Excerpts from

The Dark Art of Linear Algebra

An Intuitive Geometric Approach

The whole book is available as a paperback at Amazon.com and as a pdf at BraverNewMath.com

Seth Braver

South Puget Sound Community College

VECTOR VECTORUM BOOKS

Vector Vectorum Books

Olympia, WA

Front and back cover illustrations created by Dall-E 2
(which features much linear algebra under the hood)

ISBN-13: 9798988140207

Library of Congress Control Number: 2023906846

Contents

Only these sections are contained in "EXCERPTS from DALA"

Preface for Teachers

Of making many books there is no end;
and much study is a weariness of the flesh.
- Ecclesiastes 12:12

Linear algebra is still new. Its first devoted textbook for undergraduates, Paul Halmos's *Finite-Dimensional Vector Spaces,* was published only in 1942. As an undergraduate in the late 1990s, I was assigned a newish textbook that's now considered a classic in its own right: *Linear Algebra Done Right* by Sheldon Axler, one of Halmos's "grandstudents".* As fate would have it, I went on to become one of Axler's grandstudents. *The Dark Art of Linear Algebra* thus continues a family tradition. But the line of descent ends here. Childless and alone, I roam the manor's halls, clutching my candelabra, conscious of the ancestral portraits gazing at me from the walls.†

We know how to write calculus textbooks. We've known this for centuries, and as a result, almost all calculus texts now tend towards the same form. Naturally, there are variations in authors' expository skill, level of rigor, interest in applications, and so forth. But pick up a random calculus textbook and we know what to expect inside: not just the topics, but the order in which they will appear, and even which hoary old homework problems will accompany them.‡

Linear algebra is different. We don't yet know the right way to teach it to an audience of beginners, and I'm confident that neither of the two most common approaches today will stand the test of time. The traditional abstract approach, which starts from the vector space axioms, is, for all its mathematical elegance, opaque to beginners. On the other hand, the current fashion for introducing linear algebra via systems of linear equations is so mind-numbingly dull as to constitute a crime against art.

In *The Dark Art of Linear Algebra,* I take a third approach, emphasizing geometry and intuition and delaying systems of linear equations until after students have developed a strong grasp of linear maps (and the matrices that represent them) and have mastered linear algebra's core vocabulary: span, linear independence, basis, subspace, and so forth. This third way is rare in textbooks, but a series of beautiful videos develops its outlines: Grant Sanderson's "The Essence of Linear Algebra", available on his YouTube channel 3Blue1Brown. Sanderson's videos make excellent complementary material for students reading *The Dark Art of Linear Algebra*, since our approaches are so similar in spirit.

Market forces have produced generations of grotesquely bloated textbooks, but self-publication has mercifully freed me from them. I've had the luxury of writing a textbook that – miracle of miracles – students can read from cover to cover while taking a single college course. Completing a textbook is a good feeling. We should not deny our students this small pleasure.

Should you spot any typos or errors, please let me know. I can't offer you extra credit (as I do for my students), but I can offer you thanks, and I will praise your name as I ruthlessly expunge any imperfections you've identified. Should you use this book in a class, I'd especially appreciate hearing from you.

---

* Axler was Halmos's "grandstudent" in the sense that Axler's PhD advisor, Donald Sarason, was advised in *his* PhD by Halmos. Such mathematical begettings are fun to trace online at the Mathematics Genealogy Project. Karl Friedrich Gauss, it turns out, is one of my direct mathematical ancestors. Eleven intervening generations separate us.

† Forgive me, great-great-grandfather Halmos. My textbooks have illustrations, and I will sometimes sacrifice rigor for intuition. Forgive me, grandfather Axler, for I have used determinants. I honor you both. My book is meant to reach different audiences than either of yours do, and as such, it is animated by a somewhat different spirit, but I trust that it will extend our family's legacy of high quality linear algebraic exposition.

‡ My own calculus textbook, *Full Frontal Calculus*, is unusual in several respects (it is short, for instance, and uses infinitesimals), but even it adheres to the basic script – hoary old problems and all.

Preface for Students

> Just remember that where there's no linear
> there's no delineation. Try and stay focused.
>
> - The Thalidomide Kid, from Cormac McCarthy's *The Passenger*.

Calculus is the summit of the high school math curriculum. Few climb it, but everyone that knows it exists. But linear algebra? How many people have even heard of it? Is it like ordinary algebra? Is it harder than calculus? Does anyone actually use it in applications? What *is* this dark art?

Geometrically speaking, linear algebra is concerned only with flatness: lines, planes, and hyperplanes. This sounds limited (and indeed it is), but those very constraints are what make the subject fundamental. Compare trigonometry. Who needs a whole subject devoted to measuring triangles? Well, everyone does. Any polygon can be chopped up into triangles, so if you understand triangles, you understand polygons. Moreover, the basic trigonometric functions take on their own life, transcend their humble origins, and become central to all periodic phenomena. Similarly, who needs linear algebra? Everyone does. There's a sense, familiar from calculus, in which non-linear phenomena can be reduced (on an infinitesimal scale) to linear phenomena, so understanding the linear world helps us grasp the nonlinear world as well. Moreover, linear algebra's basic functions, linear *transformations* (and the matrices that represent them), take on their own lives, transcend their origins, and become indispensable tools throughout mathematics, science, engineering, computer science, statistics, and even economics. In today's world – especially its technological side – linear algebra is probably the single most frequently applied part of mathematics.

You should read *The Dark Art of Linear Algebra* slowly and carefully, with your pen and paper at hand. When I omit algebraic details, you should supply them. When I use a phrase such as "as you should verify", you should do so. Only *after* reading a section should you try to solve the problems at the end. When you encounter something you don't understand, mark the relevant passage and try to clear it up – first on your own, then by discussing it with your classmates or teacher.

Conceptual understanding is especially crucial in linear algebra – much more so even than in calculus. Unlike many calculus computations, those involved in linear algebra are straightforward, even childish. They can be tedious, to be sure, but there's nothing here like the difficulty of cracking a tough integral. You need to be able to do these linear algebraic computations, of course, but that's not what learning linear algebra is about. Your main job is to understand how linear algebra's many concepts fit together, which in turn will let you understand which computations to make, and why they are appropriate.

For supplementary material, I highly recommend Grant Sanderson's series of videos "The Essence of Linear Algebra", which you can find on his YouTube channel, 3Blue1Brown. Sanderson's beautiful animations bring linear transformations to life in ways that simply aren't possible on the printed page. The general approach he takes in his videos is similar to mine in this book, although he does not discuss linear algebra's computational aspects. I also recommend trawling the internet for examples of how linear algebra is applied in your fields of interest. Examples abound, and these will give you further impetus to learn linear algebra – even if you can't at first fully understand the applications.

But enough throat clearing. Let's begin.

Chapter 3

Linear Transformations and Matrices

Linear Maps and Their Matrices

> As Gregor Samsa awoke one morning from uneasy dreams, he found himself transformed in his bed into a gigantic insect.
> - Franz Kafka, "The Metamorphosis"

Calculus focuses on real-valued functions of real numbers; we visualize these functions as static graphs. Linear algebra focuses on vector-valued functions of vectors; we visualize these functions dynamically, as distortions of space itself. More specifically, linear algebra's functions are "grid transformations": they work by dragging the tips of some or all of the standard basis vectors to new locations, thus transforming the grid they generate. These special functions are called **linear transformations** (or linear *maps*).

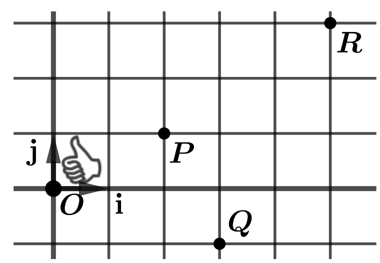

For example, the graph at right shows the standard graph-paper grid generated by $\mathbf{i}$ and $\mathbf{j}$. (Point O is the origin. Points P, Q, R and the enthusiastic right hand are all just stage props that will help us illustrate the effect of a linear map in a few moments.) Now let's go ahead and apply our first linear transformation: We will drag the tips of $\mathbf{i}$ and $\mathbf{j}$ to points $(2, 1)$ and $(1, 2)$ respectively. When we transform $\mathbf{i}$ and $\mathbf{j}$ this way, we transform the entire grid that they generate. The result is shown below, in the second figure.

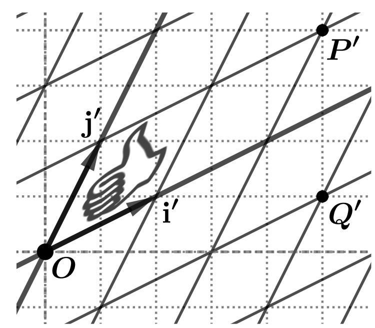

The original grid of squares (displayed now as a ghostly background) has been replaced by a grid of parallelograms generated by $\mathbf{i}'$ and $\mathbf{j}'$, the names I'll use for the transformed images of $\mathbf{i}$ and $\mathbf{j}$. Clearly, this linear transformation will distort any figure lying in the plane, such as the hand. But one crucial feature remains constant: Even though the linear map moves most of the points in the plane, it preserves their coordinates *relative to the grid*. For example, point P (first figure) corresponds to vector $2\mathbf{i} + \mathbf{j}$; its image, point P' (second figure) corresponds to vector $2\mathbf{i}' + \mathbf{j}'$. In both cases, we get to the point by following precisely the same "marching orders": Start from the origin, take two steps forward in the direction of the first "axis", and then take one step forward in the direction of the second "axis". We usually describe this by saying that *linear transformations preserve linear combinations of vectors*.

You should verify on the pictures above that the same phenomenon happens with points Q and Q'. Finally, although point R's image, R', is "offscreen", we know where it must lie: In the first figure, R corresponds to $5\mathbf{i} + 3\mathbf{j}$, so in the second, R' must "by linearity" correspond to $5\mathbf{i}' + 3\mathbf{j}'$. It would thus lie somewhere just beyond the right edge of this page.

Every linear transformation – not just the one above – preserves linear combinations of vectors. This is their defining feature, the very thing that makes linear transformations special.

To summarize, a linear map's geometric essence is a grid transformation induced by dragging the tips of the standard basis vectors to new places. A linear map's algebraic essence, which is a direct reflection of its geometric essence, is the preservation of linear combinations: Every vector is a certain linear combination of the standard basic vectors; when subjected to a linear map, its transformed image will be precisely that same linear combination of the *transformed* basis vectors. It follows that if we know how a linear map transforms the standard basis vectors, we know how it transforms *every* vector.

This means we can describe any linear transformation quite compactly. Rather than trying to illustrate its effect with figures (tedious in $\mathbb{R}^2$, difficult in $\mathbb{R}^3$, impossible in higher dimensions), we need only specify where the transformation sends the standard basis vectors. We encode this core information in a *matrix* whose i^{th} column simply records the components of the i^{th} standard basis vector's transformed image. We call the resulting matrix **the linear map's matrix** (relative to the standard basis).

Example 1. According to our definition above, the matrix

$$\begin{pmatrix} 1.5 & -1 \\ 0.5 & 1 \end{pmatrix}$$

represents the linear map that sends $\mathbf{i}$ to $\mathbf{i}' = 1.5\mathbf{i} + 0.5\mathbf{j}$ and sends $\mathbf{j}$ to $\mathbf{j}' = -\mathbf{i} + \mathbf{j}$. I've drawn the transformation at right. With or without this visual aid, we can easily determine where this linear transformation will send any given vector.

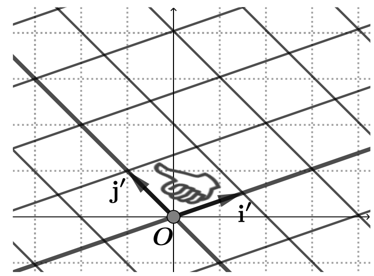

For instance, where will it send $2\mathbf{i} + 3\mathbf{j}$? Well, even without the visual aid, we know that linear maps preserve linear combinations, so this vector will be sent to $2\mathbf{i}' + 3\mathbf{j}'$, which is

$$2\begin{pmatrix} 1.5 \\ 0.5 \end{pmatrix} + 3\begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 4 \end{pmatrix}.$$

You can also confirm this by looking at the figure: Starting from the origin, two "$\mathbf{i}'$-steps" followed by three "$\mathbf{j}'$-steps" brings us to $(0, 4)$ on the original square grid. ♦

Mathematicians have invented a powerful algebraic shortcut: Instead of writing out the sentence

"The linear transformation whose matrix is $\begin{pmatrix} 1.5 & -1 \\ 0.5 & 1 \end{pmatrix}$ sends $\begin{pmatrix} 2 \\ 3 \end{pmatrix}$ to $\begin{pmatrix} 0 \\ 4 \end{pmatrix}$",

we express the same idea more concisely as an equation:

$$\begin{pmatrix} 1.5 & -1 \\ 0.5 & 1 \end{pmatrix}\begin{pmatrix} 2 \\ 3 \end{pmatrix} = \begin{pmatrix} 0 \\ 4 \end{pmatrix}.$$

We refer to the operation on the left side of this equation as **matrix-vector multiplication**. However, this "multiplication" should really make you think of *function evaluation*. This matrix (or better yet, the linear map that it represents) takes $2\mathbf{i} + \mathbf{3j}$ as its input, and produces $4\mathbf{j}$ as its output. Once you have internalized the idea that matrix-vector multiplication simply indicates where a given linear map sends a given vector, it becomes very easy to reduce the process to a few mindless turns of an algebraic crank.

Consider a perfectly general expression of matrix-vector multiplication:

$$\begin{pmatrix} a & c \\ b & d \end{pmatrix}\begin{pmatrix} v_1 \\ v_2 \end{pmatrix}.$$

What will the product be? Well, by definition, this matrix represents the linear map that sends $\mathbf{i}$ and $\mathbf{j}$ to

$$\mathbf{i}' = \begin{pmatrix} a \\ b \end{pmatrix} \quad \text{and} \quad \mathbf{j}' = \begin{pmatrix} c \\ d \end{pmatrix}.$$

Because linear transformations preserve linear combinations, it follows that our generic linear map must send the generic vector $v_1\mathbf{i} + v_2\mathbf{j}$ to the vector $v_1\mathbf{i}' + v_2\mathbf{j}'$, which is

$$v_1\begin{pmatrix} a \\ b \end{pmatrix} + v_2\begin{pmatrix} c \\ d \end{pmatrix}.$$

And thus we have our answer:

$$\begin{pmatrix} a & c \\ b & d \end{pmatrix}\begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = v_1\begin{pmatrix} a \\ b \end{pmatrix} + v_2\begin{pmatrix} c \\ d \end{pmatrix}.$$

This is a crucial result. Let us rephrase it and sanctify it in a box:

> A **matrix-vector product** is a *weighted sum of the matrix's columns*, where the i^{th} column's weight is the vector's i^{th} entry.

Example 2. $\begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}\begin{pmatrix} 3 \\ -1 \end{pmatrix} = 3\begin{pmatrix} 2 \\ 1 \end{pmatrix} - 1\begin{pmatrix} 1 \\ 2 \end{pmatrix} = \begin{pmatrix} \mathbf{5} \\ \mathbf{1} \end{pmatrix}.$

As you should verify, the matrix here represents the linear map whose "before and after" pictures appear on the first page of this section. Moreover, as you should also verify, this example's input and output vectors correspond to points Q and Q' on those figures. ♦

Every matrix-vector multiplication can be understood as mapping one vector to another via a linear transformation, a geometric operation whose effect is entirely determined by the columns of the matrix. This is linear algebra's central geometric idea, comparable in importance to calculus's key geometric ideas: all derivatives can be understood as slopes, and all definite integrals can be understood in terms of area. This isn't the only way to think of matrix-vector multiplication, but it is the fundamental one to which we'll return time and time again. It follows that this section may be the most important one in the entire book. To reinforce its ideas, let's consider a few more examples.

Example 3. Although grids and linear transformations are difficult to draw in $\mathbb{R}^3$, they are not hard to imagine. The standard basis vectors generate a grid of cubes (as the figure at right suggests), and we can induce a linear map by dragging the tips of some or all of the standard basis vectors to new locations, as I've done in the figure below, where the cubes have become parallelepipeds.

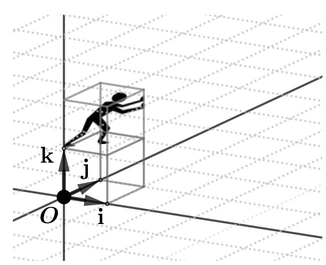

If I tell you that the matrix of this transformation is

$$\begin{pmatrix} 1.5 & 2 & 1 \\ 0 & 1 & -0.5 \\ -0.5 & 0 & 2 \end{pmatrix},$$

then you can easily determine where it sends any point/vector in $\mathbb{R}^3$. The first column of the matrix, for example, tells us that $\mathbf{i}$ is scaled a bit and pushed under the xy-plane to end up at $\mathbf{i}' = 1.5\mathbf{i} - 0.5\mathbf{k}$. Similarly, we can read the fates of $\mathbf{j}$ and $\mathbf{k}$ directly from the matrix.

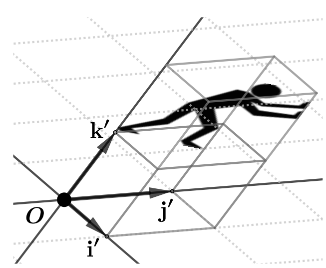

Now consider any old vector: say, $\mathbf{4i} - 3\mathbf{j} + \mathbf{6k}$. Where will our linear map send it? Following our matrix-vector multiplication recipe (in the box above), we see that the map will send the vector to

$$\begin{pmatrix} 1.5 & 2 & 1 \\ 0 & 1 & -0.5 \\ -0.5 & 0 & 2 \end{pmatrix}\begin{pmatrix} 4 \\ -3 \\ 6 \end{pmatrix} = 4\begin{pmatrix} 1.5 \\ 0 \\ -0.5 \end{pmatrix} - 3\begin{pmatrix} 2 \\ 1 \\ 0 \end{pmatrix} + 6\begin{pmatrix} 1 \\ -0.5 \\ 2 \end{pmatrix} = \begin{pmatrix} \mathbf{6} \\ \mathbf{-6} \\ \mathbf{10} \end{pmatrix}. \quad ♦$$

In $\mathbb{R}^4$ or still higher-dimensional spaces, linear maps may be impossible to visualize, and they are certainly impossible to draw, but writing down matrices that represent them is easy. You'll have the chance to do that in the exercises, but for now, let us return to $\mathbb{R}^2$ for an eminently familiar transformation.

Example 4. Rotating the plane counterclockwise around the origin through an angle of θ is a transformation, but is it a *linear* transformation? If so, how can we represent it as a matrix?

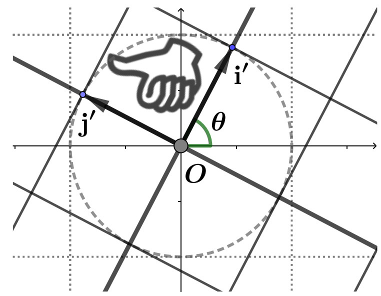

Solution. We've defined a linear map as a grid transformation induced by dragging the tips of standard basis vectors to new places. A rotation about the origin clearly qualifies, because rotating $\mathbf{i}$ and $\mathbf{j}$ through θ induces a rotation of the whole grid.

Since a rotation is a linear map, we can represent it as a matrix. As discussed above, the columns of this matrix will be the rotated images of $\mathbf{i}$ and $\mathbf{j}$, expressed as column vectors.

The dashed circle in the figure is the *unit* circle, and by the unit-circle definitions of the sine and cosine functions, we can see that $\mathbf{i}' = (\cos\theta)\,\mathbf{i} + (\sin\theta)\,\mathbf{j}$. Those same definitions yield $\mathbf{j}' = [\cos(\theta + 90°)]\mathbf{i} + [\sin(\theta + 90°)]\mathbf{j} = (-\sin\theta)\,\mathbf{i} + (\cos\theta)\,\mathbf{j}$. It follows that the rotation matrix is

$$\begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}.$$

For example, if we wanted to know where the point $(1.5, 0.2)$ ends up after rotating it $22°$ counterclockwise around the origin, we could now find the answer very easily. The point will end up at the point corresponding to the vector

$$\begin{pmatrix} \cos 22° & -\sin 22° \\ \sin 22° & \cos 22° \end{pmatrix}\begin{pmatrix} 1.5 \\ 0.2 \end{pmatrix} = 1.5\begin{pmatrix} \cos 22° \\ \sin 22° \end{pmatrix} + 0.2\begin{pmatrix} -\sin 22° \\ \cos 22° \end{pmatrix} \approx \begin{pmatrix} 1.32 \\ 0.63 \end{pmatrix}.$$

In other words, after the rotation, point $(1.5, 0.2)$ ends up at approximately $(\mathbf{1.32}, \mathbf{0.63})$.

Similarly, if we wished to rotate, say, $(0.5, -0.7)$ *clockwise* around the origin by $137°$, the same type of calculation will work if we let $\theta = -137°$ (Note that negative!):

$$\begin{pmatrix} \cos(-137°) & -\sin(-137°) \\ \sin(-137°) & \cos(-137°) \end{pmatrix}\begin{pmatrix} 0.5 \\ -0.7 \end{pmatrix} = 0.5\begin{pmatrix} \cos(-137°) \\ \sin(-137°) \end{pmatrix} - 0.7\begin{pmatrix} -\sin(-137°) \\ \cos(-137°) \end{pmatrix} \approx \begin{pmatrix} -0.84 \\ 0.17 \end{pmatrix}.$$

Thus, after the rotation, point $(0.5, -0.7)$ ends up at approximately $(\mathbf{-0.84}, \mathbf{0.17})$. ♦

So far, we've considered *local* changes wrought by linear maps. But what do such maps achieve *globally*? The linear maps of $\mathbb{R}^2$ that we've met so far have moved points around the plane, but the global output of each such map was always… the plane. The analogous story held in our one example of a map of $\mathbb{R}^3$. This raises a question: Must *every* linear map's global output simply be the space on which it is defined? No. This has happened so far only because each map's matrix has had *linearly independent* columns, which ensured that the standard basis vectors were just mapped onto a new *basis* for the same old space. But if a map were to have a matrix with linearly *de*pendent columns, then the clean standard grid with which we started would be mapped onto a tangled grid, at least one of whose generating vectors lies in the span of the others, with the obvious result that at least one dimension of the global output would collapse.

Example 5. Discuss the linear map whose matrix is $\begin{pmatrix} 2 & 4 \\ 1 & 2 \end{pmatrix}$.

Solution. The columns of this map's matrix are linearly *de*pendent: The transformed images of $\mathbf{i}$ and $\mathbf{j}$ lie along the same line. It follows that the columns' span *is* that line. The linear map crushes the standard two-dimensional grid of squares down into the line, and thus, although the linear map is defined on $\mathbb{R}^2$, its global output is nothing but this lone line, a *one*-dimensional subspace of $\mathbb{R}^2$. ♦

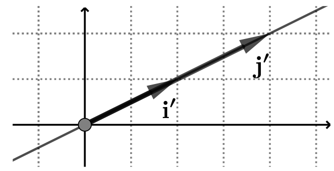

It helps to think of a linear map as a continuous transformation, with the standard basis vectors (and grid) morphing into their images. For instance, in the previous example, we might imagine $\mathbf{i}$ and $\mathbf{j}$ lengthening and approaching one another (as in the sequence of pictures below), with both coming to rest on the line $y = x/2$, at which point all of two-dimensional space has folded up, fanlike, into a line.

