pi}{3})$

2.5 Overview of linear algebra

In linear algebra, you'll learn new computational techniques and develop new ways of thinking about math. With these new tools, you'll be able to use linear algebra techniques for many applications. Let's look at what lies ahead in this book.

Computational linear algebra

The first steps toward understanding linear algebra will seem a little tedious. In Chapter 3 you'll develop basic skills for manipulating vectors and matrices. Matrices and vectors have many components and performing operations on them involves many arithmetic steps—there is no way to circumvent this complexity. Make sure you understand the basic algebra rules (how to add, subtract, and multiply vectors and matrices) because they are a prerequisite for learning more advanced material. You should be able to perform all the matrix algebra operations with pen and paper for small and medium-sized matrices.

The good news is, with the exception of your homework assignments and final exam, you won't have to carry out matrix algebra by hand. It is much more convenient to use a computer for large matrix calculations. The more you develop your matrix algebra skills, the deeper you'll be able to delve into the advanced topics.

Geometric linear algebra

So far, we've described vectors and matrices as arrays of numbers. This is fine for the purpose of doing *algebra* on vectors and matrices, but this description is not sufficient for understanding their geometric properties. The components of a vector $\vec{v} \in \mathbb{R}^n$ can be thought of as distances measured along a coordinate system with n axes. The vector $\vec{v}$ can therefore be said to “point” in a particular direction with respect to the coordinate system. The fun part of linear algebra starts when you learn about the geometric interpretation of the algebraic operations on vectors and matrices.

Consider some unit vector that specifies a direction of interest $\hat{r}$. Suppose we're given some other vector $\vec{v}$, and we're asked to find *how much of $\vec{v}$ is in the $\hat{r}$ direction*. The answer is computed using the dot product: $v_r = \vec{v} \cdot \hat{r} = \|\vec{v}\| \cos \theta$, where θ is the angle between $\vec{v}$ and $\hat{r}$. The technical term for the quantity v_r is “the length of the projection of $\vec{v}$ in the $\hat{r}$ direction.” By “projection,” I mean we ignore all parts of $\vec{v}$ that are not in the $\hat{r}$ direction. Projections are used in mechanics to calculate the x - and y -components of forces in force

diagrams. In Section 4.2 we'll learn how to calculate all kinds of projections using the dot product.

To further consider the geometric aspects of vector operations, imagine the following situation. Suppose I gave you two vectors $\vec{u}$ and $\vec{v}$, and asked you to find a third vector $\vec{w}$ that is perpendicular to both $\vec{u}$ and $\vec{v}$. A priori this sounds like a complicated question to answer, but in fact the required vector $\vec{w}$ can easily be obtained by computing the cross product $\vec{w} = \vec{u} \times \vec{v}$.

In Section 4.1 we'll learn how to describe lines and planes in terms of points, direction vectors, and normal vectors. Consider the following geometric problem: given the equations of two planes in $\mathbb{R}^3$, find the equation of the line where the two planes intersect. There is an algebraic procedure called *Gauss-Jordan elimination* we can use to find the solution.

The determinant of a matrix has a geometric interpretation (Section 3.4). The determinant tells us something about the relative orientation of the vectors that make up the rows of the matrix. If the determinant of a matrix is zero, it means the rows are not *linearly independent*, in other words, at least one of the rows can be written in terms of the other rows. Linear independence, as we'll see shortly, is an important property for vectors to have. The determinant is a convenient way to test whether vectors are linearly independent.

As you learn about geometric linear algebra, practice *visualizing* each new concept you learn about. Always keep a picture in your head of what is going on. The relationships between two-dimensional vectors can be represented in vector diagrams. Three-dimensional vectors can be visualized by pointing pens and pencils in different directions. Most of the intuition you build about vectors in two and three dimensions are applicable to vectors with more dimensions.

Theoretical linear algebra

Linear algebra will teach you how to reason about vectors and matrices in an abstract way. By thinking abstractly, you'll be able to extend your geometric intuition of two and three-dimensional problems to problems in higher dimensions. Much *knowledge buzz* awaits as you learn about new mathematical ideas and develop new ways of thinking.

You're no doubt familiar with the normal coordinate system made of two orthogonal axes: the x -axis and the y -axis. A vector $\vec{v} \in \mathbb{R}^2$ is specified in terms of its coordinates (v_x, v_y) with respect to these axes. When we say $\vec{v} = (v_x, v_y)$, what we really mean is $\vec{v} = v_x\hat{i} + v_y\hat{j}$, where $\hat{i}$ and $\hat{j}$ are unit vectors that point along the x - and y -axes. As it turns out, we can use many other kinds of coor-

dinate systems to represent vectors. A *basis* for $\mathbb{R}^2$ is any set of two vectors $\{\hat{e}_1, \hat{e}_2\}$ that allows us to express all vectors $\vec{v} \in \mathbb{R}^2$ as linear combinations of the basis vectors: $\vec{v} = v_1\hat{e}_1 + v_2\hat{e}_2$. The same vector $\vec{v}$ corresponds to two different coordinate pairs, depending on which basis is used for the description: $\vec{v} = (v_x, v_y)$ in the basis $\{\hat{i}, \hat{j}\}$ and $\vec{v} = (v_1, v_2)$ in the basis $\{\hat{e}_1, \hat{e}_2\}$. We'll learn about the properties of bases in great detail in the coming chapters. The choice of basis plays a fundamental role in all aspects of linear algebra.

In the text above, I explained that computing the product between a matrix and a vector $A\vec{x} = \vec{y}$ can be thought of as a linear transformation, with input $\vec{x}$ and output $\vec{y}$. Any linear transformation (Section 5.1) can be represented (Section 5.2) as a multiplication by a matrix A . Conversely, every $m \times n$ matrix $A \in \mathbb{R}^{m \times n}$ can be thought of as performing a linear transformation $T_A: \mathbb{R}^n \rightarrow \mathbb{R}^m$. The equivalence between matrices and linear transformations allows us to identify certain matrix properties with properties of linear transformations. For example, the *column space* $\mathcal{C}(A)$ of the matrix A (the set of vectors that can be written as a combination of the columns of A) corresponds to the image space of the linear transformation T_A (the set of possible outputs of T_A).

The eigenvalues and eigenvectors of matrices (Section 6.1) allow us to describe the actions of matrices in a natural way. The set of eigenvectors of a matrix are special input vectors for which the action of the matrix is described as a *scaling*. When a matrix acts on one of its eigenvectors, the output is a vector in the same direction as the input vector scaled by a constant. The scaling constant is the *eigenvalue* (own value) associated with this eigenvector. By specifying all the eigenvectors and eigenvalues of a matrix, it is possible to obtain a complete description of what the matrix does. Thinking of matrices in terms of their eigenvalues and eigenvectors is a powerful technique for describing their properties and has many applications.

Linear algebra is useful because linear algebra techniques can be applied to all kinds of “vector-like” objects. The abstract concept of a vector space (Section 6.3) captures precisely what it means for some class of mathematical objects to be “vector-like.” For example, the set of polynomials of degree at most two, denoted $P_2(x)$, consists of all functions of the form $f(x) = a_0 + a_1x + a_2x^2$. Polynomials are vector-like because it's possible to describe each polynomial in terms of its coefficients (a_0, a_1, a_2) . Furthermore, the sum of two polynomials and the multiplication of a polynomial by a constant both correspond to vector-like calculations of coefficients. Once you realize polynomials are vector-like, you'll be able to use notions like *linear independence*, *dimension*, and *basis* when working with polynomials.

Useful linear algebra

One of the most useful skills you'll learn in linear algebra is the ability to solve systems of linear equations. Many real-world problems are expressed as linear equations in multiple unknown quantities. You can solve for n unknowns simultaneously if you have a set of n linear equations that relate the unknowns. To solve this system of equations, eliminate the variables one by one using basic techniques such as substitution and subtraction (see Section 1.15); however, the procedure will be slow and tedious for many unknowns. If the system of equations is linear, it can be expressed as an *augmented matrix* built from the coefficients in the equations. You can then use the Gauss–Jordan elimination algorithm to solve for the n unknowns (Section 3.1). The key benefit of the augmented matrix approach is that it allows you to focus on the coefficients without worrying about the variable names. This saves time when you must solve for many unknowns. Another approach for solving n linear equations in n unknowns is to express the system of equations as a matrix equation (Section 3.2) and then solve the matrix equation by computing the matrix inverse (Section 3.5).

In Section 6.6 you'll learn how to *decompose* a matrix into a product of simpler matrices. Matrix decompositions are often performed for computational reasons: certain problems are easier to solve on a computer when the matrix is expressed in terms of its simpler constituents. Other decompositions, like the decomposition of a matrix into its eigenvalues and eigenvectors, give you valuable information about the properties of the matrix. Google's original PageRank algorithm for ranking webpages by "importance" can be explained as the search for an eigenvector of a matrix. The matrix in question contains information about all hyperlinks that exist between webpages. The eigenvector we're looking for corresponds to a vector that describes the relative importance of each page. So when I tell you eigenvectors are *valuable information*, I'm not kidding: a little 350-billion-dollar company called Google started as an eigenvector idea.

The techniques of linear algebra find applications in many areas of science and technology. We'll discuss applications such as *modelling* multidimensional real-world problems, finding *approximate solutions* to equations (curve fitting), solving constrained optimization problems using *linear programming*, and many other in Chapter 7. As a special bonus for readers interested in physics, a short introduction to quantum mechanics can be found in Chapter 9; if you have a good grasp of linear algebra, you can understand matrix quantum mechanics at no additional mental cost.

Our journey into the land of linear algebra will continue in the next chapter with the study of computational aspects of linear algebra. We'll learn how to solve large systems of linear equations, practice computing matrix products, discuss matrix determinants, and compute matrix inverses.

2.6 Introductory problems

We've been having fun learning about vector and matrix operations, and we've also touched upon linear transformations. I've summarized what linear algebra is about; now it's time for you to put in the effort and check whether you understand the definitions of the operations.

Don't cheat yourself by thinking my summaries are enough; you can't magically understand everything about linear algebra merely by reading about it. Learning doesn't work that way! The only way to truly "get" math—especially advanced math—is to solve problems using the new concepts you've learned. Indeed, the only math I remember from my university days is math that I practiced by solving lots of problems. There's no better way to test whether you understand than testing yourself. Of course, it's your choice whether you'll dedicate the next hour of your life to working through the problems in this section. All I'll say is that you'll have something to show for your efforts; and it's totally worth it.

P2.1 Which of the following functions are linear?

$$\begin{array}{ll} \mathbf{a)} q(x) = x^2 & \mathbf{b)} f(x) = g(h(x)), \text{ where } g(x) = \sqrt{3}x \text{ and } h(x) = -4x \\ \mathbf{c)} i(x) = \frac{1}{mx} & \mathbf{d)} j(x) = \frac{x-a}{x-b} \end{array}$$

P2.2 Find the sum of the vector $(1, 0, 1)$ and the vector $(0, 2, 2)$.

P2.3 Your friend is taking a quantum physics class and needs your help answering the following vectors question. "Let $|a\rangle = 1|0\rangle + 3|1\rangle$ and $|b\rangle = 4|0\rangle - 1|1\rangle$. Find $|a\rangle + |b\rangle$."

Hint: The angle-bracket notation describes vectors: $|0\rangle = \hat{i}$ and $|1\rangle = \hat{j}$.

P2.4 Given unit vectors $\hat{i} = (1, 0, 0)$, $\hat{j} = (0, 1, 0)$, and $\hat{k} = (0, 0, 1)$, find the following cross products: **a)** $\hat{i} \times \hat{i}$, **b)** $\hat{i} \times \hat{j}$, **c)** $(-\hat{i}) \times \hat{k} + \hat{j} \times \hat{i}$, **d)** $\hat{k} \times \hat{j} + \hat{i} \times \hat{i} + \hat{j} \times \hat{k} + \hat{j} \times \hat{i}$.

P2.5 Given $\vec{v} = (2, -1, 3)$ and $\vec{w} = (1, 0, 1)$, compute the following vector products: **a)** $\vec{v} \cdot \vec{w}$, **b)** $\vec{v} \times \vec{w}$, **c)** $\vec{v} \times \vec{v}$, and **d)** $\vec{w} \times \vec{w}$.

P2.6 Given the vectors $\vec{u} = (1, 1, 1)$, $\vec{v} = (2, 3, 1)$, and $\vec{w} = (-1, -1, 2)$, compute the following products:

$$\begin{array}{lll} \mathbf{a)} \vec{u} \cdot \vec{v} & \mathbf{b)} \vec{u} \cdot \vec{w} & \mathbf{c)} \vec{v} \cdot \vec{w} \\ \mathbf{d)} \vec{u} \times \vec{v} & \mathbf{e)} \vec{u} \times \vec{w} & \mathbf{f)} \vec{v} \times \vec{w} \end{array}$$

P2.7 Given the vectors $\vec{p} = (1, 1, 0, 3, 3)$ and $\vec{q} = (1, 2, 3, 4, 5)$, calculate the following expressions:

a) $\vec{p} + \vec{q}$

b) $\vec{p} - \vec{q}$

c) $\vec{p} \cdot \vec{q}$

P2.8 Find a unit vector that is perpendicular to both $\vec{u} = (1, 0, 1)$ and $\vec{v} = (1, 2, 0)$.

Hint: Use the cross product.

P2.9 Find a vector that is orthogonal to both $\vec{u}_1 = (1, 0, 1)$ and $\vec{u}_2 = (1, 3, 0)$, and whose dot product with the vector $\vec{v} = (1, 1, 0)$ is equal to 8.

P2.10 A farmer with a passion for robotics has built a prototype of a robotic tractor. The tractor is programmed to move with a speed of 0.524 km/h and follow the direction of the hour-hand on a conventional watch. Assume the tractor starts at 12:00 p.m. (noon) and is left to roam about in a field until 6:00 p.m. What is the shape of the trajectory that the tractor will follow? What is the total distance travelled by the tractor after six hours?

P2.11 Prove the geometric formula for the dot product $\vec{u} \cdot \vec{v} = \|\vec{u}\| \|\vec{v}\| \cos(\varphi)$, where φ is the angle between vectors $\vec{u}$ and $\vec{v}$.

Hint: Consider the triangle with sides $\vec{u}$, $\vec{v}$, and $\vec{u} - \vec{v}$. Connect the algebraic calculation of the length $\|\vec{u} - \vec{v}\|^2 = (\vec{u} - \vec{v}) \cdot (\vec{u} - \vec{v})$ with the geometric calculation based on the cosine rule from Section 1.10.

P2.12 Compute the product $M\vec{v}$ where $M = \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix}$ and $\vec{v} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix}$.

P2.13 Consider the following three linear transformations that take two-dimensional vectors as inputs and produce two-dimensional vectors as outputs:

$$T_A\left(\begin{bmatrix} x \\ y \end{bmatrix}\right) = \begin{bmatrix} 3x \\ y \end{bmatrix}, \quad T_B\left(\begin{bmatrix} x \\ y \end{bmatrix}\right) = \begin{bmatrix} x+y \\ y \end{bmatrix}, \quad T_C\left(\begin{bmatrix} x \\ y \end{bmatrix}\right) = \begin{bmatrix} 3x+y \\ y \end{bmatrix}.$$

Find the 2×2 matrices A , B , and C such that

$$T_A(\vec{v}) = A\vec{v}, \quad T_B(\vec{v}) = B\vec{v}, \quad T_C(\vec{v}) = C\vec{v}, \quad \text{for all } \vec{v}.$$

Use your answers to compute the matrix products AB and BA .

P2.14 Consider the following matrices of different dimensions:

$$A = \begin{bmatrix} 2 & 3 \\ 5 & 6 \end{bmatrix}, \quad B = \begin{bmatrix} 7 & 8 \\ 1 & 4 \end{bmatrix}, \quad C = \begin{bmatrix} -1 & 4 \\ 2 & 1 \\ -1 & 2 \end{bmatrix}, \quad U = \begin{bmatrix} 2 & 1 \end{bmatrix}, \quad \text{and } V = \begin{bmatrix} 1 \\ -2 \end{bmatrix}.$$

Compute the following matrix expressions:

$$\begin{array}{lllllll} \text{a)} A-B & \text{b)} 5C^T & \text{c)} AB & \text{d)} AC & \text{e)} CA & \text{f)} UA & \text{g)} AV \\ \text{h)} UV & \text{i)} VU & \text{j)} \text{Tr}(A) & \text{k)} \text{Tr}(B) & \text{l)} \det(A) & \text{m)} \det(B) \end{array}$$

P2.15 Use the determinant properties to simplify these expressions:

a) $\det(ABA^{-1})$ **b)** $\frac{\det(AB)}{\det(B)}$ **c)** $\det(AB) - \det(BA)$

P2.16 Suppose A , B , and C are unknown matrices such that $|A| = 2$, $|B| = 3$, and $|C| = 5$. Determine the values of the following expressions:

a) $|AB|$ **b)** $|A^{-1}B|$ **c)** $|BC|$ **d)** $|ABC|$

Hint: Use the properties of the determinant operation.

Chapter 3

Computational linear algebra

This chapter covers the computational aspects of performing matrix calculations. Understanding matrix computations is important because the rest of the chapters in this book depend on them. Suppose we're given a huge matrix $A \in \mathbb{R}^{n \times n}$ with $n = 1000$. Behind the innocent-looking mathematical notation of the matrix inverse A^{-1} , the matrix product AA , and the matrix determinant $\det(A)$, are hidden monster computations involving all the $1000 \times 1000 = 1$ million entries of the matrix A . Millions of arithmetic operations must be performed... so I hope you have at least a thousand pencils ready!

Okay, calm down. I won't *actually* make you calculate millions of arithmetic operations. In fact, to learn linear algebra, it is sufficient to know how to carry out calculations with 3×3 and 4×4 matrices. Yet, even for such moderately sized matrices, computing products, inverses, and determinants by hand are serious computational tasks. If you're ever required to take a linear algebra final exam, you'll need to make sure you can do these calculations quickly. And even if no exam looms in your imminent future, it's still important you practice matrix operations by hand to get a feel for them.

This chapter will introduce you to four important computational tasks involving matrices.

Gauss–Jordan elimination Suppose we're trying to solve two equations in two unknowns x and y :

$$ax + by = c,$$

$$dx + ey = f.$$

If we add α -times the first equation to the second equation, we obtain an equivalent system of equations:

$$\begin{aligned} ax + by &= c \\ (d + \alpha a)x + (e + \alpha b)y &= f + \alpha c. \end{aligned}$$

This is called a *row operation*: we added α -times the first row to the second row. Row operations change the coefficients of the system of equations, but leave the solution unchanged. Gauss–Jordan elimination is a systematic procedure for solving systems of linear equations using row operations.

Matrix multiplication The product AB between matrices $A \in \mathbb{R}^{m \times \ell}$ and $B \in \mathbb{R}^{\ell \times n}$ is the matrix $C \in \mathbb{R}^{m \times n}$ whose entries c_{ij} are defined by the formula $c_{ij} = \sum_{k=1}^{\ell} a_{ik}b_{kj}$ for all $i \in [1, \dots, m]$ and $j \in [1, \dots, n]$. In Section 3.3, we’ll unpack this formula and learn its intuitive interpretation: that computing $C = AB$ is computing all the dot products between the rows of A and the columns of B .

Determinant The determinant of a matrix A , denoted $\det(A)$, is an operation that gives us useful information about the linear independence of the rows of the matrix. The determinant is connected to many notions of linear algebra: linear independence, geometry of vectors, solving systems of equations, and matrix invertibility. We’ll discuss these aspects of determinants in Section 3.4.

Matrix inverse In Section 3.5, we’ll build upon our knowledge of Gauss–Jordan elimination, matrix products, and determinants to derive three different procedures for finding the matrix inverse A^{-1} .

3.1 Reduced row echelon form

In this section we’ll learn to solve systems of linear equations using the *Gauss–Jordan elimination* procedure. A system of equations can be represented as a matrix of coefficients. The Gauss–Jordan elimination procedure converts any matrix into its *reduced row echelon form* (RREF). We can use the RREF to easily find the solution (or solutions) of the system of equations.

Heads up: the material covered in this section requires your full-on, caffeinated attention, as the procedures you’ll learn are somewhat tedious. Gauss–Jordan elimination involves many repetitive manipulations of arrays of numbers. It’s important you follow the step-by-step manipulations, as well as verify each step I present *on your own* with pen and paper. Don’t just take my word for it—always verify the steps!

Solving equations

Suppose you're asked to solve the following system of equations:

$$\begin{aligned}1x_1 + 2x_2 &= 5 \\ 3x_1 + 9x_2 &= 21.\end{aligned}$$

The standard approach is to use one of the equation-solving tricks we learned in Section 1.15 to combine the equations and find the values of the two unknowns x_1 and x_2 .

Observe that the *names* of the two unknowns are irrelevant to the solution of the system of equations. Indeed, the solution (x_1, x_2) to the above system of equations is the same as the solution (s, t) to the system of equations

$$\begin{aligned}1s + 2t &= 5 \\ 3s + 9t &= 21.\end{aligned}$$

The important parts of a system of linear equations are the *coefficients* in front of the variables, and the constants on the right side of each equation.

Augmented matrix

The system of linear equations can be written as an *augmented matrix*:

$$\left[\begin{array}{cc|c} 1 & 2 & 5 \\ 3 & 9 & 21 \end{array} \right].$$

The first column corresponds to the coefficients of the first variable; the second column is for the second variable; and the last column corresponds to the constants on the right side. It is customary to draw a vertical line where the equal signs in the equations would normally appear. This line helps distinguish the coefficients of the equations from the column of constants on the right side.

Once we have the augmented matrix, we can simplify it by using *row operations* (which we'll discuss shortly) on its entries. After simplification by row operations, the augmented matrix will be transformed to

$$\left[\begin{array}{cc|c} 1 & 0 & 1 \\ 0 & 1 & 2 \end{array} \right],$$

which corresponds to the system of equations

$$\begin{aligned}x_1 &= 1 \\ x_2 &= 2.\end{aligned}$$

This is a *trivial* system of equations; there is nothing left to solve and we can see the solutions are $x_1 = 1$ and $x_2 = 2$. This example illustrates the general idea of the Gauss–Jordan elimination procedure for solving systems of equations by manipulating an augmented matrix.

Row operations

We can manipulate the rows of an augmented matrix without changing its solutions. We're allowed to perform the following three types of row operations:

- Add a multiple of one row to another row
- Swap the position of two rows
- Multiply a row by a constant

Let's trace the sequence of row operations needed to solve the system of equations

$$\begin{aligned} 1x_1 + 2x_2 &= 5 \\ 3x_1 + 9x_2 &= 21, \end{aligned}$$

starting from its augmented matrix:

$$\left[\begin{array}{cc|c} 1 & 2 & 5 \\ 3 & 9 & 21 \end{array} \right].$$

1. As a first step, we eliminate the first variable in the second row by subtracting three times the first row from the second row:

$$\left[\begin{array}{cc|c} 1 & 2 & 5 \\ 0 & 3 & 6 \end{array} \right].$$

We denote this row operation as $R_2 \leftarrow R_2 - 3R_1$.

2. To simplify the second row, we divide it by 3 to obtain

$$\left[\begin{array}{cc|c} 1 & 2 & 5 \\ 0 & 1 & 2 \end{array} \right].$$

This row operation is denoted $R_2 \leftarrow \frac{1}{3}R_2$.

3. Finally we eliminate the second variable from the first row by subtracting two times the second row, $R_1 \leftarrow R_1 - 2R_2$:

$$\left[\begin{array}{cc|c} 1 & 0 & 1 \\ 0 & 1 & 2 \end{array} \right].$$

We can now read off the solution: $x_1 = 1$ and $x_2 = 2$.

Note how we simplified the augmented matrix through a specific procedure: we followed the *Gauss–Jordan elimination* algorithm to bring the matrix into its *reduced row echelon form*.

The *reduced row echelon form* (RREF) is the simplest form for an augmented matrix. Each row contains a *leading one* (a numeral 1) also known as a *pivot*. Each column's pivot is used to eliminate the numbers that lie below and above it in the same column. The end result of this procedure is the reduced row echelon form:

$$\left[\begin{array}{cccc|c} 1 & 0 & * & 0 & * \\ 0 & 1 & * & 0 & * \\ 0 & 0 & 0 & 1 & * \end{array} \right].$$

Note the matrix contains only zero entries below and above the pivots. The asterisks * denote arbitrary numbers that cannot be eliminated because no leading one is present in these columns.

Definitions

- The *solution* to a system of linear equations in the variables $x_1, x_2, \dots, x_n$ is the set of values $\{(x_1, x_2, \dots, x_n)\}$ that satisfy *all* the equations.
- The *pivot* for row j of a matrix is the left-most nonzero entry in the row j . Any *pivot* can be converted into a *leading one* by an appropriate scaling of that row.
- *Gaussian elimination* is the process of bringing a matrix into *row echelon form*.
- A matrix is said to be in *row echelon form* (REF) if all entries below the leading ones are zero. This form can be obtained by adding or subtracting the row with the leading one from the rows below it.
- *Gaussian-Jordan elimination* is the process of bringing a matrix into *reduced row echelon form*.
- A matrix is said to be in *reduced row echelon form* (RREF) if all the entries below *and above* the pivots are zero. Starting from the REF, we obtain the RREF by subtracting the row containing the pivots from the rows above that row.
- $\text{rank}(A)$: the *rank* of the matrix A is the number of pivots in the RREF of A .

Gauss–Jordan elimination algorithm

The Gauss–Jordan elimination algorithm proceeds in two phases: a forward phase in which we move left to right, and a backward phase in which we move right to left.

1. Forward phase (left to right):
 - 1.1 Obtain a pivot (a leading one) in the leftmost column.
 - 1.2 Subtract the row with the pivot from all rows below it to obtain zeros in the entire column.
 - 1.3 Look for a leading one in the next column and repeat.
2. Backward phase (right to left):
 - 2.1 Find the rightmost pivot and use it to eliminate all numbers above the pivot in its column.
 - 2.2 Move one column to the left and repeat.

Example We’re asked to solve the following system of equations:

$$\begin{aligned} 1x + 2y + 3z &= 14 \\ 2x + 5y + 6z &= 30 \\ -1x + 2y + 3z &= 12. \end{aligned}$$

The first step toward the solution is to build the augmented matrix that corresponds to this system of equations:

$$\left[\begin{array}{ccc|c} 1 & 2 & 3 & 14 \\ 2 & 5 & 6 & 30 \\ -1 & 2 & 3 & 12 \end{array} \right].$$

We can now start the left-to-right phase of the algorithm:

1. Conveniently, there is a leading one at the top of the leftmost column. If a zero were there instead, a row-swap operation would be necessary to obtain a nonzero entry.
2. The next step is to clear the entries in the entire column below this pivot. The row operations we’ll use for this purpose are $R_2 \leftarrow R_2 - 2R_1$ and $R_3 \leftarrow R_3 + R_1$:

$$\left[\begin{array}{ccc|c} 1 & 2 & 3 & 14 \\ 0 & 1 & 0 & 2 \\ 0 & 4 & 6 & 26 \end{array} \right].$$

3. We now shift our attention to the second column. Using the leading one for the second column, we set the number in the column below it to zero using $R_3 \leftarrow R_3 - 4R_2$. The result is

$$\left[\begin{array}{ccc|c} 1 & 2 & 3 & 14 \\ 0 & 1 & 0 & 2 \\ 0 & 0 & 6 & 18 \end{array} \right].$$

4. Next, we move to the third column. Instead of a leading one, we find it contains a “leading six,” which we can convert to a leading one using $R_3 \leftarrow \frac{1}{6}R_3$. We thus obtain

$$\left[\begin{array}{ccc|c} 1 & 2 & 3 & 14 \\ 0 & 1 & 0 & 2 \\ 0 & 0 & 1 & 3 \end{array} \right].$$

The forward phase of the Gauss–Jordan elimination procedure is now complete. We identified three pivots and used them to systematically set all entries below each pivot to zero. The matrix is now in *row echelon form*.

Next, we perform the backward phase of the Gauss–Jordan elimination procedure, where we’ll work right-to-left to set all numbers above each pivot to zero:

5. The first row operation is $R_1 \leftarrow R_1 - 3R_3$, and it leads to

$$\left[\begin{array}{ccc|c} 1 & 2 & 0 & 5 \\ 0 & 1 & 0 & 2 \\ 0 & 0 & 1 & 3 \end{array} \right].$$

6. The final step is $R_1 \leftarrow R_1 - 2R_2$, which gives

$$\left[\begin{array}{ccc|c} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 2 \\ 0 & 0 & 1 & 3 \end{array} \right].$$

The matrix is now in *reduced row echelon form*, and we can see the solution is $x = 1$, $y = 2$, and $z = 3$.

We’ve described the general idea of the Gauss–Jordan elimination and explored some examples where the solutions to the system of equations were *unique*. There are other possibilities for the solutions of a system of linear equations. We’ll describe these other possible scenarios next.

Number of solutions

A system of three linear equations in three variables could have:

- **One solution.** If the RREF of a matrix has a pivot in each row, we can read off the values of the solution by inspection:

$$\left[\begin{array}{ccc|c} 1 & 0 & 0 & c_1 \\ 0 & 1 & 0 & c_2 \\ 0 & 0 & 1 & c_3 \end{array} \right].$$

The *unique* solution is $x_1 = c_1$, $x_2 = c_2$, and $x_3 = c_3$.

- **Infinitely many solutions.** If one of the equations is redundant, a row of zeros will appear when the matrix is brought to the RREF. This happens when one of the original equations is a linear combination of the other two. In such cases, we're really solving *two* equations in *three* variables, so we can't "pin down" one of the unknown variables. We say the solution contains a *free variable*. For example, consider the following RREF:

$$\left[\begin{array}{ccc|c} 1 & 0 & a_1 & c_1 \\ 0 & 1 & a_2 & c_2 \\ 0 & 0 & 0 & 0 \end{array} \right].$$

The column that doesn't contain a leading one corresponds to the free variable. To indicate that x_3 is a free variable, we give it a special label, $x_3 = t$. The variable t could be any number $t \in \mathbb{R}$. In other words, when we say t is free, it means t can take on *any* value from $-\infty$ to $+\infty$. The information in the augmented matrix can now be used to express x_1 and x_2 in terms of the constants and the free variable t :

$$\left\{ \begin{array}{l} x_1 = c_1 - a_1 t \\ x_2 = c_2 - a_2 t \\ x_3 = t \end{array} \right\}, \quad \forall t \in \mathbb{R} = \left\{ \begin{bmatrix} c_1 \\ c_2 \\ 0 \end{bmatrix} + t \begin{bmatrix} -a_1 \\ -a_2 \\ 1 \end{bmatrix}, \quad \forall t \in \mathbb{R} \right\}.$$

The solution corresponds to the equation of a line passing through the point $(c_1, c_2, 0)$ with direction vector $(-a_1, -a_2, 1)$. We'll discuss the geometry of lines in Chapter 4. For now, it's important you understand that a system of equations can have more than one solution; any point on the line $\ell = \{(c_1, c_2, 0) + t(-a_1, -a_2, 1), \forall t \in \mathbb{R}\}$ is a solution to the above system of equations.

- **Infinitely many solutions in two dimensions.** It's also possible to obtain a two-dimensional solution space. This happens

when two of the three equations are redundant. This results in a single leading one, and thus two free variables. For example, in the RREF

$$\left[\begin{array}{ccc|c} 0 & 1 & a_2 & c_2 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{array} \right],$$

the variables x_1 and x_3 are free. As in the previous infinitely-many-solutions case, we define new labels for the free variables $x_1 = s$ and $x_3 = t$, where $s \in \mathbb{R}$ and $t \in \mathbb{R}$ are two arbitrary numbers. The solution to this system of equations is

$$\left\{ \begin{array}{l} x_1 = s \\ x_2 = c_2 - a_2 t, \quad \forall s, t \in \mathbb{R} \\ x_3 = t \end{array} \right\} = \left\{ \begin{bmatrix} 0 \\ c_2 \\ 0 \end{bmatrix} + s \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + t \begin{bmatrix} 0 \\ -a_2 \\ 1 \end{bmatrix}, \quad \forall s, t \in \mathbb{R} \right\}.$$

This solution set corresponds to the *parametric equation* of a plane that contains the point $(0, c_2, 0)$ and the vectors $(1, 0, 0)$ and $(0, -a_2, 1)$.

The *general equation* for the solution plane is $0x + 1y + a_2z = c_2$, as we can observe from the first row of the augmented matrix. We'll learn more about the geometry of planes and how to convert between their general and parametric forms in Section 4.1.

- **No solutions.** If there are no numbers (x_1, x_2, x_3) that simultaneously satisfy all three equations, the system of equations has no solution. An example of a system of equations with no solution is the pair $s + t = 4$ and $s + t = 44$. There are no numbers (s, t) that satisfy both of these equations.

A system of equations has no solution if its reduced row echelon form contains a row of zero coefficients with a nonzero constant in the right side:

$$\left[\begin{array}{ccc|c} 1 & 0 & 0 & c_1 \\ 0 & 1 & 0 & c_2 \\ 0 & 0 & 0 & c_3 \end{array} \right].$$

If $c_3 \neq 0$, this system of equations is impossible to satisfy. There is no solution because there are no numbers (x_1, x_2, x_3) such that $0x_1 + 0x_2 + 0x_3 = c_3$.

Dear reader, we've reached a moment in this book where you'll need to update your math vocabulary. The solution to an individual equation is a finite set of points. The solution to a *system* of equations can be an entire space containing infinitely many points, such as a line or a plane. Please update your definition of the term *solution*

to include the new, more specific term *solution set*—the set of points that satisfy the system of equations. The *solution set* of a system of three linear equations in three unknowns could be either the empty set $\{\emptyset\}$ (no solution), a set with one element $\{(x_1, x_2, x_3)\}$, or a set with infinitely many elements like a line $\{p_0 + t\vec{v}, t \in \mathbb{R}\}$ or a plane $\{p_0 + s\vec{v} + t\vec{w}, s, t \in \mathbb{R}\}$. Another possible solution set is all of $\mathbb{R}^3$, where every vector $\vec{x} \in \mathbb{R}^3$ is a solution to the equation:

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

Note the distinction between the three types of infinite solution sets. A line is one-dimensional, a plane is two-dimensional, and $\mathbb{R}^3$ is three-dimensional. Describing all points on a line requires one parameter, describing all points on a plane takes two parameters, and describing a point in $\mathbb{R}^3$ takes three parameters.

Geometric interpretation

We can gain some intuition about solution sets by studying the geometry of the intersections of lines in $\mathbb{R}^2$ and planes in $\mathbb{R}^3$.

Lines in two dimensions

Equations of the form $ax + by = c$ correspond to lines in $\mathbb{R}^2$. Solving systems of equations of the form

$$\begin{aligned} a_1x + b_1y &= c_1 \\ a_2x + b_2y &= c_2 \end{aligned}$$

requires finding the point $(x, y) \in \mathbb{R}^2$ where these lines intersect. There are three possibilities for the solution set:

- **One solution** if the two lines intersect at a point
- **Infinitely many solutions** if the lines are superimposed
- **No solution** if the two lines are parallel and never intersect

See figures 1.67, 1.68, and 1.69 (pages 119–120) for an illustration.


Planes in three dimensions

Equations of the form $ax + by + cz = d$ correspond to planes in $\mathbb{R}^3$. When solving three such equations,


$$\begin{aligned} a_1x + b_1y + c_1z &= d_1, \\ a_2x + b_2y + c_2z &= d_2, \\ a_3x + b_3y + c_3z &= d_3, \end{aligned}$$

we want to find the points (x, y, z) that satisfy all three equations simultaneously. There are four possibilities for the solution set:

- **One solution.** Three non-parallel planes intersect at a point.
- **Infinitely many solutions.** If one of the plane equations is redundant, the solution corresponds to the intersection of two planes. Two non-parallel planes intersect on a line.
- **Infinitely many solutions in two dimensions.** If two of the equations are redundant, the solution space is a two-dimensional plane.
- **No solution.** If no common points exist at the intersection of all three planes, then the system of equations has no solution.




(a) One solution




(b) Infinite solutions

Figure 3.1: Three planes can intersect at a unique point, as in figure (a); or along a line, as in figure (b). In the first case, there is a unique point (x_0, y_0, z_0) common to all three planes. In the second case, all points on the line $\{p_0 + t\vec{v}, \forall t \in \mathbb{R}\}$ are shared by the planes.



(a) No solution



(b) No solution

Figure 3.2: These illustrations depict systems of three equations in three unknowns that have no solution. No common points of intersection exist.

Computer power

The computer algebra system at <http://live.sympy.org> can be used to compute the reduced row echelon form of any matrix.

Here is an example of how to create a sympy Matrix object:

```
>>> from sympy.matrices import Matrix
>>> A = Matrix([[1, 2, 5],      # use SHIFT+ENTER for newline
                [3, 9, 21]])
```

In Python, we define lists using the square brackets [and]. A matrix is defined as a list of lists.

To compute the reduced row echelon form of A, call its `rref()` method:

```
>>> A.rref()
( [1, 0, 1] # RREF of A          # locations of pivots
  [0, 1, 2],          [0, 1] )
```

The `rref()` method returns a tuple containing the RREF of A and an array that tells us the 0-based indices of the columns that contain leading ones. Usually, we'll want to find the RREF of A and ignore the pivots; to obtain the RREF without the pivots, select the first (index zero) element in the result of `A.rref()`:

```
>>> Arref = A.rref()[0]
>>> Arref
[1, 0, 1]
[0, 1, 2]
```

The `rref()` method is the fastest way to obtain the reduced row echelon form of a SymPy matrix. The computer will apply the Gauss–Jordan elimination procedure and show you the answer. If you want to see the intermediary steps of the elimination procedure, you can also manually apply row operations to the matrix.

Example Let's compute the reduced row echelon form of the same augmented matrix by using row operations in SymPy:

```
>>> A = Matrix([[1, 2, 5],
                [3, 9, 21]])

>>> A[1,:] = A[1,:] - 3*A[0,:]
>>> A
[1, 2, 5]
[0, 3, 6]
```

We use the notation $A[i, :]$ to refer to entire rows of the matrix. The number i specifies the 0-based row index: the first row of A is $A[0, :]$ and the second row is $A[1, :]$. The code example above implements the row operation $R_2 \leftarrow R_2 - 3R_1$.

To obtain the reduced row echelon form of the matrix A , we carry out two more row operations, $R_2 \leftarrow \frac{1}{3}R_2$ and $R_1 \leftarrow R_1 - 2R_2$, using the following commands:

```
>>> A[1,:] = S(1)/3*A[1,:]
>>> A[0,:] = A[0,:] - 2*A[1,:]
>>> A
[1, 0, 1]          # the same result as A.rref()[0]
[0, 1, 2]
```

Note we represent the fraction $\frac{1}{3}$ as $S(1)/3$ in order to obtain the exact rational expression `Rational(1,3)`. If we were to input $\frac{1}{3}$ as $1/3$, SymPy would interpret this either as integer or floating point division, which is not what we want. The single-letter helper function `S` is an alias for the function `sympify`, which ensures a SymPy object is produced. Another way to input the exact fraction $\frac{1}{3}$ is `S('1/3')`.

If you need to swap two rows of a matrix, you can use the standard Python tuple assignment syntax. To swap the position of the first and second rows, use

```
>>> A[0,:], A[1,:] = A[1,:], A[0,:]
>>> A
[0, 1, 2]
[1, 0, 1]
```

Using row operations to compute the reduced row echelon form of a matrix allows you to see the intermediary steps of a calculation; which is useful, for instance, when checking the correctness of your homework problems.

There are other applications of matrix methods that use row operations (see Section 7.6), so it's good idea to know how to use SymPy for this purpose.

Discussion

In this section, we learned the Gauss–Jordan elimination procedure for simplifying matrices, which just so happens to be one of the most important computational tools of linear algebra. Beyond being a procedure for finding solutions to systems of linear equations, the Gauss–Jordan elimination algorithms can be used to solve a broad range of other linear algebra problems. Later in the book, we'll use the Gauss–Jordan elimination algorithm to compute inverse matrices (Section 3.5) and to “distill” bases for vector spaces (Section 4.5).

Exercises

You've read about linear algebra—now you must apply what you've learned. Let's see if you can solve the exercises in this section. It will probably take you a few hours to complete all of them, but trust me: if you do the work, you'll reap the rewards.

Imagine you want to join a martial arts school and the master tells you to first practice horse stance for hours before you can learn any moves. Well, this is the math equivalent. You need to know the Gauss–Jordan elimination procedure inside out, since it's the base on which we'll build many other linear algebra skills.

E3.1 Consider the system of equations and its augmented matrix representation:

$$\begin{array}{rcrcrc} 3x & + & 3y & = & 6 & \\ 2x & + & \frac{3}{2}y & = & 5 & \end{array} \quad \Rightarrow \quad \left[\begin{array}{cc|c} 3 & 3 & 6 \\ 2 & \frac{3}{2} & 5 \end{array} \right].$$

Find the solution to this system of equations by bringing the augmented matrix into reduced row echelon form.

E3.2 Repeat E3.1 using the calculator at <http://live.sympy.org>. First define the augmented matrix using

```
>>> A = Matrix([
    [3,      3,  6],
    [2,  S(3)/2, 5]]) # note use of S(3)/2 to obtain 3/2
```

Then perform row operations using SymPy to bring the matrix into RREF. Confirm your answer using the direct method `A.rref()`.

E3.3 Find the solutions to the systems of equations that correspond to the following augmented matrices.

$$\text{a) } \left[\begin{array}{cc|c} 3 & 3 & 6 \\ 1 & 1 & 5 \end{array} \right] \quad \text{b) } \left[\begin{array}{cc|c} 3 & 3 & 6 \\ 2 & \frac{3}{2} & 3 \end{array} \right] \quad \text{c) } \left[\begin{array}{cc|c} 3 & 3 & 6 \\ 1 & 1 & 2 \end{array} \right]$$

Hint: The third system of equations has many solutions.

In this section we learned a practical computational algorithm for solving systems of equations by using row operations on an augmented matrix. In the next section, we'll increase the level of abstraction. By “zooming out” one level, we can view the entire system of equations as a matrix equation $A\vec{x} = \vec{b}$ and solve the problem in one step: $\vec{x} = A^{-1}\vec{b}$.

3.2 Matrix equations

We can express the problem of solving a system of linear equations as a matrix equation and obtain the solution using the matrix inverse. Consider the following system of linear equations:

$$\begin{aligned}x_1 + 2x_2 &= 5 \\ 3x_1 + 9x_2 &= 21.\end{aligned}$$

We can rewrite this system of equations as a matrix-vector product:

$$\begin{bmatrix} 1 & 2 \\ 3 & 9 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 5 \\ 21 \end{bmatrix},$$

or, more compactly, as $A\vec{x} = \vec{b}$, where A is a 2×2 matrix, $\vec{x}$ is the vector of unknowns (a 2×1 matrix), and $\vec{b}$ is a vector of constants (a 2×1 matrix).

We can solve for $\vec{x}$ in this matrix equation by multiplying both sides of the equation by the inverse A^{-1} :

$$A^{-1}A\vec{x} = \mathbb{1}\vec{x} = \vec{x} = A^{-1}\vec{b}.$$

Thus, to solve a system of linear equations, we can find the inverse of the matrix of coefficients, then compute the product:

$$\vec{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = A^{-1}\vec{b} = \begin{bmatrix} 3 & -\frac{2}{3} \\ -1 & \frac{1}{3} \end{bmatrix} \begin{bmatrix} 5 \\ 21 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}.$$

The amount of effort it takes to find A^{-1} is roughly equivalent to the effort needed to bring an augmented matrix $[A \mid \vec{b}]$ to reduced row echelon form—it's not like we're given the solution for free by simply rewriting the system of equations in matrix form. Nevertheless, expressing the system of equations as $A\vec{x} = \vec{b}$ and its solution as $\vec{x} = A^{-1}\vec{b}$ is useful in that it saves us from needing to juggle dozens of individual coefficients. The same symbolic expression $\vec{x} = A^{-1}\vec{b}$ applies whether A is a 2×2 matrix or a 1000×1000 matrix.

Introduction

It's time we had an important discussion about matrix equations and how they differ from regular equations with numbers. If a , b , and c are three numbers, and I tell you to *solve* for a in the equation

$$ab = c,$$

you'd know the answer is $a = c/b = c\frac{1}{b} = \frac{1}{b}c$, and that would be the end of it.

Now suppose A , B , and C are matrices and you want to solve for A in the matrix equation

$$AB = C.$$

The answer $A = C/B$ is not allowed. So far, we defined matrix *multiplication* and matrix *inversion*, but not matrix *division*. Instead of dividing by B , we must multiply by B^{-1} , which has the same effect as a “divide by B ” operation. The product of B and B^{-1} gives the identity matrix,

$$BB^{-1} = \mathbb{1}, \quad B^{-1}B = \mathbb{1}.$$

When applying the inverse matrix B^{-1} to the equation, we must specify whether we are multiplying from the left or from the right, because matrix multiplication is not commutative. Can you determine the correct answer for A in the above equations? Is it $A = CB^{-1}$ or $A = B^{-1}C$?

To solve matrix equations, we employ the same technique we used to solve equations in Chapter 1: undoing the operations that stand in the way of the unknown. Recall that we must always **do the same thing to both sides of an equation** for it to remain true.

With matrix equations, it's the same story all over again, but there are two new things you need to keep in mind:

- The order in which matrices are multiplied matters. Matrix multiplication is not a commutative operation, $AB \neq BA$. The expressions ABC and BAC are different despite the fact that they are the product of the same three matrices.
- When performing operations on matrix equations, you can act either *from the left* or *from the right* on the equation.

The best way to familiarize yourself with the peculiarities of matrix equations is to look at example calculations. Don't worry, there won't be anything too mathematically demanding in this section; we'll just look at some pictures.

Matrix times vector

Suppose we want to solve the equation $A\vec{x} = \vec{b}$, in which an $n \times n$ matrix A multiplies the vector $\vec{x}$ to produce a vector $\vec{b}$. Recall, we can think of vectors as “tall and skinny” $n \times 1$ matrices.

The picture corresponding to the equation $A\vec{x} = \vec{b}$ is

$$\boxed{A} \boxed{\vec{x}} = \boxed{\vec{b}} .$$

Assuming A is invertible, we can multiply by the inverse A^{-1} on the left of both sides of the equation:

$$\boxed{A^{-1}} \boxed{A} \boxed{\vec{x}} = \boxed{A^{-1}} \boxed{\vec{b}} .$$

By definition, A^{-1} times its inverse A is equal to the identity matrix $\mathbb{1}$, which is a diagonal matrix with ones on the diagonal and zeros everywhere else:

$$\boxed{I} \boxed{\vec{x}} = \boxed{A^{-1}} \boxed{\vec{b}} .$$

Any vector times the identity matrix remains unchanged, so

$$\boxed{\vec{x}} = \boxed{A^{-1}} \boxed{\vec{b}} ,$$

which is the final answer.

Note that the question “Solve for $\vec{x}$ in $A\vec{x} = \vec{b}$ ” sometimes arises in situations where the matrix A is not invertible. If the system of equations is under-specified (A is wider than it is tall), there will be a whole subspace of acceptable solutions $\vec{x}$. Recall the cases with infinite solutions (lines and planes) we saw in the previous section.

Matrix times matrix

Let’s look at some other matrix equations. Suppose we want to solve for A in the equation $AB = C$:

$$\boxed{A} \boxed{B} = \boxed{C} .$$

To isolate A , we multiply by B^{-1} *from the right* on both sides:

$$\boxed{A} \boxed{B} \boxed{B^{-1}} = \boxed{C} \boxed{B^{-1}} .$$

When B^{-1} hits B , the matrices cancel ($BB^{-1} = \mathbb{I}$) and we obtain

$$\boxed{A} = \boxed{C} \boxed{B^{-1}} .$$

Matrix times matrix variation

What if we want to solve for B in the same equation $AB = C$?

$$\boxed{A} \boxed{B} = \boxed{C} .$$

Again, we must *do the same* to both sides of the equation. To cancel A , we need to multiply by A^{-1} from the left:

$$\boxed{A^{-1}} \boxed{A} \boxed{B} = \boxed{A^{-1}} \boxed{C} .$$

When A^{-1} cancels with A , we obtain the final result

$$\boxed{B} = \boxed{A^{-1}} \boxed{C} .$$

This completes our lightning tour of matrix equations. There is really nothing new to learn here; just make sure you're aware that the *order* in which matrices are multiplied matters, and remember the general principle of "doing the same thing to both sides of the equation." Acting according to this principle is essential in all of math, and particularly important when manipulating noncommutative operations as in matrix equations.

In the next section, we'll "zoom in" on matrix equations by examining the arithmetic operations performed on matrix entries during multiplication.

Exercises

E3.4 Solve for X in the following matrix equations: **a)** $XA = B$, **b)** $ABCXD = E$, and **c)** $AC = XDC$. You can assume the matrices A , B , C , D , and E are all invertible.

3.3 Matrix multiplication

Suppose we're given the matrices

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} e & f \\ g & h \end{bmatrix},$$

and we want to compute the *matrix product* AB .

Unlike matrix addition and subtraction, matrix multiplication is *not* performed entry-wise:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} e & f \\ g & h \end{bmatrix} \neq \begin{bmatrix} ae & bf \\ cg & dh \end{bmatrix}.$$

Instead, the matrix product is computed by taking the dot product between each row of the matrix A and each column of the matrix B :

$$\begin{aligned} \vec{r}_1 &\rightarrow \begin{bmatrix} a & b \end{bmatrix} \begin{bmatrix} e & f \\ g & h \end{bmatrix} &= \begin{bmatrix} \vec{r}_1 \cdot \vec{c}_1 & \vec{r}_1 \cdot \vec{c}_2 \end{bmatrix} \\ \vec{r}_2 &\rightarrow \begin{bmatrix} c & d \end{bmatrix} \begin{bmatrix} e & f \\ g & h \end{bmatrix} &= \begin{bmatrix} \vec{r}_2 \cdot \vec{c}_1 & \vec{r}_2 \cdot \vec{c}_2 \end{bmatrix} \\ &\quad \quad \quad \uparrow \quad \uparrow &= \begin{bmatrix} ae + bg & af + bh \\ ce + dg & cf + dh \end{bmatrix}. \end{aligned}$$

Recall the dot product between two vectors $\vec{v}$ and $\vec{w}$ is computed using the formula $\vec{v} \cdot \vec{w} = \sum_i v_i w_i$.

Now, let's look at a picture that shows how to compute the product of a matrix with four rows and another matrix with five columns. To compute the top left entry of the result, take the dot product of the first row of the first matrix and the first column of the second matrix.

$$\begin{pmatrix} \boxed{r_1} \\ \boxed{r_2} \\ \boxed{r_3} \\ \boxed{r_4} \end{pmatrix} \begin{pmatrix} \boxed{c_1} & \boxed{c_2} & \boxed{c_3} & \boxed{c_4} & \boxed{c_5} \end{pmatrix} = \begin{pmatrix} \boxed{1-1} & \boxed{1-2} & \boxed{1-3} & \boxed{1-4} & \boxed{1-5} \\ \boxed{2-1} & \boxed{2-2} & \boxed{2-3} & \boxed{2-4} & \boxed{2-5} \\ \boxed{3-1} & \boxed{3-2} & \boxed{3-3} & \boxed{3-4} & \boxed{3-5} \\ \boxed{4-1} & \boxed{4-2} & \boxed{4-3} & \boxed{4-4} & \boxed{4-5} \end{pmatrix}$$

Figure 3.3: Matrix multiplication is performed rows-times-columns. The first-row, first-column entry of the product is the dot product of r_1 and c_1 .

Similarly, the entry on the third row and fourth column of the product is computed by taking the dot product of the third row of the first matrix and the fourth column of the second matrix. See Figure 3.4.

For the matrix product to work, the rows of the first matrix must have the same length as the columns of the second matrix.

Matrix multiplication rules

- Matrix multiplication is associative:

$$(AB)C = A(BC) = ABC.$$

$$\begin{pmatrix} r_1 \\ r_2 \\ r_3 \\ r_4 \end{pmatrix} \begin{pmatrix} c_1 & c_2 & c_3 & c_4 & c_5 \end{pmatrix} = \begin{pmatrix} 1 \cdot 1 & 1 \cdot 2 & 1 \cdot 3 & 1 \cdot 4 & 1 \cdot 5 \\ 2 \cdot 1 & 2 \cdot 2 & 2 \cdot 3 & 2 \cdot 4 & 2 \cdot 5 \\ 3 \cdot 1 & 3 \cdot 2 & 3 \cdot 3 & 3 \cdot 4 & 3 \cdot 5 \\ 4 \cdot 1 & 4 \cdot 2 & 4 \cdot 3 & 4 \cdot 4 & 4 \cdot 5 \end{pmatrix}$$

Figure 3.4: The third-row, fourth-column entry of the product is computed by taking the dot product of r_3 and c_4 .

- The touching dimensions of the matrices must be the same. For the triple product ABC to exist, the rows of A must have the same dimension as the columns of B , and the rows of B must have the same dimension as the columns of C .
- Given two matrices $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times k}$, the product AB is an $m \times k$ matrix.
- Matrix multiplication is not a commutative operation.

$$\boxed{A} \boxed{B} \neq \boxed{B} \boxed{A}$$

Figure 3.5: The order of multiplication matters for matrices: the product AB does not equal the product BA .

Example Consider the matrices $A \in \mathbb{R}^{2 \times 3}$ and $B \in \mathbb{R}^{3 \times 2}$. The product $AB = C \in \mathbb{R}^{2 \times 2}$ is computed as

$$\underbrace{\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}}_A \underbrace{\begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix}}_B = \begin{bmatrix} 1 + 6 + 15 & 2 + 8 + 18 \\ 4 + 15 + 30 & 8 + 20 + 36 \end{bmatrix} = \underbrace{\begin{bmatrix} 22 & 28 \\ 49 & 64 \end{bmatrix}}_C.$$

We can also compute the product $BA = D \in \mathbb{R}^{3 \times 3}$:

$$\underbrace{\begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix}}_B \underbrace{\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}}_A = \begin{bmatrix} 1 + 8 & 2 + 10 & 3 + 12 \\ 3 + 16 & 6 + 20 & 9 + 24 \\ 5 + 24 & 10 + 30 & 15 + 36 \end{bmatrix} = \underbrace{\begin{bmatrix} 9 & 12 & 15 \\ 19 & 26 & 33 \\ 29 & 40 & 51 \end{bmatrix}}_D.$$

In each case, the touching dimensions of the two matrices in the product are the same. Note that $C = AB \neq BA = D$, and in fact, the products AB and BA are matrices with different dimensions.

Applications

Why is matrix multiplication defined the way it is defined?

Composition of linear transformations

The long answer to this question will be covered in depth when we reach the chapter on linear transformations (Section 5, page 257). Since I don't want you to live in suspense, I'll give you the short answer now. We can think of multiplying a column vector $\vec{x} \in \mathbb{R}^n$ by a matrix $A \in \mathbb{R}^{m \times n}$ as analogous to applying a *linear transformation* T_A of the form:

$$T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m.$$

Applying the linear transformation T_A to the input $\vec{x}$ is the same as computing the matrix-vector product $A\vec{x}$:

$$\text{for all } \vec{x} \in \mathbb{R}^n, \quad T_A(\vec{x}) = A\vec{x}.$$

Every linear transformation from $\mathbb{R}^n$ to $\mathbb{R}^m$ can be described as a multiplication by some matrix $A \in \mathbb{R}^{m \times n}$.

What happens when we apply two linear transformations in succession? When we do this to ordinary functions, we call it *function composition*, and denote it with a little circle:

$$g \circ f(x) = g(f(x)),$$

where $g \circ f(x)$ indicates we should first apply f to x , then apply g to the output of $f(x)$. The composition of two linear transformations $T_B \circ T_A$ (applying T_A then T_B) corresponds to multiplication by matrices A and B :

$$T_B \circ T_A(\vec{x}) = T_B(T_A(\vec{x})) = B(A\vec{x}) = (BA)\vec{x}.$$

Note we can describe the composite transformation $T_B \circ T_A$ by the matrix BA , which is the product of matrices B and A . **Matrix products enable us to easily compose linear transformations.**

Understanding the connection between matrices and linear transformations helps explain why matrix multiplication is not commutative. In general, $BA \neq AB$: there's no reason to expect AB will equal BA , just as there's no reason to expect that $f \circ g$ will equal $g \circ f$ for two arbitrary functions.

Matrix multiplication is an extremely useful computational tool. At the moment, your feelings about matrix multiplication might not be

so warm and fuzzy, given it can be tedious and repetitive. Be patient and stick with it. Solve some exercises to make sure you understand. Afterward, you can let computers multiply matrices for you—because computers are good at repetitive tasks.

Row operations as matrix products

There is an important connection between row operations and matrix multiplication. Performing the row operation $\mathcal{R}$ on a matrix is equivalent to a left multiplication by an *elementary matrix* $E_{\mathcal{R}}$:

$$A' = \mathcal{R}(A) \quad \Leftrightarrow \quad A' = E_{\mathcal{R}}A.$$

There are three types of elementary matrices that correspond to the three types of row operations. Let's look at an example.

Example The row operation of adding m times the second row to the first row ($\mathcal{R} : R_1 \leftarrow R_1 + mR_2$) corresponds to the elementary matrix

$$E_{\mathcal{R}} = \begin{bmatrix} 1 & m \\ 0 & 1 \end{bmatrix}, \text{ which acts as } \begin{bmatrix} 1 & m \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} a + mc & b + md \\ c & d \end{bmatrix}.$$

We'll discuss elementary matrices in more detail in Section 3.5.

We can also perform column operations on matrices if we multiply them by elementary matrices from the right.

Exercises

E3.5 Compute the following matrix products:

$$P = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \begin{bmatrix} 5 & 6 \\ 7 & 8 \end{bmatrix} \quad \text{and} \quad Q = \begin{bmatrix} 3 & 1 & 2 & 2 \\ 0 & 2 & -2 & 1 \end{bmatrix} \begin{bmatrix} -2 & 3 \\ 1 & 0 \\ -2 & -2 \\ 2 & 2 \end{bmatrix}.$$

E3.6 Consider the following matrices of different dimensions:

$$A = \begin{bmatrix} 2 & 3 \\ 5 & 6 \end{bmatrix}, \quad B = \begin{bmatrix} 7 & 8 \\ 1 & 4 \end{bmatrix}, \quad C = \begin{bmatrix} -1 & 4 \\ 2 & 1 \\ -1 & 2 \end{bmatrix}, \quad U = [2 \quad 1], \quad \text{and} \quad V = \begin{bmatrix} 1 \\ -2 \end{bmatrix}.$$

Compute the following matrix expressions.

$$\text{a) } A^2 \quad \text{b) } B^2 \quad \text{c) } AB \quad \text{d) } BA \quad \text{e) } AC \quad \text{f) } CA \quad \text{g) } UABV$$

3.4 Determinants

What is the volume of a rectangular box of length 1 m, width 2 m, and height 3 m? It's easy to compute the volume of this box because its shape is a *right rectangular prism*. The volume of this rectangular prism is $V = \ell \times w \times h = 6 \text{ m}^3$. What if the shape of the box was a *parallelepiped* instead? A parallelepiped is a box whose opposite faces are parallel but whose sides are slanted, as shown in Figure 3.9 on page 189. How do we compute the volume of a parallelepiped? The determinant operation, specifically the 3×3 determinant, is the perfect tool for this purpose.

The determinant operation takes square matrices as inputs and produces numbers as outputs:

$$\det : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}.$$

The determinant of a matrix, denoted $\det(A)$ or $|A|$, is a particular way to multiply the entries of the matrix to produce a single number. We use determinants for all kinds of tasks: to compute areas and volumes, to solve systems of equations, to check whether a matrix is invertible or not, etc.

We can intuitively interpret the determinant of a matrix as a geometric calculation. The determinant is the “volume” of the geometric shape whose “sides” are the rows of the matrix. For 2×2 matrices, the determinant corresponds to the area of a parallelogram. For 3×3 matrices, the determinant corresponds to the volume of a parallelepiped. For dimensions $d > 3$, we say the determinant measures a d -dimensional hypervolume.

Consider the linear transformation $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined through the matrix-vector product with a matrix A_T : $T(\vec{x}) = A_T \vec{x}$. The determinant of the matrix A_T is the *scale factor* associated with the linear transformation T . The scale factor of the linear transformation T describes how the area of a unit square in the input space (a square with dimensions 1×1) is transformed by T . After passing through T , the unit square is transformed to a parallelogram with area $\det(A_T)$. Linear transformations that “shrink” areas have $\det(A_T) < 1$, while linear transformations that “enlarge” areas have $\det(A_T) > 1$. A linear transformation that is *area preserving* has $\det(A_T) = 1$.

The determinant is also used to check linear independence for a given set of vectors. We construct a matrix using the vectors as the matrix rows, and compute its determinant. If the determinant is nonzero, the vectors are linearly independent.

The determinant of a matrix tells us whether or not the matrix is invertible. If $\det(A) \neq 0$, then A is invertible; if $\det(A) = 0$, A is not invertible.




Figure 3.6: A square with side length 1 in the input space of T is transformed to a parallelogram with area $\det(A_T)$ in the output space of T . The determinant measures the *scale factor* by which the area changes.

The determinant shares a connection with the vector cross product, and is also used in the definition of the eigenvalue equation.

In this section, we'll discuss all the applications of determinants. As you read along, I encourage you to actively draw connections between the geometric, algebraic, and computational aspects of determinants. Don't worry if it doesn't all click right away—you can always review this section once you've learned more about linear transformations, the geometry of cross products, and eigenvalues.

Formulas

The determinant of a 2×2 matrix is

$$\det\left(\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}\right) = \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix} = a_{11}a_{22} - a_{12}a_{21}.$$

The formulas for the determinants of larger matrices are defined recursively. For example, the determinant of a 3×3 matrix is defined in terms of 2×2 determinants:

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

There's a neat computational trick for computing 3×3 determinants by hand. The trick consists of extending the matrix A into a 3×5 array that contains copies of the columns of A : the first column

of A is copied to the fourth column of the extended array, and the second column of A is copied to the fifth column. The determinant is then computed by summing the products of the entries on the three positive diagonals and subtracting the products of the entries on the three negative diagonals, as illustrated in Figure 3.7.




Figure 3.7: Computing the determinant using the extended array trick. The solid lines indicate the positive terms while the dashed lines indicate the negative terms in the determinant calculation.

The general formula for the determinant of an $n \times n$ matrix is

$$\det(A) = \sum_{j=1}^n (-1)^{1+j} a_{1j} M_{1j},$$

where M_{ij} is called the *minor* associated with the entry a_{ij} . The minor M_{ij} is the determinant of the submatrix obtained by removing the i^{th} row and the j^{th} column of the matrix A . Note the factor $(-1)^{i+j}$ that alternates between the values $+1$ and -1 for different terms in the formula.

In the case of 3×3 matrices, applying the determinant formula gives the correct formula,

$$\begin{aligned} \det(A) &= (+1)a_{11}M_{11} + (-1)a_{12}M_{12} + (+1)a_{13}M_{13} \\ &= a_{11} \begin{vmatrix} a_{22} & a_{23} \\ a_{32} & a_{33} \end{vmatrix} - a_{12} \begin{vmatrix} a_{21} & a_{23} \\ a_{31} & a_{33} \end{vmatrix} + a_{13} \begin{vmatrix} a_{21} & a_{22} \\ a_{31} & a_{32} \end{vmatrix}. \end{aligned}$$

The determinant of a 4×4 matrix B is

$$\det(B) = b_{11}M_{11} - b_{12}M_{12} + b_{13}M_{13} - b_{14}M_{14}.$$

The determinant formula, $\det(A) = \sum_{j=1}^n (-1)^{1+j} a_{1j} M_{1j}$, assumes we *expand* the determinant along the first row of the matrix. In fact, a determinant formula can be obtained by expanding the determinant along *any* row or column of the matrix. For example, expanding the determinant of a 3×3 matrix along the second column produces the determinant formula

$$\det(A) = \sum_{i=1}^3 (-1)^{i+2} a_{i2} M_{i2} = (-1)a_{12}M_{12} + (1)a_{22}M_{22} + (-1)a_{32}M_{32}.$$

The expand-along-any-row-or-column nature of determinants can be very handy: if you need to calculate the determinant of a matrix with one row (or column) containing many zero entries, it makes sense to expand along that row since many of the terms in the formula will be zero. If a matrix contains a row (or column) consisting entirely of zeros, we can immediately tell its determinant is zero.

Geometric interpretation

Area of a parallelogram

Suppose we're given vectors $\vec{v} = (v_1, v_2)$ and $\vec{w} = (w_1, w_2)$ in $\mathbb{R}^2$ and we construct a parallelogram with corner points $(0, 0)$, $\vec{v}$, $\vec{w}$, and $\vec{v} + \vec{w}$.

The area of this parallelogram is equal to the determinant of the matrix that contains (v_1, v_2) and (w_1, w_2) as rows:

$$\text{area} = \begin{vmatrix} v_1 & v_2 \\ w_1 & w_2 \end{vmatrix} = v_1 w_2 - v_2 w_1.$$




Figure 3.8: The determinant of a 2×2 matrix corresponds to the area of the parallelogram constructed from the rows of the matrix.

Volume of a parallelepiped

Suppose we're given three vectors: $\vec{u} = (u_1, u_2, u_3)$, $\vec{v} = (v_1, v_2, v_3)$, and $\vec{w} = (w_1, w_2, w_3)$ in $\mathbb{R}^3$. Then suppose we construct the parallelepiped with corner points $(0, 0, 0)$, $\vec{u}$, $\vec{v}$, $\vec{w}$, $\vec{v} + \vec{w}$, $\vec{u} + \vec{w}$, $\vec{u} + \vec{v}$, and $\vec{u} + \vec{v} + \vec{w}$, as illustrated in Figure 3.9.

The volume of this parallelepiped is equal to the determinant of the matrix containing the vectors $\vec{u}$, $\vec{v}$, and $\vec{w}$ as rows:

$$\begin{aligned} \text{volume} &= \begin{vmatrix} u_1 & u_2 & u_3 \\ v_1 & v_2 & v_3 \\ w_1 & w_2 & w_3 \end{vmatrix} \\ &= u_1(v_2 w_3 - v_3 w_2) - u_2(v_1 w_3 - v_3 w_1) + u_3(v_1 w_2 - v_2 w_1). \end{aligned}$$




Figure 3.9: The determinant of a 3×3 matrix corresponds to the volume of the parallelepiped constructed from the rows of the matrix.

Sign and absolute value of the determinant

Calculating determinants can produce positive or negative numbers. Consider the vectors $\vec{v} = (v_1, v_2) \in \mathbb{R}^2$ and $\vec{w} = (w_1, w_2) \in \mathbb{R}^2$ and the determinant

$$D = \det \begin{pmatrix} v_1 & v_2 \\ w_1 & w_2 \end{pmatrix} = v_1 w_2 - v_2 w_1.$$

Let's denote the value of the determinant by the variable D . The absolute value of the determinant is equal to the area of the parallelogram constructed by the vectors $\vec{v}$ and $\vec{w}$. The sign of the determinant (positive, negative, or zero) tells us information about the relative orientation of the vectors $\vec{v}$ and $\vec{w}$. If we let θ be the measure of the angle from $\vec{v}$ toward $\vec{w}$, then the following possibilities exist.

- If θ is between 0 and π (180°), the determinant is positive $D > 0$. This is the case illustrated in Figure 3.8.
- If θ is between π (180°) and 2π (360°), the determinant is negative $D < 0$.
- When $\theta = 0$ (the vectors point in the same direction), or when $\theta = \pi$ (the vectors point in opposite directions), the determinant is zero, $D = 0$.

The formula for the area of a parallelogram is $A = b \times h$, where b is the length of the parallelogram's base, and h is the parallelogram's height. In the case of the parallelogram in Figure 3.8, the length of the base is $\|\vec{v}\|$ and the height is $\|\vec{w}\| \sin \theta$, where θ is the angle measure between $\vec{v}$ and $\vec{w}$. The geometric interpretation of the 2×2 determinant is described by the formula,

$$D = \det \begin{pmatrix} v_1 & v_2 \\ w_1 & w_2 \end{pmatrix} = \underbrace{\|\vec{v}\|}_b \underbrace{\|\vec{w}\| \sin \theta}_h.$$

Observe the "height" of the parallelogram is negative when θ is between π and 2π .

Properties

Let A and B be two square matrices of the same dimension. The determinant operation has the following properties:

- $\det(AB) = \det(A) \det(B) = \det(B) \det(A) = \det(BA)$
- If $\det(A) \neq 0$, the matrix is invertible and $\det(A^{-1}) = \frac{1}{\det(A)}$
- $\det(A^T) = \det(A)$
- $\det(\alpha A) = \alpha^n \det(A)$, for an $n \times n$ matrix A
- $\det(A) = \prod_{i=1}^n \lambda_i$, where $\{\lambda_i\} = \text{eig}(A)$ are the eigenvalues of A

The effects of row operations on determinants

Recall the three row operations we used for the Gauss–Jordan elimination procedure:

- Add a multiple of one row to another row
- Swap two rows
- Multiply a row by a constant

We'll now describe the effects of these row operations on the value of the matrix determinant. In each case, we'll connect the effects of the row operation to the geometric interpretation of the determinant operation.

Add a multiple of one row to another row

Adding a multiple of one row of a matrix to another row does not change the determinant of the matrix.

$$\begin{vmatrix} r_1 \\ r_2 \\ r_3 \end{vmatrix} = \begin{vmatrix} r_1 + \alpha r_2 \\ r_2 \\ r_3 \end{vmatrix}$$

Figure 3.10: Row operations of the form $\mathcal{R}_\alpha : R_i \leftarrow R_i + \alpha R_j$ do not change the value of the matrix determinant.

This property follows from the fact that parallelepipeds with equal base enclosed between two parallel planes have the same volume even if they have different slants. This is called *Cavalieri's principle*.

It is easier to visualize Cavalieri's principle in two dimensions by considering two parallelograms with base b and different slants, enclosed between two parallel lines. The area of both parallelograms is the same: $A = b \times h$, where h is the distance between the parallel lines.

Swapping rows

Swapping two rows of a matrix changes the sign of its determinant.

$$\begin{vmatrix} r_1 \\ r_2 \\ r_3 \end{vmatrix} = -1 \begin{vmatrix} r_2 \\ r_1 \\ r_3 \end{vmatrix}$$

Figure 3.11: Swapping any two rows, $\mathcal{R}_\beta : R_i \leftrightarrow R_j$, flips the sign of the determinant.

This property is a consequence of measuring *signed volumes*. Swapping two rows changes the relative orientation of the vectors, and hence changes the sign of the volume from positive to negative or vice versa.

Multiply a row by a constant

Multiplying a row by a constant is equivalent to multiplying the determinant by the constant.

$$\begin{vmatrix} r_1 \\ \alpha r_2 \\ r_3 \end{vmatrix} = \alpha \begin{vmatrix} r_1 \\ r_2 \\ r_3 \end{vmatrix}$$

Figure 3.12: Row operations of the form $\mathcal{R}_\gamma : R_i \leftarrow \alpha R_i$ scale the value of the determinant by the factor α .

This property follows from the fact that making one side of the parallelepiped α times longer increases the volume of the parallelepiped by a factor of α .

When each entry of an $n \times n$ matrix is multiplied by the constant α , each of the n rows is multiplied by α , and so the determinant changes by a factor of α^n : $\det(\alpha A) = \alpha^n \det(A)$.

Zero-vs-nonzero determinant property

There is an important distinction between matrices with zero determinant and matrices with nonzero determinant. We can understand this distinction geometrically by considering the 3×3 determinant calculation. Recall, the volume of the parallelepiped with sides $\vec{u}$, $\vec{v}$, and $\vec{w}$ is equal to the determinant of the matrix containing the vectors $\vec{u}$, $\vec{v}$, and $\vec{w}$ as rows. If the determinant is zero, it means at least one of the rows of the matrix is a linear combination of the other rows. The volume associated with this determinant is zero because the geometric shape it corresponds to is a flattened, two-dimensional

parallelepiped; in other words, a parallelogram. We say the matrix is “deficient” if its determinant is zero.

On the other hand, if the determinant of a matrix is nonzero, the rows of the matrix are linearly independent. In this case, the determinant calculation corresponds to the volume of a real parallelepiped. We say the matrix is “full” if its determinant is nonzero.

The zero-vs-nonzero determinant property of a matrix does not change when we perform row operations on the matrix. If a matrix A has a nonzero determinant, we know its reduced row echelon form will also have nonzero determinant. The number of nonzero rows in the reduced row echelon form of the matrix is called the *rank* of the matrix. We say a matrix $A \in \mathbb{R}^{n \times n}$ has *full rank* if its RREF contains n pivots. If the RREF of the matrix A contains a row of zeros, then A is not full rank and $\det(A) = 0$. On the other hand, if $\det(A) \neq 0$, we know that $\text{rref}(A) = \mathbb{I}$.

Applications

Apart from the geometric and invertibility-testing applications of determinants described above, determinants are related to many other topics in linear algebra. We’ll briefly cover some of these in this section.

Cross product as a determinant

We can compute the cross product of the vectors $\vec{v} = (v_1, v_2, v_3)$ and $\vec{w} = (w_1, w_2, w_3)$ by computing the determinant of a special matrix. We place the symbols $\hat{i}$, $\hat{j}$, and $\hat{k}$ in the first row of the matrix then write the components of $\vec{v}$ and $\vec{w}$ in the second and third rows. After expanding the determinant along the first row, we obtain the cross product:

$$\begin{aligned}\vec{v} \times \vec{w} &= \begin{vmatrix} \hat{i} & \hat{j} & \hat{k} \\ v_1 & v_2 & v_3 \\ w_1 & w_2 & w_3 \end{vmatrix} \\ &= \hat{i} \begin{vmatrix} v_2 & v_3 \\ w_2 & w_3 \end{vmatrix} - \hat{j} \begin{vmatrix} v_1 & v_3 \\ w_1 & w_3 \end{vmatrix} + \hat{k} \begin{vmatrix} v_1 & v_2 \\ w_1 & w_2 \end{vmatrix} \\ &= (v_2 w_3 - v_3 w_2) \hat{i} - (v_1 w_3 - v_3 w_1) \hat{j} + (v_1 w_2 - v_2 w_1) \hat{k} \\ &= (v_2 w_3 - v_3 w_2, v_3 w_1 - v_1 w_3, v_1 w_2 - v_2 w_1).\end{aligned}$$

Observe that the anticommutative property of the vector cross product, $\vec{v} \times \vec{w} = -\vec{w} \times \vec{v}$, corresponds to the swapping-rows-changes-the-sign property of determinants.

The extended array trick for computing 3×3 determinants (see Figure 3.7) doubles as a trick for computing cross products:




Figure 3.13: We can quickly compute the cross product of two vectors using the extended array trick.

Using the correspondence between the cross product and the determinant, we can write the determinant of a 3×3 matrix in terms of the dot product and cross product:

$$\begin{vmatrix} u_1 & u_2 & u_3 \\ v_1 & v_2 & v_3 \\ w_1 & w_2 & w_3 \end{vmatrix} = \vec{u} \cdot (\vec{v} \times \vec{w}).$$

Cramer's rule

Cramer's rule is an approach for solving systems of linear equations using determinants. Consider the following system of equations and its representation as a matrix equation:

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + a_{13}x_3 &= b_1 \\ a_{21}x_1 + a_{22}x_2 + a_{23}x_3 &= b_2 \\ a_{31}x_1 + a_{32}x_2 + a_{33}x_3 &= b_3 \end{aligned} \quad \Leftrightarrow \quad A\vec{x} = \vec{b}.$$

We're looking for the vector $\vec{x} = (x_1, x_2, x_3)$ that satisfies this system of equations.

Begin by writing the system of equations as an augmented matrix:

$$\left[\begin{array}{ccc|c} a_{11} & a_{12} & a_{13} & b_1 \\ a_{21} & a_{22} & a_{23} & b_2 \\ a_{31} & a_{32} & a_{33} & b_3 \end{array} \right] = \left[\begin{array}{ccc|c} | & | & | & | \\ \vec{a}_1 & \vec{a}_2 & \vec{a}_3 & \vec{b} \\ | & | & | & | \end{array} \right].$$

We use the notation $\vec{a}_j$ to denote the j^{th} column of the matrix A , and $\vec{b}$ to denote the column of constants.

Cramer's rule requires computing ratios of determinants. To find x_1 , the first component of the solution vector $\vec{x}$, we compute the following ratio of determinants:

$$x_1 = \frac{\begin{vmatrix} | & | & | \\ \vec{b} & \vec{a}_2 & \vec{a}_3 \\ | & | & | \end{vmatrix}}{\begin{vmatrix} | & | & | \\ \vec{a}_1 & \vec{a}_2 & \vec{a}_3 \\ | & | & | \end{vmatrix}} = \frac{\begin{vmatrix} b_1 & a_{12} & a_{13} \\ b_2 & a_{22} & a_{23} \\ b_3 & a_{32} & a_{33} \end{vmatrix}}{\begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix}}.$$

Basically, we replace the column that corresponds to the unknown we want to solve for (in this case the first column) with the vector of constants $\vec{b}$, and compute the determinant before dividing by $\det(A)$. To find x_2 , we'll need to compute the determinant of a matrix where $\vec{b}$ replaces the second column of A . Similarly, to find x_3 , we replace the third column with $\vec{b}$.

Cramer's rule is a neat computational trick that might come in handy if you ever want to solve for one particular component of the unknown vector $\vec{x}$, without solving for the other components.

Linear independence test

Suppose you're given a set of n , n -dimensional vectors $\{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n\}$ and asked to check whether the vectors are linearly independent.

You could use the Gauss–Jordan elimination procedure to accomplish this task. Write the vectors $\vec{v}_i$ as the rows of a matrix M . Next, use row operations to find the reduced row echelon form (RREF) of the matrix M . Row operations do not change the linear independence between the rows of a matrix, so you can tell whether the rows are independent from the reduced row echelon form of the matrix M .

Alternatively, you can use the *determinant test* as a shortcut to check whether the vectors are linearly independent. If $\det(M)$ is zero, the vectors that form the rows of M are not linearly independent. On the other hand, if $\det(M) \neq 0$, then the rows of M are linearly independent.

Eigenvalues

The determinant operation is used to define the *characteristic polynomial* of a matrix. The characteristic polynomial of A is denoted $p_A(\lambda)$,

and is defined as the determinant of the expression $(A - \lambda \mathbb{1})$:

$$\begin{aligned}
 p_A(\lambda) &\stackrel{\text{def}}{=} \det(A - \lambda \mathbb{1}) \\
 &= \begin{vmatrix} a_{11} - \lambda & a_{12} \\ a_{21} & a_{22} - \lambda \end{vmatrix} \\
 &= (a_{11} - \lambda)(a_{22} - \lambda) - a_{12}a_{21} \\
 &= \lambda^2 - \underbrace{(a_{11} + a_{22})}_{\text{Tr}(A)} \lambda + \underbrace{(a_{11}a_{22} - a_{12}a_{21})}_{\det(A)}.
 \end{aligned}$$

The variable λ (the Greek letter *lambda*) is customarily used in this definition. The roots of the characteristic polynomial are the *eigenvalues* of the matrix A . Observe the coefficient of the linear term in $p_A(\lambda)$ is equal to $-\text{Tr}(A)$ and the constant term equals $\det(A)$. The name *characteristic polynomial* is indeed appropriate since $p_A(\lambda)$ encodes the information about three important properties of the matrix A : its eigenvalues (λ_1, λ_2) , its trace $\text{Tr}(A)$, and its determinant $\det(A)$.

We'll continue to discuss characteristic polynomials and eigenvalues in Section 6.1.

Exercises

E3.7 Find the determinants of the following matrices.

$$\text{a) } \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \quad \text{b) } \begin{bmatrix} 3 & 4 \\ 1 & 2 \end{bmatrix} \quad \text{c) } \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 3 \\ 1 & 2 & 7 \end{bmatrix} \quad \text{d) } \begin{bmatrix} 1 & 2 & 3 \\ 0 & 0 & 0 \\ 1 & 3 & 4 \end{bmatrix}$$

E3.8 Find the volume of the parallelepiped whose sides are the vectors $\vec{u} = (1, 2, 3)$, $\vec{v} = (2, -2, 4)$, and $\vec{w} = (2, 2, 5)$.

E3.9 Determine whether the set of vectors $\{(1, 4, 3), (2, 1, 1), (0, -2, -1)\}$ is linearly dependent or linearly independent.

Links

[More information about determinants from Wikipedia]

<http://en.wikipedia.org/wiki/Determinant>

[http://en.wikipedia.org/wiki/Minor_\(linear_algebra\)](http://en.wikipedia.org/wiki/Minor_(linear_algebra))

[The determinant explained by 3Blue1Brown]

<https://youtube.com/watch?v=Ip3X9L0h2dk>

3.5 Matrix inverse

In this section, we'll learn four different approaches for computing the inverse of a matrix. Since knowing how to compute matrix inverses is a pretty useful skill, learning several approaches is hardly overkill. Note that the matrix inverse is *unique*, so no matter which method you use to find the inverse, you'll always obtain the same answer. You can verify your calculations by computing the inverse in different ways and checking that the answers agree.

Existence of an inverse

Not all matrices are invertible. Given a matrix $A \in \mathbb{R}^{n \times n}$, we can check whether it is invertible by computing its determinant:

$$A^{-1} \text{ exists} \quad \text{if and only if} \quad \det(A) \neq 0.$$

Calculating the determinant of a matrix serves as an *invertibility test*. The exact value of the determinant is not important; it could be big or small, positive or negative; as long as the determinant is nonzero, the matrix passes the invertibility test.

Adjugate matrix approach

The inverse of a 2×2 matrix can be computed as follows:

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

The above formula is the 2×2 version of a general formula for obtaining the inverse based on the *adjugate matrix*:

$$A^{-1} = \frac{1}{\det(A)} \text{adj}(A).$$

What is the adjugate matrix, you ask? The adjugate matrix is kind of complicated, so let's proceed step by step. We'll first define a few prerequisite concepts.

In this section, we'll work on a matrix $A \in \mathbb{R}^{3 \times 3}$ and refer to its entries as a_{ij} , where i is the row index and j is the column index:

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}.$$

First we'll define three concepts associated with determinants:

1. For each entry a_{ij} , the *minor* M_{ij} is the determinant of the matrix that remains after we remove the i^{th} row and the j^{th} column of A . For example, the minor that corresponds to the entry a_{12} is given by:

$$M_{12} = \begin{vmatrix} \times & \times & \times \\ a_{21} & \times & a_{23} \\ a_{31} & \times & a_{33} \end{vmatrix} = \begin{vmatrix} a_{21} & a_{23} \\ a_{31} & a_{33} \end{vmatrix} = a_{21}a_{33} - a_{23}a_{31}.$$

2. The *sign* of each entry a_{ij} is defined as $\text{sign}(a_{ij}) \stackrel{\text{def}}{=} (-1)^{i+j}$. For example, the signs of the different entries in a 3×3 matrix are

$$\begin{pmatrix} + & - & + \\ - & + & - \\ + & - & + \end{pmatrix}.$$

3. The *cofactor* c_{ij} for the entry a_{ij} is the product of this entry's sign and its minor:

$$c_{ij} \stackrel{\text{def}}{=} \text{sign}(a_{ij})M_{ij} = (-1)^{i+j}M_{ij}.$$

The above concepts should look somewhat familiar, since they previously appeared in the formula for computing determinants. If we expand the determinant along the first row of the matrix, we obtain the formula

$$\det(A) = \sum_{j=1}^n a_{1j} \text{sign}(a_{1j})M_{1j} = \sum_{j=1}^n a_{1j}c_{1j}.$$

Here we can see where the name *cofactor* comes from: the cofactor c_{ij} is what multiplies the *factor* a_{ij} in the determinant formula.

Okay, now we're ready to describe the adjugate matrix. The adjugate matrix is defined as the transpose of the *matrix of cofactors* C . The matrix of cofactors is a matrix of the same dimensions as the original matrix A that is constructed by replacing each entry a_{ij} by its cofactor c_{ij} . The matrix of cofactors for A is

$$C = \begin{bmatrix} c_{11} & c_{12} & c_{13} \\ c_{21} & c_{22} & c_{23} \\ c_{31} & c_{32} & c_{33} \end{bmatrix} = \begin{bmatrix} + \begin{vmatrix} a_{22} & a_{23} \\ a_{32} & a_{33} \end{vmatrix} & - \begin{vmatrix} a_{21} & a_{23} \\ a_{31} & a_{33} \end{vmatrix} & + \begin{vmatrix} a_{21} & a_{22} \\ a_{31} & a_{32} \end{vmatrix} \\ - \begin{vmatrix} a_{12} & a_{13} \\ a_{32} & a_{33} \end{vmatrix} & + \begin{vmatrix} a_{11} & a_{13} \\ a_{31} & a_{33} \end{vmatrix} & - \begin{vmatrix} a_{11} & a_{12} \\ a_{31} & a_{32} \end{vmatrix} \\ + \begin{vmatrix} a_{12} & a_{13} \\ a_{22} & a_{23} \end{vmatrix} & - \begin{vmatrix} a_{11} & a_{13} \\ a_{21} & a_{23} \end{vmatrix} & + \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix} \end{bmatrix}.$$

The adjugate matrix $\text{adj}(A)$ is the transpose of the matrix of cofactors:

$$\text{adj}(A) \stackrel{\text{def}}{=} C^T.$$

The formula for the inverse matrix is $A^{-1} = \frac{1}{\det(A)} \text{adj}(A)$. In the 3×3 case, the matrix inverse formula is

$$A^{-1} = \frac{1}{\begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix}} \begin{bmatrix} + \begin{vmatrix} a_{22} & a_{23} \\ a_{32} & a_{33} \end{vmatrix} & - \begin{vmatrix} a_{12} & a_{13} \\ a_{32} & a_{33} \end{vmatrix} & + \begin{vmatrix} a_{12} & a_{13} \\ a_{22} & a_{23} \end{vmatrix} \\ - \begin{vmatrix} a_{21} & a_{23} \\ a_{31} & a_{33} \end{vmatrix} & + \begin{vmatrix} a_{11} & a_{13} \\ a_{31} & a_{33} \end{vmatrix} & - \begin{vmatrix} a_{11} & a_{13} \\ a_{21} & a_{23} \end{vmatrix} \\ + \begin{vmatrix} a_{21} & a_{22} \\ a_{31} & a_{32} \end{vmatrix} & - \begin{vmatrix} a_{11} & a_{12} \\ a_{31} & a_{32} \end{vmatrix} & + \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix} \end{bmatrix}.$$

I know this looks complicated, but I wanted you to know about the adjugate matrix approach for computing the inverse. In practice, you'll rarely have to compute inverses using this approach; nevertheless, the adjugate matrix formula represents an important theoretical concept. Note the formula fails if $\det(A) = 0$, due to a divide-by-zero error.

Using row operations

Another way to obtain the inverse of a matrix is to record all the *row operations* $\mathcal{R}_1, \mathcal{R}_2, \dots$ needed to transform A into the identity matrix:

$$\mathcal{R}_k(\dots \mathcal{R}_2(\mathcal{R}_1(A)) \dots) = \mathbb{I}.$$

Recall that we can think of the action of the matrix A as a vector transformation, $\vec{w} = A\vec{v}$. By definition, the inverse A^{-1} is the operation that undoes the effect of A : $A^{-1}\vec{w} = \vec{v}$. The combination of A followed by A^{-1} is the identity transformation, $A^{-1}A\vec{v} = \mathbb{I}\vec{v} = \vec{v}$.

The cumulative effect of the row operations required to transform A to the identity matrix is equivalent to the “undo” action of A . Applying the sequence of row operations $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k$ has the same effect as multiplying by A^{-1} :

$$A^{-1}\vec{w} = \mathcal{R}_k(\dots \mathcal{R}_2(\mathcal{R}_1(\vec{w})) \dots).$$

If you think this method of finding the inverse A^{-1} seems more complicated than useful, you're right—if it weren't for the existence of a neat procedure for recording the row operations $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k$ that makes everything simpler. We'll discuss this procedure next.

Begin by initializing an $n \times 2n$ array with the entries of the matrix A on the left side and the identity matrix on the right side: $[A \mid \mathbb{I}]$.

If we perform the Gauss–Jordan elimination procedure on this array, we’ll end up with the inverse A^{-1} in the right side of the array:

$$[A \mid \mathbb{1}] \xrightarrow{\text{Gauss–Jordan elimination}} [\mathbb{1} \mid A^{-1}].$$

Example Let’s illustrate the procedure by computing the inverse of the following matrix:

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 9 \end{bmatrix}.$$

We start by writing the matrix A next to the identity matrix $\mathbb{1}$:

$$\left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 3 & 9 & 0 & 1 \end{array} \right].$$

Next, we perform the Gauss–Jordan elimination procedure on the resulting 2×4 array:

1. Subtract three times the first row from the second row, written compactly as $\mathcal{R}_1 : R_2 \leftarrow R_2 - 3R_1$, to obtain

$$\left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 0 & 3 & -3 & 1 \end{array} \right].$$

2. Perform another row operation $\mathcal{R}_2 : R_2 \leftarrow \frac{1}{3}R_2$ to obtain a leading one in the second column:

$$\left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 0 & 1 & -1 & \frac{1}{3} \end{array} \right].$$

3. Finally, perform $\mathcal{R}_3 : R_1 \leftarrow R_1 - 2R_2$ to clear the entry above the leading one in the second column:

$$\left[\begin{array}{cc|cc} 1 & 0 & 3 & -\frac{2}{3} \\ 0 & 1 & -1 & \frac{1}{3} \end{array} \right].$$

The inverse of A now appears in the right side of the above array,

$$A^{-1} = \begin{bmatrix} 3 & -\frac{2}{3} \\ -1 & \frac{1}{3} \end{bmatrix}.$$

This algorithm works because we identify the sequence of row operations $\mathcal{R}_3(\mathcal{R}_2(\mathcal{R}_1(\cdot)))$ with the action of the inverse matrix A^{-1} :

$$\mathcal{R}_3(\mathcal{R}_2(\mathcal{R}_1(A))) = \mathbb{1} \quad \Rightarrow \quad \mathcal{R}_3(\mathcal{R}_2(\mathcal{R}_1(\mathbb{1}))) = A^{-1}.$$

The combined effect of the three row operations is to “undo” the action of A ; therefore this sequence of row operations has the same effect as the inverse operation A^{-1} . The right side of the 2×4 array serves as a record of the cumulative effect of this sequence of row operations performed. Because the right side starts from the identity matrix, it will contain A^{-1} by the end of the procedure.

Using elementary matrices

Every row operation $\mathcal{R}$ performed on a matrix is equivalent to an operation of left multiplication by an *elementary matrix* $E_{\mathcal{R}}$:

$$A' = \mathcal{R}(A) \quad \Leftrightarrow \quad A' = E_{\mathcal{R}}A.$$

There are three types of elementary matrices that correspond to the three types of row operations. We'll illustrate the three types of elementary matrices with examples from the 2×2 case:

- Adding m times the second row to the first row corresponds to

$$\mathcal{R}_{\alpha} : R_1 \leftarrow R_1 + mR_2 \quad \Leftrightarrow \quad E_{\alpha} = \begin{bmatrix} 1 & m \\ 0 & 1 \end{bmatrix}.$$

- Swapping the first and second rows corresponds to

$$\mathcal{R}_{\beta} : R_1 \leftrightarrow R_2 \quad \Leftrightarrow \quad E_{\beta} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

- Multiplying the first row by the constant k corresponds to

$$\mathcal{R}_{\gamma} : R_1 \leftarrow kR_1 \quad \Leftrightarrow \quad E_{\gamma} = \begin{bmatrix} k & 0 \\ 0 & 1 \end{bmatrix}.$$

The general rule is simple: to find the elementary matrix that corresponds to a given row operation, apply that row operation to the identity matrix $\mathbb{I}$.

Recall the procedure we used to find the inverse in the previous section. We applied the sequence of row operations $\mathcal{R}_1, \mathcal{R}_2, \dots$ to transform the array $[A \mid \mathbb{I}]$ into the reduced row echelon form:

$$\mathcal{R}_k(\dots \mathcal{R}_2(\mathcal{R}_1([A \mid \mathbb{I}])) \dots) = [\mathbb{I} \mid A^{-1}].$$

If we represent each row operation as a multiplication by an elementary matrix, we obtain the equation

$$E_k \cdots E_2 E_1 [A \mid \mathbb{I}] = [\mathbb{I} \mid A^{-1}].$$

Observe that $E_k \cdots E_2 E_1 A = \mathbb{1}$, so we've obtained an expression for the inverse matrix A^{-1} as a product of elementary matrices:

$$A^{-1} = E_k \cdots E_2 E_1.$$

We'll now illustrate how the formula $A^{-1} = E_k \cdots E_2 E_1$ applies in the case of the matrix A discussed above

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 9 \end{bmatrix}.$$

Recall the row operations we applied in order to transform $[A \mid \mathbb{1}]$ into $[\mathbb{1} \mid A^{-1}]$:

1. $\mathcal{R}_1: R_2 \leftarrow R_2 - 3R_1$
2. $\mathcal{R}_2: R_2 \leftarrow \frac{1}{3}R_2$
3. $\mathcal{R}_3: R_1 \leftarrow R_1 - 2R_2$

Let's revisit these row operations, representing each of them as a multiplication by an elementary matrix:

1. The first row operation, $\mathcal{R}_1: R_2 \leftarrow R_2 - 3R_1$, corresponds to a multiplication by the elementary matrix E_1 :

$$E_1 = \begin{bmatrix} 1 & 0 \\ -3 & 1 \end{bmatrix}, \quad E_1 A = \begin{bmatrix} 1 & 0 \\ -3 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 3 & 9 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 0 & 3 \end{bmatrix}.$$

2. The second row operation, $\mathcal{R}_2: R_2 \leftarrow \frac{1}{3}R_2$, corresponds to a matrix E_2 :

$$E_2 = \begin{bmatrix} 1 & 0 \\ 0 & \frac{1}{3} \end{bmatrix}, \quad E_2(E_1 A) = \begin{bmatrix} 1 & 0 \\ 0 & \frac{1}{3} \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 0 & 3 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix}.$$

3. The third row operation, $\mathcal{R}_3: R_1 \leftarrow R_1 - 2R_2$, corresponds to the elementary matrix E_3 :

$$E_3 = \begin{bmatrix} 1 & -2 \\ 0 & 1 \end{bmatrix}, \quad E_3(E_2 E_1 A) = \begin{bmatrix} 1 & -2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

Note that $E_3 E_2 E_1 A = \mathbb{1}$, so the expression $E_3 E_2 E_1$ must equal A^{-1} :

$$A^{-1} = E_3 E_2 E_1 = \begin{bmatrix} 1 & -2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \frac{1}{3} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -3 & 1 \end{bmatrix} = \begin{bmatrix} 3 & -\frac{2}{3} \\ -1 & \frac{1}{3} \end{bmatrix}.$$

Verify the last equation by computing the product of the three elementary matrices.

Since we know $A^{-1} = E_3 E_2 E_1$, then $A = (A^{-1})^{-1} = (E_3 E_2 E_1)^{-1} = E_1^{-1} E_2^{-1} E_3^{-1}$. We can write A as a product of elementary matrices:

$$A = E_1^{-1} E_2^{-1} E_3^{-1} = \begin{bmatrix} 1 & 0 \\ 3 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix}.$$

The inverses of the elementary matrices are easy to compute; they correspond to the elementary “undo” operations of E_1 , E_2 , and E_3 . For example, E_1 performs the row operation $R_2 \leftarrow R_2 - 3R_1$, so its “undo” operation is $R_2 \leftarrow R_2 + 3R_1$, which corresponds to $E_1^{-1} = \begin{bmatrix} 1 & 0 \\ 3 & 1 \end{bmatrix}$.

The elementary matrix approach teaches us that every invertible matrix A can be decomposed as the product of elementary matrices. The inverse matrix A^{-1} consists of the product of the inverses of the elementary matrices that make up A (in the reverse order).

Using a computer algebra system

You can use a computer algebra system to specify matrices and compute their inverses. Let’s illustrate how to find the matrix inverse using the computer algebra system at <http://live.sympy.org>.

```
>>> from sympy.matrices import Matrix
>>> A = Matrix( [ [1,2],[3,9] ] )    # define a Matrix object
>>> A.inv()                          # call the inv method on A
[ 3, -2/3]
[-1, 1/3]
```

Note SymPy returns an answer in terms of exact rational numbers. This is in contrast with numerical computer algebra systems like Octave and MATLAB, which are based on floating point arithmetic.

Discussion

We’ve explored several ways to compute matrix inverses. If you need to find the inverse of a matrix using only pen and paper, on a final exam for example, I recommend using the Gauss–Jordan elimination procedure on the extended array:

$$[A \mid \mathbb{I}] \text{ -- G-J elimination } \rightarrow [\mathbb{I} \mid A^{-1}],$$

since it is fairly easy to perform even for large matrices and it leverages your experience with the Gauss–Jordan elimination procedure.

For 2×2 matrices, the formula $\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad-bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$ is faster.

Invertibility

Not all matrices are invertible. Keep this in mind, since teachers might try to trick you by asking you to find the inverse of a non-invertible matrix. Let's analyze how each procedure for computing the inverse fails when applied to a non-invertible matrix D . The inverse formula based on the determinant and the adjugate matrix is $D^{-1} = \frac{1}{\det(D)} \text{adj}(D)$. However, if the matrix D is not invertible, then $\det(D) = 0$ and the formula fails due to a divide-by-zero error. The row operations approach to computing the inverse will also fail. Starting from the extended array $[D \mid \mathbb{1}]$, you can apply all the row operations you want, but you'll never be able to obtain the identity matrix in the left half of the extended array. This is because the reduced row echelon form of a non-invertible matrix D has at least one row of zeros: $\text{rref}(D) \neq \mathbb{1}$. We'll discuss invertible matrices and their properties in Section 5.4. For now, be sure to remember the determinant *invertibility test*: if $\det(A) = 0$, then A is non-invertible, and if $\det(A) \neq 0$, then A is invertible.

Exercises

E3.10 For what values of α is the matrix $A = \begin{bmatrix} 2 & 3 \\ 4 & \alpha \end{bmatrix}$ invertible?

E3.11 Compute A^{-1} where $A = \begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix}$.

E3.12 Solve for x and y in the equation $\begin{bmatrix} 1 & 3 \\ -1 & -2 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$.

E3.13 Show that for an $n \times n$ invertible matrix A , the determinant of the adjugate matrix is $|\text{adj}(A)| = (|A|)^{n-1}$.

Hint: Recall that $|A^{-1}| = \frac{1}{|A|}$ and $|\alpha A| = \alpha^n |A|$.

3.6 Computational problems

We've reached the problem section where you're supposed to practice all the computational techniques of linear algebra. This is not going to be the most exciting three hours of your life, but you'll get through it. You need to know how to solve computational problems by hand and apply the Gauss–Jordan elimination procedure; and you need to know how to multiply matrices, calculate determinants, and find matrix inverses. These computational techniques enable all the advanced procedures we'll develop later in the book. If you skip these practice problems, you'll have trouble later when it comes to mastering more advanced topics that rely on these basic matrix operations as building blocks. Do this important work now, and you'll be on your way to becoming fluent in linear algebra computations... plus, the rest of the book will be much more pleasant.

P3.1 Mitchell is on a new diet. His target is to eat exactly 25 grams of fat and 32 grams of protein for lunch today. There are two types of food in the fridge, x and y . One serving of food x contains one gram of fat and two grams of protein, while a serving of food y contains five grams of fat and one gram of protein. To figure out how many servings of each type of food he should eat, Mitchell writes the following system of equations:

$$\begin{array}{rcrcrcrc} x & + & 5y & = & 25 \\ 2x & + & y & = & 32 \end{array} \quad \Rightarrow \quad \left[\begin{array}{cc|c} 1 & 5 & 25 \\ 2 & 1 & 32 \end{array} \right].$$

Help Mitchell find how many servings of x and y he should eat.

Hint: Find the reduced row echelon form of the augmented matrix.

P3.2 Alice, Bob, and Charlotte are solving this system of equations:

$$\begin{array}{rcrcrcrc} 3x & + & 3y & = & 6 \\ 2x & + & \frac{3}{2}y & = & 5 \end{array} \quad \Rightarrow \quad \left[\begin{array}{cc|c} 3 & 3 & 6 \\ 2 & \frac{3}{2} & 5 \end{array} \right].$$

Alice follows the standard procedure to obtain a leading one by performing the row operation $R_1 \leftarrow \frac{1}{3}R_1$. Bob starts with a different row operation, applying $R_1 \leftarrow R_1 - R_2$ to obtain a leading one. Charlotte takes a third approach by swapping the first and second rows: $R_1 \leftrightarrow R_2$. Their respective versions of the augmented matrix are shown below.

a) $\left[\begin{array}{cc|c} 1 & 1 & 2 \\ 2 & \frac{3}{2} & 5 \end{array} \right]$

b) $\left[\begin{array}{cc|c} 1 & \frac{3}{2} & 1 \\ 2 & \frac{3}{2} & 5 \end{array} \right]$

c) $\left[\begin{array}{cc|c} 2 & \frac{3}{2} & 5 \\ 3 & 3 & 6 \end{array} \right]$

Help Alice, Bob, and Charlotte finish solving the system of equations by writing the remaining row operations each of them must perform to bring their version of the augmented matrix into reduced row echelon form.

P3.3 Find the solutions to the systems of equations that correspond to the following augmented matrices.

$$\text{a) } \left[\begin{array}{cc|c} -1 & -2 & -2 \\ 3 & 3 & 0 \end{array} \right] \quad \text{b) } \left[\begin{array}{ccc|c} 1 & -1 & -2 & 1 \\ -2 & 3 & 3 & -1 \\ -1 & 0 & 1 & 2 \end{array} \right] \quad \text{c) } \left[\begin{array}{ccc|c} 2 & -2 & 3 & 2 \\ 1 & -2 & -1 & 0 \\ -2 & 2 & 2 & 1 \end{array} \right]$$

P3.4 Find the solution sets for the systems of equations described by the following augmented matrices.

$$\text{a) } \left[\begin{array}{cc|c} -1 & -2 & -2 \\ 3 & 6 & 6 \end{array} \right] \quad \text{b) } \left[\begin{array}{ccc|c} 1 & -1 & -2 & 1 \\ -2 & 3 & 3 & -1 \\ -1 & 2 & 1 & 0 \end{array} \right] \quad \text{c) } \left[\begin{array}{ccc|c} 2 & -2 & 3 & 2 \\ 0 & 0 & 5 & 3 \\ -2 & 2 & 2 & 1 \end{array} \right]$$

P3.5 Find the solution sets for the augmented matrices.

$$\text{a) } \left[\begin{array}{ccc|c} 1 & -1 & -2 & 1 \\ -2 & 2 & 4 & -2 \\ 3 & -3 & -6 & 3 \end{array} \right] \quad \text{b) } \left[\begin{array}{ccc|c} 2 & -2 & 3 & 2 \\ 0 & 0 & 5 & 3 \\ 6 & -6 & -1 & 0 \\ 6 & -6 & 9 & 6 \end{array} \right]$$

P3.6 Find the solutions to the systems of equations.

$$\text{a) } \left[\begin{array}{ccc|c} 2 & 1 & -1 & 0 \\ 0 & 1 & 1 & 0 \\ 4 & 2 & -2 & 0 \end{array} \right] \quad \text{b) } \left[\begin{array}{ccc|c} 2 & 0 & 1 & 5 \\ 1 & 4 & 2 & 2 \\ 0 & 2 & 1 & 1 \end{array} \right] \quad \text{c) } \left[\begin{array}{ccc|c} 1 & 1 & 2 & 2 \\ 4 & -2 & -4 & 5 \\ 3 & -3 & -6 & 3 \end{array} \right]$$

P3.7 Consider an unknown matrix $A \in \mathbb{R}^{5 \times 5}$. You're told the system of linear equations $A\vec{x} = \vec{b}$ has an infinite number of solutions. What is the maximum rank of the matrix A ?

P3.8 Solve for C in the matrix equation $ABCD = AD$.

P3.9 Solve for the following matrix equations problems:

- Simplify the expression $MNB^{-1}BK^{-1}KN^{-1}M^{-2}L^{-1}S^{-1}SMK^2$.
- Simplify $J^{-3}K^2G^{-1}GK^{-3}J^2$.
- Solve for A in the equation $A^{-1}BNK = 2B^2B^{-1}NK$.
- Solve for Y in $SUNNY = SUN$.

You can assume all matrices are invertible.

P3.10 Solve for $\vec{x}$ in $A\vec{x} = \vec{b}$, where $A = \begin{bmatrix} 1 & 0 & -3 \\ 2 & -1 & 1 \\ 0 & 0 & -1 \end{bmatrix}$ and $\vec{b} = (2, 2, 3)^T$.

Hint: Express the equation $A\vec{x} = \vec{b}$ as an augmented matrix.

P3.11 Solve for $\vec{x}$ in the equation $\vec{x} = \vec{d} + A\vec{x}$, where $A = \begin{bmatrix} 0 & 0.05 & 0.3 \\ 0.01 & 0 & 0.01 \\ 0.1 & 0 & 0 \end{bmatrix}$, and $\vec{d} = (25, 10, 14)^T$. Use `live.sympy.org` to perform the calculations.

Hint: Rewrite as $\mathbb{1}\vec{x} = \vec{d} + A\vec{x}$, then bring all the $\vec{x}$ s to one side.

P3.12 Given the following two matrices,

$$A = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 4 & 2 \\ 3 & 1 & 0 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 3 & 1 \\ 2 & 0 \\ 1 & 1 \end{bmatrix},$$

compute the matrix products **a)** AB , **b)** AA , **c)** BA , and **d)** BB .

P3.13 Compute the product of three matrices:

$$\begin{bmatrix} 2 & 10 & -5 & 0 \\ 0 & 0 & 1 & 3 \end{bmatrix} \begin{bmatrix} 1 & 3 \\ 0 & 2 \\ 5 & 1 \\ -3 & -4 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}.$$

P3.14 Consider the following three matrices:

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad \text{and} \quad H = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}.$$

Show that $HXH = Z$ and that $HZH = X$.

P3.15 Given the matrices

$$L = \begin{bmatrix} -1 & 1 & 3 \\ 3 & 0 & 3 \\ 3 & 2 & 1 \end{bmatrix}, \quad M = \begin{bmatrix} -1 & 3 & 3 & 1 \\ -1 & 4 & -1 & 2 \\ -2 & 2 & 3 & -2 \end{bmatrix}, \quad \text{and} \quad N = \begin{bmatrix} 5 & 3 \\ 0 & 5 \\ 3 & -2 \end{bmatrix},$$

compute the value of the following matrix products:

$$\mathbf{a)} L^2 \quad \mathbf{b)} LM \quad \mathbf{c)} LN \quad \mathbf{d)} M^T L \quad \mathbf{e)} N^T L \quad \mathbf{f)} N^T LM$$

P3.16 Given an unknown variable $\alpha \in \mathbb{R}$ and the matrices

$$A = \begin{bmatrix} \cos(\alpha) & 1 \\ -1 & -\sin(\alpha) \end{bmatrix}; \quad B = \begin{bmatrix} \sin(\alpha) & 0 \\ 0 & -\sin(\alpha) \end{bmatrix}; \quad C = \begin{bmatrix} 1 & -\cos(\alpha) \\ \sin(\alpha) & 1 \end{bmatrix},$$

compute the value of **a)** $A^2 + B^2$, **b)** $A^2 + C$, and **c)** $A^2 + C - B^2$. Give your answer in terms of α and use the double-angle formulas as needed.

P3.17 Find the determinants of the following matrices.

$$\mathbf{a)} \begin{bmatrix} 2 & 1 \\ 3 & 0 \end{bmatrix} \quad \mathbf{b)} \begin{bmatrix} 0 & 5 & 3 \\ 0 & 1 & 1 \\ 0 & 1 & 0 \end{bmatrix} \quad \mathbf{c)} \begin{bmatrix} 1 & 2 & 0 \\ 3 & 1 & 1 \\ 4 & -2 & 0 \end{bmatrix}$$

P3.18 Find the determinants of the matrices.

$$A = \begin{bmatrix} 3 & -1 & 5 & 2 \\ 0 & 2 & 2 & -3 \\ 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & -2 \end{bmatrix} \quad B = \begin{bmatrix} 2 & -1 & 0 & -3 & 2 \\ 0 & 1 & 1 & 3 & 0 \\ 1 & 4 & 0 & 0 & -1 \\ 3 & -2 & 3 & 1 & 0 \\ -1 & 0 & -1 & 0 & 2 \end{bmatrix}$$

P3.19 Determine if the following sets of vectors are linearly dependent or linearly independent.

$$\mathbf{a)} \{(1, 1, 0), (1, 0, 1), (0, 1, 1)\} \quad \mathbf{b)} \{(1, 1, 0), (1, 0, 1), (1, 3, -2)\} \\ \mathbf{c)} \{(1, 2, 3, 4), (-1, -2, 1, 0), (0, 0, 1, 1), (1, 0, 0, 1)\}$$

P3.20 Find the area of a parallelogram that has vectors $\vec{v} = (3, -5)$ and $\vec{w} = (1, -1)$ as its sides.

Hint: Use the formula from Section 3.4 (page 188).

P3.21 Find the volume of the parallelepiped that has the vectors $\vec{u} = (2, 0, 1)$, $\vec{v} = (1, -1, 1)$, and $\vec{w} = (0, 2, 3)$ as sides. See Figure 3.9 for an illustration.

P3.22 Suppose M and N are unknown 4×4 matrices with $|M| = -2$ and $|N| = 7$. Compute the values of these determinant expressions:

a) $|M^T|$ b) $|3M|$ c) $|M^3|$ d) $|MN|$ e) $|M^T N M|$

Hint: Use the properties of the determinant operation.

P3.23 Given the matrix

$$A = \begin{bmatrix} 3 & -2 & 0 & 1 \\ 0 & 1 & 3 & -1 \\ 5 & 0 & 1 & 4 \\ 0 & 3 & -4 & 2 \end{bmatrix},$$

- Find the determinant of A .
- Find the determinant when you interchange the first and third rows.
- Find the determinant after multiplying the second row by -2 .

P3.24 Check whether the rows of the following matrices are linearly independent:

$$A = \begin{bmatrix} 1 & 3 \\ 2 & 6 \end{bmatrix}$$

$$B = \begin{bmatrix} 1 & 4 & 3 \\ 2 & 1 & 1 \\ 0 & -2 & -1 \end{bmatrix}$$

$$C = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 2 & -2 & 2 & -2 \\ -4 & 4 & -4 & 4 \\ -8 & 8 & -8 & 8 \end{bmatrix}$$

$$D = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 2 & 0 & -2 & 0 \\ -1 & 2 & 1 & -2 \\ 1 & -1 & -1 & 1 \end{bmatrix}$$

Are the *columns* of these matrices linearly independent?

P3.25 The transformation from *polar coordinates* (r, θ) to Cartesian coordinates (x, y) is given by the equations $x(r, \theta) = r \cos \theta$ and $y(r, \theta) = r \sin \theta$. Under this transformation, “a little piece of area” $dx dy$ is transformed to “a little piece of area” $\det(J) dr d\theta$, where $\det(J)$ is the *area scale factor* of the transformation from polar coordinates to Cartesian coordinates. The matrix J contains the partial derivatives of $x(r, \theta)$ and $y(r, \theta)$, and is called the *Jacobian matrix* of the transformation:

$$J = \begin{bmatrix} \frac{\partial x}{\partial r} & \frac{\partial x}{\partial \theta} \\ \frac{\partial y}{\partial r} & \frac{\partial y}{\partial \theta} \end{bmatrix}.$$

Compute the value of $\det(J)$.

P3.26 Spherical coordinates (ρ, θ, ϕ) are described by

$$\begin{aligned} x &= \rho \sin \phi \cos \theta, \\ y &= \rho \sin \phi \sin \theta, \\ z &= \rho \cos \phi. \end{aligned}$$

Small “volume chunks” transform according to $dx dy dz = \det(J_s) d\phi d\theta d\rho$, where $\det(J_s)$ is the *volume scale factor*, computed as the determinant of the Jacobian matrix of the change-of-coordinates transformation:

$$J_s = \begin{bmatrix} \frac{\partial x}{\partial \rho} & \frac{\partial x}{\partial \theta} & \frac{\partial x}{\partial \phi} \\ \frac{\partial y}{\partial \rho} & \frac{\partial y}{\partial \theta} & \frac{\partial y}{\partial \phi} \\ \frac{\partial z}{\partial \rho} & \frac{\partial z}{\partial \theta} & \frac{\partial z}{\partial \phi} \end{bmatrix}.$$

Compute the absolute value of $\det(J_s)$.

P3.27 Find the inverses of the following matrices:

a) $\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$

b) $\begin{bmatrix} 1 & 2 \\ 2 & 5 \end{bmatrix}$

c) $\begin{bmatrix} 2 & 3 \\ 2 & 4 \end{bmatrix}$

P3.28 Given the matrix equation $AB = C$, where A and C are 2×2 matrices, find the matrix B .

$$A = \begin{bmatrix} 1 & 4 \\ 2 & 7 \end{bmatrix}$$

$$C = \begin{bmatrix} 3 & 2 \\ 1 & -4 \end{bmatrix}$$

P3.29 Find the inverses of the matrices $A = \begin{bmatrix} 1 & 4 & 3 \\ 2 & 0 & -1 \end{bmatrix}$ and $B = \begin{bmatrix} 0 & -3 & 2 & 4 \\ 1 & -1 & 1 & -1 \\ 2 & 4 & 0 & -2 \\ 3 & 0 & 1 & 0 \end{bmatrix}$.

P3.30 Prove that the zero matrix A has no inverse.

P3.31 Obtain the matrices of cofactors for the following matrices.

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

$$B = \begin{bmatrix} 5 & 0 & 1 \\ 3 & -1 & -3 \\ 0 & -4 & -2 \end{bmatrix}$$

$$C = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 2 & 0 & -2 & 0 \\ -1 & 2 & 1 & -2 \\ 1 & -1 & -1 & 1 \end{bmatrix}$$

P3.32 Implement the formula $A^{-1} = \frac{1}{\det(A)} \text{adj}(A)$ for the case of 3×3 matrices using a spreadsheet application like LibreOffice, OpenOffice, Excel, or Google Sheets. Assume the entries of the matrix are specified in the top right corner of the spreadsheet: A1:C3. Start by writing the formula for computing the determinant and the matrix of cofactors, then combine these partial calculations to obtain the nine entries of the matrix inverse. Test that your formula is correct by finding the inverse of $A = \begin{bmatrix} 1 & 4 & 3 \\ 2 & 0 & 1 \\ 0 & -2 & -1 \end{bmatrix}$. Compare your formula's output with the built-in function =MINVERSE(A1:C3).

P3.33 Find the values a, b, c , and d that satisfy the equation

$$\begin{bmatrix} 1 & 3 \\ -2 & -1 \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 3 & -5 \\ 4 & 0 \end{bmatrix}.$$

P3.34 Given the constraints $a = g$, $e = b = f$, and $c = d = h$, find a choice of the variables a, b, c, d, e, f, g, h that satisfies the matrix equation:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} e & f \\ g & h \end{bmatrix} = \begin{bmatrix} -2 & -3 \\ 0 & 2 \end{bmatrix}.$$

P3.35 Given the matrix $A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$, explain how you can obtain the matrices $B = \begin{bmatrix} aa_{11} & aa_{12} \\ \beta a_{21} & \beta a_{22} \end{bmatrix}$ and $C = \begin{bmatrix} aa_{11} & \beta a_{12} \\ aa_{21} & \beta a_{22} \end{bmatrix}$.

Hint: Use matrix multiplication (from the left and from the right).

Chapter 4

Geometric aspects of linear algebra

In this section, we'll study geometric objects like lines, planes, and vector spaces. We'll use what we learned about vectors and matrices in the previous chapters to perform geometric calculations such as projections and distance measurements.

Developing your intuition about the geometric problems of linear algebra is very important: of all the things you learn in this course, your geometric intuition will stay with you the longest. Years from now, you may not recall the details of the Gauss–Jordan elimination procedure, but you'll still remember that the solution to three linear equations in three variables corresponds to the intersection of three planes in $\mathbb{R}^3$.

4.1 Lines and planes

Points, lines, and planes are the basic building blocks of geometry. In this section, we'll explore these geometric objects, the equations that describe them, and their visual representations.

Concepts

- $p = (p_x, p_y, p_z)$: a *point* in $\mathbb{R}^3$
- $\vec{v} = (v_x, v_y, v_z)$: a *vector* in $\mathbb{R}^3$
- $\hat{v} = \frac{\vec{v}}{\|\vec{v}\|}$: the *unit vector* in the same direction as the vector $\vec{v}$
- An infinite line ℓ is a one-dimensional space defined in one of several possible ways:
 - ▷ $\ell : \{p_o + t \vec{v}, t \in \mathbb{R}\}$: a *parametric equation* of a line with direction vector $\vec{v}$ passing through the point p_o

$$\triangleright \ell : \left\{ \frac{x-p_{0x}}{v_x} = \frac{y-p_{0y}}{v_y} = \frac{z-p_{0z}}{v_z} \right\}: \text{a symmetric equation}$$

- An infinite plane P is a two-dimensional space defined in one of several possible ways:
 - $\triangleright P : \{Ax + By + Cz = D\}$: a *general equation*
 - $\triangleright P : \{p_0 + s\vec{v} + t\vec{w}, s, t \in \mathbb{R}\}$: a *parametric equation*
 - $\triangleright P : \{\vec{n} \cdot [(x, y, z) - p_0] = 0\}$: a *geometric equation* of the plane that contains point p_0 and has normal vector $\hat{n}$
- $d(a, b)$: the shortest *distance* between geometric objects a and b

Points

We can specify a point in $\mathbb{R}^3$ by its coordinates $p = (p_x, p_y, p_z)$, which is similar to how we specify vectors. In fact, the two notions are equivalent: we can either talk about the destination point p or the vector $\vec{p}$ that takes us from the origin to the point p . This equivalence lets us add and subtract vectors and points. For example, $\vec{d} = q - p$ denotes the displacement vector that takes the point p to the point q .

We can also specify a point as the intersection of two lines. As an example in $\mathbb{R}^2$, let's define $p = (p_x, p_y)$ to be the intersection of the lines $x - y = -1$ and $3x + y = 9$. We must solve the two equations simultaneously to find the coordinates of the point p . We can use the standard techniques for solving equations to find the answer. The intersection point is $p = (2, 3)$. Note that for two lines to intersect at a point, the lines must not be parallel.




Figure 4.1: Two non-parallel lines in $\mathbb{R}^2$ intersect at a point.

Example 1 Find where the lines $x + 2y = 5$ and $3x + 9y = 21$ intersect. To find the point of intersection, we solve these equations

simultaneously and obtain the point (x, y) that is contained in both lines. The answer is the point $p = (1, 2)$.

In three dimensions, a point can also be specified as the intersection of three planes. This is precisely what happens when we solve equations of the form:

$$A_1x + B_1y + C_1z = D_1,$$

$$A_2x + B_2y + C_2z = D_2,$$

$$A_3x + B_3y + C_3z = D_3.$$

To solve this system of equations, we must find the point (x_0, y_0, z_0) that satisfies all three equations, which means this point is contained in all three planes.




Figure 4.2: Three non-parallel planes in $\mathbb{R}^3$ intersect at a point.

Lines

A line ℓ is a one-dimensional space that is infinitely long. There are several equivalent ways to specify a line in space.

The *parametric equation* of a line is obtained as follows. Given a direction vector $\vec{v}$ and some point p_0 on the line, we define the line as the following set:

$$\ell : \{(x, y, z) \in \mathbb{R}^3 \mid (x, y, z) = p_0 + t \vec{v}, t \in \mathbb{R}\}.$$

The line consists of all the points (x, y, z) that can be reached by starting from the point p_0 and adding any multiple of the direction vector $\vec{v}$, as illustrated in Figure 4.3. We say the line is *parametrized* by the variable t .

The *symmetric equation* is an equivalent way to describe a line that does not require an explicit parametrization. The equations that correspond to each coordinate in the parametric equation of a line are

$$x = p_{0x} + t v_x, \quad y = p_{0y} + t v_y, \quad z = p_{0z} + t v_z.$$




Figure 4.3: All points on the line ℓ can be reached by starting from the point p_0 and adding some multiple of the direction vector $\vec{v}$.

When we solve for t in these equations and equate the results, we obtain the *symmetric equation* of a line:

$$\ell : \left\{ \frac{x - p_{0x}}{v_x} = \frac{y - p_{0y}}{v_y} = \frac{z - p_{0z}}{v_z} \right\}.$$

Note the parameter t does not appear. The symmetric equation specifies the line as the relationships between the x , y , and z coordinates that hold for all points on the line.

You're probably most familiar with the symmetric equation of lines in $\mathbb{R}^2$, which does not involve the variable z . For nonvertical lines in $\mathbb{R}^2$ ($v_x \neq 0$), we can think of y as a function of x and write the equation of the line in the equivalent form:

$$\frac{x - p_{0x}}{v_x} = \frac{y - p_{0y}}{v_y} \quad \Rightarrow \quad y(x) = mx + b,$$

where $m = \frac{v_y}{v_x}$ and $b = p_{0y} - \frac{v_y}{v_x} p_{0x}$. The equation $m = \frac{v_y}{v_x}$ makes sense intuitively: the slope of a line m corresponds to how much the line "moves" in the y -direction divided by how much the line "moves" in the x -direction.

Another way to describe a line is to specify two points that are part of the line, as shown in Figure 4.4. The equation of a line that contains the points p and q can be obtained as follows:

$$\ell : \{ \vec{x} = p + t(q - p), t \in \mathbb{R} \},$$

where $(q - p)$ plays the role of the direction vector $\vec{v}$ for this line. Any vector parallel to the line can be used as the direction vector for the line.

Example 2 Find the parametric equation of the line that passes through the points $p = (1, 1, 1)$ and $q = (2, 3, 4)$. What is the symmetric equation of this line?




Figure 4.4: The line ℓ is defined by the points p and q .

Using the direction vector $\vec{v} = q - p = (1, 2, 3)$ and the point p on the line, we can write a parametric equation for the line as $\{(x, y, z) \in \mathbb{R}^3 \mid (x, y, z) = (1, 1, 1) + t(1, 2, 3), t \in \mathbb{R}\}$. Note that a parametric equation using the direction vector $(-1, -2, -3)$ would be equally valid: $\{(1, 1, 1) + t(-1, -2, -3), t \in \mathbb{R}\}$. The symmetric equation of the line is $\frac{x-1}{1} = \frac{y-1}{2} = \frac{z-1}{3}$.

Lines as intersections of planes

In three dimensions, the intersection of two non-parallel planes forms a line. For example, the intersection of the xy -plane $P_{xy} : \{(x, y, z) \in \mathbb{R}^3 \mid z = 0\}$ and the xz -plane $P_{xz} : \{(x, y, z) \in \mathbb{R}^3 \mid y = 0\}$ is the x -axis: $\{(x, y, z) \in \mathbb{R}^3 \mid (0, 0, 0) + (1, 0, 0)t, t \in \mathbb{R}\}$. For this simple case, we can imagine the two planes (use your hands) and visually establish that they intersect along the x -axis. Wouldn't it be nice if there was a general procedure for finding the line of intersection of two planes?




Figure 4.5: The x -axis is the intersection of the xy -plane and the xz -plane.

You already know such a procedure! The line of intersection between the planes $A_1x + B_1y + C_1z = D_1$ and $A_2x + B_2y + C_2z = D_2$

is the solution of the following set of linear equations:

$$A_1x + B_1y + C_1z = D_1,$$

$$A_2x + B_2y + C_2z = D_2.$$

Example 3 Find the intersection of the planes $0x + 0y + 1z = 0$ and $0x + 1y + 1z = 0$. We follow the standard Gauss–Jordan elimination procedure we learned in Chapter 3: construct an augmented matrix, perform row operations (denoted $\sim$), obtain the RREF, and interpret the solution:

$$\left[\begin{array}{ccc|c} 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 \end{array} \right] \sim \left[\begin{array}{ccc|c} 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \end{array} \right] \sim \left[\begin{array}{ccc|c} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{array} \right].$$

The first column is a free variable $t \in \mathbb{R}$. The solution is the line

$$\left\{ \begin{array}{l} x = t \\ y = 0 \\ z = 0 \end{array} \right\}, \quad \forall t \in \mathbb{R} = \left\{ \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} + t \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \forall t \in \mathbb{R} \right\},$$

which corresponds to the x -axis.

Planes

A plane P in $\mathbb{R}^3$ is a two-dimensional space with infinite extent. In general, we specify a plane through a constraint equation that must be satisfied by all points in the plane:

$$P : \{(x, y, z) \in \mathbb{R}^3 \mid Ax + By + Cz = D\}.$$

The plane P is the set of all points $(x, y, z) \in \mathbb{R}^3$ that satisfy the equation $Ax + By + Cz = D$. The equation $Ax + By + Cz = D$ is called the *general equation* of the plane. This definition represents the *algebraic view* of planes, which is useful for calculations.

There is an equally useful geometric view of planes. A plane can be specified by a *normal vector* $\vec{n}$ and some point p_o in the plane. The normal vector $\vec{n}$ is perpendicular to the plane: it sticks out at right angles to the plane like the normal force between surfaces in physics problems. All points in the plane P can be obtained by starting from the point p_o and moving in a direction orthogonal to the normal vector $\vec{n}$. The geometric formula of a plane is

$$P : \vec{n} \cdot [(x, y, z) - p_o] = 0.$$

Recall that the dot product of two vectors is zero, if and only if these vectors are orthogonal. In the above equation, the expression




Figure 4.6: All points on the plane P can be reached by starting from the point p_o and moving perpendicularly to the normal vector $\vec{n}$.

$[(x, y, z) - p_o]$ forms an arbitrary vector with one endpoint at p_o . From all these vectors, we select *only* those that are perpendicular to $\vec{n}$, and thus we obtain all the points in the plane.

The geometric equation $\vec{n} \cdot [(x, y, z) - p_o] = 0$ is equivalent to the general equation $Ax + By + Cz = D$. We can find the parameters A , B , C , and D by calculating the dot product: $A = n_x$, $B = n_y$, $C = n_z$, and $D = \vec{n} \cdot p_o = n_x p_{ox} + n_y p_{oy} + n_z p_{oz}$.

Observe that scaling the general equation of a plane by a constant factor does not change the plane: the equations $Ax + By + Cz = D$ and $\alpha Ax + \alpha By + \alpha Cz = \alpha D$ define the same plane. Similarly, the geometric equations $\vec{n} \cdot [(x, y, z) - p_o] = 0$ and $\alpha \vec{n} \cdot [(x, y, z) - p_o] = 0$ define the same plane. In each case, the direction of the normal vector $\vec{n}$ is important, but not its length.

We can also give a *parametric equation* of a plane P . If we know a point p_o in the plane and two linearly independent vectors $\vec{v}$ and $\vec{w}$ that lie in the plane, then a parametric equation for the plane can be obtained as follows:

$$P : \{(x, y, z) \in \mathbb{R}^3 \mid (x, y, z) = p_o + s\vec{v} + t\vec{w}, s, t \in \mathbb{R}\}.$$

Since a plane is a two-dimensional space, we need two parameters (s and t) to describe the location of arbitrary points in the plane.

Suppose we're given three points p , q , and r that lie in the plane. How can we find the geometric equation for this plane $\vec{n} \cdot [(x, y, z) - p_o] = 0$? We can use the point p as the reference point p_o , but how do we find the normal vector $\vec{n}$ for the plane? The trick is to use the cross product. First we build two vectors that are parallel to the plane, $\vec{v} = q - p$ and $\vec{w} = r - p$. Then compute their cross product to find a vector that is perpendicular to both of them, and hence normal to the plane:

$$\vec{n} = \vec{v} \times \vec{w} = (q - p) \times (r - p).$$

We can use the vector $\vec{n}$ to write the geometric equation of the plane $\vec{n} \cdot [(x, y, z) - p] = 0$. Recall that the cross product of two vectors

is perpendicular to both vectors. This property of the cross product makes it the perfect tool for finding normal vectors. This procedure is illustrated in Figure 4.7.




Figure 4.7: We can obtain the normal vector $\vec{n}$ for the plane P by starting from the three points p , q , and r which lie in the plane.

Example 4 Consider the plane that contains the points $p = (1, 0, 0)$, $q = (0, 1, 0)$, and $r = (0, 0, 1)$. Find a geometric equation, a general equation, and a parametric equation for this plane.

We need a normal vector for the geometric equation. We can obtain a normal vector from the cross product of the vectors $\vec{v} = q - p = (-1, 1, 0)$ and $\vec{w} = r - p = (-1, 0, 1)$, which both lie in the plane. We obtain the normal $\vec{n} = \vec{v} \times \vec{w} = (1, 1, 1)$ and write the geometric equation as $(1, 1, 1) \cdot [(x, y, z) - (1, 0, 0)] = 0$, using p as the point p_0 in the plane. To find the general equation for the plane, we compute the dot product in the geometric equation and obtain $1x + 1y + 1z - 1 = 0$, which is the same as $x + y + z = 1$. The vectors $\vec{v}$ and $\vec{w}$ obtained above can also be used in the parametric equation of the plane: $\{(1, 0, 0) + s(-1, 1, 0) + t(-1, 0, 1), s, t \in \mathbb{R}\}$.

Distance formulas

We'll now discuss three formulas for calculating distances: the distance between two points, the closest distance between a line and the origin, and the closest distance between a plane and the origin.

