

MATH 3260: LINEAR ALGEBRA I

Daily Log for Lectures and Readings

Timothy E. Faver

April 23, 2025

CONTENTS

How to Use This Daily Log	3
Day 1: Monday, January 6	4
Day 2: Wednesday, January 8	8
Day 3: Friday, January 10	14
Day 4: Monday, January 13	15
Day 5: Wednesday, January 15	19
Day 6: Friday, January 17	25
Day 7: Wednesday, January 22	29
Day 8: Friday, January 24	34
Day 9: Monday, January 27	37
Day 10: Wednesday, January 29	43
Day 11: Friday, January 31	50
Day 12: Monday, February 3	53
Day 13: Wednesday, February 5	59
Day 14: Friday, February 7	64
Day 15: Monday, February 10	68
Day 16: Wednesday, February 12	72
Day 17: Friday, February 14	74
Day 18: Monday, February 17	75
Day 19: Wednesday, February 19	80
Day 20: Friday, February 21	87
Day 21: Monday, February 24	91
Day 22: Wednesday, February 26	96
Day 23: Friday, February 28	101
Day 24: Monday, March 3	105
Day 25: Wednesday, March 5	109
Day 26: Friday, March 7	113
Day 27: Monday, March 17	118
Day 28: Wednesday, March 19	124
Day 29: Friday, March 21	129
Day 30: Monday, March 24	132
Day 31: Wednesday, March 26	136
Day 32: Friday, March 28	139
Day 33: Monday, March 31	139
Day 34: Wednesday, April 2	145
Day 35: Friday, April 4	149
Day 36: Monday, April 7	153
Day 37: Wednesday, April 9	159
Day 38: Friday, April 11	163
Day 39: Monday, April 14	166
Day 40: Wednesday, April 16	170
Day 41: Friday, April 18	175
Day 42: Monday, April 21	179

Day 43: Wednesday, April 23	185
------------------------------------	------------

How to Use This Daily Log

This document is our primary reference for the course. It contains all of the material that we discuss in class along with some supplementary remarks that may not be mentioned in a class meeting. Each individual day has references, when applicable, to relevant material from the text *Introduction to Linear Algebra (Sixth Edition)* by Gilbert Strang. These references are spread throughout a day's notes, and you should be consulting both the daily log and Strang's text more or less simultaneously.

This log contains several classes of problems.

(!) Problems marked (!) are meant to be attempted *immediately*. They will directly address or reinforce something that we covered (or perhaps omitted) in class. It will be to your great benefit to pause and work (!)-problems as you encounter them.

(★) Problems marked (★) are intentionally more challenging and deeper than (!)-problems. The (★)-problems will summarize and generalize ideas that we have discussed in class and give you broader, possibly more abstract perspectives. You should attempt the (★)-problems on a second rereading of the lecture notes, after you have completed the (!)-problems. Completing all of the (★)-problems constitutes the *minimal* preparation for exams.

(+) Problems marked (+) are meant to be more challenging than the (!)- and (★)-problems and will take you deeper into calculations and proofs and make connections to concepts across and beyond the course. It will not be necessary to do any (+)-problems to master the essential material of the course, but your experience may be richer (and more meaningful, and more fun) by considering them. If you have done all of the (!)- and (★)-problems, and the required and recommended problems from the textbook, and if you're still feeling bored or wondering if something is "missing," check out the (+)-problems.

Day 1: Monday, January 6.

Linearity pervades mathematics and science. An “operator” is **LINEAR** if (1) we can “add” its inputs and outputs in ways that “respect” all of the “usual” properties of addition of real numbers, (2) we can “multiply” its inputs and outputs by numbers in ways that, again, “respect” all of the “usual” properties of multiplication of real numbers, and (3) if the operator itself “respects” addition and “scalar” multiplication. Lots of quotes, lots of words, here are some symbols.

Let \mathcal{A} be that operator and let x and y be inputs. Then there is a notion of adding x and y so that $x + y$ is another input and behaves the way that we expect $+$ to behave; for example, $x + y = y + x$. We can also add the outputs $\mathcal{A}x$ and $\mathcal{A}y$, and this addition behaves as we expect, e.g., $\mathcal{A}x + \mathcal{A}y = \mathcal{A}y + \mathcal{A}x$. Here we are using the same symbol $+$ for addition on both the input side and the output side, even though the inputs and outputs could “live” in totally different “universes.”

And there is a notion of multiplying inputs and outputs by real numbers, which we denote by “juxtaposition.” That is, if c is a real number and x is an input, then cx is another input, and we have properties like $c(x + y) = cx + cy$ and $(cd)x = c(dx)$.

Now here is how \mathcal{A} respects addition and scalar multiplication:

$$\mathcal{A}(x + y) = \mathcal{A}x + \mathcal{A}y \quad \text{and} \quad \mathcal{A}(cx) = c(\mathcal{A}x). \quad (1.1)$$

These two identities are what we mean by the linearity of \mathcal{A} .

You already know some linear operators because you can do arithmetic and calculus. Say that we define $\mathcal{A}x := 2x$ for real numbers x . Or that we define $\mathcal{A}f := f'$, with f' as the derivative of a differentiable function f .

1.1 Problem (!). Prove that if \mathcal{A} is defined in either of these ways, then the linear identities (1.1) hold.

Many, many problems possess a linear structure. The inputs and outputs obey natural rules of addition and scalar multiplication, and the problem is either governed by or well-approximated by a linear operator. You’ve already met such problems in calculus. Many differential equations are linear; a fundamental problem of physics asks us to find functions f such that $f'' + f = g$ for a given function g . This compresses as $\mathcal{A}f = g$ with $\mathcal{A}f := f'' + f$ as a linear operator satisfying (1.1). Or maybe you want to approximate a hard problem in a nice way; you know how to do this with a local linear approximation. If you want to study a function f around a point x , and f is complicated, study instead $f(x + h) \approx f(x) + f'(x)h$. Here $\mathcal{A}h := f'(x)h$ is linear.

Much of the point of calculus is to *be wise and linearize*; the point of linear algebra is to understand the linear structure of that approximation. However, this is about as much calculus as we’ll do in this course. We will focus on a particular kind of linear operator that arises from *linear systems of equations*. These are hugely worthwhile problems in and of themselves, and many other problems that don’t look like linear systems of equations either hinge obliquely or are well-approximated by such systems. (Unfortunately, or fortunately, we will not see many, if any, “concrete” examples of *applications* of such systems—every

subdiscipline of math, and every scientific discipline allied with math, has its own favorite examples, and you probably won't be convinced of the worth of any one of those examples if you aren't already convinced.)

We can tease out a tremendous amount of structure and theory from very simple motivating examples, and here will be our favorite for the foreseeable future. Let's try to solve the **LINEAR SYSTEM**

$$\begin{cases} x - 2y = 1 \\ 3x + 2y = 11. \end{cases} \quad (1.2)$$

It's a system of equations because there is more than one equation, and it's linear because the unknowns only appear as the "linear powers" x and y , not x^2 or xy or $\cos(x + y)$.

By the way, you didn't need to get out of bed today and come to class to figure out how to solve it, but imagine if the system had 50 variables and 50 equations. You'd probably want a precise and systematic way of approaching it.

1.2 Problem (!). Try to solve (1.2). What does your gut instinct say to do? (If you're reading these notes for the first time and haven't been in class, *don't* read below this problem for our approach just yet—try it by yourself.)

Before we do anything to (1.2), here are some questions that we should ask.

1. Does it have a solution? That is, do there exist numbers x and y that make the two equalities in (1.2) true?
2. If not, why not? Can we quantify or qualify *failure* to solve a linear system?
3. Is there only one solution? Is there only one way to choose the values of x and y to make the two equalities in (1.2) true? That is, is the solution **UNIQUE**?
4. If not, why not? Can we quantify or qualify why a linear system might have more than one solution?

We will solve (1.2) by transforming it into an "equivalent" system of equations that is much easier to solve—actually, several "equivalent" systems. We'll say that two systems are **EQUIVALENT** if they have precisely the same solutions. And we'll do this via algebra.

Recall that if a and b are real numbers, then $a = b$ if and only if $ac = bc$ for all nonzero c . That is, if you know $a = b$, then you also know $ac = bc$ for all nonzero c (actually for $c = 0$, too, although that's boring). And if you know $ac = ab$ for all nonzero c (actually, for just *one* nonzero c), then you can divide to get $a = b$. In the context of linear systems, scaling *both* sides of the *same* equation by the *same nonzero* number doesn't change things. Let's multiply the first equation by the very convenient number -3 :

$$\begin{cases} x - 2y = 1 \\ 3x + 2y = 11 \end{cases} \iff \begin{cases} -3x + 6y = -3 \\ 3x + 2y = 11. \end{cases}$$

Now we'll use another property of algebra. Recall that if a and b are real numbers, then $a = b$ if and only if $a + c = b + c$ for all real numbers c . That is, if you know $a = b$, then you

can add c to both sides to get $a + c = b + c$. And if you know $a + c = b + c$ for some c , then just subtract c from both sides (or add $-c$ to both sides) to get $a = b$. (Actually, if you're watching your language, we said "for all" c , so if you know $a + c = b + c$ for *all* c , just take $c = 0$.)

Thus

$$\begin{cases} x - 2y = 1 \\ 3x + 2y = 11 \end{cases} \iff \begin{cases} -3x + 6y = -3 \\ 3x + 2y + c = 11 + c \end{cases}$$

for any c that we like. What if we take $c = -3x + 6y$ on the left and $c = -3$ on the right? The first equation says that these two versions of c have to be equal. Thus

$$\begin{aligned} \begin{cases} x - 2y = 1 \\ 3x + 2y = 11 \end{cases} &\iff \begin{cases} -3x + 6y = -3 \\ 3x + 2y + (-3x + 6y) = 11 + (-3) \end{cases} \\ &\iff \begin{cases} -3x + 6y = -3 \\ 8y = 8 \end{cases} \\ &\iff \begin{cases} -3x + 6y = -3 \\ y = 1 \end{cases} \end{aligned}$$

The second equation is extremely transparent, but the first looks worse than it originally did because of that extra factor of -3 . But our work above was redundant; there was no need to keep that -3 multiplying both sides of the first equation, and we could have divided by -3 at any time that we wanted. That is,

$$\begin{cases} x - 2y = 1 \\ 3x + 2y = 11 \end{cases} \iff \begin{cases} x - 2y = 1 \\ y = 1 \end{cases}$$

What this is saying is that x and y satisfy the first system precisely when they satisfy the second—and we know what y is from the second system. With this value of y , the first equation in the second system becomes

$$x - 2 = 1 \iff x = 3.$$

All of our work boils down to

$$\begin{cases} x - 2y = 1 \\ 3x + 2y = 11 \end{cases} \iff \begin{cases} x = 3 \\ y = 1. \end{cases} \quad (1.3)$$

This is an existence and uniqueness result for (1.2): there exists a solution ($x = 3$ and $y = 1$), and it is the only solution. Way to go.

1.3 Problem (!). What would you do if I asked you to *check* that $x = 3$ and $y = 1$ solves (1.2)? Would you repeat all of the work above, or would you just plug in these values and do arithmetic?

The preceding work illustrates two incredibly important operations in solving linear systems: multiply both sides of one equation by the same number, and subtract (or add) a multiple of one equation to another equation. There's a third operation—interchanging two equations, which sounds silly but actually is worthwhile—that we'll meet later. Eventually we will encode and view these operations at pretty high and abstract levels.

The preceding work also illustrates something that is incredibly *unimportant* about linear systems: what we call the variables. As long as you are consistent, it doesn't matter if you write x and y , or x_1 and x_2 , or α and β , and so on. What matters are the *coefficients* on the variables and the *numbers on the right*.

We are going to stack these numbers together as **COLUMN VECTORS**, which we'll just call “lists of numbers” right now. Here are the three important vectors in (1.2), and we'll also write them as ordered pairs to make typesetting easier:

$$\begin{bmatrix} 1 \\ 3 \end{bmatrix} = (1, 3), \quad \begin{bmatrix} -2 \\ 2 \end{bmatrix} = (-2, 2), \quad \text{and} \quad \begin{bmatrix} 1 \\ 11 \end{bmatrix} = (1, 11).$$

We'll do a lot of arithmetic with (column) vectors, and much of it will happen “componentwise.” We add vectors by adding their corresponding components, so

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

1.4 Problem (!). Compute

$$\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}.$$

Then we can rewrite the original problem (1.2) as

$$\begin{cases} x & - & 2y & = & 1 \\ 3x & + & 2y & = & 11 \end{cases} \iff \begin{bmatrix} x \\ 3x \end{bmatrix} + \begin{bmatrix} -2y \\ 2y \end{bmatrix} = \begin{bmatrix} 1 \\ 11 \end{bmatrix}.$$

Big deal, right? All we have done is introduced some new notation; this tells us absolutely nothing about solving (1.2) that we did not already know. Let's do one more bit of arithmetic. There are “common factors” of x and y in some of those vectors, and our gut instinct should be to factor them out.

So, we define multiplication of a vector by a number (we do *not* multiply two vectors) componentwise:

$$2 \begin{bmatrix} 1 \\ 3 \end{bmatrix} = \begin{bmatrix} 2(1) \\ 2(3) \end{bmatrix} = \begin{bmatrix} 2 \\ 6 \end{bmatrix}.$$

When multiplying a vector by a number, we always write the number first:

$$2 \begin{bmatrix} 1 \\ 3 \end{bmatrix}, \text{ not } \begin{bmatrix} 1 \\ 3 \end{bmatrix} 2 \quad \text{and} \quad c\mathbf{a}, \text{ not } \mathbf{a}c.$$

1.5 Problem (!). Compute

$$-1 \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \quad \text{and} \quad 0 \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}.$$

Content from Strang's ILA 6E. See the pictures on pp. v–vi for how to interpret vector addition and “scalar” multiplication in two dimensions. There is more componentwise arithmetic on pp. 1–2.

We rewrite (1.2) once again as

$$\begin{cases} x & - & 2y & = & 1 \\ 3x & + & 2y & = & 11 \end{cases} \iff x \begin{bmatrix} 1 \\ 3 \end{bmatrix} + y \begin{bmatrix} -2 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 \\ 11 \end{bmatrix}.$$

Again, this offers absolutely no insights into actually solving (1.2)—yet.

The expression

$$x \begin{bmatrix} 1 \\ 3 \end{bmatrix} + y \begin{bmatrix} -2 \\ 2 \end{bmatrix}$$

is something that we'll see often: it's a **LINEAR COMBINATION** of the vectors $(1, 3)$ and $(-2, 2)$. By the way, this is an example of typesetting a column vector as an ordered pair to save space. Many important ideas can be phrased in the language of linear combinations.

Content from Strang's ILA 6E. Page 3 has some pictures of linear combinations. See also a linear system on p3 that is written in vector form and then solved with elimination, as we did (1.2).

Day 2: Wednesday, January 8.**Vocabulary from today**

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Column vector of length n , linear combination of vectors, $m \times n$ matrix, matrix-vector product

Here are some more precise (well, mostly precise) definitions of concepts from our first pass at linear systems and vectors. Throughout, we use the following set-theoretic terminology as a convenient abbreviation: if S is a set and x is an element of S , then we write $x \in S$. For example, $1 \in \{1, 2, 3\}$. We denote by \mathbb{R} the set of all real numbers, so $1 \in \mathbb{R}$.

2.1 Undefinition. Let $n \geq 1$ be an integer. A **COLUMN VECTOR** of length n is an “ordered list” of n real numbers, which we call the **ENTRIES** or the **COMPONENTS** of \mathbf{v} . If \mathbf{v} is a

column vector of length n with entries v_1, \dots, v_n in that order, then we write

$$\mathbf{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \quad \text{or} \quad \mathbf{v} = (v_1, \dots, v_n).$$

The set of all column vectors of length n is \mathbb{R}^n , and we write

$$\mathbb{R}^n = \left\{ \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \mid v_1, \dots, v_n \in \mathbb{R} \right\}.$$

We typically work with $n \geq 2$, and we do not distinguish \mathbb{R}^1 and \mathbb{R} , so $\mathbb{R}^1 = \mathbb{R}$.

Two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ are **EQUAL** if and only if their corresponding entries are equal:

$$\begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} \iff v_j = w_j, \quad j = 1, \dots, n.$$

Why is this an “undefinition,” not a definition? Because we didn’t give a rigorous definition of “ordered list.” I like to think of column vectors of length n as *functions* from the set $\{1, \dots, n\}$ to \mathbb{R} . That is, if $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{R}^n$, then \mathbf{v} is the same as the function $f: \{1, \dots, n\} \rightarrow \mathbb{R}$ such that $f(j) = v_j$ for $j = 1, \dots, n$. And since functions are really sets of ordered pairs, $f = \{(j, v_j)\}_{j=1}^n$. This is probably a *useless* way to think about column vectors for day-to-day purposes, but it comforts me to know that there is deeper math behind that undefinition. If it doesn’t comfort you, it’s okay to move on.

We continue to define vector addition and multiplication by real numbers componentwise, regardless of the length of the vectors. In particular, if $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ and $c \in \mathbb{R}$, then $\mathbf{v} + \mathbf{w} \in \mathbb{R}^n$ and $c\mathbf{v} \in \mathbb{R}^n$. However, we only add vectors that have the same number of components, so something like

$$\begin{bmatrix} 1 \\ 2 \end{bmatrix} + \begin{bmatrix} 3 \\ 4 \\ 5 \end{bmatrix}$$

is not defined.

2.2 Example. We compute

$$0 \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = \begin{bmatrix} 0(1) \\ 0(2) \\ 0(3) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

I hope it’s obvious why we want to call the vector on the right the “zero vector in \mathbb{R}^3 .”

2.3 Definition. The **ZERO VECTOR** in \mathbb{R}^n is the vector $\mathbf{0}$ whose entries are all 0. Sometimes we will write $\mathbf{0}_n$ to emphasize that this is the zero vector with n entries.

For example,

$$\mathbf{0}_2 = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad \text{but} \quad \mathbf{0}_3 = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

2.4 Problem (!). (i) Let $\mathbf{v} \in \mathbb{R}^n$. What is $\mathbf{v} + \mathbf{0}_n$?

(ii) Does $\mathbf{0}_2 + \mathbf{0}_3$ make sense?

In studying our motivating toy problem, we encountered a “linear combination” of vectors. Here is that object in general.

2.5 Definition. Let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$ and $c_1, \dots, c_n \in \mathbb{R}$. The **LINEAR COMBINATION** of $\mathbf{v}_1, \dots, \mathbf{v}_n$ **WEIGHTED** by c_1, \dots, c_n is the vector $\mathbf{v} \in \mathbb{R}^m$ defined by

$$\mathbf{v} = c_1 \mathbf{v}_1 + \dots + c_n \mathbf{v}_n.$$

We may also express this in sigma notation:

$$\mathbf{v} = \sum_{j=1}^n c_j \mathbf{v}_j.$$

2.6 Problem (!). Convince yourself that, in the notation of the previous definition, we do indeed have $\mathbf{v} \in \mathbb{R}^m$. Also, what are the integers m and n encoding in that definition?

2.7 Problem (!). Let

$$\mathbf{e}_1 := \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{e}_2 := \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad \text{and} \quad \mathbf{e}_3 := \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$

Explain why any $\mathbf{v} \in \mathbb{R}^3$ is a linear combination of \mathbf{e}_1 , \mathbf{e}_2 , and \mathbf{e}_3 .

Content from Strang's ILA 6E. There are examples of linear combinations with $n = 2$ on p. vi and p. 2.

So far, none of this (mostly) more precise terminology tells us anything new about solving linear systems, and, honestly, none of the following is going to help, either. The goal is to build more terminology so that we can ask questions about linear systems *in the right language*.

Here is a major step toward that right language. Recall that our original problem (1.2) can be written as a system of linear equations or as a vector equation involving a linear

combination:

$$\begin{cases} x_1 & - & 2x_2 & = & 1 \\ 3x_1 & + & 2x_2 & = & 11 \end{cases} \iff x_1 \begin{bmatrix} 1 \\ 3 \end{bmatrix} + x_2 \begin{bmatrix} -2 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 \\ 11 \end{bmatrix}.$$

Let's put the coefficient vectors together into a matrix:

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

I hope you'll agree that this is a "square" matrix: it has 2 rows and 2 columns. We most often think of matrices in terms of columns (though rows are also useful). If we put

$$\mathbf{a}_1 := \begin{bmatrix} 1 \\ 3 \end{bmatrix} \quad \text{and} \quad \mathbf{a}_2 := \begin{bmatrix} -2 \\ 2 \end{bmatrix},$$

then we will also write A as

$$A = [\mathbf{a}_1 \quad \mathbf{a}_2].$$

This is sort of a "row vector" of column vectors.

Here is where we are going with all of this. Abbreviate $\mathbf{x} = (x_1, x_2)$ and $\mathbf{b} = (1, 11)$. Our goal is to define a notion of "matrix-vector multiplication" so that if $A\mathbf{x}$ is the "product" of A and \mathbf{x} , then our original problem compresses to

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

First, of course, we need some more terminology. We control the "sizes" or "dimensions" of matrices by counting the numbers of rows and the numbers of columns—and we always list *rows before columns*. We'll say $A \in \mathbb{R}^{2 \times 2}$ for the matrix A above, and I hope you believe that

$$\begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix} \in \mathbb{R}^{2 \times 3} \quad \text{and} \quad \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix} \in \mathbb{R}^{3 \times 2}.$$

More generally, we say the following.

2.8 Definition. Let $m, n \geq 1$ be integers. An $m \times n$ **MATRIX** is a rectangular array of numbers with m rows and n columns. We denote the set of all $m \times n$ matrices by $\mathbb{R}^{m \times n}$.

Since a matrix with m rows and 1 column is really just an ordered list of m numbers, we will not distinguish $\mathbb{R}^{m \times 1}$ and \mathbb{R}^m , so $\mathbb{R}^{m \times 1} = \mathbb{R}^m$. Also, $\mathbb{R}^{1 \times 1} = \mathbb{R}$. But we do not equate $\mathbb{R}^{1 \times n}$ and \mathbb{R}^n .

The (i, j) -**ENTRY** of a matrix is the entry in row i , column j of that matrix. Sometimes we will write A_{ij} for the (i, j) -entry of A , although with large matrices it might be clearer to write $A_{i,j}$. Two matrices are **EQUAL** if and only if they have the same number of rows and columns and if all of their corresponding entries are equal.

Regarding that last caveat, we have things like

$$\begin{bmatrix} 1 & 2 & 3 \end{bmatrix} \neq \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = (1, 2, 3).$$

2.9 Problem (!). Reread that until it makes sense.

2.10 Example. Let

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix}.$$

The (1, 2)-entry of A is 2, and the (2, 1)-entry of A is 3.

Content from Strang's *ILA* 6E. A 3×2 matrix appears on p. vi, and a larger one (what size?) on p. vii.

As with column vectors, our attempt at defining a matrix is really an undefinition because we did not rigorously define “rectangular array” of numbers. If you really want to, you can think of $A \in \mathbb{R}^{m \times n}$ as the function $f: I \rightarrow \mathbb{R}$ such that $f(i, j) = A_{ij}$, where $I = \{(i, j) \mid i = 1, \dots, m, j = 1, \dots, n\}$. Or as the function $g: \{1, \dots, n\} \rightarrow \mathbb{R}^m$ such that $g(j) = \mathbf{a}_j$, where \mathbf{a}_j is the j th column of A , i.e., $A = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n]$. Neither way of thinking will make any of the following any easier.

And as with column vectors, we add matrices and multiply them by real numbers componentwise.

2.11 Problem (!). Compute

$$\begin{bmatrix} 1 & 2 \\ 0 & 0 \\ -1 & -1 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 3 & 4 \\ 5 & 6 \end{bmatrix} \quad \text{and} \quad 2 \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}.$$

We are finally ready to think about linear systems. With

$$A := \begin{bmatrix} 1 & -2 \\ 3 & 2 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 1 \\ 11 \end{bmatrix}, \quad \text{and} \quad \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix},$$

how should we define the symbol $A\mathbf{x}$ so that

$$\begin{cases} x_1 - 2x_2 = 1 \\ 3x_1 + 2x_2 = 11 \end{cases} \iff x_1 \begin{bmatrix} 1 \\ 3 \end{bmatrix} + x_2 \begin{bmatrix} -2 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 \\ 11 \end{bmatrix} \iff A\mathbf{x} = \mathbf{b}?$$

The answer is pretty much staring us in the face:

$$A\mathbf{x} := x_1 \begin{bmatrix} 1 \\ 3 \end{bmatrix} + x_2 \begin{bmatrix} -2 \\ 2 \end{bmatrix}.$$

This is something new. This is not a componentwise definition of multiplication. *Instead, the idea behind matrix-vector multiplication is that we take a linear combination of the columns of the matrix weighted by the entries of the vector.* If we write

$$A = [\mathbf{a}_1 \ \mathbf{a}_2], \quad \mathbf{a}_1 = \begin{bmatrix} 1 \\ 3 \end{bmatrix}, \quad \mathbf{a}_2 = \begin{bmatrix} -2 \\ 2 \end{bmatrix},$$

then we are saying

$$A\mathbf{x} = [\mathbf{a}_1 \quad \mathbf{a}_2] \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = x_1\mathbf{a}_1 + x_2\mathbf{a}_2.$$

Let's do some computations with this definition of matrix-vector multiplication in words first: take the linear combination of the columns of the matrix with the weights as the entries from the vector, all appearing in order.

2.12 Problem (!). Convince yourself that for this to work, the number of columns of the matrix has to equal the number of entries of the vector.

2.13 Example. (i) $\begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} = 1 \begin{bmatrix} 1 \\ 2 \end{bmatrix} + 0 \begin{bmatrix} 3 \\ 4 \end{bmatrix} + 1 \begin{bmatrix} 5 \\ 6 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 5 \\ 6 \end{bmatrix} = \begin{bmatrix} 1+0+5 \\ 2+0+6 \end{bmatrix} = \begin{bmatrix} 6 \\ 8 \end{bmatrix}$

(ii) $\begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 3 & 0 \end{bmatrix} \begin{bmatrix} 2 \\ 4 \end{bmatrix} = 2 \begin{bmatrix} 1 \\ 0 \\ 3 \end{bmatrix} + 4 \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix} = \begin{bmatrix} 2 \\ 0 \\ 6 \end{bmatrix} + \begin{bmatrix} 0 \\ 8 \\ 0 \end{bmatrix} = \begin{bmatrix} 2+0 \\ 0+8 \\ 6+0 \end{bmatrix} = \begin{bmatrix} 2 \\ 8 \\ 6 \end{bmatrix}$

And now for the definition in symbols.

2.14 Definition. Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{v} \in \mathbb{R}^n$ with

$$A = [\mathbf{a}_1 \quad \cdots \quad \mathbf{a}_n] \quad \text{and} \quad \mathbf{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}.$$

The **MATRIX-VECTOR PRODUCT** of A and \mathbf{v} is

$$A\mathbf{v} = [\mathbf{a}_1 \quad \cdots \quad \mathbf{a}_n] \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} = v_1\mathbf{a}_1 + \cdots + v_n\mathbf{a}_n = \sum_{j=1}^n v_j\mathbf{a}_j.$$

Content from Strang's ILA 6E. Examples of matrix-vector multiplication appear on p. 1.

2.15 Problem (!). Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{v} \in \mathbb{R}^n$. How many entries does $A\mathbf{v}$ have? Use the definition of $A\mathbf{v}$ from Definition 2.14.

2.16 Problem (*). Let $A \in \mathbb{R}^{m \times n}$. Prove that $A\mathbf{0}_n = \mathbf{0}_m$. Use the definition of $A\mathbf{0}_n$ from Definition 2.14.

2.17 Problem (★). Prove that matrix-vector multiplication is **LINEAR** in the following sense: if $A \in \mathbb{R}^{m \times n}$, $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$, and $c \in \mathbb{R}$, then

$$A(\mathbf{v} + \mathbf{w}) = A\mathbf{v} + A\mathbf{w} \quad \text{and} \quad A(c\mathbf{v}) = c(A\mathbf{v}).$$

This could involve a lot of \dots that might obscure the actual arithmetic going on; if it makes things more transparent, do it for $n = 2$ or $n = 3$ first. *However you do it, use the definition of matrix-vector multiplication from Definition 2.14.*

Every linear system compresses as a matrix-vector equation. Suppose there are m equations in n unknowns. Let \mathbf{x} be the column vector of length n that contains all of these unknowns. Let A be the $m \times n$ matrix containing all of the coefficients, so the (i, j) -entry of A is the coefficient on the j th unknown in the i th equation. Let \mathbf{b} be the column vector of length m that contains the right sides of these equations. Then the problem is

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

Our original questions remain the same—how to solve it, how to understand failure to solve it. The new question is probably *Why is writing it as $A\mathbf{x} = \mathbf{b}$ any better than the original way?*

Content from Strang's ILA 6E. Read all of p. 2 right now.

Day 3: Friday, January 10.

No class due to weather. The following problems will reinforce our work on matrices and matrix-vector multiplication.

3.1 Problem (!). Rewrite each linear system below as a matrix-vector equation $A\mathbf{x} = \mathbf{b}$ for some matrix $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^n$. Specify the values of m and n in each case.

$$(i) \quad \begin{cases} x_1 + 2x_2 + 3x_4 = 1 \\ x_3 + 4x_4 = 2 \end{cases}$$

$$(ii) \quad \begin{cases} x_1 + 2x_2 + x_3 + 7x_4 = 1 \\ 2x_1 + 4x_2 + 2x_3 + 14x_4 = 2 \\ 2x_3 + 8x_4 = 3 \end{cases}$$

$$(iii) \quad \begin{cases} x_1 = 1 \\ 2x_1 = 2 \\ x_2 = 3 \\ x_3 = 4 \end{cases}$$

$$(iv) \begin{cases} x_1 + 2x_2 = 1 \\ 2x_1 + 4x_2 = 2 \\ x_1 + 2x_2 + 2x_3 = 3 \\ 7x_1 + 14x_2 + 8x_3 = 4 \end{cases}$$

3.2 Problem (★). Compute each matrix-vector product and then describe in words the effect of this multiplication. For your description in words, pretend that you are talking out loud to a classmate about this multiplication, and you do not have any paper or board to write on; try to use as few symbols as possible in your description.

$$(i) \begin{bmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 9 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

$$(ii) \begin{bmatrix} 1 & 0 & 0 \\ 0 & c & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \text{ for any } c, x_1, x_2, x_3 \in \mathbb{R}$$

$$(iii) \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ -3 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \text{ for any } x_1, x_2, x_3 \in \mathbb{R}$$

$$(iv) \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \text{ for any } x_1, x_2 \in \mathbb{R}$$

Day 4: Monday, January 13.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Dot product of vectors in \mathbb{R}^n

No class due to conflict with department chair interviews. Please read and work through the following material—word by word, line by line. Check all calculations and details.

The goal of the class is the same as always: solve $A\mathbf{x} = \mathbf{b}$, and when we can't solve it, understand why. Eventually this will take us into understanding just A , apart from any linear systems. For now, we should try to understand $A\mathbf{x}$ as best as we can. There is another way of computing matrix-vector products in addition to Definition 2.14. We'll tease it out in an example.

4.1 Example. We compute

$$\begin{bmatrix} 1 & -2 \\ 3 & 2 \end{bmatrix} \begin{bmatrix} 3 \\ 1 \end{bmatrix} = 3 \begin{bmatrix} 1 \\ 3 \end{bmatrix} + 1 \begin{bmatrix} -2 \\ 2 \end{bmatrix} = \begin{bmatrix} 3 \\ 9 \end{bmatrix} + \begin{bmatrix} -2 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 \\ 11 \end{bmatrix}.$$

This is just checking that $x_1 = 3$ and $x_2 = 1$ solves our original problem

$$\begin{cases} x_1 - 2x_2 = 1 \\ 3x_1 + 2x_2 = 11 \end{cases},$$

right?

Here is another way of looking at this arithmetic:

$$3 \begin{bmatrix} 1 \\ 3 \end{bmatrix} + 1 \begin{bmatrix} -2 \\ 2 \end{bmatrix} = \begin{bmatrix} 3(1) + 1(-2) \\ 3(3) + 1(2) \end{bmatrix}.$$

Do you see how the vectors $(3, 1)$ and $(1, -2)$ appear in the first component on the right? And how $(3, 1)$ and $(3, 2)$ appear in the second component? It's almost as though the vector by which we're multiplying the matrix, and the rows of the matrix *viewed as column vectors*, are doing all of the arithmetic.

Let's introduce a new structure: the **DOT PRODUCT** of vectors in \mathbb{R}^2 . (Just \mathbb{R}^2 now for starters.) Put

$$\begin{bmatrix} v_1 \\ v_2 \end{bmatrix} \cdot \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} := v_1 w_1 + v_2 w_2.$$

So we have

$$\begin{bmatrix} 1 \\ -2 \end{bmatrix} \cdot \begin{bmatrix} 3 \\ 1 \end{bmatrix} = 1(3) + (-2)(1) = 3 - 2 = 1$$

and

$$\begin{bmatrix} 3 \\ 2 \end{bmatrix} \cdot \begin{bmatrix} 3 \\ 1 \end{bmatrix} = 3(3) + 2(1) = 9 + 2 = 11.$$

Here is the takeaway in words: we can compute a matrix-vector product by taking the dot product of the rows of the matrix—*viewed as column vectors*—with the vector in the product.

Content from Strang's ILA 6E. Equation (1) on p. 9 defines the dot product of vectors in \mathbb{R}^2 . See the box above on p. 9 for more dot products.

Let's generalize this example.

4.2 Definition. The dot product of $\mathbf{v} = (v_1, \dots, v_n)$, $\mathbf{w} = (w_1, \dots, w_n) \in \mathbb{R}^n$ is

$$\mathbf{v} \cdot \mathbf{w} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \cdot \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} = v_1 w_1 + \dots + v_n w_n = \sum_{j=1}^n v_j w_j.$$

Content from Strang's ILA 6E. This is equation (2) on p. 9. We won't talk about anything else from Section 1.2 for quite a while. The dot product turns out to be the key to a deeper *geometric* understanding of \mathbb{R}^n , in particular an understanding of *angles*, but we won't need that for some time.

4.3 Example.
$$\begin{bmatrix} 3 \\ 4 \\ 5 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = 3(1) + 4(0) + 5(0) = 3$$

I will do my best to reserve the symbol \cdot for the dot product and use “juxtaposition” to denote multiplication of real numbers, e.g., $3(1)$, not $3 \cdot 1$. But I guess the dot product in $\mathbb{R}^1 = \mathbb{R}$ is just ordinary multiplication, so no big deal.

4.4 Problem (★). Prove that the dot product is **COMMUTATIVE** in the sense that $\mathbf{v} \cdot \mathbf{w} = \mathbf{w} \cdot \mathbf{v}$ for all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$. This is how we expect multiplication to behave, that $xy = yx$ for all numbers x and y , right?

We can use the dot product to “extract” components of a vector. This will be a hugely useful operation.

4.5 Example. Here is how this works in \mathbb{R}^3 . (I like \mathbb{R}^3 : it's big enough to be interesting but not so big that it's intimidating.) Put

$$\mathbf{e}_1 := \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{e}_2 := \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad \text{and} \quad \mathbf{e}_3 := \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$

These are the **STANDARD BASIS VECTORS** for \mathbb{R}^3 , and we will use them a lot. I claim that if $\mathbf{v} = (v_1, v_2, v_3) \in \mathbb{R}^3$, then

$$\mathbf{v} \cdot \mathbf{e}_1 = v_1, \quad \mathbf{v} \cdot \mathbf{e}_2 = v_2, \quad \text{and} \quad \mathbf{v} \cdot \mathbf{e}_3 = v_3.$$

We basically did the first equality in Example 4.3, so here is the second:

$$\mathbf{v} \cdot \mathbf{e}_2 = \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} = v_1(0) + v_2(1) + v_3(0) = v_2.$$

I'll let you check the third.

Now here is another nice identity: start with \mathbf{v} and “expand it”:

$$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \begin{bmatrix} v_1 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ v_2 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ v_3 \end{bmatrix}$$

$$\begin{aligned}
&= v_1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + v_2 \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} + v_3 \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = v_1 \mathbf{e}_1 + v_2 \mathbf{e}_2 + v_3 \mathbf{e}_3 \\
&= (\mathbf{v} \cdot \mathbf{e}_1) \mathbf{e}_1 + (\mathbf{v} \cdot \mathbf{e}_2) \mathbf{e}_2 + (\mathbf{v} \cdot \mathbf{e}_3) \mathbf{e}_3 = \sum_{j=1}^3 (\mathbf{v} \cdot \mathbf{e}_j) \mathbf{e}_j.
\end{aligned}$$

This is a really clean representation of a vector in terms of its components and some other, simpler vectors. We'll return to such representations many times in the future.

4.6 Problem (+). The **STANDARD BASIS VECTORS** IN \mathbb{R}^n are the vectors $\mathbf{e}_1, \dots, \mathbf{e}_n \in \mathbb{R}^n$ defined as follows: the components of \mathbf{e}_j are all 0, except for the component in row j , which is 1.

(i) Write out the standard basis vectors in \mathbb{R}^5 . You should make clear what all of their entries are.

(ii) Prove that

$$\mathbf{e}_j \cdot \mathbf{e}_k = \begin{cases} 1, & j = k \\ 0, & j \neq k. \end{cases}$$

(iii) Let $\mathbf{v} \in \mathbb{R}^n$. Prove that

$$\mathbf{v} = \sum_{j=1}^n (\mathbf{v} \cdot \mathbf{e}_j) \mathbf{e}_j.$$

Now that we have an understanding of the mechanics of dot product calculations, we can examine how the dot product arises in matrix-vector multiplication. All of the ideas are in Example 4.1. We'll work with a matrix with three columns to see this a little more abstractly. Let $A \in \mathbb{R}^{m \times 3}$ and write

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ * & * & * \end{bmatrix}.$$

I just want to focus on the first row of A , so I've listed that out explicitly. The symbols $*$ below denote the remaining $m - 1$ rows of A . The exact values of the entries in those rows are wholly unimportant right now. (If it makes you feel better, take $m = 2$ and replace each $*$ with 0.)

Let $\mathbf{v} = (v_1, v_2, v_3) \in \mathbb{R}^3$. Then

$$A\mathbf{v} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ * & * & * \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = v_1 \begin{bmatrix} a_{11} \\ * \end{bmatrix} + v_2 \begin{bmatrix} a_{12} \\ * \end{bmatrix} + v_3 \begin{bmatrix} a_{13} \\ * \end{bmatrix} = \begin{bmatrix} v_1 a_{11} + v_2 a_{12} + v_3 a_{13} \\ * \end{bmatrix}.$$

At the risk of being annoying, I am using the same symbol $*$ to denote rows 2 through m of the columns of A and then the vector $A\mathbf{v}$; I sincerely don't care what's going on there right

now. Here is what we have shown: the first component of $A\mathbf{v}$ is

$$v_1a_{11} + v_2a_{12} + v_3a_{13} = \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \end{bmatrix} \cdot \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \end{bmatrix} \cdot \mathbf{v},$$

which is the dot product of the first row of A viewed as a column vector with \mathbf{v} .

This generalizes substantially; the proof is just good bookkeeping and good notation.

4.7 Theorem. Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{v} \in \mathbb{R}^n$. The i th component of $A\mathbf{v}$ is the dot product of row i of A viewed as a column vector and \mathbf{v} .

4.8 Example. We compute

$$\begin{bmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 9 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1(1) + 4(0) + 7(1) \\ 2(1) + 5(0) + 8(1) \\ 3(1) + 6(0) + 9(1) \end{bmatrix} = \begin{bmatrix} 8 \\ 10 \\ 12 \end{bmatrix}.$$

What do you get if you use Definition 2.14?

Content from Strang's ILA 6E. Read about the “row picture” and the “column picture” on p. 19. (The matrix is A given on p. 18.) Strang says it best: to *compute* $A\mathbf{v}$ by hand for “small” A and \mathbf{v} , use dot products, but to *understand* $A\mathbf{v}$, use the “linear combination of columns” definition. This is morally similar to the derivative: to compute it by hand, use the product rule or chain rule or something like that, but to understand it, use the limit definition.

4.9 Problem (★). Go back and redo each of the matrix-vector products in Example 2.13 and Problem 3.2 with dot products. What do you find easier for work by hand: Definition 2.14 or Theorem 4.7?

Day 5: Wednesday, January 15.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Column space of a matrix

We started thinking about matrices *statically*: they encode data, specifically the coefficients of a linear system of equations. Now that we can multiply matrices and vectors, we can think *dynamically*: matrices act on vectors to produce new vectors. We might even associate a matrix $A \in \mathbb{R}^{m \times n}$ with a “map” (dare I say “function”?) that associates each vector $\mathbf{v} \in \mathbb{R}^n$ with a new vector $A\mathbf{v} \in \mathbb{R}^m$.

Matrix-vector multiplication tells us useful things about matrices, not just vectors. I first claim that matrix-vector multiplication can “extract” the columns of a matrix. Let’s start small. As before, we’ll write

$$\mathbf{e}_1 := \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{e}_2 := \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad \text{and} \quad \mathbf{e}_3 := \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$

Let $A = [\mathbf{a}_1 \ \mathbf{a}_2 \ \mathbf{a}_3] \in \mathbb{R}^{m \times 3}$. It’s important that A has only three columns, but here the number of rows doesn’t matter. We compute

$$A\mathbf{e}_1 = [\mathbf{a}_1 \ \mathbf{a}_2 \ \mathbf{a}_3] \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = 1\mathbf{a}_1 + 0\mathbf{a}_2 + 0\mathbf{a}_3 = \mathbf{a}_1 + \mathbf{0}_m + \mathbf{0}_m = \mathbf{a}_1.$$

In words, *multiplying by \mathbf{e}_1 extracted the first column of A .*

5.1 Problem (!). With $A \in \mathbb{R}^{m \times 3}$ as above, show that $A\mathbf{e}_2 = \mathbf{a}_2$ and $A\mathbf{e}_3 = \mathbf{a}_3$.

This generalizes nicely.

5.2 Theorem. Let $\mathbf{e}_1, \dots, \mathbf{e}_n \in \mathbb{R}^n$ be the standard basis vectors for \mathbb{R}^n : the components of \mathbf{e}_j are all 0, except for the component in row j , which is 1. Let $A \in \mathbb{R}^{m \times n}$. Then $A\mathbf{e}_j$ is the j th column of A .

5.3 Problem (★). Prove it!

5.4 Problem (★). Let $I_n \in \mathbb{R}^{n \times n}$ be the matrix whose j th column is \mathbf{e}_j . We might write $I_n = [\mathbf{e}_1 \ \cdots \ \mathbf{e}_n]$. Prove that $I_n \mathbf{v} = \mathbf{v}$ for any $\mathbf{v} \in \mathbb{R}^n$. We therefore call I_n the **IDENTITY MATRIX**: multiplying \mathbf{v} by I_n just tells you what \mathbf{v} is. [Hint: prove it for $n = 3$ first to see the pattern of the arithmetic before doing it for n arbitrary.]

Content from Strang’s ILA 6E. Look at the four matrices on p. 18: identity, diagonal, triangular, symmetric. Why are the last three called what they are?

We can go further than data extraction via matrix-vector multiplication. I like to say that *what things do defines what things are*. And what matrices do is multiply vectors! Recall that two matrices $A, B \in \mathbb{R}^{m \times n}$ are equal if their corresponding entries are all equal: $A_{ij} = B_{ij}$ for $i = 1, \dots, m$ and $j = 1, \dots, n$. That is a “static” way of viewing matrix equality (and not a bad way at all). Here is the “dynamic” way: A and B are equal if they always do the same thing to the same vector.

5.5 Theorem. Let $A, B \in \mathbb{R}^{m \times n}$. Then $A = B$ if and only if $A\mathbf{v} = B\mathbf{v}$ for all $\mathbf{v} \in \mathbb{R}^n$.

Proof. This is an “if and only if” statement, so we need to prove two things. First we want to assume that $A = B$ and then show that $A\mathbf{v} = B\mathbf{v}$ for all $\mathbf{v} \in \mathbb{R}^n$. This feels pretty silly, right? We should just be able to “substitute” A in for B . If we want to be pickier and more precise about what $=$ means here, $A = B$ means that A and B have equal entries, so also equal columns. That is, $A = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n]$ and $B = [\mathbf{b}_1 \ \cdots \ \mathbf{b}_n]$ with $\mathbf{a}_j = \mathbf{b}_j$ for all j . (And what does $\mathbf{a}_j = \mathbf{b}_j$ mean? Componentwise equality.) So, if $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{R}^n$, then

$$A\mathbf{v} = \sum_{j=1}^n v_j \mathbf{a}_j = \sum_{j=1}^n v_j \mathbf{b}_j = B\mathbf{v}.$$

Now we want to show that if $A\mathbf{v} = B\mathbf{v}$ for all $\mathbf{v} \in \mathbb{R}^n$, then $A = B$. The key words here are “for all.” We can pick any $\mathbf{v} \in \mathbb{R}^n$ that we like, and we will have the equality $A\mathbf{v} = B\mathbf{v}$. If we want to extract data about A and B , there are good, specific choices for \mathbf{v} : take $\mathbf{v} = \mathbf{e}_j$. Then $A\mathbf{e}_j = B\mathbf{e}_j$ for each j , and so the j th column of A equals the j th column of B . That means $A = B$. ■

5.6 Problem (!). Are you sure about that? If $A, B \in \mathbb{R}^{m \times n}$, and the j th column of A equals the j th column of B for $j = 1, \dots, n$, why do we have $A = B$? [Hint: *there are several versions of equality here: equality of vectors in \mathbb{R}^m , equality of matrices in $\mathbb{R}^{m \times n}$, equality of numbers in \mathbb{R} . What role does each version play in answering the question?*]

We now have as good an understanding of matrix-vector multiplication as we’re going to get without doing anything new. Remember that our goal in this course is to understand the problem $A\mathbf{x} = \mathbf{b}$ as best as we can. Our work so far has focused on understanding $A\mathbf{x}$. Now it is time to relate \mathbf{b} to A .

By definition, $A\mathbf{x}$ is a linear combination of the columns of A weighted by the entries of \mathbf{x} . To have $A\mathbf{x} = \mathbf{b}$, we therefore want to be able to express \mathbf{b} as a linear combination of the columns of A . We give this a special name.

5.7 Definition. The **COLUMN SPACE** of $A \in \mathbb{R}^{m \times n}$ is the set of all linear combinations of the columns of A . We denote it by $\mathbf{C}(A)$, and every vector in $\mathbf{C}(A)$ is a vector in \mathbb{R}^m . Equivalently,

$$\mathbf{C}(A) = \{A\mathbf{v} \mid \mathbf{v} \in \mathbb{R}^n\}.$$

Content from Strang’s ILA 6E. The column space is defined at the bottom of p. 20.

5.8 Example. Let

$$A = \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix}.$$

Then

$$\mathbf{C}(A) = \{A\mathbf{v} \mid \mathbf{v} \in \mathbb{R}^2\} = \left\{ v_1 \begin{bmatrix} 2 \\ 0 \end{bmatrix} + v_2 \begin{bmatrix} 0 \\ 3 \end{bmatrix} \mid v_1, v_2 \in \mathbb{R} \right\} = \left\{ \begin{bmatrix} 2v_1 \\ 3v_2 \end{bmatrix} \mid v_1, v_2 \in \mathbb{R} \right\}.$$

To be able to solve $A\mathbf{x} = \mathbf{b}$ for as many \mathbf{b} as possible, we want $\mathbf{C}(A)$ to be as “large” as possible. Ideally (perhaps) we would have $\mathbf{C}(A) = \mathbb{R}^m$. Then every $\mathbf{b} \in \mathbb{R}^m$ would be in $\mathbf{C}(A)$, so every $\mathbf{b} \in \mathbb{R}^m$ would be a linear combination of the columns of A , and so we could solve $A\mathbf{x} = \mathbf{b}$ for all $\mathbf{b} \in \mathbb{R}^m$. Whether or not that is true, $\mathbf{C}(A)$ consists of all the vectors in \mathbb{R}^m for which we can solve $A\mathbf{x} = \mathbf{b}$. Right now, understanding the column space does not replace solving $A\mathbf{x} = \mathbf{b}$, and understanding the column space gives us no new tools or algorithms for solving $A\mathbf{x} = \mathbf{b}$ efficiently. That’s coming.

5.9 Example. With A as in Example 5.8, we claim that $\mathbf{C}(A) = \mathbb{R}^2$. We need to take an arbitrary $\mathbf{b} = (b_1, b_2) \in \mathbb{R}^2$ and show $\mathbf{b} \in \mathbf{C}(A)$. That is, we need to find $\mathbf{v} \in \mathbb{R}^2$ such that $A\mathbf{v} = \mathbf{b}$. From Example 5.8, it suffices to find $v_1, v_2 \in \mathbb{R}$ such that

$$\begin{bmatrix} 2v_1 \\ 3v_2 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}.$$

Looking at componentwise equalities, this is equivalent to

$$2v_1 = b_1 \quad \text{and} \quad 3v_2 = b_2,$$

and that is the same as

$$v_1 = \frac{b_1}{2} \quad \text{and} \quad v_2 = \frac{b_2}{3}.$$

This tells us what \mathbf{v} should be for us to have $\mathbf{b} = A\mathbf{v}$, and we get something more: there is only one way to define \mathbf{v} in terms of \mathbf{b} , because there is only one way to define v_1 and v_2 in terms of b_1 and b_2 .

5.10 Problem (★). (i) Prove that

$$\mathbf{C}\left(\begin{bmatrix} 1 & -2 \\ 3 & 2 \end{bmatrix}\right) = \mathbb{R}^2.$$

[Hint: repeat the work that brought us from (1.2) to (1.3) but instead of having the right side of that system be $(1, 11)$, use an arbitrary $\mathbf{b} = (b_1, b_2)$.]

(ii) What is

$$\mathbf{C}\left(\begin{bmatrix} 1 & -2 & 4 \\ 3 & 2 & 5 \end{bmatrix}\right)?$$

[Hint: don’t reinvent the wheel. You know the column space from the previous part, and you know that this column space is the set of all linear combinations of the form

$$v_1 \begin{bmatrix} 1 \\ -2 \end{bmatrix} + v_2 \begin{bmatrix} 3 \\ 2 \end{bmatrix} + v_3 \begin{bmatrix} 4 \\ 5 \end{bmatrix}.$$

Is there an “easy” value that you can pick for v_3 to relate this linear combination to what would appear in the previous part?]

Failure in math and life teaches us a lot, and there is a lot to be learned from what happens when $\mathbf{C}(A) \neq \mathbb{R}^m$ for $A \in \mathbb{R}^{m \times n}$. Here are some problematic A .

5.11 Example. (i) Let

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

If $\mathbf{v} = (v_1, v_2)$, then $A\mathbf{v} = (2v_1, 0)$. That is, the second component of $A\mathbf{v}$ is always 0, so if $\mathbf{b} = (b_1, b_2) \in \mathbf{C}(A)$, then $b_2 = 0$. Surely not all vectors in \mathbb{R}^2 have 0 as their second component; for example, $(1, 1) \notin \mathbf{C}(A)$.

(ii) Let

$$A = \begin{bmatrix} 1 & -2 \\ 3 & -6 \end{bmatrix}.$$

For $v_1, v_2 \in \mathbb{R}$, we have

$$v_1 \begin{bmatrix} 1 \\ 3 \end{bmatrix} + v_2 \begin{bmatrix} -2 \\ -6 \end{bmatrix} = v_1 \begin{bmatrix} 1 \\ 3 \end{bmatrix} - 2v_2 \begin{bmatrix} 1 \\ 3 \end{bmatrix} = (v_1 - 2v_2) \begin{bmatrix} 1 \\ 3 \end{bmatrix}.$$

(You believe that multiplication distributes over addition, right? That $c_1\mathbf{v} + c_2\mathbf{v} = (c_1 + c_2)\mathbf{v}$?) This calculation says that every $\mathbf{b} \in \mathbf{C}(A)$ is a multiple of $(1, 3)$. Is every vector in \mathbb{R}^2 a multiple of $(1, 3)$? Surely not: something like $(0, 1)$ cannot be written as

$$\begin{bmatrix} 0 \\ 1 \end{bmatrix} = c \begin{bmatrix} 1 \\ 3 \end{bmatrix}.$$

What goes wrong in an equality like that?

(iii) Let

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

I think you'll agree that any $\mathbf{b} = (b_1, b_2, b_3) \in \mathbf{C}(A)$ has $b_3 = 0$; the deadly thing is that row of all 0. If not, let's use dot products for a change:

$$\begin{bmatrix} 1 & 0 & 3 \\ 0 & 2 & 4 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \begin{bmatrix} 1(v_1) + 0(v_2) + 3v_3 \\ 0(v_1) + 2v_2 + 4v_3 \\ 0(v_1) + 0(v_2) + 0(v_3) \end{bmatrix} = \begin{bmatrix} v_1 + 3v_3 \\ 2v_2 + 4v_3 \\ 0 \end{bmatrix}.$$

What really was going on in the previous example? The rows of zeros in the first and third matrices were problematic, but the column space is about *columns*.

5.12 Example. Let's take another look at those matrices. I think it's easier to start with

$$\begin{bmatrix} 1 & -2 \\ 3 & -6 \end{bmatrix}$$

and recall our arithmetic to see that the second column is -2 times the first column. Then maybe we'll recognize that the second column of

$$\begin{bmatrix} 2 & 0 \\ 0 & 0 \end{bmatrix}$$

is 0 times the first column. Even though these matrices have two columns, only one matters—somehow there is “redundant” data in the matrix!

Is there redundancy in

$$\begin{bmatrix} 1 & 0 & 3 \\ 0 & 2 & 4 \\ 0 & 0 & 0 \end{bmatrix}?$$

I claim that no column is a multiple of another—this is annoying to check, but it builds character, so you should do it. (Here's how to get started: if the first column is c times the second column, won't some of the zero and nonzero entries interact badly?) But maybe, if we're lucky, we'll notice patterns relating the third column to the first and second. Because life is short, I'll tell you those patterns:

$$\begin{bmatrix} 3 \\ 4 \\ 0 \end{bmatrix} = \begin{bmatrix} 3 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 4 \\ 0 \end{bmatrix} = 3 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + 2 \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix}.$$

The third column is a linear combination of the first two. This is redundancy again: at a “linear” level, the third column can be recovered from the first two.

In fact, the third column just disappears when looking at the column space:

$$\begin{aligned} v_1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + v_2 \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix} + v_3 \begin{bmatrix} 3 \\ 4 \\ 0 \end{bmatrix} &= v_1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + v_2 \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix} + v_3 \left(3 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + 2 \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix} \right) \\ &= (v_1 + 3v_3) \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + (v_2 + 2v_3) \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix}. \end{aligned}$$

Content from Strang's *ILA* 6E. Look at Examples 1, 2, and 3 on p. 20. Check that we can always solve $A_1 \mathbf{x} = \mathbf{b}$ for any \mathbf{b} , and that there is only one choice of \mathbf{x} that works. Then check all of the arithmetic that appears in the statements about A_2 and A_3 .

So why is this bad? Why do “redundant” columns make the column space smaller than we'd like? Do they always do that? And can we be more precise than “redundant”?

5.13 Problem (+). Here is a generalization of these issues.

(i) Let $\mathbf{a} \in \mathbb{R}^2$, $c \in \mathbb{R}$, and $A = \begin{bmatrix} \mathbf{a} & c\mathbf{a} \end{bmatrix}$. First explain why every vector in $\mathbf{C}(A)$ is a constant multiple of \mathbf{a} . Then find $\mathbf{b} \in \mathbb{R}^2$ such that $\mathbf{b} \notin \mathbf{C}(A)$. (By the way, \notin just means

“is not an element of.”) [Hint: *what happens if both $\mathbf{e}_1, \mathbf{e}_2 \in \mathbf{C}(A)$?*]

(ii) Let $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}^3$ and $c_1, c_2 \in \mathbb{R}$, and

$$A = [\mathbf{a}_1 \quad \mathbf{a}_2 \quad (c_1\mathbf{a}_1 + c_2\mathbf{a}_2)].$$

First explain why every vector in $\mathbf{C}(A)$ can be written as $v_1\mathbf{a}_1 + v_2\mathbf{a}_2$ for some $v_1, v_2 \in \mathbb{R}$. If $\mathbf{C}(A) = \mathbb{R}^3$, does it feel weird that the “three-dimensional” space \mathbb{R}^3 can be described by varying only two parameters v_1 and v_2 ? Try to find $\mathbf{b} \in \mathbb{R}^3$ such that $\mathbf{b} \notin \mathbf{C}(A)$. I expect this to be annoying, since I’m not telling you what the entries of \mathbf{a}_1 and \mathbf{a}_2 are, but go as far as you can and see if you get stuck. It’s a +-problem, after all.

Day 6: Friday, January 17.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Span of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$, matrix with dependent columns (N), matrix with independent columns (N)

Here was the problem with the matrices in Examples 5.11 and 5.12: one of their columns was a linear combination of the other columns. Informally, from the point of view of the column structure of the matrices, there was redundant data. Somehow this prevented the column space from being as large as possible. Our job is to understand why.

First, now is a good time to review and augment our vocabulary. We want to understand $A\mathbf{x} = \mathbf{b}$ for $A \in \mathbb{R}^{m \times n}$, and to be able to solve this we want $\mathbf{b} \in \mathbf{C}(A)$, where

$$\mathbf{C}(A) := \{A\mathbf{v} \mid \mathbf{v} \in \mathbb{R}^n\}.$$

If $A = [\mathbf{a}_1 \quad \dots \quad \mathbf{a}_n]$ with $\mathbf{a}_j \in \mathbb{R}^m$, we can also write

$$\mathbf{C}(A) = \{c_1\mathbf{a}_1 + \dots + c_n\mathbf{a}_n \mid c_1, \dots, c_n \in \mathbb{R}\}.$$

That is, $\mathbf{C}(A)$ is the set of all linear combinations of the columns of A . We may want to consider such sets of linear combinations not strictly in the context of columns of a matrix.

6.1 Definition. The **SPAN** of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$ is the set of all linear combinations of these vectors, and we denote it by $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_n)$.

Content from Strang’s ILA 6E. The span of a list of vectors is defined in the box on p. 21.

6.2 Example. (i) Let $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^m$. Then $\mathbf{C}(\begin{bmatrix} \mathbf{a}_1 & \cdots & \mathbf{a}_n \end{bmatrix}) = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_n)$.
(ii) Let $\mathbf{v} \in \mathbb{R}^m$. Then $\text{span}(\mathbf{v}) = \{c\mathbf{v} \mid c \in \mathbb{R}\}$.

6.3 Problem (!). Prove that $\text{span}(\mathbf{0}) = \{\mathbf{0}\}$. That is, the only vector in the span of $\mathbf{0}$ is $\mathbf{0}$ itself.

6.4 Problem (*). Let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$. Prove that $\mathbf{0}_m \in \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_n)$.

6.5 Problem (*). Let $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^m$ and $c, c_1, c_2 \in \mathbb{R}$. Prove that

$$\text{span}(\mathbf{v}_1, c\mathbf{v}_1) = \text{span}(\mathbf{v}_1) \quad \text{and} \quad \text{span}(\mathbf{v}_1, \mathbf{v}_2, c_1\mathbf{v}_1 + c_2\mathbf{v}_2) = \text{span}(\mathbf{v}_1, \mathbf{v}_2).$$

Explain how these “small” spans illustrate the following general principle: the span of a list of vectors equals the span of those vectors in the list that are not linear combinations of other vectors in the list.

The problem with the matrices in Examples 5.11 and 5.12 was twofold: these were matrices in $\mathbb{R}^{m \times m}$ but their column spaces were not all of \mathbb{R}^m (so we could not always solve $A\mathbf{x} = \mathbf{b}$ for all \mathbf{b}), and one of their columns was in the span of the others. Somehow these problems are related. We first give a name to the latter situation and then make a conjecture.

6.6 Definition. The columns of a matrix $A \in \mathbb{R}^{m \times n}$ are **DEPENDENT** if (at least) one column is in the span of the others, i.e., if (at least) one column is a linear combination of the other columns. If $n = 1$ and the matrix only has one column, we say its column is dependent if it is the zero vector.

The inclusion of the special case of the zero vector when there is only one column (and when it does not make sense to talk about “span of the others,” because there are no “other” columns when $n = 1$) is a bit of a technicality that will be helpful later. For $n \geq 2$, here is the importance of quantifiers: all that it takes for a matrix to have dependent columns is for *one* column to be “bad.” And here is our conjecture.

6.7 Conjecture. If the columns of $A \in \mathbb{R}^{m \times m}$ are dependent, then $\mathbf{C}(A) \neq \mathbb{R}^m$.

This conjecture encapsulates the situation of the matrices in Examples 5.11 and 5.12. Unfortunately, it’s only a conjecture right now, and we don’t yet have the tools to prove it. And even when we know it’s true, we probably want a way of verifying that the columns of a matrix are dependent—hopefully a more systematic way than just “getting lucky” and noticing that one column is a linear combination of the others.

6.8 Problem (!). We can talk about a nonsquare matrix with dependent columns, but

the conjecture was only for a square matrix. Here's why. Let

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

Show that $\mathbf{C}(A) = \mathbb{R}^2$ and that the columns of A are dependent.

6.9 Problem (★). Show that the columns of

$$\begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$$

are dependent. [Hint: $\mathbf{v} = 1\mathbf{v}$.] Conclude that if a matrix contains the same column two or more times, then its columns are dependent.

We probably think that the opposite of “dependent” is “independent,” and so the columns of a matrix should be “independent” if no column is in the span of the others, i.e., if no column is a linear combination of the others. This could be hard to check! We'd have to fix our attention on one column at a time and then compare it to every other column. That could take forever. Here is a better definition of “independent,” although it requires a little more technical notation.

6.10 Definition. The columns of $A = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n] \in \mathbb{R}^{m \times n}$ are **INDEPENDENT** if $\mathbf{a}_1 \neq \mathbf{0}_m$ and if \mathbf{a}_j is not a linear combination of $\mathbf{a}_1, \dots, \mathbf{a}_{j-1}$ for $j = 2, \dots, n$ (if, indeed, $n \geq 2$). That is, $\mathbf{a}_1 \neq \mathbf{0}_m$ and $\mathbf{a}_j \notin \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_{j-1})$ for $j = 2, \dots, n$.

In the particular case that $n = 1$ and A has only one column, then we say that this column is independent if it is not the zero vector.

I think we better do a concrete example right away.

6.11 Example. Let

$$A = \begin{bmatrix} 1 & 2 & 4 \\ 0 & 3 & 5 \\ 0 & 0 & 6 \end{bmatrix}.$$

Certainly

$$\mathbf{a}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \neq \mathbf{0}_3.$$

Next, we want to check that $\mathbf{a}_2 \notin \text{span}(\mathbf{a}_1)$. Otherwise, we would have

$$\begin{bmatrix} 2 \\ 3 \\ 0 \end{bmatrix} = c \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

for some $c \in \mathbb{R}$. Equating the second components, this would mean $3 = 0$. Last, we want to check that $\mathbf{a}_3 \notin \text{span}(\mathbf{a}_1, \mathbf{a}_2)$. Otherwise, we would have

$$\begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix} = c_1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + c_2 \begin{bmatrix} 2 \\ 3 \\ 0 \end{bmatrix}.$$

Equating the third components, we would have $6 = 0$.

6.12 Problem (!). Where does the algorithm for independence in Definition 6.10 break down with the matrix

$$\begin{bmatrix} 1 & 1 & 2 & 3 \\ 1 & 0 & 2 & 4 \\ 0 & 0 & 0 & 5 \\ 0 & 0 & 0 & 0 \end{bmatrix}?$$

6.13 Remark. *Why do we exclude the case $\mathbf{a}_1 = \mathbf{0}_m$ so specifically in the definition of independent columns? Honestly, “because it’s the right thing to do.” This definition of independent columns is, in a moment, going to give us the correct analogue of Conjecture 6.7. (Well, “correct” once we prove it.)*

That’s not very satisfying right now, I realize. Perhaps the better reason to exclude $\mathbf{a}_1 = \mathbf{0}_m$ is because if any column of a matrix is the zero vector, then the columns are dependent: that zero column is the linear combination of all the other columns with weights equal to 0. Think about part (i) of Example 5.11.

And why do we not say that subsequent columns of the matrix beyond the first can’t be the zero vector? Because the condition $\mathbf{a}_j \notin \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_{j-1})$ excludes that: we know $\mathbf{0}_m \in \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_{j-1})$ by Problem 6.4, so if $\mathbf{a}_j \notin \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_{j-1})$, then it’s definitely the case that $\mathbf{a}_j \neq \mathbf{0}_m$.

Content from Strang’s ILA 6E. Reread p. 20, this time paying attention to dependence and independence. Then work through the (in)dependence tests on p. 21 for the matrices A_4 and A_5 . There is one thing here that we have not yet discussed: what does it mean for only “some” of the columns of A to be independent?

Now here is the analogue of Conjecture 6.7 for a matrix with independent columns.

6.14 Conjecture. *If the columns of $A \in \mathbb{R}^{m \times m}$ are independent, then $\mathbf{C}(A) = \mathbb{R}^m$.*

6.15 Problem (!). Again, the conjecture is only for square matrices. Explain why the columns of

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}$$

are independent but $\mathbf{C}(A) \neq \mathbb{R}^3$.