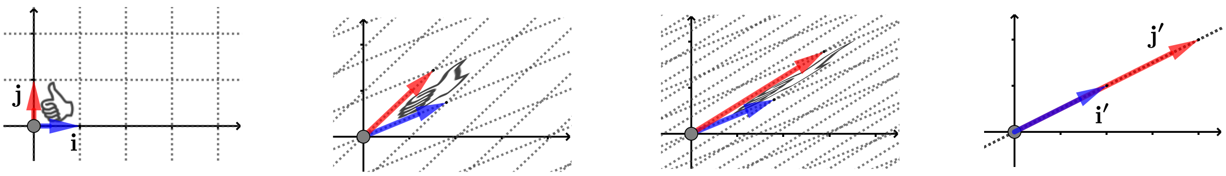

But suppose one of the vectors *didn't* stop at the line, but crossed over it. Then what would happen?

The figure at right shows such a situation. I've dragged $\mathbf{j}'$ past $\mathbf{i}'$, bringing its tip to rest at point $(3, -1)$. Vectors $\mathbf{i}'$ and $\mathbf{j}'$ are now linearly independent, so the map's global output is two-dimensional, but with a notable change: From its "folded up fan" state, the grid has now re-emerged "flipped over". We've gone through the looking glass, reversing the orientations of figures in the plane. In our original drawing, for example, the hand was a *right* hand, but it has now been transformed into a *left* hand, just as your own right hand becomes a left hand when you view it in a mirror. I'll say more about orientation-reversing maps in Chapter 5. For now, I'll note that the simplest orientation-reversing maps are ordinary reflections. Reflections are some of the most basic (and important) of all linear transformations.

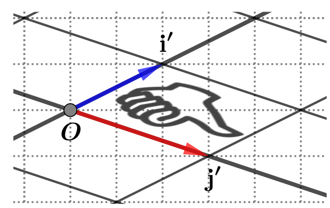

Example 6. Find the matrix for reflection across the line $y = x$.

Solution. The before-and-after figures at right give us all we need: The tips of $\mathbf{i}$ and $\mathbf{j}$ are mapped, respectively, to points $(0,1)$ and $(1,0)$. Accordingly, the reflection matrix will be

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

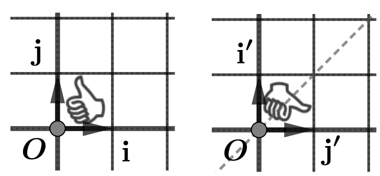

To confirm that this is correct, we note that our reflection should send any point (a, b) to (b, a). Does our matrix actually accomplish this? Yes it does:

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}\begin{pmatrix} a \\ b \end{pmatrix} = a\begin{pmatrix} 0 \\ 1 \end{pmatrix} + b\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} b \\ a \end{pmatrix}.$$ ♦

Exercises.

1. True or false. Explain your answers.

a) Every linear transformation fixes the origin (i.e. maps the origin to itself).
b) Every linear transformation moves all points other than the origin to new locations.

2. For each transformation of $\mathbb{R}^2$, find the associated matrix, and use it to determine where the map sends $(2, 3)$.

a) Reflection over $x = 0$ b) Reflection over $y = 0$ c) What about reflection over $y = 1$? (Careful.)
d) Rotation by θ counterclockwise about the origin.
[Don't just copy it down from the book or memorize it. Be sure you can explain where it comes from.]
e) Rotation by $30°$ counterclockwise about the origin. f) Rotation by $45°$ clockwise about the origin.
g) The "do nothing" map, which leaves the plane as it is. This map's matrix is called **the identity matrix**.

3. The linear maps represented by the following matrices distort the square grid generated by **i** and **j** into new forms. Sketch the new grids. On each, indicate (without doing any calculations!) the images of **i**, **j**, and $\mathbf{v} = 2\mathbf{i} + 3\mathbf{j}$.

a) $\begin{pmatrix} 1.5 & 0 \\ 0 & 2 \end{pmatrix}$ b) $\begin{pmatrix} 2 & 1 \\ -1 & 1 \end{pmatrix}$ c) $\begin{pmatrix} -1 & 2 \\ 1 & 0 \end{pmatrix}$

4. The linear map whose picture is shown at right is called a horizontal **shear**.

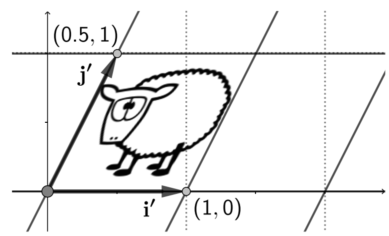

a) Find the shear's matrix.
b) To which point is $(10, 22)$ mapped by the shear?
c) Which point is mapped to $(10, 22)$ by the shear?
d) The shear transforms the original grid's squares into parallelograms. Does it change their areas in the process? Why or why not?
e) Does shearing the sheep change its area? Why or why not?

5. Some terminology: A matrix with n rows and n columns is called an $\boldsymbol{n} \times \boldsymbol{n}$ **matrix** (spoken as: "n by n matrix").

a) Give two examples of 2×2 matrices that map all of $\mathbb{R}^2$ onto the x-axis.
b) Give two examples of 2×2 matrices that map all of $\mathbb{R}^2$ onto the line $y = 2x$.
c) Give two examples of 3×3 matrices that map all of $\mathbb{R}^3$ onto the x-axis, thus crushing *two* dimensions.

6. The $n \times n$ **zero matrix** consists of nothing but zeros. What does it do geometrically?

7. Find the matrices that carry out the following transformations in $\mathbb{R}^3$:

a) Reflection across the xy-plane (i.e. the plane $z = 0$) b) Reflection over the *plane* $y = x$.
c) Rotation by θ about the z-axis, counterclockwise from the perspective of one looking down at the origin from a point on the positive z-axis.
d) The shear that fixes all points in the xy-plane but moves $(0,0,1)$ to $(0.5, 0.5, 1)$.
e) The "do nothing" transformation. [See Exercise 2g above.]

8. Find the matrix that…

a) reflects $\mathbb{R}^4$ across the 3-dimensional hyperplane $w = 0$ (if we call the four dimensions x, y, z, and w). You'll need to think about what this means in analogy with lower-dimensional cases.
b) reflects $\mathbb{R}^n$ across the $(n-1)$-dimensional hyperplane $x_i = 0$ (where the i^{th} dimension is x_i).
c) Describe the matrix that does nothing to $\mathbb{R}^n$. (**The $\boldsymbol{n} \times \boldsymbol{n}$ identity matrix**.)

9. Carry out the following matrix-vector multiplications:

a) $\begin{pmatrix} 1 & 3 \\ 2 & 4 \end{pmatrix}\begin{pmatrix} 5 \\ 6 \end{pmatrix}$ b) $\begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & -1 \\ 2 & 3 & 0 \end{pmatrix}\begin{pmatrix} 4 \\ -1 \\ 8 \end{pmatrix}$ c) $\begin{pmatrix} 1 & 1 & 2 & -1 \\ 0 & 1 & 1 & 0 \\ 2 & 1 & 0 & 1 \\ 0 & -1 & 2 & -1 \end{pmatrix}\begin{pmatrix} 1 \\ 0 \\ -1 \\ 2 \end{pmatrix}$

Non-square Matrices

Rationalists, wearing square hats,
Think, in square rooms…
- Wallace Stevens, "Six Significant Landscapes"

So far, we've only considered linear transformations that map $\mathbb{R}^n$ into *itself*. True, we've seen that some such maps can collapse dimensions (when the matrix's columns are linearly *de*pendent), but the collapse still occurs within the ambient space of $\mathbb{R}^n$. For example, the map that orthogonally projects each vector of $\mathbb{R}^2$ onto the horizontal axis crushes all of $\mathbb{R}^2$ into a single line, but this line *isn't* $\mathbb{R}$; the projection map's output vectors are still very much vectors in $\mathbb{R}^2$. They all just happen to have a second component of zero.

Some transformations, however, map $\mathbb{R}^n$ into a different space altogether, $\mathbb{R}^m$. How does this work? By following our familiar guiding idea: If we know where the map sends $\mathbb{R}^n$'s standard basis vectors, then the fates of all the other vectors in $\mathbb{R}^n$ will be determined *by linearity* (which is a shorthand phrase for "by the preservation of linear combinations"). Maps of this sort are usually defined directly by a matrix, but unlike the matrices we've seen so far, these ones aren't square. Instead, they have n columns (one for each standard basis vector of $\mathbb{R}^n$), each of which has m entries (since each column represents a vector in the target space of $\mathbb{R}^m$). We call such a matrix an $\boldsymbol{m} \times \boldsymbol{n}$ **matrix**. Note the order: m rows, n columns.

Example 1. The 3×2 matrix

$$\begin{pmatrix} 2 & -1 \\ 0 & 1 \\ 2 & 1 \end{pmatrix}$$

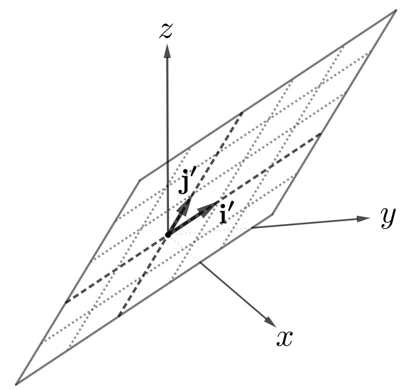

defines a linear transformation that sends vectors from $\mathbb{R}^2$ (this is clear because it has two columns) to $\mathbb{R}^3$ (clear because the columns have three entries). The figure at right indicates what the map/matrix does: It sends $\mathbb{R}^2$'s two standard basis vectors to $\mathbf{i}'$ and $\mathbf{j}'$, the columns of the matrix. Since these are linearly independent vectors, their span is a plane in $\mathbb{R}^3$.

Matrix-vector multiplication is still defined as the usual weighted sum of the matrix's columns. For example, where does this map send $7\mathbf{i} - 3\mathbf{j}$? An easy computation yields the answer:

$$\begin{pmatrix} 2 & -1 \\ 0 & 1 \\ 2 & 1 \end{pmatrix}\begin{pmatrix} 7 \\ -3 \end{pmatrix} = 7\begin{pmatrix} 2 \\ 0 \\ 2 \end{pmatrix} - 3\begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} \mathbf{17} \\ \mathbf{-3} \\ \mathbf{11} \end{pmatrix}. \qquad \blacklozenge$$

Exercises.

10. In a matrix-vector product involving an $m \times n$ matrix, how many entries must the *vector* have?

11. An $m \times n$ matrix determines a linear transformation from where to where?

12. In the example above, a 3×2 matrix maps $\mathbb{R}^2$ onto a plane in $\mathbb{R}^3$. In similar geometric terms, explain what the following matrices do:

a) $\begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$ b) $\begin{pmatrix} 1 & 0 & 2 \\ 0 & 1 & 3 \end{pmatrix}$ c) $\begin{pmatrix} 2 & -2 & 4 \\ 3 & -3 & 6 \end{pmatrix}$ d) $\begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}$ e) $\begin{pmatrix} 1 & 5 & 3 \\ 0 & 2 & 9 \\ 6 & 7 & 2 \\ 0 & 0 & 4 \end{pmatrix}$ f) $\begin{pmatrix} 1 & 2 & 4 & 1 \\ 1 & 3 & 6 & 1 \end{pmatrix}$

Another Look at the Matrix-Vector Product

We've defined matrix-vector multiplication as a weighted sum of the matrix's columns and we know what it signifies: the *transformation* of the given vector by the given matrix (i.e. the linear map that it encodes). I'll now introduce a formula that lets us quickly find any particular entry in a matrix-vector product. This simple formula is useful in proofs and speeds up our hand computations of matrix-vector products.

In the proof that follows, I'll introduce some basic matrix notation that we'll often use in the future. We symbolize a matrix with an uppercase letter such as A. To indicate a particular entry in the matrix, we use the same letter, but in *lowercase*, along with two subscripts to indicate the entry's row and column. For example, $a_{1,3}$ signifies the entry in matrix A's first row and third column. (When the context allows, we sometimes minimize eye strain by omitting the comma in the subscript, writing $a_{1,3}$ simply as a_{13}.)

With this notation in hand, let's state and prove our "i^{th}-Entry Formula".

> **Matrix-Vector Multiplication ($\boldsymbol{i}^{\text{th}}$-Entry Formula).**
> If A is a matrix and $\mathbf{v}$ is a vector, we can compute $\boldsymbol{A}\mathbf{v}$'s $\boldsymbol{i}^{\text{th}}$ entry with a *dot product*:*
>
> $$\boldsymbol{A}\mathbf{v}\text{'s } \boldsymbol{i}^{\text{th}} \text{ entry} = (i^{\text{th}} \text{ row of } A) \cdot \mathbf{v}$$

Proof. Consider the product of a general $m \times n$ matrix A and a general column vector $\mathbf{v}$:

$$\begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \cdots & \vdots \\ a_{i,1} & a_{i,2} & \cdots & a_{i,n} \\ \vdots & \vdots & \cdots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{pmatrix}.$$

By matrix-vector multiplication's definition, the product $A\mathbf{v}$ is a weighted sum of A's columns, where the weights are $\mathbf{v}$'s entries. Thinking a bit about vector addition and scalar multiplication reveals that the i^{th} entry in a weighted sum of column vectors is just… the weighted sum of the column vectors' i^{th} entries (with the same weights). It follows that the i^{th} entry of $A\mathbf{v}$ must be

$$v_1 a_{i,1} + v_2 a_{i,2} + \cdots + v_n a_{i,n} \,,$$

which is, as claimed, the dot product of A's i^{th} row with the given vector $\mathbf{v}$. ■

Example 1. In the following matrix-vector multiplication the product's 2nd entry is the dot product of the matrix's second row (in the box) and the vector:

$$\begin{pmatrix} 2 & 7 & 1 \\ \boxed{3 \quad 1 \quad -1} \\ 2 & 3 & 0 \end{pmatrix} \begin{pmatrix} 4 \\ -1 \\ 6 \end{pmatrix}.$$

Thus, it is $3(4) + 1(-1) + (-1)6 = \mathbf{5}$, as you can check by doing the full product the old way. ♦

* This is a purely formal "dot product" - mere shorthand for "sum of the products of corresponding entries in two number lists". That it works out like this is just a happy accident; it has no connection to the actual geometric definition of the dot product.

When doing matrix-vector multiplications by hand, it's quicker to repeatedly use this i^{th}-entry formula, mentally finding the product's entries one at a time, than it is to write out the weighted sums of columns. Make up a few matrix-vector multiplications on your own and do them both ways; you'll see what I mean. Because of this computational ease, I predict that you'll soon be using the dot product for all your matrix-vector multiplications, but please, please, please… never forget that matrix-vector multiplication is first and foremost a weighted sum of columns, an operation with clear geometric *meaning*, as we discussed at length in this chapter's first section. In contrast, this section's formula is just a happy algebraic accident.*

Because the i^{th}-entry formula reduces matrix-vector multiplication to dot products, whose properties *we've already proved*, it can often serve as a "bridge" in proofs, linking the unknown back to the known. For example, let's use it prove that matrix-vector multiplication can be distributed over vector addition.

Claim. For any matrix A and vectors $\mathbf{v}$ and $\mathbf{w}$ for which the following expressions are defined,

$$A(\mathbf{v}+\mathbf{w}) = A\mathbf{v} + A\mathbf{w}\,.$$

Proof. The expressions on both sides of the equals sign represent column vectors. To prove that these two column vectors are equal as claimed, we must show that all their corresponding entries are equal. To this end, we note that for all i, we have that

$A(\mathbf{v}+\mathbf{w})$'s i^{th} entry

$= (i^{\text{th}}$ row of $A) \cdot (\mathbf{v}+\mathbf{w})$	(by the i^{th}-entry formula)
$= (i^{\text{th}}$ row of $A) \cdot \mathbf{v} \,+\, (i^{\text{th}}$ row of $A) \cdot \mathbf{w}$	(distributing the dot over vector addition)
$= A\mathbf{v}$'s i^{th} entry $+$ $A\mathbf{w}$'s i^{th} entry	(by the i^{th}-entry formula)
$= (A\mathbf{v}+A\mathbf{w})$'s i^{th} entry	(by definition of column vector addition).

Since all entries of $A(\mathbf{v}+\mathbf{w})$ and $A\mathbf{v}+A\mathbf{w}$ are equal, the two vectors themselves are equal. ■

Exercises.

13. Redo each of the matrix-vector products from Exercise 9, but this time, do not write out the weighted sums of the columns. Instead, just use the i^{th}-entry formula to mentally compute the entries of the product.

14. The scalar multiple of a *matrix* is defined as you'd expect: Each entry in the matrix is multiplied by the scalar.

a) If A is any old matrix, describe the geometric relationship between the linear maps represented by A and $2A$.

b) Prove that for any matrix A and scalar c, the following holds: $A(c\mathbf{v}) = c(A\mathbf{v})$. That is, scalar multiples can be 'pulled through' matrix-vector multiplication. [*Hint*: Use the ideas from the proof of the claim above.]

c) Prove that matrix-vector multiplication preserves linear combinations: $A(c\mathbf{v}+d\mathbf{w}) = c(A\mathbf{v}) + d(A\mathbf{w})$. [*Hint*: All the hard work has been done already. No need to reinvent the wheel.]

15. If A is a 5×5 matrix, describe the vector $\mathbf{v}$ we'd need to use as an input so that the output $A\mathbf{v}$ would be…

a) A's 5^{th} column. b) The sum of A's 2^{nd} and 3^{rd} columns. c) 3 times A's 4^{th} column.

d) 3 times A's 4^{th} column minus 5 times A's 1^{st} column.

The problem's moral: We can use matrix-vector multiplication to form linear combinations of a matrix's columns.

* Regrettably, many linear algebra books *define* matrix-vector multiplication in terms of this shortcut formula – long before they even discuss linear transformations. As a result, students end up memorizing a strange rule without any hope of understanding why matrix-vector multiplication exists in the first place.

Matrix Multiplication

It's a mistake to think the practice of my art has come easily to me. I assure you, dear friend, no one has given so much care as I to the study of composition.

- Mozart, to conductor Johann Baptist Kucharz

As you learned long ago, one way to build new functions from old ones is through function *composition*. For example, if we compose the squaring function $f(x) = x^2$ with the doubling function $g(x) = 2x$, we will obtain, depending on the order of composition, either $f\big(g(x)\big) = 4x^2$, or $g\big(f(x)\big) = 2x^2$.

We can also compose linear maps. Composing two yields a third, which of course has its own matrix. We wish to understand how the matrices of the two "parent maps" relate to the matrix of their child, the composite map. We'll begin with an example of finding a composite map's matrix from first principles. Then, after we establish an algebraic theorem about matrix multiplication, we'll return to the same example and view it from another perspective.

Example 1. Let R represent a 90° counterclockwise rotation about the origin, and let F represent a flip (i.e. a reflection) across the horizontal axis. A little thought yields

$$R = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \text{ and } F = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Let us consider the composite map that first rotates, then flips. We will call this map's matrix FR. Note the "right-to-left" order: FR means reflect first, then flip. This isn't as odd as it might seem. It's analogous to function notation, where $f(g(x))$ means first apply g, then f.

What does matrix FR look like? We can determine this matrix's columns in the usual way: Rotating $\mathbf{i}$ by 90° then flipping the result over the horizontal axis turns it into $-\mathbf{j}$. Rotating $\mathbf{j}$ by 90° and flipping the result yields $-\mathbf{i}$. Recording these transformed images in a new matrix, we obtain

$$FR = \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}.$$ ♦

We human beings understand that the three matrices above are related through a series of geometric operations that we can visualize. But a computer – a soulless box of chips and wires – can understand none of this. It can only follow orders. Can we program a computer to obtain FR from F and R in a purely algebraic way, by following simple instructions that don't require *thinking* about rotations and reflections? We can. By definition of function composition, we know that $(FR)\mathbf{i} = F(\boldsymbol{R}\mathbf{i})$, and the latter expression can guide us as we program a computer to compute FR's first column mindlessly:

$$F(R\mathbf{i}) = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\left(\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}\begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ -1 \end{pmatrix}.$$

You and I can understand the *meaning* of what the computer is doing here: It first determines where $\mathbf{i}$ is mapped by the rotation matrix, and then determines where this transformed image of $\mathbf{i}$ is sent by the flip. But this means nothing at all to the computer, which just blindly follows orders, doing lots of arithmetic – adding, subtracting, and multiplying numbers – and storing the results where we tell it to store them.

Because notation such as FR looks like multiplication, we will in fact call it **matrix multiplication**. In a moment, we'll have an algorithm for doing matrix multiplication quickly. But never let the algorithm distract you from the core meaning: *Matrix multiplication corresponds to the composition of linear maps*. That is why we are interested in it in the first place. Every instance of matrix multiplication can be thought of as the composition of two linear transformations. Never forget.

Now we'll derive two formulas for the product AB of two matrices A and B. They will give us two different perspectives on a matrix product: a *column* perspective (often useful in proofs) and an *entry* perspective (useful in proofs, and also for computing matrix products by hand.)

Matrix Multiplication (Column Perspective).

If A and B are matrices, then AB looks like this:

$$AB = \begin{pmatrix} | & \cdots & | \\ A\mathbf{b_1} & \cdots & A\mathbf{b}_n \\ | & \cdots & | \end{pmatrix},$$

where $\mathbf{b}_j$ represents B's j^{th} column.

Proof.

We must show that AB's j^{th} column is $A\mathbf{b}_j$ for all j. To this end, we note that, for any j, we have

AB's j^{th} column
$= \mathbf{e}_j$'s image under the composite map* (by definition of a linear map's matrix)
$= A(B\mathbf{e}_j)$ (by definition of AB)
$= A\mathbf{b}_j$ (by definition a linear map's matrix) ■

We'll use this column perspective right away – to derive the *entry* perspective!

Matrix Multiplication (Entry Perspective).

If A and B are matrices, we can compute AB's $\boldsymbol{ij}^{\text{th}}$ **entry** (that is, its entry in row i, column j) with a *dot product*:†

$$AB\text{'s } ij^{\text{th}} \text{ entry} = (A\text{'s } i^{\text{th}} \text{ row}) \cdot (B\text{'s } j^{\text{th}} \text{ column})$$

Proof.

AB's ij^{th} entry
= the i^{th} entry in AB's j^{th} column (definition of ij^{th} entry)
= the i^{th} entry in $A\mathbf{b}_j$ (column perspective of matrix multiplication)
$= (A$'s i^{th} row$) \cdot \mathbf{b}_j$ (i^{th}-entry formula for matrix-vector multiplication)
$= (A$'s i^{th} row$) \cdot (B$'s j^{th} column$)$ (definition of $\mathbf{b}_j$) ■

* Recall from Chapter 1 that in $\mathbb{R}^n$ (when $n > 3$), we call the standard basis vectors $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$.

† This is another purely formal "dot product" – a sum of products of corresponding entries in two number lists – with no real connection to the dot product's geometric definition.

From the 'entry perspective', matrix multiplication is a mere algorithmic grind in which we crank out the product one entry at a time, doing the arithmetic in our heads. To see how this works out in practice, let's redo an earlier example. In Example 1, we composed a rotation and a flip, finding the composition's matrix from first principles: We thought geometrically about where the composition would map $\mathbf{i}$ and $\mathbf{j}$, and then we made their transformed images our composition matrix's columns. But now we can redo this problem, producing the composition's matrix in seconds. We begin by writing down our two matrices, along with an empty matrix – soon to be filled in – to represent their product:

$$FR = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} \quad & \quad \\ \quad & \quad \end{pmatrix}.$$

From here, our task is nothing but simple mental arithmetic: FR's top left entry (that is, row 1, column 1) will be, according to the entry perspective, the dot product of F's row 1 and R's column 1.

$$FR = \begin{pmatrix} \boxed{1 \quad 0} \\ 0 \quad -1 \end{pmatrix}\begin{pmatrix} \boxed{\begin{matrix} 0 \\ 1 \end{matrix}} & \begin{matrix} -1 \\ 0 \end{matrix} \end{pmatrix} = \begin{pmatrix} \quad & \quad \\ \quad & \quad \end{pmatrix}.$$

This dot product is $(1)(0) + (0)(1) = 0$, so we put a 0 in the top left corner of our product matrix:

$$FR = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & \quad \\ \quad & \quad \end{pmatrix}.$$

Doing the appropriate dot products for the three remaining entries, we end up with, as you should verify,

$$FR = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}.$$

Naturally, this is the same matrix we produced in Example 1, but now we've obtained it by mindlessly turning an algebraic crank, heedless of the *meanings* of any of the matrices involved. Reducing matrix multiplication to an algorithm frees us to concentrate on larger things, and in that sense, it is a blessing… provided that we never forget matrix multiplication's underlying meaning, which is now hidden neatly under the hood: Matrix multiplication corresponds to the *composition* of linear maps.