Distance between points

The distance between points p and q is equal to the length of the vector that goes from p to q :

$$d(p, q) = \|q - p\| = \sqrt{(q_x - p_x)^2 + (q_y - p_y)^2 + (q_z - p_z)^2}.$$

Distance between a line and the origin

The closest distance between the line with equation $\ell : \{p_o + t \vec{v}, t \in \mathbb{R}\}$ and the origin $O = (0, 0, 0)$ is given by the formula

$$d(\ell, O) = \left\| p_o - \frac{p_o \cdot \vec{v}}{\|\vec{v}\|^2} \vec{v} \right\|.$$




Figure 4.8: The closest distance between the line ℓ and the origin.

Example 5 The closest distance between the line $\ell : \{(4, 5, 6) + t(1, 0, 1), t \in \mathbb{R}\}$ and the origin $O = (0, 0, 0)$ is calculated as follows:

$$\begin{aligned} d(\ell, O) &= \left\| (4, 5, 6) - \frac{(4, 5, 6) \cdot (1, 0, 1)}{1^2 + 0^2 + 1^2} (1, 0, 1) \right\| \\ &= \left\| (4, 5, 6) - \frac{4 + 0 + 6}{2} (1, 0, 1) \right\| \\ &= \|(-1, 5, 1)\| = 3\sqrt{3}. \end{aligned}$$

Distance between a plane and the origin

The closest distance between the plane with geometric equation $P : \vec{n} \cdot [(x, y, z) - p_o] = 0$ and the origin O is given by

$$d(P, O) = \frac{|\vec{n} \cdot p_o|}{\|\vec{n}\|}.$$

For example, the distance between the plane $P : (-3, 0, -4) \cdot [(x, y, z) - (1, 2, 3)] = 0$ and the origin is computed as

$$d(P, O) = \frac{|(-3, 0, -4) \cdot (1, 2, 3)|}{\|(-3, 0, -4)\|} = \frac{|-3 - 12|}{5} = \frac{15}{5} = 3.$$




Figure 4.9: The closest distance between the plane P and the origin.

Discussion

The distance formulas $d(\ell, O)$ and $d(P, O)$ are complicated expressions that involve dot products and vector lengths. To understand the logic behind these distance formulas, we need to learn a bit about *projective geometry*. The techniques of projective geometry allow us to measure distances between arbitrary points, lines, and planes. No new math operations are required. Instead, we'll learn how to use a combination of vector subtraction, vector length, and the dot product to compute distances. Each distance function $d(\cdot, \cdot)$ corresponds to an abstract procedure with one or two steps which can be described using a vector diagram. Projections play a key role in projective geometry, so we'll learn about them in detail in the next section.

Exercises

E4.1 Define the line $\ell : \{(0, 3) + t(1, -1), t \in \mathbb{R}\}$. Find the closest distance between the line ℓ and the origin.

E4.2 Find the closest distance between the origin and the line defined by $\ell : \{(0, 0, 3) + t(1, 1, 0), t \in \mathbb{R}\}$.

E4.3 Find the distance between the plane P with geometric equation $(1, 1, 1) \cdot [(x, y, z) - (4, 5, 6)] = 0$ and the origin.

E4.4 Find the general equation of the line that passes through the points $(0, 5)$ and $(6, -7)$ in $\mathbb{R}^2$.

E4.5 Given the point $r = (1, 3, 0)$, the line $\ell : \{(0, 0, 2) + t(1, -1, 0), t \in \mathbb{R}\}$, the plane $P : x + y + z = 1$, and the origin $O = (0, 0, 0)$, compute the following closest distances:

a) $d(r, O)$ b) $d(\ell, O)$ c) $d(P, O)$ d) $d(r, \ell)$ e) $d(r, P)$ f) $d(\ell, P)$

Hint: Draw a diagram and find the closest distances visually. Label the point $p_\ell = (0, 0, 2)$, which is an arbitrary point in the line, and the point $p_P = (1, 0, 0)$, which is an arbitrary point in the plane. Don't

look for a one-size-fits-all formula for the different cases; derive the appropriate formula for each case starting from the basic projection operations Π_ℓ and Π_P .

4.2 Projections

In this section we'll learn to compute projections of vectors onto lines and planes. Given an arbitrary vector, we'll find how much of this vector points in a given direction (projection onto a line). We'll also find the part of the vector that lies in some plane (projection onto a plane). The dot product, $\vec{u} \cdot \vec{v} = u_1v_1 + u_2v_2 + u_3v_3$, will play a central role in these calculations.

Each projection formula corresponds to a vector diagram. Vector diagrams, also known as “picture proofs,” are used to describe the precise sequence of operations for computing a projection. Focussing on vector diagrams makes it much easier to understand projection and distance formulas. Indeed, the pictures in this section are a heck of a lot more important than the formulas. Be sure you understand each vector diagram, and don't worry about memorizing the corresponding formula. You can easily reproduce the formula by starting from the vector diagram.

Concepts

- $S \subseteq \mathbb{R}^n$: S is a *vector subspace* of $\mathbb{R}^n$. In this section, we assume $S \subseteq \mathbb{R}^3$. The subspaces of $\mathbb{R}^3$ are lines ℓ and planes P that pass through the origin.
- $S^\perp$: the orthogonal complement to S , $S^\perp \stackrel{\text{def}}{=} \{\vec{w} \in \mathbb{R}^n \mid \vec{w} \cdot S = 0\}$. The symbol $^\perp$ stands for *perpendicular to*.
- Π_S : the *projection* onto the subspace S .
- $\Pi_{S^\perp}$: the projection onto the subspace $S^\perp$.

Definitions

Let S be a *vector subspace* of $\mathbb{R}^n$, denoted $S \subseteq \mathbb{R}^n$. In this section, we'll focus on the subspaces of the vector space $\mathbb{R}^3$ because they are easy to visualize and understand intuitively. The vector subspaces of $\mathbb{R}^3$ are lines and planes that pass through the origin. We defer the general discussion of subspaces in n dimensions until Section 4.4.

The projection operation onto the subspace S is a linear transformation that takes as inputs vectors in $\mathbb{R}^3$, and produces outputs in the subspace S :

$$\Pi_S : \mathbb{R}^3 \rightarrow S.$$

The transformation Π_S , pronounced “projection onto S ,” cuts off all parts of the input that do not lie within the subspace S . We can understand Π_S by analyzing its action for different inputs:

- If $\vec{v} \in S$, then $\Pi_S(\vec{v}) = \vec{v}$.
- If $\vec{w} \in S^\perp$, then $\Pi_S(\vec{w}) = \vec{0}$.
- Linearity and the above two conditions imply that, for any vector $\vec{u} = \alpha\vec{v} + \beta\vec{w}$ with $\vec{v} \in S$ and $\vec{w} \in S^\perp$, we have

$$\Pi_S(\vec{u}) = \Pi_S(\alpha\vec{v} + \beta\vec{w}) = \alpha\vec{v}.$$

The *orthogonal subspace* to S is the set of vectors that are perpendicular to all vectors in S :

$$S^\perp \stackrel{\text{def}}{=} \{ \vec{w} \in \mathbb{R}^3 \mid \vec{w} \cdot \vec{s} = 0, \forall \vec{s} \in S \}.$$

The transformation Π_S “projects” to the space S in the sense that, no matter which vector $\vec{u}$ you start from, applying the projection Π_S results in a vector that is in S :

$$\forall \vec{u} \in \mathbb{R}^3, \quad \Pi_S(\vec{u}) \in S.$$

All parts of $\vec{u}$ in the *perp*-space $S^\perp$ were killed by Π_S . Meet Π_S —the *S-perp* killer.

We can split the set of all vectors $\mathbb{R}^3$ into two disjoint sets: vectors entirely contained in S and vectors perpendicular to S . We say $\mathbb{R}^3$ decomposes into the *direct sum* of the subspaces S and $S^\perp$:

$$\mathbb{R}^3 = S \oplus S^\perp.$$

Any vector $\vec{u} \in \mathbb{R}^3$ can be split into an S -part $\vec{v} = \Pi_S(\vec{u})$ and a $S^\perp$ -part $\vec{w} = \Pi_{S^\perp}(\vec{u})$, such that

$$\vec{u} = \vec{v} + \vec{w}.$$

A defining property of projections is that they are *idempotent operations*, meaning it doesn’t matter if you project a vector once, twice, or a million times; the result will always be the same:

$$\Pi_S(\vec{u}) = \Pi_S(\Pi_S(\vec{u})) = \Pi_S(\Pi_S(\Pi_S(\vec{u}))) = \dots$$

Once you project a vector onto the subspace S , any further projections to S have no effect.

In the remainder of this section, we’ll derive formulas for projections onto lines and planes that pass through the origin.

Projection onto a line

Consider the line ℓ passing through the origin with direction vector $\vec{v}$:

$$\ell : \{(x, y, z) \in \mathbb{R}^3 \mid (x, y, z) = \vec{0} + t\vec{v}, t \in \mathbb{R}\}.$$

The projection onto ℓ for an arbitrary vector $\vec{u} \in \mathbb{R}^3$ is given by the formula

$$\Pi_{\ell}(\vec{u}) = \frac{\vec{u} \cdot \vec{v}}{\|\vec{v}\|^2} \vec{v}.$$




Figure 4.10: The vector $\vec{u}$ can be decomposed into the sum of two projections defined with respect to the line ℓ . The projection $\Pi_{\ell}(\vec{u})$ is parallel to the line ℓ , while the projection $\Pi_{\ell^{\perp}}(\vec{u})$ is perpendicular to the line ℓ .

The orthogonal complement to the line ℓ consists of all vectors perpendicular to the direction vector $\vec{v}$. Mathematically speaking,

$$\ell^{\perp} : \{(x, y, z) \in \mathbb{R}^3 \mid (x, y, z) \cdot \vec{v} = 0\}.$$

Recognize that the equation $(x, y, z) \cdot \vec{v} = 0$ defines a *plane*. The orthogonal complement for a line ℓ with direction vector $\vec{v}$ is a plane with normal vector $\vec{v}$. Makes sense, yes?

We can easily find the projection operation onto $\ell^{\perp}$ as well. Any vector can be written as the sum of an ℓ part and a $\ell^{\perp}$ part: $\vec{u} = \vec{v} + \vec{w}$, where $\vec{v} = \Pi_{\ell}(\vec{u}) \in \ell$ and $\vec{w} = \Pi_{\ell^{\perp}}(\vec{u}) \in \ell^{\perp}$. To obtain $\Pi_{\ell^{\perp}}(\vec{u})$, subtract the Π_{ℓ} part from the original vector $\vec{u}$:

$$\Pi_{\ell^{\perp}}(\vec{u}) = \vec{w} = \vec{u} - \vec{v} = \vec{u} - \Pi_{\ell}(\vec{u}) = \vec{u} - \frac{\vec{u} \cdot \vec{v}}{\|\vec{v}\|^2} \vec{v}.$$

We can think of $\Pi_{\ell^{\perp}}(\vec{u}) = \vec{w}$ as the part of $\vec{u}$ that remains after we've removed the ℓ -part.

Example 1 Consider the line ℓ defined by the parametric equation $\{(x, y, z) \in \mathbb{R}^3 \mid (x, y, z) = t(1, 2, 3), t \in \mathbb{R}\}$. Find the projection of the vector $\vec{u} = (4, 5, 6)$ onto ℓ . Find the projection of $\vec{u}$ onto $\ell^{\perp}$ and verify that $\Pi_{\ell}(\vec{u}) + \Pi_{\ell^{\perp}}(\vec{u}) = \vec{u}$.

The direction vector of the line ℓ is $\vec{v} = (1, 2, 3)$, so $\Pi_{\ell}(\vec{u}) = \frac{\vec{u} \cdot \vec{v}}{\|\vec{v}\|^2} \vec{v} = \frac{16}{7} \vec{v} = (\frac{16}{7}, \frac{32}{7}, \frac{48}{7})$. Next, using the formula $\Pi_{\ell^{\perp}}(\vec{u}) = \vec{u} - \frac{\vec{u} \cdot \vec{v}}{\|\vec{v}\|^2} \vec{v}$, we find $\Pi_{\ell^{\perp}}(\vec{u}) = (\frac{12}{7}, \frac{3}{7}, \frac{-6}{7})$. Observe that $(\frac{12}{7}, \frac{3}{7}, \frac{-6}{7}) \cdot \vec{v} =$

0, which shows the vector $\Pi_{\ell^\perp}(\vec{u})$ is indeed perpendicular to ℓ . Adding the results of the two projections, we obtain the whole $\vec{u}$: $(\frac{16}{7}, \frac{32}{7}, \frac{48}{7}) + (\frac{12}{7}, \frac{3}{7}, \frac{-6}{7}) = (\frac{28}{7}, \frac{35}{7}, \frac{42}{7}) = (4, 5, 6) = \vec{u}$.

Projection onto a plane

Now consider the two-dimensional plane P passing through the origin with normal vector $\vec{n}$:

$$P : \{(x, y, z) \in \mathbb{R}^3 \mid \vec{n} \cdot (x, y, z) = 0\}.$$

The perpendicular space $P^\perp$ is a line with direction vector $\vec{n}$:

$$P^\perp : \{(x, y, z) \in \mathbb{R}^3 \mid (x, y, z) = \vec{0} + t\vec{n}, t \in \mathbb{R}\}.$$

Again, the vector space $\mathbb{R}^3$ decomposes into the direct sum of P and $P^\perp$: $\mathbb{R}^3 = P \oplus P^\perp$.

We want to find Π_P , but it will actually be easier to find $\Pi_{P^\perp}$ first and then compute $\Pi_P(\vec{u})$ using $\Pi_P(\vec{u}) = \vec{v} = \vec{u} - \vec{w}$, where $\vec{w} = \Pi_{P^\perp}(\vec{u})$. See Figure 4.11 for an illustration.




Figure 4.11: Any vector $\vec{u}$ can be written as the sum of two projections defined with respect to the plane P . The projection $\Pi_P(\vec{u})$ is parallel to the plane P , while the projection $\Pi_{P^\perp}(\vec{u})$ is perpendicular to the plane P .

Since $P^\perp$ is a line, we know the formula for projecting onto it is

$$\Pi_{P^\perp}(\vec{u}) = \frac{\vec{u} \cdot \vec{n}}{\|\vec{n}\|^2} \vec{n}.$$

We can now obtain the formula for Π_P :

$$\Pi_P(\vec{u}) = \vec{v} = \vec{u} - \vec{w} = \vec{u} - \Pi_{P^\perp}(\vec{u}) = \vec{u} - \frac{\vec{u} \cdot \vec{n}}{\|\vec{n}\|^2} \vec{n}.$$

Example 2 Consider the plane P defined by the geometric equation $(1, 1, 1) \cdot [(x, y, z) - (0, 0, 0)] = 0$. Find the projection of the vector $\vec{u} = (4, 5, 6)$ onto P and onto $P^\perp$. Verify that $\Pi_P(\vec{u}) + \Pi_{P^\perp}(\vec{u}) = \vec{u}$.

Using the formula $\Pi_P(\vec{u}) = \vec{u} - \frac{\vec{u} \cdot \vec{n}}{\|\vec{n}\|^2} \vec{n}$, we find $\Pi_P(\vec{u}) = (-1, 0, 1)$.

We also find $\Pi_{P^\perp}(\vec{u}) = \frac{\vec{u} \cdot \vec{n}}{\|\vec{n}\|^2} \vec{n} = (5, 5, 5)$, which is a vector in the same direction as $\vec{n}$. Observe that the vector $\vec{u}$ can be reconstructed by adding the two projections: $\Pi_P(\vec{u}) + \Pi_{P^\perp}(\vec{u}) = (-1, 0, 1) + (5, 5, 5) = (4, 5, 6) = \vec{u}$.

Distances formulas revisited

Suppose you want to find the distance between the line $\ell : \{(x, y, z) \in \mathbb{R}^3 \mid (x, y, z) = p_o + t \vec{v}, t \in \mathbb{R}\}$ and the origin $O = (0, 0, 0)$. This problem is equivalent to finding the distance between the line $\ell' : \{(x, y, z) \in \mathbb{R}^3 \mid (x, y, z) = \vec{0} + t \vec{v}, t \in \mathbb{R}\}$ and the point p_o , the answer to which is the length of the projection $\Pi_{\ell^\perp}(p_o)$:

$$d(\ell', p_o) = \|\Pi_{\ell^\perp}(p_o)\| = \left\| p_o - \frac{p_o \cdot \vec{v}}{\|\vec{v}\|^2} \vec{v} \right\|.$$

The distance between a plane $P : \vec{n} \cdot [(x, y, z) - p_o] = 0$ and the origin O is the same as the distance between the plane $P' : \vec{n} \cdot (x, y, z) = 0$ and the point p_o . We can obtain this distance by finding the length of the projection of p_o onto $P'^\perp$ using the formula

$$d(P', p_o) = \frac{|\vec{n} \cdot p_o|}{\|\vec{n}\|}.$$

On your own, try drawing the pictures for the above two scenarios. Make sure the formulas make sense to you.

Projections matrices

Since projections are *linear transformations*, they can be expressed as matrix-vector products:

$$\vec{v} = \Pi(\vec{u}) \quad \Leftrightarrow \quad \vec{v} = M_\Pi \vec{u}.$$

Multiplying the vector $\vec{u}$ by the matrix M_Π is the same as applying the projection Π .

We'll learn more about projection matrices later. For now, I'll show you a simple example of a projection matrix in $\mathbb{R}^3$. Let Π be the projection onto the xy -plane. This projection operation corresponds

to the following matrix-vector product:

$$\Pi(\vec{u}) = M_{\Pi}\vec{u} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} u_x \\ u_y \\ u_z \end{bmatrix} = \begin{bmatrix} u_x \\ u_y \\ 0 \end{bmatrix}.$$

Note how multiplying a vector by M_{Π} results in selecting only the x - and y -components of the vector while killing the z -component, which is precisely what the projection onto the xy -plane is supposed to do.

Discussion

In the next section we'll talk about a particular set of projections known as *coordinate projections*. We use coordinate projections to find the components of vectors $\vec{v}$ with respect to a coordinate system:

$$\begin{aligned} v_x \hat{i} &= \Pi_x(\vec{v}) = \frac{\vec{v} \cdot \hat{i}}{\|\hat{i}\|^2} \hat{i} = (\vec{v} \cdot \hat{i}) \hat{i}, \\ v_y \hat{j} &= \Pi_y(\vec{v}) = \frac{\vec{v} \cdot \hat{j}}{\|\hat{j}\|^2} \hat{j} = (\vec{v} \cdot \hat{j}) \hat{j}, \\ v_z \hat{k} &= \Pi_z(\vec{v}) = \frac{\vec{v} \cdot \hat{k}}{\|\hat{k}\|^2} \hat{k} = (\vec{v} \cdot \hat{k}) \hat{k}. \end{aligned}$$

The coordinate projection Π_x projects onto the x -axis. Similarly, Π_y and Π_z project onto the y - and z -axes.

Exercises

E4.6 Consider the points $r_1 = (1, 3, 0)$, $r_2 = (3, 1, 0)$, the vector $\vec{u} = (0, 0, 5)$, and the plane $P : x + y + z = 0$ that passes through the origin and has normal vector $\vec{n} = (1, 1, 1)$. Compute the projections.

$$\text{a) } \Pi_{\vec{n}}(r_1) \quad \text{b) } \Pi_P(r_1) \quad \text{c) } \Pi_P(r_2) \quad \text{d) } \Pi_{\vec{n}}(\vec{u}) \quad \text{e) } \Pi_P(\vec{u})$$

Hint: Your answer should be a vector in each case.

E4.7 Consider the points $p = (10, 10, 10)$ and the two parallel lines $\ell_1 : \{(3, 0, 5) + t(1, -2, 0), t \in \mathbb{R}\}$ and $\ell_2 : \{(6, 0, 0) + t(1, -2, 0), t \in \mathbb{R}\}$. Compute the following quantities.

$$\text{a) } \Pi_{\ell_1}(p) \quad \text{b) } d(p, \ell_1) \quad \text{c) } \Pi_{\ell_2}(p) \quad \text{d) } d(p, \ell_2) \quad \text{e) } d(\ell_1, \ell_2)$$

Is p closer to ℓ_1 or to ℓ_2 ?

E4.8 Find the projection of the vector $\vec{v} = (4, 5, 6)$ onto the plane that passes through the origin and contains the vectors $(2, -4, 2)$ and $(6, 1, -4)$.

4.3 Coordinate projections

In science, it's common to express vectors as components: (v_x, v_y, v_z) . Thinking of vectors as lists of components is fine for computational purposes, since vector operations require manipulating the components of vectors. However, focussing on a vector's components overlooks an important concept—the *basis* with respect to which the vector's components are expressed.

It's not uncommon for students to have misconceptions about linear algebra due to an incomplete understanding of the fundamental distinction between vectors and their components. Since I want you to have a *thorough* understanding of linear algebra, we'll review—in full detail—the notion of a basis and how to compute vector components with respect to different bases.

Before we begin, let's quickly review what we know about vectors and vector components from previous sections of the book. In Section 1.13, we described vectors in terms of their x - and y -components. Given a standard xy -coordinate system, we can decompose a vector $\vec{F}$ in terms of its components:

$$F_x = \|\vec{F}\| \cos \theta, \quad F_y = \|\vec{F}\| \sin \theta,$$

where θ is the angle the vector $\vec{F}$ makes with the x -axis. We can express the vector as *coordinates* with respect to the basis $\{\hat{i}, \hat{j}\}$ as $\vec{F} = F_x \hat{i} + F_y \hat{j} = (F_x, F_y)_{\hat{i}\hat{j}}$. The number F_x corresponds to the length of the *projection* of the vector $\vec{F}$ onto the x -axis.

In the last section, we discussed the projection operation and learned how to compute projections using the dot product or as a matrix-vector product:

$$F_x \hat{i} = \frac{\vec{F} \cdot \hat{i}}{\|\hat{i}\|^2} \hat{i} = \Pi_x(\vec{F}) = \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}}_{M_{\Pi_x}} \underbrace{\begin{bmatrix} F_x \\ F_y \end{bmatrix}}_{\vec{F}},$$

where $\Pi_x : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the projection onto the x -axis (a linear transformation) and M_{Π_x} is the matrix representation of Π_x with respect to the basis $\{\hat{i}, \hat{j}\}$.

In this section, we'll extend what we know about basic vectors coordinates with respect to the basis $\{\hat{i}, \hat{j}\}$, and formally define vector coordinates with respect to any basis.

Concepts

We can define three different types of bases for an n -dimensional vector space V :

- A generic basis $B_f = \{\vec{f}_1, \vec{f}_2, \dots, \vec{f}_n\}$ consists of any set of linearly independent vectors in V .
- An orthogonal basis $B_e = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ consists of n mutually orthogonal vectors in V obeying $\vec{e}_i \cdot \vec{e}_j = 0, \forall i \neq j$.
- An orthonormal basis $B_{\hat{e}} = \{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$ is an orthogonal basis of unit vectors: $\hat{e}_i \cdot \hat{e}_j = 0$ and $\|\hat{e}_i\| = 1, \forall i \in \{1, 2, \dots, n\}$.

A vector $\vec{v}$ is expressed as coordinates v_i with respect to any basis B :

$$\vec{v} = v_1 \vec{e}_1 + v_2 \vec{e}_2 + \dots + v_n \vec{e}_n = (v_1, v_2, \dots, v_n)_B.$$

We can use two different bases, B and B' , to express the same vector:

- $\vec{v}$: a vector
- $[\vec{v}]_B = (v_1, v_2, \dots, v_n)_B$: the vector $\vec{v}$ expressed in the basis B
- $[\vec{v}]_{B'} = (v'_1, v'_2, \dots, v'_n)_{B'}$: the same vector $\vec{v}$ expressed in a different basis B'
- ${}_{B'}[\mathbb{1}]_B$: the change-of-basis matrix that converts from B coordinates to B' coordinates: $[\vec{v}]_{B'} = {}_{B'}[\mathbb{1}]_B [\vec{v}]_B$

Components with respect to a basis

A vector's *components* describe how much of the vector lies in a given direction. For example, a vector $\vec{v} \in \mathbb{R}^3$ expressed as components *with respect to* the standard orthonormal basis $\{\hat{i}, \hat{j}, \hat{k}\}$ is denoted $\vec{v} = (v_x, v_y, v_z)_{\hat{i}\hat{j}\hat{k}}$. The *components* of a vector are also called *coordinates* (in the context of a coordinate system) and *coefficients* (in the context of a linear combination). Don't be confused by this multitude of terms because it's all the same idea: components, coordinates, and coefficients are all ways to describe vectors with respect to bases.

When the basis consists of a set of orthonormal vectors like the vectors $\hat{i}$, $\hat{j}$, and $\hat{k}$, we can compute vector components using the dot product:

$$v_x = \vec{v} \cdot \hat{i}, \quad v_y = \vec{v} \cdot \hat{j}, \quad v_z = \vec{v} \cdot \hat{k}.$$

In this section, we'll discuss how to find coordinates with respect to three different types of bases: orthonormal bases, orthogonal bases, and generic bases. First, let's precisely define what a *basis* is.

Definition of a basis

A basis $B = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ for the vector space V has the following two properties:

- **Spanning property:** Any vector $\vec{v} \in V$ can be expressed as a linear combination of the basis vectors:

$$\vec{v} = v_1\vec{e}_1 + v_2\vec{e}_2 + \cdots + v_n\vec{e}_n.$$

This property guarantees that the vectors in the basis B are *sufficient* to represent any vector in V .

- **Linear independence property:** The vectors that form the basis $B = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ are linearly independent. The linear independence of the vectors in the basis guarantees that none of the vectors $\vec{e}_i$ are redundant.

If a set of vectors $B = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ satisfies both properties, we say B is a basis for V . In other words, B can serve as a coordinate system for V .

Coordinates with respect to an orthonormal basis

An orthonormal basis $B_{\hat{e}} = \{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$ consists of a set of mutually orthogonal, unit vectors:

$$\hat{e}_i \cdot \hat{e}_j = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

The vectors are mutually orthogonal since $\hat{e}_i \cdot \hat{e}_j = 0$ for all $i \neq j$, and the vectors have length one since $\hat{e}_i \cdot \hat{e}_i = 1$ implies $\|\hat{e}_i\| = 1$.

To compute the coordinates of the vector $\vec{a}$ with respect to an orthonormal basis $B_{\hat{e}}$, we use the standard “prescription” similar to the one we used for the $\{\hat{i}, \hat{j}, \hat{k}\}$ basis:

$$(a_1, a_2, \dots, a_n)_{B_{\hat{e}}} \Leftrightarrow \underbrace{(\vec{a} \cdot \hat{e}_1)}_{a_1} \hat{e}_1 + \underbrace{(\vec{a} \cdot \hat{e}_2)}_{a_2} \hat{e}_2 + \cdots + \underbrace{(\vec{a} \cdot \hat{e}_n)}_{a_n} \hat{e}_n.$$

Coordinates with respect to an orthogonal basis

Consider a basis $B_e = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ that is orthogonal, but not orthonormal. We can compute the coordinates of any vector $\vec{b}$ with respect to the basis $\{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ as follows:

$$(b_1, b_2, \dots, b_n)_{B_e} \Leftrightarrow \left(\frac{\vec{b} \cdot \vec{e}_1}{\|\vec{e}_1\|^2} \right) \vec{e}_1 + \left(\frac{\vec{b} \cdot \vec{e}_2}{\|\vec{e}_2\|^2} \right) \vec{e}_2 + \cdots + \left(\frac{\vec{b} \cdot \vec{e}_n}{\|\vec{e}_n\|^2} \right) \vec{e}_n.$$

To find the coordinates of the vector $\vec{b}$ with respect to B_e , we use the general projection formula,

$$b_1 = \frac{\vec{b} \cdot \vec{e}_1}{\|\vec{e}_1\|^2}, \quad b_2 = \frac{\vec{b} \cdot \vec{e}_2}{\|\vec{e}_2\|^2}, \quad \dots, \quad b_n = \frac{\vec{b} \cdot \vec{e}_n}{\|\vec{e}_n\|^2}.$$

Observe that each component of the vector can be computed independently of the other components: to compute b_1 , all we need to know is $\vec{b}$ and $\vec{e}_1$; we don't need to know $\vec{e}_2, \vec{e}_3, \dots, \vec{e}_n$, because we know the other basis vectors are orthogonal to $\vec{e}_1$. The computation of the coordinates correspond to n independent *orthogonal projections*. The coordinate b_i tells us how much of the basis vector $\vec{e}_i$ we need in the linear combination to construct the vector $\vec{b}$.

Coordinates with respect to a generic basis

Suppose we're given a generic basis $\{\vec{f}_1, \vec{f}_2, \dots, \vec{f}_n\}$ for a vector space. To find the coordinates $(c_1, c_2, \dots, c_n)$ of a vector $\vec{c}$ with respect to this basis, we need to solve the equation

$$c_1\vec{f}_1 + c_2\vec{f}_2 + \dots + c_n\vec{f}_n = \vec{c}$$

for the unknowns $c_1, c_2, \dots, c_n$. Because the vectors $\{\vec{f}_i\}$ are not orthogonal, the calculation of the coefficients $c_1, c_2, \dots, c_n$ must be done simultaneously.

Example Express the vector $\vec{v} = (5, 6)_{ij}$ in terms of the basis $B_f = \{\vec{f}_1, \vec{f}_2\}$ where $\vec{f}_1 = (1, 1)_{ij}$ and $\vec{f}_2 = (3, 0)_{ij}$.

We are looking for the coefficients v_1 and v_2 such that

$$(v_1, v_2)_{B_f} = v_1\vec{f}_1 + v_2\vec{f}_2 = \vec{v} = (5, 6)_{ij}.$$

To find the coefficients we need to solve the following system of equations *simultaneously*:

$$1v_1 + 3v_2 = 5$$

$$1v_1 + 0 = 6.$$

From the second equation we find $v_1 = 6$. We substitute this answer into the first equation and find $v_2 = \frac{-1}{3}$. Thus, the vector $\vec{v}$ written with respect to the basis $\{\vec{f}_1, \vec{f}_2\}$ is $\vec{v} = 6\vec{f}_1 - \frac{1}{3}\vec{f}_2 = \left(6, \frac{-1}{3}\right)_{B_f}$.

The general case of computing a vector's coordinates with respect to a generic basis $\{\vec{f}_1, \vec{f}_2, \dots, \vec{f}_n\}$ requires solving a system of n equations in n unknowns. You know how to do this, but it will take some work. The take-away message is that computing vector coordinates with respect to generic bases is more difficult than computing vector coordinates with respect to orthogonal and orthonormal bases.

Change of basis

We often identify a vector $\vec{v}$ with its components (v_x, v_y, v_z) . It's important to always keep in mind the basis with respect to which the components are taken, and if necessary specify the basis as a subscript $\vec{v} = (v_x, v_y, v_z)_{ijk}$. When performing vector arithmetic operations like $\vec{u} + \vec{v}$, we don't really care what basis the vectors are expressed in so long as the *same* basis is used for both $\vec{u}$ and $\vec{v}$.

We sometimes need to use multiple bases, however. Consider the basis $B = \{\hat{e}_1, \hat{e}_2, \hat{e}_3\}$ and another basis $B' = \{\hat{e}'_1, \hat{e}'_2, \hat{e}'_3\}$. Suppose we're given the coordinates v_1, v_2, v_3 of some vector $\vec{v}$ in terms of the basis B :

$$\vec{v} = (v_1, v_2, v_3)_B = v_1\hat{e}_1 + v_2\hat{e}_2 + v_3\hat{e}_3.$$

How can we find the coordinates of $\vec{v}$ in terms of the basis B' ?

This is called a *change-of-basis* transformation, and it's performed as a matrix multiplication with a *change-of-basis matrix*:

$$\begin{bmatrix} v'_1 \\ v'_2 \\ v'_3 \end{bmatrix}_{B'} = \underbrace{\begin{bmatrix} \hat{e}'_1 \cdot \hat{e}_1 & \hat{e}'_1 \cdot \hat{e}_2 & \hat{e}'_1 \cdot \hat{e}_3 \\ \hat{e}'_2 \cdot \hat{e}_1 & \hat{e}'_2 \cdot \hat{e}_2 & \hat{e}'_2 \cdot \hat{e}_3 \\ \hat{e}'_3 \cdot \hat{e}_1 & \hat{e}'_3 \cdot \hat{e}_2 & \hat{e}'_3 \cdot \hat{e}_3 \end{bmatrix}}_{B'[\mathbb{1}]_B} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix}_B.$$

The columns of the change-of-basis matrix describe the vectors of the basis $\{\hat{e}_i\}$ in terms of the basis $\{\hat{e}'_i\}$.

Note that multiplying a vector by the matrix $B'[\mathbb{1}]_B$ doesn't actually *do* anything since it doesn't change the vector. The change-of-basis operation acts like the identity transformation, which is why we use the notation $B'[\mathbb{1}]_B$ to describe it. The vector $\vec{v}$ stays the same—it is simply expressed in terms of a different basis:

$$(v'_1, v'_2, v'_3)_{B'} = v'_1\hat{e}'_1 + v'_2\hat{e}'_2 + v'_3\hat{e}'_3 = \vec{v} = v_1\hat{e}_1 + v_2\hat{e}_2 + v_3\hat{e}_3 = (v_1, v_2, v_3)_B.$$

We say the vector $\vec{v}$ has two *representations*. The vector $\vec{v}$ corresponds to the coordinates triple (v_1, v_2, v_3) with respect to the basis B , and to the triple (v'_1, v'_2, v'_3) with respect to the basis B' .

The matrix $B'[\mathbb{1}]_B$ contains the information about how each vector of the old basis (B) is expressed in terms of the new basis (B'). For example, the first column of the change-of-basis matrix describes how the vector $\hat{e}_1$ is expressed in terms of the basis B' :

$$\hat{e}_1 = (\hat{e}'_1 \cdot \hat{e}_1)\hat{e}'_1 + (\hat{e}'_2 \cdot \hat{e}_1)\hat{e}'_2 + (\hat{e}'_3 \cdot \hat{e}_1)\hat{e}'_3 = (\hat{e}'_1 \cdot \hat{e}_1, \hat{e}'_2 \cdot \hat{e}_1, \hat{e}'_3 \cdot \hat{e}_1)_{B'}.$$

Note this is the generic formula for expressing vectors in the basis B' .

To find the entries of the change-of-basis matrix $B'[\mathbb{1}]_B$ between orthonormal bases B and B' , it's sufficient to compute all the dot

products $\vec{e}_i' \cdot \vec{e}_j$. To compute the entries of a change-of-basis matrix between bases B and B' (which are orthogonal but not necessarily orthonormal), we use $({}_B[1]_{B'})_{ij} = \frac{\vec{e}_i' \cdot \vec{e}_j}{\|\vec{e}_i'\| \|\vec{e}_j\|}$. Computing the change-of-basis matrix between nonorthogonal bases is more complicated.

Change of basis to the standard basis

We'll now discuss an important special case of the change-of-basis operation, where we change from a generic basis $B_f = \{\vec{f}_1, \vec{f}_2, \vec{f}_3\}$ to the standard basis $B_s = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$. In three dimensions, the vectors of the standard basis are denoted $\hat{i} \stackrel{\text{def}}{=} (1, 0, 0)$, $\hat{j} \stackrel{\text{def}}{=} (0, 1, 0)$, and $\hat{k} \stackrel{\text{def}}{=} (0, 0, 1)$. Assume the vectors of the generic basis are given to us as coordinate vectors with respect to the standard basis $\vec{f}_1 = (f_{1x}, f_{1y}, f_{1z})_{B_s}$, $\vec{f}_2 = (f_{2x}, f_{2y}, f_{2z})_{B_s}$, and $\vec{f}_3 = (f_{3x}, f_{3y}, f_{3z})_{B_s}$. The change-of-basis matrix needed to convert vectors from the generic basis B_f to the standard basis B_s is

$${}_B[1]_{B_f} = \begin{bmatrix} \hat{i} \cdot \vec{f}_1 & \hat{i} \cdot \vec{f}_2 & \hat{i} \cdot \vec{f}_3 \\ \hat{j} \cdot \vec{f}_1 & \hat{j} \cdot \vec{f}_2 & \hat{j} \cdot \vec{f}_3 \\ \hat{k} \cdot \vec{f}_1 & \hat{k} \cdot \vec{f}_2 & \hat{k} \cdot \vec{f}_3 \end{bmatrix} = \begin{bmatrix} f_{1x} & f_{2x} & f_{3x} \\ f_{1y} & f_{2y} & f_{3y} \\ f_{1z} & f_{2z} & f_{3z} \end{bmatrix}.$$

Observe the matrix ${}_B[1]_{B_f}$ contains the components of the vectors $\vec{f}_1$, $\vec{f}_2$, and $\vec{f}_3$ as its columns. This makes sense, since we're computing dot products with the vectors $\hat{i}$, $\hat{j}$, and $\hat{k}$.

To obtain the change-of-basis matrix for the opposite direction—from B_s to B_f —compute the inverse of the change-of-basis matrix:

$${}_{B_f}[1]_{B_s} = \left({}_B[1]_{B_f} \right)^{-1}.$$

We can find the inverse using one of the computational techniques we learned in Chapter 3.

Example Consider the basis $B_f = \{\vec{f}_1, \vec{f}_2\}$, where $\vec{f}_1 = (1, 2)_{B_s}^T$ and $\vec{f}_2 = (1, 1)_{B_s}^T$. Find the coordinates of the vector $\vec{v} = (5, 6)_{B_f}^T$ with respect to the standard basis. Also, find the coordinates of the vector $\hat{i}$ with respect to the basis B_f .

To answer the first part of the question, we need to find the change-of-basis matrix from B_f to the standard basis:

$${}_B[1]_{B_f} = \begin{bmatrix} \hat{i} \cdot \vec{f}_1 & \hat{i} \cdot \vec{f}_2 \\ \hat{j} \cdot \vec{f}_1 & \hat{j} \cdot \vec{f}_2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 2 & 1 \end{bmatrix}.$$

We then compute $\begin{bmatrix} 1 & 1 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} 5 \\ 6 \end{bmatrix} = \begin{bmatrix} 11 \\ 16 \end{bmatrix}$. Thus we've shown that $\vec{v} = (11, 16)_{B_s}$.

To answer the second part of the question, we compute $({}_{B_s}[\mathbb{1}]_{B_f})^{-1}$:

$$\left[\begin{array}{cc|cc} 1 & 1 & 1 & 0 \\ 2 & 1 & 0 & 1 \end{array} \right] \sim \left[\begin{array}{cc|cc} 1 & 1 & 1 & 0 \\ 0 & -1 & -2 & 1 \end{array} \right] \sim \left[\begin{array}{cc|cc} 1 & 1 & 1 & 0 \\ 0 & 1 & 2 & -1 \end{array} \right] \sim \left[\begin{array}{cc|cc} 1 & 0 & -1 & 1 \\ 0 & 1 & 2 & -1 \end{array} \right].$$

Having obtained ${}_{B_f}[\mathbb{1}]_{B_s} = ({}_{B_s}[\mathbb{1}]_{B_f})^{-1} = \begin{bmatrix} -1 & 1 \\ 2 & -1 \end{bmatrix}$, we can now compute the coordinates of $\hat{i}$ in the basis B_f using $\begin{bmatrix} -1 & 1 \\ 2 & -1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} -1 \\ 2 \end{bmatrix}$. Thus we've shown that $\hat{i} = (-1, 2)_{B_f}$.

Links

[Khan Academy video on the change-of-basis operation]

<https://youtube.com/watch?v=meibWcbGqt4>

[Change of basis explained by 3Blue1Brown]

<https://youtube.com/watch?v=P2LTAU01TdA>

Exercises

E4.9 Find the coordinates of the vector $\vec{v} = 2\hat{i} + 3\hat{j} + 4\hat{k}$ with respect to the basis $W = \{\hat{w}_1 = (0, 0, 1), \hat{w}_2 = (0, 1, 0), \hat{w}_3 = (1, 0, 0)\}$.

E4.10 Find the change-of-basis matrix ${}_{B_m}[\mathbb{1}]_{B_s}$ that transforms vectors expressed in the standard basis $B_s = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$ to vectors in the modified basis $B_m = \{(1, 1, 0), (1, -1, 0), (0, 0, 2)\}$.

Hint: Find the change-of-basis matrix ${}_{B_s}[\mathbb{1}]_{B_m}$ first.

E4.11 How many vectors exist in a basis for a four-dimensional vector space? Besides the number of vectors it contains, what other condition must a set of vectors satisfy to be a basis?

4.4 Vector spaces

We're about to shift our attention from individual vectors to entire sets of vectors. We're entering the territory of *vector spaces*. For instance, the set of all possible three-dimensional vectors is denoted $\mathbb{R}^3$, and is a type of *vector space*. A vector space consists of a set of vectors and all linear combinations of these vectors. This means if the vectors $\vec{v}_1$ and $\vec{v}_2$ are part of some vector space, then so is the vector $\alpha\vec{v}_1 + \beta\vec{v}_2$ for any α and β . A *vector subspace* consists of a subset of all possible vectors. The vector subspaces of $\mathbb{R}^3$ are lines and planes that pass through the origin.

Since vector spaces and subspaces play a central role in many areas of linear algebra, you'll want to learn about the properties of vector spaces and develop your vocabulary for describing them.

By using the language of vector spaces, you'll be able to describe certain key properties of matrices. The four *fundamental subspaces* associated with a matrix A are its *column space* $\mathcal{C}(A)$, its *row space* $\mathcal{R}(A)$, its *null space* $\mathcal{N}(A)$, and its *left null space* $\mathcal{N}(A^T)$. Let's now define these vector spaces and discuss how they help us understand the solutions to the matrix equation $A\vec{x} = \vec{b}$, and the properties of the linear transformation $T_A(\vec{x}) = A\vec{x}$.

Definitions

- V : a *vector space*
- $\vec{v}$: a *vector*. We use the notation $\vec{v} \in V$ to indicate the vector $\vec{v}$ is part of the vector space V .
- W : a *vector subspace*. We use the notation $W \subseteq V$ to indicate the vector space W is a subspace of the vector space V .
- *span*: the span of a set of vectors is the set of vectors that can be constructed as linear combinations of these vectors:

$$\text{span}(\vec{v}_1, \dots, \vec{v}_n) \stackrel{\text{def}}{=} \{\vec{v} \in V \mid \vec{v} = \alpha_1 \vec{v}_1 + \dots + \alpha_n \vec{v}_n, \alpha_i \in \mathbb{R}\}.$$

For every matrix $M \in \mathbb{R}^{m \times n}$, we define the following *fundamental subspaces* associated with the matrix M :

- $\mathcal{R}(M) \subseteq \mathbb{R}^n$: the *row space* of the matrix M consists of all possible linear combinations of the rows of the matrix M .
- $\mathcal{C}(M) \subseteq \mathbb{R}^m$: the *column space* of the matrix M consists of all possible linear combinations of the columns of the matrix M .
- $\mathcal{N}(M) \subseteq \mathbb{R}^n$: the *null space* of M is the set of vectors that go to the zero vector when multiplying M from the right: $\mathcal{N}(M) \stackrel{\text{def}}{=} \{\vec{v} \in \mathbb{R}^n \mid M\vec{v} = \vec{0}\}$.
- $\mathcal{N}(M^T)$: the *left null space* of M is the set of vectors that go to the zero vector when multiplying M from the left: $\mathcal{N}(M^T) \stackrel{\text{def}}{=} \{\vec{w} \in \mathbb{R}^m \mid \vec{w}^T M = \vec{0}^T\}$.

The dimensions of the column space and the row space of a matrix are equal. We call this dimension the *rank* of the matrix: $\text{rank}(M) \stackrel{\text{def}}{=} \dim(\mathcal{C}(M)) = \dim(\mathcal{R}(M))$.

Vector space

A vector space V consists of a set of vectors and all possible linear combinations of these vectors. The notion of *all possible linear combinations* is very powerful. In particular, it implies two useful properties. First, vector spaces are *closed under addition*: for all vectors in that space, the sum of two vectors is also a vector in that vector space. Mathematically, we write this as

$$\forall \vec{v}_1, \vec{v}_2 \in V, \quad \vec{v}_1 + \vec{v}_2 \in V.$$

Recall the symbol $\forall$ is math shorthand for the phrase “for all.”

Second, vector spaces are *closed under scalar multiplication*:

$$\forall \alpha \in \mathbb{R} \text{ and } \vec{v} \in V, \quad \alpha \vec{v} \in V.$$

Starting from any vector $\vec{v}$ in the vector space V and scaling the vector by any constant α results in a vector in the same vector space.

These two properties codify the essential nature of what a vector space is: a space of vectors that can be added together and scaled by constants.

Span

The *span* operator is a useful shorthand for denoting “the set of all linear combinations” of some set of vectors. This may seem like a weird notion at first, but it will prove very useful for describing vector spaces.

Let’s now illustrate how to define vector spaces using the span operator through some examples. Given a vector $\vec{v}_1 \in V$, define the following vector space:

$$V_1 \stackrel{\text{def}}{=} \text{span}(\vec{v}_1) = \{\vec{v} \in V \mid \vec{v} = \alpha \vec{v}_1 \text{ for some } \alpha \in \mathbb{R}\}.$$

We say V_1 is *spanned* by $\vec{v}_1$, which means any vector in V_1 can be written as a multiple of $\vec{v}_1$. The shape of V_1 is an infinite line.

Given two vectors $\vec{v}_1, \vec{v}_2 \in V$, we define the vector space spanned by these vectors as

$$V_2 \stackrel{\text{def}}{=} \text{span}(\vec{v}_1, \vec{v}_2) = \{\vec{v} \in V \mid \vec{v} = \alpha \vec{v}_1 + \beta \vec{v}_2 \text{ for some } \alpha, \beta \in \mathbb{R}\}.$$

The vector space V_2 contains all vectors that can be written as a linear combination of $\vec{v}_1$ and $\vec{v}_2$. This is a two-dimensional vector space. Geometrically speaking, the shape of V_2 is an infinite plane that passes through the origin.

Now suppose we’re given three vectors $\vec{v}_1, \vec{v}_2, \vec{v}_3 \in V$, such that $\vec{v}_3 = \vec{v}_1 + \vec{v}_2$, and we define the vector space $V_3 \stackrel{\text{def}}{=} \text{span}(\vec{v}_1, \vec{v}_2, \vec{v}_3)$.

The vector space V_3 is actually the same as V_2 ; adding the vector $\vec{v}_3$ to the span of $\vec{v}_1$ and $\vec{v}_2$ does not enlarge the vector space because the vector $\vec{v}_3$ is a linear combination of $\vec{v}_1$ and $\vec{v}_2$. Geometrically speaking, the vector $\vec{v}_3$ lies in the same plane as $\vec{v}_1$ and $\vec{v}_2$.

Consider the vector space $V'_2 = \text{span}(\vec{v}_1, \vec{v}'_2)$, where $\vec{v}'_2 = \gamma\vec{v}_1$, for some $\gamma \in \mathbb{R}$. Since $\vec{v}'_2$ is a linear combination of the vector $\vec{v}_1$, the vector space V'_2 is one-dimensional. In fact, V'_2 is the same as the vector space V_1 defined above: $V'_2 = \text{span}(\vec{v}_1, \vec{v}'_2) = \text{span}(\vec{v}_1) = V_1$.

Note that the word “span” can be used as a verb or a noun. The statement “The vectors $\vec{v}_1$ and $\vec{v}_2$ span the vector space V_2 ” is equivalent to the statement “The span of $\vec{v}_1$ and $\vec{v}_2$ is V_2 .” Both uses of the word communicate the same idea.

Vector subspaces

We use the notation $W \subseteq V$ to indicate that W is a *subspace* of V . A *subspace* is a subset of the vectors in the larger space that has a vector space structure. In other words, $W \subseteq V$ if the following conditions are satisfied:

- W is contained in V : for all $\vec{w}$, if $\vec{w} \in W$, then $\vec{w} \in V$.
- W is closed under addition: for all $\vec{w}_1, \vec{w}_2 \in W$, $\vec{w}_1 + \vec{w}_2 \in W$.
- W is closed under scalar multiplication: for all $\vec{w} \in W$, $\alpha\vec{w} \in W$.

Subspaces always contain the zero vector $\vec{0}$. This is implied by the third condition: *any* vector becomes the zero vector when multiplied by the scalar $\alpha = 0$: $\alpha\vec{w} = 0\vec{w} = \vec{0}$.

Subspaces specified by constraints

One way to define a vector subspace W is to start with a larger space V and describe a *constraint* that is satisfied by all vectors in the subspace W . For example, the xy -plane is defined as the set of vectors $(x, y, z) \in \mathbb{R}^3$ that satisfy the constraint

$$(0, 0, 1) \cdot (x, y, z) = 0.$$

More formally, we define the xy -plane as

$$P_{xy} = \{(x, y, z) \in \mathbb{R}^3 \mid (0, 0, 1) \cdot (x, y, z) = 0\}.$$

Since the vector $\hat{k} = (0, 0, 1)$ is perpendicular to all vectors in the xy -plane, we can describe the xy -plane as “the set of all vectors perpendicular to the vector $\hat{k}$.”

Subspaces specified as a span

Another way to represent the xy -plane is to describe it as the span of two linearly independent vectors in the plane:

$$P_{xy} = \text{span}((1, 0, 0), (0, 1, 0)),$$

which is equivalent to saying,

$$P_{xy} = \{\vec{v} \in \mathbb{R}^3 \mid \vec{v} = \alpha(1, 0, 0) + \beta(0, 1, 0), \forall \alpha, \beta \in \mathbb{R}\}.$$

This expression is a *parametrization* of the space P_{xy} with α and β as the two parameters. Each point in the plane is described by a unique pair of parameters (α, β) . The parametrization of an m -dimensional vector space requires m parameters.

Subsets vs. subspaces

In linear algebra, the terms *subset* and *subspace* are used somewhat interchangeably, and the same symbol is used to denote both subset ($S \subseteq V$) and subspace ($W \subseteq V$) relationships. Every subspace is also a subset, but not every subset is a subspace.

Let's clarify the distinction between the terms *subset* and *subspace*. Assume we're working in a vector space V . A subset of V can be described in the form $S = \{\vec{v} \in V \mid \langle \text{conditions} \rangle\}$ and consists of all vectors in V that satisfy the $\langle \text{conditions} \rangle$. A subspace W is a type of subset with a *vector space structure*, meaning it is closed under addition (for all $\vec{w}_1, \vec{w}_2 \in W$, $\vec{w}_1 + \vec{w}_2 \in W$), and closed under scalar multiplication (for all $\vec{w} \in W$, $\alpha\vec{w} \in W$). When mathematicians refer to some subset as a *subspace*, they're letting you know that you can take arbitrary vectors in the set, scale or add them together, and obtain a vector in the same set.

To further illustrate the difference between *subsets* and *subspaces*, let's compare the solution sets of two systems of equations: $A\vec{x} = \vec{b}$ and $A\vec{x} = \vec{0}$. The solution set of $A\vec{x} = \vec{b}$ is $\{\vec{c} + \vec{v}_n\}$, for all $\vec{v}_n \in \mathcal{N}(A)$, which is not a subspace of V unless $\vec{c} = \vec{0}$. Observe that if $\vec{x}_1 = \vec{c} + \vec{v}_1$ and $\vec{x}_2 = \vec{c} + \vec{v}_2$ are two solutions to $A\vec{x} = \vec{b}$, their sum is not a solution: $\vec{x}_1 + \vec{x}_2 = 2\vec{c} + \vec{v}_1 + \vec{v}_2 \notin \{\vec{c} + \vec{v}_n\}$. The solution set to the equation $A\vec{x} = \vec{b}$ is not a subspace. In contrast, the solution set of $A\vec{x} = \vec{0}$ is a *vector space* that is called *null space of A* and denoted $\mathcal{N}(A)$. If $\vec{x}_1$ and $\vec{x}_2$ are two solutions to $A\vec{x} = \vec{0}$, then $\alpha\vec{x}_1 + \beta\vec{x}_2$ is also a solution to $A\vec{x} = \vec{0}$.

A real-life situation

You walk into class one day and are caught completely off guard by a surprise quiz—wait, let’s make it a mini-exam for emotional effect. Although you’ve read a chapter or two in the book, you’ve been “busy” and are totally unprepared for this exam. The first question asks you to “find the solution of the *homogeneous* system of equations and the *non-homogeneous* system of equations.” You rack your brain, but the only association with homogeneity that comes to mind is the homogenized milk you had for breakfast. Oh, there’s more: the question also asks you to “state whether each of the solutions obtained is a *vector space*.” As you stare at the page, the words and equations begin to blur and panic sets in.

Don’t fear! Look at the problem again. You don’t know what the heck a homogeneous system of equations is, but you sure as heck know how to solve systems of equations. You solve the given system of equations $A\vec{x} = \vec{b}$ by building the augmented matrix $[A \mid \vec{b}]$ and computing its reduced row echelon form using row operations. You obtain the solution set $\vec{x} = \{\vec{v} \in V \mid \vec{v} = \vec{c} + t\vec{v}_n, \forall t \in \mathbb{R}\}$, where $\vec{c}$ is the particular solution and $\vec{v}_n$ is a vector that spans the null space of A .

Next, you ponder the “vector space” part of the question. You notice the solution set to the *system* of equations $A\vec{x} = \vec{b}$ isn’t a vector space since it doesn’t pass through the origin. However, the solution set to the equation $A\vec{x} = \vec{0}$ is a vector space $\{\vec{v} \in V \mid \vec{v} = t\vec{v}_n, \forall t \in \mathbb{R}\} = \text{span}(\vec{v}_n)$. Suddenly it clicks: a *homogeneous* system of equations must be the system of equations $A\vec{x} = \vec{0}$, in which the constants on the right side of the equation are all zero. The term *homogeneous* kind of makes sense; all the constants of the right side have the same value $b_1 = b_2 = \dots = 0$. The solution to the non-homogeneous system of equations $A\vec{x} = \vec{b}$ is the set $\{\vec{c} + s\vec{v}_n, \forall s \in \mathbb{R}\}$, which is not a vector space. The solution to the homogeneous system of equations $A\vec{x} = \vec{0}$ is $\{\vec{v} \in V \mid \vec{v} = t\vec{v}_n, \forall t \in \mathbb{R}\}$, which is a vector space. Well done!

Matrix fundamental subspaces

We now define four *fundamental subspaces* associated with a matrix $M \in \mathbb{R}^{m \times n}$.

- The column space $\mathcal{C}(M)$ is the span of the columns of the matrix. The column space consists of all possible output vectors the matrix can produce when multiplied by a vector from the

right:

$$\mathcal{C}(M) \stackrel{\text{def}}{=} \{\vec{w} \in \mathbb{R}^m \mid \vec{w} = M\vec{v} \text{ for some } \vec{v} \in \mathbb{R}^n\}.$$

- The null space $\mathcal{N}(M)$ of a matrix $M \in \mathbb{R}^{m \times n}$ consists of all vectors the matrix M sends to the zero vector:

$$\mathcal{N}(M) \stackrel{\text{def}}{=} \{\vec{v} \in \mathbb{R}^n \mid M\vec{v} = \vec{0}\}.$$

The null space is sometimes called the *kernel* of the matrix.

- The row space $\mathcal{R}(M)$ is the span of the rows of the matrix. We obtain linear combinations of the rows by multiplying the matrix with an m -dimensional vector from the left:

$$\mathcal{R}(M) \stackrel{\text{def}}{=} \{\vec{v} \in \mathbb{R}^m \mid \vec{v} = \vec{w}^T M \text{ for some } \vec{w} \in \mathbb{R}^n\}.$$

Note, we used the transpose ^{T} to transform $\vec{w}$ to a row vector.

- The left null space $\mathcal{N}(M^T)$ of a matrix $M \in \mathbb{R}^{m \times n}$ consists of all vectors the matrix M sends to the zero vector when multiplied from the left:

$$\mathcal{N}(M^T) \stackrel{\text{def}}{=} \{\vec{w} \in \mathbb{R}^m \mid \vec{w}^T M = \vec{0}^T\}.$$

These vector spaces are called *fundamental* because they describe important properties of the matrix M . Recall that matrix equations can be used to represent systems of linear equations, as well as linear transformations. A solid understanding of the fundamental subspaces leads to a solid understanding of linear equations and linear transformations.

Matrices and systems of linear equations

The null space $\mathcal{N}(M)$ corresponds to the solution set of the matrix equation $M\vec{x} = \vec{0}$. If a matrix has a nonempty null space, the system of equations corresponding to $M\vec{x} = \vec{b}$ has an infinite solution set. Indeed, we can write the solution of $M\vec{x} = \vec{b}$ as a *particular solution* $\vec{c}$ plus all possible vectors in the null space of M :

$$\vec{x} = \vec{c} + \text{span}(\vec{v}_1, \dots, \vec{v}_k), \quad \text{where} \quad \text{span}(\vec{v}_1, \dots, \vec{v}_k) = \mathcal{N}(M).$$

We can verify this claim as follows. Suppose $\vec{x} = \vec{c}$ is a solution to the equation $M\vec{x} = \vec{b}$. Consider the vector $\vec{x} = \vec{c} + \alpha_1\vec{v}_1 + \dots + \alpha_k\vec{v}_k$, which contains $\vec{c}$ and some arbitrary linear combination of vectors from the null space of M . Observe that $\vec{x}$ is also a solution to the equation $M\vec{x} = \vec{b}$:

$$M\vec{x} = M(\vec{c} + \alpha_1\vec{v}_1 + \dots + \alpha_k\vec{v}_k) = M\vec{c} + \alpha_1 \cancel{M\vec{v}_1}^0 + \dots + \alpha_k \cancel{M\vec{v}_k}^0 = M\vec{c} = \vec{b}.$$

If the null space of M contains only the zero vector $\{\vec{0}\}$, then the system of equations $M\vec{x} = \vec{b}$ has a unique solution.

It's worth connecting the above observations with what we learned about the Gauss–Jordan elimination procedure in Chapter 3. Suppose we want to solve the matrix equation $M\vec{x} = \vec{b}$, where M is an $(m \times n)$ -dimensional matrix, $\vec{x} = (x_1, x_2, \dots, x_n)^T$ is an n -dimensional vector of unknowns, and $\vec{b}$ is an m -dimensional vector of constants. We can construct the augmented matrix $[M \mid \vec{b}]$ and perform row operations until we bring the augmented matrix to its reduced row echelon form $[\text{rref}(M) \mid \vec{c}]$. In the case where the null space of M is k -dimensional, the reduced row echelon form of the matrix M will have k free variables and $n - k$ pivots. The solution the system of equations is $\vec{x} = \vec{c} + \alpha_1 \vec{v}_1 + \dots + \alpha_k \vec{v}_k$, where $\{\vec{v}_1, \dots, \vec{v}_k\}$ are the vectors in the null space of M .

Matrices and linear transformations

Matrices can be used to *represent* linear transformations. We postpone the detailed discussion about linear transformations and their representation as matrices until Chapter 5, but we'll discuss the subject here briefly—mainly to introduce an important connection between the column space and the row space of a matrix, and to explain why each matrix has *two* null spaces (what's up with that?).

Matrix-vector and vector-matrix products

A matrix $M \in \mathbb{R}^{m \times n}$ corresponds to not one but *two* linear transformations. Up until now we've focused on the matrix-vector product $M\vec{x} = \vec{y}$, which corresponds to a linear transformation of the form $T_M : \mathbb{R}^n \rightarrow \mathbb{R}^m$. In addition to the linear transformation we obtain by multiplication from the right, there is also the option of multiplying M by a vector from the left: $\vec{a}^T M = \vec{b}^T$, where $\vec{a}^T$ (the input) is an m -dimensional row vector, and $\vec{b}^T$ (the output) is an n -dimensional row vector.¹

The vector-matrix product $\vec{a}^T M = \vec{b}^T$ corresponds to a linear transformation of the form $T_{M^T} : \mathbb{R}^m \rightarrow \mathbb{R}^n$. We identify the output of this linear transformation, $\vec{b} = T_{M^T}(\vec{a})$, with the result of the vector-matrix product $\vec{b}^T = \vec{a}^T M$. The linear transformation $T_{M^T} : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is called the *adjoint* of the linear transformation $T_M : \mathbb{R}^n \rightarrow \mathbb{R}^m$. Adjoint linear transformations are represented by

¹Our convention is to assume vectors are column vectors by default. The transpose operation is required to obtain row vectors.

the same matrix M . T_M is defined as the multiplication of M from the right, while T_{M^T} is defined as multiplication of M from the left.

Let's clarify why we used the notation T_{M^T} to denote the adjoint operation of T_M . We previously used the notation T_A to describe the linear transformation obtained by right multiplication by A ($T_A(\vec{x}) = A\vec{x}$). Instead of creating a new notation for left multiplication, we can transform left multiplication into right multiplication by using the transpose operation:

$$\vec{a}^T A = \vec{b}^T \quad \Rightarrow \quad (\vec{a}^T A)^T = (\vec{b}^T)^T \quad \Rightarrow \quad A^T \vec{a} = \vec{b}.$$

We can think of left multiplication by A as right multiplication by A^T . This correspondence also explains why we use the notation $\mathcal{N}(M^T)$ for the left null space of M ; we can rewrite the condition $\vec{w}^T M = \vec{0}^T$ as $M^T \vec{w} = \vec{0}$, so the left null space of M is equivalent to the right null space of M^T .

Left and right input spaces

Let's call the *left space* of M the set of vectors suitable for multiplying M from the left. Similarly, we'll call *right space* of M the set of vectors suitable for multiplying M from the right. If $M \in \mathbb{R}^{m \times n}$, the left space is $\mathbb{R}^m$ and the right space is $\mathbb{R}^n$.

By combining all the vectors in the row space of M and all the vectors in its null space, we obtain the full right space:

$$\mathcal{R}(M) \oplus \mathcal{N}(M) = \mathbb{R}^n.$$

This means any vector $\vec{v} \in \mathbb{R}^n$ can be written as a sum, $\vec{v} = \vec{v}_r + \vec{v}_n$, such that $\vec{v}_r \in \mathcal{R}(M)$ and $\vec{v}_n \in \mathcal{N}(M)$. The symbol $\oplus$ stands for *orthogonal sum*, which means we can pick $\vec{v}_r$ and $\vec{v}_n$ to be orthogonal vectors, $\vec{v}_r \cdot \vec{v}_n = 0$.

If we consider the dimensions involved the above equation, we obtain the following important relation between the dimension of the row space and the null space of a matrix:

$$\dim(\mathcal{R}(M)) + \dim(\mathcal{N}(M)) = n = \dim(\mathbb{R}^n).$$

The n -dimensional right space splits into row-space dimensions and null-space dimensions.

Similar to the split in the right space, the left-space $\mathbb{R}^m$ decomposes into an orthogonal sum of the column space and the left null space of the matrix:

$$\mathcal{C}(M) \oplus \mathcal{N}(M^T) = \mathbb{R}^m.$$

If we count the dimensions in this equation, we obtain a relation between the dimension of the column space and the left null space of the matrix: $\dim(\mathcal{C}(M)) + \dim(\mathcal{N}(M^T)) = m$.

Matrix rank

The column space and the row space of a matrix have the same dimension. We call this dimension the *rank* of the matrix:

$$\text{rank}(M) \stackrel{\text{def}}{=} \dim(\mathcal{R}(M)) = \dim(\mathcal{C}(M)).$$

The *rank* of M is the number of linearly independent rows in M , which is equal to the number of linearly independent columns in M . The dimension of the null space of M is called the *nullity* of M : $\text{nullity}(M) \stackrel{\text{def}}{=} \dim(\mathcal{N}(M))$.

Applying this new terminology, we can update our earlier observation about the dimensions of the right fundamental subspaces of a matrix:

$$\text{rank}(M) + \text{nullity}(M) = n = \dim(\mathbb{R}^n).$$

This formula is called the *rank–nullity theorem*, and can be used to deduce the rank of a matrix given its nullity, or vice versa.

Summary

Together, $\mathcal{R}(M)$, $\mathcal{N}(M)$, $\mathcal{C}(M)$, and $\mathcal{N}(M^T)$ describe all aspects of the matrix M when multiplied by vectors from the left or the right. Everything we’ve learned so far about how the matrix M maps vectors between its left and right spaces can be summarized by the following observations:

$$\begin{array}{ccc} \mathcal{C}(M) & \xleftarrow{M} & \mathcal{R}(M), \\ \vec{0} & \xleftarrow{M} & \mathcal{N}(M), \\ \mathcal{N}(M^T) & \xrightarrow{M} & \vec{0}. \end{array}$$

Note the zero vector in the second row is $\vec{0} \in \mathbb{R}^m$, while the zero vector in the third row is $\vec{0} \in \mathbb{R}^n$. In Section 5.1, we’ll learn how to interpret the fundamental subspaces of the matrix M as the input and output spaces of the linear transformations $T_M : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $T_{M^T} : \mathbb{R}^m \rightarrow \mathbb{R}^n$.

Linear independence

One of the most important ideas in linear algebra is the notion of linear independence. Given a set of vectors $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$, we’re often

interested in knowing whether one of the vectors in the set can be written as a linear combination of the other vectors. For example, if $\vec{v}_3 = \vec{v}_1 + \vec{v}_2$, then the set $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ is *linearly dependent*. On the other hand, if none of the vectors in the set can be written as a linear combination of the other vectors, then the set is *linearly independent*.

Understanding the concept of linear independence of a set of vectors is a prerequisite for understanding the concept of a basis for a vector space, so it is important to give a clear definition.

Definition: The set of vectors $\{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n\}$ is *linearly independent* if the only solution to the equation

$$\alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \dots + \alpha_n \vec{v}_n = \vec{0}$$

is the zero vector $(\alpha_1, \alpha_2, \dots, \alpha_n) = (0, 0, \dots, 0)$.

When $\vec{\alpha} = \vec{0}$ is the only solution to the equation, then none of the vectors $\vec{v}_i$ can be written as a linear combination of the other vectors.

To understand the importance of the all zeros solution, let's consider an example where a nonzero solution $\vec{\alpha}$ exists. Suppose the set of vectors $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ satisfy $\alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \alpha_3 \vec{v}_3 = \vec{0}$, with $\alpha_1 = 1$, $\alpha_2 = 1$, and $\alpha_3 = -1$. Then we can write $\vec{v}_3 = \vec{v}_1 + \vec{v}_2$, which shows that $\vec{v}_3$ is a linear combination of $\vec{v}_1$ and $\vec{v}_2$, hence the vectors are linearly dependent. The strange wording of the definition in terms of an all zeros solution is required to make the definition of linear independence symmetric. An all zero alphas solution implies that *no* vector can be written as a linear combination of the other vectors.

Basis

To carry out calculations with vectors in a vector space V , we need to know a basis $B = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ for that space. Intuitively, a basis for a vector space is any set of vectors that can serve as a coordinate system for that vector space. A *basis* for an n -dimensional vector space V is a set of n linearly independent vectors in V .

Throughout this section, we've referred to the *dimension* of a vector space, which is essentially the number of vectors in a basis for that vector space. A basis for an n -dimensional vector space contains exactly n vectors. Any set of fewer than n vectors would not satisfy the spanning property. Any set of more than n vectors from V cannot be linearly independent (see page 226). To form a basis for a vector space, a set of vectors must be "just right": it must contain a sufficient number of vectors (but not too many) so that the coordinates of each vector are uniquely determined.

The rank–nullity theorem

The relationship between the *rank* and the *nullity* of a matrix is so important that it's worth formally stating the rank–nullity theorem.

Rank–nullity theorem. *For any matrix $M \in \mathbb{R}^{m \times n}$, the following statement holds:*

$$\text{rank}(M) + \text{nullity}(M) = n,$$

where the rank of M is $\text{rank}(M) = \dim(\mathcal{R}(M)) = \dim(\mathcal{C}(M))$ and its nullity is defined as $\text{nullity}(M) = \dim(\mathcal{N}(M))$.

The rank–nullity theorem is important because it “splits” the vectors in the right space of the matrix into two categories: those that lie in its row space and those that lie in its null space. We can use the rank–nullity theorem to infer the dimension of the row space of a matrix given the dimension of the null space. Vice versa, given the dimension of a null space, we can infer the dimension of the matrix's row space.

Example 1 Suppose the null space of the matrix $M \in \mathbb{R}^{m \times n}$ consists only of the zero vector, $\mathcal{N}(M) = \{\vec{0}\}$; then the nullity of the matrix is $\text{nullity}(M) = \dim(\mathcal{N}(M)) = 0$. Using the rank–nullity theorem we can conclude that $\text{rank}(M) = \dim(\mathcal{R}(M)) = \dim(\mathcal{C}(M)) = n$.

Example 2 Consider a matrix $A \in \mathbb{R}^{3 \times 6}$. After performing some calculations (which we'll discuss in the next section) we find that one of the rows of the matrix is a linear combination of the other two. Therefore, the row space of A is two-dimensional and $\text{rank}(A) = 2$. From this, we can infer $\text{nullity}(A) = 6 - 2 = 4$, meaning the null space of A is four-dimensional.

Distilling bases

A basis for an n -dimensional vector space V consists of *exactly* n vectors. Any set of vectors $\{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ can serve as a basis for V , as long as the vectors are linearly independent and there are exactly n of them.

Sometimes an n -dimensional vector space V is specified as the span of more than n vectors:

$$V = \text{span}(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_m), \quad m > n.$$

Since there are $m > n$ of the $\vec{v}$ -vectors, there are *too many* to form a basis. We say this set of vectors is *overcomplete*. They cannot all be linearly independent since there can be at most n linearly independent vectors in an n -dimensional vector space.

If we want to find a basis for the space V , we must reject some of the vectors. Given the set of vectors $\{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_m\}$, our task is to *distill* a set of n linearly independent vectors $\{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ from them. We'll learn how to do this in the next section.

Links

[Linear combinations, span, and bases explained by 3Blue1Brown]
<https://youtube.com/watch?v=k7RM-ot2NWY>

[Inverse, column space, and null space explained by 3Blue1Brown]
<https://youtube.com/watch?v=uQhTuRlWMxw>

Exercises

E4.12 Determine whether the following subsets are subspaces of $\mathbb{R}^3$.

- The single vector $(0, 0, 0)$
- The line $\ell : \{(3, 3, 3) + t(1, 1, 1), \forall t \in \mathbb{R}\}$
- The plane defined by the equation $z = 0$
- The plane defined by the equation $2x + y + z = 3$
- $\text{span}((1, 1, 0), (0, 1, 1))$

E4.13 Find the rank and the dimensions of the row space, the column space, the null space, and the left null space of each of these matrices.

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad B = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad C = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Make sure your answers are consistent with the rank-nullity theorem.

E4.14 Consider the matrix $A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \in \mathbb{R}^{3 \times 4}$. Specify bases for its row space $\mathcal{R}(A)$, its column space $\mathcal{C}(A)$, and its null space $\mathcal{N}(A)$.

E4.15 Do the matrices A and $-A$ have the same fundamental subspaces?

4.5 Vector space techniques

In this section, we'll learn how to “distill” a basis for any vector space, which is an important procedure for characterizing vector spaces. Actually, the procedure is not new—it's really an application of the Gauss–Jordan elimination procedure we saw in Section 3.1.

Starting from a set of vectors that are not linearly independent, we can write them as the rows of a matrix, and then perform *row*

operations on this matrix until we find the reduced row echelon form of the matrix. Since row operations do not change the row space of a matrix, the nonzero rows in the final RREF of the matrix will span the same space as the original set of vectors. The rows in the RREF of the matrix will be linearly independent and thus will form a basis.

The ability to distill a basis is important when characterizing any vector space. The basis serves as the coordinate system for that vector space, and the number of vectors in a basis tells us the dimension of the vector space. For this reason, we'll spend an entire section learning how to distill bases for various vector spaces.

Finding a basis

Suppose the vector subspace V is defined as the span of m vectors $\{\vec{u}_1, \vec{u}_2, \dots, \vec{u}_m\}$, which are not necessarily linearly independent:

$$V \stackrel{\text{def}}{=} \text{span}(\vec{u}_1, \vec{u}_2, \dots, \vec{u}_m).$$

Our task is to find a basis for V . We're looking for an alternate description of the vector space V as

$$V = \text{span}(\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n),$$

such that the vectors $\{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ will be linearly independent.

One way to accomplish this task is to write the vectors $\vec{u}_i$ as the rows of a matrix M . By this construction, the space V corresponds to the *row space* of the matrix M , denoted $\mathcal{R}(M)$. We can then use standard *row operations* to bring the matrix into its reduced row echelon form. Applying row operations to a matrix does not change its row space: $\mathcal{R}(M) = \mathcal{R}(\text{rref}(M))$. Transforming the matrix into its RREF allows us to see which of the rows are linearly independent, and thus can serve as basis vectors:

$$\begin{bmatrix} - & \vec{u}_1 & - \\ - & \vec{u}_2 & - \\ - & \vec{u}_3 & - \\ & \vdots & \\ - & \vec{u}_m & - \end{bmatrix} \xrightarrow{\text{Gauss-Jordan elimination}} \begin{bmatrix} - & \vec{e}_1 & - \\ & \vdots & \\ - & \vec{e}_n & - \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

The nonzero rows in the RREF of the matrix form a set of linearly independent vectors $\{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ that span the vector space V . The linearly dependent vectors have been reduced to rows of zeros.

The above process is called “finding a basis” or “distilling a basis” and it's important you understand how to carry out this procedure. Even more important is that you understand *why* we'd want to distill

a basis in the first place! By the end of the Gauss–Jordan procedure, we obtain a description of the same vector space V in terms of a new set of vectors. Why is it better to describe the vector space V in terms of the vectors $\{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$, rather than in terms of $\{\vec{u}_1, \vec{u}_2, \dots, \vec{u}_m\}$?

I'll tell you exactly why. We prefer to use the basis $\{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ to characterize the vector space V because there exists a one-to-one correspondence between each vector $\vec{v} \in V$ and the coefficients $v_1, v_2, \dots, v_n$ in the linear combination

$$\vec{v} = v_1\vec{e}_1 + v_2\vec{e}_2 + \dots + v_n\vec{e}_n.$$

Using the basis $B = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ allows us to represent each vector $\vec{v} \in V$ as a unique list of coordinates $(v_1, v_2, \dots, v_n)_B$.

Expressing every vector $\vec{v}$ as a unique list of coordinates would not be possible if we used the vectors $\{\vec{u}_1, \vec{u}_2, \dots, \vec{u}_m\}$ to describe the vector space V . Since the vectors $\{\vec{u}_1, \vec{u}_2, \dots, \vec{u}_m\}$ are not linearly independent, the same vector $\vec{v}$ can be represented by many different linear combinations of the form

$$\vec{v} = v'_1\vec{u}_1 + v'_2\vec{u}_2 + \dots + v'_m\vec{u}_m.$$

We cannot identify $\vec{v}$ with a *unique* set of coefficients $v'_1, v'_2, \dots, v'_m$, therefore vectors are not represented faithfully by the coefficients in the linear combination.

Another reason we prefer to describe V in terms of a basis is because we can immediately see the vector space V is n -dimensional, since there are n vectors in the basis for V .

Definitions

- $B = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$. A *basis* for an n -dimensional vector space S is a set of n linearly independent vectors that span S . Any vector $\vec{v} \in S$ can be written as a linear combination of the basis vectors:

$$\vec{v} = v_1\vec{e}_1 + v_2\vec{e}_2 + \dots + v_n\vec{e}_n.$$

A basis for an n -dimensional vector space contains exactly n vectors.

- $\dim(S)$: the dimension of the vector space S is equal to the number of vectors in a basis for S .

Recall the four *fundamental subspaces* of a matrix $M \in \mathbb{R}^{m \times n}$ we defined in the previous section:

- $\mathcal{R}(M) \subseteq \mathbb{R}^n$: the *row space* of the matrix M that consists of all possible linear combinations of the rows of the matrix M .

- $\mathcal{N}(M) \subseteq \mathbb{R}^n$: the *null space* of the matrix contains all the vectors that become the zero vector when multiplied by M :

$$\mathcal{N}(M) \stackrel{\text{def}}{=} \{ \vec{v} \in \mathbb{R}^n \mid M\vec{v} = \vec{0} \}.$$

- $\mathcal{C}(M) \subseteq \mathbb{R}^m$: the *column space* of the matrix M that consists of all possible linear combinations of the columns of the matrix M .
- $\mathcal{N}(M^T) \subseteq \mathbb{R}^m$: the *left null space* of the matrix contains all the vectors that become the zero vector when multiplying M from the left:

$$\mathcal{N}(M^T) \stackrel{\text{def}}{=} \{ \vec{w} \in \mathbb{R}^m \mid \vec{w}^T M = \vec{0}^T \}.$$

Bases for the fundamental subspaces of matrices

Performing the Gauss–Jordan elimination procedure on a matrix A has the effect of distilling a basis for its row space $\mathcal{R}(A)$. How do we find bases for the other fundamental subspaces of a matrix? In this section, we’ll learn about a useful shortcut for computing bases for the column space $\mathcal{C}(A)$ and the null space $\mathcal{N}(A)$ of a matrix, starting from the reduced row echelon form of the matrix. Sorry, there is no shortcut for finding the left null space—we’ll have to use the transpose operation to obtain A^T and then find its null space $\mathcal{N}(A^T)$.

Pay careful attention to the locations of the pivots (leading ones) in the RREF of A , because they play an important role in the procedures described below.

Basis for the row space

The row space $\mathcal{R}(A)$ of a matrix A is defined as the space of all vectors that can be written as linear combinations of the rows of A . To find a basis for $\mathcal{R}(A)$, we use the Gauss–Jordan elimination procedure:

1. Perform row operations to find the RREF of A .
2. The nonzero rows in the RREF of A form a basis for $\mathcal{R}(A)$.

Basis for the column space

To find a basis for the column space $\mathcal{C}(A)$ of a matrix A , we need to determine which columns of A are linearly independent. To find the linearly independent columns of A , follow these steps:

1. Perform row operations to find the RREF of A .
2. Identify the columns that contain pivots (leading ones).

3. The corresponding columns in the original matrix A form a basis for the column space of A .

This procedure works because elementary row operations do not change the independence relations between the columns of the matrix. If two columns are linearly independent in the RREF of A , then these columns are also linearly independent in the original matrix A .

Note that the column space of the matrix A corresponds to the row space of the matrix transposed A^T . From this fact, we derive another procedure for finding a basis for the column space of a matrix A : we can use the find-a-basis-for-the-row-space procedure on A^T .

Basis for the null space

The null space $\mathcal{N}(A)$ of a matrix $A \in \mathbb{R}^{m \times n}$ is

$$\mathcal{N}(A) = \{\vec{x} \in \mathbb{R}^n \mid A\vec{x} = \vec{0}\}.$$

The vectors in the null space are orthogonal to the row space of the matrix A .

The null space of A is the *solution* of the equation $A\vec{x} = \vec{0}$. You should already be familiar with the procedure for finding the solution of systems of equations from Section 3.1. The steps of the procedure are:

1. Perform row operations to find the RREF of A .
2. Identify the columns that *do not* contain a leading one. These columns correspond to *free variables* of the solution. For example, consider a matrix whose reduced row echelon form is

$$\text{rref}(A) = \begin{bmatrix} \mathbf{1} & 2 & 0 & 0 \\ 0 & 0 & \mathbf{1} & -3 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

The second column and the fourth column do not contain leading ones (pivots), so these columns correspond to free variables, which are customarily called s , t , r , etc. We're looking for a vector with two free variables: $(x_1, s, x_3, t)^T$.

3. Rewrite the null space problem as a system of equations:

$$\begin{bmatrix} 1 & 2 & 0 & 0 \\ 0 & 0 & 1 & -3 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ s \\ x_3 \\ t \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad \Rightarrow \quad \begin{array}{rcl} 1x_1 + 2s & = & 0 \\ 1x_3 - 3t & = & 0 \\ 0 & = & 0. \end{array}$$

We can express the unknowns x_1 and x_3 in terms of the free variables s and t : $x_1 = -2s$ and $x_3 = 3t$. The vectors in the null space are of the form $(-2s, s, 3t, t)^T$, for all $s, t \in \mathbb{R}$. We can rewrite this expression by splitting it into an s -part and a t -part:

$$\begin{bmatrix} -2s \\ s \\ 3t \\ t \end{bmatrix} = \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix} s + \begin{bmatrix} 0 \\ 0 \\ 3 \\ 1 \end{bmatrix} t.$$

4. The direction vectors associated with each free variable form a basis for the null space of the matrix A :

$$\mathcal{N}(A) = \left\{ \begin{bmatrix} -2s \\ s \\ 3t \\ t \end{bmatrix}, \forall s, t \in \mathbb{R} \right\} = \text{span} \left(\begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 3 \\ 1 \end{bmatrix} \right).$$

Verify that the matrix A multiplied by any vector from its null space produces a zero vector.

Examples

Let's check out some examples that illustrate the procedures for finding bases for $\mathcal{R}(A)$, $\mathcal{C}(A)$, and $\mathcal{N}(A)$. It's important you become proficient at these "find a basis" tasks because they often appear on homework assignments and exams.

Example 1 Find a basis for the row space, the column space, and the null space of the matrix:

$$A = \begin{bmatrix} 4 & -4 & 0 \\ 1 & 1 & -2 \\ 2 & -6 & 4 \end{bmatrix}.$$

The first steps toward finding the row space, column space, and the null space of a matrix all require calculating the RREF of the matrix, so this is what we'll do first.

1. Let's focus on the first column. To create a pivot in the top left corner, we divide the first row by 4, denoted $R_1 \leftarrow \frac{1}{4}R_1$:

$$\begin{bmatrix} 1 & -1 & 0 \\ 1 & 1 & -2 \\ 2 & -6 & 4 \end{bmatrix}.$$

2. We use this pivot to clear the numbers on the second and third rows by performing $R_2 \leftarrow R_2 - R_1$ and $R_3 \leftarrow R_3 - 2R_1$:

$$\begin{bmatrix} 1 & -1 & 0 \\ 0 & 2 & -2 \\ 0 & -4 & 4 \end{bmatrix}.$$

3. We can create a pivot in the second row if we divide it by 2, denoted $R_2 \leftarrow \frac{1}{2}R_2$:

$$\begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & -4 & 4 \end{bmatrix}.$$

4. We now clear the entry below the pivot using $R_3 \leftarrow R_3 + 4R_2$:

$$\begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{bmatrix}.$$

5. The final simplification step is to clear the -1 in the first row using $R_1 \leftarrow R_1 + R_2$:

$$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{bmatrix}.$$

Now that we have the RREF of the matrix, we can answer the questions like professionals.

Before we find bases for the fundamental subspaces of A , let's first do some basic dimension counting. Observe that the matrix has just two pivots. We say $\text{rank}(A) = 2$. This means both the row space and the column spaces are two-dimensional.

Recall the equation $n = \text{rank}(A) + \text{nullity}(A)$, which we saw in the previous section. The right space $\mathbb{R}^3$ splits into two types of vectors: those in the row space of A and those in the null space. Since we know the row space is two-dimensional, we can deduce the dimension of the null space: $\text{nullity}(A) = \dim(\mathcal{N}(A)) = n - \text{rank}(A) = 3 - 2 = 1$.

Now let's answer the questions posed in the problem. The row space of A consists of the two nonzero vectors in the RREF of A :

$$\mathcal{R}(A) = \text{span}((1, 0, -1), (0, 1, -1)).$$

To find the column space of A , observe that the first and second columns contain the pivots in the RREF of A . If they do, then the

first two columns of the original matrix A form a basis for the column space of A :

$$\mathcal{C}(A) = \text{span} \left(\begin{bmatrix} 4 \\ 1 \\ 2 \end{bmatrix}, \begin{bmatrix} -4 \\ 1 \\ -6 \end{bmatrix} \right).$$

Let's now find an expression for the null space of A . First, observe that the third column does not contain a pivot. No pivot indicates that the third column corresponds to a free variable; it can take on any value, which we write $x_3 = t$, $t \in \mathbb{R}$. We want to give a description of all vectors $(x_1, x_2, t)^T$ that satisfy the system of equations:

$$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ t \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad \Rightarrow \quad \begin{array}{rcl} 1x_1 - 1t & = & 0 \\ 1x_2 - 1t & = & 0 \\ 0 & = & 0. \end{array}$$

Solving the system equations for x_1 and x_2 , we find $x_1 = t$ and $x_2 = t$, and thus obtain the following final expression for the null space:

$$\mathcal{N}(A) = \left\{ \begin{bmatrix} t \\ t \\ t \end{bmatrix}, t \in \mathbb{R} \right\} = \text{span} \left(\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right).$$

The null space of A is one-dimensional and consists of all multiples of the vector $(1, 1, 1)^T$.

Example 2 Find a basis for the row space, column space, and null space of the matrix:

$$B = \begin{bmatrix} 1 & 3 & 1 & 4 \\ 2 & 7 & 3 & 9 \\ 1 & 5 & 3 & 1 \\ 1 & 2 & 0 & 8 \end{bmatrix}.$$

First, we find the reduced row echelon form of the matrix B :

$$\sim \begin{bmatrix} 1 & 3 & 1 & 4 \\ 0 & 1 & 1 & 1 \\ 0 & 2 & 2 & -3 \\ 0 & -1 & -1 & 4 \end{bmatrix} \sim \begin{bmatrix} 1 & 0 & -2 & 1 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & -5 \\ 0 & 0 & 0 & 5 \end{bmatrix} \sim \begin{bmatrix} 1 & 0 & -2 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

As in the previous example, we begin by calculating the dimensions of the subspaces. The rank of this matrix is three, so the column space and the row space will be three-dimensional. Since the right space is $\mathbb{R}^4$, this leaves one dimension for the null space. Next, let's find the fundamental subspaces for the matrix B .

The row space of B consists of the three nonzero vectors in the RREF of B :

$$\mathcal{R}(B) = \text{span}((1, 0, -2, 0), (0, 1, 1, 0), (0, 0, 0, 1)).$$

The column space of B is spanned by the first, second and fourth columns of B since these columns contain the leading ones in the RREF of B :

$$\mathcal{C}(B) = \text{span}\left(\begin{bmatrix} 1 \\ 2 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 3 \\ 7 \\ 5 \\ 2 \end{bmatrix}, \begin{bmatrix} 4 \\ 9 \\ 1 \\ 8 \end{bmatrix}\right).$$

The third column lacks a leading one, so it corresponds to a free variable $x_3 = t$, $t \in \mathbb{R}$. The null space of B is the set of vectors $(x_1, x_2, t, x_4)^T$ such that:

$$\begin{bmatrix} 1 & 0 & -2 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ t \\ x_4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad \Rightarrow \quad \begin{array}{rcl} 1x_1 - 2t & = & 0 \\ 1x_2 + 1t & = & 0 \\ x_4 & = & 0 \\ 0 & = & 0. \end{array}$$

We find the values of x_1 , x_2 , and x_4 in terms of t and obtain

$$\mathcal{N}(B) = \left\{ \begin{bmatrix} 2t \\ -t \\ t \\ 0 \end{bmatrix}, t \in \mathbb{R} \right\} = \text{span}\left(\begin{bmatrix} 2 \\ -1 \\ 1 \\ 0 \end{bmatrix}\right).$$

Discussion

Dimensions

For an $m \times n$ matrix $M \in \mathbb{R}^{m \times n}$ the row space and the null space consist of vectors with n components, while the column space and the left null space consist of vectors with m components.

Don't confuse the number of components of vectors in a vector space with the *dimension* of the space. Suppose we're given a matrix $M \in \mathbb{R}^{5 \times 10}$ with five rows and 10 columns, and the RREF of M contains three pivots. We say the *rank* of the matrix is 3, which means the row space of M is three-dimensional. A basis for the row space of M contains three vectors, each vector having 10 components. The null space of the matrix is seven-dimensional ($10 - 3 = 7$) and consist of vectors with 10 components. The column space of the matrix is also three-dimensional ($\dim(\mathcal{R}(M)) = \dim(\mathcal{C}(M))$). A basis for the column space of M consists of three vectors with five components. The left null space of M is two-dimensional ($5 - 3 = 2$) and is spanned by vectors with five components.

Importance of bases

The procedures for identifying bases are somewhat technical and potentially boring, but they are of great practical importance. To illustrate the importance of a basis, consider a scenario in which you're given a description of the xy -plane P_{xy} as the span of *three* vectors:

$$P_{xy} = \text{span}((1, 0, 0), (0, 1, 0), (1, 1, 0)).$$

The above definition of P_{xy} says that any point $p \in P_{xy}$ can be written as a linear combination of the form

$$p = a(1, 0, 0) + b(0, 1, 0) + c(1, 1, 0),$$

for some coefficients a , b , and c . This representation of P_{xy} is misleading. It might make us think (erroneously) that P_{xy} is three-dimensional, since it takes three coefficients to describe points in P_{xy} .

Do we really need linear combinations of three vectors and three coefficients to describe any point in P_{xy} ? No, we don't. Linear combinations using two vectors are sufficient: $(1, 0, 0)$ and $(0, 1, 0)$, for example. The same point p described above can be written as:

$$p = \underbrace{(a + c)}_{\alpha}(1, 0, 0) + \underbrace{(b + c)}_{\beta}(0, 1, 0) = \alpha(1, 0, 0) + \beta(0, 1, 0).$$

Note the point is described in terms of *two* coefficients α and β . The vector $(1, 1, 0)$ is not *necessary* for the description of points in P_{xy} . The vector $(1, 1, 0)$ is redundant because it can be expressed in terms of the vectors $(1, 0, 0)$ and $(0, 1, 0)$. By getting rid of the redundant vector, we obtain a description of P_{xy} in terms of a basis:

$$P_{xy} = \text{span}((1, 0, 0), (0, 1, 0)).$$

Recall that a basis B for a space V must be made of linearly independent vectors and must span the space V . The set $\{(1, 0, 0), (0, 1, 0)\}$ is sufficient to represent any point in P_{xy} , and the vectors in this set are linearly independent. We can correctly conclude that the space P_{xy} is two-dimensional. If someone asks you, "How do you know that P_{xy} is two-dimensional?" you can say, "Because its basis contains two vectors."

Exercises

E4.16 Find the null spaces $\mathcal{N}(A)$ and $\mathcal{N}(B)$ of the following matrices:

$$A = \begin{bmatrix} 1 & 0 & -1 & 0 \\ 4 & 0 & 0 & -2 \\ 0 & 2 & -2 & -1 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 1 & -1 & -2 \\ -2 & 2 & 4 \\ 3 & -3 & -6 \end{bmatrix}.$$

E4.17 Consider the matrix

$$A = \begin{bmatrix} 1 & 3 & 3 & 3 \\ 2 & 6 & 7 & 6 \\ 3 & 9 & 9 & 10 \end{bmatrix}.$$

Find the RREF of A , and bases for $\mathcal{R}(A)$, $\mathcal{C}(A)$, and $\mathcal{N}(A)$.

4.6 Geometric problems

So far, we've defined all the important linear algebra concepts like vectors and matrices, and we've learned some useful computational techniques like the Gauss–Jordan elimination procedure. It's now time to apply what you've learned to solve geometric problems.

Points, lines, and planes can be difficult to understand and conceptualize. But now that you're armed with the tools of vectors, projections, and geometric intuition, you can solve all kinds of complicated geometric analysis problems—such as those waiting for you at the end of this paragraph. Remember to always sketch a diagram before you begin to write equations. Diagrams are great for visualizing and determining the steps you'll need to solve each problem.

P4.1 Find the intersections of these pairs of lines: **a)** $\ell_1: 2x + y = 4$ and $\ell_2: 3x - 2y = -1$, **b)** $\ell_1: y + x = 2$ and $\ell_2: 2x + 2y = 4$, **c)** $\ell_1: y + x = 2$ and $\ell_2: y - x = 0$.

P4.2 Find the lines of intersection between these pairs of planes: **a)** $P_1: 3x - 2y - z = 2$ and $P_2: x + 2y + z = 0$, **b)** $P_3: 2x + y - z = 0$ and $P_4: x + 2y + z = 3$.

P4.3 Find whether the planes are parallel, perpendicular, or neither: **a)** $P_1: x - y - z = 0$ and $P_2: 2x - 2y - 2z = 4$, **b)** $P_3: 3x + 2y = 1$ and $P_4: y - z = 0$, **c)** $P_5: x - 2y + z = 5$ and $P_6: x + y + z = 3$.

P4.4 Find the distance from the point $r = (2, 3, 5)$ to the plane P defined by the equation $2x + y - 2z = 0$.

P4.5 Find the closest distance between $p = (5, 3, 5)$ and $Q: 2x + y - 2z = 1$.
Hint: Consider an arbitrary point in the plane Q , such as $q = (0, 1, 0)$.

P4.6 Find the distance between the points **a)** $p = (4, 7, 3)$ and $q = (1, 1, 1)$, **b)** $m = (4, -2, 0)$ and $n = (0, 1, 0)$, **c)** $r = (1, 0, 1)$ and $s = (-1, 1, -1)$, **d)** $i = (2, 1, 2)$ and $j = (1, -2, -1)$.

P4.7 Find the general equation of the plane that passes through the points $q = (1, 3, 0)$, $r = (0, 2, 1)$, and $s = (1, 1, 1)$.

P4.8 Find the symmetric equation of the line ℓ described by the equations

$$x = 2t - 3, \quad y = -4t + 1, \quad z = -t.$$

P4.9 Define the line ℓ_1 to be the intersection of the planes $x + 2y + z = 1$ and $2x - y - z = 2$. Define ℓ_2 to be the line with parametric equation $x = 1 + 2t$, $y = -2 + t$, $z = -1 - t$. Find the equation of the plane that contains the line ℓ_1 and is parallel to the line ℓ_2 . See Figure 4.12.




Figure 4.12: The line ℓ_1 is the intersection of the two planes. In P4.9 we want to find the plane that contains ℓ_1 and doesn't intersect ℓ_2 .

P4.10 Given two vectors $\vec{u} = (2, 1, -1)$ and $\vec{v} = (1, 1, 1)$, find the projection of $\vec{v}$ onto $\vec{u}$, and the projection of $\vec{u}$ onto $\vec{v}$.

P4.11 Find a projection of $\vec{v} = (3, 4, 1)$ onto the plane $P: 2x - y + 4z = 0$.

P4.12 Find the component of the vector $\vec{u} = (-2, 1, 1)$ that is perpendicular to the plane P , and that contains the points $m = (2, 4, 1)$, $s = (6, 4, -2)$, and $r = (6, 5, -2)$.

P4.13 Find the distance between the line $\ell: \{x = 1 + 2t, y = -3 + t, z = 2\}$ and the plane $P: -x + 2y + 2z = 4$.

P4.14 Find the coordinates of the vector $\vec{v} = 9\hat{i} + 5\hat{j} + 4\hat{k} = (9, 5, 4)_{B_s}$ with respect to the basis $W = \{\vec{w}_1, \vec{w}_2, \vec{w}_3\}$, which consists of the vectors $\vec{w}_1 = (0, 0, 2)_{B_s}$, $\vec{w}_2 = (0, 5, 0)_{B_s}$, and $\vec{w}_3 = (3, 0, 0)_{B_s}$.

P4.15 Find the change-of-basis matrix ${}_V[\mathbb{1}]_U$ that transforms vectors expressed in the basis $U = \{\vec{u}_1 = (1, 0, 0), \vec{u}_2 = (0, 1, 1), \vec{u}_3 = (0, 1, -1)\}$ to vectors in the basis $V = \{\vec{v}_1 = (1, 1, 1), \vec{v}_2 = (1, 1, 0), \vec{v}_3 = (1, 0, 1)\}$.

Hint: Start by computing the change-of-basis matrices to the standard basis: ${}_{B_s}[\mathbb{1}]_U$ and ${}_{B_s}[\mathbb{1}]_V$. Then combine the matrices to obtain ${}_V[\mathbb{1}]_U$.

P4.16 An $m \times n$ matrix A is upper triangular if all entries lying below the main diagonal are zero; that is, if $A_{ij} = 0$ whenever $i > j$. Prove that upper triangular matrices form a subspace of $\mathbb{R}^{m \times n}$.

Hint: Is the set closed under addition and scaling? Does it contain zero?

P4.17 Consider the vector space of three-dimensional vectors $V = \mathbb{R}^3$. Which of the following sets are subspaces of V ?

- a) $W_1 = \{(v_x, v_y, v_z) \in V \mid v_x + v_y = 0\}$
- b) $W_2 = \{(v_x, v_y, v_z) \in V \mid v_y v_z = 0\}$

- c) $W_3 = \{(v_x, v_y, v_z) \in V \mid v_x = v_y = v_z\}$
 d) $W_4 = \{(v_x, v_y, v_z) \in V \mid v_x \geq 0\}$
 e) $W_5 = \{(v_x, v_y, v_z) \in V \mid v_x + v_z = 3\}$

Hint: To form a subspace, a set must be closed under addition, closed under scalar multiplication, and contain the zero element.

P4.18 Find 2×3 matrices A , B , and C that satisfy the conditions.

- a) $\mathcal{R}(A) = \text{span}((1, 2, 3))$ and $\mathcal{C}(A) = \text{span}((1, 2)^T)$
 b) $\mathcal{N}(B) = \text{span}((1, 1, 1)^T)$
 c) $\mathcal{N}(C^T) = \text{span}((1, 3))$

P4.19 Consider the matrix $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ where a , b , and c are fixed constants and $a \neq 0$. Find the value of d so that A will have rank one.

In this chapter we studied geometric objects like points, lines, and planes, and we derived various distance and projection formulas. We also introduced the notion of a vector and studied the vector spaces associated with matrices. With so much new material, it's easy to miss the important new definitions found throughout this chapter. Here's a little reminder:

- The *span* of a set of vectors is the vector space that can be constructed from linear combinations of the vectors:

$$\text{span}(\vec{v}_1, \dots, \vec{v}_n) = \{\alpha_1 \vec{v}_1 + \dots + \alpha_n \vec{v}_n, \forall \alpha_i \in \mathbb{R}\}.$$

- A set of n vectors $\{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n\}$ is *linearly independent* if the equation $\alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \dots + \alpha_n \vec{v}_n = \vec{0}$ has only trivial solution $\alpha_i = 0$.
- A *basis* for a d -dimensional vector space consists of any set of d linearly independent vectors.

The problems below will ensure that you've internalized these important concepts, and that you feel comfortable using the above definitions in proofs.

P4.20 Describe geometrically the subspaces of $\mathbb{R}^3$ spanned by the following sets of vectors.

- a) $\{(1, 0, 0), (2, 0, 0)\}$ b) $\{(1, 0, 0), (0, 1, 0)\}$ c) $\{(1, 0, 0), (0, 1, 0), (1, 1, 0)\}$
 d) $\{(1, 0, 0), (0, 1, 0), (1, 1, 1)\}$ e) $\{(0, 1, 1), (0, 1, 2), (0, 1, 3)\}$

P4.21 Suppose the set of vectors $\{\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4\}$ spans the vector space V . Prove that $\{\vec{v}_1, \vec{v}_2 - \vec{v}_1, \vec{v}_3 - \vec{v}_2, \vec{v}_4 - \vec{v}_3\}$ also spans V .

P4.22 A given set of m vectors is known to span the vector space V . Can you conclude that V is m -dimensional? Explain why or why not.

P4.23 Consider the vector space V and the vectors $\vec{w}_1 \in V$ and $\vec{w}_2 \in V$. Prove that $\text{span}(\vec{w}_1, \vec{w}_2)$ is a subspace of V .

P4.24 Suppose the set of vectors $\{\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4\}$ is linearly independent. Prove that $\{\vec{v}_1, \vec{v}_2 - \vec{v}_1, \vec{v}_3 - \vec{v}_2, \vec{v}_4 - \vec{v}_3\}$ is also a linearly independent set.

P4.25 Suppose the set $\{\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4\}$ forms a linearly independent set. Show that $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ is also a linearly independent set.

Hint: Use a proof by contradiction.

P4.26 Suppose $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ is a linearly independent set, and the vector $\vec{v}_4 \notin \text{span}(\vec{v}_1, \vec{v}_2, \vec{v}_3)$. Show that $\{\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4\}$ is linearly independent.

Hint: Use a proof by contradiction.

P4.27 Suppose that $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ is a linearly independent set. Define the vectors $\vec{w}_1 = \vec{v}_3$, $\vec{w}_2 = \vec{v}_2 + \vec{v}_3$, and $\vec{w}_3 = \vec{v}_1 + \vec{v}_2 + \vec{v}_3$. Prove $\{\vec{w}_1, \vec{w}_2, \vec{w}_3\}$ is a linearly independent set.

P4.28 Let $\vec{u}$ and $\vec{v}$ be distinct vectors from the vector space V , and assume that $\{\vec{u}, \vec{v}\}$ is a basis for V . Show that $\{\vec{u} + \vec{v}, a\vec{u}\}$ and $\{a\vec{u}, b\vec{v}\}$ are also bases for V , for any choice of nonzero constants a and b .

P4.29 Consider the vectors $\vec{v}_1 = (1, 2, 3)$ and $\vec{v}_2 = (1, 2, 4)$. Find a vector $\vec{v}_3$ such that $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ is a basis for $\mathbb{R}^3$. Also, find another vector $\vec{v}_4$ such that $\{\vec{v}_1, \vec{v}_2, \vec{v}_4\}$ is not a basis for $\mathbb{R}^3$.

P4.30 Find a vector that is perpendicular to the vectors $\vec{v}_1 = (1, 2, 3, 0)$, $\vec{v}_2 = (0, 1, 1, 1)$, and $\vec{v}_3 = (1, 0, 2, 1)$.

Chapter 5

Linear transformations

Linear transformations are a central idea of linear algebra—they form the cornerstone that connects all the seemingly unrelated concepts we’ve studied so far. We previously introduced linear transformations, informally describing them as “vector functions.” In this chapter, we’ll formally define linear transformations, describe their properties, and discuss their applications.

In Section 5.2, we’ll learn how matrices can be used to *represent* linear transformations. We’ll show the matrix representations of important types of linear transformations like projections, reflections, and rotations. Section 5.3 discusses the relation between bases and matrix representations. We’ll learn how the bases chosen for the input and output spaces determine the entries of matrix representations. A single linear transformation can correspond to many different matrix representations, depending on the choice of bases for the input and output spaces.

Section 5.4 discusses and characterizes the class of *invertible linear transformations*. This section serves to connect several topics we covered previously: linear transformations, matrix representations, and the fundamental subspaces of matrices.

5.1 Linear transformations

Linear transformations take vectors as inputs and produce vectors as outputs. A transformation T that takes n -dimensional vectors as inputs and produces m -dimensional vectors as outputs is denoted $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$.

The class of linear transformations includes most of the useful transformations of analytical geometry: stretchings, projections, reflections, rotations, and combinations of these. Since linear trans-

formations describe and model many real-world phenomena in physics, chemistry, biology, and computer science, learning the theory behind them is worthwhile.

Concepts

Linear transformations are mappings between *vector inputs* and *vector outputs*. The following concepts describe the input and output spaces:

- V : the input vector space of T
- W : the output vector space of T
- $\dim(U)$: the dimension of the vector space U
- $T : V \rightarrow W$: a linear transformation that takes vectors $\vec{v} \in V$ as inputs and produces vectors $\vec{w} \in W$ as outputs. The notation $T(\vec{v}) = \vec{w}$ describes T acting on $\vec{v}$ to produce the output $\vec{w}$.




Figure 5.1: An illustration of the linear transformation $T : V \rightarrow W$.

- $\text{Im}(T)$: the *image space* of the linear transformation T is the set of vectors that T can output for some input $\vec{v} \in V$. The mathematical definition of the image space is

$$\text{Im}(T) \stackrel{\text{def}}{=} \{\vec{w} \in W \mid \vec{w} = T(\vec{v}), \text{ for some } \vec{v} \in V\} \subseteq W.$$

The image space is the vector equivalent of the *image set* of a single-variable function $\text{Im}(f) \stackrel{\text{def}}{=} \{y \in \mathbb{R} \mid y = f(x), \forall x \in \mathbb{R}\}$.

- $\text{Ker}(T)$: the *kernel* of the linear transformation T ; the set of vectors mapped to the zero vector by T . The mathematical definition of the kernel is

$$\text{Ker}(T) \stackrel{\text{def}}{=} \{\vec{v} \in V \mid T(\vec{v}) = \vec{0}\} \subseteq V.$$

The kernel of a linear transformation is the vector equivalent of the roots of a function: $\{x \in \mathbb{R} \mid f(x) = 0\}$.




Figure 5.2: Two key properties of a linear transformation $T : V \rightarrow W$; its kernel $\text{Ker}(T) \subseteq V$, and its image space $\text{Im}(T) \subseteq W$.

Example The linear transformation $T : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ is defined by the equation $T(x, y) = (x, y, x + y)$. Applying T to the input vector $(1, 0)$ produces the output vector $(1, 0, 1 + 0) = (1, 0, 1)$. Applying T to the input vector $(3, 4)$ produces the output vector $(3, 4, 7)$.

The kernel of T contains only the zero vector $\text{Ker}(T) = \{\vec{0}\}$. The image space of T is a two-dimensional subspace of the output space $\mathbb{R}^3$, namely $\text{Im}(T) = \text{span}((1, 0, 1), (0, 1, 1)) \subseteq \mathbb{R}^3$.

Matrix representations

Given bases for the input and output spaces of a linear transformation T , the transformation's action on vectors can be represented as a matrix-vector product:

- $B_V = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$: a basis for the input vector space V
- $B_W = \{\vec{b}_1, \vec{b}_2, \dots, \vec{b}_m\}$: a basis for the output vector space W
- $M_T \in \mathbb{R}^{m \times n}$: a matrix representation of the linear transformation T :

$$\vec{w} = T(\vec{v}) \quad \Leftrightarrow \quad \vec{w} = M_T \vec{v}.$$

To be precise, we denote the matrix representation as ${}_{B_W}[M_T]_{B_V}$ to show it depends on the input and output bases.

- $\mathcal{C}(M_T)$: the *column space* of the matrix M_T
- $\mathcal{R}(M_T)$: the *row space* of the matrix M_T
- $\mathcal{N}(M_T)$: the *null space* the matrix M_T

Properties of linear transformations

We'll start with the feature of linear transformations that makes them suitable for modelling a wide range of phenomena in science, engineering, business, and computing.

Linearity

The fundamental property of linear transformations is—you guessed it—their *linearity*. If $\vec{v}_1$ and $\vec{v}_2$ are two input vectors and α and β are two constants, then

$$T(\alpha\vec{v}_1 + \beta\vec{v}_2) = \alpha T(\vec{v}_1) + \beta T(\vec{v}_2) = \alpha\vec{w}_1 + \beta\vec{w}_2,$$

where $\vec{w}_1 = T(\vec{v}_1)$ and $\vec{w}_2 = T(\vec{v}_2)$.




Figure 5.3: A linear transformation T maps the linear combination of inputs $\frac{1}{4}\vec{v}_1 + \frac{3}{4}\vec{v}_2$ to the linear combination of outputs $\frac{1}{4}\vec{w}_1 + \frac{3}{4}\vec{w}_2$.

Linear transformations map any linear combination of inputs to the same linear combination of outputs. If you know the outputs of T for the inputs $\vec{v}_1$ and $\vec{v}_2$, you can deduce the output T for any linear combination of the vectors $\vec{v}_1$ and $\vec{v}_2$ by computing the appropriate linear combination of the outputs $T(\vec{v}_1)$ and $T(\vec{v}_2)$. This is perhaps the most important idea in linear algebra: it's the *linear* that we refer to when we talk about *linear algebra*. Linear algebra is not about lines, but about mathematical transformations that map linear combinations of inputs to the same linear combinations of outputs.

In this chapter, we'll study various aspects and properties of linear transformations, the abstract objects that map input vectors to output vectors. The fact that linear transformations map linear combinations of inputs to corresponding linear combinations of outputs will be of central importance in many calculations and proofs. Make a good note and store a mental image of the example shown in Figure 5.3.

Linear transformations as black boxes

Suppose someone gives you a black box that implements the linear transformation T . While you can't look inside the box to see how T acts, you can *probe* the transformation by choosing various input vectors and observing what comes out.

Assume the linear transformation T is of the form $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$. By probing this transformation with n vectors of a basis for the input space and observing the outputs, you can characterize the transformation T completely.

To see why this is true, consider a basis $\{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ for the n -dimensional input space $V = \mathbb{R}^n$. To characterize T , input each of the n basis vectors $\vec{e}_i$ into the black box and record the $T(\vec{e}_i)$ that comes out. Any input vector $\vec{v}$ can be written as a linear combination of the basis vectors:

$$\vec{v} = v_1\vec{e}_1 + v_2\vec{e}_2 + \cdots + v_n\vec{e}_n.$$

Using these observations and the linearity of T , we can predict the output of T for this vector:

$$T(\vec{v}) = v_1T(\vec{e}_1) + v_2T(\vec{e}_2) + \cdots + v_nT(\vec{e}_n).$$

This black box model of probing is used in many areas of science and is one of the most important ideas in linear algebra. The transformation T could be the description of a chemical process, an electrical circuit, or some phenomenon in biology. As long as we know that T is (or can be approximated by) a linear transformation, we can describe it completely by probing it with a small number of inputs. This is in contrast to characterizing nonlinear transformations, which correspond to arbitrarily complex input-output relationships and require significantly more probing.

Input and output spaces

Consider the linear transformation T from n -vectors to m -vectors:

$$T: \mathbb{R}^n \rightarrow \mathbb{R}^m.$$

The *input space* of the linear transformation T is $\mathbb{R}^n$ and its *output space* is $\mathbb{R}^m$. The output space is also called the *target space* and it is similar to the notion of the *target set* for functions. The input space is identical to the *domain* of the linear transformation, since $T(\vec{v})$ is defined for all inputs $\vec{v} \in \mathbb{R}^n$.

The *image space* $\text{Im}(T)$ consists of all possible outputs of the transformation T . The image space of a linear transformation is a subset of its output space, $\text{Im}(T) \subseteq \mathbb{R}^m$. A linear transformation T whose

image space is equal to its output space ($\text{Im}(T) = \mathbb{R}^m$) is called *surjective* or *onto*.

The *kernel* of T is the subspace of the domain $\mathbb{R}^n$ that is mapped to the zero vector by T : $\text{Ker}(T) \stackrel{\text{def}}{=} \{\vec{v} \in \mathbb{R}^n \mid T(\vec{v}) = \vec{0}\}$. A linear transformation with an empty kernel $\text{Ker}(T) = \{\vec{0}\}$ is called *injective*. Injective transformations map different inputs to different outputs.

If a linear transformation T is both injective and surjective, it is called *bijective*. In this case, T is a *one-to-one correspondence* between the input vector space and the output vector space.

Note the terminology used to characterize linear transformations (injective, surjective, and bijective) is the same as the terminology used to characterize functions in Section 1.8. Indeed, we can use the same terminology since linear transformations are functions. The concepts of image space and kernel are illustrated in Figure 5.4.




Figure 5.4: Pictorial representations of the image space $\text{Im}(T)$ and the kernel $\text{Ker}(T)$ of a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$. The image space is the set of all possible outputs of T . The kernel of T is the set of inputs that T maps to the zero vector.

Observation The dimensions of the input space and the output space of a bijective linear transformation must be the same. Indeed, if $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is bijective, then it is both injective and surjective. Since T is surjective, the input space must be at least as large as the output space; $n \geq m$. Since T is injective, the output space must be larger or equal to the input space; $m \geq n$. Combining these observations, we find that if $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is bijective then $m = n$.

Example 2 Consider the linear transformation $T : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ defined by the equation $T(x, y, z) = (x, z)$. Find the kernel and the image space of T . Is T injective? Is T surjective?

The action of T is to delete the y -components of inputs. Any vector that has only a y -component will be sent to the zero vector. We have $\text{Ker}(T) = \text{span}((0, 1, 0))$. The image space is $\text{Im}(T) = \mathbb{R}^2$. The transformation T is not injective. As an explicit example proving T is not injective, observe that $T(0, 1, 0) = T(0, 2, 0)$ but $(0, 1, 0) \neq (0, 2, 0)$. Since $\text{Im}(T)$ is equal to the output space $\mathbb{R}^2$, T is surjective.

Linear transformations as matrix multiplications

An important relationship exists between linear transformations and matrices. If you fix a basis for the input vector space and a basis for the output vector space, a linear transformation $T(\vec{v}) = \vec{w}$ can be represented as matrix multiplication $M_T \vec{v} = \vec{w}$ for some matrix M_T :

$$\vec{w} = T(\vec{v}) \quad \Leftrightarrow \quad \vec{w} = M_T \vec{v}.$$

Using this equivalence, we can re-interpret several properties of matrices as properties of linear transformations. The equivalence is useful in the other direction too, since it allows us to use the language of linear transformations to talk about the properties of matrices.

The idea of representing the action of a linear transformation as a matrix multiplication is extremely important since it transforms the abstract, mathematical description of the linear transformation T into a concrete, computational one: “take the input vector $\vec{v}$ and multiply it from the right by the matrix M_T .” Borrowing an idea from the field of software engineering, we can understand the relationship between a linear transformation and its matrix representation as analogous to the relationship between the abstract software specification of a computer program, and the program’s concrete implementation as code that runs on a computer.

Example 3 We’ll now illustrate the “linear transformation $\Leftrightarrow$ matrix-product” equivalence with an example. Define $\Pi_{P_{xy}}$ to be the *orthogonal projection* onto the xy -plane P_{xy} . In words, the action of this projection is to zero-out the z -component of input vectors. The matrix that corresponds to this projection is

$$T\left(\begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix}\right) = \begin{bmatrix} v_x \\ v_y \\ 0 \end{bmatrix} \quad \Leftrightarrow \quad M_T \vec{v} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix} = \begin{bmatrix} v_x \\ v_y \\ 0 \end{bmatrix}.$$

Finding the matrix

In order to find the matrix representation of any linear transformation $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$, it is sufficient to probe T with the n vectors in the standard basis for $\mathbb{R}^n$:

$$\hat{e}_1 = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \hat{e}_2 = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix}, \quad \dots, \quad \hat{e}_n = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}.$$

To obtain M_T , we combine the outputs $T(\hat{e}_1), T(\hat{e}_2), \dots, T(\hat{e}_n)$ as the *columns* of a matrix:

$$M_T = \begin{bmatrix} | & | & & | \\ T(\hat{e}_1) & T(\hat{e}_2) & \cdots & T(\hat{e}_n) \\ | & | & & | \end{bmatrix}.$$

Observe that the matrix constructed in this way has the correct dimensions: $m \times n$. We have $M_T \in \mathbb{R}^{m \times n}$ since we used n “probe vectors,” and since the outputs of T are m -dimensional column vectors.

To help visualize the column structure of M_T , let’s analyze what happens when we compute the product $M_T \hat{e}_2$. The probe vector $\hat{e}_2 = (0, 1, 0, \dots, 0)^T$ “selects” only the second column from M_T ; thus we obtain the correct output: $M_T \hat{e}_2 = T(\hat{e}_2)$. Similarly, applying M_T to other basis vectors selects the other columns of M_T .

Any input vector can be written as a linear combination of the standard basis vectors $\vec{v} = v_1 \hat{e}_1 + v_2 \hat{e}_2 + \cdots + v_n \hat{e}_n$. Therefore, by linearity, we can compute the output $T(\vec{v})$ as follows:

$$\begin{aligned} T(\vec{v}) &= v_1 T(\hat{e}_1) + v_2 T(\hat{e}_2) + \cdots + v_n T(\hat{e}_n) \\ &= v_1 \begin{bmatrix} | \\ T(\hat{e}_1) \\ | \end{bmatrix} + v_2 \begin{bmatrix} | \\ T(\hat{e}_2) \\ | \end{bmatrix} + \cdots + v_n \begin{bmatrix} | \\ T(\hat{e}_n) \\ | \end{bmatrix} \\ &= \begin{bmatrix} | & | & & | \\ T(\hat{e}_1) & T(\hat{e}_2) & \cdots & T(\hat{e}_n) \\ | & | & & | \end{bmatrix} \begin{bmatrix} | \\ \vec{v} \\ | \end{bmatrix} \\ &= M_T [\vec{v}], \end{aligned}$$

where $[\vec{v}] = (v_1, v_2, \dots, v_n)^T$ is the coordinate vector of $\vec{v}$, represented as a column vector.

Input and output spaces

We can identify correspondences between the properties of a linear transformation $T : V \rightarrow W$ and the properties of a $\dim(W) \times \dim(V)$ matrix M_T that implements T .

The outputs of the linear transformation T consist of all possible linear combinations of the columns of the matrix M_T . Thus, the *image space* of the linear transformation T is equivalent to the *column space* of the matrix M_T :

$$\text{Im}(T) = \{\vec{w} \in W \mid \vec{w} = T(\vec{v}), \text{ for some } \vec{v} \in V\} = \mathcal{C}(M_T).$$

There is also an equivalence between the kernel of the linear transformation T and the null space of the matrix M_T :

$$\text{Ker}(T) = \{\vec{v} \in V \mid T(\vec{v}) = \vec{0}\} = \{\vec{v} \in V \mid M_T \vec{v} = \vec{0}\} = \mathcal{N}(M_T).$$

The null space of a matrix $\mathcal{N}(M_T)$ consists of all vectors that are orthogonal to the rows of the matrix M_T . The vectors in the null space of M_T have a zero dot product with each of the rows of M_T . This orthogonality can also be phrased in the opposite direction: any vector in the row space $\mathcal{R}(M_T)$ is orthogonal to the null space $\mathcal{N}(M_T)$.

These observations allow us to decompose the domain of the transformation T as the *orthogonal sum* of the row space and the null space of a matrix M_T :

$$V = \mathcal{R}(M_T) \oplus \mathcal{N}(M_T).$$

This split implies the *conservation of dimensions* formula:

$$\dim(V) = \dim(\mathcal{R}(M_T)) + \dim(\mathcal{N}(M_T)),$$

which describes how the sum of the dimensions of the row space and the null space of a matrix M_T is equal to the total dimensions of the input space.

We can summarize everything we know about the input-output relationship of the linear transformation T as follows:

$$T: \mathcal{R}(M_T) \rightarrow \mathcal{C}(M_T), \quad T: \mathcal{N}(M_T) \rightarrow \{\vec{0}\}.$$

Input vectors $\vec{v} \in \mathcal{R}(M_T)$ are mapped to output vectors $\vec{w} \in \mathcal{C}(M_T)$ in a one-to-one correspondence. Input vectors $\vec{v} \in \mathcal{N}(M_T)$ are mapped to the zero vector $\vec{0} \in W$.

Composition of linear transformations

The consecutive application of two linear transformations T and S on an input vector $\vec{v}$ corresponds to the following matrix product:

$$S \circ T(\vec{v}) = S(T(\vec{v})) = M_S M_T \vec{v}.$$

The matrix M_T acts on the vector first, followed by the matrix M_S .

For this composition to be well-defined, the dimension of the output space of T must be the same as the dimension of the input space of S . This requirement corresponds to the condition that the rows of M_S and the columns of M_T must have the same dimension for the product $M_S M_T$ to exist.

Importance of the choice of bases

Above, we assumed the standard basis was used both for inputs and outputs of the linear transformation. Thus, we obtained the entries in the matrix M_T *with respect to* the standard basis.

In particular, we assumed that the outputs of T were given as column vectors in terms of the standard basis for $W = \mathbb{R}^m$. If the outputs were given in some other basis B_W , the entries of the matrix M_T would have been *with respect to* B_W .

Due to the dependence of matrix entries on the basis used, **a linear transformation does not correspond to a unique matrix**. Indeed, the same linear transformation T will correspond to different matrices when different bases are used. We say the linear transformation T corresponds to a matrix M for a *given* choice of input and output bases. We write ${}_{B_W}[M_T]_{B_V}$ to show the entries of the matrix M_T depend on the choice of left and right bases. Recall we can also use the basis-in-a-subscript notation for vectors. For example, writing $(v_x, v_y, v_z)_{B_s}$ shows the components v_x , v_y , and v_z are expressed in terms of the standard basis $B_s = \{\hat{i}, \hat{j}, \hat{k}\}$.

The choice of basis is an important technical detail: be aware of it, but don't worry about it too much. Unless otherwise specified, assume the standard basis is used for specifying vectors and matrices. When you see the product $A\vec{v}$, it means ${}_{B_s}[A]_{B_s} [\vec{v}]_{B_s}$. The only time you really need to pay attention to the choice of bases is when performing *change-of-basis* transformations, which we'll discuss in Section 5.3.

Invertible transformations

We'll now revisit the properties of invertible matrices and connect them to the notion of invertible transformations. Think of multiplication by a matrix M as "doing" something to vectors, and multiplication by M^{-1} as doing the opposite thing, restoring the original vector:

$$M^{-1}M\vec{v} = \vec{v}.$$

For example, the matrix

$$M = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}$$

corresponds to a stretching of space by a factor of 2 in the x -direction, while the y -direction remains unchanged. The inverse transformation corresponds to a shrinkage by a factor of 2 in the x -direction:

$$M^{-1} = \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{bmatrix}.$$

In general, it's hard to see exactly what the matrix M does, since it performs some arbitrary linear combination of the components of the input vector.

If M is an invertible matrix, we can start from any output vector $\vec{w} = M\vec{v}$, and go back to find the input $\vec{v}$ that produced the output $\vec{w}$. We do this by multiplying $\vec{w}$ by the inverse: $M^{-1}\vec{w} = M^{-1}M\vec{v} = \vec{v}$.

A linear transformation T is *invertible* if there exists an inverse transformation T^{-1} such that $T^{-1}(T(\vec{v})) = \vec{v}$. By the correspondence $\vec{w} = T(\vec{v}) \Leftrightarrow \vec{w} = M_T\vec{v}$, we can identify the class of invertible linear transformations with the class of invertible matrices. A linear transformation T is invertible if and only if its matrix representation M_T is an invertible matrix.

Affine transformations

An *affine transformation* is a function $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ that is the combination of a linear transformation T followed by a *translation* by a fixed vector $\vec{b}$:

$$\vec{y} = A(\vec{x}) = T(\vec{x}) + \vec{b}.$$

By the $T \Leftrightarrow M_T$ equivalence, we can write the formula for an affine transformation as

$$\vec{y} = A(\vec{x}) = M_T\vec{x} + \vec{b}.$$

To obtain the output $\vec{y}$, apply the linear transformation T (the matrix-vector product $M_T\vec{x}$), then add the vector $\vec{b}$. This is the vector generalization of a single-variable *affine function* $y = f(x) = mx + b$.

Discussion

The most general linear transformation

In this section we learned that a linear transformation *can* be represented as matrix multiplication. Are there other ways to represent linear transformations? To study this question, let's analyze the most general form a linear transformation $T: V \rightarrow W$ can take. We'll use $V = \mathbb{R}^3$ and $W = \mathbb{R}^2$ to keep things simple.

First consider the component w_1 of the output vector $\vec{w} = T(\vec{v})$ when the input vector is $\vec{v} \in \mathbb{R}^3$. The fact that T is linear means that w_1 can be an arbitrary mixture of the input vector components v_1, v_2, v_3 :

$$w_1 = \alpha_1 v_1 + \alpha_2 v_2 + \alpha_3 v_3.$$

Similarly, the component w_2 must be some arbitrary linear combination of the input components $w_2 = \beta_1 v_1 + \beta_2 v_2 + \beta_3 v_3$. Thus, the

most general linear transformation $T: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ can be written as

$$\begin{aligned}w_1 &= \alpha_1 v_1 + \alpha_2 v_2 + \alpha_3 v_3, \\w_2 &= \beta_1 v_1 + \beta_2 v_2 + \beta_3 v_3.\end{aligned}$$

This is precisely the kind of expression that can be obtained as a matrix product:

$$T(\vec{v}) = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} \alpha_1 & \alpha_2 & \alpha_3 \\ \beta_1 & \beta_2 & \beta_3 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = M_T \vec{v}.$$

Indeed, matrix multiplication is defined as rows-times-columns because it allows us to easily describe linear transformations.

Links

[Examples of linear transformations from Wikibooks]

http://wikibooks.org/wiki/Linear_Algebra/Linear_Transformations

[Linear transformations and matrices explained by 3Blue1Brown]

<https://youtube.com/watch?v=kYB8IZa5AuE>

[Three-dimensional transformations explained by 3Blue1Brown]

<https://youtube.com/watch?v=rHLEWRxRGiM>

[Nonsquare matrices explained by 3Blue1Brown]

https://youtube.com/watch?v=v8VSDg_WQ1A

Exercises

E5.1 Determine whether the following transformations are linear.

a) $T_1(x, y, z) = (x + 2y + 3z)$	b) $T_2(x, y) = (1, x, y)$
c) $T_3(x, y) = (x + 2y, y + x^2)$	d) $T_4(x, y, z) = (z, y, x)$

If the transformation is linear, find its matrix representation. If the transformation is nonlinear, show an example where linearity fails.

E5.2 Consider the transformation $T(x, y, z) = (y + z, x + z, x + y)$. Find the input space, the output space, the kernel, and the image of the linear transformation T . Is T injective, surjective, or bijective?

E5.3 Consider the matrix $M = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ and the linear transformations T_M and T_{M^T} defined through left and right multiplication by M : $T_M(\vec{v}) = M\vec{v}$ and $T_{M^T}(\vec{u}) = \vec{u}^T M$. Find the input, output, kernel, and image spaces of the linear transformations T_M and T_{M^T} .

E5.4 What linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ takes the vector $(1, 0)$ to the vector $(1, 2, 3)$, and the vector $(0, 1)$ to the vector $(4, 5, 6)$? Express your answer as a function $T(x, y) = \dots$ and as a matrix M_T .

5.2 Finding matrix representations

Every linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ can be represented as a matrix $M_T \in \mathbb{R}^{m \times n}$. Suppose you're given the following description of a linear transformation: “ T is the counterclockwise rotation of all points in the xy -plane by 30° ,” and you want to find the matrix M_T that corresponds to this transformation.

Do you know how to find the matrix representation of T ? This section describes a simple and intuitive probing procedure for finding matrix representations. Don't worry; no alien technology is involved, and we won't be probing any humans—only linear transformations! As you read, try to bridge your understanding between the abstract, mathematical specification of a transformation $T(\vec{v})$ and its concrete implementation as a matrix-vector product $M_T \vec{v}$. We'll use the probing procedure to study various linear transformations and derive their matrix representations.

Once we find the matrix representation of a given transformation, we can efficiently apply that transformation to many vectors. This is exactly how computers carry out linear transformations. For example, a black-and-white image file can be represented as a long list that contains the coordinates of the image's black pixels: $\{\vec{x}_1, \vec{x}_2, \dots, \vec{x}_\ell\}$. The image is obtained by starting with a white background and drawing a black pixel in each of the locations $\vec{x}_i$ on the screen.¹ To rotate the image, we can process the list of pixels using the matrix-vector product $\vec{y}_i = M_T \vec{x}_i$, where M_T is the matrix representation of the desired rotation. The transformed list of pixels $\{\vec{y}_1, \vec{y}_2, \dots, \vec{y}_\ell\}$ corresponds to a rotated version of the image. This is essentially the effect of using the “rotate tool” in an image editing program—the computer multiplies the image by a rotation matrix.

Concepts

The previous section covered linear transformations and their matrix representations:

- $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$: a linear transformation that takes inputs in $\mathbb{R}^n$ and produces outputs in $\mathbb{R}^m$
- $M_T \in \mathbb{R}^{m \times n}$: the matrix representation of T

The action of the linear transformation T is equivalent to multiplication by the matrix M_T :

$$\vec{w} = T(\vec{v}) \quad \Leftrightarrow \quad \vec{w} = M_T \vec{v}.$$

¹Location on a computer screen is denoted using pixel coordinates $\vec{x}_i = (h_i, v_i)$. The number h_i describes a horizontal distance measured in pixels from the left edge of the image, and v_i measures the vertical distance from the top of the image.

Theory

To find the matrix representation of the transformation $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$, it is sufficient to probe T with the n vectors of the standard basis for the input space $\mathbb{R}^n$:

$$\hat{e}_1 = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \hat{e}_2 = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix}, \quad \dots, \quad \hat{e}_n = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}.$$

The matrix M_T that corresponds to the action of T on the standard basis is

$$M_T = \begin{bmatrix} | & | & & | \\ T(\vec{e}_1) & T(\vec{e}_2) & \cdots & T(\vec{e}_n) \\ | & | & & | \end{bmatrix}.$$

This is an $m \times n$ matrix whose columns are the outputs of T for the n probe inputs.

The remainder of this section illustrates the probing procedure for finding matrix representations of linear transformations.

Projections

We'll start with a class of linear transformations you're already familiar with: *projections*. I hope you still remember what you learned in Section 4.2 (page 219).

X projection

The projection onto the x -axis is denoted Π_x . The projection Π_x acts on any vector or point by leaving the x -component unchanged and setting the y -component to zero. The action of Π_x on two sample points is illustrated in Figure 5.5.

Let's analyze how the projection Π_x transforms the two vectors of the standard basis:

$$\Pi_x\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}\right) = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad \Pi_x\left(\begin{bmatrix} 0 \\ 1 \end{bmatrix}\right) = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

The action of Π_x on the basis $\hat{i}$ is to leave it unchanged. The action of Π_x on the basis $\hat{j}$ is to send it to the zero vector.

The matrix representation of Π_x is therefore given by:

$$M_{\Pi_x} = \left[\Pi_x\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}\right) \quad \Pi_x\left(\begin{bmatrix} 0 \\ 1 \end{bmatrix}\right) \right] = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}.$$




Figure 5.5: The projection Π_x takes all points to the x -axis.

Y projection

Similar to Π_x , Π_y is defined as the projection onto the y -axis. The action of Π_y on two sample points is illustrated in Figure 5.6. Can you guess what the matrix for the projection onto the y -axis will look like?




Figure 5.6: The projection Π_y takes all points to the y -axis.