6.16 Problem (★). Here is another reason to prevent the first column from being the zero vector, beyond Remark 6.13. Suppose that at least one column of $A = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_m] \in \mathbb{R}^{m \times m}$ is the zero vector $\mathbf{0}_m$. Explain why $\text{span}(\mathbf{a}_1, \dots, \mathbf{a}_n)$ can be “described” by at most $m - 1$ different parameters and therefore $\mathbf{C}(A)$ is at most “ $(m - 1)$ -dimensional.”

As with Conjecture 6.7, we do not yet have the tools to prove Conjecture 6.14, nor do we have a systematic way of verifying that a matrix’s columns are independent. Both conjectures beg the question of which columns in a matrix really matter—which ones are redundant and which ones are essential for describing the column space. This will lead to a more general definition of dependence and independence that can be given beyond the context of matrices.

6.17 Problem (+). Let

$$A = [\mathbf{a}_1 \ \mathbf{a}_2 \ \mathbf{a}_3] = \begin{bmatrix} 1 & 0 & 3 \\ 0 & 2 & 4 \\ 0 & 0 & 0 \end{bmatrix}.$$

Of course, $\mathbf{C}(A) = \text{span}(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3)$. However, we can be more efficient. Show that we can write $\mathbf{C}(A)$ in the following ways:

$$\mathbf{C}(A) = \text{span}(\mathbf{a}_1, \mathbf{a}_2), \quad \mathbf{C}(A) = \text{span}(\mathbf{a}_2, \mathbf{a}_3), \quad \mathbf{C}(A) = \text{span}(\mathbf{a}_1, \mathbf{a}_3).$$

But we can’t beat that: explain why $\mathbf{C}(A)$ is not the span of any one column of A .

So, what columns really matter?

Content from Strang’s ILA 6E. Answer: the “independent ones,” as alluded to on pp. 20–22. This will require us to broaden the definition of independence to allow only *some* of the columns of the matrix to be independent—that is, some of the columns of a matrix with dependent columns can still be independent, if we define “independent” correctly. Now is also a good time to (re)read pp. v–vii up to, but not including, the $A = CR$ section.

Day 7: Wednesday, January 22.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Independent list of vectors (N), dependent list of vectors (N), rank of a matrix

To describe what columns really matter, we need a variation on our notions of dependence

and independence that free us from thinking solely about vectors as columns of matrices. Before we state that, we clarify our expectation that a matrix with dependent columns can't have independent columns.

7.1 Theorem. *Let $A \in \mathbb{R}^{m \times n}$. The columns of A are dependent if and only if they are not independent.*

Proof. This is an “if and only if” statement, so we need our logic to go in two directions. First suppose that the columns of A are dependent; write, as usual, $A = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n]$. We need to show that either $\mathbf{a}_1 = \mathbf{0}_m$ or $\mathbf{a}_j \in \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_{j-1})$ for some $j \geq 2$, if indeed $n \geq 2$. If $\mathbf{a}_1 = \mathbf{0}_m$, then we're done, so assume $\mathbf{a}_1 \neq \mathbf{0}_m$ (and, implicitly, $n \geq 2$). This could get messy in the abstract, so consider the very special case of $n = 4$ with \mathbf{a}_3 as a linear combination of \mathbf{a}_1 and \mathbf{a}_4 :

$$\mathbf{a}_3 = c_1 \mathbf{a}_1 + c_4 \mathbf{a}_4.$$

Our program is to look for the j th column as a linear combination of the previous $j - 1$ columns; it doesn't look like \mathbf{a}_3 is a linear combination of \mathbf{a}_1 and \mathbf{a}_2 here.

Or is it? Since $c_4 \in \mathbb{R}$, we have two options: $c_4 = 0$ or $c_4 \neq 0$. If $c_4 = 0$, then

$$\mathbf{a}_3 = c_1 \mathbf{a}_1 = c_1 \mathbf{a}_1 + 0 \mathbf{a}_2 \in \text{span}(\mathbf{a}_1, \mathbf{a}_2).$$

This violates the definition of independent columns with $j = 2$. If $c_4 \neq 0$, some algebra (which I will leave you to check, please) gives

$$\mathbf{a}_4 = \left(-\frac{c_1}{c_4}\right) \mathbf{a}_1 + 0 \mathbf{a}_2 + \left(\frac{1}{c_4}\right) \mathbf{a}_3 \in \text{span}(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3).$$

Now here is how this works in general (and I am not expecting you to read this unless you're curious). Say that the columns of $A \in \mathbb{R}^{m \times n}$ are dependent and that column ℓ is a linear combination of the other columns:

$$\mathbf{a}_\ell = \sum_{\substack{k=1 \\ k \neq \ell}}^n c_k \mathbf{a}_k.$$

Let j be the largest integer such that $c_j \neq 0$. If $j = \ell - 1$, then $\mathbf{a}_\ell \in \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_{\ell-1})$. Otherwise, assume $j \geq \ell + 1$ (since c_ℓ does not exist in the notation above). Then $c_k = 0$ for $k \geq j + 1$, so

$$\mathbf{a}_\ell = \sum_{\substack{k=1 \\ k \neq \ell}}^{j-1} c_k \mathbf{a}_k + c_j \mathbf{a}_j.$$

Since $c_j \neq 0$, this rearranges to

$$\mathbf{a}_j = \left(\frac{1}{c_j}\right) \mathbf{a}_\ell + \sum_{\substack{k=1 \\ k \neq \ell}}^{j-1} \left(-\frac{c_k}{c_j}\right) \mathbf{a}_k \in \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_{j-1}).$$

Here \mathbf{a}_ℓ is in the list $\mathbf{a}_1, \dots, \mathbf{a}_{j-1}$ since $\ell < \ell + 1 \leq j$. Glad I didn't do that in class?

Now for the other direction: if the columns of A are not independent, they are dependent. First suppose $\mathbf{a}_1 = \mathbf{0}_m$, so $\mathbf{a}_1 \in \text{span}(\mathbf{a}_2, \dots, \mathbf{a}_n)$ by Problem 6.4. Next suppose $\mathbf{a}_j \in \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_{j-1})$ for some $j \geq 2$. Then

$$\mathbf{a}_j = c_1 \mathbf{a}_1 + \dots + c_{j-1} \mathbf{a}_{j-1} = c_1 \mathbf{a}_1 + \dots + c_{j-1} \mathbf{a}_{j-1} + 0\mathbf{a}_{j+1} + \dots + 0\mathbf{a}_n,$$

and so \mathbf{a}_j is a linear combination of the other columns. ■

This should be fundamentally comforting: “dependent” should mean “not independent.” Now we free ourselves from talking about columns of matrices by saying exactly the same thing for lists of vectors.

7.2 Definition. A list of vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$ is **INDEPENDENT** if $\mathbf{v}_1 \neq \mathbf{0}_m$ and if $\mathbf{v}_j \notin \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_{j-1})$. The list is **DEPENDENT** if it is not independent. In particular, a list consisting of one vector is dependent if and only if that vector is the zero vector.

This vocabulary and the following procedure are the keys to talking *efficiently* about the column space of a matrix—to describing it with the minimal amount of data necessary. Say that we start with a list of vectors and consider its span (like we do with the columns of a matrix and the column space of that matrix). How can we remove all of the “unnecessary” vectors from that list so that we arrive at a “sublist” of just enough vectors whose span equals that of the original list? I hope an example helps. We are going to need the general principle suggested by Problem 6.5, so you should (re)read, and maybe (re)do, that problem now.

7.3 Example. Consider the following list of vectors in \mathbb{R}^4 :

$$\mathbf{v}_1 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{v}_3 = \begin{bmatrix} 2 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{v}_4 = \begin{bmatrix} 2 \\ 3 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{v}_5 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \quad \text{and} \quad \mathbf{v}_6 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 3 \end{bmatrix}.$$

The zero vector contributes nothing to the span:

$$\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_6) = \text{span}(\mathbf{v}_2, \dots, \mathbf{v}_6).$$

Since $\mathbf{v}_2 \neq \mathbf{0}_4$, we have $\text{span}(\mathbf{v}_1, \mathbf{v}_2) = \text{span}(\mathbf{v}_2)$, and also the list consisting of the single vector \mathbf{v}_2 is independent.

Next, $\mathbf{v}_3 = 2\mathbf{v}_2$, so $\text{span}(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3) = \text{span}(\mathbf{v}_2)$, and the list consisting of the single vector \mathbf{v}_2 is still independent. Nothing really new yet.

Onward: we have $\mathbf{v}_4 \notin \text{span}(\mathbf{v}_2)$. Why? So, the list $\mathbf{v}_2, \mathbf{v}_4$ is independent and $\text{span}(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4) = \text{span}(\mathbf{v}_2, \mathbf{v}_4)$.

I claim that $\mathbf{v}_5 \in \text{span}(\mathbf{v}_2, \mathbf{v}_4)$ and I will leave that for you to check. Thus $\text{span}(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4, \mathbf{v}_5) = \text{span}(\mathbf{v}_2, \mathbf{v}_4)$. Last, I claim that $\mathbf{v}_6 \notin \text{span}(\mathbf{v}_2, \mathbf{v}_4)$, and I'll also ask you to check that. Thus the list $\mathbf{v}_2, \mathbf{v}_4, \mathbf{v}_6$ is independent and $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_6) = \text{span}(\mathbf{v}_2, \mathbf{v}_4, \mathbf{v}_6)$.

In summary, the “sublist” $\mathbf{v}_2, \mathbf{v}_4, \mathbf{v}_6$ is independent and has the same span as the original list. This is efficient: we have cut the number of vectors needed to describe the span in half.

7.4 Problem (!). Check the things that I asked you to check in the previous example.

Here is the general result.

7.5 Lemma. *Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ be a list in \mathbb{R}^m with at least one nonzero vector in the list. There exists an independent sublist $\mathbf{v}_{j_1}, \dots, \mathbf{v}_{j_r}$ of $\mathbf{v}_1, \dots, \mathbf{v}_n$ such that $\text{span}(\mathbf{v}_{j_1}, \dots, \mathbf{v}_{j_r}) = \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_n)$.*

In Example 7.3 we had $r = 3$, $j_1 = 2$, $j_2 = 4$, and $j_3 = 6$.

Proof (of Lemma 7.5). We reduce the list as follows. Let \mathbf{v}_{j_1} be the first nonzero vector in the list. (At least one exists.) So $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_{j_1}) = \text{span}(\mathbf{v}_{j_1})$. Also, the list \mathbf{v}_{j_1} is independent because $\mathbf{v}_{j_1} \neq \mathbf{0}$.

Let \mathbf{v}_{j_2} be the first vector in the list that is a multiple of \mathbf{v}_{j_1} , i.e., $\mathbf{v}_{j_2} \notin \text{span}(\mathbf{v}_{j_1})$. So $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_{j_2}) = \text{span}(\mathbf{v}_{j_1}, \mathbf{v}_{j_2})$. Also, the list $\mathbf{v}_{j_1}, \mathbf{v}_{j_2}$ is independent because $\mathbf{v}_{j_1} \neq \mathbf{0}$ and $\mathbf{v}_{j_2} \notin \text{span}(\mathbf{v}_{j_1})$.

Let \mathbf{v}_{j_3} be the first vector in the list that is not in $\text{span}(\mathbf{v}_{j_1}, \mathbf{v}_{j_2})$. So $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_{j_3}) = \text{span}(\mathbf{v}_{j_1}, \mathbf{v}_{j_2}, \mathbf{v}_{j_3})$. And the list $\mathbf{v}_{j_1}, \mathbf{v}_{j_2}, \mathbf{v}_{j_3}$ is independent since $\mathbf{v}_{j_3} \notin \text{span}(\mathbf{v}_{j_1}, \mathbf{v}_{j_2})$.

Now turn the crank and keep going: eventually we run out of vectors in the list. ■

Content from Strang’s ILA 6E. Think once more about the matrices A_1 through A_5 on pp. 20–21. Apply the algorithm in the example and lemma above to extract the linearly independent columns that span the column spaces.

Just because we found one sublist that preserves the span doesn’t mean there isn’t another: reread, and maybe attempt, Problem 6.17 right now. But that problem suggests the following conjecture.

7.6 Conjecture. *Let $A \in \mathbb{R}^{m \times n}$. If $\mathbf{C}(A)$ is the span of r independent columns, then any list of r independent columns of A also spans the column space. No list of fewer than r columns can span the column space, and any list of more than r columns is dependent.*

We don’t have the tools to prove this conjecture yet, but it suggests that there is a “threshold” for independence and spans: a number of columns that is “just right” to span the column space efficiently without any redundancy. We give this number a name, even though we don’t know how to compute it yet.

7.7 Definition. *The **RANK** of a matrix $A \in \mathbb{R}^{m \times n}$ is the length of the longest list of linearly independent columns of A .*

7.8 Example. The rank of the matrix whose columns are the vectors from Example 7.3 is 3.

Actually, we do know how to compute rank if we start out knowing what it should be. Here are some of the “worst” matrices from the point of view of redundancy: they have far more columns than necessary to describe their column spaces. Actually, we only need one column to span the column space.

7.9 Example. (i) $C \left(\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \right) = \{\mathbf{0}_2\}$

(ii) $C \left(\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} \right) = \text{span} \left(\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right)$

(iii) $C \left(\begin{bmatrix} 1 & 2 & 3 \\ 1 & 2 & 3 \\ 1 & 2 & 3 \end{bmatrix} \right) = \text{span} \left(\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right)$ as each column is a multiple of the first.

What a waste of storage space!

Content from Strang's ILA 6E. Work through the example on p. 23 with the matrix A_6 . We won't talk about this for some time in class, but the “row rank = column rank” calculations for 2×2 and 3×3 rank-1 matrices are good practice, so check the details yourself.

Perhaps it would be nice if we had a more efficient way of representing such redundant matrices. Is there a way to extract only the columns that are absolutely necessary for representing the column space? (And will this help us solve, and understand, linear systems?)

The right approach is a new tool: matrix multiplication and matrix factorizations. Think about the factoring that you've already done in life before linear algebra. You've factored integers into products of powers of primes:

$$12 = 2^2(3).$$

And you've factored polynomials into simpler polynomials:

$$x^2 - 4x + 4 = (x - 2)^2.$$

Both kinds of factorizations reveal (potentially) useful information: what the essential components of an integer are, how to find zeros and maybe graph polynomials. If we know how to multiply matrices, perhaps we can factor them so that only the most important information comes out in the factorization.

For example, maybe we could have something like

$$\begin{bmatrix} 1 & 2 & 3 \\ 1 & 2 & 3 \\ 1 & 2 & 3 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} [1 \quad 2 \quad 3],$$

where the first matrix (okay, column vector) is the only column that matters, and the second matrix (okay, row vector) contains the data needed for constructing all of the columns out of this first column. This hinges on defining matrix products in such a way that the definition yields the equality above. How do we do it?

Content from Strang's ILA 6E. The real goal is to answer the questions posed at the end of p. 22. We'll get there.

Day 8: Friday, January 24.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked "N."

Matrix-matrix product (a.k.a. matrix multiplication)

As we often say, the goal of this course is to study the problem $A\mathbf{x} = \mathbf{b}$ for given $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. Ideally we could solve it and find $\mathbf{x} \in \mathbb{R}^n$ that makes this equation true. If we can't solve it, we should understand why—can we quantify our failure with further information about why \mathbf{b} doesn't work, or can we approximate the problem somehow so that we could solve a related version?

Along the way, we've picked up notation and language to manipulate this problem (linear combinations, spans, matrix-vector multiplication, dot products) and to develop alternate ways of phrasing it. In particular, we have $A\mathbf{x} = \mathbf{b}$ if and only if $\mathbf{b} \in \mathbf{C}(A)$, where $\mathbf{C}(A)$ is the column space of A : the set of all linear combinations of columns of A . Some of those columns may be redundant and contribute nothing new to the column space, so we are developing the language of dependent and independent vectors to ensure that we work with the minimal amount of data necessary to describe those \mathbf{b} for which the problem $A\mathbf{x} = \mathbf{b}$ makes sense.

Our next great leap forward will be a notion of multiplying two matrices, not just a matrix and a vector. We will, in time, reverse-engineer that multiplication to *factor* matrices to reveal further useful and meaningful data about matrices, and thus about our fundamental problem. In fact, matrix multiplication will give us an algorithm for solving $A\mathbf{x} = \mathbf{b}$, something we haven't really done yet!

So, what is a "good" definition of matrix multiplication? Starting small might help: let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{v} \in \mathbb{R}^n$. I know we've said that $\mathbb{R}^n \neq \mathbb{R}^{n \times 1}$, and we've basically never thought about $n \times 1$ matrices anyway. But any $B \in \mathbb{R}^{n \times 1}$ has the form $B = [\mathbf{b}]$ for some $\mathbf{b} \in \mathbb{R}^n$. Of course, usually we think of $\mathbf{b} \in \mathbb{R}^n$ and $[\mathbf{b}]$ as being the same object.

Let's break that pattern. For $A \in \mathbb{R}^{m \times n}$ and $\mathbf{v} \in \mathbb{R}^n$, we have $A\mathbf{v} \in \mathbb{R}^m$. Any $C \in \mathbb{R}^{m \times 1}$ has the form $C = [\mathbf{c}]$ for some $\mathbf{c} \in \mathbb{R}^m$. So, I think that we can think of matrix-vector multiplication $A\mathbf{v}$ as matrix-matrix multiplication $A[\mathbf{v}]$, and through that lens we have

$$A[\mathbf{v}] = [A\mathbf{v}].$$

What did we do? The matrix-matrix product $A[\mathbf{v}]$ is just the matrix whose only column is the vector $A\mathbf{v}$.

What if the second factor has more columns? Let $B = [\mathbf{b}_1 \ \cdots \ \mathbf{b}_p]$. If we want to compute AB and continue the pattern above, then we want to multiply each column of B by A . But if $A \in \mathbb{R}^{m \times n}$, then each column of B needs to be in \mathbb{R}^n so that we can do that multiplication. And the matrix-vector product A times column of B yields a vector in \mathbb{R}^m . We don't care how many columns of B there are, so p can be arbitrary. Thus $AB \in \mathbb{R}^{m \times p}$.

8.1 Definition. Let $A \in \mathbb{R}^{m \times n}$ and $B = [\mathbf{b}_1 \ \cdots \ \mathbf{b}_p] \in \mathbb{R}^{n \times p}$. The **MATRIX PRODUCT** AB is

$$AB := [A\mathbf{b}_1 \ \cdots \ A\mathbf{b}_p] \in \mathbb{R}^{m \times p}.$$

Content from Strang's ILA 6E. Matrix multiplication is defined in equation (1) on p. 27. Work through the examples on that page and p. 28, noting the appearance of the dot product.

Here is the first reason for this definition and the restriction on the sizes of A and B : we want this definition to return the usual definition of matrix-vector multiplication when B is a column vector. There are other good reasons (possibly better reasons). They're coming. For now, let's practice.

8.2 Example. (i) Let

$$A = \begin{bmatrix} 1 & 0 \\ -2 & 1 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 1 & 3 \\ 2 & 4 \end{bmatrix}.$$

Both $A, B \in \mathbb{R}^{2 \times 2}$, so the product AB is defined and $AB \in \mathbb{R}^{2 \times 2}$ as well.

We compute

$$A\mathbf{b}_1 = \begin{bmatrix} 1 & 0 \\ -2 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

and

$$A\mathbf{b}_2 = \begin{bmatrix} 1 & 0 \\ -2 & 1 \end{bmatrix} \begin{bmatrix} 3 \\ 4 \end{bmatrix} = \begin{bmatrix} 3 \\ -2 \end{bmatrix}.$$

Feel free to do this with the original definition of matrix-vector multiplication or dot products. Thus

$$AB = [A\mathbf{b}_1 \ A\mathbf{b}_2] = \begin{bmatrix} 1 & 3 \\ 0 & -2 \end{bmatrix}.$$

(ii) Let

$$A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}.$$

Since $A \in \mathbb{R}^{2 \times 2}$ and $B \in \mathbb{R}^{2 \times 3}$, the product AB is defined and $AB \in \mathbb{R}^{2 \times 3}$.

We compute

$$A\mathbf{b}_1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 4 \end{bmatrix} = \begin{bmatrix} 4 \\ 1 \end{bmatrix},$$

$$A\mathbf{b}_2 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 2 \\ 5 \end{bmatrix} = \begin{bmatrix} 5 \\ 2 \end{bmatrix},$$

and

$$A\mathbf{b}_3 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 3 \\ 6 \end{bmatrix} = \begin{bmatrix} 6 \\ 3 \end{bmatrix}.$$

Thus

$$AB = [A\mathbf{b}_1 \quad A\mathbf{b}_2 \quad A\mathbf{b}_3] = \begin{bmatrix} 4 & 5 & 6 \\ 1 & 2 & 3 \end{bmatrix}.$$

(iii) Let

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix}.$$

Since $A \in \mathbb{R}^{3 \times 3}$ and $B \in \mathbb{R}^{3 \times 2}$, the product AB is defined and $AB \in \mathbb{R}^{3 \times 2}$.

We compute

$$A\mathbf{b}_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = \begin{bmatrix} 1 \\ 4 \\ 9 \end{bmatrix}$$

and

$$A\mathbf{b}_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{bmatrix} \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix} = \begin{bmatrix} 4 \\ 10 \\ 18 \end{bmatrix}.$$

Thus

$$AB = [A\mathbf{b}_1 \quad A\mathbf{b}_2] = \begin{bmatrix} 1 & 4 \\ 4 & 10 \\ 9 & 18 \end{bmatrix}.$$

8.3 Problem (!). Describe in words the effects of computing the three products in the previous example. [Hint: for part (i), think about subtraction.] Compare your response to patterns that you observed in Problem 3.2.

Coming out of these examples is a nice fact that helps when computing “small” products AB by hand.

8.4 Theorem. Let $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times p}$. Then the (i, j) -entry of AB is the dot product of row i of A (considered as a column vector in \mathbb{R}^n) with column j of B .

Proof. We know what AB is at the level of columns: column j of AB is the matrix-vector product of A with column j of B . So the entry in row i of column j of AB is the dot product

of row i of A (considered as a column vector in \mathbb{R}^n) with column j of B . ■

8.5 Problem (★). Suppose that A and B are matrices such that the product AB is defined.

- (i) If a whole row of A is all 0, what do you know about AB ?
- (ii) If a whole column of B is all 0, what do you know about AB ?

Here is something less nice. We expect that the order in which we multiply real numbers doesn't matter: if $x, y \in \mathbb{R}$, then $xy = yx$. Not so for matrices.

8.6 Problem (★). (i) Explain why even if the matrix product AB is defined, the product BA may not be defined. What do you need to know about A and B for both products AB and BA to be defined?

(ii) Use the matrices A and B from part (i) of Example 8.2 to show that we may have $AB \neq BA$ even when these products are both defined.

Is this that big a deal? Is our definition of matrix multiplication wrong? Frankly, no. Example 8.2 and Problem 8.3 suggest that matrices fundamentally act not just on vectors but on other matrices: they are dynamic. We can live without commutativity of matrix multiplication but not without dynamic matrix action. Also, most actions aren't commutative—try putting your shoes on before your socks. Order matters.

Content from Strang's ILA 6E. Check the multiplication in equation (6) on p. 28 for further reinforcement that $AB \neq BA$ in general. Answer the question at the bottom of the page.

Day 9: Monday, January 27.

Here is a matrix product that may look a little weird at first glance.

9.1 Example. Let

$$A = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 1 & 0 & 2 & 3 \end{bmatrix}.$$

The columns of B are vectors in \mathbb{R}^1 , and of course we usually think of these as real numbers: $\mathbb{R} = \mathbb{R}^1 (= \mathbb{R}^{1 \times 1})$. But here it is helpful, if silly, to keep the vector point of view. That is, we think that

$$B = [\mathbf{b}_1 \quad \mathbf{b}_2 \quad \mathbf{b}_3 \quad \mathbf{b}_4], \quad \text{where} \quad \mathbf{b}_1 = [1], \quad \mathbf{b}_2 = [0], \quad \mathbf{b}_3 = [2], \quad \mathbf{b}_4 = [3].$$

Now we compute AB by multiplying A against the columns of B . Of course we are going

to have

$$AB = [Ab_1 \quad Ab_2 \quad Ab_3 \quad Ab_4],$$

so what are these matrix-vector products?

We start with

$$Ab_1 = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} [1] = 1 \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}.$$

The second equality is the definition of matrix-vector multiplication $A\mathbf{v}$ for $A \in \mathbb{R}^{n \times 1}$ and $\mathbf{v} \in \mathbb{R}^1$: it's a linear combination of one vector. Pretty silly, I know.

Let's do it again in gory detail:

$$Ab_2 = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} [0] = 0 \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

I think you'll agree that

$$Ab_3 = \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix} \quad \text{and} \quad Ab_4 = \begin{bmatrix} 3 \\ 6 \\ 9 \end{bmatrix}.$$

All together,

$$AB = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} [1 \quad 0 \quad 2 \quad 3] = \begin{bmatrix} 1 & 0 & 2 & 3 \\ 2 & 0 & 4 & 6 \\ 3 & 0 & 6 & 9 \end{bmatrix}.$$

Just look at that matrix: it's so redundant! Every column is a multiple of the first, and the entries of B tell us how to do that multiplication. Far better to keep the matrix factored as AB so that we can see the important data: the one column in A and the multipliers in B .

This pattern generalizes nicely: for any $\mathbf{a} \in \mathbb{R}^m$ and $c_1, \dots, c_{n-1} \in \mathbb{R}$

$$[\mathbf{a} \quad c_1\mathbf{a} \quad \cdots \quad c_{n-1}\mathbf{a}] = [\mathbf{a}] [1 \quad c_1 \quad \cdots \quad c_{n-1}]. \quad (9.1)$$

9.2 Problem (!). Stare at this equality until you believe it. Maybe write something, too.

Content from Strang's ILA 6E. Read and work through all of the calculations on pp. 29–30 under “Rank One Matrices and $A = CR$.”

The factorization (9.1) is the sort of factorization that we desire. It takes a matrix with a lot of redundant data and breaks it up into the important chunks. The first factor in (9.1) contains the only independent column of the product, and the second factor tells you how to build the other columns out of that one and only one necessary column. Can we get something like this for a matrix with more than one independent column and for which the “building” might be more complicated?

This is one of those times where it's helpful to know where you're going before you get there.

9.3 Example. We compute

$$\begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 3 \\ 0 & 1 & 2 \end{bmatrix}$$

and we do so by thinking very intentionally about columns this time (not dot products, please, even though that's faster):

$$\begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = 1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + 0 \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix},$$

$$\begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = 0 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + 1 \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix},$$

and

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 3 \\ 2 \end{bmatrix} = 3 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + 2 \begin{bmatrix} 0 \\ 2 \\ 0 \end{bmatrix} = \begin{bmatrix} 3 \\ 4 \\ 0 \end{bmatrix}.$$

Thus

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 3 \\ 0 & 1 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 3 \\ 0 & 2 & 4 \\ 0 & 0 & 0 \end{bmatrix}.$$

We've seen this matrix before. We know that, going from left to right, its first two columns are linearly independent, and its third column is a linear combination of the two. The product on the left makes that explicit: the first factor in the product contains the linearly independent columns, and the second factor tells you how to put those columns together.

The structure of the second factor is interesting. The first two columns of the identity matrix $I_2 \in \mathbb{R}^{2 \times 2}$ (Problem 5.4) are there, and I will block them off as follows:

$$\begin{bmatrix} 1 & 0 & 3 \\ 0 & 1 & 2 \end{bmatrix} = \left[\begin{array}{cc|c} 1 & 0 & 3 \\ 0 & 1 & 2 \end{array} \right].$$

The matrix on the right is an example of a **PARTITIONED MATRIX**. This is totally euphemistic and just a nice way of breaking matrices into “submatrices” so that you see more meaningful patterns in the data.

Let me put

$$I_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad F = \begin{bmatrix} 3 \\ 2 \end{bmatrix}.$$

Then

$$\begin{bmatrix} 1 & 0 & 3 \\ 0 & 1 & 2 \end{bmatrix} = [I_2 \quad F].$$

This is a **BLOCK MATRIX**: a matrix whose entries are other matrices. Again, totally euphemistic, just a convenient way of seeing the most important parts from a bird's-eye perspective.

The factorization above is the dream. We want to take a matrix $A \in \mathbb{R}^{m \times n}$ with r independent columns and no more than r independent columns—so the rank of A is r —and write it as a product $A = CR$. The matrix C should contain those r independent columns, so $C \in \mathbb{R}^{m \times r}$. We want the matrix $R \in \mathbb{R}^{r \times n}$ to tell us how to build the columns of A out of linear combinations of the columns of C . So C times the j th column of R should give us the j th column of A .

Content from Strang's ILA 6E. Read “ C Contains the First r Independent Columns of A ” on p. 30 and “Matrix Multiplication C times R on pp. 31–32. Check the calculations in Example 2, equation (10), equation (11), and the box on p. 32. Also jump ahead to Example 5 on pp. 34–35 (you don’t have to read about that “columns \times rows” way of multiplying matrices). For yet another example, go back to “Matrix Multiplication $A = CR$ on p. vii. You do not have to feel that you could see these CR -factorizations immediately; you should agree that the given matrix multiplication works out.

Unfortunately, we still do not have the tools to do prove that such a factorization exists or to develop an algorithm for computing it “reasonably” and effectively for all but the most trivial and obvious matrices. This is just like we don’t have efficient tools for checking dependence or independence of vectors. Let me at least state this factorization as a dream.

9.4 Conjecture. *Let $A \in \mathbb{R}^{m \times n}$ have rank r . Then there exist matrices $C \in \mathbb{R}^{m \times r}$ and $R \in \mathbb{R}^{r \times n}$ such that $A = CR$. In particular, the columns of C are r independent columns of A .*

Content from Strang's ILA 6E. If you’re curious, read pp. 32–33 to learn more about computing R . Feel free to skip that for now. We will revisit this in extensive detail in the future.

Ideally, we could write

$$R = \begin{bmatrix} I_r & F \end{bmatrix} \quad (9.2)$$

with r as the $r \times r$ identity matrix. If $r = n$ and all of the columns of A are independent, then we just have $R = I_n$, and there is no F block present, since no column is a linear combination of the others in this case. You should think of that F block being possibly fictitious.

9.5 Problem (!). If $A \in \mathbb{R}^{m \times n}$ can be written as $A = CR$ with $C \in \mathbb{R}^{m \times r}$ and R as in (9.2), and if $r < n$, what are the dimensions of that F ?

This is actually not quite the structure of R all the time. The snag can be that the first r columns of A may not be the independent ones, and maybe those r independent columns are “interspersed” throughout A .

9.6 Problem (!). Let $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}^m$ be independent and $c_1, c_2 \in \mathbb{R}$. Compute

$$\begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 \end{bmatrix} \begin{bmatrix} 1 & c_1 & 0 \\ 0 & c_2 & 1 \end{bmatrix}$$

and comment on the structure of the matrix that results (which columns are independent and why?).

Content from Strang's ILA 6E. This is what Strang means by the parenthetical remark “in correct order” on p. in the displayed equations after “ $A = CR$ becomes.”

To get around this, the reality is that we'll have to write

$$R = \begin{bmatrix} I_r & F \end{bmatrix} P$$

for a “permutation” matrix P that will reshuffle the columns appropriately. We'll get around to talking about permutation matrices later, but this will mean that A would have the factorization

$$A = C \begin{bmatrix} I_r & F \end{bmatrix} P.$$

Is that allowed? Can we multiply three matrices at once? Will it matter which two matrices we multiply first?

Nope!

9.7 Theorem. Let $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{n \times p}$, and $C \in \mathbb{R}^{p \times q}$. Then $(AB)C = A(BC)$.

This theorem says that the order in which you *group* matrices during multiplication doesn't matter: matrix multiplication is **ASSOCIATIVE**. Thus we just write ABC and eliminate the parentheses. The *order* still totally matters, and we should not expect $ABC = ACB$ or some nonsense like that.

Content from Strang's ILA 6E. Read “ AB times $C = A$ times BC ” on p. 29.

The proof of Theorem 9.7 is largely a thankless exercise in juggling parentheses, so I will leave that for you to suss out.

9.8 Problem (+). (i) Let $A \in \mathbb{R}^{m \times n}$, $B = [\mathbf{b}_1 \ \cdots \ \mathbf{b}_p] \in \mathbb{R}^{n \times p}$, and $\mathbf{v} = (v_1, \dots, v_p) \in \mathbb{R}^p$. Explain why each of the following four equalities is true:

$$\begin{aligned} A(B\mathbf{v}) &= A(v_1\mathbf{b}_1 + \cdots + v_p\mathbf{b}_p) \\ &= v_1A\mathbf{b}_1 + \cdots + v_pA\mathbf{b}_p \text{ [Hint: Problem 2.17]} \\ &= [A\mathbf{b}_1 \ \cdots \ A\mathbf{b}_p] \\ &= (AB)\mathbf{v}. \end{aligned}$$

(ii) Let $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{n \times p}$, and $C \in \mathbb{R}^{p \times q}$. Let $\mathbf{e}_1, \dots, \mathbf{e}_q$ be the standard basis vectors for \mathbb{R}^q (Problem 4.6). Explain why to prove that $(AB)C = A(BC)$, it suffices to show that

$$((AB)C)\mathbf{e}_j = (A(BC))\mathbf{e}_j.$$

(iii) Use only the fact that $(DE)\mathbf{v} = D(E\mathbf{v})$ for matrices D , E and vectors \mathbf{v} for which both sides of that equality are defined (as proved above), justify each equality below:

$$((AB)C)\mathbf{e}_j = (AB)(C\mathbf{e}_j) = A(B(C\mathbf{e}_j)) = A((BC)\mathbf{e}_j) = (A(BC))\mathbf{e}_j.$$

All this being said, we still have no idea of how to compute that “ CR -factorization” of a matrix unless we are really lucky and see the dependence relations among the columns from the get-go. There is quite a systematic way of doing that, and it is related to proving Conjectures 6.7 and 6.14, and to developing an explicit algorithm for solving $A\mathbf{x} = \mathbf{b}$ when we can actually solve it. That is, all of our dreams will come true through very related techniques.

Content from Strang’s ILA 6E. At this point we have learned all the matrix-vector mechanics that we need to actually solve linear systems (and to understand our failure when we can’t solve them). Just to be safe, read “Review of AB on p. 29 and make sure you have no doubts there. Then read “Thoughts on Chapter 1” on p. 38 for a summary of everything that we’ve done and a hint of what’s to come.

Let’s finally start solving linear systems. We’re going to take a break from matrix manipulations—very briefly—and look at three linear systems, each of which is in a very nice form, and which together illustrate the scope of possibilities for solution behavior to $A\mathbf{x} = \mathbf{b}$.

9.9 Example. (i) We consider

$$\begin{bmatrix} 1 & -2 \\ 0 & 8 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 \\ 8 \end{bmatrix}.$$

As a linear system, this reads

$$\begin{cases} x_1 - 2x_2 = 1 \\ 8x_2 = 8 \end{cases}$$

Look familiar? This was our very first problem!

Of course, we “back-solve” or “back-substitute” to get first $x_2 = 1$ and then $x_1 - 2 = 1$, so $x_1 = 3$. The problem has only one solution:

$$\mathbf{x} = \begin{bmatrix} 3 \\ 1 \end{bmatrix}.$$

(ii) We consider

$$\begin{bmatrix} 1 & -2 \\ 0 & 0 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 \\ 8 \end{bmatrix}.$$

Write it out, and don't laugh:

$$\begin{cases} x_1 - 2x_2 = 1 \\ 0 = 8. \end{cases}$$

Of course this system has no solution, because $0 \neq 8$.

(iii) We consider

$$\begin{bmatrix} 1 & -2 \\ 0 & 0 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

Write it out, keep laughing:

$$\begin{cases} x_1 - 2x_2 = 1 \\ 0 = 0. \end{cases}$$

There is really not much to do, since the second equation is both true and doesn't involve unknowns. There's not much more we can do with the first equation, since we don't know the value for x_2 .

Here is the right, if not obvious, thing to do: rewrite $x_1 = 1 + 2x_2$. This says that every choice of $x_2 \in \mathbb{R}$ gives x_1 via this formula. You can pick any x_2 that you want, so there are infinitely many solutions. At the level of vectors, we could write

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 + 2x_2 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \begin{bmatrix} 2x_2 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} + x_2 \begin{bmatrix} 2 \\ 1 \end{bmatrix}.$$

Every value of x_2 gives a different solution, and so this problem has infinitely many solutions.

Content from Strang's *ILA* 6E. Work through the three systems on p. 40, which have the same properties as the three above.

The three examples above are paradigmatic in the sense that a linear system has only one of three general solution “behaviors”: only one solution, no solution, or infinitely many solutions. This is actually very easy to prove using matrix notation—which is why we use that notation, to make our lives easier. But the other thing to take from this example is that the *structure* of the linear systems was very nice: all of the matrices were “upper-triangular” in the sense that their entries were 0 below the diagonal. This made back-solving/substituting very, very easy.

Day 10: Wednesday, January 29.

Content from Strang's *ILA* 6E. For a very broad overview of where we're going, read p. 39. It's okay if you don't understand everything on a first pass. Then read the first three paragraphs on p. 83.

We formalize the situations of Example 9.9.

10.1 Theorem. Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. Then one, and only one, of the following is true.

- (i) There exists a unique solution $\mathbf{x} \in \mathbb{R}^n$ to the problem $A\mathbf{x} = \mathbf{b}$. That is, we can solve the problem, and if $A\mathbf{x}_1 = \mathbf{b}$ and $A\mathbf{x}_2 = \mathbf{b}$ for some $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^n$, then $\mathbf{x}_1 = \mathbf{x}_2$.
- (ii) There is no solution to the problem $A\mathbf{x} = \mathbf{b}$. That is, $A\mathbf{x} \neq \mathbf{b}$ for every $\mathbf{x} \in \mathbb{R}^n$.
- (iii) There are infinitely many solutions to $A\mathbf{x} = \mathbf{b}$.

Proof. We want one, and only, one, of three possibilities to hold. One way for this to work out is to assume that the first two are false and then show why the third must be true. So, assume that $A\mathbf{x} = \mathbf{b}$ has a solution (so the second part is false) but this solution is not unique (so the first part is false). That is, there are $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^n$ such that $A\mathbf{x}_1 = \mathbf{b}$, $A\mathbf{x}_2 = \mathbf{b}$, and $\mathbf{x}_1 \neq \mathbf{x}_2$.

Our goal is to find infinitely many different $\mathbf{x} \in \mathbb{R}^n$ that satisfy $A\mathbf{x} = \mathbf{b}$. Here is the trick. Like most tricks in math, it may not be obvious at first glance, so you should reread this proof until it becomes obvious.

Put $\mathbf{z} := \mathbf{x}_1 - \mathbf{x}_2$. Then $\mathbf{z} \neq \mathbf{0}_n$, since $\mathbf{x}_1 \neq \mathbf{x}_2$. And

$$A\mathbf{z} = A(\mathbf{x}_1 - \mathbf{x}_2) = A\mathbf{x}_1 - A\mathbf{x}_2 = \mathbf{b} - \mathbf{b} = \mathbf{0}_m.$$

The second equality is the linearity of matrix-vector multiplication (Problem 2.17).

Now let $c \in \mathbb{R}$ be arbitrary and $\mathbf{x} = \mathbf{x}_1 + c\mathbf{z}$. Then

$$A\mathbf{x} = A(\mathbf{x}_1 + c\mathbf{z}) = A\mathbf{x}_1 + A(c\mathbf{z}) = A\mathbf{x}_1 + c(A\mathbf{z}) = \mathbf{b} + c\mathbf{0}_m = \mathbf{b} + \mathbf{0}_m = \mathbf{b}.$$

The second and third equalities are, again, the linearity of matrix-vector multiplication. Make sure you understand why all of the other equalities are true.

So why does this give infinitely many solutions? Maybe we should have put $\mathbf{x}_c := \mathbf{x}_1 + c\mathbf{z}$ instead to emphasize the dependence of \mathbf{x}_c on the parameter c . (I guess there might be conflicts of notation with $c = 1$ and $c = 2$?) The point is that each different $c \in \mathbb{R}$ generates a different $\mathbf{x}_1 + c\mathbf{z} \in \mathbb{R}^n$: you can, and should, check that if $c_1 \neq c_2$, then $\mathbf{x}_1 + c_1\mathbf{z} \neq \mathbf{x}_1 + c_2\mathbf{z}$. ■

10.2 Problem (★). (i) Let $A \in \mathbb{R}^{m \times n}$. Suppose that there is $\mathbf{z} \in \mathbb{R}^n$ such that $\mathbf{z} \neq \mathbf{0}_n$ and $A\mathbf{z} = \mathbf{0}_m$. Let $\mathbf{b} \in \mathbb{R}^m$. Prove that if the problem $A\mathbf{x} = \mathbf{b}$ has a solution, it is not unique.

(ii) Consider the other side of this: if the only solution to $A\mathbf{x} = \mathbf{0}_m$ is $\mathbf{x} = \mathbf{0}_n$, then solutions to $A\mathbf{x} = \mathbf{b}$ (if they exist) are unique. Here's why: if $A\mathbf{x}_1 = \mathbf{b}$ and $A\mathbf{x}_2 = \mathbf{b}$, what does $\mathbf{z} := \mathbf{x}_1 - \mathbf{x}_2$ solve? Why does that imply $\mathbf{x}_1 = \mathbf{x}_2$ and thus uniqueness?

(iii) By considering the vector $\mathbf{z} = (2, -1)$, explain how the previous part generalizes the situation in part (iii) of Example 9.9.

Content from Strang's ILA 6E. After you do the problem above, reread Example 3 on p. 40. The vector that Strang calls \mathbf{X} is what I call \mathbf{z} .

The time has come to systematically solve linear systems! We go all the way back to our very first example, in which we showed that

$$\begin{bmatrix} 1 & -2 \\ 3 & 2 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 \\ 11 \end{bmatrix} \iff \begin{bmatrix} 1 & -2 \\ 0 & 8 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 \\ 8 \end{bmatrix}.$$

The latter system was easy to solve with “back-substitution.”

What's up with the \iff ? Why are these two problems equivalent (in the sense that \mathbf{x} solves one of them precisely when it solves the other)? More abstractly, we started with $A \in \mathbb{R}^{m \times m}$ (for now A will be square), we wanted to solve $A\mathbf{x} = \mathbf{b}$, and we somehow converted or “reduced” the problem to $U\mathbf{x} = \mathbf{c}$, where U was upper-triangular. Then we back-substituted.

10.3 Definition. A matrix $U \in \mathbb{R}^{m \times m}$ is **UPPER-TRIANGULAR** if all of the entries of U below the diagonal are 0. That is, the (i, j) -entry of U is 0 when $i > j$.

10.4 Example. Each matrix below is upper-triangular:

$$\begin{bmatrix} 1 & -2 \\ 0 & 8 \end{bmatrix}, \quad \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad \text{and} \quad \begin{bmatrix} 1 & 2 & 3 \\ 0 & 4 & 5 \\ 0 & 0 & 6 \end{bmatrix}.$$

Content from Strang's ILA 6E. For a longer example of why upper-triangular matrices are nice for back-substitution, read p. 41 through the “Special note” in the box. I expect that you are comfortable with this back-substitution method for solving linear systems, and I will not do examples with it here.

How do you do this? How do you “convert” $A \in \mathbb{R}^{m \times m}$ into an upper-triangular matrix U so that we have the equivalence of the problems

$$A\mathbf{x} = \mathbf{b} \quad \text{and} \quad U\mathbf{x} = \mathbf{c}$$

for some appropriate \mathbf{c} ? The point is that the arrows go *both* ways: $A\mathbf{x} = \mathbf{b} \implies U\mathbf{x} = \mathbf{c}$ and $U\mathbf{x} = \mathbf{c} \implies A\mathbf{x} = \mathbf{b}$. Having an arrow go one way in math doesn't always mean it goes the other way.

The good news is that we already know how to do this. It's all contained in the manipulations that we did on our very first problem at the level of equations and variables. The big idea was subtracting a multiple of one equation from another. We can do all of this at the level of matrices (and cut out the variables) by subtracting a multiple of one row of a matrix from another.

Specifically, to turn

$$A = \begin{bmatrix} 1 & -2 \\ 3 & 2 \end{bmatrix}$$

into

$$U = \begin{bmatrix} 1 & -2 \\ 0 & 8 \end{bmatrix},$$

we want to subtract 3 times the first row of A from the second row of A . The revolution of linear algebra is that we can encode this via matrix multiplication. Whenever we want to “do something” in this class, you should ask yourself how we can accomplish this by multiplying by a suitable matrix.

So, what matrix E satisfies

$$EA = U?$$

At the very least we need $E \in \mathbb{R}^{m \times 2}$ since $A \in \mathbb{R}^{2 \times 2}$. And we really want $m = 2$ since $EA = U \in \mathbb{R}^{2 \times 2}$. So, $E \in \mathbb{R}^{2 \times 2}$.

Here is where it is wise to think about matrix multiplication as E times the columns of A . What is E doing to each column? We want

$$E \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} v_1 \\ v_2 - 2v_1 \end{bmatrix}. \quad (10.1)$$

How can we view the vector on the right as a linear combination with weights given by v_1 and v_2 ? The vectors in that linear combination will be the columns of E .

So, work backwards:

$$\begin{bmatrix} v_1 \\ v_2 - 3v_1 \end{bmatrix} = \begin{bmatrix} v_1 \\ -3v_1 \end{bmatrix} + \begin{bmatrix} 0 \\ v_2 \end{bmatrix} = v_1 \begin{bmatrix} 1 \\ -3 \end{bmatrix} + v_2 \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

If we put

$$E := \begin{bmatrix} 1 & 0 \\ -3 & 1 \end{bmatrix},$$

then we have the desired equality (10.1).

10.5 Problem (!). Check that. Then compute $EA = U$ with A and U as above.

Here is how we’re thinking. Assume $A\mathbf{x} = \mathbf{b}$ with $\mathbf{b} = (1, 11)$. Then $EA\mathbf{x} = E\mathbf{b}$. Compute $EA = U$ with U as above and $E\mathbf{b} = (1, 8) =: \mathbf{c}$. Then solve $U\mathbf{x} = \mathbf{c}$. That *should* give a solution to the original problem $A\mathbf{x} = \mathbf{b}$, and we can always plug it in and check that it does.

Going in reverse requires a little more thought. Why does solving $EA\mathbf{x} = E\mathbf{b}$ give a solution to $A\mathbf{x} = \mathbf{b}$? It would be nice if we could “cancel” the factor of E from both sides. We can, and that’s called inverting a matrix, and we’ll do that nice and abstractly soon.

10.6 Problem (★). Put

$$F := \begin{bmatrix} 1 & 0 \\ 3 & 1 \end{bmatrix}.$$

First explain in words the effect of multiplying $F\mathbf{w}$ for some $\mathbf{w} \in \mathbb{R}^2$. Then check that $FE\mathbf{v} = \mathbf{v}$ for all $\mathbf{v} \in \mathbb{R}^2$. Finally, suppose that $EA\mathbf{x} = E\mathbf{b}$, multiply both sides by F , and

explain why $A\mathbf{x} = \mathbf{b}$.

It feels like we're doing "elimination" twice: we multiplied EA and then $E\mathbf{b}$ separately. We can combine all of the data of our problem into one "augmented" matrix: put

$$[A \quad \mathbf{b}] = \left[\begin{array}{cc|c} 1 & -2 & 1 \\ 3 & 2 & 11 \end{array} \right].$$

I like to draw a line separating the \mathbf{b} column when I'm working with actual numbers. Then do one matrix multiplication:

$$E[A \quad \mathbf{b}] = [EA \quad E\mathbf{b}] = \left[\begin{array}{cc|c} 1 & -2 & 1 \\ 0 & 8 & 8 \end{array} \right] = [U \quad \mathbf{c}].$$

From here, solve $U\mathbf{x} = \mathbf{c}$ by back-substitution.

I'm going to tell you the path forward, even if it isn't obvious right now. Here is the cartoon for $A \in \mathbb{R}^{3 \times 3}$. We want to turn A into an upper-triangular matrix U by multiplying A by the "right" matrices. In the "ideal" case, at the level of rows, we are going to subtract multiples of row 1 to create 0 entries in rows 2 and below of column 1. Specifically, the multiples will be based on the $(1, 1)$ -entry, which for now we hope is nonzero.

So we have the conversion

$$\begin{bmatrix} \textcircled{*} & * & * \\ * & * & * \\ * & * & * \end{bmatrix} \rightarrow \begin{bmatrix} * & * & * \\ 0 & * & * \\ 0 & * & * \end{bmatrix}.$$

I've written the changed entries in blue. Now subtract a multiple of the second row from the third row to create zeros in the second column below the second row. Again, in the "ideal" case, the multiple will be based on the $(2, 2)$ -entry, which we should hope is nonzero:

$$\begin{bmatrix} * & * & * \\ 0 & * & * \\ 0 & * & * \end{bmatrix} \rightarrow \begin{bmatrix} * & * & * \\ 0 & \textcircled{*} & * \\ 0 & 0 & * \end{bmatrix}.$$

Again, the blue entries are new or changed. Because both the second and third rows had 0 in their first column, subtracting a multiple of the second row from the third row did not destroy that 0 in the first column of the third row. This is the nice upper-triangular structure that is ideal for back-solving.

How do we accomplish this multiplication? I am going to tell you the answer, which generalizes all our work with E above. Let $A \in \mathbb{R}^{m \times n}$ and $\ell \in \mathbb{R}$. To subtract ℓ times row j of A from row i of A (with $i \neq j$), multiply A by the **ELIMINATION MATRIX** $E_{ij} \in \mathbb{R}^{m \times m}$ whose entries are 1 on the diagonal, $-\ell$ in the (i, j) -position, and 0 elsewhere. So, E_{ij} is "almost" the identity matrix, except for the (i, j) -entry.

10.7 Problem (!). Prove that this formula for E_{ij} works by computing the following very

special case and explaining the effect in words:

$$E_{21}\mathbf{v}, \quad \text{where} \quad E_{21} := \begin{bmatrix} 1 & 0 & 0 \\ -\ell & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{v} := \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix}.$$

Then spend at least five minutes thinking about how using dot products could help you prove the more general result stated in the paragraph above this problem.

We do an example in glacially slow detail.

10.8 Example. Let

$$A = \begin{bmatrix} 2 & 1 & 1 \\ 4 & 3 & 3 \\ 8 & 7 & 9 \end{bmatrix}.$$

We want to multiply A by “elimination” matrices like the 2×2 situation above so that 0 appears in the second and third rows of the first column. To get 0 in the $(2,1)$ -entry, we should subtract 2 times the first row from the second. The matrix

$$E_{21} := \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

accomplishes this, and here is what we get:

$$E_{21}A = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 2 & 1 & 1 \\ 4 & 3 & 3 \\ 8 & 7 & 9 \end{bmatrix} = \begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 8 & 7 & 9 \end{bmatrix}.$$

I’ll use the idiosyncratic notation

$$\begin{bmatrix} 2 & 1 & 1 \\ 4 & 3 & 3 \\ 8 & 7 & 9 \end{bmatrix} \xrightarrow[E_{21}]{\text{R2} \mapsto \text{R2} - 2 \times \text{R1}} \begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 8 & 7 & 9 \end{bmatrix}, \quad E_{21} := \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

to represent this. Saying $\text{R2} \mapsto \text{R2} - 2 \times \text{R1}$ means that row 2 is replaced by row 2 minus 2 times row 1.

Now we want to clear out the $(3,1)$ -entry, and we can do this by subtracting 4 times row 1 from row 3. So, we multiply

$$\begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 8 & 7 & 9 \end{bmatrix} \xrightarrow[E_{31}]{\text{R3} \mapsto \text{R3} - 4 \times \text{R1}} \begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 3 & 5 \end{bmatrix}, \quad E_{31} := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -4 & 0 & 1 \end{bmatrix}.$$

Finally, we want to clear out the 3 in the $(3,2)$ -entry:

$$\begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 3 & 5 \end{bmatrix} \xrightarrow[E_{32}]{\text{R3} \mapsto \text{R3} - 3 \times \text{R2}} \begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 2 \end{bmatrix}, \quad E_{32} := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -3 & 1 \end{bmatrix}.$$

We're done! Let's abbreviate $E = E_{32}E_{31}E_{21}$. The product

$$EA = \begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 2 \end{bmatrix} =: U$$

is upper-triangular. If we wanted to solve $A\mathbf{x} = \mathbf{b}$ for some $\mathbf{b} \in \mathbb{R}^3$, it would suffice to solve $U\mathbf{x} = E\mathbf{b}$ instead.

Content from Strang's ILA 6E. Read and work through everything on p. 42 right now. This is hugely important. Then read p. 45 up to and including equation (7). This is another example of elimination. Last, read all of p. 49 (but don't worry about inverses for now).

We are going to focus on “reducing” A to an upper-triangular form, and I am going to leave practicing with back-substitution to you. It's mostly just a longer version of part (i) of Example 9.9.

10.9 Problem (!). Use the results (and the notation) of Example 10.8 to solve $A\mathbf{x} = \mathbf{b}$, where $\mathbf{b} = (0, 1, 5)$.

10.10 Problem (★). We prefer upper-triangular matrices, in part for consistency, but “lower-triangular” matrices can be equally nice. Solve

$$\begin{bmatrix} 2 & 0 & 0 \\ 1 & 1 & 0 \\ 2 & 1 & 1 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix}.$$

10.11 Problem (+). We usually expect that matrix multiplication is not commutative. However, sometimes it is.

(i) Let $\ell_1, \ell_2 \in \mathbb{R}$ and put

$$E_{21} := \begin{bmatrix} 1 & 0 & 0 \\ -\ell_1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad E_{31} := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -\ell_2 & 0 & 1 \end{bmatrix}.$$

Explain in words what E_{21} and E_{31} “do” (i.e., what is the effect of multiplying $E_{21}\mathbf{v}$ and $E_{31}\mathbf{v}$ for some $\mathbf{v} \in \mathbb{R}^3$?). Then explain why you think this means that $E_{21}E_{31} = E_{31}E_{21}$. Do the actual matrix multiplication to convince yourself that this is true.

(ii) Let $\ell_3 \in \mathbb{R}$ and

$$E_{32} := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -\ell_3 & 1 \end{bmatrix}.$$

Without doing any calculations, explain why you should expect E_{31} and E_{32} *not* to commute. Then do the multiplication to check $E_{31}E_{32} \neq E_{32}E_{31}$.

Day 11: Friday, January 31.

For larger matrices, the pattern of elimination is the same. Use the $(1, 1)$ -entry to “create zeros” in rows 2 and below of column 1 by subtracting appropriate multiples of row 1 from those lower rows. Then use the *new* $(2, 2)$ -entry to “create zeros” in the *new* rows 3 and below of the *new* column 2 by subtracting appropriate multiples of the *new* row 2 from those *new* lower rows. Keep going until you’ve reached the last row and the matrix has been “reduced” to an upper-triangular structure.

This approach to elimination can break down in two ways. The first is not so bad and just requires a new kind of matrix to correct things. The second is worse and will prevent us from solving the linear system.

11.1 Example. What if at the j th step of elimination, the (j, j) -entry is 0, but an entry further down in column j is not 0? All hope is not lost. Consider

$$\begin{bmatrix} 2 & 2 & 1 \\ 4 & 4 & 3 \\ 8 & 9 & 9 \end{bmatrix} \xrightarrow[E_{21}]{R2 \mapsto R2 - 2 \times R1} \begin{bmatrix} 2 & 2 & 1 \\ 0 & 0 & 1 \\ 8 & 9 & 9 \end{bmatrix} \xrightarrow[E_{31}]{R3 \mapsto R3 - 4 \times R1} \begin{bmatrix} 2 & 2 & 1 \\ 0 & 0 & 1 \\ 0 & 1 & 5 \end{bmatrix}.$$

The matrices E_{21} and E_{31} are the same as before in Example 10.8, so I didn’t write them out again.

The problem is that the $(2, 2)$ -entry is now 0. We can’t use that to eliminate the 3 in the $(3, 2)$ -entry. But if we could “flip” rows 2 and 3, we’d be done. (This is totally legitimate: you can interchange the order of equations in a system of equations and not change the solution structure at all.) If only there were a matrix $P \in \mathbb{R}^{3 \times 3}$ such that

$$P \begin{bmatrix} 2 & 2 & 1 \\ 0 & 0 & 1 \\ 0 & 1 & 5 \end{bmatrix} = \begin{bmatrix} 2 & 2 & 1 \\ 0 & 1 & 5 \\ 0 & 0 & 1 \end{bmatrix}.$$

What we really want is that

$$P \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \begin{bmatrix} v_1 \\ v_3 \\ v_2 \end{bmatrix}.$$

We can get P by working backwards and thinking of matrix-vector multiplication as a linear combination:

$$\begin{bmatrix} v_1 \\ v_3 \\ v_2 \end{bmatrix} = \begin{bmatrix} v_1 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ v_3 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ v_2 \end{bmatrix} = v_1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + v_3 \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} + v_2 \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix}.$$

Are you okay with how I got the last equality? Maybe it would help to rearrange the sum so that v_1 , v_2 , and v_3 come in order:

$$v_1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + v_3 \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} + v_2 \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = v_1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + v_2 \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} + v_3 \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}.$$

Here is the result:

$$\begin{bmatrix} 2 & 2 & 1 \\ 0 & 0 & 1 \\ 0 & 1 & 5 \end{bmatrix} \xrightarrow[\substack{P_{23}}]{R3 \mapsto R2, R2 \mapsto R3} \begin{bmatrix} 2 & 2 & 1 \\ 0 & 1 & 5 \\ 0 & 0 & 1 \end{bmatrix}, \quad P_{23} := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}.$$

I am calling this P_{23} to emphasize that we get it by interchanging columns 2 and 3 of the identity matrix. We'll call such a matrix formed by swapping columns of the identity a **PERMUTATION MATRIX**.

What we get is that

$$EA = U, \quad E := P_{23}E_{31}E_{21}$$

with U upper-triangular. The matrix E is now a little more complicated than in Example 10.8, as we have to include a factor of a permutation matrix, not just an elimination matrix.

In general, to interchange rows i and j of $A \in \mathbb{R}^{m \times m}$, multiply $P_{ij}A$, where $P_{ij} \in \mathbb{R}^{m \times m}$ is the matrix whose columns are those of the $m \times m$ identity matrix with columns i and j interchanged. Such a matrix P_{ij} is, again, a **PERMUTATION MATRIX**. So, if at some stage of elimination, the diagonal entry that you want to use to eliminate entries below is 0, but other entries in that column are nonzero, just “permute” the rows to bring that nonzero entry up to the row that you want. Then eliminate as usual in the remaining rows.

Content from Strang's ILA 6E. Read “Possible breakdown of elimination” on p. 43 up to but not including the “Caution!” paragraph. Then read p. 45 after equation (1) and look at the calculation in “Exchange rows 2 and 3.” These P_{ij} permutation matrices are special cases of a more general permutation matrix structure, which is the identity matrix with its columns (equivalently, rows) rearranged in various ways. See pp. 64–65. We won't need those more general permutation matrices for a while.

11.2 Problem (!). Explain in words (no need for any calculations) why $P_{ij}A = P_{ji}A$.

11.3 Problem (★). Let $P_{13} \in \mathbb{R}^{3 \times 3}$ be the permutation matrix that interchanges columns 1 and 3 of the 3×3 identity matrix. Compute $P_{13}A$ and AP_{13} for an arbitrary $A \in \mathbb{R}^{3 \times 3}$. Then conjecture about what the different effects of multiplying $P_{ij}A$ and AP_{ij} are for an arbitrary $A \in \mathbb{R}^{m \times m}$ and an arbitrary permutation matrix $P_{ij} \in \mathbb{R}^{m \times m}$ that interchanges columns i and j of the $m \times m$ identity matrix. (You do not have to prove your conjecture.)

11.4 Problem (+). Let $A \in \mathbb{R}^{m \times n}$ and let $S \in \mathbb{R}^{n \times d}$ be a matrix whose columns are some of the columns of the $n \times n$ identity matrix. Here $d \geq 1$ is any integer, and the columns of the identity may be repeated, and some columns of the identity may not appear at all. Describe in words the structure of the matrix AS . [Hint: *the letter S might stand for “selection” matrix—what is being “selected” here?*]

Here is the nastier breakdown of elimination: what if at some step, the diagonal entry that you want to use to eliminate entries below is 0 and *all* other entries in that column are 0, too? Good news is that you don’t have to do any more elimination on entries in that column, as they’re already 0. Bad news is that you won’t be able to solve $A\mathbf{x} = \mathbf{b}$ for all \mathbf{b} . Here’s a particular example of why.

11.5 Example. Here is a problematic matrix:

$$A = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \\ 0 & 0 & 5 \end{bmatrix}.$$

We eliminate:

$$\begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \\ 0 & 0 & 5 \end{bmatrix} \xrightarrow[E_{21}]{R2 \mapsto R2 - 2 \times R1} \begin{bmatrix} 1 & 2 & 3 \\ 0 & 0 & 0 \\ 0 & 0 & 5 \end{bmatrix}, \quad E_{21} := \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Maybe it doesn’t look so problematic right now. We would want to use the $(2, 2)$ -entry in $E_{21}A$ to eliminate the $(3, 2)$ -entry, but the $(3, 2)$ -entry is already 0. So, $E_{21}A$ is already upper-triangular! Why is this not enough for us to be happy?

Let’s actually try to solve $A\mathbf{x} = \mathbf{b}$ for $\mathbf{b} = (b_1, b_2, b_3) \in \mathbb{R}^3$ arbitrary. If $A\mathbf{x} = \mathbf{b}$, then $E_{21}A\mathbf{x} = E_{21}\mathbf{b} = (b_1, b_2 - 2b_1, b_3)$. Thus we want

$$\begin{bmatrix} 1 & 2 & 3 \\ 0 & 0 & 0 \\ 0 & 0 & 5 \end{bmatrix} \mathbf{x} = \begin{bmatrix} b_1 \\ b_2 - 2b_1 \\ b_3 \end{bmatrix}.$$

At the level of actual equations, this is

$$\begin{cases} x_1 + 2x_2 + 3x_3 = b_1 \\ 0 = b_2 - 2b_1 \\ 5x_3 = b_3. \end{cases}$$

Look at that second equation: it says $b_2 - 2b_1 = 0$, equivalently, $b_2 = 2b_1$. Think about the logic here. We assumed that $A\mathbf{x} = \mathbf{b}$ with $\mathbf{b} = (b_1, b_2, b_3)$, and we deduced that $b_2 = 2b_1$. This means that \mathbf{b} cannot be just any vector in \mathbb{R}^3 ; it has to satisfy this “solvability condition” of $b_2 = 2b_1$. Surely not every vector in \mathbb{R}^3 does this—for example, take $\mathbf{b} = (1, 0, 0)$. So we can’t always solve $A\mathbf{x} = \mathbf{b}$.

It's worth interpreting this in the context of the column space. Look at the structure of A : the second row is twice the first row. More precisely,

$$A\mathbf{x} = \begin{bmatrix} x_1 + 2x_2 + 3x_3 \\ 2x_1 + 4x_2 + 6x_3 \\ 5x_3 \end{bmatrix} = \begin{bmatrix} x_1 + 2x_2 + 3x_3 \\ 2(x_1 + 2x_2 + 3x_3) \\ 5x_3 \end{bmatrix}.$$

So, if $\mathbf{b} = (b_1, b_2, b_3) \in \mathbf{C}(A)$, then $b_2 = 2b_1$. This is exactly the solvability condition that we deduced from elimination.

11.6 Problem (*). Does the “arrow go the other way”? We have shown

$$\mathbf{b} \in \mathbf{C}(A) \implies b_2 = 2b_1.$$

Do we have

$$b_2 = 2b_1 \implies \mathbf{b} \in \mathbf{C}(A)?$$

Yes! If $b_2 = 2b_1$, then $A\mathbf{x} = \mathbf{b}$ is the system

$$\begin{cases} x_1 + 2x_2 + 3x_3 = b_1 \\ 2x_1 + 4x_2 + 6x_3 = 2b_1 \\ 5x_3 = b_3. \end{cases}$$

Use the third equation to solve for x_3 , take x_2 to be any number that you like, and then use the first equation to write x_1 in terms of the values forced on x_3 and chosen for x_2 . Why does this also satisfy the second equation automatically?

Content from Strang's ILA 6E. Read the rest of “Possible Breakdown of Elimination” on p. 43 starting with “Caution!”

Day 12: Monday, February 3.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Invertible matrix (N), inverse of a matrix

These results will follow and support us for the rest of the course and beyond. Here is an abstraction of our elimination procedure.

12.1 Theorem (Gaussian elimination). Let $A \in \mathbb{R}^{m \times m}$. Then there exist matrices E , $U \in \mathbb{R}^{m \times m}$ with the following properties.

- (i) $EA = U$.
- (ii) U is upper-triangular.
- (iii) E is the product of elimination matrices E_{ij} and/or permutation matrices P_{ij} .

Proof. If the $(1, 1)$ -entry of A is nonzero, multiply A by elimination matrices E_{21}, \dots, E_{m1} to subtract multiples of row 1 of A from rows 2 through m of A . Call the product of these elimination matrices E_1 . If $m = 2$, then E_1A is upper-triangular. If $m \geq 3$ and the $(2, 2)$ -entry of E_1A is nonzero, multiply E_1A by elimination matrices E_{32}, \dots, E_{m2} to subtract multiples of row 2 of E_1A from rows 3 through m of E_1A . Call the product of these elimination matrices E_2 . If $m = 3$, then E_2E_1A is upper-triangular. Otherwise, turn the crank and keep going.