While Example 1's matrices are here on the page, let's go ahead and compute RF, too:

$$RF = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Note well: $RF \neq FR$. We've discovered something important: **Matrix multiplication is *not* commutative**. This may sound exotic, but on further reflection, it should feel completely natural to you – and to anyone who knows not just how but *why* we multiply matrices. Matrix multiplication corresponds to the composition of linear maps, and everyone knows that when we compose two functions, *the order matters* (as this section's first paragraph reminded us). Hence, order must matter in matrix multiplication, too.

On the other hand, **matrix multiplication is associative**. An algebraic proof of this is unilluminating, but we don't need one. We can see why it is true simply by thinking about matrix multiplication's *meaning*. If we think of each matrix as "doing something" – namely, carrying out the corresponding map – then $A(BC)$ means "do C then B… then do A." On the other hand, $(AB)C$ means "do C… then do B then A." Obviously, the result is the same either way, so it follows that $A(BC) = (AB)C$, as claimed.

Exercises.

16. Carry out the following multiplications using the 'entry perspective' on matrix multiplication:

a) $\begin{pmatrix} 1 & 3 \\ 2 & 1 \end{pmatrix}\begin{pmatrix} 0 & 2 \\ -1 & 1 \end{pmatrix}$ b) $\begin{pmatrix} 0 & 2 \\ -1 & 1 \end{pmatrix}\begin{pmatrix} 1 & 3 \\ 2 & 1 \end{pmatrix}$ c) $\begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}\begin{pmatrix} 0 & 2 & 0 \\ 1 & 0 & 3 \\ 0 & 4 & 0 \end{pmatrix}$

17. Although matrix multiplication is noncommutative in general, certain specific pairs of matrices A and B have the property that $AB = BA$. Think of some specific examples of pairs of 2×2 matrices like this.
[***Hint***: *Think geometrically. Think of linear maps that yield the same result can be done in either order*.]

18. The $n \times n$ *identity matrix*, which you met in Exercise 8c, has 1s on its "main diagonal" (upper left to lower right) and 0s elsewhere. It is denoted by the letter I. Explain why $\boldsymbol{AI} = \boldsymbol{IA} = \boldsymbol{A}$ holds for all $n \times n$ matrices A.

19. In Exercise 11, you saw that an $m \times n$ matrix represents a linear map from $\mathbb{R}^n$ to $\mathbb{R}^m$. (Note the order!)

a) Suppose A is a 5×2 matrix, and B is a 2×3 matrix. Explain why (in terms of the underlying linear maps) the matrix AB is defined, but BA is undefined.

b) What are the dimensions of matrix AB in the previous part? It represents a linear map from where to where?

c) Now suppose C is a 3×5 matrix. Given the dimensions of A and B above, is the matrix product $BCAB$ defined? If so, find its dimensions.

20. Let $M = \begin{pmatrix} 1 & 2 \\ 3 & 0 \\ -2 & 1 \end{pmatrix}$ and $N = \begin{pmatrix} 2 & -3 & 1 \\ 1 & 1 & 0 \end{pmatrix}$. Find the following (if they are defined): $MN,\ NM,\ M^2,\ N^2,\ NMN$.

21. Consider the linear map of the composition that rotates points in the plane by 120° counterclockwise about the origin and then reflects the results over the horizontal axis.

a) Find its matrix. b) Where does this operation send the point $(3, 8)$?

c) If we call this composition's matrix M, find M^2. Why does the matrix that you found make intuitive sense?

22. Prove that scalar multiples can be 'pulled through' matrix multiplication: $(cA)(dB) = (cd)AB$.

23. The matrix that undoes the action of a square matrix A is called the **inverse matrix** of A, and is denoted A^{-1}. Thus, by definition, A^{-1} is a matrix such that $A^{-1}A = I = AA^{-1}$ (where I is the identity matrix).* In this exercise, you'll play with the idea of inverse matrices and learn two important computational facts about them.

a) Let A be the 2×2 matrix of a 90° rotation counterclockwise about the origin. Describe what A^{-1} does. Then find A^{-1}. Check your work by grinding out $A^{-1}A$ and $A^{-1}A$. Is the product what you expected?

b) Only *square* matrices can have inverses. Explain why.

c) Prove that matrix inverses are *unique*. (i.e. if B and C are inverses of A, then $B = C$.)

d) Not every square matrix has an inverse. Explain why, for example, the 2×2 zero matrix (see Exercise 6), which crushes the whole plane down into the origin, is not invertible.

e) **Important fact 1.** If A and B are two invertible matrices, then $\boldsymbol{(AB)^{-1} = B^{-1}A^{-1}}$. *Note that reversal of order!* (You can grasp this result intuitively by thinking of A as "put on your shoes" and B as "put on your socks". With those in mind, state the meanings of $A^{-1},\ B^{-1},\ AB,\ BA, (AB)^{-1},$ and $B^{-1}A^{-1}$. The equality will make sense.) Prove the result formally by showing that $B^{-1}A^{-1}$ satisfies the definition of AB's inverse.

f) Let M be the matrix in Exercise 21. Find M^{-1} by using Part D. Check your work by finding M^{-1} another way.

g) **Important fact 2**: If A is an invertible matrix and c is a scalar, then $\boldsymbol{(cA)^{-1} = c^{-1}A^{-1}}$. Explain why.

---

* If we wish to show that a matrix B is in fact A^{-1}, we must – by this definition – verify two separate things: $BA = I$ and $AB = I$. But we'll prove later (Ch. 5, Exercise 12) that each of these things *implies the other*, so to verify both, we just need to verify *one*. This is tricky to prove rigorously, but it's intuitively plausible for those (like you!) who understand that matrix multiplication is a kind of *composition*; if B is the map that "undoes" whatever A does, then it makes sense that A is the map that "undoes" B.

Matrix Addition and Transposition

> Algebra which cannot be translated into good English
> and sound common sense is bad algebra.
>
> - William Kingdon Clifford, *The Common Sense of the Exact Sciences* (Chapter 1, Section 7)

If matrices A and B have the same dimensions, we define their **matrix sum** $A + B$ as the matrix we obtain by adding the corresponding entries in A and B. Thus, for example,

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} + \begin{pmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{pmatrix} = \begin{pmatrix} 2 & 5 & 8 \\ 6 & 9 & 12 \end{pmatrix}.$$

Matrix subtraction is defined analogously. Compared to matrix multiplication, matrix addition is boring; it lacks a universal geometric meaning, and thus cannot easily be translated into "sound common sense". That said, it is not "bad algebra". Adding and subtracting matrices enriches matrix algebra in subtle ways, usually by simplifying linear algebraic algorithms, which is no small thing. We'll see, for example (in Chapter 7) that our algorithm for finding a matrix's *eigenvalues* will involve matrix subtraction. It will also involve scaling a matrix, an operation you met in Exercise 14. Moreover, there are some instances in which matrix addition does have a tangible interpretation, as you'll see in Exercise 24.

The **transpose** of any matrix M is the matrix whose *columns* are M's *rows* (taken in the usual order). The symbol for M's transpose is M^T. Thus, for example, if

$$M = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix}, \quad \text{then we have} \quad M^T = \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}.$$

Transposes assume greater prominence as one moves deeper into linear algebra. In time, we'll see how transposes relate to both inverse matrices and *symmetric matrices*, which you'll first meet in exercise 27. As with matrix addition, it's difficult to appreciate transposition's value at the outset, since it lacks a clear geometric interpretation. Still, it's good to make its acquaintance early on. In the exercises that follow, you'll develop a few simple algebraic properties of the transpose, which we'll use later.

Exercises.

24. A grayscale image, such as the one at right, is ultimately an array of pixels. This particular one, for example, is composed of over 100,000 pixels: It is 328 pixels high and 314 wide. When you or I look at it, we behold the King of Rock and Roll and the King of the Beasts. In contrast, a sufficiently tiny bug crawling on the image would see a mere patchwork of shaded squares. In still further contrast, a *computer* would "see" such an image as a matrix. Each pixel in the grayscale image is one of 256 possible shades of gray in a black-to-white spectrum. The different shades in the spectrum are assigned numerical codes, the extreme values being Black = 0 and White = 255. Throughout the range, lower codes correspond to darker shades, higher codes to lighter ones. A computer would "know" this image as a 328 × 314 matrix, whose 102,992 entries are the codes for the shades of gray of the pixels in the corresponding position. For example, the first several rows of this giant matrix would consist of 255s (or numbers close to 255) since the top of the image is black. Call this matrix M. As we'll now see, we can alter the image by doing matrix arithmetic.

a) Roughly speaking, what would the image corresponding to the matrix $2M$ look like? (Assume throughout this exercise that any code greater than 255 is interpreted as if it were 255, and any negative code as if it were 0.)
b) What about $.5M$?
c) What about M^T?
d) What about $255U - M$, where U is the 328×314 matrix consisting entirely of 1s.
e) What about MF, where F is the 314×314 matrix with 1's on its *off* diagonal (bottom left to top right) and 0's elsewhere. [**Hint**: *The 'column perspective' on matrix multiplication should help here*.]
f) What about FM?
g) What about M', the matrix that remains after removing the bottom 14 rows of M?
h) What about FM'?

25. Matrix addition obeys the expected distributive properties, as you should now prove:

a) $(A + B)\mathbf{v} = A\mathbf{v} + B\mathbf{v}$ [**Hint**: *Explain why the sides' i^{th} entries are equal*.]
b) $(A + B)C = AC + BC$ [**Hint**: *Explain why the sides' ij^{th} entries are equal*.]
c) $A(B + C) = AB + AC$

26. In this problem, you'll learn a few algebraic facts about transposes. For parts b – d, your best strategy for the proof is to explain why the sides' ij^{th} entries are equal.

a) A curious and occasionally useful fact: We can compute a dot product with matrix multiplication, one factor of which involves a transpose. Namely, if $\mathbf{v}$ and $\mathbf{w}$ are column vectors, then

$$\mathbf{v} \cdot \mathbf{w} = \mathbf{v}^T\mathbf{w}.$$

Explain why. [**Note**: *To make this work, we blur the distinction between a real number and a* 1×1 *matrix*.]
b) The transpose of a sum is the sum of the transposes: $(M + N)^T = M^T + N^T$. Prove it.
c) Scalar multiples can be pulled through a transpose: $(cM)^T = c(M^T)$. Prove it.
d) The transpose of a product has an important order-reversing property: $(MN)^T = N^T M^T$. Prove it. (The order reversal should remind you of what you learned about matrix *inverses* in Exercise 23c.)
e) Explain why the previous property extends to 3 or more matrices, so that, for example, $(ABC)^T = C^T B^T A^T$.

27. A matrix A is said to be *symmetric* if $A = A^T$.

a) Can a non-square matrix be symmetric? If so, give an example. If not, why not?
b) Write down some symmetric matrices of various sizes.
c) Make up any old 2×3 matrix, M. Compute MM^T. What do you notice? Now compute $M^T M$. Magic, no?
d) In fact, given *any* matrix whatsoever, its product with its transpose (in either order) will always be symmetric. Explain why. (The cleanest - but not the only - way to prove this uses the property you proved in Exercise 26d.)

Abstract Linear Algebra: A Trailer

There is no branch of mathematics, however abstract, which may not someday be applied to phenomena of the real world.

- Nicolai Ivanovich Lobachevski

We can extend the scope and power of linear algebra by considering it from a higher level of abstraction. Abstraction, however, comes at a cost: diminished intuition. Accordingly, too much abstraction is usually a poor choice for one's initial foray into a subject. Part of the goal of this book is to help you develop such a strong intuition for linear algebra in a concrete setting – where vectors are arrows – that you'll be able to appreciate abstract linear algebra, where "vectors" can be all sorts of unexpected things. And even though we've just begun our journey through this book, we've made enough progress already to warrant a brief ascent into the abstract skies – just enough to give you a glimpse of what linear algebra looks like from a higher perspective. Consider this section a trailer for this course's sequel. Later in the book, I may refer in passing to this section, but it is not essential for what follows, and may be skipped.

What are **vectors**? This is the question with which we began Chapter 1. In the abstract perspective, we simply say that a vector is any element of a **vector space**, which itself can be a set of any kind so long as it is closed under addition and multiplication by scalars.* Naturally, the spaces in which we've worked ($\mathbb{R}^n$ and its subspaces) are all examples of vector spaces, but many other sets qualify as well.

For example, $\mathcal{P}_3$, the set of all polynomials of degree 3 (or less) with real coefficients, is a vector space. After all, this set is clearly closed under addition (the sum of any two such polynomials is another one), and multiplication by scalars (multiplying a polynomial by a scalar never increases its degree). Thus $\mathcal{P}_3$ is indeed a vector space, in which the "vectors" are *polynomials*. Another example: $M_{3\times 3}$, the set of all 3×3 matrices, which is obviously closed under matrix addition and under multiplication by scalars. Another notable example comes from the study of differential equations: The set of all solutions to any specific linear differential equation (such as $3y'' + 2y' + 5 = 0$) is a vector space.

For the rest of this section, I'll stick with $\mathcal{P}_3$ for some concrete examples.

Defined abstractly, a **linear transformation** is a function defined on a vector space that preserves linear combinations. In symbols, T is a linear transformation of a vector space if $T(\mathbf{v} + \mathbf{w}) = T(\mathbf{v}) + T(\mathbf{w})$ and $T(c\mathbf{v}) = cT(\mathbf{v})$ for all scalars c and all vectors $\mathbf{v}, \mathbf{w}$ in the space.

One example of a linear transformation of $\mathcal{P}_3$ is *differentiation*. This might be easier to see if we signify the operation of taking a derivative with a D, rather than the usual prime. For in that case, we have $D\big(p(x) + q(x)\big) = D\big(p(x)\big) + D(q(x))$ and $D\big(cp(x)\big) = cD(p(x))$ for any polynomials $p(x)$ and $q(x)$, as you learned in your first calculus class. Indeed, you may have even learned in that class that these two properties (derivative of a sum is the sum of the derivatives, constant multiples can be pulled through the derivative) are called the derivative's "linearity properties". Now you know the source of that name.

All of Chapter 2's concepts (linear independence, span, basis, etc.) can be defined abstractly so that they can apply to any vector space. For example, consider the four simple "vectors" $x^3,\ x^2,\ x,\ 1$ in $\mathcal{P}_3$. Linear combinations of these four vectors have the form $ax^3 + bx^2 + cx + d$. Ah, but *every* vector in $\mathcal{P}_3$ has that form, so that means that these four vectors *span* $\mathcal{P}_3$. It's also clear that these four vectors are *linearly independent* of each other (each lies outside the span of the other three). Consequently, they constitute a *basis* for $\mathcal{P}_3$. In fact, we call them $\mathcal{P}_3$'s *standard* basis vectors. Moreover, by using a shorthand

* There's a bit more to the abstract definition of a vector space but the details are immaterial in this overview.

column vector notation, we can indicate vectors in $\mathcal{P}_3$ without explicitly writing out $\mathcal{P}_3$'s standard basis vectors. For example, consider $\mathbf{p} = 2x^3 - 4x^2 + 8x + 3$ in $\mathcal{P}_3$. We can rewrite this as the column vector

$$\begin{pmatrix} 2 \\ -4 \\ 8 \\ 3 \end{pmatrix}.$$

Every linear map of $\mathcal{P}_3$ (or any other vector space) can be represented by a *matrix*, whose columns indicate the transformed images of the vector space's standard basis vectors. For example, differentiation sends our first standard basis vector, x^3, to $3x^2$, so the first column of the differentiation matrix will be

$$\begin{pmatrix} 0 \\ 3 \\ 0 \\ 0 \end{pmatrix}.$$

In fact, the full differentiation matrix will be, as you should verify,

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

And just as we used rotation matrices to carry out rotations in $\mathbb{R}^2$, we can use this differentiation matrix to carry out differentiation in $\mathcal{P}_3$. For example, we could differentiate $\mathbf{p} = 2x^3 - 4x^2 + 8x + 3$ as follows:

$$D\mathbf{p} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 3 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 2 \\ -4 \\ 8 \\ 3 \end{pmatrix} = \begin{pmatrix} 0 \\ 6 \\ -8 \\ 8 \end{pmatrix}.$$

That last column vector corresponds to $6x^2 - 8x + 8$, which is of course the correct expression. We've taken a derivative by means of matrix-vector multiplication!

Since every vector in $\mathcal{P}_3$ can be indicated by a column vector with four entries, $\mathcal{P}_3$ is intimately related to $\mathbb{R}^4$, which of course has the same property. We say that those two vector spaces are **isomorphic**, meaning that although their elements look very different on the surface (arrows in four-dimensional space in one case, polynomials of degree three or lower in the other), the two vector spaces have the same basic structure. Accordingly, we can carry insights about one space over to the other. *Every* finite-dimensional real vector space, in fact, turns out to be isomorphic to $\mathbb{R}^n$ for some value of n. It is this fact that justifies concentrating exclusively on $\mathbb{R}^n$ in a first linear algebra course. Everything you learn about linear algebra in $\mathbb{R}^n$ can be used later to understand what seem to be radically different vector spaces.

Chapter 5
The Determinant

The Determinant: Definition

Once upon a midnight dreary, while I pondered, weak and weary,
over many a quaint and curious volume…

- Edgar Allen Poe, "The Raven"

The determinant's definition will initially look mysterious, but the clouds will soon clear.

> **Definition**. If A is a square $n \times n$ matrix, its **determinant** ($\det A$) is the real number whose…
>
> - *Magnitude* is the n-dimensional volume of the "box" in $\mathbb{R}^n$ that A's columns determine.
> - *Sign* is positive if A preserves orientation, and negative if A reverses orientation.

I've used the term "box" to encompass parallelograms (in $\mathbb{R}^2$), parallelepipeds (in $\mathbb{R}^3$), and their higher-dimensional analogues (parallelotopes). Similarly, I've used "volume" to cover parallelograms' *areas* (in $\mathbb{R}^2$), parallelepipeds' volumes (in $\mathbb{R}^3$), and parallelotopes "hypervolumes" (in the appropriate spaces).

As for "orientation", recall that in Chapter 3 we considered a drawing of a right hand in $\mathbb{R}^2$, and we observed that certain maps (such as rotations) preserve the right hand's right-handedness, while others (such as reflections) transform it into a *left* hand. An orientation-*reversing* map of this latter sort sends the entire plane through the looking glass. This can occur in spaces of any number of dimensions. How does such a thing arise? Recall that any linear map of $\mathbb{R}^n$ is determined by its action on the standard basis vectors, since they specify the space's axes and their associated "grid". The linear map then transforms the space by transforming the grid, rotating and stretching the basis vectors and the axes they determine. If, while undergoing this transformation, one axis crosses through the others' span, this will cause an orientation reversal of the space. If two such crossings occur, the two reversals will undo one another, with a net effect that orientation is preserved. Indeed, the orientation is preserved or reversed according to whether there are an even or odd number of such "axis crossings" during the transformation.

Example 1. The figure at right shows the map whose matrix is

$$A = \begin{pmatrix} 5/4 & 1/2 \\ -1/4 & 3/2 \end{pmatrix}.$$

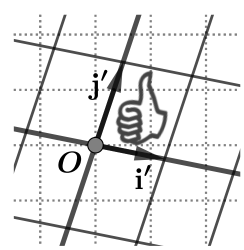

The area of the "box" in which the hand is drawn turns out to be 2 units^2. (Don't worry about how to determine that for now.) This map obviously preserves orientation (the right hand remains a right hand), so $\det A = 2$. ♦

Example 2. The figure at right shows the map whose matrix is

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

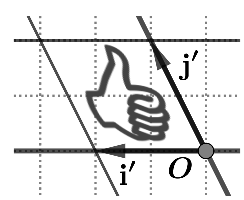

In this case, the area of the "box" is clearly 4 units^2. Moreover, this map *reverses* orientation (the right hand is now a left hand), so $\det A = -4$. ♦

Note the reason for the orientation reversal in Example 2: Dragging $\mathbf{i}$ and $\mathbf{j}$ to their new positions $\mathbf{i}'$ and $\mathbf{j}'$ requires us to pass one basis vector through the other's span – or to put it another way, it requires us to pass axis through the other.

Example 3. The puzzled man standing on the box at left and scratching his head with his left hand wonders if there are linear maps that would make him scratch with his *right* hand instead. Yes, there are: Any orientation-reversing map will do the trick. The figure at right shows one. And if the volume of the box at right is, say, 30% greater than that of the box at left, then this map's determinant must be $-\mathbf{1.3}$. ♦

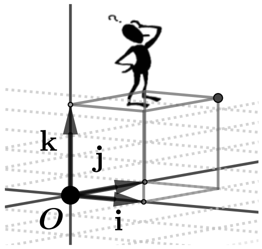

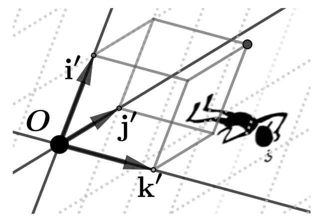

A good visualization exercise is to imagine the original axes from Example 3's left figure morphing into those at right. One way is to first imagine $\mathbf{i}$ and $\mathbf{j}$ being transformed (through rotations and stretches) into their new positions, carrying along the plane that they span. Then, once they've arrived, imagine $\mathbf{k}$ being moved down to *its* new position. At some point during this two-stage transformation, the third axis will have to pass through the plane spanned by the first two, yielding the orientation reversal.

Example 4 (Shears). A **shear** is a map that moves just *one* standard basis vector (call it the m^{th}), and in a way that its m^{th} coordinate remains 1. (It fixes all the other standard basis vectors.) The left figure shows a shear in $\mathbb{R}^2$ that moves only $\mathbf{j}$. The right figure depicts a shear in $\mathbb{R}^3$ that moves only $\mathbf{k}$, pushing the tip of $\mathbf{k}'$ over to the point $(-0.5,\ 0.2,\ 1)$. Observe that the third coordinate is still 1, as required.

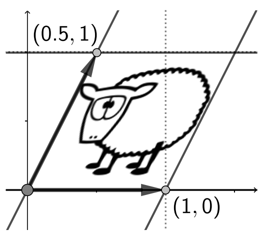

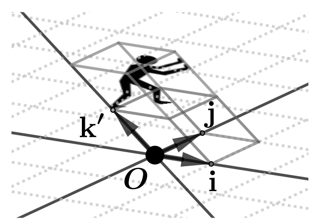

Shears will be important later in this chapter, where they'll help us find a method for computing *any* matrix's determinant. For now, however, we just want to recognize what a shear's matrix looks like, and understand what the determinant of any shear matrix must be.

The first part is easy: It follows immediately from a shear's definition that a shear matrix looks like an identity matrix in which someone has tampered with the zeros in one column, changing at least one of them to a nonzero value. Thus, the two shears depicted above are represented by these matrices:

$$\begin{pmatrix} 1 & .5 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 & -.5 \\ 0 & 1 & -.2 \\ 0 & 0 & 1 \end{pmatrix}.$$

Similarly, the following matrices represent other shears:

$$S_1 = \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix}, \qquad S_2 = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad S_3 = \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 6 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

You can draw pictures of how S_1 and S_2 transform the standard grids of their respective spaces. Shear S_3 is in $\mathbb{R}^4$, so obviously we can't draw it, but you might still enjoy thinking about it. (For example, if you were a three-dimensional being living in the hyperplane $w = 0$ with no concept of the fourth spatial dimension in which your world is embedded, would you notice it if the four-dimensional world were subjected to that shear?)

