

Math 290-1: Linear Algebra & Multivariable Calculus

Northwestern University, Lecture Notes

Written by Santiago Cañez

These are notes which provide a basic summary of each lecture for Math 290-1, the first quarter of “MENU: Linear Algebra & Multivariable Calculus”, taught by the author at Northwestern University. The book used as a reference is the 5th edition of *Linear Algebra with Applications* by Bretscher. Watch out for typos! Comments and suggestions are welcome.

Contents

Lecture 1: Introduction to Linear Systems	2
Lecture 2: Gauss-Jordan Elimination	3
Lecture 3: Solutions of Linear Systems	7
Lecture 4: More on Solutions of Systems and Vectors	11
Lecture 5: Linear Transformations	14
Lecture 6: Geometric Transformations	19
Lecture 7: Matrix Multiplication	22
Lecture 8: Invertibility and Inverses	26
Lecture 9: More on Inverses, the Amazingly Awesome Theorem	30
Lecture 10: Images and Kernels	33
Lecture 11: Subspaces of \mathbb{R}^n	38
Lecture 12: Linear Dependence/Independence and Bases	40
Lecture 13: Bases and Dimension	45
Lecture 14: Coordinates Relative to a Basis	49
Lecture 15: More on Coordinates	55
Lecture 16: Determinants	60
Lecture 17: Properties of Determinants	65
Lecture 18: Geometric Interpretation of Determinants	70
Lecture 19: Eigenvalues	75
Lecture 20: Eigenvectors	79
Lecture 21: Applications of Eigenvectors	85
Lecture 22: Diagonalization	91
Lecture 23: More on Diagonalization	96
Lecture 24: Complex Eigenvalues	102

Lecture 1: Introduction to Linear Systems

Today I gave a brief introduction to some concepts we'll be looking at this quarter, such as matrices, eigenvalues, and eigenvectors. I mentioned one or two ways in which such concepts show up in other areas.

Example 1. The system of linear equations (also known as a linear system):

$$\begin{aligned}x + 2y &= 0 \\ -3x - 2y &= 8\end{aligned}$$

has precisely one solution: $x = -4, y = 2$. Geometrically, both of these equations describe lines in the xy -plane and the existence of only one solution means that these two lines intersect in exactly one point.

Example 2. The system of linear equations:

$$\begin{aligned}x + 2y &= 0 \\ -3x - 6y &= -3\end{aligned}$$

has no solutions. Geometrically, this happens because the corresponding lines are parallel and don't intersect.

Example 3. The system of equations:

$$\begin{aligned}x + 2y &= 0 \\ -3x - 6y &= 0\end{aligned}$$

has infinitely many solutions, meaning that there are infinitely many pairs of numbers (x, y) which satisfy both equations simultaneously. Geometrically, these two equations describe the same line and so intersect everywhere.

Important. The same phenomena regarding number of solutions is true in any number of dimensions. In other words, any system of linear equations no matter how many variables or equations are involved will have exactly one solution, no solution, or infinitely many solutions.

Example 4. Consider the system:

$$\begin{aligned}x + 2y + 3z &= 0 \\ -3x - 2y - 8z &= 8 \\ 2x + 12y + z &= 2\end{aligned}$$

Geometrically, each of these equations describe planes in 3-dimensional space (we'll talk about planes a lot more when we get to multivariable calculus) and by finding the solution(s) of this system we are determining where these three planes intersect. We solve the system using what are called "row operations", and we'll describe this method in detail next time.

For now, note that multiplying the first equation by 3 gives $3x + 6y + 9z = 0$, and adding this entire equation to the second one gives $4y + z = 8$. The point is that this resulting equation no longer has an x in it, so we've "eliminated" a variable. Similarly, multiplying the first equation by -2 gives $-2x - 4y - 6z = 0$ and adding this to the third gives $8y - 5z = 2$, and again we've

eliminated x . Now consider the system keeping the first equation the same but replacing the second and third with the new ones obtained:

$$\begin{aligned}x + 2y + 3z &= 0 \\4y + z &= 8 \\8y - 5z &= 2\end{aligned}$$

The point is that this new system has precisely the same solutions as the original one! In other words, “row operations” do change the actual equations involved but do **not** change the set of solutions.

We can keep going. Now we move down to the $4y$ terms and decide we want to get rid of the $8y$ below it. We multiply the second equation by -2 and add the result to the third equation to give $-7z = -14$. Thus we get the new system

$$\begin{aligned}x + 2y + 3z &= 0 \\4y + z &= 8 \\-7z &= -14\end{aligned}$$

Now we’re in business: the third equation tells us that $z = 2$, substituting this into the second and solving for y gives $y = 3/2$, and finally substituting these two values into the first equation and solving for x gives $x = -9$. Thus this system has only solution:

$$x = -9, y = 3/2, z = 2.$$

Again, since this method does not change the solutions of the various systems of equations we use, this is also the only solution of our original system.