If at any stage the (j, j) -entry is zero and the entries in column j in rows $j + 1$ through m are zero, just proceed to the next step and consider the $(j + 1, j + 1)$ -entry. If the (j, j) -entry is zero and some entry in rows $j + 1$ through m of column j is nonzero, multiply by a permutation matrix so that this nonzero entry is now the (j, j) -entry. Then eliminate as before. Call the product of the elimination matrices and the permutation matrices E_j . ■

What this result says is that if $A\mathbf{x} = \mathbf{b}$, then $EA\mathbf{x} = E\mathbf{b}$, and so $U\mathbf{x} = E\mathbf{b}$. The upper-triangular system $U\mathbf{x} = E\mathbf{b}$ is much easier to solve, and so we like it. At least, we like it when the diagonal entries of U are nonzero.

12.2 Theorem. Let $U \in \mathbb{R}^{m \times m}$ be an upper-triangular matrix whose diagonal entries are nonzero. Then for any $\mathbf{c} \in \mathbb{R}^m$, there exists a unique $\mathbf{x} \in \mathbb{R}^m$ such that $U\mathbf{x} = \mathbf{c}$.

Proof. This is really back-substitution in the abstract. Here's the proof for $m = 3$. Take

$$U = \begin{bmatrix} u_{11} & * & * \\ 0 & u_{22} & * \\ 0 & 0 & u_{33} \end{bmatrix},$$

where u_{11} , u_{22} , and u_{33} are nonzero. So if you want to solve $U\mathbf{x} = \mathbf{c}$ with $\mathbf{c} = (c_1, c_2, c_3)$, first you'd look at

$$u_{33}x_3 = c_3.$$

Since $u_{33} \neq 0$, we can divide to find that x_3 must be

$$x_3 = \frac{c_3}{u_{33}}.$$

Go back up a step and look at

$$u_{22}x_2 + \text{stuff depending on } x_3 = c_2.$$

The point is that we know what this “stuff” is because we know x_3 exactly. Solve this as

$$x_2 = \frac{c_2 - \text{stuff}}{u_{22}}.$$

This is the only choice for x_2 . Do the same for x_1 . ■

But are we really sure that if $EA = U$, then a solution to $U\mathbf{x} = E\mathbf{b}$ is also a solution to $A\mathbf{x} = \mathbf{b}$? For small problems, we can check it by plug-and-chug, but why is this true in general?

The time has come to be sure that we can “invert” E , and this is a good reason to study matrix inverses in general. We will overall be much more concerned with properties of inverses than formulas for inverses. There’s an algorithm that will let you do that, and we’ll see it briefly, but we’ll mostly abide by the slogan “What things do defines what things are.”

Content from Strang’s ILA 6E. Read the first two paragraphs on p. 50.

Here is what we want: why does $E\mathbf{A}\mathbf{x} = E\mathbf{b}$ imply $\mathbf{A}\mathbf{x} = \mathbf{b}$? More abstractly, if $E \in \mathbb{R}^{m \times m}$ and $E\mathbf{v} = E\mathbf{b}$ for some $\mathbf{v}, \mathbf{b} \in \mathbb{R}^m$, do we necessarily have $\mathbf{v} = \mathbf{b}$? It would be nice if we could “undo” the “action” of E by multiplying by another matrix. Is there $F \in \mathbb{R}^{m \times m}$ such that $F(E\mathbf{v}) = \mathbf{v}$ for all $\mathbf{v} \in \mathbb{R}^m$? If so, then assuming $E\mathbf{v} = E\mathbf{b}$ gives $F(E\mathbf{v}) = F(E\mathbf{b})$, and thus $\mathbf{v} = \mathbf{b}$ as desired.

Look more closely at the equation $F(E\mathbf{v}) = \mathbf{v}$. This just says $(FE)\mathbf{v} = \mathbf{v}$. What does that tell us about the matrix product FE ? If $(FE)\mathbf{v} = \mathbf{v}$ for all $\mathbf{v} \in \mathbb{R}^m$, then we could take $\mathbf{v} = \mathbf{e}_j$ as the standard basis vectors. We find $(FE)\mathbf{e}_j = \mathbf{e}_j$, and so the j th column of FE must be \mathbf{e}_j : the j th column of the $m \times m$ identity matrix. That is, we want $FE = I_m$.

We are actually going to ask for a little bit more in the following definition: that $EF = I_m$ as well. Much later, we’ll see that this extra bit is delightfully redundant—this is a surprise, since matrix multiplication usually is not commutative.

12.3 Definition. A matrix $E \in \mathbb{R}^{m \times m}$ is **INVERTIBLE** if there exists a matrix $F \in \mathbb{R}^{m \times m}$ such that

$$FE = I_m \quad \text{and} \quad EF = I_m. \quad (12.1)$$

12.4 Example. (i) Let

$$E = \begin{bmatrix} 1 & 0 \\ -2 & 1 \end{bmatrix}$$

be the elimination matrix that subtracts 2 times the first row from the second row. Can we invert E ? We’re done if we find $F \in \mathbb{R}^{2 \times 2}$ such that $EF = FE = I_2$. What *should* F be?

This is where it might help to think about E *dynamically*: what does E do? We just said it: E subtracts 2 times the first row from the second row. So undoing E should add two times the first row to the second row. That is,

$$E \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} v_1 \\ v_2 - 2v_1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} v_1 \\ (v_2 - 2v_1) + 2v_1 \end{bmatrix} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}.$$

So maybe

$$F = \begin{bmatrix} 1 & 0 \\ 2 & 1 \end{bmatrix}$$

works. Check it yourself.

(ii) Let

$$P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

be the permutation matrix that interchanges rows 1 and 2. Undoing P should interchange those rows again: we want $F \in \mathbb{R}^{2 \times 2}$ such that if

$$P \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} v_2 \\ v_1 \end{bmatrix}, \quad \text{then} \quad F \begin{bmatrix} v_2 \\ v_1 \end{bmatrix} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}.$$

This looks like we should just take $F = P$. I suggest that you check that $P^2 = I_2$. By the way, this is the first time we're using "power" notation for matrix multiplication: $P^2 = PP$.

Content from Strang's ILA 6E. Read Examples 4 and 5 on p. 52 about inverting elimination matrices. Skip the remarks about the inverse of FE in Example 5 for now.

Example 12.4 should be comforting in that it suggests that elimination and permutation matrices are invertible. We'd probably like to say that their "inverses" are what we expect: invert subtracting by adding, invert permuting by permuting again. What gives us the right to say that a matrix has only one inverse? A (nonzero) real number has only one reciprocal to undo multiplication, but why is this true for matrices?

Here's why. Suppose that E has "two" inverses F_1 and F_2 , so

$$F_1 E = E F_1 = F_2 E = E F_2 = I_m. \quad (12.2)$$

We need to show that $F_1 = F_2$. Here is a great trick: multiply by 1. You know that $1x = x$ for any $x \in \mathbb{R}$, and the same is true for matrices.

12.5 Problem (!). Check that $AI_m = I_m A = A$ for any $A \in \mathbb{R}^{m \times m}$.

So,

$$F_1 = F_1 I_m = F_1 (E F_2) = (F_1 E) F_2 = I_m F_2 = F_2. \quad (12.3)$$

Here is the formal result.

12.6 Theorem. Let $E \in \mathbb{R}^{m \times m}$. There exists at most one $F \in \mathbb{R}^{m \times m}$ satisfying (12.1).

Content from Strang's ILA 6E. This is Note 2 on p. 50.

We can now talk about "the" inverse of a matrix.

12.7 Definition. Let $E \in \mathbb{R}^{m \times m}$ be invertible. The **INVERSE** of E is the unique matrix F satisfying

$$FE = EF = I_m,$$

and we write $F = E^{-1}$.

Let's generalize Example 12.4.

12.8 Theorem. (i) Let $E_{ij} \in \mathbb{R}^{m \times m}$ be the elimination matrix that subtracts ℓ times row j from row i (so 1's along the diagonal, $-\ell$ in the (i, j) -entry, and 0 everywhere else). Then E_{ij} is invertible, and E_{ij}^{-1} is the elimination matrix that adds ℓ times row j to row i (so 1's along the diagonal, ℓ in the (i, j) -entry, and 0 everywhere else).

(ii) Let $P_{ij} \in \mathbb{R}^{m \times m}$ be the permutation matrix that interchanges rows i and j (so P_{ij} is the $m \times m$ identity matrix with columns i and j interchanged). Then P_{ij} is invertible and $P_{ij}^{-1} = P_{ij}$.

Now go back and look very carefully at the calculation in (12.3). We did not use all of the equalities in (12.2). Instead, we only needed that $F_1 E = I_m$ and $E F_2 = I_m$. We might call F_1 a **LEFT INVERSE** and F_2 a **RIGHT INVERSE**. Here is what we have proved.

12.9 Corollary. Let $E \in \mathbb{R}^{m \times m}$ have left and right inverses in the sense that there are $F_1, F_2 \in \mathbb{R}^{m \times m}$ such that

$$F_1 E = I_m \quad \text{and} \quad E F_2 = I_m.$$

Then E is invertible and $F_1 = F_2 = E^{-1}$.

Proof. Okay, maybe this needs a teensy bit of proof. First, the calculation in (12.3) shows $F_1 = F_2$. Put $F = F_1$. Then the hypotheses give $F E = F_1 E = I_m$ and $E F = E F_2 = I_m$, and so F satisfies Definition 12.7. ■

12.10 Remark. We will eventually prove that the existence of a left or right inverse alone is enough to guarantee the invertibility of a matrix! That is, if $E, F \in \mathbb{R}^{m \times m}$ with $EF = I_m$, then both E and F are invertible. We will need some more technology to do that, however.

12.11 Problem (!). We probably expect that undoing the undoing of an action does that action. Totally makes sense, right? More precisely, if $E \in \mathbb{R}^{m \times m}$ is invertible, we should expect that E^{-1} is also invertible and $(E^{-1})^{-1} = E$. (That's how exponents work, right?) Prove this by showing that E satisfies the definition of inverse for E^{-1} . *What things do defines what things are.*

12.12 Problem (★). Let $E, A \in \mathbb{R}^{m \times m}$. Suppose that $EA = I_m$ and E is invertible. Prove that A is invertible, too.

We are particularly interested in inverting a matrix that is a product of elimination matrices and permutation matrices. We know that any elimination or permutation matrix is invertible. More generally, is the product of invertible matrices invertible?

Yes. Suppose that $A, B \in \mathbb{R}^{m \times m}$ are invertible. We will show that AB is invertible. Think about action: first you do B to a vector \mathbf{v} by multiplying $B\mathbf{v}$, and then you do A by multiplying $A(B\mathbf{v}) = (AB)\mathbf{v}$. To undo AB , you probably want to undo A first and then B .

(Getting dressed, socks go on first, then shoes; getting undressed, shoes come off first, then socks.) So we might guess that $(AB)^{-1} = B^{-1}A^{-1}$. The good news is that we can check this using the definition:

$$(B^{-1}A^{-1})(AB) = B^{-1}(A^{-1}A)B = B^{-1}I_m B = B^{-1}B = I_m.$$

12.13 Problem (!). Check that $(AB)(B^{-1}A^{-1}) = I_m$ as well.

Here is the formal result.

12.14 Theorem. Let $A, B \in \mathbb{R}^{m \times m}$ be invertible. Then AB is invertible and $(AB)^{-1} = B^{-1}A^{-1}$.

Content from Strang's ILA 6E. Read “The Inverse of a Product AB ” on pp. 51–52. Then go back to Example 5 on p. 52. The point for our larger story is that multiplying elimination matrices together when getting $EA = U$ is not the best of ideas, whereas computing E^{-1} is more meaningful.

This seems to be everything that we want. Theorem 12.1 tells us that for any $A \in \mathbb{R}^{m \times m}$, we can always find a product of elimination and/or permutation matrices, which we call E , such that $EA = U$ is upper-triangular. Now we know that E is invertible. Given $\mathbf{b} \in \mathbb{R}^m$, it is usually easier to solve $U\mathbf{x} = E\mathbf{b}$, and then we have $E^{-1}U\mathbf{x} = E^{-1}(E\mathbf{b})$, where

$$E^{-1}U = E^{-1}(EA) = (E^{-1}E)A = I_m A = A \quad \text{and} \quad E^{-1}(E\mathbf{b}) = (E^{-1}E)\mathbf{b} = \mathbf{b}.$$

Thus $A\mathbf{x} = \mathbf{b}$, which is what we always wanted to be sure of.

Invertibility is another way of asking about solvability of linear systems. Suppose that $A \in \mathbb{R}^{m \times m}$ is invertible. I claim that $A\mathbf{x} = \mathbf{b}$ always has a solution, and that solution is unique. For uniqueness, work backwards and assume $A\mathbf{x} = \mathbf{b}$; then $A^{-1}(A\mathbf{x}) = A^{-1}\mathbf{b}$, and so $\mathbf{x} = A^{-1}\mathbf{b}$. To check that this is actually a solution, plug in: $A(A^{-1}\mathbf{b}) = (AA^{-1})\mathbf{b} = I_m\mathbf{b} = \mathbf{b}$.

12.15 Theorem. Let $A \in \mathbb{R}^{m \times m}$ be invertible and $\mathbf{b} \in \mathbb{R}^m$. Then the problem $A\mathbf{x} = \mathbf{b}$ has the unique solution $\mathbf{x} = A^{-1}\mathbf{b}$.

Content from Strang's ILA 6E. This is Note 3 on p. 50.

12.16 Problem (!). Hugely important: convince yourself of the following for $A \in \mathbb{R}^{m \times m}$.

- (i) If there is $\mathbf{z} \in \mathbb{R}^m$ with $\mathbf{z} \neq \mathbf{0}_m$ and $A\mathbf{z} = \mathbf{0}_m$, then A is not invertible.
- (ii) If A is invertible, then $\mathbf{C}(A) = \mathbb{R}^m$.

12.17 Problem (★). Often knowing that a matrix is invertible is more useful than having a formula for that inverse. Here's a situation in which the presence of an invertible matrix "keeps things the same." Let $A \in \mathbb{R}^{m \times n}$ be any matrix and let $B \in \mathbb{R}^{n \times n}$ be invertible. Show that $\mathbf{C}(AB) = \mathbf{C}(A)$ as follows. First, explain why $AB\mathbf{v} \in \mathbf{C}(A)$ for any $\mathbf{v} \in \mathbb{R}^n$. Next, justify the equality $A\mathbf{w} = (AB)(B^{-1}\mathbf{w})$ and explain how that shows that anything in $\mathbf{C}(A)$ is in $\mathbf{C}(AB)$.

Content from Strang's ILA 6E. I am not going to talk about determinants now, or much later (I hope!), but you should read Note 6 on p. 50 and Example 2 on p. 51 and also think about the four 2×2 matrices in Example 3 on p. 51. Determinants are a quick and easy way of understanding 2×2 matrices, which arise in a lot of applications (e.g., ordinary differential equations). Try using Note 6 to solve our original problem

$$\begin{bmatrix} 1 & -2 \\ 3 & 2 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 \\ 11 \end{bmatrix}.$$

Day 13: Wednesday, February 5.

Using the solution formula $\mathbf{x} = A^{-1}\mathbf{b}$ from Theorem 12.15 in practice requires us to compute A^{-1} . This turns out to be "expensive" computationally, rather more so than elimination and back-substitution.

Content from Strang's ILA 6E. Read "The Cost of Elimination" on pp. 57–58. The following link to a section from the fifth edition elaborates on this:

https://math.mit.edu/gs/linearalgebra/ila5/linearalgebra5_11-1.pdf.

The point is that using A^{-1} to solve $A\mathbf{x} = \mathbf{b}$ for $A \in \mathbb{R}^{m \times m}$ might take around m^3 arithmetical operations, but using elimination would take only around $m^3/3$ operations. If this excites you, take a numerical linear algebra class. Read the beautiful book by Trefethen & Bau, too.

Let's go back to elimination in the context of inverses. How does being able to solve a linear system $A\mathbf{x} = \mathbf{b}$ via elimination say anything about the invertibility of A ?

We'll start with the nicest case: upper-triangular. I claim that we can eliminate "upwards" on an upper-triangular matrix with nonzero diagonal entries to find an invertible matrix E such that $EU = I_m$. Then $U = E^{-1}$, and so U is invertible. Here is how this works.

13.1 Example. Let

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

We met this matrix in Example 10.8. I want to turn U into I_3 starting from the bottom.

The first thing to do is to make that entry of 2 in the (3,3)-slot into a 1. This requires division by 2 in the third row. Of course we want to encode this, like everything else, via matrix multiplication. What matrix $D \in \mathbb{R}^{3 \times 3}$ does that? We want

$$D \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \begin{bmatrix} v_1 \\ v_2 \\ v_3/2 \end{bmatrix}.$$

I think you know what to do by now: expand the vector on the right as a linear combination weighted by v_1 , v_2 , and v_3 , and you'll see that D should be

$$D_{33} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1/2 \end{bmatrix}.$$

I'm calling it D_{33} now because the action is happening in the (3,3)-entry.

So, we have the transformation

$$\begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 2 \end{bmatrix} \xrightarrow[D_{33}]{R3 \mapsto (1/2) \times R3} \begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}, \quad D_{33} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1/2 \end{bmatrix}.$$

This **SCALING MATRIX** D_{33} , along with the elimination and permutation matrices, is the last of the so-called **ELEMENTARY MATRICES** that we need to encode “row operations” on matrices.

Now we eliminate “upwards.” We want the other entries in column 3 to be 0, so we subtract multiples of row 3 from rows 1 and 2. (Well, multiples of 1.) We get

$$\begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \xrightarrow[E_{23}]{R2 \mapsto R2 - R3} \begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad E_{23} := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 1 \end{bmatrix}$$

$$\xrightarrow[E_{13}]{R1 \mapsto R1 - R3} \begin{bmatrix} 2 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad E_{13} := \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

And then we'll subtract a multiple of row 2 from row 1 to make that (1,2)-entry 0:

$$\begin{bmatrix} 2 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \xrightarrow[E_{12}]{R1 \mapsto R1 - R2} \begin{bmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad E_{12} := \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Last, we rescale the first row:

$$\begin{bmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \xrightarrow[D_{11}]{R1 \mapsto (1/2) \times R1} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = I_3, \quad D_{11} := \begin{bmatrix} 1/2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

We conclude

$$D_{11}E_{12}E_{13}E_{23}D_{33}U = I_3,$$

so putting

$$E := D_{11}E_{12}E_{13}E_{23}D_{33}$$

gives $EU = I_3$. Certainly E is invertible, as all elimination matrices are invertible, and scaling matrices are invertible when their diagonal entries are nonzero. Then $U = E^{-1}I_3 = E^{-1}$, and so U is invertible with $U^{-1} = E$.

13.2 Problem (★). Let $D \in \mathbb{R}^{m \times m}$ be **DIAGONAL**: the (i, j) -entry of D is 0 for $i \neq j$. Prove that if all of the diagonal entries of D are nonzero, then D is invertible; give an explicit formula for D^{-1} .

The arithmetic in Example 13.1 is called **GAUSS–JORDAN ELIMINATION**. I'll state how this works in the abstract.

13.3 Theorem (Gauss–Jordan elimination). Let $U \in \mathbb{R}^{m \times m}$ be upper-triangular with nonzero diagonal entries. Then there exists an invertible matrix $E \in \mathbb{R}^{m \times m}$, which is the product of elimination and/or scaling matrices (but not permutation matrices), such that $EU = I_m$.

Proof. This should feel basically the same as the proof of Theorem 12.1. Multiply U by a scaling matrix D_{mm} to divide row m by $u_{mm} \neq 0$ so that the (m, m) -entry of $D_{mm}U$ is 1. Then subtract multiples of row m from rows $m-1$ through 1 to create zeros in rows $m-1$ through 1 of column m . Go to the $(m-1, m-1)$ -entry: rescale so that it's 1, create zeros in rows $m-2$ through 1 of column $m-1$ through elimination. Repeat. Let E be the product of all of the scaling and/or elimination matrices used, in the order that you use them from the bottom up at each stage. No need for permutation matrices because all of the diagonal entries are nonzero. ■

13.4 Remark. Previously we used “Gaussian elimination” on an arbitrary $A \in \mathbb{R}^{m \times m}$ to find an invertible matrix $E \in \mathbb{R}^{m \times m}$ such that $EA = U$ with U upper-triangular. Now, in the special case that the diagonal entries of U were nonzero, we used “Gauss–Jordan elimination” to find another invertible matrix \tilde{E} such that $\tilde{E}U = I_m$, thus $\tilde{E}EA = I_m$, so A is invertible with $A^{-1} = (\tilde{E}E)^{-1}$.

Content from Strang's ILA 6E. Page 57 offers an algorithm for computing A^{-1} by hand if you really need to do it for a small A . I will never ask you to do that, and Strang doesn't even give any problems asking you to do it in this edition—that's how deprecated the method is. Far better to *understand* A^{-1} than have a formula for it.

13.5 Problem (!). Explain why the matrix A from Example 10.8 is invertible. What is A^{-1} ? (Don't actually compute it—no one really cares—but express A^{-1} as the product of

the inverses of a bunch of elimination, scaling, and/or permutation matrices.)

What we really care about is not a formula for matrix inverses but the existence and behavior of inverses. We've seen a bunch of behaviors already: how the inverse of a product works, how the inverse is itself invertible, and, admittedly, the special formulas for inverses of elimination and permutation matrices.

13.6 Theorem. *Let $U \in \mathbb{R}^{m \times m}$ be upper-triangular. Then U is invertible if and only if all of its diagonal entries are nonzero.*

Proof. (\Leftarrow) This is easier, so I'll do it first. It's just Gauss–Jordan elimination: since the diagonal entries of U are nonzero, we can find $E \in \mathbb{R}^{m \times m}$ invertible such that $EU = I_m$, thus $U = E^{-1}$ is invertible.

(\Rightarrow) We are going to use contradiction. What if U is invertible and a diagonal entry is zero? Something has to go wrong, and I am going to spoil the surprise for you: we are going to find a nonzero vector $\mathbf{x} \in \mathbb{R}^m$ such that $U\mathbf{x} = \mathbf{0}_m$. This will contradict Theorem 12.15, which says that since U is invertible, the only solution is $\mathbf{x} = \mathbf{0}_m$.

I want to consider two possible structures of U : one where the first diagonal entry is zero and one where it isn't, but a zero diagonal entry occurs further down along the diagonal. Here is the first when $m = 4$:

$$U = \begin{bmatrix} 0 & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{bmatrix}.$$

More generally, U has the form

$$U = \begin{bmatrix} \mathbf{0}_m & \tilde{U} \end{bmatrix},$$

where \tilde{U} is “the rest” of U (columns 2 through m). Recall now how we “extract” columns from a matrix: multiply by the standard basis vectors. So $U\mathbf{e}_1$ is the first column of U , where \mathbf{e}_1 is 1 in row 1 and 0 everywhere else. That is, $U\mathbf{e}_1 = \mathbf{0}_m$ and $\mathbf{e}_1 \neq \mathbf{0}_m$. That's the contradiction.

Next case: a zero entry further down on the diagonal. That is, $u_{jj} = 0$ for some $j \geq 2$ but $u_{ii} \neq 0$ for $1 \leq i \leq j - 1$. Here is one such possibility when $m = 4$:

$$U = \begin{bmatrix} \odot & * & \odot & * \\ 0 & \odot & \odot & * \\ 0 & 0 & 0 & * \\ 0 & 0 & 0 & * \end{bmatrix}.$$

By \odot I mean nonzero entries, and in the notation above $j = 3$. Now look at the upper-triangular matrix

$$\hat{U} := \begin{bmatrix} \odot & * \\ 0 & \odot \end{bmatrix}.$$

This has nonzero diagonal entries and so we can find $x_1, x_2 \in \mathbb{R}$ such that

$$\widehat{U} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} \odot \\ \odot \end{bmatrix}.$$

But then

$$x_1 \begin{bmatrix} \odot \\ 0 \\ 0 \\ 0 \end{bmatrix} + x_2 \begin{bmatrix} * \\ \odot \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \odot \\ \odot \\ 0 \\ 0 \end{bmatrix},$$

and from that

$$U\mathbf{x} = \mathbf{0}_4, \quad \mathbf{x} := \begin{bmatrix} x_1 \\ x_2 \\ -1 \\ 0 \end{bmatrix}.$$

So what is the problem? We have found $\mathbf{x} \in \mathbb{R}^4$ such that $\mathbf{x} \neq \mathbf{0}_4$ but $U\mathbf{x} = \mathbf{0}_4$. This contradicts Theorem 12.15; since U is invertible, the only \mathbf{x} that should work there is $\mathbf{x} = \mathbf{0}_4$.

Here is the generalization of this, which I didn't do in class, and for which you are not responsible. As before, assume that there is $j \geq 2$ such that $u_{jj} = 0$ but $u_{ii} \neq 0$ for $1 \leq i \leq j-1$. Write

$$U = \left[\begin{array}{c|c} \widehat{U} & \widehat{\mathbf{u}}_j \\ \hline \mathbf{0} & \mathbf{0}_{m-j} \end{array} \middle| \widetilde{U} \right].$$

Here \widehat{U} is a $(j-1) \times (j-1)$ upper-triangular matrix with nonzero diagonal entries, $\widehat{\mathbf{u}}_j \in \mathbb{R}^{j-1}$, and \widetilde{U} contains the remaining columns of U . I am irritatingly using 0 to mean a matrix whose entries are all 0.

Since the diagonal entries of \widehat{U} are nonzero, we can find $\widehat{\mathbf{x}} \in \mathbb{R}^{j-1}$ such that $\widehat{U}\widehat{\mathbf{x}} = \widehat{\mathbf{u}}_j$. Then

$$U\mathbf{x} = \mathbf{0}_m, \quad \mathbf{x} := \begin{bmatrix} \widehat{\mathbf{x}} \\ -1 \\ \mathbf{0}_{m-(j+1)} \end{bmatrix},$$

and so the equation $U\mathbf{x} = \mathbf{0}_m$ has a nonzero solution. Thus U cannot be invertible. ■

13.7 Problem (+). What goes wrong if a diagonal entry of $U \in \mathbb{R}^{m \times m}$ is zero? We saw this in Example 11.5, but it's worth revisiting in the abstract. There are two possibilities: either the last diagonal entry is 0 or an entry further up the diagonal is 0, but the entries below it are nonzero.

(i) In the first case, the bottom row of U is 0. What does that say about the $\mathbf{c} \in \mathbb{R}^{m \times m}$ for which we can solve $U\mathbf{x} = \mathbf{c}$? And what does that say in the context of Theorem 12.15?

(ii) In the second case, U might have a structure like the following:

$$U = \begin{bmatrix} * & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & \odot & * \\ 0 & 0 & 0 & \odot \end{bmatrix}.$$

As before $\odot \neq 0$. Explain why for this particular U , there is an invertible $\tilde{E} \in \mathbb{R}^{4 \times 4}$ such that $\tilde{E}U$ has a row of all 0 entries. [Hint: *eliminate upward*.] What does this say about the $\mathbf{c} \in \mathbb{R}^4$ for which we can have $U\mathbf{x} = \mathbf{c}$?

Day 14: Friday, February 7.

The tool that we used in the second part of the proof of Theorem 13.6 is incredibly nice, and it's going to resolve some of our long-standing conjectures. First we isolate that part of the proof for future reference.

14.1 Corollary. *Let $U \in \mathbb{R}^{m \times m}$ be upper-triangular. Suppose that a diagonal entry of U is nonzero. Then there exists a nonzero vector $\mathbf{x} \in \mathbb{R}^m$ such that $U\mathbf{x} = \mathbf{0}_m$.*

Here's a blast from the past. We've said that the columns of a matrix $A = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n] \in \mathbb{R}^{m \times n}$ are independent if $\mathbf{a}_1 \neq \mathbf{0}_m$ and $\mathbf{a}_j \notin \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_{j-1})$. Here A does not have to be square. This is just a more precise way of saying that no column is a linear combination of the others, and so independence is a precise way of controlling redundancy in a matrix's data. It turns out that independence is intimately connected with the **HOMOGENEOUS** problem $A\mathbf{x} = \mathbf{0}_m$.

14.2 Theorem. *The columns of a matrix $A \in \mathbb{R}^{m \times n}$ are independent if and only if the only solution to $A\mathbf{x} = \mathbf{0}_m$ is $\mathbf{x} = \mathbf{0}_n$.*

Proof. It's actually easier to prove

$$\text{dependent columns} \iff A\mathbf{x} = \mathbf{0}_m \text{ has a nonzero solution,}$$

so we'll do that. (Being negative is much more fun than being positive.)

If the columns are dependent, then one column is a linear combination of the others. Let \mathbf{x} be the vector consisting of the weights from that combination with -1 for the "bad vector." That's the proof in words. In symbols, this is basically the proof of Theorem 7.1. For simplicity, if $A = [\mathbf{a}_1 \ \mathbf{a}_2 \ (c_2\mathbf{a}_2 + c_4\mathbf{a}_4) \ \mathbf{a}_4] \in \mathbb{R}^{m \times 4}$, then

$$A \begin{bmatrix} 0 \\ c_2 \\ -1 \\ c_4 \end{bmatrix} = c_2\mathbf{a}_2 - (c_2\mathbf{a}_2 + c_4\mathbf{a}_4) + c_4\mathbf{a}_4 = \mathbf{0}_m.$$

Now suppose that $A\mathbf{x} = \mathbf{0}_m$ has a nonzero solution. So, at least one entry of \mathbf{x} is nonzero. Pick any such nonzero entry, divide $A\mathbf{x} = \mathbf{0}_m$ by it, and rewrite the corresponding column of A as a linear combination of the rest. For example, if $A = [\mathbf{a}_1 \ \mathbf{a}_2 \ \mathbf{a}_3 \ \mathbf{a}_4] \in \mathbb{R}^{m \times 4}$ with $A\mathbf{x} = \mathbf{0}_m$, $\mathbf{x} = (x_1, x_2, x_3, x_4)$, and $x_3 \neq 0$, then

$$\mathbf{0}_m = A\mathbf{x} = x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + x_3\mathbf{a}_3 + x_4\mathbf{a}_4,$$

so

$$\frac{x_1}{x_3}\mathbf{a}_1 + \frac{x_2}{x_3}\mathbf{a}_2 + \mathbf{a}_3 + \frac{x_4}{x_3}\mathbf{a}_4 = \mathbf{0}_m.$$

and thus

$$\mathbf{a}_3 = \left(-\frac{x_1}{x_3}\right)\mathbf{a}_1 + \left(-\frac{x_2}{x_3}\right)\mathbf{a}_2 + \left(-\frac{x_4}{x_3}\right)\mathbf{a}_4$$

is a linear combination of the other columns. ■

We could have done this back when we were talking about independence and dependence. Maybe we should have?

Content from Strang's ILA 6E. Look at “Independent columns” toward the bottom of p. 30.

14.3 Problem (★). We are talking a lot right now about square systems: $A\mathbf{x} = \mathbf{b}$ with $A \in \mathbb{R}^{m \times m}$. Number of equations = number of variables. Independence and dependence make sense for all systems, not just square. Prove that if the columns of $A \in \mathbb{R}^{m \times n}$ are independent, then if a solution to $A\mathbf{x} = \mathbf{b}$ exists (no guarantee that it does!), it is unique. [Hint: reread the proof of Theorem 10.1 and then the statement of Problem 10.2.]

Remarkably, for a square system, independent columns and invertibility are the same thing.

14.4 Theorem. A matrix $A \in \mathbb{R}^{m \times m}$ is invertible if and only if its columns are independent.

Proof. (\implies) By Theorem 14.2, to prove that the columns are independent, we should assume that $A\mathbf{x} = \mathbf{0}_m$ for some $\mathbf{x} \in \mathbb{R}^m$ and show $\mathbf{x} = \mathbf{0}_m$. This is almost automatic by Theorem 12.15, as $\mathbf{x} = A^{-1}\mathbf{0}_m = \mathbf{0}_m$.

(\impliedby) Assume that the columns are independent, and watch Gaussian elimination rise and shine! Whether or not the columns of A are independent, we can find an invertible $E \in \mathbb{R}^{m \times m}$ such that $EA = U$ with U upper-triangular. Then $A = E^{-1}U$. If U is also invertible, Theorem 12.14 guarantees that A is invertible. Theorem 13.6 says that U will be invertible if its diagonal entries are nonzero.

So are they? What goes wrong if U has a nonzero diagonal entry? Corollary 14.1 goes wrong: there is $\mathbf{x} \neq \mathbf{0}_m$ such that $U\mathbf{x} = \mathbf{0}_m$. But then $A\mathbf{x} = E^{-1}U\mathbf{x} = E^{-1}\mathbf{0}_m = \mathbf{0}_m$. This contradicts the independence of the columns of A . That's wrong. ■

Let us celebrate this result with a proof of Conjecture 6.14 which conjectured that if $A \in \mathbb{R}^{m \times m}$ has independent columns, then $\mathbf{C}(A) = \mathbb{R}^m$. Theorem 14.4 says that independence implies invertibility. And if A is invertible, then $\mathbf{C}(A) = \mathbb{R}^m$ because we can always solve $A\mathbf{x} = \mathbf{b}$. (We get *unique* solvability from independence to boot.) This is the victory of independence: having just enough data in your matrix, and no more, lets you solve $A\mathbf{x} = \mathbf{b}$, and solve it uniquely.


14.5 Problem (!). How does this lead to a quick proof of Conjecture 6.7? [Hint: “ $P \iff Q$ ” is the same as “not $P \iff$ not Q .”]

We have proved a lot of things recently, and all of them are saying basically the same thing. It can get hard to keep track of all of this, so here’s a summary.

14.6 Theorem (Invertible matrix theorem). Let $A \in \mathbb{R}^{m \times m}$. The following statements are equivalent in the sense that if any one of them is true, then all of the others are true.

- (i) A is invertible.
- (ii) There is an invertible matrix E such that $U := EA$ is upper-triangular and all of the diagonal entries of U are nonzero.
- (iii) The columns of A are independent.
- (iv) The only solution to $A\mathbf{x} = \mathbf{0}_m$ is $\mathbf{x} = \mathbf{0}_m$.
- (v) For any $\mathbf{b} \in \mathbb{R}^m$, the problem $A\mathbf{x} = \mathbf{b}$ always has a unique solution.

14.7 Problem (!). We have essentially proved Theorem 14.6, but the work is spread out among a number of results. Now is a good opportunity for you to review the logical connections among those results. Here is a cartoon summarizing those connections, which you can tease out in the steps below.



- (i) To prove that parts (i) and (ii) are equivalent, use Theorems 12.1, 12.14, and 13.6.
- (ii) Review the proof of Theorem 14.4 to remind yourself why parts (i) and (iii) are equivalent.
- (iii) Review the proof of Theorem 14.2 to remind yourself why parts (iii) and (iv) are equivalent.
- (iv) Review the argument leading up to Theorem 12.15 to remind yourself why part (i)

implies part (v). Then argue that (v) implies part (iv), and we know that implies (iii) and thus part (i).

14.8 Problem (+). Corollary 12.9 told us that $A \in \mathbb{R}^{m \times m}$ is invertible if and only if A has left and right inverses, i.e., matrices $L, R \in \mathbb{R}^{m \times m}$ such that $LA = AR = I_m$. The upshot of this condition is that we don't need to verify $L = R$; all we need is $LA = I_m$ and $AR = I_m$, and ostensibly the matrices L and R don't have to talk to each other. Here we show that we only need a left inverse or a right inverse to guarantee invertibility.

(i) Suppose that $LA = I_m$. Prove that the columns of A are linearly independent and thus A is invertible. [Hint: if $A\mathbf{x} = \mathbf{0}_m$, apply L to both sides and figure out \mathbf{x} .]

(ii) Suppose that $AR = I_m$. Let $EA = U$ be upper-triangular for some invertible matrix E , so $UR = E$. If U is not invertible, then U has a zero diagonal entry. Problem 13.7 then implies that for some invertible \tilde{E} , the product $\tilde{E}U$ has a row whose entries are all 0. Use Problem 8.5 to deduce something about $\tilde{E}UR$, and thus about $\tilde{E}E$. Since \tilde{E} and E are invertible, what contradiction results?

(iii) We now can weaken part (v) of Theorem 14.6 slightly, but crucially. It turns out that $A \in \mathbb{R}^{m \times m}$ is invertible if and only if the problem $A\mathbf{x} = \mathbf{b}$ always has a solution for each $\mathbf{b} \in \mathbb{R}^m$. We do not need to require the solution to be unique. Certainly the existence of a solution is a consequence of invertibility (uniqueness, too). Now suppose that we can always solve $A\mathbf{x} = \mathbf{b}$. Explain how choosing \mathbf{b} to be the standard basis vectors produces a matrix R such that $AR = I_m$.

14.9 Remark. The nonzero diagonal entries of an upper-triangular matrix are sometimes called its **PIVOTS**. The pivots of a general $A \in \mathbb{R}^{m \times m}$ are the nonzero diagonal entries of the upper-triangular matrix to which A can always be transformed by elimination and row interchanges, i.e., by Theorem 12.1. This language is a little perilous, as we never proved that the matrix U from Theorem 12.1 was unique—could we write $E_1A = U_1$ and $E_2A = U_2$ with U_1 and U_2 both upper-triangular, $U_1 \neq U_2$, and E_1 and E_2 as the product of elimination and/or permutation matrices? What's important from the point of view of invertibility is not the exact value of these “pivots” but rather whether they are all nonzero or not.

14.10 Problem (+). Let $A \in \mathbb{R}^{m \times m}$ and suppose that $E_1, E_2 \in \mathbb{R}^{m \times m}$ are invertible with E_1A and E_2A both upper-triangular. Prove that if E_1A has no nonzero diagonal entries, then E_2A also has no nonzero diagonal entries. [Hint: what goes wrong if E_2A has some nonzero diagonal entries?] We will eventually prove that E_1A and E_2A must have the same number of nonzero diagonal entries, although we need more technology for that.

Content from Strang's ILA 6E. Page 41 introduces the terminology “pivot.” I personally feel that the phrase “nonzero pivot” is redundant. Informally, you should think of the

pivots as “the nonzero things that you multiply by when doing elimination.” Because we can permute rows even when we don’t need to avoid zero diagonal entries, we can select an “ideal” pivot at any state of elimination—see “Partial Pivoting” to Reduce Roundoff Errors” on p. 66 and think once more about taking a numerical linear algebra class after this one.

Day 15: Monday, February 10.

We’ve learned a lot about invertible matrices—in particular that we can always solve $A\mathbf{x} = \mathbf{b}$ with $\mathbf{x} = A^{-1}\mathbf{b}$ when A is invertible, but that we probably shouldn’t because computing A^{-1} is computationally expensive. The alternative is that we do elimination on A so that $U := EA$ is upper-triangular with nonzero diagonal entries, and then we solve $U\mathbf{x} = E\mathbf{b}$ via back-substitution. That requires us to compute $E\mathbf{b}$, too. There is a variation on this approach that is still computationally less expensive than computing A^{-1} and that gives us some new insights into matrix multiplication, so it’s worth learning. We start with a very concrete example.

15.1 Example. Let

$$A = \begin{bmatrix} 2 & 1 & 1 \\ 4 & 3 & 3 \\ 8 & 7 & 9 \end{bmatrix}.$$

We saw in Example 10.8 that

$$EA = \begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 2 \end{bmatrix} = U,$$

where

$$E := E_{32}E_{31}E_{21}$$

and

$$E_{21} := \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad E_{31} := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -4 & 0 & 1 \end{bmatrix}, \quad \text{and} \quad E_{32} := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -3 & 1 \end{bmatrix}.$$

We went further in Example 13.1 and found \tilde{E} such that $\tilde{E}U = I_3$, but that’s less important here. Rather, the new thing to focus on is the *factorization*

$$A = E^{-1}U.$$

Recall that we originally talked about multiplying matrices with the goal of *factoring* matrices: breaking matrices into products of simpler matrices to reveal meaningful properties. What is simpler about the matrices E^{-1} and U , and what is meaningful about the factorization $A = E^{-1}U$?

Certainly U is simpler than A because U is upper-triangular: U has a nice structure with a lot of simple data—many zero entries. What about E^{-1} ? A *bad* idea is to compute E as the product $E = E_{32}E_{31}E_{21}$ and then try to compute E^{-1} from that. Go ahead and try it and see how opaque the work is. (I mean, we haven't really computed a matrix inverse other than Remark 13.4.) But we do know that

$$E^{-1} = (E_{32}E_{31}E_{21})^{-1} = E_{21}^{-1}E_{31}^{-1}E_{32}^{-1}.$$

And we know what each of these inverses are because they are inverses of elimination matrices:

$$E_{32}^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 3 & 1 \end{bmatrix}, \quad E_{31}^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 4 & 0 & 1 \end{bmatrix}, \quad \text{and} \quad E_{21}^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Now think about what they are doing. Multiplying by E_{31}^{-1} says “Add 4 times row 1 to row 3”:

$$E_{31}^{-1}E_{32}^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 4 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 3 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 4 & 3 & 1 \end{bmatrix}.$$

says add 4 times row 1 to row 3. Multiplying by E_{21}^{-1} says “Add 2 times row 1 to row 2”:

$$E^{-1} = E_{21}^{-1}E_{31}^{-1}E_{32}^{-1} = E_{21}^{-1} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 4 & 3 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 4 & 3 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 4 & 3 & 1 \end{bmatrix} =: L.$$

Just look at that matrix L . It's **LOWER-TRIANGULAR**, because every entry above the diagonal is 0. And the entries below the diagonal are the negatives of the multipliers from the original elimination step. *This is no accident.*

How does this factorization $A = LU$ help? Let's solve $A\mathbf{x} = \mathbf{b}$ with $\mathbf{b} = (0, 1, 5)$. Ideally you did this in Problem 10.9. This problem is the same as $LU\mathbf{x} = \mathbf{b}$. Now here is the trick: abbreviate $\mathbf{c} := U\mathbf{x}$. Then we want $L\mathbf{c} = \mathbf{b}$. The clever idea is to view \mathbf{c} as an unknown; then we can solve $L\mathbf{c} = \mathbf{b}$ using back-substitution, and then we solve $U\mathbf{x} = \mathbf{c}$ with another round of back-substitution. Nowhere does elimination hit \mathbf{b} .

Let's go: $L\mathbf{c} = \mathbf{b}$ is the system

$$\begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 4 & 3 & 1 \end{bmatrix} \mathbf{c} = \begin{bmatrix} 0 \\ 1 \\ 5 \end{bmatrix},$$

equivalently

$$\begin{cases} c_1 & = 0 \\ 2c_1 + c_2 & = 1 \\ 4c_1 + 3c_2 + c_3 & = 5 \end{cases}$$

The first equation immediately gives $c_1 = 0$, so the second reduces to $c_2 = 1$, and then the third is $3 + c_3 = 5$, thus $c_3 = 2$. Hence

$$\mathbf{c} = \begin{bmatrix} 0 \\ 1 \\ 2 \end{bmatrix}.$$

Next, $U\mathbf{x} = \mathbf{c}$ is the system

$$\begin{bmatrix} 2 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 2 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 0 \\ 1 \\ 2 \end{bmatrix},$$

equivalently

$$\begin{cases} 2x_1 + x_2 + x_3 = 0 \\ x_2 + x_3 = 1 \\ 2x_3 = 2 \end{cases}$$

The third equation is $2x_3 = 2$, thus $x_3 = 1$. Then the second equation is $x_2 + 1 = 1$, so $x_2 = 0$. And the first equation is then $2x_1 + 0 + 1 = 0$, so $2x_1 = -1$, and therefore $x_1 = -1/2$. That is, $\mathbf{x} = (-1/2, 0, 1)$.

This example has a number of lessons for us. First, if we can factor $A = LU$, with L lower-triangular and U upper-triangular, and where both L and U have all nonzero entries on their diagonals, then we can solve $A\mathbf{x} = \mathbf{b}$ easily by back-substitution *and without doing any elimination calculations on \mathbf{b}* . Second, we *might* be able to achieve this “ LU -factorization” if we can reduce A to upper-triangular form using only elimination, not permutation, matrices. In particular, finding that factor of L involved inverting the product of elimination matrices that governed that reduction—but we did not multiply all those elimination matrices together and then calculate the inverse; instead, we used properties of inverses of products and *what elimination matrices do*. (It pains me to say this, but brute force isn’t always the best force.)

All of this turns out to be more generally true.

15.2 Theorem (LU -factorization). Suppose that $A \in \mathbb{R}^{m \times m}$ can be reduced to upper-triangular form using only elimination, not permutation, matrices. That is, there is $E \in \mathbb{R}^{m \times m}$ such that $EA = U$, where U is upper-triangular and E is a product of only elimination matrices. Then $L := E^{-1}$ is lower-triangular, the diagonal entries of L are all 1, and $A = LU$. Moreover, for any $\mathbf{b} \in \mathbb{R}^m$, there is $\mathbf{x} \in \mathbb{R}^m$ such that $A\mathbf{x} = \mathbf{b}$ if and only if there is $\mathbf{c} \in \mathbb{R}^m$ such that

$$\begin{cases} L\mathbf{c} = \mathbf{b} \\ U\mathbf{x} = \mathbf{c}. \end{cases} \quad (15.1)$$

Proof. We are only going to prove the last sentence. The proof that L is lower-triangular when E is a product of only elimination matrices is essentially an abstraction of the calculation in Example 15.1. (Try replacing the multipliers 2, 4, and 3 with arbitrary ℓ_{21} , ℓ_{31} ,

$\ell_{32} \in \mathbb{R}$ and watch the same lower-triangular structure appear. Or check out the readings in Strang mentioned below.)

Here is the proof of that last sentence, assuming that we have the factorization $A = LU$. First, if there is $\mathbf{x} \in \mathbb{R}^m$ such that $A\mathbf{x} = \mathbf{b}$, then $LU\mathbf{x} = \mathbf{b}$. Put $\mathbf{c} = U\mathbf{x}$ to find $L\mathbf{c} = \mathbf{b}$. So, both equations in (15.1) are true.

Now suppose that both equations in (15.1) are true. Work backwards:

$$\mathbf{b} = L\mathbf{c} = L(U\mathbf{x}) = (LU)\mathbf{x} = A\mathbf{x}.$$

By the way, the proof of that last sentence did not use *at all* the fact that L and U are triangular or that L has diagonal entries equal to 1. However, if we wanted to start by solving (15.1) and end up with a solution to $A\mathbf{x} = \mathbf{b}$, it would be necessary for L and U to have all nonzero diagonal entries. ■

15.3 Problem (★). Let

$$A = \begin{bmatrix} 1 & -2 \\ 3 & 2 \end{bmatrix}.$$

Find matrices $L, U \in \mathbb{R}^{2 \times 2}$ such that L is lower-triangular, U is upper-triangular, and $A = LU$. Let $\mathbf{b} = (1, 11)$. Solve $A\mathbf{x} = \mathbf{b}$ by first solving $L\mathbf{c} = \mathbf{b}$ for some $\mathbf{c} \in \mathbb{R}^2$ and then solving $U\mathbf{x} = \mathbf{c}$ for $\mathbf{x} \in \mathbb{R}^2$.

Content from Strang's ILA 6E. Here are sketches of the existence of the LU -factorization. First, reread Example 5 on p. 52 to see again how inverting products of elimination matrices works. Think carefully about the two bold sentences on “feels an effect” and “feels no effect.” Do you understand exactly what this means? Then read p. 53 and contrast the calculations in equations (10) and (11). Which do you like better? Read all of p. 59—and think about the last paragraph on p. 58: “A proof means that we have not just seen that pattern and believed it and liked it, but understood it.” *This is why we prove things.* Another proof of LU appears on p. 60, using the matrix multiplication technique discussed on p. 34.

So who cares? The work in Example 15.1 probably felt no more efficient than a routine back-substitution approach (which you did in Problem 10.9, right?) Maybe it felt more inefficient! That’s a valid feeling. All of our examples in this class are effectively toy problems designed so that the on-the-fly arithmetic is easy.

But what if you need to solve $A\mathbf{x} = \mathbf{b}_j$ for many \mathbf{b}_j ? If you have only a finite number of \mathbf{b}_j , maybe you could work with a large augmented matrix $[A \ \mathbf{b}_1 \ \cdots \ \mathbf{b}_p]$, do elimination on A via the matrix E , so $EA = U$, and then study $[U \ E\mathbf{b}_1 \ \cdots \ E\mathbf{b}_p]$. Then you would have to solve $U\mathbf{x} = E\mathbf{b}_j$ by back-substitution. However, it is arguably less computationally expensive to solve $LU = \mathbf{b}_j$ by the two-step process above. In particular, it may be the case* that solving $A\mathbf{x} = \mathbf{b}_j$ is part of a larger *iterative* process: at the j th step, you get a

* I found this StackExchange post really helpful:

<https://math.stackexchange.com/questions/266355/necessity-advantage-of-lu-decomposition-over-gaussian-elimination>.

new \mathbf{b}_j , but A stays the same. If you want to keep doing this *indefinitely*, the elimination calculations $E\mathbf{b}_j$ may become expensive. Doing the elimination just once to get $A = LU$, and then solving $LU\mathbf{x} = \mathbf{b}_j$ via the two-step process, may be less expensive.

The LU -factorization works when no row interchanges are needed, i.e., when we can write $EA = U$ with U upper-triangular and E as a product only of elimination matrices, not permutation matrices. Basically, it's possible to “almost” commute permutation and elimination matrices so that we have $PA = LU$ with P a product of permutation matrices, L lower-triangular, and U upper-triangular. Figuring out how to get that P factor out front is a little tricky, and I think this is better covered in a numerical linear algebra course. But once you know $PA = LU$, to solve $A\mathbf{x} = \mathbf{b}$, first permute $PA\mathbf{x} = P\mathbf{b}$, and then solve $LU\mathbf{x} = P\mathbf{b}$ as we did above.

Content from Strang's ILA 6E. See p. 65. This is wholly optional reading and requires a little more knowledge of permutation matrices than I expect or desire right now.

Day 16: Wednesday, February 12.

Our best successes in this course arguably come from square systems: $A\mathbf{x} = \mathbf{b}$ with $A \in \mathbb{R}^{m \times m}$ and $\mathbf{b} \in \mathbb{R}^m$, same number of equations as unknowns. We will see that it is with square systems alone that we have a chance (not a guarantee) for both existence and uniqueness of solutions—it is possible both to be able to solve the problem and have only one solution for it. With nonsquare systems— $A\mathbf{x} = \mathbf{b}$, $A \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$, $m \neq n$ —we will show that either existence or uniqueness always fails (maybe both). Understanding how to quantify and qualify our failures, and how to move on from them, will be the central part of our forthcoming story. We can see this happen with relatively small systems using relatively few numbers.

16.1 Example. We consider our favorite problem $A\mathbf{x} = \mathbf{b}$ for the variety of A below.

(i) It's hard to get nicer than

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

because then the unique solution to $A\mathbf{x} = \mathbf{b}$ is always just $\mathbf{x} = \mathbf{b}$ (for $\mathbf{b} \in \mathbb{R}^2$).

(ii) It's easy to get less nice, though. Take

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

Then existence fails for \mathbf{b} with $b_2 \neq 0$, while uniqueness also fails. Inspired by Problem 10.2 and the fact that $A\mathbf{e}_2 = \mathbf{0}_2$, where $\mathbf{e}_2 = (0, 1)$, we can check that

$$A \left(\begin{bmatrix} b_1 \\ 0 \end{bmatrix} + c\mathbf{e}_2 \right) = \begin{bmatrix} b_1 \\ 0 \end{bmatrix}$$

for any $b_1, c \in \mathbb{R}$. Thus solutions, when they exist, are never unique. The dependence of the columns affects both existence and uniqueness here, per the invertible matrix theorem.

(iii) With

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix},$$

existence fails for those $\mathbf{b} \in \mathbb{R}^3$ with $b_3 \neq 0$. But solutions, when they exist are unique, because the only solution to $A\mathbf{x} = \mathbf{0}_3$ is $\mathbf{x} = \mathbf{0}_2$. We saw this in Problem 10.2, and it's closely related to the independence of the columns. We'll revisit this soon.

(iv) With

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$$

both existence and uniqueness fail, since we can't solve $A\mathbf{x} = (0, 1, 0)$, while when we can solve $A\mathbf{x} = (b_1, 0, 0)$ with $\mathbf{x} = (b_1, 0)$, we can also solve it with $\mathbf{x} = (b_1, c)$ for any $c \in \mathbb{R}$. And the columns are dependent.

(v) With

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

we have existence but not uniqueness: take $\mathbf{x} = (b_1, b_2, c)$ to solve $A\mathbf{x} = (b_1, b_2)$. Again, dependent columns.

(vi) Last, existence and uniqueness fail for

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

Once again, we can only solve $A\mathbf{x} = \mathbf{b}$ for $b_2 = 0$, and when we can, $\mathbf{x} = (b_1, c_2, c_3)$ is a solution for any $c_2, c_3 \in \mathbb{R}$. Dependent columns are the worst.

Here is what the previous example suggests as we move beyond square systems.

16.2 Conjecture. Let $A \in \mathbb{R}^{m \times n}$.

(i) If $m > n$ (more equations than unknowns, more rows than columns, A is taller than it is wide), then we will always fail to solve $A\mathbf{x} = \mathbf{b}$ for some $\mathbf{b} \in \mathbb{R}^m$. That is, $\mathbf{C}(A) \neq \mathbb{R}^m$. It may or may not be possible to get unique solutions.

(ii) If $m < n$ (more unknowns than equations, more columns than rows, A is wider than it is tall), then we will never be able to solve $A\mathbf{x} = \mathbf{b}$ uniquely. Solutions may or may not exist in the first place.

Content from Strang's ILA 6E. Now is a good time to reread p. 38.

Our successes going forward will hinge in no small part on a new perspective: how vectors within a given set interact with each other. This may sound weird at first, but trust me that it will feel completely natural soon. Think about the column space. For $A \in \mathbb{R}^{m \times n}$, we have

$$\mathbf{C}(A) = \{A\mathbf{v} \mid \mathbf{v} \in \mathbb{R}^n\}.$$

Every vector in $\mathbf{C}(A)$ is a vector in \mathbb{R}^m . We know very well that controlling the column space is the key (well, a key) to existence of solutions to $A\mathbf{x} = \mathbf{b}$, for this equation is true if and only if $\mathbf{b} \in \mathbf{C}(A)$.

We say column *space*, not column *set*. The *set* of columns of $A = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n] \in \mathbb{R}^{m \times n}$ is just the set $\{\mathbf{a}_1, \dots, \mathbf{a}_n\}$ of at most n vectors (maybe fewer than n , if some of the columns of A are repeated). A *space* is more dynamic.

Specifically, the column space behaves well with respect to the fundamental objects of vector arithmetic. Let $\mathbf{w}_1, \mathbf{w}_2 \in \mathbf{C}(A)$. Then there are $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^n$ such that $\mathbf{w}_1 = A\mathbf{v}_1$ and $\mathbf{w}_2 = A\mathbf{v}_2$. So,

$$\mathbf{w}_1 + \mathbf{w}_2 = A\mathbf{v}_1 + A\mathbf{v}_2 = A(\mathbf{v}_1 + \mathbf{v}_2) \in \mathbf{C}(A)$$

since $\mathbf{v}_1 + \mathbf{v}_2 \in \mathbb{R}^n$. That is, $\mathbf{C}(A)$ is “closed under addition”: adding two vectors in $\mathbf{C}(A)$ yields another vector in $\mathbf{C}(A)$.

Similarly, if $\mathbf{w} \in \mathbf{C}(A)$ with $\mathbf{w} = A\mathbf{v}$ for some $\mathbf{v} \in \mathbb{R}^n$, and if $c \in \mathbb{R}$, then

$$c\mathbf{w} = c(A\mathbf{v}) = A(c\mathbf{v}) \in \mathbf{C}(A),$$

since $c\mathbf{v} \in \mathbb{R}^n$. That is, $\mathbf{C}(A)$ is “closed under scalar multiplication”: multiplying a vector in $\mathbf{C}(A)$ by a real number yields another vector in $\mathbf{C}(A)$.

Finally, since $A\mathbf{0}_n = \mathbf{0}_m$, we have $\mathbf{0}_m \in \mathbf{C}(A)$. Thus the column space is never empty, and in particular it contains one of the most important vectors for vector and matrix arithmetic alike.

Sets of vectors that have these properties—closure under vector addition and scalar multiplication and containing the zero vector—are among the most special and useful kinds of sets. They don’t just exist and contain things; they are *dynamic* with respect to vector operations. We’ll see just how special these sets—these *spaces*—are in the context of understanding, and maybe even solving, $A\mathbf{x} = \mathbf{b}$ for A nonsquare.

Day 17: Friday, February 14.

We took Exam 1.

Day 18: Monday, February 17.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Null space of a matrix, subspace (N)

The column space governs *existence*: we can solve $A\mathbf{x} = \mathbf{b}$ if and only if $\mathbf{b} \in \mathbf{C}(A)$. But the column space says nothing about *uniqueness*: having $\mathbf{b} \in \mathbf{C}(A)$ does not guarantee that there is *only* one \mathbf{x} such that $A\mathbf{x} = \mathbf{b}$, but rather that there is *at least* one such \mathbf{x} . To understand uniqueness, we need to study a new set—more precisely, because of its dynamism, a unique space.

We have discussed the following several times, and you proved it in Problem 10.2.

18.1 Theorem. Let $A \in \mathbb{R}^{m \times n}$.

(i) Suppose that the only $\mathbf{z} \in \mathbb{R}^n$ such that $A\mathbf{z} = \mathbf{0}_m$ is $\mathbf{z} = \mathbf{0}_n$. Then for any $\mathbf{b} \in \mathbb{R}^m$, the problem $A\mathbf{x} = \mathbf{b}$ has at most one solution. (Maybe it has none.)

(ii) Suppose that for all $\mathbf{b} \in \mathbb{R}^m$, the problem $A\mathbf{x} = \mathbf{b}$ has at most one solution. (Maybe it has none.) Then the only solution to $A\mathbf{x} = \mathbf{0}_m$ is $\mathbf{x} = \mathbf{0}_n$.

You saw another version of this in Theorem 14.2, right? Rereading that theorem right now is probably a good idea.

Here is the point: we can understand uniqueness of solutions to the problem $A\mathbf{x} = \mathbf{b}$ for *any* \mathbf{b} by studying the problem for the special case of $\mathbf{b} = \mathbf{0}_m$. This motivates the study of a new dynamic set related to A .

18.2 Definition. Let $A \in \mathbb{R}^m$. The **NULL SPACE** of A is

$$\mathbf{N}(A) := \{\mathbf{v} \in \mathbb{R}^n \mid A\mathbf{v} = \mathbf{0}_m\}.$$

18.3 Example. (i) The null space of I_2 is just $\{\mathbf{0}_2\}$, for if $I_2\mathbf{v} = \mathbf{0}_2$, then since $I_2\mathbf{v} = \mathbf{v}$, we just have $\mathbf{v} = \mathbf{0}_2$.

(ii) Let

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

The system $A\mathbf{x} = \mathbf{0}_3$ (for $\mathbf{x} \in \mathbb{R}^2$) is just

$$\begin{cases} x_1 = 0 \\ x_2 = 0 \\ 0 = 0, \end{cases}$$

so $N(A) = \{\mathbf{0}_2\}$ once again.

18.4 Problem (!). Prove that

$$N(I_n) = \{\mathbf{0}_n\} \quad \text{and} \quad N\left(\begin{bmatrix} I_n \\ 0 \end{bmatrix}\right) = \{\mathbf{0}_n\}.$$

In the second, block matrix, you should interpret the symbol 0 as representing one or more rows of zeros. [Hint: *convince yourself that*

$$\begin{bmatrix} A \\ B \end{bmatrix} \mathbf{x} = \begin{bmatrix} A\mathbf{x} \\ B\mathbf{x} \end{bmatrix}$$

whenever $A \in \mathbb{R}^{m_1 \times n}$, $B \in \mathbb{R}^{m_2 \times n}$, and $\mathbf{x} \in \mathbb{R}^n$.]

18.5 Example. Let

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

For $\mathbf{x} \in \mathbb{R}^4$, we have $A\mathbf{x} = \mathbf{0}_2$ if and only if

$$\begin{cases} x_1 & + 2x_3 + 3x_4 = 0 \\ x_2 & + 4x_4 = 0 \end{cases}$$

This is not as nice as the square upper-triangular systems that we have previously studied. There's no equation with just one variable in it!

The right, if not immediately obvious, strategy is to solve for what we can easily solve for. The unknowns x_1 and x_2 have coefficients of 1 on them, so solving for those two variables in terms of x_3 and x_4 is easier, comparatively speaking, than solving for x_3 or x_4 . Gotta solve for something, anyway. We get

$$\begin{cases} x_1 = -2x_3 - 3x_4 \\ x_2 = -4x_4, \end{cases}$$

and if we put $\mathbf{x} = (x_1, x_2, x_3, x_4)$, then

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} -2x_3 - 3x_4 \\ -4x_4 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} -2x_3 \\ 0 \\ x_3 \\ 0 \end{bmatrix} + \begin{bmatrix} -3x_4 \\ -4x_4 \\ 0 \\ x_4 \end{bmatrix} = x_3 \begin{bmatrix} -2 \\ 0 \\ 1 \\ 0 \end{bmatrix} + x_4 \begin{bmatrix} -3 \\ -4 \\ 0 \\ 1 \end{bmatrix}.$$

Think about that for a moment. We have shown that every $\mathbf{x} \in \mathbf{N}(A)$ is a linear combination of those two vectors on the right. More compactly,

$$\mathbf{N} \left(\begin{bmatrix} 1 & 0 & 2 & 3 \\ 0 & 1 & 0 & 4 \end{bmatrix} \right) = \mathbf{C} \left(\begin{bmatrix} -2 & -3 \\ 0 & -4 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \right).$$

(Strictly speaking, we have shown that *if* \mathbf{x} is in the null space of A , *then* \mathbf{x} is in the column space of that 4×2 matrix. You should check your work and show that each column in that 4×2 matrix is in $\mathbf{N}(A)$.)

Do you see the pattern here? Our original matrix A had the block structure

$$A = \begin{bmatrix} I_2 & F \end{bmatrix}, \quad F := \begin{bmatrix} 2 & 3 \\ 0 & 4 \end{bmatrix},$$

and its null space is

$$\mathbf{N}(A) = \mathbf{C} \left(\begin{bmatrix} -F \\ I_2 \end{bmatrix} \right).$$

This can't be an accident.

Content from Strang's ILA 6E. This example is basically the same as Example 1 on p. 93. Strang calls the columns of

$$\begin{bmatrix} -2 & -3 \\ 0 & -4 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}$$

the “special solutions” for $A\mathbf{x} = \mathbf{0}_2$. What is “special” about these solutions is that they are linearly independent, and every solution to $A\mathbf{x} = \mathbf{0}_2$ is in the span of these solutions.

18.6 Problem (★). Let $n > m$ and $F \in \mathbb{R}^{m \times (n-m)}$. Prove that

$$\mathbf{N} \left(\begin{bmatrix} I_m & F \end{bmatrix} \right) = \mathbf{C} \left(\begin{bmatrix} -F \\ I_m \end{bmatrix} \right).$$

[Hint: write any $\mathbf{x} \in \mathbb{R}^n$ as $\mathbf{x} = (\mathbf{x}_m, \mathbf{x}_{n-m})$ with $\mathbf{x}_m \in \mathbb{R}^m$ and $\mathbf{x}_{n-m} \in \mathbb{R}^{n-m}$, and argue that

$$\begin{bmatrix} A & B \end{bmatrix} \mathbf{x} = \begin{bmatrix} A\mathbf{x}_m & B\mathbf{x}_{n-m} \end{bmatrix}$$

when $A \in \mathbb{R}^{m \times m}$ and $B \in \mathbb{R}^{m \times (n-m)}$.]

18.7 Problem (!). Putting more zero rows into the matrix doesn't change the null space. We saw this already in Example 18.3. Adapt the work of Example 18.5 to express the null

space of

$$\begin{bmatrix} 1 & 0 & 2 & 3 \\ 0 & 1 & 0 & 4 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

as a column space.

18.8 Problem (+). Let $1 \leq r < n$, $F \in \mathbb{R}^{r \times (n-r)}$, and

$$A = \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix}.$$

Here the two occurrences of the symbol 0 are meant to represent matrices whose entries are all the number 0. (What are the dimensions of those matrices?) Prove that

$$\mathbf{N}(A) = \mathbf{C} \left(\begin{bmatrix} -F \\ I_r \end{bmatrix} \right).$$

We are doing examples and problems in these special forms for a reason, and I'll tell you what that reason is soon. Let's pause from concrete numbers and focus on the *dynamic* aspects of the null space. Spoiler: it's the same dynamism as the column space.

Let $A \in \mathbb{R}^{m \times n}$. Suppose $\mathbf{v}_1, \mathbf{v}_2 \in \mathbf{N}(A)$. Then $A\mathbf{v}_1 = \mathbf{0}_m$ and $A\mathbf{v}_2 = \mathbf{0}_m$, so

$$A(\mathbf{v}_1 + \mathbf{v}_2) = A\mathbf{v}_1 + A\mathbf{v}_2 = \mathbf{0}_m + \mathbf{0}_m = \mathbf{0}_m.$$

Thus $\mathbf{v}_1 + \mathbf{v}_2 \in \mathbf{N}(A)$. Like the column space, the null space is “closed under addition”: adding two vectors in $\mathbf{N}(A)$ yields another vector in $\mathbf{N}(A)$.