Shear matrices are thus easy to recognize. But what can we say about a shear's *determinant*?

The answer: **The determinant of every shear matrix is 1.** Proof: The original box and the sheared box have the *same base* (the one determined by the fixed standard basis vectors) and, relative to that base, they also have the *same height* (1 unit). Moreover, all the boxes' corresponding cross sections (taken at the same height and parallel to their common base) are equal, since these cross sections are both obviously equal to the boxes' shared base. Thus, by Cavalieri's Principle, the boxes have the same volume. Moreover, shears obviously preserve orientation, so it follows that the determinant of every shear is 1, as claimed. ♦

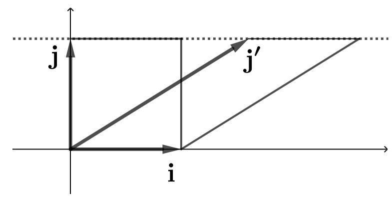

Let's make a quick geometric observation. The boxes with which we're concerned in this section should, in principle, lend themselves to simple volume calculations because any box of this sort can be understood as a slanted stack of copies of the box's base, piled up into another dimension. A parallelogram is a stack of equal line segments; a parallelepiped is a stack of equal parallelograms; a four-dimensional parallelotope is a stack of equal parallelepipeds; and so forth. This being so, any such box's "volume" is simply its height times its *base's* "volume". We can use this simple observation to derive a formula for the determinant of any *upper triangular matrix*, as we'll see next.

An **upper triangular matrix** is a square matrix whose entries below the main diagonal are all zeros. (*Lower* triangular matrices are defined analogously.) Here are some examples:

$$A = \begin{pmatrix} 5 & 6 \\ 0 & 9 \end{pmatrix}, \quad B = \begin{pmatrix} 4 & 1 & 5 \\ 0 & 2 & 1 \\ 0 & 0 & 3 \end{pmatrix}, \quad C = \begin{pmatrix} 3 & 1 & 4 & 1 \\ 0 & 5 & 9 & 2 \\ 0 & 0 & 6 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix}.$$

Every upper triangular matrix generates a grid of boxes whose volumes are easily computed. For example, consider matrix A. Its columns generate a parallelogram whose base, lying on the x-axis, is 5 units, and whose height is 9 units. Its *area* is thus $\mathbf{5} \cdot \mathbf{9} = 45$ units². (Note that the *first* entry in A's second column has no effect on the parallelogram's area; had it been 307 instead of 6, the parallelogram's area would still be 45 units².) Since this map clearly preserves orientation ($\mathbf{i}$ and $\mathbf{j}$ need not cross one another to reach their new positions $\mathbf{i}'$ and $\mathbf{j}'$), we may conclude that

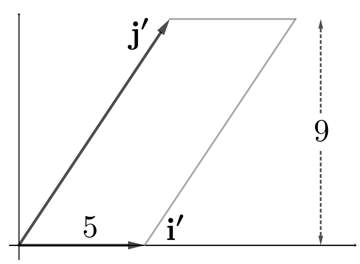

$$\det\begin{pmatrix} \mathbf{5} & 6 \\ 0 & \mathbf{9} \end{pmatrix} = \mathbf{5} \cdot \mathbf{9} = 45.$$

Will the determinant of *every* 2×2 upper triangular matrix be the product of its two diagonal entries? Yes! If both diagonal entries are positive, the same argument we just used will still hold. (Make up some examples to convince yourself of this!) If one diagonal entry is negative, then $\mathbf{i}'$ and $\mathbf{j}'$ will be oriented in such a way that the map *reverses* orientation, and hence the determinant will be negative; but the product of the diagonal entries will also be negative, so all's well (see Exercises 3a & 3b). If both diagonal entries are negative, vectors $\mathbf{i}'$ and $\mathbf{j}'$ will be oriented in such a way that the map *preserves* orientation; of course, in that case, the diagonal entries' product will be positive, so all's well in that case, too (see Exercise 3c). Finally, if either diagonal entry is zero, the "parallelogram" collapses into something *without area* (either a line segment or a point), so the determinant will be zero, as will the diagonal entries' product.

So far so simple. But can the determinant by computed so simply for *all* upper triangular matrices, or only for 2×2 ones?

Let's turn our attention to a 3×3 example, matrix B above. Its first two columns lie conveniently in the xy-plane. (It's easy to see that this will hold for *every* 3×3 upper triangular matrix.) These two columns generate a parallelogram whose area, by the logic we employed for our 2×2 matrix A, must be $4 \cdot 2$ units^2. This parallelogram is the base of the parallelepiped generated by all three columns, the third of which gives the parallelepiped's height: 3 units. This height depends, of course, exclusively on the third column's third entry. Consequently, the parallelepiped's volume must be $(4 \cdot 2) \cdot 3$ units^3, and since orientation is clearly preserved here, the determinant is positive. Thus,

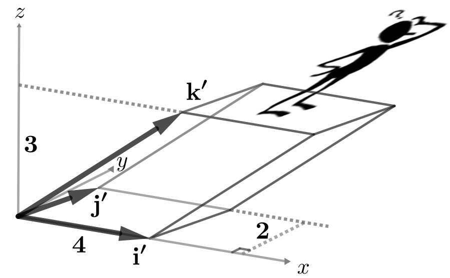

$$\det\begin{pmatrix} \mathbf{4} & 1 & 5 \\ 0 & \mathbf{2} & 1 \\ 0 & 0 & \mathbf{3} \end{pmatrix} = \mathbf{4} \cdot \mathbf{2} \cdot \mathbf{3} = 24.$$

Thus, our determinant of a 3×3 upper triangular matrix is still just the product of its diagonal entries. And indeed, the same basic argument shows that the determinant of *every* 3×3 upper triangular matrix is simply the product of its diagonal entries.

We can't depict the 4×4 case, but we can still understand it. Consider the upper triangular matrix C above. Being upper triangular, its first column lies on the x-axis, and its second lies in the xy-plane. In that plane, the first two columns generate a parallelogram of base 3 and height 5 – and thus of volume $3 \cdot 5$. That parallelogram, in turn, is the base of the parallelepiped in the xyz-hyperplane that is generated by C's first *three* columns. It has a height of 6, so its volume is $(3 \cdot 5) \cdot 6$. This parallelepiped is also the base of the *four-dimensional parallelotope* generated by all four of C's columns. The parallelotope has a height of 3, the fourth column's fourth entry. Consequently, its hypervolume is $(3 \cdot 5 \cdot 6) \cdot 3$. It follows that

$$\det\begin{pmatrix} \mathbf{3} & 1 & 4 & 1 \\ 0 & \mathbf{5} & 9 & 2 \\ 0 & 0 & \mathbf{6} & 0 \\ 0 & 0 & 0 & \mathbf{3} \end{pmatrix} = \mathbf{3} \cdot \mathbf{5} \cdot \mathbf{6} \cdot \mathbf{3} = 270,$$

so the determinant of our 4×4 matrix is indeed just the product of its diagonal entries.

Clearly, the argument above can easily be extended to 5×5 matrices, 6×6 matrices and so forth. And with minor adjustments (see exercise 4), we can use it to show that the determinant of any *lower* triangular matrix is also just the product of its diagonal entries. We summarize our hard-won results about triangular matrices (upper and lower) in a box:

Theorem. The determinant of any triangular matrix is the product of its diagonal entries.

Later in this chapter, this will help us find a method for finding the determinant of *any* square matrix M. The general method, as we'll see, will be to row-reduce M until it becomes an upper triangular matrix T, whose determinant will, thanks to our theorem, be obvious. We'll then be able to deduce $\det(M)$ from two things: $\det(T)$, and the sequence of row operations that we used to reduce M to T.

Exercises.

1. Explain geometrically (i.e. without appealing to the theorem at the end of the section) why

$$\det\begin{pmatrix} 3 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 2 \end{pmatrix} = 30.$$

2. True or false (and explain):

a) Every diagonal matrix is an upper triangular matrix.
b) Every matrix has a determinant. c) Every square matrix has a determinant.
d) Every shear matrix is a triangular matrix. e) Every triangular matrix is a shear matrix.
f) Some matrices are both upper triangular and lower triangular.
g) $\det I = 1$.

3. a) Consider the upper triangular 2×2 matrix

$$\begin{pmatrix} 2 & 4 \\ 0 & -3 \end{pmatrix},$$

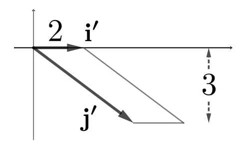

whose effect is shown at right. By our theorem on triangular matrices, its determinant should be $2(-3) = -6$. Explain geometrically (in terms of the figure) why that negative determinant makes sense. (And do you see why this will always happen for every upper triangular 2×2 matrix with a negative bottom right entry and a positive upper left one?)

b) Make up an upper triangular 2×2 matrix with a negative top left entry, but a positive bottom right one. Draw a corresponding picture and explain geometrically why the determinant should be negative. Do you see that this will be the case for every matrix of this sort?

c) Finally, consider the case when both diagonal entries are negative. Our theorem tells us that the determinant of such a matrix should be positive. That means that orientation should be preserved, so we should be able to get $\mathbf{i}$ and $\mathbf{j}$ into their new positions $\mathbf{i}'$ and $\mathbf{j}'$ in such a way that neither must cross the line containing the other. Explain, in broad terms, how to accomplish this for a specific matrix or two of your choosing.

4. a) Give a careful definition of a *lower* triangular matrix.

b) Make up your own 2×2 lower triangular matrix with positive entries on the diagonal and explain geometrically *why* its determinant is the product of its diagonal entries.
c) Same story, but with a matrix containing a negative (or two) on the diagonal.
d) Same story, but with a 3×3 lower triangular matrix. You can stick to all positive diagonal entries.
e) Again, but with a 4×4 lower triangular matrix.

5. How can we understand a 4-dimensional hypercube? One way is to build up to it dimension by dimension, through repeated perpendicular "extrusions".

Start with a 0-dimensional object (a point) and push it one unit east, tracing out a 1-dimensional object: a "stack of points", a *line segment*. We then push this segment one unit north, tracing out a 2-dimensional object: a "stack of segments", a unit *square*. When we push this a unit upwards (perpendicular to the plane it lies in), it generates a 3-dimensional object: a "stack of squares", a unit *cube*. Now we must imagine pushing that unit cube one unit in a mysterious spatial direction that is somehow perpendicular to our entire three-dimensional space. Doing so yields a "stack of cubes" (stacked along a fourth dimension, with successive cubes touching at *all* their corresponding points). This is a unit *four-dimensional hypercube*, sometimes called a *tesseract*.

a) How many 0-dimensional vertices, 1-dimensional edges, and 2-dimensional faces does a tesseract have? Explain your answers *geometrically*.

b) Four dimensional figures have not only 0-dimensional vertices, 1-dimensional edges, and 2-dimensional faces, but also 3-dimensional "cells". The tesseract's cells are cubes. How many such cells does it have?

The Determinant: Properties

For there are many properties in it that, if universally known, would habituate its use and make it more in request with us than with the Turks themselves.

- Awister, quoted by Thomas De Quincey in *Confessions of an English Opium Eater* ("To the Reader").

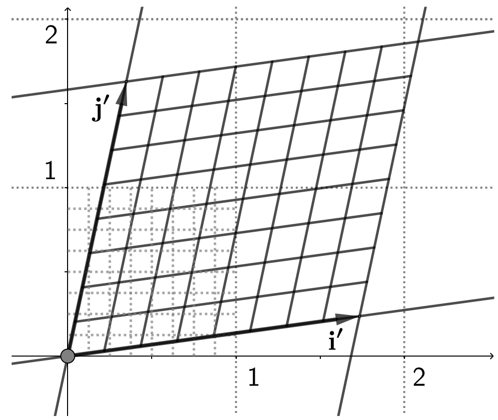

If we transform $\mathbb{R}^n$ with an $n \times n$ matrix A, it's not too hard to see that the n-dimensional volume of *every* object in $\mathbb{R}^n$ (not just the "boxes") will be scaled by a factor of $|\det A|$. To see why, we'll need to don our integral calculus glasses.

Begin by noting that each "standard box" in the standard grid (generated by the standard basis vectors) has volume 1. By the determinant's definition, A transforms each standard box into a box of volume $|\det A|$. In the figure at right, for example, the box with sides $\mathbf{i}'$ and $\mathbf{j}'$ has volume $|\det A|$. (Of course, in this two-dimensional case, "volume" is *area*.)

So far so simple. Next, imagine chopping each standard box into k equal "boxlets". (In the figure, I've made $k = 64$.) Each boxlet has volume $\mathbf{1/k}$, and since the k *transformed* boxlets (whose total area is $|\det A|$) are equal, each transformed boxlet must have volume $\mathbf{|\det A|/k}$. It follows that $|\det A|$ serves as a volume-scaling factor not just for full standard boxes, but for our boxlets, too. This will hold regardless of whether we use 64 boxlets, 1000 boxlets, 10^{100} boxlets, or – letting the spirit of calculus guides us – *infinitely many* infinitesimally small boxlets. In all cases, $|\det A|$ acts as a volume-scaling factor when A acts on the boxlets. We'll can now think of any old figure in $\mathbb{R}^n$ (a hand, a sheep, a 6-dimensional hypersphere, or whatnot) in a pixelated manner – chopped into a collection of infinitesimal boxlets. Since A scales the volume of each infinitesimal boxlet by a factor of $|\det A|$, it clearly scales *the entire figure's* volume by $|\det A|$, too.

> **Property 1.** The determinant is a volume-scaling factor: When an an $n \times n$ matrix A transforms objects in $\mathbb{R}^n$, it scales their n-dimensional volumes by $|\det A|$.

The idea of the determinant as a volume-scaling factor is tremendously important in vector calculus. In ordinary freshman calculus, we work with functions from $\mathbb{R}$ to $\mathbb{R}$. Locally (on an infinitesimal scale), such functions' graphs are straight, so locally, the functions themselves are *linear*. We can thus describe any such function's local behavior with one number: the slope of the line that it resembles at that point. And as everyone knows, we call that number the function's *derivative* at that point. In vector calculus, however, we consider functions from $\mathbb{R}^n$ to $\mathbb{R}^m$. Locally, such functions resemble *linear transformations*. Thus, describing the local behavior of a nonlinear function from $\mathbb{R}^n$ to $\mathbb{R}^m$ requires more than a number; it requires a matrix. Accordingly, the derivative of such a function is, at any point in its domain, a *matrix*. When $m = n$, the matrix is square, and thus it has a determinant. This determinant plays a role when we change variables in integration, doing the higher-dimensional analogue of the familiar "u-substitution". You can look forward to studying how this works in a future course.

This volume-scaling property helps us understand a vital algebraic property of determinants: The determinant of a product is… the product of the determinants.

> **Property 2.** If A and B are any two $n \times n$ matrices, then
> $$\det(AB) = \det(A)\det(B).$$

For example, suppose A and B are square matrices that scale volumes by factors of 5 and 3 respectively. I claim that AB must scale volumes by a factor of 15. Why? Well, multiplying matrices corresponds to composing linear maps, so AB first scales volumes by 3 (through B's action), then scales the results by 5 (through A's action). Hence, the net effect is that AB scales volumes by a factor of $3 \cdot 5 = 15$, as claimed. This same analysis clearly holds for *any* two matrices with nonnegative determinants. A little thought shows that it also holds when a negative determinant (or two) is involved. For example, if M scales volumes by 2 *and reverses orientation* (so that $\det(M) = -2$) while N scales volumes by 5 and preserves orientation (so that $\det(N) = 5$), then MN obviously scales volumes by 10 *and reverses orientation*. That is, $\det(MN) = -10$, which is, of course, the product of $\det(M)$ and $\det(N)$. I'll leave it to you to convince yourself that this property also holds when *both* matrices have negative determinants.

Our next two properties concern some relationships between determinants and **inverse matrices.** Recall from Chapter 3, Exercise 23 that the inverse of A, which we denote A^{-1}, is the matrix that "undoes" A's action. (That is, if $\boldsymbol{A}\mathbf{v} = \mathbf{w}$, then $A^{-1}\mathbf{w} = \mathbf{v}$.) It follows that the product of A and A^{-1} (in either order) is the identity matrix I. In that same exercise, we saw that not every matrix is invertible. For a linear map (or the matrix representing it) to be invertible, it must be "one-to-one"; that is, it must always take distinct points from the domain to distinct points in the range. (After all, if A were to map both $\mathbf{v_1}$ and $\mathbf{v_2}$ to $\mathbf{w}$, then what would $A^{-1}\mathbf{w}$ be?)

See if you can justify this next property on your own before reading that explanation that follows it.

> **Property 3.** If A is any invertible matrix, then
> $$\det(A^{-1}) = \frac{1}{\det A}.$$

If A scales volumes by a factor of 7, then A^{-1}, undoing that action, must scale volumes by a factor of $1/7$. Similarly, if B scales volumes by $2/3$ and reverses orientation, then B^{-1} will need to scale volumes by $3/2$ and reverse orientation. If you understand that much, you should see why Property 3 always holds.

But wait a minute. What if $\det A = 0$? In that case, our formula involves division by 0. But fear not: This can't happen. If $\det A = 0$ (and A is $n \times n$), A scales every n-dimensional volume by a factor of zero. That doesn't mean that A crushes all of $\mathbb{R}^n$ into $\mathbf{0}$, but it does mean that some dimensional collapse must occur. (For instance, a 3×3 matrix that crushes $\mathbb{R}^3$ down into a plane has determinant 0; yes, *something* comes out the other end of the mapping, but no 3-dimensional volume survives.) Whenever dimensional collapse occurs, many distinct points in the domain end up being mapped to the same point in the range.* Thus the map isn't one-to-one, so its matrix can't be invertible, so Property 3 wouldn't be applicable. We've discovered something important in passing: *If a matrix has determinant 0, it isn't invertible*.

* "Dimensional collapse" means that the map's rank less than its domain's dimension. So by the rank-nullity theorem, the kernel's dimension is *at least 1*. Hence, the map's kernel has infinitely many points, all of which get mapped to $\mathbf{0}$.

On the other hand, it's easy to see that any $n \times n$ matrix with a *non*zero determinant *is* invertible. If the determinant is nonzero, then although n-dimensional volume may be scaled, it at least *survives* as n-dimensinal volume, which implies that there's no dimensional collapse. All n dimensions remain in the map's output, which will be pervaded by a "clean grid" generated by n linearly independent vectors. It should be intuitively clear in your mind's eye that any such transformation maps distinct points to distinct points. In other words, any such transformation is one-to-one, and therefore the matrix that represents it is invertible.

Combining the results of the previous two paragraphs yields our next trophy: A matrix is invertible if and only if its determinant is nonzero.

Property 4. Matrix A is invertible $\Leftrightarrow$ $\det A \neq 0$

So far, we have built up the determinant's most important properties just by thinking geometrically. But how do we actually compute a determinant? We'll begin with one special case: The determinant of a 2×2 matrix.

Property 5. (2×2 Determinant formula)

$$\det\begin{pmatrix} a & c \\ b & d \end{pmatrix} = ad - cb$$

Proof. The figure at right shows the parallelogram determined by the 2×2 matrix's columns. The determinant is the parallelogram's area, which is the large rectangle's area minus the areas of the four right triangles and the two small rectangles in the corners. Or in symbols,

$$\det\begin{pmatrix} a & c \\ b & d \end{pmatrix} = (a+c)(d+b) - 2\left(\frac{1}{2}ab\right) - 2\left(\frac{1}{2}cd\right) - 2(bc).$$

The right-hand side reduces, as you should verify, to $ad - cb$. ■

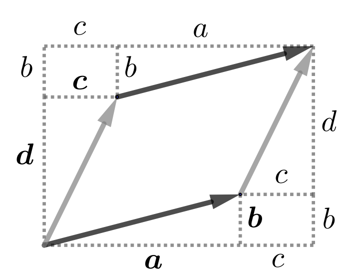

For the preceding proof to be fully rigorous, we'd need to consider some other possible configurations. (For example, what if the vectors make an obtuse angle? What if the vectors are arranged so that the determinant is negative?) You'll dispose of some such cases in Exercise 9.

Exercises.

6. We justified Property 3 geometrically. Now explain it with algebra. [**Hint**: By definition, $A^{-1}A = I$.]

7. Consider the following matrices:

$$A = \begin{pmatrix} 2 & 3 \\ 1 & 2 \end{pmatrix}, \qquad B = \begin{pmatrix} 8 & 3 \\ 7 & 2 \end{pmatrix}, \qquad C = \begin{pmatrix} 3 & 6 \\ 2 & 4 \end{pmatrix}, \qquad D = \begin{pmatrix} -4 & 3 \\ -5 & 4 \end{pmatrix}, \qquad E = \begin{pmatrix} 9 & 7 \\ 3 & 9 \end{pmatrix}.$$

a) Compute their determinants. b) Which of the preceding matrices are invertible?
c) Remind yourself *why* a matrix whose determinant is zero cannot be inverted.
d) Which of the preceding matrices represent linear transformations that preserve area?
e) Which of them reverse the orientation of $\mathbb{R}^2$?
f) Find the determinants of the *inverses* of the invertible matrices. Be sure that you understand why your answers make geometric sense. [*Note: You need not find the inverse matrices themselves, just their determinants*.]
g) Find B^2 and use it to compute $\det(B^2)$. Be sure that you understand *why* the relationship between $\det B$ and $\det(B^2)$ makes geometric sense.
h) Without finding AB, find $\det(AB)$. Then find AB and compute its determinant directly to verify your answer.

8. (Extending the **Invertible Matrix Theorem**) In Chapter 4, Exercise 40, you found that for any $n \times n$ matrix A, statements A - I below are logically equivalent, meaning that they all stand or all fall together. Convince yourself that statement J can join this growing list, too. (The list will grow again in Exercise 19.)

a) A is invertible. b) $\mathrm{rref}(A) = I$. c) $A\mathbf{x} = \mathbf{b}$ has a *unique* solution for every vector $\mathbf{b}$.
d) A's columns are linearly independent. e) A's columns span $\mathbb{R}^n$. f) A's columns constitute a basis for $\mathbb{R}^n$.
g) $\ker(A) = \mathbf{0}$. h) $\mathrm{im}(A) = \mathbb{R}^n$. i) $\mathrm{rank}(A) = n$. j) $\det A \neq 0$.

9. We justified our 2×2 determinant formula

$$\det \begin{pmatrix} a & c \\ b & d \end{pmatrix} = ad - bc$$

with a nice orientation-preserving matrix whose columns determined a parallelogram lying in the first quadrant. Given a less nice setup – such as one of those indicated in the figures below (the first of which has a parallelogram stretching into the second quadrant, the second of which corresponds to an orientation-*reversing* map) – does the 2×2 determinant formula still hold? Yes it does. Demonstrate this. Once you've grasped those two cases, you'll likely be convinced that the formula does indeed hold for all possible 2×2 matrices.

[**Hint for Part A:** *Entries in a matrix can be negative, but lengths must, of course, be positive*.]

a)

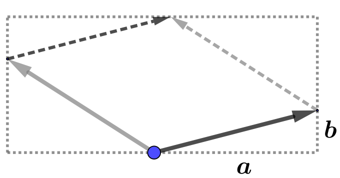

b)

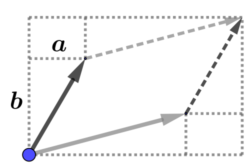

10. In Exercise 29 of Chapter 4, you established a quick formula for the inverse of a 2×2 matrix

$$A = \begin{pmatrix} a & c \\ b & d \end{pmatrix}.$$

You can now rewrite that quick formula for A^{-1} in a manner that is slightly easier to remember. Do so.