Use the standard approach to find the matrix representation of Π_y :

$$M_{\Pi_y} = \left[\Pi_y \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \quad \Pi_y \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \right] = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

The first column of the matrix representation M_{Π_y} contains the zero vector since $\Pi_y(\hat{i}) = \vec{0}$. The second column of M_{Π_y} contains $\hat{j}$ since $\Pi_y(\hat{j}) = \hat{j}$.

We can verify that the matrices M_{Π_x} and M_{Π_y} do indeed select the appropriate components from a general input vector $\vec{v} = (v_x, v_y)^\top$:

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} = \begin{bmatrix} v_x \\ 0 \end{bmatrix}, \quad \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} = \begin{bmatrix} 0 \\ v_y \end{bmatrix}.$$

Projection onto a vector

Recall that the general formula for the projection of a vector $\vec{v}$ onto another vector $\vec{a}$ is obtained as:

$$\Pi_{\vec{a}}(\vec{v}) = \left(\frac{\vec{a} \cdot \vec{v}}{\|\vec{a}\|^2} \right) \vec{a}.$$

To find the matrix representation of a projection onto an arbitrary direction $\vec{a}$, we compute

$$M_{\Pi_{\vec{a}}} = \left[\Pi_{\vec{a}} \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \quad \Pi_{\vec{a}} \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \right].$$

Projection onto a plane

We can also compute the projection of the vector $\vec{v} \in \mathbb{R}^3$ onto the plane $P : \vec{n} \cdot \vec{x} = n_x x + n_y y + n_z z = 0$ as follows:

$$\Pi_P(\vec{v}) = \vec{v} - \Pi_{\vec{n}}(\vec{v}).$$

How should we interpret the above formula? First compute the part of the vector $\vec{v}$ that is perpendicular to the plane (in the $\vec{n}$ direction); then subtract this part from $\vec{v}$ to obtain the part that lies in the plane.

To obtain the matrix representation of Π_P , calculate what Π_P does to the vectors in the standard basis $\hat{i} = \hat{e}_1 = (1, 0, 0)^T$, $\hat{j} = \hat{e}_2 = (0, 1, 0)^T$, and $\hat{k} = \hat{e}_3 = (0, 0, 1)^T$.

Projections as outer products

We can obtain the projection matrix onto any unit vector by computing the *outer product* of the vector with itself. As an example, we'll find the matrix for the projection onto the x -axis $\Pi_x(\vec{v}) = (\vec{v} \cdot \hat{i})\hat{i}$. Recall the *inner product* (dot product) between two column vectors $\vec{u}$ and $\vec{v}$ is equivalent to the matrix product $\vec{u}^T \vec{v}$, while the *outer product* is given by the matrix product $\vec{u} \vec{v}^T$. The inner product is a product between a $1 \times n$ matrix and an $n \times 1$ matrix, whose result is a 1×1 matrix—a single number. The outer product corresponds to an $n \times 1$ matrix times a $1 \times n$ matrix, making the answer an $n \times n$ matrix. The projection matrix corresponding to Π_x is $M_{\Pi_x} = \hat{i} \hat{i}^T \in \mathbb{R}^{n \times n}$.

Where did that equation come from? To derive the equation, we use the commutative law of scalar multiplication $\alpha \vec{v} = \vec{v} \alpha$, rewrite the dot product as a matrix product $\vec{u} \cdot \vec{v} = \vec{u}^T \vec{v}$, and use the *associa-*

tive law of matrix multiplication $A(BC) = (AB)C$. Check it out:

$$\begin{aligned}\Pi_x(\vec{v}) &= (\hat{i} \cdot \vec{v}) \hat{i} = \hat{i}(\hat{i} \cdot \vec{v}) = \hat{i}(\hat{i}^T \vec{v}) = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \left(\begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} \right) \\ &= (\hat{i}^T) \vec{v} = \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \begin{bmatrix} 1 & 0 \end{bmatrix} \right) \begin{bmatrix} v_x \\ v_y \end{bmatrix} \\ &= (M) \vec{v} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} = \begin{bmatrix} v_x \\ 0 \end{bmatrix}.\end{aligned}$$

The outer product $M \stackrel{\text{def}}{=} \hat{i} \hat{i}^T$ corresponds to the projection matrix M_{Π_x} we're looking for.

More generally, we obtain the projection matrix onto a line with direction vector $\vec{a}$ by constructing the unit vector $\hat{a}$, and then calculating the outer product of $\hat{a}$ with itself:

$$\hat{a} = \frac{\vec{a}}{\|\vec{a}\|}, \quad M_{\Pi_{\vec{a}}} = \hat{a} \hat{a}^T.$$

Example Find the matrix representation $M_d \in \mathbb{R}^{2 \times 2}$ for the projection Π_d onto the 45° -diagonal line with equation $y = x$.

The line with equation $y = x$ corresponds to the parametric equation $\{(x, y) \in \mathbb{R}^2 \mid (x, y) = (0, 0) + t(1, 1), t \in \mathbb{R}\}$, so the direction vector for this line is $\vec{a} = (1, 1)^T$. We want to find the matrix that corresponds to the linear transformation $\Pi_d(\vec{v}) = \left(\frac{\vec{v} \cdot (1, 1)}{2} \right) (1, 1)^T$.

The projection matrix onto $\vec{a} = (1, 1)^T$ is computed using the outer product approach. First, compute a normalized direction vector $\hat{a} = \left(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}} \right)^T$; then compute M_d using the outer product:

$$M_d = \hat{a} \hat{a}^T = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}.$$

Usually, the idea of representing projections as outer products is not covered in a first linear algebra course, so don't worry about outer products appearing on the exam. The purpose of introducing you to the equivalence between projections onto $\hat{a}$ and the outer product $\hat{a} \hat{a}^T$ is to illuminate this interesting connection between vectors and matrices. This connection plays a fundamental role in quantum mechanics, where projections in different directions are frequently used.

If you're asked a matrix representation question on an exam, keep things simple and stick to the "probing with the standard basis" approach, which gives the same answer as the one computed using the

outer product:

$$M_d = \left[\Pi_d \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \quad \Pi_d \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \right] = \left[\left(\frac{\hat{i} \cdot \vec{d}}{\|\vec{d}\|^2} \right) \vec{d} \quad \left(\frac{\hat{j} \cdot \vec{d}}{\|\vec{d}\|^2} \right) \vec{d} \right] = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}.$$

Projections are idempotent

A projection matrix M_Π satisfies $M_\Pi M_\Pi = M_\Pi$. This is one of the defining properties of projections. The technical term for this is *idempotence*, meaning the operation can be applied multiple times without changing the result beyond the initial application.

Subspaces

A projection acts differently on different sets of input vectors. While some input vectors remain unchanged, other input vectors are “killed.” This is murder! Well, murder in a mathematical sense, which means multiplication by zero.

Let Π_S be the projection onto the space S , and $S^\perp$ be the *orthogonal complement* to S defined by $S^\perp \stackrel{\text{def}}{=} \{ \vec{w} \in \mathbb{R}^n \mid \vec{w} \cdot \vec{s} = 0, \forall \vec{s} \in S \}$. The action of Π_S is completely different for the vectors from S and $S^\perp$. All vectors $\vec{v}$ in S remain unchanged,

$$\Pi_S(\vec{v}) = \vec{v},$$

whereas vectors $\vec{w}$ in $S^\perp$ are killed,

$$\Pi_S(\vec{w}) = 0\vec{w} = \vec{0}.$$

The action of Π_S on any vector from $S^\perp$ is equivalent to multiplication by zero. This is why $S^\perp$ is called the *null space* of M_{Π_S} .

Reflections

We’ll now compute matrix representations for simple *reflection* transformations.

X reflection

Reflection through the x -axis leaves the x -component unchanged and flips the sign of the y -component. Figure 5.7 illustrates the effect of reflecting two points through the x -axis.

Using the usual probing procedure, we obtain the matrix that corresponds to this linear transformation:

$$M_{R_x} = \left[R_x \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \quad R_x \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \right] = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$




Figure 5.7: The reflection through the x -axis sends every point (x, y) in the plane to its mirror point $(x, -y)$ on the other side of the x -axis.

This matrix sends $(x, y)^T$ to $(x, -y)^T$ as expected for a reflection through the x -axis.

Y reflection

The matrix associated with R_y , the reflection through the y -axis, is given by:

$$M_{R_y} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}.$$

The numbers in the above matrix tell us to change the sign of the x -component of the input and leave its y -component unchanged. In other words, every point that starts to the left of the y -axis moves to the right of the y -axis, and every point that starts to the right of the y -axis moves to the left.




Figure 5.8: The reflection through the y -axis flips the x -component of every point in the plane.

Do you see the power and simplicity of the probing procedure for finding matrix representations? In the first column, enter what you want to happen to the $\hat{e}_1$ vector; in the second column, enter what you want to happen to the $\hat{e}_2$ vector, and voila!

Diagonal reflection

Suppose we want to find the formula for the reflection through the line $y = x$. We'll call this reflection R_d . This time, dear readers, it's up to you to draw the diagram. In words, R_d is the transformation that makes x and y "swap places."

Based on this notion of swapping places, the matrix for R_d is

$$M_{R_d} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Alternatively, the usual probing procedure leads us to the same result.

Now I must point out an important property common to all reflections. The effect of a reflection is described by one of two possible actions: some points remain unchanged by the reflection, while other points flip into their exact negatives. For example, the *invariant* points under R_y are the points that lie on the y -axis—that is, the multiples of $(0, 1)^T$. The points that become their *exact negatives* are those with only an x -component—the multiples of $(1, 0)^T$. The action of R_y on all other points can be obtained as a linear combination of the actions "do not change" and "multiply by -1 ." We'll continue this line of reasoning further at the end of this section, and again in Section 6.1.

Reflections through lines and planes

What about finding reflections through arbitrary lines and planes? We can leverage the formulas for projections onto lines and planes we saw earlier to obtain formulas for reflections through any line or plane that passes through the origin.

Consider the line with parametric equation $\ell : \{\vec{0} + t\vec{a}, t \in \mathbb{R}\}$. The formula for the reflection of any vector $\vec{v}$ through the line ℓ is

$$R_{\vec{a}}(\vec{v}) = 2\Pi_{\vec{a}}(\vec{v}) - \vec{v},$$

where $\Pi_{\vec{a}}(\vec{v}) = \frac{\vec{a} \cdot \vec{v}}{\|\vec{a}\|^2} \vec{a}$ is the projection of $\vec{v}$ onto the vector $\vec{a}$.

The vector diagram in Figure 5.9 illustrates the reflection formula. Annotating the figure with a pencil can help you visualize how the

formula works. Draw the vector $2\Pi_{\vec{a}}(\vec{v})$ followed by the vector $-\vec{v}$ and confirm that $2\Pi_{\vec{a}}(\vec{v}) - \vec{v}$ results in the vector $R_{\vec{a}}(\vec{v})$.




Figure 5.9: The reflection of the vector $\vec{v}$ through the line $\ell : \{\vec{0} + t\vec{a}, t \in \mathbb{R}\}$ is computed using the formula $R_{\vec{a}}(\vec{v}) = 2\Pi_{\vec{a}}(\vec{v}) - \vec{v}$.

We can derive two equivalent formulas for computing the reflection of the vector $\vec{v}$ through the plane P defined by the equation $\vec{n} \cdot \vec{x} = 0$:

$$R_P(\vec{v}) = 2\Pi_P(\vec{v}) - \vec{v} \quad \text{or} \quad R_P(\vec{v}) = \vec{v} - 2\Pi_{\vec{n}}(\vec{v}).$$

The first formula uses a reasoning similar to the formula for the reflection through a line. The second formula can be understood as computing the projection of $\vec{v}$ in the direction of the normal vector, $\Pi_{\vec{n}}(\vec{v})$, subtracting that vector once from $\vec{v}$ to reach a point in the plane, and subtracting it a second time to move to the point $R_P(\vec{v})$ on the other side of the plane. See Figure 5.10 for the illustration.




Figure 5.10: The reflection of the vector $\vec{v}$ through the plane P with normal vector $\vec{n}$ is computed by starting with $\vec{v}$ and subtracting twice the projection of $\vec{v}$ in the direction of the normal: $R_P(\vec{v}) = \vec{v} - 2\Pi_{\vec{n}}(\vec{v})$.

Rotations

We'll now find the matrix representations for *rotation* transformations. The counterclockwise rotation by the angle θ is denoted R_θ . Figure 5.11 illustrates the action of the rotation R_θ : the point A is rotated around the origin to become the point B .




Figure 5.11: The linear transformation R_θ rotates every point in the plane by the angle θ in the counterclockwise direction. Note the effect of R_θ on the basis vectors $(1,0)$ and $(0,1)$.

To find the matrix representation of R_θ , probe it with the standard basis as usual:

$$M_{R_\theta} = \left[R_\theta \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \quad R_\theta \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \right].$$

To compute the values in the first column, observe that R_θ rotates the vector $(1,0)^T = 1\angle 0$ to the vector $1\angle \theta = (\cos \theta, \sin \theta)^T$. The second input $\hat{e}_2 = (0,1)^T = 1\angle \frac{\pi}{2}$ is rotated to $1\angle (\frac{\pi}{2} + \theta) = (-\sin \theta, \cos \theta)^T$. Therefore, the matrix for R_θ is

$$M_{R_\theta} = \left[\begin{array}{c|c} | & | \\ 1\angle \theta & 1\angle (\frac{\pi}{2} + \theta) \\ | & | \end{array} \right] = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}.$$

Finding the matrix representation of a linear transformation is like a colouring-book activity for mathematicians. Filling in the columns is just like colouring inside the lines—nothing too complicated.

Inverses

Can you determine the inverse matrix of M_{R_θ} ? You could use the formula for finding the inverse of a 2×2 matrix, *or* you could use the $[A \mid I]$ -and-RREF algorithm for finding the inverse; but using either of these approaches would be *waaaaay* too much work. Try to guess the matrix representation of the inverse without doing any calculations. If R_θ rotates points by $+\theta$, can you tell me what the inverse operation does? I'll leave a blank line here to give you some time to think....

Think you have it? The inverse operation of R_θ is $R_{-\theta}$, a rotation by $-\theta$, which corresponds to the matrix

$$M_{R_{-\theta}} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}.$$

Recall that $\cos$ is an even function, so $\cos(-\theta) = \cos(\theta)$, while $\sin$ is an odd function, so $\sin(-\theta) = -\sin(\theta)$.

For any vector $\vec{v} \in \mathbb{R}^2$ we have $R_{-\theta}(R_\theta(\vec{v})) = \vec{v} = R_\theta(R_{-\theta}(\vec{v}))$; or in terms of matrices,

$$M_{R_{-\theta}} M_{R_\theta} = \mathbb{1} = M_{R_\theta} M_{R_{-\theta}}.$$

Cool, right? This is what *representation* really means: the abstract notion of composition of linear transformations is *represented* as a matrix product.

Here's another quiz question: what is the inverse operation of the reflection through the x -axis R_x ? The “undo” action for R_x is to apply R_x again. We say R_x is a self-inverse operation.

What is the inverse matrix of a projection Π_S ? Good luck finding that one—it's a trick question. The projection Π_S sends all input vectors from the subspace $S^\perp$ to the zero vector. Projections are inherently many-to-one transformations and therefore not invertible.

Nonstandard-basis probing

At this point, you should feel confident facing any linear transformation $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and probing it with the standard basis to find its matrix representation $M_T \in \mathbb{R}^{2 \times 2}$. But what if you're not allowed to probe T with the standard basis? Instead, let's say you must find the matrix of the transformation given the outputs of T for some other basis $\{\vec{v}_1 = (v_{1x}, v_{1y})^T, \vec{v}_2 = (v_{2x}, v_{2y})^T\}$:

$$T\left(\begin{bmatrix} v_{1x} \\ v_{1y} \end{bmatrix}\right) = \begin{bmatrix} t_{1x} \\ t_{1y} \end{bmatrix}, \quad T\left(\begin{bmatrix} v_{2x} \\ v_{2y} \end{bmatrix}\right) = \begin{bmatrix} t_{2x} \\ t_{2y} \end{bmatrix}.$$

Let's test this idea. We're given the information $v_{1x}, v_{1y}, t_{1x}, t_{1y}$, and $v_{2x}, v_{2y}, t_{2x}, t_{2y}$, and must find the matrix representation of T with respect to the standard basis.

Because the vectors $\vec{v}_1$ and $\vec{v}_2$ form a basis, we can reconstruct the information about the matrix M_T from the input-output information given. We're looking for four unknowns— m_{11} , m_{12} , m_{21} , and m_{22} —that form the matrix representation of T :

$$M_T = \begin{bmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{bmatrix}.$$

We can write four equations with the input-output information provided:

$$m_{11}v_{1x} + m_{12}v_{1y} = t_{1x},$$

$$m_{21}v_{1x} + m_{22}v_{1y} = t_{1y},$$

$$m_{11}v_{2x} + m_{12}v_{2y} = t_{2x},$$

$$m_{21}v_{2x} + m_{22}v_{2y} = t_{2y}.$$

Since there are four equations and four unknowns, we can solve for the coefficients m_{11} , m_{12} , m_{21} , and m_{22} .

This system of equations differs from ones we've seen before, so we'll examine it in detail. Think of the entries of M_T as a 4×1 vector of unknowns $\vec{m} = (m_{11}, m_{12}, m_{21}, m_{22})^T$. We can rewrite the four equations as a matrix equation:

$$A\vec{m} = \vec{t} \quad \Leftrightarrow \quad \begin{bmatrix} v_{1x} & v_{1y} & 0 & 0 \\ 0 & 0 & v_{1x} & v_{1y} \\ v_{2x} & v_{2y} & 0 & 0 \\ 0 & 0 & v_{2x} & v_{2y} \end{bmatrix} \begin{bmatrix} m_{11} \\ m_{12} \\ m_{21} \\ m_{22} \end{bmatrix} = \begin{bmatrix} t_{1x} \\ t_{1y} \\ t_{2x} \\ t_{2y} \end{bmatrix}.$$

Next, solve for $\vec{m}$ by computing $\vec{m} = A^{-1}\vec{t}$.

Finding the matrix representation by probing with a nonstandard basis is more work than probing with the standard basis, but it's totally doable.

Eigenspaces

Probing the transformation T with *any* basis for the input space gives sufficient information to determine its matrix representation. We're free to choose the "probing basis," so how do we decide which basis to use? The standard basis is good for computing the matrix representation, but perhaps there's a basis that allows us to simplify the abstract description of T ; a so-called *natural* basis for probing each transformation.

Indeed, such a basis exists. Many linear transformations have a basis $\{\vec{e}_{\lambda_1}, \vec{e}_{\lambda_2}, \dots\}$ such that the action of T on the basis vector $\vec{e}_{\lambda_i}$ is equivalent to the scaling of $\vec{e}_{\lambda_i}$ by the constant λ_i :

$$T(\vec{e}_{\lambda_i}) = \lambda_i \vec{e}_{\lambda_i}.$$

Recall how projections leave some vectors unchanged (multiply by $\lambda = 1$) and send other vectors to the zero vector (multiply by $\lambda = 0$). These subspaces of the input space are specific to each transformation, and are called the *eigenspaces* (from the German "own spaces") of the transformation T .

Consider the reflection R_x and its two eigenspaces:

- The space of vectors that remain unchanged (the eigenspace corresponding to $\lambda_1 = 1$) is spanned by the vector $(1, 0)$:

$$R_x\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}\right) = 1 \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

- The space of vectors that become the exact negatives of themselves (corresponding to $\lambda_2 = -1$) is spanned by $(0, 1)$:

$$R_x\left(\begin{bmatrix} 0 \\ 1 \end{bmatrix}\right) = -1 \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

From a theoretical point of view, describing the action of a linear transformation T in its natural basis is the best way to understand it. For each of the *eigenvectors* in the various eigenspaces of T , the action of T is a simple scalar multiplication. We defer the detailed discussion on *eigenvalues* and *eigenvectors* until the next chapter.

Links

[Rotation operation as the composition of three shear operations]
<http://datagenetics.com/blog/august32013/index.html>

Exercises

E5.5 Find the matrix representation of the linear transformations: the rotation by $\frac{\pi}{4}$ radians $R_{\frac{\pi}{4}}$, and the rotation by $\frac{\pi}{2}$ radians $R_{\frac{\pi}{2}}$. Verify the equation $R_{\frac{\pi}{4}} \circ R_{\frac{\pi}{4}} = R_{\frac{\pi}{2}}$ by computing the product of the matrix representations.

Hint: The trigonometric identities from Section 1.12 might be helpful.

E5.6 Find the matrix representation of the projection $\Pi_{\vec{d}}$ that projects vectors onto the subspace spanned by the direction vector $\vec{d} = (1, 3)^T$.

Hint: Compute the unit vector $\hat{d}$, then use the outer product formula.

5.3 Change of basis for matrices

Every linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ can be represented as a matrix $M_T \in \mathbb{R}^{m \times n}$. The entries of the matrix M_T depend on the basis used to describe the input and output spaces. Note, this dependence of matrix entries on the basis is directly analogous to the dependence of vector coordinates on the basis.

In this section, we'll learn how the choice of basis affects the entries of matrix representations, and discuss how to carry out the change-of-basis operation for matrices.

Concepts

You should already be familiar with the concepts of vector spaces, bases, vector coordinates with respect to different bases, and the change-of-basis transformation:

- V : an n -dimensional vector space
- $\vec{v}$: a vector in V
- $B = \{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$: an orthonormal basis for V
- $[\vec{v}]_B = (v_1, v_2, \dots, v_n)_B$: the vector $\vec{v}$ expressed in the basis B
- $B' = \{\hat{e}'_1, \hat{e}'_2, \dots, \hat{e}'_n\}$: another orthonormal basis for V
- $[\vec{v}]_{B'} = (v'_1, v'_2, \dots, v'_n)_{B'}$: the vector $\vec{v}$ expressed in the basis B'
- ${}_B[\mathbb{1}]_{B'}$: the change-of-basis matrix that converts from B coordinates to B' coordinates, $[\vec{v}]_{B'} = {}_{B'}[\mathbb{1}]_B [\vec{v}]_B$
- ${}_B[\mathbb{1}]_{B'}$: the inverse change-of-basis matrix $[\vec{v}]_B = {}_B[\mathbb{1}]_{B'} [\vec{v}]_{B'}$ (note that ${}_B[\mathbb{1}]_{B'} = ({}_{B'}[\mathbb{1}]_B)^{-1}$)

We'll use the following concepts when describing a linear transformation $T : V \rightarrow W$:

- $B_V = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$: a basis for the input vector space V
- $B_W = \{\vec{b}_1, \vec{b}_2, \dots, \vec{b}_m\}$: a basis for the output vector space W
- ${}_{B_W}[M_T]_{B_V} \in \mathbb{R}^{m \times n}$: a matrix representation of the linear transformation T with respect to the bases B_V and B_W ,

$$\vec{w} = T(\vec{v}) \quad \Leftrightarrow \quad [\vec{w}]_{B_W} = {}_{B_W}[M_T]_{B_V} [\vec{v}]_{B_V}$$

By far, the most commonly used basis in linear algebra is the standard basis $\hat{e}_1 = (1, 0, 0, \dots)^T$, $\hat{e}_2 = (0, 1, 0, \dots)^T$, etc. It is therefore customary to denote the matrix representation of a linear transformation T simply as M_T , without an explicit reference to the input and output bases. This simplified notation causes much confusion when students later try to learn about change-of-basis operations.

In order to *really* understand the connection between linear transformations and their matrix representations, we need to have a little talk about bases and matrix entries. It's a little complicated, but the mental effort you invest is worth the overall understanding you'll gain. As the old Samurai saying goes, "Cry during training, laugh on the battlefield."

By the end of this section, you'll be able to handle any basis question your teacher may throw at you.

Matrix entries

Every linear transformation T can be represented as a matrix M_T . Consider the linear transformation $T : V \rightarrow W$. Assume the input vector space V is n -dimensional and let $B_V = \{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$ be a basis for V . Assume the output space W is m -dimensional and let $B_W = \{\vec{b}_1, \vec{b}_2, \dots, \vec{b}_m\}$ be a basis for output space of T . The entries of the matrix $M_T \in \mathbb{R}^{m \times n}$ depend on the bases B_V and B_W . We'll now analyze this dependence in detail.

To compute the matrix representation of T with respect to the input basis $B_V = \{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$, we probe T with each of the vectors in the basis and record the outputs as the columns of a matrix:

$$[M_T]_{B_V} = \begin{bmatrix} | & | & & | \\ T(\hat{e}_1) & T(\hat{e}_2) & \cdots & T(\hat{e}_n) \\ | & | & & | \end{bmatrix}_{B_V}.$$

The subscript B_V indicates the columns are built from outputs of the basis B_V . We can use the matrix $[M_T]_{B_V}$ to compute $T(\vec{v})$ for a vector $\vec{v}$ expressed in the basis B_V : $\vec{v} = (v_1, v_2, \dots, v_n)_{B_V}^T$. The matrix-vector product produces the correct linear combination of outputs:

$$\begin{aligned} [M_T]_{B_V} [\vec{v}]_{B_V} &= \begin{bmatrix} | & | & & | \\ T(\hat{e}_1) & T(\hat{e}_2) & \cdots & T(\hat{e}_n) \\ | & | & & | \end{bmatrix}_{B_V} \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix}_{B_V} \\ &= v_1 T(\hat{e}_1) + v_2 T(\hat{e}_2) + \cdots + v_n T(\hat{e}_n) \\ &= T(v_1 \hat{e}_1 + v_2 \hat{e}_2 + \cdots + v_n \hat{e}_n) \\ &= T(\vec{v}). \end{aligned}$$

So far we've treated the outputs of T as abstract vectors $T(\hat{e}_i) \in W$. Like all vectors in the space W , each output of T can be expressed as

a vector of components with respect to the basis B_W . For example, the output $T(\hat{e}_1)$ can be expressed as

$$T(\hat{e}_1) = \begin{bmatrix} c_{11} \\ c_{21} \\ \vdots \\ c_{m1} \end{bmatrix}_{B_W} = c_{11}\vec{b}_1 + c_{21}\vec{b}_2 + \cdots + c_{m1}\vec{b}_m$$

for some coefficients $c_{11}, c_{21}, \dots, c_{m1}$. Similarly, the other output vectors $T(\hat{e}_j)$ can be expressed as components with respect to the basis B_W , $T(\hat{e}_j) = (c_{1j}, c_{2j}, \dots, c_{mj})_{B_W}^T$.

We're now in a position to find the matrix representation ${}_{B_W}[M_T]_{B_V}$ of the linear transformation T , with respect to the input basis B_V and the output basis B_W :

$${}_{B_W}[M_T]_{B_V} = \begin{bmatrix} c_{11} & c_{12} & \cdots & c_{1n} \\ c_{21} & c_{22} & \cdots & c_{2n} \\ \vdots & & & \vdots \\ c_{m1} & c_{m2} & \cdots & c_{mn} \end{bmatrix}_{B_W} \in \mathbb{R}^{m \times n}.$$

The action of T on a vector $\vec{v}$ is the same as the product of ${}_{B_W}[M_T]_{B_V}$ and the vector of components $[\vec{v}]_{B_V} = (v_1, v_2, \dots, v_n)_{B_V}^T$:

$$[T(\vec{v})]_{B_W} = {}_{B_W}[M_T]_{B_V} [\vec{v}]_{B_V}.$$

You may feel this example has stretched the limits of your attention span, but bear in mind, these nitty-gritty details hold the meaning of matrix entries. If you can see how the *positions* of the entries in the matrix encode the information about T and the choice of bases B_V and B_W , you're well on your way to getting it. The entry c_{ij} in the i^{th} row and j^{th} column in the matrix ${}_{B_W}[M_T]_{B_V}$ is the i^{th} component (with respect to B_W) of the output of T when the input is $\hat{e}_j$.

Let's verify that the matrix representation ${}_{B_W}[M_T]_{B_V}$ correctly predicts the output of T for the input $\vec{v} = 5\hat{e}_1 + 6\hat{e}_2 = (5, 6, 0, \dots)_{B_V}^T$. Using the linearity of T , we know the correct output is $T(\vec{v}) = T(5\hat{e}_1 + 6\hat{e}_2) = 5T(\hat{e}_1) + 6T(\hat{e}_2)$. We can verify that the matrix-vector product ${}_{B_W}[M_T]_{B_V} [\vec{v}]_{B_V}$ leads to the same answer:

$$\begin{bmatrix} c_{11} & c_{12} & \cdots & c_{1n} \\ c_{21} & c_{22} & \cdots & c_{2n} \\ \vdots & & & \vdots \\ c_{m1} & c_{m2} & \cdots & c_{mn} \end{bmatrix} \begin{bmatrix} 5 \\ 6 \\ 0 \\ \vdots \end{bmatrix} = 5 \begin{bmatrix} c_{11} \\ c_{21} \\ \vdots \\ c_{m1} \end{bmatrix} + 6 \begin{bmatrix} c_{12} \\ c_{22} \\ \vdots \\ c_{m2} \end{bmatrix} = 5T(\hat{e}_1) + 6T(\hat{e}_2).$$

Change of basis for matrices

Given the matrix representation ${}_{B_W}[M_T]_{B_V}$ of the linear transformation $T : V \rightarrow W$, you're asked to find the matrix representation of T with respect to different bases B'_V and B'_W . This is the *change-of-basis* task for matrices.

We'll discuss the important special case where the input space and the output space of the linear transformation are the same. Let $T : V \rightarrow V$ be a linear transformation, and let $B = \{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$ and $B' = \{\hat{e}'_1, \hat{e}'_2, \dots, \hat{e}'_n\}$ be two bases for the vector space V .

Recall the change-of-basis matrix ${}_{B'}[\mathbb{1}]_B$ that converts vectors from B coordinates to B' coordinates, and its inverse ${}_B[\mathbb{1}]_{B'}$, which converts vectors from B' coordinates to B coordinates:

$$[\vec{v}]_{B'} = {}_{B'}[\mathbb{1}]_B [\vec{v}]_B \quad \text{and} \quad [\vec{v}]_B = {}_B[\mathbb{1}]_{B'} [\vec{v}]_{B'}.$$

A clarification of notation is in order. The change-of-basis matrix ${}_{B'}[\mathbb{1}]_B$ is not equal to the identity matrix $\mathbb{1}_n$. However, the change-of-basis operation is *logically* equivalent to an identity transformation: the vector $\vec{v}$ doesn't change—only its coordinates change. If you don't remember the change-of-basis operation for vectors, now's the time to flip back to Section 4.3 (page 229) and review before continuing.

Given the matrix representation ${}_B[M_T]_B$ of the linear transformation T with respect to B , we want to find the matrix ${}_{B'}[M_T]_{B'}$, which is the representation of T with respect to the basis B' . The computation is straightforward. Perform the change-of-basis operation on the input and output vectors:

$${}_{B'}[M_T]_{B'} = {}_{B'}[\mathbb{1}]_B {}_B[M_T]_B {}_B[\mathbb{1}]_{B'}.$$

This group of three matrices is interpreted as follows. Imagine an input vector $[\vec{v}]_{B'}$ multiplying the three matrices ${}_{B'}[\mathbb{1}]_B {}_B[M_T]_B {}_B[\mathbb{1}]_{B'}$ from the right. In the first step, ${}_B[\mathbb{1}]_{B'}$ converts the vector from the basis B' to the basis B so the matrix ${}_B[M_T]_B$ can be applied. In the last step, the matrix ${}_{B'}[\mathbb{1}]_B$ converts the output of ${}_B[M_T]_B$ to the basis B' . The combined effect of multiplying by this specific arrangement of three matrices is the same as applying T to the input vector $\vec{v}$:

$${}_{B'}[\mathbb{1}]_B {}_B[M_T]_B {}_B[\mathbb{1}]_{B'} [\vec{v}]_{B'} = [T(\vec{v})]_{B'},$$

which means

$${}_{B'}[M_T]_{B'} = {}_{B'}[\mathbb{1}]_B {}_B[M_T]_B {}_B[\mathbb{1}]_{B'}.$$

This formula makes sense intuitively: to obtain a matrix with respect to a different basis, we must surround it by appropriate change-of-basis matrices.

Note the touching dimensions of the matrices are expressed with respect to the same basis. Indeed, we can think of the change-of-basis matrices as adaptors we use to express vectors in different bases. The change-of-basis operation for matrices requires two adaptors; one for the input space and one for the output space.

Similarity transformation

It's interesting to note the abstract mathematical properties of the operation used above. Consider any matrix $N \in \mathbb{R}^{n \times n}$ and an invertible matrix $P \in \mathbb{R}^{n \times n}$. Define M to be the result when N is multiplied by P on the left and by the inverse P^{-1} on the right:

$$M = PNP^{-1}.$$

We say matrices N and M are related by a *similarity transformation*.

Since the matrix P is invertible, its columns form a basis for the vector space $\mathbb{R}^n$. Thus, we can interpret P as a *change-of-basis* matrix that converts the standard basis to the basis of the columns of P . The matrix P^{-1} corresponds to the inverse change-of-basis matrix. Using this interpretation, the matrix M corresponds to the *same* linear transformation as the matrix N , but is expressed with respect to the basis P .

Similarity transformations preserve certain properties of matrices:

- Trace: $\text{Tr}(M) = \text{Tr}(N)$
- Determinant: $\det(M) = \det(N)$
- Rank: $\text{rank}(M) = \text{rank}(N)$
- Eigenvalues: $\text{eig}(M) = \text{eig}(N)$

Together, the trace, the determinant, the rank, and the eigenvalues of a matrix are known as the invariant properties of the matrix because they don't depend on the choice of basis.

Exercises

E5.7 Suppose you're given the matrix representation of a linear transformation T with respect to the basis B' : ${}_{B'}[M_T]_{B'}$. What formula describes the matrix representation of T with respect to the basis B ?

E5.8 Consider the linear transformation $T(x, y) = (3x, y)$, which represents a stretching by a factor of three along the x -axis. Find the matrix representation of T with respect to the following bases.

- a) The standard basis $B_s = \{(1, 0), (0, 1)\}$
- b) The flipped basis $B_f = \{(0, 1), (1, 0)\}$
- c) The diagonal basis $B_d = \{(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}), (\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}})\}$

E5.9 You asked your lab assistant to compute the outputs of a list of vectors $\vec{x}_i$ when passed through the linear transformation T . You specified the vectors in the basis B . Unfortunately, your lab assistant made a mistake and used a different basis B' , applying the matrix ${}_{B'}[M_T]_{B'}$ to the vectors, and obtaining the wrong outputs $[\vec{y}_i^w]_* = {}_{B'}[M_T]_{B'} [\vec{x}_i]_B$. To make things worse, your assistant deleted the original data $\vec{x}_i$!

“No worries,” you say. “Since the linear transformation T is invertible, we can recover the data.” What sequence of operations will recover $\vec{x}_i$ from $[\vec{y}_i^w]_*$? What sequence of operations will convert $[\vec{y}_i^w]_*$ into the correct output $[\vec{y}_i]_B$?

5.4 Invertible matrix theorem

So far, we discussed systems of linear equations, matrices, vector spaces, and linear transformations. It's time to tie it all together! We'll now explore connections between these different contexts where matrices are used. Originally, we saw how matrices can be used to solve systems of linear equations. Later, we studied geometric properties of matrices, including their row spaces, column spaces, and null spaces. We also learned about the connection between matrices and linear transformations. In each of these domains, *invertible* matrices play a particularly important role. Lucky for us, there's a theorem that summarizes 10 important facts about invertible matrices. One theorem; 10 facts. Now that's a good deal!

Invertible matrix theorem. For an $n \times n$ matrix A , the following statements are equivalent:

- (1) A is invertible
- (2) The equation $A\vec{x} = \vec{b}$ has exactly one solution for each $\vec{b} \in \mathbb{R}^n$
- (3) The null space of A contains only the zero vector $\mathcal{N}(A) = \{\vec{0}\}$
- (4) The equation $A\vec{x} = \vec{0}$ has only the trivial solution $\vec{x} = \vec{0}$
- (5) The columns of A form a basis for $\mathbb{R}^n$:
 - The columns of A are linearly independent
 - The columns of A span $\mathbb{R}^n$; $\mathcal{C}(A) = \mathbb{R}^n$
- (6) The rank of the matrix A is n
- (7) The RREF of A is the $n \times n$ identity matrix $\mathbb{I}_n$
- (8) The transpose matrix A^T is invertible
- (9) The rows of A form a basis for $\mathbb{R}^n$:
 - The rows of A are linearly independent
 - The rows of A span $\mathbb{R}^n$; $\mathcal{R}(A) = \mathbb{R}^n$
- (10) The determinant of A is nonzero $\det(A) \neq 0$

These 10 statements are either all true or all false for any given matrix A . We can split the set of $n \times n$ matrices into two disjoint subsets: invertible matrices, for which all 10 statements are true, and non-invertible matrices, for which all statements are false.

Proving the invertible matrix theorem

It's essential that you understand the details of this theorem, including its proof. The reasoning that connects these 10 statements unites all the chunks of linear algebra we've discussed. Not being a "proof person" is not a valid excuse! Be sure to read the proof, as it will help to solidify your understanding of the material covered thus far.

Proofs by contradiction

Since our arrival at the invertible matrix theorem marks an important step, we'll first quickly review some handy proof techniques, just to make sure everyone's ready. A *proof by contradiction* starts by assuming the opposite of the fact we want to prove, and after several

derivation steps arrives at a contradiction—a mathematical inconsistency. Arriving at a contradiction implies our original premise is false, which means the fact we want to prove is true. Thus, to show that (A) implies (B)—denoted $(A) \Rightarrow (B)$ —we can show that not-(B) implies not-(A).

Review of definitions

To really make sure we're all on board before the train leaves the station, it's wise to review some definitions from previous chapters. The matrix $A \in \mathbb{R}^{n \times n}$ is *invertible* if there exists a matrix A^{-1} such that $AA^{-1} = \mathbb{I}_n = A^{-1}A$. The *null space* of A is the set of vectors that become the zero vector when multiplying A from the right: $\{\vec{v} \in \mathbb{R}^n \mid A\vec{v} = \vec{0}\}$. The *column space* $\mathcal{C}(A) \subseteq \mathbb{R}^n$ consists of all possible linear combinations of the columns of the matrix A . Similarly, the *row space* $\mathcal{R}(A) \subseteq \mathbb{R}^n$ consists of all possible linear combinations of the rows of A . The *rank* of a matrix, denoted $\text{rank}(A)$, is equal to the dimension of the row space and the column space $\text{rank}(A) = \dim(\mathcal{C}(A)) = \dim(\mathcal{R}(A))$. The rank of A is also equal to the number of pivots (leading ones) in the reduced row echelon form of A . The *determinant* of A corresponds to the *scale factor* by which the linear transformation $T_A(\vec{x}) = A\vec{x}$ transforms the n -dimensional volume of the hyper-cube in the input space when it maps it to the output space. If any of the columns of the matrix A are linearly dependent, the determinant $|A|$ will be zero.

Now tell me, do you feel ready to board the train? Don't worry too much if some of the above definitions about matrices and vector spaces are still unclear for you. The whole point of this theorem is to solidify your understanding by demonstrating how the concepts you understand less are connected to the concepts you understand well.

* * *

Proof of the invertible matrix theorem. The moment has arrived: we'll prove the equivalence of the 10 statements in the theorem by showing a closed chain of implications between statements (1) through (7). We'll separately show the equivalences (1) $\Leftrightarrow$ (8) $\Leftrightarrow$ (9) and (5) $\Leftrightarrow$ (10). Figure 5.12 shows an outline of the proof.

(1) $\Rightarrow$ (2): Assume A is invertible so there exists an inverse matrix A^{-1} such that $A^{-1}A = \mathbb{I}_n$. Therefore, for all $\vec{b} \in \mathbb{R}^n$, the expression $\vec{x} = A^{-1}\vec{b}$ is a solution to the equation $A\vec{x} = \vec{b}$. We must show the solution $\vec{x} = A^{-1}\vec{b}$ is the unique solution to this equation. Suppose that another solution $\vec{y}$ exists such that $A\vec{y} = \vec{b}$. Multiplying both




Figure 5.12: The chain of implications used to prove the invertible matrix theorem.

sides of the equation $A\vec{y} = \vec{b}$ by A^{-1} , we obtain $A^{-1}A\vec{y} = \vec{y} = A^{-1}\vec{b}$, which shows that $\vec{y} = \vec{x} = A^{-1}\vec{b}$, and so $\vec{x} = A^{-1}\vec{b}$ is the unique solution to $A\vec{x} = \vec{b}$.

(2) $\Rightarrow$ (3): We want to show that a unique solution to $A\vec{x} = \vec{b}$ implies the matrix A has a trivial null space $\mathcal{N}(A) = \{\vec{0}\}$. We start by assuming the opposite is true: that $\mathcal{N}(A)$ contains at least one nonzero vector $\vec{y} \neq \vec{0}$. If this were true, then $\vec{x}' = \vec{x} + \vec{y}$ would also be a solution to $A\vec{x} = \vec{b}$, since $A\vec{x}' = A(\vec{x} + \vec{y}) = A\vec{x} + A\vec{y} = A\vec{x} + \vec{0} = \vec{b}$, since $A\vec{y} = \vec{0}$. The fact that two solutions ($\vec{x}$ and $\vec{x}' \neq \vec{x}$) exist contradicts the statement that $A\vec{x} = \vec{b}$ has a unique solution. Thus, for $A\vec{x} = \vec{b}$ to have a unique solution, A must have a trivial null space $\mathcal{N}(A) = \{\vec{0}\}$.

(3) $\Rightarrow$ (4) Statements (3) and (4) are equivalent by definition: the condition that A 's null space contains only the zero vector, $\mathcal{N}(A) = \{\vec{0}\}$, is equivalent to the condition that the only solution to the equation $A\vec{v} = \vec{0}$ is $\vec{v} = \vec{0}$.

(4) $\Rightarrow$ (5): Analyze the equation $A\vec{v} = \vec{0}$ in the column picture of matrix multiplication, denoting the n columns of A as $\vec{c}_1, \vec{c}_2, \dots, \vec{c}_n$. We obtain $A\vec{v} = v_1\vec{c}_1 + v_2\vec{c}_2 + \dots + v_n\vec{c}_n = \vec{0}$. Since $\vec{v} = (v_1, v_2, \dots, v_n) = \vec{0}$ is the only solution to this equation, we obtain the statement in the definition of linear independence for a set of vectors. The fact that $A\vec{v} = \vec{0}$ has only $\vec{v} = \vec{0}$ as a solution implies the columns of A form a linearly independent set. Furthermore, the columns of A form a basis for $\mathbb{R}^n$ because they are a set of n linearly independent vectors in an n -dimensional vector space.

(5) $\Rightarrow$ (6): We know $\text{rank}(A)$ equals the number of linearly independent columns in A . Since the n columns of A are linearly independent, it follows that $\text{rank}(A) = n$.

(6) $\Rightarrow$ (7): The rank is equal to the number of leading ones (pivots) in the RREF of A . Since A has rank n , its reduced row echelon form

must contain n pivots. The reduced row echelon form of an $n \times n$ matrix with n pivots is the identity matrix $\mathbb{I}_n$.

(7) $\Rightarrow$ (1): We start from the assumption $\text{rref}(A) = \mathbb{I}_n$. This means it's possible to use a set of row operations $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k$ to transform A to the identity matrix: $\mathcal{R}_k(\dots \mathcal{R}_2(\mathcal{R}_1(A))\dots) = \mathbb{I}_n$. Consider the elementary matrices $E_1, E_2, \dots, E_k$ that correspond to the row operations $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k$. Rewriting the equation $\mathcal{R}_k(\dots \mathcal{R}_2(\mathcal{R}_1(A))\dots) = \mathbb{I}_n$ in terms of these elementary matrices gives us $E_k \dots E_2 E_1 A = \mathbb{I}_n$. This equation implies the inverse of A exists and is equal to the product of elementary matrices $A^{-1} = E_k \dots E_2 E_1$.

(1) $\Leftrightarrow$ (8): If A is invertible, there exists A^{-1} such that $AA^{-1} = \mathbb{I}_n$. If we apply the transpose operation to this equation, we obtain

$$(AA^{-1})^T = (\mathbb{I}_n)^T \quad \Rightarrow \quad (A^{-1})^T A^T = \mathbb{I}_n,$$

which shows the matrix $(A^{-1})^T = (A^T)^{-1}$ exists and is the inverse of A^T . Therefore, A is invertible if and only if A^T is invertible.

(8) $\Leftrightarrow$ (9): Statement (9) follows by a combination of statements (8) and (5): if A^T is invertible, its columns form a basis for $\mathbb{R}^n$. Since the columns of A^T are the rows of A , it follows that the rows of A form a basis for $\mathbb{R}^n$.

(5) $\Leftrightarrow$ (10): The determinant of an $n \times n$ matrix is nonzero if and only if the columns of the matrix are linearly independent. Thus, the columns of the matrix A form a basis for $\mathbb{R}^n$ if and only if $\det(A) \neq 0$. $\square$

Nice work! By proving the chain of implications **(1)** $\Rightarrow$ **(2)** $\Rightarrow \dots \Rightarrow$ **(7)** $\Rightarrow$ **(1)**, we've shown that the first seven statements are equivalent. If one of these statements is true, then all others are true—just follow the arrows of implication. Alternatively, if one statement is false, all statements are false, as we see by following the arrows of implication in the backward direction.

We also “attached” statements (8), (9), and (10) to the main loop of implications using “if and only if” statements. Thus, we've shown the equivalence of all 10 statements, which completes the proof.

The steps of the proof shown above cover only a small selection of all possible implications between the 10 statements. In a few pages, you'll reach the exercises and you'll be asked to prove several other implications. It's important to try these exercises on your own! By

checking whether you can obtain the proofs, you'll force your brain to truly grasp the definitions and properties of linear algebra concepts. Note the crucial difference between *one-way* implications of the form $(\mathbf{A}) \Rightarrow (\mathbf{B})$ and *if and only if* statements $(\mathbf{A}) \Leftrightarrow (\mathbf{B})$. The latter require you to prove both directions of the implication: $(\mathbf{A}) \Rightarrow (\mathbf{B})$ and $(\mathbf{A}) \Leftarrow (\mathbf{B})$.

Invertible linear transformations

We can reinterpret the statements in the invertible matrix theorem as a statement about *invertible linear transformations*:

$$T: \mathbb{R}^n \rightarrow \mathbb{R}^n \text{ is invertible} \quad \Leftrightarrow \quad M_T \in \mathbb{R}^{n \times n} \text{ is invertible.}$$

The set of linear transformations splits into two disjoint subsets; invertible linear transformations and non-invertible linear transformations.

Kernel and null space

The *kernel* of the linear transformation T is the same as the null space of its matrix representation M_T . Recall statement (3) of the invertible matrix theorem: a matrix A is invertible if and only if its null space contains only the zero vector $\mathcal{N}(A) = \{\vec{0}\}$. The equivalent condition for linear transformations is the zero kernel condition. A linear transformation T is invertible if and only if its kernel contains only the zero vector:

$$T \text{ is invertible} \quad \Leftrightarrow \quad \text{Ker}(T) = \{\vec{0}\}.$$

Invertible linear transformations $T: \mathbb{R}^n \rightarrow \mathbb{R}^n$ map different input vectors $\vec{x}$ to different output vectors $\vec{y} = T(\vec{v})$; therefore it's possible to build an inverse linear transformation $T^{-1}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ that restores every $\vec{y}$ back to the $\vec{x}$ it came from.

In contrast, a non-invertible linear transformation S sends all vectors $\vec{x} \in \text{Ker}(S) \neq \{\vec{0}\}$ to the zero vector $S(\vec{x}) = \vec{0}$. When this happens, there is no way to undo the action of S , since we can't determine the input $\vec{x}$ that S sent to $\vec{0}$.

Linear transformations as functions

In Section 1.8, we discussed the notion of *invertibility* for functions of a real variable, $f: \mathbb{R} \rightarrow \mathbb{R}$. In particular, we used the terms *injective*, *surjective*, and *bijective* to describe how a function maps different inputs from its domain to outputs in its target set (see page 38). Since

linear transformations are functions, we can apply the general terminology for functions to describe how linear transformations map different inputs to outputs.

A linear transformation is *injective* if it maps different inputs to different outputs:

$$T(\vec{v}_1) \neq T(\vec{v}_2) \text{ for all } \vec{v}_1 \neq \vec{v}_2 \quad \Leftrightarrow \quad T \text{ is injective.}$$

A linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is *surjective* if its image space equals its output space:

$$\text{Im}(T) = \mathbb{R}^m \quad \Leftrightarrow \quad T \text{ is surjective.}$$

The surjective condition for linear transformations is equivalent to the condition that the column space of the matrix M_T spans the outputs space: $\text{Im}(T) = \mathcal{C}(M_T) = \mathbb{R}^m$.

If a function is both injective and surjective then it is *bijective*. Bijective functions are *one-to-one correspondences* between elements in their input space and elements in their output space.

$$T(\vec{x}) = \vec{y} \text{ has exactly one solution for each } \vec{y} \quad \Leftrightarrow \quad T \text{ is bijective.}$$

All bijective functions are invertible since each output $\vec{y}$ in the output space corresponds to exactly one $\vec{x}$ in the input space.

Interestingly, for a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ to be invertible, the presence of either the injective or surjective property is sufficient. If T is injective, it must have a $\text{Ker}(T) = \{\vec{0}\}$ so it is invertible. If T is surjective, its matrix representation M_T has rank n . The rank-nullity theorem ($\text{rank}(M_T) + \text{nullity}(M_T) = n$) tells us M_T has an empty null space $\mathcal{N}(M_T) = \text{Ker}(T) = \{\vec{0}\}$, making T invertible.

Links

[Nice writeup about the invertible matrix theorem with proofs]
<http://bit.ly/InvMatThmProofs2>

[Visualization of two-dimensional linear transformations]
<http://ncase.me/matrix/>

Exercises

E5.10 Prove that statement (5) of the invertible matrix theorem implies statement (2).

E5.11 Prove statement (3) implies statement (6) in the invertible matrix theorem.

Hint: Use the rank-nullity theorem from Chapter 4 (page 242).

Discussion

In this chapter, we learned about linear transformations and their matrix representations. The equivalence $T(\vec{x}) = M_T \vec{x}$ is important because it forms a bridge between the abstract notion of a linear transformation and its concrete implementation as a matrix-vector product. Everything you know about matrices can be applied to linear transformations, and everything you know about linear transformations can be applied to matrices. Which is mind-blowing, if you think about it.

We say T is *represented* by the matrix ${}_{B_W}[M_T]_{B_V}$ with respect to the basis B_V for the input space, and the basis B_W for the output space. In Section 5.2, we learned about the probing procedure for finding matrix representations with respect to the standard basis, while Section 5.3 discussed the notion of change of basis for matrices. Hold tight, because in the next chapter we'll learn about the eigenvalues and eigenvectors of matrices and discuss the *eigendecomposition* of matrices, which is a type of change of basis.

Section 5.4 gave us the invertible matrix theorem along with a taste of what it takes to prove formal math statements. It's extra important that you attempt some of the proofs in the exercise section on page 293. Although proofs can be complicated, they're so worth your time because they force you to clarify the definitions and properties of all the math concepts you've encountered thus far. Attempt the proofs in the problems section to find out if you're a linear algebra amateur, or a linear algebra expert.

5.5 Linear transformations problems

Understanding linear transformations is extremely important for your overall understanding of linear algebra. This is why it's crucial for you to solve all the problems in this section. By working on these problems, you'll discover whether you really understand all the new material covered in this chapter. Remember in the book's introduction, when I mentioned that linear algebra is all about vectors and linear transformations? Well, if you can solve all the problems in this section, you're 80% of the way to understanding all of linear algebra.

P5.1 Determine whether each of the following transformations are linear.

- a)** $T_1(x, y) = (y, x + y)$ **b)** $T_2(x, y) = (x + 3, y - 3)$
c) $T_3(x, y) = (|x|, |y|)$ **d)** $T_4(x, y, z) = (3x - 2y + z, 2x + y - 4z)$
e) $T_5(x) = (x, 2x, 3x)$ **f)** $T_6(x, y, z, w) = (5x, 4y, 3z, 2w, 1)$

If the transformation is linear, find its matrix representation. If the transformation is nonlinear, show a calculation where the linear property fails.

P5.2 Find image space $\text{Im}(T)$ for the linear transformation $T : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ defined by $T(x, y) = (x, x - y, 2y)$.

P5.3 Consider the transformation defined by $T(\vec{v}) = \vec{a} \times \vec{v}$, where $\vec{a} = (a_x, a_y, a_z)$. Find the matrix representation of T . What is the kernel of T ?

P5.4 Find the matrix representation of the linear transformation T that maps the input vector $\vec{x}_1 = (1, 1)^T$ to the output vector $\vec{y}_1 = (2, -3)^T$ and the input vector $\vec{x}_2 = (1, 2)^T$ to the output vector $\vec{y}_2 = (5, 1)^T$.

Hint: This is a nonstandard basis probing question. See page 279.


P5.5 Given the linear transformation $T(x, y, z) = (x, x + y, x + y + z)$, the standard basis for $\mathbb{R}^3$ $B = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$, and the alternative basis $B' = \{(1, 0, 0), (0, 1, 1), (0, 1, -1)\}$, find the following matrix representations of T :

- a) ${}_B[M_T]_B$: the representation of T with respect to the standard basis
- b) ${}_{B'}[M_T]_{B'}$: the representation of T with respect to the basis B'
- c) ${}_B[M_T]_{B'}$: the mixed representation of T with respect to input vectors expressed in the basis B' and output vectors in the standard basis B


P5.6 Consider the linear transformation $T(x, y, z) = (2x - 5z, 2y, -3z)$. Find the matrix of T with respect to the basis $B' = \{(1, 0, 0), (0, 1, 0), (1, 0, 1)\}$.

P5.7 Find the matrix representations of each of the transformations shown in Figure 5.13. The input to each transformation is the triangle with vertices $(0, 0)$, $(2, 0)$, $(0, 1)$ that is shown in Figure 5.13 image (a).


Hint: Your answers should be 2×2 matrices. Recall that $\sin(\frac{\pi}{6}) = \frac{1}{2}$.




(a) Original




(b) Transformation 1




(c) Transformation 2



(d) Transformation 3



(e) Transformation 4



(f) Transformation 5

Figure 5.13: The effects of various linear transformations on a triangle.

P5.8 Prove that statement (2) implies statement (1) in the invertible matrix theorem.

P5.9 Prove that statements (2) and (4) of the invertible matrix theorem are equivalent.

P5.10 Prove statement (4) implies statement (6) in the invertible matrix theorem.

P5.11 Prove statement (4) implies statement (7) in the invertible matrix theorem.

P5.12 Prove the chain of implications $(1) \Rightarrow (7) \Rightarrow (10)$ in the invertible matrix theorem.

P5.13 Suppose $T : V \rightarrow W$ is an injective linear transformation, and $\{\vec{v}_1, \dots, \vec{v}_n\}$ is a linearly independent set in V . Prove that $\{T(\vec{v}_1), \dots, T(\vec{v}_n)\}$ is a linearly independent set in W .

P5.14 Suppose $T : V \rightarrow W$ is a surjective linear transformation, and the set $\{\vec{v}_1, \dots, \vec{v}_n\}$ spans V . Prove that the set $\{T(\vec{v}_1), \dots, T(\vec{v}_n)\}$ spans W .

Hint: Consider an arbitrary $\vec{w} \in W$.

P5.15 Let A and B be $n \times n$ matrices such that $AB = \mathbb{1}$. Show that B has rank n (full rank) and prove that $BA = \mathbb{1}$.

Hint: Use a proof by contradiction to show B has full rank, then apply the associative law of matrix multiplication.

Chapter 6

Theoretical linear algebra

Let's take a trip down memory lane: 170 pages ago, we embarked on a mind-expanding journey through the land of linear algebra. We encountered vector and matrix operations. We studied systems of linear equations, solving them with row operations. We covered miles of linear transformations and their matrix representations. With the skills you've acquired to reach this point, you're ready to delve into the abstract, theoretical aspects of linear algebra—that is, since you know all the useful stuff, you can officially move on to the cool stuff.

In math, we often use abstraction to find the commonalities between different mathematical objects. These parallels give us a deeper understanding of the mathematical structures we compare. This chapter extends what we know about the vector space $\mathbb{R}^n$ to the realm of abstract vector spaces of vector-like mathematical objects (Section 6.3). We'll discuss linear independence, find bases, and count dimensions for these abstract vector spaces. We'll define abstract inner product operations and use them to generalize the concept of orthogonality for abstract vectors (Section 6.4). We'll explore the Gram–Schmidt orthogonalization procedure for distilling orthonormal bases from non-orthonormal bases (Section 6.5). Finally, we'll introduce vectors and matrices with complex entries (Section 6.7). Section 6.7 also reviews everything we've learned in this book, so be sure to read it even if complex numbers are not required for your course. Along the way, we'll develop a taxonomy for the different types of matrices according to their properties and applications (Section 6.2). We'll also investigate matrix decompositions—techniques for splitting matrices into products of simpler matrices (Section 6.6). The chapter begins by discussing the most important decomposition technique of them all: the *eigendecomposition*, which is a way to uncover the “natural basis” for any matrix.

6.1 Eigenvalues and eigenvectors

The set of eigenvectors of a matrix is a special set of input vectors for which the action of the matrix is described as a simple *scaling*. In Section 5.2, we observed how linear transformations act differently on different input spaces. We also observed the special case of the “zero eigenspace,” called the *null space* of a matrix. The action of a matrix on the vectors in its null space is equivalent to a multiplication by zero. We’ll now put these eigenvalues and eigenvectors under the microscope and see what more there is to see.

Decomposing a matrix in terms of its eigenvalues and its eigenvectors gives valuable insights into the properties of the matrix. Certain matrix calculations, like computing the power of the matrix, become much easier when we use the *eigendecomposition* of the matrix. For example, suppose we’re given a square matrix A ,

$$A = \begin{bmatrix} 9 & -2 \\ -2 & 6 \end{bmatrix},$$

and we want to compute A^7 . In other words, we want to compute

$$A^7 = \begin{bmatrix} 9 & -2 \\ -2 & 6 \end{bmatrix} \begin{bmatrix} 9 & -2 \\ -2 & 6 \end{bmatrix} \begin{bmatrix} 9 & -2 \\ -2 & 6 \end{bmatrix} \begin{bmatrix} 9 & -2 \\ -2 & 6 \end{bmatrix} \begin{bmatrix} 9 & -2 \\ -2 & 6 \end{bmatrix} \begin{bmatrix} 9 & -2 \\ -2 & 6 \end{bmatrix} \begin{bmatrix} 9 & -2 \\ -2 & 6 \end{bmatrix}.$$

That’s an awful lot of matrix multiplications! Now imagine how many times we’d need to multiply the matrix if we wanted to find A^{17} or A^{77} . Too many times, that’s how many. Let’s be smart about this. Every matrix corresponds to some linear operation. This means it’s legit to ask, “What does the matrix A do?” Once we figure that out, we can compute A^{77} by simply doing what A does 77 times.

The best way to see what a matrix does is to look inside it and see what it’s made of (you may need to gradually gain the matrix’s trust before it lets you do this). To understand the matrix A , you must find its *eigenvectors* and its *eigenvalues*. The word *eigen* is the German word for “self.” The eigenvectors of a matrix are its “self vectors,” and correspond to a natural choice of basis for describing the action of the matrix. The *eigendecomposition* is a change-of-basis operation that expresses the matrix A with respect to its *eigenbasis* (“self-basis”). The eigendecomposition of the matrix A is a product of three matrices:

$$A = \begin{bmatrix} 9 & -2 \\ -2 & 6 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 2 \\ 2 & -1 \end{bmatrix}}_Q \underbrace{\begin{bmatrix} 5 & 0 \\ 0 & 10 \end{bmatrix}}_\Lambda \underbrace{\begin{bmatrix} \frac{1}{5} & \frac{2}{5} \\ \frac{2}{5} & -\frac{1}{5} \end{bmatrix}}_{Q^{-1}} = Q\Lambda Q^{-1}.$$

You can multiply the three matrices $Q\Lambda Q^{-1}$ to obtain A . Note that the middle matrix Λ (the capital Greek letter *lambda*) has entries only on the diagonal. The diagonal matrix Λ is sandwiched between the matrix Q on the left and Q^{-1} (the inverse of Q) on the right.

The eigendecomposition of A allows us to compute A^7 in a civilized manner:

$$\begin{aligned} A^7 &= A A A A A A A \\ &= Q \underbrace{\Lambda Q^{-1} Q}_{\mathbb{I}} \underbrace{\Lambda Q^{-1} Q}_{\mathbb{I}} \underbrace{\Lambda Q^{-1} Q}_{\mathbb{I}} \underbrace{\Lambda Q^{-1} Q}_{\mathbb{I}} \underbrace{\Lambda Q^{-1} Q}_{\mathbb{I}} \underbrace{\Lambda Q^{-1} Q}_{\mathbb{I}} \underbrace{\Lambda Q^{-1} Q}_{\mathbb{I}} \\ &= Q \Lambda \mathbb{I} \Lambda \mathbb{I} \Lambda \mathbb{I} \Lambda \mathbb{I} \Lambda \mathbb{I} \Lambda Q^{-1} \\ &= Q \Lambda \Lambda \Lambda \Lambda \Lambda \Lambda \Lambda Q^{-1} \\ &= Q \Lambda^7 Q^{-1}. \end{aligned}$$

All the inner Q^{-1} s cancel with the adjacent Q s. How convenient! Since the matrix Λ is diagonal, it's easy to compute its seventh power:

$$\Lambda^7 = \begin{bmatrix} 5 & 0 \\ 0 & 10 \end{bmatrix}^7 = \begin{bmatrix} 5^7 & 0 \\ 0 & 10^7 \end{bmatrix} = \begin{bmatrix} 78125 & 0 \\ 0 & 10\,000\,000 \end{bmatrix}.$$

Thus we can express our calculation of A^7 as

$$A^7 = \underbrace{\begin{bmatrix} 9 & -2 \\ -2 & 6 \end{bmatrix}}_Q^7 \underbrace{\begin{bmatrix} 78125 & 0 \\ 0 & 10\,000\,000 \end{bmatrix}}_{Q^{-1}} \underbrace{\begin{bmatrix} \frac{1}{5} & \frac{2}{5} \\ \frac{2}{5} & -\frac{1}{5} \end{bmatrix}}_{Q^{-1}}.$$

We still need to multiply these three matrices together, but we've cut the work from six matrix multiplications to two. The answer is

$$A^7 = Q \Lambda^7 Q^{-1} = \begin{bmatrix} 8015625 & -3968750 \\ -3968750 & 2062500 \end{bmatrix}.$$

With this technique, we can compute A^{17} just as easily:

$$A^{17} = Q \Lambda^{17} Q^{-1} = \begin{bmatrix} 80000152587890625 & -39999694824218750 \\ -39999694824218750 & 20000610351562500 \end{bmatrix}.$$

We could even compute $A^{777} = Q \Lambda^{777} Q^{-1}$ if we wanted to. I hope by now you get the point: once you express A in its *eigenbasis*, computing the powers of A requires computing the powers of its *eigenvalues*, which is much simpler than carrying out matrix multiplications.

Definitions

- A : an $n \times n$ square matrix. The entries of A are denoted as a_{ij} .
- $\text{eig}(A) = (\lambda_1, \lambda_2, \dots, \lambda_n)$: the list of *eigenvalues* of A . Eigenvalues are usually denoted by the Greek letter *lambda*. Note that some eigenvalues could be repeated in the list.
- $p(\lambda) = \det(A - \lambda \mathbb{1})$: the *characteristic polynomial* of A . The eigenvalues of A are the roots of the characteristic polynomial.
- $\{\vec{e}_{\lambda_1}, \vec{e}_{\lambda_2}, \dots, \vec{e}_{\lambda_n}\}$: the set of *eigenvectors* of A . Each eigenvector is associated with a corresponding eigenvalue.
- $\Lambda \stackrel{\text{def}}{=} \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$: the diagonalized version of A . The matrix Λ contains the eigenvalues of A on the diagonal:

$$\Lambda = \begin{bmatrix} \lambda_1 & \cdots & 0 \\ \vdots & \ddots & 0 \\ 0 & 0 & \lambda_n \end{bmatrix}.$$

The matrix Λ is the matrix A expressed in its eigenbasis.

- Q : a matrix whose columns are eigenvectors of A :

$$Q \stackrel{\text{def}}{=} \begin{bmatrix} | & & | \\ \vec{e}_{\lambda_1} & \cdots & \vec{e}_{\lambda_n} \\ | & & | \end{bmatrix} = {}_{B_s}[\mathbb{1}]_{B_\lambda}.$$

The matrix Q corresponds to the *change-of-basis* matrix from the eigenbasis $B_\lambda = \{\vec{e}_{\lambda_1}, \vec{e}_{\lambda_2}, \vec{e}_{\lambda_3}, \dots\}$ to the standard basis $B_s = \{\hat{i}, \hat{j}, \hat{k}, \dots\}$.

- $A = Q\Lambda Q^{-1}$: the *eigendecomposition* of the matrix A
- $\Lambda = Q^{-1}AQ$: the *diagonalization* of the matrix A

Eigenvalues

The fundamental eigenvalue equation is

$$A\vec{e}_\lambda = \lambda\vec{e}_\lambda,$$

where λ is an eigenvalue and $\vec{e}_\lambda$ is an eigenvector of the matrix A . Multiply A by one of its eigenvectors $\vec{e}_\lambda$, and the result is the same vector scaled by the constant λ .

To find the eigenvalues of a matrix, start from the eigenvalue equation $A\vec{e}_\lambda = \lambda\vec{e}_\lambda$, insert the identity $\mathbb{1}$, and rewrite the equation as a null-space problem:

$$A\vec{e}_\lambda = \lambda\mathbb{1}\vec{e}_\lambda \quad \Rightarrow \quad (A - \lambda\mathbb{1})\vec{e}_\lambda = \vec{0}.$$

This equation has a solution whenever $|A - \lambda \mathbb{1}| = 0$. The eigenvalues of $A \in \mathbb{R}^{n \times n}$, denoted $(\lambda_1, \lambda_2, \dots, \lambda_n)$, are the roots of the *characteristic polynomial*:

$$p(\lambda) \stackrel{\text{def}}{=} \det(A - \lambda \mathbb{1}) = 0.$$

Calculate this determinant and we obtain an expression involving the entries a_{ij} and the variable λ . If A is an $n \times n$ matrix, the characteristic polynomial is a polynomial of degree n in λ .

We denote the list of eigenvalues as $\text{eig}(A) = (\lambda_1, \lambda_2, \dots, \lambda_n)$. If λ_i is a repeated root of the characteristic polynomial $p(\lambda)$, it's called a *degenerate* eigenvalue. For example, the identity matrix $\mathbb{1} \in \mathbb{R}^{2 \times 2}$ has the characteristic polynomial $p_{\mathbb{1}}(\lambda) = (\lambda - 1)^2$, which has a repeated root at $\lambda = 1$. We say the eigenvalue $\lambda = 1$ is *degenerate* and has *algebraic multiplicity* 2. It's important to keep track of degenerate eigenvalues, so we specify the multiplicity of an eigenvalue by repeatedly including it in the list of eigenvalues: $\text{eig}(\mathbb{1}) = (\lambda_1, \lambda_2) = (1, 1)$.

Eigenvectors

The *eigenvectors* associated with eigenvalue λ_i of matrix A are the vectors in the *null space* of the matrix $(A - \lambda_i \mathbb{1})$.

To find the eigenvectors associated with the eigenvalue λ_i , you need to solve for the components $e_{\lambda,x}$ and $e_{\lambda,y}$ of the vector $\vec{e}_{\lambda} = (e_{\lambda,x}, e_{\lambda,y})$ that satisfies the equation

$$A\vec{e}_{\lambda} = \lambda\vec{e}_{\lambda},$$

or equivalently,

$$(A - \lambda \mathbb{1})\vec{e}_{\lambda} = 0 \quad \Leftrightarrow \quad \begin{bmatrix} a_{11} - \lambda & a_{12} \\ a_{21} & a_{22} - \lambda \end{bmatrix} \begin{bmatrix} e_{\lambda,x} \\ e_{\lambda,y} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

You previously solved this type of problem when you learned how to compute the null space of a matrix.

If λ_i is a repeated root (*degenerate* eigenvalue), the null space of $(A - \lambda_i \mathbb{1})$ could contain multiple eigenvectors. The dimension of the null space of $(A - \lambda_i \mathbb{1})$ is called the *geometric multiplicity* of the eigenvalue λ_i .

Eigendecomposition

If an $n \times n$ matrix A is *diagonalizable* (I promise I didn't make that word up) this means we can find n eigenvectors for that matrix. The eigenvectors that come from different eigenspaces are guaranteed to be linearly independent (see exercise E6.3). We can also pick

a set of linearly independent vectors *within* each of the degenerate eigenspaces. Combining the eigenvectors from all the eigenspaces gives us a set of n linearly independent eigenvectors, which form a *basis* for $\mathbb{R}^n$. This is the *eigenbasis*.

Let's place the n eigenvectors side by side as the columns of a matrix:

$$Q = \begin{bmatrix} | & & | \\ \vec{e}_{\lambda_1} & \cdots & \vec{e}_{\lambda_n} \\ | & & | \end{bmatrix}.$$

We can decompose A in terms of its eigenvalues and its eigenvectors:

$$A = Q\Lambda Q^{-1} = \begin{bmatrix} | & & | \\ \vec{e}_{\lambda_1} & \cdots & \vec{e}_{\lambda_n} \\ | & & | \end{bmatrix} \begin{bmatrix} \lambda_1 & \cdots & 0 \\ \vdots & \ddots & 0 \\ 0 & 0 & \lambda_n \end{bmatrix} \begin{bmatrix} & \\ Q^{-1} & \end{bmatrix}.$$

The matrix Λ is a diagonal matrix of eigenvalues, and the matrix Q is the change-of-basis matrix that contains the corresponding eigenvectors as columns.

Note that only the *direction* of each eigenvector is important and not the length. Indeed, if $\vec{e}_\lambda$ is an eigenvector (with eigenvalue λ), so is any $\alpha\vec{e}_\lambda$ for all $\alpha \in \mathbb{R}$. Thus we're free to use any multiple of the vectors $\vec{e}_{\lambda_i}$ as the columns of the matrix Q .

Example Find the eigendecomposition of the matrix:

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

In decreasing order, the eigenvalues of the matrix are:

$$\lambda_1 = 3, \quad \lambda_2 = 2, \quad \lambda_3 = 1.$$

The eigenvalues of A are the values that appear on the diagonal of Λ .

When a 3×3 matrix has three distinct eigenvalues, it is diagonalizable, since it has the same number of linearly independent eigenvectors as eigenvalues. We know the eigenvectors are linearly independent by the following reasoning. The matrix A has three different eigenvalues. Each eigenvalue is associated with at least one eigenvector, and these eigenvectors are linearly independent (see E6.3 on page 311). Recall that any set of n linearly independent vectors in $\mathbb{R}^n$ forms a basis for $\mathbb{R}^n$. Since the three eigenvectors of A are linearly independent, we have enough columns to construct a change-of-basis matrix of eigenvectors Q and use it to write $A = Q\Lambda Q^{-1}$.

To find the eigenvectors of A , solve for the null space of the matrices $(A - 3\mathbb{I})$, $(A - 2\mathbb{I})$, and $(A - \mathbb{I})$ respectively:

$$\vec{e}_{\lambda_1} = \begin{bmatrix} -1 \\ -1 \\ 2 \end{bmatrix}, \quad \vec{e}_{\lambda_2} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \quad \vec{e}_{\lambda_3} = \begin{bmatrix} -1 \\ 0 \\ 2 \end{bmatrix}.$$

Check that $A\vec{e}_{\lambda_k} = \lambda_k\vec{e}_{\lambda_k}$ for each of the vectors above. Let Q be the matrix constructed with the eigenvectors as its columns:

$$Q = \begin{bmatrix} -1 & 0 & -1 \\ -1 & 0 & 0 \\ 2 & 1 & 2 \end{bmatrix}, \quad \text{then compute} \quad Q^{-1} = \begin{bmatrix} 0 & -1 & 0 \\ 2 & 0 & 1 \\ -1 & 1 & 0 \end{bmatrix}.$$

Together with the matrix Λ , these matrices form the eigendecomposition of the matrix A :

$$A = Q\Lambda Q^{-1} = \begin{bmatrix} 1 & 2 & 0 \\ 0 & 3 & 0 \\ 2 & -4 & 2 \end{bmatrix} = \begin{bmatrix} -1 & 0 & -1 \\ -1 & 0 & 0 \\ 2 & 1 & 2 \end{bmatrix} \begin{bmatrix} 3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -1 & 0 \\ 2 & 0 & 1 \\ -1 & 1 & 0 \end{bmatrix}.$$

To find the diagonalization of A , move Q and Q^{-1} to the other side of the equation. More specifically, multiply the equation $A = Q\Lambda Q^{-1}$ by Q^{-1} on the left and by Q on the right to obtain the diagonal matrix:

$$\Lambda = Q^{-1}AQ = \begin{bmatrix} 0 & -1 & 0 \\ 2 & 0 & 1 \\ -1 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 2 & 0 \\ 0 & 3 & 0 \\ 2 & -4 & 2 \end{bmatrix} \begin{bmatrix} -1 & 0 & -1 \\ -1 & 0 & 0 \\ 2 & 1 & 2 \end{bmatrix} = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Have you noticed how time consuming it is to compute the eigendecomposition of a matrix? It's really incredible. Though we skipped some details of the calculation (including finding the solution of the characteristic polynomial and the three null space calculations), finding the eigendecomposition of a 3×3 matrix took more than a page of work. Don't be surprised if it takes hours to compute eigendecompositions on homework problems; you're not doing anything wrong. Eigendecompositions simply take a lot of work. We're dealing with a seven-syllable word here—did you really expect the process to be easy?

Actually, computing eigenvectors and eigenvalues *will* become easier with practice. After eigendecomposing a dozen matrices or so using only pen and paper, you'll be able to whip through the steps for a 3×3 matrix in about 20 minutes. That's a really good thing for

your GPA, because the probability of seeing an eigenvalue question on your linear algebra final exam is about 80%. It pays to have a mathematical edge with this stuff.

For readers learning linear algebra without the added motivational bonus of exam stress, I recommend you eigendecompose at least one matrix using only pen and paper to prove to yourself you can do it. In all other circumstances, you're better off using SymPy to create a Matrix object; then calling its `eigenvals()` method to find the eigenvalues, or its `eigenvects()` method to find both its eigenvalues and eigenvectors.

Explanations

Yes, we've got some explaining to do.

Eigenspaces

Recall the definition of the *null space* of a matrix A :

$$\mathcal{N}(A) \stackrel{\text{def}}{=} \{\vec{v} \in \mathbb{R}^n \mid A\vec{v} = \vec{0}\}.$$

The *dimension* of the null space is the number of linearly independent vectors in the null space.

Example If the matrix A sends exactly two linearly independent vectors $\vec{v}$ and $\vec{w}$ to the zero vector, $A\vec{v} = \vec{0}$, $A\vec{w} = \vec{0}$, then its null space is two-dimensional. We can always choose the vectors $\vec{v}$ and $\vec{w}$ to be orthogonal ($\vec{v} \cdot \vec{w} = 0$) and thus obtain an *orthogonal basis* for the null space.

Each eigenvalue λ_i is associated with an *eigenspace*. The eigenspace E_{λ_i} is the null space of the matrix $(A - \lambda_i \mathbb{I})$:

$$E_{\lambda_i} \stackrel{\text{def}}{=} \mathcal{N}(A - \lambda_i \mathbb{I}) = \{\vec{v} \in \mathbb{R}^n \mid (A - \lambda_i \mathbb{I})\vec{v} = \vec{0}\}.$$

Every eigenspace contains at least one nonzero eigenvector. For degenerate eigenvalues (repeated roots of the characteristic polynomial), the null space of $(A - \lambda_i \mathbb{I})$ can contain multiple eigenvectors.

Change-of-basis matrix

The matrix Q is a *change-of-basis* matrix. Given a vector expressed in the eigenbasis $[\vec{v}]_{B_\lambda} = (v'_1, v'_2, v'_3)^T_{B_\lambda} = v'_1 \vec{e}_{\lambda_1} + v'_2 \vec{e}_{\lambda_2} + v'_3 \vec{e}_{\lambda_3}$, we can use the matrix Q to convert it to coordinates in the standard basis $[\vec{v}]_{B_s} = (v_1, v_2, v_3)^T_{B_s} = v_1 \hat{i} + v_2 \hat{j} + v_3 \hat{k}$ as follows:

$$[\vec{v}]_{B_s} = Q[\vec{v}]_{B_\lambda} = B_s[\mathbb{I}]_{B_\lambda}[\vec{v}]_{B_\lambda}.$$

The change of basis in the other direction is given by the inverse matrix:

$$[\vec{v}]_{B_\lambda} = Q^{-1}[\vec{v}]_{B_s} = B_\lambda[\mathbb{1}]_{B_s}[\vec{v}]_{B_s}.$$

Interpretation

The eigendecomposition $A = Q\Lambda Q^{-1}$ allows us to interpret the action of A on an arbitrary input vector $\vec{v}$ as the following three steps:

$$[\vec{w}]_{B_s} = B_s[A]_{B_s}[\vec{v}]_{B_s} = Q\Lambda Q^{-1}[\vec{v}]_{B_s} = B_s[\mathbb{1}]_{B_\lambda} B_\lambda[\Lambda]_{B_\lambda} \underbrace{B_\lambda[\mathbb{1}]_{B_s}[\vec{v}]_{B_s}}_1.$$

$\underbrace{\hspace{10em}}_2$
 $\underbrace{\hspace{10em}}_3$

1. In the first step, we convert the vector $\vec{v}$ from the standard basis to the eigenbasis of A .
2. In the second step, the action of A on vectors expressed with respect to its eigenbasis corresponds to a multiplication by the diagonal matrix Λ .
3. In the third step, we convert the output $\vec{w}$ from the eigenbasis back to the standard basis.

Another way to interpret these three steps is to say that, deep down inside, the matrix A is actually the diagonal matrix Λ . To see the diagonal form of the matrix, we must express the input vectors with respect to the *eigenbasis*:

$$[\vec{w}]_{B_\lambda} = B_\lambda[\Lambda]_{B_\lambda}[\vec{v}]_{B_\lambda}.$$

It's extremely important you understand the meaning of the equation $A = Q\Lambda Q^{-1}$ intuitively in terms of the three-step procedure. To help understand the three-step procedure, we'll analyze in detail what happens when we multiply A by one of its eigenvectors. Let's pick $\vec{e}_{\lambda_1}$ and verify the equation $A\vec{e}_{\lambda_1} = Q\Lambda Q^{-1}\vec{e}_{\lambda_1} = \lambda_1\vec{e}_{\lambda_1}$ by following the vector through the three steps:

$$\begin{aligned} B_s[A]_{B_s}[\vec{e}_{\lambda_1}]_{B_s} &= Q\Lambda Q^{-1}[\vec{e}_{\lambda_1}]_{B_s} \\ &= B_s[\mathbb{1}]_{B_\lambda} B_\lambda[\Lambda]_{B_\lambda} \underbrace{B_\lambda[\mathbb{1}]_{B_s}[\vec{e}_{\lambda_1}]_{B_s}}_{(1,0,\dots)_{B_\lambda}^T} \\ &= \underbrace{(\lambda_1, 0, \dots)_{B_\lambda}^T}_{\lambda_1[\vec{e}_{\lambda_1}]_{B_s}} \end{aligned}$$

In the first step, we convert the vector $[\vec{e}_{\lambda_1}]_{B_s}$ to the eigenbasis and obtain $(1, 0, \dots, 0)_{B_\lambda}^T$. The second step results in $(\lambda_1, 0, \dots, 0)_{B_\lambda}^T$, because multiplying Λ by the vector $(1, 0, \dots, 0)_{B_\lambda}^T$ selects the value in the first column of Λ . For the third step, we convert $(\lambda_1, 0, \dots, 0)_{B_\lambda}^T = \lambda_1(1, 0, \dots, 0)_{B_\lambda}^T$ back to the standard basis to obtain $\lambda_1[\vec{e}_{\lambda_1}]_{B_s}$. Boom!

Invariant properties of matrices

The determinant and the trace of a matrix are strictly functions of the eigenvalues. The determinant of A is the product of its eigenvalues:

$$\det(A) = |A| = \prod_i \lambda_i = \lambda_1 \lambda_2 \cdots \lambda_n,$$

and the trace is the sum of the eigenvalues:

$$\text{Tr}(A) = \sum_i a_{ii} = \sum_i \lambda_i = \lambda_1 + \lambda_2 + \cdots \lambda_n.$$

The above equations are true because

$$|A| = |Q\Lambda Q^{-1}| = |Q||\Lambda||Q^{-1}| = |Q||Q^{-1}||\Lambda| = \frac{|Q|}{|Q|}|\Lambda| = |\Lambda| = \prod_i \lambda_i,$$

and

$$\text{Tr}(A) = \text{Tr}(Q\Lambda Q^{-1}) = \text{Tr}(\Lambda Q^{-1}Q) = \text{Tr}(\Lambda \mathbb{1}) = \text{Tr}(\Lambda) = \sum_i \lambda_i.$$

The first equation follows from the properties of determinants: $|AB| = |A||B|$ and $|A^{-1}| = \frac{1}{|A|}$ (see page 150). The second equation follows from the cyclic property of the trace operator $\text{Tr}(ABC) = \text{Tr}(BCA)$ (see page 149).

In fact, the above calculations are true for any *similarity transformation*. Recall that a similarity transformation is a change-of-basis calculation in which a matrix A gets multiplied by an invertible matrix P from the left and by the inverse P^{-1} from the right: $A' = PAP^{-1}$. The determinant and the trace of a matrix are *invariant* properties under similarity transformations—they don't depend on the choice of basis.

Relation to invertibility

Let's briefly revisit three of the equivalent conditions we stated in the invertible matrix theorem. For a matrix $A \in \mathbb{R}^{n \times n}$, the following statements are equivalent:

- A is invertible
- $|A| \neq 0$
- The null space contains only the zero vector $\mathcal{N}(A) = \{\vec{0}\}$

The formula $|A| = \lambda_1 \lambda_2 \cdots \lambda_n$ reveals why the last two statements are equivalent. If $|A| \neq 0$, none of the λ_i s are zero (if one of the eigenvalues is zero, the whole product is zero). We know $\lambda = 0$ is *not* an eigenvalue of A , which means there exists no vector $\vec{v}$ such that $A\vec{v} = 0\vec{v} = \vec{0}$. Therefore, there are no vectors in the null space $\mathcal{N}(A)$. We can also follow this reasoning in the other direction. If the null space of A is empty, then there is no nonzero vector $\vec{v}$ such that $A\vec{v} = 0\vec{v} = \vec{0}$, which means $\lambda = 0$ is not an eigenvalue of A ; hence the product $\lambda_1 \lambda_2 \cdots \lambda_n \neq 0$.

However, if there exists a nonzero vector $\vec{v}$ such that $A\vec{v} = \vec{0}$, then A has a non-empty null space and $\lambda = 0$ is an eigenvalue of A ; thus $|A| = \lambda_1 \lambda_2 \cdots \lambda_n = 0$.

Eigendecomposition for normal matrices

A matrix A is *normal* if it satisfies the equation $A^\top A = AA^\top$. All normal matrices are diagonalizable, and the change-of-basis matrix Q can be chosen to be an *orthogonal* matrix O .

The eigenvectors corresponding to different eigenvalues of a normal matrix are *orthogonal*. Furthermore, we can choose the eigenvectors within the same eigenspace to be orthogonal. By collecting the eigenvectors from all eigenspaces of the matrix $A \in \mathbb{R}^{n \times n}$, it is possible to obtain a basis $\{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ of orthogonal eigenvectors:

$$\vec{e}_i \cdot \vec{e}_j = \begin{cases} \|\vec{e}_i\|^2 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

Normalizing these vectors gives a set of *orthonormal* eigenvectors $\{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$ that form a basis for the space $\mathbb{R}^n$:

$$\hat{e}_i \cdot \hat{e}_j = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

Consider now the matrix O constructed by using these orthonormal vectors as the columns:

$$O = \begin{bmatrix} | & & | \\ \hat{e}_1 & \cdots & \hat{e}_n \\ | & & | \end{bmatrix}.$$

The matrix O is an *orthogonal* matrix, meaning it satisfies $OO^\top = I = O^\top O$. In other words, the inverse of O is obtained by taking

the transpose O^T . To see how this works, consider the following product:

$$O^T O = \begin{bmatrix} - & \hat{e}_1 & - \\ & \vdots & \\ - & \hat{e}_n & - \end{bmatrix} \begin{bmatrix} | & & | \\ \hat{e}_1 & \cdots & \hat{e}_n \\ | & & | \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & 1 \end{bmatrix} = \mathbb{1}.$$

Each of the ones on the diagonal arises from taking the dot product of a unit-length eigenvector with itself. The off-diagonal entries are zero because the vectors are orthogonal. By definition, the inverse O^{-1} is the matrix, which gives $\mathbb{1}$ when multiplied by O , so we have $O^{-1} = O^T$.

Using the orthogonal matrix O and its inverse O^T , we can write the eigendecomposition of a matrix A as

$$A = O \Lambda O^{-1} = O \Lambda O^T = \begin{bmatrix} | & & | \\ \hat{e}_1 & \cdots & \hat{e}_n \\ | & & | \end{bmatrix} \begin{bmatrix} \lambda_1 & \cdots & 0 \\ \vdots & \ddots & 0 \\ 0 & 0 & \lambda_n \end{bmatrix} \begin{bmatrix} - & \hat{e}_1 & - \\ & \vdots & \\ - & \hat{e}_n & - \end{bmatrix}.$$

The key advantage of using an orthogonal matrix O in the diagonalization procedure is that computing its inverse becomes a trivial task: $O^{-1} = O^T$. The class of normal matrices enjoys a special status by virtue of being diagonalizable by orthogonal matrices.

Discussion

Non-diagonalizable matrices

Not all matrices are diagonalizable. For example, the matrix

$$B = \begin{bmatrix} 3 & 1 \\ 0 & 3 \end{bmatrix}$$

has $\lambda = 3$ as a repeated eigenvalue, but the null space of the matrix $(B - 3\mathbb{1})$ contains only one eigenvector: $(1, 0)^T$. The matrix B has a single eigenvector in the eigenspace $\lambda = 3$. To describe this situation using precise mathematical terminology, we say the *algebraic multiplicity* of the eigenvalue $\lambda = 3$ is two, but the *geometric multiplicity* of the eigenvalue is one.

The matrix B is a 2×2 matrix with a single eigenvector. Since we're one eigenvector short, we can't construct the diagonalizing change-of-basis matrix Q . We say the matrix has *deficient* geometric multiplicity, meaning it doesn't have a full set of eigenvectors. Therefore, B is not diagonalizable.

Matrix power series

One of the most useful concepts of calculus is the idea that functions can be represented as Taylor series. The Taylor series of the exponential function $f(x) = e^x$ is

$$e^x = \sum_{k=0}^{\infty} \frac{x^k}{k!} = 1 + x + \frac{x^2}{2} + \frac{x^3}{3!} + \frac{x^4}{4!} + \frac{x^5}{5!} + \cdots.$$

Nothing stops us from using the same Taylor series expression to define the exponential function of a matrix:

$$e^A = \sum_{k=0}^{\infty} \frac{A^k}{k!} = 1 + A + \frac{A^2}{2} + \frac{A^3}{3!} + \frac{A^4}{4!} + \frac{A^5}{5!} + \cdots.$$

Okay, there *is* one thing stopping us—we need to compute an infinite sum of progressively larger matrix powers. Remember how we used the diagonalization of $A = Q\Lambda Q^{-1}$ to write A^{77} as $Q\Lambda^{77}Q^{-1}$? We can apply that trick here to obtain the exponential of a matrix in a much simpler form:

$$\begin{aligned} e^A &= \sum_{k=0}^{\infty} \frac{A^k}{k!} = \sum_{k=0}^{\infty} \frac{(Q\Lambda Q^{-1})^k}{k!} \\ &= \sum_{k=0}^{\infty} \frac{Q\Lambda^k Q^{-1}}{k!} \\ &= Q \left[\sum_{k=0}^{\infty} \frac{\Lambda^k}{k!} \right] Q^{-1} \\ &= Q \left(1 + \Lambda + \frac{\Lambda^2}{2} + \frac{\Lambda^3}{3!} + \frac{\Lambda^4}{4!} + \cdots \right) Q^{-1} \\ &= Q e^{\Lambda} Q^{-1} \\ &= \begin{bmatrix} & Q & \begin{bmatrix} e^{\lambda_1} & \cdots & 0 \\ \vdots & \ddots & 0 \\ 0 & 0 & e^{\lambda_n} \end{bmatrix} \\ & & Q^{-1} \end{bmatrix}. \end{aligned}$$

We can use this approach to define “matrix functions” of the form

$$F : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$$

as Taylor series of matrices. Computing the matrix function $F(A)$ on an input matrix $A = Q\Lambda Q^{-1}$ is equivalent to computing the function f of each of the eigenvalues of the matrix: $F(A) = Q f(\Lambda) Q^{-1}$.

Review

We learned how to decompose matrices in terms of their eigenvalues and eigenvectors. The fundamental equation is $A\vec{e}_\lambda = \lambda\vec{e}_\lambda$, where the vector $\vec{e}_\lambda$ is an *eigenvector* of the matrix A , and the number λ is an *eigenvalue* of A .

The characteristic polynomial is derived from a simple manipulation of the eigenvalue equation:

$$\begin{aligned} A\vec{e}_\lambda &= \lambda\vec{e}_\lambda \\ A\vec{e}_\lambda - \lambda\vec{e}_\lambda &= 0 \\ (A - \lambda\mathbb{1})\vec{e}_\lambda &= 0. \end{aligned}$$

For this equation to be satisfied, the vector $\vec{e}_\lambda$ must be in the *null space* of $(A - \lambda\mathbb{1})$. The problem of finding the eigenvalues reduces to finding the values of λ for which the matrix $(A - \lambda\mathbb{1})$ has a non-empty null space. Recall that a matrix has a non-empty null space if and only if it is not invertible. The easiest way to check if a matrix is invertible is to compute the determinant: $|A - \lambda\mathbb{1}| = 0$.

Because multiple eigenvalues and eigenvectors may satisfy this equation, we keep a list of eigenvalues $(\lambda_1, \lambda_2, \dots, \lambda_n)$ and corresponding eigenvectors $\{\vec{e}_{\lambda_1}, \vec{e}_{\lambda_2}, \dots\}$. The eigendecomposition of the matrix is $A = Q\Lambda Q^{-1}$, where Q is the matrix with eigenvectors as columns, and Λ contains the eigenvalues on the diagonal.

Applications

Many scientific methods use the eigendecomposition of a matrix as a building block. For instance:

- In statistics, the *principal component analysis* technique aims to uncover the dominant cause of the variation in datasets by eigendecomposing the *covariance matrix*—a matrix computed from the dataset.
- Google’s original PageRank algorithm for ranking webpages by “importance” can be explained as the search for an eigenvector of a matrix. The matrix contains information about all the hyperlinks that exist between webpages (see Section 8.3).
- In quantum mechanics, the energy of a system is described by the Hamiltonian operator. The eigenvalues of the Hamiltonian are the possible energy levels the system can have.

Analyzing a matrix in terms of its eigenvalues and its eigenvectors is a powerful technique to “see inside” a matrix and understand what the matrix does. In the next section, we’ll analyze several different

types of matrices and discuss their properties in terms of their eigenvalues.

Links

[Good visual examples of eigenvectors from Wikipedia]

http://en.wikipedia.org/wiki/Eigenvalues_and_eigenvectors

[Eigenvectors and eigenvalues explained by 3Blue1Brown]

<https://youtube.com/watch?v=PFDu9oVAE-g>

Exercises

E6.1 Find the characteristic polynomial and the eigenvalues of each of the following matrices:

$$\text{a) } \begin{bmatrix} 10 & -9 \\ 4 & -2 \end{bmatrix} \quad \text{b) } \begin{bmatrix} 1 & 2 \\ 4 & 3 \end{bmatrix} \quad \text{c) } \begin{bmatrix} 0 & 3 \\ 7 & 0 \end{bmatrix} \quad \text{d) } \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \quad \text{e) } \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

E6.2 Find the characteristic polynomial, the eigenvalues, and the eigenvectors of the matrix:

$$A = \begin{bmatrix} 3 & -2 & 0 \\ -2 & 3 & 0 \\ 0 & 0 & 5 \end{bmatrix}.$$

E6.3 Suppose A is a 2×2 matrix with two distinct eigenvalues, λ_1 and λ_2 . Prove that the eigenvectors that correspond to the different eigenvalues are linearly independent.

E6.4 Let λ be an eigenvalue of A and let $\vec{e}_\lambda$ be the corresponding eigenvector. Show that λ^2 is an eigenvalue of A^2 .

E6.5 Suppose λ is an eigenvalue of the invertible matrix A with corresponding eigenvector $\vec{e}_\lambda$. Show that λ^{-1} is an eigenvalue of the inverse matrix A^{-1} .

E6.6 Find the values of α and β so the matrix $A = \begin{bmatrix} 0 & \alpha \\ 1 & \beta \end{bmatrix}$ will have eigenvalues 1 and 3.

E6.7 Consider the matrix $L = \begin{bmatrix} 3 & 2 \\ 4 & 1 \end{bmatrix}$. Which of the following vectors are eigenvectors of L ?

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \quad \begin{bmatrix} 1 \\ -2 \end{bmatrix}, \quad \begin{bmatrix} 2 \\ -1 \end{bmatrix}.$$

6.2 Special types of matrices

Mathematicians just love to categorize things. Conveniently for us, they've categorized certain types of matrices. Rather than embarking on a verbose explanation of the properties of a matrix, such as

I have this matrix A whose rows are perpendicular vectors and when you multiply any vector by this matrix it doesn't change the length of the vector but just kind of rotates it---

it's much simpler to refer to the categorization by saying,

Let A be an orthogonal matrix.

Most advanced science textbooks and research papers routinely use terminology like "diagonal matrix," "symmetric matrix," and "orthogonal matrix," so make sure you're familiar with these concepts.

This section will also review and reinforce what we learned about linear transformations. Recall that we can think of the matrix-vector product $A\vec{x}$ as applying a linear transformation T_A to an input vector $\vec{x}$. Therefore, each of the special matrices discussed here also corresponds to a special type of linear transformation. Keep this dual correspondence in mind because we'll use the same terminology to describe matrices *and* linear transformations.

Notation

- $\mathbb{R}^{m \times n}$: the set of $m \times n$ matrices
- $A, B, C, O, P, Q, \dots$: typical names for matrices
- a_{ij} : the entry in the i^{th} row and j^{th} column of the matrix A
- A^T : the transpose of the matrix A
- A^{-1} : the inverse of the matrix A
- $\lambda_1, \lambda_2, \dots$: the *eigenvalues* of the matrix A . For each eigenvalue λ_i there is at least one associated *eigenvector* $\vec{e}_{\lambda_i}$ that obeys the equation $A\vec{e}_{\lambda_i} = \lambda_i\vec{e}_{\lambda_i}$. Multiplying the matrix A by its eigenvectors $\vec{e}_{\lambda_i}$ is the same as scaling $\vec{e}_{\lambda_i}$ by λ_i .

Diagonal matrices

Diagonal matrices contain entries on the diagonal and zeros everywhere else. For example:

$$A = \begin{bmatrix} a_{11} & 0 & 0 \\ 0 & a_{22} & 0 \\ 0 & 0 & a_{33} \end{bmatrix}.$$

A diagonal matrix A satisfies $a_{ij} = 0$, if $i \neq j$. The eigenvalues of a diagonal matrix are $\lambda_i = a_{ii}$.

Symmetric matrices

A matrix A is symmetric if and only if

$$A^T = A, \quad \text{or equivalently if } a_{ij} = a_{ji}, \text{ for all } i, j.$$

The eigenvalues of symmetric matrices are real numbers, and the eigenvectors can be chosen to be mutually orthogonal.

Given any matrix $B \in \mathbb{R}^{m \times n}$, the product of B with its transpose $B^T B$ is always a symmetric matrix.

Upper triangular matrices

Upper triangular matrices have zero entries below the main diagonal:

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ 0 & a_{22} & a_{23} \\ 0 & 0 & a_{33} \end{bmatrix}, \quad a_{ij} = 0, \quad \text{if } i > j.$$

For a *lower* triangular matrix, all the entries *above* the diagonal are zeros: $a_{ij} = 0$, if $i < j$.

Identity matrix

The identity matrix is denoted $\mathbb{1}$ or $\mathbb{1}_n \in \mathbb{R}^{n \times n}$ and plays the role of multiplication by the number 1 for matrices: $\mathbb{1}A = A\mathbb{1} = A$. The identity matrix is diagonal with ones on the diagonal:

$$\mathbb{1}_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Any vector $\vec{v} \in \mathbb{R}^3$ is an eigenvector of the identity matrix with eigenvalue $\lambda = 1$.

Orthogonal matrices

A matrix $O \in \mathbb{R}^{n \times n}$ is *orthogonal* if it satisfies $OO^T = \mathbb{1} = O^T O$. In other words, the inverse of an orthogonal matrix O is obtained by taking its transpose: $O^{-1} = O^T$.

Multiplication by an orthogonal matrix preserves lengths. Consider the matrix-vector product $O\vec{v} = \vec{w}$. The length of a vector before the multiplication is $\|\vec{v}\| = \sqrt{\vec{v} \cdot \vec{v}}$. The length of a vector after the multiplication is

$$\|\vec{w}\| = \sqrt{\vec{w} \cdot \vec{w}} = \sqrt{(O\vec{v})^T(O\vec{v})} = \sqrt{\vec{v}^T O^T O \vec{v}}.$$

The second equality follows from the interpretation of the dot product as a matrix product $\vec{u} \cdot \vec{v} = \vec{u}^T \vec{v}$. The third equality follows from the properties of matrix transpose $(AB)^T = B^T A^T$.

When O is an orthogonal matrix, we can substitute $O^T O = \mathbb{1}$ in the above expression to establish $\|\vec{w}\| = \sqrt{\vec{v}^T \mathbb{1} \vec{v}} = \|\vec{v}\|$, which shows that multiplication by an orthogonal matrix is a *length preserving* operation.

The eigenvalues of an orthogonal matrix have *unit* length, but can in general be complex numbers $\lambda = e^{i\theta} \in \mathbb{C}$. The determinant of an orthogonal matrix is either one or negative one $|O| \in \{-1, 1\}$.

You can visualize orthogonal matrices by thinking of their columns as a set of vectors that form an orthonormal basis for $\mathbb{R}^n$:

$$O = \begin{bmatrix} | & & | \\ \hat{e}_1 & \cdots & \hat{e}_n \\ | & & | \end{bmatrix} \quad \text{such that} \quad \hat{e}_i \cdot \hat{e}_j = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

You can verify the matrix O is orthogonal by computing $O^T O = \mathbb{1}$. The orthogonal matrix O is a change-of-basis matrix from the standard basis to the column basis $\{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$.

Everything stated above about multiplication by an orthogonal matrix also applies to orthogonal transformations $T_O : \mathbb{R}^n \rightarrow \mathbb{R}^n$ because of the equivalence $O\vec{v} = \vec{w} \Leftrightarrow T_O(\vec{v}) = \vec{w}$.

The set of orthogonal matrices contains three special cases: *rotations* matrices, *reflection* matrices, and *permutation* matrices.

Rotation matrices

A rotation matrix takes the standard basis $\{\hat{i}, \hat{j}, \hat{k}\}$ to a rotated basis $\{\hat{e}_1, \hat{e}_2, \hat{e}_3\}$. Consider an example in $\mathbb{R}^2$. The counterclockwise rotation by the angle θ is given by the matrix

$$R_\theta = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}.$$

The matrix R_θ takes $\hat{i} = (1, 0)^\top$ to $(\cos \theta, \sin \theta)^\top$, and $\hat{j} = (0, 1)^\top$ to $(-\sin \theta, \cos \theta)^\top$.

As another example, consider the rotation by the angle θ around the x -axis in $\mathbb{R}^3$:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{bmatrix}.$$

This rotation is entirely in the yz -plane, so the x -component of a vector multiplying this matrix remains unchanged.

The determinant of a rotation matrix is equal to one. The eigenvalues of rotation matrices are complex numbers with unit magnitude.

Reflections

If the determinant of an orthogonal matrix O is equal to negative one, we say it is *mirrored orthogonal*. For example, the reflection through the line with direction vector $(\cos \theta, \sin \theta)$ is given by:

$$R = \begin{bmatrix} \cos(2\theta) & \sin(2\theta) \\ \sin(2\theta) & -\cos(2\theta) \end{bmatrix}.$$

A reflection matrix always has at least one eigenvalue equal to negative one, which corresponds to the direction perpendicular to the axis of reflection.

Permutation matrices

Permutation matrices are another important class of orthogonal matrices. The action of a permutation matrix is simply to change the *order* of the vector components. For example, the permutation $\pi : \hat{e}_1 \rightarrow \hat{e}'_1, \hat{e}_2 \rightarrow \hat{e}'_3, \hat{e}_3 \rightarrow \hat{e}'_2$ can be represented as the matrix

$$M_\pi = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}.$$

An $n \times n$ permutation matrix contains n ones in n different columns and zeros everywhere else.

The *sign* of a permutation corresponds to the determinant $\det(M_\pi)$. We say that permutation π is *even* if $\det(M_\pi) = +1$ and *odd* if $\det(M_\pi) = -1$.

Positive matrices

A matrix $P \in \mathbb{R}^{n \times n}$ is *positive semidefinite* if

$$\vec{v}^T P \vec{v} \geq 0 \quad \text{for all } \vec{v} \in \mathbb{R}^n.$$

The eigenvalues of a positive semidefinite matrix are all nonnegative $\lambda_i \geq 0$.

If instead the matrix P obeys the strict inequality $\vec{v}^T P \vec{v} > 0$ for all $\vec{v} \in \mathbb{R}^n$, we say the matrix P is *positive definite*. The eigenvalues of positive definite matrices are strictly greater than zero $\lambda_i > 0$.

Projection matrices

The defining property of a projection matrix is that it can be applied multiple times without changing the result:

$$\Pi = \Pi^2 = \Pi^3 = \Pi^4 = \Pi^5 = \dots$$

A projection has two eigenvalues: one and zero. The space S that is left invariant by the projection Π_S corresponds to the eigenvalue $\lambda = 1$. The orthogonal complement $S^\perp$ corresponds to the eigenvalue $\lambda = 0$ and consists of vectors that get annihilated by Π_S . The space $S^\perp$ is the null space of Π_S .

Normal matrices

The matrix $A \in \mathbb{R}^{n \times n}$ is *normal* if it obeys $A^T A = A A^T$. If A is normal, it has the following properties:

- $\vec{v}$ is an eigenvector of A if and only if $\vec{v}$ is an eigenvector of A^T .
- For all vectors $\vec{v}$ and $\vec{w}$ and a normal transformation A , we have

$$(A\vec{v}) \cdot (A\vec{w}) = (A\vec{v})^T (A\vec{w}) = \vec{v}^T A^T A \vec{w} = \vec{v}^T A A^T \vec{w} = (A^T \vec{v}) \cdot (A^T \vec{w}).$$

- The matrix A has a full set of linearly independent eigenvectors. Eigenvectors corresponding to distinct eigenvalues are orthogonal, and eigenvectors from the same eigenspace can be chosen to be mutually orthogonal.

Every normal matrix is diagonalizable by an orthogonal matrix O . The eigendecomposition of a normal matrix is written as $A = O \Lambda O^T$, where O is orthogonal and Λ is diagonal.

Orthogonal ($O^T O = \mathbb{1}$) and symmetric ($A^T = A$) matrices are normal matrices, since $O^T O = \mathbb{1} = O O^T$ and $A^T A = A^2 = A A^T$.

Exercises

E6.8 Find the determinants and inverses of these triangular matrices:

$$A = \begin{bmatrix} 1 & 4 & 15 \\ 0 & 5 & 5 \\ 0 & 0 & 3 \end{bmatrix}, \quad B = \begin{bmatrix} x & 0 \\ y & z \end{bmatrix}, \quad C = \begin{bmatrix} \frac{1}{5} & 0 \\ 0 & 5 \end{bmatrix}.$$

E6.9 Is the matrix $\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$ symmetric?

E6.10 In this section we learned about different types of matrices: diagonal, triangular, positive (semi)definite, symmetric, and orthogonal matrices. What types of matrices are these?

$$\text{a) } A = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 4 & 1 \\ 0 & 0 & 7 \end{bmatrix} \quad \text{b) } B = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{6}} & -\frac{1}{\sqrt{3}} \\ 0 & \frac{2}{\sqrt{6}} & -\frac{1}{\sqrt{3}} \end{bmatrix} \quad \text{c) } C = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Hint: Compute the matrices' eigenvalues to analyse their properties.

6.3 Abstract vector spaces

You can apply your knowledge of vectors more generally to other vector-like mathematical objects. For example, polynomials behave similarly to vectors. To add two polynomials $P(x)$ and $Q(x)$, we add together the coefficients of each power of x —the same way vectors are added component by component.

In this section, we'll learn how to use the terminology and concepts associated with vectors to study other mathematical objects. In particular, we'll see that notions such as *linear independence*, *basis*, and *dimension* can be applied to mathematical objects like matrices, polynomials, and functions. We'll use the notation $\mathbf{v}$ for describing abstract vectors as opposed to the usual $\vec{v}$ used for ordinary vectors.

Definitions

An abstract vector space $(V, F, +, \cdot)$ consists of four things:

- A set of vector-like objects $V = \{\mathbf{u}, \mathbf{v}, \dots\}$
- A field F of scalar numbers, usually $F = \mathbb{R}$
- An addition operation “+” for elements of V that dictates how to add vectors: $\mathbf{u} + \mathbf{v}$
- A scalar multiplication operation “ $\cdot$ ” for scaling a vector by an element of the field. Scalar multiplication is usually denoted implicitly $\alpha \mathbf{u}$ (without the dot).

A vector space satisfies the following eight axioms, for all scalars $\alpha, \beta \in F$ and all vectors $\mathbf{u}, \mathbf{v}, \mathbf{w} \in V$:

1. $\mathbf{u} + (\mathbf{v} + \mathbf{w}) = (\mathbf{u} + \mathbf{v}) + \mathbf{w}$ (associativity of addition)
2. $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$ (commutativity of addition)
3. There exists a zero vector $\mathbf{0} \in V$, such that $\mathbf{u} + \mathbf{0} = \mathbf{0} + \mathbf{u} = \mathbf{u}$ for all $\mathbf{u} \in V$.
4. For every $\mathbf{u} \in V$, there exists an inverse element $-\mathbf{u}$ such that $\mathbf{u} + (-\mathbf{u}) = \mathbf{u} - \mathbf{u} = \mathbf{0}$.
5. $\alpha(\mathbf{u} + \mathbf{v}) = \alpha\mathbf{u} + \alpha\mathbf{v}$ (distributivity I)
6. $(\alpha + \beta)\mathbf{u} = \alpha\mathbf{u} + \beta\mathbf{u}$ (distributivity II)
7. $\alpha(\beta\mathbf{u}) = (\alpha\beta)\mathbf{u}$ (associativity of scalar multiplication)
8. There exists a unit scalar 1 such that $1\mathbf{u} = \mathbf{u}$.

If you know anything about vectors, the above properties should be familiar. Indeed, these are the standard properties for the vector space $\mathbb{R}^n$, where the field F is $\mathbb{R}$, and for which standard vector addition and scalar multiplication operations apply.

Theory

Believe it or not, we're actually done with all the theory for this section. Move along folks, there's nothing more to see here aside from the definitions above—which are restatements of the properties of vector addition and vector scaling that you've already seen before.

The only thing left to do is illustrate these concepts through some examples.

Examples

Matrices, polynomials, and functions are vector-like math objects. The following examples demonstrate how we can treat these math objects as abstract vector spaces $(V, F, +, \cdot)$.

Matrices

Consider the vector space of $m \times n$ matrices over the real numbers $\mathbb{R}^{m \times n}$. The addition operation for two matrices $A, B \in \mathbb{R}^{m \times n}$ is the usual rule of matrix addition: $(A + B)_{ij} = a_{ij} + b_{ij}$.

This vector space is mn -dimensional, which can be seen by constructing a basis for the space. The standard basis for $\mathbb{R}^{m \times n}$ consists

of mn matrices with zero entries everywhere except for a single 1 in the i^{th} row and the j^{th} column. Any matrix $A \in \mathbb{R}^{m \times n}$ can be written as a linear combination of the matrices in the standard basis.

Example The standard basis B_s for the vector space $\mathbb{R}^{2 \times 2}$ is

$$\mathbf{e}_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{e}_2 = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{e}_3 = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, \quad \mathbf{e}_4 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

Any matrix $A \in \mathbb{R}^{2 \times 2}$ can be written as a linear combination of the vectors in B_s :

$$\begin{aligned} A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} &= a_{11} \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} + a_{12} \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} + a_{21} \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} + a_{22} \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \\ &= a_{11}\mathbf{e}_1 + a_{12}\mathbf{e}_2 + a_{21}\mathbf{e}_3 + a_{22}\mathbf{e}_4. \end{aligned}$$

In other words, A can be expressed as a vector of coordinates with respect to the basis B_s : $A = (a_{11}, a_{12}, a_{21}, a_{22})_{B_s}$.

The abstract concept of a matrix $A \in \mathbb{R}^{2 \times 2}$ can be expressed as two equivalent representations. We can think of A either as an array of numbers with two columns and two rows, or as a four-dimensional vector of coordinates with respect to the basis B_s :

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} = (a_{11}, a_{12}, a_{21}, a_{22})_{B_s}.$$

We've arrived at a major *knowledge buzz* milestone: **matrices are vectors!** In precise mathematical terms, we just demonstrated the existence of an *isomorphism* between the set of 2×2 matrices and the set of four-dimensional vectors. We can add, subtract, and scale 2×2 matrices in their $\mathbb{R}^4$ representations. In the following exercises, we'll see how to compute the matrix trace operation $\text{Tr}(A)$ in terms of the vector representation.

Symmetric 2x2 matrices

Define the vector space consisting of 2×2 symmetric matrices

$$\mathbb{S}(2, 2) \stackrel{\text{def}}{=} \{A \in \mathbb{R}^{2 \times 2} \mid A = A^T\}$$

in combination with the usual matrix addition and scalar multiplication operations. We obtain an explicit basis for this space as follows:

$$\mathbf{v}_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \mathbf{v}_3 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

Any element of the vector space $S \in S(2, 2)$ can be written as a linear combination of the basis vectors:

$$\begin{aligned} S = \begin{bmatrix} a & b \\ b & c \end{bmatrix} &= a \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} + b \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + c \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \\ &= a\mathbf{v}_1 + b\mathbf{v}_2 + c\mathbf{v}_3. \end{aligned}$$

Since there are three vectors in a basis for $S(2, 2)$, the vector space $S(2, 2)$ is three-dimensional.

Note how we count the dimensions in this case. The space of 2×2 matrices is four-dimensional in general, but imposing the symmetry constraint $a_{12} = a_{21}$ eliminates one parameter, so we're left with a three-dimensional space.

Polynomials of degree n

Define the vector space $P_n(t)$ of polynomials with real coefficients and degree less than or equal to n . The “vectors” in this space are polynomials of the form

$$\mathbf{p} = a_0 + a_1x + a_2x^2 + \cdots + a_nx^n.$$

The coefficients of the polynomial $a_0, a_1, a_2, \dots, a_n$ are the components of the vector $\mathbf{p}$.

The addition of vectors $\mathbf{p}, \mathbf{q} \in P_n(t)$ is performed component-wise:

$$\begin{aligned} \mathbf{p} + \mathbf{q} &= (a_0 + a_1x + \cdots + a_nx^n) + (b_0 + b_1x + \cdots + b_nx^n) \\ &= (a_0 + b_0) + (a_1 + b_1)x + \cdots + (a_n + b_n)x^n. \end{aligned}$$

Similarly, scalar multiplication acts as you'd expect:

$$\alpha\mathbf{p} = \alpha \cdot (a_0 + a_1x + \cdots + a_nx^n) = (\alpha a_0) + (\alpha a_1)x + \cdots + (\alpha a_n)x^n.$$

The space $P_n(x)$ is $(n+1)$ -dimensional since each “vector” in this space has $n+1$ components.

Functions

Another interesting vector space is the set of functions $f: \mathbb{R} \rightarrow \mathbb{R}$ in combination with the point-wise addition and scalar multiplication operations:

$$\mathbf{f} + \mathbf{g} = (f + g)(x) = f(x) + g(x), \quad \alpha\mathbf{f} = (\alpha f)(x) = \alpha f(x).$$

The space of functions is *infinite*-dimensional.

Discussion

We've talked about bases, components, and dimensions of *abstract* vector spaces. Indeed, these notions are well-defined for any vector-like object. Though this section only discussed vector spaces with real components, we can apply the same techniques to vectors with components from any *field*. The notion of a *field* describes any number-like object for which the operations of addition, subtraction, multiplication, and division are defined. An example of another field is the set of complex numbers $\mathbb{C}$. We'll discuss the linear algebra of vectors with complex components in Section 6.7.

In the next section, we'll define an *abstract inner product* operation and use this definition to discuss concepts like orthogonality, length, and distance in abstract vector spaces.

Links

[Further discussion and examples on Wikipedia]

http://en.wikipedia.org/wiki/Vector_space

[Examples of vector spaces]

http://wikibooks.org/wiki/Linear_Algebra/Definition_and_Examples_of_Vector_Spaces

[Abstract vector spaces explained by 3Blue1Brown]

<https://youtube.com/watch?v=TgKwz5Ikpc8>

Exercises

E6.11 Find a basis for the space of 2×2 upper triangular matrices.

E6.12 Can every polynomial of degree at most two be written in the form $\alpha + \beta(x - 1) + \gamma(x - 1)^2$?

Hint: Try to express an arbitrary polynomial in this form.

6.4 Abstract inner product spaces

An inner product space is an abstract vector space $(V, \mathbb{R}, +, \cdot)$ for which we define an *abstract inner product* operation that takes pairs of vectors as inputs and produces numbers as outputs:

$$\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{R}.$$

We can use any inner product operation, as long as it satisfies the following criteria for all $\mathbf{u}, \mathbf{v}, \mathbf{v}_1, \mathbf{v}_2 \in V$ and $\alpha, \beta \in \mathbb{R}$. The inner product operation must be:

- Symmetric: $\langle \mathbf{u}, \mathbf{v} \rangle = \langle \mathbf{v}, \mathbf{u} \rangle$
- Linear: $\langle \mathbf{u}, \alpha \mathbf{v}_1 + \beta \mathbf{v}_2 \rangle = \alpha \langle \mathbf{u}, \mathbf{v}_1 \rangle + \beta \langle \mathbf{u}, \mathbf{v}_2 \rangle$
- Positive semidefinite: $\langle \mathbf{u}, \mathbf{u} \rangle \geq 0$ for all $\mathbf{u} \in V$ with $\langle \mathbf{u}, \mathbf{u} \rangle = 0$ if and only if $\mathbf{u} = \mathbf{0}$

These criteria are inspired by the properties of the standard inner product (dot product) for vectors in $\mathbb{R}^n$:

$$\langle \vec{u}, \vec{v} \rangle = \vec{u} \cdot \vec{v} = \sum_{i=1}^n u_i v_i = \vec{u}^T \vec{v}.$$

In this section, we'll extend the idea of the dot product by defining inner product operations $\langle \mathbf{u}, \mathbf{v} \rangle$ for abstract vectors $\mathbf{u}, \mathbf{v} \in V$. We'll define inner product operations for matrices $\langle A, B \rangle$, polynomials $\langle \mathbf{p}, \mathbf{q} \rangle$, and functions $\langle f, g \rangle$. These inner product operations will allow us to talk about *orthogonality* between abstract vectors,

$$\mathbf{u} \text{ and } \mathbf{v} \text{ are orthogonal} \quad \Leftrightarrow \quad \langle \mathbf{u}, \mathbf{v} \rangle = 0,$$

the *length* of an abstract vector,

$$\|\mathbf{u}\| = \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle},$$

and the *distance* between two abstract vectors,

$$d(\mathbf{u}, \mathbf{v}) = \|\mathbf{u} - \mathbf{v}\| = \sqrt{\langle (\mathbf{u} - \mathbf{v}), (\mathbf{u} - \mathbf{v}) \rangle}.$$

Let's get started.

Definitions

We'll work with vectors from an abstract vector space $(V, \mathbb{R}, +, \cdot)$ where:

- V is the set of vectors in the vector space.
- $\mathbb{R}$ is the *field* of real numbers. The components of the abstract vectors are taken from this field.
- $+$ is the addition operation defined for elements of V .
- $\cdot$ is the scalar multiplication operation between an element of the field $\alpha \in \mathbb{R}$ and a vector $\mathbf{u} \in V$.

We define a new operation called *abstract inner product* for that space:

$$\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{R}.$$

The abstract inner product takes as inputs two vectors $\mathbf{u}, \mathbf{v} \in V$ and produces real numbers as outputs: $\langle \mathbf{u}, \mathbf{v} \rangle \in \mathbb{R}$.

We define the following related quantities in terms of the inner product operation:

- $\|\mathbf{u}\| \stackrel{\text{def}}{=} \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle}$: the *norm* or *length* of an abstract vector
- $d(\mathbf{u}, \mathbf{v}) \stackrel{\text{def}}{=} \|\mathbf{u} - \mathbf{v}\|$: the *distance* between two vectors

Orthogonality

Recall that two vectors $\vec{u}, \vec{v} \in \mathbb{R}^n$ are said to be orthogonal if their dot product is zero. This follows from the geometric interpretation of the dot product:

$$\vec{u} \cdot \vec{v} = \|\vec{u}\| \|\vec{v}\| \cos \theta,$$

where θ is the *angle* between $\vec{u}$ and $\vec{v}$. Orthogonal means “at right angle with.” Indeed, if $\vec{u} \cdot \vec{v} = 0$, the angle between $\vec{u}$ and $\vec{v}$ must be 90° or 270° , since $\cos \theta = 0$ only for these two angles.

In analogy with the regular dot product, we define the notion of *orthogonality* between abstract vectors in terms of the abstract inner product:

$$\mathbf{u} \text{ and } \mathbf{v} \text{ are orthogonal} \quad \Leftrightarrow \quad \langle \mathbf{u}, \mathbf{v} \rangle = 0.$$

Translating the geometric intuition of “at 90° angle with” might not be possible for certain abstract vector spaces. For instance, what is the angle between two polynomials? Nevertheless, the fundamental notion of “perpendicular to” exists in all abstract inner product vector spaces.

Norm

Every definition of an inner product for an abstract vector space $(V, \mathbb{R}, +, \cdot)$ induces a *norm* for that vector space:

$$\|\cdot\| : V \rightarrow \mathbb{R}.$$

The norm is defined in terms of the inner product. The norm of a vector is the square root of the inner product of the vector with itself:

$$\|\mathbf{u}\| \stackrel{\text{def}}{=} \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle}.$$

The norm corresponds to the abstract notion of length for vectors. All norms must satisfy the following criteria:

- $\|\mathbf{v}\| \geq 0$ with equality if and only if $\mathbf{v} = \mathbf{0}$
- $\|k\mathbf{v}\| = |k| \|\mathbf{v}\|$
- The triangle inequality:

$$\|\mathbf{u} + \mathbf{v}\| \leq \|\mathbf{u}\| + \|\mathbf{v}\|$$

- Cauchy–Schwarz inequality:

$$|\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\| \|\mathbf{y}\|,$$

with equality if and only if $\mathbf{x}$ and $\mathbf{y}$ are linearly dependent

Norms defined in terms of a valid inner product automatically satisfy these criteria.

Distance

The distance between two points p and q in $\mathbb{R}^n$ is equal to the norm of the vector that goes from p to q : $d(p, q) = \|q - p\|$. We can similarly define a *distance* function between pairs of vectors in an abstract vector space V :

$$d : V \times V \rightarrow \mathbb{R}.$$

The distance between two abstract vectors is equal to the norm of their difference:

$$d(\mathbf{u}, \mathbf{v}) \stackrel{\text{def}}{=} \|\mathbf{u} - \mathbf{v}\| = \sqrt{\langle \mathbf{u} - \mathbf{v}, \mathbf{u} - \mathbf{v} \rangle}.$$

Distances defined in terms of a valid norm obey the following criteria:

- $d(\mathbf{u}, \mathbf{v}) = d(\mathbf{v}, \mathbf{u})$
- $d(\mathbf{u}, \mathbf{v}) \geq 0$ with equality if and only if $\mathbf{u} = \mathbf{v}$

Examples

Let's define some inner product functions for the aforementioned abstract vector spaces.

Matrix inner product

The Hilbert–Schmidt inner product for real matrices is defined in terms of the matrix transpose, matrix multiplication, and matrix trace operations:

$$\langle A, B \rangle_{\text{HS}} = \text{Tr}(A^T B).$$

We can use this inner product to talk about *orthogonality* properties of matrices. In the last section, we defined the set of 2×2 symmetric matrices

$$S(2, 2) = \{A \in \mathbb{R}^{2 \times 2} \mid A = A^T\},$$

and gave an explicit basis for this space:

$$\mathbf{v}_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \mathbf{v}_3 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

It's easy to show that these vectors are all *mutually orthogonal* with respect to the Hilbert–Schmidt inner product $\langle \cdot, \cdot \rangle_{\text{HS}}$:

$$\langle \mathbf{v}_1, \mathbf{v}_2 \rangle_{\text{HS}} = 0, \quad \langle \mathbf{v}_1, \mathbf{v}_3 \rangle_{\text{HS}} = 0, \quad \langle \mathbf{v}_2, \mathbf{v}_3 \rangle_{\text{HS}} = 0.$$

Verify these three equations by computing each inner product. Try this by hand on a piece of paper ... like right now.

The three inner product calculations of the last equation indicate that the set $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ forms an orthogonal basis for the vector space $S(2, 2)$ with respect to the inner product $\langle \cdot, \cdot \rangle_{\text{HS}}$.

Hilbert–Schmidt norm

The Hilbert–Schmidt inner product induces the Hilbert–Schmidt norm:

$$\|A\|_{\text{HS}} \stackrel{\text{def}}{=} \sqrt{\langle A, A \rangle_{\text{HS}}} = \sqrt{\text{Tr}(A^T A)} = \left[\sum_{i,j=1}^n |a_{ij}|^2 \right]^{\frac{1}{2}}.$$

We can use this norm to define an abstract notion of length for matrices. Continuing with the above example, we can obtain an **orthonormal** basis $\{\hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2, \hat{\mathbf{v}}_3\}$ for $S(2, 2)$ as follows:

$$\hat{\mathbf{v}}_1 = \mathbf{v}_1, \quad \hat{\mathbf{v}}_2 = \frac{\mathbf{v}_2}{\|\mathbf{v}_2\|_{\text{HS}}} = \frac{1}{\sqrt{2}}\mathbf{v}_2, \quad \hat{\mathbf{v}}_3 = \mathbf{v}_3.$$

Verify that $\|\hat{\mathbf{v}}_2\|_{\text{HS}} = 1$.

Function inner product

Consider two functions, $\mathbf{f} = f(t)$ and $\mathbf{g} = g(t)$, and define their inner product as follows:

$$\langle \mathbf{f}, \mathbf{g} \rangle \stackrel{\text{def}}{=} \int_{-\infty}^{\infty} f(t)g(t) dt.$$

This formula is the continuous-variable version of the inner product formula for vectors $\vec{u} \cdot \vec{v} = \sum_i u_i v_i$. Instead of a summation, we have an integral; otherwise the idea is the same since we're measuring the *overlap* between $\mathbf{f}$ and $\mathbf{g}$. The integral passes along the real line from $-\infty$ until ∞ like a zipper that brings together $f(t)$ times $g(t)$ at each point.

Example Consider the function inner product on the interval $[-1, 1]$ as defined by the formula:

$$\langle \mathbf{f}, \mathbf{g} \rangle = \int_{-1}^1 f(t)g(t) dt.$$

Verify that the following polynomials, known as the Legendre polynomials $P_n(x)$, are mutually orthogonal with respect to the above inner product:

$$\begin{aligned} P_0(x) &= 1, & P_1(x) &= x, \\ P_2(x) &= \frac{1}{2}(3x^2 - 1), & P_3(x) &= \frac{1}{2}(5x^3 - 3x). \end{aligned}$$

Generalized dot product

We can think of the regular dot product for vectors as the following vector-matrix-vector product:

$$\vec{u} \cdot \vec{v} = \vec{u}^T \vec{v} = \vec{u}^T \mathbb{1} \vec{v}.$$

More generally, we can insert any *symmetric, positive semidefinite* matrix M between the vectors and obtain a valid inner product:

$$\langle \vec{x}, \vec{y} \rangle_M \stackrel{\text{def}}{=} \vec{x}^T M \vec{y}.$$

The matrix M is called the *metric* for this inner product, and it encodes the relative contributions of the different components of the vectors to the inner product.

The requirement that M be symmetric stems from the symmetric requirement for inner products: $\langle \mathbf{u}, \mathbf{v} \rangle = \langle \mathbf{v}, \mathbf{u} \rangle$. The requirement that the matrix be positive semidefinite comes from the positive semidefinite requirement for inner products: $\langle \mathbf{u}, \mathbf{u} \rangle = \vec{u}^T M \vec{u} \geq 0$, for all $\mathbf{u} \in V$.

We can always obtain a symmetric and positive semidefinite matrix M by setting $M = A^T A$ for some matrix A . To understand why we might want to construct M in this way, recall that each matrix A implements some linear transformation $T_A(\vec{u}) = A\vec{u}$. An inner product $\langle \vec{u}, \vec{v} \rangle_M$ can be interpreted as the regular dot product in the output space of T_A :

$$\langle \vec{u}, \vec{v} \rangle_M = \vec{u}^T M \vec{v} = \vec{u}^T A^T A \vec{v} = (A\vec{u})^T (A\vec{v}) = T_A(\vec{u}) \cdot T_A(\vec{v}).$$

The notion of a generalized inner product with metric matrix M is a powerful idea with applications in advanced math areas like analysis and differential geometry. The concept of a metric also shows up in physics: when Einstein talks about masses causing space to become “curved,” he’s talking about the curvature of the metric of space-time.

Valid and invalid inner product spaces

A standard question profs like to ask on exams is to check whether a given vector space and some weird definition of an inner product operation form a valid inner product space. Recall that *any* operation can be used as the inner product, as long as it satisfies the *symmetry*, *linearity*, and *positive semidefinite* criteria. To prove an inner product operations is valid, you must show it satisfies the three criteria.

Alternatively, you can prove the vector space $(V, \mathbb{R}, +, \cdot)$ with inner product $\langle \mathbf{u}, \mathbf{v} \rangle$ is *not* a valid inner product space if you find an example of one or more $\mathbf{u}, \mathbf{v} \in V$ that don't satisfy the axioms.

Discussion

This has been another one of those sections where we learn no new linear algebra, but simply generalize notions we already know about standard vectors $\vec{v} \in \mathbb{R}^n$ to abstract vector-like objects $\mathbf{v} \in V$. You can now talk about orthogonality and norms for matrices, polynomials, and functions.

Exercises

E6.13 Prove the set of matrices $\{\mathbf{e}_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \mathbf{e}_2 = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \mathbf{e}_3 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}\}$ is an orthonormal set under the inner product $\langle A, B \rangle_{\text{HS}} = \text{Tr}(A^T B)$.

E6.14 Compute the norm of the functions $P_0(x) = 1$ and $P_1(x) = x$ with respect to the function inner product $\langle \mathbf{f}, \mathbf{g} \rangle = \int_{-1}^1 f(x)g(x) dt$ defined on the interval $[-1, 1]$. Also compute the distance $d(P_0, P_1)$.

Hint: Recall that a vector's norm is defined as $\|\mathbf{u}\| = \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle}$.

6.5 Gram–Schmidt orthogonalization

Recall what we learned in Section 4.3 about the three “quality grades” for bases: orthonormal, orthogonal, and generic, with orthonormal bases being the easiest to work with. In this section, we'll learn how to take a generic basis for an n -dimensional vector space V —that is, a set of n linearly independent vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ —and transform it into an orthonormal basis $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \dots, \hat{\mathbf{e}}_n\}$ that satisfies the conditions

$$\langle \hat{\mathbf{e}}_i, \hat{\mathbf{e}}_j \rangle = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

This procedure is known as *Gram–Schmidt orthogonalization* and is based on a sequence of projection and subtraction operations.

The discussion and procedures in this section are described in terms of vectors in an abstract inner product space. Thus, the Gram–Schmidt algorithm applies to ordinary vectors $\vec{v} \in \mathbb{R}^n$, matrices $A \in \mathbb{R}^{m \times n}$, and polynomials $\mathbf{p} \in P_n(x)$. Indeed, we can talk about orthogonality for any set of mathematical objects for which we’ve defined an inner product operation.

Definitions

- V : an n -dimensional vector space
- $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$: a generic basis for the space V
- $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$: an *orthogonal basis* for V . Each vector $\mathbf{e}_i$ is orthogonal to all other vectors: $\mathbf{e}_i \cdot \mathbf{e}_j = 0$, for $i \neq j$.
- $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \dots, \hat{\mathbf{e}}_n\}$: an *orthonormal basis* for V . An orthonormal basis is an orthogonal basis of unit vectors.