Lecture 2: Gauss-Jordan Elimination

Today we started talking about Gauss-Jordan Elimination, which gives us a systematic way of solving systems of linear equations. This technique is going to be the most useful computational tool we’ll have the entire quarter, and it will be very beneficial to get to the point where you can carry it out fairly quickly and without errors. Practice makes perfect! We’ll continue with examples on Monday.

Warm-Up 1. Solve the system of equations:

$$\begin{aligned}2x + 3y + z &= 0 \\x - y + z &= 2\end{aligned}$$

We use the technique of “eliminating” variables. We first multiply the second row by -2 and add the first row to it, giving $5y - z = -4$. So now we have the system

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

Now there are multiple ways we could proceed. First, we could add these two equations together and use the result to replace the first equation, giving:

$$\begin{aligned}2x + 8y &= 0 \\5y - z &= -4\end{aligned}$$

Compared to our original set of equations, these are simpler to work with. The question now is: what do we do next? Do we keep trying to eliminate variables, or move on to trying to find the solution(s)? Note that any further manipulations we do cannot possibly eliminate any more variables, since such operations will introduce a variable we've already eliminated into one of the equations. We'll see later how we can precisely tell that this is the best we can do. So, let's move towards finding solutions.

For now, we actually go back to equations we had after our first manipulations, namely:

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

We could instead try to eliminate the y term in the first equation instead of the z term as we did. This illustrates a general point: there are often multiple ways of solving these systems, and it would be good if we had a systematic way of doing so. This is what Gauss-Jordan elimination will do for us. Here, let's just stick with the above equations.

We will express the values of x and y in terms of z . The second equation gives

$$y = \frac{z - 4}{5}.$$

Plugging this in for y in the first equation and solving for x gives:

$$x = \frac{-3y - z}{2} = \frac{-3\left(\frac{z-4}{5}\right) - z}{2} = \frac{12 - 8z}{10}.$$

These equations we've derived imply that our system in fact has infinitely many solutions: for any value we assign to z , setting x equal to $\frac{12-8z}{10}$ and y equal to $\frac{z-4}{5}$ gives a triple of numbers (x, y, z) which form a solution of the original equation. Since z is "free" to take on any value, we call it a "free" variable. Thus we can express the solution of our system as

$$x = \frac{12 - 8z}{10}, \quad y = \frac{z - 4}{5}, \quad z \text{ free.}$$

Warm-Up 2. Find the polynomial function of the form $f(x) = a + bx + cx^2$ satisfying the condition that its graph passes through $(1, 1)$ and $(2, 0)$ and such that $\int_1^2 f(x) dx = -1$.

The point of this problem is understanding what this has to do with linear algebra, and the realization that systems of linear equations show up in many places. In particular, this problem boils down to solving a system of three equations in terms of the three unknown "variables" a, b , and c . The condition that the graph of $f(x)$ pass through $(1, 1)$ means that $f(1)$ should equal 1 and the condition that the graph pass through $(2, 0)$ means that $f(2)$ should equal 0. Writing out what this means, we get:

$$f(1) = 1 \text{ means } a + b + c = 1$$

and

$$f(2) = 0 \text{ means } a + 2b + 4c = 0.$$

Finally, since

$$\int_1^2 (a + bx + cx^2) dx = \left(ax + \frac{bx^2}{2} + \frac{cx^3}{3} \right) \Big|_1^2 = a + \frac{3}{2}b + \frac{7}{3}c,$$

the condition that $\int_1^2 f(x) dx = -1$ gives

$$a + \frac{3}{2}b + \frac{7}{3}c = -1.$$

In other words, the unknown coefficients a, b, c we are looking for must satisfy the system of equations:

$$\begin{aligned} a + b + c &= 1 \\ a + 2b + 4c &= 0 \\ a + \frac{3}{2}b + \frac{7}{3}c &= -1 \end{aligned}$$

Thus to find the function we want we must solve this system. We'll leave this for now and come back to it in a bit.

Augmented Matrices. From now on we will work with the “augmented matrix” of a system of equations rather than the equations themselves. The augmented matrix encodes the coefficients of all the variables as well as the numbers to the right of the equals sign. For instance, the augmented matrix of the system in the first Warm-Up is

$$\left(\begin{array}{ccc|c} 2 & 3 & 1 & 0 \\ 1 & -1 & 1 & 2 \end{array} \right).$$

The first column encodes the x coefficients, the second the y coefficients, and so on. The vertical lines just separate the values which come from coefficients of variables from the values which come from the right side of the equals sign.