Similarly, if $\mathbf{v} \in \mathbf{N}(A)$ and $c \in \mathbb{R}$, then since $A\mathbf{v} = \mathbf{0}_m$, we have

$$A(c\mathbf{v}) = c(A\mathbf{v}) = c\mathbf{0}_m = \mathbf{0}_m,$$

so $c\mathbf{v} \in \mathbf{N}(A)$. That is, the null space is “closed under scalar multiplication”: multiplying a vector in $\mathbf{N}(A)$ by a real number yields another vector in $\mathbf{N}(A)$.

Finally, since $A\mathbf{0}_n = \mathbf{0}_m$, we have $\mathbf{0}_n \in \mathbf{N}(A)$. Thus the null space is never empty, and it contains one of the most important vectors for vector and matrix arithmetic.

Content from Strang's ILA 6E. These properties of the null space appear in the very last paragraph of p. 88.

Subsets of \mathbb{R}^n (or \mathbb{R}^m , or whatever) that have these three properties—closure under vector addition, closure under scalar multiplication, presence of the zero vector—are just particularly “nice” for linear algebra. They respect the fundamental arithmetic and algebra that we do, and they arise often in connection with our fundamental problem of solving and understanding and approximating $A\mathbf{x} = \mathbf{b}$. So, they deserve a special name that reflects their dynamism—they are not merely sets but *spaces* of vectors that interact well together.

18.9 Definition. A subset \mathcal{V} of \mathbb{R}^p is a **SUBSPACE** of \mathbb{R}^p if the following are true.

- (i) **[Closure under vector addition]** If $\mathbf{v}, \mathbf{w} \in \mathcal{V}$, then $\mathbf{v} + \mathbf{w} \in \mathcal{V}$.
- (ii) **[Closure under scalar multiplication]** If $\mathbf{v} \in \mathcal{V}$ and $c \in \mathbb{R}$, then $c\mathbf{v} \in \mathcal{V}$.
- (iii) **[Presence of the zero vector]** $\mathbf{0}_p \in \mathcal{V}$.

Content from Strang's ILA 6E. Page 86 discusses the axioms for a subspace. Examples 1 and 2 on p. 87 present concrete (non)examples of subspaces of \mathbb{R}^p .

18.10 Example. Let $A \in \mathbb{R}^{m \times n}$.

- (i) $\mathbf{C}(A) = \{A\mathbf{v} \mid \mathbf{v} \in \mathbb{R}^n\}$ is a subspace of \mathbb{R}^m . We proved this some time ago; the important thing here is that every vector in $\mathbf{C}(A)$ has the form $A\mathbf{v} \in \mathbb{R}^m$.
- (ii) $\mathbf{N}(A) = \{\mathbf{v} \in \mathbb{R}^n \mid A\mathbf{v} = \mathbf{0}_m\}$ is a subspace of \mathbb{R}^n . It should be obvious from the definition of $\mathbf{N}(A)$ that every vector in the null space is a vector in \mathbb{R}^n .

We will eventually show that every subspace is both a column space and a null space (probably for different matrices). This is a miracle of definitions and algebra: the abstract conditions of the definition of subspace realize themselves concretely in matrices. For the purposes of this course, the only important subspaces that we will study will eventually be column and null spaces. However, there will be times when working with the three axioms for a subspace will be more convenient than representing the subspace as a particular column or null space.

18.11 Problem (!). Let

$$\mathcal{V} = \left\{ \begin{bmatrix} x_1 \\ x_2 \\ 1 \end{bmatrix} \in \mathbb{R}^3 \mid x_1, x_2 \in \mathbb{R} \right\}.$$

Explain how each of the three conditions for a subspace fails for \mathcal{V} .

Content from Strang's ILA 6E. Section 3.1 discusses the much more general, and hugely important, concept of a **VECTOR SPACE**. This is a set of elements called **VECTORS** that we can add together and multiply by scalars (real or complex numbers), and for which these operations of **VECTOR ADDITION** and **SCALAR MULTIPLICATION** basically behave the way that we expect arithmetic to behave. See the eight axioms on p. 89.

Maybe the two most important vector spaces are the column vectors with n entries, which, of course, is \mathbb{R}^n , and, from calculus, the space of continuous functions on an interval $I \subseteq \mathbb{R}$, which we denote by $\mathcal{C}(I)$. You know from calculus that if f and g are continuous on I , then so are $f + g$ and cf for any real c . (The space $\mathcal{C}(I)$ has the additional algebraic

operation of function multiplication, fg , whereas we cannot multiply vectors in \mathbb{R}^n in any “natural” way to get another vector in \mathbb{R}^n .) The r -times continuously differentiable functions (functions whose first r derivatives exist and are continuous) form the subspace $\mathcal{C}^r(I)$ of $\mathcal{C}(I)$, which is a natural player in differential equations.

The structure of vector spaces transcends matrix problems and provide the “right” framework for understanding the linear structure that pervades calculus. See pp. 84–85 for just a little on this. We will focus mostly on subspaces of \mathbb{R}^n , not general vector spaces, in this course.

Day 19: Wednesday, February 19.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Reduced row echelon form (RREF) (N), pivot column (of a matrix in RREF), free column (of a matrix in RREF)

We return to the problem of concretely describing null spaces with a new wrinkle.

19.1 Example. Let

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

We proceed as in Example 18.5: assume $A\mathbf{x} = \mathbf{0}_2$ and write this as the linear system

$$\begin{cases} x_1 + 2x_2 + 3x_4 = 0 \\ x_3 + 4x_4 = 0. \end{cases}$$

We solve for the variables with the simplest coefficients of 1; these are now x_1 and x_3 :

$$\begin{cases} x_1 = -2x_2 - 3x_4 \\ x_3 = -4x_4. \end{cases}$$

Vectorizing, we have

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} -2x_2 - 3x_4 \\ x_2 \\ -4x_4 \\ x_4 \end{bmatrix} = \begin{bmatrix} -2x_2 \\ x_2 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} -3x_4 \\ 0 \\ -4x_4 \\ x_4 \end{bmatrix} = x_2 \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix} + x_4 \begin{bmatrix} -3 \\ 0 \\ -4 \\ 1 \end{bmatrix}.$$

Thus

$$\mathbf{N}\left(\begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \end{bmatrix}\right) = \mathbf{C}\left(\begin{bmatrix} -2 & -3 \\ 1 & 0 \\ 0 & -4 \\ 0 & 1 \end{bmatrix}\right). \quad (19.1)$$

All of our previous examples and problems about finding null spaces had the identity matrix show up in a pretty obvious way. It looks like the 2×2 identity is jumbled here. How can we sort it out?

We've handled "jumbled" matrices before. Recall that a permutation matrix $P \in \mathbb{R}^{m \times m}$ is a matrix formed by reordering the columns of the $m \times m$ identity matrix. If $B \in \mathbb{R}^{m \times n}$, then PB reorders the rows of B per the ordering of the columns in P . With our A in this example, however, it's a matter of reordering *columns*. We'd be happier if the columns of the 2×2 identity matrix appeared first in A .

How can we make this happen? If multiplying on the left by a permutation matrix reorders rows, multiplying on the right reorders columns. Here $A \in \mathbb{R}^{2 \times 4}$, so if we multiply on the right by a permutation matrix P , we better have $P \in \mathbb{R}^{4 \times 4}$. What we want is

$$AP = \begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \end{bmatrix} P = \begin{bmatrix} 1 & 0 & 2 & 3 \\ 0 & 1 & 0 & 4 \end{bmatrix} = [\mathbf{a}_1 \quad \mathbf{a}_3 \quad \mathbf{a}_2 \quad \mathbf{a}_4],$$

and we know that if $P = [\mathbf{p}_1 \quad \mathbf{p}_2 \quad \mathbf{p}_3 \quad \mathbf{p}_4]$, then

$$AP = [A\mathbf{p}_1 \quad A\mathbf{p}_2 \quad A\mathbf{p}_3 \quad A\mathbf{p}_4].$$

So, we better have

$$[\mathbf{a}_1 \quad \mathbf{a}_3 \quad \mathbf{a}_2 \quad \mathbf{a}_4] = [A\mathbf{p}_1 \quad A\mathbf{p}_2 \quad A\mathbf{p}_3 \quad A\mathbf{p}_4].$$

We know that the columns of P are going to be columns of I_4 , and we know that $A\mathbf{e}_j = \mathbf{a}_j$, where \mathbf{e}_j is the j th column of I_4 , i.e., the j th standard basis vectors. All together, this says that we should take

$$P = [\mathbf{e}_1 \quad \mathbf{e}_3 \quad \mathbf{e}_2 \quad \mathbf{e}_4] = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

to find

$$\begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \end{bmatrix} P = \begin{bmatrix} 1 & 0 & 2 & 3 \\ 0 & 1 & 0 & 4 \end{bmatrix}$$

Inverting P , we have

$$\mathbf{N} \left(\begin{bmatrix} 1 & 0 & 2 & 3 \\ 0 & 1 & 0 & 4 \end{bmatrix} \right) = \mathbf{N} \left(\begin{bmatrix} 1 & 0 & 2 & 3 \\ 0 & 1 & 0 & 4 \end{bmatrix} P^{-1} \right).$$

How does this compare to what we already know from (19.1)? Example 18.5 taught us that

$$\mathbf{N} \left(\begin{bmatrix} 1 & 0 & 2 & 3 \\ 0 & 1 & 0 & 4 \end{bmatrix} \right) = \mathbf{C} \left(\begin{bmatrix} -2 & -3 \\ 0 & -4 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \right),$$

and now we have shown that

$$\mathbf{N} \left(\begin{bmatrix} 1 & 0 & 2 & 3 \\ 0 & 1 & 0 & 4 \end{bmatrix} P^{-1} \right) = \mathbf{C} \left(\begin{bmatrix} -2 & -3 \\ 1 & 0 \\ 0 & -4 \\ 0 & 1 \end{bmatrix} \right).$$

So where is the permutation on the right?

In the rows! Multiplying on the left by P interchanges rows 2 and 3, and we have

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

So here is the conclusion:

$$\mathbf{N} \left(\begin{bmatrix} 1 & 0 & 2 & 3 \\ 0 & 1 & 0 & 4 \end{bmatrix} P^{-1} \right) = \mathbf{C} \left(P \begin{bmatrix} -2 & -3 \\ 0 & -4 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \right).$$

Content from Strang's ILA 6E. This was basically Example 2 on p. 94.

19.2 Problem (+). Let $m < n$ and $F \in \mathbb{R}^{m \times (n-m)}$. Let $P \in \mathbb{R}^{n \times n}$ be invertible. Prove that

$$\mathbf{N}([I_m \ F] P) = \mathbf{C} \left(P^{-1} \begin{bmatrix} -F^T \\ I_m \end{bmatrix} \right).$$

[Hint: you want to solve $[I_m \ F] P \mathbf{x} = \mathbf{0}_m$. Put $\mathbf{y} = P \mathbf{x}$. Now you want to solve $[I_m \ F] \mathbf{y} = \mathbf{0}_m$. You know how to do this from Problem 18.6. Along the way, check that the matrix product in the column space above is actually defined.]

19.3 Problem (!). Adapt the work of Example 19.1 to express the null space of

$$\begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

as a column space.

19.4 Problem (+). (i) The previous examples have been a little too perfect in that the identity matrix that showed up in the RREF was I_2 , and the “ F -block” that showed up

was also 2×2 . Express the null spaces of

$$\begin{bmatrix} 1 & 0 & 2 & 3 & 0 \\ 0 & 1 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 & 2 & 0 & 3 & 0 \\ 0 & 0 & 1 & 4 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

as column spaces, and comment on the identity matrix that shows up in those column spaces.

(ii) Let $1 \leq r < n$ and $F \in \mathbb{R}^{r \times (n-r)}$. Let $P \in \mathbb{R}^{n \times n}$ be invertible. Prove that

$$\mathbf{N} \left(\begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} P \right) = \mathbf{C} \left(P^{-1} \begin{bmatrix} -F \\ I_{n-r} \end{bmatrix} \right).$$

As before, the symbols 0 denote matrices whose entries are all 0.

The conclusion from all of the recent problems and examples should be that the null space is “easy” to describe when the matrix under consideration has one of the following special forms:

$$I_n, \quad \begin{bmatrix} I_n \\ 0 \end{bmatrix}, \quad \begin{bmatrix} I_m & F \end{bmatrix}, \quad \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} I_m & F \end{bmatrix} P, \quad \text{or} \quad \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} P. \quad (19.2)$$

Above, P is in practice a permutation matrix (and actually a rather specific kind of permutation matrix), although the only thing that we really required in the null space calculations was the invertibility of P .

Certainly not every matrix has one of these six forms—common to all of these forms is the appearance of an identity matrix within the columns of overall matrix. But every nonzero matrix can be *reduced* to one of these forms by the elementary row operations that we already know and love—by Gauss–Jordan elimination. Once again, the major technique in computational linear algebra is putting zeros in matrices.

19.5 Example. Let

$$A = \begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix}.$$

Previously we only performed elementary row operations on square matrices, but they certainly work on nonsquare matrices, too. We compute

$$\begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix} \xrightarrow[E_{21}]{R2 \mapsto R2 - 2 \times R1} \begin{bmatrix} 1 & 2 & 1 & 7 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 8 \end{bmatrix}, \quad E_{21} := \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$\xrightarrow[P_{23}]{R2 \mapsto R3, R3 \mapsto R2} \begin{bmatrix} 1 & 2 & 1 & 7 \\ 0 & 0 & 2 & 8 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad P_{23} := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

$$\xrightarrow[D_{22}]{R2 \mapsto (1/2) \times R2} \begin{bmatrix} 1 & 2 & 1 & 7 \\ 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad D_{22} := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1/2 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$\xrightarrow[E_{12}]{R1 \mapsto R1 - R2} \begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad E_{12} := \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

That is,

$$EA = \begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad E := E_{12}D_{22}P_{23}E_{21}.$$

Now put

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad F = \begin{bmatrix} 2 & 3 \\ 0 & 4 \end{bmatrix}.$$

Then

$$EA = \begin{bmatrix} I_2 & F \\ 0 & 0 \end{bmatrix} P =: R_0$$

where the symbol 0 denotes the matrix $\begin{bmatrix} 0 & 0 \end{bmatrix}$.

We saw in Problem 19.3 that

$$\mathbf{N}(R_0) = \mathbf{N} \left(\begin{bmatrix} I_2 & F \\ 0 & 0 \end{bmatrix} P \right) = \mathbf{C} \left(P^{-1} \begin{bmatrix} -F \\ I_2 \end{bmatrix} \right) = \mathbf{C} \left(\begin{bmatrix} -2 & -3 \\ 1 & 0 \\ 0 & -4 \\ 0 & 1 \end{bmatrix} \right).$$

This is helpful here because we have $EA = R_0$ with E invertible. If $\mathbf{v} \in \mathbf{N}(A)$, then $A\mathbf{v} = \mathbf{0}_2$, so $E(A\mathbf{v}) = \mathbf{0}_m$. And then

$$\mathbf{0}_2 = (EA)\mathbf{v} = R_0\mathbf{v},$$

so $\mathbf{v} \in \mathbf{N}(R_0)$. Conversely, if $\mathbf{v} \in \mathbf{N}(R_0)$, then $R_0\mathbf{v} = \mathbf{0}_2$, so

$$A\mathbf{v} = E^{-1}R_0\mathbf{v} = E^{-1}\mathbf{0}_2 = \mathbf{0}_2.$$

Thus $\mathbf{N}(A) = \mathbf{N}(R_0)$. This is a nice auxiliary fact: multiplying on the left by an invertible matrix does not change the kernel!

All together, we conclude

$$\mathbf{N} \left(\begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix} \right) = \mathbf{C} \left(\begin{bmatrix} -2 & -3 \\ 0 & -4 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \right).$$

The example above is prototypical: Gauss–Jordan elimination “reduces” any $A \in \mathbb{R}^{m \times n}$ to a matrix of the following structure.

19.6 Definition. A matrix $R \in \mathbb{R}^{m \times n}$ is in **REDUCED ROW ECHELON FORM (RREF)** if it has the following four properties.

Row Property 1. Any nonzero row of R is below any row with nonzero entries.

Row Property 2. If a row contains nonzero entries, the first nonzero entry of that row is 1, called the **LEADING 1** or the **PIVOT** for that row.

Column Property 1. The other entries of any column containing a leading 1 are 0. That is, a column containing a leading 1 is a column of the $m \times m$ identity matrix I_m , equivalently, a standard basis vector for \mathbb{R}^m . A column containing a leading 1 is called a **PIVOT COLUMN**. A column that is not a pivot column is called a **FREE COLUMN**.

Column Property 2. If \mathbf{e}_i and \mathbf{e}_j are columns of R with $i < j$, then the first appearance of \mathbf{e}_i must occur before any appearance of \mathbf{e}_j .

19.7 Problem (!). Explain *all* of the reasons why

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$$

is not in RREF.

19.8 Problem (!). Explain why

$$\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

is in RREF and comment on the role of the adjective “first” in Column Property 2 of Definition 19.6.

Here is the fruit of Gauss–Jordan elimination.

19.9 Theorem. Let $A \in \mathbb{R}^{m \times n}$ be nonzero (i.e., A has at least one nonzero entry). There exists an invertible matrix $E \in \mathbb{R}^{m \times m}$ such that EA is in RREF with one of the following forms:

(i) I_n , in which case A is square and invertible;

- (ii) $\begin{bmatrix} I_n \\ 0 \end{bmatrix}$, in which case $n < m$ (more rows than columns);
- (iii) $\begin{bmatrix} I_m & F \end{bmatrix}$, in which case $m < n$ (more columns than rows) and $F \in \mathbb{R}^{m \times (n-m)}$;
- (iv) $\begin{bmatrix} I_m & F \end{bmatrix} P$, with the same conditions as in form (iii) and P a permutation matrix;
- (v) $\begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix}$, in which case $1 \leq r \leq \min\{m, n\}$ and $F \in \mathbb{R}^{r \times (n-r)}$;
- (vi) $\begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} P$, with the same conditions as in form (v) and P a permutation matrix.

This form is unique in the sense that if $\tilde{E} \in \mathbb{R}^{m \times m}$ is invertible with $\tilde{E}A$ in RREF, then $EA = \tilde{E}A$. We write $EA = \text{rref}(A)$ and call $\text{rref}(A)$ the **RREF** of A .

We are not going to prove this theorem in detail. Existence, again, is just Gauss–Jordan elimination. Uniqueness is surprisingly more annoying.

19.10 Problem (★). Example 19.5 constructs $E \in \mathbb{R}^{3 \times 3}$ such that

$$E \begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

(i) By revisiting the elementary row operations in that example, explain why E in Theorem 19.9 might not be unique. [Hint: *could P_{23} or D_{33} have appeared earlier or later?*]

(ii) With E from Example 19.5, find a permutation matrix \tilde{P} such that

$$E \begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 3 & 2 \\ 0 & 1 & 4 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tilde{P}.$$

Contrast this with the result of Example 19.5 and explain how this shows that P and F from Theorem 19.9 may not be unique.

(iii) Explain why there cannot exist a matrix $A \in \mathbb{R}^{3 \times 4}$ such that

$$\text{rref}(A) = \begin{bmatrix} 1 & 0 & 3 & 2 \\ 0 & 1 & 4 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Conclude that not every permutation matrix P can appear in the forms of Theorem 19.9.

(iv) Give examples of two matrices $A \neq B$ that have the same RREF. [Hint: *look no further than the first form in Theorem 19.9.*]

Content from Strang's ILA 6E. A reduction to RREF is given at the top of p. 95 and another is done in Example 2 at the bottom of the page. A third is Example 3 on pp. 97–98, and this also includes a null space calculation and remarks on the CR -factorization (which we will revisit shortly). Page 96 gives the algorithm for computing the RREF column by column. Read p. 142 up to but not including the “Factorization” box.

19.11 Problem (★). Find a matrix $A \in \mathbb{R}^{3 \times 4}$ whose entries are all nonzero such that

$$\text{rref}(A) = \begin{bmatrix} 1 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

Provide a matrix $E \in \mathbb{R}^{3 \times 3}$ such that $EA = \text{rref}(A)$; you may express E as a product of elementary matrices, and you do not have to multiply that product out.

19.12 Problem (★). For each of the six RREF forms in Theorem 19.9, find a matrix whose RREF has that form. Construct your matrix so that it has at least two rows and at least two columns and that all of its entries are nonzero. For the forms with a permutation matrix, ensure that a permutation matrix is actually needed in your form (don't just let P be the identity, which is a permutation matrix, but a boring one). Give the exact RREF of each matrix, not just the general form that it has.

Day 20: Friday, February 21.

20.1 Example. For practice with the axioms of the RREF from Definition 19.6, we construct all matrices $R \in \mathbb{R}^{3 \times 4}$ that are in RREF and that have pivot columns in columns 2 and 4 only. We proceed via the following steps.

1. Start with the first column (a very good place to start). If any entry is nonzero, that entry is the leading nonzero entry in its row (can't start earlier than the first column), and so column 1 is a pivot column. This is not allowed under the rules of our current game, so the first column is $\mathbf{0}_3$, and therefore

$$R = \begin{bmatrix} 0 & ? & ? & ? \\ 0 & ? & ? & ? \\ 0 & ? & ? & ? \end{bmatrix}.$$

2. The second column is a pivot column, so it is either \mathbf{e}_1 , \mathbf{e}_2 , or \mathbf{e}_3 . Column Property 2 basically tells us that it's \mathbf{e}_1 . Otherwise, there would be no first appearance of \mathbf{e}_1 before \mathbf{e}_2 or \mathbf{e}_3 .

Here is another way to see this. If column 2 is \mathbf{e}_2 , then

$$R = \begin{bmatrix} 0 & 0 & ? & ? \\ 0 & 1 & ? & ? \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

The other two entries in row 1 (in columns 3 and 4) can't both be 0, as that would violate Row Property 1. So, at least one of them is nonzero, thus a leading nonzero entry. But then \mathbf{e}_1 appears in column 3 or 4, again contradicting Column Property 2.

3. We now know

$$R = \begin{bmatrix} 0 & 1 & ? & ? \\ 0 & 0 & ? & ? \\ 0 & 0 & ? & ? \end{bmatrix}.$$

Look at the third column. If it has a nonzero entry in rows 2 or 3, that is the leading nonzero entry in that row, and so column 3 is a pivot column. This is not allowed in our game. However, it doesn't look like there are any restrictions on the $(1, 3)$ -entry of R , since that would not be a leading nonzero entry in row 1. Let's write

$$R = \begin{bmatrix} 0 & 1 & * & ? \\ 0 & 0 & 0 & ? \\ 0 & 0 & 0 & ? \end{bmatrix}.$$

We have upgraded the $(1, 3)$ -entry from $?$ to $*$ to emphasize that it can be any number right now, zero or not.

4. The fourth column is a pivot column, so it is \mathbf{e}_1 , \mathbf{e}_2 , or \mathbf{e}_3 . If it's \mathbf{e}_1 , then

$$\begin{bmatrix} 0 & 1 & * & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

but then the 1 in the $(1, 4)$ -entry is not the leading nonzero entry in row 1, so column 4 is not a pivot column after all. If column 4 is \mathbf{e}_3 , then

$$R = \begin{bmatrix} 0 & 1 & * & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

and that contradicts Row Property 1. The only choice left is that column 4 is \mathbf{e}_2 .

We conclude that all matrices $R \in \mathbb{R}^{3 \times 4}$ that are in RREF with pivot columns in columns 2 and 4 only have the form

$$R = \begin{bmatrix} 0 & 1 & * & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

where the $(1, 3)$ -entry is arbitrary. This is a pretty restricted family of matrices.

While our initial interest in the RREF was for the purposes of null spaces (and thus the question of uniqueness of solutions to $A\mathbf{x} = \mathbf{b}$), the RREF has many other virtues. First, it helps us prove our longstanding Conjecture 9.4 on the CR -factorization. Recall that this conjecture says that if $A \in \mathbb{R}^{m \times n}$ has r independent columns, then $A = CR$ for some

$C \in \mathbb{R}^{m \times r}$ and $R \in \mathbb{R}^{r \times n}$, where the columns of C are r independent columns of A . We'll give a "proof by example" first.

20.2 Example. Let $A \in \mathbb{R}^{3 \times 4}$ be any matrix whose RREF is

$$R_0 := \begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

Here we are writing R_0 , not R , so we can save R for the factor in the CR -factorization to come. Then there is an invertible $E \in \mathbb{R}^{3 \times 3}$ such that $EA = R_0$, and so $A = E^{-1}R_0$. Write

$$E^{-1} = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3],$$

and be aware that the columns of E^{-1} are independent.

Then

$$A = E^{-1}R_0 = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3] \begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 \end{bmatrix} = [\mathbf{v}_1 \quad 2\mathbf{v}_1 \quad \mathbf{v}_3 \quad (3\mathbf{v}_1 + 4\mathbf{v}_2)].$$

Importantly, \mathbf{v}_1 and \mathbf{v}_2 are independent columns of A ! And no other columns of A are independent along with \mathbf{v}_1 and \mathbf{v}_2 .

If we stare at this representation of A long enough, hopefully we'll see

$$A = [\mathbf{v}_1 \quad 2\mathbf{v}_1 \quad \mathbf{v}_3 \quad (3\mathbf{v}_1 + 4\mathbf{v}_2)] = [\mathbf{v}_1 \quad \mathbf{v}_2] \begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \end{bmatrix}.$$

That's our CR -factorization! Put

$$C = [\mathbf{v}_1 \quad \mathbf{v}_2] \quad \text{and} \quad R = \begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \end{bmatrix}$$

to see that $A = CR$, that C contains the independent columns of A , and that the dimensions of C and R check out. And it was no accident that we got R by chopping off the zero rows of R_0 .

20.3 Theorem (CR -factorization—Strang). Let $A \in \mathbb{R}^{m \times n}$. There exist $C \in \mathbb{R}^{m \times r}$ and $R \in \mathbb{R}^{r \times n}$ with the following properties.

- (i) $A = CR$.
- (ii) The columns of C are independent columns of A .
- (iii) The column space of A equals the column space of C : $\mathbf{C}(A) = \mathbf{C}(C)$. That is, the columns of C together with any columns of A not in C are dependent.

Proof. Let $E \in \mathbb{R}^{m \times m}$ be invertible with EA in RREF. Write

$$EA = \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} P.$$

This is the most general possible form of the RREF. We have $1 \leq r \leq \min\{m, n\}$. We allow the zero blocks to be absent (in which case $r = m$), or the F block to be absent (in which case $r = n$), or both. The matrix P is a permutation matrix, and maybe $P = I_n$.

Write $E^{-1} = \begin{bmatrix} C & \tilde{C} \end{bmatrix}$, where C contains the first r columns of E^{-1} and \tilde{C} the last $m - r$ columns. In particular, the columns of C are independent, since they are columns of the invertible matrix E^{-1} . If $m = r$, then there is no \tilde{C} block. Then

$$A = \begin{bmatrix} C & \tilde{C} \end{bmatrix} \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} P = \begin{bmatrix} C & CF \end{bmatrix} P = C \begin{bmatrix} I_r & F \end{bmatrix} P. \quad (20.1)$$

Put $R := \begin{bmatrix} I_r & F \end{bmatrix} P$ to find $A = CR$.

The second equality in (20.1) shows that the columns of A are the columns of C and of CF , since the permutation matrix P only rearranges the ordering of the columns in $\begin{bmatrix} C & CF \end{bmatrix}$ when multiplying on the right. Thus $\mathbf{C}(A) = \mathbf{C}(\begin{bmatrix} C & CF \end{bmatrix})$. Now, any column of CF has the form $C\mathbf{f}$, where \mathbf{f} is a column of F , and vector in the form $C\mathbf{f}$ is therefore a linear combination of the columns of C , and so $C\mathbf{f} \in \mathbf{C}(C)$. Thus the columns of A that are not columns of C are columns of CF , which themselves are linear combinations of the columns of C , and so $\mathbf{C}(\begin{bmatrix} C & CF \end{bmatrix}) = \mathbf{C}(C)$. In particular, $\mathbf{C}(A) = \mathbf{C}(C)$, and we do not obtain any more independent columns of A by including columns from CF . ■

Content from Strang's ILA 6E. The second half of p. 96 revisits Example 2 (from p. 95) in the context of CR . Example 3 does the same. Now read the “Review” paragraph toward the bottom of p. 97, p. 142 up to and now including the “Factorization” box, and pp. 410–411 up to but not including #3.

The RREF fundamentally tells us about null spaces: for $A \in \mathbb{R}^{m \times n}$, if

$$\text{rref}(A) = \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} P,$$

with (as in the proof above), maybe the zero blocks absent ($r = m$), or the F -block absent ($r = n$), or $P = I_n$, then

$$\mathbf{N}(A) = \mathbf{C} \left(P^{-1} \begin{bmatrix} -F \\ I_{n-r} \end{bmatrix} \right).$$

Interpret this, maybe bizarrely, as $\mathbf{N}(A) = \{\mathbf{0}_n\}$ when $r = n$, as then the F -block is absent and I_0 just doesn't make sense. This affords us superb control over the null space in the abstract (although for toy problems that appear on problem sets, quizzes, or exams, this expression for $\mathbf{N}(A)$ may not be worth memorizing).

Content from Strang's ILA 6E. This expression for the null space appears in the box on p. 97 and the subsequent “Review” paragraph. You should now be able to understand and appreciate *all* of p. 142 and *all* of pp. 410–411, including #3. Interpret P^T as P^{-1} for now, as we haven't talked about transposes.

Because of the CR -factorization, it seems that the RREF can also tell us about a matrix's independent columns. The downside to computing the CR -factorization and extracting those independent columns is that everything hinges on that matrix E that performs the Gauss–Jordan elimination. The independent columns of A are some columns from E^{-1} . Computing E^{-1} could be nasty, and in the past (with the LU -factorization), we always preferred *not* to compute it but leave it factored as a product of elementary matrices. There is a more transparent way to extract the independent columns of A directly from knowledge of the RREF.

Day 21: Monday, February 24.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Pivot column (of a matrix in general), free column (of a matrix in general), rank of a matrix (in terms of number of pivot columns)

Our question is now what $\text{rref}(A)$ tells us about $\mathbf{C}(A)$ for $A \in \mathbb{R}^{m \times n}$. We know how to use $\text{rref}(A)$ to control $\mathbf{N}(A)$ and therefore understand uniqueness of solutions to $A\mathbf{x} = \mathbf{b}$. How can we go back to our fundamental question of *existence*?

To talk sensibly, we need some new vocabulary, which relies on Definition 19.6.

21.1 Definition. Column j of a matrix is a **PIVOT COLUMN** of that matrix if column j of the RREF is a pivot column, and otherwise column j is a **FREE COLUMN** if column j in the RREF is a free column.

21.2 Example. Let

$$A = \begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix}.$$

A byproduct of Example 19.5 was that

$$\text{rref}(A) = \begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \\ 0 & 0 & 0 & 0 \end{bmatrix} = R_0.$$

The pivot columns of R_0 are columns 1 and 3, because they contain the leading 1's; the

free columns are columns 2 and 4, because they are not pivot columns. The pivot columns of A are therefore columns 1 and 3 as well; the free columns are columns 2 and 4.

I think it's clear that the pivot columns of R_0 are independent, since they are \mathbf{e}_1 and \mathbf{e}_2 . It may not be too hard to see that the pivot columns of A are independent, thanks to that 0 in the $(3, 1)$ -entry of A . I claim that we can see the independence of the pivot columns of A directly from the RREF without using any particular knowledge of the entries of A .

Write $A = [\mathbf{a}_1 \ \mathbf{a}_2 \ \mathbf{a}_3 \ \mathbf{a}_4]$. We want to show that \mathbf{a}_1 and \mathbf{a}_3 are independent, so assume $c_1\mathbf{a}_1 + c_2\mathbf{a}_3 = \mathbf{0}_3$. The goal is $c_1 = c_2 = 0$. We know $EA = R_0$ for some invertible $E \in \mathbb{R}^{3 \times 3}$, so we have

$$E(c_1\mathbf{a}_1 + c_2\mathbf{a}_3) = E\mathbf{0}_3,$$

and therefore

$$c_1E\mathbf{a}_1 + c_2E\mathbf{a}_3 = \mathbf{0}_3,$$

and therefore

$$c_1\mathbf{e}_1 + c_2\mathbf{e}_2 = \mathbf{0}_3,$$

since $E\mathbf{a}_1 = \mathbf{e}_1$ and $E\mathbf{a}_3 = \mathbf{e}_2$. The independence of \mathbf{e}_1 and \mathbf{e}_2 therefore forces $c_1 = c_2 = 0$.

I also claim that \mathbf{a}_2 and \mathbf{a}_4 are dependent columns of A in the sense that $\mathbf{a}_2 \in \text{span}(\mathbf{a}_1)$ and $\mathbf{a}_4 \in \text{span}(\mathbf{a}_1, \mathbf{a}_3)$. That is, the free columns are in the span of the *preceding* pivot columns. Okay, that $\mathbf{a}_2 \in \text{span}(\mathbf{a}_1)$ is pretty obvious from looking at the entries of A , but, again, say that we didn't know the exact entries of A , just the RREF. With

$$R_0 = [\mathbf{e}_1 \ 2\mathbf{e}_1 \ \mathbf{e}_2 \ (3\mathbf{e}_1 + 4\mathbf{e}_2)],$$

we have

$$\mathbf{a}_2 = E^{-1}(2\mathbf{e}_1) = 2E^{-1}\mathbf{e}_1 = 2\mathbf{a}_1$$

and

$$\mathbf{a}_4 = E^{-1}(3\mathbf{e}_1 + 4\mathbf{e}_2) = 3E^{-1}\mathbf{e}_1 + 4E^{-1}\mathbf{e}_2 = 3\mathbf{a}_1 + 4\mathbf{a}_3.$$

That $\mathbf{a}_4 = 3\mathbf{a}_1 + 4\mathbf{a}_3$ may not have been so obvious from the entries of A .

Since $\mathbf{a}_2, \mathbf{a}_4 \in \text{span}(\mathbf{a}_1, \mathbf{a}_3)$, we conclude

$$\mathbf{C}(A) = \text{span}(\mathbf{a}_1, \mathbf{a}_3).$$

And since \mathbf{a}_1 and \mathbf{a}_3 are independent, we shouldn't be able to pare this span down any further—this *is* the most efficient way to write $\mathbf{C}(A)$ as a span of (some) columns of A .

21.3 Problem (!). Use Example 21.2 to explain why $\mathbf{C}(A) \neq \mathbf{C}(\text{rref}(A))$ in general.

21.4 Remark. In the problem $A\mathbf{x} = \mathbf{b}$, call the unknown x_j a **PIVOT VARIABLE** of A if column j of A is a pivot column, and call x_j a **FREE VARIABLE** if column j of A is a free variable. We can phrase the computational procedures of Examples 18.5, 19.1, and 19.5 as follows. To determine $\mathbf{N}(A)$, put A in RREF as R_0 and use the equation $R_0\mathbf{x} = \mathbf{0}_m$ to solve for the pivot variables in terms of the free variables.

Example 21.2 illustrates a number of important truths about pivot and free columns, which we summarize in the next theorem. We will not prove this theorem, as the strategy of proof is essentially that of the example. The results here mostly fall under what people mean when they say “Elementary row operations preserve linear (in)dependent relations among the columns of a matrix”—reducing A to its RREF doesn’t change how columns are dependent or independent.

21.5 Theorem. *Let $A \in \mathbb{R}^{m \times n}$.*

- (i) *The pivot columns of A are independent.*
- (ii) *The free columns of A are in the span of the pivot columns. More precisely, if the j th column of A is a free column, then it is in the span of the pivot columns that appear in the first $j - 1$ columns of A . (In fact, the weights in the expression of column j of A as a linear combination of preceding pivot columns are the same as the weights in the expression of column j of $\text{rref}(A)$ as a linear combination of the preceding pivot columns.) If the first column of A ($j = 1$) is free, then it is zero.*
- (iii) *The column space of A is the span of the pivot columns of A .*

21.6 Remark. *It’s not really fair to say that the independent columns of A are the pivot columns of A . The second and third columns of A from Example 21.2 are independent, as are the third and fourth, and the first and fourth. In fact, any pair of columns from that A is independent except for columns 1 and 2. Rather, what we are looking for is the simplest way to find independent columns of A —and that comes from taking just the pivot columns.*

The number of pivot columns in a matrix is a key piece of data for that matrix. Previously (Definition 7.7) we defined the **RANK** of a matrix as the length of the longest list of linearly independent columns in that matrix. Nothing wrong with that, but that’s not very efficient. Here is a more meaningful definition.

21.7 Definition. *The **RANK** of a matrix A , denoted $\text{rank}(A)$, is the number of pivot columns in A .*

We will check in a moment that this definition of rank agrees with the old one, but first here are some useful bounds on rank.

21.8 Theorem. *Let $A \in \mathbb{R}^{m \times n}$. Then $0 \leq \text{rank}(A) \leq \min\{m, n\}$, with $\text{rank}(A) = 0$ only when A is the zero matrix.*

Proof. If $\text{rank}(A) = 0$, then A has no pivot columns, so $\text{rref}(A)$ has no pivot columns and therefore is the zero matrix. Otherwise, $\text{rref}(A)$ would have a leading nonzero entry in some row and thus a pivot column. Then since $A = E^{-1}\text{rref}(A)$ for some invertible $E \in \mathbb{R}^{m \times m}$, A is also the zero matrix.

Next, A has n columns, so at most n of them can be pivot columns. Thus $\text{rank}(A) \leq n$.

And A has m rows, so $\text{rref}(A)$ has m rows, and therefore $\text{rref}(A)$ can have at most m leading 1's, thus at most m pivot columns. And so $\text{rank}(A) \leq m$. ■

21.9 Problem (!). Let $R \in \mathbb{R}^{m \times n}$ be one of the “canonical” RREF forms from Theorem 19.9. What is the rank of R ? (Your answer will involve the numbers m , n , and/or r .)

We should probably check that our new definition of rank agrees with the old. Suppose that $A \in \mathbb{R}^{m \times n}$ has r pivot columns. If $r = 0$, then A is the zero matrix and so has no pivot columns.

Next suppose that $1 \leq r \leq \min\{m, n\}$. If $r = n$, then any list of more than r columns in A will contain at least one repeated column. Such a list is dependent.

21.10 Problem (★). Explain why. To get yourself started, let $\mathbf{v}, \mathbf{w} \in \mathbb{R}^m$ and explain why the matrix $A = [\mathbf{v} \ \mathbf{w} \ \mathbf{v}]$ has dependent columns. [Hint: *think about* $A\mathbf{x} = \mathbf{0}_m$.]

Now suppose $1 \leq r < n$. Consider any list of columns from A of length $r + 1$ or greater. Here are some subcases. If any column appears two or more times in the list, then the list is dependent, whether or not the list contains any pivot columns. If r columns in the list are the r distinct pivot columns, then since the list has at least one more column, that column must be a free column, thus a linear combination of those r pivot columns, thus the list is dependent. But what if not all of the pivot columns are in the list? (This could happen, say, for a matrix with only one pivot column and three free columns.) What if none of the pivot columns are in the list? I think there's still a gap in our argument.

Let's leave it as a conjecture and do something concrete for a change.

21.11 Conjecture. Let $A \in \mathbb{R}^{m \times n}$ have rank r . Then any list of $r + 1$ or more columns of A is dependent.

21.12 Example. Let

$$A = \begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix}.$$

We studied the null space of A in Example 19.5 and the pivot and free columns of A in Example 21.2. In particular, the latter example showed that columns 1 and 3 are the only pivot columns, so $\text{rank}(A) = 2$.

Now we study the general problem $A\mathbf{x} = \mathbf{b}$ for $\mathbf{b} = (b_1, b_2, b_3) \in \mathbb{R}^3$ arbitrary. For a small problem like this, the most efficient thing to do is to put the augmented matrix $[A \ \mathbf{b}]$ into RREF in the form $[R_0 \ E\mathbf{b}]$, where $R_0 = \text{rref}(A)$ and $EA = R_0$ with E invertible.

We basically repeat the steps of Example 19.5, where we were taking $\mathbf{b} = \mathbf{0}_3$ throughout. This time, however, we don't repeat the elementary matrices that do all the elimination.

We have

$$\begin{aligned}
 \left[\begin{array}{cccc|c} 1 & 2 & 1 & 7 & b_1 \\ 2 & 4 & 2 & 14 & b_2 \\ 0 & 0 & 2 & 8 & b_3 \end{array} \right] &\xrightarrow{R2 \mapsto R2 - 2 \times R1} \left[\begin{array}{cccc|c} 1 & 2 & 1 & 7 & b_1 \\ 0 & 0 & 0 & 0 & b_2 - 2b_1 \\ 0 & 0 & 2 & 8 & b_3 \end{array} \right] \\
 &\xrightarrow{R2 \mapsto R3, R3 \mapsto R2} \left[\begin{array}{cccc|c} 1 & 2 & 1 & 7 & b_1 \\ 0 & 0 & 2 & 8 & b_3 \\ 0 & 0 & 0 & 0 & b_2 - 2b_1 \end{array} \right] \\
 &\xrightarrow{R2 \mapsto (1/2) \times R2} \left[\begin{array}{cccc|c} 1 & 2 & 1 & 7 & b_1 \\ 0 & 0 & 1 & 4 & b_3/2 \\ 0 & 0 & 0 & 0 & b_2 - 2b_1 \end{array} \right] \\
 &\xrightarrow{R1 \mapsto R1 - R2} \left[\begin{array}{cccc|c} 1 & 2 & 0 & 3 & b_1 - b_3/2 \\ 0 & 0 & 1 & 4 & b_3/2 \\ 0 & 0 & 0 & 0 & b_2 - 2b_1 \end{array} \right].
 \end{aligned}$$

Then $A\mathbf{x} = \mathbf{b}$ is equivalent to

$$\begin{cases} x_1 + 2x_2 + 3x_4 = b_1 - b_3/2 \\ x_3 + 4x_4 = b_3/2 \\ 0 = b_2 - 2b_1. \end{cases}$$

The third equation is a “solvability condition”: if $A\mathbf{x} = \mathbf{b}$, then we must have $b_2 = 2b_1$. This is not the first time that we’ve seen this condition, and hopefully it’s apparent from the row structure of A (the second row is twice the first row). If this condition is met, then the first two equations allow us to solve for x_1 and x_3 easily in terms of x_2 and x_4 . We did this in Example 19.5, but now we have the extra baggage of \mathbf{b} . Anyway, we get

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} (b_1 - b_3/2) - 2x_2 - 3x_4 \\ x_2 \\ b_3/2 - 4x_4 \\ x_4 \end{bmatrix} = \begin{bmatrix} b_1 - b_3/2 \\ 0 \\ b_3/2 \\ 0 \end{bmatrix} + x_2 \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix} + x_4 \begin{bmatrix} -3 \\ 0 \\ -4 \\ 1 \end{bmatrix}.$$

Assuming $b_2 = 2b_1$ and taking $x_2 = x_4 = 0$, we conclude that one solution to $A\mathbf{x} = \mathbf{b}$ is $\mathbf{x}_\star := (b_1 - b_3/2, 0, b_3/2, 0)$, while all other solutions are $\mathbf{x} = \mathbf{x}_\star + c_1\mathbf{z}_1 + c_2\mathbf{z}_2$, where we recognize $\mathbf{N}(A) = \mathbf{C}([\mathbf{z}_1 \ \mathbf{z}_2])$ from Example 19.5. By the way, taking $\mathbf{b} = \mathbf{0}$ (i.e., $b_1 = b_2 = b_3 = 0$), we recover the null space calculation from that example. This structural pattern in the solution is, like so many other things in this course, no accident.

21.13 Theorem. Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. Suppose that $\mathbf{x}_\star \in \mathbb{R}^n$ satisfies $A\mathbf{x}_\star = \mathbf{b}$. Then any other solution $\mathbf{x} \in \mathbb{R}^n$ to $A\mathbf{x} = \mathbf{b}$ has the form $\mathbf{x} = \mathbf{x}_\star + \mathbf{z}$ for some $\mathbf{z} \in \mathbf{N}(A)$.

21.14 Problem (!). Prove it. [Hint: *what does $\mathbf{x} - \mathbf{x}_\star$ do?*]

It looks like we have a “decomposition” from Theorem 21.13 for solutions to $A\mathbf{x} = \mathbf{b}$. Any solution \mathbf{x} is the sum of one “particular” solution and a vector in the null space. As in so many other places in the course, we need to build some more tools, but eventually we will be able to say a bit more about what that “particular” solution is doing, and maybe how to choose it best when we have many options.

Content from Strang’s *ILA 6E*. Read all of p. 104–105.

Day 22: Wednesday, February 26.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Matrix with full column rank (N), matrix with full row rank (N)

It probably won’t hurt to do another concrete computational example.

22.1 Example. Let

$$A = \begin{bmatrix} 1 & 2 & 0 \\ 2 & 4 & 0 \\ 1 & 2 & 2 \\ 7 & 14 & 8 \end{bmatrix}.$$

This is, of course, the transpose of the matrix from Examples 19.5 and 21.2, and we’ll talk about transposes in detail someday. For now, we study $A\mathbf{x} = \mathbf{b}$, $\mathbf{b} = (b_1, b_2, b_3, b_4)$:

$$[A \quad \mathbf{b}] = \left[\begin{array}{ccc|c} 1 & 2 & 0 & b_1 \\ 2 & 4 & 0 & b_2 \\ 1 & 2 & 2 & b_3 \\ 7 & 14 & 8 & b_4 \end{array} \right] \xrightarrow{R_2 \mapsto R_2 - 2 \times R_1} \left[\begin{array}{ccc|c} 1 & 2 & 0 & b_1 \\ 0 & 0 & 0 & b_2 - 2b_1 \\ 1 & 2 & 2 & b_3 \\ 7 & 14 & 8 & b_4 \end{array} \right]$$

$$\xrightarrow{R_3 \mapsto R_3 - R_1} \left[\begin{array}{ccc|c} 1 & 2 & 0 & b_1 \\ 0 & 0 & 0 & b_2 - 2b_1 \\ 0 & 0 & 2 & b_3 - b_1 \\ 7 & 14 & 8 & b_4 \end{array} \right]$$

$$\xrightarrow{R_4 \mapsto R_4 - 7 \times R_1} \left[\begin{array}{ccc|c} 1 & 2 & 0 & b_1 \\ 0 & 0 & 0 & b_2 - 2b_1 \\ 0 & 0 & 2 & b_3 - b_1 \\ 0 & 0 & 8 & b_4 - 7b_1 \end{array} \right]$$

$$\begin{aligned}
& \xrightarrow{R4 \mapsto (1/8) \times R4} \left[\begin{array}{ccc|c} 1 & 2 & 0 & b_1 \\ 0 & 0 & 0 & b_2 - 2b_1 \\ 0 & 0 & 2 & b_3 - b_1 \\ 0 & 0 & 1 & (b_4 - 7b_1)/8 \end{array} \right] \\
& \xrightarrow{R3 \mapsto R3 - 2 \times R4} \left[\begin{array}{ccc|c} 1 & 2 & 0 & b_1 \\ 0 & 0 & 0 & b_2 - 2b_1 \\ 0 & 0 & 0 & b_3 - b_1 - (b_4 - 7b_1)/4 \\ 0 & 0 & 1 & (b_4 - 7b_1)/8 \end{array} \right] \\
& \xrightarrow{R2 \mapsto R4, R4 \mapsto R2} \left[\begin{array}{ccc|c} 1 & 2 & 0 & b_1 \\ 0 & 0 & 1 & (b_4 - 7b_1)/8 \\ 0 & 0 & 0 & b_3 - b_1 - (b_4 - 7b_1)/4 \\ 0 & 0 & 0 & b_2 - 2b_1 \end{array} \right].
\end{aligned}$$

A byproduct of this calculation is that

$$\text{rref}(A) = \begin{bmatrix} 1 & 2 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

and so columns 1 and 3 of A are pivot columns, and $\text{rank}(A) = 2$.

The problem $A\mathbf{x} = \mathbf{b}$ is then equivalent to

$$\left\{ \begin{array}{rcl} x_1 + 2x_2 & = & b_1 \\ & x_3 & = (b_4 - 7b_1)/8 \\ & 0 & = b_3 - b_1 - (b_4 - 7b_1)/4 \\ & 0 & = b_2 - 2b_1 \end{array} \right.$$

We now have *two* solvability conditions:

$$b_3 - b_1 - \frac{b_4 - 7b_1}{4} = 0 \quad \text{and} \quad b_2 = 2b_1.$$

The second cleans up nicely to just $b_3 + 3b_1/4 - b_4/4 = 0$ (check it), and so the two solvability conditions are

$$b_2 = 2b_1 \quad \text{and} \quad b_4 = 3b_1 + 4b_3.$$

If these are met, then the solution \mathbf{x} has the form

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} b_1 - 2x_2 \\ x_2 \\ (b_4 - 7b_1)/8 \end{bmatrix} = \begin{bmatrix} b_1 \\ 0 \\ (b_4 - 7b_1)/8 \end{bmatrix} + \begin{bmatrix} -2x_2 \\ x_2 \\ 0 \end{bmatrix} = \begin{bmatrix} b_1 \\ 0 \\ (b_4 - 7b_1)/8 \end{bmatrix} + x_2 \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix}.$$

A byproduct of this calculation is that

$$\mathbf{N}(A) = \mathbf{C} \left(\begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix} \right) = \text{span} \left(\begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix} \right).$$

22.2 Problem (!). Interpret this null space calculation in light of Problem 19.4. [Hint: here $n - r = 1$ and the 1×1 identity matrix is just 1, or $[1]$, if you must.]

The previous two examples involved matrices $A \in \mathbb{R}^{m \times n}$ with $1 \leq \text{rank}(A) < \min\{m, n\}$. Some interesting things happen in the “extreme” case of $\text{rank}(A) = \min\{m, n\}$.

22.3 Example. (i) Here are two matrices $A \in \mathbb{R}^{2 \times n}$ with $n \geq 2$ and $\text{rank}(A) = 2$:

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}.$$

I hope you see that each matrix has enough columns to span \mathbb{R}^2 , so in each case $\mathbf{C}(A) = \mathbb{R}^2$. That is, we can always solve $A\mathbf{x} = \mathbf{b}$. However, in the second case, the third column is not a pivot column, so $\mathbf{N}(A) \neq \{\mathbf{0}_3\}$, and in that case solutions to $A\mathbf{x} = \mathbf{b}$ aren’t unique. We knew that anyway from Corollary 22.7 since the second matrix has more columns than rows.

(ii) Here are two matrices $A \in \mathbb{R}^{m \times 2}$ with $m \geq 2$ and $\text{rank}(A) = 2$:

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}.$$

Same deal as before with I_2 , of course, but in the second case, while we can’t always solve $A\mathbf{x} = \mathbf{b}$, we can always do so uniquely, because every column is a pivot column, and so $\mathbf{N}(A) = \{\mathbf{0}_2\}$.

Content from Strang’s ILA 6E. Read Example 1 on pp. 105–106 and Example 2 on p. 107.

Here is what these examples teach us.

22.4 Theorem. Let $A \in \mathbb{R}^{m \times n}$.

(i) Suppose that $m \leq n$ and $\text{rank}(A) = m$. Then $\mathbf{C}(A) = \mathbb{R}^m$. That is, we can always solve $A\mathbf{x} = \mathbf{b}$ for any $\mathbf{b} \in \mathbb{R}^m$. In this case, we say that A has **FULL ROW RANK**.

(ii) Suppose that $n \leq m$ and $\text{rank}(A) = n$. Then $\mathbf{N}(A) = \{\mathbf{0}_n\}$. That is, if we can solve

$A\mathbf{x} = \mathbf{b}$ for some $\mathbf{b} \in \mathbb{R}^m$, then the solution \mathbf{x} is unique. In this case, we say that A has **FULL COLUMN RANK**.

Proof. (i) Let $R_0 = \text{rref}(A)$ and let $E \in \mathbb{R}^{m \times m}$ be invertible with $EA = R_0$. Let $\mathbf{b} \in \mathbb{R}^m$. Then $A\mathbf{x} = \mathbf{b}$ if and only if $EA\mathbf{x} = E\mathbf{b}$, so if and only if $R_0\mathbf{x} = E\mathbf{b}$. Abbreviate $\mathbf{c} = E\mathbf{b}$, so the goal is to solve $R_0\mathbf{x} = \mathbf{c}$.

For example, suppose that $A \in \mathbb{R}^{3 \times 6}$ with $\text{rank}(A) = 3$ and

$$R_0 = [\mathbf{0} \quad \mathbf{e}_1 \quad \mathbf{f}_1 \quad \mathbf{e}_2 \quad \mathbf{e}_3 \quad \mathbf{f}_2]$$

with the first column and \mathbf{f}_1 and \mathbf{f}_2 all free. Abbreviate $\mathbf{c} = E\mathbf{b}$. Then we want to solve

$$[\mathbf{0} \quad \mathbf{e}_1 \quad \mathbf{f}_1 \quad \mathbf{e}_2 \quad \mathbf{e}_3 \quad \mathbf{f}_2] \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} = \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}.$$

One way to do this is just to let the standard basis vectors show up in the end and assign those vectors the weights c_1, c_2, c_3 . We can do this by taking $x_2 = c_1$, $x_4 = c_2$, and $x_5 = c_3$ and putting $x_1 = x_3 = x_6 = 0$. That is, $\mathbf{x} = (0, c_1, 0, c_2, c_3, 0)$.

More generally, if $R_0 \in \mathbb{R}^{m \times n}$ has m pivot columns and thus contains all m columns of I_m at least once, define \mathbf{x} by taking x_j to be the j th entry of $E\mathbf{b}$ if the j th column of R_0 is the pivot column \mathbf{e}_j , and otherwise let x_j be 0. Then $R_0\mathbf{x} = \mathbf{c}$. By the work above, $A\mathbf{x} = \mathbf{b}$.

(ii) In this case, every column of A is a pivot column, so all of the columns of A are independent, and therefore $\mathbf{N}(A) = \{\mathbf{0}_n\}$. ■

Content from Strang's ILA 6E. Read the rest of p. 106 starting from “This example is typical. . .” Then read all of p. 108.

22.5 Problem (★). Make and fill in a table with the following five columns. The first column contains the six forms for the RREF from Theorem 19.9. The second column contains two matrices: one a matrix whose RREF has that form (follow the guidelines of—and feel free to reuse your matrices from—Problem 19.12) and the other the exact RREF of that matrix. The third column is “Existence”; put “Always” or “Sometimes” depending on whether solutions to $A\mathbf{x} = \mathbf{b}$ always exist or only sometimes exist when A has that RREF form. The fourth column is “Uniqueness”; put “Always” or “Never” depending on whether solutions to $A\mathbf{x} = \mathbf{b}$ are unique or not when A has that RREF form. (Why is there no “Sometimes” for uniqueness?) The fifth column is “Rank”; specify how the rank relates to m and/or n (be more precise than $\text{rank}(A) \leq \min\{m, n\}$), and in particular indicate any RREF form corresponding to full column or row rank.

22.6 Problem (★). Here is a happy example of how rank is “stable.” Let $A \in \mathbb{R}^{m \times n}$ be any matrix and let $B \in \mathbb{R}^{m \times m}$ be invertible. Explain why $\text{rank}(BA) = \text{rank}(A)$, but give an example to show that $\mathbf{C}(BA) \neq \mathbf{C}(A)$ in general. [Hint: *how do the pivot columns of A become the pivot columns of BA when B is invertible? Think about spans and independence.*]

Here is one more nasty little consequence of rank. Let $A \in \mathbb{R}^{m \times n}$ with $n > m$, so there are more columns than rows, and the problem $A\mathbf{x} = \mathbf{b}$ has more unknowns than equations. We probably expect that if we can solve this problem, the solution should not be unique—there is too much “freedom” in the problem since the unknowns outnumber the equations telling them what to do.

And there is. The issue is that with $n > m$, not every column of A can be a pivot column, since $\text{rank}(A) \leq \min\{m, n\} = m < n$. So, some columns of A are free, and therefore the columns of A are dependent. Thus $\mathbf{N}(A) \neq \{\mathbf{0}_m\}$. This is disappointingly robust enough to stand on its own.

22.7 Theorem. If $A \in \mathbb{R}^{m \times n}$ with $n > m$ (more columns than rows), then $\mathbf{N}(A) \neq \{\mathbf{0}_m\}$. (Equivalently, and importantly, any list of $n > m$ vectors in \mathbb{R}^m is dependent.) In particular, if a solution to $A\mathbf{x} = \mathbf{b}$ exists, then it is never unique.

Content from Strang’s ILA 6E. This is discussed in the “Important” box on p. 98 and the two paragraphs preceding that. Now read Example 4 on p. 98.

It looks like we have accomplished the major goal of the course: solve $A\mathbf{x} = \mathbf{b}$ and understand when we can’t. We make the augmented matrix $[A \ \mathbf{b}]$ and use Gauss–Jordan elimination to reduce A to $\text{rref}(A) = R_0$ with \mathbf{b} transforming into \mathbf{c} along the way. Then we study $[R_0 \ \mathbf{c}]$. That is, $A\mathbf{x} = \mathbf{b}$ and $R_0\mathbf{x} = \mathbf{c}$ have the same solutions (if any). For the problem to have a solution, if the i th row of R_0 is all zero, then the i th entry of \mathbf{c} must be 0. Assuming those “solvability conditions” to be true, we then rewrite the system $R_0\mathbf{x} = \mathbf{c}$ as a system of equations and solve for the “free variables” in terms of the “pivot variables” (recall Remark 21.4). In the special case that A is square, we could just do Gaussian elimination to convert A to its upper-triangular form U ; if all of the diagonal entry of U are nonzero, then we can back-substitute. If a diagonal entry is 0, then it’s probably best to go all the way to RREF to have some control over the null space.

Content from Strang’s ILA 6E. Read all of the “Worked Examples” on pp. 109–110.

Considering all of this good work, I claim that we are now pretty good at solving $A\mathbf{x} = \mathbf{b}$ (especially when A is square and invertible), but we could still be better at *understanding* $A\mathbf{x} = \mathbf{b}$, particularly at understanding *failure*

Question 1. If a solution is not unique, and therefore there are infinitely many solutions, can we quantify how many “different” solutions there are beyond “infinitely many?”

Question 2. If $\mathbf{N}(A) \neq \{\mathbf{0}_n\}$, can we quantify how many “degrees of freedom” the null space gives to the problem?

Question 3. If there is no solution for some \mathbf{b} , and therefore $\mathbf{C}(A) \neq \mathbb{R}^m$, can we quantify and qualify how much of \mathbb{R}^m the column space “misses”?

In short, can we understand more about the structure of $\mathbf{N}(A)$ and $\mathbf{C}(A)$?

We basically have the answer with $\mathbf{C}(A)$ already. We know that the pivot columns of A are independent and span $\mathbf{C}(A)$. If we remove a pivot column from consideration, then the remaining pivot columns won’t span $\mathbf{C}(A)$, because otherwise the missing pivot column would be in the span of those remaining pivot columns. Then all of the pivot columns together would be dependent. This begs some other questions related to the rank, one of which we raised in Conjecture 21.11. Suppose that $A \in \mathbb{R}^{m \times n}$ has rank r .

Question 4. Do *any* r independent columns of A (pivot or not) span $\mathbf{C}(A)$? In other words, do we really have to use the pivot columns to control $\mathbf{C}(A)$ efficiently?

Question 5. Can fewer than r columns of A span $\mathbf{C}(A)$? We know that fewer than r *pivot* columns of A can’t span $\mathbf{C}(A)$. Can we somehow beat the pivot columns?

Question 6. Can we describe $\mathbf{N}(A)$ as a span of some independent vectors (in \mathbb{R}^n) or, equivalently, as a column space of a matrix with full column rank? We know well that the null space is *always* a column space by now, just a column space of a slightly complicated block matrix. If A doesn’t have full column rank and therefore $\mathbf{N}(A) \neq \{\mathbf{0}_n\}$, how does the rank of A show up in understanding $\mathbf{N}(A)$?

To answer these questions (and prove Conjecture 21.11), we need some new tools.

Day 23: Friday, February 28.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Basis (for a subspace of \mathbb{R}^p)

Content from Strang’s ILA 6E. You should be very comfortable with the notion of independence by now. Read pp. 115–117 thoroughly. Think carefully about the “guilty” remark at the end of p. 117. Which way of saying that the columns of $A \in \mathbb{R}^{m \times n}$ are independent feels easier to you—that $A\mathbf{x} = \mathbf{0}_m$ forces $\mathbf{x} = \mathbf{0}_n$ (the “democratic” way) or that one column of A is a combination of other columns (the “guilty” way)?

You should also be very comfortable with the notion of span by now. Read “Vectors that Span a Subspace” at the start of p. 118.

23.1 Example. Let

$$A = \begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix}.$$

We know from long experience (Examples 19.5, 21.2, 21.12) that columns 1 and 3 are pivot columns and therefore are independent and span $\mathbf{C}(A)$ and that

$$\mathbf{N}(A) = \mathbf{C} \left(\begin{bmatrix} -2 & -3 \\ 1 & 0 \\ 0 & -4 \\ 0 & 1 \end{bmatrix} \right).$$

The columns of this matrix giving $\mathbf{N}(A)$ are also independent, for if

$$x_1 \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix} + x_2 \begin{bmatrix} -3 \\ 0 \\ -4 \\ 1 \end{bmatrix} = \mathbf{0}_4,$$

then

$$\begin{bmatrix} * \\ x_1 \\ * \\ x_2 \end{bmatrix} = \mathbf{0}_4,$$

thus $x_1 = x_2 = 0$. Even though the arithmetic was simple, I intentionally wrote $*$ to show how much I don't care about the other entries of that linear combination beside the "special" ones that just reveal x_1 and x_2 .

The result is the same. While $\mathbf{C}(A)$ and $\mathbf{N}(A)$ are very different spaces, we have described them in the same way: as column spaces of matrices with independent columns, as spans of lists of independent vectors. The vectors in these lists are not unique. Example 21.12 showed that $\mathbf{b} = (b_1, b_2, b_3) \in \mathbf{C}(A)$ if and only if $b_2 = 2b_1$. That is,

$$\mathbf{b} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ 2b_1 \\ b_3 \end{bmatrix} = b_1 \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix} + b_3 \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$

Thus

$$\mathbf{C}(A) = \text{span} \left(\begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right) = \mathbf{C} \left(\begin{bmatrix} 1 & 0 \\ 2 & 0 \\ 0 & 1 \end{bmatrix} \right).$$

The vectors in this span are also independent (check it), but only one is a column of A .

The new result is that we have written $\mathbf{C}(A)$ as a span of independent vectors which were not all columns of A . But the old result is that we still only needed two vectors to do it.

Here is the pattern that we should be seeing: writing column spaces and null spaces as spans of independent vectors is an efficient way of describing them. It turns out that we can always do this, which should not surprise us. The RREF teaches us that the column space is always the span of the pivot columns, which are always independent. The RREF also teaches us that we can write the null space as the column space of a “special” kind of matrix, whose columns always turn out to be independent. (Depending on the form of the RREF, there are six forms that the null space representation can take. It’s okay if this feels annoying. But bear in mind that saying “the column space is always the span of the pivot columns” is actually much vaguer than those forms for the null space!)

23.2 Definition. Let \mathcal{V} be a subspace of \mathbb{R}^p . A list of vectors $\mathbf{v}_1, \dots, \mathbf{v}_d \in \mathcal{V}$ is a **BASIS** for \mathcal{V} if the following hold.

- (i) The vectors $\mathbf{v}_1, \dots, \mathbf{v}_d$ are independent.
- (ii) $\mathcal{V} = \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_d)$.

23.3 Problem (!). Let \mathcal{V} be a subspace of \mathbb{R}^p . Prove that $\mathbf{v}_1, \dots, \mathbf{v}_d \in \mathcal{V}$ are a basis for \mathcal{V} if and only if the following are both true.

- (i) The matrix $[\mathbf{v}_1 \ \cdots \ \mathbf{v}_d] \in \mathbb{R}^{p \times d}$ has full column rank.
- (ii) $\mathcal{V} = \mathbf{C}([\mathbf{v}_1 \ \cdots \ \mathbf{v}_d])$.

23.4 Example. (i) The standard basis vectors for \mathbb{R}^n are a basis for \mathbb{R}^n . (If they weren’t, it would be a pretty awful use of the word “basis.”) Here’s the proof for $n = 3$, which I think you know by now already. If $\mathbf{v} \in \mathbb{R}^3$, then

$$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = v_1 \mathbf{e}_1 + v_2 \mathbf{e}_2 + v_3 \mathbf{e}_3.$$

Thus $\mathbb{R}^3 = \text{span}(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$. If $x_1 \mathbf{e}_1 + x_2 \mathbf{e}_2 + x_3 \mathbf{e}_3 = \mathbf{0}_3$, then with $\mathbf{x} = (x_1, x_2, x_3)$, we have $\mathbf{x} = \mathbf{0}_3$, thus $x_1 = x_2 = x_3 = 0$. This is the linear independence of $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$.

(ii) The pivot columns of a matrix are a basis for that matrix’s column space. No surprises here: the pivot columns are independent and they span that matrix’s column space.

(iii) The columns of an invertible matrix $A \in \mathbb{R}^{m \times m}$ are a basis for \mathbb{R}^m . Again, unsurprising: these columns are independent and $\mathbf{C}(A) = \mathbb{R}^m$.