We assume the vector space V is equipped with an inner product operation:

$$\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{R}.$$

The following operations are defined in terms of the inner product:

- The *length* of a vector $\|\mathbf{v}\| = \langle \mathbf{v}, \mathbf{v} \rangle$
- The *projection* operation. The projection of the vector $\mathbf{u}$ onto the subspace spanned by the vector $\mathbf{e}$ is denoted $\Pi_{\mathbf{e}}(\mathbf{u})$ and is computed using

$$\Pi_{\mathbf{e}}(\mathbf{u}) = \frac{\langle \mathbf{u}, \mathbf{e} \rangle}{\|\mathbf{e}\|^2} \mathbf{e}.$$

- The *projection complement* of the projection $\Pi_{\mathbf{e}}(\mathbf{u})$ is the vector $\mathbf{w}$ that we must add to $\Pi_{\mathbf{e}}(\mathbf{u})$ to recover the original vector $\mathbf{u}$:

$$\mathbf{u} = \Pi_{\mathbf{e}}(\mathbf{u}) + \mathbf{w} \quad \Rightarrow \quad \mathbf{w} = \mathbf{u} - \Pi_{\mathbf{e}}(\mathbf{u}).$$

The vector $\mathbf{w}$ is orthogonal to the vector $\mathbf{e}$, $\langle \mathbf{w}, \mathbf{e} \rangle = 0$.

Orthonormal bases are nice

Recall that a *basis* for an n -dimensional vector space V is any set of n linearly independent vectors in V . The choice of basis is a big deal because we express the coordinates of vectors and matrices with respect to the basis. From a theoretical standpoint, all bases are equally good; but from a practical standpoint, orthogonal and orthonormal bases are much easier to work with.

An orthonormal basis $B = \{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3\}$ is the most useful kind of basis because the coordinates c_1, c_2 , and c_3 of a vector $\mathbf{c} = (c_1, c_2, c_3)_B$

with respect to B are obtained using three independent inner product calculations:

$$c_1 = \langle \mathbf{c}, \hat{\mathbf{e}}_1 \rangle, \quad c_2 = \langle \mathbf{c}, \hat{\mathbf{e}}_2 \rangle, \quad c_3 = \langle \mathbf{c}, \hat{\mathbf{e}}_3 \rangle.$$

We can express any vector $\mathbf{v}$ as

$$\mathbf{v} = \langle \mathbf{v}, \hat{\mathbf{e}}_1 \rangle \hat{\mathbf{e}}_1 + \langle \mathbf{v}, \hat{\mathbf{e}}_2 \rangle \hat{\mathbf{e}}_2 + \langle \mathbf{v}, \hat{\mathbf{e}}_3 \rangle \hat{\mathbf{e}}_3.$$

This formula is a generalization of the usual formula for coordinates with respect to the standard basis $\{\hat{i}, \hat{j}, \hat{k}\}$: $\vec{v} = (\vec{v} \cdot \hat{i})\hat{i} + (\vec{v} \cdot \hat{j})\hat{j} + (\vec{v} \cdot \hat{k})\hat{k}$.

Orthogonalization

The “best” kind of basis for computational purposes is an orthonormal basis like $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \dots, \hat{\mathbf{e}}_n\}$. How can we *upgrade* some general set of n linearly independent vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ into an orthonormal basis $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \dots, \hat{\mathbf{e}}_n\}$? The vectors $\{\hat{\mathbf{e}}_i\}$ must be linear combinations of the vectors $\{\mathbf{v}_i\}$, but which linear combinations should we choose?

Note the vector space V remains the same:

$$\text{span}(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n) = V = \text{span}(\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \dots, \hat{\mathbf{e}}_n).$$

However, the basis $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \dots, \hat{\mathbf{e}}_m\}$ is easier to work with.

The technical term for distilling a high-quality orthonormal basis from a low-quality basis of arbitrary vectors is called *orthogonalization*. Most of the work lies in obtaining the set of vectors $\{\mathbf{e}_i\}$ that are *orthogonal* to each other:

$$\langle \mathbf{e}_i, \mathbf{e}_j \rangle = 0, \quad \text{for all } i \neq j.$$

To convert an orthogonal basis into an orthonormal basis, divide each vector by its length: $\hat{\mathbf{e}}_i = \frac{\mathbf{e}_i}{\|\mathbf{e}_i\|}$.

It's now time to see how orthogonalization works; get ready for some Gram–Schmidt.

Gram–Schmidt orthogonalization procedure

The Gram–Schmidt orthogonalization procedure converts a basis of arbitrary vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ into an orthonormal basis $\{\hat{\mathbf{e}}_1, \dots, \hat{\mathbf{e}}_n\}$. The main idea is to take the vectors $\mathbf{v}_i$ one at a time, each time defining a new vector $\mathbf{e}_i$ as the *orthogonal complement* of $\mathbf{v}_i$ to all the previously chosen vectors $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_{i-1}$. Recall we can use the projection formula $\Pi_{\hat{\mathbf{e}}}(\mathbf{v}) = \langle \hat{\mathbf{e}}, \mathbf{v} \rangle \hat{\mathbf{e}}$ to compute the component of any vector $\mathbf{v}$ in the direction $\hat{\mathbf{e}}$.

The orthogonalization algorithm consists of n steps:

$$\begin{aligned}
 \mathbf{e}_1 &= \mathbf{v}_1 & \hat{\mathbf{e}}_1 &= \frac{\mathbf{v}_1}{\|\mathbf{v}_1\|}, \\
 \mathbf{e}_2 &= \mathbf{v}_2 - \Pi_{\hat{\mathbf{e}}_1}(\mathbf{v}_2), & \hat{\mathbf{e}}_2 &= \frac{\mathbf{e}_2}{\|\mathbf{e}_2\|}, \\
 \mathbf{e}_3 &= \mathbf{v}_3 - \Pi_{\hat{\mathbf{e}}_1}(\mathbf{v}_3) - \Pi_{\hat{\mathbf{e}}_2}(\mathbf{v}_3), & \hat{\mathbf{e}}_3 &= \frac{\mathbf{e}_3}{\|\mathbf{e}_3\|}, \\
 \mathbf{e}_4 &= \mathbf{v}_4 - \Pi_{\hat{\mathbf{e}}_1}(\mathbf{v}_4) - \Pi_{\hat{\mathbf{e}}_2}(\mathbf{v}_4) - \Pi_{\hat{\mathbf{e}}_3}(\mathbf{v}_4), & \hat{\mathbf{e}}_4 &= \frac{\mathbf{e}_4}{\|\mathbf{e}_4\|}, \\
 \vdots & & \vdots & \\
 \mathbf{e}_n &= \mathbf{v}_n - \sum_{i=1}^{n-1} \Pi_{\hat{\mathbf{e}}_i}(\mathbf{v}_n), & \hat{\mathbf{e}}_n &= \frac{\mathbf{e}_n}{\|\mathbf{e}_n\|}.
 \end{aligned}$$

In the j^{th} step of the procedure, we compute a vector $\mathbf{e}_j$ by starting from $\mathbf{v}_j$ and subtracting all the projections of $\mathbf{v}_j$ onto the previous vectors $\mathbf{e}_i$ for all $i < j$. In other words, $\mathbf{e}_j$ is the part of $\mathbf{v}_j$ that is orthogonal to all the vectors $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_{j-1}$.

This procedure is known as orthogonalization because it splits the vector space V into orthogonal subspaces $V_1, V_2, \dots, V_n$:

$$V_j = \text{span}\left(\mathbf{v} \in V \mid \mathbf{v} = \sum_{i=1}^j \alpha_i \mathbf{v}_i\right) \setminus \text{span}\left(\mathbf{v} \in V \mid \mathbf{v} = \sum_{i=1}^{j-1} \alpha_i \mathbf{v}_i\right).$$

Recall that the symbol $\setminus$ denotes the *set difference* operation. The set $A \setminus B$ consists of all elements that are in set A but not in set B .

Observe that the subspaces $V_1, V_2, \dots, V_n$ are, by construction, mutually orthogonal. Given any vector $\mathbf{u} \in V_i$ and another vector $\mathbf{v} \in V_j, j \neq i$, then $\mathbf{u} \cdot \mathbf{v} = 0$.

The vector space V is the sum of these subspaces:

$$V = V_1 \oplus V_2 \oplus V_3 \oplus \dots \oplus V_n.$$

The notation $\oplus$ means *orthogonal sum*. Each space V_j is spanned by a vector $\mathbf{e}_j$ which is orthogonal to all the V_i s, for $i < j$.

Discussion

The main point you must remember about orthogonalization is simply that it can be done. Any “low-quality” basis (a set of n linearly independent vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ in an n -dimensional space) can be converted into a “high quality” orthonormal basis $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \dots, \hat{\mathbf{e}}_n\}$ using the Gram–Schmidt procedure.

You can also perceive the Gram–Schmidt procedure as a technique for creating structure in an arbitrary vector space V . The initial description $V = \text{span}(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n)$ lacks structure. It’s just some amorphous vector space spanned by an arbitrary set of vectors. After the orthogonalization procedure, we obtain the equivalent description $V = V_1 \oplus V_2 \oplus V_3 \oplus \dots \oplus V_n$ that shows V is the direct sum of orthogonal subspaces.

In the next section, we’ll continue on our mathematical quest for structure by discussing procedures that uncover hidden structure in matrices. For example, when phrased in terms of matrices, the Gram–Schmidt orthogonalization procedure is called *QR decomposition*—stay tuned! And meanwhile, try the following exercises.

Exercises

E6.15 Convert the vectors $\vec{v}_1 = (4, 2)$ and $\vec{v}_2 = (1, 3)$ into an orthogonal basis for $\mathbb{R}^2$.

E6.16 Perform the Gram–Schmidt orthogonalization procedure on the vectors $\vec{v}_1 = (1, 1, 0)$, $\vec{v}_2 = (1, 0, 1)$, and $\vec{v}_3 = (0, 1, 1)$ to obtain an orthonormal basis $\{\hat{e}_1, \hat{e}_2, \hat{e}_3\}$.

E6.17 Consider the vector space of polynomials of degree at most two $P_2(x)$, and consider the inner product $\langle \mathbf{p}, \mathbf{q} \rangle = \int_0^1 p(x)q(x) dx$. Perform the Gram–Schmidt orthogonalization procedure on the polynomials $1, x$, and x^2 to obtain an orthonormal basis for $P_2(x)$.

6.6 Matrix decompositions

It’s often useful to express a given matrix as the product of other, simpler matrices. These matrix decompositions (also known as factorizations) can help us understand the structure of matrices by revealing their constituents. In this section, we’ll discuss various matrix factorizations and specify the types of matrices they apply to.

Most of the material covered here is not usually included in a first-year linear algebra course. Nevertheless, knowing about the different matrix decompositions is quite helpful, as many linear algebra applications depend on these decompositions. Got that? Good. Onward!

Eigendecomposition

The eigendecomposition breaks a matrix into its eigenvalues and eigenvectors. The eigenbasis, when it exists, is the most “natural”

basis for looking at a matrix. A diagonalizable matrix A can be written as

$$A = Q\Lambda Q^{-1},$$

where Q is a matrix whose columns are eigenvectors of A , and Λ is a diagonal matrix containing the eigenvalues of A .

The eigendecomposition of a matrix is a similarity transformation (a change of basis) where the new basis matrix consists of eigenvectors of the matrix.

If A is positive semidefinite then its eigenvalues are nonnegative. If the matrix A is symmetric then its eigenvalues are real numbers.

When the matrix A is *normal*, meaning it satisfies $AA^T = A^T A$, we can choose Q to be an orthogonal matrix O that satisfies $O^T O = \mathbb{I}$. Calculating the inverse of an orthogonal matrix is easy: $O^{-1} = O^T$. The eigendecomposition for normal matrices is $A = O\Lambda O^T$.

Singular value decomposition

We can generalize the concepts of eigenvalues and eigenvectors to non-square matrices. Consider a matrix $A \in \mathbb{R}^{m \times n}$. Since the matrix A is not a square matrix, we can't use the standard eigendecomposition. However, there is a trick for turning a non-square matrix into a square matrix while preserving some of its properties: multiply the matrix by its transpose. The matrix $AA^T \in \mathbb{R}^{n \times n}$ has the same column space as the matrix A . Similarly, $A^T A \in \mathbb{R}^{m \times m}$ has the same row space as the matrix A .

The *singular value decomposition* breaks a matrix into the product of three matrices: an $m \times m$ orthogonal matrix U which consists of *left singular vectors*, an $m \times n$ matrix Σ with the *singular values* σ_i on the diagonal, and an $n \times n$ orthogonal matrix V^T of *right singular vectors*:

$$A = \underbrace{\begin{bmatrix} | & & | \\ \hat{u}_1 & \cdots & \hat{u}_m \\ | & & | \end{bmatrix}}_U \underbrace{\begin{bmatrix} \sigma_1 & 0 & \cdots \\ 0 & \sigma_2 & \cdots \\ 0 & 0 & \cdots \end{bmatrix}}_\Sigma \underbrace{\begin{bmatrix} - & \hat{v}_1 & - \\ & \vdots & \\ - & \hat{v}_n & - \end{bmatrix}}_{V^T} = U\Sigma V^T.$$

To find the matrices U , Σ , and V , perform eigendecomposition on the matrix products AA^T and $A^T A$.

Consider first the matrix AA^T . Since AA^T is a square matrix, we can compute its eigendecomposition $AA^T = U\Lambda_\ell U^T$. The eigenvectors of AA^T span the same space as the column space of the matrix A . We call these vectors the *left singular vectors* of A .

The left singular vectors of A (the columns of U) are the eigenvectors of the matrix AA^T :

$$U = \begin{bmatrix} | & & | \\ \hat{u}_1 & \cdots & \hat{u}_m \\ | & & | \end{bmatrix}, \quad \text{where } \{(\lambda_i, \hat{u}_i)\} = \text{eigenvecs}(AA^T).$$

To find the right singular vectors of A (the rows of V^T), perform the eigendecomposition on the matrix $A^T A$, denoted $A^T A = V\Lambda_r V^T$. Build the orthogonal matrix V^T by stacking the eigenvectors of $A^T A$ as rows:

$$V^T = \begin{bmatrix} - & \hat{v}_1 & - \\ & \vdots & \\ - & \hat{v}_n & - \end{bmatrix}, \quad \text{where } \{(\lambda_i, \hat{v}_i)\} = \text{eigenvecs}(A^T A).$$

The eigenvalues of the matrix $A^T A$ are the same as the eigenvalues of the matrix AA^T . In both cases, the eigenvalues λ_i correspond to the squares of the singular values of the matrix A .

On its diagonal, the matrix of singular values $\Sigma \in \mathbb{R}^{m \times n}$ contains the singular values σ_i , which are the positive square roots of the eigenvalues λ_i of the matrix AA^T (or the matrix $A^T A$):

$$\sigma_i = \sqrt{\lambda_i}, \quad \text{where } \{\lambda_i\} = \text{eigenvals}(AA^T) = \text{eigenvals}(A^T A).$$

The singular value decomposition shows the inner structure of the matrix A . We can interpret the operation $\vec{y} = A\vec{x} = U\Sigma V^T \vec{x}$ as a three-step process:

1. Convert the input $\vec{x}$ to the basis of right singular vectors $\{\vec{v}_i\}$.
2. Scale each component by the corresponding singular value σ_i .
3. Convert the output from the $\{\vec{u}_i\}$ basis to the standard basis.

This three-step procedure is analogous to the three-step procedure we used to understand the eigendecomposition of square matrices in Section 6.1 (see page 305).

The singular value decomposition (SVD) has numerous applications in statistics, machine learning, and computer science. Applying the SVD to a matrix is like looking inside it with X-ray vision, since you can see its σ_i s. The action of $A = U\Sigma V^T$ occurs in n parallel streams: the i^{th} stream consists of multiplying the input vector

by the right singular vector $\vec{v}_i^T$, scaling by the weight σ_i , and finally multiplying by the left singular vector $\vec{u}_i$. Each singular value σ_i corresponds to the “strength” of A on the i^{th} subspace—the subspace spanned by its i^{th} left and right singular vectors.

Example Suppose you need to calculate the product $M\vec{v}$ where $M \in \mathbb{R}^{1000 \times 2000}$ and $\vec{v} \in \mathbb{R}^{2000}$. Suppose furthermore the matrix M has only three large singular values, $\sigma_1 = 6$, $\sigma_2 = 5$, $\sigma_3 = 4$, and many small singular values, $\sigma_4 = 0.0002$, $\sigma_5 = 0.0001$, $\dots$, $\sigma_{1000} = 1.1 \times 10^{-13}$. Observe that most of the “weight” of the matrix Σ is concentrated in the first three singular values, σ_1 , σ_2 , and σ_3 . We can obtain a *low-rank approximation* to the matrix M by keeping only the large singular values and their associated singular vectors. Construct the matrix $\tilde{\Sigma} \in \mathbb{R}^{3 \times 3}$ which contains only the first three singular values, and surround $\tilde{\Sigma}$ with matrices $\tilde{U} \in \mathbb{R}^{1000 \times 3}$ and $\tilde{V}^T \in \mathbb{R}^{3 \times 2000}$ which contain the singular vectors associated with the first three singular values. Despite the significant reduction in the size of the matrices used in the decomposition, the matrix $\tilde{M} = \tilde{U}\tilde{\Sigma}\tilde{V}^T \in \mathbb{R}^{1000 \times 2000}$ represents a good approximation to the original matrix M . We cut some small, insignificant singular values, which doesn’t change the matrix too much. We can quantify the difference between the original M and its low-rank approximation $\tilde{M}$ using the Hilbert–Schmidt norm: $\|M - \tilde{M}\|_{\text{HS}} = \sqrt{\sum_{i=4}^{1000} \sigma_i^2}$. Since $\sigma_4, \sigma_5, \dots, \sigma_{1000}$ are tiny numbers, we can say $\tilde{M} \approx M$.

Links

[Singular value decomposition on Wikipedia]

http://en.wikipedia.org/wiki/Singular_value_decomposition

[Excellent blog post series that explains the SVD intuitively]

Part 1, motivation: <https://jeremykun.com/?p=5946>

Part 2, computations: <https://jeremykun.com/?p=8329>

[Understanding the SVD and its applications]

<http://www.math.umn.edu/~lerman/math5467/svd.pdf>

[Principal component analysis in statistics is based on SVD]

http://en.wikipedia.org/wiki/Principal_component_analysis

LU decomposition

Computing the inverse of a triangular matrix is far easier than computing the inverse of a general matrix. Thus, for computational purposes, it’s sometimes useful to write a matrix as the product of two

triangular matrices. We call this factorization the *LU* decomposition:

$$A = LU,$$

where U is an *upper triangular* matrix and L is a *lower triangular* matrix.

The main application of this decomposition is to obtain more efficient solutions to equations of the form $A\vec{x} = \vec{b}$. Because $A = LU$, we can solve this equation in two steps. Starting from $LU\vec{x} = \vec{b}$, first multiply by L^{-1} and then by U^{-1} :

$$LU\vec{x} = \vec{b} \Rightarrow L^{-1}LU\vec{x} = U\vec{x} = L^{-1}\vec{b} \Rightarrow \vec{x} = U^{-1}L^{-1}\vec{b}.$$

We've split the work of finding the inverse A^{-1} into two simpler subtasks: finding L^{-1} and U^{-1} , which are easier to compute.

The *LU* decomposition is mainly used for linear algebra calculations on computers, but it's also possible to find the *LU* decomposition of a matrix by hand. Recall the algorithm for finding the inverse of a matrix, in which we start from the array $[A \mid \mathbb{1}]$ and perform row operations to bring the array into reduced row echelon form $[\mathbb{1} \mid A^{-1}]$. Consider the midpoint of the algorithm when the left side of the array is the row echelon form (REF). Since the matrix A in its REF is upper triangular, the array will contain $[U \mid L^{-1}]$. The U part of the decomposition is on the left side, and the L part is obtained by finding the inverse of the right side of the array.

Note the *LU* decomposition exists only for matrices that can be brought to RREF without using row-swap operations. If a matrix A requires row-swap operations to be transformed to RREF, we can decompose it as $A = PLU$, where P is a permutation matrix. The *PLU* decomposition has the same computational advantages of splitting the inverse computation into three simpler subtasks: $A^{-1} = U^{-1}L^{-1}P^{-1}$.

Cholesky decomposition

For a *symmetric, positive semidefinite* matrix A , the *LU* decomposition can take on a simpler form. Such matrices can be written as the product of a triangular matrix and its transpose:

$$A = LL^T \quad \text{or} \quad A = U^T U,$$

where U is an *upper triangular* matrix and L is a *lower triangular* matrix. This is called the *Cholesky decomposition* of a matrix, and like the *LU* decomposition, it has applications for faster numerical linear algebra calculations, nonlinear optimization, and machine learning.

QR decomposition

Any real square matrix $A \in \mathbb{R}^{n \times n}$ can be decomposed as a product of an orthogonal matrix O and an upper triangular matrix U :

$$A = OU.$$

For historical reasons, the orthogonal matrix is denoted Q instead of O , and the upper triangular matrix is denoted R (think “right-triangular” since it contains entries only to the *right* of the main diagonal). Using the conventional names, the decomposition becomes

$$A = QR,$$

which is why it’s known as the QR decomposition.

The QR decomposition is equivalent to the Gram–Schmidt orthogonalization procedure applied to the columns of the matrix. The matrix Q records the orthonormal basis while the matrix R contains the coefficients required to express the columns of A as linear combinations of the columns of Q .

Example Consider the decomposition

$$A = \begin{bmatrix} 12 & -51 & 4 \\ 6 & 167 & -68 \\ -4 & 24 & -41 \end{bmatrix} = QR.$$

We’re looking for an orthogonal matrix Q and an upper triangular matrix R such that $A = QR$. We can obtain the orthogonal matrix Q by performing the Gram–Schmidt procedure on the columns of A .

Let’s illustrate the procedure by computing the factorization A . Begin by changing the second column in A so it becomes orthogonal to the first (by subtracting a multiple of the first column). Next, change the third column in A so it is orthogonal to both of the first columns (by subtracting multiples of the first two columns). We obtain a matrix with the same column space as A , but which has orthogonal columns:

$$\begin{bmatrix} | & | & | \\ \mathbf{u}_1 & \mathbf{u}_2 & \mathbf{u}_3 \\ | & | & | \end{bmatrix} = \begin{bmatrix} 12 & -69 & -\frac{58}{5} \\ 6 & 158 & \frac{6}{5} \\ -4 & 30 & -33 \end{bmatrix}.$$

To obtain an orthogonal matrix, we must normalize each column to have length one:

$$Q = \begin{bmatrix} | & | & | \\ \frac{\mathbf{u}_1}{\|\mathbf{u}_1\|} & \frac{\mathbf{u}_2}{\|\mathbf{u}_2\|} & \frac{\mathbf{u}_3}{\|\mathbf{u}_3\|} \\ | & | & | \end{bmatrix} = \begin{bmatrix} \frac{6}{7} & -\frac{69}{175} & -\frac{58}{175} \\ \frac{3}{7} & \frac{158}{175} & \frac{6}{175} \\ -\frac{2}{7} & \frac{6}{35} & -\frac{33}{35} \end{bmatrix}.$$

We can obtain the matrix R from Q^T and A :

$$Q^T A = \underbrace{Q^T Q}_I R = R \quad \Rightarrow \quad R = Q^T A = \begin{bmatrix} 14 & 21 & -14 \\ 0 & 175 & -70 \\ 0 & 0 & 35 \end{bmatrix}.$$

The columns of R contain the mixture of coefficients required to obtain the columns of A from the columns of Q . For example, the second column of A is equal to $21 \frac{\mathbf{u}_1}{\|\mathbf{u}_1\|} + 175 \frac{\mathbf{u}_2}{\|\mathbf{u}_2\|}$. Verify that QR equals A .

Discussion

The last several pages have only scratched the surface of matrix decompositions. There are countless applications for matrix methods, and matrix factorizations play key roles in many of them.

Machine learning techniques often use matrix decompositions to uncover useful structure within data matrices. Two examples include *nonnegative matrix factorization* (used for recommender systems) and *latent Dirichlet allocation* (used for document classification). I encourage you to research this subject further on your own—it's quite an interesting wormhole to get sucked into.

Links

[Cool retro video showing the steps of the SVD procedure]

<http://www.youtube.com/watch?v=R9UoFyqJca8>

[More info and examples on Wikipedia]

http://en.wikipedia.org/wiki/Matrix_decomposition

http://en.wikipedia.org/wiki/Cholesky_decomposition

[A detailed example of the QR factorization of a matrix]

<http://www.math.ucla.edu/~yanovsky/Teaching/Math151B/handouts/GramSchmidt.pdf>

Exercises

E6.18 Compute the QR factorization of the matrix $A = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$.

6.7 Linear algebra with complex numbers

So far we've discussed the math of vectors and matrices with real components. In fact, the linear algebra techniques you've learned apply to any *field* F . The term *field* applies to any mathematical object for which the operations of addition, subtraction, multiplication, and division are defined.

Since the complex numbers $\mathbb{C}$ are a field, we can perform linear algebra over the field of complex numbers. In this section, we'll define vectors and matrices with complex components, and discover that they behave similarly to their real counterparts. You'll see that complex linear algebra is no more complex than real linear algebra. It's the same, in fact, except for one small difference: instead of matrix transpose A^T , we use the Hermitian transpose $A^\dagger$, which is the combination of the transpose and an entry-wise complex conjugate operation.

Complex vectors are not just an esoteric mathematical concept intended for specialists. Complex vectors can arise as answers for problems involving ordinary real matrices. For example, the rotation matrix

$$R_\theta = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

has complex eigenvalues $\lambda_1 = e^{i\theta}$ and $\lambda_2 = e^{-i\theta}$ and its eigenvectors have complex components. If you want to know how to calculate the eigenvalues and eigenvectors of rotation matrices, you need to understand how to do linear algebra calculations with complex numbers.

This section serves as a review of all the important linear algebra concepts we've learned in this book. I recommend you read this section, even if you're not required to know about complex matrices for your course. As your guide through the land of linear algebra, it's my duty to make sure you understand linear algebra in the complex field. It's good stuff; I guarantee there's *knowledge buzz* to be had in this section.

Definitions

Recall the basic notions of complex numbers introduced in Section 1.14:

- i : the unit imaginary number; $i = \sqrt{-1}$ and $i^2 = -1$
- $z = a + bi$: a complex number z whose real part is a and whose imaginary part is b
- $\mathbb{C}$: the set of complex numbers $\mathbb{C} = \{a + bi \mid a, b \in \mathbb{R}\}$

- $\operatorname{Re}\{z\} = a$: the *real* part of $z = a + bi$
- $\operatorname{Im}\{z\} = b$: the *imaginary* part of $z = a + bi$
- $\bar{z}$: the *complex conjugate* of z . If $z = a + bi$ then $\bar{z} = a - bi$
- $|z| = \sqrt{z\bar{z}} = \sqrt{a^2 + b^2}$: the *magnitude* or *length* of $z = a + bi$
- $\arg(z) = \tan^{-1}(\frac{b}{a})$: the *phase* or *argument* of $z = a + bi$. Note that $\tan^{-1}(\frac{b}{a})$ and $\arg(z)$ coincide only if $a \geq 0$. A correction is necessary to the output of $\tan^{-1}(\frac{b}{a})$ when $a < 0$. The computer function `atan2(b, a)` returns the correct phase for all $z = a + bi$.

Complex vectors

A complex vector $\vec{v} \in \mathbb{C}^n$ is an array of n complex numbers:

$$\vec{v} = (v_1, v_2, \dots, v_n) \in \mathbb{C}^n.$$

Complex matrices

A complex matrix $A \in \mathbb{C}^{m \times n}$ is a table of numbers with m rows and n columns. An example of a 3×2 matrix with complex entries is

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} \in \mathbb{C}^{3 \times 2}.$$

Hermitian transpose

The *Hermitian transpose* operation, denoted † , consists of the combination of the regular transpose ($A \rightarrow A^\top$) and the complex conjugation of each entry in the matrix ($a_{ij} \rightarrow \bar{a}_{ij}$):

$$A^\dagger \stackrel{\text{def}}{=} \overline{(A^\top)} = (\bar{A})^\top.$$

Expressed in terms of the entries of the matrix a_{ij} , the Hermitian transpose corresponds to the transformation $a_{ij} \rightarrow \bar{a}_{ji}$. There are *many* mathematical terms that refer to this operation, including *Hermitian conjugate*, *complex transpose*, “*dagger*” operation, *conjugate transpose*, and *adjoint*.

The term *adjoint* is preferred by mathematicians and the notation A^* is used consistently in mathematics research papers. The dagger notation † is preferred by physicists and engineers, but shunned by mathematicians. Mathematicians prefer to stick with the star superscript because they feel they invented the concept. We use the notation † in this book because at some point the author had to make an allegiance with one of the two camps, and because the symbol † looks a bit like the transpose symbol $^\top$.

The Hermitian transpose applied to a 3×2 matrix acts as follows:

$$\text{if } A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} \quad \text{then} \quad A^\dagger = \begin{bmatrix} \overline{a_{11}} & \overline{a_{21}} & \overline{a_{31}} \\ \overline{a_{12}} & \overline{a_{22}} & \overline{a_{32}} \end{bmatrix}.$$

Recall that vectors are special types of matrices. We can identify a vector $\vec{v} \in \mathbb{C}^n$ with a column matrix $\vec{v} \in \mathbb{C}^{n \times 1}$ or with a row matrix $\vec{v} \in \mathbb{C}^{1 \times n}$. We apply the complex conjugation operation to transform column vectors into conjugate row vectors:

$$\vec{v}^\dagger \stackrel{\text{def}}{=} \overline{(\vec{v}^T)} = (\overline{\vec{v}})^T.$$

The Hermitian transpose of a column vector is a row vector in which each component has been complex-conjugated:

$$\text{if } \vec{v} = \begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix} \quad \text{then} \quad \vec{v}^\dagger = \begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix}^\dagger = [\bar{\alpha} \quad \bar{\beta} \quad \bar{\gamma}].$$

The Hermitian transpose for vectors is important because it's related to the definition of the inner product for complex vectors.

Complex inner product

The inner product for vectors with complex components $\vec{u}, \vec{v} \in \mathbb{C}^n$ is defined as the following operation:

$$\langle \vec{u}, \vec{v} \rangle \stackrel{\text{def}}{=} \sum_{i=1}^n \overline{u_i} v_i = \vec{u}^\dagger \vec{v}.$$

In this expression, complex conjugation is applied to the components of the first vector. This corresponds to the notion of applying the Hermitian transpose to the first vector to turn it into a row vector of complex conjugates, then using the matrix multiplication rule for a $1 \times n$ matrix $\vec{u}^\dagger$ times an $n \times 1$ matrix $\vec{v}$.

For real vectors $\vec{u}, \vec{v} \in \mathbb{R}^n$, the complex inner product formula reduces to the inner product formula we used previously: $\vec{u} \cdot \vec{v} = \vec{u}^T \vec{v}$. Rather than thinking of the inner product for complex vectors as a new operation, we can say the inner product has always been defined as $\langle \vec{u}, \vec{v} \rangle = \vec{u}^\dagger \vec{v}$ —we just never noticed until now because complex conjugation has no effect on vectors with real components. Specifically, $\vec{u}^\dagger = \vec{u}^T$ if all the u_i s are real numbers.

Linear algebra over the complex field

One of the fundamental linear algebra ideas we've learned is how to use *linear transformations* to model input-output phenomena in which input vectors $\vec{v}$ are linearly transformed to output vectors: $\vec{w} = T(\vec{v})$. Linear transformations are functions of the form $T : \mathbb{R}^m \rightarrow \mathbb{R}^n$. We can *represent* these linear transformations as $m \times n$ matrices *with respect to* some choice of input and output bases.

These linear algebra ideas also apply to complex vectors and complex matrices. For example, a linear transformation from $T : \mathbb{C}^2 \rightarrow \mathbb{C}^2$ can be represented as the matrix multiplication

$$\begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}.$$

Each linear transformation $T : \mathbb{C}^2 \rightarrow \mathbb{C}^2$ corresponds to some 2×2 matrix $\begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix}$ with entries $\alpha, \beta, \gamma, \delta \in \mathbb{C}$.

The change from real entries to complex entries has the effect of doubling the number of parameters required to describe the transformation. A 2×2 complex matrix has eight parameters, not four. Where are those eight parameters, you ask? Here:

$$\begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix} = \begin{bmatrix} \operatorname{Re}\{\alpha\} & \operatorname{Re}\{\beta\} \\ \operatorname{Re}\{\gamma\} & \operatorname{Re}\{\delta\} \end{bmatrix} + \begin{bmatrix} \operatorname{Im}\{\alpha\} & \operatorname{Im}\{\beta\} \\ \operatorname{Im}\{\gamma\} & \operatorname{Im}\{\delta\} \end{bmatrix} i.$$

Each of the four entries of the matrix has a real part and an imaginary part, making for a total of eight parameters to pick when specifying the matrix.

Similarly, to specify a vector $\vec{v} \in \mathbb{C}^2$ you need to specify four parameters:

$$\begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} \operatorname{Re}\{v_1\} \\ \operatorname{Re}\{v_2\} \end{bmatrix} + \begin{bmatrix} \operatorname{Im}\{v_1\} \\ \operatorname{Im}\{v_2\} \end{bmatrix} i.$$

In practice, this doubling of dimensions doesn't play a role in calculations because we usually perform algebra steps with the complex entries and rarely split the matrices into their real and imaginary parts.

All the linear algebra techniques you've learned also work with complex numbers, as you'll see in the following examples.

Example 1: Solving systems of equations Suppose you're solving a problem that involves complex numbers and a system of two linear equations in two unknowns:

$$\begin{aligned} z_1 + 2z_2 &= 3 + i \\ 3z_1 + (9 + i)z_2 &= 6 + 2i. \end{aligned}$$

You're asked to find the values of the unknowns z_1 and z_2 .

The solutions z_1 and z_2 will be complex numbers, but apart from that, there's nothing special about this problem—keep in mind, linear algebra with complex numbers is the same as linear algebra with real numbers, so the techniques you learned for real numbers work just as well for complex numbers. Now let's solve this system of equations.

First observe that the system of equations can be written as a matrix-vector product:

$$\underbrace{\begin{bmatrix} 1 & 2 \\ 3 & 9+i \end{bmatrix}}_A \underbrace{\begin{bmatrix} z_1 \\ z_2 \end{bmatrix}}_{\vec{z}} = \underbrace{\begin{bmatrix} 3+i \\ 6+2i \end{bmatrix}}_{\vec{b}}.$$

We've expressed the system as a 2×2 matrix A multiplying the vector of unknowns $\vec{z}$ (a 2×1 matrix) to produce a vector of constants $\vec{b}$ (another 2×1 matrix). We can solve for $\vec{z}$ by multiplying both sides of the equation by the inverse matrix A^{-1} . The inverse matrix of A is

$$A^{-1} = \begin{bmatrix} 1 + \frac{6}{3+i} & -\frac{2}{3+i} \\ -\frac{3}{3+i} & \frac{1}{3+i} \end{bmatrix}.$$

We can now compute the answer $\vec{z}$ using the equation $\vec{z} = A^{-1}\vec{b}$:

$$\begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} 1 + \frac{6}{3+i} & -\frac{2}{3+i} \\ -\frac{3}{3+i} & \frac{1}{3+i} \end{bmatrix} \begin{bmatrix} 3+i \\ 6+2i \end{bmatrix} = \begin{bmatrix} 3+i+6-4 \\ -3+2 \end{bmatrix} = \begin{bmatrix} 5+i \\ -1 \end{bmatrix}.$$

Example 2: Finding the inverse We learned several approaches for computing matrix inverses in Section 3.5. Here we'll review the procedure for computing the inverse using row operations.

Given the matrix

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 9+i \end{bmatrix},$$

first build a 2×4 array that contains A on the left side and the identity matrix $\mathbb{I}$ on the right side:

$$\left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 3 & 9+i & 0 & 1 \end{array} \right].$$

Now perform the Gauss–Jordan elimination procedure on the resulting 2×4 array.

1. Subtract three times the first row from the second row ($R_2 \leftarrow R_2 - 3R_1$) to obtain

$$\left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 0 & 3+i & -3 & 1 \end{array} \right].$$

2. Perform $R_2 \leftarrow \frac{1}{3+i}R_2$ to create a pivot in the second row:

$$\left[\begin{array}{cc|cc} 1 & 2 & 1 & 0 \\ 0 & 1 & \frac{-3}{3+i} & \frac{1}{3+i} \end{array} \right].$$

3. Finally, perform $R_1 \leftarrow R_1 - 2R_2$ to obtain the RREF:

$$\left[\begin{array}{cc|cc} 1 & 0 & 1 + \frac{6}{3+i} & -\frac{2}{3+i} \\ 0 & 1 & \frac{-3}{3+i} & \frac{1}{3+i} \end{array} \right].$$

The inverse of A appears on the right side of the array,

$$A^{-1} = \left[\begin{array}{cc} 1 + \frac{6}{3+i} & -\frac{2}{3+i} \\ -\frac{3}{3+i} & \frac{1}{3+i} \end{array} \right].$$

Example 3: Linear transformations as matrices The effect of multiplying a vector $\vec{v} \in \mathbb{C}^n$ by a matrix $M \in \mathbb{C}^{m \times n}$ is the same as applying a linear transformation $T_M : \mathbb{C}^n \rightarrow \mathbb{C}^m$ to the vector:

$$\vec{w} = M\vec{v} \quad \Leftrightarrow \quad \vec{w} = T_M(\vec{v}).$$

The opposite is also true—any linear transformation T can be *represented* as a multiplication by some matrix M_T :

$$\vec{w} = T(\vec{v}) \quad \Leftrightarrow \quad \vec{w} = M_T\vec{v}.$$

We'll use a simple example to review the procedure for finding the matrix representation of a linear transformation.

Consider the linear transformation $T : \mathbb{C}^2 \rightarrow \mathbb{C}^2$, which produces the input-output pairs

$$T\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}\right) = \begin{bmatrix} 3 \\ 2i \end{bmatrix} \quad \text{and} \quad T\left(\begin{bmatrix} 0 \\ 1 \end{bmatrix}\right) = \begin{bmatrix} 2 \\ 4 + 4i \end{bmatrix}.$$

How can we use the information provided above to find the matrix representation of the linear transformation T ?

To obtain the matrix representation of T with respect to a given basis, we need to combine, as columns, the outputs of T for the n vectors of that basis:

$$M_T = \left[\begin{array}{ccc|c} | & | & & | \\ T(\vec{e}_1) & T(\vec{e}_2) & \cdots & T(\vec{e}_n) \\ | & | & & | \end{array} \right],$$

where the set $\{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$ is a basis for the input space.

The problem statement gives us the information needed for the first column of M_T , but we're not given the output of T for $\hat{e}_2$. However, we can work around this limitation since we know T is *linear*. The property $T(\alpha\vec{v}) = \alpha T(\vec{v})$ implies

$$T\left(2\begin{bmatrix} 0 \\ 1 \end{bmatrix}\right) = 2\begin{bmatrix} 1 \\ 2+2i \end{bmatrix} \quad \Rightarrow \quad T\left(\begin{bmatrix} 0 \\ 1 \end{bmatrix}\right) = \begin{bmatrix} 1 \\ 2+2i \end{bmatrix}.$$

Combining the information for $T(\hat{e}_1)$ and $T(\hat{e}_2)$, we obtain the matrix representation of T :

$$M_T = \begin{bmatrix} 3 & 1 \\ 2i & 2+2i \end{bmatrix}.$$

Complex eigenvalues

I want you to learn about linear algebra with complex vectors so that we can complete the classification of linear transformations in terms of their eigenvalues. Recall that projections obey $\Pi = \Pi^2$ and have eigenvalues zero or one, and reflections have at least one eigenvalue equal to -1 .

What are the eigenvalues of rotation matrices? The eigenvalues of a matrix A are the roots of its characteristic polynomial $p_A(\lambda) = \det(A - \lambda\mathbb{I})$. To find the eigenvalues of the rotation matrix R_θ we defined in Section 5.2 (page 277), we must find the solutions of the equation $p_{R_\theta}(\lambda) = 0$:

$$\begin{aligned} 0 &= p_{R_\theta}(\lambda) \\ &= \det(R_\theta - \lambda\mathbb{I}) \\ &= \det\left(\begin{bmatrix} \cos\theta - \lambda & -\sin\theta \\ \sin\theta & \cos\theta - \lambda \end{bmatrix}\right) \\ &= (\cos\theta - \lambda)^2 + \sin^2\theta. \end{aligned}$$

To solve for λ , first move $\sin^2\theta$ to the other side of the equation,

$$-\sin^2\theta = (\cos\theta - \lambda)^2,$$

then take the square root on both sides:

$$\cos\theta - \lambda = \pm\sqrt{-\sin^2\theta} = \pm\sqrt{-1}\sin\theta = \pm i\sin\theta.$$

The eigenvalues of R_θ are $\lambda_1 = \cos\theta + i\sin\theta$ and $\lambda_2 = \cos\theta - i\sin\theta$. Using Euler's formula (see page 102) we can express the eigenvalues more compactly as $\lambda_1 = e^{i\theta}$ and $\lambda_2 = e^{-i\theta}$. What's interesting here is that complex numbers emerge as answers to a matrix problem that was originally stated in terms of real variables.

This is not a coincidence: complex exponentials are in many ways the natural way to talk about rotations, periodic motion, and waves. If you pursue a career in math, physics, or engineering, you'll use complex numbers and Euler's formula on a daily basis.

Special types of matrices

We'll now define a few special types of matrices with complex entries. These matrices are analogous to the special matrices we defined in Section 6.2, but their definitions are adapted to use the Hermitian transpose operation † .

Unitary matrices

A matrix U is *unitary* if it obeys $U^\dagger U = \mathbb{1}$. The norm of the determinant of a unitary matrix is 1, $|\det(U)| = 1$. For an $n \times n$ matrix U , the following statements are equivalent:

- U is unitary.
- The columns of U form an orthonormal basis.
- The rows of U form an orthonormal basis.
- The inverse of U is $U^\dagger$.

Unitary matrices are the complex analogues of orthogonal matrices. Indeed, if a unitary matrix U has real entries, then $U^\dagger = U^T$ and we have $U^T U = \mathbb{1}$, which is the definition of an orthogonal matrix.

Hermitian matrices

A Hermitian matrix H is equal to its own Hermitian transpose:

$$H^\dagger = H \quad \Leftrightarrow \quad h_{ij} = \overline{h_{ji}}, \quad \text{for all } i, j.$$

Hermitian matrices are complex-number analogues of symmetric matrices.

A Hermitian matrix H can be freely moved from one side to the other in a complex inner product:

$$\langle H\vec{x}, \vec{y} \rangle = (H\vec{x})^\dagger \vec{y} = \vec{x}^\dagger H^\dagger \vec{y} = \vec{x}^\dagger (H\vec{y}) = \langle \vec{x}, H\vec{y} \rangle.$$

The eigenvalues of Hermitian matrices are real numbers.

Normal matrices

Previously, we defined the set of real normal matrices to be matrices that satisfy $A^T A = A A^T$. For matrices with complex entries, the definition of a normal matrix uses the dagger: $A^\dagger A = A A^\dagger$.

* * *

Consulting the concept map in Figure 6.1 on page 317 will help you see the parallels between the different types of special matrices. I realize there's a lot of new terminology to absorb all at once, so don't worry about remembering everything. The main idea is to know that these special types of matrices exist—not to know *everything* about them.

Inner product for complex vectors

The complex inner product is an operation of the form

$$\langle \cdot, \cdot \rangle : \mathbb{C}^n \times \mathbb{C}^n \rightarrow \mathbb{C}.$$

The inner product $\langle \vec{u}, \vec{v} \rangle$ for real vectors is equivalent to the matrix multiplication between the row vector $\vec{u}^T$ and the column vector $\vec{v}$. Extending the notion of inner product to work with complex vectors requires a modification to the inner product formula. The inner product for vectors $\vec{u}, \vec{v} \in \mathbb{C}^n$ is defined as

$$\langle \vec{u}, \vec{v} \rangle \stackrel{\text{def}}{=} \sum_{i=1}^n \bar{u}_i v_i = \vec{u}^\dagger \vec{v}.$$

The formula is similar to the inner product formula for real vectors, but uses the Hermitian transpose † instead of the regular transpose ^{T} . The inner product of two vectors $\vec{u}, \vec{v} \in \mathbb{C}^3$ is

$$\langle \vec{u}, \vec{v} \rangle = \bar{u}_1 v_1 + \bar{u}_2 v_2 + \bar{u}_3 v_3 = \begin{bmatrix} \bar{u}_1 & \bar{u}_2 & \bar{u}_3 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \vec{u}^\dagger \vec{v}.$$

This dagger thing is very important. Using the definition of the inner product with a dagger on the first input ensures the complex inner product will obey the positive semidefinite criterion (see page 322). The inner product of a vector $\vec{v} \in \mathbb{C}^3$ with itself is

$$\langle \vec{v}, \vec{v} \rangle = \vec{v}^\dagger \vec{v} = \begin{bmatrix} \bar{v}_1 & \bar{v}_2 & \bar{v}_3 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = |v_1|^2 + |v_2|^2 + |v_3|^2,$$

where $|v_i|^2 = \bar{v}_i v_i$ is the magnitude-squared of the component $v_i \in \mathbb{C}$. The magnitudes of the complex components are nonnegative real numbers, so the sum of their squares is also a nonnegative real number. Therefore, the complex inner product satisfies the positive semidefinite requirement $\langle \vec{v}, \vec{v} \rangle \geq 0$ for inner products.

Length of a complex vector

The complex inner product induces the following complex norm:

$$\|\vec{v}\| \stackrel{\text{def}}{=} \sqrt{\langle \vec{v}, \vec{v} \rangle} = \sqrt{\vec{v}^\dagger \vec{v}} = \sqrt{|v_1|^2 + |v_2|^2 + \cdots + |v_n|^2}.$$

The norm for complex vectors satisfies the positive semidefinite requirement $\|\vec{v}\| \geq 0$ for norms (see page 324).

Example Calculate the norm of the vector $\vec{v} = (2 + i, 3, 5i)$.

The Hermitian transpose of the row vector $\vec{v}$ is the column vector $\vec{v}^\dagger = (2 - i, 3, -5i)^\top$. The norm of $\vec{v}$ is equal to the square root of $\langle \vec{v}, \vec{v} \rangle$ so $\|\vec{v}\| = \sqrt{\langle \vec{v}, \vec{v} \rangle} = \sqrt{(2 - i)(2 + i) + 3^2 + (-5i)(5i)} = \sqrt{4 + 1 + 9 + 25} = \sqrt{39}$.

Complex inner product spaces

A real inner product space is an abstract vector space $(V, \mathbb{R}, +, \cdot)$ for which we've defined an inner product operation $\langle \mathbf{u}, \mathbf{v} \rangle$ which obeys (1) the symmetric property, (2) the linearity property, and (3) the positive semidefinite property.

Similarly, a complex inner product space is an abstract vector space $(V, \mathbb{C}, +, \cdot)$ with an inner product operation $\langle \mathbf{u}, \mathbf{v} \rangle$ that satisfies the following criteria for all $\mathbf{u}, \mathbf{v}, \mathbf{v}_1, \mathbf{v}_2 \in V$ and $\alpha, \beta \in \mathbb{C}$:

- Conjugate symmetric: $\langle \mathbf{u}, \mathbf{v} \rangle = \overline{\langle \mathbf{v}, \mathbf{u} \rangle}$
- Linear: $\langle \mathbf{u}, \alpha \mathbf{v}_1 + \beta \mathbf{v}_2 \rangle = \alpha \langle \mathbf{u}, \mathbf{v}_1 \rangle + \beta \langle \mathbf{u}, \mathbf{v}_2 \rangle$
- Positive semidefinite: $\langle \mathbf{u}, \mathbf{u} \rangle \geq 0$ for all $\mathbf{u} \in V$ with $\langle \mathbf{u}, \mathbf{u} \rangle = 0$ if and only if $\mathbf{u} = \mathbf{0}$

The conjugate symmetry property $\langle \mathbf{u}, \mathbf{v} \rangle = \overline{\langle \mathbf{v}, \mathbf{u} \rangle}$ ensures the inner product of a vector with itself is a real number: $\langle \mathbf{u}, \mathbf{u} \rangle = \overline{\langle \mathbf{u}, \mathbf{u} \rangle} \in \mathbb{R}$.

Example The Hilbert–Schmidt inner product for matrices $A, B \in \mathbb{C}^{m \times n}$ is defined as

$$\langle A, B \rangle_{\text{HS}} \stackrel{\text{def}}{=} \text{Tr}(A^\dagger B) = \sum_{i=1}^n \langle A \vec{e}_i, B \vec{e}_i \rangle.$$

The product $A\vec{e}_i$ has the effect of “selecting” the i^{th} column of the matrix A ; we can consider the Hilbert–Schmidt inner product of matrices A and B as the sum of the vector inner products of the columns of the two matrices.

We can also define the Hilbert–Schmidt norm for matrices:

$$\|A\|_{\text{HS}} \stackrel{\text{def}}{=} \sqrt{\langle A, A \rangle_{\text{HS}}} = \sqrt{\text{Tr}(A^\dagger A)} = \left[\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2 \right]^{\frac{1}{2}}.$$

The Hilbert–Schmidt norm is the square root of the sum of the squared magnitudes of the entries of the matrix.

The Hilbert–Schmidt inner product and norm are sometimes called *Frobenius inner product* and *Frobenius norm*, respectively.

Singular value decomposition

The singular value decomposition we introduced for real matrices in Section 6.6 also applies to matrices with complex entries.

The singular value decomposition of a matrix $A \in \mathbb{C}^{m \times n}$ is a way to express A as the product of three matrices:

$$A = U\Sigma V^\dagger.$$

The $m \times m$ unitary matrix U consists of *left singular vectors* of A . The $m \times n$ matrix Σ contains the *singular values* σ_i on the diagonal. The $n \times n$ unitary matrix $V^\dagger$ consists of *right singular vectors*.

The singular values σ_i of A are the positive square roots of the eigenvalues of the matrix $AA^\dagger$. To find the matrix of left singular vectors U , calculate the eigenvectors of $AA^\dagger$ and pack them as columns. To find the matrix of right singular vectors $V^\dagger$, calculate the eigenvectors $A^\dagger A$, pack them as columns in a matrix V , then take the Hermitian transpose of this matrix.

The Hilbert–Schmidt norm of a matrix $A \in \mathbb{C}^{m \times n}$ is equal to the square root of the sum of the squares of its singular values:

$$\|A\|_{\text{HS}} = \sqrt{\text{Tr}(A^\dagger A)} = \sqrt{\sum_{i=1}^n \sigma_i^2}.$$

This equation shows an important connection between the matrix’s norm and the size of its singular values. The singular value σ_i corresponds to the multiplier that A applies between the i^{th} left singular vector and the i^{th} right singular vector.

The singular value decomposition is used in many algorithms and procedures to uncover the inner structure of matrices. The machine learning technique called *principal component analysis* (PCA) corresponds to applying the SVD to a data matrix. Alternatively, you can think of the PCA as applying an eigendecomposition of the *covariance matrix* of the data.

Explanations

Complex eigenvectors

The characteristic polynomial of the rotation matrix R_θ is $p(\lambda) = (\cos \theta - \lambda)^2 + \sin^2 \theta = 0$. The eigenvalues of R_θ are $\lambda_1 = \cos \theta + i \sin \theta = e^{i\theta}$ and $\lambda_2 = \cos \theta - i \sin \theta = e^{-i\theta}$. What are its eigenvectors?

Before we get into the eigenvector calculation, I want to show you a useful trick for rewriting $\cos$ and $\sin$ expressions in terms of complex exponential functions. Recall Euler's formula, $e^{i\theta} = \cos \theta + i \sin \theta$. Using this equation and the analogous expression for $e^{-i\theta}$, we can obtain the following expressions for $\cos \theta$ and $\sin \theta$:

$$\cos \theta = \frac{1}{2} (e^{i\theta} + e^{-i\theta}), \quad \sin \theta = \frac{1}{2i} (e^{i\theta} - e^{-i\theta}).$$

Try calculating the right side in each case to verify the accuracy of each expression. These formulas are useful because they allow us to rewrite expressions of the form $e^{i\theta} \cos \phi$ as $e^{i\theta} \frac{1}{2} (e^{i\phi} + e^{-i\phi}) = \frac{1}{2} (e^{i(\theta+\phi)} + e^{i(\theta-\phi)})$, which is simpler.

Let's now see how to find the eigenvector $\vec{e}_{\lambda_1} = (\alpha, \beta)^T$ associated with the eigenvalue $\lambda_1 = e^{i\theta}$. The eigenvalue equation for the eigenvalue $\lambda_1 = e^{i\theta}$ is

$$R_\theta \vec{e}_{\lambda_1} = e^{i\theta} \vec{e}_{\lambda_1} \quad \Leftrightarrow \quad \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = e^{i\theta} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}.$$

We're looking for the components α and β .

Do you remember how to find eigenvectors? Don't worry if you've forgotten—this is why we have this review chapter! We'll go through the problem in detail. Brace yourself though, because the calculation is quite long.

The “finding the eigenvector(s) of A for the eigenvalue λ_1 ” task is carried out by finding the *null space* of the matrix $(A - \lambda_1 \mathbb{1})$. We rewrite the eigenvalue equation stated above as

$$(R_\theta - e^{i\theta} \mathbb{1}) \vec{e}_{\lambda_1} = 0 \quad \Leftrightarrow \quad \begin{bmatrix} \cos \theta - e^{i\theta} & -\sin \theta \\ \sin \theta & \cos \theta - e^{i\theta} \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

It's now clear that the finding-the-eigenvectors procedure corresponds to a null space calculation.

Let's use the cos-rewriting trick to simplify $\cos \theta - e^{i\theta}$:

$$\begin{aligned}\cos \theta - e^{i\theta} &= \frac{1}{2} \left(e^{i\theta} + e^{-i\theta} \right) - e^{i\theta} \\ &= \frac{1}{2} e^{i\theta} + \frac{1}{2} e^{-i\theta} - e^{i\theta} = \frac{-1}{2} e^{i\theta} + \frac{1}{2} e^{-i\theta} = \frac{-1}{2} \left(e^{i\theta} - e^{-i\theta} \right) \\ &= -i \frac{1}{2i} \left(e^{i\theta} - e^{-i\theta} \right) \\ &= -i \sin \theta.\end{aligned}$$

We substitute this simplified expression in the two places where it appears, and do some row operations to simplify the matrix:

$$\begin{bmatrix} -i \sin \theta & -\sin \theta \\ \sin \theta & -i \sin \theta \end{bmatrix} \sim \begin{bmatrix} \sin \theta & -i \sin \theta \\ \sin \theta & -i \sin \theta \end{bmatrix} \sim \begin{bmatrix} \sin \theta & -i \sin \theta \\ 0 & 0 \end{bmatrix} \sim \begin{bmatrix} 1 & -i \\ 0 & 0 \end{bmatrix}.$$

We can now solve the null space problem. Observe that the second column of the matrix does not contain a pivot, so β is a free variable, which we'll denote $s \in \mathbb{R}$. We thus obtain the equations:

$$\begin{bmatrix} 1 & -i \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ s \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad \Rightarrow \quad \begin{aligned} 1\alpha + (-i)s &= 0, \\ 0 &= 0. \end{aligned}$$

Solving for α in terms of s , we find $\alpha = is$, and therefore the solution is $(\alpha, \beta)^T = (is, s)$. The eigenspace that corresponds to the eigenvalue $\lambda_1 = e^{i\theta}$ is the null space of the matrix $(R_\theta - e^{i\theta} \mathbb{1})$:

$$\mathcal{N}(R_\theta - e^{i\theta} \mathbb{1}) = \left\{ \begin{bmatrix} is \\ s \end{bmatrix}, \forall s \in \mathbb{R} \right\} = \text{span} \left(\begin{bmatrix} i \\ 1 \end{bmatrix} \right).$$

After all this work, we've finally obtained an eigenvector $\vec{e}_{\lambda_1} = (i, 1)^T$ that corresponds to the eigenvalue $\lambda_1 = e^{i\theta}$. Let's verify that the vector we obtained satisfies the eigenvalue equation $R_\theta \vec{e}_{\lambda_1} = e^{i\theta} \vec{e}_{\lambda_1}$:

$$\begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} i \\ 1 \end{bmatrix} = \begin{bmatrix} i \cos \theta - \sin \theta \\ i \sin \theta + \cos \theta \end{bmatrix} = \begin{bmatrix} i(\cos \theta + i \sin \theta) \\ \cos \theta + i \sin \theta \end{bmatrix} = e^{i\theta} \begin{bmatrix} i \\ 1 \end{bmatrix}.$$

The eigenvector for the eigenvalue $\lambda_2 = e^{-i\theta}$ is $\vec{e}_{\lambda_2} = (i, -1)^T$. Verify that it satisfies the eigenvalue equation $R_\theta \vec{e}_{\lambda_2} = e^{-i\theta} \vec{e}_{\lambda_2}$.

I know it was quite a struggle to find the eigenvectors of this rotation matrix, but this is the case in general when finding eigenvectors. You must complete the null space calculation steps for each eigenspace, and this takes a long time. Be sure to practice finding eigenvectors by hand—I can pretty much guarantee you'll need this skill on your linear algebra final. And don't forget to give yourself a pat on the back when you're done!

Properties of the Hermitian transpose operation

The Hermitian transpose obeys the following properties:

- $(A + B)^\dagger = A^\dagger + B^\dagger$
- $(AB)^\dagger = B^\dagger A^\dagger$
- $(ABC)^\dagger = C^\dagger B^\dagger A^\dagger$
- $(A^\dagger)^{-1} = (A^{-1})^\dagger$

Note these are the same properties as the regular transpose operation from Section 2.3 (see page 147).

Conjugate linearity in the first input

The complex inner product we defined is linear in the second input and *conjugate-linear* in the first input:

$$\begin{aligned}\langle \vec{v}, \alpha \vec{a} + \beta \vec{b} \rangle &= \alpha \langle \vec{v}, \vec{a} \rangle + \beta \langle \vec{v}, \vec{b} \rangle, \\ \langle \alpha \vec{a} + \beta \vec{b}, \vec{w} \rangle &= \bar{\alpha} \langle \vec{a}, \vec{w} \rangle + \bar{\beta} \langle \vec{b}, \vec{w} \rangle.\end{aligned}$$

Keep this in mind every time you deal with complex inner products. The complex inner product is not symmetric since it requires that the complex conjugation be performed on the first input. Remember, instead of $\langle \vec{v}, \vec{w} \rangle \neq \langle \vec{w}, \vec{v} \rangle$, we have $\langle \vec{v}, \vec{w} \rangle = \overline{\langle \vec{w}, \vec{v} \rangle}$.

The choice of complex conjugation in the first input is a matter of convention. In this text, we *defined* the inner product $\langle \cdot, \cdot \rangle$ with the † operation on the first input, which is known as the *physics convention*. Some old mathematics texts define the inner product of complex vectors using the complex conjugation on the second input, which makes the inner product linear in the first input and conjugate-linear in the second input. This convention is fine, too. The choice of convention doesn't matter, as long as one of the inputs is conjugated to ensure the inner product obeys the positive semidefinite requirement $\langle \vec{u}, \vec{u} \rangle \geq 0$.

Function inner product

In the section on inner product spaces, we discussed the notion of the vector space of all real-valued functions of a real variable $f: \mathbb{R} \rightarrow \mathbb{R}$, and defined an inner product between functions:

$$\langle f, g \rangle = \int_{-\infty}^{\infty} f(x)g(x) dx.$$

Suppose we have two complex-valued functions $f(x)$ and $g(x)$:

$$f: \mathbb{R} \rightarrow \mathbb{C}, \quad g: \mathbb{R} \rightarrow \mathbb{C}.$$

We define the inner product for complex-valued functions as

$$\langle f, g \rangle = \int_{-\infty}^{\infty} \overline{f(x)} g(x) dx.$$

The complex conjugation of one of these functions ensures that the inner product of a function with itself results in a real number. The function inner product measures the *overlap* between $f(x)$ and $g(x)$.

Linear algebra over other fields

We can carry out linear algebra calculations over any *field*. A field is a set of numbers for which an addition, subtraction, multiplication, and division operations are defined. The addition and multiplication operations we define must be associative and commutative, and multiplication must distribute over addition. Furthermore, a field must contain an additive identity element (denoted 0) and a multiplicative identity element (denoted 1). The properties of a field are essentially all the properties of the numbers you're familiar with: $\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{C}$.

The focus of our discussion in this section was to show that the linear algebra techniques we learned for manipulating real numbers work equally well with complex numbers. This shouldn't be too surprising since, after all, linear algebra manipulations boil down to arithmetic manipulations of the components of vectors and matrices. Since both real numbers and complex numbers can be added, subtracted, multiplied, and divided, we can study linear algebra over both $\mathbb{R}$ and $\mathbb{C}$.

We can also perform linear algebra over *finite fields*. A *finite field* is a set $\mathbb{F}_q \stackrel{\text{def}}{=} \{0, 1, 2, \dots, q-1\}$, where q is a prime number or the power of a prime number. All the arithmetic operations in this field are performed *modulo* the number q , which means all arithmetic operations must result in answers in the set $\{0, 1, 2, \dots, q-1\}$. If the result of an operation falls outside this set, we either add or subtract q until the number falls in the set $\mathbb{F}_q$. Consider the finite field $\mathbb{F}_5 \stackrel{\text{def}}{=} \{0, 1, 2, 3, 4\}$. To add two numbers in $\mathbb{F}_5$, proceed as follows:

$$\begin{aligned} (3 + 3) \bmod 5 &= 6 \bmod 5 && \text{(too big, so subtract 5)} \\ &= 1 \bmod 5. \end{aligned}$$

Similarly, for subtraction,

$$\begin{aligned} (1 - 4) \bmod 5 &= (-3) \bmod 5 && \text{(too small, so add 5)} \\ &= 2 \bmod 5. \end{aligned}$$

The field of binary numbers $\mathbb{F}_2 \stackrel{\text{def}}{=} \{0, 1\}$ is an important finite field used in many areas of communications, engineering, and cryptography. In the next chapter we'll discuss the one-time cryptosystem, which allows for secure communication of messages encoded in binary (Section 7.9). We'll also discuss error-correcting codes that enable the reliable transmission of information over noisy communication channels (Section 7.10). For example, the data packets that your cell phone sends via radio waves are first linearly encoded using a matrix-vector product operation carried out over the finite field $\mathbb{F}_2$.

At first, thinking of linear algebra over the finite field $\mathbb{F}_q$ may seem complicated, but don't worry about it. It's the same stuff we've been practicing—you just need to apply “mod q ” after every arithmetic calculation. All of your intuition about dimensions and orthogonality, and all the computational procedures you know, are still applicable.

The field of rational numbers $\mathbb{Q}$ is another example of a field that's often used in practice. Solving systems of equations using rational numbers on computers is interesting because the answers obtained are exact—using rational numbers allows us to avoid many of the numerical accuracy problems associated with floating point numbers.

Discussion

The adjoint operator

Though we used the term *Hermitian transpose* and the notation $A^\dagger$ throughout this section, it's worth commenting that mathematicians prefer the term *adjoint* for the same operation, and denote it A^* . Recall we previously discussed the concept of an *adjoint linear transformation* $T_{M^T} : \mathbb{R}^m \rightarrow \mathbb{R}^n$, which corresponds to the multiplication of a matrix M by a row vector from the left $T_{M^T}(\vec{a}) \stackrel{\text{def}}{=} \vec{a}^T M$ (see page 238). We didn't use the term “transpose” then because transposing is something you do to matrices. Instead, we used the math term *adjoint*, which precisely describes the notion of the “transpose of a linear transformation.” Since we're on the topic of math terminology, it should be noted that some mathematicians use the term *adjoint operator* instead of *adjoint linear transformation*, since they call operators what we call *linear transformations*.

Matrix quantum mechanics

Guess what? Understanding linear algebra over the complex field means you understand quantum mechanics! Quantum mechanics unfolds in a complex inner product space (called a Hilbert space).

If you understood the material in this section, you should be able to understand the axioms of quantum mechanics at no additional mental cost. If you're interested in this kind of stuff you should read Chapter 9.

Exercises

E6.19 Calculate **a)** $(2 + 5i) - (3 + 4i)$, **b)** $(2 + 5i)(3 + 4i)$, and **c)** $(2 + 5i)/(3 + 4i)$.

E6.20 Find the characteristic polynomial, the eigenvalues and their associated eigenvectors for the matrix $B = \begin{bmatrix} -2 & -1 \\ 5 & 2 \end{bmatrix}$.

E6.21 Find the change-of-basis matrix ${}_F[\mathbb{1}]_B$ that transforms vectors expressed in the standard basis $B = \{(1, 0, 0, 0), (0, 1, 0, 0), (0, 0, 1, 0), (0, 0, 0, 1)\}$ to vectors in the *Fourier basis* $F = \{(1, 1, 1, 1), (1, -i, -1, i), (1, -1, 1, -1), (1, i, -1, -i)\}$.

Hint: Find the change-of-basis matrix ${}_B[\mathbb{1}]_F$ first.

E6.22 Is the matrix $\begin{bmatrix} 1 & 2+i \\ 3-i & 4 \end{bmatrix}$ Hermitian?

6.8 Theory problems

It's now time to test your understanding of the theoretical concepts we discussed in this chapter. The eigenvector equation $A\vec{e}_\lambda = \lambda\vec{e}_\lambda$ is one of the deepest ideas in linear algebra. I've prepared several problems so you can challenge yourself and test your understanding of eigenvalues and eigenvectors. The problems will test your theoretical understanding as well as your stamina, because computing eigenvectors requires many steps of arithmetic and takes a long time. The first eigenvector problem you'll solve might take you up to an hour. Don't be alarmed by this—that's totally normal. After solving a few eigenvector problems, your problem-solving time will drop to 30 minutes; and quickly after that you'll be able to solve eigenvalue problems easily in 15 minutes.

It's up to you how fluent you want to become. Certainly if you have a linear algebra exam coming up, it would be good to solve all the problems and maybe even solve problems in other books, too. If you're just reading about linear algebra for fun, you probably don't need to suffer through the steps of finding eigenvalues using only pen and paper. Solve the problems using SymPy instead—you can't say no to that!

In this chapter we also learned about abstract vector spaces, another important theoretical idea in linear algebra. All the techniques you've learned about vectors can be applied to polynomials, matrices, functions, and other vector-like objects. That's all nice in theory, but we're going to move beyond passive appreciation and get into the nitty gritty by solving problems that involve bases, linear independence, dimensions, and orthogonality in abstract vector spaces. It might seem like crazy stuff, but if you trust the idea of equivalent representations and the abstract notion of a linear transformation, you'll see it's all good and that you can work with abstract vectors.

Finally, the problems that involve linear algebra over the complex field will serve as the final review of what you've learned in this book. This is the final boss. You'll be asked to review and combine your computational, geometric, and theoretical linear algebra skills, applying them to vectors and matrices with complex coefficients. Are you ready for this?

I'm not going to lie to you and say the problems are easy, but this is the final push, so hang in there and you'll be done with all the linear algebra theory in just a few hours. After finishing the problems in this chapter, the rest of the book winds down with three chapters of cool applications, which are much lighter reading. So grab a pen, pull out some paper and kick some problem ass!

P6.1 Yuna wants to cheat on her exam and she needs your help. Please help her compute the eigenvalues of the following matrices, and slip her the piece of paper carefully so the teacher doesn't notice. Yuna will give you a chocolate bar to thank you.

$$\text{a) } \begin{bmatrix} 3 & 1 \\ 12 & 2 \end{bmatrix} \quad \text{b) } \begin{bmatrix} 0 & 1 & 0 \\ 2 & 0 & 2 \\ 0 & 1 & 0 \end{bmatrix}$$

P6.2 Find the eigenvalues of the following matrices.

$$\text{a) } \begin{bmatrix} 4 & 2 \\ 0 & 5 \end{bmatrix} \quad \text{b) } \begin{bmatrix} 3 & 1 \\ 1 & 2 \end{bmatrix} \quad \text{c) } \begin{bmatrix} 2 & 0 & 1 \\ 1 & 2 & 0 \\ 0 & 4 & -1 \end{bmatrix} \quad \text{d) } \begin{bmatrix} -3 & 0 & 0 \\ 4 & 1 & 0 \\ 2 & 1 & -1 \end{bmatrix}$$

P6.3 Compute the eigenvalues of the matrix $A = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$.

P6.4 Show that the vectors $\vec{e}_1 = (1, \frac{1}{\varphi})^T$ and $\vec{e}_2 = (1, -\varphi)^T$ are eigenvectors of the matrix $A = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$. What are the eigenvalues associated with these eigenvectors?

Hint: Compute $A\vec{e}_1$ and $A\vec{e}_2$ to see what happens. Use the fact that φ satisfies the equation $\varphi^2 - \varphi - 1 = 0$ to simplify expressions.

P6.5 We can write the matrix $A = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$ as the product of three matrices $Q\Lambda X$, where Q contains the eigenvectors of A , and Λ contains its eigenvalues:

$$\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 1 \\ \frac{1}{\varphi} & -\varphi \end{bmatrix}}_Q \underbrace{\begin{bmatrix} \varphi & 0 \\ 0 & -\frac{1}{\varphi} \end{bmatrix}}_\Lambda \underbrace{\begin{bmatrix} ? & ? \\ ? & ? \end{bmatrix}}_X.$$

Find the matrix X .

P6.6 Compute the eigenvalues and eigenvectors of these matrices:

$$\text{a) } A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \text{b) } B = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ -6 & -1 & 4 \end{bmatrix}$$

P6.7 Given $A = \begin{bmatrix} 2 & 2 \\ 5 & -1 \end{bmatrix}$, find A^{10} .

P6.8 Consider the sequence of triples $\{(x_n, y_n, z_n)\}_{n=0,1,2,\dots}$ produced according to the formula:

$$\underbrace{\begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 \\ \frac{1}{8} & \frac{3}{4} & \frac{1}{8} \\ 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}}_M \begin{bmatrix} x_n \\ y_n \\ z_n \end{bmatrix} = \begin{bmatrix} x_{n+1} \\ y_{n+1} \\ z_{n+1} \end{bmatrix}.$$

Give a formula for $(x_\infty, y_\infty, z_\infty)$ in terms of (x_0, y_0, z_0) . This recurrence relation is related to “surface smoothing” algorithms used in 3D graphics; see <https://youtu.be/mX0NB9IyYpU> for more explanations.

Hint: Compute the eigenvalues λ_1 , λ_2 , and λ_3 of the matrix M . What will happen to the eigenvalues if you raise them to the power ∞ ?

P6.9 Explain why an $n \times n$ matrix can have at most n different eigenvalues.

P6.10 Prove that $T : V \rightarrow V$ is an invertible linear transformation if and only if $\lambda = 0$ is not an eigenvalue of T .

P6.11 An unknown matrix $A \in \mathbb{R}^{3 \times 3}$ has eigenvalues $\lambda_1 = 2$, $\lambda_2 = -3$, and $\lambda_3 = 5$. Calculate the value of the following expressions:

$$\mathbf{a)} \det(2A) \quad \mathbf{b)} \det(A^2) \quad \mathbf{c)} \det(A^{-1}) \quad \mathbf{d)} \operatorname{Tr}(A + 15A^{-1} + A^T)$$

P6.12 Prove that diagonal matrices are symmetric matrices.

P6.13 Check whether the following matrices are orthogonal or not:

$$\mathbf{a)} \begin{bmatrix} \frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} \\ 0 & 1 & 0 \\ \frac{1}{\sqrt{2}} & 0 & -\frac{1}{\sqrt{2}} \end{bmatrix} \quad \mathbf{b)} \begin{bmatrix} 1 & -1 & 1 \\ 1 & -1 & -1 \\ 0 & 1 & 0 \end{bmatrix} \quad \mathbf{c)} \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$

P6.14 Given a normal matrix $M \in \mathbb{R}^{n \times n}$ ($M^T M = M M^T$), show that $\mathcal{C}(M) = \mathcal{C}(M^k)$ and $\mathcal{N}(M) = \mathcal{N}(M^k)$ for all k .

Hint: Think of the eigendecomposition for normal matrices.

P6.15 Given A and B are two positive semidefinite matrices, show that the sum $A + B$ is also a positive semidefinite matrix.

P6.16 Two friends are arguing over a matrix question. Jane claims that a matrix is orthogonal if and only if its columns are an orthonormal basis. John says that a matrix is orthogonal if and only if its rows are an orthonormal basis. Use the rows-times-columns interpretation of the matrix product to figure out who is right.

P6.17 Given that O is an orthogonal matrix, find the inverse of $2O$.

P6.18 Prove that the set of polynomials of degree two $P_2(t)$ is a vector space. Consider arbitrary elements $\mathbf{p} = a_0 + a_1 t + a_2 t^2$ and $\mathbf{q} = b_0 + b_1 t + b_2 t^2$, arbitrary constants α and β , and verify that all eight axioms stated in Section 6.3 hold.

P6.19 Let V be the set of two-dimensional vectors of real numbers, with addition defined as $(a_1, a_2) + (b_1, b_2) = (a_1 + b_1, a_2 + b_2)$ and scalar multiplication defined as $c \cdot (a_1, a_2) = (ca_1, ca_2)$. Is $(V, \mathbb{R}, +, \cdot)$ a vector space? Justify your answer.

Hint: Check whether scaling by zero obeys the vector space axioms.

P6.20 Let $V = \{(a_1, a_2)\}$, with $a_1, a_2 \in \mathbb{R}$. Define vector addition as $(a_1, a_2) + (b_1, b_2) = (a_1 + 2b_1, a_2 + 3b_2)$ and scalar multiplication as $c \cdot (a_1, a_2) = (ca_1, ca_2)$. Is $(V, \mathbb{R}, +, \cdot)$ a vector space? Justify your answer.

Hint: Check that the associative property holds.

P6.21 Determine whether the following subsets of $\mathbb{R}^3$ are subspaces:

- a) $\{(x, y, z) \in \mathbb{R}^3 \mid x + y + z = 3\}$
- b) $\{(x, y, z) \in \mathbb{R}^3 \mid x + y + z = 0\}$
- c) $\{(x, y, z) \in \mathbb{R}^3 \mid x = 2y = 3z\}$

P6.22 Consider an arbitrary matrix $A \in \mathbb{R}^{2 \times 2}$ and its representation as a vector of coordinates with respect to B_S : $\vec{A} = (a_{11}, a_{12}, a_{21}, a_{22})_{B_S}$. Suppose we want to compute the matrix trace operation in terms of the vector dot product. What vector $\vec{v} \in \mathbb{R}^4$ makes this equation true $\text{Tr}(A) = \vec{v} \cdot \vec{A}$?

P6.23 Repeat the previous question, but now think of $\vec{A}$ as a 4×1 matrix. Find the matrix V that implements the trace operation: $\text{Tr}(A) = V\vec{A}$. Assume the standard matrix-matrix product is used.

P6.24 Prove that the set of polynomials of degree two $P_2(t)$ is a vector subspace of the vector space $P_3(t)$.

Hint: A subspace of $\mathbb{R}^3$ must be closed under addition and scalar multiplication, and contain the zero element.

P6.25 Give an example of a subset of $\mathbb{R}^2$ that is closed under scalar multiplication, but is not a subspace.

P6.26 Give an example of a subset of $\mathbb{R}^2$ that is closed under addition, but is not a subspace.

P6.27 Consider the linear transformation $T : P_2(x) \rightarrow P_2(x)$ defined as $T(a_0 + a_1x + a_2x^2) = a_2x^2$. Find the matrix representation of T with respect to the basis $\{1, x, x^2\}$, and compute the eigenvalues of T .

P6.28 Find the dimension of the vector space of functions that satisfy the differential equation $f'(t) + f(t) = 0$.

Hint: Which function is equal to a multiple of its own derivative?

P6.29 Let V be the vector space consisting of all functions of the form $\alpha e^{2x} \cos x + \beta e^{2x} \sin x$. Consider the linear transformation $L : V \rightarrow V$, $L(f) = f' + f$. Find the matrix representing L with respect to the basis $\{e^{2x} \cos x, e^{2x} \sin x\}$.

P6.30 Find the matrix representation of the derivative operator $Dp(x) = \frac{d}{dx}p(x)$ for the vector space of polynomials of degree three $p(x) = a_0 + a_1x + a_2x^2 + a_3x^3$, represented as coordinates $(a_0, a_1, a_2, a_3)^T$.

P6.31 Give a basis for the vector space of 3×3 diagonal matrices.

P6.32 What is the dimension of the vector space of 3×3 symmetric matrices?

Hint: See page 313 for definition.

P6.33 How many matrices are there in a basis for the vector space of 3×3 Hermitian matrices?

Hint: See page 346 for definition.

P6.34 The linear operator $T : \mathbb{R}^{2 \times 2} \rightarrow \mathbb{R}^{2 \times 2}$ is defined by the equation

$$T\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) = \begin{bmatrix} 2c & a+c \\ b-2c & d \end{bmatrix}.$$

Find the eigenvalues and eigenvectors of this linear operator.

Hint: The eigenvectors of T are 2×2 matrices.

P6.35 Show that the following functions, called Laguerre polynomials, are orthogonal with respect to the inner product $\langle f(x), g(x) \rangle = \int_0^\infty f(x)g(x)e^{-x}dx$.

$$L_0(x) = 1, \quad L_1(x) = 1 - x, \quad L_2(x) = \frac{1}{2}(x^2 - 4x + 2).$$

Hint: Use the formula $\int_0^\infty f(x)g'(x)dx = f(x)g(x)|_0^\infty - \int_0^\infty f'(x)g(x)dx$.

P6.36 Let $\vec{v}_1, \vec{v}_2, \vec{v}_3$ be vectors from a vector space V . Given $\langle \vec{v}_1, \vec{v}_2 \rangle = 3$, $\langle \vec{v}_2, \vec{v}_3 \rangle = 2$, $\langle \vec{v}_1, \vec{v}_3 \rangle = 1$, $\langle \vec{v}_1, \vec{v}_1 \rangle = 1$, and $\langle \vec{v}_2, \vec{v}_1 + \vec{v}_2 \rangle = 13$, find:

$$\mathbf{a}) \langle \vec{v}_1, 2\vec{v}_2 + 3\vec{v}_3 \rangle \quad \mathbf{b}) \langle 2\vec{v}_1 - \vec{v}_2, \vec{v}_1 + \vec{v}_3 \rangle \quad \mathbf{c}) \|\vec{v}_2\|$$

P6.37 Prove the Cauchy-Schwarz inequality $|\langle \mathbf{u}, \mathbf{v} \rangle| \leq \|\mathbf{u}\| \|\mathbf{v}\|$.

Hint: It is true that $\|\mathbf{a}\| > 0$ for any vector $\mathbf{a}$. Use this fact to expand the expression $\|\mathbf{u} - c\mathbf{v}\| > 0$, choosing $c = \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle}$.

P6.38 Prove the triangle inequality $\|\mathbf{u} + \mathbf{v}\| \leq \|\mathbf{u}\| + \|\mathbf{v}\|$.

Hint: Compute $\|\mathbf{u} + \mathbf{v}\|$ as an inner product and simplify the expression using the fact $\langle \mathbf{a}, \mathbf{b} \rangle \leq \|\mathbf{a}\| \|\mathbf{b}\|$ for all vectors $\mathbf{a}$ and $\mathbf{b}$.

P6.39 Perform the Gram-Schmidt orthogonalization procedure on the following basis for $\mathbb{R}^2$: $\{(0, 1), (-1, 0)\}$.

P6.40 Perform Gram-Schmidt orthogonalization on vectors $\vec{v}_1 = (1, 1)$ and $\vec{v}_2 = (0, 1)$ to obtain an orthonormal basis.

P6.41 Convert the vectors $(3, 1)$ and $(-1, 1)$ into an orthonormal basis.

P6.42 Consider the vector space $P_2(x)$ of polynomials of degree two in combination with the inner product $\langle \mathbf{f}, \mathbf{g} \rangle = \int_{-1}^1 f(x)g(x)dx$. The functions $\mathbf{f}_1(x) = 1$, $\mathbf{f}_2(x) = x$, and $\mathbf{f}_3(x) = x^2$ are linearly independent and form a basis for $P_2(x)$. Transform $\{\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3\}$ into an orthonormal basis for $P_2(x)$.

P6.43 Find the eigendecomposition of the matrix $A = \begin{bmatrix} 2 & 0 & -5 \\ 0 & 2 & 0 \\ 0 & 0 & -3 \end{bmatrix}$.

P6.44 Compute the following expressions:

$$\mathbf{a}) |3i - 4| \quad \mathbf{b}) \overline{2 - 3i} \quad \mathbf{c}) (3i - 1) + \overline{3 - 2i} \quad \mathbf{d}) |\overline{-3i} - \overline{-4i + 5}|$$

P6.45 Given complex matrices $A = \begin{bmatrix} 2+i & -1+2i \\ 3+2i & -2i \end{bmatrix}$, $B = \begin{bmatrix} 2-i & 3-2i \\ 5+i & -5+5i \end{bmatrix}$, and $C = \begin{bmatrix} 1+2i & i \\ 3-i & 8 \\ 4+2i & 1-i \end{bmatrix}$, find $A + B$, CB , and $(2 + i)B$.

P6.46 Find the eigenvalues of the following matrices:

a) $\begin{bmatrix} 3 & -2 \\ 1 & 1 \end{bmatrix}$

b) $\begin{bmatrix} 3 & -9 \\ 4 & -3 \end{bmatrix}$

c) $\begin{bmatrix} 3 & -13 \\ 5 & 1 \end{bmatrix}$

P6.47 Determine all the eigenvalues of $A = \begin{bmatrix} 1+i & 1 \\ 2 & 1-i \end{bmatrix}$. For each eigenvalue λ of A , find the set of eigenvectors corresponding to λ . Determine whether or not A is diagonalizable; if so, find an invertible matrix Q and a diagonal matrix Λ such that $Q^{-1}AQ = \Lambda$.

P6.48 Show that the set $B_a = \{1 + ix, 1 + x + ix^2, 1 + 2ix\}$ is a basis for the space of polynomials with complex coefficients of degree at most two.

P6.49 Given the matrix $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, show that A has a real eigenvalue if and only if $(a - d)^2 + 4bc \geq 0$.

P6.50 Given a normal matrix $M \in \mathbb{C}^{n \times n}$, prove that $\mathcal{C}(M) = \mathcal{C}(M^\dagger)$.

Hint: Think of the eigendecomposition for normal matrices.

P6.51 Prove that the eigenvalues of Hermitian matrices are real.

Hint: Consider the equation $A\vec{e} = \lambda\vec{e}$ and the expression $\vec{e}^\dagger A\vec{e}$.

P6.52 Show that eigenvalues of unitary matrices have magnitude one.

Hint: Use the fact that unitary matrices are length preserving.

P6.53 A matrix is *nilpotent* if it becomes the zero matrix when repeatedly multiplied by itself. We say A is nilpotent if $A^k = 0$ for some power k . A nilpotent matrix has only the eigenvalue zero, hence its trace and determinant are zero. Are the following matrices nilpotent?

a) $\begin{bmatrix} -2 & 4 \\ -1 & 2 \end{bmatrix}$ b) $\begin{bmatrix} 3 & 1 \\ 1 & 3 \end{bmatrix}$ c) $\begin{bmatrix} -3 & 2 & 1 \\ -3 & 2 & 1 \\ -3 & 2 & 1 \end{bmatrix}$ d) $\begin{bmatrix} 1 & 1 & 4 \\ 3 & 0 & -1 \\ 5 & 2 & 7 \end{bmatrix}$

e) $\begin{bmatrix} 45 & -22 & -19 \\ 33 & -16 & -14 \\ 69 & -34 & -29 \end{bmatrix}$ f) $\begin{bmatrix} 5 & -3 & 2 \\ 15 & -9 & 6 \\ 10 & -6 & 4 \end{bmatrix}$

P6.54 Given M is a normal matrix, show that the matrix $\sqrt{M}$ exists.

Hint: Does an eigendecomposition of M exist?

Chapter 7

Applications

In this chapter, we'll learn about applications of linear algebra. We'll cover a wide range of topics from different areas of science, business, and technology to give you an idea of the spectrum of possible calculations based on vector and matrix algebra. Don't worry if you're not able to follow all the details in each section—we're taking a broad approach here, covering many different topics in the hope that some will interest you. Note that most of the material covered in this chapter is not likely to show up on your linear algebra final, so no pressure—this is just for fun.

Before we start, I want to say a few words about scientific ethics. Linear algebra is a powerful tool for solving problems and modelling the real world. But with great power comes great responsibility. I hope you'll make an effort to think about the ethical implications when you use linear algebra to solve problems. Certain applications of linear algebra, like building weapons, interfering with crops, and building mathematically-complicated financial scams are clearly evil, so you should avoid them. Other areas where linear algebra can be applied are not so clear-cut: perhaps you're building a satellite localization service to find missing people in emergency situations, but the same technology could be used by governments to spy on and persecute your fellow citizens. Do you want to be the person responsible for bringing about an Orwellian state? All I ask of you is to run a quick "System check" before you set to work on a project: ask yourself "Am I working for the System?" Don't just say "It's my job" and proceed without caution. If you find what you're doing for your employer is unethical, then maybe you should find a different job. There are a lot of jobs out there for people who know math, and if the bad guys can't hire qualified people like you, their power will decrease—and that's a good thing.


Our System check is complete. On to the applications!

7.1 Balancing chemical equations

Suppose you're given the chemical equation $\text{H}_2 + \text{O}_2 \rightarrow \text{H}_2\text{O}$, which indicates that hydrogen molecules (H_2) and oxygen molecules (O_2) can combine to produce water molecules (H_2O). Chemical equations describe how a set of *reactants* are transformed into a set of *products*. In this case, the reactants are hydrogen and oxygen molecules and the products are water molecules.


The equation $\text{H}_2 + \text{O}_2 \rightarrow \text{H}_2\text{O}$ is misleading since it doesn't tell us the correct *stoichiometric ratios*: how much of each type of molecule is consumed and produced. We say the equation is not *balanced*. To *balance* the equation, we must add coefficients in front of each reactant and each product, so that the total number of atoms on both sides of the reaction is the same: $2\text{H}_2 + \text{O}_2 \rightarrow 2\text{H}_2\text{O}$. Two hydrogen molecules are required for each oxygen molecule, since water molecules contain one oxygen and two hydrogen atoms.

Let's look at another example. The combustion of methane gas is described by the following chemical equation:



We want to answer the following two questions. How many molecules of oxygen will be consumed during the combustion of 1000 molecules of methane? How many CO_2 molecules will be produced as a result?

Before we can answer such questions, we must find the coefficients a , b , c , and d that balance the methane-combustion equation:



For the equation to be balanced, the same number of atoms of each type must appear on each side of the equation. For the methane combustion reaction to be balanced, the following equations must be satisfied:

$$\begin{array}{ll} a = c & \text{for C atoms to be balanced,} \\ 4a = 2d & \text{for H atoms to be balanced,} \\ 2b = 2c + d & \text{for O atoms to be balanced.} \end{array}$$


We can move the c and d terms to the left side of each equation and rewrite the system of equations as a matrix equation:

$$\begin{array}{rcl} a & -c & = 0 \\ 4a & & -2d = 0 \\ 2b & -2c & -d = 0 \end{array} \quad \Rightarrow \quad \underbrace{\begin{bmatrix} 1 & 0 & -1 & 0 \\ 4 & 0 & 0 & -2 \\ 0 & 2 & -2 & -1 \end{bmatrix}}_A \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$


We're looking for the vector of coefficients $\vec{x} = (a, b, c, d)$, which is the solution to the null space problem $A\vec{x} = \vec{0}$. See E4.16 on page 252 for details. The RREF of A contains three pivots and one free variable. The solution to the null space problem is one-dimensional:

$$\{(a, b, c, d)\} = \text{span}\left(\left(\frac{1}{2}, 1, \frac{1}{2}, 1\right)\right).$$

The solution is the span of the vector $(\frac{1}{2}, 1, \frac{1}{2}, 1)$. The solution space is infinite since any balanced equation will remain balanced if we double or triple the amount of reactants and products. Choosing the coefficients as suggested by the solution to the null space problem gives $\frac{1}{2}\text{CH}_4 + \text{O}_2 \rightarrow \frac{1}{2}\text{CO}_2 + \text{H}_2\text{O}$, which is a balanced equation. The convention in chemistry is to choose integer coefficients for reactants and products, so we'll multiply the equation by two to obtain the final answer:



Balancing chemical equations may not seem like the most exciting technique ever, but it's a useful skill for many chemistry calculations. It's a good start to know that substances A and B can transform into substances C and D , but it's better to know *how much* of each reactant is consumed and how much of each product is produced per "unit of reaction." Once we've identified one "unit of reaction," we can calculate other quantities in terms of it, such as measuring the energy released per unit of reaction. The combustion of 1[mol] ($= 6.022 \times 10^{23}$ molecules) of methane produces 890[kJ] of heat:



Now *this* is cool. If you're heating your chalet with methane gas, and you know how many joules of heat you'll need, then the balancing chemical equation will help you calculate how many litres of methane you need to stock in order to survive this winter.

Exercises

The exercises below aren't difficult, so you should totally try to solve them. Plus, they'll give you some extra practice with the Gauss-Jordan elimination procedure. It's been *ages* since Chapter 3, so a refresher can't hurt.

Watch out, chemistry is coming at ya! Can you handle it?

E7.1 Balance the chemical equation $\text{Al} + \text{O}_2 \rightarrow \text{Al}_2\text{O}_3$.

E7.2 Balance the equation $\text{Fe}(\text{OH})_3 + \text{HCl} \rightarrow \text{FeCl}_3 + \text{H}_2\text{O}$.

7.2 Input–output models in economics

Suppose you're the top economic official of a small country and you want to make a production plan for the coming year. For the sake of simplicity, let's assume your country produces only three commodities: electric power, wood, and aluminum. Your job is to choose the production rates of these commodities: x_e , x_w , and x_a . Your country must produce enough to satisfy both the internal demand and the external demand for these commodities. The problem is complicated because the production rates in one industry may affect the production rates of other industries. For instance, it takes some electric power to produce each unit of aluminum, so your production plan must account for both external demand for electric power, as well as *internal demand* for electric power for aluminum production. When complex interdependences exist between the different internal industries, as is often the case, it makes the process of picking the right production rates more complex.

In reality, most high-ranking government officials make their decisions about which industry to sponsor based on the dollar amounts of the kickbacks and bribes they received during the previous year. Let's ignore reality for a moment and assume you're an honest economist interested in using math to do what is right for the country instead of abusing your position of power like a blood-thirsty leech.

Let's assume the electric production x_e must satisfy an external demand of 25 units, plus an additional 0.05 units for each unit of wood produced (electricity needed for saw mill operations) and an additional 0.3 units for each unit of aluminum produced. The wood production must be 10 units plus additional small amounts that depend on x_e and x_a (wood for construction). The production of aluminum must match 14 units of external demand plus an additional 0.1 units for each unit of electric power (for repairs of electric cables). We can model the interdependence between the industries using the following system of equations:

$$\begin{array}{rclcl}
 x_e & = & 25 & + & 0.05x_w + 0.3x_a \\
 x_w & = & 10 & + & 0.01x_e + 0.01x_a \\
 x_a & = & \underbrace{14}_{\text{external demand}} & + & \underbrace{0.1x_e}_{\text{internal demand}}.
 \end{array}$$

You can use linear algebra to solve this complicated industry interdependence problem and choose appropriate production rates. Ex-

press the system of equations as a matrix equation:

$$\underbrace{\begin{bmatrix} x_e \\ x_w \\ x_a \end{bmatrix}}_{\vec{x}} = \underbrace{\begin{bmatrix} 25 \\ 10 \\ 14 \end{bmatrix}}_{\vec{d}} + \underbrace{\begin{bmatrix} 0 & 0.05 & 0.3 \\ 0.01 & 0 & 0.01 \\ 0.1 & 0 & 0 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_e \\ x_w \\ x_a \end{bmatrix}}_{\vec{x}}.$$

This is known as a *Leontief input-output model* in honour of Wassily Leontief, who first applied linear algebra techniques to economics, and was awarded the Nobel prize for this contribution.

To find the appropriate production rates, we must solve for the unknown $\vec{x} = (x_e, x_w, x_a)$ in the above equation. The equation $\vec{x} = \vec{d} + A\vec{x}$ is a little unusual, but we can solve it using standard techniques:

$$\mathbb{1}\vec{x} = \vec{d} + A\vec{x} \Rightarrow (\mathbb{1} - A)\vec{x} = \vec{d} \Rightarrow \vec{x} = (\mathbb{1} - A)^{-1}\vec{d}.$$

The solution to the electricity, wood, and aluminum production scenario is $\vec{x} = (x_e, x_w, x_a) = (30.64, 10.48, 17.06)$. See P3.11 for the details of the solution.

Note the electricity production rate is significantly higher than the external demand, in order to account for the internal demand of electricity for the aluminum production.

Links

[History of the Leontief input-output model in economics]
https://en.wikipedia.org/wiki/Input-output_model

Exercises

E7.3 A group of farmers has formed a cooperative that specializes in three crops: arugula (*eruca sativa*), broccoli, and carrots. They need to produce 10 tons of arugula, 200 tons of broccoli, and 300 tons of carrots for sale. The farmers also need some additional production to feed themselves and the work-for-food tourists that come to help them scale up production during the harvest season. To feed the workers needed to grow 100 kg of arugula, they need to produce an extra 10 kg of broccoli and 10 kg of carrots. To produce 100 kg of broccoli, it takes 1 kg of arugula and 15 kilograms of carrots. To produce 100 kg of carrots, it takes 1 kg of arugula and 20 kilograms of broccoli. How much of each vegetable should the farmers produce?

7.3 Electric circuits

We can use Ohm's law to solve many circuit problems. To *solve* a circuit is to find the current flowing in each wire in the circuit, and the value of the voltage at every point in the circuit. Ohm's law, $V = IR$, tells us the voltage V required to "push" a current I through a resistor with resistance R . This simple equation is all we need to find the currents and voltages of any electric circuit involving wires, batteries, and resistors.

Since we're in Math Land, the units of the circuit quantities won't play a direct role in our analysis; but we'll still introduce them, since units will help us perform dimensional analysis of the equations. Voltages are measured in volts [V], currents are measured in Amperes [A], and resistance is measured in Ohms [Ω]. Intuitively, the resistance of a resistor measures how difficult it is to "push" current through it. Indeed, the units for resistance have the dimensions of Volt per Ampere: $[\Omega] = [V/A]$. The equation $V = RI$ tells us how much current $I[A]$ flows through a resistor with resistance $R[\Omega]$ connected to a voltage source with potential $V[V]$. Alternatively, if we know the current $I[A]$ and the resistance $R[\Omega]$, we can find V , the voltage applied to the resistor. A third way to use the equation $V = RI$ is to solve for the resistance R in cases when we know both V and I .

Example Your friend gives you a $121[\Omega]$ light bulb (a resistor) and asks you to connect it outdoors on the backyard porch to provide some extra lighting for a summer party. You run to the basement and find three different spools of electric wire: a green wire rated for currents of up to $0.5[A]$, a blue wire rated for currents of up to $1[A]$, and a red wire rated for currents of up to $2[A]$. Knowing that the voltage coming out of the wall socket¹ is $110[V]$, what is the lowest-rating wire you can use to connect the light bulb?

A simple calculation using $V = RI$ shows us that the current flowing through the lightbulb is $I = \frac{V}{R} = \frac{110[V]}{121[\Omega]} = 0.909[A]$. Thanks to your calculation, you choose the blue wire rated for $1[A]$ knowing you won't have problems with wires overheating and causing a fire.

* * *

It can be difficult to solve for all the voltages and currents in electric circuits that contain multiple voltage sources (batteries) and resistors

¹The voltage coming out of wall outlets is actually not a constant $110[V]$ but a sine-wave oscillating between $+155[V]$ and $-155[V]$, but in this example we'll treat the socket's output as a constant voltage of $110[V]$.

(light bulbs), like the circuit shown in Figure 7.1. Using the equation $V = RI$ for each resistor leads to several equations that must be solved simultaneously to find the unknowns. Did someone say “system of linear equations”? Linear algebra to the rescue!




Figure 7.1: An electric circuit with three batteries and three resistors.

Knowing linear algebra will enable you to solve even the most complicated circuit using row operations (Gauss–Jordan elimination) in one or two minutes. We’ll illustrate this application of linear algebra by solving the example circuit shown on the right, which involves three batteries (the parallel lines labelled + and –) and three resistors (the wiggly lines).

Background

Before we get started, let me introduce the minimum information you need to know about circuits: *Kirchhoff’s voltage law* and *Kirchhoff’s current law*.

The voltages in a circuit are related to the electric potential energy of the electrons flowing in the circuit. The electric potential is analogous to the concept of gravitational potential: a battery raises the electric potential of electrons like an elevator raises the gravitational potential of objects by increasing their height. Starting from this heightened potential, electrons flow through the circuit and lose potential when they pass through resistors. *Kirchhoff’s voltage law* (KVL) states that the sum of the voltage gains and losses along any *loop* in the circuit must sum to zero. Intuitively, you can think of KVL as a manifestation of the *conservation of energy* principle: the potential gained by electrons when they pass through batteries is lost in the resistors (in the form of heat). By the time the electrons complete their journey around any loop in the circuit, they must return to their initial potential.

Kirchhoff’s current law states that the total current flowing into a

wire junction must equal the total current flowing out of the junction. You can think of this as a manifestation of the *conservation of charge* principle: the total charge coming into a junction equals the total charge flowing out of the junction, because charges cannot be created or destroyed.

* * *

Recall that to solve a circuit is to find the currents that flow in each wire and the voltage across each resistor. We'll now illustrate how to solve a complicated circuit by introducing current variables, applying Kirchhoff's laws to obtain systems of linear equations, and solving these equations using linear algebra.

Using linear algebra to solve circuits

The first step is to define variables for each of the quantities of interest in the circuit as shown in Figure 7.2. We'll call I_1 the current that flows down through the middle wire of the circuit, which then splits into the current I_2 in the left branch and the current I_3 going to the right. Next we follow the currents in the circuit and label the terminals of each resistor “+” and “−” to indicate its *polarity*—the direction of the voltage drop across it. The rule to follow is simple: the label “+” goes on the side where the current enters the resistor, and the label “−” goes on the side where the current leaves the resistor. We write the “+” on the side where the current enters because electric potential always *drops* when passing through a resistor. It's important to keep track of the polarity of resistors when writing the KVL equations for the circuit.




Figure 7.2: The circuit with branch currents labelled. Each resistor is assigned a *polarity* relative to the current flowing through it.

We're now in a position to apply Kirchhoff's voltage and current laws to this circuit and obtain a set of equations that relate the unknown currents. Let's first apply Kirchhoff's voltage law to the loop

along the path A-B-C-D-A, computing the total of the voltage gains along this path:

$$+10 - R_1 I_2 + 5 - R_2 I_1 = 0.$$

Each battery produces a gain in potential for the electrons flowing through it. Each resistor leads to a drop in potential (negative gain) proportional to the current flowing through the resistor (recall $V = RI$).

Similarly, we obtain a second KVL equation by following the path A-F-E-D-A in the circuit:

$$+20 - R_3 I_3 - R_2 I_1 = 0.$$

We're measuring voltage *gains* along this loop, therefore the contribution of the battery is positive while the contributions of the resistors are negative.

We obtain a third equation linking the unknown currents from Kirchhoff's current law applied to junction A:

$$I_1 = I_2 + I_3.$$

Combining the three circuit equations, we obtain a system of three linear equations in three unknowns:

$$\begin{aligned} +10 - R_1 I_2 + 5 - R_2 I_1 &= 0, \\ +20 - R_3 I_3 - R_2 I_1 &= 0, \\ I_1 &= I_2 + I_3. \end{aligned}$$

Do you see where this is going? Perhaps rewriting the equations into the standard form we discussed in Section 3.1 will help you see what is going on:

$$\begin{aligned} R_2 I_1 + R_1 I_2 &= 15, \\ R_2 I_1 + R_3 I_3 &= 20, \\ I_1 - I_2 - I_3 &= 0. \end{aligned}$$

Rewriting the system of equations as an augmented matrix, we obtain:

$$\left[\begin{array}{ccc|c} R_2 & R_1 & 0 & 15 \\ R_2 & 0 & R_3 & 20 \\ 1 & -1 & -1 & 0 \end{array} \right].$$

Assume the values of the resistors are $R_1 = 1[\Omega]$, $R_2 = 2[\Omega]$, and $R_3 = 1[\Omega]$. We substitute these values into the augmented matrix, and then find its reduced row echelon form:

$$\left[\begin{array}{ccc|c} 2 & 1 & 0 & 15 \\ 2 & 0 & 1 & 20 \\ 1 & -1 & -1 & 0 \end{array} \right] \xrightarrow{\text{RREF}} \left[\begin{array}{ccc|c} 1 & 0 & 0 & 7 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 6 \end{array} \right].$$

The currents are $I_1 = 7[\text{A}]$, $I_2 = 1[\text{A}]$, and $I_3 = 6[\text{A}]$. The voltage drop across R_1 is $1[\text{V}]$ with polarity as indicated in Figure 7.2. The voltage drops across R_2 and R_3 are $14[\text{V}]$ and $6[\text{V}]$ respectively. Verify that these currents and voltages are consistent with the KVL equations we started from.

* * *

The combination of Kirchhoff's laws and linear algebra techniques allows us to solve any electric circuit. Analyzing complicated circuits will require solving large systems of equations; but, by using matrix methods, you should be able to handle even the most complex circuit. And you can always use SymPy to perform a RREF computation if it's ever too hairy to calculate by hand.


Other network flows

The same techniques we used for finding the flow of currents in circuits can also be applied to many other problems involving flow: the flow of cars on city streets, the flow of goods and services between economies, and the flow of information (data packets) through the internet. In each case, a different law governs how the flows travel through the network, yet we can use the same matrix techniques to solve the resulting systems of linear equations.

Exercises

E7.4 Obtain a third KVL equation for the path A-B-C-D-E-F-A in the circuit from Figure 7.2. Combine this third equation with the KVL equations for the paths A-B-C-D-A and A-F-E-D-A given on page 371. Are these three equations in three unknowns sufficient to solve the circuit without using the KCL equation?

E7.5 Write the three KVL equations for the three inner loops of the circuit. Obtain the KCL equations for junctions B and C.



Combine the five equations to form a matrix equation $R\vec{I} = \vec{V}$, where R is a 5×5 matrix, $\vec{I} = (I_1, I_2, I_3, I_4, I_5)^T$ is the vector of unknown currents, and $\vec{V} = (V_1, V_2, V_3, 0, 0)^T$ is a vector of constants.

7.4 Graphs

A *graph* is an abstract mathematical model that describes connections between a set of nodes. We call the nodes *vertices* and the connections *edges*. The graph is defined as a pair of sets $G = (V, E)$, where V is the set of vertices and E is the set of edges in the graph. We can also describe the edges by specifying the *adjacency matrix* of the graph.

Rather than define graphs formally and in detail, we'll look at a simple graph example to give you an idea of the main concepts and introduce graph notation. Figure 7.3 shows a small graph with five vertices and seven edges. This abstract link structure could represent many real-world scenarios: five websites and the hyperlinks between them, five Twitter accounts and their “following” relationships, or seven financial transactions between five businesses.




Figure 7.3: A simple graph with five vertices and seven edges.

The graph in Figure 7.3 is represented mathematically as $G = (V, E)$, where $V = \{1, 2, 3, 4, 5\}$ is the set of vertices, and $E = \{(1, 2), (1, 3), (2, 3), (3, 5), (4, 1), (4, 5), (5, 1)\}$ is the set of edges. Note the edge from vertex i to vertex j is represented as the pair (i, j) .

Adjacency matrix

The *adjacency matrix* representation of the graph in Figure 7.3 is a 5×5 matrix A that contains information about the edges in the graph. Specifically, $A_{ij} = 1$ if the edge (i, j) exists, otherwise $A_{ij} = 0$ if the edge doesn't exist:

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

Each row contains ones in the positions where edges exist. The adjacency matrix representation works in tandem with the integer labels of the vertices—Vertex 1 corresponds to the first row of A , Vertex 2 to the second row, and so on for the other rows. We don't need labels for the vertices since the labels can be deduced from their position in the matrix A .

Applications

The adjacency matrix of a graph can be used to answer certain questions about the graph's properties. For example, powers of the adjacency matrix A tell us information about the connectivity in the graph. The number of ways to get from vertex i to vertex j in two steps is $(A^2)_{ij}$, which is the ij^{th} entry of A^2 :

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

The number of ways to get from vertex i to vertex j in zero, one, or two steps is $\mathbb{1} + A + A^2$:

$$\mathbb{1} + A + A^2 = \begin{bmatrix} 1 & 1 & 2 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 \\ 2 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 \end{bmatrix}.$$

Observe that most of the graph G is well connected, except for Vertex 4, which has no inbound edges. The only way to get to Vertex 4 is if we start there.

The fact that we can discuss graph connectivity by doing matrix algebra is amazing when you think about it. The entry $(\mathbb{1} + A + A^2)_{41} = 2$ tells us there are two ways to get from Vertex 4 to Vertex 1 in two steps or less. Indeed, we can either transition directly through the edge $(4, 1)$ in one step, or indirectly via Node 5 in two steps by passing through the edges $(4, 5)$ and $(5, 1)$. Rather than manually counting all possible paths between vertices, we can compute all possible paths at once using matrix algebra on the adjacency matrix A .

Discussion

The analysis of connectivity between vertices is useful in many domains. Graphs are used to describe network flows, matching problems, social networks, webpages, and many other applications. In all these domains, the adjacency matrix plays a key role in graph representation. In Section 8.3, we'll study the graph of all webpages on the web, and we'll discuss Google's PageRank algorithm, which uses information in the adjacency matrix of the web to compute an "importance" rank for each webpage.

Links

[Graph theory and its applications]


https://en.wikipedia.org/wiki/Graph_theory

[More details about adjacency matrices with examples]

https://en.wikipedia.org/wiki/Adjacency_matrix

Exercises

E7.6 Find the adjacency matrix representation of the following graphs:



E7.7 For each of the graphs in E7.6, find the number of ways to get from vertex 1 to vertex 3 in two steps or less.

Hint: You can obtain the answer by computing the matrix $\mathbb{1} + A + A^2$.

7.5 Fibonacci sequence

We'll now look at a neat trick for computing the N^{th} term in the Fibonacci sequence $(0, 1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144, \dots)$. The terms in the Fibonacci sequence $(a_0, a_1, a_2, \dots)$ start with $a_0 = 0$, $a_1 = 1$, and then each subsequent term is computed as the sum of the two terms preceding it:

$$a_0 = 0, \quad a_1 = 1, \quad a_n = a_{n-1} + a_{n-2}, \text{ for all } n \geq 2.$$

The technical term for this type of formula is a *recurrence relation*. To compute the 1000^{th} term in the sequence using the formula, you'll need to perform 1000 steps of arithmetic. But do we really need N steps to compute a_N ? In this section we'll learn an eigenvalue trick that allows us to compute a_N in just five steps of symbolic computation, no matter how big N is.

First we express the recurrence relation as a matrix product:

$$\begin{aligned} a_{n+1} &= a_n + a_{n-1} \\ a_n &= a_n \end{aligned} \quad \Rightarrow \quad \underbrace{\begin{bmatrix} a_{n+1} \\ a_n \end{bmatrix}}_{\vec{a}_n} = \underbrace{\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}}_A \underbrace{\begin{bmatrix} a_n \\ a_{n-1} \end{bmatrix}}_{\vec{a}_{n-1}}.$$

We can compute the N^{th} term in the Fibonacci sequence by starting from the initial column vector $\vec{a}_0 = (a_1, a_0)^T$, and repeatedly multiplying by the matrix A :

$$\begin{bmatrix} a_{n+1} \\ a_n \end{bmatrix} = A^N \begin{bmatrix} a_1 \\ a_0 \end{bmatrix}.$$

We can “extract” a_N from the vector $\vec{a}_N$ by computing the dot product of $\vec{a}_N$ with the vector $(0, 1)$. This dot product operation has the effect of “selecting” the second component of the vector $\vec{a}_N$.

Thus, we obtain the following compact formula for computing the N^{th} term in the Fibonacci sequence in terms of the N^{th} power of the matrix A :

$$a_N = (0, 1)A^N(1, 0)^T.$$

Do you remember the eigendecomposition trick for computing powers of matrices by only computing powers of their eigenvalues? We can use the eigendecomposition trick to compute A^N very efficiently. The first step is to find the eigendecomposition of the matrix A :

$$\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 1 \\ \frac{1}{\varphi} & -\varphi \end{bmatrix}}_Q \underbrace{\begin{bmatrix} \varphi & 0 \\ 0 & \frac{-1}{\varphi} \end{bmatrix}}_\Lambda \underbrace{\begin{bmatrix} \frac{5+\sqrt{5}}{10} & \frac{\sqrt{5}}{5} \\ \frac{5-\sqrt{5}}{10} & -\frac{\sqrt{5}}{5} \end{bmatrix}}_{Q^{-1}}.$$

The two eigenvalues of A are $\lambda_1 = \varphi = \frac{1+\sqrt{5}}{2} \approx 1.618 \dots$ (the golden ratio), and $\lambda_2 = \frac{-1}{\varphi} = \frac{1-\sqrt{5}}{2} \approx -0.618 \dots$ (the negative inverse of the golden ratio). The columns of Q contain the corresponding eigenvectors of A .

We can compute A^N using the following formula:

$$A^N = \underbrace{AA \cdots A}_{N \text{ times}} = \underbrace{Q\Lambda Q^{-1} Q\Lambda Q^{-1} \cdots Q\Lambda Q^{-1}}_{N \text{ times}} = Q\Lambda^N Q^{-1}.$$

To compute A^N it is sufficient to compute the N^{th} powers of the eigenvalues $\lambda_1 = \varphi$ and $\lambda_2 = \frac{-1}{\varphi}$. For example, to compute a_5 , the fifth element in the Fibonacci sequence, we compute A^5 using $Q\Lambda^5 Q^{-1}$, then use the formula $a_5 = (0, 1)A^5(1, 0)^T$:

$$\begin{aligned} a_5 &= [0 \quad 1] \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}^5 \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\ &= [0 \quad 1] Q \begin{bmatrix} \varphi^5 & 0 \\ 0 & \frac{-1}{\varphi^5} \end{bmatrix} Q^{-1} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\ &= [0 \quad 1] \begin{bmatrix} 8 & 5 \\ 5 & 3 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = [0 \quad 1] \begin{bmatrix} 8 \\ 5 \end{bmatrix} = 5. \end{aligned}$$

We can similarly compute A^{55} :

$$A^{55} = Q\Lambda^{55}Q^{-1} = \begin{bmatrix} 225851433717 & 139583862445 \\ 139583862445 & 86267571272 \end{bmatrix},$$

then compute a_{55} using the formula,

$$a_{55} = (0, 1)A^{55}(1, 0)^T = 139583862445.$$

The eigendecomposition trick allows us to take a mathematical shortcut by obtaining the answer a_N using a constant number of math operations—regardless of the size of N :

$$a_N = (0, 1)Q\Lambda^N Q^{-1}(1, 0)^T.$$

First compute the N^{th} power of Λ . Then multiply Λ^N by $(0, 1)Q$ on the left, and by $Q^{-1}(1, 0)^T$ on the right. This is interesting since other algorithms for computing the Fibonacci numbers usually require a number of steps proportional to the size of N .

There are some caveats to the approach outlined above. We assumed computing φ^N and $(-1/\varphi)^N$ could be done in a single operation, which is not a realistic assumption on any computer. Computing the power of a number usually takes several multiplication steps, so we're kind of cheating when we count computing Λ^N as a single step. Also, infinite-precision (symbolic) manipulations are not realistic either since computers work with finite-precision approximations to real numbers, so the eigenvalue trick will not work for large values of N . Basically, the eigenvalue trick is only a theoretical result and can't be used in practical programs.

Links

[See the Wikipedia page for more on the Fibonacci numbers]
https://en.wikipedia.org/wiki/Fibonacci_number

Exercises

E7.8 Recall the Fibonacci sequence $a_n = (0, 1, 1, 2, 3, 5, 8, 13, 21, 34, \dots)$. Describe what happens to the ratio $\frac{a_{n+1}}{a_n}$ as $n \rightarrow \infty$.

Hint: Compare the size of $(\lambda_1)^n$ and $(\lambda_2)^n$ for large values of n .

E7.9 Compute the matrix products in the expression $(0, 1)Q\Lambda^N Q^{-1}(1, 0)^T$ to obtain a closed-form expression for a_N .

Hint: You can simplify expressions that contain square roots in the denominator as follows: $\frac{1}{a+b\sqrt{5}} = \frac{1}{a+b\sqrt{5}} \cdot 1 = \frac{1}{a+b\sqrt{5}} \frac{a-b\sqrt{5}}{a-b\sqrt{5}} = \frac{a-b\sqrt{5}}{a^2-5b^2}.$

7.6 Linear programming

In the early days of computing, computers were primarily used to solve optimization problems, so the term “programming” is often used to describe optimization problems. *Linear programming* is the study of linear optimization problems that involve linear constraints. These types of optimization problems play an important role in business: the whole point of corporations is to constantly optimize profits subject to time, energy, and legal constraints.

Many optimization problems can be expressed as *linear programs*. To solve an optimization problem is to find the optimal value, which is either the maximum or the minimum of some function, called the *objective function*. In a linear program, the objective function is a multivariable linear function $g(x_1, x_2, \dots, x_n)$ and the constraints on the variables of the problem are also linear. A linear program with n variables and m constraints is expressed as a maximization problem,

$$\max_{x_1, x_2, \dots, x_n} g(x_1, x_2, \dots, x_n) = c_1x_1 + c_2x_2 + \cdots + c_nx_n,$$

subject to m linear constraints,

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

The solution to this linear program is the vector $(x_1, x_2, \dots, x_n)$ that makes $g(x_1, x_2, \dots, x_n)$ as large as possible, and also satisfies all the constraints. For example, the variables $x_1, x_2, \dots, x_n$ could represent the production rates of n different products made by a company. If the coefficients $c_1, c_2, \dots, c_n$ represent the selling price for each of the n products, then $c_i x_i$ represents the revenue generated by selling x_i units of product i , and $g(x_1, x_2, \dots, x_n) = c_1x_1 + c_2x_2 + \cdots + c_nx_n$ represents the total revenue for a given choice of production rates. The m inequalities could represent various limitations of human resources, production capacity, or logistics constraints. To solve the linear program is to find the production rates $x_1, x_2, \dots, x_n$ that maximize revenue, subject to the constraints.

The *simplex algorithm* is a systematic procedure for finding solutions to linear programming problems. The simplex algorithm is somewhat similar to the Gauss–Jordan elimination procedure since it uses row operations on a matrix-like structure called a *tableau*. For this reason, linear programming and the simplex algorithm are often forced upon students taking a linear algebra course, especially busi-

ness students. I'm not going to lie to you and tell you the simplex algorithm is very exciting, but it is very powerful. For this reason, you should know it exists, and develop a general intuition about how it works. And, as with all things corporate-related, it's worth learning about so you'll know the techniques of the enemy.

Since the details of the simplex algorithm might not be of interest to all readers of the book, I split the topic of linear programming into a separate tutorial, which you can read online at the link below.

[Linear programming tutorial]

https://minireference.github.io/linear_programming/tutorial.pdf

7.7 Least squares approximate solutions

An equation of the form $A\vec{x} = \vec{b}$ could have exactly one solution (if A is invertible), infinitely many solutions (if A has a null space), or no solution at all (if $\vec{b}$ is not in the column space of A). In this section, we'll discuss the case with no solution, and describe an approach for computing an *approximate* solution $\vec{x}^*$ such that the vector $A\vec{x}^*$ is as close as possible to $\vec{b}$.

We could jump right away to the formula for the least squares approximate solution ($\vec{x}^* = (A^T A)^{-1} A^T \vec{b}$), but this would hardly be enlightening or useful for your understanding. Instead, let's learn about the least squares approximate solution in the context of a machine learning problem in which we'll try to predict some unknown quantities based on a linear model learned from past observations. This is called *linear regression* and it's one of the most useful applications of linear algebra.

Your company's database of current clients contains all the information about the frequency of purchases f , value of purchases V , promptness of payment P , and other useful information. You know what is *really* useful information, though? Knowing the customer lifetime value (CLV)—the total revenue this customer will generate during their entire relationship with your company. You have data on the CLVs of existing customers and you want to leverage this data to predict the CLVs of new customers.

Suppose you have data for N existing customers in the form of vectors, expressed as $\vec{a}_i = (f_i, V_i, P_i, \dots)$, and you calculate a customer lifetime value (CLV) for each existing customer, $b_i = CLV_i$.

The resulting dataset consists of observations $\vec{a}_i$ and outcomes b_i :

$$D = \left\{ \begin{bmatrix} - & \vec{a}_1 & - \\ - & \vec{a}_2 & - \\ - & \vec{a}_3 & - \\ & \vdots & \\ - & \vec{a}_N & - \end{bmatrix}, \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ \vdots \\ b_N \end{bmatrix} \right\} = \{A, \vec{b}\}.$$

The clients' observational data is stored in an $N \times n$ matrix A , and the corresponding CLVs are stored as an $N \times 1$ column vector $\vec{b}$.

Statement of the problem Given the $\vec{a}_k$ of a new customer, predict the customer's b_k , based on the information in the dataset D .

Linear model

A simple way to model the dependence of the label b_i on the observational data $\vec{a}_i = (a_{i1}, a_{i2}, \dots, a_{in})$ is to use a linear model with n parameters $m_1, m_2, \dots, m_n$:

$$y_{\vec{m}}(x_1, x_2, \dots, x_n) = m_1 x_1 + m_2 x_2 + \dots + m_n x_n = \vec{m} \cdot \vec{x}.$$

If your model is accurate, then $y_{\vec{m}}(\vec{a}_i)$ will closely approximate b_i . But how can we measure the accuracy of the approximation?

Let's define *error terms* that measure how the model's predictions $y_{\vec{m}}(\vec{a}_i)$ differ from the observed values b_i . The error term for the i^{th} customer is

$$e_i(\vec{m}) = |y_{\vec{m}}(\vec{a}_i) - b_i|^2.$$

The expression $e_i(\vec{m})$ measures the *squared error* between the model's prediction and the known value. Our goal is to choose a model that makes the sum of all the error terms as small as possible:

$$S(\vec{m}) = \sum_{i=1}^N e_i(\vec{m}) = \sum_{i=1}^N |y_{\vec{m}}(\vec{a}_i) - b_i|^2.$$

We'll refer to the function $S(\vec{m})$ as the *objective function*, which is the usual terminology for optimization problems. Our goal is to choose the value of $\vec{m}$ that minimizes $S(\vec{m})$, the *sum of squared errors* of the model's predictions.

Intuitively, it makes sense to minimize the objective function $S(\vec{m})$ because it captures how the predictions of the model $y_{\vec{m}}(\vec{a}_i)$ differ from the values b_i . The objective function will be zero ($S(\vec{m}) = 0$) if the model *perfectly* predicts the data. On the other hand, any model prediction $y_{\vec{m}}(\vec{a}_i)$ that overshoots or undershoots the correct

b_i value will lead to a large error term: $|y_{\vec{m}}(\vec{a}_i) - b_i|^2$. Observe that the objective function $S(\vec{m})$ is a quadratic function: the error terms grow quadratically with the size of the discrepancy between the model and the actual values.

Linear algebra formulation

When faced with an unfamiliar problem (such as finding the approximate solution to a system of equations $A\vec{x} = \vec{b}$ using a quadratic error model) don't become alarmed. Stand your ground and try relating the problem to something you're more familiar with. Thinking about problems in terms of linear algebra can often unlock your geometric intuition and show you a path toward the solution.

Using a matrix-vector product, we can express the "vector prediction" of the model $\vec{y}_{\vec{m}}$ for the whole dataset in one shot: $\vec{y}_{\vec{m}} = A\vec{m}$. The total squared error function for the model $\vec{m}$ on the dataset $D = \{A, \vec{b}\}$ can be written as the following expression:

$$\begin{aligned} S(\vec{m}) &= \sum_{i=1}^N |y_{\vec{m}}(\vec{a}_i) - b_i|^2 \\ &= \|\vec{y}_{\vec{m}} - \vec{b}\|^2 \\ &= \|A\vec{m} - \vec{b}\|^2. \end{aligned}$$

The total squared error for the model $\vec{m}$ on the dataset $\{A, \vec{b}\}$ is the squared length of the vector $A\vec{m} - \vec{b}$.

In the ideal case, if the model were to perfectly match the observations, the total squared error would be zero, and the equation $A\vec{m} = \vec{b}$ would have a solution:

$$A\vec{m} - \vec{b} = 0 \quad \Rightarrow \quad A\vec{m} = \vec{b}.$$

In practice, however, the model predictions $A\vec{m}$ will never perfectly match the data $\vec{b}$, so we must be content with the approximate solution $\vec{m}$:

$$A\vec{m} \approx \vec{b}.$$

The *least squares approximate solution* to this equation is chosen so as to minimize the total squared error function:

$$\min_{\vec{m}} S(\vec{m}) = \min_{\vec{m}} \|A\vec{m} - \vec{b}\|^2.$$

In other words, of all the possible approximate solutions $\vec{m}$, we must pick the one that makes the length of the vector $A\vec{m} - \vec{b}$ the smallest.

Finding the least squares approximate solution

There are two possible approaches for finding the least squares solution, denoted $\vec{m}^*$. We can either use calculus techniques to minimize the total squared error $S(\vec{m})$, or geometry techniques to find the shortest vector $(A\vec{m} - \vec{b})$.

Regardless of the approach chosen, the trick to finding the least squares approximate solution to $A\vec{m} \approx \vec{b}$ is to multiply the equation by $A^\top$ to obtain

$$A^\top A \vec{m} = A^\top \vec{b}.$$

The matrix $A^\top A$ will be invertible if the columns of A are linearly independent, which is the case for most tall-and-skinny matrices. We can therefore solve for $\vec{m}$ by multiplying both sides of the above equation by the matrix $(A^\top A)^{-1}$:

$$\vec{m} = (A^\top A)^{-1} A^\top \vec{b}.$$

Indeed, this expression is the *least squares solution* to the optimization problem we set out to solve in the beginning of this section:

$$\vec{m}^* = \underset{\vec{m}}{\operatorname{argmin}} S(\vec{m}) = \underset{\vec{m}}{\operatorname{argmin}} \|A\vec{m} - \vec{b}\|^2 = (A^\top A)^{-1} A^\top \vec{b}.$$

We can refer to the linear model $y_{\vec{m}^*}$ as the *best fit* for the dataset $\{A, \vec{b}\}$ since it minimizes the sum of the prediction errors.

Moore–Penrose inverse When we combine the particular combination of A , its transpose $A^\top$, and the inverse operation that we used to find the approximate solution $\vec{m}^*$, we get what's called the *Moore–Penrose inverse* or *pseudoinverse* of the matrix A . We use the shorthand notation A^+ (not to be confused with $A^\dagger$), for the entire expression:

$$A^+ \stackrel{\text{def}}{=} (A^\top A)^{-1} A^\top.$$

Applying A^+ to both sides of the approximate equation $A\vec{m} \approx \vec{b}$ is analogous to solving the equation by applying the inverse,

$$\begin{aligned} A^+ A \vec{m} &= A^+ \vec{b}, \\ \vec{m} &= A^+ \vec{b}. \end{aligned}$$

However, the solution is approximate, since $A\vec{m} \neq \vec{b}$. The solution $A^+ \vec{b}$ is optimal according to the sum of squared errors criterion $S(\vec{m})$.

Example 1 Given the dataset with six samples (x_i, b_i) : $(105, 45)$, $(113, 63)$, $(125, 86)$, $(137, 118)$, $(141, 112)$, $(153, 169)$, what is the best linear model $y_m(x) = mx$ for predicting b given x ? The first step is to express the given samples in the standard form $\{A, \vec{b}\}$, where $A \in \mathbb{R}^{6 \times 1}$ and $\vec{b} \in \mathbb{R}^{6 \times 1}$. We then calculate the best parameter m using the Moore–Penrose inverse formula:

$$\{A, \vec{b}\} = \left\{ \begin{bmatrix} 105 \\ 113 \\ 125 \\ 137 \\ 141 \\ 153 \end{bmatrix}, \begin{bmatrix} 45 \\ 63 \\ 86 \\ 118 \\ 112 \\ 169 \end{bmatrix} \right\} \Rightarrow m^* = (A^T A)^{-1} A^T \vec{b} = 0.792.$$

The best-fitting linear model for the data is $y(x) = 0.792x$. For the details of the calculation, see http://bit.ly/leastsq_ex1. Figure 7.4 shows a scatter plot of the dataset $\{(x_i, b_i)\}$ and the graph of the linear model that best fits the data.




Figure 7.4: The line of the form $y_m(x) = mx$ that best fits the set of data points $(x_i, b_i) \in \mathbb{R}^2$. The plot shows a scatter plot of the data points, and the line of best fit $y(x) = 0.792x$. The line passes through the middle of the dataset and minimizes the sum of the squared vertical differences between the model output $y_m(x_i)$ and the true value b_i .

Geometric interpretation

The solution to the least squares optimization problem,

$$\vec{m}^* = \operatorname{argmin}_{\vec{m}} \|A\vec{m} - \vec{b}\|^2,$$

can be understood geometrically as the search for the vector in the column space of A that is *closest* to the vector $\vec{b}$, as illustrated in Fig-

ure 7.5. As we vary the parameter vector $\vec{m}$, we obtain different vectors $A\vec{m} \in \mathcal{C}(A)$. Of all the points $A\vec{m}$ in the column space of A , the point $A\vec{m}^*$ is the closest to the point $\vec{b}$.

Let's define the *error vector* that corresponds to the difference between the model prediction $A\vec{m}$ and the actual value $\vec{b}$:

$$\vec{e} \stackrel{\text{def}}{=} A\vec{m} - \vec{b}.$$

Using the geometric intuition and looking at Figure 7.5, we see that the optimal solution $A\vec{m}^*$ occurs when the error vector is perpendicular to the column space of A . Recall the left fundamental subspaces of the matrix A : its column space $\mathcal{C}(A)$ and its orthogonal complement, the left null space $\mathcal{N}(A^\top)$. Thus, if we want an error vector that is perpendicular to $\mathcal{C}(A)$, we must find an error vector that lies in the left null space of A : $\vec{e}^* \in \mathcal{N}(A^\top)$. Using the definition of the left null space,

$$\mathcal{N}(A^\top) = \{\vec{w} \in \mathbb{R}^N \mid \vec{w}^\top A = \vec{0}^\top\},$$

we obtain the following equation that defines $\vec{m}^*$:

$$(\vec{e}^*)^\top A = 0 \quad \Rightarrow \quad (A\vec{m}^* - \vec{b})^\top A = 0.$$

Taking the transpose of the last equation, we obtain $A^\top(A\vec{m}^* - \vec{b}) = 0$, which is equivalent to condition $A^\top A\vec{m}^* = A^\top \vec{b}$ used to find $\vec{m}^*$.




Figure 7.5: Linear regression can be seen as a “search” problem restricted to the column space of the data matrix A . The least squares approximate solution $\vec{m}^* = A^+ \vec{b}$ corresponds to the point in $\mathcal{C}(A)$ that is closest to $\vec{b}$. For this $\vec{m}^*$, the error vector $\vec{e}^*$ is perpendicular to the column space of A .

Using geometric intuition about vector spaces and orthogonality proves helpful in solving this complex optimization problem. Choosing $\vec{e}^*$ orthogonal to $\mathcal{C}(A)$ leads to the shortest vector, $A\vec{m}^* - \vec{b}$, and produces the smallest total squared error, $S(\vec{m}^*) = \|A\vec{m}^* - \vec{b}\|^2$.

Affine models

We can also apply the Moore–Penrose inverse formula to more complicated models. A simple extension of a linear model is the *affine model* $y_{\vec{m},c}(\vec{x}) = \vec{m} \cdot \vec{x} + c$, which adds a constant term c to the output of a linear model. The model parameters correspond to an $(n+1)$ -vector $\vec{m}' = (c, m_1, m_2, \dots, m_n)$:

$$y_{\vec{m}'}(x_1, x_2, \dots, x_n) = m_0 + m_1 x_1 + m_2 x_2 + \dots + m_n x_n = \vec{m}' \cdot (1, \vec{x}).$$

For uniformity of notation, we'll refer to the constant term as m_0 instead of c . You can think of the parameter m_0 as the y -intercept of the model.

To accommodate this change in the model, we must preprocess the dataset $D = \{A, \vec{b}\}$ by adding a column of ones to the matrix A , turning it into $N \times (n+1)$ matrix A' . The new dataset looks like this:

$$D' = \left\{ \begin{bmatrix} 1 & - & \vec{a}_1 & - \\ 1 & - & \vec{a}_2 & - \\ 1 & - & \vec{a}_3 & - \\ \vdots & & \vdots & \\ 1 & - & \vec{a}_N & - \end{bmatrix}, \begin{bmatrix} b_0 \\ b_1 \\ b_2 \\ \vdots \\ b_N \end{bmatrix} \right\} = \{A', \vec{b}\}.$$

Except for the preprocessing step, which is commonly called “data-massaging,” the remainder of the steps for finding a least squares solution are the same as in the case of linear regression. We find the optimal parameter vector $\vec{m}'^*$ by applying the Moore–Penrose inverse formula:

$$\vec{m}'^* = (A'^T A')^{-1} A'^T \vec{b}.$$

Example 2 Find the affine model $y_{\vec{m}'}(x) = m_0 + m_1 x$ that best fits the data points from Example 1 (page 383). To find the best parameter vector $\vec{m}' = (m_0, m_1)^T$, we first preprocess the dataset by adding a column of ones to the data matrix. We then apply the inverse formula:

$$\{A', \vec{b}\} = \left\{ \begin{bmatrix} 1 & 105 \\ 1 & 113 \\ 1 & 125 \\ 1 & 137 \\ 1 & 141 \\ 1 & 153 \end{bmatrix}, \begin{bmatrix} 45 \\ 63 \\ 86 \\ 118 \\ 112 \\ 169 \end{bmatrix} \right\} \Rightarrow \vec{m}'^* = A'^+ \vec{b} = \begin{bmatrix} -210.4 \\ 2.397 \end{bmatrix}.$$

The affine model $y(x) = -210.4 + 2.397x$ is the best fit. I've omitted the details of the matrix calculation for brevity, but you can verify everything is legit here: http://bit.ly/leastsq_ex2. See Figure 7.6 for the graph of the best fitting affine model.