11. Property 2 extends to products of three or more matrices. For example, it's also true that

$$\det(ABC) = \det(A)\det(B)\det(C).$$

However many matrices there are, *the determinant of a product is the product of the determinants*. Explain why. Then use this idea to compute $\det(D^{1000})$, where D is the matrix of that name in Exercise 7.

12. Determinants often give us surprising leverage in proofs - even when we are proving things unrelated to volume. In this exercise, for instance, determinants will help us prove that a square matrix A's "**left inverse"** (i.e. a matrix B such that $BA = I$) necessarily works as a **"right inverse"** as well (that is, $AB = I$).

The preceding fact about "one-sided inverses" might seem obvious, but recall that when I introduced inverse matrices (Chapter 3, Exercise 23), I did so as follows:

*The matrix that undoes the action of a square matrix A is called the **inverse matrix** of A, and is denoted A^{-1}. Thus, by definition, $A^{-1}A = I = AA^{-1}$.*

In a footnote for that sentence, I added the following:

If we wish to show that a matrix B is in fact A^{-1}, we must – by this definition – verify two separate things: $BA = I$ and $AB = I$. But we'll prove later (Ch. 5, Exercise 12) that each of these things implies the other, so to verify both, we just need to verify one.

Well, here we are. Let's get to work.

a) Convince yourself that each step in the following argument holds, justifying each step where appropriate.

Claim. If A is an $n \times n$ matrix and $BA = I$, then $AB = I$, too.

Proof. If $BA = I$, then $\det(BA) = \det I$. Thus, $\det(B)\det(A) = 1$. From this it follows that $\det B \neq 0$. Hence, B doesn't collapse any dimensions. In particular, $\ker(B) = \mathbf{0}$, a fact that we'll use shortly.

Since $BA = I$, it follows that $BAB = B$, or equivalently, $BAB - B = Z$, where Z is the zero matrix, representing the "zero map". Rewriting the last equation as $B(AB - I) = Z$, we deduce that $B(AB - I)$ sends *all* vectors in $\mathbb{R}^n$ to $\mathbf{0}$. But since B sends only $\mathbf{0}$ to $\mathbf{0}$, it must be the case that all outputs of $(AB - I)$ are already $\mathbf{0}$. In other words, matrix $(AB - I)$ must represent the zero map. Accordingly, $AB - I = Z$. Adding I to both sides of this equation, we find that $AB = I$, as claimed. ■

b) Prove the converse claim: If A is an $n \times n$ matrix and $AB = I$, then $BA = I$, too. (And hence, $B = A^{-1}$.)

The moral of this exercise is that if we ever show that some square matrix A has a "one-sided inverse", then it must in fact be the full inverse A^{-1}. We are therefore mercifully free from the burden of having to distinguish between "left inverses" and "right inverses" of square matrices.

Elementary Matrices (and their Determinants)

Elementary, my dear Watson!

- not Sherlock Holmes*

An **elementary matrix** is one that can be obtained from the *identity* matrix by applying one row operation. Accordingly, the following are all examples of elementary matrices:

$$\begin{pmatrix}1&0&0\\0&0&1\\0&1&0\end{pmatrix}, \qquad \begin{pmatrix}1&0&0\\0&3&0\\0&0&1\end{pmatrix}, \qquad \begin{pmatrix}1&0&0\\0&1&0\\0&2&1\end{pmatrix}.$$

The first matrix was obtained by swapping I's first two rows. The second was obtained by scaling I's middle row by 3. The third was obtained by adding two copies of I's middle row to its bottom row.

Row operations are violent but effective surgery in which we slice open a matrix to alter its innards. Elementary matrices can deliver all the results of "row operation surgery" but in a less invasive manner. The idea: Instead of doing a row operation on a given matrix, we left-multiply it by the elementary matrix obtained from I by that row operation. For example, instead of scaling a matrix's 2nd row by 3 like this,

$$\begin{pmatrix}2&5&7\\1/3&2/3&3\\4&1&4\end{pmatrix} \quad \times 3 \quad \begin{pmatrix}2&5&7\\1&2&9\\4&1&4\end{pmatrix},$$

we can left-multiply the given matrix by the elementary matrix obtained by scaling I's 2nd row by 3:

$$\begin{pmatrix}1&0&0\\0&3&0\\0&0&1\end{pmatrix}\begin{pmatrix}2&5&7\\1/3&2/3&3\\4&1&4\end{pmatrix} = \begin{pmatrix}2&5&7\\1&2&9\\4&1&4\end{pmatrix},$$

Obviously, doing the row operation directly is simpler in a narrow pragmatic sense, but the elementary matrix approach has a crucial advantage for theoretical work: It replaces a messy ad hoc procedure with pristine matrix *algebra*. It replaces bookkeeping scratchwork with an *equation* – an object to which we can then apply algebraic rules. And this algebra, in turn, will lead us to a simple algorithm for computing the determinant of *any* square matrix, not just those special few whose determinants we've found so far (shear, triangular, and 2×2 matrices). But before we can do that, we'll need to learn one last thing about elementary matrices: their determinants. This turns out to be easy, since, as we'll show in the next three paragraphs, every elementary matrix represents one of three simple geometric operations.

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

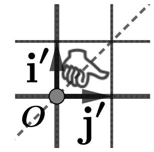

First, any elementary matrix obtained by swapping rows of I represents a *reflection*. Why? Swapping I's j^{th} and k^{th} *rows* is equivalent to swapping its j^{th} and k^{th} *columns* – an operation with clear geometric meaning: It swaps the j^{th} and k^{th} standard basis vectors' positions, *reflecting* them (in their common plane) across the line that bisects the right angle between them. Hence, all elementary matrices of this first type are reflections, as claimed.† Hence, the determinants of all "row swap elementary matrices" are $-\mathbf{1}$, since reflections preserve volumes but reverse orientation.

* This famous expression never occurs in any of Arthur Conan Doyle's many Sherlock Holmes stories and novels. One steeped in Sherlockiana might call it "the curious incident of the Holmes exclamation in the night".

† Strictly speaking, the linear map corresponding to an $n \times n$ matrix of this sort is a reflection in an $(n-1)$-dimensional "mirror". If we swap columns j and k, the "mirror" is the subspace spanned by the vector $\mathbf{e}_j + \mathbf{e}_k$ *and* the $(n-2)$ fixed basis vectors.

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

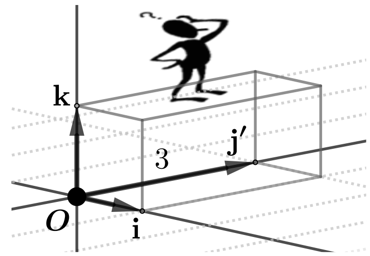

Second, any elementary matrix that we obtain by scaling one of I's rows is a *stretch along an axis* (with a reflection if the scalar is negative). Why? Scaling I's j^{th} row by c is equivalent to scaling its j^{th} *column* by c – an operation with a clear geometric interpretation: It alters the standard grid by stretching the j^{th} basis vector by a factor of $|c|$ (and reversing its direction if $c < 0$), while the others stay put. The transformed grid thus consists of rectangular boxes, each of volume $(|c| \cdot 1 \cdot 1 \cdots 1) = |c|$. Orientation will be preserved or reversed according to c's algebraic sign. Thus, all elementary matrices of this second type represent stretches (sometimes with an accompanying reflection). Hence, the determinant of any "row scale elementary matrix" is the row scale factor $\boldsymbol{c}$ itself.

Third and last, any elementary matrix that we obtain by adding a multiple of one of I's rows to another of its rows represents a *shear*. To see why, suppose that we add c copies of I's j^{th} row to its k^{th} row. The resulting elementary matrix will look like an identity matrix in which someone has tampered with one column, changing one of its zeros to some other number. As we saw in Example 4 in this chapter's first section, a matrix of that form is a *shear matrix*. And as we discussed in that same example, the determinant of any shear matrix is 1. Thus, any elementary matrix of this third kind has a determinant of $\mathbf{1}$.

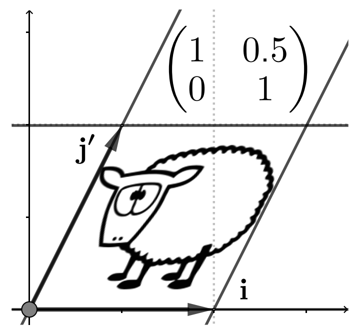

To sum up: Any elementary matrix that we can use to…

- *Swap rows* represents a **reflection**. Hence, its determinant is $-\mathbf{1}$.
- *Scale a row* (by c) represents a **stretch** (by a factor of c; if $c < 0$, there's a **reflection**, too). Hence, its determinant is $\boldsymbol{c}$.
- *Add a multiple of one row to another* is a **shear***.* Hence, its determinant is $\mathbf{1}$.

After a few exercises, we'll finally be ready to turn to our last major problem of the chapter: deriving a method for computing the determinant of *any* square matrix.

Exercises.

13. For each of the following row operations, find the elementary matrix that carries it out. Verify your matrices by left-multiplying them against the telephone matrix at right.

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

a) Swap rows 2 and 3.
b) Scale row 2 by 4.
c) Subtract 4 copies of row 1 from row 2.

14. What are the determinants of the three elementary matrices you found in the previous problem?

15. Suppose E_1, E_2, E_3 are elementary matrices; E_1 swaps rows, E_2 scales a row by -3, and E_3 adds a multiple of one row to another. Now suppose that $E_3E_2E_1A$ is a triangular matrix T, whose determinant is 5. What, if anything, can we conclude about $\det(A)$? [**Hint**: *Recall Exercise 11.*]

Computing Determinants (by Row Reduction)

At long last, we can establish a simple algorithm for computing any square matrix's determinant.

Determinant Algorithm.

Use Gaussian elimination to reduce the given matrix A to a *triangular* matrix T, and keep track of…

- the number s of row swaps you use
- the product p of all the factors by which you multiply rows.

We then have

$$\det A = \frac{\det T}{(-1)^s p}$$

(As discussed earlier in this chapter, $\det T$ is the product of T's diagonal entries.)

If you solved Exercise 15 on the previous page, you've already understood, in essence, why this algorithm works. But let's spell out the details.

Proof. To reduce A to a triangular matrix T, we perform row operations. Let $E_1, E_2, \dots, E_m$ be their corresponding elementary matrices so that $E_m \cdots E_2 E_1 A = T$. Taking determinants of both sides (and recalling Exercise 11) yields $\det(E_m) \cdots \det(E_2)\det(E_1)\det(A) = \det(T)$. Or equivalently,

$$\det A = \frac{\det T}{\det(E_m) \cdots \det(E_2)\det(E_1)}.$$

By our work in the previous section, we can evaluate that denominator. There are s "row swap elementary matrices" down there, so we know that the product of their determinants is $(-1)^s$. Next, the "row scale elementary matrices": We know that the product of their determinants is p. This leaves the elementary matrices that add multiples of one row to another. These are shear matrices, so their determinants are all 1. It follows that $\det A = \det T/(-1)^s p$, as claimed. ■

An example will make the idea clearer.

Example 1. Find the determinant of $\begin{pmatrix} 2 & -3 & 7 \\ 4 & 9 & -3 \\ 2 & 4 & -2 \end{pmatrix}$.

Solution. We'll reduce the matrix to triangular form, keeping track of row swaps and the product of any factors we use to scale rows. Here's one way to accomplish this row reduction:

$$\begin{pmatrix} 2 & -3 & 7 \\ 4 & 9 & -3 \\ 2 & 4 & -2 \end{pmatrix} \begin{matrix} R_1 \leftrightarrow R_3 \end{matrix} \begin{pmatrix} 2 & 4 & -2 \\ 4 & 9 & -3 \\ 2 & -3 & 7 \end{pmatrix} \begin{matrix} \times 1/2 \\ \\ \\ \end{matrix} \begin{pmatrix} 1 & 2 & -1 \\ 4 & 9 & -3 \\ 2 & -3 & 7 \end{pmatrix} \begin{matrix} \\ -4R_1 \\ -2R_1 \end{matrix} \begin{pmatrix} 1 & 2 & -1 \\ 0 & 1 & 1 \\ 0 & -7 & 9 \end{pmatrix} \begin{matrix} \\ \\ +7R_2 \end{matrix} \begin{pmatrix} 1 & 2 & -1 \\ 0 & 1 & 7 \\ 0 & 0 & 16 \end{pmatrix}.$$

This row reduction entailed 1 row swap. The product of all scaling factors (only one) was $1/2$. Since the process yielded a triangular matrix whose determinant is 16, we conclude that

$$\det A = \frac{16}{(-1)^1(1/2)} = \mathbf{-32}.$$ ♦

And that's that. With that technique in hand, you can now compute the determinant of any square matrix whatsoever. All you need is Gaussian elimination and some careful bookkeeping. Incidentally, we can use this algorithm to recover the quick formula for 2×2 determinants that we derived geometrically.

Example 2. Use row-reduction to rederive the "quick formula" for a 2×2 determinant.

Solution. Start with an expression for a general 2×2 matrix, and row reduce it to triangular form:

$$\begin{pmatrix} a & c \\ b & d \end{pmatrix} \underset{}{\times (1/a)} \begin{pmatrix} 1 & c/a \\ b & d \end{pmatrix} \underset{-bR_1}{} \begin{pmatrix} 1 & c/a \\ 0 & (ad-bc)/a \end{pmatrix}.$$

This reduction required no row swaps, and the product of all scaling factors was $1/a$. It yielded a triangular matrix whose determinant is $(ad - bc)/a$, so we conclude that

$$\det \begin{pmatrix} a & c \\ b & d \end{pmatrix} = \frac{(ad-bc)/a}{1/a} = ad - bc. \qquad \blacklozenge$$

We'll end with a surprising theorem: Any matrix and its transpose have the same determinant.

Theorem. $\det(A^T) = \det(A)$ for every square matrix A.

Proof. Reduce A to triangular form T. Let $E_1, E_2, \dots, E_m$ be the elementary matrices corresponding to the row operations, so that $E_m \cdots E_2 E_1 A = T$. Taking determinants of both sides (and recalling Exercise 11) yields $\det(E_m) \cdots \det(E_2) \det(E_1) \det(A) = \det(T)$. Equivalently,

$$\det A = \frac{\det T}{\det(E_m) \cdots \det(E_2) \det(E_1)}.$$

We'll now demonstrate that $\det A^T$ has this same form. Go back to $E_m \cdots E_2 E_1 A = T$ and take the *transpose* of both sides. By a property you proved in Chapter 3 (Exercise 26e), this yields

$$A^T E_1^T E_2^T \cdots E_m^T = T^T.$$

Now take determinants of both sides, using the property that the determinant of a product is the product of the determinants, then solve for $\det(A^T)$. Doing so, we find that

$$\det(A^T) = \frac{\det(T^T)}{\det(E_1^T) \det(E_2^T) \cdots \det(E_m^T)}.$$

In fact, we'll soon be able to erase all those transpose superscripts. To see why, first observe that since T is an upper triangular matrix, T^T is *lower* triangular. The determinants of both T and T^T are thus the products of their diagonal entries. Moreover, as transpose "mates", T and T^T have the *same* diagonal entries. Hence, $\det(T^T) = \det T$. Next, consider those elementary matrices. Transposition leaves the first two types (row scale and row swap) unchanged, so it doesn't change their determinants. As for the third type of elementary matrix, these are triangular with 1s on their main diagonal, and transposing any such matrix turns it into a matrix of the same sort. Since all such matrices have determinants of 1, transposition clearly preserves their determinants, too. We've now shown that $\det(E_i^T) = \det(E_i)$ for all i. The two highlighted equations indicate that we can indeed erase all the transpose superscripts from the right-hand side of our expression for $\det(A^T)$. Doing so and then rearranging the denominator's factors, we obtain the right-hand side of our earlier expression for $\det A$. Thus, $\det(A^T) = \det A$, as claimed. ■

Exercises.

16. Use row-reduction to compute the determinants of the following matrices.

a) $\begin{pmatrix} 3 & 2 & 1 \\ 0 & 3 & -6 \\ 0 & 2 & -2 \end{pmatrix}$ b) $\begin{pmatrix} 5 & 3 & 7 \\ 1 & 0 & 0 \\ 0 & 1 & 2 \end{pmatrix}$ c) $\begin{pmatrix} 2 & 4 & 6 \\ 3 & -4 & 8 \\ 0 & 2 & 5 \end{pmatrix}$ d) $\begin{pmatrix} 1 & 7 & 1 \\ -1 & 0 & 2 \\ 3 & -1 & -3 \end{pmatrix}$ e) $\begin{pmatrix} 3 & 2 & 6 \\ 4 & 8 & 4 \\ 1 & 0 & 2 \end{pmatrix}$

f) $\begin{pmatrix} 0 & 3 & 2 & 4 \\ 1 & 1 & 2 & 1 \\ 1 & 1 & 5 & 2 \\ 1 & 1 & 3 & 3 \end{pmatrix}$ g) $\begin{pmatrix} 1 & -1 & 2 & -2 \\ -1 & 2 & 1 & 6 \\ 2 & 1 & 14 & 10 \\ -2 & 6 & 10 & 33 \end{pmatrix}$ h) $\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 2 & 5 \\ 1 & 1 & 3 & 6 \\ 1 & 1 & 4 & 7 \end{pmatrix}$

17. Explain geometrically why the matrix in Exercise 16h has a determinant of 0.

18. (Quick Formula for 3×3 determinants) In Example 2, we re-derived the quick formula for a 2×2 determinant. As it happens, we can use the same idea that we used there to derive a quick formula for 3×3 determinants. The algebra involved is basic but tedious (try it), so I'll dispense with the details and just tell you the punchline:

$$\det \begin{pmatrix} a & d & g \\ b & e & h \\ c & f & i \end{pmatrix} = aei + dhc + gbf - ahf - dbi - gec.$$

To find a mnemonic device for this result, imagine lines through the matrix, acting like objects in old video games: When a line exits the matrix through one side, it re-enters on the opposite side, proceeding in the same direction. Thus, a line beginning at d and travelling southeast passes through h and then… c. Bearing this in mind, look again at the determinant's six terms: The first three (aei, dhc, gbf) are "spelled out" by three lines starting at each top-row entry, sloping southeast. The next three terms (which are *subtracted*) are spelled out by three lines starting at each top-row entry, but now sloping south*west*. Having linked those six terms to those six lines, we can quickly compute a 3×3 matrix's determinant by mentally following the lines, writing down the six products (with plusses and minuses in the right places), and then adding them up. For example,

$$\det \begin{pmatrix} 1 & 3 & -1 \\ 2 & 4 & 0 \\ 5 & 1 & 6 \end{pmatrix} = 24 + 0 + (-2) - 0 - 36 - (-20) = \mathbf{6}.$$

Note well: This quick formula works only for the special case of 3×3 determinants! Don't try to use an analogue of it for a matrix of any other size. It won't work.

a) Verify that the determinant we just computed really is 6 by *re*computing it via row reduction.
b) Use this quick formula to re-compute the determinants of the 3×3 matrices in Exercise 16.

19. (Extending the **Invertible Matrix Theorem**) In Exercise 8, you saw that statements A - J in the list below were equivalent statements about an $n \times n$ matrix A. Explain why we can add K, L, and M to the list:

a) A is invertible. b) $\text{rref}(A) = I$. c) $A\mathbf{x} = \mathbf{b}$ has a *unique* solution for every vector $\mathbf{b}$.
d) A's columns are linearly independent. e) A's columns span $\mathbb{R}^n$. f) A's columns constitute a basis for $\mathbb{R}^n$.
g) $\ker(A) = \mathbf{0}$. h) $\text{im}(A) = \mathbb{R}^n$. i) $\text{rank}(A) = n$. j) $\det A \neq 0$.
k) A's *rows* are linearly independent. l) A's rows span $\mathbb{R}^n$. m) A's rows constitute a basis for $\mathbb{R}^n$.

To reiterate, the moral of the invertible matrix theorem is that square matrices come in two types: *invertible* matrices (which satisfy all twelve of those conditions) and *noninvertible* matrices (which satisfy none of them).

20. Do the *rows* of $\begin{pmatrix} 1 & 2 & 3 \\ 0 & 1 & 4 \\ 1 & 1 & 2 \end{pmatrix}$ span $\mathbb{R}^3$? [**Hint**: *Use the previous exercise*.]

21. Suppose a square matrix has a column consisting entirely of zeros. What can we say about its determinant? Why? (And what, if anything, could we say about a matrix containing a *row* of all zeros?)

Computing Determinants (by Cofactor Expansion)

In this last section, we'll discuss *cofactor expansion*, an alternate algorithm for computing determinants. Though significantly less efficient than the row-reduction algorithm, cofactor expansion is, bizarrely, given pride of place in most linear algebra textbooks. Despite its relative inefficiency (which I'll discuss at the end of this section), cofactor expansion is at least algebraically interesting, and it's useful in certain special conditions – particularly when a matrix has a row or column consisting mainly of zeros.

We'll sneak up on it by reexamining the quick formula for 3×3 determinants we met in Exercise 18:

$$\det\begin{pmatrix} a & d & g \\ b & e & h \\ c & f & i \end{pmatrix} = aei + dhc + gbf - ahf - dbi - gec.$$

First, let's rearrange the expression on the right as follows: $\boldsymbol{a}(ei - hf) - \boldsymbol{d}(bi - hc) + \boldsymbol{g}(bf - ec)$. Consider those three terms. Each has two factors. What can we say about the two factors of each term? After some hard staring, we can discern a pattern: Each term's first factor (in boldface) is an element from the matrix's top row. Much more subtly, each term's second factor (in parentheses) is the determinant of the 2×2 matrix that we "expose" by mentally deleting the row and the column containing the *first* factor. (Verify this.) Interesting! But why is the middle term alone ***subtracted***? Is this part of a pattern as well? We can discover the answer by experimenting a bit more.

Our first experiment is inspired by a nagging aesthetic blemish: Why should a matrix's *first* row have special status with respect to the determinant? Could we rearrange our original six-term expression for the determinant to emphasize the *second* row? Well, let's play around (a cherished mathematical activity) and see what we discover. Given enough play time, you'd find that we can rearrange the original six-term determinant expression from our "quick formula" as follows: $-\boldsymbol{b}(di - gf) + \boldsymbol{e}(ai - gc) - \boldsymbol{h}(af - dc)$. The same basic pattern holds: Each term's first factor is drawn from a particular row, and its second factor is the determinant of the 2×2 matrix we get by crossing out the row and column of the first factor. Nice! But curiously, the sign pattern has been reversed. Now it's the two *outer* terms that are subtracted. Why? Maybe we'll gain further insight if we play this same game with the *third* row? Only one way to find out.

Rearranging the original six-term expression again, we get $\boldsymbol{c}(dh - ge) - \boldsymbol{f}(ah - gb) + \boldsymbol{i}(ae - db)$. Reassuringly, our basic pattern continues to hold: Each term's second factor is still the determinant of the matrix that we get by nixing the first factor's row and column. But the signs have reverted to the original $+ \; - \; +$ pattern that we saw when we "expanded" our determinant along the first row. Why?

After pondering this for a while, you may begin to feel another source of aesthetic unease. We've now seen, with satisfying symmetry, that one row is as good as another as far as our basic pattern is concerned, and yet… why should *rows* matter more than *columns* with respect to determinants? After all, we proved in the previous section that transposing a matrix – turning its rows into columns and vice versa – doesn't change its determinant. Well, can we do unto the *columns* what we've done to the rows? Let's see.