Important. Gauss-Jordan elimination takes a matrix and puts it into a specialized form known as “reduced echelon form”, using “row operations” such as multiplying a row by a nonzero number, swapping rows, and adding rows together. The key point is to eliminate (i.e. turn into a 0) entries above and below “pivots”.

Example 1. We consider the system

$$\begin{aligned} x + 2y - z + 3w &= 0 \\ 2x + 4y - 2z &= 3 \\ -2x + 4z - 2w &= 1 \\ 3x + 2y + 5w &= 1 \end{aligned}$$

with augmented matrix:

$$\left(\begin{array}{cccc|c} \mathbf{1} & 2 & -1 & 3 & 0 \\ 2 & 4 & -2 & 0 & 3 \\ -2 & 0 & 4 & -2 & 1 \\ 3 & 2 & 0 & 5 & 1 \end{array} \right).$$

The *pivot* (i.e. first nonzero entry) of the first row is in red. Our first goal is to turn every entry below this pivot into a 0. We do this using the row operations:

$$-2I + II \rightarrow II, 2I + III \rightarrow III, \text{ and } -3I + IV \rightarrow IV,$$

where the roman numerals denote row numbers and something like $-3I + IV \rightarrow IV$ means multiply the first row by -3 , add that to the fourth row, and put the result into the fourth row. These operations produce

$$\left(\begin{array}{cccc|c} \mathbf{1} & 2 & -1 & 3 & 0 \\ 0 & 0 & 0 & -6 & 3 \\ 0 & 4 & 2 & 4 & 1 \\ 0 & -4 & 3 & -4 & 1 \end{array} \right),$$

where we have all zeros below the first pivot.

Now we move to the second row. Ideally we want the pivot in the second row to be diagonally down from the pivot in the first row, but in this case it's not—the -6 is further to the right. So, here a row swap is appropriate in order to get the pivot of the second row where we want it. Swapping the second and fourth rows gives

$$\left(\begin{array}{cccc|c} 1 & 2 & -1 & 3 & 0 \\ 0 & -4 & 3 & -4 & 1 \\ 0 & 4 & 2 & 4 & 1 \\ 0 & 0 & 0 & -6 & 3 \end{array} \right).$$

Our next goal is to get rid of the entries above and below the pivot -4 of the second row. For this we use the row operations:

$$II + III \rightarrow III \text{ and } 2I + II \rightarrow I.$$

This gives

$$\left(\begin{array}{cccc|c} 2 & 0 & 1 & 2 & 1 \\ 0 & -4 & 3 & -4 & 1 \\ 0 & 0 & 5 & 0 & 2 \\ 0 & 0 & 0 & -6 & 3 \end{array} \right).$$

Now onto the third row and getting rid of entries above and below its pivot 5. Note that the point of swapping the second and fourth rows earlier as opposed to the second and third is that now we already have a zero below the 5, so we only have to worry about the entries above the 5. The next set of row operations ($5II - 3III \rightarrow II$ and $-5I + III \rightarrow I$) give

$$\left(\begin{array}{cccc|c} -10 & 0 & 0 & -10 & -3 \\ 0 & -20 & 0 & -20 & -1 \\ 0 & 0 & 5 & 0 & 2 \\ 0 & 0 & 0 & -6 & 3 \end{array} \right).$$

Finally, we move to the final pivot -6 in the last row and make all entries above it using the operations

$$-3I + (5)IV \rightarrow I \text{ and } -3II + (10)IV \rightarrow II.$$

This gives

$$\left(\begin{array}{cccc|c} 30 & 0 & 0 & 0 & 24 \\ 0 & 60 & 0 & 0 & 33 \\ 0 & 0 & 5 & 0 & 2 \\ 0 & 0 & 0 & -6 & 3 \end{array} \right).$$

As we wanted, all entries above and below pivots are zero. The final step to get to so-called “reduced echelon form” is to make all pivots one, by dividing each row by the appropriate value. So, we divide the first row by 30, the second by 60, third by 5, and fourth by -6 to get:

$$\left(\begin{array}{cccc|c} 1 & 0 & 0 & 0 & 24/30 \\ 0 & 1 & 0 & 0 & 33/60 \\ 0 & 0 & 1 & 0 & 2/5 \\ 0 & 0 & 0 & 1 & -1/2 \end{array} \right).$$

This matrix is now in reduced echelon form. Looking at the corresponding system of equations, the point is that we've now eliminated all variables but one in each equation. Right away, writing

down this corresponding system we get that

$$x = \frac{24}{30}, y = \frac{33}{60}, z = \frac{2}{5}, w = -\frac{1}{2}$$

is the only solution to our original system of equations.