(iv) Let’s do one more null space calculation. Let $A \in \mathbb{R}^{3 \times 5}$ with

$$\text{rref}(A) = \tilde{R}_0 P, \quad \tilde{R}_0 = \begin{bmatrix} 1 & 0 & f_3 & f_4 & f_5 \\ 0 & 1 & g_3 & g_4 & g_5 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

for some $f_3, f_4, f_5, g_3, g_4, g_5 \in \mathbb{R}$ and a permutation matrix $P \in \mathbb{R}^{5 \times 5}$. We know $A\mathbf{x} = \mathbf{0}_3$ if and only if $R_0 P\mathbf{x} = \mathbf{0}_3$. The permutation matrix is annoying, so here's a trick: put $\mathbf{y} = P\mathbf{x}$. Then $A\mathbf{x} = \mathbf{0}_3$ if and only if $\tilde{R}_0 \mathbf{y} = \mathbf{0}_3$. This turns into the system

$$\begin{cases} y_1 & + f_3 y_3 + f_4 y_4 + f_5 y_5 = 0 \\ & y_2 + g_3 y_3 + g_4 y_4 + g_5 y_5 = 0 \\ & & 0 = 0, \end{cases}$$

which we easily solve as

$$\mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \end{bmatrix} = \begin{bmatrix} -f_3 y_3 - f_4 y_4 - f_5 y_5 \\ -g_3 y_3 - g_4 y_4 - g_5 y_5 \\ y_3 \\ y_4 \\ y_5 \end{bmatrix} = y_3 \begin{bmatrix} -f_3 \\ -g_3 \\ 1 \\ 0 \\ 0 \end{bmatrix} + y_4 \begin{bmatrix} -f_4 \\ -g_4 \\ 0 \\ 1 \\ 0 \end{bmatrix} + y_5 \begin{bmatrix} -f_5 \\ g_5 \\ 0 \\ 0 \\ 1 \end{bmatrix}.$$

Then

$$\mathbf{x} = P^{-1}\mathbf{y} = y_3 P^{-1} \begin{bmatrix} -f_3 \\ -g_3 \\ 1 \\ 0 \\ 0 \end{bmatrix} + y_4 P^{-1} \begin{bmatrix} -f_4 \\ -g_4 \\ 0 \\ 1 \\ 0 \end{bmatrix} + y_5 P^{-1} \begin{bmatrix} -f_5 \\ g_5 \\ 0 \\ 0 \\ 1 \end{bmatrix},$$

and so

$$\mathbf{N}(A) = \mathbf{C} \left(P^{-1} \begin{bmatrix} -f_3 & -f_4 & -f_5 \\ -g_3 & -g_4 & -g_5 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right).$$

How does this give a basis for the null space? We've already expressed the null space as a column space, so that takes care of the span; as for independence, if we abbreviate $\mathbf{N}(A) = \mathbf{C}(P^{-1}B)$, then we want to show that the columns of $P^{-1}B$ are independent. So, suppose $P^{-1}B\mathbf{v} = \mathbf{0}_5$; then $B\mathbf{v} = \mathbf{0}_5$. But this says

$$\begin{bmatrix} * \\ * \\ * \\ v_1 \\ v_2 \\ v_3 \end{bmatrix} = \mathbf{0}_5,$$

from which we get $\mathbf{v} = \mathbf{0}_3$.

More generally, if for $A \in \mathbb{R}^{m \times n}$

$$\text{rref}(A) = \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} P,$$

this argument shows that the columns of

$$P^{-1} \begin{bmatrix} -F \\ I_{n-r} \end{bmatrix}$$

are a basis for $\mathbf{N}(A)$. The great thing is that we didn't need to know anything about P ! All that mattered was the *invertibility* of P , not even that P was a permutation matrix.

Content from Strang's ILA 6E. Read all of “A Basis for a Vector Space” on pp. 118–119. Every single thing here is important. Then read Worked Example 3.4 A on p. 122. This is a very important example that you should know how to prove.

Day 24: Monday, March 3.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Trivial subspace of \mathbb{R}^p , dimension (of a subspace of \mathbb{R}^p)

A basis is fundamentally a “coordinate system” for a subspace: we can “reach” every vector in the subspace in a unique way “via” the basis. We know this in our hearts with the standard basis $\mathbf{e}_1, \mathbf{e}_2$ for \mathbb{R}^2 : any $\mathbf{x} = (x, y) \in \mathbb{R}^2$ has the unique form $\mathbf{x} = x\mathbf{e}_1 + y\mathbf{e}_2$.

24.1 Theorem. Let \mathcal{V} be a subspace of \mathbb{R}^p and let $\mathbf{v}_1, \dots, \mathbf{v}_d \in \mathcal{V}$ be a basis for \mathcal{V} . For any $\mathbf{v} \in \mathcal{V}$, there are unique $c_1, \dots, c_d \in \mathbb{R}$ such that

$$\mathbf{v} = c_1\mathbf{v}_1 + \dots + c_d\mathbf{v}_d.$$

Proof. Let $A = [\mathbf{v}_1 \ \dots \ \mathbf{v}_d]$, so $\mathcal{V} = \mathbf{C}(A)$. Then any $\mathbf{v} \in \mathcal{V}$ can be written as a linear combination of the columns of A ; this is the *existence* result for the coefficients c_k . We need to show *uniqueness*. If there are two sets of coefficients, then we have $\mathbf{v} = A\mathbf{x}$ and $\mathbf{v} = A\mathbf{y}$ for some $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. But then $\mathbf{x} - \mathbf{y} \in \mathbf{N}(A) = \{\mathbf{0}_d\}$ since A has independent columns, thus $\mathbf{x} = \mathbf{y}$. ■

But this presumes that a subspace has a basis. We know this to be true for the most important subspaces, column spaces and null spaces, but it turns out that the three subspace axioms alone guarantee the existence of a basis. To prove that, we need an unsurprising auxiliary result.

24.2 Lemma. Let \mathcal{V} be a subspace. If $\mathbf{v}_1, \dots, \mathbf{v}_d \in \mathcal{V}$ and $c_1, \dots, c_d \in \mathbb{R}$, then $c_1\mathbf{v}_1 + \dots + c_d\mathbf{v}_d \in \mathcal{V}$. More generally, $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_d)$ is contained in \mathcal{V} .

Proof. This is really an induction argument on d . Here's why it's true for $d = 3$. Since \mathbf{v}_1 ,

$\mathbf{v}_2, \mathbf{v}_3 \in \mathcal{V}$, the subspace axioms guarantee $c_1\mathbf{v}_1, c_2\mathbf{v}_2, c_3\mathbf{v}_3 \in \mathcal{V}$. Then the axioms guarantee $c_1\mathbf{v}_1 + c_2\mathbf{v}_2 \in \mathcal{V}$, and so $(c_1\mathbf{v}_1 + c_2\mathbf{v}_2) + c_3\mathbf{v}_3 \in \mathcal{V}$. ■

Actually, we need to exclude one kind of subspace from the following existential result about bases.

24.3 Definition. The **TRIVIAL SUBSPACE** of \mathbb{R}^p is $\{\mathbf{0}_p\}$.

24.4 Problem (!). Check that $\{\mathbf{0}_p\}$ is indeed a subspace of \mathbb{R}^p . Why can't it have a basis?

24.5 Theorem. Let $\mathcal{V} \neq \{\mathbf{0}_p\}$ be a subspace of \mathbb{R}^p . Then \mathcal{V} has a basis.

Proof. This is a proof by exhaustion, which means that we'll exhaust all possible cases and also ourselves. Since $\mathcal{V} \neq \{\mathbf{0}_p\}$, there is $\mathbf{v}_1 \in \mathcal{V}$ such that $\mathbf{v}_1 \neq \mathbf{0}_p$. One of two things is true: either $\mathcal{V} = \text{span}(\mathbf{v}_1)$ or $\mathcal{V} \neq \text{span}(\mathbf{v}_1)$. In the first case, \mathbf{v}_1 by itself is a basis for \mathcal{V} since it's nonzero and therefore independent and also spans \mathcal{V} .

In the second case, there is $\mathbf{v}_2 \in \mathcal{V}$ such that $\mathbf{v}_2 \notin \text{span}(\mathbf{v}_1)$. Since $\mathbf{v}_1 \neq \mathbf{0}_p$, the list $\mathbf{v}_1, \mathbf{v}_2$ is independent. Once again, one of two things is true: either $\mathcal{V} = \text{span}(\mathbf{v}_1, \mathbf{v}_2)$ or $\mathcal{V} \neq \text{span}(\mathbf{v}_1, \mathbf{v}_2)$. In the first case, $\mathbf{v}_1, \mathbf{v}_2$ is a basis for \mathcal{V} .

In the second case, there is $\mathbf{v}_3 \in \mathcal{V}$ such that $\mathbf{v}_3 \notin \text{span}(\mathbf{v}_1, \mathbf{v}_2)$. I think you know what to do...

Assuming $\mathcal{V} \neq \text{span}(\mathbf{v}_1)$, one of two things has to happen in the end. First, we could have a list $\mathbf{v}_1, \dots, \mathbf{v}_d \in \mathcal{V}$ such that $\mathcal{V} = \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_d)$, $\mathbf{v}_1 \neq \mathbf{0}_p$ and $\mathbf{v}_j \notin \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_{j-1})$ for $j = 2, \dots, d$. In this case, $\mathbf{v}_1, \dots, \mathbf{v}_d$ is a basis for \mathcal{V} .

Or, we've turned the crank far enough to have an independent list $\mathbf{v}_1, \dots, \mathbf{v}_p \in \mathcal{V}$ of p (necessarily distinct!) vectors. If we try to push this one step further, either we'd have $\mathcal{V} = \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_p)$, and there's our basis, or we'd find $\mathbf{v}_{p+1} \in \mathcal{V}$ such that $\mathbf{v}_{p+1} \notin \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_p)$. I claim the second situation won't happen. Here's why.

Put $A = [\mathbf{v}_1 \ \cdots \ \mathbf{v}_p]$, so $A \in \mathbb{R}^{p \times p}$ has independent columns and therefore is invertible. Thus $\mathbf{C}(A) = \mathbb{R}^p$. Lemma 24.2 guarantees that $\mathbf{C}(A)$ is contained in \mathcal{V} , so all of \mathbb{R}^p is contained in \mathcal{V} . But \mathcal{V} is contained in \mathbb{R}^p , so the only possibility is that $\mathcal{V} = \mathbb{R}^p$ and $\mathcal{V} = \mathbf{C}(A)$. Then $\mathbf{v}_1, \dots, \mathbf{v}_p$ is a basis for \mathbb{R}^p after all. ■

24.6 Problem (!). Reread the preceding proof and identify exactly where we used the assumption that \mathcal{V} was a subspace.

The point of a basis is efficient representation. Theorem 24.1 gives us part of that efficiency: there is only one way to represent vectors with respect to a basis. And now we know there is always a basis. One more big thing remains: the amount of data in a basis is effectively always the same in that any basis contains the same number of vectors. (Deeper question: is there a "best" basis for a subspace? What more could we want? Think about it...)

24.7 Theorem. Let $\mathcal{V} \neq \{\mathbf{0}_p\}$ be a subspace of \mathbb{R}^p . Any basis for \mathcal{V} contains the same number of vectors.

Proof. Say that your basis is $\mathbf{a}_1, \dots, \mathbf{a}_r$ and mine is $\mathbf{b}_1, \dots, \mathbf{b}_s$. So, both $\mathbf{a}_1, \dots, \mathbf{a}_r$ and $\mathbf{b}_1, \dots, \mathbf{b}_s$ are independent and $\mathcal{V} = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_r)$ and $\mathcal{V} = \text{span}(\mathbf{b}_1, \dots, \mathbf{b}_s)$. We want $r = s$; otherwise, one of us is more efficient than the other, and that's not right. This will be a proof by contradiction. We assume $r \neq s$, and we figure out something wrong. If $r \neq s$, then either $r < s$ or $s < r$. We'll do the proof for $s < r$ and get a contradiction, and I claim that "flipping things around" basically shows $r < s$ leads to a contradiction, too.

I want to do two small cases of $s < r$ before giving the abstract proof. This will show us the pattern. Remember throughout that since $\mathcal{V} = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_r)$ and each $\mathbf{b}_j \in \mathcal{V}$, we have $\mathbf{a}_j \in \text{span}(\mathbf{b}_1, \dots, \mathbf{b}_s)$. Here we go.

1. $s = 1$ and $r = 2$. So $\mathbf{a}_1 = x_1 \mathbf{b}_1$ and $\mathbf{a}_2 = x_2 \mathbf{b}_1$. This looks suspiciously like \mathbf{a}_1 and \mathbf{a}_2 are both multiples of the same vector \mathbf{b}_1 , which suggests dependence. If we think about dependence as "nontrivial linear combination adding to zero," we might "cross-multiply" to get

$$x_2 \mathbf{a}_1 + (-x_1) \mathbf{a}_2 = x_2 x_1 \mathbf{b}_1 - x_1 x_2 \mathbf{b}_1 = \mathbf{0}_p.$$

Neither of x_1 nor x_2 can be 0, as otherwise $\mathbf{a}_1 = \mathbf{0}_p$ or $\mathbf{a}_2 = \mathbf{0}_p$ (which contradicts the independence of \mathbf{a}_1 and \mathbf{a}_2), so this is a nontrivial linear combination that adds to the zero vector. That contradicts the independence of \mathbf{a}_1 and \mathbf{a}_2 .

2. $s = 2$ and $r = 3$. Then we can write

$$\mathbf{a}_1 = x_1 \mathbf{b}_1 + x_2 \mathbf{b}_2 = [\mathbf{b}_1 \quad \mathbf{b}_2] \begin{bmatrix} x_1 \\ x_2 \end{bmatrix},$$

$$\mathbf{a}_2 = y_1 \mathbf{b}_1 + y_2 \mathbf{b}_2 = [\mathbf{b}_1 \quad \mathbf{b}_2] \begin{bmatrix} y_1 \\ y_2 \end{bmatrix},$$

and

$$\mathbf{a}_3 = z_1 \mathbf{b}_1 + z_2 \mathbf{b}_2 = [\mathbf{b}_1 \quad \mathbf{b}_2] \begin{bmatrix} z_1 \\ z_2 \end{bmatrix}.$$

Put it all together to get

$$A = BM, \quad A = [\mathbf{a}_1 \quad \mathbf{a}_2 \quad \mathbf{a}_3], \quad B = [\mathbf{b}_1 \quad \mathbf{b}_2], \quad \text{and} \quad C = \begin{bmatrix} x_1 & y_1 & z_1 \\ x_2 & y_2 & z_2 \end{bmatrix}.$$

Here's the problem: $C \in \mathbb{R}^{2 \times 3}$. This matrix C has more columns than rows. So, there is $\mathbf{w} \in \mathbb{R}^3$ such that $\mathbf{w} \neq \mathbf{0}_3$ and $C\mathbf{w} = \mathbf{0}_2$. But then

$$A\mathbf{w} = BC\mathbf{w} = B(C\mathbf{w}) = B\mathbf{0}_2 = \mathbf{0}_p,$$

thus $\mathbf{w} \in \mathbf{N}(A)$. Since A has independent columns, $\mathbf{N}(A) = \{\mathbf{0}_p\}$. This is the contradiction.

3. We are assuming $\text{span}(\mathbf{a}_1, \dots, \mathbf{a}_r) = \text{span}(\mathbf{b}_1, \dots, \mathbf{b}_s)$ with $s < r$. Put $A = [\mathbf{a}_1 \quad \dots \quad \mathbf{a}_r]$ and $B = [\mathbf{b}_1 \quad \dots \quad \mathbf{b}_s]$. Then $\mathbf{C}(A)$ is contained in $\mathbf{C}(B)$. Write each \mathbf{a}_j as $\mathbf{a}_j = B\mathbf{c}_j$ for

some $\mathbf{c}_j \in \mathbb{R}^s$. Then $A = BC$ with $C = [\mathbf{c}_1 \ \cdots \ \mathbf{c}_r]$, so $C \in \mathbb{R}^{s \times r}$. Since $r > s$, there is $\mathbf{z} \in \mathbb{R}^r$ such that $\mathbf{z} \neq \mathbf{0}_r$ and $C\mathbf{z} = \mathbf{0}_s$. But then $A\mathbf{z} = BC\mathbf{z} = B(C\mathbf{z}) = B\mathbf{0}_s = \mathbf{0}_r$. This contradicts the independence of the columns of A . ■

24.8 Problem (*). Reread the third, general step in the previous proof and explain how it proves the following. If $\mathbf{a}_1, \dots, \mathbf{a}_r, \mathbf{b}_1, \dots, \mathbf{b}_s \in \mathbb{R}^p$ with $\text{span}(\mathbf{a}_1, \dots, \mathbf{a}_r)$ contained in $\text{span}(\mathbf{b}_1, \dots, \mathbf{b}_s)$, and if $r > s$, then $\mathbf{a}_1, \dots, \mathbf{a}_r$ is dependent.

24.9 Problem (+). Complete the proof of Theorem 24.7 by showing that $r < s$ leads to a contradiction. [Hint: use Problem 24.8 to interchange the role of $\mathbf{a}_1, \dots, \mathbf{a}_r$ and $\mathbf{b}_1, \dots, \mathbf{b}_s$.]

24.10 Problem (+). Here is a chance to fill in some gaps, answer some old questions, and prove some lingering conjectures about spans and independence. Let \mathcal{V} be a subspace of \mathbb{R}^p with dimension $\dim(\mathcal{V}) = r \geq 1$.

- (i) Use Problem 24.8 to prove that any list of more than r vectors in \mathcal{V} is dependent.
- (ii) Use the previous part to prove that $\dim(\mathcal{V}) \leq p$.
- (iii) Prove that any list of r independent vectors in \mathcal{V} is a basis for \mathcal{V} . [Hint: if not, why is there a list of $r + 1$ independent vectors in \mathcal{V} ?]
- (iv) Prove that any list of vectors that spans \mathcal{V} contains an independent “sublist” that spans \mathcal{V} (and thus is a basis for \mathcal{V}). By “sublist” I mean that if the original list has the form, say, $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4, \mathbf{v}_5$, then the vectors $\mathbf{v}_1, \mathbf{v}_3, \mathbf{v}_5$ are a sublist. [Hint: put the original list into a matrix and think about the pivot columns.]
- (v) Prove that any list of r vectors that spans \mathcal{V} is a basis for \mathcal{V} . [Hint: if such a list is dependent, use the previous part to conclude $\dim(\mathcal{V}) < r$.]

Content from Strang’s ILA 6E. Page 120 proves Theorem 24.7. Read the paragraphs after the boxed “Definition.” Then read Worked Example 3.4 B on p. 122.

Since any (nontrivial) subspace has a basis, and any basis for a (nontrivial) subspace has the same length, it’s fair to give a name to that length.

24.11 Definition. Let \mathcal{V} be a subspace of \mathbb{R}^p . The **DIMENSION** of \mathcal{V} , denoted $\dim(\mathcal{V})$, is the length of any basis for that subspace. We define $\dim(\{\mathbf{0}_p\}) := 0$.

24.12 Example. (i) Since the standard basis $\mathbf{e}_1, \dots, \mathbf{e}_n$ for \mathbb{R}^n contains n vectors, $\dim(\mathbb{R}^n) = n$.

(ii) Let $A \in \mathbb{R}^{m \times n}$ have rank $r \geq 1$. Then $\dim[\mathbf{C}(A)] = r$, since A has r pivot columns and those pivot columns form a basis for A .

(iii) Let $A \in \mathbb{R}^{m \times n}$ have rank $r \geq 1$. If $r = n$, then $\mathbf{N}(A) = \{\mathbf{0}_n\}$, so in that case $\dim[\mathbf{N}(A)] = 0$. If $1 \leq r < n$, then A has $n - r$ free variables, and we should expect that this means $\dim[\mathbf{N}(A)] = n - r$. (If $r = n$, we recover the $r = n$ result.) More precisely, we worked out in part (iv) of Example 23.4 that if

$$\text{rref}(A) = \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} P,$$

then the columns of

$$P^{-1} \begin{bmatrix} -F \\ I_{n-r} \end{bmatrix}$$

are a basis for $\mathbf{N}(A)$. There are $n - r$ such columns.

Content from Strang's ILA 6E. Look at the matrix R_0 on p. 130. Then read #2 and #3 on pp. 130–131 on the dimensions of its column and null spaces. Next, look at the matrix A at the bottom of p. 131 and read about its column and null spaces in #2 and #3 on pp. 132–133. We'll come back to the row space and left null space shortly. Read Worked Example 3.5 B on p. 137.

Day 25: Wednesday, March 5.

Part (iii) of Example 24.12 deserves to stand on its own.

25.1 Theorem (Rank-nullity). Let $A \in \mathbb{R}^{m \times n}$. Then

$$\dim[\mathbf{N}(A)] + \dim[\mathbf{C}(A)] = n.$$

The point is that if you know one of these dimensions, then you know the other. It's interesting, and maybe a bit weird, that even though $\mathbf{C}(A)$ is not a subspace of \mathbb{R}^n (it's a subspace of \mathbb{R}^m), its dimension still talks to the dimension of $\mathbf{N}(A)$ (which is a subspace of \mathbb{R}^n) and the dimension of \mathbb{R}^n itself.

25.2 Example. We finally need a larger matrix than our most frequently used, beloved example. Let

$$A = \begin{bmatrix} 1 & 2 & 1 & 7 & 0 \\ 2 & 4 & 2 & 14 & 0 \\ 0 & 0 & 0 & 2 & 8 \end{bmatrix}.$$

I'll leave it to you to check that

$$\text{rref}(A) = \begin{bmatrix} 1 & 2 & 0 & 3 & 0 \\ 0 & 0 & 1 & 4 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

We see that A has 2 pivot columns, so $\text{rank}(A) = \dim[\mathbf{C}(A)] = 2$, and therefore

$\dim[\mathbf{N}(A)] = 5 - 2 = 3$. We can verify this explicitly by computing (I'll let you do that)

$$\mathbf{N}(A) = \mathbf{C} \left(\begin{bmatrix} -2 & -3 & 0 \\ 1 & 0 & 0 \\ 0 & -4 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right)$$

and seeing the three independent columns spanning $\mathbf{N}(A)$ right there.

The notion of dimension allows us to quantify how much of \mathbb{R}^m the column space of $A \in \mathbb{R}^{m \times n}$ misses: it misses $m - \dim[\mathbf{C}(A)]$ “dimensions” of \mathbb{R}^m . But what is going on elsewhere in \mathbb{R}^m beyond $\mathbf{C}(A)$? The deeper question is not just *how much* of \mathbb{R}^m does $\mathbf{C}(A)$ miss but rather *what exactly* in \mathbb{R}^m does $\mathbf{C}(A)$ miss. Is there a simpler way to characterize and describe $\mathbf{C}(A)$ than just its definition?

We could also ask about \mathbb{R}^n and $\mathbf{N}(A)$. When $\mathbf{v} \in \mathbf{N}(A)$, we have $A\mathbf{v} = \mathbf{0}_m$. For what $\mathbf{w} \in \mathbb{R}^n$ does A “act nontrivially” with $A\mathbf{w} \neq \mathbf{0}_m$? On what parts of \mathbb{R}^n is A “more interesting”?

It turns out that these questions are “dual” to each other in that if we know how to handle one of them, we can understand the other pretty quickly. And it also turns out (I think) that asking how $\mathbf{N}(A)$ interacts with the rest of \mathbb{R}^n is an easier thing to control.

25.3 Example. Let

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

Then

$$\mathbf{N}(A) = \mathbf{C} \left(\begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \right).$$

I hope it's glaringly obvious what's missing from $\mathbf{N}(A)$: \mathbf{e}_1 .

More precisely, let

$$\mathcal{W} = \text{span} \left(\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right).$$

Then any $\mathbf{x} \in \mathbb{R}^3$ can be written as

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} x_1 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ x_2 \\ x_3 \end{bmatrix}.$$

That is, any $\mathbf{x} \in \mathbb{R}^3$ can be written (or, more evocatively, “decomposed”) in the form $\mathbf{x} = \mathbf{v} + \mathbf{w}$ for $\mathbf{v} \in \mathbf{N}(A)$ and $\mathbf{w} \in \mathcal{W}$. Think for a moment about why those \mathbf{v} and \mathbf{w} are unique; that is, why there is only one way to achieve this “decomposition” of \mathbf{x} .

Can we see \mathcal{W} directly from A itself, without passing to the null space? Sure: \mathbf{e}_1 is the first row of A . We're more used to thinking about columns than rows, so let's flip every row of A to a column and every column of A to a row by taking the **TRANSPOSE**:

$$A^T = \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

Then we see

$$\mathcal{W} = \mathbf{C}(A^T).$$

And so we have written any $\mathbf{x} \in \mathbb{R}^3$ uniquely as a sum of the form $\mathbf{x} = \mathbf{v} + \mathbf{w}$ with $\mathbf{v} \in \mathbf{N}(A)$ and $\mathbf{w} \in \mathbf{C}(A^T)$.

Did we just get really lucky, since the bases for $\mathbf{N}(A)$ and $\mathbf{C}(A^T)$ in the previous example just involved the coordinate axes for three-dimensional space? Here's a more complicated situation. Let

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

This was the RREF of the original matrix from Example 25.2. Then

$$\mathbf{N}(A) = \mathbf{C} \left(\begin{bmatrix} -2 & -3 & 0 \\ 1 & 0 & 0 \\ 0 & -4 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right)$$

and (flipping columns to rows and rows to columns)

$$A^T = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 0 & 0 \\ 0 & 1 & 0 \\ 3 & 4 & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

so (ignoring that third column, but keeping those first two independent columns)

$$\mathbf{C}(A^T) = \text{span} \left(\begin{bmatrix} 1 \\ 2 \\ 0 \\ 3 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 4 \\ 0 \end{bmatrix} \right).$$

Can we write each $\mathbf{x} \in \mathbb{R}^5$ as $\mathbf{x} = \mathbf{v} + \mathbf{w}$ for some $\mathbf{v} \in \mathbf{N}(A)$ and $\mathbf{w} \in \mathbf{C}(A^T)$, and is such a decomposition of \mathbf{x} unique?

Some notation will compress things helpfully. Let

$$\mathbf{v}_1 = \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} -3 \\ 0 \\ -4 \\ 1 \\ 0 \end{bmatrix}, \quad \mathbf{v}_3 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \quad \mathbf{w}_1 = \begin{bmatrix} 1 \\ 2 \\ 0 \\ 3 \\ 0 \end{bmatrix}, \quad \text{and} \quad \begin{bmatrix} 0 \\ 0 \\ 1 \\ 4 \\ 0 \end{bmatrix}, \quad (25.1)$$

so \mathbf{v}_1 , \mathbf{v}_2 , and \mathbf{v}_3 form a basis for $\mathbf{N}(A)$, and \mathbf{w}_1 and \mathbf{w}_2 form a basis for $\mathbf{C}(A^T)$. Let

$$M = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3 \quad \mathbf{w}_1 \quad \mathbf{w}_2], \quad (25.2)$$

so $M \in \mathbb{R}^{5 \times 5}$. If M is invertible, then we can write each $\mathbf{x} \in \mathbb{R}^5$ as $\mathbf{x} = M\mathbf{y}$ for some (unique!) $\mathbf{y} \in \mathbb{R}^5$. Then

$$\mathbf{x} = (y_1\mathbf{v}_1 + y_2\mathbf{v}_2 + y_3\mathbf{v}_3) + (y_4\mathbf{w}_1 + y_5\mathbf{w}_2) \quad (25.3)$$

with $y_1\mathbf{v}_1 + y_2\mathbf{v}_2 + y_3\mathbf{v}_3 \in \mathbf{N}(A)$ and $y_4\mathbf{w}_1 + y_5\mathbf{w}_2 \in \mathbf{C}(A^T)$. That gives the decomposition of \mathbf{x} that we want, and we could probably push it further with an independence argument to get uniqueness.

So, is M invertible? Do you really want to do the row operations to find out? This matrix M is a beast! There's a less obvious approach that will teach us some valuable new things, and it all hinges on a deeper notion of the dot product. First, I encourage you to review some hopefully unsurprising dot product arithmetic.

25.4 Problem (★). Show that the dot product has the following properties. All vectors below are in the same space, e.g., \mathbb{R}^p . (If it makes things more concrete for you, do it for $p = 3$.)

(i) $\mathbf{v} \cdot \mathbf{w} = \mathbf{w} \cdot \mathbf{v}$.

(ii) $\mathbf{v} \cdot (\mathbf{w}_1 + \mathbf{w}_2) = (\mathbf{v} \cdot \mathbf{w}_1) + (\mathbf{v} \cdot \mathbf{w}_2)$.

(iii) $\mathbf{v} \cdot (c\mathbf{w}) = c(\mathbf{v} \cdot \mathbf{w})$.

(iv) $\mathbf{v} \cdot \mathbf{v} \geq 0$.

(v) If $\mathbf{v} \cdot \mathbf{v} = 0$, then $\mathbf{v} = \mathbf{0}$. [Hint: if $\mathbf{v} \neq \mathbf{0}$, explain why $\mathbf{v} \cdot \mathbf{v} > 0$.]

Next, take a look at the dot products of the vectors in (25.1). I'll get you started:

$$\mathbf{v}_1 \cdot \mathbf{w}_1 = \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 2 \\ 0 \\ 3 \\ 0 \end{bmatrix} = (-2 \cdot 1) + (1 \cdot 2) + (0 \cdot 0) + (0 \cdot 3) + (0 \cdot 0) = -2 + 2 + 0 + 0 + 0 = 0.$$

25.5 Problem (!). Check that

$$\mathbf{v}_i \cdot \mathbf{w}_j = 0$$

for $i = 1, 2, 3$ and $j = 1, 2$.

Why would we think to look for these relations between the bases of $\mathbf{N}(A)$ and $\mathbf{C}(A^\top)$? One reason might be the transparent relations among the bases from Example 25.3. Those bases talk to each other so well via dot products—maybe other bases do, too.

Now we'll show that M from (25.2) is invertible. Assume $M\mathbf{y} = \mathbf{0}_5$; we'll find $\mathbf{y} = \mathbf{0}_5$. We rewrite this as

$$\mathbf{v} + \mathbf{w} = \mathbf{0}_5, \quad \mathbf{v} = y_1\mathbf{v}_1 + y_2\mathbf{v}_2 + y_3\mathbf{v}_3, \quad \mathbf{w} = y_4\mathbf{w}_1 + y_5\mathbf{w}_2.$$

I claim that $\mathbf{v} \cdot \mathbf{w} = 0$ and that this is a consequence of the dot product arithmetic in Problem 25.4 and the interaction of \mathbf{v}_i and \mathbf{w}_j from Problem 25.5. For example, if $y_1 = 1$ and $y_2 = y_3 = 0$, then

$$\mathbf{v}_1 \cdot (y_4\mathbf{w}_1 + y_5\mathbf{w}_2) = y_4(\mathbf{v}_1 \cdot \mathbf{w}_1) + y_5(\mathbf{v}_1 \cdot \mathbf{w}_2) = (y_4 \cdot 0) + (y_5 \cdot 0) = 0.$$

You could check this in general, but I don't think it's hugely worth your time.

But here is why this matters. We are assuming $\mathbf{v} + \mathbf{w} = \mathbf{0}_5$ with $\mathbf{v} \cdot \mathbf{w} = 0$. The great trick is to take the dot product of both sides with \mathbf{v} (although \mathbf{w} would also work):

$$(\mathbf{v} + \mathbf{w}) \cdot \mathbf{v} = \mathbf{0}_5 \cdot \mathbf{v}.$$

Then we get

$$(\mathbf{v} \cdot \mathbf{v}) + (\mathbf{w} \cdot \mathbf{v}) = 0$$

and since $\mathbf{w} \cdot \mathbf{v} = 0$ this simplifies to

$$\mathbf{v} \cdot \mathbf{v} = 0.$$

Dot product properties tell us $\mathbf{v} = \mathbf{0}_5$, and so

$$y_1\mathbf{v}_1 + y_2\mathbf{v}_2 + y_3\mathbf{v}_3 = \mathbf{0}_5.$$

The independence of \mathbf{v}_1 , \mathbf{v}_2 , and \mathbf{v}_3 implies $y_1 = y_2 = y_3 = 0$, thus $\mathbf{v} = \mathbf{0}_5$, and so the equation $\mathbf{v} + \mathbf{w} = \mathbf{0}_5$ reduces to $\mathbf{w} = \mathbf{0}_5$. That is, $y_4\mathbf{w}_1 + y_5\mathbf{w}_2 = \mathbf{0}_5$. Another independence argument, now for \mathbf{w}_1 and \mathbf{w}_2 , yields $y_4 = y_5 = 0$.

All of this is to say that if $M\mathbf{y} = \mathbf{0}_5$, then $\mathbf{y} = \mathbf{0}_5$, so M is invertible, and so we get the desired representation (25.3). Why is this representation unique? We could use independence, but we can also use dot products.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Orthogonal vectors, transpose of a matrix

We know that we can write any $\mathbf{x} \in \mathbb{R}^5$ as $\mathbf{x} = \mathbf{v} + \mathbf{w}$ for some $\mathbf{v} \in \mathbf{N}(A)$ and $\mathbf{w} \in \mathbf{C}(A^\top)$. What if we can do so in two ways? Say $\mathbf{x} = \tilde{\mathbf{v}} + \tilde{\mathbf{w}}$ for some $\tilde{\mathbf{v}} \in \mathbf{N}(A)$ and $\tilde{\mathbf{w}} \in \mathbf{C}(A^\top)$ as well. Then

$$\mathbf{0}_5 = \mathbf{x} - \mathbf{x} = (\mathbf{v} + \mathbf{w}) - (\tilde{\mathbf{v}} + \tilde{\mathbf{w}}) = (\mathbf{v} - \tilde{\mathbf{v}}) + (\mathbf{w} - \tilde{\mathbf{w}})$$

and so

$$\mathbf{v} - \tilde{\mathbf{v}} = \tilde{\mathbf{w}} - \mathbf{w}.$$

Unfortunately for all of us, $\mathbf{N}(A)$ and $\mathbf{C}(A^\top)$ are subspaces. Since $\mathbf{v}, \tilde{\mathbf{v}} \in \mathbf{N}(A)$, we have $\mathbf{v} - \tilde{\mathbf{v}} \in \mathbf{N}(A)$, and likewise $\mathbf{w} - \tilde{\mathbf{w}} \in \mathbf{C}(A^\top)$. Hence $\mathbf{v} - \tilde{\mathbf{v}} \in \mathbf{N}(A)$ and $\mathbf{v} - \tilde{\mathbf{v}} = \tilde{\mathbf{w}} - \mathbf{w} \in \mathbf{C}(A^\top)$.

What happens if there is $\mathbf{y} \in \mathbb{R}^5$ such that both $\mathbf{y} \in \mathbf{N}(A)$ and $\mathbf{y} \in \mathbf{C}(A^\top)$? Recall that if $\mathbf{v} \in \mathbf{N}(A)$ and $\mathbf{w} \in \mathbf{C}(A^\top)$, then $\mathbf{v} \cdot \mathbf{w} = 0$. Taking $\mathbf{v} = \mathbf{y}$ and $\mathbf{w} = \mathbf{y}$, we get $\mathbf{y} \cdot \mathbf{y} = 0$, thus $\mathbf{y} = \mathbf{0}_5$.

In the situation above, this means $\mathbf{v} - \tilde{\mathbf{v}} = \mathbf{0}_5$, so $\mathbf{v} = \tilde{\mathbf{v}}$, but then also $\tilde{\mathbf{w}} - \mathbf{w} = \mathbf{0}_5$, so $\mathbf{w} = \tilde{\mathbf{w}}$. This is our desired uniqueness.

Now here is the great thing: all of this generalizes far beyond the specific matrices A just considered. First, let's name the feature of dot products that we've been using.

26.1 Definition. The vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ are **ORTHOGONAL** if $\mathbf{v} \cdot \mathbf{w} = 0$.

Next, let's formalize the notion of transpose.

26.2 Definition (What the transpose is). The **TRANSPOSE** of $A \in \mathbb{R}^{m \times n}$ is the matrix $A^\top \in \mathbb{R}^{n \times m}$ such that the (i, j) -entry of A^\top is the (j, i) -entry of A . We write

$$A_{ij}^\top = A_{ji}.$$

26.3 Example. If

$$A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix},$$

then

$$A^\top = \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix}.$$

Now we can generalize how $\mathbf{N}(A)$ and $\mathbf{C}(A^\top)$ interact in general.

26.4 Theorem. Let $A \in \mathbb{R}^{m \times n}$, $\mathbf{v} \in \mathbf{N}(A)$, and $\mathbf{w} \in \mathbf{C}(A^\top)$. Then $\mathbf{v} \cdot \mathbf{w} = 0$.

Proof. If $\mathbf{v} \in \mathbf{N}(A)$, then $A\mathbf{v} = \mathbf{0}_m$. One way to compute $A\mathbf{v}$ is by taking dot products of \mathbf{v} with the rows of A viewed as columns in \mathbb{R}^n . That $A\mathbf{v} = \mathbf{0}_m$ says that each such dot product is 0. Say that the i th row of A , viewed as a column in \mathbb{R}^n , is $\mathbf{b}_i \in \mathbb{R}^n$. Then $\mathbf{v} \cdot \mathbf{b}_i = 0$ for all i . Any vector in $\mathbf{C}(A^\top)$ is a linear combination of the rows of A viewed as columns in \mathbb{R}^n . Say that $\mathbf{w} \in \mathbf{C}(A^\top)$ has the form

$$\mathbf{w} = c_1 \mathbf{b}_1 + \cdots + c_m \mathbf{b}_m.$$

Then

$$\mathbf{v} \cdot \mathbf{w} = \mathbf{v} \cdot (c_1 \mathbf{b}_1 + \cdots + c_m \mathbf{b}_m) = c_1(\mathbf{v} \cdot \mathbf{b}_1) + \cdots + c_m(\mathbf{v} \cdot \mathbf{b}_m) = 0. \quad \blacksquare$$

This is a perfectly adequate proof based on what A^\top is: the matrix formed by swapping the rows and columns of A . This is a “static” way to think about A^\top —it’s an array of data. Nothing wrong with that.

But we can be dynamic: *what things do defines what things are*. Here is what A^\top does. Let $\mathbf{e}_1, \dots, \mathbf{e}_n$ be the standard basis vectors for \mathbb{R}^n and $\tilde{\mathbf{e}}_1, \dots, \tilde{\mathbf{e}}_m$ be the standard basis vectors for \mathbb{R}^m . (I’m putting tildes on the vectors for \mathbb{R}^m because the notation \mathbf{e}_j doesn’t otherwise indicate what space it’s in. I guess we could stack extra subscripts or superscripts to indicate n and m but, ew.) So, if $n = 4$ and $m = 3$, then

$$\mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{e}_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \quad \text{and} \quad \mathbf{e}_4 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix},$$

while

$$\tilde{\mathbf{e}}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \tilde{\mathbf{e}}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad \text{and} \quad \tilde{\mathbf{e}}_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$

Now recall that multiplying a matrix by standard basis vectors extracts its columns, while taking the dot product of a vector with standard basis vectors extracts its entries. Since $A \in \mathbb{R}^{m \times n}$, the j th column of A is $A\mathbf{e}_j \in \mathbb{R}^m$, and then $A\mathbf{e}_j \cdot \tilde{\mathbf{e}}_i$ is the i th entry in that column. That is,

$$A_{ij} = A\mathbf{e}_j \cdot \tilde{\mathbf{e}}_i.$$

26.5 Problem (!). Suppose that $A, B \in \mathbb{R}^{m \times n}$. Certainly if $A = B$, then

$$A\mathbf{v} \cdot \mathbf{w} = B\mathbf{v} \cdot \mathbf{w} \tag{26.1}$$

for all $\mathbf{v} \in \mathbb{R}^n$ and $\mathbf{w} \in \mathbb{R}^m$. Just substitute A for B and plug and chug. Prove that if (26.1) holds for all $\mathbf{v} \in \mathbb{R}^n$ and $\mathbf{w} \in \mathbb{R}^m$, then $A = B$. [Hint: *get the standard basis vectors to show up.*]

Likewise, since $A^\top \in \mathbb{R}^{n \times m}$, the j th column of A^\top is $A^\top \tilde{\mathbf{e}}_j \in \mathbb{R}^n$, and then $A^\top \tilde{\mathbf{e}}_j \cdot \mathbf{e}_i$ is the i th entry in that column. That is,

$$A_{ij}^\top = A^\top \tilde{\mathbf{e}}_j \cdot \mathbf{e}_i.$$

Thus

$$A^T \tilde{\mathbf{e}}_j \cdot \mathbf{e}_i = A_{ij}^T = A_{ji}^T = A \mathbf{e}_i \cdot \tilde{\mathbf{e}}_j,$$

and so

$$A \mathbf{e}_i \cdot \tilde{\mathbf{e}}_j = A^T \tilde{\mathbf{e}}_j \cdot \mathbf{e}_i.$$

The commutativity of the dot product ($\mathbf{v} \cdot \mathbf{w} = \mathbf{w} \cdot \mathbf{v}$) gives

$$A \mathbf{e}_i \cdot \tilde{\mathbf{e}}_j = \mathbf{e}_i \cdot A^T \tilde{\mathbf{e}}_j. \quad (26.2)$$

This is what A^T does: it pops across the dot product.

26.6 Remark. *There are actually two dot products in (26.2). The one on the left is in \mathbb{R}^m , since $A \mathbf{e}_i, \tilde{\mathbf{e}}_j \in \mathbb{R}^m$. The one on the right is in \mathbb{R}^n , since $\mathbf{e}_i, A^T \tilde{\mathbf{e}}_j \in \mathbb{R}^n$.*

This “popping” behavior of the transpose is not limited to standard basis vectors.

26.7 Theorem (What the transpose does). *Let $A \in \mathbb{R}^{m \times n}$, $\mathbf{v} \in \mathbb{R}^n$, and $\mathbf{w} \in \mathbb{R}^m$. Then*

$$A \mathbf{v} \cdot \mathbf{w} = \mathbf{v} \cdot A^T \mathbf{w}.$$

Moreover, the transpose is the only matrix in $\mathbb{R}^{n \times m}$ to do this: if there is $B \in \mathbb{R}^{n \times m}$ such that

$$A \mathbf{v} \cdot \mathbf{w} = \mathbf{v} \cdot B \mathbf{w} \quad (26.3)$$

for all $\mathbf{v} \in \mathbb{R}^n$ and $\mathbf{w} \in \mathbb{R}^m$, then $B = A^T$.

26.8 Problem (+). Prove it.

(i) First, with $\tilde{\mathbf{e}}_1, \dots, \tilde{\mathbf{e}}_m$ as the standard basis vectors for \mathbb{R}^m , show

$$A \mathbf{v} \cdot \tilde{\mathbf{e}}_i = \mathbf{v} \cdot A^T \tilde{\mathbf{e}}_i$$

by expanding $\mathbf{v} = v_1 \mathbf{e}_1 + \dots + v_n \mathbf{e}_n$ and using linearity of matrix-vector multiplication and dot product arithmetic from Problem 25.4. Then show the general result by expanding $\mathbf{w} = w_1 \tilde{\mathbf{e}}_1 + \dots + w_m \tilde{\mathbf{e}}_m$.

(ii) Use Problem 26.5 to show that (26.3).

Here are some nice properties of the transpose that can easily be deduced from what it does, rather than what it is. These are important, and I expect that you’re going to know them, but I think they’ll be more meaningful if you prove them yourself.

26.9 Problem (★). (i) Let $A \in \mathbb{R}^{m \times n}$. First explain why $(A^T)^T \in \mathbb{R}^{m \times n}$, too. Then use Problem 26.5 to prove that $(A^T)^T = A$ by showing that

$$(A^T)^T \mathbf{v} \cdot \mathbf{w} = A \mathbf{v} \cdot \mathbf{w}$$

for all $\mathbf{v} \in \mathbb{R}^n$ and $\mathbf{w} \in \mathbb{R}^m$.

(ii) Let $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times p}$. First explain why $(AB)^\top, B^\top A^\top \in \mathbb{R}^{p \times m}$. Then show that

$$(B^\top A^\top) \mathbf{v} \cdot \mathbf{w} = \mathbf{v} \cdot AB\mathbf{w}$$

for all $\mathbf{v} \in \mathbb{R}^m$ and $\mathbf{w} \in \mathbb{R}^p$. Use the uniqueness result of Theorem 26.7 to conclude $(AB)^\top = B^\top A^\top$.

(iii) Let $A \in \mathbb{R}^{m \times m}$ be invertible. Prove that A^\top is also invertible with inverse $(A^\top)^{-1} = (A^{-1})^\top$ by computing $(AA^{-1})^\top = I_m^\top$ and $(AA^{-1})^\top = A^\top(A^{-1})^\top$. What does this tell you?

(iv) Let $P \in \mathbb{R}^{m \times m}$ be a permutation matrix, so P contains all of the columns of the identity matrix I_m (each column appearing once, and only once) in some order. Argue that $P\mathbf{v} \cdot P\mathbf{w} = \mathbf{v} \cdot \mathbf{w}$ for all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^m$. [Hint: *maybe do this for the only non-identity 2×2 permutation to get a feel for what's going on, then generalize.*] Conclude $P^\top P\mathbf{v} \cdot \mathbf{w} = \mathbf{v} \cdot \mathbf{w}$ for all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^m$. Why does this imply that P is invertible with $P^{-1} = P^\top$? (Good news: no more writing P^{-1} when doing calculations with the abstract form of the RREF!)

Content from Strang's ILA 6E. Pages 67–68 discuss fundamental properties of the transpose. Pages 68–69 show how the transpose interacts with dot products (I wholly disagree that \cdot is “unprofessional”—I like that it emphasizes how the dot product takes in two inputs and how it’s “linear in each input.” I like dot products.). If you have seen integration by parts in calculus, read Example 2 on p. 69.

Now we can give another proof of Theorem 26.4 that relies on what the transpose does, rather than is. Let $A \in \mathbb{R}^{m \times n}$, $\mathbf{v} \in \mathbf{N}(A)$, and $\mathbf{w} \in \mathbf{C}(A^\top)$. Then $A\mathbf{v} = \mathbf{0}_m$ and there is $\mathbf{y} \in \mathbb{R}^m$ such that $\mathbf{w} = A\mathbf{y}$. We compute

$$\mathbf{v} \cdot \mathbf{w} = \mathbf{v} \cdot A^\top \mathbf{y} = A\mathbf{v} \cdot \mathbf{y} = \mathbf{0}_m \cdot \mathbf{y} = 0. \quad (26.4)$$

So slick! The first two dot products were dot products in \mathbb{R}^n , but the second two were in \mathbb{R}^m .

Content from Strang's ILA 6E. Read p. 144 starting with the box “The **nullspace** of $A \dots$ ” Then read the “Important” paragraph on p. 145 about the orthogonality of $\mathbf{C}(A)$ and $\mathbf{N}(A^\top)$, which we will discuss in greater detail later, and Example 1. Note that Strang typically likes to write the dot product as $\mathbf{x}^\top \mathbf{y}$, not $\mathbf{x} \cdot \mathbf{y}$.

26.10 Problem (!). Let

$$A = \begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix}.$$

Use results from Examples 19.5 and 22.1 to give bases for $\mathbf{N}(A)$ and $\mathbf{C}(A^\top)$, and check directly that the vectors in the basis for $\mathbf{N}(A)$ are orthogonal to the vectors in the basis for $\mathbf{C}(A^\top)$.

26.11 Problem (+). Here is a much less slick way to show the orthogonality of vectors in $\mathbf{N}(A)$ and $\mathbf{C}(A^\top)$ that reinforces properties of the RREF. Let $A \in \mathbb{R}^{m \times n}$ with $r = \text{rank}(A)$.

- (i) If $r = n$, why do you have very little work to do? Do it.
- (ii) From now on suppose $r < n$. Suppose $EA = R_0$ with $E \in \mathbb{R}^{m \times m}$ invertible and $R_0 = \text{rref}(A)$. Use Problem 12.17 to explain why $\mathbf{C}(A^\top) = \mathbf{C}(R_0)$. (Don't get too excited: remember that in general $\mathbf{C}(A) \neq \mathbf{C}(R_0)$.)
- (iii) Write R_0 in the very general form

$$R_0 = \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} P,$$

where maybe the F - and/or zero blocks are not present, and P is a permutation matrix that could be the identity. Then compute

$$R_0^\top = P^\top \begin{bmatrix} I_r & 0 \\ F^\top & 0 \end{bmatrix}. \quad (26.5)$$

- (iv) Combine some old ideas with part (iv) of Problem 26.9 to show

$$\mathbf{N}(A) = \mathbf{C} \left(P^\top \begin{bmatrix} -F \\ I_{n-r} \end{bmatrix} \right).$$

Since $r < n$, that block I_{n-r} will always genuinely be present.

- (v) Conclude that if $\mathbf{v} \in \mathbf{N}(A)$ and $\mathbf{w} \in \mathbf{C}(A^\top)$, then there are $\mathbf{x} \in \mathbb{R}^{n-r}$ and $\mathbf{y} \in \mathbb{R}^r$ such that

$$\mathbf{v} = P^\top \begin{bmatrix} -F \\ I_{n-r} \end{bmatrix} \mathbf{x} \quad \text{and} \quad \mathbf{w} = P^\top \begin{bmatrix} I_r & 0 \\ F^\top & 0 \end{bmatrix} \mathbf{y}.$$

Use this to show $\mathbf{v} \cdot \mathbf{w} = 0$. [Hint: an identity from part (iv) of Problem 26.9 will be help with those common factors of P^\top .]

Day 27: Monday, March 17.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked "N."

Row space of a matrix, orthogonal complement of a subspace

It will save us some time (like 3 seconds) to give a different name to $\mathbf{C}(A^\top)$ other than "column space of A transpose."

27.1 Definition. The **ROW SPACE** of $A \in \mathbb{R}^{m \times n}$ is $\mathbf{C}(A^\top)$.

The orthogonality of vectors in $\mathbf{N}(A)$ and $\mathbf{C}(A^\top)$ is the key to generalizing the decomposition of \mathbb{R}^n that we did for some special matrices and some special n . Let $A \in \mathbb{R}^{m \times n}$. For each $\mathbf{x} \in \mathbb{R}^n$, we want to find unique $\mathbf{v} \in \mathbf{N}(A)$ and $\mathbf{w} \in \mathbf{C}(A^\top)$ such that $\mathbf{x} = \mathbf{v} + \mathbf{w}$. We need one more fact.

27.2 Theorem. Let $A \in \mathbb{R}^{m \times n}$. Then $\text{rank}(A) = \text{rank}(A^\top)$.

Informally, “row rank = column rank.” This is a little technical to prove precisely, so we’ll punt that to a wholly optional problem and give a proof by example. The rank of A is the number of pivot columns and the number of pivot rows of A . If you look at the RREF, pivot rows are independent. I mean, just look at

$$R_0 = \begin{bmatrix} 1 & * & 0 & * & * \\ 0 & 0 & 1 & * & * \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Two pivot columns, two pivot rows. Then

$$R_0^\top = \begin{bmatrix} 1 & 0 & 0 \\ * & 0 & 0 \\ 0 & 1 & 0 \\ * & * & 0 \\ * & * & 0 \end{bmatrix}.$$

Those first two columns in R_0^\top are definitely independent (even though R_0^\top isn’t in RREF anymore). So $\text{rank}(R_0) = 2$ and $\text{rank}(R_0^\top) = 2$. Basically, you can see from the structure of the RREF that

$$\text{rank}[\text{rref}(A)] = \text{rank}[\text{rref}(A)^\top].$$

Then you use the fact that $A = E^{-1}\text{rref}(A)$ for some invertible matrix E , compute

$$A^\top = \text{rref}(A)^\top (E^{-1})^\top,$$

and use the fact (which needs proof) that multiplying on the right by an invertible matrix doesn’t change rank. Donezo.

27.3 Problem (+). Here’s how you make all of this much more painfully precise. Let $A \in \mathbb{R}^{m \times n}$ with $\text{rank}(A) = r$ and $R_0 = \text{rref}(A)$.

- (i) Use part (ii) of Problem 26.11 to show that $\text{rank}(A^\top) = \text{rank}(R_0^\top)$.
- (ii) Suppose that R_0 has the very general form

$$R_0 = \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} P,$$

where maybe the F - and/or zero blocks are not present, and P is a permutation matrix that could be the identity. Use Problem 22.6 and the invertibility of P from part (iv) of Problem 26.9 to show that

$$\text{rank}(R_0^\top) = \text{rank} \left(\begin{bmatrix} I_r & 0 \\ F^\top & 0 \end{bmatrix} \right).$$

(iii) Put all of this together to conclude that $\text{rank}(A) = \text{rank}(A^\top)$.

Content from Strang's ILA 6E. Read #1 on p. 130, #4 on p. 131, #1 on p. 132, and #4 on p. 133. Actually, probably best to reread all of pp. 130–133 and see all four subspaces talk to each other.

Now we can build a basis for \mathbb{R}^n . Suppose that $A \in \mathbb{R}^{m \times n}$ has rank r . Then $\dim[\mathbf{N}(A)] = n - r$ by rank-nullity and $\dim[\mathbf{C}(A^\top)] = r$ by the result above. Since $(n - r) + r = n$, this should make us feel optimistic.

27.4 Problem (!). (i) Prove that if A has full column rank ($r = n$), then every $\mathbf{x} \in \mathbb{R}^n$ can be written uniquely in the form $\mathbf{x} = \mathbf{v} + \mathbf{w}$ for $\mathbf{v} \in \mathbf{N}(A)$ and $\mathbf{w} \in \mathbf{C}(A^\top)$. [Hint: *you don't have many choices for \mathbf{v} .*]

(ii) What happens if $r = 0$?

Going forward, suppose $1 \leq r < n$. Let $\mathbf{v}_1, \dots, \mathbf{v}_{n-r}$ be a basis for $\mathbf{N}(A)$ and let $\mathbf{w}_1, \dots, \mathbf{w}_r$ be a basis for $\mathbf{C}(A^\top)$.

27.5 Problem (!). Suppose you know that $\mathbf{v}_1, \dots, \mathbf{v}_{n-r}, \mathbf{w}_1, \dots, \mathbf{w}_r$ are independent. Why do they form a basis for \mathbb{R}^n ? [Hint: *what is $\mathbf{C}(\begin{bmatrix} \mathbf{v}_1 & \cdots & \mathbf{v}_{n-r} & \mathbf{w}_1 & \cdots & \mathbf{w}_r \end{bmatrix})$? Or look at part (iii) of Problem 24.10.*]

I claim that we can check independence using the orthogonality argument that we used on the matrix M from (25.2), and I think you'll learn more by doing that yourself.

27.6 Problem (*). Suppose that

$$y_1 \mathbf{v}_1 + \cdots + y_{n-r} \mathbf{v}_{n-r} + z_1 \mathbf{w}_1 + \cdots + z_r \mathbf{w}_r = \mathbf{0}_n$$

for some $y_j, z_j \in \mathbb{R}$. Put

$$\mathbf{v} = y_1 \mathbf{v}_1 + \cdots + y_{n-r} \mathbf{v}_{n-r} \quad \text{and} \quad \mathbf{w} = z_1 \mathbf{w}_1 + \cdots + z_r \mathbf{w}_r.$$

Explain why $\mathbf{v} \cdot \mathbf{w} = 0$ and $\mathbf{v} + \mathbf{w} = \mathbf{0}_n$. Obtain $\mathbf{v} \cdot \mathbf{v} = 0$, thus $\mathbf{v} = \mathbf{0}$, and therefore $y_j = 0$ for all j . From this, obtain $\mathbf{w} = \mathbf{0}_n$, thus $z_j = 0$ for all j .

This implies that $\mathbf{v}_1, \dots, \mathbf{v}_{n-r}, \mathbf{w}_1, \dots, \mathbf{w}_r$ form a basis for \mathbb{R}^n , and so we can write each

$\mathbf{x} \in \mathbb{R}^n$ as

$$\mathbf{x} = y_1 \mathbf{v}_1 + \cdots + y_{n-r} \mathbf{v}_{n-r} + z_1 \mathbf{w}_1 + \cdots + z_r \mathbf{w}_r$$

for some $y_j, z_j \in \mathbb{R}$. With

$$\mathbf{v} = y_1 \mathbf{v}_1 + \cdots + y_{n-r} \mathbf{v}_{n-r} \quad \text{and} \quad \mathbf{w} = z_1 \mathbf{w}_1 + \cdots + z_r \mathbf{w}_r,$$

this shows that $\mathbf{x} = \mathbf{v} + \mathbf{w}$ for some $\mathbf{v} \in \mathbf{N}(A)$ and $\mathbf{w} \in \mathbf{C}(A^\top)$.

27.7 Problem (★). Prove uniqueness of this decomposition by generalizing the argument that preceded Definition 26.1.

27.8 Remark. *This is one of those times when having a basis for a subspace in the abstract was very useful. Without knowing precisely the forms of the bases for $\mathbf{N}(A)$ (which we could extract from the RREF of A) or $\mathbf{C}(A^\top)$ (which we could extract from the pivot rows of the RREF of A), we built a basis for \mathbb{R}^n and used that to get our desired decomposition.*

Content from Strang's ILA 6E. “Combining Bases from Subspaces” on p. 147 contains these “counting” arguments that lead to a basis for all of \mathbb{R}^n out of bases for $\mathbf{N}(A)$ and $\mathbf{C}(A^\top)$. Read Examples 3 and 4. Then go back to the box on p. 145 with the inequality $\dim(\mathcal{V}) + \dim(\mathcal{W}) \leq n$. Can you prove this? [Hint: *start with bases for \mathcal{V} and \mathcal{W} , and show that together, the vectors in both bases are still independent.*] Can you give an example of orthogonal subspaces for which the inequality is strict? [Hint: *look at some, but not all, of the standard basis vectors.*]

And so we (mostly you, but also me) have proved a pretty big result.

27.9 Theorem. Let $A \in \mathbb{R}^{m \times n}$. For each $\mathbf{x} \in \mathbb{R}^n$ there exist unique $\mathbf{v} \in \mathbf{N}(A)$ and $\mathbf{w} \in \mathbf{C}(A^\top)$ such that $\mathbf{x} = \mathbf{v} + \mathbf{w}$. Also, $\mathbf{v} \cdot \mathbf{w} = 0$. We summarize this symbolically by writing

$$\mathbb{R}^n = \mathbf{N}(A) \oplus \mathbf{C}(A^\top),$$

and we call \mathbb{R}^n the **ORTHOGONAL DIRECT SUM** of $\mathbf{N}(A)$ and $\mathbf{C}(A^\top)$.

27.10 Problem (!). Let $A \in \mathbb{R}^{m \times n}$. Prove that we can write any $\mathbf{b} \in \mathbb{R}^m$ uniquely in the form $\mathbf{b} = \mathbf{v} + \mathbf{w}$ for some $\mathbf{v} \in \mathbf{C}(A)$ and $\mathbf{w} \in \mathbf{N}(A^\top)$, and that $\mathbf{v} \cdot \mathbf{w} = 0$. We summarize this symbolically by writing

$$\mathbb{R}^m = \mathbf{C}(A) \oplus \mathbf{N}(A^\top).$$

[Hint: *replace A in Theorem 27.9 with A^\top .*]

We should not interpret the dual results $\mathbb{R}^n = \mathbf{N}(A) \oplus \mathbf{C}(A^\top)$ and $\mathbb{R}^m = \mathbf{C}(A) \oplus \mathbf{N}(A^\top)$ as saying that any vector in \mathbb{R}^n or \mathbb{R}^m is in one or another of these **FOUR FUNDAMENTAL SUBSPACES** $\mathbf{N}(A)$, $\mathbf{C}(A^\top)$, $\mathbf{C}(A)$, and $\mathbf{N}(A^\top)$ associated with A . Rather, we can build \mathbb{R}^n and \mathbb{R}^m out of the four fundamental subspaces.

27.11 Problem (★). Let

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

Give an example of a vector $\mathbf{v} \in \mathbb{R}^4$ such that $\mathbf{v} \notin \mathbf{N}(A)$ and $\mathbf{v} \notin \mathbf{C}(A^T)$. Then give an example of $\mathbf{w} \in \mathbb{R}^3$ such that $\mathbf{w} \notin \mathbf{C}(A)$ and $\mathbf{w} \notin \mathbf{N}(A^T)$. Feel free to cite prior examples/problems that specify the four fundamental subspaces for this matrix.

The decomposition $\mathbb{R}^m = \mathbf{C}(A) \oplus \mathbf{N}(A^T)$ will tell us something valuable about the solvability of $A\mathbf{x} = \mathbf{b}$, if we work at it. I claim that

$$\mathbf{N}(A^T) = \{\mathbf{v} \in \mathbb{R}^m \mid \mathbf{v} \cdot \mathbf{w} = 0 \text{ for all } \mathbf{w} \in \mathbf{C}(A)\}. \quad (27.1)$$

We basically saw this when we proved that all vectors in $\mathbf{N}(A)$ and $\mathbf{C}(A^T)$ are orthogonal, but let's do it again for practice.

First suppose $A^T\mathbf{v} = \mathbf{0}_n$ and let $\mathbf{w} \in \mathbf{C}(A)$. Then $\mathbf{w} = A\mathbf{x}$ for some $\mathbf{x} \in \mathbb{R}^n$, and we compute

$$\mathbf{v} \cdot \mathbf{w} = \mathbf{v} \cdot A\mathbf{x} = A^T\mathbf{v} \cdot \mathbf{x} = \mathbf{0}_n \cdot \mathbf{x} = 0.$$

Now suppose $\mathbf{v} \cdot \mathbf{w} = 0$ for all $\mathbf{w} \in \mathbf{C}(A)$. We want to show $A^T\mathbf{v} = \mathbf{0}_n$. When all else fails, rewrite what you know. Since $\mathbf{C}(A) = \{A\mathbf{x} \mid \mathbf{x} \in \mathbb{R}^n\}$, we have

$$0 = \mathbf{v} \cdot A\mathbf{x} = A^T\mathbf{v} \cdot \mathbf{x}$$

for all $\mathbf{x} \in \mathbb{R}^n$. Gloriously, this is enough to get us $A^T\mathbf{v} = \mathbf{0}_n$.

Here's why.

27.12 Problem (!). Suppose that $\mathbf{v} \in \mathbb{R}^n$ satisfies $\mathbf{v} \cdot \mathbf{w} = 0$ for all $\mathbf{w} \in \mathbb{R}^n$. Prove that $\mathbf{v} = \mathbf{0}_n$. [Hint: take advantage of that generous quantifier “for all” and let \mathbf{w} be one of the standard basis vectors.]

I view this problem as another way that the dot product extracts information about vectors. If you test or measure a given vector against all vectors under the lens of the dot product and you always get 0, then that given vector is the zero vector. (This makes me feel like a real scientist using lab instruments.)

We conclude the equality (27.1), which motivates a new kind of structure.

27.13 Definition. Let \mathcal{V} be a subset of \mathbb{R}^p (not necessarily a subspace). The **ORTHOGONAL COMPLEMENT** of \mathcal{V} in \mathbb{R}^p is

$$\mathcal{V}^\perp := \{\mathbf{w} \in \mathbb{R}^p \mid \mathbf{v} \cdot \mathbf{w} = 0 \text{ for all } \mathbf{v} \in \mathcal{V}\}.$$

We pronounce the symbol \mathcal{V}^\perp as “vee perp.”

Content from Strang's ILA 6E. The last paragraph on p. 145 defines orthogonal complements.

27.14 Problem (★). Let \mathcal{V} be a subset of \mathbb{R}^p . Prove that \mathcal{V}^\perp is a subspace of \mathbb{R}^p . Convince yourself that you didn't need \mathcal{V} to be a subspace.

27.15 Example. (i) The equality (27.1) says that

$$\mathbf{C}(A)^\perp = \mathbf{N}(A^\top) \quad (27.2)$$

for any $A \in \mathbb{R}^{m \times n}$.

(ii) Let $\mathcal{V} = \mathbb{R}^p$ and suppose that $\mathbf{w} \in \mathbb{R}^p$ with $\mathbf{v} \cdot \mathbf{w} = 0$ for all $\mathbf{v} \in \mathbb{R}^p$. Problem 27.12 says that $\mathbf{w} = \mathbf{0}_p$, so $(\mathbb{R}^p)^\perp = \{\mathbf{0}_p\}$.

(iii) Let $\mathcal{V} = \{\mathbf{0}_p\}$. Then $\mathbf{0}_p \cdot \mathbf{w} = 0$ for any $\mathbf{w} \in \mathbb{R}^p$, so $\{\mathbf{0}_p\}^\perp = \mathbb{R}^p$.

27.16 Problem (!). For $A \in \mathbb{R}^{m \times n}$, prove that $\mathbf{N}(A) = \mathbf{C}(A^\top)^\perp$.

We just saw the extreme cases of $(\mathbb{R}^p)^\perp = \{\mathbf{0}_p\}$ and $\{\mathbf{0}_p\}^\perp = \mathbb{R}^p$. Thus

$$((\mathbb{R}^p)^\perp)^\perp = \{\mathbf{0}_p\}^\perp = \mathbb{R}^p \quad \text{and} \quad (\{\mathbf{0}_p\}^\perp)^\perp = (\mathbb{R}^p)^\perp = \{\mathbf{0}_p\}.$$

27.17 Problem (!). Here is a less extreme case in \mathbb{R}^2 . Let $\mathcal{V} = \text{span}(\mathbf{e}_1)$, $\mathbf{e}_1 = (1, 0)$. Draw pictures to convince yourself that $(\mathcal{V}^\perp)^\perp = \mathcal{V}$ and then prove it.