We can in fact rewrite our original determinant in the form $\boldsymbol{a}(ei - hf) - \boldsymbol{b}(di - gf) + \boldsymbol{c}(dh - ge)$. Now each term's first factor is from the first *column*, and the second is the expected 2×2 determinant. In fact, we can play this same game with the second column, $-\boldsymbol{d}(bi - hc) + \boldsymbol{e}(ai - gc) - \boldsymbol{f}(ah - gb)$, or the third, $\boldsymbol{g}(bf - ec) - \boldsymbol{h}(af - dc) + \boldsymbol{i}(ae - db)$. Each time, everything works out as we'd expect. Not only have we vindicated our aesthetic sense of symmetry, but we've now also gathered enough data (our six different "expansions" of the determinant along each row and column) to crack the code of the alternating signs… after some more hard staring and thinking.

We've seen that the following are all equivalent to our six-term expression for a 3×3 determinant:

$$\begin{aligned}
&\boldsymbol{a}(ei - hf) - \boldsymbol{d}(bi - hc) + \boldsymbol{g}(bf - ec) \\
&-\boldsymbol{b}(di - gf) + \boldsymbol{e}(ai - gc) - \boldsymbol{h}(af - dc) \\
&\boldsymbol{c}(dh - ge) - \boldsymbol{f}(ah - gb) + \boldsymbol{i}(ae - db) \\
&\boldsymbol{a}(ei - hf) - \boldsymbol{b}(di - gf) + \boldsymbol{c}(dh - ge) \\
&-\boldsymbol{d}(bi - hc) + \boldsymbol{e}(ai - gc) - \boldsymbol{f}(ah - gb) \\
&\boldsymbol{g}(bf - ec) - \boldsymbol{h}(af - dc) + \boldsymbol{i}(ae - db)
\end{aligned}$$

$$\begin{pmatrix} \boldsymbol{a} & \boldsymbol{d} & \boldsymbol{g} \\ \boldsymbol{b} & \boldsymbol{e} & \boldsymbol{h} \\ \boldsymbol{c} & \boldsymbol{f} & \boldsymbol{i} \end{pmatrix} \qquad \begin{pmatrix} + & - & + \\ - & + & - \\ + & - & + \end{pmatrix}$$

In the 18 terms above, we see that the sign is intimately connected with the first factor. Observe that both terms whose first factor is $\boldsymbol{a}$ are *added*. Both terms whose first factor is $\boldsymbol{b}$ are *subtracted*. Both terms whose first factor is $\boldsymbol{c}$ are *added*. And so on and so forth, all the way down to $\boldsymbol{i}$. In fact, what we see is that a term's sign is determined by its first factor's *position* in the matrix. At right, I've rewritten the original matrix, supplemented by another containing $+$ and – signs in the corresponding slots. The result, a simple checkerboard pattern, captures the sign that goes with each position in the matrix. Or, more formally, we can say that if the term's first factor is in row j, column k, then the term is added if $(j + k)$ is even, and subtracted if $(i + j)$ is odd.

We've now unearthed the full pattern, which generalizes to determinants of every size. With patience, we could follow the strategy that led us to our "quick formulas" for 2×2 and 3×3 determinants (Example 2 in the previous section & Exercise 18) and find formulas for 4×4, 5×5, or higher order determinants. Doing so, we'd see that any $n \times n$ determinant can be expressed as a sum of $n!$ terms. With some algebraic shenanigans, we can rearrange those $n!$ terms into a sum of just n terms, each of which has two factors (and a choice of sign). We can arrange matters so that the first factors come from any chosen row or column of the given matrix. The corresponding second factors will then be the determinants of the $(n-1) \times (n-1)$ matrices that we get by deleting the first factor's row and column. Finally, the sign is given by the first factor's position in the $n \times n$ "checkerboard matrix" with a $+$ in its top left corner.

This leads us to a recursive procedure for computing determinants. A 5×5 determinant, for example, reduces to a computation involving 4×4 determinants, each of which reduces to computations involving 3×3 determinants, which then reduce to computations involving 2×2 determinants, which are easy.

This process of reducing a determinant to determinants of lower degree is called **cofactor expansion**.* Since it involves picking a row or column along which to "expand", we refer more specifically to, say, cofactor expansion *along the 1st row*, or cofactor expansion *along the 5th column*, or what have you.

In practice, cofactor expansion is most convenient when a row or column consists mainly of zeros, because in that case, each zero entry yields a zero term in the cofactor expansion. A few examples will make the idea clear.

* The name is explained by some terminology traditionally associated with the technique. First, the matrix we get by crossing out row j and column k is called the **minor** associated with the j, k^{th} entry; next, the minor's determinant multiplied by $(-1)^{j+k}$ (i.e. the $+$ or $-$ dictated by the "checkerboard matrix") is called the the j, k^{th} entry's **cofactor**. With that terminology in place, we can summarize the technique of cofactor expansion as follows:

A matrix's determinant is a weighted sum of any row's (or column's) entries, where the weights are the entries' cofactors.

Example 1. Use cofactor expansion to compute the determinant of $A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 0 \\ 1 & 2 & -2 \end{pmatrix}$.

Solution. Let's expand along the second row to take advantage of that zero:

$$\det\begin{pmatrix} 1 & 2 & 3 \\ \mathbf{4} & \mathbf{5} & \mathbf{0} \\ 1 & 2 & -2 \end{pmatrix} = -\mathbf{4}\det\begin{pmatrix} 2 & 3 \\ 2 & -2 \end{pmatrix} + \mathbf{5}\det\begin{pmatrix} 1 & 3 \\ 1 & -2 \end{pmatrix} - \mathbf{0}$$

$$= -4(-10) + 5(-5) = \mathbf{15}.$$

Observe that the first and third terms in the initial expansion are *subtracted* owing to the positions of entries 4 and 0 in the checkerboard matrix of plusses and minuses. Also, I didn't bother writing down the determinant associated with the 0 in the third term, because there would be no point: whatever it is, it will be multiplied by zero. ♦

Now let's try this with a larger matrix.

Example 2. Use cofactor expansion to compute the determinant of the matrix

$$B = \begin{pmatrix} -8 & 2 & 3 & 7 \\ 1 & 2 & 0 & 3 \\ 4 & 5 & 0 & 0 \\ 1 & 2 & 0 & -2 \end{pmatrix}.$$

Solution. Using cofactor expansion on the third column to profit from all those zeros, we find that

$$\det\begin{pmatrix} -8 & 2 & \mathbf{3} & 7 \\ 1 & 2 & \mathbf{0} & 3 \\ 4 & 5 & \mathbf{0} & 0 \\ 1 & 2 & \mathbf{0} & -2 \end{pmatrix} = \mathbf{3}\det\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 0 \\ 1 & 2 & -2 \end{pmatrix} - \mathbf{0} + \mathbf{0} - \mathbf{0} = \mathbf{3}(15) = \mathbf{45}.$$

Why was the first term *added*? Well, the "checkerboard matrix" of plusses and minuses always has a $+$ in the upper left corner, so moving two spots to the right brings us to another $+$. Incidentally, the one 3×3 matrix in our cofactor expansion was matrix A from Example 1. ♦

Computer programmers turn up their noses at cofactor expansion – with reason. Roughly speaking, computing an $n \times n$ determinant by cofactor expansion along on a random row or column requires about $n!$ arithmetic operations, whereas computing it by means of row reduction requires about n^3 operations. Hence, cofactor expansion is a much more "expensive" way to compute a determinant (when $n > 5$). A computer program using row-reduction to compute a 20×20 determinant matrix will grind through the roughly $20^3 = 8000$ arithmetic operations in no time, but another that uses cofactor expansion will never reach the end of its task, which requires $20! \approx 2$ million trillion operations.

Exercises.

22. Verify the result of Example 1 by recomputing that determinant in the following ways:
a) With the 3×3 "quick formula" b) cofactor expansion on column 3 c) cofactor expansion on row 1

23. Use cofactor expansion to find the determinants of the matrices in Exercise 16.

Chapter 7
Eigenstuff

Eigenvectors and their Eigenvalues

Keep on a straight line… I don't believe I can.
Trying to find a needle in a haystack –
Chilly wind, you're piercing like a dagger, it hurts me so.
Nobody needs to discover me. I'm back again.

- Peter Gabriel as a world-weary eigenvector in "Looking for Someone" (from Genesis's album *Trespass*).

German-English dictionaries will tell you that *eigen* means "characteristic", "particular", or "[one's] own", but these don't fully convey *eigen*'s mathematical significance. In linear algebra, *eigen* means something like "essence-revealing". A linear map's *eigen* things (eigenvectors, eigenvalues, eigenbasis, eigenspaces) reveal its geometric essence.

When blasted by the chilly wind of a linear map, most vectors get blown off the line on which they lie, but a few, hidden like needles in the vector haystack, manage to "keep on a straight line"; the map merely *scales* them. We call these the linear map's *eigenvectors*. Whenever we can discover a basis for a space made up of eigenvectors relative to a map, we have grasped the map's geometric action on the space: It simply scales the space by various factors along various axes. (As for vectors that *don't* lie on the axes, the map just sums their scaled *components* relative to those axes.) Now for our formal definitions.

Definitions. A nonzero vector $\mathbf{v}$ is an **eigenvector** of a linear map T (or matrix A) if there is some scalar λ such that $T(\mathbf{v}) = \lambda\mathbf{v}$. (Or $A\mathbf{v} = \lambda\mathbf{v}$.)

The scalar λ is called an **eigenvalue** of the map (or matrix).
(We also say it is the eigenvalue corresponding to the eigenvector $\mathbf{v}$.)

Some examples will clarify these two simple definitions.

Example 1. Let $A = \begin{pmatrix} 3 & 1 \\ 0 & 2 \end{pmatrix}$, a matrix whose action is depicted below.

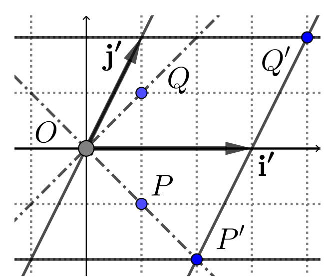

First, let's see what A does to the vector corresponding to point Q:

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

Since the output vector is *not* a scalar multiple of the input vector (equivalently, since A maps point Q to Q', knocking it off the line OQ), we conclude that

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix} \text{ is } \textit{not} \text{ an eigenvector of } A.$$

Now let's consider A's effect on the vector corresponding to point P:

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

Here, the output vector is just the input vector scaled by 2. Thus,

$$\begin{pmatrix} 1 \\ -1 \end{pmatrix} \text{ is an eigenvector of } A, \text{ with eigenvalue } 2.$$

Another eigenvector of A is $\mathbf{i}$ itself; as the figure shows, it gets scaled by 3. Thus, $\mathbf{i}$ is an eigenvector of A with eigenvalue 3. ♦

Example 2. Let T be a counterclockwise rotation of $\mathbb{R}^2$ about the origin by $90°$. Obviously, *every* nonzero vector changes its direction when subjected to this rotation (none of them manages to "keep on a straight line") so this rotation has no eigenvectors. ♦

Example 3. Let R be a reflection of $\mathbb{R}^3$ across the xy-plane. Thinking about this geometric operation, you should be able to convince yourself of the following three things:

- All vectors in the xy-plane are eigenvectors of R (with eigenvalue 1).
- All vectors lying along the z-axis are eigenvectors of R (with eigenvalue -1).
- No other vectors in $\mathbb{R}^3$ are eigenvectors.

Be sure you can understand the three preceding statements by visualizing the reflection. ♦

Now that you know what eigenvectors and eigenvalues are, here is our next big definition:

> **Definition**. An **eigenbasis** relative to an $n \times n$ matrix (or linear transformation of $\mathbb{R}^n$) is a basis of $\mathbb{R}^n$ consisting entirely of eigenvectors of the matrix/map.

We love eigenbases because if we represent a map relative to an eigenbasis, we get a *diagonal* matrix.* We love diagonal matrices because their geometric action is intuitive (mere scaling along the basis "axes") and their algebraic properties are nice (their inverses, determinants, and powers are all easy to compute).

Looking back at Example 1, we see that $\mathbb{R}^2$ has - relative to matrix A - the following eigenbasis:

$$\text{Eigenbasis } \mathcal{B}: \quad \mathbf{b}_1 = \begin{pmatrix} 1 \\ -1 \end{pmatrix} \quad \text{and} \quad \mathbf{b}_2 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

Matrix A merely scales these two linearly independent vectors by factors of 2 and 3 respectively. If we let T be the map whose matrix representation relative to the *standard* basis is A, then T's representation relative to the *eigen*basis $\mathcal{B}$, will - as promised - be diagonal:

$$[T]_{\mathcal{B}} = \begin{pmatrix} 2 & 0 \\ 0 & 3 \end{pmatrix}.$$

We were able to get this diagonal representation only because the map had enough linearly independent eigenvectors to form an eigenbasis. Not all maps are like this.

When a matrix or map does have enough linearly independent eigenvectors to form an eigenbasis, we say that it is **diagonalizable**, since we can then use that eigenbasis to represent it as a diagonal matrix. Thus, A from Example 1 is diagonalizable.

In contrast, the $90°$ rotation in Example 2 has no eigenvectors at all, so there clearly can't be an eigenbasis for $\mathbb{R}^2$ relative to it. Equivalently, a $90°$ rotation in the plane is *not* a diagonalizable map.

In Example 3, reflection R (reflection in the xy-plane) *is* diagonalizable, since it admits an eigenbasis. In fact, the *standard* basis is an eigenbasis here, so R's standard matrix representation is diagonal:

$$[R]_{\mathcal{E}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}.$$

---

* **Proof**: Suppose $\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n$ is an eigenbasis $\mathcal{B}$ (with corresponding eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_n$) relative to some linear map T. Then for each i, we have $T\mathbf{b}_i = \lambda_i \mathbf{b}_i$, so $[T\mathbf{b}_i]_{\mathcal{B}}$ is the column vector whose ith entry is λ_i and whose other entries are all zeros. Since $[T]_{\mathcal{B}}$ consists of these column vectors, it is diagonal as claimed. Moreover, the diagonal entries are T's eigenvalues.

Discovering an eigenvector is like finding gold. When we find one eigenvector, we take it as a sign that a whole "eigenvector vein" is near, spurring us to the happy task of uncovering our treasure's full extent. Each eigenvector belongs to a *subspace* consisting entirely of eigenvectors, all with the same eigenvalue. To see why, first observe that if $\mathbf{v}$ is an eigenvector of A with eigenvalue λ, then so are all its nonzero scalar multiples. (Proof: For any scalar c, we have $A(c\mathbf{v}) = cA(\mathbf{v}) = c(\lambda\mathbf{v}) = (c\lambda)\mathbf{v} = (\lambda c)\mathbf{v} = \lambda(c\mathbf{v})$.*) Moreover, if $\mathbf{v}$ and $\mathbf{w}$ are any *two* eigenvectors with the same eigenvalue λ, then their sum will be a third such eigenvector. (Proof: $A(\mathbf{v}+\mathbf{w}) = A(\mathbf{v}) + A(\mathbf{w}) = \lambda\mathbf{v} + \lambda\mathbf{w} = \lambda(\mathbf{v}+\mathbf{w})$.) The set of eigenvectors with eigenvalue λ is thus closed under scalar multiplication and vector addition. Hence, we may conclude that this is a subspace, as claimed.† Indeed, it has a special name, which I'll bet you can guess:

Definition. Each eigenvalue λ of a map/matrix has a corresponding **eigenspace**, denoted E_λ, which is a subspace consisting of all eigenvectors with eigenvalue λ (and $\mathbf{0}$, too).

Example 1's matrix has two eigenvalues, so relative to it, $\mathbb{R}^2$ has two eigenspaces, depicted at left.

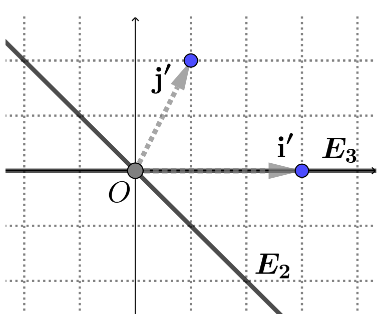

Example 2's map (90° rotation) lacks eigenvalues, so it has no associated eigenspaces. Example 3's reflection in the xy-plane has two eigenvalues, so relative to that reflection, $\mathbb{R}^3$ has two eigenspaces, depicted at right.

The figures also make it clear that we can find an eigenbasis relative to each of these two maps. For example, relative to the reflection, we'll get an eigenbasis for $\mathbb{R}^3$ by taking any two linearly independent vectors in E_1 plus any one nonzero vector in E_{-1}.

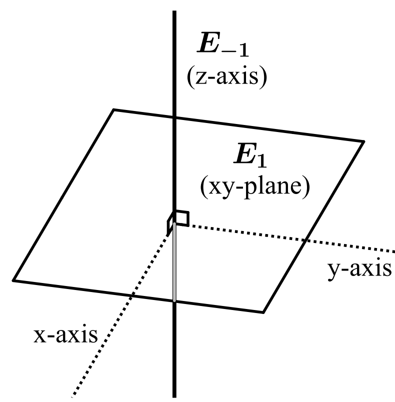

* The Footnote Pedant wishes to point out that the first equals sign in this chain is justified by a basic linearity property (Exercise 14B in Chapter 3); the second is justified because we were given that $\mathbf{v}$ is an eigenvector with eigenvalue λ; the third and fifth are justified by an associative property of scalar multiplication (mentioned in a footnote near the beginning of Chapter 1); the fourth is the commutativity of multiplication of real numbers.

† The Footnote Pedant is now standing athwart my last sentence, yelling STOP. "You said," he declares, "that the set is closed under scalar multiplication, but you've only proved that it is closed under multiplication by *nonzero* scalars! Aren't you sweeping something under the rug? Scaling any eigenvector by 0 turns it into $\mathbf{0}$, which *isn't* an eigenvector because we *defined* eigenvectors to be nonzero. Therefore, the set of eigenvectors with eigenvalue λ is *not* closed under scalar multiplication!" Strictly speaking, he's right. I could satisfy him by replacing the offending phrase "the set of all eigenvectors with eigenvalue λ" with this: "the set of all vectors $\mathbf{v}$ with the property that $A\mathbf{v} = \lambda\mathbf{v}$", which includes all the eigenvectors and *also* the zero vector, since $A(\mathbf{0}) = \lambda\mathbf{0}$, which holds, of course, because every linear map sends $\mathbf{0}$ to $\mathbf{0}$. I'll not do that, though – except down here in the footnotes – since that reformulation is much less memorable, and the discrepancy is minor.

A more interesting thinker than the Footnote Pedant would ask this: "Why did we define eigenvectors to be nonzero in the first place? Why don't we just drop that requirement?" It's a good question. The answer is that if we allowed $\mathbf{0}$ to count as an eigenvector, then *every* real number would be an eigenvalue of *every* matrix, since $A(\mathbf{0}) = \mathbf{0} = r\mathbf{0}$ for all real r. That won't do. Moreover, if $\mathbf{0}$ were considered an eigenvector, then it wouldn't have a particular eigenvalue the way every other eigenvector does; it would have infinitely many eigenvalues. This is a mess. Hence, we exclude $\mathbf{0}$ from the eigenvectors to preserve the statement "every eigenvector has a unique eigenvalue", much as we exclude 1 from the primes to preserve the statement "every number can be factored into a unique product of primes."

Exercises.

1. Is $\begin{pmatrix} 4 \\ -3 \\ 1 \end{pmatrix}$ an eigenvector of $\begin{pmatrix} 3 & 7 & 9 \\ -4 & -5 & 1 \\ 2 & 4 & 4 \end{pmatrix}$? If so, find the corresponding eigenvalue.

2. Describe the eigenvalues and eigenspaces of the following linear transformations of $\mathbb{R}^2$.

a) The zero map (which maps everything to the origin).
b) $180°$ rotation about the origin.
c) Orthogonal projection onto the line $y = 2x$.
d) Reflection across the y-axis.

3. Let A be an invertible $n \times n$ matrix. Suppose $\mathbf{v}$ is an eigenvector of A with eigenvalue λ. Determine whether $\mathbf{v}$ is also an eigenvector of each of the following matrices – and if so, with what eigenvalue. Explain your answers.

a) $2A$ b) nA c) A^2 d) A^n e) A^{-1} f) I g) $2A + 3I$.

4. "Eigenstuff" is defined only for *square* matrices. Explain why.

5. What are the only possible eigenvalues of an *orthogonal* matrix? Explain your answer.

6. (Extending the **Invertible Matrix Theorem**) In Exercise 19 of Chapter 5, you saw that statements A - M in the list below were equivalent statements about an $n \times n$ matrix A. Explain why we can add statement N to the list:

a) A is invertible. b) $\text{rref}(A) = I$. c) $A\mathbf{x} = \mathbf{b}$ has a unique solution for every vector $\mathbf{b}$.
d) A's columns are linearly independent. e) A's columns span $\mathbb{R}^n$. f) A's columns constitute a basis for $\mathbb{R}^n$.
g) $\ker(A) = \mathbf{0}$. h) $\text{im}(A) = \mathbb{R}^n$. i) $\text{rank}(A) = n$. j) $\det A \neq 0$.
k) A's *rows* are linearly independent. l) A's rows span $\mathbb{R}^n$. m) A's rows constitute a basis for $\mathbb{R}^n$.
n) A *doesn't* have 0 as an eigenvalue.

To reiterate, the moral of the invertible matrix theorem is that square matrices come in two types: *invertible* matrices (which satisfy all 14 of those conditions) and *noninvertible* matrices (which satisfy none of them).

7. To appreciate why expressing everything an eigenbasis makes computations so simple…

a) Suppose $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3$ is an eigenbasis $\mathcal{B}$ (with corresponding eigenvalues 3, 8, -2) relative to a linear map T on $\mathbb{R}^3$. Suppose we have a vector $\mathbf{v}$ whose $\mathcal{B}$-coordinates are $(4, -2, 5)$. What will the $\mathcal{B}$-coordinates of $T\mathbf{v}$ be?

b) Given the previous part's setup, if $\mathbf{x}$'s $\mathcal{B}$-coordinates are (x_1, x_2, x_3), what will the $\mathcal{B}$-coordinates of $T\mathbf{x}$ be?

b) More generally, suppose $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n$ is an eigenbasis $\mathcal{B}$ (with corresponding eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$) relative to some linear map T. If $\mathbf{x}$ is a vector with $\mathcal{B}$-coordinates $(x_1, x_2, \dots, x_n)$, then what will $T\mathbf{x}$'s $\mathcal{B}$-coordinates be?

8. As you now know, if $\mathcal{B}$ is an eigenbasis relative to a map T, then $[T]_{\mathcal{B}}$ is a *diagonal* matrix representation of T. This is important, so explain in your own words *why* this is true and what the diagonal entries of $[T]_{\mathcal{B}}$ will be.

9. **(Eigendecomposition)** We can sometimes decompose a given matrix into a product of several simpler factors. (Such a matrix decomposition is analogous to factoring a polynomial or factoring an ordinary whole number.) Of the various types of matrix decompositions, one of the most common and useful is called *eigendecomposition*. An eigendecomposition is like an X-ray of a matrix, exposing its normally hidden eigenstuff to plain view. As you'll learn in this exercise, a matrix can be eigendecomposed if and only if it admits an eigenbasis. Let's get to work.

We begin by noting that we can view any matrix A as the standard matrix of some linear map T (so that $A = [T]_{\mathcal{E}}$). Moreover, if this map has enough linearly independent eigenvectors $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n$ to form an eigenbasis $\mathcal{B}$, we know that $[T]_{\mathcal{B}}$ is a *diagonal* matrix Λ, whose diagonal entries are eigenvalues corresponding to the $\mathbf{b}_i$.* Now…

* The pointy symbol Λ is a Greek letter: *capital* lambda. We use it in this context to remind ourselves that this diagonal matrix's entries are eigenvalues, which, of course, we also represent by lambdas, albeit lower-case ones: λ.

a) Let $V = \begin{pmatrix} | & | & & | \\ \mathbf{b}_1 & \mathbf{b}_2 & \cdots & \mathbf{b}_n \\ | & | & & | \end{pmatrix}$, the columns of which "store" the standard coordinates of A's eigenbasis vectors.