Figure 7.6: The affine model $(y_{\vec{m}'}(x) = m_0 + m_1x)$ that best fits the data points is the line $y(x) = -210.4 + 2.397x$. Allowing for one extra parameter, m_0 (the y -intercept), results in a model that fits the data much better, as compared to the fit in Figure 7.4.

Note the terms *linear regression* and *linear least squares* apply to fitting both linear models $\vec{y}_{\vec{m}} = \vec{m} \cdot \vec{x}$, and affine models $\vec{y}_{\vec{m}'} = \vec{m}' \cdot (1, \vec{x})$. After all, an affine model is just a linear model with a y -intercept, so it makes sense to refer to them by the same name.

Quadratic models

The mathematical approach based on the Moore–Penrose inverse we used in linear regression can also be used to find approximate solutions for nonlinear models. For the sake of simplicity, let's assume there are only two observed quantities in the dataset ($n = 2$). The most general quadratic model for two variables is:

$$y_{\vec{m}}(x, y) = m_0 + m_1x + m_2y + m_3xy + m_4x^2 + m_5y^2.$$

The terms with coefficients m_0 , m_1 , and m_2 describe an affine model for the variables x and y , and we add additional terms with parameters m_3 , m_4 , and m_5 to describe quadratic components in x and y . The parameter vector for the most general quadratic model of two variables is six-dimensional: $\vec{m} = (m_0, m_1, m_2, m_3, m_4, m_5)$.

Assume we start from the following dataset:

$$D = \left\{ \begin{bmatrix} x_1 & y_1 \\ x_2 & y_2 \\ \vdots & \vdots \\ x_N & y_N \end{bmatrix}, \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_N \end{bmatrix} \right\} = \{A, \vec{b}\}.$$

To use the Moore–Penrose inverse formula, we must preprocess A by adding several new columns:

$$D' = \left\{ \begin{bmatrix} 1 & x_1 & y_1 & x_1 y_1 & x_1^2 & y_1^2 \\ 1 & x_2 & y_2 & x_2 y_2 & x_2^2 & y_2^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & x_N & y_N & x_N y_N & x_N^2 & y_N^2 \end{bmatrix}, \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_N \end{bmatrix} \right\} = \{A', \vec{b}\}.$$

This preprocessing step allows us to compute a vector prediction of the model with parameters $\vec{m}$ applied to the entire dataset: $\vec{y}_{\vec{m}} = A' \vec{m}$. The total squared error for this model is $S(\vec{m}) = \|\vec{y}_{\vec{m}} - \vec{b}\|^2$. The least squares approximate solution $\vec{m}^*$ is obtained as usual: $\vec{m}^* = (A'^T A')^{-1} A'^T \vec{b}$.

Note the number of parameters in a quadratic model grows—surprise, surprise—quadratically. There were $n = 2$ variables in the above example, and the parameters vector $\vec{m}$ is six-dimensional. The number of parameters of a general quadratic model in n variables is $\frac{1}{2}(n+1)(n+2) = \frac{1}{2}(n^2 + 3n + 2)$. The number of model parameters is an important consideration to take into account when fitting models to large datasets.

Example 3 Imagine you're a data-friendly business person and you want to pick which model to fit to a dataset $\{A, b\}$ with $n = 1000$ features. The observational data matrix A has N rows and 1000 columns. You have the option of choosing either an affine model, $y_a(\vec{x}) = m_0 + m_1 x_1 + \cdots + m_n x_n$, or a quadratic model, $y_q(\vec{x}) = m_0 + m_1 x_1 + \cdots + m_n x_n + m_{n+1} x_1 x_1 + m_{n+2} x_1 x_2 + \cdots$, for the dataset.

To fit an affine model, you'll need to preprocess $A \in \mathbb{R}^{N \times 1000}$ into $A' \in \mathbb{R}^{N \times 1001}$, and then find the inverse of $(A'^T A') \in \mathbb{R}^{1001 \times 1001}$. That's totally doable.

In contrast, the number of parameters in a quadratic model with $n = 1000$ dimensions is $\frac{1}{2}(1000+1)(1000+2) = 501501$. To fit a quadratic model, you'll need to preprocess the data matrix to obtain $A' \in \mathbb{R}^{N \times 501501}$, and then find the inverse of $(A'^T A') \in \mathbb{R}^{501501 \times 501501}$. That's a big matrix. You'll need one terabyte of memory just to store all the entries of the matrix $A'^T A'$, and computing its inverse will take a *really* long time.

So the choice is made; affine model it is. The affine model may be less accurate than a quadratic model, but if it works and provides value, you can use it. This is the reason why scientists, engineers, statisticians, and business folk are so crazy about building linear models, even though more advanced models are available. Linear models are a *sweet spot* for prediction applications: they have great

modelling power, they're easy to implement, and they lead to computational problems that are easy to solve.

Links

[Further discussion about least squares problems on Wikipedia]

https://en.wikipedia.org/wiki/Linear_regression

[https://en.wikipedia.org/wiki/Linear_least_squares_\(mathematics\)](https://en.wikipedia.org/wiki/Linear_least_squares_(mathematics))

[More about the Moore–Penrose inverse]

https://en.wikipedia.org/wiki/Moore-Penrose_inverse

Exercises

E7.10 You want to determine whether a coin is fair, so you toss it repeatedly and record the number of times it lands heads. On the first trial, you flip the coin eight times and obtain four heads, which is a heads-to-flips ratio of $\frac{4}{8}$. On subsequent trials you obtain heads-to-flips ratios of $\frac{9}{16}$, $\frac{13}{24}$, $\frac{17}{32}$, and $\frac{20}{40}$. Find the best-fitting linear model $h(x) = mx$ to describe the number of heads in a trial with x flips.

E7.11 Find the best-fitting affine model $y = b + mx$ to the (x, y) data points $(0, 3.9)$, $(1, 3.2)$, and $(2, 1.9)$. Perform all the calculations by hand.

Hint: Find the Moore–Penrose inverse.

E7.12 Calculate the total squared error $S(m^*) = \|Am^* - \vec{b}\|^2$ of the best-fit linear model obtained in Example 1 (page 383). Use the SymPy calculation at `bit.ly/leastsq_ex1` as your starting point.

Hint: The Matrix method `.norm()` might come in handy.

E7.13 Revisit Example 2 (page 385) and find the total squared error of the best-fit affine model $S(\vec{m}^*) = \|A\vec{m}^* - \vec{b}\|^2$. You can start from the calculation provided here `bit.ly/leastsq_ex2` and extend it.

7.8 Computer graphics

Linear algebra is the mathematical language of computer graphics. Whether you're building a simple two-dimensional game with stick figures, or a fancy three-dimensional visualization, knowing linear algebra will help you understand the graphics operations that draw pixels on the screen.

In this section, we'll discuss some basic computer graphics concepts. In particular, we'll introduce *homogeneous coordinates*, a

representation for vectors and matrices that uses an extra dimension. Instead of representing a three-dimensional vector as a triple of Cartesian coordinates $(x, y, z)_c$, we'll use a four-dimensional homogeneous coordinates vector $(x, y, z, 1)_h$. Homogeneous coordinates allow us to represent any computer graphics transformation as a matrix-vector product. We've already seen that scalings, rotations, reflections, and orthogonal projections can be represented as matrix-vector products. With homogeneous coordinates, we can represent translations and perspective projections as matrix products, too. That's very convenient, since it enables us to understand all computer graphics operations in terms of matrix multiplications.

Computer graphics is a vast subject, far more extensive than the pages allotted in this book. To keep things simple, we'll focus on two-dimensional graphics, and only briefly touch upon three-dimensional graphics. The goal is not to teach you the commands of computer graphics APIs like OpenGL and WebGL, but to give you the basic math tools for understanding what's happening under the hood.

Affine transformations

In Chapter 5 we studied various linear transformations and their representations as matrices. We also briefly discussed the class of *affine transformations*, which consist of a linear transformation followed by a translation:

$$\vec{w} = T(\vec{v}) + \vec{d}.$$

In the above equation, the input vector $\vec{v}$ is first acted upon by a linear transformation T , and then the output of T is translated by the displacement vector $\vec{d}$ to produce the output vector $\vec{w}$.

In this section, we'll use *homogeneous coordinates* for vectors and transformations, which allow us to express affine transformations as matrix-vector products in a larger vector space:

$$\vec{w} = T(\vec{v}) + \vec{d} \quad \Leftrightarrow \quad \vec{W} = A\vec{V}.$$

If $\vec{v}$ is an n -dimensional vector, then its representation in homogeneous coordinates $\vec{V}$ is an $(n + 1)$ -dimensional vector. The $(n + 1) \times (n + 1)$ matrix A contains the combined information about both the linear transformation T and the translation $\vec{d}$.

Homogeneous coordinates

Instead of using a triple of Cartesian coordinates to represent points $p = (x, y, z)_c \in \mathbb{R}^3$, we'll use the quadruple $P = (x, y, z, 1)_h \in \mathbb{R}^4$,

which is a representation of the same point in *homogeneous coordinates*. Similarly, the vector $\vec{v} = (v_x, v_y, v_z)_c \in \mathbb{R}^3$ in Cartesian coordinates corresponds to the vector $\vec{V} = (v_x, v_y, v_z, 1)_h \in \mathbb{R}^4$ in homogeneous coordinates. Though there is no mathematical difference between points and vectors, we'll stick to the language of points as it is more natural for graphics problems.

An interesting property of homogeneous coordinates is that they're not unique: the vector $\vec{v} = (v_x, v_y, v_z)_c$ corresponds to a set of points in homogeneous coordinates: $\vec{V} = \{(\alpha v_x, \alpha v_y, \alpha v_z, \alpha)_h\}$, for $\alpha \in \mathbb{R}$. This makes homogeneous coordinates invariant to scaling:

$$\begin{bmatrix} a \\ b \\ c \end{bmatrix}_c \quad \Leftrightarrow \quad \begin{bmatrix} a \\ b \\ c \\ 1 \end{bmatrix}_h = \begin{bmatrix} 5a \\ 5b \\ 5c \\ 5 \end{bmatrix}_h = \begin{bmatrix} 500a \\ 500b \\ 500c \\ 500 \end{bmatrix}_h.$$

This is kind of weird, but this extra freedom to rescale vectors arbitrarily leads to many useful applications.

To convert from homogeneous coordinates $(X, Y, Z, W)_h = (a, b, c, d)_h$ to Cartesian coordinates, we divide each component by the W -component to obtain the equivalent vector $(X, Y, Z, W)_h = (\frac{a}{d}, \frac{b}{d}, \frac{c}{d}, 1)_h$, which corresponds to the point $(x, y, z)_c = (\frac{a}{d}, \frac{b}{d}, \frac{c}{d})_c \in \mathbb{R}^3$.

When the underlying Cartesian space is two-dimensional, the point $p = (x, y)_c \in \mathbb{R}^2$ is written as $P = (X, Y, W)_h = (x, y, 1)_h$ in homogeneous coordinates. The homogeneous coordinates $(X, Y, W)_h = (a, b, d)_h$ (where $d \neq 0$) represent the point $(x, y)_c = (\frac{a}{d}, \frac{b}{d})_c \in \mathbb{R}^2$.

This conversion between homogeneous coordinates and Cartesian coordinates can also be understood geometrically, as illustrated in Figure 7.7. The point $(X, Y, W)_h$ in homogeneous coordinates corresponds to an infinite line in the three-dimensional XYW -coordinate space. To obtain the Cartesian coordinates for this point, find the intersection of the infinite line with the plane $W = 1$.

To distinguish Cartesian vectors from homogeneous vectors, we'll use capital letters like $P, \vec{A}, \vec{B}, \dots$ for points and vectors in homogeneous coordinates, and lowercase letters like $p, \vec{a}, \vec{b}, \dots$ when referring to points and vectors in Cartesian coordinates.

Affine transformations in homogeneous coordinates

Consider the affine transformation that consists of the transformation $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ followed by a translation by $\vec{d} = (d_1, d_2)$. If T is




Figure 7.7: A two-dimensional Cartesian point $(x, y)_c$ corresponds to an infinite line in a three-dimensional homogeneous coordinates space $\{(\alpha x, \alpha y, \alpha)_h, \alpha \in \mathbb{R}\}$. To convert a point from homogeneous coordinates to Cartesian coordinates, we must find where the infinite line intersects the plane $W = 1$. The homogeneous coordinates of the intersection point will have the form $(x, y, 1)_h$. The first two components of this point are the Cartesian coordinates $(x, y)_c$.

represented by the matrix $M_T = \begin{bmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{bmatrix}$, then the affine transformation as a whole can be represented as follows:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} \quad \Leftrightarrow \quad \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} m_{11} & m_{12} & d_1 \\ m_{21} & m_{22} & d_2 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}.$$

In the homogeneous coordinates representation, the input has an additional component that contains the constant value one, and the matrix has an additional column that operates on this constant component. The result of the matrix-vector product is to add the constants d_1 and d_2 to the first and second coordinates:

$$x' = m_{11}x + m_{12}y + d_1, \quad y' = m_{21}x + m_{22}y + d_2,$$

which is exactly what translation by $\vec{d}$ is all about. As you can see, there's nothing fancy about homogeneous coordinates. Make sure you understand the above matrix-vector products and that you're convinced there is no new math here—just the good old matrix-vector product you're familiar with.

Graphics transformations in 2D

When points and vectors are represented in homogeneous coordinates, we can express most useful geometric transformations of the form $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ as 3×3 matrices in homogeneous coordinates. In addition to affine transformations we discussed above, homogeneous coordinates also allow us to perform *perspective transformations*, which are of central importance in computer graphics. The

most general transformation we can perform using homogeneous coordinates is

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} m_{11} & m_{12} & d_1 \\ m_{21} & m_{22} & d_2 \\ p_1 & p_2 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix},$$

where $M = \begin{bmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{bmatrix}$ corresponds to any linear transformation, $\vec{d} = (d_1, d_2)$ corresponds to a translation, and (p_1, p_2) is a perspective transformation.

Linear transformations

Let R_θ be the clockwise rotation by the angle θ of all points in $\mathbb{R}^2$. In homogeneous coordinates, this rotation is represented as

$$p' = R_\theta(p) \quad \Leftrightarrow \quad P' = M_{R_\theta} P,$$

where M_{R_θ} is the following 3×3 matrix:

$$\begin{bmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

The 2×2 submatrix in the top-left corner of matrices in homogeneous coordinates can represent any linear transformation: projections, reflections, scalings, and shear transformations. The following equation shows the homogeneous-coordinates matrix representations of three linear transformations: the reflection through the x -axis M_{R_x} , an arbitrary scaling M_S , and a shear along the x -axis M_{SH_x} .

$$M_{R_x} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad M_S = \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad M_{SH_x} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Orthogonal projections

Projections can also be represented as 3×3 matrices. The projection onto the x -axis corresponds to the following representation in homogeneous coordinates:

$$p' = \Pi_x(p) \quad \Leftrightarrow \quad \begin{bmatrix} x' \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}.$$

Translation

The translation by the displacement vector $\vec{d} = (d_x, d_y)$ corresponds to the matrix

$$M_{T_{\vec{d}}} = \begin{bmatrix} 1 & 0 & d_x \\ 0 & 1 & d_y \\ 0 & 0 & 1 \end{bmatrix}.$$

Note the top-left submatrix of $M_{T_{\vec{d}}}$ contains the identity matrix; we perform the identity linear transformation and a translation by $\vec{d}$.




Figure 7.8: Illustration of the different transformations on a sample shape. Source: wikipedia File:2D_affine_transformation_matrix.svg

Rotations, reflections, shear transformations, and translations can all be represented as multiplications by 3×3 matrices. Figure 7.8 shows examples of transformations we can perform using the matrix-vector product in homogeneous coordinates. I hope by

now you're convinced that this idea of adding an extra "constant" dimension to vectors is useful. But wait, there's more! We haven't yet seen what happens when we put coefficients in the last row of the matrix.

Perspective projections

The notion of perspective is very important in the world of visual art. For a painting to look realistic, the relative size of the objects pictured must convey their relative distance from the viewer. Distant objects in the scene are painted smaller than objects in the foreground. To transform a three-dimensional scene into a realistic two-dimensional painting, we need to apply a *perspective transformation*. Rather than give you the general formula for perspective transformations, we'll derive the matrix representation for perspective transformations from scratch. Trust me, it will be a lot more interesting.

We can understand perspective transformations by tracing imaginary light rays that originate from each corner of an object and travel to the eye of an observer O in a straight line (see Figure 7.9). The image of the perspective projections is created along the path of the light ray, where the ray intersects the *projective plane*, which represents the screen where the image will form. Since we're working in $\mathbb{R}^2$, we can draw a picture of what's going on.




Figure 7.9: The perspective projection of a square onto a screen. Note how the side of the square that is closer to the observer appears longer than the farther side.

Let's assume the observer is placed at the origin, $O = (0,0)$, and let's compute the *perspective transformation* that projects points onto the line with equation $y = d$. Under this perspective projection, every point $p = (x, y)$ in the plane maps to a point $p' = (x', y')$ on the line $y = d$. Figure 7.10 illustrates the situation.

The only math prerequisite you need to remember is the general principle for *similar* triangles: if a triangle with sides a, b, c is similar to a triangle with sides a', b', c' , then the ratios of the lengths of the




Figure 7.10: The point $p' = (x', y')$ is the projection of the point $p = (x, y)$ onto the line with equation $y = d$. The ratio of lengths d/y must equal the ratio of lengths x'/x .

triangles' sides will be equal:

$$\frac{a'}{a} = \frac{b'}{b} = \frac{c'}{c}.$$

Since the two triangles in Figure 7.10 are similar, we know $\frac{x'}{x} = \frac{d}{y}$, and therefore we directly obtain the expression for (x', y') as follows:

$$x' = \frac{d}{y}x,$$

$$y' = \frac{d}{y}y = d.$$

This doesn't look very promising, however, since the expression for x' contains a division by y . The set of equations is not linear in y , and therefore cannot be expressed as a matrix-product. If only there were some way to represent vectors and transformations that also allowed division by coefficients too.

Let's analyze the *perspective transformation* in terms of homogeneous coordinates, and see if we find anything useful:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} \frac{d}{y}x \\ d \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ \frac{y}{d} \end{bmatrix}.$$

The second equality holds because vectors in homogeneous coordinates are invariant to scalings: $(a, b, c)_h = \alpha(a, b, c)_h$ for all α . We can shift the factor $\frac{d}{y}$ as we please: $(\frac{d}{y}x, d, 1) = \frac{d}{y}(x, y, \frac{y}{d}) = (x, y, \frac{y}{d})$. In the alternate homogeneous coordinates expression, we're no longer dividing by y . This means we can represent the perspective transformation as a matrix-vector product:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ \frac{y}{d} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & \frac{1}{d} & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}.$$

Now that's interesting. By preparing the vector (X', Y', W') with a third component $W' \neq 1$, we can force each coefficient to be scaled by $\frac{1}{W'}$, which is exactly what we need for perspective transformations. Depending on how the coefficient W' is constructed, different perspective transformations can be obtained.

A *perspective projection* transformation is a perspective transformation followed by an orthogonal projection that removes some of the vector's components. The perspective projection onto the line with equation $y = d$ is the composition of a *perspective transformation* $P : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ followed by an orthogonal projection $\Pi_x : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ which simply discards the y' coordinate:

$$\begin{bmatrix} x' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & \frac{1}{d} & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \Rightarrow \begin{bmatrix} x' \\ 1 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{d} & 0 \end{bmatrix}}_{\Pi_x P} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}.$$

Note that only the x' coordinate remains after the projection. This is the desired result since we want only the “local” coordinates of the projection plane $y = d$ to remain after the projection.

Certain textbooks on computer graphics discuss only the combined perspective-plus-projection transformation $\Pi_x P$, as described on the right side of the above equation. I prefer to treat the perspective transformation separately from the projection, since it makes the math easier to understand. Also, by including the “depth information” (the y -coordinates) of the objects we're projecting we can determine which objects appear in front of others.

General perspective transformation

Let's now look at the general case of a perspective transformation that projects arbitrary points $p = (x, y)$ onto the line $ax + by = d$. Again, we assume the observer is located at the origin $O = (0, 0)$. We want to calculate the coordinates of the projected point $p' = (x', y')$, as illustrated in Figure 7.11.

To obtain the general perspective transformation, we'll follow a logical reasoning similar to the special case we considered above based on the properties of similar triangles. Define α to be the projection of the point p onto the line with direction vector $\vec{n} = (a, b)$ passing through the origin. Using the general formula for distances (see Section 4.1), we can obtain the length ℓ from O to α :

$$\ell = d(O, \alpha) = \frac{\vec{n} \cdot p}{\|\vec{n}\|} = \frac{ax + by}{\|\vec{n}\|}.$$




Figure 7.11: The point $p' = (x', y')$ is the projection of the point $p = (x, y)$ onto the line with equation $ax + by = d$. We define points α' and α in the direction of line's normal vector $\vec{n} = (a, b)$. The distances from the origin to these points are ℓ' and ℓ respectively. We have $\ell'/\ell = x'/x = y'/y$.

Similarly, we define ℓ' to be the length of the projection of p' onto $\vec{n}$:

$$\ell' = d(O, \alpha') = \frac{\vec{n} \cdot p'}{\|\vec{n}\|} = \frac{ax' + by'}{\|\vec{n}\|} = \frac{d}{\|\vec{n}\|}.$$

The last equation holds true because the point p' is located on the line with equation $ax + by = d$.

By the similarity of triangles, we know the ratio of lengths x'/x and y'/y must equal the ratio of orthogonal distances ℓ'/ℓ :

$$\frac{x'}{x} = \frac{y'}{y} = \frac{\ell'}{\ell} = \frac{d}{ax + by}.$$

We can use this fact to express the coordinates x' and y' in terms of the original x and y coordinates. As in the previous case, expressing points in homogeneous coordinates allows us to arbitrarily shift scale factors:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} \frac{\ell'}{\ell} x \\ \frac{\ell'}{\ell} y \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ \frac{\ell}{\ell'} \end{bmatrix} = \begin{bmatrix} x \\ y \\ \frac{ax+by}{d} \end{bmatrix}.$$

The last expression is linear in the variables x and y , therefore it has a matrix representation:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} X' \\ Y' \\ W' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{a}{d} & \frac{b}{d} & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}.$$

This is the most general example of a perspective transformation. The "scale factor" W' is a linear combination of the input coordi-

nates: $W' = \frac{a}{d}x + \frac{b}{d}y$. Observe that by setting $a = 0$ and $b = 1$, we recover the perspective transformation for the projection onto the line $y = d$.

Graphics transformations in 3D

Everything we saw in the previous section about two-dimensional transformations also applies to three-dimensional transformations. A three-dimensional Cartesian coordinate triple $(x, y, z)_c \in \mathbb{R}^3$ is represented as $(x, y, z, 1)_h \in \mathbb{R}^4$ in homogeneous coordinates. Transformations in homogeneous coordinates are represented by 4×4 matrices. An affine transformation is represented by the matrix

$$A = \begin{bmatrix} m_{11} & m_{12} & m_{13} & d_1 \\ m_{21} & m_{22} & m_{23} & d_2 \\ m_{31} & m_{32} & m_{33} & d_3 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

and the perspective transformation onto the line $ax + by + cz = d$ with the observer at the origin is expressed as the matrix

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ \frac{a}{d} & \frac{b}{d} & \frac{c}{d} & 0 \end{bmatrix}.$$

These should look somewhat familiar to you. Affine transformations, projections, and perspective transformations are the same in 3D as in 2D, except we're now working in a four-dimensional homogeneous coordinates space.

The best part about being able to apply the homogeneous coordinates representation to all transformations is that we can *compose* transformations together, in a way that's similar to building with LEGOs. We'll learn about this in the next section.

3D graphics programming

Inside every modern computer is a special-purpose processor, called the *graphics processing unit* (GPU), which is dedicated to computer graphics operations. Modern GPUs can have thousands of individual graphics processing units called *shaders*, and each shader can perform millions of linear algebra operations per second. Think about it; thousands of processors working in parallel to calculate matrix-vector products for you—that's a lot of linear algebra calculating power!

The reason we need so much processing power is because 3D models are made of thousands of little polygons. Drawing a 3D scene (also known as *rendering*) involves performing linear algebra manipulations on all these polygons. This is where the GPU comes in. The job of the GPU is to translate, rotate, and scale the polygons of the 3D models by placing them into the scene, and then computing what the scene looks like when projected to the two-dimensional window (the screen) through which you're observing the virtual world. This transformation—from the model coordinates to world coordinates, and then to screen coordinates (pixels)—is carried out in a *graphics processing pipeline*.




Figure 7.12: A graphics processing pipeline for drawing 3D objects on the screen. A 3D model is composed of polygons expressed with respect to a coordinate system centred on the object. The model matrix positions the object in the scene, the view matrix positions the camera in the scene, and finally the projection matrix computes what should appear on the screen.

We can understand the graphics processing pipeline as a sequence of matrix transformations: the model matrix M , the view matrix V , and the projection matrix Π_s . The GPU applies this sequence of operations to each of the object's vertices, $(x, y, z, 1)_o$, to obtain the pixel coordinates, $(x', y')_s$, of the vertices on the screen:

$$\begin{bmatrix} x' \\ y' \end{bmatrix}_s = \Pi_s V M \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}_m \quad \Rightarrow \quad (x, y, z, 1)_m M^T V^T \Pi_s^T = (x', y')_s.$$

In the context of computer graphics, it is customary to represent the graphics processing pipeline in the “transpose picture,” so that vertex data flows from left to right as in Figure 7.12. Instead of representing vertices as column vectors multiplied by matrices M , V , and Π_s from the left, we represent vertices as row vectors multiplied by matrices M^T , V^T , and Π_s^T from the right. All the reasoning remains the same, and all the transformation matrices described above still work; you just might need to transpose them if you're using them in a program.

Finally, a comment on efficiency. It is not necessary to compute the matrix-vector products with M^T , V^T , and Π_s^T for each vertex. It's much more efficient to pre-compute a combined transformation matrix, $C^T = M^T V^T \Pi_s^T$, and apply C^T to each of the coordinates of

the 3D objects. Similarly, when adding another object to the scene, only the model matrix needs to be modified, while the view and projection matrices remain the same.

Practical considerations

We discussed homogeneous coordinates and the linear algebra transformations used for computer graphics. This is the essential theory you'll need in order to get started with computer graphics programming. The graphics pipelines used in modern 3D software involve a few more steps than the simplified version shown in Figure 7.12. We'll now discuss some points and practical considerations for programming computer graphics on real GPUs:

- There are actually *two* graphics pipelines at work in the GPU. The *geometry pipeline* handles the transformation of polygons that form the 3D objects. A separate *texture pipeline* controls the graphics patterns that fill the polygons. The final step in the rendering process combines the outputs of the two pipelines.
- The model and view matrices can be combined to form a single model-view matrix that converts object coordinates to camera coordinates.
- Not all objects in the scene need to be rendered. We don't need to render objects that fall outside the camera's viewing angle. We can also skip objects that are closer than the *near plane* or farther than the *far plane* of the camera. Though the scene could extend infinitely, we're only interested in rendering the subset of the scene that we want displayed on the screen, which we call the *view frustum*. See the illustration in Figure 7.13.




Figure 7.13: The perspective transformation that projects a three-dimensional model onto a two-dimensional screen surface. Note only a subset of the scene (the *view frustum*) is drawn, which lies between the *near plane* and the *far plane* within the field of view.

I encourage you to pursue the subject of computer graphics on your own. There are excellent free resources for learning OpenGL and WebGL on the web.

Discussion

Homogeneous coordinates and projective geometry are powerful mathematical techniques with deep connections to many advanced math subjects. For the purpose of this section, we explored the subject from an engineering perspective—as a handy trick for working with computer graphics using matrix-vector products. Homogeneous coordinates have many other applications that we did not have time to discuss. We'll briefly mention some of these additional applications here, with the hope of inspiring you to research the subject further.

Homogeneous coordinates are convenient for representing planes in $\mathbb{R}^3$. The plane with general equation $ax + by + cz = d$ is represented by the vector $\vec{N} = (a, b, c, -d)$ in homogeneous coordinates. This is a natural way to express planes. Instead of describing the plane's normal vector $\vec{n}$ separately from the constant d , we can represent the plane by an “enhanced” normal vector $\vec{N} \in \mathbb{R}^4$:

$$\text{Plane with } \vec{n} = (a, b, c) \text{ and constant } d \quad \Leftrightarrow \quad \vec{N} = (a, b, c, -d).$$

This is really cool because we can now use the same representation for both vectors and planes, and perform vector operations between them. Yet again, we find a continuation of the “everything is a vector” theme we've encountered throughout this book.

What good is representing planes in homogeneous coordinates? For starters, it allows us to verify whether any point $\vec{P}$ lies in the plane $\vec{N}$ by computing the dot product $\vec{N} \cdot \vec{P}$. The point $\vec{P}$ lies inside the plane with normal vector $\vec{N}$ if and only if $\vec{N} \cdot \vec{P} = 0$. Consider the point $\vec{P}_1 = (0, 0, \frac{d}{c}, 1)$ that lies in the plane $ax + by + cz = d$; we can verify that

$$\vec{N} \cdot \vec{P}_1 = (a, b, c, -d) \cdot (0, 0, \frac{d}{c}, 1) = 0.$$

It's also possible to easily obtain the homogeneous coordinates of the plane $\vec{N}$ that passes through any three points, $\vec{P}$, $\vec{Q}$, and $\vec{R}$. We're looking for a vector $\vec{N}$ that is perpendicular to all three points. We can obtain $\vec{N}$ using a generalization of the cross product, computed using a four-dimensional determinant:

$$\vec{N} = \begin{vmatrix} \hat{e}_1 & \hat{e}_2 & \hat{e}_3 & \hat{e}_4 \\ p_x & p_y & p_z & p_w \\ q_x & q_y & q_z & q_w \\ r_x & r_y & r_z & r_w \end{vmatrix},$$

where $\{\hat{e}_1, \hat{e}_2, \hat{e}_3, \hat{e}_4\}$ is the standard basis for $\mathbb{R}^4$. I bet you haven't seen a four-dimensional cross product before—this stuff is wild! See

the links below if you want to learn more about homogeneous coordinates, projective spaces, or computer graphics.

Links

[A detailed tutorial series on WebGL]

<https://github.com/greggman/webgl-fundamentals/>

[Visualization of polygon drawing in a 3D scene]

<http://orbides.org/apps/superslow.html>

7.9 Cryptography

Cryptography is the study of secure communication. The two main tasks that cryptographers aim to achieve are private communication (no eavesdroppers) and authenticated communication (no impersonators). Using algebraic operations over finite fields $\mathbb{F}_q$, it's possible to achieve both of these goals. Math is the weapon for privacy!

The need for private communication between people has existed long before the development of modern mathematics. Thanks to modern mathematical techniques, we can now perform cryptographic operations with greater ease, and build cryptosystems with security guaranteed by mathematical proofs. In this section, we'll discuss the famous *one-time pad* encryption technique invented by Gilbert Vernam. One-time pad encryption is important because Claude Shannon proved it is *absolutely secure*. In order to understand what that means precisely, we'll first need some context about the concepts studied in the field of cryptography.

Context

The secure communication scenarios we'll discuss in this section involve three parties:

- Alice is the message sender
- Bob is the message receiver
- Eve is the eavesdropper

Alice wants to send a private message to Bob, but Eve has the ability to see all communication between Alice and Bob. You can think of Eve as a Facebook administrator, or an employee of the Orwellian, privacy-invading web application *du jour*. To defend against Eve, Alice will *encrypt* her messages before sending them to Bob, using a secret key only Alice and Bob have access to. Eve will be able to capture the encrypted messages (called *ciphertexts*) but they will be

unintelligible to her because of the encryption. Assuming Bob receives the messages from Alice, he'll be able to *decrypt* them using his copy of the secret key. Encryption allows Alice and Bob to have *private* communication despite Eve's eavesdropping.

Cryptography is an interesting subject that is very pertinent in the modern age of pervasive network surveillance. Learning cryptography is a bit like learning kung fu for self defence. You don't need to go around beating people up, but you should know how to defend yourself in case something happens. I encourage you to learn more about practical cryptography for your communications. Knowing some basics about passwords, secret keys, digital signatures, and certificates will help you keep pace with the forces out there.

The material in this section only scratches the surface of all there is to learn about cryptography; it serves to illustrate the main concepts through a simple example: the *one-time pad* encryption protocol. This simple encryption scheme is provably unconditionally secure, given some assumptions about the secret key $\vec{k}$.

Definitions

A cryptographic protocol consists of an encryption function, a decryption function, and a procedure for generating the secret key $\vec{k}$. For simplicity we assume messages and keys are all binary strings:

- $\vec{m} \in \{0,1\}^n$: the *message* or *plaintext* is a bitstring of length n
- $\vec{k} \in \{0,1\}^n$: the *key* is a shared secret between Alice and Bob
- $\vec{c} \in \{0,1\}^n$: the *ciphertext* is the encrypted message
- $\text{Enc}(\vec{m}, \vec{k})$: the encryption function that takes as input a message $\vec{m} \in \{0,1\}^n$ and the key $\vec{k} \in \{0,1\}^n$ and produces a ciphertext $\vec{c} \in \{0,1\}^n$ as output
- $\text{Dec}(\vec{c}, \vec{k})$: the decryption function that takes as input a ciphertext $\vec{c} \in \{0,1\}^n$ and the key $\vec{k} \in \{0,1\}^n$ and produces the decrypted message $\vec{m} \in \{0,1\}^n$ as output

We consider the protocol to be secure if Eve cannot gain any information about the messages $\vec{m}_1, \vec{m}_2, \dots$ from the ciphertexts $\vec{c}_1, \vec{c}_2, \dots$ she intercepts.

Before we describe the one-time pad encryption protocol and discuss its security, we must introduce some background about binary numbers.

Binary

The cryptographic operations we'll discuss in this section apply to information encoded *in binary*. A *bit* is an element of the binary field, a finite field with two elements $\mathbb{F}_2 \stackrel{\text{def}}{=} \{0, 1\}$. A *bitstring* of size n is an n -dimensional vector of bits $\vec{v} \in \{0, 1\}^n$. For notational simplicity, we denote bitstrings as $v_1 v_2 \cdots v_n$ instead of using the usual vector notation $(v_1, v_2, \dots, v_n)$. For example, 0010 is a bitstring of length 4.

All information in modern digital systems is encoded in binary. Every file on your computer consists of long sequences of bits that are copied between disk memory, RAM, caches, and CPU registers. For example, using the ASCII encoding convention, the character `a` is encoded as the bitstring 01100001, and the character `b` is encoded as 01100010. We can express a message that consists of several characters by concatenating the bitstrings that correspond to each character, like so:

$$\text{"baba"} \quad \Leftrightarrow \quad 01100010 \ 01100001 \ 01100010 \ 01100001.$$

For example, if the secret message $\vec{m}$ is the text `beer@18h`, we'll represent it as a sequence of eight chunks of eight bits each, making for a bitstring of total length 64:

$$\vec{m} = \underbrace{01100010}_b \underbrace{01100101}_e \dots \dots \dots \underbrace{01101000}_h \in \{0, 1\}^{64}.$$

Representing the text `beer@18h` as a 64-dimensional binary vector allows us to manipulate it using linear algebra operations. This may sound like a little thing—a convenience at best—but it's actually a big deal since it opens the door to many applications.

In this section we'll describe how basic vector operations on bitstrings can be used for cryptography. In the next section, we'll learn about *error-correcting codes*, which apply matrix operations on binary information vectors to protect them from noise.

The XOR operation

The XOR operation, usually denoted $\oplus$, corresponds to addition in the finite field $\mathbb{F}_2$:

$$c = a \oplus b \quad \Leftrightarrow \quad c = (a + b) \bmod 2.$$

Table 7.1 shows the results of the XOR operation on all possible combinations of inputs.

$\oplus$	0	1
0	0	1
1	1	0

Table 7.1: The XOR operation $\oplus$ applied to all possible binary inputs.

The XOR operation acting on two bitstrings of the same length is the XOR operation applied to each of their elements:

$$\vec{c} = \vec{a} \oplus \vec{b} \quad \Leftrightarrow \quad c_i = (a_i + b_i) \pmod{2}.$$

The XOR operation isn't anything new—it's just the normal vector addition operation for vectors with components in the finite field $\mathbb{F}_2$.

Intuitively, you can think of XOR as a “conditional toggle” operation. Computing the XOR of a bit $a \in \mathbb{F}_2$ with 0 leaves the bit unchanged, while computing the XOR with 1 toggles the bit, changing 0 to 1 or 1 to 0. To illustrate this “toggle” effect, consider the bitstring $\vec{v} = 0010$ and the results of XOR-ing it with the all-zeros bitstring 0000, the all-ones bitstring 1111, and a random-looking bitstring 0101:

$$0010 \oplus 0000 = 0010, \quad 0010 \oplus 1111 = 1101, \quad 0010 \oplus 0101 = 0111.$$

In the first case, none of the bits are flipped; in the second case all the bits are flipped; and in the third case only the second and fourth bits are flipped.

An important property of the XOR operation is that it is its own inverse: $\vec{z} \oplus \vec{z} = \vec{0}$, for all $\vec{z}$. In particular, if XOR-ing the bitstring $\vec{x}$ with some bitstring $\vec{k}$ produces $\vec{y}$, then XOR-ing $\vec{y}$ with $\vec{k}$ recovers the original vector $\vec{x}$:

$$\text{if } \vec{x} \oplus \vec{k} = \vec{y}, \quad \text{then } \vec{y} \oplus \vec{k} = \vec{x}.$$

This statement follows from the self-inverse property of the XOR operation and the associative law for vector addition:

$$\begin{aligned} \vec{y} \oplus \vec{k} &= (\vec{x} \oplus \vec{k}) \oplus \vec{k} \\ &= \vec{x} \oplus (\vec{k} \oplus \vec{k}) \\ &= \vec{x} \oplus \vec{0} \\ &= \vec{x}. \end{aligned}$$

The self-inverse property of the XOR operation is the basis for the one-time pad encryption and decryption operations.

One-time pad cryptosystem

The one-time pad encryption scheme is based on computing the XOR of the message $\vec{m}$ and the secret key $\vec{k}$, which has the same length as the message. The security of the protocol depends crucially on how the secret key $\vec{k}$ is generated, and on how the key is used:

- The key must be kept secret. Only Alice and Bob know $\vec{k}$.
- The key must be random: the value of each bit of the key $k_i \in \vec{k}$ is random; $k_i = 1$ with probability 50%, and $k_i = 0$ with probability 50%.
- Alice and Bob must *never reuse* the secret key.

One-time pad encryption

To obtain the ciphertext $\vec{c}$ for the message $\vec{m}$, Alice computes the XOR of $\vec{m}$ and the secret key $\vec{k}$:

$$\text{Enc}(\vec{m}, \vec{k}) \stackrel{\text{def}}{=} \vec{m} \oplus \vec{k} = \vec{c}.$$

One-time pad decryption

Upon receiving the ciphertext $\vec{c}$, Bob decrypts it by computing the XOR with the secret key:

$$\text{Dec}(\vec{c}, \vec{k}) \stackrel{\text{def}}{=} \vec{c} \oplus \vec{k} = \vec{m}.$$

From the self-inverse property of the XOR operation, we know XOR-ing any bitstring with the secret key $\vec{k}$ twice acts like the identity operation: $\vec{m} \oplus \vec{k} \oplus \vec{k} = \vec{m}$.

Discussion

Observe that using this encryption method to send an n -bit message requires n bits of secret key. Since bits of the secret key cannot be reused, we say that sending n bits securely “consumes” n bits of secret key. This notion of secret keys as a “consumable resource” is an important general theme in cryptography. Indeed, coming up with encryption and decryption functions is considered the easy part of cryptography, and the hard parts of cryptography are *key management* and *key distribution*.

Consider a practical use of the one-time pad encryption scheme. Sending multiple messages $m_1, m_2, \dots$ requires multiple secret keys $\vec{k}_1, \vec{k}_2, \dots$ that Alice uses when she wants to communicate with Bob.

Imagine a pad of paper; each page of the pad contains one secret key $\vec{k}_i$. Since keys cannot be reused, Alice must keep flipping through the pad, using the key from each page only once. This is where the name “one-time pad” comes from.

One-time pad security

Security definitions in cryptography are based on different assumptions about the powers of the eavesdropper Eve. We need to formally define what *secure* means, before we can evaluate the security of any cryptosystem. Modern cryptography is a subfield of mathematics, so we shouldn’t be surprised if the definition of security is stated in mathematical terms.

Definition: Indistinguishability under chosen-plaintext attack (IND-CPA) A cryptosystem is considered secure in terms of indistinguishability if no eavesdropper can distinguish the ciphertexts of two messages $\vec{m}_a$ and $\vec{m}_b$ chosen by the eavesdropper, with a probability greater than guessing randomly.

This is one of the strongest security standards we can expect from a cryptosystem. It is also the simplest to understand. This definition of security can be understood as a game played between Alice and Eve. Suppose Eve can force Alice to send only one of two possible messages $\vec{m}_a$ or $\vec{m}_b$. Every time Eve sees a ciphertext $\vec{c}$, she must guess whether Alice sent $\vec{m}_a$ or $\vec{m}_b$. We want to compare an Eve that has access to $\vec{c}$, with a “control Eve” that only has access to a random string $\vec{r}$ that is completely uncorrelated with the message or the ciphertext. A cryptosystem is secure, according to the *indistinguishability under chosen-plaintext* security definition, if the Eve with access to $\vec{c}$ is no better at distinguishing between messages $\vec{m}_a$ and $\vec{m}_b$ than she is at guessing randomly (which is the best that “control Eve” can do).

We won’t go into the formal details of the math, but we need to specify exactly what we mean by “with a probability greater than guessing randomly.” Control Eve has a completely random string $\vec{r}$, which contains no information about the message $\vec{m}$, so control Eve must guess randomly and her probability of success is $\frac{1}{2}$. An Eve that can distinguish the ciphertext $\vec{c}_a = \text{Enc}(\vec{m}_a)$ from the ciphertext $\vec{c}_b = \text{Enc}(\vec{m}_b)$ with a probability significantly greater than $\frac{1}{2}$ is considered to have an advantage in distinguishing the ciphertext. Any such scheme is not considered secure in terms of IND-CPA. Intuitively, the IND-CPA definition of security captures the notion that Eve should learn no information about the message $\vec{m}$ after seeing its ciphertext $\vec{c}$.

Sketch of security proof The one-time pad encryption system is secure according to IND-CPA because of the assumption we make about the shared, secret key $\vec{k}$ —namely that it is generated randomly. Each bit $k_i \in \vec{k}$ takes on the binary value 1 with probability 50%, and the value 0 with probability 50%.

Eve knows the plaintext of the two possible messages m_a and m_b , but she can't tell which message was sent from the ciphertext $\vec{c}$ since she can't distinguish between two equally likely alternative scenarios. In Scenario 1, Alice sent $\vec{m}_a$, and the secret key is $\vec{k}_a$, where $\vec{c} = \vec{m}_a \oplus \vec{k}_a$. In Scenario 2, Alice sent $\vec{m}_b$, and the secret key is $\vec{k}_b$, where $\vec{c} = \vec{m}_b \oplus \vec{k}_b$. The XOR operation combines the randomness of $\vec{m}$ with the randomness in the secret key $\vec{k}$, so trying to distinguish whether $\vec{m}_a$ or $\vec{m}_b$ was sent is just as difficult as distinguishing $\vec{k}_a$ and $\vec{k}_b$. Since $\vec{k}$ is completely random, Eve is forced to guess randomly. Thus the probability of determining the correct message is no better than guessing, which is precisely the requirement for the definition of security.

The randomness of the shared secret key is crucial to the security of the one-time pad encryption scheme. In general, we can think of *shared randomness* (shared secret key) as a communication resource that allows for mathematically-secure private communication between two parties. But what if Alice and Bob don't have access to shared randomness (that is, to some shared secret)? In the next section, we'll introduce a different type of cryptosystem that doesn't depend on a shared secret between Alice and Bob.

Public key cryptography

Assume Alice and Bob are dissidents in two neighbouring countries who want to communicate with each other, but they can't trust the network that connects them. If the dissidents cannot meet in person, obtaining a shared secret key to use for encryption will be difficult. They can't send the secret key $\vec{k}$ over the network because the eavesdropper will see it, and thus be able to decrypt all subsequent encrypted communications between the dissidents. This is a major limitation of all cryptosystems in which the same secret is used to encrypt and decrypt.

Do not despair dear readers, the System hasn't won yet: the dissidents can use *public-key* cryptography techniques to share a secret key over the untrusted network, in plain view of all state and non-state sponsored eavesdroppers. The security of public-key cryptosystems comes from some clever math operations (which are performed in a finite field), and the computational difficulty

of “reversing” these mathematical operations for very large numbers. For example, the security of the *Rivest–Shamir–Adleman* (RSA) cryptosystem depends on the difficulty of factoring integers. It’s easy to compute the product of two integers d and e ; but given only the product de , it is computationally difficult to find the factors e and d . For very large prime numbers e and d , even the letter agencies will have a difficult time finding the factors of de . This means ordinary citizens can use the power of math and cryptography to defend against corporate spying and the police state. Understanding cryptography will become increasingly important in the 21st century.

Definitions

In a public-key cryptosystem, the secret key is actually a pair of keys, $\vec{k} = \{e_{\vec{k}}, d_{\vec{k}}\}$, where $e_{\vec{k}}$ is the public encryption key and $d_{\vec{k}}$ is the private decryption key. The same encryption function, called Enc , is used for both encryption and decryption, but with different keys.

To use public-key cryptography, each communicating party must generate their own public-private key pairs. We’ll focus on Alice, and assume Bob performs analogous steps. Alice generates a public-private key pair $\{e_{\vec{k}}, d_{\vec{k}}\}$, and then shares the public part of the key with Bob. Note the public key can be shared openly, in plain sight, and it’s not a problem if Eve intercepts it—this is why it’s called a public key—everyone is allowed to know it.

Encryption

Bob uses Alice’s public encryption key whenever he wants to send Alice a secret message. To encrypt message $\vec{m}$, Bob uses the function Enc and Alice’s public encryption key $e_{\vec{k}}$ as follows:

$$\vec{c} = \text{Enc}(\vec{m}, e_{\vec{k}}).$$

When Alice receives the ciphertext $\vec{c}$, she uses her private key $d_{\vec{k}}$ (that only she knows) to decrypt the message:

$$\vec{m} = \text{Dec}(\vec{c}, d_{\vec{k}}).$$

Observe that public-key cryptosystems are inherently many-to-one: anyone who knows Alice’s public key $e_{\vec{k}}$ can create encrypted messages that only she can decode.

Digital signatures

Alice can also use her public-private key pair to broadcast one-to-many *authenticated* statements $\vec{s}$, meaning receivers can be sure the

statements they receive were sent by Alice. The math is the same; we just use the keys in the opposite order. Alice encrypts the statement $\vec{s}$ to produce a ciphertext,

$$\vec{c} = \text{Enc}(\vec{s}, d_{\vec{k}}),$$

then publishes the encrypted post $\vec{c}$ to her blog or a public forum. Everyone who knows Alice's public key $e_{\vec{k}}$ can decrypt the post $\vec{c}$ to obtain the statement $\vec{s}$:

$$\vec{s} = \text{Dec}(\vec{c}, e_{\vec{k}}).$$

The interesting property here is that we can be sure the statement $\vec{s}$ was sent by Alice, since only she controls the private key $d_{\vec{k}}$. This digital signature scheme makes it difficult for any third parties to impersonate Alice, since they don't know Alice's private key $d_{\vec{k}}$. This is the principle behind *digital signatures* used in the delivery of software updates.

We don't have space in this section to delve deeper into public-key cryptography, but we'll illustrate the main ideas through an example.

Example: ssh keys for remote logins

The secure shell protocol allows users to login to remote hosts. To login as user on the server `remotehost.com`, run the command `ssh user@remotehost.com` on the command prompt. The remote host will usually ask you for a password before it lets you login.

Passwords are the weakest form of authentication. Given enough effort, it's always possible for an attacker to guess your password. Furthermore, people tend to use the same password for multiple accounts, meaning that if one account is compromised then all accounts will be compromised. In this section, we'll describe a more secure approach for ssh logins based on public key authentication.

Setting up key-based authentication for ssh requires that steps are performed on your machine and on the remote server. Commands prefixed with `laptop>` should be typed on your local machine, while commands that start with `remotehost>` are to be executed on the remote server. Note these commands assume a UNIX environment. If you're using a Windows machine, I recommend you install `cygwin` to obtain the same functionality.

The first step is to generate an ssh private-public key pair in the subdirectory called `.ssh` inside your home directory:

```
laptop> cd ~ # go to your home dir ~
laptop> mkdir .ssh # create a .ssh subdir
laptop> ssh-keygen -t rsa -b 4096 # generate an RSA key pair
```


Accept the defaults for the questions you're asked, and be sure to set a password to protect your ssh private key. A new public key pair will be created in the directory `~/.ssh/`. The private key is contained in the file `~/.ssh/id_rsa`, and the corresponding public key is in the file `~/.ssh/id_rsa.pub`. You can confirm that these files exist by running the command `ls .ssh`. Additionally, you can ensure that these files have restrictive access permissions by issuing the command,

```
laptop> chmod -R go-rwx .ssh
```

This command has the effect of recursively (`-R`) changing permissions for users in the same group (`g`) as you and other (`o`) users, by removing (`-`), read (`r`), write (`w`), and execute (`x`) permissions. Basically, nobody other than you should be allowed to touch these files.

Another useful command to know involves adding the private key to a "key chain," which is a temporary store of credentials tied to your current login session on your laptop:

```
laptop> ssh-add -K ~/.ssh/id_rsa    # add priv. key to keychain
```

This command will prompt you to enter the password you used when generating the ssh key pair. Now that the private key is on your keychain, you won't be prompted for a password for a little while.

Next, copy the public key `~/.ssh/id_rsa.pub` from your laptop to the remote server. You can do this via the command `scp` as follows:

```
laptop> scp id_rsa.pub user@remotehost:~/key_from_laptop.pub
```

You'll need to substitute `user` and `remotehost` with the actual username and hostname (or IP address) of the remote server for which you want to set up the ssh-key login.

The final step is to place the public key information in a special file called `~/.ssh/authorized_keys`, located on the remote host. You can do this using the following commands:

```
laptop> ssh user@remotehost
remotehost> cd ~
remotehost> mkdir .ssh
remotehost> cat key_from_laptop.pub >> .ssh/authorized_keys
remotehost> chmod -R go-rwx .ssh
```

Inspect the resulting file and you'll see it contains a copy of the public key from your laptop. Now, if you log out of the remote host and try reconnecting to it using `ssh user@remotehost.com`, you'll be logged in automatically without needing to enter a password.

Once you've set up your login using ssh keys for your server, it's a good idea to completely disable password logins. In general, this is the best approach for ensuring security, and many hosting services only support logins using ssh keys.

Discussion

We'll conclude with some basic advice for programmers who want to use cryptography in their software. Rather than trying to roll your own crypto functions, keep in mind it's best to use established libraries. Libraries are good for you—there are many ways cryptosystems can fail, and experts have thought about defending against many potential attacks. Don't be a cowboy programmer—use the libraries.

Links

[Signal is a secure messaging app for mobile]

<https://whispersystems.org/>

[Public-key cryptography general concepts]

https://en.wikipedia.org/wiki/Public-key_cryptography

Exercises

E7.14 Alice wants to send the message $\vec{m} = 0110\ 1000\ 0110\ 1001$ to Bob. They have pre-shared the secret key $\vec{k} = 1010\ 0111\ 0010\ 0111$. Compute the ciphertext $\vec{c} = \text{Enc}(\vec{m}, \vec{k}) = \vec{m} \oplus \vec{k}$ that Alice will send to Bob. Verify that Bob will obtain the correct message after decrypting.

7.10 Error-correcting codes

The raw information-carrying capacity of a DVD is roughly 5.64GB; which is about 20% more than the 4.7GB of data that your computer will let you write to it. Why this overhead? Are DVD manufacturers trying to cheat you? Actually, they're looking out for you; the extra space is required for the *error-correcting code* that is applied to your data before writing it to the disk. Without the error-correcting code, even the tiniest scratch on the surface of the disk would make the disk unreadable, destroying your precious data. In this section, we'll learn how error-correcting codes work.

Error-correcting codes play an essential part in the storage, the transmission, and the processing of digital information. Even the slightest change to a computer program will make it crash—computer programs simply don't like it when you fiddle with their bits. Crashing was the norm back in the 1940s as illustrated by this quote:

"Two weekends in a row I came in and found that all my stuff had been dumped and nothing was done. I was really annoyed because I wanted

those answers and two weekends had been lost. And so I said, Dammit, if the machine can detect an error, why can't it locate the position of the error and correct it?" —Richard Hamming

Richard Hamming was a researcher at Bell in the 1940s. He ran into the problem of digital data corruption, and decided to do something to fix it. As a solution, he figured out a clever way to encode k bits of information into n bits of storage, such that it's possible to recover the information even if some errors occurred on the storage medium. An *error-correcting code* is a mathematical strategy for defending against erasures and errors. Hamming's invention of error-correcting codes became a prerequisite for the modern age of computing—after all, reliable computation is much more useful than unreliable computation.

Definitions

An *error-correcting code* is a prescription for encoding *binary* information. Recall that bits are elements of the binary field, $\mathbb{F}_2 = \{0, 1\}$. A bitstring of length n is an n -dimensional vector of bits $\vec{v} \in \{0, 1\}^n$. For example, 0010 is a bitstring of length 4.

We use several parameters to characterize error-correcting codes:

- k : the size, or length, of the messages for the code.
- $\vec{x}_i \in \{0, 1\}^k$: a *message*. Any bitstring of length k is a valid message.
- n : the size of the codewords in the code.
- $\vec{c}_i \in \{0, 1\}^n$: the *codeword* that corresponds to message $\vec{x}_i$.
- A *code* consists of 2^k codewords $\{\vec{c}_1, \vec{c}_2, \dots\}$, one for each of the possible messages $\{\vec{x}_1, \vec{x}_2, \dots\}$.
- $d(\vec{c}_i, \vec{c}_j)$: the *Hamming distance* between codewords $\vec{c}_i$ and $\vec{c}_j$.
- An (n, k, d) code is a procedure for encoding messages into codewords; $\text{Enc} : \{0, 1\}^k \rightarrow \{0, 1\}^n$, which guarantees the *minimum distance* between any two codewords is at least d .

The *Hamming distance* between two bitstrings $\vec{x}, \vec{y} \in \{0, 1\}^n$ counts the number of bits where the two bitstrings differ:

$$d(\vec{x}, \vec{y}) \stackrel{\text{def}}{=} \sum_{i=1}^n \delta(x_i, y_i), \quad \text{where } \delta(x_i, y_i) = \begin{cases} 0 & \text{if } x_i = y_i, \\ 1 & \text{if } x_i \neq y_i. \end{cases}$$

Intuitively, the Hamming distance between two bitstrings measures the minimum number of substitutions required to transform one bitstring into the other. For example, the Hamming distance between

codewords $\vec{c}_1 = 0010$ and $\vec{c}_2 = 0101$ is $d(\vec{c}_1, \vec{c}_2) = 3$, because it takes three substitutions (also called *bit flips*) to convert $\vec{c}_1$ to $\vec{c}_2$ or vice versa.

An (n, k, d) code is defined by a function $\text{Enc} : \{0, 1\}^k \rightarrow \{0, 1\}^n$ that encodes messages $\vec{x}_i \in \{0, 1\}^k$ into codewords $\vec{c}_i \in \{0, 1\}^n$. Usually the encoding procedure Enc is paired with a decoding procedure, $\text{Dec} : \{0, 1\}^n \rightarrow \{0, 1\}^k$, which recovers messages from (possibly corrupted) codewords.




Figure 7.14: An error-correcting scheme using the encoding function Enc and the decoding function Dec to protect against the effect of noise (denoted ζ). Each message $\vec{x}$ is encoded into a codeword $\vec{c}$. The codeword $\vec{c}$ is transmitted through a *noisy channel* that can corrupt the codeword by transforming it into another bitstring $\vec{c}'$. The decoding function Dec looks for a valid codeword $\vec{c}$ that is close in Hamming distance to $\vec{c}'$. If the protocol is successful, the decoded message will match the transmitted message $\vec{x}' = \vec{x}$, despite the noise (ζ).

Linear codes

A code is *linear* if its encoding function Enc is a linear transformation:

$$\text{Enc}(\vec{x}_i + \vec{x}_j) = \text{Enc}(\vec{x}_i) + \text{Enc}(\vec{x}_j), \text{ for all messages } \vec{x}_i, \vec{x}_j.$$

An (n, k, d) linear code encodes k -bit messages into n -bit codewords with minimum inter-codeword distance d . Linear codes are interesting because their encoding function Enc can be implemented as a matrix multiplication. We use the following terms when defining linear codes as matrices:

- $G \in \mathbb{F}_2^{k \times n}$: the *generating matrix* of the code. Each codeword $\vec{c}_i$ is produced by multiplying the message $\vec{x}_i$ by G from the right:

$$\text{Enc}(\vec{x}_i) = \vec{c}_i = \vec{x}_i G.$$

- $\mathcal{R}(G)$: the row space of the generator matrix is called the *code space*. We say a codeword $\vec{c}$ is valid if $\vec{c} \in \mathcal{R}(G)$, which means there exists some message $\vec{x} \in \{0, 1\}^k$ such that $\vec{x}G = \vec{c}$.
- $H \in \mathbb{F}_2^{(n-k) \times n}$: the *parity check matrix* of the code. The *syndrome* vector $\vec{s}$ of any bitstring $\vec{c}'$ is obtained by multiplying $\vec{c}'^T$ by H

from the left:

$$\vec{s} = H\vec{c}'^T.$$

If $\vec{c}'$ is a valid codeword (no error occurred), then $\vec{s} = \vec{0}$. If $\vec{s} \neq \vec{0}$, we know an error has occurred. The syndrome information helps us correct the error.

We can understand linear codes in terms of the input and output spaces of the encoding function $\text{Enc}(\vec{x}) = \vec{x}G$. Left multiplication of G by a k -dimensional row vector produces a linear combination of the rows of G . Thus, the set of all possible codewords (called the *code space*) corresponds to the row space of G .

Every vector in the null space of G is orthogonal to every codeword $\vec{c}_i$. We can construct a parity-check matrix H by choosing any basis for the null space for G . We call H the orthogonal complement of G , which means $\mathcal{N}(G) = \mathcal{R}(H)$. Alternately, we can say the space of n -dimensional bitstrings decomposes into orthogonal subspaces of valid and invalid codewords: $\mathbb{F}_2^n = \mathcal{R}(G) \oplus \mathcal{R}(H)$. We know $H\vec{c}^T = \vec{0}$ for all valid codewords $\vec{c}$. Furthermore, the *syndrome* obtained by multiplying an invalid codeword $\vec{c}'$ with the parity check matrix $\vec{s} = H\vec{c}'^T$ can help us characterize the error that occurred, and correct it.

Coding theory

The general idea behind error-correcting codes is to choose the 2^k codewords so they are placed far apart from each other in the space $\{0,1\}^n$. If a code has minimum distance $d \geq 2$ between codewords, then this code is robust to one-bit errors. To understand why, imagine a bubble of radius one (in Hamming distance) around each codeword. When a one-bit error occurs, a codeword will be displaced from its position, but it will remain within the bubble of radius one. In other words, if a one-bit error occurs, we can still find the correct codeword by looking for the closest valid codeword. See Figure 7.15 for an illustration of a set of codewords that are $d > 2$ distance apart. Any bitstring that falls within one of the bubbles will be decoded as the codeword at the centre of the bubble. We cannot guarantee this decoding procedure will succeed if more than one errors occur.

Observation 1 An (n, k, d) -code can correct up to $\lfloor \frac{d}{2} \rfloor$ errors.

The notation $\lfloor x \rfloor$ describes the *floor* function, which computes the closest integer value smaller than x . For example, $\lfloor 2 \rfloor = 2$ and $\lfloor \frac{3}{2} \rfloor = \lfloor 1.5 \rfloor = 1$. We can visualize Observation 1 using Figure 7.15 by imagining the radius of each bubble is $\lfloor \frac{d}{2} \rfloor$ instead of 1.




Figure 7.15: The rectangular region represents the space of binary strings of length n . Each codeword c_i is denoted with a black dot. A “bubble” of Hamming distance one around each codeword is shown. Observe that the distance between any two codewords is greater than two ($d > 2$). By Observation 1, we know this code can correct any one-bit error ($\lfloor \frac{d}{2} \rfloor \geq 1$).

Repetition code

The simplest possible error-correcting code is the *repetition code*, which protects information by recoding multiple copies of each message bit. For instance, we can construct a $(3, 1, 3)$ code by repeating each message bit three times. The encoding procedure Enc is defined as follows:

$$\text{Enc}(0) = 000 = \vec{c}_0, \quad \text{Enc}(1) = 111 = \vec{c}_1.$$

Three bit flips are required to change the codeword $\vec{c}_0$ into the codeword $\vec{c}_1$, and vice versa. The Hamming distance between the codewords of this repetition code is $d = 3$.

Encoding a string of messages $x_1 x_2 x_3 = 010$ results in a string of codewords 000111000 . We can apply the “majority vote” decoding strategy using the following decoding function Dec , defined by

$$\begin{aligned} \text{Dec}(000) &= 0, \text{Dec}(100) = 0, \text{Dec}(010) = 0, \text{Dec}(001) = 0, \\ \text{Dec}(111) &= 1, \text{Dec}(011) = 1, \text{Dec}(101) = 1, \text{Dec}(110) = 1. \end{aligned}$$

Observe that any one-bit error is corrected. For example, the message $x = 0$ is encoded as the codeword $\vec{c} = 000$. If an error occurs on the first bit during transmission, the received codeword will be $\vec{c}' = 100$, and majority-vote decoding will correctly output $x = 0$. Since $d > 2$ for this repetition code, the code can correct all one-bit errors.

The Hamming code

The $(7, 4, 3)$ *Hamming code* is a linear code that encodes four-bit messages into seven-bit codewords with minimum Hamming distance

of $d = 3$ between any two codewords. The generator matrix for the Hamming code is

$$G = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix}.$$

Note that other possibilities for the matrix G exist. Any permutation of the columns and rows of the matrix will be a generator matrix for a $(7, 4, 3)$ Hamming code. We have chosen this particular G because of the useful structure in its parity-check matrix H , which we'll discuss shortly.

Encoding

We'll now look at how the generating matrix is used to encode four-bit messages into seven-bit codewords. Recall that all arithmetic operations are performed in the finite field $\mathbb{F}_2$. The message $(0, 0, 0, 1)$ is encoded as the codeword

$$(0, 0, 0, 1)G = (0, 0, 0, 0, 1, 1, 1),$$

similarly $(0, 0, 1, 0)$ is encoded into $(0, 0, 1, 0)G = (0, 0, 1, 1, 0, 0, 1)$. Now consider the message $(0, 0, 1, 1)$, which is a linear combination of the messages $(0, 0, 1, 0)$ and $(0, 0, 0, 1)$. To obtain the codeword for this message, we can multiply it with G as usual to find $(0, 0, 1, 1)G = (0, 0, 1, 1, 1, 1, 0)$. Another approach is to use the linearity of the code and add the codewords for the messages $(0, 0, 1, 0)$ and $(0, 0, 0, 1)$: $(0, 0, 1, 1, 0, 0, 1) + (0, 0, 0, 0, 1, 1, 1) = (0, 0, 1, 1, 1, 1, 0)$.

Decoding with error correction

The minimum distance for this Hamming code is $d = 3$, which means it can correct one-bit errors. In this section, we'll look at some examples of bit-flip errors that can occur, and discuss the decoding procedure we can follow to extract messages—even from a corrupted codeword $\vec{c}'$.

The *parity-check matrix* for the $(7, 4, 3)$ Hamming code is

$$H = \begin{bmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{bmatrix}.$$

This matrix is the orthogonal complement of the generating matrix G . Every valid codeword $\vec{c}$ is in the row space of G , since $\vec{c} = \vec{x}G$

for some message $\vec{x}$. Since the rows of H are orthogonal to $\mathcal{R}(G)$, the product of H with any valid codeword will be zero: $H\vec{c}^T = 0$.

On the other hand, if the codeword $\vec{c}'$ contains an error, then multiplying it with H will produce a nonzero *syndrome* vector $\vec{s}$:

$$H\vec{c}'^T = \vec{s} \neq 0.$$

The decoding procedure Dec uses the information in the syndrome vector $\vec{s}$ to correct the error. In general, the decoding function can be a complex procedure that involves $\vec{s}$ and $\vec{c}'$. In the case of the Hamming code, the decoding procedure is very simple because the syndrome vector $\vec{s} \in \{0, 1\}^3$ contains the binary representation of the location where the bit-flip error occurred. Let's look at an example to illustrate how error correction works.

Example Suppose we send the message $\vec{x} = (0, 0, 1, 1)$ encoded as the codeword $\vec{c} = (0, 0, 1, 1, 1, 1, 0)$. If an error on the last bit occurs in transit, the received codeword will be $\vec{c}' = (0, 0, 1, 1, 1, 1, 1)$. Computing the syndrome for $\vec{c}'$, we obtain

$$\vec{s} = H\vec{c}'^T = \begin{bmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$

The syndrome vector $\vec{s} = (0, 0, 1)$ corresponds to the binary string 001, which is the number one. Note we count bits from right to left when interpreting the syndrome of an error: bit one is the rightmost bit of the codeword, bit two is the second to last, and so on. This syndrome 001 tells us the location of the error is on bit one of the codeword, which is the rightmost bit. After correcting the error by flipping the rightmost bit, we obtain the correct codeword $\vec{c} = (0, 0, 1, 1, 1, 1, 0)$, which decodes to the message $\vec{x} = (0, 0, 1, 1)$ that was sent.

Let's check the error-correcting ability of the Hamming code with another single-bit error. If a bit-flip error occurs on bit four (counting from the right), the received codeword will be $\vec{c}'' = (0, 0, 1, 0, 1, 1, 0)$. The syndrome for $\vec{c}''$ is $H\vec{c}''^T = (1, 0, 0)$, which corresponds to the number four when interpreted in binary. Again, we're able to obtain the position where the error has occurred from the syndrome.

The fact that the syndrome tells us where the error has occurred is not a coincidence, but a consequence of deliberate construction of the

matrices G and H of the Hamming code. Let's analyze the possible received codewords $\vec{c}'$, when the transmitted codeword is $\vec{c}$:

$$\vec{c}' = \begin{cases} \vec{c} & \text{if no error occurs} \\ \vec{c} + \vec{e}_i & \text{if bit-flip error occurs in position } i, \end{cases}$$

where $\vec{e}_i$ is a vector that contains a single one in position i . Indeed a bit flip in the i^{th} position is the same as adding one in that position, since we're working in the finite field $\mathbb{F}_2$.

In the case when no error occurs, the syndrome will be zero $H\vec{c} = \vec{0}$, because H is defined as the orthogonal complement of the code space (the row space of G). In the case when a single error occurs, the syndrome calculation only depends on the error:

$$\vec{s} = H\vec{c}'^T = H(\vec{c} + \vec{e}_i)^T = H\vec{c} + H\vec{e}_i = H\vec{e}_i.$$

If you look carefully at the structure in the parity-check matrix H , you'll see its columns contain the binary encoding of the numbers between seven and one. With this clever construction of the matrix H , we're able to obtain a syndrome that tells us the binary representation of where an error has occurred.

Discussion

Throughout this section, we referred to "the" $(7, 4, 3)$ Hamming code, but in fact there exists much freedom when defining a Hamming code with these dimensions. For example, we're free to perform any permutation of the columns of the generator matrix G and the parity check matrix H , and the resulting code will have the same properties as the $(7, 4, 3)$ Hamming code discussed in this section.

The term *Hamming code* actually applies to a whole family of linear codes. For any $r > 2$, there exists a $(2^r - 1, 2^r - r - 1, 3)$ Hamming code that has similar structure and properties as the *Hamming (7, 4, 3) code*. The ability to "read" the location of the error directly from the syndrome is truly a marvellous mathematical construction particular to the Hamming code. Other types of error-correcting codes that infer the error from the syndrome vector $\vec{s}$ may require more complicated procedures.

Note that all Hamming codes have minimum distance $d = 3$, which means they allow us to correct $\lfloor \frac{3}{2} \rfloor = 1$ bit errors. Hamming codes are therefore not appropriate for use with communication channels on which multi-bit errors are likely to occur. There exist other code families, like the *Reed-Muller codes* and *Reed-Solomon codes*, which can be used in noisier scenarios. For example, Reed-Solomon codes are used by NASA for deep-space communications and for error correction on DVDs.

Error-detecting codes

Another approach for dealing with errors is to focus on *detecting* the errors, rather than trying to correct them. Error-detecting codes, like the *parity-check code*, are used in scenarios where it is possible to retransmit messages. If the receiver detects a transmission error has occurred, she can ask the sender to retransmit the corrupted message. The receiver will be like, “Yo, Sender, I got your message, but its *parity* was *odd*, so I know there was an error and I want you to send that message again.” Error detection and retransmission is how internet protocols work (TCP/IP).

The *parity-check code* is a simple example of an error-detecting code. The *parity* of a bitstring describes whether the number of 1s in the string is odd or even. The bitstring 0010 has odd parity, while the bitstring 1100 has even parity. We can compute the parity of any bitstring by taking the sum of its bits—the sum being performed in the finite field $\mathbb{F}_2$.

A simple $(k + 1, k, 2)$ parity-check code is created by appending a single bit p (the parity bit) to the end of every message to indicate the parity of the message bitstring $x_1x_2 \cdots x_k$. We append $p = 1$ if the message has odd parity, and $p = 0$ if the message has even parity. The resultant message-plus-parity-check bitstring $\vec{c} = x_1x_2 \cdots x_kp$ will always have even parity.

If a single bit-flip error occurs during transmission, the received codeword $\vec{c}'$ will have odd parity, which tells us the message data has been affected by noise. More advanced error-detecting schemes can detect multiple errors, at the cost of appending more parity-check bits at the end of messages.

Links

[The Hamming distance between bitstrings]

https://en.wikipedia.org/wiki/Hamming_distance

[More examples of linear codes on Wikipedia]

https://en.wikipedia.org/wiki/Linear_code

https://en.wikipedia.org/wiki/Hamming_code

https://en.wikipedia.org/wiki/Reed-Muller_code

Exercises

E7.15 Find the codeword $\vec{c}$ that corresponds to the message $\vec{x} = (1, 0, 1, 1)$ for the $(7, 4, 3)$ Hamming code, which has the generator matrix G as given on page 417.

E7.16 Construct the $(5, 4, 2)$ parity check code’s encoding matrix G .

7.11 Fourier analysis

Way back in the 17th century, Isaac Newton carried out a famous experiment using light beams and glass prisms. He showed that when a beam of white light passes through a prism, it splits into a rainbow of colours: the rainbow is red at one end, followed by orange, yellow, green, blue, and finally violet at the other end. This experiment showed that white light is made of *components* with different colours. Using the language of linear algebra, we can say that white light is a “linear combination” of different colours.

Today we know that different colours of light correspond to electromagnetic waves with different frequencies: red light has a frequency around 450 THz, while violet light has a frequency around 730 THz. We can therefore say that white light is made of components with different frequencies. The notion of describing complex phenomena in terms of components with different frequencies is the main idea behind *Fourier analysis*.

Fourier analysis is used to describe sounds, vibrations, electric signals, radio signals, light signals, and many other phenomena. The Fourier transform allows us to represent all these “signals” in terms of components with different frequencies. Indeed, the Fourier transform can be understood as a change-of-basis operation that converts a signal from a time basis to a frequency basis:

$$[\mathbf{v}]_t \quad \Leftrightarrow \quad [\mathbf{v}]_f.$$

For example, if $\mathbf{v}$ represents a musical vibration, then $[\mathbf{v}]_t$ corresponds to the vibration as a function of time, while $[\mathbf{v}]_f$ corresponds to the frequency content of the vibration. Depending on the properties of the signal in the time domain and the choice of basis for the frequency domain, different Fourier transformations are possible.

We’ll study three different bases for the frequency domain based on orthonormal sets of sinusoidal and complex exponential functions. The *Fourier series* is a representation for continuous periodic functions $f(t) \in \{\mathbb{R} \rightarrow \mathbb{R}\}$; that is, functions that satisfy $f(T + t) = f(t)$. The Fourier basis used in the Fourier series is the set of sines and cosines of the form $\sin(\frac{2\pi n}{T}t)$ and $\cos(\frac{2\pi n}{T}t)$, which form an orthogonal set. The *Fourier transform* is the continuous version of the Fourier series. Instead of a countable set of frequency components, the frequency representation of the signal is described by a complex-valued continuous function $f(\omega) \in \{\mathbb{R} \rightarrow \mathbb{C}\}$. Instead of a continuous time parameter $t \in \mathbb{R}$, certain signals are described in terms of N samples from the time signal: $\{f[t]\}_{t \in [0,1,\dots,N-1]}$. The *discrete Fourier transform* is a version of the Fourier transform for signals defined at discrete time samples. Table 7.2 shows a summary of these three

Fourier-type transformations. The table indicates the class of functions for which the transform applies, the Fourier basis for the transform, and the frequency-domain representation used.

Fourier transformations

Name	Time domain	Fourier basis	Frequency domain
FS	$f(t) \in \{\mathbb{R} \rightarrow \mathbb{R}\}$ s.t. $f(t) = f(t + T)$	$1, \{\cos\left(\frac{2\pi n}{T}\right)\}_{n \in \mathbb{N}_+},$ $\{\sin\left(\frac{2\pi n}{T}\right)\}_{n \in \mathbb{N}_+}$	$(a_0, a_1, b_1, \dots)$
FT	$f(t) \in \{\mathbb{R} \rightarrow \mathbb{R}\}$ s.t. $\int_{-\infty}^{\infty} f(t) ^2 dt < \infty$	$\{e^{i\omega t}\}_{\omega \in \mathbb{R}}$	$f(\omega) \in \{\mathbb{R} \rightarrow \mathbb{C}\}$
DFT	$f[t] \in \{[N] \rightarrow \mathbb{R}\}$	$\{e^{\frac{i\omega t}{N}}\}_{w \in [N]}$	$f[w] \in \{[N] \rightarrow \mathbb{C}\}$

Table 7.2: This table shows the time domains, Fourier bases, and frequency domains for the three Fourier transformations that we'll discuss. The *Fourier series* (**FS**) converts periodic continuous time signals into Fourier coefficients. The *Fourier transform* (**FT**) converts finite-power continuous signal into continuous functions of frequency. The *discrete Fourier transform* (**DFT**) is the discretized version of the Fourier transform.

Fourier transforms are normally studied by physics and electrical engineering students during their second or third year at university. You, dear readers, thanks to your understanding of linear algebra, are ready for a 10-page sneak-peek of the Fourier transformations course right now. In this section, we'll study the mathematics behind the Fourier transforms and discuss how they're used in practical signal-processing applications. Before we jump into signal processing, let's look at something much simpler: the vibration of a guitar string.

Example 1: Describing the vibrations of a string

Imagine a string of length L that is stretched between two attachment points, like a guitar string. If you pluck this string it will vibrate. As it vibrates, we can describe the displacement of the string from its resting (straight, or non-vibrating) position as a function $f(x)$, where $x \in [0, L]$. The longest vibration along the string's length is called the *fundamental*, while other, shorter vibrations are called *overtone*s. See Figure 7.16 for an illustration. When you pluck the guitar string, your ears perceive the fundamental and the overtones combined as a single tone. The relative prominence of the fundamental vs. over-

tone frequencies varies among instruments.

When strings vibrate, only certain *modes* of vibration continue to vibrate over time, while others decay (fade out and stop) more quickly. Once the string is plucked and the initial vibration begins, the energy in the string “bounces around,” reflecting from the fixed attachments at both ends, and many vibrations cancel out. After some time, the only vibrations remaining on the string are the following sine-like vibrations:

$$\mathbf{e}_1(x) = \sin\left(\frac{\pi}{L}x\right), \quad \mathbf{e}_2(x) = \sin\left(\frac{2\pi}{L}x\right), \quad \mathbf{e}_3(x) = \sin\left(\frac{3\pi}{L}x\right), \quad \dots$$

Vibrations of the form $\mathbf{e}_n(x) = \sin\left(\frac{n\pi}{L}x\right)$ persist because they are stable, which means they satisfy the physics constraints imposed on the string.² One constraint is that the string’s endpoints are clamped down: $\mathbf{e}_n(0) = 0$ and $\mathbf{e}_n(L) = 0$. Another constraint requires that the vibration of the string $f(x)$ must satisfy the equation $f(x) = f(L - x)$.




Figure 7.16: Standing waves on a string with length $L = 1$. The longest vibration is called the *fundamental*. Other vibrations are called *overtones*.

The functions $\{\mathbf{e}_n(x)\}_{n \in \mathbb{N}^+}$ form a basis for the set of vibrations that satisfy $f(0) = 0$, $f(L) = 0$, and $f(x) = f(L - x)$ on $[0, L]$. Any vibration on the string can be described in terms of a sequence of coefficients $(a_1, a_2, a_3, \dots)$:

$$f(x) \text{ s.t. } f(x) = f(L - x) \quad \Leftrightarrow \quad (a_1, a_2, a_3, a_4, \dots).$$

The coefficient a_i represents how much of the i^{th} vibration exists on the string.

Depending how you pluck the string, the shape of the vibrating string $f(x)$ will be some linear combination of the vibrations $\mathbf{e}_n(x)$:

$$\begin{aligned} f(x) &= a_1 \sin\left(\frac{\pi}{L}x\right) + a_2 \sin\left(\frac{2\pi}{L}x\right) + a_3 \sin\left(\frac{3\pi}{L}x\right) + \dots \\ &= a_1 \mathbf{e}_1(x) + a_2 \mathbf{e}_2(x) + a_3 \mathbf{e}_3(x) + \dots \end{aligned}$$

²You can learn more about the physics of standing waves from this tutorial: <http://www.phy.duke.edu/~rgb/Class/phy51/phy51/node34.html>.

That's quite crazy if you think about it: rather than describing the exact shape of the vibrating string $f(x)$, it's sufficient to specify a list of coefficients to describe the vibration. The coefficient a_1 tells us the prominence of the fundamental vibration, and the other coefficients tell us the prominence of the overtones. Since the laws of physics restrict the possible vibrations to linear combinations of the basis vibrations $\{\mathbf{e}_n(x)\}_{n \in \mathbb{N}^+}$, we can represent any vibration on the string through its sequence of coefficients $(a_1, a_2, a_3, \dots)$.

The main idea

The guitar string example shows it's possible to understand a complex real-world system in terms of a succinct description—the coefficients $(a_1, a_2, a_3, \dots)$. This general pattern occurs when solving many science problems stated in terms of differential equations. Often, the solutions can be expressed as linear combinations of an orthonormal basis of functions. Using this example, we're now positioned to explain the main idea behind Fourier transformations and their connection to the concept of inner product space, which we studied in Section 6.4.

The fundamental starting point of all Fourier-type reasoning is to define the inner product space for the solution space, and a basis of orthonormal functions for that space. In the case of the vibrating string, the space of string vibrations consists of functions $f(x)$ on the interval $[0, L]$ that satisfy $f(0) = 0$, $f(L) = 0$, and $f(x) = f(L - x)$. The inner product operation we'll use for this space is

$$\langle f(x), g(x) \rangle \stackrel{\text{def}}{=} \int_{x=0}^{x=L} f(x)g(x) dx,$$

and the functions $\{\mathbf{e}_n(x)\}_{n \in \mathbb{N}^+}$ form an orthogonal basis for this space. We can verify that $\langle \mathbf{e}_m(x), \mathbf{e}_n(x) \rangle = 0$ for all $m \neq n$. See Exercise E7.17 for a hands-on experience.

Since $\{\mathbf{e}_n(x)\}_{n \in \mathbb{N}^+}$ forms a basis, we can express any function as a linear combination of the basis functions. The complicated-looking Fourier transform formulas that we'll discuss shortly can therefore be understood as applications of the general change-of-basis formula.

Change-of-basis review Let's quickly review the change-of-basis formula. You recall the change-of-basis operation, right? Given two orthonormal bases, $B = \{\hat{e}_1, \hat{e}_2, \hat{e}_3\}$ and $B' = \{\hat{e}'_1, \hat{e}'_2, \hat{e}'_3\}$, you should know how to convert a vector $[\vec{v}]_B$ expressed with respect to basis B , into coordinates with respect to basis B' . In case you've forgotten, look back to page 229, where we define the change-of-basis matrix

${}_{B'}[\mathbb{1}]_B$ that converts from B coordinates to B' coordinates, and its inverse ${}_B[\mathbb{1}]_{B'}$ that converts from B' coordinates to B coordinates:

$$[\vec{v}]_{B'} = {}_{B'}[\mathbb{1}]_B [\vec{v}]_B \quad \Leftrightarrow \quad [\vec{v}]_B = {}_B[\mathbb{1}]_{B'} [\vec{v}]_{B'}.$$

When the bases are orthonormal, the change-of-basis operation depends only on the inner product between the vectors of the two bases:

$$\vec{v}'_i = \sum_{j=1}^3 \langle \hat{e}'_i, \hat{e}_j \rangle v_j \quad \Leftrightarrow \quad \vec{v}_i = \sum_{j=1}^3 \langle \hat{e}_i, \hat{e}'_j \rangle v'_j.$$

Orthonormal bases are nice like that—if you know how to compute inner products, you can perform the change of basis.

To apply the general formulas for change of basis to the case of the vibrating string, we must first define the two bases and compute the inner product between them. We can think of the function $f(x) \in \{[0, L] \rightarrow \mathbb{R}\}$ as described in the “default basis” $\mathbf{e}_x$, which is equal to one at x and zero everywhere else. The Fourier basis for the problem consists of the functions $\mathbf{e}_n = \sin(\frac{n\pi}{L}x)$ for $n \in \mathbb{N}^+$. The inner product between $\mathbf{e}_x$ and $\mathbf{e}_n$ is

$$\langle \mathbf{e}_x, \mathbf{e}_n \rangle = \int_0^L \mathbf{e}_x(y) \sin(\frac{n\pi}{L}y) dy = \sin(\frac{n\pi}{L}x),$$

since the function $\mathbf{e}_x(y)$ is equal to zero everywhere except at $y = x$.

We can obtain the change-of-basis transformations by extending the notion of a matrix-vector appropriately. We use the integral on the interval $[0, L]$ for the change of basis to the Fourier coefficients, and the infinite sum over coefficients to change from the space of Fourier coefficients to functions:

$$a_n = \int_0^L \sin(\frac{n\pi}{L}x) f(x) dx \quad \Leftrightarrow \quad f(x) = \sum_{n=1}^{\infty} \sin(\frac{n\pi}{L}x) a_n.$$

Compare these formulas with the basic change-of-basis formula between bases B and B' .

$$\begin{bmatrix} f(0) \\ \vdots \\ f(x) \\ \vdots \\ f(L) \end{bmatrix} = \begin{bmatrix} \text{orange sine wave} & \text{yellow sine wave} & \text{green sine wave} & \text{blue sine wave} & \cdots \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ \vdots \end{bmatrix}$$

Figure 7.17: Any string vibration $f(x)$ can be represented as coefficients $(a_1, a_2, a_3, a_4, \dots)$ with respect to the basis of functions $\mathbf{e}_n(x) = \sin(\frac{n\pi}{L}x)$.

Figure 7.17 illustrates the right side of the change-of-basis formula. A vibration expressed in the Fourier basis $(a_1, a_2, a_3, a_4, \dots)$ is transformed to a function $f(x) \in \{[0, L] \rightarrow \mathbb{R}\}$ in the default basis.

Analysis and synthesis

The terminology for describing Fourier transformations borrows from the world of music. A *frequency analyzer* is a device used to measure the sonic frequencies contained in audio signals. A *synthesizer* is an instrument that generates combinations of electronic signals to produce various musical sounds.

Similarly, in the jargon of Fourier transformations, we refer to the change of basis from the default basis to the Fourier basis $({}_f[1]_x)$ as *Fourier analysis*, and the transformation from the Fourier basis to the default basis $({}_x[1]_f)$ as *Fourier synthesis*. This terminology, in addition to sounding extra fancy, gives us some useful intuition about the purpose of the Fourier transform. Given a vibration on the string, we use the Fourier transform to analyze the vibration, decomposing it into its constituent vibrations. The *Fourier analysis equations* describe how to calculate the Fourier coefficients of the different frequencies.

The opposite process of the Fourier transformation—starting with the Fourier coefficients and combining them to obtain a vibration in the default basis—is called *synthesis*, in analogy with a synthesizer that can generate myriad sounds. The synthesis of a vibration as a function of x is expressed mathematically as a linear combination of vibrations with different frequencies multiplied by frequency-domain coefficients. For example, you can pick any set of coefficients $(a_1, a_2, a_3, a_4, \dots)$, and form the linear combination $\sum_{n=1}^N a_n \sin(\frac{2\pi n}{L}x)$; then you can check to hear what that vibration sounds like. This is how synthesizers work: they can mimic the sounds of various musical instruments and natural phenomena, as well as create totally new sounds, by producing different linear combinations of a basis of vibrations.

Fourier series

When discussing Fourier transformations, it is natural to work with functions of time $f(t)$, which we call *signals*. A signal in the “time basis” is specified by its values $f(t)$ for all times t . This is the “time domain” representation for signals. The *Fourier series* corresponds to the following change-of-basis operation:

$$f(t) \text{ s.t. } f(t) = f(T + t) \quad \Leftrightarrow \quad (a_0, a_1, b_1, a_2, b_2, \dots).$$

The coefficients $(a_0, a_1, b_1, a_2, b_2, \dots)$ are called the *Fourier coefficients* of $f(t)$. The *Fourier series* applies to all signals $f(t)$ that satisfy the constraint $f(t) = f(t + T)$ for all $t \in [0, T)$. We say such signals are *periodic* with period T .

The basis used for the Fourier series consists of the set of cosine and sine functions with frequencies that are multiples of $\frac{2\pi}{T}$:

$$\left\{ \sin\left(\frac{2\pi n}{T}t\right) \right\}_{n \in \mathbb{N}_+} \quad \text{and} \quad \left\{ \cos\left(\frac{2\pi n}{T}t\right) \right\}_{n \in \mathbb{N}}.$$

This family of functions forms an orthogonal set with respect to the standard inner product:

$$\langle f(t), g(t) \rangle = \frac{1}{T} \int_{t=0}^{t=T} f(t)g(t) dt.$$

Every periodic signal $f(t)$ can be represented as a *Fourier series* of the form:

$$f(t) = a_0 + \sum_{n=1}^{\infty} a_n \cos\left(\frac{2\pi n}{T}t\right) + \sum_{n=1}^{\infty} b_n \sin\left(\frac{2\pi n}{T}t\right). \quad (\text{FSS})$$

This is the Fourier series *synthesis* equation, hence the label (FSS). The coefficient a_i represents how much of the i^{th} cosine-like vibration is contained in $f(t)$, and the coefficient b_i represents how much of the i^{th} sine-like vibration is contained in $f(t)$. See the illustration in Figure 7.18 for a visual representation of the synthesis equation. This representation of a periodic signal is directly analogous to the analysis we performed for the vibrating string, but this time we use both sines and cosines.

We compute the Fourier coefficients using the following formulas:

$$a_n = \frac{1}{T} \int_0^T f(t) \cos\left(\frac{2\pi n}{T}t\right) dt, \quad b_n = \frac{1}{T} \int_0^T f(t) \sin\left(\frac{2\pi n}{T}t\right) dt. \quad (\text{FSA})$$

These are the Fourier series *analysis* equations, hence the label (FSA). These transformations correspond to the standard change of basis for the function $f(t)$ in the time domain to the Fourier basis of cosines and sines.

For the Fourier series representation of a periodic time signal to be exact, we must compute an infinite number of coefficients in the Fourier series $(a_0, a_1, b_1, a_2, b_2, \dots)$. However, we're often interested in obtaining an approximation to a $f(t)$ using only a finite set of Fourier coefficients:

$$f(t) \approx a_0 + \sum_{n=1}^N a_n \cos\left(\frac{2\pi n}{T}t\right) + \sum_{n=1}^N b_n \sin\left(\frac{2\pi n}{T}t\right).$$

$$\begin{bmatrix} f(0) \\ \vdots \\ f(t) \\ \vdots \\ f(T) \end{bmatrix} = \begin{bmatrix} \text{red solid} & \text{orange dashed} & \text{yellow dashed} & \text{green dashed} & \text{blue dashed} & \dots \\ \text{orange solid} & \text{yellow dashed} & \text{green dashed} & \text{blue dashed} & \dots & \dots \\ \text{yellow solid} & \text{green dashed} & \text{blue dashed} & \dots & \dots & \dots \\ \text{green solid} & \text{blue dashed} & \dots & \dots & \dots & \dots \\ \text{blue solid} & \dots & \dots & \dots & \dots & \dots \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ b_1 \\ a_2 \\ b_2 \\ a_3 \\ b_3 \\ \vdots \end{bmatrix}$$

Figure 7.18: The Fourier series synthesis equation (FSS) represented as a matrix-vector product. A periodic signal $f(t)$ is represented as a series of Fourier coefficients $(a_0, a_1, b_1, a_2, b_2, a_3, b_3, \dots)$. The first column of the change-of-basis matrix corresponds to the constant component $1 = \cos(0)$. The remaining columns correspond to cosines and sines with different frequencies.

This is called a *Fourier series approximation* since the frequency representation does not contain the components with frequencies $\frac{N+1}{T}$, $\frac{N+2}{T}$, and higher. Nevertheless, these finite-series approximations of signals are used in many practical scenarios; it's much easier to compute a finite number of Fourier coefficients instead of an infinite number.

Example For an example calculation of the Fourier coefficients of the *square wave* signal, see bit.ly/fourier_series_square_wave by Joshua Vaughan. Note the square wave analyzed is an *odd* function, so its coefficients a_n are all zero.




Figure 7.19: The square-wave signal can be approximated by a linear combination of sine functions with different frequencies.

In the next section, we'll describe the Fourier transform, which is a continuous-frequency version of the Fourier series.

Fourier transform

The Fourier series representation applies only to periodic functions $f(t)$, but not every function of interest in signal processing is periodic. The *Fourier transform* applies to all signals $f(t)$ (periodic or not) that obey the *finite-power* constraint:

$$\int_{t=-\infty}^{t=+\infty} |f(t)|^2 dt \leq \infty.$$

This class of functions includes most signals of practical importance in communication scenarios.

The result of the *Fourier transform* is a complex-valued continuous function in the frequency domain:

$$f(t) \in \{\mathbb{R} \rightarrow \mathbb{R}\} \quad \Leftrightarrow \quad f(\omega) \in \{\mathbb{R} \rightarrow \mathbb{C}\}.$$

The basis used in the Fourier transforms is

$$\mathbf{e}_\omega = e^{i\omega t}, \text{ for } \omega \in \mathbb{R},$$

parametrized by the continuous parameters $\omega \in \mathbb{R}$. The functions $\{\mathbf{e}_\omega\}$ form an orthogonal basis with respect to the standard inner product for functions:

$$\langle f(t), g(t) \rangle = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \overline{f(t)} g(t) dt.$$

The change-of-basis from the time domain to the frequency domain is performed using the integral:

$$f(\omega) = \int \langle \mathbf{e}_\omega, \mathbf{e}_t \rangle f(t) dt = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-i\omega t} f(t) dt. \quad (\text{FTA})$$

This is the Fourier transform *analysis* equation. We can understand this formula as an instance of the general change-of-basis formula from the “default basis” for the time domain $\{\mathbf{e}_t\}_{t \in \mathbb{R}}$ to the basis $\{\mathbf{e}_\omega\}_{\omega \in \mathbb{R}}$. The function $f(\omega)$ is called the *spectrum* or *Fourier transform* of $f(t)$; it tells us all the information about the frequency content of the signal $f(t)$. Note *Fourier transform* can be used as a verb when referring to the change-of-basis transformation, or as a noun when referring to the function $f(\omega)$, which is the result of the Fourier transform.

The *Inverse Fourier transform* is the change of basis from the frequency domain back to the time domain. Given the frequency representation of a function $f(\omega)$, we can reconstruct the time representation of the signal using the Fourier transform *synthesis* integral:

$$f(t) = \int \langle \mathbf{e}_t, \mathbf{e}_\omega \rangle f(\omega) d\omega = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{i\omega t} f(\omega) d\omega. \quad (\text{FTS})$$

Note the similarity between the forward (analysis) and the inverse (synthesis) Fourier transform formulas. Indeed, these are representations of the general change-of-basis formula using the inner product for functions. Compare the Fourier transform change-of-basis formulas and the basic change-of-basis formulas between bases B and B' discussed on page 424. The conjugate-symmetric property of the inner product $\langle \mathbf{e}_t, \mathbf{e}_\omega \rangle = \overline{\langle \mathbf{e}_\omega, \mathbf{e}_t \rangle}$ explains the change from $e^{-i\omega t}$ to $e^{i\omega t}$.

Further discussion about the Fourier transform and its many applications is beyond the scope of the current section. If you take a signal-processing course (the *best* course in the electrical engineering curriculum), you'll learn all about the Fourier transform.

Discrete Fourier transform

Continuous-time signals are important in sound engineering, radio communications, and other communication scenarios. Signals can also be digitized and represented as discrete samples rather than continuous functions. A continuous-time signal $f(t)$ can be approximated by taking a finite set of N samples from the signal. This sampling results in a discrete-time signal $f[t] \in \{[N] \rightarrow \mathbb{R}\}$, where the shorthand $[N]$ denotes the sequence $[0, 1, \dots, N-1]$. We'll use square brackets around the function input to distinguish the discrete-time signal $f[t]$ from its continuous-time counterpart $f(t) \in \{\mathbb{R} \rightarrow \mathbb{R}\}$.

The *discrete Fourier transform* converts N samples of a discrete-time signal into a frequency representation with N frequency components:

$$f[t] \in \mathbb{R}, \text{ for } t \in [N] \quad \Leftrightarrow \quad f[w] \in \mathbb{C}, \text{ for } w \in [N].$$

The basis for the frequency domain consists of complex exponentials:

$$\mathbf{e}_w = e^{i\frac{2\pi w}{N}t}, \text{ for } t \in [N] \text{ and } w \in [N].$$

These functions form an orthonormal set with respect to the inner product $\langle f[n], g[n] \rangle \stackrel{\text{def}}{=} \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} \overline{f[n]} g[n]$. You can understand this inner product and this basis as a discrete version of the inner product and basis used for the Fourier transform; hence the name.

To convert from the time domain to the frequency domain, we use the discrete Fourier transform analysis equation:

$$f[w] = \sum_{t=0}^{N-1} \langle \mathbf{e}_w, \mathbf{e}_t \rangle f[t] = \frac{1}{\sqrt{N}} \sum_{t=0}^{N-1} f[t] e^{-i\frac{2\pi w}{N}t}. \quad (\text{DFTA})$$

This is the usual change-of-basis formula from the standard basis in the time domain $\{\mathbf{e}_t\}$ to the discrete Fourier basis $\{\mathbf{e}_w\}$.

To convert a signal from the frequency domain back to the time domain, we use the discrete Fourier transform synthesis equation:

$$f[t] = \sum_{w=0}^{N-1} \langle \mathbf{e}_t, \mathbf{e}_w \rangle f[w] = \frac{1}{\sqrt{N}} \sum_{w=0}^{N-1} f[w] e^{i \frac{2\pi w}{N} t}. \quad (\text{DFTS})$$

Again, this inverse change-of-basis transformation can be understood as a special case of the standard change-of-basis formula we learned in Section 4.3 (see page 229).

Sampling signals The *sample rate* is an important practical consideration to take into account when converting analog signals to digital form. Usually the samples of the continuous signal $f(t)$ are taken at uniform time intervals spaced by time Δt . The *sample rate* or *sampling frequency* f_s is defined as the inverse of the time between samples: $f_s = \frac{1}{\Delta t}$. For good performance, we must choose the sampling frequency f_s to be at least double that of the highest frequency of interest in the signal we're sampling.

For example, the upper limit of the range of human hearing is 20 kHz; therefore CD digital audio recordings are sampled at $f_s = 44.1$ kHz. This means one second of audio recording corresponds to 44 100 discrete audio samples, which we'll denote as a sequence $f[0], f[1], f[2], \dots, f[44099]$. The sample $f[0]$ corresponds to the value of the signal at $t = 0$, $f[1]$, which corresponds to the value of $f(t)$ at the next sampling time $t = \frac{1}{f_s}$, and so on for the rest of the samples $f[t] = f(t \frac{1}{f_s})$.

Digital signal processing

Fourier transforms play a central role in digital signal processing. Many image compression and sound compression algorithms depend on the discrete Fourier transform. For example, mp3 compression works by cutting an audio file into short time segments, transforming these segments to the frequency domain using the discrete Fourier transform, and then omitting the high-frequency components. Below, we give a high-level overview of the pipeline of transformations used for mp3 encoding and playback. We'll assume we start from a song recorded in the digital audio wav format. A wav file describes the song's sound intensity as a function of time; in other words, the song is represented in the time basis $[\text{song}]_t$.

The idea behind mp3 encoding is to leave out certain "non-essential" aspects of audio recordings. The field of psychoacoustics

develops models of human hearing that tell us which aspects of audio signals can and cannot be detected by humans. For example, we can delete the sound frequencies that are beyond the range of human hearing without noticeably degrading a recording's playback quality. Let's consider a simplified model of mp3 encoding that consists of a projection in the frequency domain ${}_f[\Pi]_f$. The overall mp3 encoding and playback pipeline can be described as a sequence of transformations:

$$\underbrace{[\text{song}']_t = {}_t[\mathbb{1}]_f}_\text{mp3 playback} \underbrace{{}_f[\Pi]_f {}_f[\mathbb{1}]_t}_{\text{mp3 encoding}} [\text{song}]_t.$$

Note we're modelling the mp3 pipeline using standard matrix-vector products, so the sound information flows from right to left in the above equation. First we use the Fourier transform, denoted ${}_f[\mathbb{1}]_t$, to transform the song from the time-basis to the frequency-basis, where the psychoacoustic projection ${}_f[\Pi]_f$ is applied. An mp3 file is the frequency-domain representation of the song after applying the projection: $[\text{song}]'_f = {}_f[\Pi]_f {}_f[\mathbb{1}]_t [\text{song}]_t$. Compression ratios up to a factor of 10 can be achieved using this approach: a 50 MB wav file can be compressed to a 5 MB mp3 file. During playback, the inverse change of basis, ${}_t[\mathbb{1}]_f$, transforms the song from the frequency-basis back to the time-basis.

Obviously some level of sound degradation occurs in this process, and the song you hear during playback $[\text{song}']_t$ sounds different from the original $[\text{song}]_t$. If good models of human hearing are used, the differences should be mostly imperceptible. Crucial to the mp3 encoding process, the psychoacoustic models are applied in the frequency domain. This is an example of the Fourier transform acting as a building block for a more complex procedure.

Discussion

In this section we learned about three different Fourier transformations: the Fourier series, the Fourier transform, and the discrete Fourier transform. All the scary-looking formulas we saw can be understood as special cases of the same idea: Fourier transformations are change-of-basis operations between a time basis and a frequency basis, which consists of orthogonal functions of the form $e^{i\omega t}$.

It may seem like the Fourier series, with its sine and cosine terms, is a different kind of beast; but it only looks different until you recall Euler's formula, $e^{i\theta} = \cos \theta + i \sin \theta$ (see page 102). Using Euler's formula, we can express the Fourier series of a function as a sequence of complex coefficients $c_n \in \mathbb{C}$ that contain the combined information

of both cosine- and sine-like components of a periodic signal $f(t)$. The complex Fourier coefficients c_n are obtained using the formula

$$c_n = \int_0^T \left\langle e^{i\frac{2\pi n}{T}t}, \mathbf{e}_t \right\rangle f(t) dt = \frac{1}{T} \int_0^T e^{-i\frac{2\pi n}{T}t} f(t) dt.$$

Problem P7.14 will ask you to verify the connection between the Fourier series of complex exponentials and the Fourier series of sines and cosines.

The Fourier change of basis is a general principle that can be applied in many contexts. All the transforms we discussed in this section use the family of orthogonal functions $e^{i\omega t}$, but there are other sets of orthogonal functions that can serve as alternate bases for different applications. To mention a few, we have the Bessel functions J_λ , the spherical harmonics $Y_{mn}(\theta, \psi)$, the Legendre polynomials $P_n(x)$, the Hermite polynomials $H_n(x)$, and the Laguerre polynomials $L_n(x)$. All these families of functions form orthogonal sets on different intervals with respect to different inner products.

Links

[Visualizations of Fourier synthesis of a square wave signal]

<http://codepen.io/anon/pen/jPGJMK/>

<http://bgrawi.com/Fourier-Visualizations/>

[Excellent video tutorial about digital audio processing]

<http://xiph.org/video/vid2.shtml>

[Website with animations that explain signal processing concepts]

<http://jackschaedler.github.io/circles-sines-signals/>

[A nice discussion on math.stackexchange.com]

<https://math.stackexchange.com/questions/1002/>

[Wikipedia pages of the three Fourier transforms we described]

https://en.wikipedia.org/wiki/Fourier_series

https://en.wikipedia.org/wiki/Fourier_transform

https://en.wikipedia.org/wiki/Discrete_Fourier_transform

[Orthogonal polynomials and generalized Fourier series]

<http://math24.net/orthogonal-polynomials.html>

Exercises

E7.17 Recall the functions $\mathbf{e}_n(x) = \sin(\frac{n\pi}{L}x)$ that can be used to describe all vibrations of a guitar string of length L . Verify that $\mathbf{e}_1(x)$ and $\mathbf{e}_2(x)$ are orthogonal functions by computing the inner product $\langle \mathbf{e}_1(x), \mathbf{e}_2(x) \rangle$. Use the inner product definition from page 424.

Hint: You might find the double-angle formula from page 78 useful.

E7.18 Explain why the Fourier series $(a_0, a_1, b_1, a_2, b_2, \dots)$ of a periodic function $f(t)$ contains a coefficient a_0 but not a coefficient b_0 .

Discussion

We've only scratched the surface of all the possible problem domains that can be modelled using linear algebra. The topics covered in this chapter are a small sample of the range of scientific, computing, and business activities that benefit from the use of matrix methods and understanding of vector spaces. Linear algebra allows you to perform complex numerical procedures using a high level of abstraction.

Normally, a linear algebra class should end here. We've covered all the information you need to know about vectors, matrices, linear transformations, and vector spaces; and also discussed several applications. Feel free to close this book now, content that your valiant effort to learn linear algebra is done. But perhaps you'd like to learn some further topics and make use of your linear algebra skills? If this sounds interesting, I encourage you to keep reading, as there are two more chapters of cool "optional material" ahead.

Chapter 8 covers the basics of probability theory, Markov chains, and the idea behind Google's PageRank algorithm for classifying web pages. Probability theory is not directly related to linear algebra, but the applications we'll discuss make heavy use of linear algebra concepts.

The laws of quantum mechanics govern physics phenomena at the femto-scale—think individual atoms. It's a common misconception to assume quantum laws are somehow mysterious or counter-intuitive, and that perhaps they can only be understood by people with a solid math background. Chapter 9 contains a concise introduction to the principles of quantum mechanics specifically tailored to people who know linear algebra. The material is adapted from lectures the author prepared for a graduate-level introductory course, so there will be no dumbing down or oversimplification. With your background in linear algebra, you can handle the real stuff.

More linear algebra applications

You'll find the following resources useful if you want to learn more about linear algebra applications.

[Three compilations of linear algebra applications]

<http://aix1.uottawa.ca/~jkhoury/app.htm>

<https://medium.com/@jeremyjkun/633383d4153f>

<http://people.math.harvard.edu/~knill/teaching/math21b2018/handouts/use.pdf>

[A document that describes many applications in detail]

http://www.geverstine.com/reprints/Everstine_linearalgebra.pdf

[33 miniatures: algorithmic applications of linear algebra]

<http://kam.mff.cuni.cz/~matousek/stml-53-matousek-1.pdf>

[A book about linear algebra applications to data science]

<http://amazon.com/Data-Science-from-Scratch/dp/149190142X>


7.12 Applications problems

It would be easy to think of all the applications of linear algebra presented in this chapter as a TV program, designed to entertain rather than teach. Certainly you can continue to the next chapter without solving any problems, but do you really want to do that to yourself?

Presented next are a number of practice problems that will test your understanding of the new concepts and give you a great opportunity to practice your linear algebra skills. The linear algebra techniques we learned in previous chapters are key building blocks for applications. So don't sit on your laurels thinking, "Yay, I'm in Chapter 7 and I know linear algebra now, I'm so good." Prove it.

P7.1 Consider the following chemical equation that describes how your body burns fat molecules: $\text{C}_{55}\text{H}_{104}\text{O}_6 + \text{O}_2 \rightarrow \text{CO}_2 + \text{H}_2\text{O}$. Balance this chemical equation.

P7.2 Check out this circuit containing two batteries and five resistors:



- Label the polarity of each resistor in the circuit.
- Write three KVL equations and two KCL equations.
- Rewrite the equations in the form $R\vec{I} = \vec{V}$, where R is a 5×5 matrix, $\vec{I} = (I_1, I_2, I_3, I_4, I_5)^T$, and $\vec{V}$ is a vector of constants.
- Find the value of the currents I_1 and I_5 given $V_1 = 15[\text{V}]$, $V_2 = 10[\text{V}]$, $R_1 = 1[\Omega]$, $R_2 = 1[\Omega]$, $R_3 = 4[\Omega]$, $R_4 = 2[\Omega]$, $R_5 = 2[\Omega]$.

Hint: The direction of the voltage drop across a resistor depends on the direction of the current flowing through it.

P7.3 Given the (x, y) pairs $(0, 0.9)$, $(1, 1.6)$, $(2, 2.1)$, and $(3, 2.4)$, find the best-fitting affine model $y = b + mx$ for this data.

Hint: Find the Moore–Penrose inverse.

P7.4 You just moved to a new city and you're looking for a new place to live. To get an idea of the rent prices per square foot, you check out the classifieds and find the following offers: a 200 sq ft mini-studio for \$500, a 300 sq ft studio for \$620, a 400 sq ft small apartment for \$750, a 500 sq ft one bedroom condo for \$890, a 900 sq ft two bedroom apartment for \$1250, and a 1000 sq ft apartment for \$1300. Find the best-fitting curve $p(x) = b + mx$ to the above data, where $p(x)$ represents the price for x square feet. What is the estimated price for a 700 sq ft apartment?

Hint: Use a computer algebra system like SymPy for the calculations. To get you started, I set up the problem data here: bit.ly/apt_rent_data.


P7.5 Redo P7.4, but this time find the best-fitting quadratic model $q(x) = c + bx + ax^2$ for the price per square foot. What is the estimated price for a 700 sq ft apartment predicted by the quadratic model?

Hint: After preprocessing, the data matrix must have three columns.

P7.6 Describe the possible types of solutions to the equation $A\vec{x} = \vec{b}$, where A is an $m \times n$ matrix. Discuss the following three cases: when A is wide ($m < n$), when A is square ($m = n$), and when A is tall ($m > n$).

Hint: Describe the column space, row space, and null space of the matrix.

P7.7 Find the adjacency matrix representation of the following graphs:



P7.8 For each of the graphs in P7.7, find the number of ways to get from vertex 1 to vertex 3 in two steps or less.

Hint: Obtain the answer by inspection or by looking at the appropriate entry of the matrix $\mathbb{1} + A + A^2$.


P7.9 Draw the graphs that correspond to these adjacency matrices:

a) $A = \begin{bmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix}$


b) $A = \begin{bmatrix} 1 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$

P7.10 Find the homogeneous coordinates representations of each of the transformations shown in Figure 7.20. The input to each transformation is the triangle with vertices $(0, 0)$, $(2, 0)$, $(0, 1)$ shown in Figure 7.20 (a).


Hint: Your answers should be 3×3 matrices. Recall that $\sin(\frac{\pi}{6}) = \frac{1}{2}$.




(a) Original




(b) Transformation 1




(c) Transformation 2



(d) Transformation 3



(e) Transformation 4




(f) Transformation 5

Figure 7.20: The effects of various linear transformations on a triangle.

P7.11 Find the homogeneous coordinates representations of each of the transformations shown in Figure 7.21. The input to each transformation is the triangle with vertices $(0,0)$, $(2,1)$, $(0,1)$ shown in Figure 7.20 (a).

Hint: Use your answers from parts (b) and (d) to answer parts (e) and (f).

P7.12 Find the homogeneous coordinates representation of the perspective transformation illustrated in the following figure. The observer is located at the origin $(0,0)$ and the projection line is given by the equation $2x + y = 2$.



Use the matrix to find the coordinates of the projected points p' and q' .

P7.13 Alice and Bob share the key $\vec{k} = 10010111 \ 01010011 \ 10011110$. Bob receives the ciphertext $\vec{c} = 11100100 \ 00100110 \ 11101110$ sent by Alice. What is the message sent by Alice? Find an ASCII lookup table on the web and use it to convert the binary message into characters.




Figure 7.21: The effects of various affine transformations on the triangular shape shown in Figure 7.20 (a).

P7.14 Consider the signal $f(t)$ that is periodic with period T . The coefficients of the Fourier series with complex coefficients are defined using the formula

$$c_n = \int_0^T \left\langle e^{i\frac{2\pi n}{T}t}, \mathbf{e}_t \right\rangle f(t) dt = \frac{1}{T} \int_0^T f(t) e^{-i\frac{2\pi n}{T}t} dt.$$

Show that $c_n = a_n - ib_n$, where a_n and b_n are the coefficients of the regular Fourier series for $f(t)$ defined in terms of cosines and sines.

Hint: Obtain the real and imaginary parts of c_n using Euler's formula.

P7.15 Find the Fourier transform of the function $f(t) = e^{-\frac{t^2}{2\sigma^2}}$.

Hint: Use the complete-the-square technique from page 28, and remove from the exponent a factor that does not contain t . The change of variable $t' = t + \sigma^2 i\omega$ might also come in handy. Recall $\int_{-\infty}^{\infty} e^{-\frac{x^2}{2\sigma^2}} dx = \sqrt{2\pi}\sigma$.

P7.16 Find the Fourier transform of the function $f(t) = e^{-|at|}$.

Hint: Split $e^{-i\omega t}$ into sine and cosine terms and use symmetry to convert the two-sided integral $\int_{-\infty}^{\infty} \cdots dt$ into a one-sided integral $\int_0^{\infty} \cdots dt$.

P7.17 Consider the function $f(t)$, which has Fourier transform $f(\omega)$. Find the Fourier transform of the function's derivative $f'(t)$, and express your answer in terms of $f(\omega)$.

Hint: Start from the formula for $f(t)$ in equation FTS on page 429.

Chapter 8

Probability theory

In this chapter, we'll use linear algebra concepts to explore the world of probability theory. Think of this as bonus material because the topics we'll discuss are not normally part of a linear algebra course. Given the general usefulness of probabilistic reasoning and the fact that you have already covered all the prerequisites, it would be a shame *not* to learn a bit about probability theory and its applications.

The chapter is structured as follows. In Section 8.1, we'll discuss probability distributions, which are mathematical models for describing random events. Section 8.2 introduces the concept of a *Markov chain*, which can be used to characterize the random transitions between different states of a system. Of the myriad of topics in probability theory, we've chosen to discuss probability distributions and Markov chains because they correspond one-to-one with vectors and matrices. This means you should feel right at home. In Section 8.3, we'll describe Google's PageRank algorithm for ranking webpages, which is an interesting application of Markov chains.

8.1 Probability distributions

Many phenomena in the world are inherently unpredictable. When you throw a six-sided die, one of the outcomes $\{1, 2, 3, 4, 5, 6\}$ will result, but you don't know which one. Similarly, when you toss a coin, you know the outcome will be either heads or tails but you can't predict which outcome will result. Probabilities are used to describe events where uncertainty plays a role. We can assign probabilities to the different outcomes of a dice roll, the outcomes of a coin toss, and also to many real-world systems. For example, we can build a probabilistic model of hard drive failures using past observations.

We can then calculate the probability that your family photo albums will survive the next 10 or 20 years. Backups my friends, backups.

Probabilistic models can help us better understand random events. The fundamental concept in probability theory is that of a *probability distribution*, which describes the likelihood of different outcomes of a random event. For example, the probability distribution for the roll of a fair die is $p_X = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6})^T$, and the probability distribution for a coin toss is $p_Y = (\frac{1}{2}, \frac{1}{2})^T$. Each entry of a probability distribution corresponds to the *probability mass* of a given outcome. This terminology borrows from the concept of mass distribution used in physics. The entries of a probability distribution satisfy the following conditions: each entry is a nonnegative number, and the sum of the entries is one. These two conditions are known as the *Kolmogorov axioms* of probability.

Strictly speaking, understanding linear algebra is not required for understanding probability theory. However, vector notation is very effective for describing probability distributions. Your existing knowledge of vectors and the rules for matrix multiplication will allow you to quickly understand many concepts in probability theory. Probabilistic reasoning is highly useful, so it's totally worth taking the time to learn about it.

Random variables

A random variable X is described by a probability distribution p_X . Before we formally define the notion of a probability distribution, we must introduce some formalism. We denote by $\mathcal{X}$ (calligraphic X) the *sample space* of the random variable X , which is the set of all possible outcomes of the random variable. A *discrete* random variable has a finite sample space. For example, we can describe the outcome of rolling a six-sided die using the random variable $X \in \mathcal{X}$, where the sample space is $\mathcal{X} = \{1, 2, 3, 4, 5, 6\}$. The number of possible outcomes is six: $|\mathcal{X}| = 6$.

We can describe the random outcome of a coin toss as a random variable $Y \in \{\text{heads}, \text{tails}\}$. The possible outcomes of the coin toss are $\mathcal{Y} = \{\text{heads}, \text{tails}\}$. The number of possible outcomes is two: $|\mathcal{Y}| = 2$.

In the case of the hard drive failure model, we can define the random variable $L \in \mathbb{N}$ as the years of a hard drive's lifetime before it fails. Using the random variable L we can describe interesting scenarios using probabilistic reasoning. For example, the condition that a hard drive will function correctly for at least eight years can be described as $L \geq 8$.

Probability distributions

The probability distribution of a discrete random variable $X \in \mathcal{X}$ is a vector of $|\mathcal{X}|$ nonnegative numbers that sum to one. Using mathematically precise notation, we write the definition of p_X as follows:

$$p_X \in \mathbb{R}^{|\mathcal{X}|} \quad \text{such that} \quad p_X(x) \geq 0, \forall x \in \mathcal{X} \quad \text{and} \quad \sum_{x \in \mathcal{X}} p_X(x) = 1.$$

A probability distribution is a vector in $\mathbb{R}^{|\mathcal{X}|}$ that satisfies two special requirements: its entries must be nonnegative and the sum of the entries must be one.

Events

In addition to the individual outcomes of a random variable, we can define *events* that consist of compositions of individual outcomes. Events correspond to different subsets of the sample space and are usually defined using words like “{event description},” which is the standard curly-bracket notation used to denote sets and subsets.

Recall the random variable X that describes the outcome of rolling a six-sided die. The sample space for this random variable is $\mathcal{X} = \{1, 2, 3, 4, 5, 6\}$. We can define various events that involve the six-sided die and describe them as subsets of the sample space. For example, the event $\{X \text{ is odd}\}$ corresponds to the subset $\{1, 3, 5\} \subset \mathcal{X}$.

Events are useful because they allow us to describe specific combinations of outcomes that are of practical interest. We can use the logical operators OR and AND to define composite events. The logical OR operator for events corresponds to the union operation for sets. For example, X resulting in an odd number,

$$\{X \text{ is odd}\} = \{X = 1\} \text{ OR } \{X = 3\} \text{ OR } \{X = 5\},$$

corresponds to the union of sets $\{1\} \cup \{3\} \cup \{5\} = \{1, 3, 5\} \subset \mathcal{X}$.

The logical AND operator corresponds to the intersection of sets, and is often used to describe composite events that involve multiple random variables.

With these simple principles for describing probabilistic phenomena, we can build rich combinations of random events, expectations, and statistics—and develop methods for predicting the likelihood of future events. In the remainder of this section, we’ll explore concepts of probabilistic reasoning through a series of examples.

Example 1 Consider again the random variable X that describes the outcome of rolling a six-sided die. Assuming the die is fair, the probability distribution of the random variable X is

$$p_X = \begin{bmatrix} \frac{1}{6} \\ \frac{1}{6} \\ \frac{1}{6} \\ \frac{1}{6} \\ \frac{1}{6} \\ \frac{1}{6} \end{bmatrix} \begin{matrix} \leftarrow p_X(1) \\ \leftarrow p_X(2) \\ \leftarrow p_X(3) \\ \leftarrow p_X(4) \\ \leftarrow p_X(5) \\ \leftarrow p_X(6) \end{matrix}.$$

The probability distribution p_X is a six-dimensional vector and the probability of each outcome is $\frac{1}{6}$. The sum of the entries is one.

Let's take a moment to discuss the different notations for referring to the entries of the distribution p_X . The x^{th} element of the vector p_X can be written in three equivalent ways:

$$p_{Xx} = p_X(x) = \Pr(\{X = x\}).$$

These three different ways of denoting the probability mass associated with outcome x in the distribution p_X represent three different ways of thinking about probability distributions. We can use the normal vector notation based on subscripts, p_{Xx} , but the appearance of two subscripts tends to be confusing. When the meaning of the random variable X is clearly apparent from the context, we can use the simpler notation p_x . The notation $p_X(x)$ emphasizes the idea that a probability distribution can be understood as a *probability mass function* that assigns probability mass for each of the possible outcomes. We can think of p_X as a function of the form $p_X : \{1, 2, 3, 4, 5, 6\} \rightarrow \mathbb{R}$. The most precise notation for probabilities is $\Pr(\{X = x\})$, which describes the probability, denoted $\Pr$, of the event that a random draw from X results in the outcome x , denoted $\{X = x\}$. Note we use the capital letter X when referring to the random variable, and the lowercase letter x to refer to particular outcomes of X . In the standard terminology for probability theory, we refer to the particular outcome x as a *realization* of the random variable X .

Let's use the distribution p_X to compute the probability of rolling a number five or greater, which corresponds to the event $\{X \geq 5\}$. Note the outcome "greater than or equal to five" occurs when we roll a five *or* a six. Thus, the complex event $\{X \geq 5\}$ is the logical OR of these two basic events:

$$\{X \geq 5\} = \{X = 5\} \text{ OR } \{X = 6\}.$$

The probability of the event $\{X \geq 5\}$ is

$$\begin{aligned}\Pr(\{X \geq 5\}) &= \Pr(\{X = 5\} \text{ OR } \{X = 6\}) \\ &= p_X(5) + p_X(6) = \frac{1}{6} + \frac{1}{6} = \frac{1}{3}.\end{aligned}$$

Since the outcomes five and six are mutually exclusive, the logical OR of these events corresponds to the sum of the probabilities of the outcomes.

Example 2 The probability distribution of the random variable Y that describes the outcome of a coin toss is

$$p_Y = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix} \begin{array}{l} \leftarrow p_Y(\text{heads}) \\ \leftarrow p_Y(\text{tails}). \end{array}$$

The probability distribution p_Y corresponds to a function of the form $p_Y: \{\text{heads}, \text{tails}\} \rightarrow \mathbb{R}$, but for the sake of notational convenience, we associate the probability of heads and tails with the first and second entries of a two-dimensional vector p_Y . After all, this is a book about linear algebra, so everything must be turned into a vector!

What is the probability of getting heads three times in a row? We need three variables, Y_1 , Y_2 , and Y_3 , to represent the outcomes of the three coin tosses. Each random variable Y_i is an independent copy of the random variable Y . The probability of the outcome “three heads in a row” corresponds to the outcome where all three coin tosses come out heads: the first coin toss results in heads ($\{Y_1 = \text{heads}\}$), the second coin toss results in heads ($\{Y_2 = \text{heads}\}$), and the third coin toss also results in heads ($\{Y_3 = \text{heads}\}$). The probability of throwing three heads in a row is given by:

$$\begin{aligned}\Pr(\{\text{three heads in a row}\}) &= \Pr(\{Y_1 = \text{heads}\} \text{ AND } \{Y_2 = \text{heads}\} \text{ AND } \{Y_3 = \text{heads}\}) \\ &= p_Y(\text{heads}) \cdot p_Y(\text{heads}) \cdot p_Y(\text{heads}) = \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{8}.\end{aligned}$$

Since each coin toss is an *independent* random event, the probability of the logical AND of the three events is the product of the probabilities of the three events.

* * *

The examples above only scratch the surface of possibilities for modelling real-world situations in terms of random events, and for reasoning mathematically about the probabilities of different outcomes. I encourage you to research the topic of probability theory further on your own.

Expectations

Consider a function $f : \mathcal{X} \rightarrow \mathbb{R}$ that assigns values to each of the possible outcomes of a random variable X . You can think of f as the payout function in a game of chance based on the random variable X . You obtain $f(x)$ dollars when the outcome of X is x . Using the usual notation for functions, we can write $f(X)$ to describe the payout of this game. Since the input X is a random variable, the output value $f(X)$ is also a random variable. We can obtain the *expected payout* of this game if we multiply each of the possible payouts by its probability of occurrence:

$$\mathbb{E}_X[f(X)] \stackrel{\text{def}}{=} \sum_{x \in \mathcal{X}} f(x)p_X(x).$$

The above equation introduces the *expectation operator* $\mathbb{E}_X$, which computes the weighted sum of the payouts for all possible outcomes. The payout amount for each outcome is multiplied by the probability of the outcome. The expected amount you'll take home after playing the game of chance n times is $n\mathbb{E}_X[f(X)]$. Sometimes you win big, sometimes you win less—by weighting each payout by its probability, we obtain the correct estimate of your expected winnings or losses.

Example 3 Consider the following game involving a six-sided die. You pay \$1 to roll the die and the payout for the game is as follows. If you roll a $\square$, a $\blacksquare$, a $\circ$, or a $\boxtimes$, you win nothing. If you roll a $\boxdot$, you win \$1. If you roll a $\boxplus$, you win \$4. Should you play this game?

The payout function for this game is defined as follows:

$$f(\square) = f(\blacksquare) = f(\circ) = f(\boxtimes) = \$0, \quad f(\boxdot) = \$1, \quad f(\boxplus) = \$4.$$

We'll model the die roll as a random variable X with distribution $p_X = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6})^T$. The expected payout for this game is

$$\begin{aligned} \mathbb{E}_X[f(X)] &= \sum_x f(x)p_X(x) \\ &= f(\square)\frac{1}{6} + f(\blacksquare)\frac{1}{6} + f(\circ)\frac{1}{6} + f(\boxtimes)\frac{1}{6} + f(\boxdot)\frac{1}{6} + f(\boxplus)\frac{1}{6} \\ &= (\$0)\frac{1}{6} + (\$0)\frac{1}{6} + (\$0)\frac{1}{6} + (\$0)\frac{1}{6} + (\$1)\frac{1}{6} + (\$4)\frac{1}{6} \\ &= \frac{\$1+\$4}{6} = \frac{\$5}{6} \approx 83 \text{ cents.} \end{aligned}$$

The expected payout of this game is less than the cost of playing, so it doesn't make sense to play. Computing the expected value for any casino game will quickly convince you not to play—the expected payout is always less than what you must pay to play.

Expected value and variance of random variables

We can use the expectation operator to compute certain characteristic properties of the random variable. The *expected value* and the *variance* are two properties of any random variable X that capture important aspects of its behaviour.

The *expected value* of the random variable X is computed using the formula

$$\mu_X \stackrel{\text{def}}{=} \mathbb{E}_X[X] = \sum_x x p_X(x).$$

The expected value is a single number that tells us what value of X we can expect to obtain on average from the random variable X . The expected value is also called the *average* or the *mean* of the random variable X .

The *variance* of the random variable X is defined as follows:

$$\sigma_X^2 \stackrel{\text{def}}{=} \mathbb{E}_X[(X - \mu_X)^2] = \sum_x (x - \mu_X)^2 p_X(x).$$

The variance formula computes the expectation of the squared distance of the random variable X from its expected value. The variance σ_X^2 , also denoted $\text{var}(X)$, gives us an indication of how clustered or spread the values of X are. A small variance indicates the outcomes of X are tightly clustered near the expected value μ_X , while a large variance indicates the outcomes of X are widely spread.

The expected value μ and the variance σ^2 are two central concepts in probability theory and statistics because they allow us to characterize any random variable. The expected value is a measure of the *central tendency* of the random variable, while the variance σ^2 measures its *dispersion*. Readers familiar with concepts from physics can think of the expected value as the *centre of mass* of the distribution, and the variance as the *moment of inertia* of the distribution.

Computing expected values and variances allows us to compare different random variables. Suppose you want to compare two random variables, X and Y . It can be difficult to compare their probability distributions p_X and p_Y directly, but we can compute their expected values and their variances and use these properties as the basis for the comparison. If we find $\mu_X > \mu_Y$, then we know X takes on larger values than Y on average. If $\sigma_X^2 > \sigma_Y^2$, then we know the probability mass function of X is more spread out than that of Y .

Conditional probability distributions

Probability theory also allows us to model dependencies between random variables. We use *conditional probability distributions* to describe situations where one random variable depends on another.

Consider the random variable X whose random outcomes depend on another variable Y . To describe this probability dependence, we must specify the probability distributions of X for each of the possible values of Y . This information is expressed as a conditional probability distribution:

$$p_{X|Y}(x|y) \stackrel{\text{def}}{=} \Pr(\{X = x\} \text{ given } \{Y = y\}).$$

The vertical bar is pronounced “given,” and it separates the unknown random variable from the random variable whose value is known. The distribution $p_{X|Y}$ satisfies the conditions for a probability distribution for all $y \in \mathcal{Y}$:

$$p_{X|Y} \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{Y}|} \text{ s.t. } p_{X|Y}(x|y) \geq 0 \text{ and } \sum_{x \in \mathcal{X}} p_{X|Y}(x|y) = 1, \forall y \in \mathcal{Y}.$$

It is natural to represent $p_{X|Y}$ as a $|\mathcal{X}| \times |\mathcal{Y}|$ matrix, with each column representing the probability of X for a given value of Y .

Example 4 Consider a probability distribution with six possible outcomes obtained by rolling one of two different dice—one fair and one biased. The probability distribution when rolling the fair die is $p_{X_f} = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6})^\top$. The probability distribution when rolling the biased die is $p_{X_b} = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, 0, 0)^\top$. Introducing the conditioning variable Y that describes which die is rolled, we can express the situation with the two dice as the following conditional probability distribution:

$$p_{X|Y} = \begin{bmatrix} \frac{1}{6} & \frac{1}{4} \\ \frac{1}{6} & \frac{1}{4} \\ \frac{1}{6} & \frac{1}{4} \\ \frac{1}{6} & \frac{1}{4} \\ \frac{1}{6} & 0 \\ \frac{1}{6} & 0 \end{bmatrix}.$$

The first column corresponds to the fair die $Y = f$, and the second column corresponds to the biased die $Y = b$.

Combining the probability distributions for the two dice, p_{X_f} and p_{X_b} , into a single conditional distribution, $p_{X|Y}$, allows us to construct more complex probabilistic structures, as illustrated in the next example.

Example 5 Consider an experiment in which the outcome of a coin toss Y decides which die we throw—the fair die or the biased die. Suppose the coin is biased with $p_Y = (\frac{3}{4}, \frac{1}{4})^T$. If the outcome of the coin toss is heads, we roll the fair die. If the coin toss gives tails, we roll the biased die. We’re interested in describing the random variable X , which corresponds to die rolls in which the fair die is used $\frac{3}{4}$ of the time and the biased die is used $\frac{1}{4}$ of the time. What is the probability distribution p_X for X ?

To model this situation, we combine the “which die” probability distribution $p_Y = (\frac{3}{4}, \frac{1}{4})^T$ with the conditional probability distribution $p_{X|Y}$ obtained in the previous example:

$$\begin{aligned} p_X(x) &= \sum_{y \in \mathcal{Y}} p_{X|Y}(x|y) p_Y(y) \\ &= p_{X|Y}(x|f) p_Y(\text{heads}) + p_{X|Y}(x|b) p_Y(\text{tails}) \\ &= p_{X|Y}(x|f) \frac{3}{4} + p_{X|Y}(x|b) \frac{1}{4} \\ &= \frac{3}{4} \left(\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6} \right)^T + \frac{1}{4} \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, 0, 0 \right)^T \\ &= \left(\frac{3}{16}, \frac{3}{16}, \frac{3}{16}, \frac{3}{16}, \frac{1}{8}, \frac{1}{8} \right)^T. \end{aligned}$$

The probabilistic mixture of two random events corresponds to a linear combination—and when you think “linear combination,” you immediately think “matrix-vector product representation,” right? Indeed, we can also express p_X as a matrix-vector product between the conditional probability distribution of the matrix $p_{X|Y}$ and the vector p_Y :

$$p_X = \begin{bmatrix} \frac{1}{6} & \frac{1}{4} \\ \frac{1}{6} & \frac{1}{4} \\ \frac{1}{6} & \frac{1}{4} \\ \frac{1}{6} & \frac{1}{4} \\ \frac{1}{6} & 0 \\ \frac{1}{6} & 0 \end{bmatrix} \begin{bmatrix} \frac{3}{4} \\ \frac{1}{4} \end{bmatrix} = \begin{bmatrix} \frac{3}{16} \\ \frac{3}{16} \\ \frac{3}{16} \\ \frac{3}{16} \\ \frac{1}{8} \\ \frac{1}{8} \end{bmatrix}.$$

You can verify that $\sum_x p_X(x) = 1$ as expected for all probability distributions.