These examples might make us wonder the following.

27.18 Conjecture. $(\mathcal{V}^\perp)^\perp = \mathcal{V}$ for any subset \mathcal{V} of \mathbb{R}^p .

In particular, if true Conjecture 27.18 would imply

$$\mathbf{C}(A) = (\mathbf{C}(A)^\perp)^\perp = \mathbf{N}(A^\top)^\perp. \quad (27.3)$$

This would give us a new way of deciding solvability of $A\mathbf{x} = \mathbf{b}$: check that \mathbf{b} is orthogonal to everything in $\mathbf{N}(A^\top)$. Or that \mathbf{b} is orthogonal to a basis for $\mathbf{N}(A^\top)$.

27.19 Problem (!). Let \mathcal{V} be a subspace of \mathbb{R}^p and let $\mathbf{v}_1, \dots, \mathbf{v}_d$ be a basis for \mathcal{V} . Suppose that $\mathbf{v} \in \mathbb{R}^p$ satisfies $\mathbf{v} \cdot \mathbf{v}_j = 0$ for $j = 1, \dots, d$. Prove that $\mathbf{v} \in \mathcal{V}^\perp$.

The point is that if (27.3) is true, then it would give us a different way of describing the column space. In particular, we might get an easier way of checking that a vector is *not* in the column space than doing elementary row operations and going to the RREF.

27.20 Problem (★). Assuming Conjecture 27.18 to be true, prove the **FREDHOLM ALTERNATIVE**: if $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$, then one, and only one, of the following is true.

- (i) $\mathbf{b} \in \mathbf{C}(A)$ and so the problem $A\mathbf{x} = \mathbf{b}$ has a solution.
- (ii) There is $\mathbf{v} \in \mathbb{R}^m$ such that $A^T \mathbf{v} = \mathbf{0}_n$ and $\mathbf{b} \cdot \mathbf{v} \neq 0$.

In the second case, $\mathbf{b} \notin \mathbf{C}(A)$. So who cares? If you want to consider the problem $A\mathbf{x} = \mathbf{b}$ for a bunch of \mathbf{b} , I think it might be easier to figure out $\mathbf{N}(A^T)$ and then see if the \mathbf{b} are orthogonal to every vector in a basis for $\mathbf{N}(A^T)$. Those orthogonality relations give “solvability conditions” for $A\mathbf{x} = \mathbf{b}$.

Day 28: Wednesday, March 19.

We are going to prove Conjecture 27.18. First, it is not all that hard to show that any vector in \mathcal{V} is also in $(\mathcal{V}^\perp)^\perp$. Let $\mathbf{v} \in \mathcal{V}$. We want to show $\mathbf{v} \cdot \mathbf{w} = 0$ for all $\mathbf{w} \in \mathcal{V}^\perp$. But that is exactly what it means for \mathbf{w} to be a vector in \mathcal{V}^\perp !

Now let $\mathbf{x} \in (\mathcal{V}^\perp)^\perp$. Why do we have $\mathbf{x} \in \mathcal{V}$? This is a bit harder.

28.1 Problem (!). I want to convince you that this is true in a comfortingly familiar particular case. Let

$$A = \begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix}.$$

Show that if $\mathbf{w} \in \mathbf{N}(A^T)^\perp$, then $\mathbf{w} \in \mathbf{C}(A)$. Feel free to refer to Example 21.12 (which tells you about $\mathbf{C}(A)$) and Example 22.1 (which tells you about $\mathbf{N}(A^T)$). Also,

Going forward, we need a bit of a trick. Any subspace \mathcal{V} of \mathbb{R}^p has the form $\mathcal{V} = \mathbf{C}(B)$ for some matrix $B \in \mathbb{R}^{p \times d}$, where $d = \dim(\mathcal{V})$. Let $A = B^T$, so $A \in \mathbb{R}^{d \times p}$ and $\mathcal{V} = \mathbf{C}(A^T)$. Then

$$\mathcal{V}^\perp = \mathbf{C}(A^T)^\perp = \mathbf{N}(A) \tag{28.1}$$

by Problem 27.16. And by Theorem 27.9, we can write any $\mathbf{x} \in \mathbb{R}^p$ as $\mathbf{x} = \mathbf{v} + \mathbf{w}$ for unique $\mathbf{v} \in \mathbf{C}(A^T) = \mathcal{V}$ and $\mathbf{w} \in \mathbf{N}(A) = \mathcal{V}^\perp$. (I realize I’m flipping the roles of \mathbf{v} and \mathbf{w} from that theorem, but I want to keep \mathbf{v} as the label for things in \mathcal{V} , which is $\mathbf{C}(A^T)$. Sue me.) Here is what we have proved.

28.2 Lemma. Let \mathcal{V} be a subspace of \mathbb{R}^p . For each $\mathbf{x} \in \mathbb{R}^p$, there exist unique $\mathbf{v} \in \mathcal{V}$ and $\mathbf{w} \in \mathcal{V}^\perp$ such that $\mathbf{x} = \mathbf{v} + \mathbf{w}$.

28.3 Problem (!). If \mathcal{V} is a subspace of \mathbb{R}^p , prove that $\dim(\mathcal{V}^\perp) = p - \dim(\mathcal{V})$.

Here is why all of this jumping around between subspaces and matrices matters. Start

with a subspace \mathcal{V} of \mathbb{R}^p and $\mathbf{x} \in (\mathcal{V}^\perp)^\perp$. We want to show $\mathbf{x} \in \mathcal{V}$. The key thing is that \mathcal{V}^\perp is also a subspace of \mathbb{R}^p . This was Problem 27.14.

Now apply Lemma 28.2 with \mathcal{V}^\perp in place of \mathcal{V} . So, we can write any $\mathbf{x} \in (\mathcal{V}^\perp)^\perp$ as $\mathbf{x} = \mathbf{v} + \mathbf{w}$ for (unique) $\mathbf{v} \in \mathcal{V}$ and $\mathbf{w} \in \mathcal{V}^\perp$. If we can show $\mathbf{w} = \mathbf{0}_p$, then we'll have $\mathbf{x} = \mathbf{v} \in \mathcal{V}$. The trick is subtraction:

$$\mathbf{w} = \mathbf{x} - \mathbf{v} \in (\mathcal{V}^\perp)^\perp.$$

This is because we're assuming $\mathbf{x} \in (\mathcal{V}^\perp)^\perp$, and just above we showed that if $\mathbf{v} \in \mathcal{V}$, then $\mathbf{v} \in (\mathcal{V}^\perp)^\perp$. But then $\mathbf{w} \in \mathcal{V}^\perp$ and $\mathbf{w} \in (\mathcal{V}^\perp)^\perp$. I claim this means $\mathbf{w} = \mathbf{0}_p$.

28.4 Problem (★). Prove that if \mathcal{V} is a subspace of \mathbb{R}^p and $\mathbf{v} \in \mathbb{R}^p$ with both $\mathbf{v} \in \mathcal{V}$ and $\mathbf{v} \in \mathcal{V}^\perp$, then $\mathbf{v} = \mathbf{0}_p$. Draw a picture illustrating this phenomenon in \mathbb{R}^2 .

This completes the proof of Conjecture 27.18.

28.5 Problem (★). Did that feel hard? It always feels hard to me. I wish I could just start with $\mathbf{v} \in \mathbf{N}(A^\top)^\perp$ and show $\mathbf{v} \in \mathbf{C}(A)$. Try doing that. Where do you get stuck?

Let's upgrade the conjecture to a theorem.

28.6 Theorem. Let \mathcal{V} be a subspace of \mathbb{R}^p . Then $(\mathcal{V}^\perp)^\perp = \mathcal{V}$.

28.7 Corollary. Let $A \in \mathbb{R}^{m \times n}$. Then $\mathbf{C}(A) = \mathbf{N}(A^\top)^\perp$.

28.8 Problem (!). A ton of machinery went into proving Corollary 28.7. Spend at least 15 minutes reviewing its proof. Do you understand all of the vocabulary and symbols involved? How does the logic feel? Did you do all of the (!)- and (★)-problems cited in the proof?

Here is a summary of all of our work. This answers the question “What is missing beyond the null space or the column space?” and provides a complete overview of how a matrix in $\mathbb{R}^{m \times n}$ determines the structure of \mathbb{R}^n and \mathbb{R}^m .

28.9 Theorem (Fundamental theorem of linear algebra). Let $A \in \mathbb{R}^{m \times n}$.

- (i) $\mathbb{R}^n = \mathbf{N}(A) \oplus \mathbf{C}(A^\top)$
- (ii) $\mathbb{R}^m = \mathbf{C}(A) \oplus \mathbf{N}(A^\top)$
- (iii) $\mathbf{N}(A) = \mathbf{C}(A^\top)^\perp$
- (iv) $\mathbf{C}(A) = \mathbf{N}(A^\top)^\perp$
- (v) $\dim[\mathbf{N}(A)] = n - \text{rank}(A)$

$$(vi) \text{ rank}(A) = \text{rank}(A^T)$$

Content from Strang's ILA 6E. Figure 4.1 on p. 146 says all of this. Study the figure carefully and read the paragraph following its caption. I like to start reading the figure by beginning with \mathbf{b} , then tracking it back to \mathbf{x}_r and \mathbf{x}_n . (The subscript n there is for “null space,” not the n in \mathbb{R}^n .)

28.10 Problem (+). Here is a more general way to view that symbol \oplus . Suppose that \mathcal{V} and \mathcal{W} are subspaces of \mathbb{R}^p . Write $\mathbb{R}^p = \mathcal{V} \oplus \mathcal{W}$ and say that \mathbb{R}^p is the orthogonal direct sum of \mathcal{V} and \mathcal{W} if for all $\mathbf{x} \in \mathbb{R}^p$, there are $\mathbf{v} \in \mathcal{V}$ and $\mathbf{w} \in \mathcal{W}$ such that $\mathbf{x} = \mathbf{v} + \mathbf{w}$, and if $\mathbf{v} \cdot \mathbf{w} = 0$ for all $\mathbf{v} \in \mathcal{V}$ and $\mathbf{w} \in \mathcal{W}$. These properties are enough to recover first uniqueness of the decomposition of \mathbf{x} and second $\mathcal{W} = \mathcal{V}^\perp$.

(i) Use orthogonality and the strategy of the argument preceding Definition 26.1 to prove that the decomposition is unique.

(ii) We show here that $\mathcal{V}^\perp = \mathcal{W}$. Since $\mathbf{v} \cdot \mathbf{w} = 0$ for all $\mathbf{v} \in \mathcal{V}$ and $\mathbf{w} \in \mathcal{W}$, we know that if $\mathbf{w} \in \mathcal{W}$, then $\mathbf{w} \in \mathcal{V}^\perp$. If $\mathbf{x} \in \mathcal{V}^\perp$, write $\mathbf{x} = \mathbf{v} + \mathbf{w}$ for some $\mathbf{v} \in \mathcal{V}$, $\mathbf{w} \in \mathcal{W}$, so, since $\mathbf{w} \in \mathcal{V}^\perp$, too, we get $\mathbf{v} = \mathbf{x} - \mathbf{w} \in \mathcal{V}^\perp$. Why does this prove $\mathbf{x} \in \mathcal{W}$?

There is just one major problem with our fundamental theorem: all of these results are highly existential. We developed those existential results by starting with the null space and asking “What else is missing from \mathbb{R}^n ?” Now we’ll start with the column space. Specifically, it’s great at a theoretical level to say that $\mathbb{R}^m = \mathbf{C}(A) \oplus \mathbf{N}(A^T)$ in the sense that each $\mathbf{b} \in \mathbb{R}^m$ can be written uniquely as $\mathbf{b} = \mathbf{v} + \mathbf{w}$ for some $\mathbf{v} \in \mathbf{C}(A)$ and $\mathbf{w} \in \mathbf{N}(A^T)$ and that $\mathbf{v} \cdot \mathbf{w} = 0$, but how do we find those \mathbf{v} and \mathbf{w} explicitly and easily?

First, we only need one of them. For if we know $\mathbf{b} = \mathbf{v} + \mathbf{w}$, then $\mathbf{w} = \mathbf{b} - \mathbf{v}$. So how do we get \mathbf{v} ?

Well, how do we do anything in this course? *We multiply by a matrix.* Can we find $P \in \mathbb{R}^{m \times m}$ such that if $\mathbf{b} \in \mathbb{R}^m$, then $P\mathbf{b} \in \mathbf{C}(A)$ and $\mathbf{b} - P\mathbf{b} \in \mathbf{N}(A^T)$. Then we have $\mathbf{b} = P\mathbf{b} + (\mathbf{b} - P\mathbf{b})$ as our decomposition.

28.11 Problem (!). Explain why you expect $P^2 = P$. (Have we talked about matrix powers? Just in case: $P^2 = PP$.) [Hint: *what does P do?* We have $P\mathbf{b} \in \mathbf{C}(A)$ and any \mathbf{b} can be written uniquely as $\mathbf{b} = P\mathbf{b} + \mathbf{w}$ with $\mathbf{w} \in \mathbf{N}(A^T)$. If $\mathbf{b} \in \mathbf{C}(A)$ already, what is that \mathbf{w} , and so what should $P\mathbf{b}$ be? Then what is $P^2\mathbf{b}$?]

Content from Strang's ILA 6E. Read all of pp. 151–152 up to, but not including, “Projection Onto a Line.” This is the mission statement of Section 4.2, and it’s a very helpful overview of where we’re going.

It turns out to be very helpful to assume that A has full column rank (= all of its columns are independent = all of its columns are pivot columns). This is not as huge a restriction


as you might initially think. After all, $\mathbf{C}(A) = \mathbf{C}(\tilde{A})$, where \tilde{A} is the matrix containing just the pivot columns of A , and \tilde{A} has full column rank. So, if we are going to understand the decomposition $\mathbb{R}^m = \mathbf{C}(A) \oplus \mathbf{N}(A^\top)$, we may as well do it when A has full column rank.

We'll do this first in the case that A has only one column, in which case $\mathbf{C}(A) = \text{span}(\mathbf{a})$ for some $\mathbf{a} \neq \mathbf{0}_m$. This has some transparent geometry and will give a useful auxiliary result for later.

So here is what we want: given $\mathbf{b} \in \mathbb{R}^m$, there are (necessarily unique) $\mathbf{v} \in \mathbf{C}(A) = \text{span}(\mathbf{a})$ and $\mathbf{w} \in \mathbf{N}(A^\top)$ such that $\mathbf{b} = \mathbf{v} + \mathbf{w}$. Since $\mathbf{v} \in \text{span}(\mathbf{a})$, we can write $\mathbf{v} = c\mathbf{a}$ for some $c \in \mathbb{R}$. Then

$$\mathbf{b} = c\mathbf{a} + \mathbf{w}.$$

Here is a picture of what's going on when $m = 2$ and \mathbf{a} is a multiple of $\mathbf{e}_1 = (1, 0)$.



The two unknowns $c \in \mathbb{R}$ and $\mathbf{w} \in \mathbf{N}(A^\top)$ have to satisfy this one equation—not a recipe for success—but remember that we have an orthogonality condition:

$$0 = \mathbf{v} \cdot \mathbf{w} = (c\mathbf{a}) \cdot \mathbf{w} = c(\mathbf{a} \cdot \mathbf{w}).$$

Actually, since $\mathbf{a} \in \mathbf{C}(A)$ and $\mathbf{w} \in \mathbf{N}(A^\top)$, we always have

$$\mathbf{a} \cdot \mathbf{w} = 0,$$

so forget about c there. Now we have two equations and two unknowns:

$$\begin{cases} \mathbf{b} = c\mathbf{a} + \mathbf{w} \\ \mathbf{a} \cdot \mathbf{w} = 0. \end{cases} \quad (28.2)$$

A little algebraic trickery will reduce this to one equation: rewrite

$$\mathbf{w} = \mathbf{b} - c\mathbf{a}$$

and plug in to get

$$0 = \mathbf{a} \cdot (\mathbf{b} - c\mathbf{a}).$$

Rearrange a little:

$$0 = (\mathbf{a} \cdot \mathbf{b}) - c(\mathbf{a} \cdot \mathbf{a}),$$

and a little more:

$$c(\mathbf{a} \cdot \mathbf{a}) = \mathbf{a} \cdot \mathbf{b},$$

and divide:

$$c = \frac{\mathbf{a} \cdot \mathbf{b}}{\mathbf{a} \cdot \mathbf{a}}.$$

This division is perfectly legal since $\mathbf{a} \neq \mathbf{0}_m$, and therefore $\mathbf{a} \cdot \mathbf{a} \neq 0$.

We worked backwards, so we should check our work. Certainly

$$\left(\frac{\mathbf{a} \cdot \mathbf{b}}{\mathbf{a} \cdot \mathbf{a}} \right) \mathbf{a} \in \text{span}(\mathbf{a}).$$

28.12 Problem (!). Let

$$\mathbf{w} = \mathbf{b} - \left(\frac{\mathbf{a} \cdot \mathbf{b}}{\mathbf{a} \cdot \mathbf{a}} \right) \mathbf{a}.$$

Check that $\mathbf{a} \cdot \mathbf{w} = 0$, so $\mathbf{w} \in \mathbf{C}(A)^\perp = \mathbf{N}(A^\top)$.

28.13 Problem (★). Where is P ? This requires a bit of sleight-of-hand. We want

$$P\mathbf{b} = \left(\frac{\mathbf{a} \cdot \mathbf{b}}{\mathbf{a} \cdot \mathbf{a}} \right) \mathbf{a}.$$

Here it is helpful to think of column vectors as $m \times 1$ matrices and the dot product as the matrix product

$$\mathbf{a} \cdot \mathbf{b} = \mathbf{a}^\top \mathbf{b}.$$

It's also helpful to break our usual convention of how we write scalar multiplication and allow $\mathbf{c}\mathbf{a} = \mathbf{a}\mathbf{c}$. I don't like it, either. Show then that

$$\left(\frac{\mathbf{a} \cdot \mathbf{b}}{\mathbf{a} \cdot \mathbf{a}} \right) \mathbf{a} = \frac{1}{\mathbf{a} \cdot \mathbf{a}} (\mathbf{a}\mathbf{a}^\top) \mathbf{b},$$

so

$$P = \frac{1}{\mathbf{a} \cdot \mathbf{a}} (\mathbf{a}\mathbf{a}^\top).$$

28.14 Example. Computations with this sort of “projection” onto $\text{span}(\mathbf{a})$ can become bulky, so let's check how this respects our intuition. Say $\mathbf{a} = \mathbf{e}_1$ in \mathbb{R}^2 . Of course, we expect

$$\mathbf{b} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = b_1 \mathbf{e}_1 + \begin{bmatrix} 0 \\ b_2 \end{bmatrix}.$$

We compute

$$\left(\frac{\mathbf{e}_1 \cdot \mathbf{b}}{\mathbf{e}_1 \cdot \mathbf{e}_1} \right) \mathbf{e}_1 = \frac{b_1}{1} \mathbf{e}_1 = b_1 \mathbf{e}_1.$$

How nice it was that $\mathbf{e}_1 \cdot \mathbf{e}_1 = 1$.

Content from Strang's *ILA* 6E. Read “Projection Onto a Line” from pp. 152–154. Check Examples 1 and 2.

Day 29: Friday, March 21.

Now we consider the general case in which A has an arbitrary number of columns. Remember, though, that A still has full column rank, and so those columns are independent. Again, we start with $\mathbf{b} \in \mathbb{R}^m$, and we want to find $\mathbf{v} \in \mathbf{C}(A)$ and $\mathbf{w} \in \mathbf{N}(A^\top)$ such that

$$\mathbf{b} = \mathbf{v} + \mathbf{w} \quad \text{and} \quad \mathbf{v} \cdot \mathbf{w} = 0.$$

Before, when A had just one column, we rewrote \mathbf{v} as a scalar multiple of that column. Now we can say $\mathbf{v} = A\hat{\mathbf{x}}$ for some $\hat{\mathbf{x}} \in \mathbb{R}^n$. (The hat is sort of traditional.) The analogue of (28.2) is now

$$\begin{cases} \mathbf{b} = A\hat{\mathbf{x}} + \mathbf{w} \\ \mathbf{a}_j \cdot \mathbf{w} = 0, \quad j = 1, \dots, n. \end{cases}$$

That second (set of) equation(s) is the orthogonality of \mathbf{w} to everything in $\mathbf{C}(A)$, equivalently, to the columns of A .

This reduces to n equations:

$$0 = \mathbf{a}_j \cdot \mathbf{w} = \mathbf{a}_j \cdot (\mathbf{b} - A\hat{\mathbf{x}}), \quad j = 1, \dots, n.$$

And now for the trick: rewrite $\mathbf{a}_j = A\mathbf{e}_j$, so

$$0 = (A\mathbf{e}_j) \cdot (\mathbf{b} - A\hat{\mathbf{x}}) = \mathbf{e}_j \cdot A^\top(\mathbf{b} - A\hat{\mathbf{x}}), \quad j = 1, \dots, n.$$

The powerful Problem 27.12 implies that

$$A^\top(\mathbf{b} - A\hat{\mathbf{x}}) = \mathbf{0}_n,$$

which rearranges to

$$A^\top A\hat{\mathbf{x}} = A^\top \mathbf{b}.$$

If only $A^\top A$ were invertible, we could peel it off to solve for $\hat{\mathbf{x}}$:

$$\hat{\mathbf{x}} = (A^\top A)^{-1} A^\top \mathbf{b}.$$

If only. Then we would have

$$\mathbf{v} = A\hat{\mathbf{x}} = A(A^\top A)^{-1} A^\top \mathbf{b},$$

and so putting

$$P := A(A^\top A)^{-1} A^\top$$

would give $P\mathbf{b} \in \mathbf{C}(A)$, $\mathbf{b} - P\mathbf{b} \in \mathbf{N}(A^\top)$, and $\mathbf{b} = P\mathbf{b} + (\mathbf{b} - P\mathbf{b})$.

Good news: $A^\top A$ is invertible here.

29.1 Lemma. *Let $A \in \mathbb{R}^{m \times n}$ have full column rank. Then $A^\top A$ is invertible.*

Proof. We might initially think $A^T A =$ independent rows in A^T dotted with independent columns in A has to give us something good. It does, but the trick is to show $\mathbf{N}(A^T A) = \{\mathbf{0}_n\}$. For if $A^T A \mathbf{x} = \mathbf{0}_n$, then

$$0 = \mathbf{x} \cdot \mathbf{0}_n = \mathbf{x} \cdot (A^T A \mathbf{x}) = (A \mathbf{x}) \cdot (A \mathbf{x}),$$

and so $A \mathbf{x} = \mathbf{0}_n$, thus $\mathbf{x} \in \mathbf{N}(A)$. Since A has full column rank, $\mathbf{N}(A) = \{\mathbf{0}_n\}$, so $\mathbf{x} = \mathbf{0}_n$. This is just one of those classical tricks that I never would have thought of myself if someone else hadn't shown it to me, but now it feels like instinct. ■

Content from Strang's ILA 6E. This lemma is proved on p. 157. Read the warning at the top of the page and then the calculations at the bottom of the page of how this breaks when A has dependent columns.

Bad news: this was a lot of working backward.

29.2 Problem (★). Let $A \in \mathbb{R}^{m \times n}$ have full column rank and set

$$P_A := A(A^T A)^{-1} A^T.$$

(i) Explain why, just from looking at P_A , every vector in $\mathbf{C}(P_A)$ is in $\mathbf{C}(A)$.

(ii) Show that $P_A^2 = P_A$.

(iii) Conclude that if $\mathbf{b} \in \mathbf{C}(A)$, then $P_A \mathbf{b} \in \mathbf{C}(A)$ and thus $\mathbf{C}(A) = \mathbf{C}(P_A)$.

(iv) Show that $P_A^T = P_A$.

(v) Justify each of the following equalities:

$$\mathbf{N}(P_A) = \mathbf{N}(P_A^T) = \mathbf{C}(P_A)^\perp = \mathbf{C}(A)^\perp = \mathbf{N}(A^T). \quad (29.1)$$

(vi) Explain why $\mathbf{b} - P_A \mathbf{b} \in \mathbf{N}(P_A)$ for each $\mathbf{b} \in \mathbb{R}^m$. [Hint: *just compute it.*] So, it's also true that $\mathbf{b} - P_A \mathbf{b} \in \mathbf{N}(A^T)$.

The fruit of this problem is that we can write any $\mathbf{b} \in \mathbb{R}^m$ as

$$\mathbf{b} = \mathbf{v} + \mathbf{w}, \quad \mathbf{v} = P_A \mathbf{b}, \quad \mathbf{w} = \mathbf{b} - P_A \mathbf{b},$$

and we'll have $\mathbf{v} \in \mathbf{C}(A)$ and $\mathbf{w} \in \mathbf{N}(A^T)$. Problem 28.10 assures us that this decomposition is unique.

29.3 Problem (★). If you wanted to find the decomposition of some $\mathbf{x} \in \mathbb{R}^n$ as a sum of vectors in $\mathbf{N}(A)$ and $\mathbf{C}(A^T)$, discuss how you would use P_{A^T} . Is there any relation between P_{A^T} and P_A^T , or is that just wishful, and inappropriate, juggling of the symbol T?

Content from Strang's ILA 6E. Pages 155–156 develop all of this. I don't think memorizing equations (5), (6), and (7) is a good idea, or even memorizing the structure of our P_A above. I think it's more important to be able to *replicate* the derivation of P_A on your own. Check Worked Example 4.2 A on p. 158.

29.4 Example. None of this will work for my favorite matrix

$$A = \begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix},$$

since it does not have full column rank. However, we saw in Example 23.1 that $\mathbf{C}(A) = \mathbf{C}(B)$, where

$$B = \begin{bmatrix} 1 & 0 \\ 2 & 0 \\ 0 & 1 \end{bmatrix},$$

and B does have full column rank. Then we can write any $\mathbf{b} \in \mathbb{R}^3$ uniquely as $\mathbf{b} = P_B \mathbf{b} + \mathbf{w}$, $\mathbf{w} = \mathbf{b} - P_B \mathbf{b} \in \mathbf{N}(A^\top)$. We compute

$$B^\top B = \begin{bmatrix} 1 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 2 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 5 & 0 \\ 0 & 1 \end{bmatrix},$$

$$(B^\top B)^{-1} = \begin{bmatrix} 1/5 & 0 \\ 0 & 1 \end{bmatrix},$$

and

$$\begin{aligned} P_B &= B(B^\top B)^{-1}B^\top \\ &= \begin{bmatrix} 1 & 0 \\ 2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1/5 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 \\ 2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1/5 & 2/5 & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} 1/5 & 2/5 & 0 \\ 2/5 & 4/5 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \end{aligned}$$

I think that was the only time we've ever needed an inverse explicitly. How nice that it was diagonal—what would have happened if I had selected the pivot columns of A instead? (Think carefully about the fact that the columns of B were orthogonal...)

29.5 Problem (!). For A in the previous example, we know from past experience that $\mathbf{b} \in \mathbf{C}(A)$ if and only if $b_2 = 2b_1$. Check that

$$P_B \mathbf{b} \cdot \mathbf{e}_2 = 2(P_B \mathbf{b} \cdot \mathbf{e}_1).$$

The matrix P_A that we cooked up deserves a special name.

29.6 Definition. Let $P \in \mathbb{R}^{m \times m}$.

- (i) P is a **PROJECTION** if $P^2 = P$.
- (ii) P is an **ORTHOGONAL PROJECTION** if $P^2 = P$ and $P^\top = P$.
- (iii) Let \mathcal{V} be a subspace of \mathbb{R}^m . The matrix P is an **ORTHOGONAL PROJECTION ONTO** \mathcal{V} if P is an orthogonal projection with $\mathbf{C}(P) = \mathcal{V}$.

29.7 Problem (!). Without doing any matrix calculations, in \mathbb{R}^3 , what do you expect an orthogonal projection onto $\text{span}(\mathbf{e}_1, \mathbf{e}_2)$ to be? Now do those calculations.

Day 30: Monday, March 24.

30.1 Example. (i) The $m \times m$ matrix whose entries are all 0 and I_m are both orthogonal projections.

(ii) Problem 29.2 shows that P_A is an orthogonal projection onto $\mathbf{C}(A)$ when $A \in \mathbb{R}^{m \times n}$ has full column rank.

(iii) Let \mathcal{V} be a subspace of \mathbb{R}^m . If $\mathcal{V} = \{\mathbf{0}_m\}$, then the matrix $P \in \mathbb{R}^{m \times m}$ whose entries are all 0 is an orthogonal projection onto \mathcal{V} . Otherwise, let $d = \dim(\mathcal{V})$ and let $A \in \mathbb{R}^{m \times d}$ be a matrix whose columns are a basis for \mathcal{V} . Then P_A is an orthogonal projection onto \mathcal{V} .

After all of our work, we'd probably like to call P_A *the* orthogonal projection *onto* $\mathbf{C}(A)$. Is it unique? What an insult it would be if it weren't. This is asking if there is only one $P \in \mathbb{R}^{m \times m}$ such that $P^2 = P$, $P^\top = P$, and $\mathbf{C}(P) = \mathbf{C}(A)$. And this is true. You can have the fun of proving this over the course of two problems.

30.2 Problem (!). Here is a generalization of our orthogonal decomposition results for the case when you start with an orthogonal projection, whether or not that orthogonal projection has full rank. Let $P \in \mathbb{R}^{m \times m}$ be an orthogonal projection and let $\mathbf{b} \in \mathbb{R}^m$. Prove that there exists a unique $\mathbf{w} \in \mathbf{N}(P)$ such that

$$\mathbf{b} = P\mathbf{b} + \mathbf{w} \quad \text{and} \quad P\mathbf{b} \cdot \mathbf{w} = 0.$$

[Hint: the desired decomposition forces $\mathbf{w} = \mathbf{b} - P\mathbf{b}$, which gives uniqueness. Check that $\mathbf{w} \in \mathbf{N}(P)$. Then use the orthogonality of $\mathbf{C}(P)$ and $\mathbf{N}(P^\top)$, or brute-force compute

$$Pb \cdot (\mathbf{b} - P\mathbf{b}).]$$

30.3 Problem (★). Let $P_1, P_2 \in \mathbb{R}^{m \times m}$ be orthogonal projections with $\mathbf{C}(P_1) = \mathbf{C}(P_2)$. We'll show $P_1 = P_2$ by showing $P_1\mathbf{b} = P_2\mathbf{b}$ for all $\mathbf{b} \in \mathbb{R}^m$. (*What things do defines what things are.*)

(i) Use Problem 30.2 to explain why we can write any $\mathbf{b} \in \mathbb{R}^m$ as both

$$\mathbf{b} = P_1\mathbf{b} + \mathbf{w}_1 \quad \text{and} \quad \mathbf{b} = P_2\mathbf{b} + \mathbf{w}_2$$

for some $\mathbf{w}_1 \in \mathbf{N}(P_1)$ and $\mathbf{w}_2 \in \mathbf{N}(P_2)$.

(ii) Show that $\mathbf{N}(P_1) = \mathbf{N}(P_2)$. [Hint: $\mathbf{N}(P) = \mathbf{C}(P^\top)^\perp$ for any $P \in \mathbb{R}^{m \times m}$.]

(iii) Conclude that

$$P_1\mathbf{b} - P_2\mathbf{b} = \mathbf{w}_2 - \mathbf{w}_1$$

and so $P_1\mathbf{b} - P_2\mathbf{b} \in \mathbf{C}(P_1)$ and $P_1\mathbf{b} - P_2\mathbf{b} \in \mathbf{N}(P_1)$. Invoke Problem 28.4.

We now possess a much deeper understanding of how a matrix induces *structure* from the decompositions $\mathbb{R}^n = \mathbf{N}(A) \oplus \mathbf{C}(A^\top)$ and $\mathbb{R}^m = \mathbf{C}(A) \oplus \mathbf{N}(A^\top)$ for $A \in \mathbb{R}^{m \times n}$, and how to perform those decompositions via matrix multiplication. We also have the characterization $\mathbf{C}(A) = \mathbf{N}(A^\top)^\perp$ and the resulting “solvability conditions” from the Fredholm alternative (Problem 27.20). I think this is a great piece of *narrative* and incorporates some geometry to boot into our very algebraic world.

What else do we gain from these results? There's been something of a dichotomy in our approach to linear systems. Either we can solve $A\mathbf{x} = \mathbf{b}$ (uniquely or not) or we can't. We've focused on the solving part. Otherwise, if $A\mathbf{x} = \mathbf{b}$ has no solution, what's the point in talking about it?

Very often in life, mathematically or otherwise, we can't solve the problems that we face. The next best thing is to solve an easier problem. (If your question is too hard, give up and ask a different question.) If we can't solve $A\mathbf{x} = \mathbf{b}$, could we solve a related problem $A\hat{\mathbf{x}} = \mathbf{p}$ and view that related problem as an approximation to our desired problem? Yes! If we pick the right problem.

If we're going to solve $A\hat{\mathbf{x}} = \mathbf{p}$, we need $\mathbf{p} \in \mathbf{C}(A)$. Is there some “ideal” \mathbf{p} to pick relative to the \mathbf{b} that won't work? Again, yes!

Approximating requires a new concept: the notion of *size*, which is really a notion of *length*. The following definition generalizes the notion that the length of the line segment in two dimensions from the origin $(0, 0)$ to a point (x, y) is $\sqrt{x^2 + y^2}$.

30.4 Definition. The **NORM** of $\mathbf{v} = (v_1, \dots, v_m) \in \mathbb{R}^m$ is


$$\|\mathbf{v}\| := (v_1^2 + \dots + v_m^2)^{1/2} = \sqrt{\mathbf{v} \cdot \mathbf{v}}.$$

30.5 Example. If $\mathbf{v} = (1, 2, 3)$, then

$$\|\mathbf{v}\| = \sqrt{1^2 + 2^2 + 3^2} = \sqrt{1 + 4 + 9} = \sqrt{14}.$$

Content from Strang's ILA 6E. Reread all of p. 9 right now. There are plenty of other meaningful ways of measuring the length of a vector in \mathbb{R}^m that we won't need. You might enjoy reading pp. 355–356 up to and including Figure 9.8.

Length and orthogonality interact in a helpful way. You know this already because you believe the Pythagorean theorem, which the definition of $\|\cdot\|$ is basically designed to respect.



30.6 Theorem (Pythagorean theorem). Let $\mathbf{v}, \mathbf{w} \in \mathbb{R}^m$ with $\mathbf{v} \cdot \mathbf{w} = 0$. Then

$$\|\mathbf{v} + \mathbf{w}\|^2 = \|\mathbf{v}\|^2 + \|\mathbf{w}\|^2.$$

30.7 Problem (!). Prove it! [Hint: use the definition $\|\mathbf{x}\|^2 = \mathbf{x} \cdot \mathbf{x}$ to compute $\|\mathbf{v} + \mathbf{w}\|^2$ and get $\mathbf{v} \cdot \mathbf{w} = 0$ to show up somewhere.]

30.8 Problem (★). Let $\mathbf{v}, \mathbf{w} \in \mathbb{R}^m$ with $\mathbf{v} \cdot \mathbf{w} = 0$. Prove that

$$\|\mathbf{v}\| \leq \|\mathbf{v} + \mathbf{w}\|.$$

[Hint: use the Pythagorean theorem to get an expression for $\|\mathbf{v} + \mathbf{w}\|$, and then use the facts that the square root function is increasing and $\|\mathbf{w}\| \geq 0$.]

Here is how we use this new tool of the norm. We'll think that two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^m$ are "close" if the difference $\|\mathbf{v} - \mathbf{w}\|$ is "small."


30.9 Remark. And what exactly does "small" mean? Say $\|\mathbf{v}\| < \epsilon$ for some $\epsilon > 0$. Then since the square root is increasing,

$$|v_j| = \sqrt{v_j^2} \leq \sqrt{v_1^2 + \cdots + v_m^2} = \|\mathbf{v}\| < \epsilon.$$

So if $\|\mathbf{v}\|$ is "small" in the sense that it's less than some threshold $\epsilon > 0$, then each component v_j is "small" in the same way: $|v_j| < \epsilon$ for all j . I think that a vector with small components is a small vector.

Given $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$ with $\mathbf{b} \notin \mathbf{C}(A)$, can we find $\mathbf{p} \in \mathbf{C}(A)$ such that \mathbf{b} and \mathbf{p} are “close”? Then maybe solving $A\hat{\mathbf{x}} = \mathbf{p}$ will be an adequate substitute for failing to solve $A\mathbf{x} = \mathbf{b}$.

Gloriously, it’s quite easy to find this \mathbf{p} , and it probably won’t surprise you that it involves the orthogonal projection $P_A\mathbf{b}$. Here’s a picture that we drew before, more or less.



Here I’m thinking that

$$A = [\mathbf{a}] \in \mathbb{R}^{m \times 1}, \quad \mathbf{a} \neq \mathbf{0}_m, \quad \text{and} \quad P = P_A = P_{[\mathbf{a}]}.$$

I think the picture makes it clear that the closest vector in $\mathbf{C}([\mathbf{a}]) = \text{span}(\mathbf{a})$ to \mathbf{b} is $P\mathbf{b}$.

Let’s prove it. Going forward, we will that $A \in \mathbb{R}^{m \times n}$ has full column rank: $\text{rank}(A) = n$. This is what makes the projection P_A exist. We can think about the case when A doesn’t have full column rank later. That is slightly harder, and the following will be hard enough to start.

Take any $\mathbf{v} \in \mathbf{C}(A)$. We’re going to prove an *inequality*:

$$\|\mathbf{b} - P_A\mathbf{b}\| \leq \|\mathbf{b} - \mathbf{v}\|. \quad (30.1)$$

This says that you’ll never make the distance between \mathbf{b} and a vector $\mathbf{v} \in \mathbf{C}(A)$ smaller than when you take $\mathbf{v} = P_A\mathbf{b}$.

This inequality is equivalent to

$$\|\mathbf{b} - \mathbf{v}\|^2 \geq \|\mathbf{b} - P_A\mathbf{b}\|^2, \quad (30.2)$$

and so that’s what we’ll prove. We can make $P_A\mathbf{b}$ show up on the left by adding and subtracting:

$$\|\mathbf{b} - \mathbf{v}\|^2 = \|\mathbf{b} - P_A\mathbf{b} + P_A\mathbf{b} - \mathbf{v}\|^2. \quad (30.3)$$

Now group things and pay attention:

$$\|\mathbf{b} - P_A\mathbf{b} + P_A\mathbf{b} - \mathbf{v}\|^2 = \|(\mathbf{b} - P_A\mathbf{b}) + (P_A\mathbf{b} - \mathbf{v})\|^2. \quad (30.4)$$

We know $\mathbf{b} - P_A\mathbf{b} \in \mathbf{N}(A^\top)$, $P_A\mathbf{b} \in \mathbf{C}(A)$, and $\mathbf{v} \in \mathbf{C}(A)$. So, $P_A\mathbf{b} - \mathbf{v} \in \mathbf{C}(A) = \mathbf{N}(A^\top)^\perp$. That is, $\mathbf{b} - P_A\mathbf{b}$ and $P_A\mathbf{b} - \mathbf{v}$ are orthogonal. Problem 30.8 implies

$$\|(\mathbf{b} - P_A\mathbf{b}) + (P_A\mathbf{b} - \mathbf{v})\|^2 = \|\mathbf{b} - P_A\mathbf{b}\|^2 + \|P_A\mathbf{b} - \mathbf{v}\|^2 \geq \|\mathbf{b} - P_A\mathbf{b}\|^2. \quad (30.5)$$

Combine (30.3), (30.4), and (30.5) to get (30.2), and we’re done.

30.10 Theorem (Least squares). Let $A \in \mathbb{R}^{m \times n}$ have full column rank: $\text{rank}(A) = n$ and $m \geq n$. Then for any $\mathbf{b} \in \mathbb{R}^m$, the closest vector in $\mathbf{C}(A)$ to \mathbf{b} is $P_A \mathbf{b}$:

$$\|\mathbf{b} - P_A \mathbf{b}\| \leq \|\mathbf{b} - \mathbf{v}\| \quad (30.6)$$

for any $\mathbf{v} \in \mathbf{C}(A)$. Moreover, with $\hat{\mathbf{x}} := (A^\top A)^{-1} A^\top \mathbf{b}$, we have

$$\|A\hat{\mathbf{x}} - \mathbf{b}\| \leq \|A\mathbf{x} - \mathbf{b}\| \quad (30.7)$$

for any $\mathbf{x} \in \mathbb{R}^n$, and so the **LEAST SQUARES SOLUTION** $\hat{\mathbf{x}}$ is the best “approximate solution” to the (possibly unsolvable) problem $A\mathbf{x} = \mathbf{b}$.

Proof. The first inequality (30.6) is just a restatement of our goal (30.1). And the second inequality (30.7) is just (30.6) with

$$P_A \mathbf{b} = A\hat{\mathbf{x}}, \quad \hat{\mathbf{x}} := (A^\top A)^{-1} A^\top \mathbf{b},$$

and $\mathbf{v} \in \mathbf{C}(A)$ replaced by $\mathbf{v} = A\mathbf{x}$ for any $\mathbf{x} \in \mathbb{R}^n$. ■

The moral is that if you want to solve $A\mathbf{x} = \mathbf{b}$ but can't, since $\mathbf{b} \notin \mathbf{C}(A)$, and if A has full column rank, content yourself with solving $A\hat{\mathbf{x}} = P_A \mathbf{b}$. We use the phrase “least squares solution” because the sum of the squares in $\|A\hat{\mathbf{x}} - \mathbf{b}\|$ is the smallest of all sums of squares of the form $\|A\mathbf{x} - \mathbf{b}\|$.

30.11 Problem (!). If $A \in \mathbb{R}^{m \times m}$ is invertible, what is $\hat{\mathbf{x}}$? Are you surprised?

Content from Strang's ILA 6E. Read p. 163 up to and including the box before Example 1. Then read “Minimizing the Error” on pp. 164–165. Skip the “By calculus” section on pp. 165–166 if you haven't taken multivariable calculus. Then read “The Big Picture for Least Squares” on pp. 166–167. Spend some time contrasting Figure 4.7 on p. 166 with Figure 4.1 back on p. 146. How is \mathbf{b} behaving differently between the two figures?

Day 31: Wednesday, March 26.

The crux of least squares is that when we can't solve $A\mathbf{x} = \mathbf{b}$, we first find the best approximation to $\mathbf{b} \in \mathbf{C}(A)$, which we call \mathbf{p} , and then we solve $A\hat{\mathbf{x}} = \mathbf{p}$. So far, success requires A to have full column rank.

31.1 Example. Let

$$A = \begin{bmatrix} 1 & 0 \\ 2 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \mathbf{b} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}.$$

If $\mathbf{y} \in \mathbf{C}(A)$, then $y_2 = 2y_1$; here $\mathbf{b} \notin \mathbf{C}(A)$. We could use the formula from Theorem 30.10 to find the least squares solution $\hat{\mathbf{x}} \in \mathbb{R}^2$ that makes $\|A\hat{\mathbf{x}} - \mathbf{b}\|$ as small as possible, but here I think it might be enlightening to see how the structure of $P_A \mathbf{b}$ allows us to solve

$A\hat{\mathbf{x}} = P_A\mathbf{b}$ directly.

From Example 29.4, the orthogonal projection onto $\mathbf{C}(A)$ is

$$P_A = \begin{bmatrix} 1/5 & 2/5 & 0 \\ 2/5 & 4/5 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

We compute

$$P_A\mathbf{b} = \begin{bmatrix} 1/5 & 2/5 & 0 \\ 2/5 & 4/5 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 1/5 \\ 2/5 \\ 0 \end{bmatrix}.$$

Then the problem $A\hat{\mathbf{x}} = P_A\mathbf{b}$ becomes

$$\begin{cases} \hat{x}_1 &= 1/5 \\ 2\hat{x}_1 &= 2/5 \\ \hat{x}_2 &= 0, \end{cases}$$

which gives $\hat{x}_1 = 1/5$ and $\hat{x}_2 = 0$, so the least squares solution is

$$\hat{\mathbf{x}} = \begin{bmatrix} 1/5 \\ 0 \end{bmatrix}.$$

Let's overthink this a little. Saying that this is the least squares solution means $\|A\hat{\mathbf{x}} - \mathbf{b}\| \leq \|A\mathbf{x} - \mathbf{b}\|$ for all $\mathbf{x} \in \mathbb{R}^2$. We're never going to make $\|A\mathbf{x} - \mathbf{b}\|$ smaller than when we choose $\mathbf{x} = \hat{\mathbf{x}}$. Let's compute it—I'll square to get rid of the square root (always a good idea):

$$\|A\mathbf{x} - \mathbf{b}\|^2 = \left\| \begin{bmatrix} x_1 \\ 2x_1 \\ x_2 \end{bmatrix} - \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right\|^2 = \left\| \begin{bmatrix} x_1 - 1 \\ 2x_1 \\ x_2 \end{bmatrix} \right\|^2 = (x_1 - 1)^2 + (2x_1)^2 + x_2^2 \geq (x_1 - 1)^2 + 4x_1^2.$$

That last inequality holds because $x_2^2 \geq 0$. What this says is that $\|A\mathbf{x} - \mathbf{b}\|^2$ is always at least as large as $(x_1 - 1)^2 + 4x_1^2$. And what is that? A function of x_1 alone! A little calculus, or graphing the parabola, will convince you that the minimum of $f(x_1) = (x_1 - 1)^2 + 4x_1^2$ occurs at $x_1 = 1/5$. Thus


$$\|A\mathbf{x} - \mathbf{b}\|^2 \geq (1/5 - 1)^2 + 4(1/5)^2 + 0^2 = \|A\hat{\mathbf{x}} - \mathbf{b}\|^2,$$

exactly as least squares predicts.

31.2 Example. This is probably the most legitimate “application” of linear algebra that we'll ever do in this class. You and I know that you can find a line running between any two points in the plane. Three or more points, maybe, maybe not.

Say that you have m sample points of data: $(x_1, y_1), \dots, (x_m, y_m)$. You probably can't find a line that passes through all of them, but can you find the line that is “closest” to all

of them? What does “closest” even mean here? Here’s a picture with $m = 4$.



I claim “closest” should mean the line whose “perpendicular distance” from each point is the smallest. This line has the form $y = mx + b$. Ideally, we’d have $y_k = mx_k + b$ for $k = 1, \dots, 4$. This really becomes the system of equations

$$\begin{cases} mx_1 + b = y_1 \\ mx_2 + b = y_2 \\ mx_3 + b = y_3 \\ mx_4 + b = y_4, \end{cases}$$

and that is the matrix-vector equation

$$\begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ x_3 & 1 \\ x_4 & 1 \end{bmatrix} \begin{bmatrix} m \\ b \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix}.$$

Remember that in the notation of this problem, x_k and y_k are given, while m and b are unknown.

Let

$$A = \begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ x_3 & 1 \\ x_4 & 1 \end{bmatrix}.$$

We want to do least squares, so A better have independent columns. For that, the first column shouldn’t be a multiple of the second. The first column *is* a multiple of the second precisely when all of the x_k ’s are the same number. But in that case, all of the data points have the same x -coordinate, in which case they all lie on the same vertical line. Boring!

So, assume that at least one of the x_k ’s is not equal to the other. Then you can do least squares and say that the best choice of slope and y -intercept is

$$\begin{bmatrix} \widehat{m} \\ \widehat{b} \end{bmatrix} = (A^T A)^{-1} A^T \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix}.$$

You could compute what \hat{m} and \hat{b} are explicitly, or you could go to a computer and replace thinking with typing.

The takeaway for me from this example is that an extremely natural, and important, application results in a matrix A that transparently has full column rank but not full row rank. This justifies our emphasis on A having full column rank in the least squares developments so far.

Content from Strang's ILA 6E. Least squares for data fitting to lines appears in Example 1 on pp. 163–164, Figure 4.6, and pp. 167–168. Pay careful attention to the utility of orthogonal columns in A in Example 2 on p. 168. There's no reason to stop with lines. What if you wanted to find the “best” parabola approximating a set of data? Add one more column to A to account for the extra coefficient in the parabola and read p. 170.

Day 32: Friday, March 28.

You took Exam 2.

Day 33: Monday, March 31.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Orthogonal list of vectors, orthonormal list of vectors, orthogonal matrix

Nonetheless, a careful review of the work leading to Theorem 30.10 will convince you that we didn't need A to have full column rank to find a best approximation to \mathbf{b} .

33.1 Problem (★). Let \mathcal{V} be a subspace of \mathbb{R}^m and let $\mathbf{b} \in \mathbb{R}^m$. Find $\mathbf{v}_\star \in \mathcal{V}$ such that

$$\|\mathbf{b} - \mathbf{v}_\star\| \leq \|\mathbf{b} - \mathbf{v}\|$$

for any $\mathbf{v} \in \mathcal{V}$. [Hint: if $\mathcal{V} = \{\mathbf{0}_m\}$, there isn't much to do. Otherwise, start by writing $\mathcal{V} = \mathbf{C}(A)$ for some $A \in \mathbb{R}^{m \times d}$ with $\text{rank}(A) = d$.]

This vector \mathbf{v}_\star is a best approximation to \mathbf{b} . If we're going to say “best,” we probably want only one “best,” and this is one of those times when we get just that.

33.2 Problem (!). Use the definition of the norm $\|\cdot\|$ to prove the **PARALLELOGRAM LAW**:

$$\|\mathbf{v} + \mathbf{w}\|^2 + \|\mathbf{v} - \mathbf{w}\|^2 = 2\|\mathbf{v}\|^2 + 2\|\mathbf{w}\|^2$$

for any $\mathbf{v}, \mathbf{w} \in \mathbb{R}^m$.

33.3 Problem (+). Let \mathcal{V} be a subspace of \mathbb{R}^m and $\mathbf{b} \in \mathbb{R}^m$. Problem 33.1 tells us that there exists a closest point $\mathbf{v}_\star \in \mathcal{V}$ to \mathbf{b} satisfying

$$\|\mathbf{b} - \mathbf{v}_\star\| \leq \|\mathbf{b} - \mathbf{v}\| \quad (33.1)$$

for all $\mathbf{v} \in \mathcal{V}$. Now we show that there is only one such closest point.

(i) Draw a picture in \mathbb{R}^2 that convinces you that there should only be one such closest point. [Hint: *copy the pictures in \mathbb{R}^2 that we've been drawing all along.*]

(ii) Prove it! Suppose that $\tilde{\mathbf{v}}_\star \in \mathcal{V}$ also satisfies

$$\|\mathbf{b} - \tilde{\mathbf{v}}_\star\| \leq \|\mathbf{b} - \mathbf{v}\| \quad (33.2)$$

for all $\mathbf{v} \in \mathcal{V}$. First combine (33.1) and (33.2) to show

$$\|\mathbf{b} - \mathbf{v}_\star\| = \|\mathbf{b} - \tilde{\mathbf{v}}_\star\|.$$

(iii) Explain why you want to show $\|\mathbf{v}_\star - \tilde{\mathbf{v}}_\star\| = 0$. Add and subtract \mathbf{b} and use the parallelogram law (Problem 33.2) to show

$$\|\mathbf{v}_\star - \tilde{\mathbf{v}}_\star\|^2 = 2\|\mathbf{b} - \mathbf{v}_\star\|^2 + 2\|\mathbf{b} - \tilde{\mathbf{v}}_\star\|^2 - 4\left\|\mathbf{b} - \frac{\mathbf{v}_\star + \tilde{\mathbf{v}}_\star}{2}\right\|^2. \quad (33.3)$$

(iv) Abbreviate

$$\epsilon := \|\mathbf{b} - \mathbf{v}_\star\| = \|\mathbf{b} - \tilde{\mathbf{v}}_\star\|.$$

Use the fact that $\mathbf{v}_\star, \tilde{\mathbf{v}}_\star \in \mathcal{V}$ to obtain $(\mathbf{v}_\star + \tilde{\mathbf{v}}_\star)/2 \in \mathcal{V}$ and thus

$$\epsilon \leq \left\|\mathbf{b} - \frac{\mathbf{v}_\star + \tilde{\mathbf{v}}_\star}{2}\right\|. \quad (33.4)$$

(v) Combine (33.3) and (33.4) to conclude

$$\|\mathbf{v}_\star - \tilde{\mathbf{v}}_\star\|^2 \leq 2\epsilon^2 + 2\epsilon^2 - 4\epsilon^2 = 0.$$

So, if A doesn't have full column rank, we could still find the closest point $\mathbf{p} \in \mathbf{C}(A)$ to \mathbf{b} and then try to solve $A\hat{\mathbf{x}} = \mathbf{p}$. We will definitely succeed in solving this because $\mathbf{p} \in \mathbf{C}(A)$! The challenge is that because A doesn't have full column rank, we will succeed with too many degrees of freedom: $\mathbf{N}(A) \neq \{\mathbf{0}_n\}$, and so we have many choices for $\hat{\mathbf{x}}$. Which is best?

33.4 Remark. Here is one way of proceeding, motivated by the notion that less complicated data is probably better than complicated data.

Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$, and, by Problems 33.1 and 33.3, let $\mathbf{p} \in \mathbf{C}(A)$ be the closest

point in $\mathbf{C}(A)$ to \mathbf{b} . Let $\hat{\mathbf{x}} \in \mathbb{R}^n$ satisfy $A\hat{\mathbf{x}} = \mathbf{p}$. By Theorem 27.9, write $\hat{\mathbf{x}} = \hat{\mathbf{v}} + \hat{\mathbf{w}}$, where $\hat{\mathbf{v}} \in \mathbf{C}(A^\top)$ and $\hat{\mathbf{w}} \in \mathbf{N}(A)$. Then $A\hat{\mathbf{v}} = \mathbf{p}$, and so by Theorem 21.13 any other solution $\hat{\mathbf{y}}$ to $A\hat{\mathbf{y}} = \mathbf{p}$ also has the form $\hat{\mathbf{y}} = \hat{\mathbf{v}} + \hat{\mathbf{z}}$ for some $\hat{\mathbf{z}} \in \mathbf{N}(A)$.

By Problem 30.8, $\|\hat{\mathbf{v}}\| \leq \|\hat{\mathbf{y}}\|$. That is, $\hat{\mathbf{v}}$ has the smallest norm of any solution $\hat{\mathbf{y}}$ to $A\hat{\mathbf{y}} = \mathbf{p}$. We might call $\hat{\mathbf{v}}$ the **MINIMUM-NORM LEAST SQUARES SOLUTION**.

But how do we find $\hat{\mathbf{v}}$? This requires first finding that $\hat{\mathbf{x}}$ that solves $A\hat{\mathbf{x}} = \mathbf{p}$, which requires knowing \mathbf{p} ; this requires an orthogonal projection onto $\mathbf{C}(A)$, which I guess we could get from writing $\mathbf{C}(A) = \mathbf{C}(B_1)$ for some $B_1 \in \mathbb{R}^{m \times r}$ with full column rank. Then to get $\hat{\mathbf{v}}$ from $\hat{\mathbf{x}}$, we'd need the orthogonal projection onto $\mathbf{C}(A^\top)$, for which we'd start by writing $\mathbf{C}(A^\top) = \mathbf{C}(B_2)$, where $B_2 \in \mathbb{R}^{n \times r}$. (I don't think $B_2 = B_1^\top$, but am I wrong?)

This seems like a lot of work. It would be nice if there were a simpler formula for $\hat{\mathbf{v}}$ in terms of A and \mathbf{b} , and experience teaches us that such a formula probably involves multiplying \mathbf{b} by a special matrix. This turns out to be true: there is a matrix $A^+ \in \mathbb{R}^{n \times m}$ such that $\hat{\mathbf{v}} = A^+\mathbf{b}$, and this A^+ is the **PSEUDOINVERSE** of A .

Content from Strang's ILA 6E. Page 169 gives a concrete example of what to do when A doesn't have full column rank. The construction of the pseudoinverse is best resolved via the glorious tool of the singular value decomposition. Read the comment at the bottom of p. 169 for a nice review of the three possibilities for solutions to linear systems.


Optionally (this is wholly, totally optional), read Section 4.5, which details what the pseudoinverse does. You can skip the example from the "Incidence Matrix of a Graph" on p. 194. Ideally we will develop the SVD, so the formula for A^+ on p. 195 will eventually make sense.

33.5 Problem (★). Here is the opposite question: what is the best solution when we have too many solutions? Suppose that $A \in \mathbb{R}^{m \times n}$ has full row rank, so we can always solve $A\mathbf{x} = \mathbf{b}$. However, perhaps A is not square, in which case A won't have full column rank as well, and so solutions won't be unique. This sort of arose above in Remark 33.4, and the idea was to choose the "minimum norm solution."


First reread that remark carefully; since $\mathbf{b} \in \mathbf{C}(A)$ here, we can assume $\mathbf{p} = \mathbf{b}$ throughout, and we may as well dispense with the hats since there is actually a solution to $A\mathbf{x} = \mathbf{b}$ now. Use Problem 29.2 to get the orthogonal projection P_{A^\top} onto $\mathbf{C}(A^\top)$. Write $\mathbf{v} = P_{A^\top}\mathbf{x}$ and conclude $\mathbf{v} = A^\top(AA^\top)^{-1}\mathbf{b}$. This is our formula for \mathbf{v} in terms of A and \mathbf{b} .

Here is a summary, in pictures, of everything we've done. Literally: these three pictures encapsulate most of the ideas of the course. In these pictures, which are really fake cartoons, I'm imagining that all of the four fundamental subspaces are one-dimensional (except in the second, where the null space of A is trivial), and so we can imagine them as coordinate axes in a two-dimensional plane.


Best case is that we can solve $A\mathbf{x} = \mathbf{b}$, although maybe $\mathbf{N}(A) \neq \{\mathbf{0}_n\}$ and we have infinitely many solutions.



Next best case is that while we can't solve $A\mathbf{x} = \mathbf{b}$, since $\mathbf{b} \notin \mathbf{C}(A)$, A does have full column rank, so we can do a least squares approximation.



Worst case is that we can't solve $A\mathbf{x} = \mathbf{b}$ and A doesn't have full column rank. Then while we can approximate $A\mathbf{x} = \mathbf{b}$ with the problem $A\hat{\mathbf{x}} = \mathbf{p}$, where \mathbf{p} is the projection of \mathbf{b} onto $\mathbf{C}(A)$, this new approximate problem won't have a unique solution, since $\mathbf{N}(A) \neq \{\mathbf{0}_n\}$. This is why you may want to go learn about the pseudoinverse on your own.



We're going to switch our focus a little and go back to some old stuff and ask how can we do better. Actually calculating the orthogonal projection onto a column space can be annoying, because we have to invert $A^T A$. This could make solving a least squares problem hard; we know that if $A \in \mathbb{R}^{m \times n}$ has full column rank and $\mathbf{b} \notin \mathbf{C}(A)$, the best thing to do is solve $A\hat{\mathbf{x}} = P_A \mathbf{b}$ with $P_A = A(A^T A)^{-1} A^T$. We do have a solution formula: $\hat{\mathbf{x}} = (A^T A)^{-1} A^T \mathbf{b}$. But we know that numerically computing inverses is rarely a good idea. We could think instead about $(A^T A)\hat{\mathbf{x}} = A^T \mathbf{b}$. This is the **NORMAL EQUATION** for our least squares problem. Maybe an *LU*-factorization or Gaussian elimination would be more efficient than doing the inverse.

Content from Strang's ILA 6E. Reread the first three paragraphs on p. 163.

It turns out that if we ask a little more of A , the projection P_A becomes much nicer. Equivalently, if we have a really good basis for $\mathbf{C}(A)$ and stack that basis in a matrix Q , then since $\mathbf{C}(A) = \mathbf{C}(Q)$, we just need to compute P_Q instead of P_A . (This is the uniqueness of the orthogonal projection onto $\mathbf{C}(A) = \mathbf{C}(Q)$ from Problem 30.3.)

The right thing to do is exploit geometry further. Long ago (in Problem 4.6) we saw why the standard basis vectors in \mathbb{R}^m were so nice. Since

$$\mathbf{e}_j \cdot \mathbf{e}_k = \begin{cases} 1, & j = k \\ 0, & j \neq k \end{cases}$$

and $\mathbb{R}^m = \text{span}(\mathbf{e}_1, \dots, \mathbf{e}_m) = \mathbf{C}(I_m)$, we have the representation

$$\mathbf{v} = (\mathbf{v} \cdot \mathbf{e}_1)\mathbf{e}_1 + \dots + (\mathbf{v} \cdot \mathbf{e}_m)\mathbf{e}_m$$

for any $\mathbf{v} \in \mathbb{R}^m$. The actual formulas for the standard basis vectors aren't what's special here; rather, it's how they interact under the dot product. What's most important is their mutual orthogonality.

33.6 Definition. A list $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathbb{R}^m$ is **ORTHOGONAL** if

$$\mathbf{u}_j \cdot \mathbf{u}_k = 0$$

for $j \neq k$.

To keep things simple, let's look at an orthogonal list $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3 \in \mathbb{R}^m$ and say $\mathbf{v} \in \text{span}(\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3)$. Then $\mathbf{v} = c_1\mathbf{u}_1 + c_2\mathbf{u}_2 + c_3\mathbf{u}_3$ for some $c_1, c_2, c_3 \in \mathbb{R}$. Here's the trick:

$$\begin{aligned} \mathbf{v} \cdot \mathbf{u}_1 &= (c_1\mathbf{u}_1 + c_2\mathbf{u}_2 + c_3\mathbf{u}_3) \cdot \mathbf{u}_1 \\ &= ((c_1\mathbf{u}_1) \cdot \mathbf{u}_1) + ((c_2\mathbf{u}_2) \cdot \mathbf{u}_1) + ((c_3\mathbf{u}_3) \cdot \mathbf{u}_1) \\ &= c_1(\mathbf{u}_1 \cdot \mathbf{u}_1) + c_2(\mathbf{u}_2 \cdot \mathbf{u}_1) + c_3(\mathbf{u}_3 \cdot \mathbf{u}_1) \\ &= c_1(\mathbf{u}_1 \cdot \mathbf{u}_1). \end{aligned}$$

If $\mathbf{u}_1 \neq \mathbf{0}_m$, then $\mathbf{u}_1 \cdot \mathbf{u}_1 = \|\mathbf{u}_1\|^2 \neq 0$, and so we have

$$c_1 = \frac{\mathbf{v} \cdot \mathbf{u}_1}{\|\mathbf{u}_1\|^2}.$$

Let's assume that none of the \mathbf{u}_j are $\mathbf{0}_m$; otherwise, they contribute nothing worthwhile to the span. Taking dot products of \mathbf{v} against the other \mathbf{u}_j then yields

$$c_j = \frac{\mathbf{v} \cdot \mathbf{u}_j}{\|\mathbf{u}_j\|^2}.$$

This generalizes to an arbitrary orthogonal list.

33.7 Theorem. Let $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathbb{R}^m$ be orthogonal and let $\mathbf{v} \in \text{span}(\mathbf{u}_1, \dots, \mathbf{u}_n)$. Then

$$\mathbf{v} = \left(\frac{\mathbf{v} \cdot \mathbf{u}_1}{\|\mathbf{u}_1\|^2} \right) \mathbf{u}_1 + \dots + \left(\frac{\mathbf{v} \cdot \mathbf{u}_n}{\|\mathbf{u}_n\|^2} \right) \mathbf{u}_n.$$

A nice consequence is that any orthogonal list of *nonzero* vectors is independent. For if $c_1\mathbf{u}_1 + \dots + c_n\mathbf{u}_n = \mathbf{0}_m$, then each c_j must be 0. (In the theorem above, take $\mathbf{v} = \mathbf{0}_m$, so the dot products collapse to 0.)

33.8 Problem (!). What is the maximum length of any list of orthogonal vectors in \mathbb{R}^m ?

All that division, however, gets annoying. It's much more efficient to assume $\|\mathbf{u}_j\| = 1$ for all j .

33.9 Definition. A list $\mathbf{q}_1, \dots, \mathbf{q}_n \in \mathbb{R}^m$ is **ORTHONORMAL** if

$$\mathbf{q}_j \cdot \mathbf{q}_k = \begin{cases} 1, & j = k \\ 0, & j \neq k. \end{cases}$$

The $j \neq k$ condition means that an orthonormal list is orthogonal, while the $j = k$ condition gives $\|\mathbf{q}_j\| = \sqrt{\mathbf{q}_j \cdot \mathbf{q}_j} = 1$.

33.10 Problem (★). Let $\mathbf{q}_1, \mathbf{q}_2 \in \mathbb{R}^m$ be orthonormal and let $\mathbf{v} \in \text{span}(\mathbf{q}_1, \mathbf{q}_2)$. Use the Pythagorean theorem (Theorem 30.6) to show that

$$\|\mathbf{v}\|^2 = |\mathbf{v} \cdot \mathbf{q}_1|^2 + |\mathbf{v} \cdot \mathbf{q}_2|^2.$$

This generalizes: if $\mathbf{q}_1, \dots, \mathbf{q}_n \in \mathbb{R}^m$ are orthonormal and $\mathbf{v} \in \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_n)$, then

$$\|\mathbf{v}\|^2 = |\mathbf{v} \cdot \mathbf{q}_1|^2 + \dots + |\mathbf{v} \cdot \mathbf{q}_n|^2.$$

We work with matrices and column spaces as much as we do with lists of vectors and spans, so let's stuff those orthonormal vectors into a matrix and get an unfortunate definition.

33.11 Definition. A matrix $Q \in \mathbb{R}^{m \times n}$ is **ORTHOGONAL** if the columns of Q are orthonormal.

I'm sorry that we don't just say "orthonormal matrix." Math just isn't hard enough.