I claim that matrix V is invertible. Explain why this is so.

b) I claim that any matrix A that admits an eigenbasis can be "eigendecomposed" into a product of three matrices:

$$A = V\Lambda V^{-1},$$

where the diagonal matrix Λ stores the eigenvalues, and V stores an eigenvector corresponding to each one (in the corresponding column). Your problem: Justify my claim.

c) For instance, consider matrix A from Example 1. We found that it has an eigenbasis consisting of the vectors $\mathbf{b}_1 = \mathbf{i} - \mathbf{j}$ and $\mathbf{b}_2 = \mathbf{i}$ with corresponding eigenvalues $\lambda_1 = 2$ and $\lambda_2 = 3$. In this case, what are the matrices V, Λ, and V^{-1}? Check your eigendecomposition by multiplying out $V\Lambda V^{-1}$ and verifying that the product is A.

d) **(Eigendecomposition helps us raise a matrix to a *power*)**
Eigendecomposition is useful when we want to apply a linear map to something, then apply the same linear map to the output, then apply the map to *that* output, and so forth. (Such iterative procedures are common in statistics, numerical analysis, machine learning, and even in Google's Page Rank algorithm.) The situation I've just described yields a computation of the form

$$A\left(A\left(A\left(A\cdots\left(A(\mathbf{v})\right)\cdots\right)\right)\right),$$

and since function composition corresponds to matrix multiplication, this reduces to $AAAA\cdots A(\mathbf{v})$, or more compactly, to $A^n(\mathbf{v})$. Alas, raising a matrix to a high power, especially if the matrix is large (as they so often are in applications) is computationally "expensive", even for a very fast computer.

For example, suppose A is a 100×100 matrix, and we wish to raise it to some high power. How many computations must a computer do to carry such a computation out by brute force? Well, just to find $A^2 = AA$ involves doing 10,000 dot products (one for each entry in the matrix product), and each dot product involves 100 products (one for each component) and 99 sums. That works out to almost 2 million arithmetic operations to multiply A by itself just *once*. Grinding out A^n by brute force would require $(n-1)$ of these computationally expensive matrix multiplications. Clearly, we'll want to minimize the number of such matrix multiplications that our computer will have to do. A way to reduce expenses is to find an eigendecomposition of A (if it has one) and raise *that* to the n^{th} power, because if $A = V\Lambda V^{-1}$, it follows that

$$A^n = (V\Lambda V^{-1})^n = \underbrace{(V\Lambda V^{-1})(V\Lambda V^{-1})(V\Lambda V^{-1})\cdots(V\Lambda V^{-1})}_{n \text{ of these trios}}.$$

The associativity of matrix multiplication lets us regroup those parentheses, pairing each trio's concluding V^{-1} with the V from the trio that follows. Such pairs cancel each other out, leaving us with this:

$$A^n = V\Lambda^n V^{-1}.$$

Your problem: Explain why this expression for A^n is much less computationally expensive (when n is large) than the brute force approach to A^n described above. [*Hint*: Raising a *diagonal* matrix to the n^{th} power is easy. Recall Exercise 30e from Chapter 4.]

e) Given matrix A from Example 1, compute A^{10} by hand using the eigendecomposition of A you found in Part C. Then try computing A by brute force, too. (It's ok if you eventually give up. The point is to *feel* the difference between the two approaches so that eigendecomposition's superiority will be palpable to you.)

10. (A fun curiosity) If a square matrix has the property that the entries on each row add up to the same number s, then s is an eigenvalue of the matrix. Explain why this is so.

Finding the Eigenstuff

Mary had a little λ, little λ, little λ...

- Mrs. Traditional

Now that you know what a matrix's eigenthings are, it's time to discuss how we *find* a matrix's eigenthings. Since each eigenvector $\mathbf{v}$ is scaled by some particular eigenvalue λ, we'll want to seek eigenthings in pairs. Namely, given any square matrix A, we want to find all pairs of *nonzero vectors* $\mathbf{v}$ and *scalars* λ such that

$$A\mathbf{v} = \lambda\mathbf{v}.$$

This is the second most important equation in all linear algebra (behind only the ubiquitous $\boldsymbol{A}\mathbf{x} = \mathbf{b}$). Let's rewrite it in an equivalent form where $\mathbf{v}$ appears only once:

$$A\mathbf{v} = \lambda\mathbf{v} \quad \Leftrightarrow \quad A\mathbf{v} - \lambda\mathbf{v} = \mathbf{0} \quad \Leftrightarrow \quad A\mathbf{v} - (\lambda I)\mathbf{v} = \mathbf{0} \quad \Leftrightarrow \quad (A - \lambda I)\mathbf{v} = \mathbf{0}.^{*}$$

It follows from this last reformulation that a given nonzero vector $\mathbf{v}$ and a given scalar λ will be one of A's eigenvector/eigenvalue pairs if and only if $(A - \lambda I)$ maps the given *nonzero* vector $\mathbf{v}$ to $\mathbf{0}$.

This reformulation of the original equation may seem odd, but it delivers crucial geometric insights: First, recall that a *nonzero* vector $\mathbf{v}$ can get mapped to $\mathbf{0}$ by a matrix (here, $A - \lambda I$) only when the matrix *collapses at least one dimension* of the space on which it acts. (Otherwise, the matrix would be invertible, in which case it would only map $\mathbf{0}$ to $\mathbf{0}$.) Second, recall that a matrix induces a dimensional collapse precisely when *the matrix's determinant is zero*.

Putting this all together, we see that our original equation for eigenvector/eigenvalue pairs will be satisfied by a *nonzero* vector $\mathbf{v}$ (and its scalar mate λ) precisely when

$$\det(A - \lambda I) = 0.$$

Notably, this equation doesn't explicitly refer to the eigen*vector* half $\mathbf{v}$ of the pairs that we are seeking. It's actually a good thing that $\mathbf{v}$ is hiding, since this temporarily narrows our focus to a single unknown, λ. If we can solve the equation $\det(A - \lambda I) = 0$ for our unknown λ, we'll have all the eigen*values* of A. Once we have them, they'll confess – after some algebraic coaxing – the locations of the eigen*spaces* where their eigen*vector* mates live. We can then see if those eigenspaces contain enough linearly independent eigenvectors to build an eigenbasis, and if so, we can use it to carry out an eigendecomposition of A.

But before we go that far, let's begin with a quick example that concentrates on eigenvalues alone. That way, we'll see how the first – and most crucial – piece of the puzzle plays out in a concrete instance.

Example 1. Find the eigenvalues of $A = \begin{pmatrix} 1 & 2 \\ 4 & 3 \end{pmatrix}$.

Solution. As discussed above, A's eigenvalues are the solutions of $\det(A - \lambda I) = 0$.
Working out that determinant reveals that the left side of the equation is just a polynomial in λ:

$$\det(A - \lambda I) = \det\left(\begin{pmatrix} 1 & 2 \\ 4 & 3 \end{pmatrix} - \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix}\right) = \det\begin{pmatrix} 1-\lambda & 2 \\ 4 & 3-\lambda \end{pmatrix} = \lambda^2 - 4\lambda - 5.$$

The roots of that polynomial, and hence A's eigenvalues, are $\mathbf{5}$ and $-\mathbf{1}$. ♦

---

[*] The only potentially mysterious step here is rewriting $\lambda\mathbf{v}$ as $(\lambda I)\mathbf{v}$. You should convince yourself that this substitution is valid (i.e. that matrix λI's effect on $\mathbf{v}$ is the same as merely scaling $\mathbf{v}$ by λ). We took that step so that both terms in the expression would be *matrix*-vector multiplications, from which we could then factor out the common vector $\mathbf{v}$.

If A is an $n \times n$ matrix, $\det(A - \lambda I)$ always turns out to be a n^{th}-degree polynomial. A natural name for this eigenvalue-laden polynomial would be A's "eigenpolynomial", but alas, it's called something else:*

Definition. If A is any square matrix, $\det(A - \lambda I)$ is called A's **characteristic polynomial**.

To repeat, A's characteristic polynomial is important because its roots are A's eigenvalues. And once we have its eigen*values*, finding the other eigenstuff is as simple as following your nose, as we'll now see.

Example 2. (Continuing Example 1.) Find the eigenspaces corresponding to the matrix

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

Then, if A admits an eigenbasis, find one, and use it to create an eigendecomposition of A.†

Solution. In Example 1, we already found A's eigenvalues: 5 and -1.

By definition, the eigenspace E_5 consists of all vectors $\mathbf{v}$ such that $A\mathbf{v} = 5\mathbf{v}$. Using an algebraic trick from the previous page, we'll rewrite this equation in an equivalent form:

$$(A - 5I)\mathbf{v} = \mathbf{0}.$$

Now we're in familiar territory. We'll just rewrite this last equation as an augmented matrix,

$$\left(\begin{array}{cc|c} -4 & 2 & 0 \\ 4 & -2 & 0 \end{array}\right),$$

and solve the system. Its solutions, as you should verify, are all the points on the line $y = 2x$. Thus, E_5 is the line $y = 2x$.

The eigenspace E_{-1} consists of all vectors $\mathbf{v}$ such that $A\mathbf{v} = -\mathbf{v}$. Or equivalently, all vectors $\mathbf{v}$ such that $(A + I)\mathbf{v} = \mathbf{0}$. We can find them by solving

$$\left(\begin{array}{cc|c} 2 & 2 & 0 \\ 4 & 4 & 0 \end{array}\right).$$

When we do so, we find, as you should verify, that every point on the line $y = -x$ is a solution. Thus, E_{-1} is the line $y = -x$.

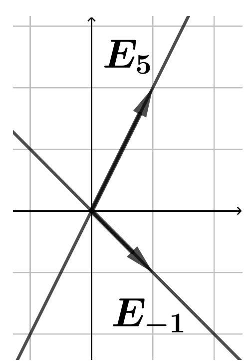

Any two linearly independent eigenvectors of A will form an eigenbasis, so we can clearly get one by taking one eigenvector from each eigenspace – such as these ones (expressed standard coordinates):

$$\begin{pmatrix} 1 \\ 2 \end{pmatrix} \text{ with eigenvalue 5} \qquad \text{and} \qquad \begin{pmatrix} 1 \\ -1 \end{pmatrix} \text{ with eigenvalue } -1.$$

It follows from Exercise 9B that A's corresponding eigendecomposition is

$$\begin{pmatrix} 1 & 2 \\ 4 & 3 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 2 & -1 \end{pmatrix} \begin{pmatrix} 5 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} 1/3 & 1/3 \\ 2/3 & -1/3 \end{pmatrix},$$

as you can verify by multiplying out the right side. ♦

---

* Well, sort of. "Characteristic" is a translation of *eigen*, so it might as well be "eigenpolynomial". But perhaps this word, cobbled together from German (*eigen*), Greek (*poly*), and Latin (*nomen*), was deemed too outré for use by some dreary linguistic purist.

† The notion of an *eigendecomposition* was introduced in Exercise 9.

If you watch someone (your professor, a friend, or someone online) computing a matrix's eigenspaces, it can look like a mysterious algorithm, especially if he or she compresses some of the preliminary steps. Indeed, I've seen some students learn to carry out the process in a mindless algorithmic manner, which runs something like this: "If one of A's eigenvalues is 42 (or whatever), then to find E_{42}, we subtract 42 from each diagonal entry of A, then we turn the result into an augmented matrix by appending a column of zeros, and finally we solve that system. The solutions are the eigenvectors that make up E_4." That recipe will indeed produce the desired eigenspace, and there's nothing wrong with using it – provided you understand *why* it works. If you use it without understanding why it works, you aren't learning linear algebra; you are just following orders. On the other hand, if you understand why it works, there's no reason to memorize that sequence of steps in the first place. After all, it takes only a few seconds to reason through the entire process, justifying every step: "By definition, E_{42} consists of all vectors that satisfy $A\mathbf{v} = 42\mathbf{v}$. To find these vectors, we'll rewrite that equation as $(A - 42I)\mathbf{v} = \mathbf{0}$, since we can easily turn this into an augmented matrix whose underlying system we can then solve with Gaussian elimination." I encourage you to think through the process this way rather than memorize an algorithm.

Once you have the eigenspaces of an $n \times n$ matrix, it's easy to see whether the underlying map is diagonalizable: If the eigenspaces' dimensions add up to n, we can gather up enough linearly independent eigenvectors to form an eigenbasis, which we can then use to represent the map as a diagonal matrix. But if sum of the eigenspaces' dimensions falls short of n, we're out of luck: The map doesn't admit an eigenbasis, so we can't represent it as a diagonal matrix. (Incidentally, matrices that can't be diagonalized are sometimes called **defective** matrices.)

Let's do one another example of finding eigenstuff.

Example 3. Find the eigenspaces of

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

Solution. A's characteristic polynomial is

$$\det(A - \lambda I) = \det\begin{pmatrix} 1-\lambda & 0 & 2 \\ 0 & 3-\lambda & 0 \\ 2 & 0 & 1-\lambda \end{pmatrix}.$$

Doing cofactor expansion on the second row and massaging the resulting expression reveals the characteristic polynomial to be $(-1)(3-\lambda)^2(\lambda+1)$. Thus, A's eigenvalues are $\mathbf{3}$ and $-\mathbf{1}$.

The eigenspace E_3 consists of all vectors $\mathbf{v}$ such that $A\mathbf{v} = 3\mathbf{v}$. We can rewrite this equation in the equivalent form $(A - 3I)\mathbf{v} = \mathbf{0}$, so we can solve it by row-reducing an augmented matrix:

$$\left(\begin{array}{ccc|c} -2 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 \\ 2 & 0 & -2 & 0 \end{array}\right) \begin{matrix} \\ \\ +R_1 \end{matrix} \left(\begin{array}{ccc|c} -2 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{array}\right) \begin{matrix} \div(-2) \\ \\ \\ \end{matrix} \left(\begin{array}{ccc|c} 1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{array}\right).$$

From this, we see that $x = z$, while there are no constraints on y or z, so these are free variables. Hence, solutions have the form

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} t \\ s \\ t \end{pmatrix} = t\begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + s\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}.$$

That is, the eigenspace E_3 is a plane in $\mathbb{R}^3$: the span of the two independent vectors $\mathbf{i} + \mathbf{k}$ and $\mathbf{j}$.

As for the eigenspace E_{-1}, this consists of all vectors $\mathbf{v}$ such that $A\mathbf{v} = -\mathbf{v}$. We'll rewrite this in the form $(A + I)\mathbf{v} = \mathbf{0}$, and then solve it by row-reducing an augmented matrix:

$$\left(\begin{array}{ccc|c} 2 & 0 & 2 & 0 \\ 0 & 4 & 0 & 0 \\ 2 & 0 & 2 & 0 \end{array}\right) \begin{array}{c} \div 2 \\ \div 4 \\ \div 2 \end{array} \left(\begin{array}{ccc|c} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \end{array}\right) \begin{array}{c} \\ \\ -R_1 \end{array} \left(\begin{array}{ccc|c} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{array}\right).$$

From this we see that $z = -x$, and that $y = 0$. Hence, the solutions have the form

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} t \\ 0 \\ -t \end{pmatrix} = t\begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}.$$

That is, the eigenspace E_{-1} is a line in $\mathbb{R}^3$: the span of the vector $\mathbf{i} - \mathbf{k}$.

To sum up, A's eigenvalues are 3 and -1, and the corresponding eigenspaces are:

$$E_3\text{: span of } \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} \text{ and } \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}; \qquad E_{-1}\text{: span of } \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}.$$ ♦

Although I didn't ask for it in the previous example, you should be able to see that A admits an eigenbasis: any two linearly independent vectors from E_3 plus any nonzero vector from E_{-1} will do. For example, vectors $\mathbf{i} + \mathbf{k}$, $\mathbf{j}$, and $\mathbf{i} - \mathbf{k}$ make up one eigenbasis. With respect to that basis, the linear map whose standard matrix is A will have a diagonal representation, and its diagonal entries will be 3, 3, and -1. You should be able to write down the related eigendecomposition of A, too.

Once we've found a matrix's eigen*values*, the rest of its eigenstuff comes easily. But there's a catch: Finding the eigenvalues usually isn't easy. Well… finding the eigenvalues of a 2×2 matrix is easy, since 2×2 matrices have *quadratic* characteristic polynomials, whose roots we can always find with the quadratic formula. But alas, 3×3 and 4×4 matrices have *cubic* and *quartic* characteristic polynomials. Cubic and quartic analogs of the quadratic formula do exist, but they are so appallingly complicated that no one in his or her right mind knows them.* And for larger matrices, the situation is considerably worse: No algebraic formula for the roots of 5th (or higher) degree polynomials exists.† So how, in applications, do we find large matrices' eigenvalues? We approximate them. Laborers in the numerical linear algebra mines have worked out *eigenvalue algorithms* to approximate matrices' eigenvalues as closely as we like. These sophisticated iterative algorithms – which are built into computer programs that scientists and engineers use every day – typically *don't* involve the characteristic polynomial. Yet even for large matrices, the characteristic polynomial remains a vital theoretical touchstone. For example, in Exercise 13, you'll use it to prove an important fact about the eigenvalues of triangular matrices – of any size.

---

* The cubic and quartic formulas are *historically* important. Their discovery in 16th-century Italy was the first original achievement of European mathematics since the ancient Greeks and initiated the *mathematical* Renaissance. Remarkably, their discovery involved the first real use of complex numbers in mathematics – and this was at a time when even negative numbers, much less their square roots, were considered hopelessly absurd fictions. The story of the cubic, which includes mathematical duels, flashes of insight, vows of secrecy made and broken, and a colorful cast of characters (above all Niccolò Tartaglia and Gerolamo Cardano), is engagingly told by Paul Nahin in the first chapter of his semi-popular book on complex numbers, *An Imaginary Tale.*

† To be clear, I do not mean that no one has found such a formula yet. I mean that there *isn't* one and there never will be one. This was proved in the 19th century in yet another fascinating episode of mathematical history – one that ultimately gave birth to large portions of what we now call *abstract algebra* – group theory and Galois theory in particular.

Exercises.

11. a) Remind yourself *why* a matrix A's eigenvalues are the zeros of its characteristic polynomial, $\det(A - \lambda I)$.
b) If, say, 8 is an eigenvalue of matrix A, explain how to find the eigenspace E_8, and why your method works.

12. For each of the following, find the eigenvalues, describe the corresponding eigenspaces, and determine whether the matrix admits an eigenbasis. If so, state an eigenbasis, and give the corresponding eigendecomposition.

a) $\begin{pmatrix} 3 & 2 \\ 3 & 8 \end{pmatrix}$ b) $\begin{pmatrix} -2 & 6 \\ -2 & 5 \end{pmatrix}$ c) $\begin{pmatrix} 2 & 3 \\ -4 & -2 \end{pmatrix}$ d) $\begin{pmatrix} 2 & 2 \\ -8 & -6 \end{pmatrix}$

e) $\begin{pmatrix} 5 & 2 & 3 \\ 0 & 5 & 4 \\ 0 & 0 & 1 \end{pmatrix}$ f) $\begin{pmatrix} 3 & -2 & 5 \\ 1 & 0 & 7 \\ 0 & 0 & 2 \end{pmatrix}$ g) $\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$

13. (Triangular matrices' eigenvalues) In Part E of the previous problem, you saw that the eigenvalues of a *triangular* matrix turned out to be the entries on its main diagonal. Was that a coincidence? It was not! Prove the following:

Theorem. If M is a *triangular* matrix, then M's eigenvalues are the entries on its main diagonal.

This gives us another reason to like triangular matrices: We can see their eigenvalues at a glance.

14. (The transpose's eigenvalues)
a) Justify each equals sign in the following chain, which pertains to every square matrix A:

$$\det(A - \lambda I) = \det((A - \lambda I)^T) = \det(A^T - (\lambda I)^T) = \det(A^T - \lambda I).$$

You've just shown that every square matrix and its transpose share the same characteristic polynomial.

b) Can we conclude that A and A^T will have the same eigenvalues? What about their eigenvectors?

c) Suppose that a square matrix A has the property the entries on each *column* sum up to the same number s.*
What, if anything, can we conclude about A's "eigenstuff"? (Compare Exercise 10.)

15. *Invertibility and diagonalizability are completely unrelated concepts*.
To reinforce this fact, I've presented four matrices below that cover all possible combinations of invertibility and diagonalizability. Verify this, and fill in the table below.

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \qquad B = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad C = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \qquad D = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

	Invertible?	Diagonalizable?
A		
B		
C		
D		

16. You've seen 2×2 matrices with no real eigenvalues, such as the matrix in 12C above, or the matrix that effects a 90° rotation of $\mathbb{R}^2$ about the origin. However, *every* 3×3 matrix has at least one real eigenvalue. Explain why.†

* Such matrices arise naturally in probability calculations, for example as *transition matrices* (also known as *stochastic matrices*) in Markov chains. The entries in each column of such a matrix represent probabilities that must add up to 1.

† There are such things as *complex* eigenvalues, but not in this introductory textbook.

17. (Algebraic and Geometric Multiplicities of Eigenvalues)

You might or might not be familiar with "the Fundamental Theorem of Algebra" (FTA), which states that every n^{th}-degree polynomial $p(x)$ can be factored as follows:

$$p(x) = c(x - r_1)(x - r_2) \cdots (x - r_n),$$

where c is a constant, and where some or all of the constants r_i may be complex numbers. A few examples:

$$x^2 - 6x + 9 = (x - 3)(x - 3)$$
$$2x^2 + 3x - 2 = 2(x - 1/2)(x + 2)$$
$$x^3 + x = (x - 0)(x - i)(x + i)$$
$$x^4 - 2x^3 - 7x^2 + 20x - 12 = (x - 2)^2(x - 1)(x + 3).$$

The FTA can't tell you how to find a factorization – only that a factorization *exists* (in the mind of God, as it were). The FTA is a vital theorem for theoretical work; we'll use it below - and again in exercise 18. The constants r_i that appear in the n^{th}-degree polynomial's factorization are its n roots. If the same root appears k times in this factorization, we say that root has **multiplicity** k. (When $k > 1$, we call the root a **repeated root**.)

a) State the roots – and the multiplicity of each repeated root – of each of the four polynomials above.

Now for some definitions:
The ***algebraic*** **multiplicity** of an eigenvalue of A is its multiplicity as a root of A's characteristic polynomial.
The ***geometric*** **multiplicity** of an eigenvalue of A is the dimension of its associated eigenspace.

b) Exercise 12F concerned a matrix whose characteristic polynomial was $(-1)(\lambda - 2)^2(\lambda - 1)$. Its eigenvalues were thus 2 and 1, and you found that the respective eigenspaces were both *lines* in $\mathbb{R}^3$.
That being so, state the algebraic and geometric multiplicities of each eigenvalue of that matrix.

One can show that *each eigenvalue's **geometric** multiplicity is less than or equal to its algebraic multiplicity*. The proof is somewhat involved, and would involve a serious detour, so we'll take it for granted in this exercise. With that in mind, explain why each of the following statements about an $n \times n$ matrix A's *real* eigenvalues hold.

c) The sum of the algebraic multiplicities is at most n.
d) The sum of the geometric multiplicities is at most n.
e) The matrix is diagonalizable ⇔ the geometric multiplicities add up to n.
f) If the algebraic multiplicities add up to n, then the matrix might - or might not - be diagonalizable.
g) The matrix is diagonalizable ⇔ (1) the algebraic multiplicities add up to n *and* (2) each eigenvalue's geometric multiplicity equals its algebraic multiplicity.
h) If the matrix has even one eigenvalue whose geometric multiplicity is strictly less than its algebraic multiplicity, then the matrix is defective (i.e. it *can't* be diagonalized).
i) If an eigenvalue's algebraic multiplicity is 1, the eigenvalue's corresponding eigenspace is a *line*.