Note we were able to use conditional probability distributions to describe complicated random events in terms of the matrix-vector product. The notion of a conditional probability distribution is a fundamental building block of probability theory. Many modern machine learning techniques use probabilistic models constructed from conditional probability distributions.

Interpretations of probability theory

One approach to understanding probability theory is to think about probabilities as describing relative frequencies of occurrence for different possible outcomes. The quantity $p_X(a)$ represents the proportion of outcomes of the event $\{X = a\}$ among all possible outcomes of X .

This suggests an approach for estimating the probability mass of each of the possible outcomes. To characterize some real-world phenomenon as the random variable X , we estimate the probability distribution of X by repeating the real-world phenomenon n times, where n grows to infinity. The probability mass of each outcome $a \in \mathcal{X}$ is defined as the following limit:

$$p_X(a) \stackrel{\text{def}}{=} \frac{N(\{X = a\}, n)}{n}, \quad n \rightarrow \infty,$$

where $N(\{X = a\}, n)$ is the number of times the outcome a occurs during n repetitions. If I tell you the probability mass of outcome $\{X = a\}$ is 0.3, it means that if you take 1000 draws from the random variable X , you can expect approximately 300 of those draws to result in the outcome $\{X = a\}$. Defining probabilities by counting the occurrences of different outcomes is called the *frequentist* approach to probability theory. In practice, we don't have time to actually repeat events an infinite number of times in order to obtain their exact probabilities; so think of this frequentist approach as a thought experiment, rather than applicable in practice.

Another conceptual interpretation of probability theory is to consider probability distributions as representations of our state of *knowledge* or *beliefs* about reality. Instead of describing some objective reality, the random variable X and its probability distribution p_X can represent our state of knowledge about the real-world phenomenon that X describes. Since p_X represents our state of knowledge about the random variable X , it makes sense to update the distribution p_X as we learn new facts. This is called the *Bayesian* approach to probability theory, named after Thomas Bayes who was an 18th century statistician. Consider the following example of Bayesian-style reasoning.

You're given a coin and asked to come up with a probability distribution that describes its chances of landing heads or tails. Initially you have no information about whether the coin is biased in favour of heads or tails, so it makes sense to start with the initial belief that the coin is fair. This is called the *prior belief*, or simply *prior*—it's the belief you hold prior to gathering data. If you toss the coin several times and observe significantly more heads than tails,

you can update your belief about the coin's probability distribution to account for your observations of the data. The specific technique for updating a prior belief in light of new observations is called *Bayes' rule*.

Both the frequentist and Bayesian perspectives lead to useful techniques for modelling random events. Frequentist methods are concerned with drawing principled conclusions given a set of empirical observations. Bayesian models are generally more flexible and allow us to combine empirical observations with priors that encode the domain knowledge provided by experts. The frequentist approach is useful when you want to analyze data, prepare reports, and come up with hard numbers about the data. The Bayesian approach is more useful for building machine learning applications like voice recognition.

Discussion

We described the basic notions in probability theory, such as random variables, probability distributions, and conditional probability distributions. These concepts are the bread-and-butter of probabilistic reasoning. Let's take a minute to recap and summarize this new material, linking it back to vectors and linear transformations—the main subjects of this book. The probability distribution of a discrete random variable is a vector of real numbers. For instance, the probability distribution describing the outcome of rolling a six-sided die is a six-dimensional vector, $p_X \in \mathbb{R}^6$. Probability distributions are vectors that obey some extra constraints: each entry must be a non-negative number, and the sum of the entries in the vector must equal one.

Conditional probability distributions are mappings that describe how a set of random “given” variables influence the probabilities of a set of random “outcome” variables. Conditional probability distributions can be represented as matrices, where each column of the matrix contains the outcome distribution for one of the values of the given variable. A conditional probability distribution with five possible outcomes, conditioned on a given variable with 10 possible states, is represented as a 5×10 matrix that contains probability distributions in each of its columns.

Conditional probability distributions are powerful tools for modelling real-world scenarios. Many machine learning algorithms involve the characterization, estimation, and exploitation of conditional probability distributions. We can also describe a noisy communication channel as a conditional probability distribution of the channel's outputs given each of the channel's possible inputs.

In the next section, we'll learn about an application of conditional probability distributions that describes the state of a system undergoing random transitions between a number of possible states. This is a special case of a conditional probability distribution for which the random conditioning variable and random outcome variable are of the same dimension. Indeed, the outcome variable and the conditional space both represent states of the same system at different times.

Links

[Discussion on the Bayesian way of thinking about probabilities]
https://en.wikipedia.org/wiki/Bayesian_probability

[Detailed discussion about Bayes rule]
<http://yudkowsky.net/rational/bayes>
<http://yudkowsky.net/rational/technical>

Exercises

E8.1 Do the following vectors represent probability distributions?

$$\text{a)} \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right)^T \quad \text{b)} \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)^T \quad \text{c)} (0.3, 0.3, -0.1, 0.5)^T$$

E8.2 Compute the expected value μ_X and the variance σ_X^2 of the random variable X that describes the outcome of rolling a six-sided die. The probability distribution of X is $p_X = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6})^T$.

8.2 Markov chains

So far, we've talked about random events without any reference to the flow of time. In this section we'll combine the idea of random variables with the notion of time. A *random process* is a model of a system that undergoes transitions between states over time. The state of the system at time t is described by a random variable X_t . We obtain the next state of the system X_{t+1} from the following conditional probability distribution:

$$p_{X_{t+1}|X_t X_{t-1} \dots X_0}(x_{t+1}|x_t x_{t-1} \dots x_0).$$

The random variable X_{t+1} depends on the previous states of the system: $X_t, X_{t-1}, X_{t-2}, \dots, X_0$. Studying such history-dependent processes is a formidable task due to the myriad of possible influences

from past states. Faced with this complexity, we can make a useful simplifying assumption by considering *memoryless* random processes.

A *Markov process* or *Markov chain* is a random process for which the transition probabilities for the system's next state depend only on its current state, regardless of previous states:

$$p_{X_{t+1}|X_t X_{t-1} \dots X_0}(x_{t+1}|x_t x_{t-1} \dots x_0) = p_{X_{t+1}|X_t}(x_{t+1}|x_t).$$

The system's next state X_{t+1} depends only on its current state X_t , and not on its prior history: X_{t-1} , X_{t-2} , $\dots$, X_0 . Markov chains forget everything that happened in the past, which is why we refer to them as *memoryless* random processes. A Markov chain is fully described by the conditional probability distribution $p_{X_{t+1}|X_t}(x_{t+1}|x_t)$, which describes the probability of the system's next state, given its current state. We also refer to $p_{X_{t+1}|X_t}$ as the *transition matrix*, since it describes the probabilities that the system will transition from one state to another.

Markov chains are extremely versatile models for analyzing many real-world systems. We don't have the space to cover the topic in full, but I'll introduce the basic notions for your general knowledge. First, we'll look at the connection between Markov chains and matrix multiplication. Then, we'll see how your eigenvalue-finding skills can help compute an important property of Markov chains. Understanding Markov chains is also necessary background material for understanding Google's PageRank algorithm, which we'll discuss in Section 8.3.

Example

Three friends are kicking a football in the park: Alice, Bob, and Charlie. Whenever Alice gets the ball, she passes it to Bob 40% of the time, passes it to Charlie 40% of the time, or holds onto the ball 20% of the time. Bob is kind of a greedy dude: when he gets the ball, he holds onto it 80% of the time, and is equally likely to pass it to Alice or Charlie 20% of the time. When Charlie gets the ball, he's equally likely to pass the ball to Alice or Bob, or keep it for himself. Assume these friends kick the ball around for a very long time (hundreds of passes), and you observe them at some point. What is the probability that each player will be in possession of the ball at the instant you observe them?

We can model the ball possession as a Markov process with three possible states, $\mathcal{X} = \{A, B, C\}$, where each state describes moments when Alice, Bob, or Charlie has the ball. The transition probabilities

$p_{X_{t+1}|X_t}(x_{t+1}|x_t)$ describe how the next state of the ball's possession x_{t+1} depends on the previous state of the ball's possession x_t . The transition matrix of the Markov chain in our current example is

$$p_{X_{t+1}|X_t} = \begin{bmatrix} 0.2 & 0.1 & 0.\bar{3} \\ 0.4 & 0.8 & 0.\bar{3} \\ 0.4 & 0.1 & 0.\bar{3} \end{bmatrix} = M.$$

To maintain consistency with the notation for conditional probability distributions, we refer to the entries of M as $p_{X_{t+1}|X_t}(x_{t+1}|x_t)$. The “given” variable x_t selects the column of the matrix M , and the different entries in this column represent the transition probabilities for that state. Using the matrix M and some basic linear algebra techniques, we can calculate the probability of finding the ball in any given player's possession after many iterations of the “pass the ball” Markov process.




Figure 8.1: A representation of the transition probabilities between three states: “Alice has the ball,” “Bob has the ball,” and “Charlie has the ball.”

Let's walk through an example calculation in which we assume the ball starts in Alice's possession. Since we know that Alice has the ball at $t = 0$, we can describe the initial state of the system by the probability distribution $p_{X_0} = (1, 0, 0)^T$, with 100% of the probability on Alice. We obtain the probability of finding the ball in each player's possession after one time step, or unit of time, by multiplying the initial probability vector p_{X_0} by the matrix M :

$$\begin{aligned} p_{X_1} &= Mp_{X_0} \\ &= \begin{bmatrix} 0.2 & 0.1 & 0.\bar{3} \\ 0.4 & 0.8 & 0.\bar{3} \\ 0.4 & 0.1 & 0.\bar{3} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0.2 \\ 0.4 \\ 0.4 \end{bmatrix}. \end{aligned}$$

All of the “where is the ball” probability starts with Alice, but after one time step it spreads to Bob and Charlie, according to Alice’s expected passing behaviour (the first column in the transition matrix).

We can continue this process to obtain the probability of ball possession after two time steps. We simply multiply the probability vector p_{X_1} by the transition matrix M , which is the same as multiplying the initial state p_{X_0} by M twice:

$$\begin{aligned} p_{X_2} &= MMp_{X_0} = Mp_{X_1} \\ &= \begin{bmatrix} 0.2 & 0.1 & 0.\bar{3} \\ 0.4 & 0.8 & 0.\bar{3} \\ 0.4 & 0.1 & 0.\bar{3} \end{bmatrix} \begin{bmatrix} 0.2 & 0.1 & 0.\bar{3} \\ 0.4 & 0.8 & 0.\bar{3} \\ 0.4 & 0.1 & 0.\bar{3} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0.21\bar{3} \\ 0.5\bar{3} \\ 0.25\bar{3} \end{bmatrix}. \end{aligned}$$

Let’s focus on the first entry of this vector. The probability that Alice is holding the ball after two time steps is $0.21\bar{3}$. This number represents the combination of three different “paths” of starting from the state $x_0 = A$ and coming to the state $x_2 = A$ at $t = 2$. Either Alice kept the ball for two time steps ($x_1 = A$); or Alice passed the ball to Bob ($x_1 = B$) and then he passed it back to her; or Alice passed the ball to Charlie ($x_1 = C$) and Charlie passed it back to her. By computing the sum of the products of the probabilities of the events on each path, we arrive at the value

$$\begin{aligned} p_{X_2|X_0}(A|A) &= p_{X_2|X_1}(A|A) p_{X_1|X_0}(A|A) \\ &\quad + p_{X_2|X_1}(A|B) p_{X_1|X_0}(B|A) \\ &\quad + p_{X_2|X_1}(A|C) p_{X_1|X_0}(C|A) \\ &= 0.2 \times 0.2 + 0.1 \times 0.4 + 0.\bar{3} \times 0.4 \\ &= 0.21\bar{3}. \end{aligned}$$

This is a lot of work to do manually. Compare the above calculation with the multiplication-by-the-transition-matrix approach, and the compact expression $p_{X_2} = MMp_{X_0}$ which “takes care” of these three different paths automatically: there’s no need to manually consider all possible paths—instead, just multiply the state by the transition matrix to obtain the state probabilities in the next time step.

Thanks to the power of the Markov chain construction, we can carry the ball possession probabilities vector forward through time as far as we want to. To calculate the random state $p_{X_{t+1}}$, multiply the previous state vector p_{X_t} by the Markov transition probabilities matrix M . The probability distribution of ball possession after three

time steps is

$$p_{X_3} = Mp_{X_2} = MMp_{X_1} = MMMp_{X_0} = M^3p_{X_0} = \begin{bmatrix} 0.180\bar{4} \\ 0.596\bar{4} \\ 0.223\bar{1} \end{bmatrix}.$$

Continuing this process, we can obtain the probability state vector after 10 and 20 time steps of the Markov chain:

$$p_{X_{10}} = M^{10}p_{X_0} = \begin{bmatrix} 0.161\dots \\ 0.645\dots \\ 0.193\dots \end{bmatrix} \quad \text{and} \quad p_{X_{20}} = M^{20}p_{X_0} = \begin{bmatrix} 0.1612903\dots \\ 0.6451612\dots \\ 0.193548\dots \end{bmatrix}.$$

Observe how the weights of the probability distribution seem to change less and less as the Markov chain advances. Recall we originally set out to calculate the probability that each player will be in possession of the ball after a very long time. We seem to be on the right track to find this long-term probability of ball possession.

Stationary distribution

If the evolution of a Markov chain continues for long enough, the probability vector will converge to a stable distribution p_{X_∞} that remains unchanged when multiplied by M :

$$Mp_{X_\infty} = p_{X_\infty}.$$

This is called the *stationary distribution* of the Markov chain. Observe that p_{X_∞} is an eigenvector of the matrix M with eigenvalue $\lambda = 1$.

The convergence to a unique stationary distribution is a fundamental property of Markov chains. Assuming the Markov chain represented by M satisfies some technical conditions (which we won't go into), it will converge to a stationary distribution p_{X_∞} . Thus, if we want to find p_{X_∞} , we just need to keep repeatedly multiplying by M until the distribution stabilizes:

$$p_{X_\infty} = \begin{bmatrix} 0.161290322580645\dots \\ 0.645161290322581\dots \\ 0.193548387096774\dots \end{bmatrix} = M^\infty p_{X_0}.$$

The Markov chain will converge to the same stationary distribution p_{X_∞} regardless of the starting point p_{X_0} . The ball could start with Bob $p_{X_0} = (0, 1, 0)^T$ or with Charlie $p_{X_0} = (0, 0, 1)^T$; and after running

the Markov chain for long enough, we'd still arrive at the stationary distribution p_{X_∞} .

Since we know the stationary distribution is an eigenvector of M with eigenvalue $\lambda = 1$, we can use the usual eigenvector-finding techniques to obtain p_{X_∞} directly: just solve for $\vec{v}$ in $(M - \mathbb{1})\vec{v} = \vec{0}$. This approach gives us an exact analytic expression for the answer. Note we must normalize the eigenvector $\vec{v}$ so the sum of its components equals one. The final answer is $p_{X_\infty} = \frac{\vec{v}}{\|\vec{v}\|_1} = (\frac{5}{31}, \frac{20}{31}, \frac{6}{31})^\top$, where $\|\vec{v}\|_1 = |v_1| + |v_2| + |v_3|$ is the ℓ^1 -norm of the eigenvector.

Discussion

Markov chains have countless applications in physics, speech recognition, information processing, machine learning, and many other areas. Their simple, memoryless structure and their intuitive representation as matrices make them easy to understand and easy to fit to many situations.

In the next section, we'll describe a Markov chain model for people's web browsing behaviour. The stationary distribution of this Markov chain serves to quantify the relative importance of webpages.

Links

[Awesome visual representation of states and transitions]

<http://setosa.io/blog/2014/07/26/markov-chains/index.html>

[More details about Markov chain applications from Wikipedia]

https://en.wikipedia.org/wiki/Markov_chain

Exercises

E8.3 After reading the section on Markov chains, you decide to research the subject further by checking out a book on Markov chains from the library. In the book's notation, Markov chains are represented using *right*-multiplication of the state vector: $\vec{v}' = \vec{v}B$, where $\vec{v}$ is the state of the system at time t , $\vec{v}'$ is the state at time $t + 1$, and the matrix B represents the Markov chain transition probabilities.

Find the matrix B that corresponds to the transition probabilities discussed in the Markov chain example on page 451. How is this matrix B related to the matrix M that we used in the Markov chain example?

E8.4 Go to <https://live.sympy.org> and create a `Matrix` object describing the passing-the-ball Markov chain transition probabilities. Use the `Matrix` methods `.eigenvecs()` and `.nullspace()` to confirm the stationary distribution of the Markov chain is $(\frac{5}{31}, \frac{20}{31}, \frac{6}{31})$.

Hint: You can create a `Matrix` object using

```
>>> M = Matrix([[ 2/10, 1/10, 1/3 ],
                 [ 4/10, 8/10, 1/3 ],
                 [ 4/10, 1/10, 1/3 ]])
```

If you start from a matrix with exact rational entries, you'll obtain the exact answer in terms of rational numbers. You can also use the function `S` (short for `sympify`) to create rationals: $S(1/3) = \frac{1}{3}$.

Hint: The eigenspace that corresponds to the eigenvalue $\lambda = 1$ is the null space of the matrix $(M - \mathbb{I}_3)$. Use the SymPy command `eye(3)` to create a 3×3 identity matrix $\mathbb{I}_3$, then apply the `nullspace` method.

Hint: Use `vec.norm(1)` to compute the ℓ^1 -norm of the vector `vec`.

E8.5 Find the stationary distribution of the following Markov chain:

$$C = \begin{bmatrix} 0.8 & 0.3 & 0.2 \\ 0.1 & 0.2 & 0.6 \\ 0.1 & 0.5 & 0.2 \end{bmatrix}.$$

8.3 Google's PageRank algorithm

Consider the information contained in the **links between webpages**. Each link from Page A to Page B can be interpreted as a recommendation by Page A's author for the contents of Page B. In web-speak we say links from Page A to Page B are "sending eyeballs" to Page B, presumably because there is something interesting to see on Page B. These observations about "eyeball worthiness" are the inspiration behind Google's PageRank algorithm. We find a good summary of the idea behind PageRank in Google's 2001 patent application:

A method assigns importance ranks to nodes in [...] the world wide web or any other hypermedia database. The rank assigned to a document is calculated from the ranks of documents citing it. In addition, the rank of a document is calculated from a constant representing the probability that a browser through the database will randomly jump to the document. — Patent US6285999

You may notice there is a weird self-referential thing going on in that definition. The rank of a document depends on the ranks of documents that link to it, but how do you discover the ranks of these documents? By calculating the ranks of documents that link to them, and so on and so forth. Don't worry if this sounds confusing, it will all start to make sense very soon.

The random surfer model

Imagine someone who browses the web by randomly clicking on links. We'll call this person Randy. Every time Randy opens his web browser he follows the following two strategies:

1. When visiting a webpage, he'll make a list of all the outbound links on that page and click on one of the links at random.
2. Randy randomly selects a page on the web and goes to it.

Randy follows Strategy 1 90% of the time and Strategy 2 the remaining 10% of the time. In the unlikely event that he reaches a page with no outbound links, he'll follow Strategy 2 and jump to a random page.

You have to agree that Randy's behaviour is pretty random! This is a simple model for the behaviour of users on the web that assumes people either follow links blindly, or randomly jump to webpages. Simple as it is, this model allows us to capture an important aspect of web browsing behaviour: it allows us to measure the relative importance of different webpages. In the next section, we'll describe the random surfer model using the machinery of Markov chains.

The PageRank Markov chain

We want to construct a Markov chain that models the behaviour of Randy, the random web surfer. The *state* in the Markov chain corresponds to the webpage Randy is currently viewing, and the transition matrix describes Randy's link-clicking behaviour. Modelling Randy's behaviour as a Markov chain will help us calculate the PageRank vector, which describes the "importance" of each webpage on the web.

We'll now revisit the two random strategies Randy uses to browse the web, and construct the appropriate Markov chain matrices to describe them. Note the dimension of the matrices is $n \times n$, where n is the number of webpages on the web.

1. Suppose Randy visits webpage j , which has N_j outbound links. The probability he'll click on one of these links is $\frac{1}{N_j}$. Thus, the

column for Page j should contain entries $\frac{1}{N_j}$ in the rows that correspond to the pages that Page j links to. In the special case that Page j has zero outbound links, we'll fill column j with the uniform distribution $(\frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n})^T$. Thus, we can describe the ij^{th} entry of the transition matrix for Strategy 1 as follows:

$$(M_1)_{ij} = \begin{cases} \frac{1}{N_j} & \text{if there's a link from Page } j \text{ to Page } i \\ \frac{1}{n} & \text{if Page } j \text{ has no outbound links} \\ 0 & \text{otherwise} \end{cases}$$

2. The transition matrix for Strategy 2 is much simpler. No matter which page Randy visited previously, we want him to jump to a random page on the web; thus each column of the transition matrix M_2 will contain the uniform distribution over all pages $(\frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n})^T$. We can express this succinctly as:

$$M_2 = \frac{1}{n} \mathbb{J},$$

where $\mathbb{J}$ is a matrix with ones in all its entries.

Recall that Randy uses Strategy 1 90% of the time and Strategy 2 10% of the time. The Markov chain M that describes his behaviour is a linear combination of the Markov chains for the two strategies:

$$M = (1 - \alpha)M_1 + \alpha M_2 = (1 - \alpha)M_1 + \frac{\alpha}{n} \mathbb{J},$$

where in this instance, $\alpha = 0.1$ gives us a 90-10 mix of strategies. Note that in general, α is a parameter we can adjust for different applications.

The PageRank vector is the stationary distribution of the Markov chain M , which is defined through the equation $Mp_{X_\infty} = p_{X_\infty}$. We can obtain p_{X_∞} by finding the eigenvector of M in the eigenspace $\lambda_1 = 1$, which is equivalent to solving the null space problem $(M - 1\mathbb{I})\vec{e} = \vec{0}$. We can also obtain the stationary distribution p_{X_∞} by running the Markov chain for many iterations, until we see the probability distribution converge:

$$\{p_{X_\infty}\} = \mathcal{N}(M - \mathbb{I}) \quad \Leftrightarrow \quad p_{X_\infty} = M^\infty p_{X_0}.$$

Both approaches for finding p_{X_∞} are useful in different contexts. Solving the null space problem gives us the answer directly, whereas the second approach of iterating the Markov chain is more scalable for large graphs.

The PageRank value of a webpage is the probability that Randy will navigate to a given page after running the Markov chain for a sufficiently long time. The PageRank vector p_{X_∞} captures an important aspect of the well-connectedness or importance of webpages. Let's illustrate the entire procedure for calculating the PageRank vector for the simple network shown in Figure 8.2.

Example: micro-web

We'll now study the micro-web illustrated in Figure 8.2. This is a *vastly* simplified version of the link structure between webpages on the web. Rather than include billions of webpages, the micro-web contains only eight webpages $\{1, 2, 3, 4, 5, 6, 7, 8\}$. Instead of trillions of links between webpages, the micro-web contains only fourteen links $\{(1, 2), (1, 5), (2, 3), (2, 5), (3, 1), (3, 5), (4, 5), (5, 6), (5, 7), (6, 3), (6, 7), (7, 5), (7, 6), (7, 8)\}$. Simple as it may be, this example is sufficient to illustrate the main idea of the PageRank algorithm. Scaling the solution from the case $n = 8$ to the case $n = 1\,000\,000\,000$ is left as an exercise for the reader.




Figure 8.2: A graph showing the links between the pages on the micro-web. Page 5 seems to be an important page because many pages link to it. Since Page 5 links to pages 6 and 7, these pages will probably get a lot of eyeballs, too. Page 4 is the least important, since no links lead to it. Page 8 is an example of the unlikely case of a webpage with no outbound links.

We'll first construct the transition matrix M_1 that corresponds to Strategy 1 (follow a random outbound link), and the matrix M_2 for Strategy 2 (teleport to a random page), then construct a 90-10 weighted linear combination of M_1 and M_2 , which is the PageRank Markov chain matrix.

The state of the Markov chain we want to construct corresponds to the probability distribution of Randy visiting one of the eight pages. Let's say he is currently on Page 1, so the state of the Markov

Note the new notation $\mathbb{J}_n$, which corresponds to an $n \times n$ matrix with ones in all its entries.

The PageRank Markov chain M is defined in terms of the matrices M_1 and M_2 and a mixture parameter, which we denote α . Choosing $\alpha = 0.1$ models a Randy who uses Strategy 1 90% of the time and Strategy 2 10% of the time:

$$\begin{aligned} M &= (1 - \alpha)M_1 + \frac{\alpha}{8} \mathbb{J}_8 \\ &= \frac{9}{10} M_1 + \frac{1}{80} \mathbb{J}_8. \end{aligned}$$

Or, written out in full detail:

$$M = \begin{bmatrix} \frac{1}{80} & \frac{1}{80} & \frac{37}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{8} \\ \frac{37}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{8} \\ \frac{1}{80} & \frac{37}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{37}{80} & \frac{1}{80} & \frac{1}{8} \\ \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{8} \\ \frac{37}{80} & \frac{37}{80} & \frac{37}{80} & \frac{73}{80} & \frac{1}{80} & \frac{1}{80} & \frac{5}{16} & \frac{1}{8} \\ \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{37}{80} & \frac{1}{80} & \frac{5}{16} & \frac{1}{8} \\ \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{37}{80} & \frac{37}{80} & \frac{1}{80} & \frac{1}{8} \\ \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{1}{80} & \frac{5}{16} & \frac{1}{8} \end{bmatrix}.$$

Notice the sprinkling of $\frac{1}{80}$ s that fill the regions which had only zeros in M_1 . The addition of these weights is important for technical reasons. The *mixing time* of a Markov chain is how long it takes for it to reach its steady state distribution. If we were to use just the matrix M_1 for the Markov chain, the probability distribution that describes Randy's state might get stuck in one part of the graph, and not explore all the nodes. Adding a bit of $\mathbb{J}_n$ creates connections between all the nodes, and this guarantees the Markov chain will quickly converge to its stationary distribution.

Since we built the Markov chain for the micro-web, we can now find its stationary distribution p_{X_∞} , which is the PageRank vector:

$$p_{X_\infty} = (0.08152, 0.05868, 0.1323, 0.02199, 0.2268, 0.1864, 0.2079, 0.08437)^\top.$$

We obtain p_{X_∞} by finding the vector in the $\lambda_1 = 1$ eigenspace of M . Use SymPy to check this result. See bit.ly/microWebPR for the commands required to solve the null space problem $(M - \mathbb{I})\vec{v} = \vec{0}$.

We can now use the entries of the PageRank vector to sort the pages of the micro-web by their relative importance, as shown in Table 8.2.

Page ID	PageRank
Page 5	0.22678
Page 7	0.20793
Page 6	0.18642
Page 3	0.13229
Page 8	0.08437
Page 1	0.08152
Page 2	0.05868
Page 4	0.02199

Table 8.2: The PageRanks of the pages from the graph in Figure 8.2.

According to their PageRank score, the top two pages in the micro-web are Page 5 with PageRank 0.22678 and Page 7 with PageRank 0.20793. Page 6 is not far behind with PageRank 0.18642. Looking at Figure 8.2, we can confirm this ranking makes sense, since Page 5 has the most links pointing to it, and since Page 5 links to Page 6 and Page 7. As expected, Page 4 ranks as the least important page on the micro-web since no pages link to it.

Discussion

The example we discussed above involved only eight webpages and an 8×8 Markov chain matrix. Imagine the size of the matrix needed to represent all the links between webpages on the web. The web has upward of a billion webpages and contains trillions of hyperlinks. Imagine an $n \times n$ Markov chain matrix, where $n = 1\,000\,000\,000$. You'd need a ridiculous amount of memory to store this matrix, and a huge amount of computing power to work with it.

The computation of the PageRank vector for the entire web can be performed using the *power iteration* method, which finds the eigenvector for the largest eigenvalue of the matrix. Additionally, careful analysis of the data required to perform the algorithm can avoid the need to construct the impossibly-large $n \times n$ matrix. Thus the contributions of PageRank inventors Larry Page and Sergey Brin go beyond basic linear algebra, since they offer a clever approach for splitting the computation among a cluster of computers. This is a lesson to keep in mind when you use your linear algebra knowledge to build stuff. You need to cultivate a healthy mix of intuition, mathematical tools, *and* engineering prowess to get things done. Go find the right linear combination of these resources and build your own \$300B company. I'm betting on you.

Links

[The original PageRank paper]

<http://ilpubs.stanford.edu/422/1/1999-66.pdf>

[Further discussion about the PageRank algorithm]

<https://en.wikipedia.org/wiki/PageRank>

[The *power iteration* method for finding the PageRank eigenvector]

https://en.wikipedia.org/wiki/Power_iteration

Exercises

E8.6 Compute the PageRank vector for the network of webpages shown in Figure 7.3 on page 373.

Hint: Use <https://live.sympy.org> and define the matrix M_1 using the syntax `Matrix([[...], ...])`, then mix it with an all-ones matrix `ones(n,n)`. Use the nullspace method from E8.4 to obtain the eigenvector in the $\lambda = 1$ eigenspace.

8.4 Probability problems

To better understand random variables and probability distributions, you need to practice using these concepts to solve real-world problems. It just so happens there are some practice problems on this very topic in this section—how convenient is that? Don't skip them!

Solving practice problems will help you understand probability theory and Markov chains. If you haven't played with SymPy yet, now is a great chance to get to know this powerful computer algebra system because Markov chain calculations are difficult to do by hand.

P8.1 Given a random variable X with three possible outcomes $\{1, 2, 3\}$ and probability distribution $p_X = (p_1, p_2, p_3)$, prove that $p_1 \leq 1$.

Hint: Use the Kolmogorov's axioms and build a proof by contradiction.

P8.2 The probability of heads for a fair coin is $p = \frac{1}{2}$. The probability of getting heads n times in a row is given by the expression p^n . What is the probability of getting heads four times in a row?

P8.3 You have a biased coin that lands on heads with probability p , and consequently lands on tails with probability $(1 - p)$. Suppose you want to flip the coin until you get heads. Define the random variable N as the number of tosses required until the first heads outcome. What is the probability mass function $P_N(n)$ for success on the n^{th} toss? Confirm that the formula is a valid probability distribution by showing $\sum_{n=1}^{\infty} P_N(n) = 1$.

Hint: Find the probabilities for cases $n = 1, 2, 3, \dots$ and look for a pattern.

P8.4 The probability mass function for the geometric distribution with success probability p is $p_X(x) = (1 - p)^{x-1}p$, where X describes the number of trials until the first success. Compute the expected value $\mathbb{E}_X[X]$.

Hint: The formula for the sum of the geometric series is $\sum_{k=0}^{\infty} r^k = \frac{1}{1-r}$, and taking its derivative with respect to r gives $\sum_{k=0}^{\infty} k r^{k-1} = \frac{1}{(1-r)^2}$.

P8.5 A mathematician walks over to a roulette table in a casino. The roulette wheel has 101 numbers: 50 are black, 50 are red, and the number zero is green. If the mathematician bets \$1 on black and the roulette ball stops on a black number, the payout is \$2, otherwise the bet is lost. Calculate the expected payout for playing this game, and determine whether it's worth playing.

P8.6 Consider the following variation of the six-sided die game. You pay \$1 to play one round of the game and the payout for the game is as follows. If you roll a $\square$, a $\square$, or a $\square$, you win nothing. If you roll a $\square$ or a $\square$, you win \$1. If you roll a $\square$, you win \$5. Should you play this game?

P8.7 Show that variance of a random variable X with distribution p_X is given by the formula $\text{var}(X) = \sum_{x \in \mathcal{X}} x^2 p_X(x) - \mu_X^2$.

Hint: Start from the definition $\text{var}(X) \stackrel{\text{def}}{=} \mathbb{E}_X[(X - \mu_X)^2]$ and simplify it.

P8.8 Consider the weather in a city which has “good” and “bad” years. Suppose the weather conditions over the years form a Markov chain where a good year is equally likely to be followed by a good or a bad year, while a bad year is three times as likely to be followed by a bad year as by a good year. Given that last year, call it Year 0, was a good weather year, find the probability distribution that describes the weather in Year 1, Year 2, and Year ∞ .

P8.9 Consider the network of webpages shown in Figure 8.3. Find the Markov chain transition matrices M_1 and M_2 for Randy's two browsing strategies; then combine the strategies using $\alpha = 0.1$ to obtain the PageRank matrix M . Compute the PageRank vector. Which pages are the most important?




Figure 8.3: Graph showing five webpages and the links between them.

Hint: Use SymPy for the calculation. Consult the micro-web calculation example (bit.ly/microWebPR) for the SymPy commands you'll need.

Chapter 9

Quantum mechanics

By the end of the 19th century, physicists thought they had figured out most of what there is to know about the laws of nature. Newton's laws of mechanics described the motion of objects in space, and Maxwell's equations described electricity and magnetism. Wave phenomena—including the propagation of sound, light, and waves on the surface of liquids—were also well understood. Only a few small inconsistencies between theory and experiments with atoms and radiation remained unsolved.

“[...] it seems probable that most of the grand underlying principles have now been firmly established and that further advances are to be sought chiefly in the rigorous application of these principles to all the phenomena which come under our notice.”

—Albert A. Michelson in 1894

Physicists like Michelson were worried about the future of physics research. It was as if they were wondering, “What are we going to do now that we’ve figured everything out?” Little did they know about the quantum storm that was about to hit physics, and with it, the complete rewrite of our understanding of nature at the smallest scale.

Understanding the structure of atoms—the smallest constituents of matter known at the time—was no trivial task. Describing the absorption of electromagnetic radiation by metals also turned out to be quite complicated. In both cases, the physical theories of the time predicted that the energy of physical systems could take on any value; yet experimental observations showed *discrete* energy levels. Imagine you throw a (very-very tiny) ball, and the laws of physics force you to choose an initial velocity for the ball from a list of “allowed” values: 0 m/s, 1 m/s, 2 m/s, 3 m/s, and so forth. That would

be weird, no? Weird indeed, and this is the situation physicists were facing in the beginning of the 20th century: their theories described the energy levels of atoms as real numbers $E \in \mathbb{R}$, but experiments showed that only a discrete set of energy levels exist. For example, the energy levels that the electrons of the hydrogen atom can take on are:

$$E \in \{ 21.8 \times 10^{-19} \text{ J}, 5.44 \times 10^{-19} \text{ J}, 2.42 \times 10^{-19} \text{ J}, \\ 13.6 \times 10^{-20} \text{ J}, 8.71 \times 10^{-20} \text{ J}, 6.05 \times 10^{-20} \text{ J}, \dots \}.$$

Other experimental observations suggested that electromagnetic radiation is not a continuous wave, but comes in discrete “wave packets,” which we call *photons* today. The theory of *quantum mechanics* was born out of a need to explain these observations. The term *quantum*, from the Latin *quantus* for quantity, was coined to describe the discrete nature of the phenomena that physicists were trying to explain.

During the first half of the 20th century, in experiment after experiment, quantum principles were used to correctly predict many previously-unexplained observations. During the second half of the 20th century, biologists, chemists, engineers, and physicists applied quantum principles to all areas of science. This process of “upgrading” classical models to quantum models led to a better understanding of the laws of nature, and the discovery of useful things like transistors and lasers.

The fundamental principles of quantum mechanics can be explained in the space on the back of an envelope. Understanding quantum mechanics is a matter of combining a little knowledge of linear algebra (vectors, inner products, projections) with some probability theory (Chapter 8). In this chapter, we’ll take a little excursion to the land of physics to learn about the ideas of great scientists like Bohr, Planck, Dirac, Heisenberg, and Pauli. Your linear algebra skills will allow you to learn about some fascinating 20th-century discoveries. This chapter is totally optional reading, reserved for readers who *insist* on learning about the quantum world. If you’re not interested in quantum mechanics, it’s okay to skip this chapter, but I recommend you check out Section 9.3 on *Dirac notation* for vectors and matrices. Learning Dirac notation serves as an excellent review of the core concepts of linear algebra.

9.1 Introduction

The principles of quantum mechanics have far-reaching implications for many areas of science: physics, chemistry, biology, engineering,

philosophy, and many other fields of study. Each field of study has its own view on quantum mechanics, and has developed a specialized language for describing quantum concepts. We'll formally introduce the postulates of quantum mechanics in Section 9.5, but before we get there, let's look at some of the disciplines where quantum principles are used.

Physics Physicists use the laws of quantum mechanics as a toolbox to understand and predict the outcomes of atomic-scale physics experiments. By “upgrading” classical physics models to reflect the ideas of quantum mechanics, physicists (and chemists) obtain more accurate models that lead to better predictions.

For example, in a *classical* physics model, the motion of a particle is described by its position $x(t)$ and velocity $v(t)$ as functions of time:

$$\text{classical state} = (x(t), v(t)), \text{ for all times } t.$$

At any given time t , the particle is at position $x(t)$ and moving with velocity $v(t)$. Using Newton's laws of motion and calculus, we can predict the position and the velocity of a particle at all times.

In a quantum description of the motion of a particle in one dimension, the state of a particle is represented by a *wave function* $|\psi(x, t)\rangle$, which is a complex-valued function of position x and time t :

$$\text{quantum state} = |\psi(x, t)\rangle, \text{ for all times } t.$$

At any given time t , the state of the particle corresponds to a complex-valued function of a real variable $|\psi(x)\rangle \in \{\mathbb{R} \rightarrow \mathbb{C}\}$. The wave function $|\psi(x)\rangle$ is also called the *probability-amplitude* function. The probability of finding the particle at position x_a is proportional to the value of the squared norm of the wave function:

$$\Pr(\{\text{particle position} = x_a\}) \propto ||\psi(x_a)\rangle|^2.$$

Instead of having a definite position $x(t)$ as in the classical model, the position of the particle in a quantum model is described by a probability distribution calculated from its wave function $|\psi(x)\rangle$. Instead of having a definite momentum $p(t)$, the momentum of a quantum particle is another function calculated based on its wave function $|\psi(x)\rangle$.

Classical models provide accurate predictions for physics problems involving macroscopic objects, but fail to predict the physics of atomic-scale phenomena. Much of 20th-century physics research efforts were dedicated to the study of quantum concepts like ground states, measurements, spin angular momentum, polarization, uncertainty, entanglement, and non-locality.

Computer science Computer scientists understand quantum mechanics using principles of information. Quantum principles impose a fundamental change to the “data types” used to represent information. Classical information is represented as *bits*, elements of the finite field of size two $\mathbb{Z}_2$:

$$\text{bit: } x = 0 \text{ or } x = 1.$$

In the quantum world, the fundamental unit of information is the *qubit*, which is a two-dimensional unit vector in a complex inner product space:

$$\text{qubit: } |x\rangle = \alpha|0\rangle + \beta|1\rangle.$$

This change to the underlying information model requires reconsidering fundamental information processing tasks like computation, data compression, encryption, and communication.

Philosophy Philosophers have also updated their conceptions of the world to incorporate the laws of quantum mechanics. Observations of physics experiments forced them to reconsider the fundamental question, “What are things made of?” Another interesting question philosophers have considered is whether the quantum state $|\psi\rangle$ of a physical system really exists, or if $|\psi\rangle$ is a representation of our knowledge about the system.

A third central philosophy concept that quantum mechanics calls into question is *determinism*—the clockwork-model of the universe, where each effect has a cause we can trace, like the connections between gears in a mechanical clock. The laws of physics tell us that the next state of the universe is determined by the current state of the universe, and the state changes according to the equations of physics. However, representing the universe as a quantum state has implications for our understanding of how the universe “ticks.” Clockwork (deterministic) models of the universe are not wrong—they just require a quantum upgrade.

Many scientists are also interested in the philosophical aspects of quantum mechanics. Physicists call these types of questions *foundations* or *interpretations*. Since different philosophical interpretations of quantum phenomena cannot be tested experimentally, these questions are considered outside the scope of physics research. Nevertheless, these questions are so deep and fascinating that physicists continue to pursue them, and contribute interesting philosophical work.

[Philosophical issues in quantum theory]

<http://plato.stanford.edu/entries/qt-issues/>

Physical models of the world

We'll situate our discussion of quantum mechanics within two conceptual worlds:

- The **real world** is where physical experiments are performed.
- The **mathematical world** is a purely theoretical construct that aims to model certain aspects of the real world.

The better the mathematical model, the more closely its predictions correspond to the behaviour of real-world systems. Sill, no math or physics model can ever predict real-world outcomes with 100% accuracy. When we say that a certain mathematical model is better than another, we mean it can predict the outcomes of controlled experiments with greater accuracy. Physicists are very open to new theories—anyone can be a physicist! You can start by constructing any crazy mathematical model for describing nature, and if your model correctly predicts the outcomes of experiments, other physicists will start using it.

We can make a further distinction among mathematical models, classifying them into two categories depending on the type of math they use:

- **Classical models** describe the world in terms of real variables like positions and velocities.
- **Quantum models** describe systems in terms of vectors in complex vector spaces.

Table 9.1 compares the objects used in the two types of mathematical models of the real world. In physics, classical models describe the motion of particles using trajectories $\vec{r}(t)$, whereas quantum models use wave functions $|\psi(\vec{r}, t)\rangle$. In computer science, classical information is stored in bits $i \in \{0, 1\}$, whereas quantum information is stored in qubits $|x\rangle \in \mathbb{C}^2$.

Example Let's analyze the difference between classical and quantum models of the real world using an example. Consider a photon (a particle of light) going through an optical circuit that consists of several lenses, mirrors, and other optical instruments. A photon detector is placed at position x_f at the end of the circuit. The objective of the experiment is to predict if the photon will arrive at the detector and cause it to “click.” The two possible outcomes of the experiment are `click` (photon arrives at detector) or `noclick` (photon doesn't arrive at detector).¹

¹We're assuming the detector has 100% efficiency (detects every photon that arrives at it) and generates zero noise (no false-positive clicks).

Real world:

- The motion of a ball thrown in the air
- The motion of an electron through space
- The paths of light particles moving through optical circuits
- The electric current flowing through a superconducting loop

Classical models:

- $x(t) \in \{\mathbb{R} \rightarrow \mathbb{R}\}$
- $\vec{r}(t) \in \{\mathbb{R} \rightarrow \mathbb{R}^3\}$
- $i \in \mathbb{Z}_2 = \{0, 1\}$ (a bit)
- $j \in \mathbb{Z}_d$

Quantum models:

- $|\psi(x, t)\rangle \in \{\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}\}$
- $|\psi(\vec{r}, t)\rangle \in \{\mathbb{R}^3 \times \mathbb{R} \rightarrow \mathbb{C}\}$
- $|x\rangle \in \mathbb{C}^2$ (a qubit)
- $|y\rangle \in \mathbb{C}^d$

Table 9.1: Examples of the math used in classical and quantum models.

A classical model of the motion of the photon calculates the photon's position at all times $x(t)$ and leads to the prediction $i = 1$ (click) if $x_f = x(t)$, for some t . On the other hand, if the detector does not lie on the photon's trajectory, then the classical model will predict $i = 0$ (noclick).

A quantum model would describe the photon's trajectory through the circuit as a linear combination of two different possible paths:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \quad \text{where} \quad |\alpha|^2 + |\beta|^2 = 1.$$

Here $|1\rangle$ describes paths that arrive at the detector, and $|0\rangle$ describes paths that don't. The coefficients α and β describe the relative "weights" of the different paths. Using the quantum model, we can obtain a probabilistic prediction of whether the detector will click or not:

$$\text{Pr}(\text{noclick}) = |\alpha|^2 \quad \text{and} \quad \text{Pr}(\text{click}) = |\beta|^2.$$

For this example, both the classical and the quantum models describe the same real-world phenomenon. We can test the validity of both models by comparing the models' predictions with what happens in reality.

Note that the two models make very different assumptions about reality. The classical model assumes the photon follows a single path through the circuit, whereas the quantum model assumes the photon can take multiple paths through the circuit. Despite the difference in the mathematical substrate of the models and their fundamentally different views of reality, we can compare the two models' predictions on the same footing. Note it doesn't make sense to say one

model is more real than the other. The only thing that is real is the photon in the optical circuit, and the photon doesn't care whether you use classical or quantum models to describe its path.

Quantum model peculiarities

We'll now comment on the relative “intuitiveness” of classical and quantum models and introduce the concept of *quantum measurement*, which is of central importance in quantum mechanics.

Classical models have the advantage of being more intuitively understandable than quantum models. The variables in classical models often correspond to measurable aspects of real-world systems. We can identify the position variable in a classical model with the position of a particle in the real world. Velocity and momentum are harder to understand intuitively, but we have some general intuition about motion and collisions from everyday life. In general, we can understand classical models more readily because it's easier for us to think about a mechanistic, clockwork-like universe, in which objects push on each other with clearly defined cause and effect, like a clock that goes click, click, click.

In contrast, we do not enjoy such intuitive interpretation of quantum models, since our senses cannot directly perceive movement and interaction at the quantum level. Because quantum models' states do not directly correspond to observable aspects in the real world, quantum models are often described as mysterious and counterintuitive. Quantum models are harder to understand in part because they use complex vector quantities to represent systems, and complex numbers are more difficult to visualize. For example, visualizing the complex-valued state of a photon $|\psi\rangle$ is difficult, since you must think about both the real part and the imaginary part of $|\psi\rangle$. Even though we can't see what $|\psi\rangle$ looks like, we can describe it using an equation, and do math calculations with it. In particular, we can compare the predictions obtained from calculations based on the quantum state $|\psi\rangle$ to measurements performed in the real world.

The process of *quantum measurement* is how we map the predictions of the quantum model to observable quantities. A quantum measurement acts on a particle's wave function $|\psi\rangle$ to produce a classical outcome. **Performing quantum measurements is like asking questions of particles, and the measurement outcomes are the answers to these questions.**

What is your position? $\Leftrightarrow$ position($|\psi\rangle$) = $x \in \mathbb{R}$

What is your momentum? $\Leftrightarrow$ momentum($|\psi\rangle$) = $p \in \mathbb{R}$

What is your spin momentum? $\Leftrightarrow$ spin _{$\uparrow\downarrow$} ($|\psi\rangle$) = $s \in \{\uparrow, \downarrow\}$

Since measurement outcomes correspond to real-world quantities that can be measured, we can judge the merits of quantum models the same way we judge the merits of classical models—in terms of the quality of their predictions.

Chapter overview

In the next section, we'll describe a tabletop experiment involving lasers and polarization lenses, with an outcome that's difficult to explain using classical physics. The remainder of the chapter will introduce the tools needed to explain the outcome of this experiment in terms of quantum mechanics. We'll start by introducing a special notation for vectors that is used to describe quantum phenomena (Section 9.3).

In Section 9.5, we'll formally define the “rules” of quantum mechanics, also known as the *postulates* of quantum mechanics. We'll learn the “rules of the game” using the simplest possible quantum systems (qubits), and define how quantum systems are prepared, how we manipulate them using *quantum operations*, and how we extract information from them using *quantum measurements*. This part of the chapter is based on the notes from the introductory lectures of a graduate-level quantum information course, so don't think you'll be getting some watered-down, hand-wavy version of quantum mechanics. You'll learn the real stuff, because I know you can handle it.

In Section 9.6 we'll apply the quantum formalism to the polarizing lenses experiment, showing that a quantum model leads to the correct qualitative and quantitative prediction for the observed outcome. We'll close the chapter with short explanations of different applications of quantum mechanics with pointers for further exploration about each topic.

Throughout the chapter, we'll focus on *matrix* quantum mechanics and use computer science language to describe quantum phenomena. A computer science approach allows us to discuss the fundamental aspects of quantum theory without introducing all the physics required to understand atoms. Finally, I just might throw in a sample calculation using the wave function of the hydrogen atom, to give you an idea of what that's like.

9.2 Polarizing lenses experiment

Let's run through a simple tabletop experiment that illustrates the limitations of classical, deterministic reasoning. The outcome of the

experiment will highlight the need for careful consideration of the measurements used in scientific experiments.

We'll describe the experiment using words and diagrams, but you can easily reproduce the experiment in your own "lab," since it requires only simple equipment. I encourage you to try it yourself. You'll need three polarizing lenses, a laser pointer, a piece of paper, and three binder clips for holding the lenses upright. You can buy polarizing lenses on the cheap from a second-hand camera shop—any polarizing lens will do.

Background

In photography, polarizing lenses are used to filter out undesirable light reflections, like reflections that occur from water surfaces or glass windows. To better understand the experiment, we need to introduce some basic notions about the physics of light, specifically the concept of light polarization.

Light consists of photons. Photons are travelling pulses of electromagnetic energy. Electromagnetic energy can travel through space in the form of a wave. Polarization refers to the orientation of the electric field $\vec{E}$ of a propagating electromagnetic wave.

Light is normally unpolarized, meaning it corresponds to a mixture of photons that have electric and magnetic components of random orientation. A light beam is *polarized* if all its photons have the same orientation of their electric field.




Figure 9.1: A photon is a pulse of electromagnetic energy. The energy of a photon travels in the form of a wave that has an electric component $\vec{E}$ and a magnetic component $\vec{B}$. The figure shows a photon travelling in the positive x -direction with its electric component along the z -axis.

Light reflected from flat surfaces, like the surface of a lake or a glass window, becomes polarized, which means the electric components of all the reflected photons become aligned.

Photographers use this fact to selectively filter out light with a particular polarization. A *polarizing filter* or *polarizing lens* has a special coating which conducts electricity in one direction, but not in the other. You can think of a polarizing lens as a surface covered by

tiny conductive bands that interact with the electric component of incoming light particles. Light rays that hit the polarizing lens will either pass through or be reflected depending on their polarization. Light particles with a polarization perpendicular to the conductive bands pass through the lens, while light particles with polarization parallel to the conductive bands are reflected. This is because the surface of the lens has different conductive properties in the parallel and perpendicular directions.




Figure 9.2: Incoming photons interact with the horizontal conductive bands of a polarizing lens. The horizontal bands of the lens reflect the horizontal component of the photons's electric field. Vertically-polarized photons pass through the lens because the conductive bands are perpendicular to their electric field. Thus, a vertically polarizing lens denoted V allows only vertically polarized light to pass through.

Consider the illustration in Figure 9.3. The effect of a vertically polarizing lens on a beam of light is to only allow vertically polarized light to pass through.




Figure 9.3: A vertically polarizing lens (V) allows only vertically polarized light particles to pass through.

In Figure 9.4 we see another aspect of polarizing lenses. If the light is already vertically polarized, adding a second vertically polarizing lens will not affect the beam. All light that passes through the first lens will also pass through the second.

Taking a vertically polarizing lens and rotating it by 90 degrees turns it into a horizontally polarizing lens. See Figure 9.5.

Note that horizontally polarizing lenses and vertically polarizing lenses are complementary: vertically polarized light will not pass




Figure 9.4: A second vertically polarizing lens has no further effect since light is already vertically polarized.




Figure 9.5: A horizontally polarizing lens (H) allows only horizontally polarized light particles to pass through.

through a horizontally polarizing lens. This situation is illustrated in Figure 9.6.




Figure 9.6: Placing a horizontally polarizing lens after the vertically polarizing lens has the effect of filtering all light. Zero photons make it through both filters, which we indicate with the empty set symbol $\emptyset$.

The previous examples can familiarize you with the properties of polarizing lenses, in case you don't have actual lenses to play with. If you do have polarizing lenses, you can shine a laser pointer through them, observing when light passes through and when light is filtered out. Use binder clips to position the lenses on a flat surface, and reproduce the setup in Figure 9.4. Don't worry about finding the exact orientation for "vertical." Any orientation of the lenses will do, as long as the first and the second polarizing lens have the same orientation. Next, you can rotate the second lens by 90° to obtain the setup shown in Figure 9.6, where the second lens has a perpendicular orientation and thus rejects all light.

Example Polarized sunglasses leverage the properties of light polarization to make outdoor activities more enjoyable. When a beam of light bounces off the surface of a lake, it becomes horizontally polarized. This polarization effect is due to the interaction of light's electric field at the surface of the water. A person wearing vertically polarizing lenses (polarized sunglasses) cannot see the sun's reflection off the water surface because the V-polarizing lenses filter out the horizontally polarized light reflected from the surface of the lake. This effect is useful for people who are often outdoors in bright sunlight, as it reduces the blinding effect of the sun's reflection.

Classical physics paradigm

Before we describe the outcome of the polarizing lenses experiment, let's take a moment to describe the assumptions about the world that 19th-century physicists held. Understanding this classical world view will explain why the outcomes of the polarizing lenses experiment are so surprising.

The classical laws of physics are deterministic, meaning they do not allow randomness. According to the classical school, experimental outcomes depend on *definite* variables, like the properties of particles. Physicists assume they can predict the outcome of any experiment given they know the properties of the particles involved. If some outcome cannot be predicted, it must be because the value of some property of the particles is unknown. In other words, everything happens for a reason. Another key assumption classical physicists make is that a photon's properties are immutable, meaning we cannot change them. Classical physicists assume their **experiments correspond to passive observations** that cannot change the system's properties.

A 19th-century physicist would expect the outcomes of polarizing lenses experiments to be fully determined by the polarization property of photons. Each photon carries a tag H or V that describes its polarization. In the setup shown in Figure 9.3, each photon that passes through the lens must have tag="V", because we know by definition that a V-polarizing lens only allows vertically polarized photons to pass through. Readers familiar with SQL syntax will recognize the action of the vertically polarizing lens as the following query:

```
SELECT photon FROM photons WHERE tag="V";
```

In other words, out of all the incoming photons, only the vertically polarized photons pass through the lens. Similarly, for the

H-polarizing lens shown in Figure 9.5, the filtering process can be understood as the query:

```
SELECT photon FROM photons WHERE tag="H";
```

In both cases, classical physicists would assume that whether or not a photon passes through a lens is predetermined, and is dependent only on the photon's tag.

For the purpose of our discussion, we'll restrict our attention to photons with either horizontal ($\text{tag}=\text{"H"}$) or vertical ($\text{tag}=\text{"V"}$) polarization. There are other possible polarization directions, but we'll focus on the tags *H* and *V* because they're mutually exclusive—if a photon is horizontally polarized, we know a vertically polarizing lens will reject it. We can assert that photons that pass through an *H*-polarizing lens are not vertically polarized; if they were, these photons would reflect off the lens instead of passing through.

Polarizing lenses experiment

The physics experiment we'll describe consists of sending photons through an optical circuit made of polarizing lenses and observing how many photons pass through the circuit. We describe the number of photons that reach any point in the circuit in qualitative terms, by referring to the *optical power* reaching that point, denoted P . We choose the light intensity of the beam after it passes through the first polarizing lens as our reference, and call it $P = 1$ (full brightness). You can think of optical power as brightness. If you were to insert a piece of paper somewhere in the optical circuit, the optical power would measure the brightness of the spot of light on the paper at that particular location in the circuit. When power $P = 1$, the spot of light is fully bright. If $P = 0.5$ the spot is half as bright as when $P = 1$. The case $P = 0$ corresponds to zero brightness and occurs when no photons hit the piece of paper.

The initial setup for the experiment consists of an *H*-polarizing lens followed by a *V*-polarizing lens, as shown in Figure 9.7.

We know the photons that pass through the first lens are horizontally polarized. It's no surprise that when this light hits the second lens, none of the photons make it through, since a *V*-polarizing lens rejects *H*-polarized photons.

Adding a third lens We now introduce a third lens between the first two lenses, and we orient the middle lens differently from the other two—in the diagonal direction, for example. The result is shown in Figure 9.8. Suddenly, light appears at the end of the




Figure 9.7: The initial setup for the polarizing lenses experiment consists of an H -polarizing lens followed by a V -polarizing lens. Only photons with tag="H" can pass through the first lens, so no photons with tag="V" pass through the first lens. No photons can pass through both lenses since the V -polarizing lens accepts only photons with tag="V".

circuit! How does this make sense? You tell me if this is crazy or not. Intuitively, adding more filtering only reduces the amount of light passing through the circuit; yet the amount of light that passes through the circuit increases when we add the middle filter. How can adding more filtering increase light intensity? What is going on?




Figure 9.8: Adding an additional polarizing lens in the middle of the circuit causes light to appear at the end of the optical circuit.

We pick the middle lens to be a diagonally polarizing lens D . A diagonally polarizing lens is obtained by rotating any polarizing lens by 45° . The exact choice for the middle lens is not crucial for the experiment to work; so long as its polarization is different from the H - and V -polarizing lenses that surround it.

For a demonstration of the three-polarizing lenses experiment, see this Youtube video: <https://youtu.be/PJHCADY-Bio?t=6m14s>.

Classical analysis

The experimental observations illustrated in Figure 9.8 are difficult to explain using the classical way of thinking, in which particle properties are immutable tags, and experiments are passive observations.

We have evidence that a particle's state can change during the process of measurement. Let's examine this more closely. We'll trace the path of the photons through the optical circuit in Figure 9.8, keeping track of what we know about the photons at each stage. First, all

photons that pass through the first H -polarizing lens are known to be horizontally polarized (tag="H"). We're sure no V -polarized photons pass through the first lens, because an H -polarizing lens is guaranteed to reject all V -polarized photons. Yet, after passing through the second lens (the D lens), these same photons appear to be vertically polarized since they pass through the third lens. Is something wrong with the tagging system? It seems the photons' tag states are affected by the measurements performed on them. This fact is difficult to explain for classical physicists since they assume measurements correspond to passive observations. In the classical paradigm, measuring a photon's D -polarization using the middle lens should not affect its H and V tags.

In Section 9.5 we'll revisit the polarizing lenses experiment after we have learned the postulates of quantum mechanics. We'll see we can explain the outcome of the experiment by describing the photon's polarization in terms of vector-like states; and that our understanding of vectors can even predict the final light intensity of $P = 0.25$ we observed during the experiment. Before we discuss the postulates of quantum mechanics (Section 9.5), we'll need to introduce some new notation for describing quantum states.

9.3 Dirac notation for vectors

This section is a quick primer on Dirac notation, which is a precise and concise language for talking about vectors and matrices. The Dirac notation for vectors $|v\rangle$ is an alternative to the usual notations for vectors like $\vec{v}$ or $\mathbf{v}$.

This new notation will look weird initially, but once you get the hang of it, I guarantee you'll like it. Learning Dirac notation is an excellent way to review essential linear algebra concepts like bases, vectors, inner products, and matrices. Understanding Dirac notation is essential if you're interested in learning quantum mechanics, but it's also worth learning "just because"—it's *that* cool.

We'll now discuss several vector topics you're familiar with, and compare standard vector notation $\vec{v}$ with the equivalent Dirac notation $|v\rangle$.

The standard basis

Consider a d -dimensional complex vector space $\mathbb{C}^d$. We refer to complex vector spaces as Hilbert spaces, in honour of David Hilbert, who contributed prolific developments to math and physics.

To understand any vector space, it is essential to construct a basis for the space. A natural choice for a basis is the standard basis, which we'll denote $\{|0\rangle, |1\rangle, |2\rangle, \dots, |d-1\rangle\}$. The basis vectors are defined as:

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad |1\rangle = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix}, \quad \dots, \quad |d-1\rangle = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}.$$

Note the indices are shifted by one so the first basis vector has index 0, not index 1. This zero-based indexing is chosen to make certain links between quantum theory and computer science more apparent.

One benefit of Dirac notation is that it doesn't require writing subscripts. To refer to a vector associated with properties a , b , and c , we can write $|a, b, c\rangle$, instead of the more convoluted expression $\vec{v}_{a,b,c}$.

We'll now focus solely on the two-dimensional complex vector space $\mathbb{C}^2$; however, the results and definitions presented below also apply to vectors of any dimension.

Vectors

In Dirac notation, a vector in $\mathbb{C}^2$ is denoted as a *ket*:

$$|v\rangle = \alpha|0\rangle + \beta|1\rangle \quad \Leftrightarrow \quad \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \alpha \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \beta \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

where $\alpha \in \mathbb{C}$ and $\beta \in \mathbb{C}$ are the *components* of $|v\rangle$ and $\{|0\rangle, |1\rangle\}$ is the standard basis for $\mathbb{C}^2$:

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

Why do we call the angle-bracket thing a “ket,” you ask? Let me tell you about the *bra* part, and then it will start to make sense.

The Hermitian transpose of the ket-vector $|v\rangle = \alpha|0\rangle + \beta|1\rangle$ is the *bra*-vector $\langle v|$:

$$\langle v| = \bar{\alpha}\langle 0| + \bar{\beta}\langle 1| \quad \Leftrightarrow \quad [\bar{\alpha}, \bar{\beta}] = \bar{\alpha}[1, 0] + \bar{\beta}[0, 1].$$

Recall that the Hermitian transpose, also called the complex transpose, is the combination of the regular transpose ($\vec{v} \rightarrow \vec{v}^T$) and the complex conjugation of each component of the vector ($v_i \rightarrow \bar{v}_i$), and is denoted as the dagger operator “ $\dagger$ ” (see page 340 for a refresher on the dagger operator). Now observe how much simpler the *bra*

notation for the Hermitian transpose of a vector is compared to the other notations we've seen so far:

$$\langle v| \quad \Leftrightarrow \quad \vec{v}^\dagger = \overline{(\vec{v}^\top)} = (\vec{v})^\top.$$

When you put a bra $\langle v|$ next to a ket $|w\rangle$, they become a bra-ket $\langle v|w\rangle$, which looks very similar to the brackets used to denote the inner product between two vectors. Observe the relative notational simplicity of calculating the inner product between the vectors $|v\rangle = v_0|0\rangle + v_1|1\rangle$ and $|w\rangle = w_0|0\rangle + w_1|1\rangle$ in Dirac notation:

$$\langle v|w\rangle = \overline{v_0}w_0 + \overline{v_1}w_1 \quad \Leftrightarrow \quad \vec{v}^\dagger \vec{w} = \vec{v} \cdot \vec{w}.$$

Complex conjugation was already applied to the coordinates of $|v\rangle$ when transforming it into a bra-vector $\langle v|$, thus we can simply “put together” the bra and the ket to compute the inner product. Because the bra notation $\langle v|$ contains the Hermitian transpose, it removes the need for the dagger symbol. For instance, the inner product of $|v\rangle = \alpha|0\rangle + \beta|1\rangle$ with itself is

$$\begin{aligned} \langle v|v\rangle &= (\overline{\alpha}\langle 0| + \overline{\beta}\langle 1|, \alpha|0\rangle + \beta|1\rangle) \\ &= \overline{\alpha}\alpha\langle 0|0\rangle + \overline{\alpha}\beta \underbrace{\langle 0|1\rangle}_0 + \overline{\beta}\alpha \underbrace{\langle 1|0\rangle}_0 + \overline{\beta}\beta\langle 1|1\rangle \\ &= |\alpha|^2 + |\beta|^2. \end{aligned}$$

This small simplification of the notation for inner products is quite useful, since inner products are the workhorse for calculating vector coordinates, finding projections, and performing change-of-basis transformations. The concise notation enables us to dig deeper into these aspects of linear algebra without getting overwhelmed by notational complexity. Moreover, Dirac's bra-ket notation is sufficiently simple to use in equations without needing to define new variables v_i for vector coordinates. We'll look at this more closely in the next section.

Vector coordinates

The *coordinates* v_i of a vector $\vec{v}$ with respect to an orthonormal basis $\{\hat{e}_i\}$ are computed using the inner product $v_i = \hat{e}_i^\dagger \vec{v}$. In Dirac notation, the coordinates of $|v\rangle$ with respect to the standard basis $\{|0\rangle, |1\rangle\}$ can be written as $\langle i|v\rangle$. We can write any vector $|v\rangle$ as a linear combination of kets, with bra-kets as coefficients:

$$|v\rangle = \underbrace{\langle 0|v\rangle}_{v_0} |0\rangle + \underbrace{\langle 1|v\rangle}_{v_1} |1\rangle.$$

The expression $\langle i|v\rangle$ explicitly defines the i^{th} coordinate of $|v\rangle$; therefore, we don't need to define the variable v_i .

Another basis for the vector space $\mathbb{C}^2$ is the *Hadamard basis*, which corresponds to the standard basis rotated by 45° in the counter-clockwise direction:

$$\begin{aligned}|+\rangle &= \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle, \\ |-\rangle &= \frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle.\end{aligned}$$

The Hadamard basis, henceforth denoted $B_h = \{|+\rangle, |-\rangle\}$, is an orthonormal basis:

$$\langle +|+\rangle = 1, \quad \langle +|-\rangle = 0, \quad \langle -|+\rangle = 0, \quad \langle -|-\rangle = 1.$$

Since the Hadamard basis B_h is an orthonormal basis, we can express any vector $|v\rangle \in \mathbb{C}^2$ as a linear combination of kets:

$$|v\rangle = \underbrace{\langle +|v\rangle}_{v_+} |+\rangle + \underbrace{\langle -|v\rangle}_{v_-} |-\rangle.$$

Note that the coefficients of the linear combination are computed using the inner product with the corresponding basis vector $\langle +|v\rangle$ and $\langle -|v\rangle$. The bra-ket notation allows us to refer to the coordinates of $|v\rangle$ with respect to $\{|+\rangle, |-\rangle\}$, without the need to define variables v_+ and v_- .

It's not uncommon to see vector coordinates with respect to different bases used in calculations. With the usual vector notation, we must specify which basis we're using as a subscript. For example, the same vector $\vec{v}$ can be expressed as a coordinate vector $\vec{v} = (v_0, v_1)_{B_s}^T$ with respect to the standard basis B_s , or as a coordinate vector $\vec{v} = (v_+, v_-)_{B_h}^T$ with respect to the Hadamard basis. In the bra-ket notation, the coordinates with respect to B_s are $\langle 0|v\rangle$ and $\langle 1|v\rangle$, and the coordinate with respect to B_h are $\langle +|v\rangle$ and $\langle -|v\rangle$, making the choice of basis evident.

Change of basis

Consider the task of finding the *change-of-basis* matrix $_{B_h}[\mathbb{1}]_{B_s}$ that converts vectors from the standard basis $B_s = \{(1,0), (0,1)\}$ to the Hadamard basis $B_h = \{(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}), (\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}})\}$.

Using the standard approach for finding change-of-basis matrices discussed in Section 4.3 (page 229), we know the columns of $_{B_h}[\mathbb{1}]_{B_s}$ contain the coordinates of $(1,0)$ and $(0,1)$ as expressed with respect

the basis B_h :

$${}_{B_h}[\mathbb{1}]_{B_s} = {}_{B_h} \begin{bmatrix} \langle +|0\rangle & \langle +|1\rangle \\ \langle -|0\rangle & \langle -|1\rangle \end{bmatrix} = {}_{B_h} \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}_{B_s}.$$

We can now compute the coordinates of any vector $\vec{v} = (v_0, v_1)_{B_s}^T$ with respect to the Hadamard basis, by multiplying $(v_0, v_1)_{B_s}^T$ by the change-of-basis matrix:

$$\begin{aligned} \begin{bmatrix} v_+ \\ v_- \end{bmatrix}_{B_h} &= {}_{B_h}[\mathbb{1}]_{B_s} \begin{bmatrix} v_0 \\ v_1 \end{bmatrix}_{B_s} \\ &= \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}_{B_h} \begin{bmatrix} v_0 \\ v_1 \end{bmatrix}_{B_s} \\ &= \begin{bmatrix} \frac{1}{\sqrt{2}}(v_0 + v_1) \\ \frac{1}{\sqrt{2}}(v_0 - v_1) \end{bmatrix}_{B_h}. \end{aligned}$$

This is the usual approach for computing the coordinates of a vector in one basis in terms of the coordinates of another basis using matrix notation.

Consider now the same change-of-basis operation, but with calculations carried out in Dirac notation. Given the vector $\vec{v} = (v_0, v_1)_{B_s} = v_0|0\rangle + v_1|1\rangle$ expressed as coordinates with respect to the standard basis B_s , we want to find $\vec{v} = (v_+, v_-)_{B_h} = v_+|+\rangle + v_-|-\rangle$. Starting from the definition of v_+ and v_- , we obtain

$$\begin{aligned} \vec{v} &= (v_+, v_-)_{B_h}^T \\ &= \langle +|v\rangle|+\rangle + \langle -|v\rangle|-\rangle \\ &= \langle +|(v_0|0\rangle + v_1|1\rangle)|+\rangle + \langle -|(v_0|0\rangle + v_1|1\rangle)|-\rangle \\ &= (v_0\langle +|0\rangle + v_1\langle +|1\rangle)|+\rangle + (v_0\langle -|0\rangle + v_1\langle -|1\rangle)|-\rangle \\ &= \underbrace{\frac{1}{\sqrt{2}}(v_0 + v_1)}_{v_+}|+\rangle + \underbrace{\frac{1}{\sqrt{2}}(v_0 - v_1)}_{v_-}|-\rangle. \end{aligned}$$

Working from the definitions of vectors and their coordinates, and using only basic algebra rules, we can perform the change-of-basis operation without explicitly constructing the change-of-basis matrix.

Outer products

Recall the *outer product* operation for vectors that we introduced in Section 5.2 (see page 272). The expression $|0\rangle\langle 0|$ is equivalent to the

projection onto the subspace spanned by the vector $|0\rangle$:

$$|0\rangle\langle 0| \quad \Leftrightarrow \quad \begin{bmatrix} 1 \\ 0 \end{bmatrix} \begin{bmatrix} 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}.$$

We can verify this by considering the product of $|0\rangle\langle 0|$ and an arbitrary vector $|v\rangle = \alpha|0\rangle + \beta|1\rangle$:

$$|0\rangle\langle 0| (\alpha|0\rangle + \beta|1\rangle) = \alpha|0\rangle \underbrace{\langle 0|0\rangle}_1 + \beta|0\rangle \underbrace{\langle 0|1\rangle}_0 = \alpha|0\rangle.$$

The ability to easily express outer products is another win for Dirac notation. For example, the projection onto the $|+\rangle$ -subspace is $|+\rangle\langle +|$.

Matrices

Now get ready for some crazy stuff. It turns out outer-product expressions are useful not only for projections, but can in fact represent any matrix. Consider the linear operator $A : \mathbb{C}^2 \rightarrow \mathbb{C}^2$ and its matrix representation with respect to the standard basis:

$${}_{B_s}[A]_{B_s} = {}_{B_s} \begin{bmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \end{bmatrix}_{B_s}.$$

Instead of positioning the entries in an array, we can represent A as a linear combination of outer products:

$$A = a_{00}|0\rangle\langle 0| + a_{01}|0\rangle\langle 1| + a_{10}|1\rangle\langle 0| + a_{11}|1\rangle\langle 1|.$$

Consider the matrix vector product $A|x\rangle = |y\rangle$. The matrix entry a_{10} describes the multiplication factor that connects the $|0\rangle$ -component in the input $|x\rangle$ to the $|1\rangle$ -component of the output vector $|y\rangle$. The expression $a_{10}|1\rangle\langle 0|$ is a concise description of the same story. The $\langle 0|$ in this expression will “select” only the $|0\rangle$ component of the input, and the $|1\rangle$ indicates that this term contributes to the $|1\rangle$ component of the output.

The entries of the matrix representation ${}_{B_s}[A]_{B_s}$ depend on the choice of bases for the input and output spaces. The value of the entry a_{ij} in the matrix representation is computed by “probing” the matrix with the j^{th} basis vector of the input basis, and observing the value of the i^{th} component in the resulting output. We can express the entire “probing procedure” easily in Dirac notation: $a_{ij} = \langle i|A|j\rangle$. Thus, we can write the matrix entries as follows:

$${}_{B_s} \begin{bmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \end{bmatrix}_{B_s} = {}_{B_s} \begin{bmatrix} \langle 0|A|0\rangle & \langle 0|A|1\rangle \\ \langle 1|A|0\rangle & \langle 1|A|1\rangle \end{bmatrix}_{B_s}.$$

In fact, we don't need matrix notation or the entries a_{ij} at all. Instead, we can express A as a linear combination of outer products, with appropriately chosen coefficients:

$$A = \langle 0|A|0\rangle|0\rangle\langle 0| + \langle 0|A|1\rangle|0\rangle\langle 1| + \langle 1|A|0\rangle|1\rangle\langle 0| + \langle 1|A|1\rangle|1\rangle\langle 1|.$$

Let's verify the formula for the $a_{10} = \langle 1|A|0\rangle$ -entry of A , to see how this linear combination of outer products thing works. We start from the definition $A = a_{00}|0\rangle\langle 0| + a_{01}|0\rangle\langle 1| + a_{10}|1\rangle\langle 0| + a_{11}|1\rangle\langle 1|$, and multiply A by $|0\rangle$ from the right and $\langle 1|$ from the left:

$$\begin{aligned}\langle 1|A|0\rangle &= \langle 1|\left(a_{00}|0\rangle\langle 0| + a_{01}|0\rangle\langle 1| + a_{10}|1\rangle\langle 0| + a_{11}|1\rangle\langle 1|\right)|0\rangle \\ &= \langle 1|\left(a_{00}|0\rangle \underbrace{\langle 0|0\rangle}_1 + a_{01}|0\rangle \underbrace{\langle 1|0\rangle}_0 + a_{10}|1\rangle \underbrace{\langle 0|0\rangle}_1 + a_{11}|1\rangle \underbrace{\langle 1|0\rangle}_0\right) \\ &= \langle 1|\left(a_{00}|0\rangle + a_{10}|1\rangle\right) \\ &= a_{00} \underbrace{\langle 1|0\rangle}_0 + a_{10} \underbrace{\langle 1|1\rangle}_1 = a_{10}.\end{aligned}$$

Indeed, $\langle 1|A|0\rangle$ is the same as a_{10} . In fact, we'll rarely use the notation a_{10} , since $\langle 1|A|0\rangle$ is just as easy to write, and much more intuitive: the a_{10} -entry of A is what you obtain when you "sandwich" the matrix A between the vectors $\langle 1|$ on the left and $|0\rangle$ on the right.

In Dirac notation, the basis appears explicitly in expressions for entries of a matrix. We can define the entries of A in any other basis easily and precisely. The representation of A with respect to the Hadamard basis $B_h = \{|+\rangle, |-\rangle\}$ is

$${}_{B_h}[A]_{B_h} = {}_{B_h}\begin{bmatrix} \langle +|A|+\rangle & \langle +|A|-\rangle \\ \langle -|A|+\rangle & \langle -|A|-\rangle \end{bmatrix}_{B_h},$$

or equivalently,

$$A = \langle +|A|+\rangle|+\rangle\langle +| + \langle +|A|-\rangle|+\rangle\langle -| + \langle -|A|+\rangle|-\rangle\langle +| + \langle -|A|-\rangle|-\rangle\langle -|.$$

The coefficient $\langle +|A|+\rangle$ is the a_{++} -entry of the matrix representation of A with respect to the Hadamard basis B_h .

Summary Dirac notation is a convenient way to represent linear algebra concepts: vectors $|v\rangle$ and their Hermitian transposes $\langle v|$, vector coordinates $\langle i|v\rangle$, inner products $\langle v|w\rangle$, outer products $|v\rangle\langle w|$, and matrix entries $\langle i|A|j\rangle$. Because of this expressiveness, Dirac notation is widely used when discussing quantum mechanical topics in modern chemistry, physics, and computer science. In particular, Dirac notation is a core component in the study of quantum mechanics. In fact, we could say that if you understand Dirac notation, you already understand half of quantum mechanics.

Exercises

E9.1 Consider the vectors $|u\rangle = \alpha|0\rangle + \beta|3\rangle$ and $|v\rangle = a|1\rangle + b|2\rangle$, which are elements of $\mathbb{C}^4$. Express the ket $|u\rangle$ and the bra $\langle v|$ as four-dimensional coordinate vectors. Indicate whether your answers correspond to row vectors or column vectors.

E9.2 Express the vector $\vec{w} = (1, i, -1)_{B_s}^T \in \mathbb{C}^3$ as a ket $|w\rangle$ and as a bra $\langle w|$. Compute its length $\|\vec{w}\| = \sqrt{\langle w|w\rangle}$.

E9.3 Find the determinant of the following matrix:

$$A = 1|0\rangle\langle 0| + 2|0\rangle\langle 1| + 3|1\rangle\langle 0| + 4|1\rangle\langle 1|.$$

E9.4 Consider the vectors $|u\rangle = |0\rangle$, $|v\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{i}{\sqrt{2}}|1\rangle$, and $|w\rangle = \frac{1}{\sqrt{5}}|0\rangle - \frac{2}{\sqrt{5}}|1\rangle$. Compute the following expressions:

$$\begin{array}{llllll} \text{a)} |u\rangle\langle v| & \text{b)} \||v\rangle\| + \||w\rangle\| & \text{c)} \langle u|v\rangle & \text{d)} \langle v|w\rangle & \text{e)} \langle w|v\rangle & \text{f)} |u\rangle\langle u| \\ \text{g)} \Pi_{|u\rangle}(|v\rangle) & \text{h)} |v\rangle\langle v| & \text{i)} \Pi_{|v\rangle}(|u\rangle) & \text{j)} \Pi_{|v\rangle}(|w\rangle) \end{array}$$

Hint: The projection matrix onto a unit vector $\hat{a}$ is equal to $\hat{a}\hat{a}^\dagger$.

E9.5 Given the matrix $A = 1|0\rangle\langle 0| + 3|0\rangle\langle 1| + 4|1\rangle\langle 0| + 5|1\rangle\langle 1|$ and the vectors $|v\rangle = 1|0\rangle + 3|1\rangle$ and $|w\rangle = -3|0\rangle - 2i|1\rangle$, compute:

$$\text{a)} A|v\rangle \quad \text{b)} \langle v|A \quad \text{c)} A|w\rangle \quad \text{d)} \langle v|A|v\rangle \quad \text{e)} \langle w|A \quad \text{f)} \langle w|A|w\rangle$$

E9.6 Express the linear transformation $T(x, y) = (2x + y, -3y)$ as a matrix that acts on an input vector of the form $x|0\rangle + y|1\rangle$.

9.4 Quantum information processing

Digital technology is sought after because of the computational, storage, and communication advantages of manipulating digital information instead of continuous-time signals. Similarly, quantum technology enables certain new advances for computational and communication tasks. This section will equip you with a mental model for thinking about quantum information processing tasks in analogy to a digital information processing pipeline you're already familiar with: the digitization, compression, and playback of sound recording that we discussed in Section 7.11 (page 431).

The use of quantum technology for information processing tasks is no more mysterious than the use of digital technology for information processing tasks. Playing a digital recording on your mp3 player involves a number of processing, conversion, and signal amplification steps. Similarly, using a quantum computer involves several conversion, processing, and measurement steps. In both cases you

input some data into a machine, and wait for the machine to process the data and output the answer.

We can think of both digital and quantum technology as black box processes, with internal workings that we can't access directly. In both cases, the intermediate representation of data is in a format that is unintelligible: we can't understand what is encoded in quantum states any more than we can understand what is encoded in digital data. For instance, an mp3 file contains ones and zeros; but, unless we're unusually gifted, it's impossible to tell which artist plays the song 010100101010101000111... just by looking at the raw, digital data. To understand information processing in the digital and quantum worlds, we must study the “adaptors”—the processes used to convert the internal data representation into signals we can intelligibly perceive.

To further highlight the parallel structure between digital information processing and quantum information processing, we'll now review the mp3 compression task as an example of a digital information processing pipeline.

Digital signal processing

A *sound card* is a computer component that converts between analog signals that we can hear and digital signals that computers understand. The sound card digitizes sound using an analog-to-digital converter (ADC). For music playback, the sound card uses a digital-to-analog converter (DAC), which transforms digital sounds into analog sound vibrations to be played through speakers. The ADC receives signal via the sound card's line-in and microphone jacks; the DAC outputs sound via the sound card's line-out and headphone jacks.




Figure 9.9: A digital information processing pipeline for sound recording and playback. Sound vibrations are captured by a microphone and converted to digital form using an analog-to-digital converter (ADC). Next the digital wav file is converted to the more compact mp3 format using digital processing. In the last step, sound is converted back into analog sound vibrations by a digital-to-analog converter (DAC).

Figure 9.9 illustrates a full digital information processing pipeline for sound. We use an *analog-to-digital converter* (ADC) to transform

the analog sound into digital form; the sound's digital representation is then processed, and finally a *digital-to-analog converter* (DAC) transforms the digital signal into analog form. If the music encoding, processing, and playback steps are successful, the final output will sound like the original sound recorded.

The grey-shaded region in Figure 9.9 corresponds to digital data. The example of mp3 compression is just one of many uses of digital processing that become possible once we convert analog information into digital form. Information stored in digital form allows for countless other affordances: the long-term storage of the mp3 file, the file's transmission over the internet, or the computation of a digital fingerprint of the song. Mathematically speaking, we can describe any digital information processing task as a function, $f : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2^k$, that takes n -dimensional binary vectors as inputs and produces k -dimensional binary vectors as outputs.

Quantum information processing

In order to understand the context where quantum information processing has applications, we'll describe a hypothetical scenario. Suppose you're trying to calculate the output of some mathematical function $f(x)$ for the input x . Your first approach is to express the math function f as a function in code $\mathbf{f}$, but you quickly notice that running the code of the function $\mathbf{f}$ takes too long for the types of inputs you're interested in. Computing $\mathbf{f}(x)$ would take a century on your laptop, and even if you ran your code on a powerful cluster of computers, it might still take years to compute the output. Normally people quit at this point.

You're not a quitter though, so instead of waiting for the classical computer to finish computing $\mathbf{f}(x)$, you decide to use a quantum computer. Figure 9.10 shows the sequence of steps needed to operate a quantum computer. Starting from the input x expressed as a classical bitstring $x \in \mathbb{Z}_2^k$, you use state preparation to encode the input into a quantum vector $|x\rangle \in \mathbb{C}^m$ inside the quantum computer. You can then perform quantum operations on the input state to transform it to an output state $|y\rangle \in \mathbb{C}^n$. Finally, you perform a measurement on the output state $|y\rangle$ to obtain a classical bitstring $y \in \mathbb{Z}_2^\ell$, which is hopefully the answer you're looking for.

The main difference between classical and quantum computation is the underlying data types used in the computation. Instead of working with *bits* and using digital functions $\mathbf{f} : \mathbb{Z}_2^k \rightarrow \mathbb{Z}_2^\ell$, quantum computers work with *qubits* and use *quantum operations* $T : \mathbb{C}^m \rightarrow \mathbb{C}^n$ to perform computations. Quantum computation is more general than classical computation since quantum computers can make use

of quantum effects like constructive interference.

The other big difference between classical and quantum computation is that the quantum computer outputs the state $|y\rangle$ only once. You can think of **quantum measurement as asking a question** about the state $|y\rangle$. You're free to perform any measurement, but **you can ask only one question**, since quantum measurements disrupt the state of a system.




Figure 9.10: A quantum information processing pipeline. A classical bitstring x of length k is used as instructions for preparing an m -dimensional quantum state $|x\rangle$. Next, quantum operations are performed on the state $|x\rangle$ to convert it to the output state $|y\rangle$. Finally, the state $|y\rangle$ is *measured* to obtain the classical bitstring y as output.

Quantum processing pipelines are analogous to digital information processing pipelines. The process of *state preparation* in quantum processing is analogous to the analog-to-digital conversion step in digital processing. In both cases we convert the input to the format required for the processing step. Similarly, quantum *measurements* correspond to the digital-to-analog conversion step. In both cases we must convert the output to a format we can understand.

In the next section, we'll discuss the components of the quantum information processing pipeline in more detail. We'll introduce the four *postulates* of quantum mechanics, which specify how quantum systems are represented and what we can do with them. The postulates of quantum mechanics roughly correspond to the conversion steps illustrated in Figure 9.10. One postulate defines how quantum states are prepared, another postulate describe the types of operations we can perform on quantum states, and a third postulate formally defines the process of quantum measurement. The next section is the “quantum mechanics explained in the space on the back of an envelope” part alluded to in the introduction of this chapter. We've set the scene, introduced Dirac notation, and now we can finally discuss the details of the quantum formalism.

9.5 Postulates of quantum mechanics

The *postulates* of quantum mechanics dictate the rules for working within the “quantum world.” The four postulates define:

- What quantum states are
- Which quantum operations can be performed on quantum states
- How to extract information from quantum systems by measuring them
- How to represent composite quantum systems

These postulates specify the structure that all quantum theories must have. Together, the four postulates are known as the *quantum formalism*, and describe the math structure common to all fields that use quantum mechanics: physics, chemistry, engineering, and quantum information. Note the postulates are not provable or derivable from a more basic theory: scientists simply take the postulates as facts and make sure their theories embody these principles.

Quantum states

Quantum states are modelled as special types of vectors. The *state* of a d -dimensional quantum system is a unit vector $|\psi\rangle$, in a d -dimensional complex inner product vector space $\mathbb{C}^d$, which we call the Hilbert space. This is the first postulate that belongs on the back of the envelope.

Postulate 1. To every isolated quantum system is associated a complex inner product space (Hilbert space) called the *state space*. A state is described by a unit vector in state space.

The following comments apply to the description of quantum systems:

- The requirement that state vectors must have length one will become important when we discuss the probabilistic nature of quantum measurements.
- The *global phase* of a quantum state vector doesn't matter, which means $|\psi\rangle = e^{i\theta}|\psi\rangle, \forall \theta \in \mathbb{R}$. The vectors $|\psi\rangle$, $-|\psi\rangle$, and $i|\psi\rangle$ all represent the same quantum state.
- Each physical system corresponds to its own Hilbert space, usually denoted with the same label as the system, $|\psi\rangle_1 \in V_1$ and $|\psi\rangle_2 \in V_2$.

In general, the quantum states of physical systems are represented as vectors in d -dimensional or sometimes even infinite-dimensional Hilbert spaces. To keep things simple, we'll focus on two-dimensional quantum systems.

The qubit In analogy with two-state classical bits $\in \{0, 1\}$ we call two-dimensional quantum systems *qubits* $\in \mathbb{C}^2$, which is short for *quantum bit*. Many physical systems, like the polarization of a photon or the spin of an electron, can be represented as qubits. A qubit is a unit vector in a two-dimensional Hilbert space $\mathbb{C}^2$:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle,$$

where

$$\alpha \in \mathbb{R}, \quad (\text{global phase doesn't matter})$$

$$\beta \in \mathbb{C},$$

$$|\alpha|^2 + |\beta|^2 = 1.$$

Recall that $|0\rangle$ and $|1\rangle$ are the vectors of the standard basis for $\mathbb{C}^2$:

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

The restriction that α must be a real number follows from the fact that the global phase of a quantum state can be ignored. The condition that a quantum state must have length one is equivalent to the constraint $|\alpha|^2 + |\beta|^2 = 1$.




Figure 9.11: A photon encounters a half-silvered mirror, which is also known as a beam splitter. The photon can take one of the two possible paths, so we describe it as the superposition $|\gamma\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$, where $|0\rangle$ describes the photon passing through the mirror, and $|1\rangle$ describes the photon being reflected.

Though the notion of a qubit is an abstract concept, many physical systems can embody it. In Figure 9.11, the information of a qubit is encoded in the path a photon takes after it encounters a half-silvered mirror. Qubits are a device-independent representation of quantum information, similar to how classical bits are device-independent representations for classical information. A bit is a bit, regardless of whether it is transmitted over the network, stored in RAM, or stored on a hard drive. Similarly, a qubit is a qubit, regardless of whether it's encoded in the polarization of a photon, an electron's spin, or in the direction of magnetic flux of a superconducting loop.

Quantum state preparation The operation of encoding some classical information into a quantum system is called *state preparation*. Imagine an apparatus that prepares quantum systems in one of several possible quantum states, depending on the position of the “control switch” x of the machine.

$$x \longrightarrow \boxed{\text{state preparation}} \longrightarrow |x\rangle$$

Figure 9.12: The classical input x is used to prepare a quantum system. The quantum state $|x\rangle$ is produced when the classical input is x .

An example of quantum state preparation is a machine that can produce photons in two different polarizations $|H\rangle$ and $|V\rangle$. If the input $x = 0$ is specified, the machine will produce the state $|H\rangle$, and if $x = 1$ is specified as the input, the machine will produce $|V\rangle$.

The quantum state preparation step is analogous to the analog-to-digital conversion step; it’s how we get classical information into a quantum computer. Once we have converted the input x into a quantum state $|x\rangle$, the natural next question to ask is what we can *do* with $|x\rangle$. What quantum operations are we allowed to perform on quantum states?

Quantum operations

The second definition we can write on the back of the envelope is about identifying *quantum operations* with unitary transformations acting on quantum states. Every quantum operation can be represented as a unitary operation U applied to the input state $|\psi\rangle$ to produce the output state $|\psi'\rangle$.

$$|\psi\rangle \longrightarrow \boxed{U} \longrightarrow |\psi'\rangle$$

The requirement that all quantum operations be unitary is codified in the second postulate of quantum mechanics.

Postulate 2. Time evolution of an isolated quantum system is unitary. If the state at time t is $|\psi\rangle$ and at time t' is $|\psi'\rangle$, then there exists a unitary operator U such that $|\psi'\rangle = U|\psi\rangle$.

This is a major piece of the quantum puzzle, so let’s analyze what it means: mathematically, physically, and computationally.

First let’s get the math out of the way. Recall that a unitary matrix U obeys $U^\dagger U = U U^\dagger = \mathbb{1}$. Postulate 2 ensures that quantum states will maintain their unit-length property after quantum operations are performed on them. Assume the quantum system starts

from a state $|\psi\rangle$ that has length one, $\| |\psi\rangle \|^2 = \langle \psi | \psi \rangle = 1$. After the unitary U is applied, the state after evolution will be $|\psi'\rangle = U|\psi\rangle$ and its squared length will be $\| |\psi'\rangle \|^2 = \| U|\psi\rangle \|^2 = \langle \psi | U^\dagger U | \psi \rangle = \langle \psi | \mathbb{I} | \psi \rangle = \langle \psi | \psi \rangle = 1$. In other words, **quantum operations are length-preserving**.

The second postulate refers explicitly to a time variable and the state of the system at different times, which is the physics way of thinking about quantum operations. For example, applying a quantum operation to a real physical system could correspond to applying a certain magnetic field to that system. If applying the magnetic field for one second applies the rotation operation R_θ , then applying the magnetic field for two seconds performs the rotation $R_{2\theta}$; therefore the time variable is of central importance to the experiment.

The computer science perspective on the second postulate is more abstract. Instead of counting seconds, computer scientists model quantum operations as discrete “gates” that can be applied to quantum states. Computer scientists describe their quantum algorithms in terms of unitary operators U , and they assume there exists a specific physical operation that performs the quantum operation U on a real physical qubit system.

In this chapter, we’ll take the computer science approach and think of quantum information processing tasks in terms of applying different gates to quantum states. We’ll now define several quantum gates that perform useful unitary operations on qubits.

Example 1: phase gate The Z operator is defined by its action on the vectors of the standard basis.

$$Z|0\rangle = |0\rangle, \quad Z|1\rangle = -1|1\rangle.$$

The Z operator leaves the $|0\rangle$ unchanged but flips the phase of $|1\rangle$.

Given knowledge of the actions of the Z operator on the vectors of the standard basis, we can construct its matrix representations:

$$Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}_{B_s} = |0\rangle\langle 0| - |1\rangle\langle 1|.$$

Example 2: NOT gate The X operator is defined by the following actions on the vectors of the standard basis:

$$X|0\rangle = |1\rangle, \quad X|1\rangle = |0\rangle.$$

The X operator acts as a “quantum NOT gate,” changing $|0\rangle$ s into $|1\rangle$ s, and $|1\rangle$ s into $|0\rangle$ s. The matrix representation of the X operator is:

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}_{B_s} = |0\rangle\langle 1| + |1\rangle\langle 0|.$$

Example 3: Hadamard gate The Hadamard operator takes the vectors of the standard basis to the vectors of the Hadamard basis $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$:

$$H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) = |+\rangle, \quad H|1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) = |-\rangle.$$

You can also think of the H operator as a 45° counter-clockwise rotation. The matrix representation of the Hadamard gate is

$$H = {}_{B_s} \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}_{B_s}.$$

By linearity, we can deduce the effects of the operators Z , X , and H on an arbitrary qubit $\alpha|0\rangle + \beta|1\rangle$:

$$\begin{aligned} Z(\alpha|0\rangle + \beta|1\rangle) &= \alpha|0\rangle - \beta|1\rangle, \\ X(\alpha|0\rangle + \beta|1\rangle) &= \beta|0\rangle + \alpha|1\rangle, \\ H(\alpha|0\rangle + \beta|1\rangle) &= \frac{\alpha+\beta}{\sqrt{2}}|0\rangle + \frac{\alpha-\beta}{\sqrt{2}}|1\rangle. \end{aligned}$$

Example 4 The effect of the operator XZ corresponds to the combination of the effects of the Z and X operators. We can understand the action of XZ either by applying it to an arbitrary qubit $\alpha|0\rangle + \beta|1\rangle$:

$$XZ(\alpha|0\rangle + \beta|1\rangle) = -\beta|0\rangle + \alpha|1\rangle,$$

or by multiplying together the operator's matrix representations:

$$XZ = {}_{B_s} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}_{B_s} {}_{B_s} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}_{B_s} = {}_{B_s} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}_{B_s}.$$

In general, it is possible to perform many quantum operations and combinations of operations when working with quantum states. The examples presented above represent the most commonly used operations. Note that unitary time evolution is invertible: for every quantum gate G there exists an inverse gate G^{-1} such that $G^{-1}G = \mathbb{1}$.

Links

[Wikipedia article on quantum gates]

https://en.wikipedia.org/wiki/Quantum_gate

Exercises

E9.7 The Hadamard gate is defined as $H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$. Compute the effect of the operator HH on the vectors of the standard basis $\{|0\rangle, |1\rangle\}$.

E9.8 Compute XX , XZ , and ZX .

* * *

So far, we've filled half the space on the back of the envelope. It's time to talk about the third fundamental idea in quantum mechanics: quantum measurements.

Quantum measurements

A quantum measurement performed on a quantum system corresponds to a collection of projection operators $\{\Pi_i\}$ that act on the Hilbert space. A measurement with n possible outcomes is represented by n projection operators $\{\Pi_1, \Pi_2, \dots, \Pi_n\}$. The projection operators form a *decomposition of the identity*, meaning their sum is equal to the identity matrix:

$$\sum_{i=1}^n \Pi_i = \mathbb{1}.$$

Intuitively, the n different projection operators correspond to n different alternatives for the evolution of the quantum system. Performing the measurement is like asking "which way is it going to be?" and letting the system decide which path it wants to take. *Born's rule* is used to assign probabilities to different measurement outcomes.

Postulate 3. A quantum measurement is modelled by a collection of projection operators $\{\Pi_i\}$ that act on the state space of the system being measured and satisfy $\sum_i \Pi_i = \mathbb{1}$. The index i labels the different measurement outcomes.

The probability of outcome i when performing measurement $\{\Pi_i\}$ on a quantum system in the state $|\psi\rangle$ is given by the squared length of the state after applying the i^{th} projection operator:

$$\Pr(\{\text{outcome } i \text{ given state } |\psi\rangle\}) \stackrel{\text{def}}{=} \|\Pi_i|\psi\rangle\|^2 \quad (\text{Born's rule}).$$

When outcome i occurs, the post-measurement state of the system is

$$|\psi_i'\rangle \stackrel{\text{def}}{=} \frac{\Pi_i|\psi\rangle}{\|\Pi_i|\psi\rangle\|}.$$

Let's unpack this definition to see what is going on.

Born's rule For the measurement defined by the projection operators $\{\Pi_1, \Pi_2, \dots, \Pi_n\}$, Born's rule states that the probability of outcome i is $\langle \psi | \Pi_i | \psi \rangle$.

This expression for the squared norm of the overlap between $|\psi\rangle$ and Π_i can be written in several equivalent ways:

$$\|\Pi_i|\psi\rangle\|^2 = (\Pi_i|\psi\rangle, \Pi_i|\psi\rangle) = \langle \psi | \Pi_i \Pi_i | \psi \rangle = \langle \psi | \Pi_i | \psi \rangle,$$

where the last equality follows from the idempotence property of projectors $\Pi_i = \Pi_i^2$. The last expression, where the projection operator is "sandwiched" by two copies of the quantum state, is the physicist's usual way of expressing Born's rule by defining $\Pr(\{\text{outcome } i\}) \stackrel{\text{def}}{=} \langle \psi | \Pi_i | \psi \rangle$. For the class of projective measurements we're discussing here, the two definitions are equivalent.

The set of projection operators $\{\Pi_1, \Pi_2, \dots, \Pi_n\}$ forms a decomposition of the identity $\mathbb{1} = \sum_i \Pi_i$. This guarantees that the probability distribution of the different outcomes is normalized:

$$\begin{aligned} 1 &= \|\mathbb{1}|\psi\rangle\|^2 = \|(\Pi_1 + \dots + \Pi_n)|\psi\rangle\|^2 \\ &\stackrel{\text{py}}{=} \|\Pi_1|\psi\rangle\|^2 + \dots + \|\Pi_n|\psi\rangle\|^2 \\ &= \Pr(\{\text{outcome } 1\}) + \dots + \Pr(\{\text{outcome } n\}). \end{aligned}$$

That's good to check; otherwise Kolmogorov would be angry with us. Note the equality labelled $\stackrel{\text{py}}{=}$ follows from Pythagoras' theorem; we're using the fact that the operators $\{\Pi_i\}$ are mutually orthogonal.

Post-measurement state When outcome i occurs, Postulate 3 tells us that the state of the quantum system becomes $|\psi'_i\rangle = \frac{\Pi_i|\psi\rangle}{\|\Pi_i|\psi\rangle\|}$, which is the result of applying the projection Π_i to obtain $\Pi_i|\psi\rangle$, and then normalizing the state so that $\|\psi'_i\rangle\| = 1$.

Measurements are not passive observations! Quantum measurement is an invasive procedure that typically changes the state of the system being measured. In general, the quantum state after the measurement will not be the same as the state before the measurement:

$$|\psi'_i\rangle \neq |\psi\rangle,$$

and we say the state is *disturbed* by the measurement, (though it's possible that $|\psi'_i\rangle = |\psi\rangle$ when the input state lives entirely within the image space of one of the projection operators $|\psi'_i\rangle = \Pi_i|\psi\rangle = |\psi\rangle$).

A quantum measurement is an interaction that creates classical information and destroys quantum information. By measuring, we obtain the classical information i that tells us which outcome occurred, but we disturb the initial state $|\psi\rangle$, forcing the quantum state into the “aligned with Π_i ”-state $|\psi'_i\rangle$. We can still carry out further experiments with the post-measurement state $|\psi'_i\rangle$, but it’s not the same as the initial state $|\psi\rangle$. Specifically, we’ve lost all the information about $|\psi\rangle$ that used to exist in the subspace $(\mathbb{1} - \Pi_i)$.

Another way to describe what happens during a quantum measurement is to say the state $|\psi\rangle$ *collapses* into the state $|\psi'_i\rangle$. This is the terminology used by physicists to describe the effects of the projection that occurs during a measurement. Before the measurement, the quantum state $|\psi\rangle$ could be any vector, but after we observe the outcome i , we know the state has “collapsed” and is now confined to the image subspace of the projection operator Π_i . We won’t use the terminology of state collapse in this chapter, because the term “projection” more accurately describes what’s happening.

Historically, the founders of quantum mechanics chose some rather poor terminology to describe quantum measurements. In addition to “collapse,” they also used the term “observer” to describe whomever observes the outcome i of a quantum measurement experiment. Trying to understand quantum mechanics using this terminology often leads to magical thinking and anthropomorphizing of the underlying physical events. Proponents of pop-psychology ideas have latched onto the term “observer” to suggest that the collapse of quantum states that occurs during quantum measurements is caused by the presence of human consciousness. Watch out for that fluff.

A more grounded way to think about quantum measurement is in terms of *interaction* between the quantum state of a particle and a classical measurement apparatus that can be in one of n possible states. Due to the relative size of the two interacting systems (tiny quantum system and large measurement apparatus), the state $|\psi\rangle$ is forced to “align” with one of the n possible states of the measurement apparatus.

Example 4 In Figure 9.13 a state vector $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ is measured with photo detectors modelled as projectors given by

$$\begin{aligned}\Pi_0 &= |0\rangle\langle 0| & \sum_j \Pi_j &= \mathbb{1} \\ \Pi_1 &= |1\rangle\langle 1|\end{aligned}$$

$$\begin{aligned}
\Pr(\{0\}|\psi) &= \Pr(\{\text{outcome } 0 \text{ given state } |\psi\rangle\}) \\
&= \langle\psi| \Pi_0 |\psi\rangle \\
&= \langle\psi| |0\rangle\langle 0| |\psi\rangle \\
&= (\bar{\alpha}\langle 0| + \bar{\beta}\langle 1|) |0\rangle\langle 0| (\alpha|0\rangle + \beta|1\rangle) \\
&= \bar{\alpha}\alpha \\
&= |\alpha|^2.
\end{aligned}$$

The probability of outcome 1 is $\Pr(\{1\}|\psi) = \langle\psi|\Pi_1|\psi\rangle = |\beta|^2$. After the measurement, the quantum system exists in one of two possible states: $|\psi'_0\rangle = |0\rangle$ or $|\psi'_1\rangle = |1\rangle$.

Example 5 Consider the measurement $\{\Pi_+, \Pi_-\}$ that consists of the projectors onto the Hadamard basis:

$$\begin{aligned}
\Pi_+ &= |+\rangle\langle +| & \Pi_+ + \Pi_- &= \mathbb{I} \\
\Pi_- &= |-\rangle\langle -|
\end{aligned}$$

Given the quantum state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, the probability of outcome “+” is given by

$$\begin{aligned}
\Pr(\{+\}|\psi) &= \|\Pi_+|\psi\rangle\|^2 \\
&= \| |+\rangle\langle +| \psi \|^2 \\
&= \| |+\rangle\langle +| (\alpha|0\rangle + \beta|1\rangle) \|^2 \\
&= \| |+\rangle (\alpha\langle +|0\rangle + \beta\langle +|1\rangle) \|^2 \\
&= \| |+\rangle \left(\alpha \frac{1}{\sqrt{2}} + \beta \frac{1}{\sqrt{2}} \right) \|^2 \\
&= \frac{|\alpha + \beta|^2}{2} \| |+\rangle \|^2 \\
&= \frac{|\alpha + \beta|^2}{2}.
\end{aligned}$$

The probability of outcome “−” is $\Pr(\{-\}|\psi) = \frac{|\alpha - \beta|^2}{2}$. After the measurement, the quantum system is in one of two possible states: $|\psi'_+\rangle = |+\rangle$ or $|\psi'_-\rangle = |-\rangle$.

The measurement process is a fundamental aspect of quantum models. You’ll need to acclimate to the idea that measurements change systems’ states. It’s not magic; it’s a phenomenon that occurs due to the relative size of the systems (tiny quantum particles and huge measurement apparatus), and the fact that measurements force quantum and classical systems to interact.




Figure 9.13: The state of a photon after encountering a $(1 - \alpha)$ -silvered mirror is $|\gamma\rangle = \alpha|0\rangle + \beta|1\rangle$. The probability that the horizontal photodetector “clicks” is $|\alpha|^2$, and is obtained by projecting $|\gamma\rangle$ on the subspace $|0\rangle\langle 0|$. The probability that the top photodetector clicks is equal to $|\beta|^2$, and is obtained by projecting $|\gamma\rangle$ on the subspace $|1\rangle\langle 1|$.

Composite quantum systems

So far we discussed state preparation, quantum operations, and quantum measurements of individual qubits. There’s just enough room on the back of our envelope to discuss quantum models for systems made of multiple qubits.

Classically, if we have two bits $b_1 \in \{0,1\}$ and $b_2 \in \{0,1\}$, we can concatenate them to obtain a bit string $b_1b_2 \in \{0,1\}^2$, which can have one of four possible values: 00, 01, 10, and 11. The combined state of two qubits $|\varphi_1\rangle \in \mathbb{C}^2$ and $|\varphi_2\rangle \in \mathbb{C}^2$ is the *tensor product state* $|\varphi_1\rangle \otimes |\varphi_2\rangle$ in the four-dimensional *tensor product space* $\mathbb{C}^2 \otimes \mathbb{C}^2 = \mathbb{C}^4$. A basis for the tensor product space can be obtained by taking all possible combinations of the basis vectors for the individual qubits: $\{|0\rangle \otimes |0\rangle, |0\rangle \otimes |1\rangle, |1\rangle \otimes |0\rangle, |1\rangle \otimes |1\rangle\}$.

Postulate 4. The state space of a composite quantum system is equal to the tensor product of the state spaces of the individual systems. If systems $1, 2, \dots, n$ exist in states $|\varphi_1\rangle, |\varphi_2\rangle, \dots, |\varphi_n\rangle$, then the state of the composite system is $|\varphi_1\rangle \otimes |\varphi_2\rangle \otimes \dots \otimes |\varphi_n\rangle$.

Postulate 4 tells us how we can combine the state spaces of different quantum systems to describe a composite system. Many interesting quantum applications involve operations on multiple qubits and are described by vectors in a tensor product space, so let’s look more closely at this “ $\otimes$ ”-thing.

Tensor product space If you’d never heard of *tensor products* before, don’t worry—the only scary part is the tensor product symbol “ $\otimes$,” which we’ll explain in this section. A *tensor product space* consists of all possible combinations of the basis vectors for the two subspaces.

For example, consider two qubits $|\varphi_1\rangle \in V_1 = \mathbb{C}^2$ and $|\varphi_2\rangle \in V_2 = \mathbb{C}^2$. We'll denote the standard basis for V_1 as $B_1 = \{|0\rangle_1, |1\rangle_1\}$ and the standard basis for V_2 as $B_2 = \{|0\rangle_2, |1\rangle_2\}$. The tensor product space $B_{12} = V_1 \otimes V_2$ is four-dimensional and has the following basis:

$$B_{12} = \{|0\rangle_1 \otimes |0\rangle_2, |0\rangle_1 \otimes |1\rangle_2, |1\rangle_1 \otimes |0\rangle_2, |1\rangle_1 \otimes |1\rangle_2\}.$$

This level of subscripts and the explicit use of the symbol $\otimes$ hurts the eyes (and the hand if you must use this notation to solve problems). It's therefore customary to drop the subscripts, omit the tensor product symbol, and draw a single ket that contains a "string" of indices:

$$|a\rangle_1 \otimes |b\rangle_2 = |a\rangle \otimes |b\rangle = |a\rangle |b\rangle = |ab\rangle.$$

The basis for the tensor product space $B_{12} = V_1 \otimes V_2$ looks much nicer in the simplified notation:

$$B_{12} = \{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}.$$

Tensor product of two vectors Suppose we're given two qubits with states described by the following vectors:

$$|\varphi_a\rangle_1 = \alpha_1|0\rangle_1 + \beta_1|1\rangle_1, \quad |\varphi_b\rangle_2 = \alpha_2|0\rangle_2 + \beta_2|1\rangle_2.$$

Note the subscripts indicate which system we're describing: $|0\rangle_1$ is the state $|0\rangle$ for the first qubit, while $|0\rangle_2$ is the state $|0\rangle$ of the second qubit.

The state of the combined system is the tensor product state $|\varphi_{ab}\rangle_{12} = |\varphi_a\rangle_1 \otimes |\varphi_b\rangle_2$, which is computed by combining all possible combinations of the components of $|\varphi_a\rangle_1$ and the components of $|\varphi_b\rangle_2$:

$$(\alpha_1, \beta_1)_{B_1} \otimes (\alpha_2, \beta_2)_{B_2} = (\alpha_1\alpha_2, \alpha_1\beta_2, \beta_1\alpha_2, \beta_1\beta_2)_{B_{12}}.$$

The notion of "all possible combinations" is easier to see by considering the tensor product operation in terms of the basis vectors:

$$\begin{aligned} |\varphi_{ab}\rangle_{12} &= |\varphi_a\rangle_1 \otimes |\varphi_b\rangle_2 \\ &= (\alpha_1|0\rangle_1 + \beta_1|1\rangle_1) \otimes (\alpha_2|0\rangle_2 + \beta_2|1\rangle_2) \\ &= \alpha_1\alpha_2|0\rangle_1|0\rangle_2 + \alpha_1\beta_2|0\rangle_1|1\rangle_2 + \beta_1\alpha_2|1\rangle_1|0\rangle_2 + \beta_1\beta_2|1\rangle_1|1\rangle_2 \\ &= \alpha_1\alpha_2|00\rangle + \alpha_1\beta_2|01\rangle + \beta_1\alpha_2|10\rangle + \beta_1\beta_2|11\rangle \\ &= (\alpha_1\alpha_2, \alpha_1\beta_2, \beta_1\alpha_2, \beta_1\beta_2)_{B_{12}}, \end{aligned}$$

where $B_{12} = \{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ is the standard basis for the tensor product space.

State spaces and dimension counting A quantum state that consists of n qubits can represent any unit vector in $\mathbb{C}^{2^n}$. That's an insanely big state space—a huge 2^n -dimensional playground. In comparison, a classical bitstring of length n can take on one of 2^n values. Let's compare the state space of a two-bit classical register R with the state space of a two-qubit quantum register $|R\rangle$. The classical register can take on one of four possible values:

$$R \in \{00, 01, 10, 11\},$$

whereas the quantum register $|R\rangle$ can be any unit vector in $\mathbb{C}^4$. Similar to a classical registers, the standard basis for the state space of a two-qubit quantum register also consists of four basis vectors:

$$|00\rangle = (1, 0, 0, 0), \quad |01\rangle = (0, 1, 0, 0), \quad |10\rangle = (0, 0, 1, 0), \quad |11\rangle = (0, 0, 0, 1),$$

but the quantum register can also represent superpositions of the basis states like $|R\rangle = (\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}, 0, 0)$ or $|R\rangle = (\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, 0)$.

Using a very large vector space to represent states does not necessarily make a model more powerful, but the large dimension of the tensor product space suggests many new possibilities. Much of the recent excitement in the area of quantum computing is based on the promise of using the qubits of a quantum computer to perform computations in very large quantum state spaces. We shouldn't get carried away with enthusiasm, because with great state space comes great noise. It's easy to imagine n qubits in a row in a mathematical model, but building a physical system that can store n qubits and protect them from noise is a much more difficult task. Another bottleneck in quantum computing is the difficulty of extracting information from quantum systems. The quantum state space of n qubits is $\mathbb{C}^{2^n}$, but projective measurements of the form $\{\Pi_1, \Pi_2, \dots, \Pi_m\}$ can only obtain *one* answer to a question with m possible classical outcomes ($m \leq 2^n$). We'll learn more about theoretical and practical considerations for quantum computing in Section 9.8.

Exercises

E9.9 Show that the quantum state $|0\rangle|1\rangle - |1\rangle|0\rangle$ is equal to the quantum state $|+\rangle|-\rangle - |-\rangle|+\rangle$.

Hint: Express $|0\rangle$ and $|1\rangle$ in the basis $\{|+\rangle, |-\rangle\}$.

Quantum entanglement

At the risk of veering further off-topic for a linear algebra book, we'll now briefly describe *entangled* quantum states. In particular, we'll

discuss the properties of the *entangled state* $|\Psi_-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$. Entanglement is some really crazy stuff, and some of the most fascinating results in quantum information science make use of pre-shared entangled states.

In Section 7.9 we discussed how a secret key shared between two parties, Alice and Bob, can function as a *communication resource* used to achieve private communication (via the one-time pad cryptosystem). You can think of entanglement as another type of communication resource: a stronger-than-classical correlation between two parts of a quantum system. One half of the system is controlled by Alice, the other half is controlled by Bob. When the collective state of a two-qubit quantum system is in the entangled state $\frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$, measuring the individual qubits will produce anti-correlated results *in any basis*.

Example 7 The Einstein-Podolsky-Rosen (EPR) state is a two-qubit quantum state with interesting *nonlocal* properties. We assume Alice controls one half of the quantum state, and Bob controls the other half of the following state:

$$|\Psi_-\rangle^{AB} = \frac{1}{\sqrt{2}}|0\rangle^A|1\rangle^B - \frac{1}{\sqrt{2}}|1\rangle^A|0\rangle^B.$$

Note the use of superscripts to denote which party controls each part of the system.

Let's analyze the different measurements Alice and Bob can perform on this state. If Alice measures her system in the basis $\{|0\rangle, |1\rangle\}$, the projection operators that correspond to the two outcomes are

$$\Pi_0^{AB} = |0\rangle\langle 0|^A \otimes \mathbb{1}^B \quad \text{and} \quad \Pi_1^{AB} = |1\rangle\langle 1|^A \otimes \mathbb{1}^B.$$

Since only Alice's half of the state is measured, the measurement acts like the identity operator on Bob's half of the state. There's a 50-50 chance of outcomes 0 and 1:

$$\Pr(0|\Psi_-) = \langle \Psi_- | \Pi_0^{AB} | \Psi_- \rangle = \frac{1}{2}, \quad \Pr(1|\Psi_-) = \langle \Psi_- | \Pi_1^{AB} | \Psi_- \rangle = \frac{1}{2}.$$

Depending on the outcome, the post-measurement state of the system will be either $|0\rangle^A|1\rangle^B$ or $|1\rangle^A|0\rangle^B$. If Bob then measures his half of the system, he'll obtain the outcome opposite Alice's. In other words, the measurement outcomes that Alice and Bob obtain are perfectly anti-correlated.

What if Alice and Bob choose to measure their respective halves of the EPR state $|\Psi_-\rangle^{AB}$ in the basis $\{|+\rangle, |-\rangle\}$? Using some basic calculations (see Exercise E9.9), we can express the EPR state $|\Psi_-\rangle^{AB}$

in terms of the basis $\{|+\rangle, |-\rangle\}$ as follows:

$$\frac{1}{\sqrt{2}}|0\rangle^A|1\rangle^B - \frac{1}{\sqrt{2}}|1\rangle^A|0\rangle^B = |\Psi_{-}\rangle^{AB} = \frac{1}{\sqrt{2}}|+\rangle^A|-\rangle^B - \frac{1}{\sqrt{2}}|-\rangle^A|+\rangle^B.$$

Observe that the state $|\Psi_{-}\rangle^{AB}$ has the same structure in the Hadamard basis as in the standard basis. Thus, Alice and Bob's measurement outcomes will also be perfectly anti-correlated when measuring in the Hadamard basis.

A state $|\psi\rangle^{AB}$ is called *entangled* if it cannot be written as a tensor product of quantum states $|\phi\rangle^A \otimes |\varphi\rangle^B$, where $|\phi\rangle^A$ describes the state held by Alice and $|\varphi\rangle^B$ the state held by Bob. The EPR state $|\Psi_{-}\rangle^{AB}$ is *entangled*, which means it cannot be written as a tensor product of the quantum states of individual qubits:

$$|0\rangle \otimes |1\rangle - |1\rangle \otimes |0\rangle \neq (a|0\rangle + b|1\rangle) \otimes (c|0\rangle + d|1\rangle),$$

for any $a, b, c, d \in \mathbb{C}$. Since we cannot describe the EPR state $|\Psi_{-}\rangle^{AB}$ as the tensor product of two local states $|\phi\rangle^A$ and $|\varphi\rangle^B$, we say it requires a *nonlocal* description, which is another way of saying $|\Psi_{-}\rangle^{AB}$ is entangled.

There is something strange about the EPR state. If Alice measures her half of the state and finds $|0\rangle$, then we know immediately that Bob's state will be $|1\rangle$. The collapse in the superposition on Alice's side immediately causes a collapse of the superposition on Bob's side. Note that Bob's collapse will occur *immediately*, no matter how distant Bob's system is from Alice's. This is what the authors Einstein, Podolsky, and Rosen called "spooky action at a distance." How could Bob's system "know" to always produce the opposite outcome even at the other end of the universe?

Now imagine you have a whole bunch of physical systems prepared in the EPR state. Alice controls one half of each of the EPR pairs, while Bob controls the other half. This is a communication resource called *shared entanglement*. Many of the quantum information protocols make use of shared entanglement between Alice and Bob, to achieve novel communication tasks.

A useful perspective for thinking about quantum entanglement is to consider the information encoded in the quantum state. This is called the *information theory* perspective on quantum mechanics. From the point of view of information, we can say a system is entangled whenever we know more about the system as a whole than about its parts. We have complete certainty about the state of the composite system $|\Psi_{-}\rangle^{AB}$, and complete uncertainty about the states of the individual subsystems controlled by Alice and Bob. When Alice measures her half of the entangled state, the two outcomes of the measurement are equally likely. Her measurement outcome essentially corresponds to a completely random bit, which encodes zero

information. I would like to tell you more about quantum information theory, but let's not go further off course.

Instead of discussing theory, let's describe a practical scenario in which it is possible to know more about a whole system than about its constituent parts.

Physics example We'll now describe a physical process that leads to the creation of an entangled quantum state. Consider a quantum particle p that *decays* into two quantum subparticles, p_a and p_b . The decay process obeys various physics conservation laws; in particular, the total spin angular momentum before and after the decay must be conserved. Supposing the particle p has zero spin angular momentum before the decay, then the law of conservation of angular momentum dictates the subparticles p_a and p_b must have opposite spin. The spin angular momentum of each particle is in an unknown direction (up, down, left, right, in, or out), but whichever spin direction we measure for p_a , the spin angular momentum of p_b immediately takes on the opposite direction. This way, regardless of the direction of their individual spins, their combined spin angular momentum will be zero, as required by the law of conservation of angular momentum.

The general scenario discussed above describes what would happen if a Helium atom were to explode, the two electrons in its ground state flying off to distant sides of the universe. The two electrons have opposite spins, but we don't know the directions of the individual spins. The only thing we know is that their total spin equals zero, since the ground state of the Helium atom has spin zero. This is how the "anti-correlation in *any* basis" aspect of quantum entanglement arises.

Summary

We can summarize the concepts of quantum mechanics we learned in this chapter, and relate them to the concepts of linear algebra:

$$\begin{aligned} \text{quantum state} &\Leftrightarrow \text{vector } |v\rangle \in \mathbb{C}^d \\ \text{evolution} &\Leftrightarrow \text{unitary operations} \\ \text{measurement} &\Leftrightarrow \text{projections} \\ \text{composite system} &\Leftrightarrow \text{tensor product} \end{aligned}$$

The quantum formalism embodied in the four postulates of quantum mechanics has been applied in describing many physical phenomena. Using complex vectors to represent quantum states leads to useful models and predictions for experimental outcomes. In the

next section, we'll use the quantum formalism to analyze the outcomes of the polarizing lenses experiment.

In addition to the applications of quantum principles, studying the structure of quantum states and operations is an interesting field on its own. An example of a quantum idea that is fundamentally new is the existence of *entangled* quantum states. Later in this chapter, we'll discuss an interesting application of quantum entanglement as part of the *quantum teleportation* protocol, illustrated in Figure 9.23 (page 527).

Exercises

E9.10 Find the matrix representation of the projection matrices $\Pi_+ = |+\rangle\langle +|$ and $\Pi_- = |-\rangle\langle -|$. Show that $\Pi_+ + \Pi_- = \mathbb{1}$.

E9.11 Compute the probability of outcome “−” for the measurement $\{\Pi_+, \Pi_-\} = \{|+\rangle\langle +|, |-\rangle\langle -|\}$ performed on the quantum state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$.

E9.12 Given the state $|\theta\rangle = (\frac{1}{\sqrt{2}}, \frac{e^{i\theta}}{\sqrt{2}})^T$, find a quantum state $|\theta^\perp\rangle$ that is orthogonal to $|\theta\rangle$. Find the projection operators Π_θ and $\Pi_{\theta^\perp}$ that correspond to the measurements in the basis $\{|\theta\rangle, |\theta^\perp\rangle\}$. Verify that $\Pi_\theta + \Pi_{\theta^\perp} = \mathbb{1}$. Compute the probability of outcome θ when performing the measurement $\{\Pi_\theta, \Pi_{\theta^\perp}\}$ on the state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$.

Hint: The state $|\theta^\perp\rangle$ satisfies $\langle\theta^\perp|\theta\rangle = 0$.

E9.13 Given the two qubits $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ and $|\phi\rangle = \gamma|0\rangle + \delta|1\rangle$, compute the tensor product state $|\psi\rangle \otimes |\phi\rangle$.

Links

[Compact set of notes on QM written by a physicist]
http://graybits.biz/notes/quantum_mechanics/preface

[Lecture series on QM written by a computer scientist]
<http://scottaaronson.com/democritus/lec9.html>

[Quantum mechanics summary written by a philosopher]
<http://plato.stanford.edu/entries/qm/>

9.6 Polarizing lenses experiment revisited

Let's revisit the polarizing lenses experiment, this time modelling photons' polarization states as two-dimensional complex vectors. We define the state of a horizontally polarized photon as $|H\rangle \stackrel{\text{def}}{=} |0\rangle = (1, 0)^T$, and the state of a vertically polarized photon as $|V\rangle \stackrel{\text{def}}{=} |1\rangle = (0, 1)^T$. This choice corresponds to the observation that horizontal and vertical polarizations are complementary. We shall interpret the polarizing lenses in the optical circuit both as state preparation and measurement steps. The measurement outcome we observe is whether photons pass through the complete optical circuit.

Figure 9.14 illustrates the state preparation step we can use in experiments with photons. We can prepare photons in the states $|H\rangle$ or $|V\rangle$ by starting from unpolarized photons, and passing them through either an H -polarizing lens or a V -polarizing lens.

$$\text{light} \rightarrow \boxed{H} \rightarrow \begin{bmatrix} 1 \\ 0 \end{bmatrix} = |H\rangle, \quad \text{light} \rightarrow \boxed{V} \rightarrow \begin{bmatrix} 0 \\ 1 \end{bmatrix} = |V\rangle.$$

Figure 9.14: State preparation procedure for photons with horizontal or vertical polarization. Depending on our choice of H - or V -polarizing lens, we can prepare different quantum states.

Placing a polarizing lens in the path of a photon performs the measurement $\{\Pi_{\leftarrow}, \Pi_{\rightarrow}\}$ of the photon's polarization state. Using the lens to measure a photon's polarization state yields two possible outcomes: "photon passes through lens" or "photon is reflected by lens." The polarizing lens is the measurement apparatus, and the measurement corresponds to whether each photon that hits the lens is horizontally or vertically polarized.

Figure 9.15 shows the projection operators that correspond to a measurement using a horizontally polarizing lens. The outcome "passes through" corresponds to the projection operator $\Pi_{\rightarrow} = |H\rangle\langle H| = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$. The outcome "is reflected" corresponds to the projection matrix $\Pi_{\leftarrow} = |V\rangle\langle V| = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$.

Let's now use the quantum formalism to analyze the results of the three-lenses experiment we discussed earlier in the chapter. Figure 9.17 shows the optical circuit that consists of a state preparation step and two measurement steps. The diagonally polarizing lens placed in the middle of the circuit only allows photons with 45° -diagonal polarization to pass through: $|D\rangle = (\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})^T$. The projection operator associated with the “passes through” outcome of the diagonal polarization measurement is

$$\Pi_{\rightarrow} = |D\rangle\langle D| = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}.$$

The probability of $|H\rangle$ photons passing through the diagonal lens is

$$\begin{aligned} \Pr(\{\text{pass through } D \text{ lens given state } |H\rangle\}) &= \| |D\rangle\langle D| |H\rangle \|^2 \\ &= \left\| \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right\|^2 \\ &= \left\| \begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix} \right\|^2 \\ &= \left(\frac{1}{2}\right)^2 + \left(\frac{1}{2}\right)^2 = \frac{1}{2}. \end{aligned}$$

The post-measurement state of photons that make it through the diagonally polarizing lens is $|D\rangle = (\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})^T$.




Figure 9.17: Photons prepared in the state $|H\rangle$ are subjected to two sequential measurements: a diagonal polarizing measurement D followed by a vertical polarizing measurement V . The projection operators for the “is reflected” and “passes through” outcomes are indicated in each step.

Photons that pass through the middle lens are in the state $|D\rangle$. The probability of these photons passing through the V -polarizing lens

is

$$\begin{aligned}
 \Pr(\{\text{pass through } V \text{ lens given state } |D\rangle\}) &= \|V \times V |D\rangle\|^2 \\
 &= \left\| \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix} \right\|^2 \\
 &= \left\| \begin{bmatrix} 0 \\ \frac{1}{\sqrt{2}} \end{bmatrix} \right\|^2 \\
 &= 0^2 + \left(\frac{1}{\sqrt{2}}\right)^2 = \frac{1}{2}.
 \end{aligned}$$

The overall probability of a photon passing through both measurement lenses is $\frac{1}{4}$.

This probability is consistent with the light intensity observations we see in Figure 9.8 (page 478). Note the new interpretation that quantum mechanics provides for the same observations. We previously referred to the optical power $P = 0.5$ after the first measurement and optical power $P = 0.25$ after the second measurement. The quantum formalism predicts the probability $p = 0.5$ for photons to pass through the first lens and the probability $p = 0.25$ for the photons to reach the end of the circuit.

Discussion

The polarizing lenses experiment illustrates some key aspects of the quantum formalism: the modelling of quantum states as vectors, the process of state preparation, and the effects of measurements on systems. Mainly, the experiment shows how a probabilistic approach for describing the light intensity in the circuit successfully predicts the experiment's outcome.

Let's compare the quantum interpretation of this experiment with a classical interpretation. It's possible to explain all the observations of the polarizing lenses experiment using the classical theory of electromagnetic waves. Modelling the light beam as a classical wave allows us to understand the projections to the "passes through" orientation of the polarizing lenses. Because the light beam consists of many photons, it behaves like a continuous quantity that can be split: part of the wave passes through the lens while another part is reflected. Classical wave theory correctly predicts the qualitative aspects of the experiment shown in Figure 9.17, but the light intensity is described in terms of optical power $P = \frac{1}{4}$, which is the average squared amplitude of an electromagnetic wave. In contrast, quantum theory models photons as discrete packets of energy, and explains the outcome of the polarizing lenses experiment as the

probability $p = \frac{1}{4}$ that a photon will pass through the circuit. For this tabletop experiment, both the classical model of light (as a wave that can be infinitely subdivided) and the quantum model of light (as discrete particles) predict the same outcome. Therefore, the polarizing lenses experiment performed with a laser pointer does not serve as proof for the necessity of a quantum mechanical description of reality.

As we discussed in the introduction to this chapter, to really reveal quantum effects, we need to look at the very small-scale or very low-power regimes. It's possible to reproduce the polarizing lenses experiment using a *single photon source*. A single photon source behaves like a super-weak laser pointer that emits only one photon at a time. When discussing the single photon regime, the classical theory of electromagnetic waves hits a wall, since a single photon cannot subdivide into parts—it is a quantum of light; a small, indivisible bundle of energy. Classical wave theory can correctly predict average optical power, but it can't provide a clear picture of what exactly happens when individual photons hit a polarizing lens.

The polarizing lenses experiment is inspired by the famous *Stern–Gerlach experiment*, which is performed with the magnetic spin of silver atoms, and which involves a similar demonstration, comparable observed outcomes, and analogous reasoning. I encourage you to learn more about the original Stern–Gerlach experiment.

[The Stern–Gerlach experiment]

https://en.wikipedia.org/wiki/Stern-Gerlach_experiment

<https://youtube.com/watch?v=rg4Fnag4V-E>

9.7 Quantum mechanics is not that weird

Without a doubt, you've heard that quantum mechanics is weird, mysterious, and generally “magical.” Well, unless vector operations count as magic, it's not *that* magical. In this section, we'll single out three so-called “weird” aspects of quantum mechanics: superposition, interference, and the fact that quantum measurements affect the states of systems being measured.

Quantum superposition

Classical binary variables (bits) can have one of two possible values: 0 or 1. Examples of physical systems that behave like bits are electric switches that can be either open or closed, digital transistors that either conduct or don't conduct electricity, and capacitors that are either charged or discharged.

A quantum bit (qubit), can be both 0 and 1 *at the same time*. Wow! Said this way, it surely sounds impressive and mystical, no? But if we use the term *linear combination* instead of “at the same time,” the quantum reality doesn’t seem so foreign. A quantum state is a linear combination of the basis states. This isn’t so crazy. The *superposition principle* is a general notion in physics that is not specific to quantum phenomena, but applies to all systems described by differential equations. Indeed, superpositions exist in many classical physics problems, too.

Example Consider a mass attached to a spring that undergoes simple harmonic motion. The differential equation that governs the motion of the mass is $x''(t) + \omega^2 x(t) = 0$. This equation has two solutions: $x_0(t) = \sin(\omega t)$ and $x_1(t) = \cos(\omega t)$, corresponding to two different starting points of the oscillation. Since both $x_0(t)$ and $x_1(t)$ satisfy the equation $x''(t) + \omega^2 x(t) = 0$, any linear combination of $x_0(t)$ and $x_1(t)$ is also a solution. Thus, the most general solution to the differential equation is of the form:

$$x(t) = \alpha x_0(t) + \beta x_1(t) = \alpha \sin(\omega t) + \beta \cos(\omega t).$$

Usually we combine the sin and cos terms and describe the equation of motion for the mass-spring system in the equivalent form $x(t) = A \cos(\omega t + \phi)$, where A and ϕ are computed from α and β . The mass-spring system might be described as undergoing both sin motion and cos motion “at the same time,” but do you see how ridiculous and incomplete this sounds?