33.12 Example. (i) The identity matrix is always orthogonal.

(ii) Let $\theta \in \mathbb{R}$ and

$$Q = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}.$$

Then trig and the Pythagorean identity $\sin^2(\theta) + \cos^2(\theta) = 1$ say that Q is orthogonal.

Content from Strang's ILA 6E. Everything on pp. 176–178 is important. I am being a little more general and calling any matrix, square or not, with orthogonal columns an "orthogonal matrix."

Day 34: Wednesday, April 2.

Here is a nice consequence of definitions. Let $Q = [\mathbf{q}_1 \ \dots \ \mathbf{q}_n] \in \mathbb{R}^{m \times n}$ be orthogonal. Then

$$\mathbf{q}_i \cdot \mathbf{q}_j = \begin{cases} 1, & i = j \\ 0, & i \neq j. \end{cases}$$

Remember that row i of Q^T is just column i of Q , and that the (i, j) -entry of $Q^T Q$ is the dot product of row i of Q^T and column j of Q . That is, the (i, j) -entry of $Q^T Q$ is $\mathbf{q}_i \cdot \mathbf{q}_j$, and so this (i, j) -entry is 1 when $i = j$ (= on the diagonal) and 0 otherwise (= off the diagonal). This sounds a lot like an identity matrix, and it is! Since $Q \in \mathbb{R}^{m \times n}$, we have $Q^T \in \mathbb{R}^{n \times m}$, and so $Q^T Q \in \mathbb{R}^{n \times n}$.

34.1 Theorem. Let $Q \in \mathbb{R}^{m \times n}$ be orthogonal. Then $Q^T Q = I_n$.

34.2 Problem (!). So is every orthogonal matrix invertible?

34.3 Problem (★). State and prove an analogue of Theorem 34.1 for the case when the columns of Q are only orthogonal, not orthonormal.

34.4 Problem (!). Is every orthogonal projection (Definition 29.6) an orthogonal matrix?

We can use this property of orthogonal matrices to recover our slick representation from an orthonormal basis. Let $Q \in \mathbb{R}^{m \times n}$ be orthogonal and $\mathbf{b} \in \mathbf{C}(Q)$. Then $\mathbf{b} = Q\mathbf{x}$ for some $\mathbf{x} \in \mathbb{R}^n$, and so $Q^T \mathbf{b} = Q^T Q \mathbf{x} = \mathbf{x}$. Thus

$$\mathbf{b} = Q\mathbf{x} = QQ^T \mathbf{b}.$$

Now let's think about multiplication. One way to compute the entries of $Q^T \mathbf{b}$ is to take the dot product of the rows of Q^T with \mathbf{b} . (This is probably how we usually compute matrix-vector products by hand.) And the rows of Q^T are the columns of Q , so

$$Q^T \mathbf{b} = \begin{bmatrix} \mathbf{q}_1 \cdot \mathbf{b} \\ \vdots \\ \mathbf{q}_n \cdot \mathbf{b} \end{bmatrix}.$$

Next, one way to compute $QQ^T \mathbf{b}$ is to take the linear combination of the columns of Q weighted by the entries of $Q^T \mathbf{b}$:

$$\mathbf{b} = QQ^T \mathbf{b} = [\mathbf{q}_1 \ \cdots \ \mathbf{q}_n] \begin{bmatrix} \mathbf{q}_1 \cdot \mathbf{b} \\ \vdots \\ \mathbf{q}_n \cdot \mathbf{b} \end{bmatrix} = (\mathbf{q}_1 \cdot \mathbf{b}) + \cdots + (\mathbf{q}_n \cdot \mathbf{b}).$$

I love how this brings together two ways of looking at matrix-vector multiplication: the dot product way for quick and dirty calculations by hand, the linear combination of columns way to actually understand what's happening. I'd be scared if this doesn't make you happy, too.

Enough old stuff. Here's the new: orthonormality and orthogonal matrices make least squares so much easier. Suppose that $Q \in \mathbb{R}^{m \times n}$ is orthogonal and we want to solve $Q\mathbf{x} = \mathbf{b}$, but $\mathbf{b} \notin \mathbf{C}(Q)$. Then we'd solve the least squares problem

$$Q\hat{\mathbf{x}} = P_Q \mathbf{b},$$

where

$$P_Q = Q(Q^T Q)^{-1} Q^T = Q I_n^{-1} Q^T = QQ^T.$$

Look at that: the orthogonal projection onto $\mathbf{C}(Q)$ collapses to QQ^T . No inverses needed. This is so nice that I want to emphasize it by itself.

34.5 Theorem. Let $Q = [\mathbf{q}_1 \ \cdots \ \mathbf{q}_n] \in \mathbb{R}^{m \times n}$ be orthogonal. Then the orthogonal projection onto $\mathbf{C}(Q)$ is QQ^T , and every $\mathbf{v} \in \mathbf{C}(Q)$ has the form

$$\mathbf{v} = QQ^T \mathbf{v} = (\mathbf{v} \cdot \mathbf{q}_1) + \cdots + (\mathbf{v} \cdot \mathbf{q}_n) \mathbf{q}_n.$$

Content from Strang's ILA 6E. You should hold the answer to Worked Example 4.4 B on p. 185 deep within your heart.

Then the least squares problem is just

$$Q\hat{\mathbf{x}} = QQ^T \mathbf{b},$$

and that says

$$\hat{\mathbf{x}} = Q^T \mathbf{b}.$$

Again, no inverses, just transposing and multiplying.

34.6 Problem (!). Reread Example 29.4 and explain how orthonormality made calculating the projection operator easier. How would things have been more complicated there if we used the pivot columns of A as the basis for the column space, not the columns of B ?

Content from Strang's ILA 6E. Page 179 through the top of p. 180 discuss least squares with orthogonal matrices.

It looks like we really win if the important vectors in our problem are orthonormal. But often they aren't. If we start with a matrix A , how can we get an “orthonormal basis” for its column space? If $A \in \mathbb{R}^{m \times n}$, is there an orthogonal matrix $Q \in \mathbb{R}^{m \times n}$ with $\mathbf{C}(A) = \mathbf{C}(Q)$?

Yes, and we can construct it explicitly. We'll go through the procedure in detail for $n = 3$. Let $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \in \mathbb{R}^m$ be independent. (Why independent? We want to transform any old basis into the best kind of basis—the orthonormal basis—and so we may as well start with independent vectors.) We want to find orthonormal $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3 \in \mathbb{R}^m$ that “preserve the span”:

$$\begin{cases} \text{span}(\mathbf{v}_1) = \text{span}(\mathbf{q}_1) \\ \text{span}(\mathbf{v}_1, \mathbf{v}_2) = \text{span}(\mathbf{q}_1, \mathbf{q}_2) \\ \text{span}(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3) = \text{span}(\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3). \end{cases} \quad (34.1)$$

The first equality suggests that we just define \mathbf{q}_1 directly in terms of \mathbf{v}_1 . Since we want $\|\mathbf{q}_j\| = 1$ for each j , and since we know $\mathbf{v}_1 \neq \mathbf{0}_m$ (by independence), we put

$$\mathbf{q}_1 = \frac{\mathbf{v}_1}{\|\mathbf{v}_1\|}.$$

This immediately gives $\|\mathbf{q}_1\| = 1$, and since \mathbf{q}_1 is just a (nonzero) multiple of \mathbf{v}_1 , we preserve the span: $\text{span}(\mathbf{v}_1) = \text{span}(\mathbf{q}_1)$.

Next, we want $\mathbf{q}_2 \in \mathbb{R}^m$ such that $\mathbf{q}_2 \cdot \mathbf{q}_1 = 0$ and $\|\mathbf{q}_2\| = 1$. If we have a vector $\mathbf{u}_2 \in \mathbb{R}^m$ with $\mathbf{u}_2 \cdot \mathbf{q}_1 = 0$ and $\mathbf{u}_2 \neq \mathbf{0}_m$, then we can always normalize and put $\mathbf{q}_2 = \mathbf{u}_2 / \|\mathbf{u}_2\|$. Then

we'll still have $\mathbf{q}_2 \cdot \mathbf{q}_1 = 0$. Also, if we know $\text{span}(\mathbf{v}_1, \mathbf{v}_2) = \text{span}(\mathbf{q}_1, \mathbf{u}_2)$, then $\text{span}(\mathbf{v}_1, \mathbf{v}_2) = \text{span}(\mathbf{q}_1, \mathbf{q}_2)$.

34.7 Problem (!). Check that: show $\mathbf{C} \left(\begin{bmatrix} A & \mathbf{v} \end{bmatrix} \right) = \mathbf{C} \left(\begin{bmatrix} A & c\mathbf{v} \end{bmatrix} \right)$ for any $A \in \mathbb{R}^{m \times n}$, $\mathbf{v} \in \mathbb{R}^m$, and $c \in \mathbb{R}$. (I'm sure you did this at some point in the course already, but do it again.)

So, we focus on getting $\mathbf{u}_2 \cdot \mathbf{q}_1 = 0$, $\text{span}(\mathbf{v}_1, \mathbf{v}_2) = \text{span}(\mathbf{q}_1, \mathbf{u}_2)$, and $\mathbf{u}_2 \neq \mathbf{0}_m$. One of the big ways to make orthogonality show up is to project. Problem 28.12 tells us that \mathbf{q}_1 and $\mathbf{v} - (\mathbf{v} \cdot \mathbf{q}_1)\mathbf{q}_1$ are orthogonal for any $\mathbf{v} \in \mathbb{R}^m$. (Or you could just check that right now yourself.) Since we want to get \mathbf{v}_2 to show up, it's worth trying

$$\mathbf{u}_2 = \mathbf{v}_2 - (\mathbf{v}_2 \cdot \mathbf{q}_1)\mathbf{q}_1.$$

This definitely gives $\mathbf{u}_2 \cdot \mathbf{q}_1 = 0$.

Now we check spans. Anything in $\text{span}(\mathbf{v}_1, \mathbf{v}_2)$ is

$$\begin{aligned} x_1\mathbf{v}_1 + x_2\mathbf{v}_2 &= c_1\mathbf{q}_1 + x_2\mathbf{v}_2 \text{ for some } c_1 \in \mathbb{R}, \text{ since } \text{span}(\mathbf{v}_1) = \text{span}(\mathbf{q}_1) \\ &= c_1\mathbf{q}_1 + x_2(\mathbf{v}_2 - (\mathbf{v}_2 \cdot \mathbf{q}_1)\mathbf{q}_1) + x_2(\mathbf{v}_2 \cdot \mathbf{q}_1)\mathbf{q}_1 \\ &= (c_1 + x_2(\mathbf{v}_2 \cdot \mathbf{q}_1))\mathbf{q}_1 + x_2\mathbf{u}_2 \\ &\in \text{span}(\mathbf{q}_1, \mathbf{u}_2). \end{aligned}$$

And anything in $\text{span}(\mathbf{q}_1, \mathbf{u}_2)$ is

$$\begin{aligned} y_1\mathbf{q}_1 + y_2\mathbf{u}_2 &= y_1\mathbf{q}_1 + y_2(\mathbf{v}_2 - (\mathbf{v}_2 \cdot \mathbf{q}_1)\mathbf{q}_1) \\ &= (y_1 - (\mathbf{v}_2 \cdot \mathbf{q}_1))\mathbf{q}_1 + y_2\mathbf{v}_2 \\ &= z_1\mathbf{v}_1 + y_2\mathbf{v}_2 \text{ since } \mathbf{q}_1 \in \text{span}(\mathbf{v}_1) \\ &\in \text{span}(\mathbf{v}_1, \mathbf{v}_2). \end{aligned}$$

This proves $\text{span}(\mathbf{q}_1, \mathbf{u}_2) = \text{span}(\mathbf{v}_1, \mathbf{v}_2)$.

Finally, we check that $\mathbf{u}_2 \neq \mathbf{0}_m$. Otherwise, we'd have

$$\mathbf{0}_m = \mathbf{v}_2 - (\mathbf{v}_2 \cdot \mathbf{q}_1)\mathbf{q}_1 = \mathbf{v}_2 + c\mathbf{v}_1,$$

since $\mathbf{q}_1 \in \text{span}(\mathbf{v}_1)$. This contradicts the linear independence of \mathbf{v}_1 and \mathbf{v}_2 . We have therefore met all of our goals with \mathbf{u}_2 and so put $\mathbf{q}_2 = \mathbf{u}_2 / \|\mathbf{u}_2\|$. Onwards.

We finally want $\mathbf{q}_3 \in \mathbb{R}^m$ such that $\mathbf{q}_3 \cdot \mathbf{q}_2 = 0$, $\mathbf{q}_3 \cdot \mathbf{q}_1 = 0$, and $\text{span}(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3) = \text{span}(\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3)$. Projections get us the orthogonality once again: if $B \in \mathbb{R}^{m \times n}$ has full column rank, then $\mathbf{v} - P_B\mathbf{v}$ is orthogonal to every vector in $\mathbf{C}(B)$, since $\mathbf{v} - P_B\mathbf{v} \in \mathbf{N}(B^\top)$. (Don't believe me? Calculate it right now.) Take $A = Q_2 = [\mathbf{q}_1 \ \mathbf{q}_2]$. Then

$$\mathbf{u}_3 = \mathbf{v}_3 - P_{Q_2}\mathbf{v}_3 = \mathbf{v}_3 - Q_2Q_2^\top$$

will satisfy $\mathbf{u}_3 \cdot \mathbf{q}_2 = \mathbf{u}_3 \cdot \mathbf{q}_1 = 0$. Here we are using the fact that if Q is orthogonal, then the orthogonal projection onto its column space is $P_Q = QQ^\top$.

Now we'll check $\text{span}(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3) = \text{span}(\mathbf{q}_1, \mathbf{q}_2, \mathbf{u}_3)$ and $\mathbf{u}_3 \neq \mathbf{0}_m$. Then $\mathbf{q}_3 = \mathbf{u}_3 / \|\mathbf{u}_3\|$ will finish the job.

34.8 Problem (★). Here's a general way of getting the spans (which, by the way, also takes care of the spans from constructing \mathbf{u}_2): show that if $A, B \in \mathbb{R}^{m \times n}$ with $\mathbf{C}(A) = \mathbf{C}(B)$ with $\text{rank}(B) = n$, then

$$\mathbf{C}\left(\begin{bmatrix} A & \mathbf{v} \end{bmatrix}\right) = \mathbf{C}\left(\begin{bmatrix} B & (\mathbf{v} - P_B \mathbf{v}) \end{bmatrix}\right)$$

for any $\mathbf{v} \in \mathbb{R}^m$. [Hint: Start with $A\mathbf{x} + c\mathbf{v} = A\mathbf{x} + c(\mathbf{v} - P_B \mathbf{v}) + cP_B \mathbf{v}$. Explain why $A\mathbf{x} + cP_B \mathbf{v} \in \mathbf{C}(B)$. Next, $B\mathbf{y} + z(\mathbf{v} - P_B \mathbf{v}) = (B\mathbf{y} - zP_B \mathbf{v}) + z\mathbf{v}$. Explain why $B\mathbf{y} - zP_B \mathbf{v} \in \mathbf{C}(A)$.]

Take $A = [\mathbf{v}_1 \ \mathbf{v}_2]$, $B = [\mathbf{q}_1 \ \mathbf{q}_2]$, and $\mathbf{v} = \mathbf{v}_3$ to conclude that $\text{span}(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3) = \text{span}(\mathbf{q}_1, \mathbf{q}_2, \mathbf{u}_3)$. Now we check $\mathbf{u}_3 \neq \mathbf{0}_m$. Otherwise, we'd have

$$\mathbf{v}_3 = Q_2 Q_2^\top \mathbf{v}_3 \in \mathbf{C}(Q_2) = \text{span}(\mathbf{v}_1, \mathbf{v}_2).$$

This contradicts the independence of $\mathbf{v}_1, \mathbf{v}_2$, and \mathbf{v}_3 . So, we conclude with $\mathbf{q}_3 = \mathbf{u}_3 / \|\mathbf{u}_3\|$.

Content from Strang's ILA 6E. Pages 180–181 do Gram–Schmidt for three vectors. See in particular the 3D drawings in Figure 4.10 on p. 181.

Day 35: Friday, April 4.

Here is the general result.

35.1 Theorem (Gram–Schmidt procedure). Let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$ be independent. There exist orthonormal $\mathbf{q}_1, \dots, \mathbf{q}_n \in \mathbb{R}^m$ such that $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_j) = \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_j)$ for $j = 1, \dots, n$. Specifically,

$$\mathbf{q}_1 = \frac{\mathbf{v}_1}{\|\mathbf{v}_1\|} \quad \text{and} \quad \mathbf{q}_j = \frac{\mathbf{u}_j}{\|\mathbf{u}_j\|}, \quad j \geq 2,$$

where

$$\mathbf{u}_j := \mathbf{v}_j - Q_{j-1} Q_{j-1}^\top \mathbf{v}_j = \mathbf{v}_j - ((\mathbf{v}_j \cdot \mathbf{q}_1)\mathbf{q}_1 + \dots + (\mathbf{v}_j \cdot \mathbf{q}_{j-1})\mathbf{q}_{j-1}), \quad Q_{j-1} := [\mathbf{q}_1 \ \dots \ \mathbf{q}_{j-1}]. \quad (35.1)$$

Proof. This is really a proof by induction, but the key ideas are outlined in the $n = 3$ case above. The point is that you know how to construct \mathbf{q}_1 , and then you assume that you've constructed through \mathbf{q}_j with $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_j) = \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_j)$ and $\mathbf{q}_1, \dots, \mathbf{q}_j$ orthonormal. Put $\mathbf{u}_{j+1} = \mathbf{v}_{j+1} - Q_j Q_j^\top \mathbf{v}_{j+1}$. This immediately gives $\mathbf{u}_{j+1} \cdot \mathbf{q}_k = 0$ for $k = 1, \dots, j$, since $\mathbf{v} - Q_j Q_j^\top \mathbf{v}$ is orthogonal to anything in $\mathbf{C}(Q_j)$. For spans, use Problem 34.8 with $A = [\mathbf{v}_1 \ \dots \ \mathbf{v}_j]$, $B = Q_j$, and $\mathbf{v} = \mathbf{u}_{j+1}$. To be sure that $\mathbf{u}_{j+1} \neq \mathbf{0}_m$, suppose otherwise and get $\mathbf{v}_{j+1} = Q_j Q_j^\top \mathbf{v}_{j+1} \in \mathbf{C}(Q_j) = \mathbf{C}(A)$, which contradicts the independence of $\mathbf{v}_1, \dots, \mathbf{v}_n$. ■

Content from Strang's ILA 6E. Page 183 presents some pseudocode for computing Gram–Schmidt. See also the confession on p. 184. Read it and take a numerical linear algebra class.

My feeling is that to write things out concisely, it's nice to use the projection $Q_j Q_j^\top$. To do calculations by hand, it's easier to use the second formula for \mathbf{u}_j in (35.1).

35.2 Problem (★). Using the hypotheses and notation of the Gram–Schmidt procedure, prove that $\mathbf{v}_j \cdot \mathbf{q}_j > 0$ as follows.

(i) First explain why we just need $\mathbf{v} \cdot \mathbf{u}_j > 0$, and check that this is true in the case $j = 1$.

(ii) For $j \geq 2$, rewrite

$$\mathbf{v}_j = \mathbf{u}_j + Q_{j-1} Q_{j-1}^\top \mathbf{v}_j.$$

Explain why $Q_{j-1} Q_{j-1}^\top \mathbf{v}_j \cdot \mathbf{u}_j = 0$. [Hint: $Q_{j-1} Q_{j-1}^\top \mathbf{v}_j = Q_{j-1} (Q_{j-1}^\top \mathbf{v}_j)$ and $\mathbf{u}_j \in \mathbf{C}(Q_{j-1})^\perp$.]

(iii) Conclude that $\mathbf{v}_j \cdot \mathbf{u}_j = \|\mathbf{u}_j\|^2$.

35.3 Example. Let

$$\mathbf{v}_1 = \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}, \quad \text{and} \quad \mathbf{v}_3 = \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix}.$$

I claim that these vectors are independent, and so we can do Gram–Schmidt on them.

Start by computing

$$\|\mathbf{v}_1\| = \sqrt{1^2 + 2^2 + 0^2} = \sqrt{5}$$

and put

$$\mathbf{q}_1 = \frac{\mathbf{v}_1}{\|\mathbf{v}_1\|} = \frac{1}{\sqrt{5}} \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix} = \begin{bmatrix} 1/\sqrt{5} \\ 2/\sqrt{5} \\ 0 \end{bmatrix}.$$

Then we want to set

$$\mathbf{u}_2 = \mathbf{v}_2 - (\mathbf{v}_2 \cdot \mathbf{q}_1) \mathbf{q}_1.$$

We compute

$$\mathbf{v}_2 \cdot \mathbf{q}_1 = \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix} \cdot \begin{bmatrix} 1/\sqrt{5} \\ 2/\sqrt{5} \\ 0 \end{bmatrix} = \frac{1}{\sqrt{5}} + \frac{4}{\sqrt{5}} = \frac{5}{\sqrt{5}} = \sqrt{5},$$

so

$$\mathbf{u}_2 = \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix} - \sqrt{5} \begin{bmatrix} 1/\sqrt{5} \\ 2/\sqrt{5} \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix} - \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 2 \end{bmatrix}.$$

Then

$$\|\mathbf{u}_2\| = \sqrt{0^2 + 0^2 + 2^2} = 2,$$

so we put

$$\mathbf{q}_2 = \frac{\mathbf{v}_2}{\|\mathbf{v}_2\|} = \frac{1}{2} \begin{bmatrix} 0 \\ 0 \\ 2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$

Last, we want to set

$$\mathbf{u}_3 = \mathbf{v}_3 - (\mathbf{v}_3 \cdot \mathbf{q}_1)\mathbf{q}_1 - (\mathbf{v}_3 \cdot \mathbf{q}_2)\mathbf{q}_2.$$

We have

$$\mathbf{v}_3 \cdot \mathbf{q}_1 = \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 1/\sqrt{5} \\ 2/\sqrt{5} \\ 0 \end{bmatrix} = -\frac{2}{\sqrt{5}} + \frac{2}{\sqrt{5}} = 0$$

and

$$\mathbf{v}_3 \cdot \mathbf{q}_2 = \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = 0,$$

so there's not much work to do here. Just take

$$\mathbf{u}_3 = \mathbf{v}_3,$$

compute

$$\|\mathbf{u}_3\| = \|\mathbf{v}_3\| = \sqrt{(-2)^2 + 1^2 + 0^2} = \sqrt{5},$$

and set

$$\mathbf{q}_3 = \frac{\mathbf{u}_3}{\|\mathbf{u}_3\|} = \frac{1}{\sqrt{5}} \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} -2/\sqrt{5} \\ 1/\sqrt{5} \\ 0 \end{bmatrix}.$$

The result is that the list $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3$ is orthonormal and preserves spans in the sense that (34.1) holds. In particular, since there are three vectors in the list $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3$, it is an orthonormal basis for \mathbb{R}^3 . The best basis.

If the original vectors look familiar, that's because \mathbf{v}_1 and \mathbf{v}_2 are the pivot columns of

$$A = \begin{bmatrix} 1 & 2 & 1 & 7 \\ 2 & 4 & 2 & 14 \\ 0 & 0 & 2 & 8 \end{bmatrix},$$

as discussed in Examples 19.5, 21.2, and 21.12, while $\mathbf{N}(A^T) = \text{span}(\mathbf{v}_3)$, per Example 22.1. We already expect that \mathbf{v}_3 is orthogonal to \mathbf{v}_1 and \mathbf{v}_2 , and that is shown in the calculations above. In particular, since $\text{span}(\mathbf{v}_1, \mathbf{v}_2) = \text{span}(\mathbf{q}_1, \mathbf{q}_2)$, the list $\mathbf{q}_1, \mathbf{q}_2$ is an orthonormal basis for $\mathbf{C}(A)$. Again, the best basis.

35.4 Problem (★). Let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$.

(i) Suppose that for some integer j with $1 \leq j \leq n-1$, the vectors $\mathbf{v}_1, \dots, \mathbf{v}_j$ are inde-

pendent but $\mathbf{v}_{j+1} \in \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_j)$, so $\mathbf{v}_1, \dots, \mathbf{v}_{j+1}$ are dependent. What happens at the $(j+1)$ st step in the Gram–Schmidt process? [Hint: *reread the proof of Theorem 35.1.*]

(ii) Let j be an integer with $1 \leq j \leq n-1$, and now suppose that $\mathbf{v}_1, \dots, \mathbf{v}_{j+1}$ are independent, so you can do Gram–Schmidt through the $(j+1)$ st step. Suppose that $\mathbf{v}_{j+1} \cdot \mathbf{v}_k = 0$ for $k = 1, \dots, j$. What now happens at this $(j+1)$ st step in Gram–Schmidt? [Hint: *think about the third step in Example 35.3.*]

That the Gram–Schmidt procedure “preserves spans” is probably not a consequence that we expected when we originally started out with an independent list and wanted to get an orthonormal list with the same span as the whole list. (Okay, we expected that *one* span would be preserved.) Sometimes accidental consequences are nice. Look at the $n = 3$ situation.

We have independent vectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \in \mathbb{R}^m$ and orthonormal vectors $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3 \in \mathbb{R}^m$ such that the spans are preserved:

$$\begin{cases} \mathbf{v}_1 \in \text{span}(\mathbf{q}_1) \\ \mathbf{v}_2 \in \text{span}(\mathbf{q}_1, \mathbf{q}_2) \\ \mathbf{v}_3 \in \text{span}(\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3). \end{cases}$$

Since the \mathbf{q}_k are orthonormal, we have the expansions

$$\begin{cases} \mathbf{v}_1 = (\mathbf{v}_1 \cdot \mathbf{q}_1)\mathbf{q}_1 \\ \mathbf{v}_2 = (\mathbf{v}_2 \cdot \mathbf{q}_1)\mathbf{q}_1 + (\mathbf{v}_2 \cdot \mathbf{q}_2)\mathbf{q}_2 \\ \mathbf{v}_3 = (\mathbf{v}_3 \cdot \mathbf{q}_1)\mathbf{q}_1 + (\mathbf{v}_3 \cdot \mathbf{q}_2)\mathbf{q}_2 + (\mathbf{v}_3 \cdot \mathbf{q}_3)\mathbf{q}_3. \end{cases}$$

Do you see a “triangular” structure here in my very intentional typesetting?

Let’s work backwards:

$$\mathbf{v}_1 = (\mathbf{v}_1 \cdot \mathbf{q}_1)\mathbf{q}_1 = [\mathbf{q}_1 \quad \mathbf{q}_2 \quad \mathbf{q}_3] \begin{bmatrix} \mathbf{v}_1 \cdot \mathbf{q}_1 \\ 0 \\ 0 \end{bmatrix},$$

$$\mathbf{v}_2 = (\mathbf{v}_2 \cdot \mathbf{q}_1)\mathbf{q}_1 + (\mathbf{v}_2 \cdot \mathbf{q}_2)\mathbf{q}_2 = [\mathbf{q}_1 \quad \mathbf{q}_2 \quad \mathbf{q}_3] \begin{bmatrix} \mathbf{v}_2 \cdot \mathbf{q}_1 \\ \mathbf{v}_2 \cdot \mathbf{q}_2 \\ 0 \end{bmatrix},$$

and

$$\mathbf{v}_3 = (\mathbf{v}_3 \cdot \mathbf{q}_1)\mathbf{q}_1 + (\mathbf{v}_3 \cdot \mathbf{q}_2)\mathbf{q}_2 + (\mathbf{v}_3 \cdot \mathbf{q}_3)\mathbf{q}_3 = [\mathbf{q}_1 \quad \mathbf{q}_2 \quad \mathbf{q}_3] \begin{bmatrix} \mathbf{v}_3 \cdot \mathbf{q}_1 \\ \mathbf{v}_3 \cdot \mathbf{q}_2 \\ \mathbf{v}_3 \cdot \mathbf{q}_3 \end{bmatrix}.$$

Then

$$[\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3] = [\mathbf{q}_1 \quad \mathbf{q}_2 \quad \mathbf{q}_3] \begin{bmatrix} (\mathbf{v}_1 \cdot \mathbf{q}_1) & (\mathbf{v}_2 \cdot \mathbf{q}_1) & (\mathbf{v}_3 \cdot \mathbf{q}_1) \\ 0 & (\mathbf{v}_2 \cdot \mathbf{q}_2) & (\mathbf{v}_3 \cdot \mathbf{q}_2) \\ 0 & 0 & (\mathbf{v}_3 \cdot \mathbf{q}_3) \end{bmatrix}.$$

Put

$$A = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3], \quad Q = [\mathbf{q}_1 \quad \mathbf{q}_2 \quad \mathbf{q}_3], \quad \text{and} \quad R = \begin{bmatrix} (\mathbf{v}_1 \cdot \mathbf{q}_1) & (\mathbf{v}_2 \cdot \mathbf{q}_1) & (\mathbf{v}_3 \cdot \mathbf{q}_1) \\ 0 & (\mathbf{v}_2 \cdot \mathbf{q}_2) & (\mathbf{v}_3 \cdot \mathbf{q}_2) \\ 0 & 0 & (\mathbf{v}_3 \cdot \mathbf{q}_3) \end{bmatrix}$$

to see that we have factored the matrix A (which has independent columns) into the product $A = QR$, with Q orthogonal and R upper-triangular. In fact, the diagonal entries of R are positive (not just nonzero) by Problem 35.2.

We have a lot of recent knowledge about why orthogonal matrices are nice, and we have a lot of past knowledge about why upper-triangular matrices with nonzero diagonal entries are nice. We'll put all of that together with this “ QR -factorization” to obtain the ultimate form of least squares. Remember, that's what you do when you can't do what you want (which is solve $A\mathbf{x} = \mathbf{b}$, of course).

Content from Strang's ILA 6E. Page 182 develops the QR -factorization for a matrix with three independent columns.

Day 36: Monday, April 7.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Eigenvalue, eigenvector, eigenspace, geometric multiplicity

We formalize our third major matrix factorization.

36.1 Theorem (QR -factorization). Let $A \in \mathbb{R}^{m \times n}$ have independent columns (so A has full column rank: $\text{rank}(A) = n$). There exist an orthogonal matrix $Q \in \mathbb{R}^{m \times n}$ and an upper-triangular matrix $R \in \mathbb{R}^{m \times n}$ such that $A = QR$. Specifically, the columns of Q are the vectors constructed from the columns of A by the Gram–Schmidt procedure, and the (i, j) -entry of R is $\mathbf{a}_j \cdot \mathbf{q}_i$, where \mathbf{a}_j is the j th column of A , and \mathbf{q}_i is the i th column of Q .

We proved the $n = 3$ case of this. The general proof just hinges on (1) the orthonormality of the vectors produced by Gram–Schmidt, (2) the “span preservation property” that $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_j) = \text{span}(\mathbf{q}_1, \dots, \mathbf{q}_j)$ for each j , not just $j = n$, and (3) Problem 35.2 to get the positive diagonal entries in R .

36.2 Example. Let

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

We performed Gram–Schmidt on the columns of A in Example 35.3. Collect the Gram–Schmidt output in

$$Q = \begin{bmatrix} 1/\sqrt{5} & 0 & -2/\sqrt{5} \\ 2/\sqrt{5} & 0 & 1/\sqrt{5} \\ 0 & 1 & 0 \end{bmatrix}.$$

If we don't remember all of the coefficients from the Gram–Schmidt work, we can compute them quickly (and we only do this for R_{ij} with $j \geq i$):

$$R_{11} = \mathbf{a}_1 \cdot \mathbf{q}_1 = \sqrt{5},$$

$$R_{12} = \mathbf{a}_2 \cdot \mathbf{q}_1 = \sqrt{5},$$

$$R_{13} = \mathbf{a}_3 \cdot \mathbf{q}_1 = 0,$$

$$R_{22} = \mathbf{a}_2 \cdot \mathbf{q}_2 = 2,$$

$$R_{23} = \mathbf{a}_3 \cdot \mathbf{q}_2 = 0,$$

$$R_{33} = \mathbf{a}_3 \cdot \mathbf{q}_3 = \sqrt{5}.$$

Then

$$R = \begin{bmatrix} \sqrt{5} & \sqrt{5} & 0 \\ 0 & 2 & 0 \\ 0 & 0 & \sqrt{5} \end{bmatrix},$$

and we have

$$\begin{bmatrix} 1 & 1 & -2 \\ 2 & 2 & 1 \\ 0 & 2 & 0 \end{bmatrix} = \begin{bmatrix} 1/\sqrt{5} & 0 & -2/\sqrt{5} \\ 2/\sqrt{5} & 0 & 1/\sqrt{5} \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \sqrt{5} & \sqrt{5} & 0 \\ 0 & 2 & 0 \\ 0 & 0 & \sqrt{5} \end{bmatrix}$$

36.3 Problem (!). Let $A \in \mathbb{R}^{m \times n}$ have independent columns. What is the CR -factorization of A ? When is the R in that factorization the same as the R in the QR -factorization?

Here is how the QR -factorization is useful for least squares. Start with $A \in \mathbb{R}^{m \times n}$ with independent columns and factor $A = QR$ with $Q \in \mathbb{R}^{m \times n}$ orthogonal and $R \in \mathbb{R}^{n \times n}$ upper-triangular with positive diagonal entries.

If $\mathbf{b} \notin \mathbf{C}(A)$, then solving $A\hat{\mathbf{x}} = P_A\mathbf{b}$ is the next best thing to solving the unsolvable problem $A\mathbf{x} = \mathbf{b}$. While we do have a formula for $\hat{\mathbf{x}}$, the annoying thing is that it requires computing the inverse $(A^T A)^{-1}$. Better to solve a linear system than compute an inverse, and here is the system you want to solve.

First, expand $P_A = A(A^T A)^{-1}A^T$, so $A\hat{\mathbf{x}} = P_A\mathbf{b}$ becomes

$$A\hat{\mathbf{x}} = A(A^T A)^{-1}A^T\mathbf{b}.$$

Since A has full column rank, if $A\hat{\mathbf{x}} = A\mathbf{y}$, then $\hat{\mathbf{x}} = \mathbf{y}$. So, $\hat{\mathbf{x}} = (A^T A)^{-1}A^T\mathbf{b}$. This is the formula for $\hat{\mathbf{x}}$ that we previously developed—nothing new here, but I wanted to review it with you. The right idea is not to stay with this formula but to make life seemingly more complicated:

$$(A^T A)\hat{\mathbf{x}} = A^T\mathbf{b}. \quad (36.1)$$

We could always use Gaussian elimination to solve (36.1). Again, better to solve a linear system than compute an inverse.

But this system (36.1) is pretty nice after the QR -factorization. Since $A = QR$, we have $A^\top = (QR)^\top = R^\top Q^\top$, and since Q is orthogonal, we have $Q^\top Q = I_n$. Then

$$A^\top A = (R^\top Q^\top)(QR) = R^\top(Q^\top Q)R = R^\top I_n R = R^\top R.$$

The problem (36.1) now reads

$$R^\top R \hat{\mathbf{x}} = R^\top Q^\top \mathbf{b}. \quad (36.2)$$

This is great! Since R is upper-triangular with positive diagonal entries, R^\top is lower-triangular with positive diagonal entries.

36.4 Problem (★). Prove that. [Hint: recall that $R_{ij}^\top = R_{ji}$. Since R has positive diagonal entries, $R_{ii} > 0$. Since R is upper-triangular, $R_{ij} = 0$ for $i > j$. To show that R^\top is lower-triangular, you want $R_{ij}^\top = 0$ for $j > i$. Is this true?]

Any triangular matrix with nonzero diagonal entries is invertible, so R^\top is invertible. Then (36.2) is just

$$R \hat{\mathbf{x}} = Q^\top \mathbf{b}.$$

Again, R is upper-triangular with positive diagonal entries, so we can solve this system by back-substitution (no need even for Gaussian elimination!) No inverses anywhere in the actual calculations, just in the theory.

Content from Strang's ILA 6E. Pages 182–183 discuss how to use the QR -factorization in least squares. Read the second half of p. 185, which summarizes everything. You don't need to read about the pseudoinverses.

Then read about the “victory of orthogonality” on p. 197. Stop with #5 for now, and there just be able to explain why if $Q \in \mathbb{R}^{n \times n}$ is orthogonal, so is any power Q^k for $k \geq 1$.

In the paragraph after that, note the “sum of squares definition of length.” There are many valid, meaningful ways of defining the length of a vector (pp. 355–356), but the way that interacts best with the dot product is saying length is $\|\mathbf{v}\| = \sqrt{\mathbf{v} \cdot \mathbf{v}}$. As you think about least squares, keep in mind how length and dot product interact so nicely.

I like to think that courses tell stories, and sometimes stories have abrupt plot twists. Most of the story of this course has been solving, then understanding, and finally approximating $A\mathbf{x} = \mathbf{b}$. We've introduced a significant amount of vocabulary, notation, and technology to do this. Now the time has come to ask a different question, mostly about A , and less about $A\mathbf{x} = \mathbf{b}$ (although we will still think about that). I've said many times, and I hope you agree, that matrices are *static* and *dynamic*: they *encode* data and they *act* on data.

Namely, matrices act by multiplying other matrices and vectors (matrix-matrix multiplication is, of course, matrix-vector multiplication done repeatedly). When we have to choose between what is right and what is easy in math, we always want what's easy. What is the easiest action of a matrix? This is a little subjective, but I'd say it's when there's as little multiplication involved as possible. What if $A = \lambda I_n$ for some $\lambda \in \mathbb{R}$? (The Greek letter λ

is a traditional cultural thing here.) Then multiplying by A is just scalar multiplication by λ .

The matrix λI_n is diagonal with constant diagonal. Maybe the next simplest matrix is still diagonal but has a nonconstant diagonal, say,

$$A = \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix}.$$

If $\mathbf{v} = c\mathbf{e}_j$ for some $c \in \mathbb{R}$ and $j = 1, 2, 3$, then $A\mathbf{v} = c\lambda_j\mathbf{e}_j$. For these special vectors, A still acts like scalar multiplication.

The right question to ask in our search for simple matrix operations is where/when/how does the matrix act just as scalar multiplication? If $A \in \mathbb{R}^{n \times n}$, are there $\mathbf{v} \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}$ such that

$$A\mathbf{v} = \lambda\mathbf{v}?$$

36.5 Problem (!). Explain why the equality above only makes sense for square A .

The boring answer is yes when $\mathbf{v} = \mathbf{0}_n$. For this reason, we make the following restriction.

36.6 Definition. Let $A \in \mathbb{R}^{n \times n}$. A nonzero vector $\mathbf{v} \in \mathbb{R}^n$ is an **EIGENVECTOR** of A corresponding to the **EIGENVALUE** $\lambda \in \mathbb{R}$ if

$$A\mathbf{v} = \lambda\mathbf{v}.$$

36.7 Remark. Here's some linguistic commentary from a footnote on p. 69 of the excellent *Linear Algebra* by Meckes and Meckes: “[The words ‘eigenvector’ and ‘eigenvalue’] are halfway translated from the German words “Eigenvektor” and “Eigenwert.” The German adjective ‘eigen’ can be translated as ‘own’ or ‘proper,’ so an eigenvector of [a matrix] is something like ‘the matrix’s very own vector.’” Because of these German origins, Tefethen and Bau’s excellent *Numerical Linear Algebra* suggests abbreviating “eigenvector” by “ev” and “eigenvalue” by “ew.”

The eigenvalues and eigenvectors of a square matrix encode a huge amount of information about it, and this will become apparent over time—trust me. Here is one quick application. Say that we want to compute “matrix powers”: A^k for $k \geq 2$. Here $A^2 = AA$, $A^3 = A^2A = AAA$, and so on. If \mathbf{v} is an eigenvector corresponding to λ , then

$$A^2\mathbf{v} = A(A\mathbf{v}) = A(\lambda\mathbf{v}) = \lambda(A\mathbf{v}) = \lambda(\lambda\mathbf{v}) = \lambda^2\mathbf{v},$$

and, more generally,

$$A^k\mathbf{v} = \lambda^k\mathbf{v}.$$

This is vastly easier than computing A^k and then multiplying $A^k\mathbf{v}$.

Content from Strang's ILA 6E. Read the first two paragraphs on p. 216 and then all of p. 217.

To warm up, we'll think computationally and just try to find eigenvalues and eigenvectors. Our computations will not exactly be profound. To be fair, all of our matrices in this course have been “toys”—I've selected (or stolen) the entries so that the arithmetic is very easy to do by hand. But, in the end, it's all been arithmetic. All of the elementary row operations and the calculations in the Gram–Schmidt procedure have been, fundamentally, *arithmetic*: adding and multiplying numbers. Finding eigenvalues, however, will fundamentally be a *transcendental* operation that we won't be able to resolve, except in very special cases, with a neat and finite algorithm.

The first thing that you might think is that the “eigenproblem” $A\mathbf{v} = \lambda\mathbf{v}$ is too hard because it contains two kinds of unknowns: the vector $\mathbf{v} \in \mathbb{R}^n$ and the scalar $\lambda \in \mathbb{R}$. What usually happens is that we find the eigenvalues first and then $A\mathbf{v} = \lambda\mathbf{v}$ becomes the matrix-vector equation

$$(A - \lambda I_n)\mathbf{v} = \mathbf{0}_n. \quad (36.3)$$

We know how to check if this has a nonzero solution. The greater utility of (36.3) is that it tells us what happens when λ is an eigenvalue: $\mathbf{N}(A - \lambda I_n) \neq \{\mathbf{0}_n\}$, so $A - \lambda I_n$ is not invertible.

36.8 Example. Let

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

We want to find $\lambda \in \mathbb{R}$ such that $A - \lambda I_2$ is not invertible. We have

$$A - \lambda I_2 = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix} - \lambda \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix} - \begin{bmatrix} \lambda & 0 \\ 0 & \lambda \end{bmatrix} = \begin{bmatrix} (1 - \lambda) & 1 \\ 0 & (2 - \lambda) \end{bmatrix}.$$

This last matrix is upper-triangular, so we know it's not invertible (precisely) when a diagonal entry is zero. This happens when $1 - \lambda = 0$ or $2 - \lambda = 0$. Thus the eigenvalues are $\lambda = 1$ and $\lambda = 2$. Surely it's no coincidence that the diagonal entries are the eigenvalues.

Just for practice, we find eigenvectors corresponding to the eigenvalue 1. We want to find $\mathbf{v} \in \mathbb{R}^2$ such that $(A - I_2)\mathbf{v} = \mathbf{0}_2$, and we want $\mathbf{v} \neq \mathbf{0}_2$ to keep things interesting. We look at

$$\begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix} \mathbf{v} = \mathbf{0}_2.$$

This becomes the linear system

$$\begin{cases} v_2 = v_1 \\ v_2 = v_2. \end{cases}$$

The second equation tells us nothing useful (of course $v_2 = v_2$, what else would it equal?), while the first tells us

$$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} v_2 \\ v_2 \end{bmatrix} = v_2 \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

So, the eigenvectors corresponding to the eigenvalue 1 are all scalar multiples of $\mathbf{v} = (1, 1)$.

Here is the generalization of the eigenvalue result from this example.

36.9 Theorem. *Let $A \in \mathbb{R}^{n \times n}$ be triangular. Then the eigenvalues of A are the diagonal entries of A .*

36.10 Problem (★). Prove it.

The eigenvector calculation in Example 36.8 revealed the eigenvectors corresponding to a particular eigenvalue as the vectors in a certain span. This generalizes nicely.

36.11 Problem (★). Suppose that $\lambda \in \mathbb{R}$ is an eigenvalue of $A \in \mathbb{R}^{n \times n}$. The **EIGENSPACE** of A corresponding to λ is

$$\mathbf{E}(A, \lambda) := \{\mathbf{v} \in \mathbb{R}^n \mid A\mathbf{v} = \lambda\mathbf{v}\}.$$

The **GEOMETRIC MULTIPLICITY** of λ as an eigenvalue of A is $g(A, \lambda) := \dim[\mathbf{E}(A, \lambda)]$.

(i) Prove that $\mathbf{E}(A, \lambda)$ is a subspace of \mathbb{R}^n . [Hint: for extra practice, try proving this in two ways: from Definition 36.6 alone and by thinking about null spaces.]

(ii) Is every vector in $\mathbf{E}(A, \lambda)$ an eigenvector of A ?

Now for some bad news.

36.12 Problem (★). (i) By considering

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix},$$

explain why A and its RREF need not have the same eigenvalues.

(ii) Suppose that $A, E \in \mathbb{R}^{n \times n}$ and $U = EA$ is upper-triangular. Do you expect A and U to have the same eigenvalues?

Content from Strang's ILA 6E. Read Example 1 on pp. 218–219. Convince yourself that if λ is an eigenvalue of A with eigenvector \mathbf{v} , then λ^2 is an eigenvalue of A^2 , still with eigenvector \mathbf{v} .

Day 37: Wednesday, April 9.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Diagonalizable matrix

37.1 Example. Let

$$A = \begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix}.$$

We could think about $A - \lambda I_2$, or we could think about what an eigenvector \mathbf{v} and an eigenvalue λ *do*. They satisfy $A\mathbf{v} = \lambda\mathbf{v}$. Let's compute

$$A\mathbf{v} = \begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = 2(v_1 + v_2) \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

This almost looks like $A\mathbf{v} = \lambda\mathbf{v}$, if we pick λ and \mathbf{v} correctly. The vector on the right is $(1, 1)$, so if we try $\mathbf{v} = (1, 1)$, then we get $A\mathbf{v} = 4\mathbf{v}$. This says that 4 is an eigenvalue with eigenvector $(1, 1)$.

Is this the only one? Let's think about (36.3) now. We want $A - \lambda I_2$ to fail to be invertible, and it looks like A is already not invertible, since its columns are dependent. We can get just A to show up in $A - \lambda I_2$ by taking $\lambda = 0$. Will this be an eigenvalue? Does $A\mathbf{v} = \mathbf{0}_2$ have a nontrivial solution? Sure: $\mathbf{v} = (1, -1)$. This says that 0 is an eigenvalue with eigenvector $(1, -1)$.

Here are more generalizations of the previous example for you to consider.

37.2 Problem (★). Let $A \in \mathbb{R}^{n \times n}$ such that the sum of the entries in any row of A is always the same value $s \in \mathbb{R}$.

(i) Prove that s is an eigenvalue of A . [Hint: for what $\mathbf{v} \in \mathbb{R}^n$ does $A\mathbf{v}$ involve adding the entries in each row?]

(ii) Is 0 always an eigenvalue of A ?

37.3 Problem (★). Prove that $A \in \mathbb{R}^{n \times n}$ is not invertible if and only if 0 is an eigenvalue of A .

Content from Strang's ILA 6E. Read Examples 2 and 3 on pp. 221–222.

We've seen some special cases of eigenexistence. (Every word is better when you put “eigen” in front of it.) It turns out that every matrix has at least one eigenvalue...it just

may not be real.

This is easiest to see at the 2×2 level. Let

$$A = \begin{bmatrix} a & c \\ b & d \end{bmatrix} \in \mathbb{R}^{2 \times 2}.$$

Then $\lambda \in \mathbb{R}$ is an eigenvalue of A if and only if $A - \lambda I_2$ is not invertible. While we haven't stressed it all that much, I know you know that a 2×2 matrix is not invertible if and only if its determinant is 0:

$$\det \left(\begin{bmatrix} p & q \\ r & s \end{bmatrix} \right) = ps - rq.$$

So, $A - \lambda I_2$ is not invertible if and only if

$$0 = \det(A - \lambda I_2) = \det \left(\begin{bmatrix} (a - \lambda) & c \\ b & (d - \lambda) \end{bmatrix} \right) = (a - \lambda)(d - \lambda) - cb.$$

I claim that $(a - \lambda)(d - \lambda) - cb$ is just a quadratic in λ , and we have years of experience studying that:

$$(a - \lambda)(d - \lambda) - cb = ad - a\lambda - d\lambda + \lambda^2 - cb = \lambda^2 - (a + d)\lambda + (ad - bc).$$

This has a nice structure: the coefficient $a + d$ is the sum of the diagonal entries of A , which is its **TRACE**, denoted $\text{tr}(A)$. And $ad - bc$ is its determinant. So, λ is an eigenvalue of A if and only if

$$\lambda^2 - \text{tr}(A)\lambda + \det(A) = 0.$$

This is the **CHARACTERISTIC EQUATION** of A , and the quadratic on the left is the **CHARACTERISTIC POLYNOMIAL** of A .

37.4 Problem (!). Revisit the matrices in Examples 36.8 and 37.1 and compute their eigenvalues by finding the roots of their characteristic polynomials.

Content from Strang's ILA 6E. Read pp. 220–222 on determinants. For now, just assume that A is 2×2 throughout. Then read Worked Example 6.1 A on pp. 224–225. Can you prove in general the statements about the eigenvalues and eigenvectors of A^2 , A^{-1} , and $A + cI$ (with $c \in \mathbb{R}$), relative to the eigenvalues and eigenvectors of A ?

We have years of painful experience with solving quadratic equations. In particular, they don't always have two distinct real solutions, and so we shouldn't expect a 2×2 matrix always to have two distinct real eigenvalues.

37.5 Example. Let

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

This is upper-triangular with only 1's on the diagonal, so the only eigenvalue is 1. If

$A\mathbf{v} = \mathbf{v}$, then

$$\begin{cases} v_1 + v_2 = v_1 \\ v_2 = v_2 \end{cases}$$

The second equation tells us nothing useful, but the first collapses to $v_2 = 0$, so

$$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} v_1 \\ 0 \end{bmatrix} = v_1 \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

We can see this at the level of the characteristic equation by computing

$$\det(A - \lambda I_2) = \det \left(\begin{bmatrix} (1 - \lambda) & 1 \\ 0 & (1 - \lambda) \end{bmatrix} \right) = (1 - \lambda)^2,$$

which has only $\lambda = 1$ as its (repeated real) root.

37.6 Problem (!). For an arbitrary integer $n \geq 1$, give an example of a matrix $A \in \mathbb{R}^{n \times n}$ with only one eigenvalue $\lambda \in \mathbb{R}$ but such that $\dim[\mathbf{E}(A, \lambda)] = n$.

And quadratic equations don't always have real roots, even when the coefficients are real.

37.7 Example. Let

$$A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}.$$

Then

$$\det(A - \lambda I_2) = \det \left(\begin{bmatrix} -\lambda & -1 \\ 1 & -\lambda \end{bmatrix} \right) = \lambda^2 + 1.$$

The quadratic equation $\lambda^2 + 1 = 0$ has $\lambda = \pm i$ as solutions. Here i is the **COMPLEX NUMBER** such that $i^2 = -1$.

We can find an eigenvector for i just as before: solve $A\mathbf{v} = i\mathbf{v}$. This becomes the system

$$\begin{cases} -v_2 = iv_1 \\ v_1 = iv_2. \end{cases}$$

The second equation is slightly easier and gives

$$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} iv_2 \\ v_2 \end{bmatrix} = v_2 \begin{bmatrix} i \\ 1 \end{bmatrix}.$$

So, all eigenvectors of A corresponding to i are scalar multiples of $(i, 1)$.

37.8 Problem (!). With A from the previous example, find all eigenvectors corresponding to $-i$.

We need a more generous notion of eigenvalue and eigenvector than afforded by Definition

36.6. Let

$$\mathbb{C} = \{x + iy \mid x, y \in \mathbb{R}, i^2 = -1\}.$$

We add and multiply numbers in \mathbb{C} by following our noses, combining like terms, and using $i^2 = -1$:

$$(1 + 2i) + (3 + 4i) = (1 + 3) + (2i + 4i) = 4 + 6i$$

and

$$(1 + 2i)(3 + 4i) = 1 + 4i + 6i + 8i^2 = (1 - 8) + 10i = -7 + 10i.$$

Let \mathbb{C}^n be the set of all column vectors with entries in \mathbb{C} . We add vectors in \mathbb{C}^n componentwise and multiply by scalars in \mathbb{C} componentwise, too. *Nothing changes in the arithmetic.*

37.9 Definition (Improvement of Definition 36.6). A scalar $\lambda \in \mathbb{C}$ is an **EIGENVALUE** of the matrix $A \in \mathbb{R}^{n \times n}$ corresponding to the **EIGENVECTOR** $\mathbf{v} \in \mathbb{C}^n$ if

$$A\mathbf{v} = \lambda\mathbf{v}.$$

Content from Strang's ILA 6E. For a refresher on complex numbers, look at pp. 262–263. Then Read “Imaginary Eigenvalues on p. 223 and Worked Example 6.1 B on p. 225.

Now we can state more precisely an eigenexistence result.

37.10 Theorem. Let $A \in \mathbb{R}^{n \times n}$. There exists at least one $\lambda \in \mathbb{C}$ such that λ is an eigenvalue of A .

We will eventually sketch a proof of this. That proof, however, will not exactly be constructive—we’re not going to get a simple algorithm for finding eigenvalues. (This is a really. big. deal. in numerical linear algebra.) More accessible is an *upper* bound on the number of eigenvalues that a matrix can have: at most n . This has been achieved in most of our examples so far.

Here’s why.

37.11 Theorem. Eigenvectors corresponding to distinct eigenvalues are independent. More precisely, let $A \in \mathbb{C}^{n \times n}$ and suppose that $\mathbf{v}_1, \dots, \mathbf{v}_j \in \mathbb{C}^n$ are eigenvectors of A corresponding to the distinct eigenvalues $\lambda_1, \dots, \lambda_j \in \mathbb{C}$. (That is, $A\mathbf{v}_k = \lambda_k\mathbf{v}_k$ and $\lambda_k \neq \lambda_\ell$ for $k \neq \ell$.) Then $\mathbf{v}_1, \dots, \mathbf{v}_j$ are linearly independent.

Proof. This is really an induction argument on j . For $j = 1$, there is only one eigenvector, and that can’t be $\mathbf{0}_n$.

I’ll just do the $j = 2$ case and stop there. For $j = 2$, suppose that $A\mathbf{v}_1 = \lambda_1\mathbf{v}_1$, $A\mathbf{v}_2 = \lambda_2\mathbf{v}_2$, $\mathbf{v}_1 \neq \mathbf{0}_n$, $\mathbf{v}_2 \neq \mathbf{0}_n$, and $\lambda_1 \neq \lambda_2$. Start with $c_1\mathbf{v}_1 + c_2\mathbf{v}_2 = \mathbf{0}_n$. We want to prove $c_1 = c_2 = 0$. We should probably get the eigenvalues to show up, and one way to do that is to make A show up. We have

$$A(c_1\mathbf{v}_1 + c_2\mathbf{v}_2) = A\mathbf{0}_n,$$

and so

$$\mathbf{0}_n = c_1 A \mathbf{v}_1 + c_2 A \mathbf{v}_2 = c_1 \lambda_1 \mathbf{v}_1 + c_2 \lambda_2 \mathbf{v}_2.$$

We therefore have a sort of system for c_1 and c_2 :

$$\begin{cases} c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2 = \mathbf{0}_n \\ c_1 \lambda_1 \mathbf{v}_1 + c_2 \lambda_2 \mathbf{v}_2 = \mathbf{0}_n. \end{cases}$$

If \mathbf{v}_1 and \mathbf{v}_2 were scalars, we'd try Gaussian elimination, and that sort of still works here. Multiply the first equation by λ_1 and subtract it from the second equation:

$$\mathbf{0}_n = (c_1 \lambda_1 \mathbf{v}_1 + c_2 \lambda_2 \mathbf{v}_2) - \lambda_1 (c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2) = c_2 (\lambda_2 - \lambda_1) \mathbf{v}_2.$$

Since $\mathbf{v}_2 \neq \mathbf{0}_n$ and $\lambda_1 \neq \lambda_2$ and $c_2 (\lambda_2 - \lambda_1) \mathbf{v}_2 = \mathbf{0}_n$, the only possibility is that $c_2 = 0$. And since we're assuming $c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2 = \mathbf{0}_n$, this collapses to $c_1 \mathbf{v}_1 = \mathbf{0}_n$. Because $\mathbf{v}_1 \neq \mathbf{0}_n$, we get $c_1 = 0$, as desired. ■

37.12 Problem (+). Prove the $j = 3$ case, using the $j = 2$ case. [Hint: assume $c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2 + c_3 \mathbf{v}_3 = \mathbf{0}_n$, multiply by A , and subtract a multiple by the right λ_k to conclude $c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2 = \mathbf{0}_n$. The $j = 2$ case says that $\mathbf{v}_1, \mathbf{v}_2$ are independent.]

37.13 Problem (!). Give an example of a matrix A with eigenvalue λ and eigenvectors $\mathbf{v}_1, \mathbf{v}_2$ that both correspond to λ such that $\mathbf{v}_1, \mathbf{v}_2$ are dependent.

Content from Strang's ILA 6E. Independence of eigenvectors corresponding to distinct eigenvalues is proved in the "Fact" at the bottom of p. 234 through the top of p. 235.

Now we can explain why a matrix has at most n distinct eigenvalues. If $A \in \mathbb{R}^{n \times n}$ has any more, $n + 1$ or higher, then there would be a list of at least $n + 1$ independent vectors in \mathbb{R}^n . This contradicts stuff about dimension.

Day 38: Friday, April 11.

This is all of the groundwork that we need to appreciate eigenvalues—hopefully not too much prep. The best situation is when a matrix has n independent eigenvectors. This definitely happens if the matrix has n independent eigenvalues, but that's not strictly necessary (I_n has only one eigenvalue but $\mathbf{e}_1, \dots, \mathbf{e}_n$ are all eigenvectors).

Let's start small. Say that $A \in \mathbb{R}^{3 \times 3}$ has 3 independent eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \in \mathbb{C}^3$. This means $A \mathbf{v}_k = \lambda_k \mathbf{v}_k$ for some $\lambda_k \in \mathbb{C}$, and also $\mathbf{v}_k \neq \mathbf{0}_3$. We're not requiring λ_1, λ_2 , and λ_3 to be distinct, although that doesn't hurt. Slap it all together in a matrix:

$$[A \mathbf{v}_1 \quad A \mathbf{v}_2 \quad A \mathbf{v}_3] = [\lambda_1 \mathbf{v}_1 \quad \lambda_2 \mathbf{v}_2 \quad \lambda_3 \mathbf{v}_3].$$

You've made it this far in life, so you know how to factor:

$$A [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3] = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3] \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix}.$$

(Multiplying by a diagonal matrix on the right = scaling the columns. I always mix this up!)

Now abbreviate

$$V = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3] \quad \text{and} \quad \Lambda = \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix}$$

so that we have

$$AV = V\Lambda.$$

Here's the great thing: $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ are independent, so V is invertible. We conclude

$$A = V\Lambda V^{-1}.$$

Nothing here was special about $n = 3$, and the following holds more generally.

38.1 Theorem (Diagonalization). Let $A \in \mathbb{R}^{n \times n}$ have n linearly independent eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{C}^n$ corresponding to the eigenvalues $\lambda_1, \dots, \lambda_n \in \mathbb{C}$. (That is, $A\mathbf{v}_k = \lambda_k \mathbf{v}_k$. We do not require the λ_k to be distinct.) Put

$$V = [\mathbf{v}_1 \quad \cdots \quad \mathbf{v}_n] \quad \text{and} \quad \Lambda = \text{diag}(\lambda_1, \dots, \lambda_n) = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}.$$

Then

$$A = V\Lambda V^{-1}.$$

38.2 Definition. A matrix $A \in \mathbb{R}^{n \times n}$ is **DIAGONALIZABLE** if there are an invertible matrix $V \in \mathbb{C}^{n \times n}$ and a diagonal matrix $\Lambda \in \mathbb{C}^{n \times n}$ such that $A = V\Lambda V^{-1}$.

38.3 Problem (★). Prove that if A is diagonalizable, then the columns of V are eigenvectors of A corresponding to the eigenvalues given by the diagonal elements of Λ . [Hint: $A = V\Lambda V^{-1}$ means $AV = \Lambda V$.]

38.4 Example. Let

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

We saw in Example 36.8 that the eigenvalues of A are 1 and 2 and that $\mathbf{v}_1 = \mathbf{e}_1$ is an eigenvector corresponding to the eigenvalue 1. I claim that $\mathbf{v}_2 = (1, 1)$ is an eigenvector corresponding to the eigenvalue 2. The vectors $(1, 0)$ and $(1, 1)$ are independent—we expect this from eigentheory and we can check this in any number of ways.

Put

$$V = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \Lambda = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}.$$

Then

$$V^{-1} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix}.$$

By the way, V is an elimination matrix, so we know how to find its inverse without any fancy calculations. We should then have

$$\begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix},$$

and you can check this by hand.

Here is one major upshot of a diagonalizable matrix $A = V\Lambda V^{-1}$: computing its powers is so easy. For example,

$$A^2 = (V\Lambda V^{-1})(V\Lambda V^{-1}) = V\Lambda(V^{-1}V)\Lambda V^{-1} = V\Lambda^2 V^{-1},$$

and, more generally,

$$A^k = V\Lambda^k V^{-1}.$$

And Λ^k is so easy to compute: just raise each diagonal entry to the power k .

38.5 Problem (★). Suppose that the eigenvalues of $A \in \mathbb{R}^{4 \times 4}$ are $1/2$, $1/3$, $1/4$, and $1/5$. How do the powers A^k behave when k is very large?

Here is another upshot. If your problem revolves around a diagonalizable matrix $A = V\Lambda V^{-1}$, then the columns of V might form the “best basis” of \mathbb{R}^n for that problem. (I’m assuming that the eigenvalues of A are real and that the eigenvectors of A are vectors in \mathbb{R}^n , not \mathbb{C}^n here. It’s worthwhile to ask what conditions on A guarantee this.) First, they *are* a basis: there are n of them, and they’re independent. Second, they’re eigenvectors of A by Problem 38.3.

Here is how this helps. I’ll do this at the level of $n = 3$ first. Any $\mathbf{v} \in \mathbb{R}^3$ has the form $\mathbf{v} = c_1\mathbf{v}_1 + c_2\mathbf{v}_2 + c_3\mathbf{v}_3$ for some $c_k \in \mathbb{R}$, and so

$$A\mathbf{v} = A(c_1\mathbf{v}_1 + c_2\mathbf{v}_2 + c_3\mathbf{v}_3) = c_1A\mathbf{v}_1 + c_2A\mathbf{v}_2 + c_3A\mathbf{v}_3 = c_1\lambda_1\mathbf{v}_1 + c_2\lambda_2\mathbf{v}_2 + c_3\lambda_3\mathbf{v}_3.$$

So, computing $A\mathbf{v}$ is really just a matter of multiplying by the λ_k . What really mattered wasn’t \mathbf{v} itself but rather its “coordinates” relative to the basis $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$.

Here is a way to see this more broadly for an arbitrary n . Suppose $A = V\Lambda V^{-1}$ with $V, \Lambda \in \mathbb{R}^{n \times n}$. Any $\mathbf{x} \in \mathbb{R}^n$ has the form $\mathbf{x} = V\mathbf{c}$ for some $\mathbf{c} \in \mathbb{R}^n$, and then

$$A\mathbf{v} = (V\Lambda V^{-1})(V\mathbf{c}) = V\Lambda(V^{-1}V)\mathbf{c} = V\Lambda\mathbf{c} = \begin{bmatrix} \mathbf{v}_1 & \cdots & \mathbf{v}_n \end{bmatrix} \begin{bmatrix} \lambda_1 c_1 \\ \vdots \\ \lambda_n c_n \end{bmatrix} = \lambda_1 c_1 \mathbf{v}_1 + \cdots + \lambda_n c_n \mathbf{v}_n.$$

Again, the dynamic thing is \mathbf{c} , the coordinates of \mathbf{v} relative to the basis from V , not \mathbf{v} itself.

This begs a good question. What conditions on A guarantee that the eigenvectors can be chosen to be real (in \mathbb{R}^n)? What conditions on A guarantee that we can form a basis

for \mathbb{R}^n from its eigenvectors? (Don't say A has n distinct real eigenvalues.) And what conditions on A guarantee that it's easy to extract the coordinates of an arbitrary vector relative to the basis that ideally comes from its eigenvectors? (You know the answer to this one: orthonormal eigenvectors!)

Content from Strang's ILA 6E. Read all of pp. 232–233. Then think about the four remarks on pp. 233–234.

Day 39: Monday, April 14.

We don't really do applications in this class (other than that brief mention in least squares of fitting a line to data), because the application really *is* the class: the application of linear algebra to solving, understanding, and approximating $A\mathbf{x} = \mathbf{b}$, which is a universal equation. Or maybe the class is the application—I don't know. Still, I grant that our work so far on eigenvalues may seem more esoteric and theoretical and less connected to $A\mathbf{x} = \mathbf{b}$ than anything from before.

That's because it is. Trefethen and Bau's magnificent *Numerical Linear Algebra* (pp. 181–182) says the following about eigenvalues:

“Eigenvalue problems have a very different character from the problems involving square or rectangular systems of linear equations. . . To ask about the eigenvalues of a [nonsquare matrix] A would be meaningless. Eigenvalue problems make sense only when the [matrix is square]. This reflects the fact that in applications, eigenvalues are generally used when a matrix is to be compounded iteratively. . .

Broadly speaking, eigenvalues and eigenvectors are useful for two reasons, one algorithmic, the other physical. Algorithmically, eigenvalue analysis can simplify solutions of certain problems by reducing a coupled system to a collection of scalar problems. Physically, eigenvalue analysis can give insight into the behavior of evolving systems governed by linear equations. The most familiar examples in this latter class are the study of *resonance* (e.g., of musical instruments when struck or plucked or bowed) and of *stability* (e.g., of fluid flows subjected to small perturbations). In such cases eigenvalues tend to be particularly useful for analyzing behavior for large times t .”