18. (The Trace of a Matrix)

If A is a square matrix, then its **trace** (which we denote $\text{tr}(A)$) is the sum of the entries on A's main diagonal. Surprisingly, the trace turns up in some interesting places, a few of which you'll meet in this problem.

a) Prove the following: If A is a 2×2 matrix, then A's characteristic polynomial is

$$\lambda^2 - \big(\text{tr}(A)\big)\lambda + \det(A).$$

b) Use the preceding result to find the characteristic polynomials of the 2×2 matrices in Exercise 12 more quickly than you found them originally.

c) Prove the following: If A is a 2×2 matrix with eigenvalues λ_1 and λ_2 (which may be equal or complex!), then

$$\operatorname{tr}(A) = \lambda_1 + \lambda_2 \quad \text{and} \quad \det(A) = \lambda_1 \lambda_2.$$

That is, the eigenvalues' sum is the trace; the eigenvalues' product is the determinant.
[*Hint*: *Combine the result of Part A with the Fundamental Theorem of Algebra from Exercise 17*.]

d) Use Part C to find the following matrices' eigenvalues *without* finding their characteristic polynomials:

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

e) In fact, the results from Part C extend to square matrices of all sizes, not just 2×2 ones. That is, the sum and product of every square matrix's eigenvalues turn out to be its trace and determinant respectively.* This result, perhaps combined with others, can be used to impress your friends at parties by finding a matrix's eigenvalues without having to think about its characteristic polynomial. For example, what are the eigenvalues of

$$M = \begin{pmatrix} 2 & 4 & 3 \\ 1 & 2 & 6 \\ 3 & 6 & 0 \end{pmatrix}?$$

Well, all the rows add up to the same constant, so that constant, **9**, must be an eigenvalue (by Exercise 10). Next, observe that the columns are linearly *de*pendent, since column two is double column one. Hence, by the Invertible Matrix Theorem (see Exercise 6), this matrix is *not* invertible, and accordingly, it must have **0** as one of its eigenvalues. What's the third eigenvalue? Well, $\operatorname{tr}(A) = 4$, so the three eigenvalues must sum to 4, which means that the third eigenvalue must be **-5**. Hence, M's eigenvalues are 9, 0, and -5.

Use such "party tricks" to find the eigenvalues of these matrices without computing their characteristic polynomials, carefully justifying each of your steps:

$$A = \begin{pmatrix} 1 & 3 & 4 \\ 2 & 2 & 4 \\ 3 & 1 & 4 \end{pmatrix}, \qquad B = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 5 & 1 & 3 \end{pmatrix}, \qquad C = \begin{pmatrix} 1 & 2 & 3 \\ 1 & 2 & 3 \\ 4 & 2 & 0 \end{pmatrix}.$$

* But note that in this eigenvalue sum, one must include each eigenvalue a number of times equal to its own algebraic multiplicity. Thus, if A's characteristic polynomial is $(\lambda - 5)^2(\lambda + 2)(\lambda - 1)^3$, then its trace would be $5 + 5 + (-2) + 1 + 1 + 1 = 11$, and its determinant would be $(5)^2(-2)(1)^3 = -50$.

Eigenstuff and Long Run Behavior

In the long run we are all dead.

- John Maynard Keynes, *A Tract on Monetary Reform*

One way to grasp a mathematical system's long run behavior is to express the system in linear algebraic terms, and then analyze its eigenstuff, which often turns out to be related to *limits*, your old friends from calculus. To explain the idea, I'll dwell for a bit on the famous **Fibonacci sequence**: $0, 1, 1, 2, 3, 5, 8, 13$ ..., a numerical system that unfolds from a simple seed (the two initial terms, 0 and 1) plus a simple rule: Each successive term is the sum of the previous two terms.

Given *two* consecutive terms in the Fibonacci sequence, we produce the next term by adding them. But given just *one* number from the sequence (say, 832040), is there some mathematical way to find, or even approximate, its successor without having to reconstruct the entire sequence up to that point?

There is. The trick involves grasping the sequence in algebraic, and ultimately *linear* algebraic terms. If we let F_n be the Fibonacci sequence's n^{th} term, we can specify the full sequence recursively as follows:

$$F_0 = 0, \quad F_1 = 1; \qquad F_n = F_{n-2} + F_{n-1} \text{ for all } n \geq 2.$$

Next, to inject some geometry, we will form vectors in $\mathbb{R}^2$ whose top component is a Fibonacci number (i.e. any number from the sequence), and whose bottom component is the *previous* Fibonacci number:

$$\begin{pmatrix} F_n \\ F_{n-1} \end{pmatrix}.$$

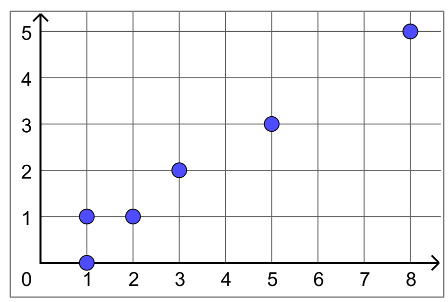

At right, I've plotted points corresponding to the first six of these "Fibonacci vectors". For example, the point at $(5, 3)$ tells us that 5 is a Fibonacci number… whose predecessor in the sequence is 3. The *second* component of a Fibonacci vector might initially seem pointless, but tucking it into a vector along with the first is clever, opening the door not only to geometry, but to a matrix as well. The key observation that leads us to it is that a certain *linear map* will take us from any one Fibonacci vector/point,

$$\begin{pmatrix} F_{n-1} \\ F_{n-2} \end{pmatrix}, \text{ to the next one, } \begin{pmatrix} F_n \\ F_{n-1} \end{pmatrix},$$

since the latter's components are linear combinations of the former's components:

$$\begin{aligned} F_n &= 1F_{n-1} + 1F_{n-2} \\ F_{n-1} &= 1F_{n-1} + 0F_{n-2}. \end{aligned}$$

We can bundle these linear relationships into a single matrix-vector equation:

$$\begin{pmatrix} F_n \\ F_{n-1} \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} F_{n-1} \\ F_{n-2} \end{pmatrix}.$$

Thus, we've found a 2×2 matrix that sends each Fibonacci vector/point to its *successor* on the graph. Accordingly, we can now see the entire graph as generated from the initial point $(1, 0)$ and the matrix: We feed $(1,0)$ into the matrix, which maps it to $(2, 1)$. We feed *that* point back into the matrix, which maps it to $(3, 2)$… and so on forever. The points generated this way are the Fibonacci points.

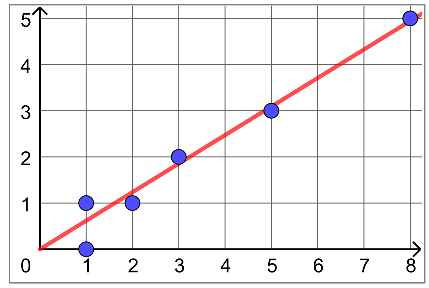

Observe that the graph's points come close to falling on a line. Or rather, the line is the limit towards which successive points tend. In the long run, points in the sequence are indistinguishable from points that lie *on* the line. Once we're far enough into the sequence, the matrix is essentially just *scaling* Fibonacci vectors, stretching them along the line without rotating them. In other words, in the long run, the Fibonacci vectors are eigenvectors, and the line is an eigenspace of the matrix. Its associated eigen*value* is, of course, the scaling factor by which the matrix, in the long run, stretches each Fibonacci vector.

Computing the matrix's eigenvalues is easy. You should do that now, using Exercise 18a's shortcut. You'll find that it has only one eigenvalue greater than 1. That's the relevant eigenvalue here, since the Fibonacci vectors are stretched *away* from the origin. This eigenvalue turns out to be the **golden ratio**,

$$\phi = \frac{1+\sqrt{5}}{2} \approx 1.618,$$

an irrational number famous for turning up in all sorts of mathematical places. You'll have the chance to explore some of ϕ's properties in the next exercise set. At any rate, we've established that if $\mathbf{f}_n$ is the n^{th} Fibonacci vector, then $\mathbf{f}_{n+1} \approx \phi \mathbf{f}_n$, with the approximation improving as n gets larger. Plucking the first components from each side of this approximation yields

$$F_{n+1} \approx \phi F_n,$$

and so, to return to our original question, if someone hands us a large Fibonacci number such as 832040 and asks us what the next one in the sequence will be, we need not reconstruct the whole sequence up to that point to find it. We can just multiply 832040 by ϕ. The product, rounded to the nearest integer, turns out to be 1346269. This integer should be close to the next term in the Fibonacci sequence. Remarkably, it turns out to be the next term exactly!* In fact, when we round ϕF_n to the nearest integer, we obtain F_{n+1} exactly for all $n \geq 2$. In other words, this "long run" tendency for the Fibonacci sequence takes hold almost immediately. This happens because the initial Fibonacci point $(1, 0)$ already lies so close to the magic line. This line, the eigenspace E_ϕ, attracts vectors like a magnet as we run Fibonacci vectors (and the resulting sequence of successive outputs) through the matrix.

Interestingly, the line's magnetic quality persists even if we initially feed the matrix a vector that *isn't* a Fibonacci vector. For example, the point/vector $(7, 1)$ isn't a Fibonacci vector, but if we feed it to the matrix, then feed its output back in, and so on, the resulting sequence of points tends towards the line, even though we're no longer working with Fibonacci numbers. These points' first components constitute a new sequence of numbers, still governed by Fibonacci's rule ("add the previous two to get the next"), but developing from a different seed:

$$1, 7, 8, 15, 23, 38, 61, 99, \ldots .$$

In the long run, this new sequence still grows at each step by a factor of ϕ. But since our initial point $(7, 1)$ now sits considerably further from the "magnetic" line (the eigenspace corresponding to ϕ), it takes a few more steps before the scaling factor of ϕ nails down the next term flawlessly. For example, 23 is the new

* Either patiently (by generating the sequence) or impatiently (by looking it up), you can confirm that 832040 and 1346269 are, respectively, F_{30} and F_{31}.

sequence's fifth term, and $23\phi \approx 37$, which is close… but not quite right. The sixth term is actually 38. However, by the next term, all's well: $38\phi \approx 61$, which is indeed the sequence's seventh term.

It helps to think of the matrix as generating a force field that pervades the plane, capable of pushing points (i.e. vectors) towards the eigenspace. The figure below gives the idea:*

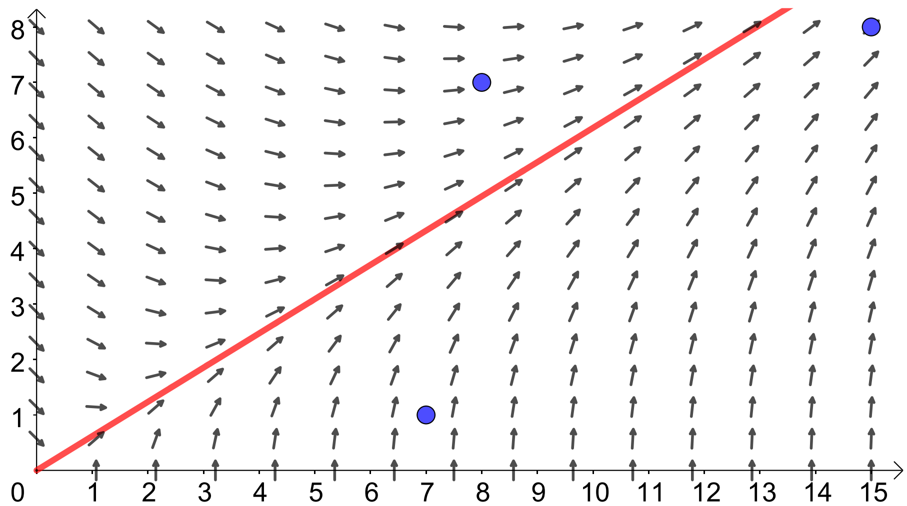

Initially, we "dropped" a point into the field at $(7, 1)$, where it was subjected to forces pushing it north and very slightly east (the arrows in the picture give only the *direction* in which the "wind" blows at each point; they don't indicate its strength). This brought the point to $(8, 7)$, where the wind blows easterly, with a slight push to the north. These forces then pushed the point to $(15, 8)$, quite close to the line. Here, the wind blows nearly parallel to the line, but it still pushes points a bit closer to it. The overall effect, as the figure makes clear, is that no matter where we initially drop our point, the "winds" will ultimately push the point towards the line.

"But," you might reasonably ask, "What about the matrix's *other* eigenvalue,

$$\frac{1 - \sqrt{5}}{2} \approx 0.618,$$

and *its* eigenspace? Don't they have any effect on the geometric picture above?" They do indeed, but they will be visible only if we widen our field of view to encompass the entire plane – not just its first quadrant. When we do so, our picture will contain a pair of one-dimensional eigenspaces, like so:

* A computer generates the figure by taking many points (x, y) in the plane, and at each such point, computing the quantity

$$\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}\begin{pmatrix} x \\ y \end{pmatrix} - \begin{pmatrix} x \\ y \end{pmatrix},$$

which is the vector pointing *from* (x, y) to the point where the matrix sends (x, y). The computer then scales this direction vector down dramatically (to avoid traffic jams in the picture) and places it at (x, y).

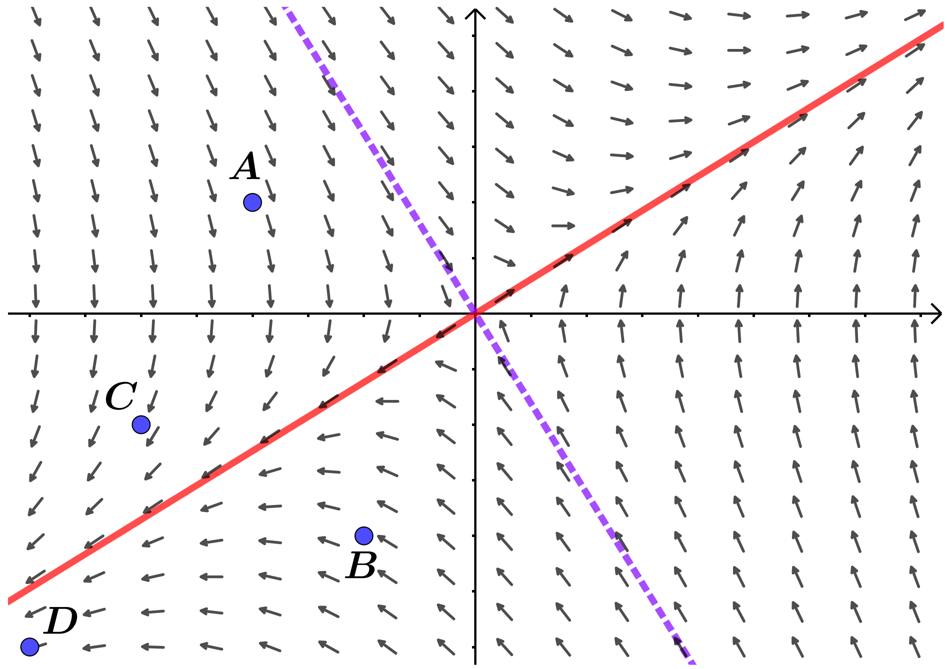

The dotted eigenspace's eigenvalue is approximately 0.618. Since this is less than 1, points on it are drawn *towards* the origin, as the "force field" indicates. Even in this expanded view, it's easy to see that a point dropped anywhere in the plane will eventually find itself moving towards the first eigenspace, E_ϕ. For example, if we drop a point at $(-4, 2)$, it first gets mapped southeast to B $(-2, -4)$, then pushed northwest to C $(-6, -2)$, and then to D $(-8, -6)$. From there, the point is clearly locked forevermore into the stream that flows along E_ϕ, albeit in its "negative direction", flowing into the 3rd quadrant.

The ideas above have applications far beyond the Fibonacci sequence. We are now trespassing in the vast domain of *dynamical systems*, which you can explore in an introductory differential equations course – should you be lucky enough to take one employing a dynamical systems point of view. Still, even without the machinery of differential equations, you can sample the topic's flavor by considering the following example of a *discrete* dynamical system.

Living in the remote depths of a certain forest are some wild dachshunds and their preferred prey, feral mailmen. If D_n and M_n represent the population of each species in this region at the beginning of year n, then a simple **predator-prey model** of how these populations change in time might look like this:

$$\begin{aligned} D_{n+1} &= .86D_n + .08M_n \\ M_{n+1} &= -.12D_n + 1.14M_n. \end{aligned}$$

To see why this might make sense, consider each equation in turn. The first indicates that the number of dachshunds *next* year depend on how many dachshunds and mailmen there are this year. If for, example, there are no mailmen in the region this year, then the dachshund population will, according to this model, be reduced by 14% as they are forced to endure a year without their favorite food. On the other hand, for every 100 mailmen in the area this year, there will be 8 new dachshunds next year – a perfectly reasonable assumption, since mailman meat nourishes the local dachshunds and attracts more dachshunds from afar.

As for the model's second equation, it tells us that in the absence of dachshunds, the unchecked mailman population would grow by 14% each year.* However, for every hundred dachshunds, twelve mailmen will succumb to their ruthless predation. Incidentally, some have speculated that dachshund-on-mailman violence was precisely what inspired Alfred Lord Tennyson's phrase "Nature, red in tooth and claw".

Naturally, the model above can be translated into a matrix-vector equation:

$$\begin{pmatrix} D_{n+1} \\ M_{n+1} \end{pmatrix} = \begin{pmatrix} .86 & 0.08 \\ -.12 & 1.14 \end{pmatrix} \begin{pmatrix} D_n \\ M_n \end{pmatrix},$$

and an "eigenanalysis" like the one we employed in our Fibonacci problem will reveal much about this system's long run behavior. Carrying this out, we find that the matrix has two eigenvalues: 1.1 and 0.9. Their corresponding eigenspaces are shown in the figure below.

The figure reveals how the fate of this predator-prey system will depend on its initial conditions. For example, if we start with 800 dachshunds and 200 mailmen (i.e. if we drop a point into the "force field" at $(800, 200)$), the arrows point out a dire fate for the mailmen: They'll be hunted to extinction, driven down to the D-axis of doom. However, had those 800 dachshunds been paired initially with 600 mailmen, our moving point would be drawn to a happier fate: eigenspace $E_{1.1}$, where a stable ratio of 3 mailmen for every 1 dachshund eventually holds sway, with both populations then growing by 10% in each subsequent year. Of course, that growth rate can't be sustained forever in practice, which is one weakness of this crude predator-prey model.† More refined models exist, but this example was only meant to convey the basic idea.

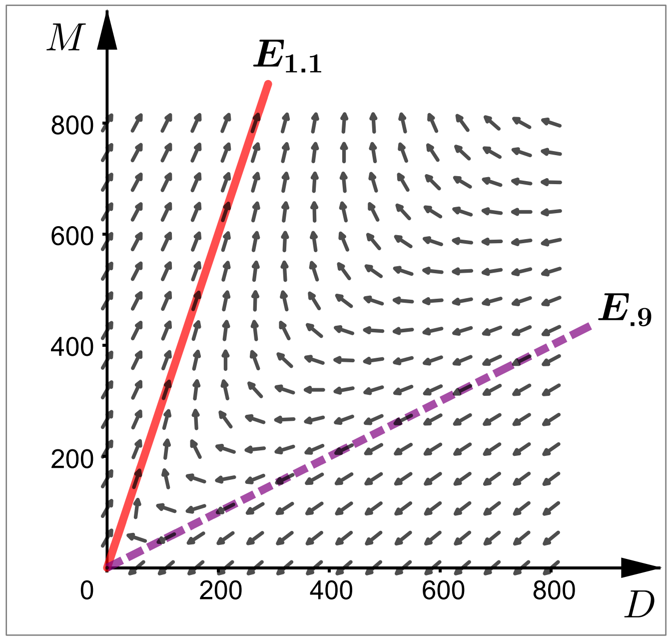

At any rate, you'll have the chance to dig more deeply into these ideas in future classes. My intention here was just to gesture in their general direction, pointing towards the mathematical horizon – towards an educational eigenspace that you may find yourself drawn to bit by bit should your own initial conditions happen to predispose you to becoming sucked into the world of dynamical systems.

* Astute readers may now be wondering if these feral mailmen have evolved to reproduce asexually. Not so. Although I've used the gendered term "mailmen", it should be understood as including mail*women*, known sometimes as *femailmen*.

† A better model would account for the area's "carrying capacity" – a measure of how many dachshunds and mailmen the area can support before overcrowding leads to population decline.

Exercises.

19. Historically, the golden ratio arose from the following problem: Given any line segment, cut it into two pieces so that that the whole segment is to the longer piece as the longer piece is to the shorter piece. If we can do this, we define the numerical value of this ratio (whole/long, or equivalently, long/short) to be the **golden ratio**, ϕ.

Your problem: From the golden ratio's definition, deduce its exact numerical value.

[**Hint**: *Define the short piece of our divided line segment to be* 1 *unit. Call the long piece* x*, and use* ϕ*'s definition to deduce* x*'s length. Then by* ϕ*'s definition (again), it follows that* $\phi = x/1 = x$*, giving us* ϕ*'s numerical value.*]

20. The polynomial $x^2 - x - 1$ is intimately linked to the golden ratio: Its positive root is ϕ, while its negative root, which I'll call ψ, is ϕ's "conjugate". That is, the polynomial's two roots are:

$$\phi = \frac{1+\sqrt{5}}{2}, \qquad \psi = \frac{1-\sqrt{5}}{2}.$$

Demonstrate each of the following facts, some of which may prove useful to you in the next exercise:

a) $\phi^2 = 1 + \phi$ b) $\phi^{-1} = -\psi$ c) $\psi = 1 - \phi$ d) $\phi \approx 1.618$ e) $\psi \approx -0.618$

21. (The n^{th} Fibonacci number) Eigenstuff can lead us to an exact formula for F_n, the n^{th} Fibonacci number. In this exercise, I'll walk you through the argument, letting you fill in the details.

a) Let $A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$, the matrix we encountered in this section's Fibonacci example. Explain why

$$\begin{pmatrix} F_{n+1} \\ F_n \end{pmatrix} = A^n \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

b) Explain why it follows that

$$F_n = (0 \quad 1) A^n \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

c) Although the preceding equation gives us a recipe for F_n, it's not very satisfying, as it involves raising matrix A to a power, which is computationally expensive. But as you saw in Exercise 8D, we know a trick for raising a matrix to a power: *eigendecomposition*. Recall the eigendecomposition punchline (spelled out in Exercise 9B): If A admits an eigenbasis (as it does here), we can write $A = V\Lambda V^{-1}$, where Λ is a diagonal matrix whose diagonal entries are A's eigenvalues, while matrix V stores - in its corresponding columns - an eigenvector for each eigenvalue. Accordingly, eigendecomposition will turn Part B's formula into something of the form

$$F_n = (0 \quad 1) V \Lambda^n V^{-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

This still contains a matrix raised to a power, but now it is a *diagonal* matrix, which is good, because raising a diagonal matrix to the n^{th} power is easy: we just raise each of its diagonal entries to the n^{th} power.

Your problems: Show that A's characteristic polynomial is $\lambda^2 - \lambda - 1$, and thus (by Exercise 20), A's eigenvalues are ϕ and ψ. Find a corresponding eigenvector for each eigenvalue. Having done this, conclude that the following expansion of our previous equation is valid:

$$F_n = (0 \quad 1) \begin{pmatrix} \phi & \psi \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \phi & 0 \\ 0 & \psi \end{pmatrix}^n \begin{pmatrix} \phi & \psi \\ 1 & 1 \end{pmatrix}^{-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

d) Simplify the preceding until you obtain the following remarkable formula:

$$F_n = \frac{\phi^n - \psi^n}{\sqrt{5}}.^*$$

e) Use the preceding formula to compute F_{40} directly with a scientific calculator.

* Remarkably, a formula involving two irrational numbers always yields Fibonacci numbers, which are *integers*.