* * *

The notion of *quantum superposition* is simply a consequence of the general superposition principle for differential equations. If the quantum states $|0\rangle$ and $|1\rangle$ both represent valid solutions to a quantum differential equation, then the state of the system can be described as a linear combination of these two solutions:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle.$$

The observation that “ $|\psi\rangle$ is both $|0\rangle$ and $|1\rangle$ at the same time” is not wrong; it’s just not very useful. It’s much more precise to describe the quantum state $|\psi\rangle$ as a linear combination.

Interference

Unlike particles that bounce off each other, waves can co-exist in the same place. The resulting wave pattern is the sum of the constituent waves. Quantum particles behave similarly to waves in certain experiments, and this can lead to *interference* between quantum systems.

The prototypical example of interference is Young's double-slit experiment, in which particles passing through two thin slits interact with each other, causing an interference pattern of alternating bright and dark spots on a screen. Classical physics models assume particles behave like tiny point-like balls that bounce off each other whenever they come in contact. A classical model predicts that particles will appear on the screen in two bright peaks, directly facing the two slits.

In contrast, the quantum model of a particle describes it as a travelling energy pulse that exhibits wave-like properties.² In a quantum model, the particles passing through the slits behave like waves and can combine constructively or destructively, depending on the relative distances travelled by the particles. Similar interference patterns occur whenever waves combine, as in the example of waves on the surface of a liquid, or sound waves.




Figure 9.18: The waves emitted by two synchronized sources form an interference pattern. Observe the stripes of destructive interference where the waves meet “out of sync” (peak to trough) and cancel each other out.

Performing Young's double-slit experiment reveals a pattern of bright and dark stripes (called fringes) on the screen in support of the quantum model. The locations of the dark fringes correspond

²This is where the name *wave function* comes from.

exactly to the places where particles passing through the two slits arrive “out of sync,” and combine destructively:

$$|\psi\rangle - |\psi\rangle = 0.$$

The locations where destructive interference occurs correspond to the dark fringes on the screen, where no particles arrive.

The idea that one wave can cancel another wave is not new. What *is* new is the observation that particles behave like waves that can interfere with each other. That’s definitely new. Indeed, interference was one of the first puzzling effects of quantum systems that was observed. Observations from interference experiments forced physicists to attribute wave-like properties to particles.

[Video demonstration of Young’s double-slit experiment]

<https://youtube.com/watch?v=qCmtegdq00A>

Measurement of a system affects the system’s state

Another seemingly weird aspect of quantum mechanics is the notion that quantum measurements can affect the states of the systems being measured. This phenomenon is not attributable to some sort of quantum magic, but is rather due to the energy scale and the size of systems where quantum physics comes into play. Let’s see why.

When we think about physical systems on the scale of individual atoms, we can no longer consider ourselves (and our physical measurement apparatuses) as passive observers of these systems. Instead, we need to account for the *interactions* between quantum systems and the measurement apparatuses used to observe them. The fact that measurements affect the state of the very systems they measure is not some magical process, but rather a consequence of the natural properties of the particles themselves. The particles we observe are affected by our measurement methods.

Wave functions

The quantum mechanics techniques we discussed in this chapter are useful for modelling physical systems that have discrete sets of states. In *matrix quantum mechanics*, quantum states are described by vectors in finite-dimensional, complex inner product spaces. Other physics problems require the use of *wave function quantum mechanics*, in which quantum states are represented as complex-valued functions of space coordinates $\vec{r} = (x, y, z)$. Instead of the dot product between vectors, the inner product for wave functions is $\langle f(\vec{r}), g(\vec{r}) \rangle = \iiint_{\mathbb{R}^3} \overline{f(\vec{r})} g(\vec{r}) d^3\vec{r}$. This may seem like a totally new ball game, but

actually calculations using wave functions are not too different from the inner product calculations we used to compute Fourier transformations in Section 7.11.

It's beyond the scope of the current presentation to discuss *wave functions* in detail, but I want to show you an example of a calculation with wave functions, so you won't say I didn't show you some proper physics stuff. The ground state of the hydrogen atom is described by the wave function $\psi(\vec{r}) = \frac{1}{\sqrt{\pi a^3}} \exp(-r/a)$, where $r = \|\vec{r}\|$. The probability of finding the electron at position $\vec{r}$ from the proton is described by the inner product $\psi(\vec{r})\psi(\vec{r}) = |\psi(\vec{r})|^2$:

$$\Pr(\{\text{finding electron at } \vec{r} \}) = |\psi(\vec{r})|^2.$$

Since $\psi(\vec{r})$ depends only on the distance r , we know the wave function has a spherically symmetric shape, as illustrated in Figure 9.19.




Figure 9.19: The s orbital of an electron is spherically symmetric.

We'll now check whether $|\psi(\vec{r})|^2$ is a properly normalized probability density function. Integrating the probability density function $|\psi(\vec{r})|^2$ over all of $\mathbb{R}^3$ should give a total probability of one. Instead of Cartesian coordinates (x, y, z) , we'll use spherical coordinates (r, ϕ, θ) to solve this problem. In spherical coordinates, the volume of a thin slice from the surface of a sphere of width $d\theta$, height $d\phi$, and thickness dr is given by $r^2 \sin \phi \, d\phi \, d\theta \, dr$. If you haven't seen spherical coordinates before, don't worry about this expression too much. The conversion factor $r^2 \sin \phi$ is a trick for converting the "small piece of volume" $d^3\vec{r} = dx \, dy \, dz$ to an equivalent small piece of volume in spherical coordinates $d^3\vec{r} = r^2 \sin \phi \, d\phi \, d\theta \, dr$. See P3.26 for the derivation.

We split the triple integral into two parts: an integral that depends only on the radius r , and double integral over the angles ϕ and θ . The total probability of finding the electron somewhere in

space is

$$\begin{aligned}
 p_{\text{total}} &= \iiint_{\mathbb{R}^3} |\psi(\vec{r})|^2 d^3\vec{r} \\
 &= \int_0^\infty \int_0^{2\pi} \int_0^\pi |\psi(r)|^2 r^2 \sin \phi \, d\phi \, d\theta \, dr \\
 &= \left(\int_0^\infty \frac{1}{\pi a^3} \exp(-2r/a) r^2 \, dr \right) \left(\int_0^{2\pi} \int_0^\pi \sin \phi \, d\phi \, d\theta \right) \\
 &= \left(\int_0^\infty \frac{1}{\pi a^3} \exp(-2r/a) r^2 \, dr \right) (4\pi) \\
 &= \int_0^\infty \underbrace{\frac{4}{a^3} \exp\left(-\frac{2r}{a}\right) r^2}_{p(r)} \, dr.
 \end{aligned}$$

The expression $p(r) = \frac{4}{a^3} \exp(-2r/a) r^2$ describes the probability of finding the electron at a distance r from the centre of the nucleus. We can complete the calculation of total probability by taking the integral of $p(r)$ from $r = 0$ to $r = \infty$:

$$\begin{aligned}
 p_{\text{total}} &= \int_0^\infty p(r) \, dr \\
 &= \frac{4}{a^3} \int_0^\infty \exp\left(-\frac{2r}{a}\right) r^2 \, dr \\
 &= 1.
 \end{aligned}$$

The purpose of working through this wave function calculation is to give you an idea of the complex calculations physicists regularly perform using the wave function formalism. In comparison, the matrix formalism for quantum mechanics is much simpler, involving only basic linear algebra calculations.

9.8 Quantum mechanics applications

What can we accomplish using quantum physics that we can't do using classical physics? What can we compute with qubits that we can't compute with bits? You've learned about the quantum formalism; but how useful is it? In this section, we'll explore some areas of physics and computer science that wouldn't exist without the laws of quantum mechanics.

Particle physics

The basic quantum mechanics formalism we learned in this chapter is not appropriate for describing the behaviour of high energy particles. The best current model for describing high energy physics is called *quantum field theory*. Quantum field theory is a generalization of basic quantum mechanics that models fundamental particles as disturbances in various *particle fields*. Recall that photons, the quanta of light, are described as disturbances of the electromagnetic field. The equations of quantum field theory describe how elementary particles and antiparticles are created and destroyed, and the various interactions between different particle fields. Quantum field theory is a bit like chemistry, where different combinations of atoms are transformed into other combinations of atoms—only instead of atoms, we have elementary particles like quarks and leptons. In the same way *Mendeleev's periodic table* is a catalogue of all available atoms, the *Standard Model* of particle physics is a catalogue of all available elementary particles. These elementary particles combine to form other, more complex particles, like protons and neutrons. The transformations of these particles are only observed in high energy physics experiments performed in particle accelerators.

High energy physics becomes important at the extreme end of an energy continuum, where quantum field theory is most applicable. In contrast, the rules of chemistry dominate the low-energy end of the energy continuum. Chemical reactions describe how molecules transform into other molecules; and molecules essentially represent the various ways electrons are shared between groups of atoms. At higher energies, atoms are “stripped” of their electrons; the atoms have so much energy, they’re no longer bound to the nucleus. At this point, the laws of chemistry are no longer relevant, since electrons and molecules move freely. Enter nuclear physics, which studies the combinations of protons and neutrons that form the nuclei of different atoms. A nuclear reaction is like a chemical reaction, but instead of chemical molecules, the reactant and products are various types of nuclei. An example of nuclear reaction is the fusion of two heavy hydrogen nuclei to form a helium nucleus. At higher energy still, even protons and neutrons can break apart, and the analysis shifts to interactions between elementary particles like leptons, bosons, neutrinos, quarks, and photons. This is the domain of high energy physics.

The basic postulates of quantum mechanics still apply in quantum field theory, but the models become more complicated since we assume even the interactions between particles are quantized. You can think of the basic quantum mechanics described in this chapter as learning the alphabet, and quantum field theory as studying Shakespeare, including the invention of new words. Studying quan-

tum field theory requires new math tools like path integrals, new intuitions like symmetry observations, and new computational tricks like *renormalization*.

The essential way of thinking about photons, electrons, and the interactions between them can be obtained by reading Richard Feynman's short book titled *QED*, which stands for *quantum electrodynamics*. In this tiny book, Feynman uses the analogy of a "tiny clock" attached to each particle to explain the phase $e^{i\theta}$ of a wave function. From this simple analogy, the author builds to explain complex concepts (path integrals, for instance) at the graduate level of quantum field theory. I highly recommended this book; it's a wonderful chance to learn from one of the great scientists in the field and one of the best physics teachers of all times.

[The Standard Model of particle physics]

https://en.wikipedia.org/wiki/Standard_Model

[BOOK] Richard P. Feynman. *QED: The strange theory of light and matter*. Princeton University Press, 2006, ISBN 0691125759.

Solid state physics

Physicists have sought to understand the inner structure of materials since the first days of physics. As they've learned, they've developed numerous applications, from semiconductors to lasers, photovoltaic batteries (solar panels), light emitting diodes (LEDs). These applications all depend on materials with specially engineered conductivity properties. Indeed, working with the conductivity of materials gives us insight into their other properties. We can classify materials into the following general conductivity-type groups: insulators, metals, and semi-conductors. These categories correspond to materials with different *energy band structures*.

Insulators are the most boring type of material, because their energy band structure doesn't permit any interesting chemical interactions. Take glass, for instance—just a clump of silica (SiO_2). The term *glass* is used in physics to describe any material made of randomly oriented molecules that lack a specific crystal structure.

Conductors are more interesting. A hand-wavy explanation of conductivity would be to say the electrons in conductors like aluminum and copper are "free to move around." Solid state physics allows for a more precise understanding of the phenomenon. Using quantum mechanical models, we can determine the energy levels that electrons can occupy, and predict how many electrons will be available to conduct electricity.

Semiconductors are the most interesting type of material since they can switch between conductive and non-conductive states. The

transistor, the invention that makes all electronics possible, consists of a sandwich of three different types of semiconductors. The voltage applied to the middle section of a transistor is called the *gate voltage*, and it controls the amount of current that can flow through the transistor. If the gate voltage is set to ON (think 1 in binary), then semiconducting material is biased; free electrons are available in its conduction band, and current can flow through. If the gate voltage is set to OFF (think 0 in binary), then the conduction band is depleted and the transistor won't conduct electricity. The improvements in semiconductor technologies—specifically the ability to pack billions of transistors into a tiny microprocessor chip—have been fuelling the ongoing computer revolution pretty much since transistors were first commercialized. In summary, no solid state physics = no mobile phones.

Quantum mechanics is used so much in solid state physics that we could suitably nickname the field “applied quantum physics.”

[Simple explanation of energy band structure and conductivity]
<https://w.wiki/inQ>

Superconductors

Certain materials exhibit surprising physical properties at very low temperatures. By low temperatures, I mean *really low*, like -272°C ; a temperature close to *absolute zero*, the temperature at which all molecules stop. You'd exhibit surprising properties too if you were placed in an environment this cold! Take for example the properties of different metals, which are generally considered to have low electric resistance. There are regular conductors like aluminum that have low resistance, and high-end conductors like silver that have even lower resistance, and then there are *superconductors* which have zero resistance. Superconductors are an example of a purely quantum phenomenon that cannot be explained by classical physics.

Some of the most iconic landmarks of modern scientific progress, like magnetic resonance imaging (MRI) machines and magnetically levitating bullet trains, are made possible by superconductor technology. Superconductors offer zero resistance to electric current, which means they can support much stronger currents than regular conductors like aluminum and silver. All these applications require extensive refrigeration equipment to keep the materials at a temperature where they behave like superconductors.

[Superconductivity]
<https://en.wikipedia.org/wiki/Superconductivity>

Quantum optics

Classical optics deals with beams of light that contain quintillions of photons. A *quintillion* is 10^{18} , which is more than a lot. When working with this many photons, it's possible to model light beams as continuous electromagnetic waves, and use classical electromagnetic theory and optics to understand experiments. Quantum optics comes into play when we perform optics experiments that involve far fewer photons, including experiments with single photons. When a single photon travels through an optical circuit, it cannot “split” like a continuous wave. For example, when a beam of light hits a half-silvered mirror, we say the beam is partially reflected. However, we can't say the same for a single photon, since the photon cannot be split. Instead, the state of the photon after it has encountered a half-silvered mirror is best described as a superposition of the “passed through” and “reflected” states, as shown in Figure 9.11 (page 491).

An example of a quantum optics effect is the *spontaneous down-conversion effect*, in which a single photon is absorbed by a material and then reemitted as two photons with an entangled polarization state:

$$|\Psi_-\rangle = \frac{1}{\sqrt{2}}|H\rangle|V\rangle - \frac{1}{\sqrt{2}}|V\rangle|H\rangle.$$

By the crystal's properties, we know one of the two emitted photons has horizontal polarization and the other has vertical polarization, but we don't know which is which. Such entangled photons can be used as starting points for other experiments that involve entanglement. Another interesting aspect of quantum optics are the so-called *squeezed states* that can be detected more accurately than regular (unsqueezed) photons.

Quantum optics is a field of active research. Scientists in academia and industry study exotic photon generation, advanced photon detection schemes, and generally explore how photons can most efficiently carry information.

[Basic principles in physics of light]

<https://materialford.wordpress.com/introduction-to-research-light/>

Quantum cryptography

Performing a quantum measurement on the state $|\psi\rangle$ tends to disturb the state. From the perspective of experimental physics, this disturbance is an obstacle since it gives us limited, one-time access to the quantum state $|\psi\rangle$, making the study of quantum states more difficult. From the perspective of cryptography however, the state-disturbing aspect of quantum measurement is an interesting and po-

tentially valuable property. If Alice transmits a secret message to Bob encoded in the state of a quantum system, it would be impossible for an eavesdropper Eve to “listen in” on the state unnoticed because Eve’s measurement would disturb the state. The *BB84 protocol*, named after its inventors Charles Bennett and Gilles Brassard, is based on this principle.

The standard basis $B_s = \{|0\rangle, |1\rangle\}$ and the Hadamard basis $B_h = \{|+\rangle, |-\rangle\}$ are *mutually unbiased bases*, which means that a basis vector from one basis lies exactly halfway between the vectors from the other basis. If we measure the state $|+\rangle$ (or $|-\rangle$) in the basis B_s , the outcomes 0 and 1 are equally unlikely. Similarly, measuring $|0\rangle$ (or $|1\rangle$) in the basis B_h produces outcomes “+” and “−” with equal probability. The use of mutually unbiased bases is central to the security of the BB84 protocol, which we’ll describe step by step:

1. Alice starts with $2n$ random candidate bits which she sends to Bob. Roughly half of the candidate bits will live on to become the final shared secret key, while the other half will be discarded. She chooses one of the bases B_s or B_h at random when encoding each bit of information she wants to send. Bob chooses to perform his measurement randomly, either in the standard basis or in the Hadamard basis. The information is transmitted correctly whenever Bob happens to pick the same basis as Alice, which happens about half the time. Whenever Bob’s basis for measurement is different from the basis Alice uses for encoding, Bob’s output is completely random.
2. Alice and Bob publicly announce the basis they used for each transmission and discard the bits where different bases were used. This leaves Alice and Bob with roughly n candidate bits of secret key.
3. Alice and Bob then publicly reveal αn of the candidate bits, which we’ll call the check bits. Assuming the quantum communication channel between Alice and Bob does not introduce any noise, Alice and Bob’s copies of the check bits should be identical, since they used the same basis. If they observe many check bits that disagree, Alice and Bob will immediately abort the protocol.
4. If the αn check bits agree, then Alice and Bob can be sure the remaining $(1 - \alpha)n$ bits they share are known only to them.

Consider what happens if the eavesdropper Eve tries to intercept the messages between Alice and Bob. Eve can measure the quantum state $|\psi\rangle$ sent by Alice, then forward to Bob the post measurement state $|\psi'\rangle$. Eve is forced to choose a basis for the measurement

she performs, and her measurement disturbs the state $|\psi\rangle$ whenever she picks a basis different from the one used by Alice. Since it's not possible to measure quantum systems without disturbing them, the eavesdropper Eve reveals her presence by introducing errors in the transmitted data. Some of the *check bits* Alice and Bob compare in Step 3 will disagree, and thus Alice and Bob will know that someone is eavesdropping on them. Though quantum mechanics does not prevent eavesdropping, it does give Alice and Bob the ability to detect when an eavesdropper is present.

The BB84 protocol established the beginning of a new field at the intersection of computer science and physics that studies *quantum key distribution* protocols. The field has developed rapidly from theory to research, and today there are even commercial quantum cryptography systems. It's interesting to compare quantum cryptography with the public key cryptography systems discussed in Section 7.9. The security of the RSA public-key encryption is based on the computational difficulty of factoring large numbers. The security of quantum cryptography is guaranteed by the laws of quantum mechanics.

[Bennett–Brassard quantum cryptography protocol from 1984]
<https://en.wikipedia.org/wiki/BB84>

[Using quantum phenomena to distribute secret keys]
https://en.wikipedia.org/wiki/Quantum_key_distribution

Quantum computing

The idea of quantum computing has existed since the early days of quantum mechanics. Richard Feynman originally proposed the idea of a *quantum simulator* in 1982, which is a quantum apparatus that can simulate the quantum behaviour of another physical system. Imagine a device that can simulate the behaviour of physical systems that would otherwise be too difficult and expensive to build. The quantum simulator would be much better at simulating quantum phenomena than any simulation of quantum mechanics on a classical computer.

Another possible application of a quantum simulator could be to encode classical mathematical optimization problems as constraints in a quantum system, then let the quantum evolution of the system “search” for good solutions. Using a quantum simulator in this way, it might be possible to find solutions to optimization problems much faster than any classical optimization algorithm could.

Once computer scientists started thinking about quantum computing, they weren't satisfied with studying optimization problems

alone, and they set out to qualify and quantify all the computational tasks that are possible with qubits. A quantum computer stores and manipulates information that is encoded as quantum states. It's possible to perform certain computational tasks on a quantum computer much faster than on any classical computer. We'll discuss *Grover's search algorithm* and *Shor's factoring algorithm* below, but first let's introduce the basic notions of quantum computing.

Quantum circuits Computer scientists like to think of quantum computing tasks as series of "quantum gates," in analogy with the logic gates used to construct classical computers. Figure 9.20 shows an example of a quantum circuit that takes two qubits as inputs and produces two qubits as outputs.




Figure 9.20: A quantum circuit that applies the Hadamard gate to the first qubit, then applies the controlled-NOT gate from the first qubit to the second qubit.

This circuit in Figure 9.20 is the combination of two quantum gates. The first operation is to apply the Hadamard gate H on the first qubit, leaving the second qubit untouched. This operation is equivalent to multiplying the input state by the matrix $H \otimes \mathbb{1}$. The second operation is called the *controlled-NOT* (or controlled- X) gate, which applies the X operator (also known as the NOT gate) to the second qubit whenever the first qubit is $|1\rangle$, and does nothing otherwise:

$$\text{CNOT}(|0\rangle \otimes |\varphi\rangle) = |0\rangle \otimes |\varphi\rangle, \quad \text{CNOT}(|1\rangle \otimes |\varphi\rangle) = |1\rangle \otimes X|\varphi\rangle.$$

The circuit illustrated in Figure 9.20 can be used to create entangled quantum states. If we input the quantum state $|00\rangle = |0\rangle \otimes |0\rangle$ into the circuit, we obtain the maximally entangled state $|\Phi_+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ as output, as depicted in Figure 9.21.

Quantum circuits can also represent quantum measurements. Figure 9.22 shows how a quantum measurement in the standard basis is represented.




Figure 9.21: Inputting $|0\rangle \otimes |0\rangle$ into the circuit produces an EPR state $|\Phi_+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ on the two output wires of the circuit.




Figure 9.22: Measurement in the standard basis $B_s = \{|0\rangle, |1\rangle\}$. The projectors of this measurement are $\Pi_0 = |0\rangle\langle 0|$ and $\Pi_1 = |1\rangle\langle 1|$.

We use double lines to represent the flow of classical information in the circuit.

Quantum registers Consider a quantum computer with a single register $|R\rangle$ that consists of three qubits. The quantum state of this quantum register is a vector in $\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2$:

$$|R\rangle = (\alpha_1|0\rangle + \beta_1|1\rangle) \otimes (\alpha_2|0\rangle + \beta_2|1\rangle) \otimes (\alpha_3|0\rangle + \beta_3|1\rangle),$$

where the tensor product $\otimes$ is used to combine the quantum states of the individual qubits. We'll call this the "physical representation" of the register and use 0-based indexing for the qubits. Borrowing language from classical computing, we'll call the rightmost qubit the *least significant* qubit, and the leftmost qubit the *most significant* qubit.

The tensor product of three vectors with dimension two is a vector with dimension eight. The quantum register $|R\rangle$ is thus a vector in an eight-dimensional vector space. The quantum state of a three-qubit register can be written as:

$$|R\rangle = a_0|0\rangle + a_1|1\rangle + a_2|2\rangle + a_3|3\rangle + a_4|4\rangle + a_5|5\rangle + a_6|6\rangle + a_7|7\rangle,$$

where a_i are complex components. We'll call this eight-dimensional vector space the "logical representation" of the quantum register. Part of the excitement about quantum computing is the huge size of the "logical space" where quantum computations take place. The logical space of a 10-qubit quantum register has dimension $2^{10} = 1024$. That's 1024 complex components we're talking about. That's a big state space for just a 10-qubit quantum register. Compare this with a 10-bit classical register, which can store one of $2^{10} = 1024$ discrete values.

We won't discuss quantum computing further here, but I still want to show you some examples of single-qubit quantum operations and their effect on the tensor product space, so you'll have an idea of the craziness that is possible.

Quantum gates Let's say you've managed to construct a quantum register; what can you do with it? Recall the single-qubit quantum operations Z , X , and H we described earlier. We can apply any of these operations on individual qubits in the quantum register. For example, applying the $X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ gate to the first (most significant) qubit of the quantum register corresponds to the following quantum operation:

$$\begin{array}{c} \text{---} \boxed{X} \text{---} \\ \text{---} \\ \text{---} \end{array} \Leftrightarrow X \otimes \mathbb{1} \otimes \mathbb{1} = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

The operator $X \otimes \mathbb{1} \otimes \mathbb{1}$ "toggles" the first qubit in the register while leaving all other qubits unchanged.

Yes, I know the tensor product operation is a bit crazy, but that's the representation of composite quantum systems and operations so please get used to it. What if we apply the X operator to the middle qubit?

$$\begin{array}{c} \text{---} \\ \text{---} \boxed{X} \text{---} \\ \text{---} \end{array} \Leftrightarrow \mathbb{1} \otimes X \otimes \mathbb{1} = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}.$$

Compare the structure of the operators $X \otimes \mathbb{1} \otimes \mathbb{1}$ and $\mathbb{1} \otimes X \otimes \mathbb{1}$. See how the action of X s affects different parts of the tensor product space $\mathbb{C}^8$?

To complete the picture, let's also see the effects of applying the X gate to the third (least significant) qubit in the register:

$$\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \boxed{X} \text{---} \end{array} \Leftrightarrow \mathbb{1} \otimes \mathbb{1} \otimes X = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}.$$

Crazy stuff, right? Don't worry, in time you'll get used to the space-within-a-space structure concept.

Okay, so what?

Quantum computers give us access to a very large state space. The fundamental promise of quantum computing is that a small set of simple quantum operations (quantum gates) can be used to perform interesting computational tasks. Sure it's difficult to interact with and manipulate quantum systems, but the space is so damn big that it's worth checking out what kind of computing you can do in there. It turns out there are already several useful things you can do using a quantum computer. The two flagship applications for quantum computing are Grover's search algorithm and Shor's factoring algorithm.

Grover's search algorithm Suppose you're given an unsorted list of n items and you want to find a particular item in that list. This is called an *unstructured search problem*. This is a hard problem to solve for a classical computer since the algorithm must parse through the entire list, which takes roughly n steps. In contrast, the unstructured problem can be solved in roughly $\sqrt{n}$ steps on a quantum computer using *Grover's search algorithm*.

The quantum speed for the unstructured search problem sure is nice, but it's really nothing to get excited about. The real money-maker for the field of quantum computing has been Shor's factoring algorithm for factoring products of prime numbers.

Shor's factoring algorithm The security of the RSA cryptosystem we discussed in Section 7.9 is based on the assumption that factoring products of large prime numbers is computationally intractable. Given the product de of two unknown prime numbers d and e , it is computationally difficult to find the factors e and d . No classical algorithm is known that can factor large numbers; even the letter agencies will have a hard time finding the factors of de when d and e are chosen to be sufficiently large prime numbers. Thus, if an algorithm that could quickly factor large numbers existed, attackers would be able to break many of the current security systems. *Shor's factoring algorithm* fits the bill, theoretically speaking.

Shor's algorithm reduces the factoring problem to the problem of *period finding*, which can be solved efficiently using the quantum Fourier transform. Shor's algorithm can factor large numbers efficiently (in polynomial time). This means RSA encryption would be easily hackable by running Shor's algorithm on a sufficiently large, and sufficiently reliable quantum computer. The letter agencies are excited about this development since they'd love to be able to hack all present-day cryptography. Can you imagine not being able to log

in securely to any website because Eve is listening in, hacking your crypto using her quantum computer?

Currently, Shor's algorithm is only a *theoretical* concern. Despite considerable effort, no quantum computers exist today that can manipulate quantum registers with thousands of qubits.

Discussion

Quantum computing certainly presents interesting possibilities, but it's a little early to imagine a quantum computing revolution in tomorrow's newspaper. As with startup ventures, it's the implementation that counts—not the idea. The current status of quantum computing as a technology is mixed. On one hand, certain quantum algorithms performed in logical space are very powerful; on the other hand, the difficulty of building a quantum computer is not to be underestimated.

It's also important to keep in mind that quantum computers are not better at solving arbitrary computational problems than the computers we already use. The problems that may benefit from a quantum speedup have a particular structure, which can be tackled with a choreographed pattern of constructive and destructive interference in quantum registers. Yet not all computationally hard problems have this structure. Quantum computing technology is at a crossroad: it could become a revolutionary development, or it could turn out that building a large-scale quantum computer is not worth the engineering challenge. So although it's cool we can execute certain tasks faster on a quantum computer, don't throw out your classical computer just yet.

Even if the quest to build a quantum computer doesn't pan out, we're certain to learn many interesting things about fundamental physics along the way. Besides, learning about the fundamental nature of quantum information is more scientifically valuable than trying to hack people's email. In the next section, we'll give an example of a new communication task that was discovered through the study of quantum information science.

Quantum teleportation Figure 9.23 illustrates a surprising aspect of quantum information: we can “teleport” a quantum state $|\psi\rangle$ from one lab to another. The quantum state $|\psi\rangle$ starts in the first qubit of the register, which is held by Alice, and ends in the third qubit, which is in Bob's lab, but there is no quantum communication channel between the two labs. This is why the term “quantum teleportation” was coined to describe this communication task, since the state $|\psi\rangle$

seems to materialize in Bob's lab like the teleportation machines used in Star Trek.

The communication resources required for the quantum teleportation protocol are one maximally entangled state shared between Alice's and Bob's labs, and two bits of classical communication from Alice to Bob. We can express the quantum teleportation protocol as a quantum circuit.




Figure 9.23: The first two qubits are in Alice's lab. The state of the first qubit $|\psi\rangle_1$ is transferred into the third qubit $|\psi\rangle_3$, which Bob controls. We say ψ is "teleported" from Alice's lab to Bob's lab because the quantum state ends up in Bob's lab, but there is no quantum communication channel connecting the labs. The state teleportation happens thanks to the pre-shared entanglement and the two bits of classical information.

The quantum teleportation protocol requires that Alice and Bob pre-share a maximally-entangled state $|\Phi_+\rangle$. By meeting in a central location, Alice and Bob can produce an entangled state using the circuit shown in Figure 9.21 (page 523). Alice and Bob then bring their respective halves of the entangled state to their labs. Note Bob's lab could be very far from Alice's lab; in another building, another city, or even on the other side of the world.

The initial state for the quantum teleportation protocol is

$$|\psi\rangle_1 \otimes |\Phi_+\rangle_{23} = (\alpha|0\rangle_1 + \beta|1\rangle_1) \otimes \left(\frac{1}{\sqrt{2}}|00\rangle_{23} + \frac{1}{\sqrt{2}}|11\rangle_{23} \right).$$

Alice has two qubits in her lab, the state $|\psi\rangle_1 = \alpha|0\rangle_1 + \beta|1\rangle_1$ and half of the entangled state, and Bob has the third qubit, which is the other half of the entangled state. At the end of the teleportation protocol, information about the state ψ appears in Bob's lab: $|\psi\rangle_3 = \alpha|0\rangle_3 + \beta|1\rangle_3$.

Without quantum communication, it seems impossible for Alice to communicate the components α and β to Bob. The pre-shared entanglement between Alice and Bob somehow enables this feat. The quantum information about the state ψ becomes available in Bob's lab as soon as Alice performs the measurement of her two qubits. But before Bob obtains the state information, he must apply a *recovery operation*, which is one of $\mathbb{1}$, X , Z , or ZX . The state information will remain unknown until Bob learns which of the four recovery operations he must perform on the state. The information about which

operation Bob should perform can be transmitted by classical means: Alice can shout the result to Bob if he's next door, tell him the results by phone, or send him a text message. After applying the needed recovery operation, Bob obtains the state $|\psi\rangle_3 = \alpha|0\rangle_3 + \beta|1\rangle_3$, which is the state that originated in the first qubit controlled by Alice.

The need for pre-shared entanglement $|\Phi_+\rangle$ between Alice and Bob is analogous to how Alice and Bob needed to pre-share a secret key $\vec{k}$ in order to use the one-time pad encryption protocol. Indeed, pre-shared entangled states are a prime resource in quantum information science. The *superdense coding protocol* is another surprising application of quantum entanglement. With this protocol, Alice can communicate two bits of classical information to Bob by sending him a single qubit and consuming one pre-shared entangled state.

Links

[Quantum simulators and practical implementations]

https://en.wikipedia.org/wiki/Quantum_simulator

[Some data about the difficulty of RSA factoring]

https://en.wikipedia.org/wiki/RSA_numbers

[Video tutorials on quantum computing by Michael Nielsen]

<http://michaelnielsen.org/blog/quantum-computing-for-the-determined/>

[Grover's algorithm for unstructured search]

https://en.wikipedia.org/wiki/Grover's_algorithm

[Shor's algorithm for factoring products of prime integers]

https://en.wikipedia.org/wiki/Shor's_algorithm

[Emerging insights on limitations of quantum computing]

<https://archive.siam.org/pdf/news/100.pdf>

Quantum error-correcting codes

Quantum states are finicky things. Every interaction of a qubit with its environment corrupts the quantum information the qubit stores. In the previous section we talked about quantum computing in the abstract, assuming the existence of an ideal noiseless quantum computer. Since the real world is a noisy place, constructing a practical quantum computer is a much greater challenge.

Recall that errors caused by noise are also a problem for classical computers. If classical computers can be made robust to errors using error-correcting codes, can we use error-correcting codes on quantum computers too? Indeed it's possible to use *quantum error-correcting codes* to defend against the effects of quantum noise. Keep

in mind, quantum error-correcting codes are more complicated to build than their classical counterparts, so it's not an obvious thing to do, but it can be done.

We won't go into too much detail, but it's worth pointing out the following interesting fact about quantum error correction. Building quantum error-correcting codes that can defend against a finite set of errors is sufficient to defend against all possible types of errors. The use of quantum error-correcting schemes is analogous to the classical error-correcting schemes we saw in Section 7.10. We encode k qubits of data that we want to protect from noise into a larger n -qubit state. The encoded state can support some number of errors before losing the data. The error-correcting procedure involves a syndrome measurement on a portion of the state, and "correction" operators applied to the remaining portion. I encourage you to follow the links below to learn more about this topic.

Building reliable quantum gates is a formidably complicated task due to the difficulty of protecting qubits from noise while simultaneously enabling quantum operations and strong interactions between qubits. It is the author's opinion that Feynman's original idea of building quantum simulators for physical systems will be the first useful applications in quantum computing.

[More on quantum error-correcting codes]

https://en.wikipedia.org/wiki/Quantum_error_correction

Quantum information theory

Classical information theory studies problems like the compression of information and the transmission of information through noisy communication channels. Quantum information theory studies the analogous problems of compression of quantum information and communication over noisy quantum channels.

The appearance of the word "theory" in "quantum information theory" should indicate that this is a mostly theoretical area of research that studies problems in the abstract. The main results of information theory are abstract theorems that may not have direct bearing on practical communication scenarios. For now, applications of quantum information theory remain in the far-off future, but that's how it is with theory subjects in general. The classical information theorems proven in the 1970s probably looked a bit useless, too; but these theorems serve as the basis of all modern wireless communications. Perhaps the purely theoretical quantum information theorems of today will solve the practical communication problems of the future.

Current efforts in quantum information theory aim to establish capacity results for quantum channels. Some of the existing results are directly analogous to classical capacity results. Other problems in quantum information theory, like the use of entanglement-assisted codes, have no classical counterparts and require a completely new way of thinking about communication problems. The book *From Classical to Quantum Shannon Theory* by Mark M. Wilde is an excellent guide to the field.

Recently, quantum theory has been applied to novel communication systems, and there is a growing interest from the communications industry to develop applications that push optical communication channels to the bounds of their theoretical efficiency. Essentially, quantum networks are being invented in parallel with quantum computers, so that when we finally build quantum computers, we'll be able to connect them together, presumably so they can share funny cat videos. What else would we use them for?

[BOOK] Mark M. Wilde. *From Classical to Quantum Shannon Theory*, Cambridge University Press, Second edition, ISBN 1107176166, <http://arxiv.org/abs/1106.1445>.

Conclusion

With this chapter, I wanted to bring you closer to the fascinating subject of quantum mechanics. I hope the material helped you understand the basic principles of quantum mechanics and clarified some of the sensational mythology surrounding the “mysteries” of the quantum world. While there's still much to discover, there's nothing too counterintuitive about quantum mechanics; it's just linear algebra, right?

One hundred years ago, quantum mechanics was seen as a foreign subject not to be trusted. In time, physicists developed good models, found better ways to explain experiments, wrote good books, and even started teaching the subject to undergraduate students. This gives me hope for humanity that we can handle even the most complex and uncertain topics when we put our minds to it.

Today we face many complex problems: consolidated corporate control of innovation, cartels, corruption, eroding democratic government systems, the militarization of everything, and conflicting ideologies. We have Sunni and Shia brothers shooting at each other and red gang versus blue gang brothers shooting at each other, and all of this for no good reason. We have all kinds of other bullshit divisions between us.

Let's hope that one hundred years from now, we'll have learned to limit the violent and corrupt aspects of human nature, so that we

can realize the potential of every child born anywhere in the world. Right now it seems like it won't be an easy change, but this is how it seemed when people were trying to figure out quantum mechanics, too. All it takes is a critical mass of people who realize and truly internalize the fact we're all on the same team, and all the divisions we see between us are pure bullshit. If we make sure that two generations of kids grow up without economic strife or bullets flying by their heads, then—I guarantee you—they will be able to figure out the rest. Together, an educated citizenry armed with the knowledge of math, history, science, and technology is more powerful than a dozen Systems combined. All the System has is an outdated hierarchical power structure and capital to pay people and make them do as they're told. We've got six billion people, the internet, and printing presses on our side. Who do you think will win?

9.9 Quantum mechanics problems

Let's recap what just happened here. Did we really cover all the topics of an introductory quantum mechanics course? Yes, we did! Thanks to your solid knowledge of linear algebra, learning the postulates of quantum mechanics took only a few dozen pages. Sure we went quickly and skipped the more physics-y topics, but we covered all the core ideas of quantum theory.

But surely it's impossible to learn quantum mechanics in such a short time? Well, you tell me. You're here. The problems are here. Prove to me you've really learned quantum mechanics by tackling the practice problems presented in this section like a boss. It's the end of the book, so don't be saving your energy. Solve these problems and then you're done.

P9.1 You work in a quantum computing startup and your boss asks you to implement the quantum gate $Q = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ i & 1 \end{bmatrix}$. Can you do it?

Hint: Recall the requirements for quantum gates.

P9.2 The Y gate is defined as $Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}$. Compute the effect of the operator YY on the vectors of the standard basis $\{|0\rangle, |1\rangle\}$.

P9.3 Compute $XHHY(\alpha|0\rangle + \beta|1\rangle)$.

Hint: Use the Hadamard gate's properties to simplify the calculation.

P9.4 Specifying an arbitrary vector $\alpha|0\rangle + \beta|1\rangle \in \mathbb{C}^2$ requires four parameters: the real and imaginary parts of α and β . Thus one might think that qubits have four *degrees of freedom*. However, the unit-length requirement and the fact that we can ignore the global phase of a qubit correspond to

additional constraints that reduce the number of degrees of freedom. How many parameters are required to specify a general quantum state $|\psi\rangle \in \mathbb{C}^2$?

P9.5 We can write any qubit using only two real parameters:

$$|\psi\rangle = \alpha|0\rangle + \sqrt{1-\alpha^2}e^{i\varphi}|1\rangle,$$

where $\alpha \in \mathbb{R}$ and $\varphi \in \mathbb{R}$. What are the ranges of values for α and φ such that all qubits can be represented?

P9.6 Consider the parametrization for qubits using two angles θ and φ :

$$|\psi\rangle = \cos(\theta/2)|0\rangle + \sin(\theta/2)e^{i\varphi}|1\rangle.$$

What values should θ and φ have for all qubits to be represented?

P9.7 Compute the products of the quantum gates HXH and HZH .

P9.8 Consider the state $|v\rangle = (a, b)^T$ and its orthogonal complement $|v^\perp\rangle = (\bar{b}, -\bar{a})^T$. The projection operators Π_v and $\Pi_{v^\perp}$ correspond to the measurements in the basis $\{|v\rangle, |v^\perp\rangle\}$. Compute the probability of outcome v in the measurement $\{\Pi_v, \Pi_{v^\perp}\}$ on the state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$.

P9.9 When we measure a quantum system $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ in the basis $\{|0\rangle, |1\rangle\}$, the Born rule tells us the probability of outcome 0 is equal to $\langle\psi|\Pi_0|\psi\rangle$. Consider this calculation that involves the trace operation:

$$\begin{aligned} \Pr(\{0\}|\psi) &= \langle\psi|\Pi_0|\psi\rangle = \langle\psi|\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}|\psi\rangle = \text{Tr}\left\{\langle\psi|\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}|\psi\rangle\right\} \\ &\stackrel{(c)}{=} \text{Tr}\left\{\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}|\psi\rangle\langle\psi|\right\} = \text{Tr}\left\{\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}\underbrace{\begin{bmatrix} |\alpha|^2 & \bar{\beta}\alpha \\ \bar{\alpha}\beta & |\beta|^2 \end{bmatrix}}_{\text{density matrix}}\right\} \\ &= \text{Tr}\left\{\begin{bmatrix} |\alpha|^2 & \bar{\beta}\alpha \\ 0 & 0 \end{bmatrix}\right\} = |\alpha|^2. \end{aligned}$$

The equality labelled (c) follows from the cyclic property of the trace operation $\text{Tr}\{ABC\} = \text{Tr}\{BCA\}$. The above calculation suggests an alternative approach for computing the probabilities of different outcomes of quantum measurements, $\Pr(\{x\}|\psi) = \text{Tr}\{\Pi_x\rho\}$, where ρ is the *density matrix* representation of the quantum state. The density matrix of the quantum state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ corresponds to the outer product $\rho = |\psi\rangle\langle\psi|$.

Calculate the probability of outcome $\{1\}$ using $\Pr(\{1\}|\psi) = \text{Tr}\{\Pi_1\rho\}$. Calculate the probabilities of the two outcomes of a measurement in the Hadamard basis $\Pr(\{+\}|\psi) = \text{Tr}\{\Pi_+\rho\}$ and $\Pr(\{-\}|\psi) = \text{Tr}\{\Pi_-\rho\}$.

P9.10 This problem explores the operation of the quantum teleportation circuit shown in Figure 9.23 (see page 527). The initial state of the three-qubits register is $|\psi\rangle_1 \otimes |\Phi_+\rangle_{23}$, where $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ is a quantum state Alice wants to send to Bob, and where $|\Phi_+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ is a maximally entangled state shared between Alice and Bob.


a) Show the following equation holds by expanding the tensor product:

$$|\psi\rangle_1 \otimes |\Phi_+\rangle_{23} = \frac{1}{\sqrt{2}} [\alpha|000\rangle_{123} + \beta|100\rangle_{123} + \alpha|011\rangle_{123} + \beta|111\rangle_{123}].$$

- b) The expression from part a) can be written as a linear combination of the four Bell states: $|\Phi_+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, $|\Phi_-\rangle = \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle)$, $|\Psi_+\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$, and $|\Psi_-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$. Verify the equation

$$\begin{aligned} & \frac{1}{\sqrt{2}}[\alpha|000\rangle_{123} + \beta|100\rangle_{123} + \alpha|011\rangle_{123} + \beta|111\rangle_{123}] \\ &= \frac{1}{2} \left[|\Phi_+\rangle_{12} |\psi\rangle_3 + |\Phi_-\rangle_{12} Z |\psi\rangle_3 + |\Psi_+\rangle_{12} X |\psi\rangle_3 + |\Psi_-\rangle_{12} XZ |\psi\rangle_3 \right]. \end{aligned}$$

- c) A *Bell measurement* consists of the combination of a controlled-NOT gate and a Hadamard gate on the first qubit, followed by measurements of both qubits in the standard basis, as illustrated in the circuit below.



Using the definition of the controlled-NOT gate (page 522) and the Hadamard gate (page 494), show that the Bell measurement performed on the state $|\Phi_+\rangle$ produces the classical measurement outcome 00. Similarly, show that measuring $|\Phi_-\rangle$ produces 10, measuring $|\Psi_+\rangle$ produces 01, and measuring $|\Psi_-\rangle$ produces 11.

- d) After Alice performs the Bell measurement on the two qubits under her control, the state of Bob's qubit will be one of the following: $|\psi\rangle_3$ if the measurement outcome is 00, $Z|\psi\rangle_3$ if outcome is 10, $X|\psi\rangle_3$ if outcome is 01, or $XZ|\psi\rangle_3$ if outcome is 11. Indicate the *recovery operation* Bob must apply in order to recover the state $|\psi\rangle_3$ in each case.

P9.11 The wave function of the electron of the hydrogen atom is $\psi(\vec{r}) = \frac{1}{\sqrt{\pi a^3}} \exp(-r/a)$. The electron's distance from the centre is described by the random variable R with probability distribution $p_R(r) = \frac{4}{a^3} \exp(-2r/a)r^2$. Calculate the expected distance of the electron $\mathbb{E}_R[R] = \int_{r=0}^{\infty} r p_R(r) dr$.

Hint: You can solve this problem using integration by parts once.

P9.12 Show that the functions $\psi_1(x) = 2x - 1$ and $\psi_2(x) = 6x^2 - 6x + 1$ are orthogonal with respect to the inner product $\langle f, g \rangle = \int_0^1 f(x)g(x) dx$.

P9.13 Consider a model of a particle in a one-dimensional box of width one. The state of the particle is described by the wave function $\psi(x)$, where $x \in [0, 1]$. Find the probability of observing x in the first quarter of the box (x between 0 and $\frac{1}{4}$) for the following wave functions: **a)** $\psi_a(x) = \sqrt{3}(2x - 1)$, **b)** $\psi_b(x) = \sqrt{5}(6x^2 - 6x + 1)$, **c)** a constant wave function ψ_c .

Hint: The probability of finding the particle somewhere in the interval $[a, b]$ is computed using the integral $\Pr(\{a \leq x \leq b\}|\psi) = \int_a^b |\psi(x)|^2 dx$.

End matter

Conclusion

By tackling the linear algebra concepts in this book, you've proven you can handle computational complexity, develop geometric intuition, and understand abstract math ideas. These are precisely the types of skills you'll need in order to understand more advanced math concepts, build scientific models, and develop useful applications. Congratulations on taking this important step toward your mathematical development. Throughout this book, we learned about vectors, linear transformations, matrices, abstract vector spaces, and many other math concepts that are useful for building math models.

Mathematical models serve as a highly useful common core for all sciences, and the techniques of linear algebra are some of the most versatile modelling tools that exist. Every time you use an equation to characterize a real-world phenomenon, you're using your math modelling skills. Whether you're applying some well-known scientific model to describe a phenomenon or developing a new model specifically tailored to a particular application, the deeper your math knowledge, the better the math models you'll be able to leverage. Let's review and catalogue some of the math modelling tools we've learned about, and see how linear algebra fits into a wider context.

To learn math modelling, you must first understand basic math concepts such as numbers, equations, and functions $f : \mathbb{R} \rightarrow \mathbb{R}$. Once you know about functions, you can use different formulas $f(x)$ to represent, model, and predict the values of real-world quantities. Working with functions is the first modelling superpower conferred on people who become knowledgeable in math. For example, understanding the properties of the function $f(x) = Ae^{-x/B}$ in the abstract enables you to describe the expected number of atoms remaining in a radioactive reaction $N(t) = N_0 e^{-\gamma t}$, predict the voltage of a discharging capacitor over time $v(t) = V_0 e^{-\frac{t}{RC}}$, and understand the ex-

ponential probability distribution $p_X(x) = \lambda e^{-\lambda x}$.

To further develop your math modelling skills, the next step is to generalize the concepts of inputs x , outputs y , and functions f to other input-output relationships. In linear algebra, we studied functions of the form $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ that obey the linear property:

$$T(\alpha \vec{x}_1 + \beta \vec{x}_2) = \alpha T(\vec{x}_1) + \beta T(\vec{x}_2).$$

This linear structure enables us to study the properties of many functions, solve equations involving linear transformations, and build useful models for many applications (some of which we discussed in Chapter 7). The mathematical structure of a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ can be *represented* as multiplication by a matrix $M_T \in \mathbb{R}^{m \times n}$. The notion of matrix representations ($T \Leftrightarrow M_T$) was central throughout this book. Even if you forget the computational procedures we learned, the idea of representations should stick with you, and you should be able to recognize representations in many contexts. That's a big deal, because most advanced math topics involve studying the parallels between different abstract notions. Understanding linear transformations and their representations is an important first step toward advanced math topics.

The computational skills you learned in Chapter 3 are also useful; though you probably won't be solving any problems by hand using row operations from this point forward, since computers outclass humans on matrix arithmetic tasks. Good riddance. Until now, you did all the work and used SymPy to check your answers. From now on, you can let SymPy do all the calculations and your job will be to chill.

If you didn't skip the sections on abstract vector spaces, you know about the parallels between the vector space $\mathbb{R}^4$ and the abstract vector spaces of third-degree polynomials $a_0 + a_1x + a_2x^2 + a_3x^3$ and 2×2 matrices $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$. This is another step up the ladder of abstraction, as it deepens your understanding of all math objects with vector-like structure.

It was my great pleasure to be your guide through the subject of linear algebra. I hope you walk away from this book with a solid understanding of how the concepts of linear algebra fit together. In the book's introduction, I likened linear algebra to playing with LEGOs. Indeed, if you feel comfortable manipulating vectors and matrices, performing change-of-basis operations, and using the matrix decomposition techniques to see inside matrices, you'll be able to "play" with all kinds of complex systems and problems. For example, consider the linear transformation T that you want to apply to an input vector $\vec{v}$. Suppose the linear transformation T is most easily described in the basis B' , but the vector $\vec{v}$ is expressed with respect to

the basis B . “No problem,” you can say, and proceed to build the following chain of matrices that compute the output vector $\vec{w}$:

$$[\vec{w}]_B = {}_B[\mathbb{1}]_{B' \ B'} [A_T]_{B' \ B'} [{}_B[\mathbb{1}]_B [\vec{v}]_B.$$

Do you see how matrices and vectors fit together neatly like LEGOs?

I can’t tell you what the next step on your journey will be. With your new linear algebra modelling skills, a thousand doors have opened for you; now you must explore and choose. Will you learn how to code and start a software company? Maybe you’ll use your analytical skills to go to Wall Street and destroy the System from the inside. Or perhaps you’ll apply your modelling skills to revolutionize energy generation, thus making human progress sustainable. Regardless of your choice of career, I hope you’ll stay on good terms with math and continue learning whenever you have the chance. Good luck with your studies!

Social stuff

Be sure to contact me if you have any feedback about this book. It helps to hear which parts of the book readers like, hate, or don’t understand. I consider all feedback in updating and improving future editions of this book. This is how the book got good in the first place—lots of useful feedback from readers. You can reach me by email at ivan@minireference.com.

Another thing you can do to help us is to review the book on Amazon.com, Goodreads, Google Books, or otherwise spread the word about the **No Bullshit Guide** textbook series. Talk to your friends and let them in on the knowledge buzz.

If you want to know what Minireference Co. has been up to, check out our blog at minireference.com/blog/. The blog is a mix of 30% publishing technology talk, 50% startup business talk, and 20% announcements. Checking the blog is the easiest way to follow the progress of our revolution in the textbook industry. For real; we’re totally serious about making education accessible through affordable textbooks. You can also connect via Twitter [@minireference](https://twitter.com/minireference) and Facebook [fb.me/noBSguide](https://facebook.com/noBSguide).

Acknowledgements

This book would not have been possible without the support of my parents, and my teachers: Claude Crépeau, Patrick Hayden, Paul Kenton, Benoit Larose, and Ioannis Psaromiligkos. It’s by listening to

them that I gained the confidence to learn and later teach advanced math topics. Patrick Hayden deserves a particular mention in this regard since his ability to explain concepts clearly inspired me to tackle complicated problems using a steady, methodical approach. Patrick also showed me that it's possible to trick students into learning even the most advanced topics by presenting the material in a logical and structured manner. Many thanks also to David Avis, Arlo Breault, Michael Hilke, Igor Khavkine, Felix Kwok, Juan Pablo Di Lelle, Ivo Panayotov, and Mark M. Wilde for their support with the book.

The errorlessness and consistency of the text would not have been possible without the help of my editor Sandy Gordon, who did a great job at polishing the text until it flowed. Truly no bullshit is allowed into the book when Sandy's on watch. Many thanks to Polina Anis'kina who helped me to create the problem sets for the book.

General linear algebra links

Below are some useful links to resources where you can learn more about linear algebra. We covered a lot of ground, but linear algebra is endless. Don't sit on your laurels and think you're the boss now that you've completed this book and its problem sets. You have the tools, but you need to practice using them. Try reading about the same topics from some other sources. See if you can do the problem sets in another linear algebra textbook. Try to use linear algebra in the coming year and further solidify your understanding of the material.

[Video lectures of Gilbert Strang's linear algebra class at MIT]

<http://ocw.mit.edu/courses/mathematics/18-06-linear-algebra-spring-2010>

[The essence of linear algebra video playlist by 3Blue1Brown]

http://bit.ly/essence_of_LA

[A free online textbook with amazing interactive visualizations]

<http://immersivemath.com/ila/index.html>

[Lecture notes by Terrence Tao]

<http://www.math.ucla.edu/~tao/resource/general/115a.3.02f/>

[Wikipedia overview on matrices]

[https://en.wikipedia.org/wiki/Matrix_\(mathematics\)](https://en.wikipedia.org/wiki/Matrix_(mathematics))

[Linear algebra wikibook (with solved problems)]

https://en.wikibooks.org/wiki/Linear_Algebra

[Proofs involving linear algebra]

http://proofwiki.org/wiki/Category:Linear_Algebra

[Linear algebra from first principles using diagrams only]

<https://graphicallylinearalgebra.net/>

Appendix A

Answers and solutions

Chapter 1 solutions

Answers to exercises

E1.1 a) $x = 3$; b) $x = 30$; c) $x = 2$; d) $x = -3$. **E1.2** a) $\mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$; b) $\mathbb{C}$; c) $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$; d) $\mathbb{Q}, \mathbb{R}, \mathbb{C}$; e) $\mathbb{R}, \mathbb{C}$. **E1.3** a) 21; b) 0; c) $\frac{2}{27}$. **E1.4** a) $x = 2$; b) $x = 25$; c) $x = 100$. **E1.5** a) $f^{-1}(x) = x^2$, $x = 16$. b) $g^{-1}(x) = -\frac{1}{2} \ln(x)$, $x = 0$. **E1.6** a) $(x-1)(x-7)$; b) $(x+2)^2$; c) $(x+3)(x-3)$. **E1.7** a) $x^2 + 2x - 15 = (x+1)^2 - 16 = 0$, which has solutions $x = 3$ and $x = -5$; b) $x^2 + 4x + 1 = (x+2)^2 - 3 = 0$, with solutions $x = -2 + \sqrt{3}$ and $x = -2 - \sqrt{3}$. **E1.8** $x_1 = \frac{3}{2}$ and $x_2 = -1$. **E1.9** $x = \pm\sqrt{2}$. **E1.10** Domain: $\mathbb{R}$. Image: $[-2, 2]$. Roots: $\{\dots, -\frac{\pi}{2}, \frac{\pi}{2}, \frac{3\pi}{2}, \frac{5\pi}{2}, \dots\}$. **E1.11** a) $p(x)$ is even and has degree 4. b) $q(x)$ is odd and has degree 7. **E1.12** a) $x = 5$ and $x = -3$; b) $x = 1 + \sqrt{3}$ and $x = 1 - \sqrt{3}$. **E1.13** $x = \sqrt{21}$. **E1.14** $V = 33.51$ and $A = 50.26$. **E1.15** Length of track $= 5C = 5\pi d = 11.47$ m. **E1.16** $x = 5 \cos(45^\circ) = 3.54$, $y = 5 \sin(45^\circ) = 3.54$; $C = 10\pi$. **E1.17** a) $\frac{\pi}{6}$ rad; b) $\frac{\pi}{4}$ rad; c) $\frac{\pi}{3}$ rad; d) $\frac{3\pi}{2}$ rad. **E1.18** a) -1 ; b) 1 ; c) 0 . **E1.19** a) 0 ; b) 1 ; c) $\frac{1}{2}$; d) 1 . **E1.20** a) $(4, 0)$. b) $(-2, -3)$. c) $(7, 3)$. **E1.21** a) $\vec{v}_1 = (5\sqrt{3}, 5) = (8.66, 5)$. b) $\vec{v}_2 = (0, -12)$. c) $\vec{v}_3 = (-2.95, 0.52)$. **E1.22** a) $\vec{u}_1 = 4\angle 0^\circ$. b) $\vec{u}_2 = \sqrt{2}\angle 45^\circ$. c) $\vec{u}_3 = \sqrt{10}\angle 108.43^\circ$. **E1.23** a) $(1, \frac{1}{2})$. b) $(1, 2)$. c) $(-2, 2)$. **E1.24** $x = 2$, $y = 3$. **E1.25** $x = 5$, $y = 6$, and $z = -3$. **E1.26** $p = 7$ and $q = 3$. **E1.27** a) $\{2, 4, 6, 7\}$; b) $\{1, 2, 3, 4, 5, 6\}$; c) $\{1, 3, 5\}$; d) $\emptyset$; e) $\{1, 2, 3, 4, 5, 6, 7\}$; f) $\{7\}$; g) $\{2, 4, 6, 7\}$; h) $\emptyset$. **E1.28** a) $(-\infty, \frac{3}{2})$; b) $(-\infty, -5]$; c) $(-1, 4)$; d) $(4, \infty)$; e) $[\frac{14}{3}, \infty)$; f) $(-\infty, -4] \cup [2, \infty)$.

Solutions to selected exercises

E1.12 a) Rewrite the equation putting all terms on the right-hand side: $0 = x^2 - 2x - 15$. We can factor this quadratic by inspection. Are there numbers a and b such that $a + b = -2$ and $ab = -15$? Yes, $a = -5$ and $b = 3$, so $0 = (x-5)(x+3)$. b) Rewrite the equation so all terms are on the left-hand side: $3x^2 - 6x - 6 = 0$. Nice, the cubic terms cancel! We'll use the quadratic formula to solve this equation $x = \frac{6 \pm \sqrt{(-6)^2 - 4(3)(-6)}}{6} = \frac{6 \pm 6\sqrt{3}}{6} = 1 \pm \sqrt{3}$.

E1.13 The cosine rule tells us $x^2 = 4^2 + 5^2 - 2(4)(5) \cos(60^\circ) = 21$. Therefore $x = \sqrt{21}$.

E1.14 The volume of the sphere with radius $r = 2$ is $V = \frac{4}{3}\pi 2^3 = 33.51$. Its surface area is $A = 4\pi 2^2 = 50.26$.

E1.17 To convert an angle measure from degrees to radians we must multiply it by the conversion ratio $\frac{\pi}{180}$.

E1.23 See <https://www.desmos.com/calculator/ocedywekcl> for plots.

E1.28 a) Dividing both sides of the inequality by two gives $x < \frac{3}{2}$. **b)** Divide both sides by negative four to obtain $x \leq -5$. Note the " $\geq$ " changed to " $\leq$ " since we divided by a negative number. **c)** If the absolute value of $(2x - 3)$ is less than five, then $(2x - 3)$ must lie in the interval $(-5, 5)$. We can therefore rewrite the inequality as $-5 < 2x - 3 < 5$, then add three to both sides to obtain $-2 < 2x < 8$, and divide by two to obtain the final answer $-1 < x < 4$. **d)** Let's collect all the x -terms on the right and all the constants on the left: $8 < 2x$, which leads to $4 < x$. **e)** To simplify, add two to both sides of the inequality to obtain $\frac{1}{2}x \geq \frac{1}{3} + 2$. You remember how to add fractions right? We have $\frac{1}{3} + 2 = \frac{1}{3} + \frac{6}{3} = \frac{7}{3}$, and therefore $\frac{1}{2}x \geq \frac{7}{3}$. Multiply both sides by two to obtain $x \geq \frac{14}{3}$. **f)** The first step is to get rid of the square by taking the square root operation on both sides: $\sqrt{(x+1)^2} \geq \sqrt{9}$. Recall that $\sqrt{x^2} = |x|$, so we have $|x+1| \geq 3$. There are two ways for the absolute value of $(x+1)$ to be greater than three. Either $x+1 \geq 3$ or $x+1 \leq -3$. We subtract one in each of these inequalities to find $x \geq 2$ or $x \leq -4$. The solution to this inequality is the union of these two intervals.

Answers to problems

P1.1 $x = \pm 4$. **P1.2** $x = A \cos(\omega t + \phi)$. **P1.3** $x = \frac{ab}{a+b}$. **P1.4 a)** 2.2795. **b)** 1024. **c)** -8.373. **d)** 11. **P1.5 a)** $\frac{3}{4}$. **b)** $-\frac{141}{35}$. **c)** $3\frac{23}{32}$. **P1.6 a)** c. **b)** 1. **c)** $\frac{9|a|}{|b|}$. **d)** a. **e)** $\frac{b}{ac}$. **f)** $x^2 + ab$. **P1.7 a)** $x^2 + (a-b)x - ab$. **b)** $2x^2 - 7x - 15$. **c)** $10x^2 + 31x - 14$. **P1.8 a)** $(x-4)(x+2)$. **b)** $3x(x-3)(x+3)$. **c)** $(x+3)(6x-7)$. **P1.9 a)** $(x-2)^2 + 3$. **b)** $2(x+3)^2 + 4$. **c)** $6(x + \frac{11}{12})^2 - \frac{625}{24}$. **P1.10** \$0.05. **P1.11** 5 years later. **P1.12** girl = 80 nuts, boy = 40 nuts. **P1.13** Alice is 15. **P1.14** 18 days. **P1.15** After 2 hours. **P1.16** $\phi = \frac{1+\sqrt{5}}{2}$. **P1.17** $x = \frac{-5+\sqrt{41}}{2}$. **P1.18 a)** $x = \sqrt[3]{2}$. **b)** $x = (\frac{\pi}{2} + 2\pi n)$ for $n \in \mathbb{Z}$. **P1.19** No real solutions if $0 < m < 8$. **P1.20 a)** e^z . **b)** $\frac{x^3 y^{15}}{z^3}$. **c)** $\frac{1}{4x^4}$. **P1.21** For $n > 250$, Algorithm Q is faster. **P1.22** 10 cm. **P1.23** 22.52 in. **P1.24** $d = \frac{1800 \tan 20^\circ - 800 \tan 25^\circ}{\tan 25^\circ - \tan 20^\circ}$, $h = 1658.46$ m. **P1.25** $x = \tan \theta \sqrt{a^2 + b^2 + c^2}$. **P1.26** $\sin^2 \theta \cos^2 \theta = \frac{1 - \cos 4\theta}{8}$. **P1.27** $c = \frac{a \sin 75^\circ}{\sin 41^\circ} \approx 14.7$. **P1.28 a)** $h = a \sin \theta$. **b)** $A = \frac{1}{2}ba \sin \theta$. **c)** $c = \sqrt{a^2 + b^2 - 2ab \cos(180^\circ - \theta)}$. **P1.29** 1.06 cm. **P1.30** $A_{\text{rect}} = 5c + 10$. **P1.31** $V_{\text{box}} = 1.639$ L. **P1.32** $V = 300\,000$ L. **P1.33** 315 000 L. **P1.34** 4000 L. **P1.35** A rope of length $\sqrt{2}\ell$. **P1.36** 20 L of water. **P1.37** $h = 7.84$ inches. **P1.38** $1 + 2 + \cdots + 100 = 50 \times 101 = 5050$. **P1.39** $x = -2$ and $y = 2$. **P1.40** $x = 1$, $y = 2$, and $z = 3$. **P1.41 a)** $\vec{u}_1 = 5\angle 90^\circ$. **b)** $\vec{u}_2 = \sqrt{5}\angle 63.4^\circ$. **c)** $\vec{u}_3 = \sqrt{5}\angle 243.4^\circ$ or $\sqrt{5}\angle -116.6^\circ$. **P1.42 a)** $\vec{v}_1 = (17.32, 10)$. **b)** $\vec{v}_2 = (0, -10)$. **c)** $\vec{v}_3 = (-4.33, 2.5)$. **P1.43 a)** $\vec{w}_1 = 9.06\hat{i} + 4.23\hat{j}$. **b)** $\vec{w}_2 = -7\hat{j}$. **c)** $\vec{w}_3 = 3\hat{i} - 2\hat{j} + 3\hat{k}$. **P1.44 a)** (3, 4). **b)** (0, 1). **c)** (7.33, 6.5). **P1.45** $Q = (5.73, 4)$. **P1.46 a)** 2i. **b)** $\frac{1}{4}(5 + i)$. **c)** $2 + i$. **P1.47 a)** $x = \pm 2i$. **b)** $x = -16$. **c)** $x = -1 - i$ and $x = -1 + i$. **d)** $x = i$, $x = -i$, $x = \sqrt{3}i$, and $x = -\sqrt{3}i$. **P1.48 a)** $\sqrt{5}$. **b)** $\frac{1}{2}(-3 + i)$. **c)** $-5 - 5i$.

Solutions to selected problems

P1.5 For **c)**, $1\frac{3}{4} + 1\frac{31}{32} = \frac{7}{4} + \frac{63}{32} = \frac{56}{32} + \frac{63}{32} = \frac{119}{32} = 3\frac{23}{32}$.

P1.9 The solutions for **a)** and **b)** are fairly straightforward. To solve **c)**, we first factor out 6 from the first two terms to obtain $6(x^2 + \frac{11}{6}x) - 21$. Next we choose half of the

coefficient of the linear term to go inside the square and add the appropriate correction to maintain equality: $6[x^2 + \frac{11}{6}x] - 21 = 6[(x + \frac{11}{12})^2 - (\frac{11}{12})^2] - 21$. After expanding the rectangular brackets and simplifying, we obtain the final expression: $6(x + \frac{11}{12})^2 - \frac{625}{24}$.

P1.11 We must solve for x in $35 + x = 4(5 + x)$. We obtain $35 + x = 20 + 4x$, then $15 = 3x$, so $x = 5$.

P1.13 Let A be Alice's age and B be Bob's age. We're told $A = B + 5$ and $A + B = 25$. Substituting the first equation into the second we find $(B + 5) + B = 25$, which is the same as $2B = 20$, so Bob is 10 years old. Alice is 15 years old.

P1.14 The first shop can bind $4500/30 = 150$ books per day. The second shop can bind $4500/45 = 100$ books per day. The combined production capacity rate is $150 + 100 = 250$ books per day. It will take $4500/250 = 18$ days to bind the books when the two shops work in parallel.

P1.15 Let t_m denote the time when the two planes meet, as measured from the moment the second plane departs. Since it left one hour earlier, the slower plane will have travelled a distance $600(t_m + 1)$ km when they meet. The faster plane will have travelled the distance $900t_m$ km when they meet. Combining the two expressions we find $600(t_m + 1) = 900t_m$. The time when the planes meet is $t_m = 2$ hours after the departure of the second plane.

P1.19 Using the quadratic formula, we find $x = \frac{m \pm \sqrt{m^2 - 8m}}{4}$. If $m^2 - 8m \geq 0$, the solutions are real. If $m^2 - 8m < 0$, the solutions will be complex numbers. Factoring the expressions and plugging in some numbers, we observe that $m^2 - 8m = m(m - 8) < 0$ for all $m \in (0, 8)$.

P1.21 The running time of Algorithm Q grows linearly with the size of the problem, whereas Algorithm P's running time grows quadratically. To find the size of the problem when the algorithms take the same time, we solve $P(n) = Q(n)$, which is $0.002n^2 = 0.5n$. The solution is $n = 250$. For $n > 250$, the linear-time algorithm (Algorithm Q) will take less time.

P1.24 Observe the two right-angle triangles drawn in Figure 1.70. From the triangle with angle 25° we know $\tan 25^\circ = \frac{h}{800+d}$. From the triangle with angle 20° we know $\tan 20^\circ = \frac{h}{1800+d}$. We isolate h in both equations and eliminate h by equating $(1800 + d) \tan 25^\circ = \tan 20^\circ (800 + d)$. Solving for d we find $d = \frac{1800 \tan 20^\circ - 800 \tan 25^\circ}{\tan 25^\circ - \tan 20^\circ} = 2756.57$ m. Finally we use $\tan 25^\circ = \frac{h}{800+d}$ again to obtain $h = \tan 25^\circ (800 + d) = 1658.46$ m.

P1.26 We know $\sin^2(\theta) = \frac{1}{2}(1 - \cos(2\theta))$ and $\cos^2(\theta) = \frac{1}{2}(1 + \cos(2\theta))$, so their product is $\frac{1}{4}(1 - \cos(2\theta)\cos(2\theta))$, and $\cos(2\theta)\cos(2\theta) = \cos^2(2\theta)$. Using the power-reduction formula on the term $\cos^2(2\theta)$ gives $\sin^2 \theta \cos^2 \theta = \frac{1}{4} \left(1 - \frac{1}{2}(1 + \cos(4\theta))\right)$.

P1.29 The volume of the water stays constant and is equal to 1000 cm^3 . Initially the height of the water h_1 can be obtained from the formula for the volume of a cylinder $1000 \text{ cm}^3 = h_1 \pi (8.5 \text{ cm})^2$, so $h_1 = 4.41$ cm. After the bottle is inserted, the water has the shape of a cylinder with a cylindrical part missing. The volume of water is $1000 \text{ cm}^3 = h_2 (\pi (8.5 \text{ cm})^2 - \pi (3.75 \text{ cm})^2)$. We find $h_2 = 5.47$ cm. The change in height is $h_2 - h_1 = 5.47 - 4.41 = 1.06$ cm.

P1.30 The rectangle's area is equal to its length times its height $A_{\text{rect}} = \ell h$.

P1.31 The box's volume is $V = w \times h \times \ell = 10.5 \times 7 \times 22.3 = 1639 \text{ cm}^3 = 1.639 \text{ L}$.

P1.33 The tank's total capacity is $15 \times 6 \times 5 = 450 \text{ m}^3$. If 30% of its capacity is spent, then 70% of the capacity remains: 315 m^3 . Knowing that $1 \text{ m}^3 = 1000 \text{ L}$, we find there are 315 000 L in the tank.

P1.34 The first tank contains $\frac{1}{4} \times 4000 = 1000$ L. The second tank contains three times more water, so 3000 L. The total is 4000 L.

P1.35 The amount of wood in a pack of wood is proportional to the area of a circle $A = \pi r^2$. The circumference of this circle is equal to the length of the rope $C = \ell$. Note the circumference is proportional to the radius $C = 2\pi r$. If we want double the area, we need the circle to have radius $\sqrt{2}r$, which means the circumference needs to be $\sqrt{2}$ times larger. If we want a pack with double the wood, we need to use a rope of length $\sqrt{2}\ell$.

P1.36 In 10 L of a 60% acid solution there are 6 L of acid and 4 L of water. A 20% acid solution will contain four times as much water as it contains acid, so 6 L acid and 24 L water. Since the 10 L we start from already contains 4 L of water, we must add 20 L.

P1.37 The document must have a 768/1004 aspect ratio, so its height must be $6 \times \frac{1004}{768} = 7.84375$ inches.

P1.38 If we rewrite $1 + 2 + 3 + \cdots + 98 + 99 + 100$ by pairing numbers, we obtain the sum $(1 + 100) + (2 + 99) + (3 + 98) + \cdots$. This list has 50 terms and each term has the value 101. Therefore $1 + 2 + 3 + \cdots + 100 = 50 \times 101 = 5050$.

P1.49 There exists at least one banker who is not a crook. Another way of saying the same thing is “not all bankers are crooks”—just *most* of them.

P1.50 Everyone steering the ship at Monsanto ought to burn in hell, forever.

P1.51 **a)** Investors with money but without connections. **b)** Investors with connections but no money. **c)** Investors with both money and connections.

Chapter 2 solutions

Answers to exercises

E2.1 $A^{-1} = \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & \frac{1}{2} \end{bmatrix}$. **E2.2** **a)** $A\vec{v} = \begin{bmatrix} 7 \\ 14 \end{bmatrix}$; **b)** $B\vec{v} = \begin{bmatrix} -1 \\ 9 \end{bmatrix}$; **c)** $A(B\vec{v}) = \begin{bmatrix} 26 \\ 41 \end{bmatrix}$; **d)** $B(A\vec{v}) = \begin{bmatrix} -7 \\ 63 \end{bmatrix}$; **e)** $A\vec{w} = \begin{bmatrix} -15 \\ -32 \end{bmatrix}$; **f)** $B\vec{w} = \begin{bmatrix} 3 \\ -21 \end{bmatrix}$. **E2.3** $v_1 = -2$, $v_2 = 3$. **E2.4** **a)** $(1, 1, 3)$; **b)** $(1, 1, -3)$; **c)** $(3, 3, 3)$; **d)** $\sqrt{2}$. **E2.5** **a)** 5; **b)** $(-1, -1, 1)$; **c)** $(1, 1, -1)$; **d)** $(0, 0, 0)$. **E2.6** **a)** This part has been omitted for brevity. **b)** $\vec{v}_1 = (9.848, 1.736)$; $\vec{v}_2 = (8.66, 5)$; $\vec{v}_3 = (5, 8.66)$; $\vec{v}_4 = (-5, 8.66)$. **c)** $\Pi_i(\vec{v}_1) = 9.848$; $\Pi_i(\vec{v}_2) = 8.66$; $\Pi_i(\vec{v}_3) = 5$; $\Pi_i(\vec{v}_4) = -5$. **d)** $\Pi_j(\vec{v}_1) = 1.736$; $\Pi_j(\vec{v}_2) = 5$; $\Pi_j(\vec{v}_3) = 8.66$; $\Pi_j(\vec{v}_4) = 8.66$. **e)** $\Pi_{\vec{d}}(\vec{v}_1) = (5.79, 5.79)$ and $\|\Pi_{\vec{d}}(\vec{v}_1)\| = 8.19$; $\Pi_{\vec{d}}(\vec{v}_2) = (6.83, 6.83)$ and $\|\Pi_{\vec{d}}(\vec{v}_2)\| = 9.66$; $\Pi_{\vec{d}}(\vec{v}_3) = (6.83, 6.83)$ and $\|\Pi_{\vec{d}}(\vec{v}_3)\| = 9.66$; $\Pi_{\vec{d}}(\vec{v}_4) = (1.83, 1.83)$ and $\|\Pi_{\vec{d}}(\vec{v}_4)\| = 2.59$. **E2.7** **a)** $\begin{bmatrix} 3 & 2 \\ 4 & 1 \end{bmatrix}$; **b)** $\begin{bmatrix} -2 & 2 \\ 3 & -2 \\ 0 & 1 \end{bmatrix}$; **c)** $\begin{bmatrix} 17 & 16 \\ 8 & 9 \end{bmatrix}$; **d)** $\begin{bmatrix} 13 & 12 & 11 & 10 \\ 2 & 3 & 4 & 5 \end{bmatrix}$; **e)** $\begin{bmatrix} 2 & 1 & 4 \\ -2 & 4 & 1 \end{bmatrix}$; **f)** Doesn't exist; **g)** $\begin{bmatrix} -2 & -6 \\ 5 & 10 \\ 2 & 1 \end{bmatrix}$; **h)** -5; **i)** Doesn't exist; **j)** Doesn't exist; **k)** -5; **l)** 1; **m)** 4; **n)** 4. **E2.8** **a)** 14; **b)** 5; **c)** 0; **d)** $\begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \\ 3 & 6 & 9 \end{bmatrix}$; **e)** $\begin{bmatrix} 4 & -2 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$; **f)** $\begin{bmatrix} 2 & -1 & 0 \\ 4 & -2 & 0 \\ 6 & -3 & 0 \end{bmatrix}$. **E2.9** $\alpha = -\frac{1}{2}$ and $\beta = 12$. **E2.10** **a)** Yes; **b)** No; **c)** Yes.

Solutions to selected exercises

E2.1 To find A^{-1} we must consider the action of $A = \begin{bmatrix} 7 & 0 \\ 0 & 2 \end{bmatrix}$ on an arbitrary vector $\vec{v} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}$, and perform the inverse action. Since A multiplies the first component by 7, A^{-1} must divide the first component by 7. Since A multiplies the second component by 2, A^{-1} must divide the second component by 2. Thus $A^{-1} = \begin{bmatrix} \frac{1}{7} & 0 \\ 0 & \frac{1}{2} \end{bmatrix}$.

E2.10 An expression is linear in the variable v if it contains v raised only to the first power. This is the case for the first and third expressions but not the second, since it contains $\sqrt{x} = x^{\frac{1}{2}}$.

Answers to problems

P2.1 a) $q(x)$ is nonlinear; b) $f(x)$, $g(x)$, and $h(x)$ are all linear; c) $i(x)$ is nonlinear; d) $j(x)$ is nonlinear. **P2.2** $(1, 2, 3)$. **P2.3** $|a\rangle + |b\rangle = 5|0\rangle + 2|1\rangle$. **P2.4** a) 0; b) $(0, 0, 1)$; c) $(0, 1, -1)$; d) $(0, 0, -1)$. **P2.5** a) 5; b) $(-1, 1, 1)$; c) $(0, 0, 0)$; d) $(0, 0, 0)$. **P2.6** a) 6. b) 0. c) -3 . d) $(-2, 1, 1)$. e) $(3, -3, 0)$. f) $(7, -5, 1)$. **P2.7** a) $(2, 3, 3, 7, 8)$. b) $(0, -1, -3, -1, -2)$. c) 30. **P2.8** $(-\frac{2}{3}, \frac{1}{3}, \frac{2}{3})$ or $(\frac{2}{3}, -\frac{1}{3}, -\frac{2}{3})$. **P2.9** $(12, -4, -12)$. **P2.10** The tractor's trajectory is a half-circle. The total distance travelled is 3.14 km. **P2.12** $M\vec{v} = \begin{bmatrix} az_1 + \beta z_2 \\ \gamma z_1 + \delta z_2 \end{bmatrix}$. **P2.13** $A = \begin{bmatrix} 3 & 0 \\ 0 & 1 \end{bmatrix}$; $B = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$; $C = \begin{bmatrix} 3 & 1 \\ 0 & 1 \end{bmatrix}$. $AB = \begin{bmatrix} 3 & 3 \\ 0 & 1 \end{bmatrix}$; $BA = \begin{bmatrix} 3 & 1 \\ 0 & 1 \end{bmatrix} = C$. **P2.14** a) $\begin{bmatrix} -5 & -5 \\ 4 & 2 \end{bmatrix}$; b) $\begin{bmatrix} -5 & 10 & -5 \\ 20 & 5 & 10 \end{bmatrix}$; c) $\begin{bmatrix} 17 & 28 \\ 41 & 64 \end{bmatrix}$; d) Doesn't exist; e) $\begin{bmatrix} 18 & 21 \\ 9 & 12 \\ 8 & 9 \end{bmatrix}$; f) $\begin{bmatrix} 9 & 12 \end{bmatrix}$; g) $\begin{bmatrix} -4 \\ -7 \end{bmatrix}$; h) 0; i) $\begin{bmatrix} 2 & 1 \\ -4 & -2 \end{bmatrix}$; j) 8; k) 11; l) -3 ; m) 20. **P2.15** a) $\det(B)$; b) $\det A$; c) 0. **P2.16** a) 6; b) $\frac{3}{2}$; c) 15; d) 30.

Solutions to selected problems

P2.1 A function is linear in x if it contains x raised only to the first power. Basically, $f(x) = mx$ (for some constant m) is the only possible linear function of one variable.

P2.8 See `bit.ly/1c0a8yo` for calculations.

P2.9 Any multiple of the vector $\vec{u}_1 \times \vec{u}_2 = (-3, 1, 3)$ is perpendicular to both $\vec{u}_1$ and $\vec{u}_2$. We must find a multiplier $t \in \mathbb{R}$ such that $t(-3, 1, 3) \cdot (1, 1, 0) = 8$. Computing the dot product we find $-3t + t = 8$, so $t = -4$. The vector we're looking for is $(12, -4, -12)$. See `bit.ly/1nmYH8T` for calculations.

P2.10 The direction of the tractor changes constantly throughout the day, and the overall trajectory has the shape of a half-circle. The total distance travelled by the tractor is equal to half the circumference of a circle of radius R . Since it took the tractor six hours of movement at $v = 0.524$ km/h to travel half the circumference of the circle, we have $\frac{1}{2}C = \pi R = v(t_f - t_i) = 0.524(6)$, from which we find $R = 1$ km. The total distance travelled by the tractor is $\pi R = 3.14$ km.

P2.11 Using algebra we find $\|\vec{u} - \vec{v}\|^2 = \|\vec{u}\|^2 + \|\vec{v}\|^2 - 2\vec{u} \cdot \vec{v}$. Using the cosine rule we find $\|\vec{u} - \vec{v}\|^2 = \|\vec{u}\|^2 + \|\vec{v}\|^2 - 2\|\vec{u}\|\|\vec{v}\|\cos(\varphi)$. Equating these two expressions, we can obtain the geometric formula for the cosine product.

P2.13 Using the definition of the matrix vector product, we can imitate the action of each linear transformation T by choosing appropriate entries in the matrix. The top row corresponds to the first component of the output; the bottom row corresponds to the second component of the output. Observe that $BA = C$. The composite transformation of applying T_A followed by T_B (denoted $T_B \circ T_A$), is equivalent to the transformation T_C . Note $AB \neq BA$: the matrix product AB corresponds to composing the composition of the linear transformations in the opposite order $T_A \circ T_B$.

Chapter 3 solutions

Answers to exercises

E3.1 $x = 4$, $y = -2$. **E3.3** a) No solution; b) $(0, 2)$; c) $\{(2, 0) + s(-1, 1), \forall s \in \mathbb{R}\}$. **E3.4** a) $X = BA^{-1}$; b) $X = C^{-1}B^{-1}A^{-1}ED^{-1}$; c) $X = AD^{-1}$. **E3.5** $P = \begin{bmatrix} 19 & 22 \\ 43 & 50 \end{bmatrix}$;

$Q = \begin{bmatrix} -5 & 9 \\ 8 & 6 \end{bmatrix}$. **E3.6** **a)** $\begin{bmatrix} 19 & 24 \\ 40 & 51 \end{bmatrix}$; **b)** $\begin{bmatrix} 57 & 88 \\ 11 & 24 \end{bmatrix}$; **c)** $\begin{bmatrix} 17 & 28 \\ 41 & 64 \end{bmatrix}$; **d)** $\begin{bmatrix} 54 & 69 \\ 22 & 27 \end{bmatrix}$; **e)** Doesn't exist; **f)** $\begin{bmatrix} 18 & 21 \\ 9 & 12 \\ 8 & 9 \end{bmatrix}$; **g)** -165. **E3.7** **a)** -2; **b)** 2; **c)** 4; **d)** 0. **E3.8** $V = 2$. **E3.9** Linearly independent. **E3.10** A^{-1} exists for all $\alpha \neq 6$. **E3.11** $A^{-1} = \begin{bmatrix} 2 & -1 \\ -1 & 1 \end{bmatrix}$. **E3.12** $x = -8$ and $y = 3$.

Solutions to selected exercises

E3.2 The row operations required to bring **A** to reduced row echelon form are: $R_1 \leftarrow \frac{1}{3}R_1$, $R_2 \leftarrow R_2 - 2R_1$, $R_2 \leftarrow -2R_2$, $R_1 \leftarrow R_1 - R_2$. Using SymPy these operations are implemented as follows:

```
>>> A[0,:] = A[0,:]/3
>>> A[1,:] = A[1,:] - 2*A[0,:]
>>> A[1,:] = -2*A[1,:]
>>> A[0,:] = A[0,:] - A[1,:]
```

Try displaying the matrix **A** after each operation to watch the progress.

E3.8 See bit.ly/181ugMm for the calculations.

E3.9 We can determine if the vectors are linearly independent by constructing a 3×3 matrix from them and computing the determinant. In this case, we have $\det\left(\begin{bmatrix} 1 & 4 & 3 \\ 2 & 1 & 1 \\ 0 & -2 & -1 \end{bmatrix}\right) = -3 \neq 0$ so the vectors are linearly independent.

E3.12 We can solve the equation $A \begin{bmatrix} x \\ y \end{bmatrix} = \vec{b}$ by multiplying the equation by the inverse A^{-1} . We find $A^{-1} = \begin{bmatrix} -2 & -3 \\ 1 & 1 \end{bmatrix}$ and $\begin{bmatrix} x \\ y \end{bmatrix} = A^{-1}\vec{b} = \begin{bmatrix} -8 \\ 3 \end{bmatrix}$.

Answers to problems

P3.1 $x = 15$ and $y = 2$. **P3.2** **a)** Alice: $R_2 \leftarrow R_2 - 2R_1$, $R_2 \leftarrow -2R_2$, $R_1 \leftarrow R_1 - R_2$; **b)** Bob: $R_2 \leftarrow R_2 - 2R_1$, $R_2 \leftarrow -\frac{2}{3}R_2$, $R_1 \leftarrow R_1 - \frac{3}{2}R_2$; **c)** Charlotte: $R_1 \leftarrow \frac{1}{2}R_1$, $R_2 \leftarrow R_2 - 3R_1$, $R_2 \leftarrow \frac{4}{3}R_2$, $R_1 \leftarrow R_1 - \frac{3}{4}R_2$. **P3.3** **a)** $(-2, 2)$; **b)** $(-4, -1, -2)$; **c)** $(\frac{-2}{5}, \frac{-1}{2}, \frac{3}{5})$. **P3.4** **a)** $\{(2, 0) + s(-2, 1), \forall s \in \mathbb{R}\}$; **b)** $\{(2, 1, 0) + s(3, 1, 1), \forall s \in \mathbb{R}\}$; **c)** $\{(\frac{1}{10}, 0, \frac{3}{5}) + \alpha(1, 1, 0), \forall \alpha \in \mathbb{R}\}$. **P3.5** **a)** $\{(1, 0, 0) + s(1, 1, 0) + t(2, 0, 1), \forall s, t \in \mathbb{R}\}$; **b)** $\{(\frac{1}{10}, 0, \frac{3}{5}, 0) + \alpha(1, 1, 0, 0) + \beta(-\frac{1}{10}, 0, -\frac{3}{5}, 1), \forall \alpha, \beta \in \mathbb{R}\}$. **P3.6** **a)** $\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in$

$\left\{ \begin{bmatrix} 0 \\ 0 \end{bmatrix} + t \begin{bmatrix} -1 \\ 1 \end{bmatrix}, \forall t \in \mathbb{R} \right\}$. **b)** $\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0 \\ -2 \\ 5 \end{bmatrix}$. **c)** $\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \left\{ \begin{bmatrix} \frac{3}{2} \\ 0 \\ 0 \end{bmatrix} + t \begin{bmatrix} 0 \\ -2 \\ 1 \end{bmatrix}, \forall t \in \mathbb{R} \right\}$. **P3.7**

$\text{rank}(A) \leq 4$. **P3.8** $C = B^{-1}$. **P3.9** **a)** $M^{-1}L^{-1}MK^2$; **b)** $J^{-3}K^{-1}J^2$; **c)** $A = \frac{1}{2}I$; **d)** $Y = N^{-1}$. **P3.10** $\vec{x} = (-7, -19, -3)^T$. **P3.11** $\vec{x} = (30.64, 10.48, 17.06)$.

P3.12 **a)** $AB = \begin{bmatrix} 6 & 2 \\ 10 & 2 \\ 12 & 3 \end{bmatrix}$. **b)** $AA = \begin{bmatrix} 4 & 6 & 3 \\ 6 & 18 & 8 \\ 3 & 7 & 5 \end{bmatrix}$. **c)** BA doesn't exist. **d)** BB doesn't exist. **P3.13** $\begin{bmatrix} -2 & -2 \\ -15 & -15 \end{bmatrix}$. **P3.15** **a)** $\begin{bmatrix} 13 & 5 & 3 \\ 6 & 9 & 12 \\ 6 & 5 & 16 \end{bmatrix}$; **b)** $\begin{bmatrix} -6 & 7 & 5 & -5 \\ -9 & 15 & 18 & -3 \\ -7 & 19 & 10 & 5 \end{bmatrix}$; **c)** $\begin{bmatrix} 4 & -4 \\ 24 & 3 \\ 18 & 17 \end{bmatrix}$;

d) $\begin{bmatrix} -8 & -5 & -8 \\ 15 & 7 & 23 \\ 3 & 9 & 9 \\ -1 & -3 & 7 \end{bmatrix}$; **e)** $\begin{bmatrix} 4 & 11 & 18 \\ 6 & -1 & 22 \end{bmatrix}$; **f)** $\begin{bmatrix} -51 & 92 & 55 & -10 \\ -49 & 58 & 85 & -40 \end{bmatrix}$. **P3.16** **a)** $\begin{bmatrix} 0 & \cos(\alpha) - \sin(\alpha) \\ \sin(\alpha) - \cos(\alpha) & -\cos(2\alpha) \end{bmatrix}$;

b) $\begin{bmatrix} \cos^2(\alpha) & -\sin(\alpha) \\ 2\sin(\alpha) - \cos(\alpha) & \sin^2(\alpha) \end{bmatrix}$; **c)** $\begin{bmatrix} \cos(2\alpha) & -\sin(\alpha) \\ 2\sin(\alpha) - \cos(\alpha) & 0 \end{bmatrix}$. **P3.17** **a)** -3; **b)** 0; **c)** 10. **P3.18**

$\det(A) = -48$; $\det(B) = 13$. **P3.19** **a)** Independent; **b)** Dependent; **c)** Dependent. **P3.20** Area = 2. **P3.21** Volume = 8. **P3.22** **a)** -2; **b)** -162; **c)** -8; **d)** -14; **e)** 28. **P3.23** **a)** 86; **b)** -86; **c)** -172. **P3.24** For both rows and columns: **A**: not independent; **B**: independent; **C**: not independent; **D**: not independent. **P3.25** $\det(J) = r$. **P3.26**

$|\det(J_s)| = \rho^2 \sin \phi$. **P3.27 a)** The inverse doesn't exist; **b)** $\begin{bmatrix} 5 & -2 \\ -2 & 1 \end{bmatrix}$; **c)** $\begin{bmatrix} 2 & -\frac{3}{2} \\ -1 & 1 \end{bmatrix}$.

P3.28 $B = \begin{bmatrix} -17 & -30 \\ 5 & 8 \end{bmatrix}$. **P3.29** $A^{-1} = \begin{bmatrix} -\frac{1}{3} & \frac{2}{3} & -\frac{1}{3} \\ -\frac{2}{3} & \frac{1}{3} & -\frac{2}{3} \\ \frac{4}{3} & -\frac{2}{3} & \frac{7}{3} \end{bmatrix}$ and $B^{-1} = \frac{1}{21} \begin{bmatrix} -3 & -5 & -\frac{21}{6} & 11 \\ 3 & -2 & 7 & -4 \\ 9 & 15 & \frac{21}{2} & -12 \\ 3 & -9 & 0 & 3 \end{bmatrix}$.

P3.31 $C_A = \begin{bmatrix} 1 & 2 & -4 \\ -2 & -1 & 2 \\ 1 & 5 & -7 \end{bmatrix}$; $C_B = \begin{bmatrix} -10 & 6 & -12 \\ -4 & -10 & 20 \\ 1 & 18 & -5 \end{bmatrix}$; $C_C = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 2 & -2 & 2 & -2 \\ -4 & 4 & -4 & 4 \\ -8 & 8 & -8 & 8 \end{bmatrix}$. **P3.32**

$A^{-1} = \begin{bmatrix} -1 & 1 & -2 \\ -1 & \frac{1}{2} & -\frac{5}{2} \\ 2 & -1 & 4 \end{bmatrix}$. **P3.33** $a = -3, b = 1, c = 2, d = -2$. **P3.34** $\begin{bmatrix} 1 & -1 \\ 2 & 2 \end{bmatrix} \begin{bmatrix} -1 & -1 \\ 1 & 2 \end{bmatrix}$ or $\begin{bmatrix} -1 & 1 \\ -2 & -2 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ -1 & 2 \end{bmatrix}$. **P3.35** $B = SA$ and $C = AS$, where $S = \begin{bmatrix} \alpha & 0 \\ 0 & \beta \end{bmatrix}$.

Solutions to selected problems

P3.7 Since the system of equations $A\vec{x} = \vec{b}$ has an infinite number of solutions, the RREF of A must contain at least one row of zeros. Therefore, the rank of A can be at most $5 - 1 = 4$.

P3.8 First simplify the equation by multiplying with A^{-1} from the left, and with D^{-1} from the right, to obtain $BC = \mathbb{1}$. Now we can isolate C by multiplying with B^{-1} from the left. We obtain $B^{-1} = C$.

P3.11 Start by rewriting the matrix equations as $(\mathbb{1} - A)\vec{x} = \vec{d}$, then solve for $\vec{x}$ by hitting the equation with the appropriate inverse: $\vec{x} = (\mathbb{1} - A)^{-1}\vec{d}$. See bit.ly/1hg44Ys for the details of the calculation.

P3.14 First rewrite H as $\frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$ to simplify calculations. Then computing HXH gives us $HXH = \frac{1}{2} \left(\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \right) \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 2 & 0 \\ 0 & -2 \end{bmatrix} = Z$. The calculation for HZH is similar and it leads to X .

P3.19 We can determine if the sets of vectors are linearly independent by combining them to form a matrix then computing this matrix's determinant. If the determinant of the matrix is nonzero, the vectors are linearly independent.

P3.23 The answers in **a)** and **b)** have different signs because interchanging rows in a matrix changes the sign of the determinant. For part **c)**, we use the fact that multiplying one row of a matrix by a constant has the effect of multiplying the determinant by the same constant.

P3.24 We can calculate the determinant of the matrix to check if its rows are independent. If the determinant is not zero then vectors are independent, otherwise vectors are dependent. The columns of a square matrix are linearly independent if and only if the rows of the matrix are linearly independent.

P3.25 The determinant of J is

$$\begin{vmatrix} \cos \theta & -r \sin \theta \\ \sin \theta & r \cos \theta \end{vmatrix} = r \cos^2 \theta + r \sin^2 \theta = r.$$

P3.26 The determinant of J_s is

$$\begin{aligned} \det(J_s) &= \begin{vmatrix} \sin \phi \cos \theta & -\rho \sin \phi \sin \theta & \rho \cos \phi \cos \theta \\ \sin \phi \sin \theta & \rho \sin \phi \cos \theta & \rho \cos \phi \sin \theta \\ \cos \phi & 0 & -\rho \sin \phi \end{vmatrix} \\ &= \rho^2 \left(\cos^2 \phi \sin \phi (-1) - \sin^3 \phi (1) \right) \\ &= -\rho^2 \sin \phi \left(\cos^2 \phi + \sin^2 \phi \right) = -\rho^2 \sin \phi. \end{aligned}$$

Since we're only interested in finding the *volume factor* we can ignore the sign of the Jacobian's determinant: $|\det(J_s)| = \rho^2 \sin \phi$.

P3.28 To solve for the matrix B in the equation $AB = C$, we must get rid of the matrix A on the left side. We do this by multiplying the equation $AB = C$ by the inverse A^{-1} . We find the inverse of A by starting from the array $[A \mid \mathbf{1}]$, and performing the row operations $\mathcal{R}_1 : R_2 \leftarrow R_2 - 2R_1$, $\mathcal{R}_2 : R_2 \leftarrow -R_2$, and $\mathcal{R}_3 : R_1 \leftarrow R_1 - 4R_2$, to find the matrix $A^{-1} = \begin{bmatrix} -7 & 4 \\ 2 & -1 \end{bmatrix}$. Applying A^{-1} to both sides of the equation, we find $B = A^{-1}C = \begin{bmatrix} -17 & -30 \\ 5 & 8 \end{bmatrix}$.

P3.30 Zero matrix has the $\det(A) = 0$. We have $A^{-1} = \frac{1}{\det(A)}\text{adj}(A)$. We cannot divide by zero, so the zero matrix has no inverse.

P3.32 See `bit.ly/matinvxls` for the solution.

P3.33 Find the inverse of $\begin{bmatrix} 1 & 3 \\ -2 & -1 \end{bmatrix}$, then multiply both sides of the equation by the inverse to isolate the matrix of unknowns $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$.

P3.34 First we calculate $\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} e & f \\ g & h \end{bmatrix} = \begin{bmatrix} ae+bg & af+bh \\ ce+dg & cf+dh \end{bmatrix} = \begin{bmatrix} 2ae & ae+ec \\ ce+ca & c^2+ec \end{bmatrix}$. We proceed by comparing the two matrices entry by entry. Observe $2ae = -2$, which implies $ae = -1$, which allows us to simplify $ae + ec = -3$ to $ec = -2$. Substituting $ec = -2$ into $ce + c^2 = 2$, we find $c^2 = 4$, so there are two possibilities: $c = 2$ or $c = -2$. When $c = 2$, we obtain the answer $c = d = h = 2$, $a = g = 1$, and $e = b = -1$. When $c = -2$, we find $c = d = h = -2$, $a = g = -1$, and $e = b = 1$.

Chapter 4 solutions

Answers to exercises

E4.1 $d(\ell, O) = \frac{3\sqrt{2}}{2} \approx 2.121$. **E4.2** $d(\ell, O) = 3$. **E4.3** $d(P, O) = 5\sqrt{3} \approx 8.66$. **E4.4** $2x + y = 5$. **E4.5** a) $\sqrt{10}$; b) 2; c) $\frac{\sqrt{3}}{3}$; d) $2\sqrt{3}$; e) $\sqrt{3}$; f) $\frac{\sqrt{3}}{3}$. **E4.6** a) $(\frac{4}{3}, \frac{4}{3}, \frac{4}{3})$; b) $(-\frac{1}{3}, \frac{5}{3}, -\frac{4}{3})$; c) $(\frac{5}{3}, -\frac{1}{3}, -\frac{4}{3})$; d) $(\frac{5}{3}, \frac{5}{3}, \frac{5}{3})$; e) $(-\frac{5}{3}, -\frac{5}{3}, \frac{10}{3})$. **E4.7** a) $(\frac{2}{5}, \frac{26}{5}, 5)$; b) 11.841; c) $(\frac{14}{5}, \frac{32}{5}, 0)$; d) 12.837; e) 5.6745. The point p is closer to ℓ_1 than to ℓ_2 .

E4.8 $(\frac{30}{53}, \frac{5}{53}, -\frac{20}{53})$. **E4.9** $\vec{v} = (4, 3, 2)_W$. **E4.10** $B_m[\mathbf{1}]_{B_s} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 \\ \frac{1}{2} & -\frac{1}{2} & 0 \\ 0 & 0 & \frac{1}{2} \end{bmatrix}$. **E4.11** There

must be four vectors in the set and the vectors must be linearly independent. **E4.12** a) Yes; b) Yes; c) Yes; d) No; e) Yes. **E4.13** $\text{rank}(A) = \dim(\mathcal{R}(A)) = \dim(\mathcal{C}(A)) = 3$, $\dim(\mathcal{N}(A)) = 1$, and $\dim(\mathcal{N}(A^T)) = 0$. $\text{rank}(B) = \dim(\mathcal{R}(B)) = \dim(\mathcal{C}(B)) = 3$, $\dim(\mathcal{N}(B)) = 2$, and $\dim(\mathcal{N}(B^T)) = 1$. $\text{rank}(C) = \dim(\mathcal{R}(C)) = \dim(\mathcal{C}(C)) = 0$, $\dim(\mathcal{N}(C)) = 4$, and $\dim(\mathcal{N}(C^T)) = 3$. **E4.14** $\mathcal{R}(A) = \text{span}((1, 0, 0, 0), (0, 1, -1, 0))$, $\mathcal{C}(A) = \text{span}((1, 0, 0)^T, (0, 1, 0)^T)$, and $\mathcal{N}(A) = \text{span}((0, 1, 1, 0)^T, (0, 0, 0, 1)^T)$. **E4.15** Yes. **E4.16** $\mathcal{N}(A) = \text{span}((\frac{1}{2}, 1, \frac{1}{2}, 1))$; $\mathcal{N}(B) = \text{span}((1, 1, 0), (2, 0, 1))$. **E4.17** $\mathcal{R}(A) = \text{span}((1, 3, 0, 0), (0, 0, 1, 0), (0, 0, 0, 1))$, $\mathcal{C}(A) = \text{span}((1, 2, 3)^T, (3, 7, 9)^T, (3, 6, 10)^T)$, and $\mathcal{N}(A) = \text{span}((-3, 1, 0, 0)^T)$.

Solutions to selected exercises

E4.1 Using the formula from page 217, we find

$$\begin{aligned} d(\ell, O) &= \left\| (0, 3) - \frac{(0, 3) \cdot (1, -1)}{1^2 + (-1)^2} (1, -1) \right\| \\ &= \left\| (0, 3) - \frac{-3}{2} (1, -1) \right\| \\ &= \left\| (\frac{3}{2}, \frac{3}{2}) \right\| = \frac{3\sqrt{2}}{2}. \end{aligned}$$

E4.2 Using the formula from page 217, we find

$$d(\ell, O) = \left\| (0, 0, 3) - \frac{(0, 0, 3) \cdot (1, 1, 0)}{1^2 + 1^2 + 0^2} (1, 1, 0) \right\| = 3.$$

E4.5 a) This is the norm of $\|r\|$. **b)** To find a vector from the origin to the closest point on the line ℓ , find a vector to any point on the line and subtract the part of the vector that is parallel to the line. The length of this perpendicular-only vector is the closest distance $d(\ell, O)$. **c)** Start with an arbitrary point in the plane, say $p_P = (1, 0, 0)$. Then compute the length of the projection of p_P in the normal direction to obtain a vector to the point closest to the origin in the plane P . The length of this vector is equal to $d(P, O)$. **d)** The procedure is analogous to part (b), but we'll use a vector $\vec{u} = (0, 0, 2) - (1, 3, 0)$, which starts at r and ends on a point on the line ℓ . The part of the vector $\vec{u}$ that is perpendicular to the line ℓ is computed as $\Pi_{\ell^\perp}(\vec{u}) = \vec{u} - \Pi_\ell(\vec{u})$. The closest distance between r and ℓ is the length of this vector, $\|\Pi_{\ell^\perp}(\vec{u})\|$. **e)** The procedure is analogous to part (c), but we start with the vector $\vec{v} = (1, 0, 0) - (1, 3, 0)$, which starts at r and ends on a point in the plane P . The closest distance between r and P is $\|\Pi_{\bar{n}_P}(\vec{v})\|$. **f)** The line ℓ is parallel to the plane P . Define a vector that starts at some point on the line and ends at some point in the plane $\vec{w} = (1, 0, 0) - (0, 0, 2)$; then compute $\|\Pi_{\bar{n}_P}(\vec{w})\|$.

E4.6 See bit.ly/projxcrcis for the solution using SymPy.

E4.7 See bit.ly/dist_lines_exrcis for the solution.

E4.8 The vectors $(2, -4, 2)$ and $(6, 1, -4)$ define a plane P . We're looking for the projection of $\vec{v}$ onto P . First, use the cross product to find a normal vector to the plane P : $\vec{n} = (2, -4, 2) \times (6, 1, -4) = (14, 20, 26)$. Then, compute the projection using the formula $\Pi_P(\vec{v}) = \vec{v} - \Pi_{\vec{n}}(\vec{v}) = (\frac{30}{53}, \frac{5}{53}, \frac{-20}{53})$.

E4.10 Find the inverse change of basis matrix $B_s[\mathbb{1}]_{B_m} = \begin{bmatrix} 1 & 1 & 0 \\ 1 & -1 & 0 \\ 0 & 0 & 2 \end{bmatrix}$, whose columns are the vectors of the basis B_m . Then compute $B_m[\mathbb{1}]_{B_s} = (B_s[\mathbb{1}]_{B_m})^{-1} = \frac{1}{2} \begin{bmatrix} 1 & 1 & 0 \\ 1 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$.

E4.12 All the given subsets are subspaces of $\mathbb{R}^3$ except for **d)** because the plane $2x + y + z = 3$ does not contain the zero vector $(0, 0, 0)$.

E4.15 The null space of A and $-A$ are identical since $A\vec{v} = \vec{0}$ if and only if $-A\vec{v} = \vec{0}$. The same goes for the left null spaces. The column spaces are the same since any vector $\vec{w} \in \mathcal{C}(A)$ corresponds to some choice of coefficients $\vec{v}$ such that $\vec{w} = A\vec{v}$, and since using $-A$ and $-\vec{v}$ gives the same vector $(-A)(-\vec{v}) = \vec{w}$. The same goes for the row spaces.

E4.16 The null space of A consists of all vectors $\vec{x} = (x_1, x_2, x_3, x_4)^T$ that satisfy $A\vec{x} = \vec{0}$. To solve this equation we first compute the RREF of A :

$$\text{rref}(A) = \begin{bmatrix} 1 & 0 & 0 & -\frac{1}{2} \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -\frac{1}{2} \end{bmatrix},$$

and observe it pivots in the first three columns, while $x_4 = s$ is a free variable. The null space of A is

$$\begin{bmatrix} 1 & 0 & 0 & -\frac{1}{2} \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -\frac{1}{2} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ s \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \Rightarrow \begin{aligned} 1x_1 - \frac{1}{2}s &= 0 \\ 1x_2 - s &= 0 \\ x_3 - \frac{1}{2}s &= 0 \end{aligned}$$

The null space is $\mathcal{N}(A) = \text{span}((\frac{1}{2}, 1, \frac{1}{2}, 1))$.

Answers to problems

P4.1 a) (1, 2). **b)** Lines overlap so every point on the lines is an intersection point.
c) (1, 1). **P4.2 a)** $\{(\frac{1}{2}, -\frac{1}{4}, 0) + s(0, -\frac{1}{2}, 1), \forall s \in \mathbb{R}\}$; **b)** $\{(-1, 2, 0) + t(1, -1, 1), \forall t \in \mathbb{R}\}$.
P4.3 a) parallel; **b)** neither; **c)** perpendicular. **P4.4** $d(r, P) = 1$. **P4.5** $d(p, Q) = \frac{2}{3}$. **P4.6**
a) $d(p, q) = 7$; **b)** $d(m, n) = 5$; **c)** $d(r, s) = 3$; **d)** $d(i, j) = \sqrt{19}$. **P4.7** $x + y + 2z = 4$.
P4.8 $\ell : \left\{ \frac{x+3}{2} = \frac{y-1}{-4} = \frac{z-1}{-1} \right\}$. **P4.9** $2x - 11y - 7z = 2$. **P4.10** $\Pi_{\vec{u}}(\vec{v}) = \frac{\vec{v} \cdot \vec{u}}{\|\vec{u}\|^2} \vec{u} =$
 $\frac{(1,1,1) \cdot (2,1,-1)}{2^2+1^2+(-1)^2} (2,1,-1) = \frac{1}{3} (2,1,-1)$; $\Pi_{\vec{v}}(\vec{u}) = \frac{\vec{u} \cdot \vec{v}}{\|\vec{v}\|^2} \vec{v} = \frac{(1,1,1) \cdot (2,1,-1)}{1^2+1^2+1^2} (1,1,1) =$
 $\frac{2}{3} (1,1,1)$. **P4.11** $\Pi_P(\vec{v}) = \frac{1}{7} (17, 30, -1)$. **P4.12** $\Pi_{P^\perp}(\vec{u}) = (-\frac{6}{25}, 0, -\frac{8}{25})$. **P4.13**
 $d(\ell, P) = \frac{7}{3}$. **P4.14** $\vec{v} = (2, 1, 3)_W$. **P4.15** ${}_V[1]_U = \begin{bmatrix} -1 & 2 & 0 \\ 1 & -1 & 1 \\ 1 & -1 & -1 \end{bmatrix}$. **P4.17 a)** Yes; **b)** No;
c) Yes; **d)** No; **e)** No. **P4.18 a)** $A = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \end{bmatrix}$; **b)** $B = \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{bmatrix}$; **c)** $C = \begin{bmatrix} 3 & 6 & 9 \\ -1 & -2 & -3 \end{bmatrix}$.
Other answers are possible. **P4.19** $d = \frac{bc}{a}$. **P4.20 a)** The x -axis; **b)** The xy -plane;
c) The xy -plane; **d)** All of $\mathbb{R}^3$; **e)** The yz -plane. **P4.22** No, but you can conclude that
 $\dim V \leq m$. **P4.29** $\vec{v}_3 = (0, 1, 0)$, $\vec{v}'_3 = (2, 4, 7)$. **P4.30** $\vec{v}_4 = (5, 2, -3, 1)$.

Solutions to selected problems

P4.1 To find the intersection point, solve the two equations simultaneously.

P4.3 We can solve this problem by comparing the plane normals. **a)** The normals to the planes are $\vec{n}_1 = (1, -1, -1)$ and $\vec{n}_2 = (2, -2, -2) = 2\vec{n}_1$. The normals are multiples of each other so the planes are parallel. **b)** $\vec{n}_1 = (3, 2, 0)$, $\vec{n}_2 = (0, 1, -1)$, and $\vec{n}_1 \cdot \vec{n}_2 = 2 \neq 0$. The planes are neither parallel nor perpendicular. **c)** $\vec{n}_1 = (1, -2, 1)$, $\vec{n}_2 = (1, 1, 1)$, and $\vec{n}_1 \cdot \vec{n}_2 = 0$; therefore the planes are perpendicular.

P4.4 Use the formula $d(r, P) = \frac{|\vec{n} \cdot \vec{r}|}{\|\vec{n}\|} = \frac{|2(2) + 1(3) - 2(5)|}{\sqrt{2^2 + 1^2 + (-2)^2}} = \frac{3}{3} = 1$.

P4.5 Construct the vector $\vec{v} = p - q = (5, 3, 5) - (0, 1, 0) = (5, 2, 5)$, which starts at q (an arbitrary point in the plane Q) and ends at the point p . The shortest distance between the point p and the plane Q is equal to the length of the projection of the vector $\vec{v}$ in the direction of the plane's normal vector $\vec{n}$, which is given by the formula $\|\Pi_{\vec{n}}(\vec{v})\|$. We find $d(p, Q) = \left\| \frac{\vec{v} \cdot \vec{n}}{\|\vec{n}\|^2} \vec{n} \right\| = \left\| \frac{10 + 2 - 10}{2^2 + 1^2 + (-2)^2} (2, 1, -2) \right\| = \left\| \frac{2}{9} (2, 1, -2) \right\| = \frac{2}{3}$.

P4.7 First we find two vectors in the plane; for example $\vec{u} = r - q = (-1, -1, 1)$ and $\vec{v} = s - q = (0, -2, 1)$. Then we must find the normal vector $\vec{n}$, $\vec{n} = \vec{u} \times \vec{v} = (-1, -1, 1) \times (0, -2, 1) = (1, 1, 2)$. We can use any of the three points as the point p_0 in the geometric equation of the plane $\vec{n} \cdot [(x, y, z) - p_0] = 0$. Using $q = (1, 3, 0)$, we obtain the equation $(1, 1, 2) \cdot [(x, y, z) - (1, 3, 0)] = 1(x - 1) + 1(y - 3) + 2z = 0$. Computing the expression gives $x - 1 + y - 3 + 2z = 0$, which simplifies to $x + y + 2z = 4$.

P4.9 We use the Gauss-Jordan elimination procedure to find the intersection of the two planes. The line of intersection is $\ell_1 : \{(1, 0, 0) + (1, -3, 5)t \mid \forall t \in \mathbb{R}\}$, where $(1, 0, 0)$ is a point on the line of intersection and $\vec{v}_1 = (1, -3, 5)$ is its direction vector. The direction vector of ℓ_2 is $\vec{v}_2 = (2, 1, -1)$. We want to find the equation of a plane $\vec{n} \cdot [(x, y, z) - p_0] = 0$ that has a normal perpendicular to both $\vec{v}_1 = (1, -3, 5)$ and $\vec{v}_2 = (2, 1, -1)$. Using $\vec{n} = \vec{v}_1 \times \vec{v}_2 = (-2, 11, 7)$, and choosing the point $p_0 = (1, 0, 0)$ from the line ℓ_1 , we obtain $(-2, 11, 7) \cdot [(x, y, z) - (1, 0, 0)] = 0$.

P4.11 Let's first find $\Pi_{P^\perp}(\vec{v}) = \frac{\vec{v} \cdot \vec{n}}{\|\vec{n}\|^2} \vec{n} = \frac{(3, 4, 1) \cdot (2, -1, 4)}{2^2 + (-1)^2 + 4^2} (2, -1, 4) = \frac{2}{7} (2, -1, 4)$. Then $\Pi_P(\vec{v}) = \vec{v} - \Pi_{P^\perp}(\vec{v}) = (3, 4, 1) - \frac{2}{7} (2, -1, 4) = \frac{1}{7} (17, 30, -1)$. We can verify $\Pi_P(\vec{v}) + \Pi_{P^\perp}(\vec{v}) = \frac{1}{7} (17, 30, -1) + \frac{2}{7} (2, -1, 4) = (3, 4, 1) = \vec{v}$. This shows that the projection we found is correct.

P4.12 First we find the normal $\vec{n}$ of the plane P using the cross-product trick, $\vec{n} = (s - m) \times (r - m) = (4, 0, -3)$ and $r - m = (4, 1, -3)$, we find $\vec{n} = (4, 0, -3) \times (4, 1, -3) = (3, 0, 4)$. Now we want to find the projection of $\vec{u}$ onto the space perpendicular to P , which is given by the formula $\Pi_{P^\perp}(\vec{u}) = \frac{\vec{u} \cdot \vec{n}}{\|\vec{n}\|^2} \vec{n} = \frac{(-2, 1, 1) \cdot (3, 0, 4)}{3^2 + 4^2} (3, 0, 4) = \frac{-2}{25} (3, 0, 4) = (-\frac{6}{25}, 0, -\frac{8}{25})$.

P4.13 We'll compute the distance by first finding a vector $\vec{v}$ that connects an arbitrary point in the plane P with an arbitrary point on the line ℓ , and then computing the component of $\vec{v}$ that is perpendicular to the plane. The point that lies on the line is $p_\ell = (1, -3, 2)$, and the point in the plane is $p_P = (0, 1, 1)$. The vector between them is $\vec{v} = p_P - p_\ell = (1, -4, 1)$. To compute $d(\ell, P)$ we must find the length of the projection of $\vec{v}$ in the direction of $\vec{n}$: $d(\ell, P) = \left\| \frac{\vec{v} \cdot \vec{n}}{\|\vec{n}\|^2} \vec{n} \right\| = \frac{|\vec{v} \cdot \vec{n}|}{\|\vec{n}\|} = \frac{(1, -4, 1) \cdot (-1, 2, 2)}{\sqrt{(-1)^2 + 2^2 + 2^2}} = \frac{7}{3}$.

P4.15 First we find the change-of-basis transformations to the standard basis

$$B_s[\mathbb{1}]_U = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & -1 \end{bmatrix} \text{ and } B_s[\mathbb{1}]_V = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}. \text{ Next we compute } {}_V[\mathbb{1}]_{B_s} = ({}_B[\mathbb{1}]_V)^{-1} = \begin{bmatrix} -1 & 1 & 1 \\ 1 & 0 & -1 \\ 1 & -1 & 0 \end{bmatrix}. \text{ Finally, we compute } {}_V[\mathbb{1}]_U = {}_V[\mathbb{1}]_{B_s} {}_{B_s}[\mathbb{1}]_U.$$

P4.16 We can check that the sum of two upper triangular matrices results in an upper triangular matrix. Scaling an upper triangular matrix also preserves its nature. The zero matrix is upper triangular. Since all three subspace conditions are satisfied, the upper triangular matrices form a subspace.

P4.17 In each case, we must either recognize the set as being a subspace or explain which of the three subset conditions fails. **a)** W_1 corresponds to the plane $x + y = 0$ which is a subspace of $\mathbb{R}^3$. **b)** Consider the vectors $(0, 1, 0) \in W_2$ and $(0, 0, 1) \in W_2$. The sum $(0, 1, 0) + (0, 0, 1) = (0, 1, 1) \notin W_2$, so W_2 is not closed under addition. **c)** W_3 is equivalent to the line $\{(0, 0, 0) + t(1, 1, 1), \forall t \in \mathbb{R}\}$, which is a subspace of $\mathbb{R}^3$. **d)** Consider the vector $(1, 0, 0) \in W_4$. Multiplying this vector by -3 results in $(-3, 0, 0) \notin W_4$, so the set W_4 is not closed under scalar multiplication. **e)** The set W_5 does not contain the zero element $(0, 0, 0)$.

P4.19 Imagine performing the first step in the Gauss-Jordan elimination procedure. Subtracting $\frac{c}{a}$ times the first row from the second row results in $\begin{bmatrix} a & b \\ 0 & d - \frac{c}{a}b \end{bmatrix}$. If $d = \frac{bc}{a}$, the matrix will have rank one.

P4.21 Define $W = \text{span}(\vec{v}_1, \vec{v}_2 - \vec{v}_1, \vec{v}_3 - \vec{v}_2, \vec{v}_4 - \vec{v}_3)$. Every vector $\vec{w} \in W$ can be written in the following two equivalent ways:

$$\begin{aligned} \vec{w} &= \alpha_1 \vec{v}_1 + \alpha_2 (\vec{v}_2 - \vec{v}_1) + \alpha_3 (\vec{v}_3 - \vec{v}_2) + \alpha_4 (\vec{v}_4 - \vec{v}_3) \\ &= (\alpha_1 - \alpha_2) \vec{v}_1 + \alpha_2 \vec{v}_2 + \alpha_3 (\vec{v}_3 - \vec{v}_2) + \alpha_4 (\vec{v}_4 - \vec{v}_3). \end{aligned}$$

This implies $W = \text{span}(\vec{v}_1, \vec{v}_2, \vec{v}_3 - \vec{v}_2, \vec{v}_4 - \vec{v}_3)$. Using this approach we can show $W = \text{span}(\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4 - \vec{v}_3)$, and then $W = \text{span}(\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4) = V$.

P4.22 Recall that the dimension of a vector space is equal to the number of vectors in a basis for the vector space. For a set of vectors to be a basis, the vectors must span the space *and* be linearly independent. Since the problem does not specify whether the set of vectors are linearly independent, we cannot conclude that $\dim V = m$.

P4.23 Define $W = \text{span}(\vec{w}_1, \vec{w}_2)$, which means any vector $\vec{w} \in W$ can be written in the form $\vec{w} = \alpha_1 \vec{w}_1 + \alpha_2 \vec{w}_2$. To show W is a subspace, we must show it is closed under addition and scalar multiplication, and that it contains the zero element. Consider two arbitrary vectors chosen from W , $\vec{w}_a = \alpha_1 \vec{w}_1 + \alpha_2 \vec{w}_2$ and $\vec{w}_b = \beta_1 \vec{w}_1 + \beta_2 \vec{w}_2$. The sum of these two vectors is $\vec{w}_a + \vec{w}_b = (\alpha_1 + \beta_1) \vec{w}_1 + (\alpha_2 + \beta_2) \vec{w}_2$, which is also a vector in W . Similarly, scaling $\vec{w}_a$ by a constant produces a vector in W , and choosing the coefficients $\alpha_1 = \alpha_2 = 0$ gives the zero vector. Therefore $W \subseteq V$.

P4.24 We are told the set $S = \{\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4\}$ is linearly independent, which implies $\alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \cdots + \alpha_n \vec{v}_n = \vec{0}$ has only trivial solution $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) = (0, 0, 0, 0)$. Define the set $S' = \{\vec{v}_1, \vec{v}_2 - \vec{v}_1, \vec{v}_3 - \vec{v}_2, \vec{v}_4 - \vec{v}_3\}$. To find whether S' is a linearly independent set, consider the following equation:

$$\beta_1 \vec{v}_1 + \beta_2 (\vec{v}_2 - \vec{v}_1) + \beta_3 (\vec{v}_3 - \vec{v}_2) + \beta_4 (\vec{v}_4 - \vec{v}_3) = 0.$$

Rearranging the terms, we obtain the equivalent expression

$$\underbrace{(\beta_1 - \beta_2)}_{\alpha_1} \vec{v}_1 + \underbrace{(\beta_2 - \beta_3)}_{\alpha_2} \vec{v}_2 + \underbrace{(\beta_3 - \beta_4)}_{\alpha_3} \vec{v}_3 + \underbrace{\beta_4}_{\alpha_4} \vec{v}_4 = 0.$$

We recognize this form of equation from the definition of linear independence. Since we're told $\{\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4\}$ is a linearly independent set, we know the only solution to the above equation is $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) = (0, 0, 0, 0)$, or

$$0 = \beta_1 - \beta_2, \quad 0 = \beta_2 - \beta_3, \quad 0 = \beta_3 - \beta_4, \quad 0 = \beta_4.$$

Solving for the unknowns in the order $\beta_4, \beta_3, \beta_2, \beta_1$, we find the only solution is $(\beta_1, \beta_2, \beta_3, \beta_4) = (0, 0, 0, 0)$; therefore, $\{\vec{v}_1, \vec{v}_2 - \vec{v}_1, \vec{v}_3 - \vec{v}_2, \vec{v}_4 - \vec{v}_3\}$ is a linearly independent set.

P4.25 We'll use a proof by contradiction. If $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ is not a linearly independent set then there must be some $(\beta_1, \beta_2, \beta_3) \neq (0, 0, 0)$ such that $\beta_1 \vec{v}_1 + \beta_2 \vec{v}_2 + \beta_3 \vec{v}_3 = 0$. If this is true, then the equation $\alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \alpha_3 \vec{v}_3 + \alpha_4 \vec{v}_4 = 0$ will have a non-trivial solution $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) \neq (0, 0, 0, 0)$: $\alpha_1 = \beta_1, \alpha_2 = \beta_2, \alpha_3 = \beta_3$, and $\alpha_4 = 0$. However, this contradicts the fact that $\{\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4\}$ is a linearly independent set. Since we've arrived at a contradiction, it must be true that $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ is a linearly independent set.

P4.26 We know $\vec{v}_1, \vec{v}_2, \vec{v}_3$ are linearly independent, and $\vec{v}_4 \notin \text{span}(\vec{v}_1, \vec{v}_2, \vec{v}_3)$. To set up the proof by contradiction, assume $\{\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4\}$ is a linearly dependent set. This implies the equation $\alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \alpha_3 \vec{v}_3 + \alpha_4 \vec{v}_4 = 0$ has a nontrivial solution: $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) \neq (0, 0, 0, 0)$. We'll analyze the case $\alpha_4 = 0$ separately from the case $\alpha_4 \neq 0$. In the case $\alpha_4 = 0$, at least one of the coefficients $\alpha_1, \alpha_2, \alpha_3$ must be nonzero, which means $\alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \alpha_3 \vec{v}_3 = 0$ has a non-trivial solution; but this is a contradiction since $\vec{v}_1, \vec{v}_2, \vec{v}_3$ are linearly independent. In the case of $\alpha_4 \neq 0$, we can rewrite the equation as follows:

$$\alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \alpha_3 \vec{v}_3 + \alpha_4 \vec{v}_4 = 0 \quad \Rightarrow \quad \vec{v}_4 = \frac{-\alpha_1}{\alpha_4} \vec{v}_1 + \frac{-\alpha_2}{\alpha_4} \vec{v}_2 + \frac{-\alpha_3}{\alpha_4} \vec{v}_3,$$

which shows that $\vec{v}_4 \in \text{span}(\vec{v}_1, \vec{v}_2, \vec{v}_3)$ and contradicts $\vec{v}_4 \notin \text{span}(\vec{v}_1, \vec{v}_2, \vec{v}_3)$. Since we've arrived at a contradiction in both cases, we conclude that $\{\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4\}$ must be a linearly independent set.

P4.27 To check whether $\{\vec{v}_3, \vec{v}_2 + \vec{v}_3, \vec{v}_1 + \vec{v}_2 + \vec{v}_3\}$ is a linearly independent set, we must consider the solutions to the equation $c_1 \vec{v}_3 + c_2 (\vec{v}_2 + \vec{v}_3) + c_3 (\vec{v}_1 + \vec{v}_2 + \vec{v}_3) = 0$. We can rewrite this equation as follows:

$$\begin{aligned} 0 &= c_1 \vec{v}_3 + c_2 (\vec{v}_2 + \vec{v}_3) + c_3 (\vec{v}_1 + \vec{v}_2 + \vec{v}_3) \\ &= \underbrace{c_3}_{\alpha_1} \vec{v}_1 + \underbrace{(c_2 + c_3)}_{\alpha_2} \vec{v}_2 + \underbrace{(c_1 + c_2 + c_3)}_{\alpha_3} \vec{v}_3. \end{aligned}$$

This latter equation has the same form as the definition of linear independence; and since $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ is a linearly dependent set, we know the only solution to this equation is $(\alpha_1, \alpha_2, \alpha_3) = (0, 0, 0)$, which implies

$$c_3 = 0, \quad c_2 + c_3 = 0, \quad \text{and} \quad c_1 + c_2 + c_3 = 0.$$

Solving these equations, we obtain the trivial solution $(c_1, c_2, c_3) = (0, 0, 0)$ and conclude that the set $\{\vec{v}_3, \vec{v}_2 + \vec{v}_3, \vec{v}_1 + \vec{v}_2 + \vec{v}_3\}$ is linearly independent.

P4.28 Since the set $\{\vec{u}, \vec{v}\}$ is a basis for V , the dimension of V is two. The set $\{\vec{u} + \vec{v}, a\vec{u}\}$ contains two vectors so it is sufficient to check whether the vectors in the set are linearly independent. Consider the following equation:

$$\beta_1(\vec{u} + \vec{v}) + \beta_2 a\vec{u} = 0 = \underbrace{(\beta_1 + \beta_2 a)}_{\alpha_1} \vec{u} + \underbrace{\beta_1}_{\alpha_2} \vec{v}.$$

Given $\{\vec{u}, \vec{v}\}$ is a linearly independent set, the only solution to $\alpha_1 \vec{u} + \alpha_2 \vec{v} = 0$ is $(\alpha_1, \alpha_2) = (0, 0)$; therefore $\beta_1 + \beta_2 a = \beta_1 = 0$ and $\beta_1 = \beta_2 = 0$. This implies $\{\vec{u} + \vec{v}, a\vec{u}\}$ is a basis for V . The proof for $\{a\vec{u}, b\vec{v}\}$ is analogous.

P4.29 To extend the set $\{\vec{v}_1, \vec{v}_2\}$ to a basis for $\mathbb{R}^3$ we must pick a vector that is linearly independent from $\vec{v}_1$ and $\vec{v}_2$. We can see the structure of $\text{span}(\vec{v}_1, \vec{v}_2)$ more clearly if we construct the matrix $V = \begin{bmatrix} -\vec{v}_1 \\ -\vec{v}_2 \end{bmatrix}$ and transform it to its reduced row echelon form $\sim \begin{bmatrix} 1 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. We can add the vector $\vec{v}_3 = (0, 1, 0)$ to make this matrix full rank. On the other hand, if we want to add a third vector $\vec{v}_4$ that is not linearly independent, we can pick any linear combination of $\vec{v}_1$ and $\vec{v}_2$, like $\vec{v}_4 = \vec{v}_1 + \vec{v}_2$, for example.

P4.30 Construct the matrix $V = \begin{bmatrix} -\vec{v}_1 \\ -\vec{v}_2 \\ -\vec{v}_3 \end{bmatrix}$ and find a vector in its null space.

Chapter 5 solutions

Answers to exercises

E5.1 **a)** T_1 is linear, $M_{T_1} = \begin{bmatrix} 1 & 2 & 3 \end{bmatrix}$; **b)** T_2 is not linear, $T_2(2, 0) \neq 2T_2(1, 0)$; **c)** T_3 is not linear, $T_3(2, 0) \neq 2T_3(1, 0)$; **d)** T_4 is linear, $M_{T_4} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$. **E5.2** $T : \mathbb{R}^3 \rightarrow \mathbb{R}^3$,

$\text{Ker}(T) = \{\vec{0}\}$, $\text{Im}(T) = \mathbb{R}^3$. T is injective and surjective; therefore it is bijective.

E5.3 $T_M : \mathbb{R}^3 \rightarrow \mathbb{R}^2$, $\text{Ker}(T_M) = \text{span}((0, 1, 0))$, $\text{Im}(T_M) = \mathbb{R}^2$. $T_{M^T} : \mathbb{R}^2 \rightarrow \mathbb{R}^3$, $\text{Ker}(T_{M^T}) = \{\vec{0}\}$, $\text{Im}(T_{M^T}) = \text{span}((1, 0, 0), (0, 0, 1))$. **E5.4** $T(x, y) = (x + 4y, 2x +$

$5y, 3x + 6y)$ and $M_T = \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix}$. **E5.5** $M_{R_{\frac{\pi}{4}}} = \begin{bmatrix} \frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} \\ \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \end{bmatrix}$, $M_{R_{\frac{\pi}{2}}} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$. **E5.6** $M_{\Pi_d} =$

$\hat{d}\hat{d}^T = \begin{bmatrix} \frac{1}{10} & \frac{3}{10} \\ \frac{3}{10} & \frac{9}{10} \end{bmatrix}$. **E5.7** ${}_B[M_T]_B = {}_B[\mathbb{1}]_{B'} {}_{B'}[M_T]_{B'} {}_{B'}[\mathbb{1}]_B$. **E5.8** **a)** ${}_{B_s}[M_T]_{B_s} = \begin{bmatrix} 3 & 0 \\ 0 & 1 \end{bmatrix}$;

b) ${}_{B_f}[M_T]_{B_f} = \begin{bmatrix} 1 & 0 \\ 0 & 3 \end{bmatrix}$; **c)** ${}_{B_d}[M_T]_{B_d} = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$. **E5.9** $[\vec{x}_i]_B = {}_{B'}[M_T^{-1}]_{B'} [\vec{y}_i^w]_{*'} \quad [\vec{y}_i]_B = {}_B[M_T]_B {}_{B'}[M_T^{-1}]_{B'} [\vec{y}_i^w]_{*}$.

Solutions to selected exercises

E5.5 The general formula for the rotation matrix is $R_\theta = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$. Substitute the values $\theta = \frac{\pi}{4}$ and $\theta = \frac{\pi}{2}$ in the general formula to obtain the matrix representations of $R_{\frac{\pi}{4}}$ and $R_{\frac{\pi}{2}}$.

E5.6 A unit vector in the direction $\vec{d}$ is given by $\hat{d} = \frac{\vec{d}}{\|\vec{d}\|} = (\frac{\sqrt{10}}{10}, \frac{3\sqrt{10}}{10})^T$. To find the projection matrix, compute the outer product $\hat{d}\hat{d}^T$.