Here is one physical situation in which eigenvalues arise and actually simplify the analysis. (This is why you took a calculus course before taking linear algebra.) Newton's second law (force = mass \times acceleration) says that the displacement x of a harmonic oscillator (a coupled mass-spring system) from its equilibrium location in the absence of external forces is governed by the second-order constant-coefficient linear homogeneous ordinary differential equation

$$\begin{cases} \ddot{x} + b\dot{x} + \kappa x = 0 \\ x(0) = x_0 \\ \dot{x}(0) = y_0. \end{cases} \quad (39.1)$$

Here the oscillator has mass 1 (which is why the coefficient on the second derivative \ddot{x} = acceleration is 1), and the spring exerts a restoring force proportional to $\kappa > 0$ on the oscillator. The parameter $b \geq 0$ measures the friction that the oscillator experiences as it moves (and we encode friction into a model as a term proportional to the first derivative \dot{x} —the faster you bicycle, the more wind resistance you encounter). This ODE arises because force = $-b\dot{x} - \kappa x$ equals mass \times acceleration = \ddot{x} . If $b = 0$, then we're in the magical world of no friction. That the right side of the ODE is 0 represents the absence of external forces; the spring force and the friction force (if the latter is present) count as “internal” forces that arise from the mass-spring system itself. Last, x_0 and y_0 are the initial displacement and velocity, respectively, of the oscillator.

One way to understand this ODE is to think about perhaps the simplest nontrivial ODE that there is and try to learn from that. That ODE is exponential growth:

$$\dot{x} = rx.$$

Here $r \in \mathbb{R}$ is a proportionality constant that says that the growth rate \dot{x} of a quantity is proportional to the amount x present. I hope we all agree that every solution to this ODE has the form $x(t) = Ce^{rt}$ for some $C \in \mathbb{R}$.

We can make (39.1) look like exponential growth by rewriting it. Here's the trick: put $y = \dot{x}$, so

$$\dot{y} = \ddot{x} = -b\dot{x} - \kappa x = -by - \kappa x.$$

Now stuff everything into a vector and grind away:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} y \\ -by - \kappa x \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -\kappa & -b \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}.$$

Put

$$\mathbf{v} = \begin{bmatrix} x \\ y \end{bmatrix}, \quad \dot{\mathbf{v}} = \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix}, \quad \text{and} \quad A = \begin{bmatrix} 0 & 1 \\ -\kappa & -b \end{bmatrix}, \quad \text{and} \quad \mathbf{v}_0 = \begin{bmatrix} x_0 \\ y_0 \end{bmatrix}.$$

Then (39.1) is equivalent to

$$\begin{cases} \dot{\mathbf{v}} = A\mathbf{v} \\ \mathbf{v}(0) = \mathbf{v}_0. \end{cases}$$

Content from Strang's ILA 6E. Page 270 reviews the exponential growth ODE $\dot{x} = rx$ and introduces systems of linear ODE in matrix-vector form. Pages 271–272 discuss how to find solutions using exponentials and eigenvalues. Pages 272–274 turn second-order linear ODE into linear systems.

If we're sneaky, we can make eigenvalues and diagonalization show up and reduce this to a pair of exponential growth ODE. Here's how. The eigenvalues λ of A must satisfy

$$\lambda^2 + b\lambda + \kappa = 0,$$

which looks a lot like the ODE (39.1).

39.1 Problem (!). Check this.

The quadratic formula tells us that the eigenvalues are

$$\lambda_{\pm} = \frac{-b \pm \sqrt{b^2 - 4\kappa}}{2}.$$

To keep things reasonably simple, I'm going to assume $b > 0$ and $b^2 - 4\kappa > 0$. That is, friction is present ($b > 0$), and friction is sort of stronger than the spring force ($b^2 - 4\kappa > 0$). Then the eigenvalues λ_{\pm} are real, distinct, and negative. (Other choices of b and κ lead to other eigensituations.)

39.2 Problem (*). Check this. [Hint: since $b > 0$, we have $b = \sqrt{b^2}$, and since $b^2 - 4\kappa > 0$ and $\kappa > 0$, we have $\sqrt{b^2} > \sqrt{b^2 - 4\kappa}$.]

This means that $A \in \mathbb{R}^{2 \times 2}$ has two distinct eigenvalues and therefore is diagonalizable. Write $A = V\Lambda V^{-1}$ with $\Lambda = \text{diag}(\lambda_+, \lambda_-)$. Then the ODE $\dot{\mathbf{v}} = A\mathbf{v}$ is

$$\dot{\mathbf{v}} = V\Lambda V^{-1}\mathbf{v}. \quad (39.2)$$

We will make this ODE “diagonal” with a clever change of variables. Multiply both sides by V^{-1} to get

$$V^{-1}\dot{\mathbf{v}} = \Lambda V^{-1}\mathbf{v}. \quad (39.3)$$

Now put

$$\mathbf{w} = V^{-1}\mathbf{v}. \quad (39.4)$$

By $\dot{\mathbf{w}}$ I mean the componentwise derivative of \mathbf{w} : if $\mathbf{w} = (w_1, w_2)$, then $\dot{\mathbf{w}} = (\dot{w}_1, \dot{w}_2)$. Then we get

$$\dot{\mathbf{w}} = V^{-1}\dot{\mathbf{v}}. \quad (39.5)$$

39.3 Problem (*). Check this. Assume

$$V^{-1} = \begin{bmatrix} a & c \\ b & d \end{bmatrix},$$

compute $V^{-1}\mathbf{v}$, take its componentwise derivatives, and show that's equal to $V^{-1}\dot{\mathbf{v}}$.

Finally, set

$$\mathbf{w}_0 = \begin{bmatrix} w_{01} \\ w_{02} \end{bmatrix} = V^{-1}\mathbf{v}_0. \quad (39.6)$$

Sorry for the weird subscripts. Put (39.2), (39.3), (39.4), (39.5), and (39.6) together to find

$$\begin{cases} \dot{\mathbf{w}} = \Lambda \mathbf{w} \\ \mathbf{w}(0) = \mathbf{w}_0. \end{cases}$$

Now write this out as a system:

$$\begin{cases} \dot{w}_1 = \lambda_+ w_1 \\ w_1(0) = w_{01} \end{cases} \quad \text{and} \quad \begin{cases} \dot{w}_2 = \lambda_- w_2 \\ w_2(0) = w_{02}. \end{cases}$$

I hope we agree that

$$w_1(t) = w_{01}e^{\lambda_+ t} \quad \text{and} \quad w_2(t) = w_{02}e^{\lambda_- t}.$$

Now let's vectorize:

$$\mathbf{w}(t) = \begin{bmatrix} w_1(t) \\ w_2(t) \end{bmatrix} = \begin{bmatrix} w_{01}e^{\lambda_+ t} \\ w_{02}e^{\lambda_- t} \end{bmatrix} = \begin{bmatrix} e^{\lambda_+ t} & 0 \\ 0 & e^{\lambda_- t} \end{bmatrix} \begin{bmatrix} w_{01} \\ w_{02} \end{bmatrix} = \begin{bmatrix} e^{\lambda_+ t} & 0 \\ 0 & e^{\lambda_- t} \end{bmatrix} \mathbf{w}_0.$$

Let's go back to our original unknown $\mathbf{v} = (x, y)$. Abbreviate

$$E(t) = \begin{bmatrix} e^{\lambda_+ t} & 0 \\ 0 & e^{\lambda_- t} \end{bmatrix}.$$

This says

$$\mathbf{w}(t) = E(t)\mathbf{w}_0,$$

and so we recover from (39.4) that

$$\mathbf{v}(t) = V\mathbf{w}(t) = VE(t)\mathbf{w}_0 = VE(t)V^{-1}\mathbf{v}_0.$$

If you chase through the multiplication $VE(t)V^{-1}\mathbf{v}_0$ and look at the first component, I claim this means that the solution $x(t)$ to our original problem (39.1) is a linear combination of $e^{\lambda_+ t}$ and $e^{\lambda_- t}$. The coefficients would depend, unsurprisingly, on b , κ , x_0 , and y_0 .

39.4 Problem (★). Here are two more ways to see this.

(i) Guess that $x(t) = e^{\lambda t}$ solves (39.1), plug this guess into the ODE, and conclude that $\lambda^2 + b\lambda + \kappa = 0$.

(ii) All solutions to exponential growth $\dot{x} = rx$ are multiples of e^{rt} . Guess that $\mathbf{v}(t) = e^{\lambda t}\mathbf{x}$ solves $\dot{\mathbf{v}} = A\mathbf{v}$ for some $\mathbf{x} \in \mathbb{R}^2$, compute $\dot{\mathbf{v}}(t) = \lambda\mathbf{v}(t)$, and conclude $A\mathbf{x} = \lambda\mathbf{x}$.

Also, this is the *only* solution to (39.1): if x solves (39.1), then $(x, \dot{x}) = (x, y) = \mathbf{v}$ solves (39.2), and that has the solution above. But *having a formula for something is not the same as understanding that thing*. What is the solution *doing* over long times? If you believe the claim of the previous paragraph, since $x(t) = c_1e^{\lambda_+ t} + c_2e^{\lambda_- t}$ for some $c_1, c_2 \in \mathbb{R}$, and since $\lambda_{\pm} < 0$, we have $\lim_{t \rightarrow \infty} x(t) = 0$. This is physically reasonable: x is the displacement of an oscillator from its equilibrium position, the oscillator is slowed by friction, and no other forces are contributing to the oscillator's motion. Eventually it should barely be moving.

Content from Strang's ILA 6E. Page 276 discusses the question of *stability*: that solutions to $\ddot{x} + b\dot{x} + \kappa x = 0$ converge to the solution $x = 0$ as $t \rightarrow \infty$, at least for suitably chosen b and κ . This page, and all of Section 6.5, operate in the more general framework in which A from $\dot{\mathbf{v}} = A\mathbf{v}$ has complex, nonreal eigenvalues. I've skipped that here just to avoid dealing with the complex exponential. Such a framework shows that solutions to $\ddot{x} + b\dot{x} + \kappa x = 0$ always vanish as $t \rightarrow \infty$ when $\kappa \geq 0$ and when $b > 0$ (i.e., when friction is present), not just in the more special case of $b^2 - 4\kappa > 0$. If you want to think more about differential equations, read p. 285.

But we didn't need that formula for x to figure this out. Just the formula $\mathbf{v}(t) = VE(t)V^{-1}\mathbf{v}_0$ tells us that. The entries of $E(t)$ all approach 0 as $t \rightarrow \infty$, so the product $VE(t)V^{-1}\mathbf{v}_0$ should approach $\mathbf{0}_2$ as $t \rightarrow \infty$. Then both $x(t)$ and $\dot{x}(t) = y(t)$ should go to 0 as $t \rightarrow \infty$.

You could get all of this using methods from an ODE course just for the second-order linear ODE (39.1). I claim linear algebra gets you those answers more quickly and more transparently, though.

Day 40: Wednesday, April 16.

We are going to switch focus a bit and think about determinants. We've said several times that the **DETERMINANT** of

$$A = \begin{bmatrix} a & c \\ b & d \end{bmatrix} \in \mathbb{C}^{2 \times 2}$$

is

$$\det(A) = ad - bc.$$

The chief virtue of the determinant is that A is invertible if and only if $\det(A) \neq 0$, and also $\det(A)$ shows up in the formula for A^{-1} . Our goal now is to extend this notion of determinant to an arbitrary square matrix with the goal that, still, the matrix is invertible if and only if its determinant is nonzero.

Determinants were historically somewhat more important than they are today. Carl Meyer observes the following in *Applied Linear Algebra and Matrix Analysis* (pp. 459–460):

[M]uch was written [on determinants] between 1750 and 1900. During this era, determinants became the major tool used to analyze and solve linear systems, while the theory of matrices remained relatively undeveloped. But mathematics, like a river, is everchanging in its course, and major branches can dry up to become minor tributaries while small trickling brooks can develop into raging torrents.

Nonetheless, determinants still have some perennially interesting theoretical properties, and every student of linear algebra should be able to compute the determinant of at least a 3×3 matrix without their brains totally melting. There are many ways to express a formula for the determinant, and all of them are a bit painful, and none of those formulas tell you why you'd think of that formula in the first place.

I firmly believe that *what things do defines what things are*. Look at the 2×2 determinant formula above. I think you'll believe immediately that

$$\det(I_2) = 1.$$

Next, suppose we interchange rows:

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

With $P_{12} = P_{21}$ as a permutation matrix (that's not I_2), we have

$$\det(P_{12}A) = -\det(A).$$

Finally, and this is a little less obvious, suppose that we scale a row by the same $t \in \mathbb{C}$:

$$\det \begin{pmatrix} ta & tc \\ b & d \end{pmatrix} = tad - tcb = t(ad - bc) = t \det(A).$$

Or suppose that a row is really a sum of two different rows:

$$\det \begin{pmatrix} (a_1 + a_2) & (c_1 + c_2) \\ b & d \end{pmatrix} = \det \begin{pmatrix} a_1 & c_1 \\ b & d \end{pmatrix} + \det \begin{pmatrix} a_2 & c_2 \\ b & d \end{pmatrix}.$$

40.1 Problem (!). Okay, that last identity was too annoying for me to do—you can do it.

What this says is that the determinant is “linear in each row”—it respects linear combinations of rows (just like matrix multiplication respects linear combinations). These three properties really *characterize* the determinant. By that I mean that if f is a function from $\mathbb{C}^{2 \times 2}$ to \mathbb{C} , and by function I just mean a rule that associates each $A \in \mathbb{C}^{2 \times 2}$ with a unique $f(A) \in \mathbb{C}$, and if f satisfies those properties, then $f(A) = \det(A)$.

40.2 Theorem. Suppose that f is a function from $\mathbb{C}^{2 \times 2}$ to \mathbb{C} such that the following hold.

(i) $f(I_2) = 1$.

(ii) $f(P_{12}A) = -f(A)$ for all $A \in \mathbb{C}^{2 \times 2}$.

(iii) If

$$A_1 = \begin{bmatrix} a_1 & c_1 \\ b & d \end{bmatrix} \quad \text{and} \quad A_2 = \begin{bmatrix} a_2 & c_2 \\ b & d \end{bmatrix}$$

and $t \in \mathbb{R}$, then

$$f \left(\begin{bmatrix} (a_1 + a_2) & (c_1 + c_2) \\ b & d \end{bmatrix} \right) = f(A_1) + f(A_2)$$

and

$$f \left(\begin{bmatrix} ta_1 & tc_1 \\ b & d \end{bmatrix} \right) = tf(A_1).$$

Then $f(A) = \det(A)$ for all $A \in \mathbb{C}^{2 \times 2}$.

Content from Strang's ILA 6E. Page 199 talks about sign reversal for determinants when interchanging rows. One consequence is that if a matrix has two identical rows, its determinant is 0, since $A = P_{ij}A$, and thus $\det(A) = -\det(A)$. In the “3 by 3 Determinants,” start with $\det(I_3) = +1$, then start interchanging rows to flip between ± 1 . On p. 200, if you believe that the determinant of a diagonal matrix is the product of its entries, do more row interchanges to get the six determinants there. Finally, start with the full 3×3 matrix and use linearity in the rows to get equation (1). I think that is a bit hard.

These properties open the way to constructing determinants on $\mathbb{C}^{n \times n}$. We want a function f from $\mathbb{C}^{n \times n}$ to \mathbb{C} such that $f(I_2) = 1$, $f(P_{ij}A) = -f(A)$ for any permutation matrix P_{ij} that interchanges two rows (so $P_{ij} \neq I_n$), and such that f is “linear in each row.” I claim that there can be at most one function f that does these three things—if there are two such functions, call them f and g , then we can prove $f(A) = g(A)$ for all $A \in \mathbb{C}^{n \times n}$. And I claim that these three properties lead to a specific formula for f . (I guess that’s kind of uniqueness, too, right?)

Content from Strang's ILA 6E. These three properties appear in equation (5) on p. 206. Various consequences of these properties are sketched on pp. 205–206.

Best of all, I claim that these three properties mean that A is invertible if and only if $f(A) \neq 0$. Here’s a sketch of how you get that last key result. Gaussian elimination lets us write $EA = U$, where U is upper-triangular and E is a product of elimination and permutation matrices. This is Gaussian elimination, not Gauss–Jordan elimination, and U may not be $\text{rref}(A)$, so there are no scaling matrices in E . You first show that f is “multiplicative”: $f(EA) = f(E)f(A)$ for any E , not just a product of elimination and permutation matrices. Then you show that $|f(E)| = 1$. Really, $f(P_{ij}) = -1$ for any permutation matrix P_{ij} that interchanges two rows, and $f(E_{ij}) = 1$ for any elimination matrix. Then rewrite $f(E)$ as a product of f evaluated at the factors of E . Finally, you show that $f(U)$ is the product of the diagonal entries of U whenever U is upper-triangular. All of this takes some work.

Put it all together to show that

$$|f(A)| = 1 \cdot |f(A)| = |f(E)| \cdot |f(A)| = |f(E)f(A)| = |f(EA)| = |f(U)|.$$

Since $f(A) = 0$ if and only if $|f(A)| = 0$, we have $f(A) = 0$ if and only if $f(U) = 0$, which happens if and only if one of the diagonal entries of U is 0. (Because a product of numbers is 0 if and only if at least one of them is 0. Is that true for matrices?) And A is not invertible if and only if at least one diagonal entry of U is 0, right? So A is invertible if and only if $f(A) \neq 0$.

You can also use those three properties of f to back out, with effort, a formula for f . It turns out that there are many options for this formula; here is just one, and just for the 3×3

case:

$$\det \begin{pmatrix} a & p & x \\ b & q & y \\ c & r & z \end{pmatrix} = a \det \begin{pmatrix} q & y \\ r & z \end{pmatrix} - b \det \begin{pmatrix} p & x \\ r & z \end{pmatrix} + c \det \begin{pmatrix} p & x \\ q & y \end{pmatrix}.$$

Bizarre, right?

One way to remember this formula is that we multiply by entries of the first column, alternate the sign (+, −, +), and, when multiplying by the $(i, 1)$ -entry, multiply against the determinant of the matrix that results when you remove column 1 and row i from A :

$$\begin{aligned} \det \begin{pmatrix} a & p & x \\ b & q & y \\ c & r & z \end{pmatrix} &= a \det \begin{pmatrix} q & y \\ r & z \end{pmatrix} && \begin{bmatrix} a & p & x \\ b & q & y \\ c & r & z \end{bmatrix} \\ &- b \det \begin{pmatrix} p & x \\ r & z \end{pmatrix} && \begin{bmatrix} a & p & x \\ b & q & y \\ c & r & z \end{bmatrix} \\ &+ c \det \begin{pmatrix} p & x \\ q & y \end{pmatrix} && \begin{bmatrix} a & p & x \\ b & q & y \\ c & r & z \end{bmatrix}. \end{aligned}$$

The point is that you define the determinant of an $n \times n$ matrix *recursively*, in terms of determinants of $(n - 1) \times (n - 1)$ matrices, down to the simplest case of the 2×2 . I don't think starting with this formula would ever help you show that A is invertible precisely when $\det(A) \neq 0$. *What things do defines what things are.*

40.3 Example. We compute

$$\begin{aligned} \det \begin{pmatrix} 1 & 0 & 0 \\ 4 & 2 & 0 \\ 5 & 6 & 3 \end{pmatrix} &= 1 \det \begin{pmatrix} 2 & 0 \\ 6 & 3 \end{pmatrix} && \begin{bmatrix} 1 & 0 & 0 \\ 4 & 2 & 0 \\ 5 & 6 & 3 \end{bmatrix} \\ &- 4 \det \begin{pmatrix} 0 & 0 \\ 6 & 3 \end{pmatrix} && \begin{bmatrix} 1 & 0 & 0 \\ 4 & 2 & 0 \\ 5 & 6 & 3 \end{bmatrix} \\ &+ 5 \det \begin{pmatrix} 0 & 0 \\ 2 & 0 \end{pmatrix} && \begin{bmatrix} 1 & 0 & 0 \\ 4 & 2 & 0 \\ 5 & 6 & 3 \end{bmatrix} \\ &= (1 \cdot 6) - (4 \cdot 0) + (5 \cdot 0) \end{aligned}$$

$$= 6.$$

And 6 is the product of the diagonal entries of this lower-triangular matrix, and $6 \neq 0$, and this lower-triangular matrix with nonzero diagonal entries is invertible.

Content from Strang's ILA 6E. Equation (2) on p. 201 does this recursive definition of the determinant for a 4×4 matrix. Equation (5) generalizes this formula to a “cofactor” expansion in which you can recursively expand the determinant along any row or down any column.

Last thing: an application to eigenvalues. We know that $\lambda \in \mathbb{C}$ is an eigenvalue of $A \in \mathbb{C}^{n \times n}$ if and only if A is not invertible. Assuming that we have the determinant defined on $\mathbb{C}^{n \times n}$, this means that λ is an eigenvalue of A if and only if $\det(A - \lambda I_n) = 0$. When $n = 2$, we know that $\det(A - \lambda I_2)$ is a quadratic polynomial.

40.4 Problem (★). Show that

$$\det \left(\begin{bmatrix} 1 & 0 & 0 \\ 4 & 2 & 0 \\ 5 & 6 & 3 \end{bmatrix} - \lambda I_3 \right)$$

is a cubic polynomial and that its roots are $\lambda = 1, 2, 3$.

It turns out to be the case that $\det(A - \lambda I_n)$ is an n th degree polynomial—actually, it's monic with

$$\det(A - \lambda I_n) = \lambda^n + \text{a linear combination of lower powers of } \lambda (= 1, \lambda, \lambda^2, \dots, \lambda^{n-1}).$$

We call $\det(A - \lambda I_n)$ the **CHARACTERISTIC POLYNOMIAL** of A and $\det(A - \lambda I_n) = 0$ the **CHARACTERISTIC EQUATION** of A .

Here's a vague explanation of why this is true: you compute $\det(A - \lambda I_n)$ by taking a linear combination of the entries of column 1 of $A - \lambda I_n$ against determinants of the matrices formed by removing column 1 and row i (for $i = 1, \dots, n$) of $A - \lambda I_n$. Multiplying by the $(1, 1)$ -entry, which is $a_{11} - \lambda$, against the determinant of $A - \lambda I_n$ with row 1 and column 1 removed gets λ^n to show up. Every other term in the sum will have a prefactor from rows 2 through n of column 1, and that prefactor doesn't have a λ in it, so the other terms have at most λ^{n-1} in them. All we are doing is adding, subtracting, and multiplying, so the only way that λ could show up is as a nonnegative integer power (nothing weird like $\cos(\lambda)$ or e^λ in this arithmetic).

Here is why this matters for eigenvalues. It's a fact, called the fundamental theorem of algebra, that every polynomial has at least one root in \mathbb{C} (maybe not \mathbb{R}), and at most n . This gives the *existence* of eigenvalues. While it's nice for theoretical purposes, it's effectively useless computationally, as there are no formulas for the roots of a degree n polynomial when $n \geq 5$, and for $n \leq 4$ only the quadratic formula is really tractable. Numerically computing the roots of polynomials is hard, and it's actually better to take a polynomial, find a matrix

whose characteristic polynomial is that polynomial (this matrix is called the “companion matrix” to the polynomial that you started with), and then find the eigenvalues of that matrix using some eigenvalue algorithm. Go take a numerical linear algebra class!

Content from Strang’s ILA 6E. Pages 220–221 discuss the characteristic polynomial for matrices of arbitrary size. Then read “Determinant and Trace” on p. 222. See p. 265 for a matrix with four distinct eigenvalues, some of which are purely imaginary. Last, go to p. 284 for a sketch of one of those numerical eigenvalue algorithms.

Day 41: Friday, April 18.

Here is another application of matrices and diagonalization that has nothing to do with $A\mathbf{x} = \mathbf{b}$ but some things to do with predicting the future. The ODE application did that, too, but in a “continuous” way: it gave results for every time t . Here we are going to consider “discretized” time in “chunks.”

The following very idealized toy problem will motivate everything. Suppose that a population is divided into sick and well people. At the end of each week, 90% of the well people stay well and 10% get sick, while 80% of the sick people become well and 20% of the sick people stay sick. Say that in a given week, the fraction of the population that is well is v_1 and the fraction that is sick is v_2 . So, $0 \leq v_1 \leq 1$, $0 \leq v_2 \leq 1$, and $v_1 + v_2 = 1$. Then at the end of that week, the fraction that is well is

$$.9v_1 + .8v_2,$$

while the fraction that is sick is

$$.1v_1 + .2v_2.$$

If we rewrite this in linear algebra language, the well and sick fractions are encoded as

$$\begin{bmatrix} .9v_1 + .8v_2 \\ .1v_1 + .2v_2 \end{bmatrix} = \begin{bmatrix} .9 & .8 \\ .1 & .2 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}.$$

Let

$$A = \begin{bmatrix} .9 & .8 \\ .1 & .2 \end{bmatrix}. \quad (41.1)$$

If at the start the well/sick fractions are $\mathbf{v} = (v_1, v_2)$, then after k weeks, the well/sick fractions are $A^k \mathbf{v}$. How does this matrix-vector product behave for k large? That will tell us the eventual fate of this population under this very idealized toy model.

Don’t let the decimals fool you: it is not too hard to calculate that the eigenvalues of A are 1 and $-.1$ with corresponding eigenvectors

$$\begin{bmatrix} 8 \\ 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 \\ -1 \end{bmatrix}.$$

We can therefore diagonalize A as $A = V\Lambda V^{-1}$, where $\Lambda = \text{diag}(1, -.1)$ and the columns of V are these eigenvectors. Then $A^k = V\Lambda^k V^{-1}$, and as $k \rightarrow \infty$ we have

$$\Lambda^k \rightarrow \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix},$$

since $(.1)^k \rightarrow 0$ as $k \rightarrow \infty$. Thus the well/sick fractions of the population tend to

$$V \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} V^{-1} \mathbf{v}, \quad (41.2)$$

and we could calculate this if we had to.

However, this formulaic result is not hugely illuminating, and I think it misses a lot of the deeper magic here. The first important thing to observe is that A is a matrix with nonnegative entries (actually, positive entries) whose columns all sum to 1. Then the rows of A^T all sum to 1. We have observed (Example 37.1, Problem 37.2) that such a matrix always has 1 as an eigenvalue. Turns out this guarantees that 1 is an eigenvalue of A .

41.1 Theorem. *Let $A \in \mathbb{C}^{n \times n}$. Then A and A^T have the same eigenvalues.*

Proof. Any $\lambda \in \mathbb{C}$ is an eigenvalue of A if and only if $A - \lambda I_n$ is not invertible. A matrix is invertible if and only if its transpose is invertible, so $A - \lambda I_n$ is not invertible if and only if $(A - \lambda I_n)^T$ is not invertible. And $(A - \lambda I_n)^T = A^T - \lambda I_n$, so $(A - \lambda I_n)^T$ is not invertible if and only if λ is an eigenvalue of A^T . ■

41.2 Problem (!). With A from (41.1), explain why A and A^T do not have the same eigenvectors corresponding to 1.

So, if $A \in \mathbb{C}^{n \times n}$ is a matrix whose columns sum to 1, an eigenvalue of A is 1. If we also assume that the entries of A are nonnegative real numbers, then we can get a better bound on the eigenvalues of A . To do this, we need to introduce a notion of size for complex numbers.

41.3 Definition. *The MODULUS of $x + iy \in \mathbb{C}$ is*

$$|x + iy| := \sqrt{x^2 + y^2}.$$

Basically the modulus satisfies all the familiar properties of absolute value, including the multiplicative identity

$$|zw| = |z||w|$$

and the TRIANGLE INEQUALITY

$$|z + w| \leq |z| + |w|.$$

41.4 Lemma. *If $A \in \mathbb{R}^{n \times n}$ has nonnegative entries, and if the entries of each column of A sum to 1, then any eigenvalue $\lambda \in \mathbb{C}$ of A satisfies $|\lambda| \leq 1$.*

Proof. Since A and A^T have the same eigenvalues, we prove that any eigenvalue λ of A^T satisfies $|\lambda| \leq 1$. We'll see this most transparently at the level of $n = 2$. The point is that

the rows of A^\top have nonnegative entries and sum to 1, so

$$A^\top = \begin{bmatrix} a & (1-a) \\ b & (1-b) \end{bmatrix},$$

where $0 \leq a \leq 1$ and $0 \leq b \leq 1$. Assume $A^\top \mathbf{v} = \lambda \mathbf{v}$.

The first row gives

$$av_1 + (1-a)v_2 = \lambda v_1.$$

Then

$$|\lambda v_1| = |av_1 + (1-a)v_2|.$$

On the left, the multiplicativity of the modulus implies

$$|\lambda v_1| = |\lambda| |v_1|,$$

while on the right, the triangle inequality implies

$$|av_1 + (1-a)v_2| \leq |av_1| + |(1-a)v_2| = |a||v_1| + |1-a||v_2| = a|v_1| + (1-a)|v_2|.$$

Now let

$$v_{\max} = \max\{|v_1|, |v_2|\}.$$

Then

$$|\lambda| |v_1| = |av_1 + (1-a)v_2| \leq a|v_1| + (1-a)|v_2| \leq av_{\max} + (1-a)v_{\max} = (a+1-a)v_{\max} = v_{\max}.$$

Exactly the same argument in the second row shows

$$|\lambda| |v_2| \leq v_{\max}.$$

Since either $|v_1| = v_{\max}$ or $|v_2| = v_{\max}$, we have

$$|\lambda| v_{\max} = v_{\max}.$$

And since $\mathbf{v} = (v_1, v_2)$ is an eigenvector, we have $\mathbf{v} \neq \mathbf{0}_2$, so $v_1 \neq 0$ or $v_2 \neq 0$, thus $|v_1| > 0$ or $|v_2| > 0$, and so $v_{\max} > 0$. Then we may divide to find $|\lambda| \leq 1$. ■

This result confirms what we observed about our A from (41.1): it has nonnegative entries, its columns sum to 1, and its eigenvalues are 1 and .1.

41.5 Problem (★). Here is another approach that almost proves Lemma 41.4. Let $A \in \mathbb{R}^{n \times n}$ have columns whose entries sum to 1. Let $\mathbf{1} = (1, \dots, 1)$, i.e., $\mathbf{1} \in \mathbb{C}^n$ is the vector whose entries are all 1.

(i) Explain why $A^\top \mathbf{1} = \mathbf{1}$.

(ii) Let $\mathbf{v} \in \mathbb{C}^n$. For $k = 1, 2, 3$, compute $A^k \mathbf{v} \cdot \mathbf{1}$. Use these results to conjecture the value of $A^k \mathbf{v} \cdot \mathbf{1}$ for any k .

(iii) Let λ be an eigenvalue for A with eigenvector \mathbf{v} . Compute $A^k \mathbf{v} \cdot \mathbf{1}$ in two ways and argue that $\lambda^k(\mathbf{v} \cdot \mathbf{1}) = \mathbf{v} \cdot \mathbf{1}$.

(iv) Now suppose $\lambda \neq 1$. Since $\mathbf{v} = \lambda^{-1}A\mathbf{v}$, compute $\mathbf{v} \cdot \mathbf{1} = \lambda^{-1}(\mathbf{v} \cdot \mathbf{1})$ and conclude

The last thing we'll do is improve on the behavior of $A^k \mathbf{v}$ with A from (41.1). The eigenvector corresponding to 1 that (I said) we found was $(8, 1)$, and the nice thing here is that it has nonnegative entries. We can rescale this eigenvector to $\mathbf{v}_1 = (8/9, 1/9)$, and you'll note $8 + 1 = 9$. Keep $\mathbf{v}_2 = (1, -1)$ as the eigenvector for $.1$ as before. Now let $\mathbf{u} = (u_1, u_2)$ be the well/sick fraction of the population, so $u_2 = 1 - u_1$. We know that \mathbf{v}_1 and \mathbf{v}_2 form a basis for \mathbb{R}^2 , so we can write $\mathbf{u} = c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2$ for some $c_1, c_2 \in \mathbb{R}$. But we can do better than that: $c_1 = 1$. I'll explain why in a moment.

With $c_1 = 1$, we have

$$A^k \mathbf{u} = A^k (\mathbf{v}_1 + c_2 \mathbf{v}_2) = A^k \mathbf{v}_1 + c_2 A^k \mathbf{v}_2 = \mathbf{v}_1 + c_2 (.1)^k \mathbf{v}_2 \rightarrow \mathbf{v}_1$$

as $k \rightarrow \infty$. So, no matter the starting fractions of well and sick people in the population, after many weeks, the well/sick fractions always settle down to $8/9$ and $1/9$, respectively. I think this is much more evocative than what plain old diagonalization got us in (41.2).

41.6 Problem (!). Now here's why $c_1 = 1$. Suppose $u \in \mathbb{R}$ and $\mathbf{u} = (u, 1 - u)$. With $\mathbf{v}_1 = (8/9, 1/9)$ and $\mathbf{v}_2 = (1, -1)$, prove that $\mathbf{u} - \mathbf{v}_1 \in \text{span}(\mathbf{v}_2)$.

Content from Strang's ILA 6E. Example 2 on p. 235 discusses powers of a Markov matrix. For an application of diagonalization and matrix powers to a (very mathy) situation, read about the "Fibonacci numbers" on pp. 236–237. Pages 238–239 generalize this to "iterative" equations $\mathbf{x}_{k+1} = A\mathbf{x}_k$, where you start with a vector \mathbf{x}_0 and then define subsequent vectors in this sequence by multiplying their immediate predecessor by A . These examples repeat the key idea: if $A = V\Lambda V^{-1}$, then $A^k = V\Lambda^k V^{-1}$. Pages 408–409 elaborate on Markov matrices with a "population" application to rental cars.

41.7 Problem (+). Here is the generalization of Problem 41.6. Let $\mathcal{V} = \{\mathbf{v} \in \mathbb{R}^n \mid \mathbf{v} \cdot \mathbf{1} = 0\}$.

(i) Prove that \mathcal{V} is a subspace of \mathbb{R}^n of dimension $n - 1$. [Hint: *do this for $n = 3$ first to see the pattern. Or think of \mathcal{V} as the null space of a matrix with very few rows or columns.*]

(ii) Now let $\mathbf{v}_2, \dots, \mathbf{v}_n$ be a basis for \mathcal{V} . Suppose that $\mathbf{v}, \mathbf{u} \in \mathbb{R}^n$ both satisfy $\mathbf{v} \cdot \mathbf{1} = \mathbf{u} \cdot \mathbf{1} = 1$. Show that $\mathbf{v} - \mathbf{u} \in \mathcal{V}$ and therefore $\mathbf{v} = \mathbf{u} + c_2 \mathbf{v}_2 + \dots + c_n \mathbf{v}_n$ for some $c_2, \dots, c_n \in \mathbb{R}$.

41.8 Problem (+). And here is the generalization of this example, with some unnecessary hypotheses. Let $A \in \mathbb{R}^{n \times n}$ satisfy the following.

- The entries of A are nonnegative.

- The entries in each column of A sum to 1.
- The eigenvalue 1 of A has an eigenvector with all nonnegative entries.
- A has $n - 1$ other independent eigenvectors $\mathbf{w}_2, \dots, \mathbf{w}_n$, each of which corresponds to an eigenvalue λ with $|\lambda| < 1$.

Suppose that $\mathbf{u} \in \mathbb{R}^n$ satisfies $\mathbf{u} \cdot \mathbf{1} = 1$ and let $A\mathbf{w}_1 = \mathbf{w}_1$ with $\mathbf{w}_1 \cdot \mathbf{1} = 1$ as well. Prove that $A^k \mathbf{u} \rightarrow \mathbf{w}_1$ as $k \rightarrow \infty$ via the following steps.

- Explain why 1 has an eigenvector \mathbf{w}_1 with $\mathbf{w}_1 \cdot \mathbf{1} = 1$.
- Suppose that $A\mathbf{x} = \lambda\mathbf{x}$ with $\lambda \neq 0$ and $\lambda \neq 1$. Show that $\mathbf{x} \cdot \mathbf{1} = (\mathbf{x} \cdot \mathbf{1})/\lambda$ and conclude $\mathbf{x} \cdot \mathbf{1} = 0$. [Hint: since $\lambda \neq 0$, $\mathbf{x} = \lambda^{-1}A\mathbf{x}$.]
- With \mathcal{V} from Problem , show that $A\mathbf{v} \in \mathcal{V}$ if $\mathbf{v} \in \mathcal{V}$.
- Explain why we can write both $\mathbf{u} = \mathbf{w}_1 + \mathbf{v}$ and $\mathbf{u} = x_1\mathbf{w}_1 + \tilde{\mathbf{v}} + \mathbf{z}$, where $\mathbf{v}, \tilde{\mathbf{v}} \in \mathcal{V}$, $\mathbf{z} \in \mathcal{N}(A)$, and $A^k \tilde{\mathbf{v}} \rightarrow \mathbf{0}_n$ as $k \rightarrow \infty$.
- Conclude $(1 - x_1)\mathbf{w}_1 \in \mathcal{V}$ and therefore $x_1 = 1$.

Day 42: Monday, April 21.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Unitary matrix, Hermitian (or self-adjoint) matrix, symmetric matrix, unitarily diagonalizable matrix, orthogonally diagonalizable matrix

I think by now you are convinced that having an “eigenbasis” for \mathbb{R}^n relative to a matrix $A \in \mathbb{R}^{n \times n}$ is nice, both theoretically and practically. That is, we’d like to have n linearly independent eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^n$ for A ; then we can diagonalize A . But we can do even better than this eigenbasis. What’s the best kind of basis (if you’re not thinking about eigenvalues)? Probably an *orthonormal* basis! If $\mathbf{v}_1, \dots, \mathbf{v}_n$ are orthonormal, and if each \mathbf{v}_j is an eigenvector for λ_j , then we have expansions like

$$\mathbf{v} = (\mathbf{v} \cdot \mathbf{v}_1)\mathbf{v}_1 + \dots + (\mathbf{v} \cdot \mathbf{v}_n)\mathbf{v}_n \quad \text{and so} \quad A\mathbf{v} = \lambda_1(\mathbf{v} \cdot \mathbf{v}_1)\mathbf{v}_1 + \dots + \lambda_n(\mathbf{v} \cdot \mathbf{v}_n)\mathbf{v}_n.$$

42.1 Problem (★). Suppose that $A \in \mathbb{R}^{n \times n}$ has n independent eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^n$. Why *doesn’t* the Gram–Schmidt procedure guarantee that we can produce n orthonormal *eigenvectors* from these n independent eigenvectors?

We need to be a bit careful here: eigenvectors may have complex, nonreal entries, and our prior notion of dot product just doesn’t cut it.

42.2 Example. Here we are still thinking that if $\mathbf{v}, \mathbf{w} \in \mathbb{C}^n$, then $\mathbf{v} \cdot \mathbf{w} = v_1 w_1 + \cdots + v_n w_n$. This is wrong: let $\mathbf{v} = (i, 1)$, so $\mathbf{v} \neq \mathbf{0}_2$, but

$$\mathbf{v} \cdot \mathbf{v} = i^2 + 1 = -1 + 1 = 0.$$

Thus $\mathbf{v} \neq \mathbf{0}_2$ but $\|\mathbf{v}\| = 0$. The “norm” is not behaving the way it should...

We need a refinement of our notion of dot product. And for that, we need a new operation on complex numbers.

42.3 Definition. The **COMPLEX CONJUGATE** of $x + iy \in \mathbb{C}$ is

$$\overline{x + iy} := x - iy.$$

42.4 Example. (i) $\overline{1 + 2i} = 1 - 2i$.

(ii) $\overline{2} = 2$.

(iii) $\overline{i} = -i$.

Content from Strang's ILA 6E. If you need a refresher on complex arithmetic, see p. 262.

42.5 Definition. The **DOT PRODUCT** of $\mathbf{v}, \mathbf{w} \in \mathbb{C}^n$ is

$$\mathbf{v} \cdot \mathbf{w} = v_1 \overline{w_1} + \cdots + v_n \overline{w_n},$$

and the **NORM** of $\mathbf{v} \in \mathbb{C}^n$ is (still)

$$\|\mathbf{v}\| = \sqrt{\mathbf{v} \cdot \mathbf{v}}.$$

Content from Strang's ILA 6E. This generalized dot product appears on p. 263.

We are using the same symbols for dot product and norm in \mathbb{C}^n as in \mathbb{R}^n ; you should always think of them as coming with the complex conjugate from now on. The upshot (tacitly assumed in the preceding definition) is that $\mathbf{v} \cdot \mathbf{v} \geq 0$ for all $\mathbf{v} \in \mathbb{C}^n$, and so the square root is defined. Here's why.

42.6 Lemma. Let $z = x + iy \in \mathbb{C}$. Then $z\overline{z} \geq 0$ with equality if and only if $x = y = 0$.

42.7 Problem (!). Prove it. Then prove that $\|\mathbf{v}\| = 0$ if and only if $\mathbf{v} = \mathbf{0}_n$.

All of the other familiar properties of dot product arithmetic from Problem 25.4 hold for the dot product on \mathbb{C}^n , and we won't belabor the changes here. There's one key difference,

though:

$$\mathbf{v} \cdot \mathbf{w} = \overline{\mathbf{w} \cdot \mathbf{v}}.$$

The dot product is no longer “commutative.” This also means that

$$\mathbf{v} \cdot (c\mathbf{w}) = \bar{c}(\mathbf{v} \cdot \mathbf{w}).$$

With this new definition of dot product, we still keep the definitions of orthogonal and orthonormal the same, and we retain the key fact that a list of orthogonal nonzero vectors are independent. However, we also need to think about how matrices interact with our new definition of the dot product. Before, the transpose was the key way that a matrix talked to the dot product: for $A \in \mathbb{R}^{m \times n}$, $\mathbf{v} \in \mathbb{R}^n$, and $\mathbf{w} \in \mathbb{R}^m$, we had the essential “popping” feature

$$A\mathbf{v} \cdot \mathbf{w} = \mathbf{v} \cdot A^T \mathbf{w}.$$

With the new definition of dot product, the transpose is not enough to make things pop.

42.8 Problem (!). Let

$$A = \begin{bmatrix} 1 & -i \\ i & 1 \end{bmatrix}, \quad \mathbf{v} = \mathbf{e}_1, \quad \text{and} \quad \mathbf{w} = \begin{bmatrix} i \\ 1 \end{bmatrix}.$$

Compute $A\mathbf{v} \cdot \mathbf{w}$ and $\mathbf{v} \cdot A^T \mathbf{w}$. What do you observe?

The right thing to do is to conjugate along with transposing.

42.9 Definition. The **CONJUGATE TRANSPOSE** of $A \in \mathbb{C}^{m \times n}$ is the matrix $A^* \in \mathbb{C}^{n \times m}$ whose (i, j) -entry is the complex conjugate of the (j, i) -entry of A . That is,

$$A_{ij}^* = \overline{A_{ji}}.$$

Content from Strang's ILA 6E. Page 263 presents the complex transpose. Strang uses the evocative notation $(\overline{A})^T$ to emphasize the two operations in forming A^* : conjugating and transposing. Sometimes people also write A^H instead of A^* to emphasize that the conjugate transpose is called the “Hermitian adjoint” of A .

42.10 Example. If

$$A = \begin{bmatrix} 1 & 2 \\ i & (3 + 4i) \end{bmatrix},$$

then

$$A^* = \begin{bmatrix} \overline{1} & \overline{i} \\ \overline{2} & \overline{3 + 4i} \end{bmatrix} = \begin{bmatrix} 1 & -i \\ 2 & 3 - 4i \end{bmatrix}.$$

The conjugate transpose retains all of the essential properties of the ordinary transpose, updated for the new notion of the dot product. Honestly, we could have proved all of these

from the start for the conjugate transpose—we could have run the entire course starting in \mathbb{C} , not in \mathbb{R} . Here's a summary.

42.11 Theorem. Let $A \in \mathbb{C}^{m \times n}$.

- (i) $A\mathbf{v} \cdot \mathbf{w} = \mathbf{v} \cdot A^*\mathbf{w}$ for all $\mathbf{v} \in \mathbb{C}^n$ and $\mathbf{w} \in \mathbb{C}^m$.
- (ii) $(A^*)^* = A$.
- (iii) Suppose $m = n$. Then A is invertible if and only if A^* is invertible, in which case $(A^*)^{-1} = (A^*)^{-1}$.
- (iv) If $B \in \mathbb{C}^{n \times p}$, then $(AB)^* = B^*A^*$.
- (v) $\mathbb{C}^n = \mathbf{N}(A) \oplus \mathbf{C}(A^*)$ and $\mathbb{C}^m = \mathbf{C}(A) \oplus \mathbf{N}(A^*)$. (Here I mean $\mathbb{C}^p = \mathcal{V} \oplus \mathcal{W}$ to say that each $\mathbf{x} \in \mathbb{C}^p$ can be written uniquely as $\mathbf{x} = \mathbf{v} + \mathbf{w}$ for $\mathbf{v} \in \mathcal{V}$ and $\mathbf{w} \in \mathcal{W}$, and also $\mathbf{v} \cdot \mathbf{w} = 0$.)
- (vi) $\mathbf{C}(A) = \mathbf{N}(A^*)^\perp$.
- (vii) If $m = n$, then $\lambda \in \mathbb{C}$ is an eigenvalue of A if and only if $\bar{\lambda} \in \mathbb{C}$ is an eigenvalue of A^* .

42.12 Problem (!). Redo the calculations in Problem 42.8 with A^* in place of A^\top .

42.13 Problem (+). This shouldn't be your highest priority right now, but proving the results in Theorem 42.11 might be a good refresher for you.

Content from Strang's ILA 6E. Page 264 discusses the fundamental subspaces for a matrix with complex entries.

Now we have the tools to get back to thinking about eigenstuff. Almost.

42.14 Definition. A matrix $Q \in \mathbb{C}^{n \times n}$ is **UNITARY** if its columns are orthonormal, equivalently, if $Q^*Q = I_n$, or if Q is invertible with $Q^{-1} = Q^*$.

42.15 Problem (!). Check that those three statements in the preceding definition are, in fact, equivalent. Then explain why an orthogonal matrix $Q \in \mathbb{R}^{n \times n}$ is unitary. (We aren't going to need nonsquare complex matrices whose columns are orthonormal, so we don't bother to extend the definition of unitary matrices to nonsquare ones.)

The dream is that a matrix $A \in \mathbb{C}^{n \times n}$ has n orthogonal eigenvectors. After rescaling by their norms, they become n orthonormal eigenvectors. Then A is diagonalizable (orthogonal eigenvectors are independent) with $A = V\Lambda V^{-1}$ as usual. But here we can take V to be

the matrix of orthonormal eigenvectors, so $V^{-1} = V^*$. Then $A = V\Lambda V^*$. Even better: not only is A diagonalizable, not only do we have an orthonormal basis for \mathbb{C}^n consisting of eigenvectors of A , we also don't have to manage a matrix inverse in the diagonalized form of A .

Here are more words.

42.16 Definition. (i) A matrix $A \in \mathbb{C}^{n \times n}$ is **UNITARILY DIAGONALIZABLE** if there is a unitary matrix $V \in \mathbb{C}^{n \times n}$ and a diagonal matrix $\Lambda \in \mathbb{C}^{n \times n}$ such that $A = V\Lambda V^*$.

(ii) If $A \in \mathbb{R}^{n \times n}$ has the form $A = V\Lambda V^T$ for some orthogonal $V \in \mathbb{R}^{n \times n}$ and diagonal $\Lambda \in \mathbb{R}^{n \times n}$, then A is **ORTHOGONALLY DIAGONALIZABLE**.

The goal is now to learn more about when the dream is a reality. When can we unitarily diagonalize $A \in \mathbb{C}^{n \times n}$? Fooling around, we might seize on the role of the conjugate transpose. If $A = V\Lambda V^*$, then maybe we could learn about A^* . We can:

$$A^* = (V\Lambda V^*)^* = (V^*)^* \Lambda^* V^* = V\Lambda^* V^*.$$

This looks almost exactly like A , except for the Λ^* .

What if $A = A^*$? Then

$$V\Lambda V^* = V\Lambda^* V^*.$$

Cancel the factors of V and V^* (they're invertible, after all) to get

$$\Lambda = \Lambda^*.$$

Here's what this looks like at the 2×2 level:

$$\begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} = \begin{bmatrix} \overline{\lambda_1} & 0 \\ 0 & \overline{\lambda_2} \end{bmatrix}.$$

More generally, $\Lambda = \Lambda^*$ says that any eigenvalue λ of A satisfies $\lambda = \overline{\lambda}$. This means that $\lambda \in \mathbb{R}$.

42.17 Problem (!). Show that if $x + iy = \overline{x + iy}$, then $y = 0$, so $x + iy = x \in \mathbb{R}$.

We need yet more words.

42.18 Definition. (i) A matrix $A \in \mathbb{C}^{n \times n}$ is **HERMITIAN** or **SELF-ADJOINT** if $A^* = A$. (The conjugate transpose is sometimes called the “Hermitian adjoint” of A .)

(ii) If $A \in \mathbb{R}^{n \times n}$, then A is **SYMMETRIC** if $A^T = A$.

So, the eigenvalues of a Hermitian, unitarily diagonalizable matrix are real.

42.19 Problem (★). Does the argument above still work if A is Hermitian and diagonalizable but not necessarily unitarily diagonalizable? (So $A = A^*$ and $A = V\Lambda V^{-1}$, but maybe $V^{-1} \neq V^*$.)

It turns out that assuming that A is unitarily diagonalizable is overkill.

42.20 Theorem. *The eigenvalues of any Hermitian matrix are real.*

Proof. We're assuming two things: $A \in \mathbb{C}^{n \times n}$ satisfies $A^* = A$, and there are $\lambda \in \mathbb{C}$ and $\mathbf{v} \in \mathbb{C}^n$ with $A\mathbf{v} = \lambda\mathbf{v}$ and $\mathbf{v} \neq \mathbf{0}_n$. (Okay, maybe that was three things.) We want to show $\bar{\lambda} = \lambda$.

This is a classical trick, and by “classical” I mean it’s something that I probably wouldn’t have thought of myself but that you should know how to do. We somehow want to use the fact that $A^* = A$ along with the equation $A\mathbf{v} = \lambda\mathbf{v}$. But A^* really shines when we introduce a dot product. I think the only natural thing to take a dot product with is \mathbf{v} itself. So, take the dot product of both sides of $A\mathbf{v} = \lambda\mathbf{v}$ with \mathbf{v} :

$$A\mathbf{v} \cdot \mathbf{v} = (\lambda\mathbf{v}) \cdot \mathbf{v}.$$

The left side is

$$A\mathbf{v} \cdot \mathbf{v} = \mathbf{v} \cdot A^*\mathbf{v} = \mathbf{v} \cdot A\mathbf{v} = \mathbf{v} \cdot (\lambda\mathbf{v}) = \bar{\lambda}(\mathbf{v} \cdot \mathbf{v}).$$

The right side is just

$$(\lambda\mathbf{v}) \cdot \mathbf{v} = \lambda(\mathbf{v} \cdot \mathbf{v}).$$

Equating things, we get

$$\lambda(\mathbf{v} \cdot \mathbf{v}) = \bar{\lambda}(\mathbf{v} \cdot \mathbf{v}).$$

Since $\mathbf{v} \neq \mathbf{0}_n$, we have $\mathbf{v} \cdot \mathbf{v} \neq 0$, so we can divide to find

$$\lambda = \bar{\lambda}.$$

Thus $\lambda \in \mathbb{R}$. ■

Part of the dream has been orthonormal eigenvectors. If the matrix is Hermitian, we don’t have to work too hard to get that.

42.21 Theorem. *Eigenvectors of a Hermitian matrix corresponding to distinct eigenvalues are orthogonal. More precisely, let $A \in \mathbb{C}^{n \times n}$ be Hermitian and let $\lambda_1, \lambda_2 \in \mathbb{C}$ be distinct eigenvalues of A with corresponding eigenvectors $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{C}^n$. (So $\lambda_1 \neq \lambda_2$, $A\mathbf{v}_1 = \lambda_1\mathbf{v}_1$, and $A\mathbf{v}_2 = \lambda_2\mathbf{v}_2$.) Then $\mathbf{v}_1 \cdot \mathbf{v}_2 = 0$.*

Proof. We want to learn about $\mathbf{v}_1 \cdot \mathbf{v}_2$, and we know about $A\mathbf{v}_1$ and $A\mathbf{v}_2$ and how A interacts with the dot product. The trick is to get A into the dot product. As in the previous proof, we might take the dot product of both sides of $A\mathbf{v}_1 = \lambda_1\mathbf{v}_1$ with \mathbf{v}_2 :

$$A\mathbf{v}_1 \cdot \mathbf{v}_2 = (\lambda_1\mathbf{v}_1) \cdot \mathbf{v}_2.$$

The left side is

$$A\mathbf{v}_1 \cdot \mathbf{v}_2 = \mathbf{v}_1 \cdot A^*\mathbf{v}_2 = \mathbf{v}_1 \cdot A\mathbf{v}_2 = \mathbf{v}_1 \cdot (\lambda_2\mathbf{v}_2) = \overline{\lambda_2}(\mathbf{v}_1 \cdot \mathbf{v}_2) = \lambda_2(\mathbf{v}_1 \cdot \mathbf{v}_2).$$

Here we are using the prior result that $\lambda_2 \in \mathbb{R}$. The right side is just

$$(\lambda_1\mathbf{v}_1) \cdot \mathbf{v}_2 = \lambda_1(\mathbf{v}_1 \cdot \mathbf{v}_2).$$

Equating things, we get

$$\lambda_2(\mathbf{v}_1 \cdot \mathbf{v}_2) = \lambda_1(\mathbf{v}_1 \cdot \mathbf{v}_2),$$

and a little algebra gives

$$(\lambda_1 - \lambda_2)(\mathbf{v}_1 \cdot \mathbf{v}_2) = 0.$$

Since $\lambda_1 \neq \lambda_2$, we have $\mathbf{v}_1 \cdot \mathbf{v}_2 = 0$. ■

Content from Strang's ILA 6E. Pages 246–247 give slightly different proofs of these two results for symmetric matrices (i.e., A has real entries, so $A^* = A^T = A$). Page 264 revisits these proofs for complex matrices (where you have to use $A^* = (\overline{A})^T$).

We already know that eigenvectors corresponding to distinct eigenvalues are independent; this result recovers that with more geometry in the special case of a Hermitian matrix. The point of this excursion into Hermitian matrices is to convince you that they are nice—that if the dream is n orthonormal eigenvectors, then Hermitian matrices seem close to doing that. We need one more tool to see just how close they are to the dream, and it involves a nice resurgence of an old friend: the upper-triangular matrix.

Day 43: Wednesday, April 23.

Vocabulary from today

You should memorize the definition of each term, phrase, or concept below and be able to provide a concrete example of each and a nonexample for those marked “N.”

Normal matrix

Before we do that, a recap: a matrix $A \in \mathbb{C}^{n \times n}$ is diagonalizable if we can write $A = V\Lambda V^{-1}$ with $V \in \mathbb{C}^{n \times n}$ invertible and $\Lambda \in \mathbb{C}^{n \times n}$ diagonal, in which case the columns of V are eigenvectors of A corresponding to the eigenvalues on the diagonal of Λ . This factorization helps us compute powers easily: $A^k = V\Lambda^k V^{-1}$. Even better is when A is unitarily diagonalizable and V is unitary: then $V^{-1} = V^*$.

A unitarily diagonalizable matrix is particularly nice from the point of view of computation and structure. If $A = V\Lambda V^*$, then

$$A\mathbf{v} = \lambda_1(\mathbf{v} \cdot \mathbf{v}_1)\mathbf{v}_1 + \cdots + \lambda_n(\mathbf{v} \cdot \mathbf{v}_n)\mathbf{v}_n.$$

This “separates the action” of A along n eigenvectors and helps us see how each eigenvalue-eigenvector pair contributes to the total behavior of A . Additionally, doing the matrix

multiplication in $V\Lambda V^*$ and thinking of each column of V or V^* as a matrix in $\mathbb{C}^{n \times 1}$, we have

$$A = \lambda_1 \mathbf{v}_1 \mathbf{v}_1^* + \cdots + \lambda_n \mathbf{v}_n \mathbf{v}_n^*.$$

Remember that each “outer product” $\mathbf{v}\mathbf{v}^* \in \mathbb{C}^{n \times n}$ is a rank 1 matrix. We have therefore decomposed A as a linear combination of rank 1 matrices.

So, when is a matrix unitarily diagonalizable? Maybe working backwards will help. If $A = V\Lambda V^*$, what insights does this give us into the behavior of A ? We know $A^* = V\Lambda^* V^*$, which looks so much like A , and we know that V and V^* cancel each other out. This might motivate us to compute

$$AA^* = (V\Lambda V^*)(V\Lambda^* V^*) = V\Lambda(V^*V)\Lambda^* V^* = V\Lambda\Lambda^* V^*.$$

By the way, multiplying $\Lambda\Lambda^*$ is easy as both factors are diagonal.

43.1 Problem (!). Check that if $A = V\Lambda V^*$, then $AA^* = A^*A$. [Hint: *why do Λ and Λ^* commute?*]

We conclude that if A is unitarily diagonalizable, then $AA^* = A^*A$. A matrix that commutes with its conjugate transpose has a special name.

43.2 Definition. A matrix $A \in \mathbb{C}^{n \times n}$ is **NORMAL** if $AA^* = A^*A$.

43.3 Problem (★). (i) Show that every Hermitian matrix is normal.

(ii) Is every normal matrix Hermitian?

So, every unitarily diagonalizable matrix is normal. It also turns out that every normal matrix is unitarily diagonalizable. This is great! It’s much easier to check if a matrix is normal—check if the matrix and its conjugate transpose commute—than it is to show that the matrix has n orthonormal eigenvectors (which would mean it’s unitarily diagonalizable).

We can prove this by appealing to yet another result about matrix factorizations, which this time holds for *all* matrices.

43.4 Theorem (Schur). Any matrix can be **UNITARILY TRIANGULARIZED**: if $A \in \mathbb{C}^{n \times n}$, then there are a unitary matrix $V \in \mathbb{C}^{n \times n}$ and an upper-triangular matrix $T \in \mathbb{C}^{n \times n}$ such that

$$A = VTV^*. \quad (43.1)$$

Proof. This requires induction on n . We prove only the case $n = 2$.

Let $A \in \mathbb{C}^{2 \times 2}$ and let $\lambda \in \mathbb{C}$ be an eigenvalue of A with eigenvector $\mathbf{v}_1 \in \mathbb{C}^2$; we assume $\|\mathbf{v}_1\| = 1$. Then $\mathbf{v}_1 \neq \mathbf{0}_2$, so we can find $\mathbf{w}_2 \in \mathbb{C}^2$ such that $\mathbf{v}_1, \mathbf{w}_2$ is a basis for \mathbb{C}^2 . (Why can we find \mathbf{w}_2 ? This is the “exhaustion” argument from Theorem 24.5.) Use Gram–Schmidt to convert this list to an orthonormal basis $\mathbf{v}_1, \mathbf{v}_2$ for \mathbb{C}^2 ; we keep the first vector \mathbf{v}_1 the same, since it has length 1 and that’s how Gram–Schmidt works.

Now put $V = [\mathbf{v}_1 \ \mathbf{v}_2]$. We're going to show that $AV = VT$ for some upper-triangular matrix T . We have

$$AV = [A\mathbf{v}_1 \ A\mathbf{v}_2] = [\lambda\mathbf{v}_1 \ A\mathbf{v}_2].$$

What's going on with $A\mathbf{v}_2$? There's not much that we can say other than $A\mathbf{v}_2 = c_1\mathbf{v}_1 + c_2\mathbf{v}_2$ for some $c_1, c_2 \in \mathbb{C}$, since $\mathbf{v}_1, \mathbf{v}_2$ are a basis for \mathbb{C}^2 . (Actually, $c_1 = A\mathbf{v}_2 \cdot \mathbf{v}_1$ and $c_2 = A\mathbf{v}_2 \cdot \mathbf{v}_2$ since $\mathbf{v}_1, \mathbf{v}_2$ are an orthonormal basis for \mathbb{C}^2 , but that's not hugely helpful here.) Now think about it:

$$\lambda\mathbf{v}_1 = \lambda\mathbf{v}_1 + 0\mathbf{v}_2 = [\mathbf{v}_1 \ \mathbf{v}_2] \begin{bmatrix} \lambda \\ 0 \end{bmatrix} = V \begin{bmatrix} \lambda \\ 0 \end{bmatrix}$$

and

$$A\mathbf{v}_2 = c_1\mathbf{v}_1 + c_2\mathbf{v}_2 = [\mathbf{v}_1 \ \mathbf{v}_2] \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} = V \begin{bmatrix} c_1 \\ c_2 \end{bmatrix}.$$

Put it together:

$$AV = [\lambda\mathbf{v}_1 \ A\mathbf{v}_2] = V \begin{bmatrix} \lambda & c_1 \\ 0 & c_2 \end{bmatrix}.$$

That's our T :

$$T = \begin{bmatrix} \lambda & c_1 \\ 0 & c_2 \end{bmatrix}. \quad \blacksquare$$

43.5 Problem (*). Schur's theorem didn't say anything about the eigenvalues of A , but they're right there.

(i) Matrices $A, B \in \mathbb{C}^{n \times n}$ are **SIMILAR** if there is an invertible matrix $E \in \mathbb{C}^{n \times n}$ such that $A = EBE^{-1}$. Prove that similar matrices have the same eigenvalues.

(ii) Do similar matrices necessarily have the same eigenvectors? (That is, if A and B are similar with common eigenvalue λ , and \mathbf{v} is an eigenvector of A corresponding to λ , is \mathbf{v} also an eigenvector of B corresponding to λ ?)

(iii) How do the eigenvalues of A appear in the Schur factorization (43.1)?

(iv) In the Schur factorization (43.1), are the columns of V always eigenvectors of A ?

Content from Strang's ILA 6E. Go back to pp. 235–236 and read about similar matrices. The website has a proof of Schur's theorem that works for a *real* matrix A and gets a *real* factor V :

https://math.mit.edu/gs/linearalgebra/ila6/lafe_schur03.pdf.

Now we show that if A is normal, then A is unitarily diagonalizable. We know that since A is normal, we have $AA^* = A^*A$. And Schur tells us $A = VTV^*$ for some unitary V and upper-triangular T . Put it together:

$$(VTV^*)(VTV^*)^* = (VTV^*)^*(VTV^*).$$

Compute:

$$VT(V^*V)T^*V^* = VT^*(V^*V)TV^*.$$

And compute again:

$$VTT^*V^* = VT^*TV^*.$$

And use the invertibility of V and V^* to get

$$TT^* = T^*T.$$

This is enough to show that T is diagonal, not just upper-triangular. (In words, a normal upper-triangular matrix is diagonal.)

43.6 Problem (★). This is not so hard to check in the 2×2 case but requires a thoughtful analysis of the individual entries (use dot products) of TT^* and T^*T in the general case. Let

$$T = \begin{bmatrix} a & c \\ 0 & d \end{bmatrix}.$$

Compute TT^* and T^*T , compare entries, and conclude $c = 0$, so T is diagonal.

So, when A is normal, the upper-triangular factor T from the Schur factorization really is diagonal: with $T = \Lambda$, we have $A = V\Lambda V^*$. Thus A is unitarily diagonalizable. This characterization of unitarily diagonalizable matrices has a special name.

43.7 Theorem (Spectral theorem). *A matrix is unitarily diagonalizable if and only if it is normal.*

In particular, all Hermitian matrices are unitarily diagonalizable.

Content from Strang's ILA 6E. Page 258 outlines the spectral theorem with an emphasis on real matrices.

We have shown that every Hermitian matrix has real eigenvalues. There are still lots of possibilities for real numbers. In our ODE application, it was helpful that the eigenvalues were negative real numbers (that's what made the displacement of the oscillator slow down to 0 as time went on—more exactly, it was the negativity of the *real* part of the eigenvalues. . .). If we want the matrix to be invertible, we better not have 0 as an eigenvalue.

What guarantees that the eigenvalues of a Hermitian matrix are nonnegative or maybe positive? Once again, we can learn some stuff by working backwards. Suppose that $A \in \mathbb{C}^{n \times n}$ is Hermitian with nonnegative eigenvalues. The spectral theorem says that A is unitarily diagonalizable, so $A = V\Lambda V^*$, and the diagonal entries of Λ are nonnegative.

Here's the trick: in the past, we learned stuff about Hermitian matrices by working with dot products of the form $A\mathbf{v} \cdot \mathbf{v}$. Let's do that again:

$$A\mathbf{v} \cdot \mathbf{v} = V\Lambda V^*\mathbf{v} \cdot \mathbf{v} = \Lambda V^*\mathbf{v} \cdot V^*\mathbf{v}.$$

Now put $\mathbf{w} = V^*\mathbf{v}$. I claim that $\Lambda\mathbf{w} \cdot \mathbf{w} \geq 0$.

43.8 Problem (★). Show it. [Hint: *do it for the 2×2 case first—think about what the product $\Lambda \mathbf{w}$ is when Λ is diagonal, then what the dot product $\Lambda \mathbf{w} \cdot \mathbf{w}$ is, and then why the nonnegativity of the diagonal entries of Λ matters.*]

Here is what we have shown: if A is Hermitian with nonnegative eigenvalues, then $A\mathbf{v} \cdot \mathbf{v} \geq 0$ for all $\mathbf{v} \in \mathbb{C}^n$. It turns out that the reverse is true. And everything is true if you replace “nonnegative” with “positive” and \geq with $>$.