E5.7 Start with the formula for ${}_{B'}[M_T]_{B'}$ and multiply it by ${}_B[\mathbb{1}]_{B'}$ from the left and by ${}_{B'}[\mathbb{1}]_B$ from the right.

E5.8 Find the answers to parts **a)** and **b)** by inspection. For part **c)**, obtain ${}_{B_s}[\mathbb{1}]_{B_d} = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}$ and find its inverse: ${}_{B_d}[\mathbb{1}]_{B_s} = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}$. Next, calculate ${}_{B_d}[M_T]_{B_d}$ by computing ${}_{B_d}[\mathbb{1}]_{B_s} {}_{B_s}[M_T]_{B_s} {}_{B_s}[\mathbb{1}]_{B_d}$.

E5.9 Since $[\vec{y}_i^w]_* = {}_{B'}[M_T]_{B'}[\vec{x}_i]_{B'}$, you can recover $[\vec{x}_i]_B$ by hitting $[\vec{y}_i^w]_*$ with the inverse of M_T in the basis B' : $[\vec{x}_i]_B = {}_{B'}[M_T^{-1}]_{B'}[\vec{y}_i^w]_*$. To find the correct output, compute $[\vec{y}_i]_B = {}_B[M_T]_B {}_{B'}[M_T^{-1}]_{B'}[\vec{y}_i^w]_*$.

E5.10 Let $\vec{c}_1, \vec{c}_2, \dots, \vec{c}_n$ be the n columns of A . Since the columns of A form a basis for $\mathbb{R}^n$, any vector $\vec{b} \in \mathbb{R}^n$ can be written as a unique linear combination of the columns of A : $\vec{b} = x_1\vec{c}_1 + x_2\vec{c}_2 + \dots + x_n\vec{c}_n$ for some coefficients $x_1, x_2, \dots, x_n$. Reinterpreting the last equation as a matrix product in the column picture, we conclude $A\vec{x} = \vec{b}$ has a unique solution $\vec{x}$.

E5.11 The rank-nullity theorem tells us $\text{rank}(A) + \text{nullity}(A) = n$ for any $n \times n$ matrix A . Given $\mathcal{N}(A) = \{\vec{0}\}$, this means $\text{nullity}(A) = 0$ and therefore $\text{rank}(A) = n$.

Answers to problems

P5.1 **a)** T_1 is linear, $M_{T_1} = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}$; **b)** T_2 is not linear, $T_2(2, 2) \neq 2T_2(1, 1)$; **c)** T_3 is not linear, $T_3(-1, -1) + T_3(1, 1) \neq T_2(0, 0)$; **d)** T_4 is linear, $M_{T_4} = \begin{bmatrix} 3 & -2 & 1 \\ 2 & 1 & -4 \end{bmatrix}$; **e)** T_5 is linear, $M_{T_5} = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$; **f)** T_6 is not linear, $T_6(2, 0, 0, 0) \neq 2T_6(1, 0, 0, 0)$. **P5.2** $\text{Im}(T) = \text{span}((1, 1, 0), (0, -1, 2))$. **P5.3** $M_T = \begin{bmatrix} 0 & -a_z & a_y \\ a_z & 0 & -a_x \\ -a_y & a_x & 0 \end{bmatrix}$; $\text{Ker}(T) = \text{span}((a_x, a_y, a_z)^T)$. **P5.4** $M_T = \begin{bmatrix} -1 & 3 \\ -7 & 4 \end{bmatrix}$, with respect to the standard basis. **P5.5** **a)** ${}_B[M_T]_B = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix}$; **b)** ${}_{B'}[M_T]_{B'} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & \frac{3}{2} & \frac{1}{2} \\ 0 & -\frac{1}{2} & \frac{1}{2} \end{bmatrix}$; **c)** ${}_B[M_T]_{B'} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 2 & 0 \end{bmatrix}$. **P5.6** ${}_{B'}[M_T]_{B'} = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & -3 \end{bmatrix}$. **P5.7** **a)** $\mathbb{1}$, which is the identity transformation; **b)** $M_1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$; **c)** $M_2 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$; **d)** $M_3 = \begin{bmatrix} \cos(\frac{\pi}{3}) & -\sin(\frac{\pi}{3}) \\ \sin(\frac{\pi}{3}) & \cos(\frac{\pi}{3}) \end{bmatrix}$; **e)** $M_4 = \begin{bmatrix} \cos(-\frac{\pi}{6}) & -\sin(-\frac{\pi}{6}) \\ \sin(-\frac{\pi}{6}) & \cos(-\frac{\pi}{6}) \end{bmatrix}$; **f)** $M_5 = \begin{bmatrix} 1 & \frac{1}{2} \\ 0 & 1 \end{bmatrix}$.

Solutions to selected problems

P5.2 Applying T to the input vector $(1, 0)$ produces $(1, 1 - 0 \cdot 2 \cdot 0) = (1, 1, 0)$, and the input vector $(0, 1)$ produces the output $(0, 0 - 1 \cdot 2 \cdot 1) = (0, -1, 2)$. Thus, $\text{Im}(T) = \text{span}((1, 1, 0), (0, -1, 2)) \subseteq \mathbb{R}^3$.

P5.3 Use the standard probing-with-the-standard-basis approach to obtain the matrix representation. See youtu.be/BaM70CEm3G0 for an interesting discussion about the cross product viewed as a linear transformation.

P5.5 The change-of-basis transformation from B' to the standard basis is ${}_B[\mathbb{1}]_{B'} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & -1 \end{bmatrix}$. We find ${}_B[M_T]_B = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix}$ directly by observing the formula for T . Next we compute ${}_{B'}[M_T]_{B'}$ using the “sandwich” formula ${}_{B'}[\mathbb{1}]_B {}_B[M_T]_B {}_B[\mathbb{1}]_{B'}$. Finally, we compute ${}_B[M_T]_{B'}$ by changing only the right basis of the transformation: ${}_B[M_T]_{B'} = {}_B[M_T]_B {}_B[\mathbb{1}]_{B'}$.

P5.6 The change-of-basis matrix from the basis B' to the standard basis B is ${}_B[\mathbb{1}]_{B'} = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. We then compute ${}_{B'}[M_T]_{B'} = {}_{B'}[\mathbb{1}]_B {}_B[M_T]_B {}_B[\mathbb{1}]_{B'}$.

P5.7 First, you must visually recognize the type of transformation that is acting in each case; then, use the appropriate matrix formula. Transformations 1 and 2 are reflections. Transformations 3 and 4 are rotations. Transformation 5 is a *shear* that maps $T_5(\hat{i}) = \hat{i}$ and $T_5(\hat{j}) = (\frac{1}{2}, 1)^T$. You can obtain the matrix representation from these two observations.

P5.8 Let $B_s = \{\hat{e}_1, \hat{e}_2, \dots, \hat{e}_n\}$ be the standard basis for $\mathbb{R}^n$. Since $A\vec{x} = \vec{b}$ has a solution $\vec{x}$ for every possible $\vec{b}$, it is possible to find the solutions $\vec{x}_i$ in n equations of the form $A\vec{x}_i = \hat{e}_i$, for $i \in \{1, 2, \dots, n\}$. Now construct the matrix B that contains the solutions $\vec{x}_i$ as columns: $B = [\vec{x}_1, \dots, \vec{x}_n]$. Observe that $AB = A[\vec{x}_1, \dots, \vec{x}_n] = [A\vec{x}_1, \dots, A\vec{x}_n] = [\hat{e}_1, \dots, \hat{e}_n] = \mathbb{1}_n$. The equation $BA = \mathbb{1}_n$ implies A is invertible.

P5.9 We want to show there is a unique solution to $A\vec{x} = \vec{b}$ if and only if $A\vec{y} = \vec{0}$ has only the trivial solution $\vec{y} = \vec{0}$. To show **(2)** $\Rightarrow$ **(4)**, let's assume the opposite of **(4)** is true: that $A\vec{x} = \vec{0}$ has a nontrivial solution $\vec{y} \neq \vec{0}$. If this is true, then $\vec{x}' = \vec{x} + \vec{y}$ is a solution to $A\vec{x} = \vec{b}$ since $A\vec{x}' = A(\vec{x} + \vec{y}) = A\vec{x} + A\vec{y} = A\vec{x} = \vec{b}$. Since $\vec{y} \neq \vec{0}$, $\vec{x}' \neq \vec{x}$, and therefore $A\vec{x} = \vec{b}$ does not have a unique solution. If we want $A\vec{x} = \vec{b}$ to have a unique solution, $A\vec{y} = \vec{0}$ must have only the trivial solution $\vec{y} = \vec{0}$.

We can use similar reasoning to show **(4)** $\Rightarrow$ **(2)**. Assume $A\vec{x}_1 = \vec{b}$ and $A\vec{x}_2 = \vec{b}$ are two solutions, and assume $\vec{x}_1 \neq \vec{x}_2$. Then $A(\vec{x}_1 - \vec{x}_2) = \vec{0}$; and, since $\vec{x}_1 - \vec{x}_2 \neq \vec{0}$, we've shown that $A\vec{x} = \vec{0}$ has a nontrivial solution. If we want $A\vec{y} = \vec{0}$ to have only the trivial solution $\vec{y} = \vec{0}$, then $A\vec{x} = \vec{b}$ must have a unique solution.

P5.10 If $A\vec{x} = \vec{0}$ has only the trivial solution $\vec{x} = \vec{0}$, then the reduced row echelon form of A contains no free parameters and A has full rank.

P5.11 To solve the system of equations $A\vec{x} = \vec{0}$, we build the augmented matrix $[A \mid \vec{0}]$ and compute its RREF. Since $A\vec{x} = \vec{0}$ has a unique solution, the RREF of $[A \mid \vec{0}]$ doesn't contain any free variables, so the reduced echelon form of A must be $\mathbb{1}_n$.

P5.12 Assume the matrix A is invertible. Consider the procedure for obtaining the matrix inverse that starts from the extended array $[A \mid \mathbb{1}]$. The sequence of row operations $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k$ is used to obtain the RREF of the extended array:

$$[A \mid \mathbb{1}] \xrightarrow{\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k} [\mathbb{1} \mid A^{-1}].$$

Thus, $\text{rref}(A) = \mathbb{1}$. Note the determinant of the final step is nonzero; $|\mathbb{1}| \neq 0$. Since the row operations $\mathcal{R}_i$ only act to change the sign or the size of the determinant, the overall sequence of row operations doesn't change the zero vs. nonzero nature of the determinant. This implies the determinant of the original matrix A is also nonzero; $|A| \neq 0$.

P5.13 If T is injective, then its kernel contains only the zero vector $\text{Ker}(T) = \{\vec{0}\}$. Since we're interested in determining the linear independence of the set $\{T(\vec{v}_1), \dots, T(\vec{v}_n)\}$, we must consider the solutions to the equation

$$\alpha_1 T(\vec{v}_1) + \dots + \alpha_n T(\vec{v}_n) = \vec{0}.$$

Knowing T is linear, we can rewrite this as $T(\alpha_1 \vec{v}_1 + \dots + \alpha_n \vec{v}_n) = \vec{0}$. Since $\text{Ker}(T) = \{\vec{0}\}$, the equation we started with is equivalent to the equation $\alpha_1 \vec{v}_1 + \dots + \alpha_n \vec{v}_n = \vec{0}$. But we're told $\{\vec{v}_1, \dots, \vec{v}_n\}$ is a linearly independent set, so the only solution to this equation is the trivial solution $\alpha_i = 0$. Therefore $\{T(\vec{v}_1), \dots, T(\vec{v}_n)\}$ is a linearly independent set.

P5.14 Consider an arbitrary vector $\vec{w} \in W$. Since T is surjective, there must exist a vector $\vec{v} \in V$ such that $T(\vec{v}) = \vec{w}$. We know $\{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n\}$ is a spanning set for V , so it must be that $\vec{v}$ can be written as a linear combination of this set of vectors:

$$\vec{v} = \alpha_1 \vec{v}_1 + \alpha_2 \vec{v}_2 + \dots + \alpha_n \vec{v}_n.$$

Now apply the linear transformation T to both sides of this equation:

$$\vec{w} = T(\vec{v}) = \alpha_1 T(\vec{v}_1) + \alpha_2 T(\vec{v}_2) + \dots + \alpha_n T(\vec{v}_n).$$

We've shown that an arbitrary vector $\vec{w} \in W$ can be written as a linear combination of the vectors $T(\vec{v}_1), T(\vec{v}_2), \dots, T(\vec{v}_n)$; therefore, the vector space W is spanned by $\{T(\vec{v}_1), T(\vec{v}_2), \dots, T(\vec{v}_n)\}$.

P5.15 Assume B doesn't have full rank. Then there must exist a nonzero vector $\vec{y}$ such that $B\vec{y} = \vec{0}$. Now multiply both sides of the equation $AB = \mathbb{1}$ by the vector $\vec{y}$ to obtain $AB\vec{y} = \mathbb{1}\vec{y} = \vec{y}$. The left side of this equation is $AB\vec{y} = A\vec{0} = \vec{0}$, but the right side is $\mathbb{1}\vec{y} = \vec{y} \neq \vec{0}$, so we've arrived at a contradiction and B must have full rank.

For the second part of the question, we start from the equation $AB = \mathbb{1}$ and multiply both sides by B to obtain $BAB = B$, which simplifies to $(BA - \mathbb{1})B = \vec{0}$. We can multiply this equation by any nonzero vector $\vec{x} \in \mathbb{R}^n$ and obtain $(BA - \mathbb{1})B\vec{x} = \vec{0}$. Since we know B has full rank $B\vec{x} \neq \vec{0}$, it must be that $(BA - \mathbb{1})$ is the zero matrix. Therefore $BA = \mathbb{1}$.

Chapter 6 solutions

Answers to exercises

E6.1 **a)** $p(\lambda) = (\lambda - 4)^2$; $\lambda_1 = 4$ (repeated). **b)** $p(\lambda) = \lambda^2 - 4\lambda - 5$; $\lambda_1 = 5$ and $\lambda_2 = -1$. **c)** $p(\lambda) = \lambda^2 - 21$; $\lambda_1 = \sqrt{21}$, $\lambda_2 = -\sqrt{21}$. **d)** $p(\lambda) = \lambda^2$; $\lambda_1 = 0$ (repeated). **e)** $p(\lambda) = \lambda^2 - 2\lambda + 1$; $\lambda_1 = 1$ (repeated). **E6.2** $\lambda_1 = 1$ with eigenvector $\vec{e}_1 = (1, 1, 0)^T$ and $\lambda_2 = 5$ with eigenvectors $\vec{e}_2 = (-1, 1, 0)^T$ and $\vec{e}_3 = (0, 0, 1)^T$. **E6.6** $\alpha = -3$ and $\beta = 4$. **E6.7** $(1, 1)^T$ and $(1, -2)^T$ are eigenvectors of L . **E6.8** $\det(A) = 15$, $A^{-1} =$

$$\begin{bmatrix} 1 & -\frac{4}{5} & -\frac{11}{3} \\ 0 & \frac{1}{5} & -\frac{1}{3} \\ 0 & 0 & \frac{1}{3} \end{bmatrix}; \det(B) = xz, B^{-1} = \begin{bmatrix} \frac{1}{x} & 0 \\ -\frac{y}{xz} & \frac{1}{z} \end{bmatrix}; \det(C) = 1, C^{-1} = \begin{bmatrix} 5 & 0 \\ 0 & \frac{1}{5} \end{bmatrix}. \quad \text{E6.9 No.}$$

E6.10 **a)** A is upper triangular and positive definite. **b)** B is orthogonal and normal. Specifically, B is a rotation. **c)** C is symmetric, orthogonal, and normal. Also, C is a reflection and a permutation. **E6.11** $\{\mathbf{e}_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \mathbf{e}_2 = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \mathbf{e}_3 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}\}$. **E6.12** Yes.

E6.14 $\|P_0\| = \sqrt{2}$, $\|P_1\| = \frac{\sqrt{6}}{3}$, and $d(P_0, P_1) = \frac{2\sqrt{6}}{3}$. **E6.15** $\vec{e}_1 = (4, 2)$, $\vec{e}_2 = (-1, 2)$.

E6.16 $\hat{e}_1 = (\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}, 0)$, $\hat{e}_2 = (\frac{1}{\sqrt{6}}, -\frac{1}{\sqrt{6}}, \frac{2}{\sqrt{6}})$, and $\hat{e}_3 = (-\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}})$. **E6.17** $\hat{e}_1 = 1$,

$$\hat{e}_2 = \sqrt{12}(x - \frac{1}{2}), \text{ and } \hat{e}_3 = 6\sqrt{5}(x^2 - x + \frac{1}{6}). \quad \text{E6.18 } Q = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} & -\frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{6}} & \frac{1}{\sqrt{3}} \\ 0 & \frac{2}{\sqrt{6}} & \frac{1}{\sqrt{3}} \end{bmatrix},$$

$$\text{and } R = \begin{bmatrix} \sqrt{2} & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & \frac{3}{\sqrt{6}} & \frac{1}{\sqrt{6}} \\ 0 & 0 & \frac{2}{\sqrt{3}} \end{bmatrix}. \quad \text{E6.19 a) } -1 + i; \text{ b) } -14 + 23i; \text{ c) } \frac{1}{25}(26 + 7i). \quad \text{E6.20}$$

$\lambda_1 = i$ with $\vec{e}_1 = (-\frac{1}{2+i}, 1)^T$ and $\lambda_2 = -i$ with $\vec{e}_2 = (-\frac{1}{2-i}, 1)^T$. **E6.21** $F[\mathbb{1}]_B =$

$$\frac{1}{4} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix}. \quad \text{E6.22 No.}$$

Solutions to selected exercises

E6.2 The characteristic polynomial is $p_A(\lambda) = \lambda^3 - 11\lambda^2 + 35\lambda - 25 = (\lambda - 1)(\lambda - 5)^2$. The eigenvalues of A are $\lambda_1 = 1$ and the repeated eigenvalue $\lambda_2 = 5$. To find eigenvectors that correspond to λ_1 , we must find the null space of the matrix $(A - \mathbb{1})$. We find the eigenvector $\vec{e}_1 = (1, 1, 0)^T$. For $\lambda_2 = 5$ we find $\vec{e}_2 = (-1, 1, 0)^T$ and $\vec{e}_3 = (0, 0, 1)^T$.

E6.3 Let $\vec{e}_1$ be the eigenvector associated with λ_1 and $\vec{e}_2$ be the λ_2 eigenvector. We want to show that the linear independence equation $\alpha_1\vec{e}_1 + \alpha_2\vec{e}_2 = \vec{0}$ has only the trivial solution; $(\alpha_1, \alpha_2) = (0, 0)$. Multiplying the linear independence equation by A gives us

$$A(\alpha_1\vec{e}_1 + \alpha_2\vec{e}_2) = \alpha_1\lambda_1\vec{e}_1 + \alpha_2\lambda_2\vec{e}_2 = \vec{0}.$$

Separately, multiply the linear independence equation by λ_1 to obtain

$$\lambda_1(\alpha_1\vec{e}_1 + \alpha_2\vec{e}_2) = \lambda_1\alpha_1\vec{e}_1 + \lambda_1\alpha_2\vec{e}_2 = \vec{0}.$$

Subtracting the latter two equations, we find

$$0\vec{e}_1 + \alpha_2(\lambda_2 - \lambda_1)\vec{e}_2 = \alpha_2(\lambda_2 - \lambda_1)\vec{e}_2 = \vec{0}.$$

Since $\lambda_2 - \lambda_1 \neq 0$ and $\vec{e}_2 \neq \vec{0}$, it must be that $\alpha_2 = 0$. But if $\alpha_2 = 0$, then the equation $\alpha_1\vec{e}_1 + \alpha_2\vec{e}_2 = \vec{0}$ and $\vec{e}_1 \neq \vec{0}$ implies $\alpha_1 = 0$; thus we've shown $(\alpha_1, \alpha_2) = (0, 0)$ is the only solution. Therefore, $\{\vec{e}_1, \vec{e}_2\}$ is a linearly independent set.

E6.4 Multiply both sides of the eigenvalue equation $A\vec{e}_\lambda = \lambda\vec{e}_\lambda$ by A to obtain $AA\vec{e}_\lambda = \lambda A\vec{e}_\lambda = \lambda^2\vec{e}_\lambda$. Thus λ^2 is an eigenvalue of A^2 .

E6.5 Multiply both sides of the eigenvalue equation $A\vec{e}_\lambda = \lambda\vec{e}_\lambda$ by A^{-1} to obtain $A^{-1}A\vec{e}_\lambda = \lambda A^{-1}\vec{e}_\lambda$. After A^{-1} cancels with A , we're left with the equation $\vec{e}_\lambda = \lambda A^{-1}\vec{e}_\lambda$. Dividing both sides of the equation by λ we obtain $\frac{1}{\lambda}\vec{e}_\lambda = A^{-1}\vec{e}_\lambda$, which shows that λ^{-1} is an eigenvalue of A^{-1} .

E6.6 The characteristic polynomial of A is $p_A(\lambda) = \lambda^2 - \beta\lambda - \alpha$. If we want the eigenvalues of A to be 1 and 3, we must choose α and β so that $p_A(\lambda) = (\lambda - 1)(\lambda - 3)$. Expanding the factored expression, we find $(\lambda - 1)(\lambda - 3) = \lambda^2 - 4\lambda + 3$, so $\alpha = -3$ and $\beta = 4$.

E6.7 First we compute $L(1, 1)^T = (5, 5)^T = 5(1, 1)^T$ so $(1, 1)^T$ is an eigenvector, with eigenvalue $\lambda = 5$. Next we compute $L(1, -1)^T = (1, 3)^T \neq \alpha(1, -1)^T$ so $(1, -1)^T$ is not an eigenvector. We can also compute $L(1, -2)^T = (-1, 2)^T = -1(1, -2)^T$, which implies $(1, -2)^T$ is an eigenvector with eigenvalue -1 . Since a 2×2 matrix can have at most two eigenvectors, we don't need to check $(2, -1)^T$ since we know it's not an eigenvector.

E6.9 A matrix A is symmetric if it satisfies $A^T = A$, which is not the case for the given matrix.

E6.11 First of all we must determine dimensionality of the vector space in question. The general vector space of 2×2 matrices has four dimensions, but an upper triangular matrix A satisfies $a_{ij} = 0$ for all $i > j$, which corresponds to the constraint $a_{21} = 0$. The space of 2×2 upper triangular matrices is a subspace of a four-dimensional vector space defined by one constraint; therefore it is three-dimensional. Any matrix $\begin{bmatrix} a & b \\ 0 & c \end{bmatrix}$ can be written as $a\mathbf{e}_1 + b\mathbf{e}_2 + c\mathbf{e}_3$.

E6.12 Consider an arbitrary second-degree polynomial $p(x) = a_0 + a_1x + a_2x^2$. We can rewrite it as

$$\begin{aligned} p(x) &= a_0 + a_1[(x-1) + 1] + a_2[(x-1) + 1]^2 \\ &= a_0 + a_1(x-1) + a_1 + a_2[(x-1)^2 + 2(x-1) + 1] \\ &= (a_0 + a_1 + a_2)1 + (a_1 + 2a_2)(x-1) + a_2(x-1)^2. \end{aligned}$$

Since we've expressed an arbitrary polynomial of degree two in the desired form, the answer is yes. In other words, $\{1, x-1, (x-1)^2\}$ is a basis for the vector space of polynomials of degree at most two.

E6.13 First we compute $\langle \mathbf{e}_1, \mathbf{e}_2 \rangle = 0$, $\langle \mathbf{e}_1, \mathbf{e}_3 \rangle = 0$, and $\langle \mathbf{e}_2, \mathbf{e}_3 \rangle = 0$ to show the vectors of the basis are orthogonal. Then compute $\langle \mathbf{e}_1, \mathbf{e}_1 \rangle = 1$, $\langle \mathbf{e}_2, \mathbf{e}_2 \rangle = 1$, and $\langle \mathbf{e}_3, \mathbf{e}_3 \rangle = 1$ to show the vectors have norm one. The set $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ is orthonormal.

E6.15 The exercise does not require us to normalize the vectors, so we can leave the first vector as-is: $\vec{v}_1 = \vec{e}_1 = (4, 2)$. Next, we calculate $\vec{e}_2$ using the formula $\vec{e}_2 = \vec{v}_2 - \Pi_{\vec{e}_1}(\vec{v}_2)$, which corresponds to removing the component of $\vec{v}_2$ that lies in the direction of $\vec{e}_1$. Using the projection formula we obtain $\Pi_{\vec{e}_1}(\vec{v}_2) = \frac{(4, 2) \cdot (1, 3)}{\|(4, 2)\|^2} (4, 2) = \frac{4+6}{16+4} (4, 2) = \frac{1}{2} (4, 2) = (2, 1)$. Thus $\vec{e}_2 = \vec{v}_2 - \Pi_{\vec{e}_1}(\vec{v}_2) = (1, 3) - (2, 1) = (-1, 2)$. Verify that $\vec{e}_1 \cdot \vec{e}_2 = 0$.

E6.17 We're given the functions $\mathbf{f}_1 = 1$, $\mathbf{f}_2 = x$, and $\mathbf{f}_3 = x^2$ and we want to transform them into an orthonormal set $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3\}$. Since $\|\mathbf{f}_1\| = 1$, we set $\hat{\mathbf{e}}_1 = \mathbf{f}_1$. Next we find $\mathbf{e}_2 = \mathbf{f}_2 - \langle \mathbf{f}_2, \hat{\mathbf{e}}_1 \rangle \hat{\mathbf{e}}_1 = x - \frac{1}{2}$. After normalization we obtain $\hat{\mathbf{e}}_2 = \sqrt{12}(x - \frac{1}{2})$. Continuing with the Gram-Schmidt procedure, we compute $\mathbf{e}_3 = \mathbf{f}_3 - \langle \mathbf{f}_3, \hat{\mathbf{e}}_1 \rangle \hat{\mathbf{e}}_1 - \langle \mathbf{f}_3, \hat{\mathbf{e}}_2 \rangle \hat{\mathbf{e}}_2 = x^2 - \frac{1}{3} - \frac{\sqrt{3}}{6}\sqrt{12}(x - \frac{1}{2})$ and find $\hat{\mathbf{e}}_3 = 6\sqrt{5}(x^2 - x + \frac{1}{6})$.

E6.20 The characteristic polynomial is $p_B(\lambda) = \begin{vmatrix} -2-\lambda & -1 \\ 5 & 2-\lambda \end{vmatrix} = \lambda^2 + 1$. The roots of the characteristic polynomial are $\lambda_1 = i$ and $\lambda_2 = -i$. To find the eigenvector associated with the eigenvalue $\lambda_1 = i$, we calculate $\mathcal{N}(B - i\mathbb{I})$. We obtain the eigenvector $\vec{e}_1 = (-\frac{1}{2+i}, 1)^T$. Similarly we calculate $\mathcal{N}(B + i\mathbb{I})$ and find $\vec{e}_2 = (-\frac{1}{2-i}, 1)^T$, which is the eigenvector associated with $\lambda_2 = -i$.

E6.21 First we find the inverse change of basis ${}_B[\mathbb{I}]_F$ with columns that are the vectors of the Fourier basis F . We then compute its inverse to find ${}_F[\mathbb{I}]_B = ({}_B[\mathbb{I}]_F)^{-1} = \frac{1}{4} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & -1 & -i \\ 1 & -1 & 1 & -1 \\ 1 & -i & -1 & i \end{bmatrix}$.

E6.22 A matrix A is Hermitian if it satisfies $A^\dagger = A$, which is not the case for the given matrix.

Answers to problems

P6.1 a) $\lambda_1 = 6$, $\lambda_2 = -1$; b) $\lambda_1 = -2$, $\lambda_2 = 2$, $\lambda_3 = 0$. **P6.2** a) $\lambda_1 = 5$, $\lambda_2 = 4$; b) $\lambda_1 = \frac{1}{2}(5 + \sqrt{5})$, $\lambda_2 = \frac{1}{2}(5 - \sqrt{5})$; c) $\lambda_1 = 3$, $\lambda_2 = 0$, $\lambda_3 = 0$; d) $\lambda_1 = -3$, $\lambda_2 = -1$, $\lambda_3 = 1$. **P6.3** $\lambda_1 = \varphi = \frac{1+\sqrt{5}}{2} = 1.6180339\dots$; $\lambda_2 = -\frac{1}{\varphi} = \frac{1-\sqrt{5}}{2} =$

$-0.6180339\dots$ **P6.4** $\lambda_1 = \varphi$ and $\lambda_2 = -\frac{1}{\varphi}$. **P6.5** $X = Q^{-1} = \begin{bmatrix} \frac{5+\sqrt{5}}{10} & \frac{\sqrt{5}}{5} \\ \frac{5-\sqrt{5}}{10} & -\frac{\sqrt{5}}{5} \end{bmatrix}$.

P6.6 a) $\lambda_1 = 1$, $\vec{e}_{\lambda_1} = (1, 1)^T$, $\lambda_2 = -1$, $\vec{e}_{\lambda_2} = (1, -1)^T$; b) $\lambda_1 = 3$, $\vec{e}_{\lambda_1} = (1, 3, 9)^T$, $\lambda_2 = 2$, $\vec{e}_{\lambda_2} = (1, 2, 4)^T$, $\lambda_3 = -1$, $\vec{e}_{\lambda_3} = (1, -1, 1)^T$. **P6.7** $A^{10} = \begin{bmatrix} 2 & 2 \\ 5 & -1 \end{bmatrix}^{10} =$

$\begin{bmatrix} 765854 & 282722 \\ 706805 & 341771 \end{bmatrix}$. **P6.8** $(x_\infty, y_\infty, z_\infty)^T = \frac{1}{6}(x_0 + 4y_0 + z_0, x_0 + 4y_0 + z_0, x_0 + 4y_0 + z_0)^T$. **P6.11** a) -240; b) 900; c) $\frac{-1}{30}$; d) $\frac{27}{2}$. **P6.13** a) Yes; b) No; c) Yes. **P6.16** Both are right. **P6.17** $(2O)^{-1} = \frac{1}{2}O^T$. **P6.18** Yes, $(P_2(t), \mathbb{R}, +, \cdot)$ is a vector space. **P6.19** No. **P6.20** No. **P6.21** a) No; b) Yes; c) Yes. **P6.22** $\vec{v} = (1, 0, 0, 1)$. **P6.23**

$V = [1, 0, 0, 1]$. **P6.25** The subset of $\mathbb{R}^2$ that consists of the x - and y -axes. **P6.26** The subset of the x -axis that corresponds to the integers $\mathbb{Z}$. **P6.27** $M_T = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. Eigenvalues: $\lambda = 1$ with multiplicity one, and $\lambda = 0$ with multiplicity two. **P6.28** The solution space is one-dimensional and spanned by the function e^{-t} . **P6.29** $\begin{bmatrix} 3 & 1 \\ -1 & 3 \end{bmatrix}$.

P6.30 $D = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$. **P6.31** $\{\mathbf{e}_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \mathbf{e}_2 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \mathbf{e}_3 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}\}$. **P6.32** $d =$

6. **P6.33** $d = 9$. **P6.34** $\lambda_1 = -2$; $\mathbf{e}_1 = \begin{bmatrix} -1 & 0 \\ 1 & 0 \end{bmatrix}$. $\lambda_2 = -1$; $\mathbf{e}_2 = \begin{bmatrix} -2 & 1 \\ 1 & 0 \end{bmatrix}$. $\lambda_3 = 1$ (repeated); $\mathbf{e}_3 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$ and $\mathbf{e}_4 = \begin{bmatrix} 2 & 3 \\ 1 & 0 \end{bmatrix}$. **P6.35** $\langle L_0(x), L_1(x) \rangle = 0$, $\langle L_0(x), L_2(x) \rangle = 0$, and $\langle L_1(x), L_2(x) \rangle = 0$. **P6.36** a) 9; b) 3; c) $\sqrt{10}$. **P6.39** $\hat{\mathbf{e}}_1 = (0, 1)$, $\hat{\mathbf{e}}_2 = (-1, 0)$. **P6.40** $\hat{\mathbf{e}}_1 = (\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$, $\hat{\mathbf{e}}_2 = (-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$. **P6.41** $\hat{\mathbf{e}}_1 = (\frac{3}{\sqrt{10}}, \frac{1}{\sqrt{10}})$, $\hat{\mathbf{e}}_2 = (\frac{-1}{\sqrt{10}}, \frac{3}{\sqrt{10}})$.

P6.42 $\hat{\mathbf{e}}_1 = \frac{1}{\sqrt{2}}$, $\hat{\mathbf{e}}_2 = \sqrt{\frac{3}{2}}x$, and $\hat{\mathbf{e}}_3 = \sqrt{\frac{5}{8}}(3x^2 - 1)$. **P6.43** $A = Q\Lambda Q^{-1} = \begin{bmatrix} 2 & 0 & -5 \\ 0 & 0 & -3 \end{bmatrix} =$

$\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} -3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix} \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & -1 \end{bmatrix}$. **P6.44** a) 5; b) $2 + 3i$; c) $2 + 5i$; d) $\sqrt{26}$. **P6.45** $A + B =$

$\begin{bmatrix} 4 & 3+8i \\ 8+3i & -5+3i \end{bmatrix}$; $CB = \begin{bmatrix} 3+8i & 2-i \\ 45+3i & -33+31i \end{bmatrix}$; $(2+i)B = \begin{bmatrix} 5 & 8-i \\ 9+7i & -15+5i \end{bmatrix}$. **P6.46** a) $\lambda_1 = 2 + i$

and $\lambda_2 = 2 - i$; b) $\lambda_1 = 3\sqrt{3}i$ and $\lambda_2 = 3\sqrt{3}i$; c) $\lambda_1 = 2 + 8i$ and $\lambda_2 = 2 - 8i$. **P6.47** $Q = \begin{bmatrix} -1+i & 1+i \\ 2 & 2 \end{bmatrix}$, $\Lambda = \begin{bmatrix} 0 & 0 \\ 0 & 2 \end{bmatrix}$. **P6.53** a) Yes; b) No; c) Yes; d) No; e) Yes; f) Yes.

Solutions to selected problems

P6.3 To find the eigenvalues of the matrix A we must find the roots of its characteristic polynomial:

$$p(\lambda) = \det(A - \lambda \mathbb{1}) = \det \left(\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} - \begin{bmatrix} \lambda & 0 \\ 0 & \lambda \end{bmatrix} \right) = \det \left(\begin{bmatrix} 1-\lambda & 1 \\ 1 & -\lambda \end{bmatrix} \right).$$

Using the determinant formula $\det \begin{pmatrix} a & b \\ c & d \end{pmatrix} = ad - bc$, we find the characteristic polynomial of A is $p(\lambda) = \lambda^2 - \lambda - 1$. The eigenvalues of A are the roots λ_1 and λ_2 of this equation, which we can find using the formula for solving quadratic equations we saw in Section 1.6 (see page 33).

P6.4 The vector $\vec{e}_1$ is an eigenvector of A because $A\vec{e}_1 = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ \frac{1}{\varphi} \end{bmatrix} = \begin{bmatrix} 1 + \frac{1}{\varphi} \\ 1 \end{bmatrix}$. Now observe the following interesting fact: $\frac{1}{\varphi}(1 + \frac{1}{\varphi}) = \frac{1}{\varphi} + \frac{1}{\varphi^2} = \frac{\varphi+1}{\varphi^2} = 1$. This means we can write $A\vec{e}_1 = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ \frac{1}{\varphi} \end{bmatrix} = \varphi \begin{bmatrix} 1 \\ \frac{1}{\varphi} \end{bmatrix}$, which shows that $\vec{e}_1$ is an eigenvector of A and it corresponds to eigenvalue $\lambda_1 = \varphi$. Similar reasoning shows $A\vec{e}_2 = -\frac{1}{\varphi}\vec{e}_2$ so $\vec{e}_2$ is an eigenvector of A that corresponds to eigenvalue $\lambda_2 = -\frac{1}{\varphi}$.

P6.5 The eigendecomposition of matrix A is $A = Q\Lambda Q^{-1}$. The unknown matrix X is the inverse matrix of the matrix $Q = \begin{bmatrix} 1 & 1 \\ \frac{1}{\varphi} & -\varphi \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ \frac{2}{1+\sqrt{5}} & -\frac{1+\sqrt{5}}{2} \end{bmatrix}$. To find Q^{-1} we can start from the array $[Q \mid \mathbb{1}]$ and perform row operations until we obtain $[\mathbb{1} \mid Q^{-1}]$.

P6.6 First we obtain the characteristic polynomial $p_A(\lambda) = \begin{vmatrix} -\lambda & 1 \\ 1 & -\lambda \end{vmatrix} = \lambda^2 - 1$. The eigenvalues of A are $\lambda_1 = 1$ and $\lambda_2 = -1$. To find the eigenvector that corresponds to $\lambda_1 = 1$, we must solve the null space problem $(A - \mathbb{1})\vec{v} = \vec{0}$. We start from $(A - \mathbb{1}) = \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix}$ and after two row operations find the reduced row echelon form $\begin{bmatrix} 1 & -1 \\ 0 & 0 \end{bmatrix}$. The second column is a free variable, which we'll call s , and we find the solution to the null space problem is $\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} s \\ s \end{bmatrix}$. The $\lambda = 1$ eigenspace is spanned by the eigenvector $\vec{e}_{\lambda_1} = (1, 1)^T$. The procedure for finding the λ_2 eigenvector is similar.

P6.7 First we decompose A as the product of three matrices $A = Q\Lambda Q^{-1}$, where Q is a matrix of eigenvectors, and Λ contains the eigenvalues of A . $A = \begin{bmatrix} 2 & 2 \\ 5 & -1 \end{bmatrix} = \begin{bmatrix} -2 & 1 \\ 5 & 1 \end{bmatrix} \begin{bmatrix} -3 & 0 \\ 0 & 4 \end{bmatrix} \begin{bmatrix} -1/7 & 1/7 \\ 5/7 & 2/7 \end{bmatrix}$. Since the matrix Λ is diagonal, we can compute its 10th power: $\Lambda^{10} = \begin{bmatrix} 59049 & 0 \\ 0 & 1048576 \end{bmatrix}$. Thus the calculation of A^{10} is

$$A^{10} = \begin{bmatrix} -2 & 1 \\ 5 & 1 \end{bmatrix} \begin{bmatrix} 59049 & 0 \\ 0 & 1048576 \end{bmatrix} \begin{bmatrix} -\frac{1}{7} & \frac{1}{7} \\ \frac{5}{7} & \frac{2}{7} \end{bmatrix} = \begin{bmatrix} 765854 & 282722 \\ 706805 & 341771 \end{bmatrix}.$$

P6.8 The eigenvalues of M are $\frac{1}{4}$, $\frac{1}{2}$, and 1, and its eigendecomposition is $M = Q\Lambda Q^{-1}$. We can compute $(x_\infty, y_\infty, z_\infty)^T$ using $M^\infty(x_0, y_0, z_0)^T$. To compute M^∞ , we can compute Λ^∞ . The $\frac{1}{4}$ and $\frac{1}{2}$ eigenspaces will disappear, leaving us only with the subspace of the eigenvalue 1. $M^\infty = Q\Lambda^\infty Q^{-1}$, and each row of this matrix has the form $[\frac{1}{6}, \frac{4}{6}, \frac{1}{6}]$. See `bit.ly/eigenex001` for the details.

P6.9 The characteristic polynomial of the matrix has degree n , and an n^{th} -degree polynomial has at most n distinct roots.

P6.10 If T is invertible, then T is bijective and there is a one-to-one correspondence between every input vector $\vec{v}$ in the domain of T and some vector $\vec{w}$ in the image space of T . Since T is invertible, the equation $T(\vec{v}) = \vec{0}$ is satisfied by the unique vector $\vec{v} = \vec{0}$ so $\lambda = 0$ is not an eigenvalue of T . To show the other direction, observe that if $\lambda = 0$

is an eigenvalue of T there must exist a vector $\vec{v} \neq 0$ such that $T(\vec{v}) = 0\vec{v} = \vec{0}$. Since $S(\vec{0}) = \vec{0}$ for all linear maps, it is impossible to construct a map S that can take $\vec{0}$ to $\vec{v}$ as required for the inverse of T . Therefore T^{-1} doesn't exist.

P6.11 Use the facts that $\text{Tr}(A) = \lambda_1 + \lambda_2 + \lambda_3$ and $\det(A) = \lambda_1\lambda_2\lambda_3$, and the properties of trace and determinant (page 150) to compute the expressions.

P6.12 If A is a diagonal matrix, we have $A_{ij} = 0 = A_{ji}$ when $i \neq j$. Therefore diagonal matrices are symmetric.

P6.13 To check whether the matrix is orthogonal, check whether the transpose of the matrix is its inverse so that $O^T O = \mathbb{1}$.

P6.14 If M is normal, we can write it as $M = O\Lambda O^T$, where O is orthogonal and Λ is diagonal. Compute M raised to the power k as follows: $M^k = O\Lambda O^T O\Lambda O^T \cdots O\Lambda O^T = O\Lambda^k O^T$, since the inner O s cancel with the O^T s. We see $\mathcal{C}(M) = \mathcal{C}(O) = \mathcal{C}(M^k)$ and $\mathcal{N}(M) = \mathcal{N}(O^T) = \mathcal{N}(M^k)$.

P6.15 Show this directly: $\vec{u}^T(A+B)\vec{u} = \vec{u}^T A\vec{u} + \vec{u}^T B\vec{u} \geq 0$, for all $\vec{u}$.

P6.16 An orthogonal matrix O satisfies both $O^T O = \mathbb{1}$ and $OO^T = \mathbb{1}$.

P6.17 We're looking for a matrix B such that $B^2 O = \mathbb{1}$. Since O is orthogonal, we know $OO^T = O^T O = \mathbb{1}$, so the inverse matrix B is $\frac{1}{2}O^T$.

P6.18 This isn't a difficult question; you just need to go through the motions.

P6.19 A vector space would obey $0 \cdot (a_1, a_2) = (0, 0)$ (the zero vector), but we have $0 \cdot (a_1, a_2) = (0, a_2) \neq (0, 0)$, so $(V, \mathbb{R}, +, \cdot)$ is not a vector space.

P6.20 The vector addition operation is not associative: we have $((a_1, a_2) + (b_1, b_2)) + (c_1, c_2) = (a_1 + 2b_1 + 2c_1, a_2 + 3b_2 + 3c_2)$ but $(a_1, a_2) + ((b_1, b_2) + (c_1, c_2)) = (a_1 + 2b_1 + 4c_1, a_2 + 3b_2 + 9c_2)$.

P6.21 A subspace of $\mathbb{R}^3$ must be closed under addition and scalar multiplication, and must contain the zero element.

P6.24 Start with two arbitrary elements of $P_2(t)$ —like $\mathbf{p} = a_0 + a_1 t + a_2 t^2$ and $\mathbf{q} = b_0 + b_1 t + b_2 t^2$ —and show that $P_2(t)$ is closed under addition and scalar multiplication.

P6.25 Consider the union of the points of the x -axis and the y -axis. Computing the sum of vectors taken from different axes results in a vector between the axes, which exists outside the subset. Therefore, the subset is not closed under addition, and is therefore not a vector subspace.

P6.26 The integers are closed under addition: $m + n \in \mathbb{Z}$ for all $m, n \in \mathbb{Z}$. However, choosing a non-integer scale factor α will result in $\alpha m \notin \mathbb{Z}$.

P6.27 Given an arbitrary input $\vec{v} = v_0 + v_1 x + v_2 x^2$, the effect of T is to select the quadratic term, so it corresponds to a projection matrix onto the third dimension. The eigenvalue $\lambda = 1$ corresponds to the subspace spanned by $(0, 0, 1)$. The eigenvalue $\lambda = 0$ corresponds to the two-dimensional subspace spanned by $(1, 0, 0)$ and $(0, 1, 0)$.

P6.28 The solutions to the differential equation $f'(t) + f(t) = 0$ are of the form $f(t) = Ce^{-t}$, where C is an arbitrary constant. Since any solution in the *solution space* can be written as a multiple of the function e^{-t} , we say the solution space is *spanned* by e^{-t} . Since one function is sufficient to span the solution space, the solution space is one-dimensional.

P6.29 We apply L to both $e^{2x} \cos x$ and $e^{2x} \sin x$ to obtain

$$L(e^{2x} \cos x) = 2e^{2x} \cos x - e^{2x} \sin x + e^{2x} \cos x = 3e^{2x} \cos x - e^{2x} \sin x$$

$$L(e^{2x} \sin x) = 2e^{2x} \sin x + e^{2x} \cos x + e^{2x} \sin x = e^{2x} \cos x + 3e^{2x} \sin x.$$

When the input is $e^{2x} \cos x$, the output of L corresponds to the vector $(3, -1)^T$ with respect to the basis $\{e^{2x} \cos x, e^{2x} \sin x\}$. The second output corresponds to the vector $(1, 3)^T$ with respect to this basis. These vectors are columns of the matrix representing L . Thus, the matrix representing L with respect to the given basis is $\begin{bmatrix} 3 & 1 \\ -1 & 3 \end{bmatrix}$.

P6.31 First of all we must determine dimensionality of the vector space in question. The general vector space of 3×3 matrices has nine dimensions, but a diagonal matrix A satisfies $a_{ij} = 0$ for all $i \neq j$, which corresponds to the following six constraints: $\{a_{12} = 0, a_{13} = 0, a_{21} = 0, a_{23} = 0, a_{31} = 0, a_{32} = 0\}$. The space of diagonal matrices is a subspace of a nine-dimensional vector space defined by six constraints; therefore it is three-dimensional. The answer given is the standard basis. Other answers are possible: any set of three vectors that span the same space would do.

P6.32 A matrix $A \in \mathbb{R}^{3 \times 3}$ is symmetric if and only if $A^T = A$. This means we can pick the entries on the diagonal arbitrarily, but the symmetry requirement leads to the constraints $a_{12} = a_{21}$, $a_{13} = a_{31}$, and $a_{23} = a_{32}$. Thus the space of 3×3 symmetric matrices is six-dimensional.

P6.33 A Hermitian matrix H is a matrix with complex entries that satisfies $H = H^\dagger$, or equivalently $h_{ij} = \overline{h_{ji}}$, for all i, j . A priori the space of 3×3 matrices with complex entries is 18-dimensional, for the real and imaginary parts of each of the nine entries. The Hermitian property imposes the following constraints. First, diagonal entries must be real if we want $h_{ii} = \overline{h_{ii}}$ to be true, which introduces three constraints. Second, once we pick the real and imaginary parts of an off-diagonal entry a_{ij} , we're forced to choose $a_{ji} = \overline{a_{ij}}$, leading to six more constraints. Thus, the vector space of 3×3 Hermitian matrices is $18 - 3 - 6 = 9$ -dimensional.

P6.36 a) Using linearity: $\langle \vec{v}_1, 2\vec{v}_2 + 3\vec{v}_3 \rangle = 2\langle \vec{v}_1, \vec{v}_2 \rangle + 3\langle \vec{v}_1, \vec{v}_3 \rangle = 6 + 3 = 9$.

b) Using linearity in both entries: $\langle 2\vec{v}_1 - \vec{v}_2, \vec{v}_1 + \vec{v}_3 \rangle = 2\langle \vec{v}_1, \vec{v}_1 \rangle + 2\langle \vec{v}_1, \vec{v}_3 \rangle - \langle \vec{v}_2, \vec{v}_1 \rangle - \langle \vec{v}_2, \vec{v}_3 \rangle = 2 + 2 - 3 + 2 = 3$.

c) We start with $13 = \langle \vec{v}_2, \vec{v}_1 + \vec{v}_2 \rangle = \langle \vec{v}_2, \vec{v}_1 \rangle + \langle \vec{v}_2, \vec{v}_2 \rangle$, and since we know $\langle \vec{v}_1, \vec{v}_2 \rangle = 3$, we obtain $3 + \|\vec{v}_2\|^2 = 13$, so $\|\vec{v}_2\|^2 = 10$ and $\|\vec{v}_2\| = \sqrt{10}$.

P6.37 If $\mathbf{v} = \mathbf{0}$, then the inequality holds trivially. If $\mathbf{v} \neq \mathbf{0}$, we can start from the following, which holds for any $c \in \mathbb{C}$:

$$\begin{aligned} 0 &\leq \langle \mathbf{u} - c\mathbf{v}, \mathbf{u} - c\mathbf{v} \rangle \\ &= \langle \mathbf{u}, \mathbf{u} - c\mathbf{v} \rangle - c\langle \mathbf{v}, \mathbf{u} - c\mathbf{v} \rangle \\ &= \langle \mathbf{u}, \mathbf{u} \rangle - \overline{c}\langle \mathbf{u}, \mathbf{v} \rangle - c\langle \mathbf{v}, \mathbf{u} \rangle + |c|^2\langle \mathbf{v}, \mathbf{v} \rangle. \end{aligned}$$

This is true in particular when $c = \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle}$, so we continue:

$$\begin{aligned} 0 &\leq \langle \mathbf{u}, \mathbf{u} \rangle - \frac{\langle \mathbf{u}, \mathbf{v} \rangle \overline{\langle \mathbf{u}, \mathbf{v} \rangle}}{\langle \mathbf{v}, \mathbf{v} \rangle} - \frac{\langle \mathbf{u}, \mathbf{v} \rangle \langle \mathbf{u}, \mathbf{v} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle} + \frac{|\langle \mathbf{u}, \mathbf{v} \rangle|^2 \langle \mathbf{v}, \mathbf{v} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle^2} \\ 0 &\leq \langle \mathbf{u}, \mathbf{u} \rangle - \frac{|\langle \mathbf{u}, \mathbf{v} \rangle|^2}{\langle \mathbf{v}, \mathbf{v} \rangle} - \cancel{\frac{|\langle \mathbf{u}, \mathbf{v} \rangle|^2}{\langle \mathbf{v}, \mathbf{v} \rangle}} + \cancel{\frac{|\langle \mathbf{u}, \mathbf{v} \rangle|^2}{\langle \mathbf{v}, \mathbf{v} \rangle}} \\ 0 &\leq \langle \mathbf{u}, \mathbf{u} \rangle - \frac{|\langle \mathbf{u}, \mathbf{v} \rangle|^2}{\langle \mathbf{v}, \mathbf{v} \rangle} \\ 0 &\leq \langle \mathbf{u}, \mathbf{u} \rangle \langle \mathbf{v}, \mathbf{v} \rangle - |\langle \mathbf{u}, \mathbf{v} \rangle|^2, \end{aligned}$$

from which we conclude $|\langle \mathbf{u}, \mathbf{v} \rangle|^2 \leq \|\mathbf{u}\|^2 \|\mathbf{v}\|^2$. Taking the square root on both sides of this inequality, we obtain the statement of the Cauchy-Schwarz inequality, $|\langle \mathbf{u}, \mathbf{v} \rangle| \leq \|\mathbf{u}\| \|\mathbf{v}\|$.

P6.38 We proceed using the following chain of inequalities:

$$\begin{aligned} \|\mathbf{u} + \mathbf{v}\|^2 &= \langle \mathbf{u} + \mathbf{v}, \mathbf{u} + \mathbf{v} \rangle \\ &= \langle \mathbf{u}, \mathbf{u} \rangle + \langle \mathbf{u}, \mathbf{v} \rangle + \langle \mathbf{v}, \mathbf{u} \rangle + \langle \mathbf{v}, \mathbf{v} \rangle \\ &\leq \|\mathbf{u}\|^2 + 2\langle \mathbf{u}, \mathbf{v} \rangle + \|\mathbf{v}\|^2 \\ &\leq \|\mathbf{u}\|^2 + 2\|\mathbf{u}\| \|\mathbf{v}\| + \|\mathbf{v}\|^2 \\ &= (\|\mathbf{u}\| + \|\mathbf{v}\|)^2. \end{aligned}$$

Therefore we obtain the equation $\|\mathbf{u} + \mathbf{v}\|^2 \leq (\|\mathbf{u}\| + \|\mathbf{v}\|)^2$, and since $\|\mathbf{u}\|$ and $\|\mathbf{v}\|$ are nonnegative numbers, we can take the square root on both sides of the inequality to obtain $\|\mathbf{u} + \mathbf{v}\| \leq \|\mathbf{u}\| + \|\mathbf{v}\|$.

P6.39 This is a trick question: we don't really need to perform the Gram-Schmidt procedure since $\vec{v}_2$ is already perpendicular to $\vec{v}_1$ and both vectors have length one.

P6.40 We're given the vectors $\vec{v}_1 = (1, 1)$ and $\vec{v}_2 = (0, 1)$ and want to perform the Gram-Schmidt procedure. We pick $\vec{e}_1 = \vec{v}_1 = (1, 1)$, and after normalization we have $\hat{e}_1 = (\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$. Next we compute $\vec{e}_2 = \vec{v}_2 - \Pi_{\hat{e}_1}(\vec{v}_2) = \vec{v}_2 - (\hat{e}_1 \cdot \vec{v}_2)\hat{e}_1 = (\frac{-1}{2}, \frac{1}{2})$. Normalizing $\vec{e}_2$ we obtain $\hat{e}_2 = (-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$.

P6.41 Let $\vec{v}_1 = (3, 1)$ and $\vec{v}_2 = (-1, 1)$. We start by identifying $\vec{e}_1 = \vec{v}_1$, then perform Gram-Schmidt process to find $\vec{e}_2$ from $\vec{v}_2$:

$$\vec{e}_2 = \vec{v}_2 - \Pi_{\hat{e}_1}(\vec{v}_2)\hat{e}_1 = \vec{v}_2 - \left(\frac{\vec{e}_1}{\|\vec{e}_1\|} \cdot \vec{v}_2 \right) \frac{\vec{e}_1}{\|\vec{e}_1\|} = \left(\frac{-2}{5}, \frac{6}{5} \right).$$

Now we have two orthogonal vectors and we can normalize them to make them unit vectors. We obtain the vectors $(\frac{3}{\sqrt{10}}, \frac{1}{\sqrt{10}})$ and $(\frac{-1}{\sqrt{10}}, \frac{3}{\sqrt{10}})$, which form an orthogonal basis.

P6.43 First find the eigenvalues of the matrix. Then find an eigenvector for each eigenvalue and construct a matrix Q composed of the three eigenvectors. Compute Q^{-1} and write the eigendecomposition as $A = Q\Lambda Q^{-1}$.

P6.44 a) $\sqrt{3^2 + 4^2} = 5$. b) Complex conjugation changes the sign of the imaginary part: $\overline{2 - 3i} = 2 + 3i$. c) $3i - 1 + 3 + 2i = 2 + 5i$. d) $|3i - 4i - 5| = |-5 - i| = \sqrt{26}$.

P6.47 The characteristic polynomial of A is $p_A(\lambda) = \lambda^2 - 2\lambda = \lambda(\lambda - 2)$, so the eigenvalues are 0 and 2. The eigenvector for $\lambda = 0$ is $\vec{e}_0 = (-1 + i, 2)^T$. The eigenvector for $\lambda = 2$ is $\vec{e}_2 = (1 + i, 2)^T$. As A has two one-dimensional eigenspaces, an eigenbasis is given by $\{(-1 + i, 2)^T, (1 + i, 2)^T\}$. So A is diagonalizable with $Q = \begin{bmatrix} -1+i & 1+i \\ 2 & 2 \end{bmatrix}$ and $\Lambda = \begin{bmatrix} 0 & 0 \\ 0 & 2 \end{bmatrix}$.

P6.48 Observe $P_2(x)$ is three-dimensional; therefore, if B_a is a linearly independent set, then it is a basis as it has three elements. If $a(1 + ix) + b(1 + x + ix^2) + c(1 + 2ix) = 0$ then $(a + b + c)1 + (ia + b + 2ic)x + ibx^2 = 0$. But the standard basis of $P_2(x)$ is linearly independent, so we must have $a + b + c = 0$, $ia + b + 2ic = 0$, and $ib = 0$. The last equation implies $b = 0$, and the first two imply both a and c are zero. As the only linear combination of distinct elements of B_a which sums to zero is the trivial sum, B_a is linearly independent.

P6.49 Computing the determinant of $A - \lambda\mathbb{1}$, we obtain $p_A(\lambda) = \lambda^2 - \lambda(a + d) + ad - bc$. We obtain the condition by computing the discriminant of this quadratic equation.

P6.50 Since M is normal, we can write it as $M = U\Lambda U^\dagger$, where U is a unitary matrix and Λ is diagonal. Taking the transpose of this equation we find $M^\dagger = (U\Lambda U^\dagger)^\dagger = U\Lambda^\dagger U^\dagger$. Thus, $\mathcal{C}(M) = \mathcal{C}(U) = \mathcal{C}(M^\dagger)$.

P6.51 If A is Hermitian, we know $A^\dagger = A$. Now suppose $\vec{e}$ is an eigenvector of A associated with a nonzero eigenvalue λ . The eigenvalue equation is $A\vec{e} = \lambda\vec{e}$. Taking the Hermitian transpose of both sides of the eigenvalue equation, we obtain $(A\vec{e})^\dagger = \vec{e}^\dagger A^\dagger = \vec{e}^\dagger A = \bar{\lambda}\vec{e}^\dagger$. Now suppose we multiply the eigenvalue equation by $\vec{e}^\dagger$ from the left. We obtain $\vec{e}^\dagger A\vec{e} = \lambda\vec{e}^\dagger\vec{e} = \lambda\|\vec{e}\|^2$. We can also interpret the triple product $\vec{e}^\dagger A\vec{e}$ as $(\vec{e}^\dagger A)\vec{e}$. We know the expression in the brackets equals $\bar{\lambda}\vec{e}^\dagger$, and overall we have $(\vec{e}^\dagger A)\vec{e} = \bar{\lambda}\vec{e}^\dagger\vec{e} = \bar{\lambda}\|\vec{e}\|^2$. We've thus obtained the equation $\bar{\lambda}\|\vec{e}\|^2 = \lambda\|\vec{e}\|^2$. Since $\|\vec{e}\|^2 \neq 0$, it must be that $\lambda = \bar{\lambda}$, which means λ is real. The realness of Hermitian eigenvalues is a useful property used in quantum physics. To ensure the energy of a physical system can be expressed as a real number, physicists require that energy operators be Hermitian matrices.

P6.52 Assume $\vec{e}$ is an eigenvector associated with a nonzero eigenvalue λ of a unitary matrix U . The eigenvalue equation tells us $U\vec{e} = \lambda\vec{e}$. Computing the squared length of both sides of this equation, we find $\|U\vec{e}\|^2 = \vec{e}^\dagger U^\dagger U \vec{e} = \vec{e}^\dagger \mathbb{1} \vec{e} = \vec{e}^\dagger \vec{e} = \|\vec{e}\|^2$ for the left side, and $\|\lambda\vec{e}\|^2 = \vec{e}^\dagger \bar{\lambda} \lambda \vec{e} = |\lambda|^2 \|\vec{e}\|^2$ for the right side. Combining these results leads us to the equation $\|\vec{e}\|^2 = |\lambda|^2 \|\vec{e}\|^2$; and since $\|\vec{e}\|^2 \neq 0$, this means $|\lambda|^2 = 1$, which implies $|\lambda| = 1$.

P6.53 To prove a matrix is nilpotent, you can compute its powers $A^2, A^3, A^4, \dots$ to see if you obtain the zero matrix. To prove a matrix is not nilpotent, you can show it has a nonzero eigenvalue.

- a) This matrix is nilpotent because its square is the zero matrix.
- b) The matrix is not nilpotent because its characteristic polynomial $p(\lambda) = \lambda^2 - 6\lambda + 8$ has a nonzero root.
- c) This matrix is nilpotent because it squares to the zero matrix.
- d) The matrix is not nilpotent because it has nonzero eigenvalues.
- e) Yes; the cube of this matrix is the zero matrix.
- f) Yes; the square of this matrix is the zero matrix.

P6.54 The key fact we need to remember is that the eigenvectors of normal matrices form a full basis. Since M is normal, it can be diagonalized: $M = Q\Lambda Q^{-1}$. The “square-root of M ” operator is obtained by taking the square root of the eigenvalues matrix: $\sqrt{M} = Q\sqrt{\Lambda}Q^{-1}$.

Chapter 7 solutions

Answers to exercises

E7.1 $4\text{Al} + 3\text{O}_2 \rightarrow 2\text{Al}_2\text{O}_3$. **E7.2** $\text{Fe}(\text{OH})_3 + 3\text{HCl} \rightarrow \text{FeCl}_3 + 3\text{H}_2\text{O}$. **E7.3** $x_a = 106.3$ tons, $x_b = 281.2$ tons, and $x_c = 352.8$ tons. **E7.4** $+10 - R_1 I_2 + 5 - R_3 I_3 - 20 = 0$. No, because the three KVL equations are not linearly independent. **E7.5** See solution.

E7.6 a) $A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}$; b) $A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix}$; c) $A = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}$. **E7.7** a) 1; b) 1; c) 2. **E7.8**

$\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n} = \lambda_1 = \varphi = \frac{1+\sqrt{5}}{2}$. **E7.9** $a_N = \frac{1}{\sqrt{5}} \left(\frac{1+\sqrt{5}}{2} \right)^N - \frac{1}{\sqrt{5}} \left(\frac{1-\sqrt{5}}{2} \right)^N$. **E7.10**

$h(x) = 0.52x$. **E7.11** $y = 4 - x$. **E7.12** $S(m^*) = 4704.63$. **E7.13** $S(\bar{m}^*) = 433.54$.

E7.14 $\vec{c} = 1100 \ 1111 \ 0100 \ 1110$. **E7.15** $\vec{c} = \vec{x}G = (1, 0, 1, 0, 1, 0, 1)$. **E7.16** $G =$

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

Solutions to selected exercises

E7.3 See <http://bit.ly/vegonomics> for the calculation steps.

E7.4 This exercise shows how the abstract notion of linear independence applies to electric circuits. We must choose the KVL equations that describe the current flowing in linearly independent loops. There are three loops in the circuit from which we can obtain KVL equations: the left inner loop, the right inner loop, and the perimeter of the circuit as a whole. The KVM equation for the outer loop is a linear combination of the KVL equations of the two inner loops.

E7.5 The KVL equation for the clockwise loop starting at junction A is $V_1 - R_2 I_2 - R_1 I_1 = 0$. The KVL equation starting at B is $-R_3 I_3 + V_2 - R_4 I_4 + R_2 I_2 = 0$. The equation for the loop starting at C is $V_3 - R_5 I_5 + R_4 I_4 = 0$. Observe the voltage drop across R_2 and R_4 appears as gains and drops in the KVL equations. The KCL equation for junction B is $I_1 = I_2 + I_3$ and that for C is $I_3 = I_4 + I_5$. After combining

the equations and performing some simplifications, we obtain the matrix equation

$$R\vec{I} = \vec{V}, \text{ where } R = \begin{bmatrix} R_1 & R_2 & 0 & 0 & 0 \\ 0 & -R_2 & R_3 & R_4 & 0 \\ 0 & 0 & 0 & -R_4 & R_5 \\ 1 & -1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 & -1 \end{bmatrix}.$$

E7.6 The i^{th} row of the adjacency matrix contains the information about the outgoing edges for vertex i in the graph. If the edge (i, j) exists, then $A_{ij} = 1$, otherwise $A_{ij} = 0$.

E7.8 Since $|\lambda_2| < 1$, as n increases the number λ_2^n tends to zero, and only the effects of λ_1^n remain. The ratio $\frac{a_{n+1}}{a_n}$ approaches λ_1 as $n \rightarrow \infty$.

E7.9 Compute the product $Q\Lambda^N Q^{-1}$ in the order $(Q\Lambda^N)Q^{-1}$. To find a_N , we're only interested in computing the bottom left entry of this matrix. We obtain the formula $a_N = \frac{5+\sqrt{5}}{10} \left(\frac{1+\sqrt{5}}{2}\right)^{N-1} + \frac{5-\sqrt{5}}{10} \left(\frac{1-\sqrt{5}}{2}\right)^{N-1}$, which simplifies to $a_N = \frac{\varphi^N - (-\varphi)^{-N}}{\sqrt{5}}$.

E7.10 The data you're given consists of pairs (a_i, b_i) , where a_i is the total number of flips in the trial and b_i is the number of heads. We represent the data as a matrix of flip counts $A \in \mathbb{R}^{5 \times 1}$ and a vector of heads counts $\vec{b} \in \mathbb{R}^5$. Applying the standard approach for finding the Moore–Penrose inverse, we find $(A^T A')^{-1} = [\frac{1}{3520}]$. Next we find the least squares approximate solution to $A\vec{m} = \vec{b}$ by computing $\vec{m} = (A^T A')^{-1} A^T \vec{b} = [0.52]$. The slope of the line of best fit is 0.52.

E7.11 The dataset consists of a matrix of inputs x , $A = (0, 1, 2)^T$, and a vector of outputs y , $\vec{b} = (3.9, 3.2, 1.9)^T$. Since we're interested in fitting an affine model $y = b + mx$, we must augment the matrix A with a column of ones to obtain A' , and then compute $A'^T A'$ and its inverse:

$$A' = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix}, \quad A'^T A' = \begin{bmatrix} 3 & 3 \\ 3 & 5 \end{bmatrix} \quad \Rightarrow \quad (A'^T A')^{-1} = \begin{bmatrix} \frac{5}{6} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{bmatrix}.$$

We can now compute the Moore–Penrose inverse A'^+ and obtain the approximate solution as follows:

$$A'^+ = (A'^T A')^{-1} A'^T = \begin{bmatrix} \frac{5}{6} & \frac{1}{3} & -\frac{1}{6} \\ -\frac{1}{2} & 0 & \frac{1}{2} \end{bmatrix} \quad \Rightarrow \quad A'^+ \vec{b} = \begin{bmatrix} 4 \\ -1 \end{bmatrix}.$$

Thus the best-fitting affine model for the dataset is $y = 4 - x$.

E7.16 The procedure for encoding each four-bit message $\vec{x} = x_1 x_2 x_3 x_4$ into a five-bit codeword $\vec{c} = c_1 c_2 c_3 c_4 c_5$ is as follows: set $c_i = x_i$ for the first four bits, and set the fifth bit to $c_5 = x_1 + x_2 + x_3 + x_4$.

E7.17 To find $\langle \mathbf{e}_1(x), \mathbf{e}_2(x) \rangle$, we must compute the integral $\int_0^L \sin(\frac{\pi}{L}x) \sin(\frac{2\pi}{L}x) dx$. Change variables to $y = \frac{\pi}{L}x$ to simplify. The integral becomes $\langle \mathbf{e}_1(x), \mathbf{e}_2(x) \rangle = k \int_0^\pi \sin(y) \sin(2y) dy$, where $k = \frac{L}{\pi}$. Using the double-angle formula, we find $\sin(y) \sin(2y) = \sin(y) 2 \sin(y) \cos(y) = 2 \sin^2(y) \cos(y)$. Next, use the substitution $u = \sin(y)$, $du = \cos(y) dy$ to obtain $2k \int \sin^2(y) \cos(y) dy = 2k \int u^2 du = \frac{2k}{3} [u^3] = \frac{2k}{3} \sin^3(y)$. Finally, evaluating the expression at the endpoints, we find $\langle \mathbf{e}_1(x), \mathbf{e}_2(x) \rangle = \frac{2k}{3} [\sin^3(\pi) - \sin^3(0)] = 0$. This confirms that $\mathbf{e}_1(x)$ and $\mathbf{e}_2(x)$ are orthogonal.


E7.18 The coefficient b_0 corresponds to the basis function $\sin(\frac{0\pi}{T}t)$, which is zero everywhere. Thus, the integral $b_0 = \frac{1}{T} \int_0^T f(t) \sin(\frac{2\pi}{T}t) dt$, always produces a zero result $b_0 = 0$. On the other hand, the zero-frequency cos function is constant— $\cos(\frac{2\pi}{T}t) = 1$ —and the coefficient $a_0 = \frac{1}{T} \int_0^T f(t) 1 dt$ corresponds to the average value of the function $f(t)$ on the interval $[0, T]$.

Answers to problems

P7.1 $\text{C}_{55}\text{H}_{104}\text{O}_6 + 78\text{O}_2 \rightarrow 55\text{CO}_2 + 52\text{H}_2\text{O}$. **P7.2** $I_1 = \frac{85}{11} \approx 7.7[\text{A}]$ and $I_5 = \frac{25}{11} \approx 2.27[\text{A}]$. **P7.3** $y = 1 + \frac{1}{2}x$. **P7.4** $p(x) = 334 + 1.001x$; $p(700) = \$1035$.

P7.5 $q(x) = 174 + 1.69x - 0.000558x^2$; $q(700) = \$1084$. **P7.7 a)** $A = \begin{bmatrix} 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$;

b) $A = \begin{bmatrix} 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$; **c)** $A = \begin{bmatrix} 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 \end{bmatrix}$. **P7.8 a) 1; b) 1; c) 3.** **P7.9 a)**



P7.10 a) $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$; **b)** $M_1 = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$; **c)** $M_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$;

d) $M_3 = \begin{bmatrix} \cos(\frac{\pi}{3}) & -\sin(\frac{\pi}{3}) & 0 \\ \sin(\frac{\pi}{3}) & \cos(\frac{\pi}{3}) & 0 \\ 0 & 0 & 1 \end{bmatrix}$; **e)** $M_4 = \begin{bmatrix} \cos(-\frac{\pi}{6}) & -\sin(-\frac{\pi}{6}) & 0 \\ \sin(-\frac{\pi}{6}) & \cos(-\frac{\pi}{6}) & 0 \\ 0 & 0 & 1 \end{bmatrix}$; **f)** $M_5 = \begin{bmatrix} 1 & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & 0 \\ 0 & 0 & 1 \end{bmatrix}$. **P7.11**

a) $M_6 = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$; **b)** $M_7 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{bmatrix}$; **c)** $M_8 = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{bmatrix}$; **d)** $M_4 = \begin{bmatrix} \cos(-\frac{\pi}{6}) & -\sin(-\frac{\pi}{6}) & 0 \\ \sin(-\frac{\pi}{6}) & \cos(-\frac{\pi}{6}) & 0 \\ 0 & 0 & 1 \end{bmatrix}$;

e) $M_9 = \begin{bmatrix} \cos(-\frac{\pi}{6}) & -\sin(-\frac{\pi}{6}) & 0 \\ \sin(-\frac{\pi}{6}) & \cos(-\frac{\pi}{6}) & 2 \\ 0 & 0 & 1 \end{bmatrix}$; **f)** $M_{10} = \begin{bmatrix} \cos(-\frac{\pi}{6}) & -\sin(-\frac{\pi}{6}) & -2\sin(-\frac{\pi}{6}) \\ \sin(-\frac{\pi}{6}) & \cos(-\frac{\pi}{6}) & 2\cos(-\frac{\pi}{6}) \\ 0 & 0 & 1 \end{bmatrix}$. **P7.12**

$P = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & \frac{1}{2} & 0 \end{bmatrix}$; $p' = (\frac{2}{5}, \frac{6}{5})$ and $q' = (\frac{3}{4}, \frac{1}{2})$. **P7.13** Alice sent $\vec{m} = 01110011 \ 01110101$

01110000. After ASCII decoding each byte of this binary string, we find it corresponds to the message “sup.” **P7.15** $f(\omega) = \sigma e^{-\frac{\sigma^2 \omega^2}{2}}$. **P7.16** $f(\omega) = \frac{1}{\sqrt{2\pi}} \frac{2|a|}{|a|^2 + \omega^2}$. **P7.17** Fourier transform of $f'(t)$ is $(i\omega)f(\omega)$.

Solutions to selected problems

P7.1 The reaction we’re studying is the burning of triglycerides. Read more about that here: <http://bmj.com/content/349/bmj.g7257>.

P7.2 The KVL equation for the clockwise loop starting at junction A is $V_1 - R_1 I_1 - R_2 I_2 = 0$. The equation for the loop starting at C is $-R_3 I_3 - R_4 I_4 + R_2 I_2 = 0$. The equation starting at D is $+R_5 I_5 - V_2 + R_4 I_4 = 0$. The KCL equation for junction C is $I_1 = I_2 + I_3$ and that for D is $I_3 + I_5 = I_4$. After combining the equations, we obtain

the matrix equation $R\vec{I} = \vec{V}$, where $\vec{V} = \begin{bmatrix} 15 \\ 0 \\ 10 \\ 0 \\ 0 \end{bmatrix}$ and $R = \begin{bmatrix} R_1 & R_2 & 0 & 0 & 0 \\ 0 & -R_2 & R_3 & R_4 & 0 \\ 0 & 0 & 0 & R_4 & R_5 \\ 1 & -1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 & 1 \end{bmatrix}$, which

becomes $R = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & -1 & 4 & 2 & 0 \\ 0 & 0 & 0 & 2 & 2 \\ 1 & -1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 & 1 \end{bmatrix}$. We compute $\vec{I} = R^{-1}\vec{V}$ and the answer is $I_1 = \frac{85}{11} \approx 7.7$

and $I_5 = \frac{25}{11} \approx 2.27$.

P7.3 The dataset consists of a matrix of inputs x , $A = (0, 1, 2, 3)^T$, and a vector of outputs y , $\vec{b} = (0.9, 1.6, 2.1, 2.4)^T$. Since we’re interested in fitting an affine model $y = b + mx$, we must augment the matrix A with a column of ones to obtain A' , and then compute $A'^T A'$ and its inverse:

$$A' = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{bmatrix}, \quad A'^T A' = \begin{bmatrix} 4 & 6 \\ 6 & 14 \end{bmatrix} \Rightarrow (A'^T A')^{-1} = \begin{bmatrix} \frac{7}{10} & -\frac{3}{10} \\ -\frac{3}{10} & \frac{1}{5} \end{bmatrix}.$$

We can now compute the Moore–Penrose inverse A'^+ and obtain the approximate solution as follows:

$$A'^+ = (A'^T A')^{-1} A'^T = \begin{bmatrix} \frac{7}{10} & \frac{2}{5} & \frac{1}{10} & -\frac{1}{5} \\ -\frac{3}{10} & -\frac{1}{10} & \frac{1}{10} & \frac{3}{10} \end{bmatrix} \Rightarrow A'^+ \vec{b} = \begin{bmatrix} 1 \\ \frac{1}{2} \end{bmatrix}.$$

Thus the best-fitting affine model for the dataset is $y = 1 + \frac{1}{2}x$.

P7.4 This is a standard least squares problem with an affine model. Visit this link to see the calculations, bit.ly/apt_rent_affine; and this link to see the graph of the best-fitting affine model, bit.ly/apt_rent_fits.

P7.5 This is a least squares problem with a quadratic model. For the solution obtained using SymPy, see bit.ly/apt_rent_quadratic. See bit.ly/apt_rent_fits for the graph of the best-fitting model.

P7.6 The matrix equation can be either consistent ($\exists \vec{x}$ such that $A\vec{x} = \vec{b}$), or inconsistent ($\nexists \vec{x}$ such that $A\vec{x} = \vec{b}$). If the matrix equation is consistent, then it either has a unique solution $\vec{x} = \vec{x}_p$ if $\mathcal{N}(A) = \{\vec{0}\}$, or an infinite solutions set $\vec{x} = \vec{x}_p + \mathcal{N}(A)$ if $\mathcal{N}(A)$ is non-empty. If the matrix equation is inconsistent, we can still find a least squares approximate solution.

P7.7 The i^{th} row of the adjacency matrix contains the information of the outgoing edges for vertex i in the graph. If the edge (i, j) exists, then $A_{ij} = 1$, otherwise $A_{ij} = 0$.

P7.9 The number of rows (columns) of the adjacency matrix tells us how many vertices the graph contains. The i^{th} row of the adjacency matrix contains the information of the outgoing edges for vertex i in the graph. If you see $A_{ij} = 1$ then you must draw the edge (i, j) in the graph.

P7.10 Your job is to recognize visually what type of transformation is acting in each case, and then use the appropriate matrix formula from Chapter 5.

P7.11 Recall that the third column of homogeneous coordinates matrix representations serves to perform translations. Use the matrix product M_7M_4 (first rotate then translate) to obtain the answer to part (e). Use the matrix product M_4M_7 (first translate then rotate) to obtain the answer to part (f).

P7.12 See bit.ly/persp_proj_prob for the calculations.

P7.14 Decompose the formula for c_n into a its real part and imaginary parts:

$$\begin{aligned} c_n &= \frac{1}{T} \int_0^T f(t) e^{-i \frac{2\pi n}{T} t} dt \\ &= \operatorname{Re} \left\{ \frac{1}{T} \int_0^T f(t) e^{-i \frac{2\pi n}{T} t} dt \right\} + \operatorname{Im} \left\{ \frac{1}{T} \int_0^T f(t) e^{-i \frac{2\pi n}{T} t} dt \right\} i \\ &= \frac{1}{T} \int_0^T f(t) \operatorname{Re} \left\{ e^{-i \frac{2\pi n}{T} t} \right\} dt + \frac{1}{T} \int_0^T f(t) \operatorname{Im} \left\{ e^{-i \frac{2\pi n}{T} t} \right\} i dt \\ &= \frac{1}{T} \int_0^T f(t) \cos \left(\frac{2\pi n}{T} t \right) dt - \frac{1}{T} \int_0^T f(t) \sin \left(\frac{2\pi n}{T} t \right) i dt. \end{aligned}$$

We recognize the real part of c_n as the cosine coefficients a_n of the Fourier series, and the imaginary part of c_n as the negative of the sine coefficients b_n of the Fourier series:

$$\begin{aligned} \operatorname{Re} \{c_n\} &= \operatorname{Re} \left\{ \frac{1}{T} \int_0^T f(t) e^{-i \frac{2\pi n}{T} t} dt \right\} = \frac{1}{T} \int_0^T f(t) \cos \left(\frac{2\pi n}{T} t \right) dt = a_n, \\ \operatorname{Im} \{c_n\} &= \operatorname{Im} \left\{ \frac{1}{T} \int_0^T f(t) e^{-i \frac{2\pi n}{T} t} dt \right\} = \frac{1}{T} \int_0^T f(t) \sin \left(-\frac{2\pi n}{T} t \right) dt = -b_n. \end{aligned}$$

Thus we have shown $c_n = a_n - ib_n$.

Using the simple definition of the coefficients c_n above, the synthesis equation for the complex Fourier transform is the somewhat awkward expression $f(t) = c_0 +$

$\frac{1}{2} \sum_{n=1}^{\infty} e^{-i\frac{2\pi n}{T}t} c_n + \frac{1}{2} \sum_{n=-\infty}^{-1} e^{-i\frac{2\pi n}{T}t} \overline{c_n}$. Many textbooks describe complex Fourier series in terms of the two-sided coefficients $c'_n, n \in \mathbb{Z}$ defined as

$$\begin{aligned} c'_0 &= c_0, \\ c'_n &= \frac{1}{2} c_n = \frac{1}{2} (a_n - ib_n) & \text{for } n \geq 1, \\ c'_n &= \frac{1}{2} \overline{c_{-n}} = \frac{1}{2} (a_{-n} + ib_{-n}) & \text{for } n \leq -1. \end{aligned}$$

Using the coefficients c'_n , the synthesis equation for the complex Fourier series is the simpler expression $f(t) = \sum_{n=-\infty}^{\infty} e^{-i\frac{2\pi n}{T}t} c'_n$.

P7.15 This is a long calculation but straightforward if you write things clearly. The key idea is to rewrite $t^2 + 2\sigma^2 i\omega t$ in the form $(t+h)^2 + k$ so that the constant term $k = \sigma^4 \omega^2$ can be removed from the integral.

$$\begin{aligned} f(\omega) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-i\omega t} e^{-\frac{t^2}{2\sigma^2}} dt = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{1}{2\sigma^2} (t^2 + 2\sigma^2 i\omega t)} dt \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{1}{2\sigma^2} ([t + \sigma^2 i\omega]^2 + \sigma^4 \omega^2)} dt = \frac{1}{\sqrt{2\pi}} e^{-\frac{\sigma^2 \omega^2}{2}} \int_{-\infty}^{\infty} e^{-\frac{1}{2\sigma^2} ([t + \sigma^2 i\omega]^2)} dt \\ &= \frac{1}{\sqrt{2\pi}} e^{-\frac{\sigma^2 \omega^2}{2}} \int_{-\infty}^{\infty} e^{-\frac{1}{2\sigma^2} (t')^2} dt', \quad \text{letting } t' = t + \sigma^2 i\omega \\ &= \frac{1}{\sqrt{2\pi}} e^{-\frac{\sigma^2 \omega^2}{2}} \sqrt{2\pi} \sigma = \sigma e^{-\frac{\sigma^2 \omega^2}{2}}. \end{aligned}$$

P7.16 We start from the definition $f(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-i\omega t} e^{-|at|} dt$. Next we rewrite $e^{-i\omega t}$ as $(\cos(\omega t) - i \sin(\omega t))$ and observe that sine is an odd function so this term vanishes. The integral becomes $f(\omega) = \frac{2}{\sqrt{2\pi}} \int_0^{\infty} \cos(\omega t) e^{-|at|} dt$, which we can tackle using integration by parts.

P7.17 Starting from the formula $f(t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{i\omega t} f(\omega) d\omega$, take the derivative with respect to t to obtain $f'(t) = \frac{d}{dt} \left(\int_{-\infty}^{\infty} e^{i\omega t} f(\omega) d\omega \right) = \int_{-\infty}^{\infty} e^{i\omega t} (i\omega) f(\omega) d\omega = (i\omega) \int_{-\infty}^{\infty} e^{i\omega t} f(\omega) d\omega$. This means the Fourier transform of $f'(t)$ is $(i\omega)f(\omega)$.

The Fourier transform converts differential equations involving $f(t)$ into algebraic equations involving $f(\omega)$. For example, the Fourier transform of the second-order differential equation $\alpha f''(t) + \beta f'(t) + \gamma f(t) = 0$ is a second-order (quadratic) algebraic equation: $\alpha (i\omega)^2 f(\omega) + \beta (i\omega) f(\omega) + \gamma f(\omega) = 0$. I encourage you to learn more about differential equations from the web, or email me if you think I should write a book on this topic.

Chapter 8 solutions

Answers to exercises

E8.1 a) No; weights don't add to one. b) Yes. c) No; contains a negative number.

E8.2 $\mu_X = 3.5$, $\sigma_X^2 = \frac{35}{12} \approx 2.92$. **E8.3** $B = M^T$. **E8.4** $p_{X_{\infty}} = \left(\frac{5}{31}, \frac{20}{31}, \frac{6}{31} \right)$. **E8.5**

$p_{X_{\infty}} = \left(\frac{34}{61}, \frac{14}{61}, \frac{13}{61} \right)$. **E8.6** $p_{X_{\infty}} = (0.2793, 0.1457, 0.2768, 0.02, 0.2781)^T$.

Solutions to selected exercises

E8.4 You can either use `vec = M.eigenvecs()[0][2][0]` to extract the first eigenvector, or `vec = (M-eye(3)).nullspace()[0]`, which is more efficient. To obtain a probability distribution, normalize the eigenvector by dividing it by its ℓ^1 -norm `pinf = vec/vec.norm(1)`.

E8.5 Use `evvec = (C-eye(3)).nullspace()[0]` to extract the eigenvector for $\lambda = 1$, then normalize the vector by its 1-norm to make it a probability distribution `pinf = evvec/evvec.norm(1) =  $(\frac{34}{61}, \frac{14}{61}, \frac{13}{61})$` .

E8.6 You can see the solution at this URL: bit.ly/21G0UCe.

Answers to problems

P8.2 $\Pr(\{\text{heads, heads, heads, heads}\}) = \frac{1}{16}$. **P8.4** $\mathbb{E}_X[X] = \frac{1}{p}$. **P8.5** $\mathbb{E}_X[W] = \$ \frac{100}{101} < \$1$. Not worth it. **P8.6** $\mathbb{E}_X[f(X)] = \frac{\$7}{6} > \$1$ so it's worth playing this game. **P8.8** $p_{X_1} = (\frac{1}{2}, \frac{1}{2})^T$, $p_{X_2} = (\frac{3}{8}, \frac{5}{8})^T$, $p_{X_\infty} = (\frac{1}{3}, \frac{2}{3})^T$. **P8.9** $p_{X_\infty} = (0.1721, 0.169, 0.245, 0.245, 0.169)^T$.

Solutions to selected problems

P8.1 Assume, for a contradiction, that $p_1 > 1$. We know from Kolmogorov's axioms that $p_i \geq 0$ for all i . Observe that $\sum_i p_i = p_1 + p_2 + p_3 \geq p_1 > 1$, which means the vector (p_1, p_2, p_3) is not a valid probability distribution. Therefore, it must be that $p_1 > 1$ is false, and $p_1 \leq 1$ is true.

P8.2 Substitute $p = \frac{1}{2}$ and $n = 4$ into the expression p^n .

P8.3 The biased coin flip is modelled by a random variable Y , and different coin flips correspond to random variables $Y_1, Y_2, Y_3, \dots$ which are independent copies of Y . The probability of getting heads on the first flip is $P_N(1) = \Pr(\{Y_1 = \text{heads}\}) = p$. The probability of getting heads on the second flip corresponds to the event $\{Y_1 = \text{tails}\} \text{ AND } \{Y_2 = \text{heads}\}$. We assumed the coin flips are independent, so $P_N(2) = (1-p)p$. Similarly $P_N(3) = (1-p)^2p$. The general formula is $P_N(n) = (1-p)^{n-1}p$.

P8.4 Starting from the definition, we write $\mathbb{E}_X[X] = \sum_{x=0}^{\infty} x(1-p)^{x-1}p$. Using the formula from the hint to compute the infinite summation, we obtain the answer $\mathbb{E}_X[X] = p \sum_{x=0}^{\infty} x(1-p)^{x-1} = p \frac{1}{(1-(1-p))^2} = \frac{p}{p^2} = \frac{1}{p}$.

P8.5 We'll model each spin of the roulette wheel as a random variable X with sample space $\{\text{red, black, green}\}$ and probability distribution $p_X = (\frac{50}{101}, \frac{50}{101}, \frac{1}{101})$. When placing a bet on black, the payout is $W = \$2$ if the outcome is black, and zero for other outcomes. The expected value of the payout is $\mathbb{E}_X[W] = 0 \cdot \frac{50}{101} + \$2 \cdot \frac{50}{101} + 0 \cdot \frac{1}{101} = \$ \frac{100}{101}$. Since $\frac{100}{101} < 1$, the house has an advantage, so the mathematician shouldn't play.

P8.6 The payout function for this game is defined as follows:

$$f(\square) = f(\blacksquare) = f(\boxdot) = \$0, \quad f(\boxminus) = f(\boxplus) = \$1, \quad f(\boxtimes) = \$5.$$

The die is described by the distribution $p_X = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6})^T$. The expected payout is $\mathbb{E}_X[f(X)] = \sum_x f(x) p_X(x) = \frac{\$1 + \$1 + \$5}{6} = \frac{\$7}{6} = \$1.1\bar{6}$. The expected payout is greater than the cost to play, so you'll win on average.

P8.7 Starting from $\text{var}(X) = \sum_x (x - \mu_X)^2 p_X(x)$ and expanding the bracket, we obtain $\sum_x (x^2 p_X(x) - 2x\mu_X p_X(x) + \mu_X^2 p_X(x))$. Since $\sum_x x p_X(x) = \mu_X$ and $\sum_x p_X(x) = 1$, this variance expression simplifies to $\sum_x x^2 p_X(x) - \mu_X^2$. An analogous equations relating the moments of inertia of solids exists in physics: $I_{\text{cm}} = I - md^2$. This is known as the parallel axis theorem and states that the moment of inertia of a solid around its centre of mass is equal to its moment of inertia around a different parallel axis, minus a factor md^2 proportional to the mass of the object and the squared distance of the axis to the centre of mass.

P8.8 Define X_i to be the probability distribution of the weather in Year i . The transition matrix for the weather Markov chain is $M = \begin{bmatrix} \frac{1}{2} & \frac{1}{4} \\ \frac{1}{2} & \frac{3}{4} \end{bmatrix}$. We obtain the weather in Year 1 using $p_{X_1} = M(1,0)^T = (\frac{1}{2}, \frac{1}{2})^T$. The weather in Year 2 is $p_{X_2} = M^2(1,0)^T = (\frac{3}{8}, \frac{5}{8})^T$. The long term (stationary) distribution is $p_{X_\infty} = M^\infty(1,0)^T = (\frac{1}{3}, \frac{2}{3})^T$.

P8.9 Construct M_1 by counting the outbound links for each webpage, then mix in 0.9 of it with 0.1 of $\frac{1}{5}\mathbb{J}$ to obtain the Markov chain matrix M :

$$M_1 = \begin{bmatrix} 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ \frac{1}{4} & 0 & \frac{1}{2} & 0 & 0 \\ \frac{1}{4} & \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ \frac{1}{4} & 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{4} & 0 & 0 & \frac{1}{2} & 0 \end{bmatrix}, \quad M_2 = \frac{1}{5}\mathbb{J}_5, \quad M = \begin{bmatrix} \frac{1}{50} & \frac{47}{100} & \frac{1}{50} & \frac{1}{50} & \frac{47}{100} \\ \frac{49}{200} & \frac{1}{50} & \frac{47}{100} & \frac{1}{50} & \frac{1}{50} \\ \frac{49}{200} & \frac{47}{100} & \frac{1}{50} & \frac{47}{100} & \frac{1}{50} \\ \frac{49}{200} & \frac{1}{50} & \frac{47}{100} & \frac{1}{50} & \frac{47}{100} \\ \frac{49}{200} & \frac{1}{50} & \frac{1}{50} & \frac{47}{100} & \frac{1}{50} \end{bmatrix}.$$

Solving $(M - \mathbb{1})\vec{e} = \vec{0}$, we find $p_{X_\infty} = \frac{\vec{e}}{\|\vec{e}\|_1} = (0.1721, 0.169, 0.245, 0.245, 0.169)^T$.

Chapter 9 solutions

Answers to exercises

E9.1 $\vec{u} = (\alpha, 0, 0, \beta)^T$ (a column vector); $\vec{v}^T = (0, \bar{a}, \bar{b}, 0)$ (a row vector). **E9.2**

$|w\rangle = |10\rangle + i|1\rangle - |2\rangle$; $\langle w| = 1\langle 0| - i\langle 1| - \langle 2|$; $\|\vec{w}\| = \sqrt{3}$. **E9.3** $\det(A) = -2$. **E9.4** a) $(1 + \frac{1}{\sqrt{2}})|0\rangle + \frac{i}{\sqrt{2}}|1\rangle$; b) 2; c) $\frac{\sqrt{2}}{2}$; d) $\frac{\sqrt{10}}{10} + \frac{\sqrt{10}i}{5}$; e) $\frac{\sqrt{10}}{10} - \frac{\sqrt{10}i}{5}$;

f) $\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$; g) $\frac{\sqrt{2}}{2}|0\rangle$; h) $\begin{bmatrix} \frac{1}{2} & -\frac{i}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}$; i) $\frac{1}{2}|0\rangle + \frac{i}{2}|1\rangle$; j) $(\frac{\sqrt{5}}{10} + \frac{\sqrt{5}i}{10})|0\rangle - (\frac{\sqrt{5}}{5} - \frac{\sqrt{5}i}{10})|1\rangle$.

E9.5 a) $A|v\rangle = 10|0\rangle + 19|1\rangle$; b) $\langle v|A = 13\langle 0| + 18\langle 1|$; c) $A|w\rangle = (-3 - 6i)|0\rangle - (12 + 10i)|1\rangle$; d) $\langle v|A|v\rangle = 67$; e) $\langle w|A = (-3 + 8i)\langle 0| + (-9 + 10i)\langle 1|$; f) $\langle w|A|w\rangle = 29 - 6i$. **E9.6** $M_T = 2|0\rangle\langle 0| + 1|0\rangle\langle 1| - 3|1\rangle\langle 1|$. **E9.7** $HH|0\rangle = |0\rangle$ and $HH|1\rangle = |1\rangle$. **E9.8** $XX = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, $XZ = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$, $ZX = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$. **E9.9** See solution.

E9.10 $M_{\Pi_+} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}$ and $M_{\Pi_-} = \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{bmatrix}$. **E9.11** $\Pr(\{-\}|\psi) = \frac{|\alpha - \beta|^2}{2}$. **E9.12**

$|\theta^\perp\rangle = (\frac{1}{\sqrt{2}}, \frac{e^{-i(\pi-\theta)}}{\sqrt{2}})^T$. $M_{\Pi_\theta} = \begin{bmatrix} \frac{1}{2} & \frac{e^{-i\theta}}{2} \\ \frac{e^{i\theta}}{2} & \frac{1}{2} \end{bmatrix}$ and $M_{\Pi_{\theta^\perp}} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2}e^{i(\pi-\theta)} \\ \frac{1}{2}e^{-i(\pi-\theta)} & \frac{1}{2} \end{bmatrix}$.

$\Pr(\{-\}|\psi) = \sqrt{\left|\frac{\alpha}{2} + \frac{\beta}{2}e^{-i\theta}\right|^2 + \left|\frac{\alpha}{2}e^{i\theta} + \frac{\beta}{2}\right|^2}$. **E9.13** $|\psi\rangle \otimes |\phi\rangle = \alpha\gamma|00\rangle + \alpha\delta|01\rangle + \beta\gamma|10\rangle + \beta\delta|11\rangle$.

Solutions to selected exercises

E9.7 Let's first see what happens to $|0\rangle$ when we apply the operator HH . The result of the first H applied to $|0\rangle$ is $H|0\rangle = |+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. Applying the second H operator, we get $HH|0\rangle = H|+\rangle = \frac{1}{\sqrt{2}}(H|0\rangle + H|1\rangle)$. Applying the H operation gives $\frac{1}{\sqrt{2}}(\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) + \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle))$, which simplifies to $\frac{1}{\sqrt{2}}(\frac{2}{\sqrt{2}}|0\rangle) = |0\rangle$. So $HH|0\rangle = |0\rangle$. A similar calculation shows that $HH|1\rangle = |1\rangle$.

E9.8 Use matrix multiplication. Note XZ is not the same as ZX .

E9.9 Using the definition of $|+\rangle$ and $|-\rangle$, we obtain expressions for the vectors of the standard basis written in terms of the Hadamard basis: $|0\rangle = \frac{1}{\sqrt{2}}|+\rangle + \frac{1}{\sqrt{2}}|-\rangle$,

and $|1\rangle = \frac{1}{\sqrt{2}}|+\rangle - \frac{1}{\sqrt{2}}|-\rangle$. The remainder of the calculations require arithmetic and simplifications:

$$\begin{aligned} |01\rangle - |10\rangle &= |0\rangle \otimes |1\rangle - |1\rangle \otimes |0\rangle \\ &= \left(\frac{1}{\sqrt{2}}|+\rangle + \frac{1}{\sqrt{2}}|-\rangle\right) \left(\frac{1}{\sqrt{2}}|+\rangle - \frac{1}{\sqrt{2}}|-\rangle\right) - \left(\frac{1}{\sqrt{2}}|+\rangle - \frac{1}{\sqrt{2}}|-\rangle\right) \left(\frac{1}{\sqrt{2}}|+\rangle + \frac{1}{\sqrt{2}}|-\rangle\right) \\ &= \frac{1}{2} \left(\cancel{|++\rangle} - |+-\rangle + |-+\rangle - \cancel{|--\rangle} - \cancel{|++\rangle} + \cancel{|+-\rangle} + |-\rangle + \cancel{|-+\rangle} + \cancel{|--\rangle} \right) \\ &= \frac{1}{2} \left(-2|+-\rangle + 2|-+\rangle \right) = -(|+\rangle|-\rangle - |-\rangle|+\rangle) = |+\rangle|-\rangle - |-\rangle|+\rangle. \end{aligned}$$

The last equality holds since the state's global phase can be ignored.

Answers to problems

P9.1 No, since Q is not unitary. **P9.2** $YY|0\rangle = |0\rangle$ and $YY|1\rangle = |1\rangle$. **P9.3** $i\alpha|0\rangle - i\beta|1\rangle$. **P9.4** Two parameters. **P9.5** $\alpha \in [0, 1]$ and $\varphi \in [0, 2\pi]$. **P9.6** $\theta \in [0, \pi]$ and $\varphi \in [0, 2\pi]$. **P9.7** $HXH = Z$ and $HZH = X$. **P9.8** $\Pr(\{v\}|\psi) = |\alpha\bar{a} + \beta\bar{b}|^2$. **P9.9** $\Pr(\{0\}|\psi) = |\beta|^2$; $\Pr(\{+\}|\psi) = \frac{1}{2}|\alpha + \beta|^2$; $\Pr(\{-\}|\psi) = \frac{1}{2}|\alpha - \beta|^2$. **P9.11** $\mathbb{E}_R[R] = \frac{3}{2}a$. **P9.12** $\langle\psi_1, \psi_2\rangle = 0$. **P9.13** $\Pr(\{0 \leq x \leq \frac{1}{4}\}|\psi_a) = \frac{7}{16}$; $\Pr(\{0 \leq x \leq \frac{1}{4}\}|\psi_b) = \frac{79}{256}$; $\Pr(\{0 \leq x \leq \frac{1}{4}\}|\psi_c) = \frac{1}{4}$.

Solutions to selected problems

P9.1 Quantum gates correspond to unitary operators. Since $Q^\dagger Q \neq 1$, Q is not unitary and it cannot be implemented by any physical device. The boss is *not* always right!

P9.2 Let's first see what happens to $|0\rangle$ when we apply the operator YY . The result of the first Y applied to $|0\rangle$ is $Y|0\rangle = i|1\rangle$. Then, applying the second Y operator, we get $YY|0\rangle = Y(i|1\rangle) = iY|1\rangle = i(-i)|0\rangle = |0\rangle$. So $YY|0\rangle = |0\rangle$. A similar calculation shows that $YY|1\rangle = |1\rangle$.

P9.3 Since $HH = 1$, we obtain $XHHY = XY = \begin{bmatrix} i & 0 \\ 0 & -i \end{bmatrix}$, then multiply this operator by $\alpha|0\rangle + \beta|1\rangle$ to obtain the answer.

P9.4 Starting from the four degrees of freedom for general two-dimensional complex vectors, we must subtract one degree of freedom for each of the constraints: one because we're ignoring global phase, and one because we require $|\alpha|^2 + |\beta|^2 = 1$:

$$4 \text{ d.f.} - \alpha \text{ real} - \{\|\psi\| = 1\} = 2 \text{ d.f.}$$

A qubit $|\psi\rangle$ has only two degrees of freedom. In other words, two parameters are sufficient to describe any qubit.

P9.5 Since the squared-magnitudes of the components of $|\psi\rangle$ must be a probability distribution, α is restricted in the range $[0, 1]$. The phase of the second component can be chosen arbitrarily; $\varphi \in [0, 2\pi]$.

P9.6 These are the angles of the *Bloch sphere*, which is a useful way to visualize qubit quantum states.

P9.9 All three calculations require computing the product of the appropriate projection matrix with the density matrix $\rho = \begin{bmatrix} |\alpha|^2 & \bar{\alpha}\beta \\ \bar{\alpha}\beta & |\beta|^2 \end{bmatrix}$, then taking the trace. The density matrix representation is used in many domains of physics, as well as in quantum information theory.

P9.10 a) Expand the expressions for $|\psi\rangle_1$ and $|\Phi_+\rangle_{23}$ and compute the tensor product. **b)** Using the definitions of the Bell states $|\Phi_+\rangle$, $|\Phi_-\rangle$, $|\Psi_+\rangle$, $|\Psi_-\rangle$, and the X and Z operators, we expand the expression on the right side of the equation and show it equals the expression on the left side, which in turn equals $|\psi\rangle_1 \otimes |\Phi_+\rangle_{23}$. **c)** Applying the CNOT gate to the state $|\Phi_+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ results in $\frac{1}{\sqrt{2}}(|00\rangle + |10\rangle) = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |0\rangle$. Next, applying the Hadamard gate on the first qubit leaves us with the state $|00\rangle$, which leads to the measurement outcome 00 when both qubits are measured in the standard basis. The analysis for the other bell states is similar. **d)** The recovery operation that Bob must perform is determined by the measurement outcome. Since the X and Z gates are self-inverse, the recovery operations that Bob must perform are described by the following mapping: $00 \rightarrow 1$, $10 \rightarrow Z$, $01 \rightarrow X$, $11 \rightarrow ZX$. See this video for more details about the quantum teleportation protocol: youtu.be/3wZ35c3oYUE.

P9.11 Using integration by parts with $u = r^3$ and $dv = \exp(-2r/a)r^2 dr$ simplifies the integral to the form $\frac{3a}{2} \left(\frac{4}{a^3} \int_0^\infty \exp(-\frac{2r}{a})r^2 dr \right)$. Instead of continuing with two more steps of integration by parts, we can recognize the expression inside the brackets to be equal to one, since $p_R(r) = \frac{4}{a^3} \exp(-2r/a)r^2$ is a probability distribution. This calculation supports the classical chemical viewpoint, which describes electrons as living in an “electron cloud” or “orbital” of radius roughly $\frac{3}{2}a$. Note this is somewhat misleading since the actual wave function $\psi(\vec{r})$ drops off as $\exp(-r)$.

P9.12 The inner product $\langle\psi_1, \psi_2\rangle = \int_0^1 \overline{\psi_1(x)}\psi_2(x) dx$ corresponds to the integral $\int_0^1 (2x-1)(6x^2-6x+1) dx = \int_0^1 (12x^3-18x^2+8x-1) dx$. Using the formula $\int x^n dx = \frac{x^{n+1}}{n+1}$, we find $\langle\psi_1, \psi_2\rangle = \left[\frac{12}{4}x^4 - \frac{18}{3}x^3 + \frac{8}{2}x^2 - x \right]_0^1 = 0$. The functions $\psi_1(x)$ and $\psi_2(x)$ are called the *Shifted Legendre polynomials*.

P9.13 You can perform the required integrals by hand or use the SymPy command `integrate((sqrt(3)*(2*x-1))**2, (x, 0, 1/4))` for part **a)**. For part **b)**, use the command `integrate((sqrt(5)*(6*x**2-6*x+1))**2, (x, 0, 1/4))`. For part **c)**, the constant wave function is $\psi_c(x) = 1$, and the interval $[0, \frac{1}{4}]$ contains exactly $\frac{1}{4}$ of its probability mass.

Appendix B

Notation

This appendix contains a summary of the notation used in this book.

Math notation

Expression	Read as	Used to denote
a, b, x, y		variables
$=$	is equal to	expressions that have the same value
$\stackrel{\text{def}}{=}$	is defined as	new variable definitions
$a + b$	a plus b	the combined lengths of a and b
$a - b$	a minus b	the difference in lengths between a and b
$a \times b = ab$	a times b	the area of a rectangle
$a^2 \stackrel{\text{def}}{=} aa$	a squared	the area of a square of side length a
$a^3 \stackrel{\text{def}}{=} aaa$	a cubed	the volume of a cube of side length a
a^n	a exponent n	a multiplied by itself n times
$\sqrt{a} \stackrel{\text{def}}{=} a^{\frac{1}{2}}$	square root of a	the side length of a square of area a
$\sqrt[3]{a} \stackrel{\text{def}}{=} a^{\frac{1}{3}}$	cube root of a	the side length of a cube with volume a
$a/b = \frac{a}{b}$	a divided by b	a parts of a whole split into b parts
$a^{-1} \stackrel{\text{def}}{=} \frac{1}{a}$	one over a	division by a
$f(x)$	f of x	the function f applied to input x
f^{-1}	f inverse	the inverse function of $f(x)$
$f \circ g$	f compose g	function composition; $f \circ g(x) \stackrel{\text{def}}{=} f(g(x))$
e^x	e to the x	the exponential function base e
$\ln(x)$	natural log of x	the logarithm base e
a^x	a to the x	the exponential function base a
$\log_a(x)$	log base a of x	the logarithm base a
θ, ϕ	<i>theta, phi</i>	angles
sin, cos, tan	sin, cos, tan	trigonometric ratios
%	percent	proportions of a total; $a\% \stackrel{\text{def}}{=} \frac{a}{100}$

Set notation

You don't need a lot of fancy notation to do math, but it really helps if you know a little bit of set notation.

Symbol	Read as	Denotes
$\{ \dots \}$	the set ...	definition of a set
$ $	such that	describe or restrict the elements of a set
$\mathbb{N}$	the naturals	the set $\mathbb{N} \stackrel{\text{def}}{=} \{0, 1, 2, \dots\}$. Also $\mathbb{N}_+ \stackrel{\text{def}}{=} \mathbb{N} \setminus \{0\}$.
$\mathbb{Z}$	the integers	the set $\mathbb{Z} \stackrel{\text{def}}{=} \{\dots, -2, -1, 0, 1, 2, 3, \dots\}$
$\mathbb{Q}$	the rationals	the set of fractions of integers
$\mathbb{R}$	the reals	the set of real numbers
$\mathbb{C}$		the set of complex numbers
$\mathbb{F}_q$	finite field	the set $\{0, 1, 2, 3, \dots, q-1\}$
$\subset$	subset	one set strictly contained in another
$\subseteq$	subset or equal	containment or equality
$\cup$	union	the combined elements from two sets
$\cap$	intersection	the elements two sets have in common
$S \setminus T$	S set minus T	the elements of S that are not in T
$a \in S$	a in S	a is an element of set S
$a \notin S$	a not in S	a is not an element of set S
$\forall x$	for all x	a statement that holds for all x
$\exists x$	there exists x	an existence statement
$\nexists x$	there doesn't exist x	a non-existence statement
$S \times T$	Cartesian product	all pairs (s, t) where $s \in S$ and $t \in T$

An example of a condensed math statement that uses set notation is " $\nexists m, n \in \mathbb{Z}$ such that $\frac{m}{n} = \sqrt{2}$," which reads "there don't exist integers m and n whose fraction equals $\sqrt{2}$." Since we identify the set of fractions of integers with the rationals, this statement is equivalent to the shorter " $\sqrt{2} \notin \mathbb{Q}$," which reads " $\sqrt{2}$ is irrational."

Vectors notation

Expression	Denotes
$\mathbb{R}^n$	the set of n -dimensional real vectors
$\vec{v}$	a vector
(v_x, v_y)	vector in component notation
$v_x \hat{i} + v_y \hat{j}$	vector in unit vector notation
$\ \vec{v}\ \angle \theta$	vector in length-and-direction notation
$\ \vec{v}\ $	length of the vector $\vec{v}$
θ	angle the vector $\vec{v}$ makes with the x -axis
$\hat{v} \stackrel{\text{def}}{=} \frac{\vec{v}}{\ \vec{v}\ }$	unit vector in the same direction as $\vec{v}$
$\vec{u} \cdot \vec{v}$	dot product of the vectors $\vec{u}$ and $\vec{v}$
$\vec{u} \times \vec{v}$	cross product of the vectors $\vec{u}$ and $\vec{v}$

Complex numbers notation

Expression	Denotes
$\mathbb{C}$	the set of complex numbers $\mathbb{C} \stackrel{\text{def}}{=} \{a + bi \mid a, b \in \mathbb{R}\}$
i	the unit imaginary number $i \stackrel{\text{def}}{=} \sqrt{-1}$ and $i^2 = -1$
$\operatorname{Re}\{z\} = a$	real part of $z = a + bi$
$\operatorname{Im}\{z\} = b$	imaginary part of $z = a + bi$
$ z \angle \varphi_z$	polar representation of $z = z \cos \varphi_z + i z \sin \varphi_z$
$ z = \sqrt{a^2 + b^2}$	magnitude of $z = a + bi$
$\varphi_z = \tan^{-1}(b/a)$	phase or argument of $z = a + bi$
$\bar{z} = a - bi$	complex conjugate of $z = a + bi$
$\mathbb{C}^n$	the set of n -dimensional complex vectors

Vector space notation

Expression	Denotes
U, V, W	vector spaces
$W \subseteq V$	vector space W subspace of vector space V
$\{\vec{v} \in V \mid \langle \text{cond} \rangle\}$	subspace of vectors in V satisfying condition $\langle \text{cond} \rangle$
$\operatorname{span}(\vec{v}_1, \dots, \vec{v}_n)$	span of vectors $\vec{v}_1, \dots, \vec{v}_n$
$\dim(U)$	dimension of vector space U
$\mathcal{R}(M)$	row space of M
$\mathcal{N}(M)$	null space of M
$\mathcal{C}(M)$	column space of M
$\mathcal{N}(M^T)$	left null space of M
$\operatorname{rank}(M)$	rank of M ; $\operatorname{rank}(M) \stackrel{\text{def}}{=} \dim(\mathcal{R}(M)) = \dim(\mathcal{C}(M))$
$\operatorname{nullity}(M)$	nullity of M ; $\operatorname{nullity}(M) \stackrel{\text{def}}{=} \dim(\mathcal{N}(M))$
B_s	the standard basis
$\{\vec{e}_1, \dots, \vec{e}_n\}$	an orthogonal basis
$\{\hat{e}_1, \dots, \hat{e}_n\}$	an orthonormal basis
$B'[\mathbb{I}]_B$	the change-of-basis matrix from basis B to basis B'
Π_S	projection onto subspace S
$\Pi_{S^\perp}$	projection onto the orthogonal complement of S

Abstract vector spaces notation

Expression	Denotes
$(V, F, +, \cdot)$	abstract vector space of vectors from the set V , whose coefficients are from the field F , addition operation “+” and scalar-multiplication operation “.”
$\mathbf{u}, \mathbf{v}, \mathbf{w}$	abstract vectors
$\langle \mathbf{u}, \mathbf{v} \rangle$	inner product of vectors $\mathbf{u}$ and $\mathbf{v}$
$\ \mathbf{u}\ $	norm of $\mathbf{u}$
$d(\mathbf{u}, \mathbf{v})$	distance between $\mathbf{u}$ and $\mathbf{v}$

Notation for matrices and matrix operations

Expression	Denotes
$\mathbb{R}^{m \times n}$	the set of $m \times n$ matrices with real entries
A	a matrix
a_{ij}	entry in the i^{th} row and j^{th} column of A
$ A $	determinant of A , also denoted $\det(A)$
A^{-1}	matrix inverse
$A^{\top}$	matrix transpose
$\mathbb{1}$	identity matrix; $\mathbb{1}A = A\mathbb{1} = A$ and $\mathbb{1}\vec{v} = \vec{v}$
AB	matrix-matrix product
$A\vec{v}$	matrix-vector product
$\vec{w}^{\top}A$	vector-matrix product
$\vec{u}^{\top}\vec{v}$	vector-vector inner product; $\vec{u}^{\top}\vec{v} \stackrel{\text{def}}{=} \vec{u} \cdot \vec{v}$
$\vec{u}\vec{v}^{\top}$	vector-vector outer product
$\text{ref}(A)$	row echelon form of A
$\text{rref}(A)$	reduced row echelon form of A
$\text{rank}(A)$	rank of $A \stackrel{\text{def}}{=} \text{number of pivots in } \text{rref}(A)$
$A \sim A'$	matrix A' obtained from matrix A by row operations
$\mathcal{R}_1, \mathcal{R}_2, \dots$	row operations, of which there are three types: $\rightarrow R_i \leftarrow R_i + kR_j$: add k-times row j to row i $\rightarrow R_i \leftrightarrow R_j$: swap rows i and j $\rightarrow R_i \leftarrow mR_i$: multiply row i by constant m
$E_{\mathcal{R}}$	elementary matrix for row operation $\mathcal{R}$; $\mathcal{R}(M) \stackrel{\text{def}}{=} E_{\mathcal{R}}M$
$[A \mid \vec{b}]$	augmented matrix containing matrix A and vector $\vec{b}$
$[A \mid B]$	augmented matrix array containing matrices A and B
M_{ij}	minor associated with entry a_{ij} . See page 196.
$\text{adj}(A)$	adjugate matrix of A . See page 198.
$(A^{\top}A)^{-1}A^{\top}$	Moore–Penrose inverse of A . See page 379.
$\mathbb{C}^{m \times n}$	the set of $m \times n$ matrices with complex entries
$A^{\dagger}$	Hermitian transpose; $A^{\dagger} \stackrel{\text{def}}{=} (\bar{A})^{\top}$

Notation for linear transformations

Expression	Denotes
$T : \mathbb{R}^n \rightarrow \mathbb{R}^m$	linear transformation T from the input space $\mathbb{R}^n$ to the output space $\mathbb{R}^m$
$M_T \in \mathbb{R}^{m \times n}$	matrix representation of T
$\text{Im}(T) = \mathcal{C}(M_T)$	the image space of T = column space of M_T
$\text{Ker}(T) = \mathcal{N}(M_T)$	the kernel of T = null space of M_T
$S \circ T(\vec{x})$	composition of linear transformations; $S \circ T(\vec{x}) \stackrel{\text{def}}{=} S(T(\vec{x})) = M_S M_T \vec{x}$
$M \in \mathbb{R}^{m \times n}$	an $m \times n$ matrix
$T_M : \mathbb{R}^n \rightarrow \mathbb{R}^m$	the linear transformation defined as $T_M(\vec{v}) \stackrel{\text{def}}{=} M\vec{v}$
$T_{M^T} : \mathbb{R}^m \rightarrow \mathbb{R}^n$	the adjoint linear transformation $T_{M^T}(\vec{a}) \stackrel{\text{def}}{=} \vec{a}^T M$

Matrix decompositions

Expression	Denotes
$A \in \mathbb{R}^{n \times n}$	a matrix (assume diagonalizable)
$p_A(\lambda) \stackrel{\text{def}}{=} A - \lambda \mathbb{I} $	characteristic polynomial of A
$\lambda_1, \dots, \lambda_n$	eigenvalues of A = roots of $p_A(\lambda)$
$\Lambda \in \mathbb{R}^{n \times n}$	diagonal matrix of eigenvalues of A
$\vec{e}_{\lambda_1}, \dots, \vec{e}_{\lambda_n}$	eigenvectors of A
$Q \in \mathbb{R}^{n \times n}$	matrix whose columns are eigenvectors of A
$A = Q\Lambda Q^{-1}$	eigendecomposition of A
$A = O\Lambda O^T$	eigendecomposition of a normal matrix
$B \in \mathbb{R}^{m \times n}$	a generic matrix
$\sigma_1, \sigma_2, \dots$	singular values of B
$\Sigma \in \mathbb{R}^{m \times n}$	matrix of singular values of B
$\vec{u}_1, \dots, \vec{u}_m$	left singular vectors of B
$U \in \mathbb{R}^{m \times m}$	matrix of left singular vectors of B
$\vec{v}_1, \dots, \vec{v}_n$	right singular vectors of B
$V \in \mathbb{R}^{n \times n}$	matrix of right singular vectors of B
$B = U\Sigma V^T$	singular value decomposition of B

Index

The page numbers shown in **bold** point to the concept definitions.

- absolute value, 11, **55**, 98
- abstract vector space, **318**, 348
- adjacency matrix, **373**, 460
- adjoint, **238**, 340, 354
- adjugate matrix, **196**, 208
- affine transformation, 51, 152, 385, **389**, 393
- algebraic multiplicity, **301**, 308
- anticommutative, 141, 192
- arc length, 68
- area, 14, 66, 67, 150, 188
- associative, 14, **24**, 181, 319
- augmented matrix, **165**, 214, 371
- axis, 35, 98

- basis, **94**, **226**, 241, 244, 252, 282
 - orthogonal, **226**, 304, 326
 - orthonormal, **226**, 329, 424
 - standard, 92, 225, 230, 263, 270, 300, 320, 479, 482
- Bayes' rule, 449
- Bayesian, 448
- best fit solution, 382, 384
- bijjective, **41**, 43, **262**, 293
- binary, **404**
- Born's rule, **495**
- bra, **480**, *see also* ket

- Cartesian plane, **35**, 38, 85
- Cauchy–Schwarz inequality, 325
- change of basis, 95, **229**, 230, 304, 421, 424, 482
 - for matrices, 282, 485
- characteristic polynomial, 194, **301**, 304, 345, 350
- Cholesky decomposition, **336**
- ciphertext, **403**, 407
- circle, 67
- closed under
 - addition, **233**, 234
 - scalar multiplication, **233**
- codomain, *see* output space
- coefficients, 134, 165, 226, 321
 - Fourier, 425, **427**
- cofactor, **197**, 208
- column picture, 133, 137, **146**, 290
- column space, 136, **232**, 236, 239, 246, 259, 264, 384
- column vector, 134, **147**, 238, 341
- commutative, 14, **24**, 141, 319
- completing the square, **28**, 31
- complex conjugate, **97**, 340, 480
- complex number, 13, **96**, 112
- complex transpose, *see* Hermitian transpose
- components, 76, 82, 95, **131**, 225, 323, *see also* coordinates
 - complex, 339, 480
- cone, **70**
- conjugate, **97**, 340, 480
- conjugate transpose, *see* Hermitian transpose
- conjugate-linear, **352**
- coordinate system, *see also* basis
 - Cartesian, **35**, 82, 92
 - polar, 97
- coordinates, 2, 83, 93, **131**, 210,

- 225, 245, 264, 329, 481,
see also components
- homogeneous, **390**
- cosine, 20, **62**, 72
- cosine rule, **67**
- Cramer's rule, **193**
- cross product, 83, **141**, 192, 215
- cryptography, **402**
- cylinder, **69**
- dagger, *see* Hermitian transpose
- De Moivre's formula, 102
- decryption, **403**, 406, 409
- degenerate eigenvalue, **301**, 304
- determinant, 150, 164, **185**, 188,
196, 286, 288, 301, 306
- determinant test, 191, **194**, 288
- diagonalization, **300**, 303, 308
- digital signature, **409**
- dimension, 172, 232, 239, 241,
245, 251, 265, 304, 318
- Dirac notation, **479**, 483
- direction vector, 170, **211**, 273
- distance, 216, 223, 325, 397
 - Hamming, **413**, 416
- distributive, **24**
- domain, **39**, 123, 261, 292, 421
- dot product, 83, **140**, 146, 181,
214, 219, 225, 323, 327
- eigenbasis, **298**, 302, 304, 306, 333
- eigendecomposition, 298, **301**,
305, 307, 310, 332, 376
- eigenspace, 280, 302, **304**, 307,
316, 351, 458
- eigenvalue, 190, 194, **298**, 306,
334, 339, 345, 375, 454
- eigenvector, **298**, 301, 307, 333,
339, 350, 376, 454
- encryption, **403**, 406, 409
- entries, *see* matrix entries
- error-correcting code, **412**, 528
- Euler's formula, **102**, 345, 350
- expand, **24**, **25**
- expected value, **445**
- exponent, 15
- exponential, 20, **64**
- factor, 15, **23**, 26
- factoring, 26, 34
- field, 318, **322**, 323, 339, 353, 354
- finite field, **353**, 404, 417, 468
- Fourier basis, 421, 425, 427, 431
- Fourier coefficients, 425, **427**
- Fourier series, **426**
- Fourier transformation, **421**, 429
- fraction, 13, 112
- free variable, **170**, 214, 238, 248
- Frobenius inner prod., *see* inner
prod. for matrices
- Frobenius norm, *see*
Hilbert–Schmidt norm
- function, **38**, 50
 - even, **59**
 - odd, **59**
- fundamental subspace, **232**, 236,
240, 261
 - basis, **246**
- Gauss–Jordan elimination, **168**,
174, 194, 199, 214, 238,
243, 343, 365, 369
- geometric multiplicity, **301**, 308
- golden ratio, **33**
- Gram–Schmidt
 - orthogonalization, **328**, 337
- graph, **373**, 458
- graphics processing unit, 398
- Grover's search algorithm, **525**
- Hadamard basis, **482**, 485, 498
- Hadamard gate, **494**, 522
- Hamming code, **416**, 419
- Hamming distance, **413**, 416
- Hermitian matrix, **346**
- Hermitian transpose, **340**, 480
- Hilbert–Schmidt norm, **326**, 349
- homogeneous coordinates, **390**
- homogeneous equations, 236
- idempotence, **274**
- identity matrix, 135, 149, 178,
198, **313**, 343, 393, 495
- image, **39**, 123, 136, 158, 258
- image space, 136, 158, **258**, 261,
264
- imaginary number, *see* complex
number

- infinity, 170, 209, 321, 425, 448
- information theory, 503, **529**
- injective, **41**, **262**, 293
- inner product, 143, **148**, 272, *see also* dot product
 - abstract, **322**, 329
 - complex, 341, **347**, 352, 481
 - for functions, 326, 352, 424
 - for matrices, 325, 348
- inner product space, **322**, 348, 490
- input set, *see* source set
- input space, **261**, 282, 293
- intersection, 172, 213, 390
- interval, **110**, 121
- inverse, **11**
 - function, 14, 20, 43, 136
 - matrix, 134, 136, 149, 196, 202, 287, 305, 308, 343
 - Moore–Penrose, **382**, 388
 - transformation, 4, **267**, 292
- invertibility test, **196**, 203
- invertible, 149, 185, 196, 203, 287
 - matrix, 185, 267, 287
 - transformation, **267**, 292
- isolate, 11, 20, 103

- kernel, 4, 136, 237, **258**, 262, 265, 292
- ket, **480**, 483, *see also* bra
- Kirchhoff’s laws, 369

- ℓ^1 -norm, 455
- ℓ^2 -norm, 143
- leading one, *see* pivot
- least squares approximate
 - solution, **379**
- left null space, **232**, 237, 239, 384
- length, 13, 72, 84, 98
- length preserving, *see* orthogonal matrix
- Leontief input-output model, 366
- line, 154, 173, **209**, 217, 221, 394
 - parametric equation, **211**
 - symmetric equation, **211**
- linear, 1, 56, **152**, 257, 323, 380
 - combination, 1, 134, 227, 231, 241, 320, 421, 511
 - equation, 103, 118, **153**, 164
 - independence, 194, 227, **240**, 243, 288, 301, 328
 - transformation, 2, 134, 183, 238, **257**, 269, 282, 312
- linear programming, **378**
- linear regression, **379**
- logarithm, 20, **65**
- low-rank approximation, **335**
- LU decomposition, **335**

- Markov chain, **450**, 457
- matrix, **131**, 136, 145, 263, 269
 - diagonal, 179, 302, **312**, 333
 - elementary, 184, **200**, 291, *see also* row operations
 - entries, **131**, 134, 145, 165, 185, 196, 266, **283**, 484
 - Hermitian, **346**
 - identity, *see* identity matrix
 - invariant properties, 286
 - lower triangular, **313**, 336
 - multiplication, 146, **147**, 181, 263, 389, 414, 452, 492
 - normal, 316
 - operations, 133, 145
 - orthogonal, **313**, 333, 346
 - permutation, **315**, 336
 - positive definite, 316
 - positive semidefinite, **316**, 327, 336
 - product, *see* matrix multiplication
 - projection, 148, 263, **270**, 316
 - rank, 167, 232, **240**, 286, 288
 - reflection, **274**, 280, 315, 392
 - representation, 136, **263**, 269, 282, 392, 446, 484
 - rotation, **277**, 314, 339, 392
 - symmetric, **313**, 320, 325, *see also* Hermitian matrix
 - trace, **149**, 286, 306
 - transpose, **147**, 238, 246, 382, *see also* Hermitian tr.
 - unitary, **346**, 492
 - upper triangular, **313**, 336
- matrix-vector product, **133**, 146, 177, 183, 238, 269, 381, 389, 425, 447

- metric, 327
- minor, **187**, 197
- nilpotent, **361**
- non-homogeneous, 236
- noncommutative, **178**, 182, 183
- nonnegative, 39, 55, **110**, 123
- norm, **143**, **324**, 348, 467, 496
 - Euclidian, **143**
 - Hilbert–Schmidt, **326**, 349
 - ℓ^1 -norm, 455
 - ℓ^2 -norm, 143
- normal vector, 210, **214**, 277, 401
- null space, 136, **232**, 235, 237, 239, 242, 247, 259, 265, 274, 288, 292, 301, 415
- nullity, **240**, 242, 293
- number line, 35
- objective function, 378, 380
- one-time pad, **406**
- one-to-one, *see* injective
- one-to-one correspondence, *see* bijective
- onto, *see* surjective
- origin, **36**, 55, 73
- orthogonal, 214, 227, 247, 265, **323**, 328, 384, 415, 421
 - basis, *see* basis
 - complement, **219**, 221, 274, 384, 415
 - matrix, **313**, 333, 337, 346
 - projection, 228, **270**, 392
 - subspace, **220**, 331, 415
 - sum, **239**, 265, 331
- orthogonalization, **328**, 330, 337
- orthonormal basis, *see* basis
- outer product, **148**, 272, 483, 484
- output set, *see* target set
- output space, **261**, 282, 293
- PageRank, **456**, 459, 463
- parabola, **53**
- parallelepiped, 185, 188
- parallelogram, 150, 185, 188
- parity check, **420**
- particular solution, 236, 237
- permutation, **315**, 336
- perp, *see* orthogonal complement
- perspective projection, 389, **394**
- pivot, **167**, 168, 174, 246, 251
- plaintext, **403**, 407
- plane, 154, 171, 173, **210**, 214, 217, 222, 390, 401
 - general equation, **214**, 401
 - geometric equation, **214**
 - parametric equation, **215**
- point, 173, **209**, 216, 384, 389
- polar coordinates, 97, 207
- polarizing lens, **473**, 506
- polynomial, **56**, 101
- probability distribution, **441**, 496
 - conditional, **445**, 451
- projection, 148, 219, 263, **270**, 495
- pyramid, **70**
- QR decomposition, **337**, 338
- quadratic, 20, 26, 30, **53**, 58, 381
 - formula, 31, 57
- quantum, **465**, 506
 - bit, *see* qubit
 - cryptography, **519**
 - entanglement, **501**, 522, 527
 - information theory, 504, **529**
 - measurement, **495**
 - mechanics, **465**, 489, 510
 - operations, **492**
 - state, 467, **490**, 499, 523
 - superposition, **510**, 511
 - teleportation, **526**
 - wave function, 467, 469, **513**
- qubit, 468, 488, **491**
- radian, **68**, 72, 77
- random variable, **440**
- range, *see* image
- rank, 167, 192, 232, **240**, 286, 288
- rank–nullity theorem, 240, **242**
- rational, 13, 112
- reduced row echelon form, 164, **167**, 171, 174, 194, 200, 238, 244, 289
- reflection, **274**, 280, 315, 392
- relation, **49**, 52, *see also* function
- repetition code, **416**
- representation, 136, 165, 184, **263**, 269, 392, 421, 446, 484

- Rivest–Shamir–Adleman (RSA) cryptosystem, 409
- roots, *see* solution set
- rotation, 277, 314, 339, 392
- row echelon form, 167, 336
- row operations, 166, 174, 184, 190, 198, 238, 244, 336
- row picture, 133, 137, 146
- row space, 232, 237, 239, 242, 244, 246, 259, 265, 414
- row vector, 147, 237, 341, 399, 415
- RREF, *see* reduced row echelon form
- RSA cryptosystem, 409
- sample space, 440
- scalar product, *see* dot product
- scale factor, 185, 207, 289
- set, 12, 38, 97, 109
 - difference, 109, 331
 - intersection, 109, 119
 - subset, 12, 109
 - union, 109, 110
- Shor’s factoring algorithm, 525
- similarity transformation, 286
- simplex algorithm, 378
- sine, 20, 60, 72
- sine rule, 67
- singular value decomposition, 333, 338, 349
- solution set, 11, 20, 31, 56, 118, 136, 172, 235, 237
- solution space, 170, 173
- source set, 39
- source space, *see* input space
- span, 227, 232, 233, 235, 288, 331
- sphere, 69
- spherical coordinates, 207, 514
- squared error, 380, 387
- stationary distribution, 454, 458
- subset, 235, 441
- substitution, 18, 59, 104
- surjective, 41, 262, 293
- SVD, 333, 338, 349
- swap (row op.), 166, 175, 190
- tangent, 63, 72
- target set, 39, 261, 292
- target space, *see* output space
- Taylor series, 154, 309
- term, 23, 56, 57
- trace, 149, 286, 306
- transition matrix, 451, 457
- transpose, 147, 238, 246, 340, 382
- trigonometric identities, 78
- undo, *see* inverse
- unit circle, 68, 73
- unit vector, 87, 227, 272, 468
- vector, 36, 80
 - direction, 170, 211, 273
 - linearly independent, *see* linear independence
 - normal, 210, 214, 277, 401
 - operations, 132, 138
 - unit, *see* unit vector
- vector product, *see* cross product
- vector space, 225, 231, 243, 328
 - abstract, 318, 348
 - inner product, 322, 348, 490
- vector subspace, 219, 231, 234
- view frustum, 400
- volume, 69, 188, 514
- wave function, 467, 469, 513
- XOR, 404

Linear algebra is the foundation of science and engineering. Knowledge of linear algebra is a prerequisite for studying statistics, machine learning, computer graphics, signal processing, chemistry, economics, quantum mechanics, and countless other applications. Indeed, linear algebra offers a powerful toolbox for modelling the real world.

The NO BULLSHIT GUIDE TO LINEAR ALGEBRA shows the connections between the computational techniques of linear algebra, their geometric interpretations, and the theoretical foundations. This university-level textbook contains lessons on linear algebra written in a style that is precise and concise. Each concept is illustrated through definitions, formulas, diagrams, explanations, and examples of real-world applications. Readers build their math superpowers by solving practice problems and learning to use the computer algebra system SymPy to speed up tedious matrix arithmetic tasks.

“The book explains the concepts in a way that gives a strong intuitive understanding.” — Joe Nestor, student

“It’s very well written and a fun read!” — Felix Kwok, professor

“I used this book in multiple big data courses when I needed a deeper understanding of the material.” — Zane Zakraisek, student

The author, Ivan Savoy, combines 15 years of tutoring experience with a B.Eng. in electrical engineering, an M.Sc. in physics, and a Ph.D. in computer science from McGill University.

Contents:

HIGH SCHOOL MATH REVIEW

VECTORS AND MATRICES

LINEAR TRANSFORMATIONS

LINEAR ALGEBRA THEORY

LINEAR ALGEBRA APPLICATIONS