Important. The characteristic properties of a matrix in “reduced echelon form” are: all entries above and below pivots are 0, each pivot occurs strictly to the right of any pivot above it, and all pivots are 1. This is what we aim for when performing Gauss-Jordan elimination.

Back to Warm-Up 2. Let’s finish up the second Warm-Up problem. We are left with solving

$$\begin{aligned} a + b + c &= 1 \\ a + 2b + 4c &= 0 \\ a + \frac{3}{2}b + \frac{7}{3}c &= -1 \end{aligned}$$

for a, b , and c . We “row reduce” the augmented matrix:

$$\left(\begin{array}{ccc|c} 1 & 1 & 1 & 1 \\ 1 & 2 & 4 & 0 \\ 1 & 3/2 & 7/3 & -1 \end{array} \right).$$

To avoid dealing with fractions, we first multiply the third row by 6. Performing various row operations gives:

$$\begin{aligned} \left(\begin{array}{ccc|c} 1 & 1 & 1 & 1 \\ 1 & 2 & 4 & 0 \\ 6 & 9 & 14 & -6 \end{array} \right) &\rightarrow \left(\begin{array}{ccc|c} 1 & 1 & 1 & 1 \\ 0 & -1 & -3 & 1 \\ 0 & 3 & 8 & -12 \end{array} \right) \rightarrow \left(\begin{array}{ccc|c} 1 & 0 & -2 & 2 \\ 0 & -1 & -3 & 1 \\ 0 & 0 & -1 & -9 \end{array} \right) \\ &\rightarrow \left(\begin{array}{ccc|c} 1 & 0 & 0 & 20 \\ 0 & -1 & 0 & 28 \\ 0 & 0 & -1 & -9 \end{array} \right) \rightarrow \left(\begin{array}{ccc|c} 1 & 0 & 0 & 20 \\ 0 & 1 & 0 & -28 \\ 0 & 0 & 1 & 9 \end{array} \right). \end{aligned}$$

The corresponding system of equations is

$$a = 20, b = -28, c = 9$$

and have found our desired unknown value. The conclusion is that the function $f(x) = 20 - 28x + 9x^2$ is the one satisfying the properties asked for in the second Warm-Up.

Lecture 3: Solutions of Linear Systems

Today we continued talking about solving systems of linear equations, and started talking about vectors and how they provide an alternate way to think about systems.

Warm-Up 1. We solve the following system of linear equations:

$$\begin{aligned} -2x_1 - 4x_2 - 2x_3 - 3x_4 - 3x_5 &= -5 \\ x_1 + 2x_2 + x_3 + 4x_4 - x_5 &= 5 \\ 3x_1 + 6x_2 + 5x_3 + 10x_4 - 4x_5 &= 14 \\ -x_1 - 2x_2 + x_3 - 2x_4 - 4x_5 &= 2 \end{aligned}$$

using Gauss-Jordan elimination. First, we switch the first two rows in the augmented matrix in order to have 1 in the uppermost position instead of -2 —this will help with computations. The augmented matrix is

$$\left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ -2 & -4 & -2 & -3 & -3 & -5 \\ 3 & 6 & 5 & 10 & -4 & 14 \\ -1 & -2 & 1 & -2 & -4 & 2 \end{array} \right).$$

Performing the row operations $2I + II \rightarrow II$, $-3I + III \rightarrow III$, and $I + IV \rightarrow IV$ gives:

$$\left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ -2 & -4 & -2 & -3 & -3 & -5 \\ 3 & 6 & 5 & 10 & -4 & 14 \\ -1 & -2 & 1 & -2 & -4 & 2 \end{array} \right) \rightarrow \left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ 0 & 0 & 0 & 5 & -5 & 5 \\ 0 & 0 & 2 & -2 & -1 & -1 \\ 0 & 0 & 2 & 2 & -5 & 7 \end{array} \right).$$

Now, there can be no pivot in the second column since the entries in the second, third, and fourth rows are 0. The best place for the next pivot would be the third entry of the second row, so to get a pivot here we switch the second and fourth rows:

$$\left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ 0 & 0 & 2 & 2 & -5 & 7 \\ 0 & 0 & 2 & -2 & -1 & -1 \\ 0 & 0 & 0 & 5 & -5 & 5 \end{array} \right).$$

We perform the row operation $II - III \rightarrow III$:

$$\left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ 0 & 0 & 2 & 2 & -5 & 7 \\ 0 & 0 & 2 & -2 & -1 & -1 \\ 0 & 0 & 0 & 5 & -5 & 5 \end{array} \right) \rightarrow \left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ 0 & 0 & 2 & 2 & -5 & 7 \\ 0 & 0 & 0 & 4 & -4 & 8 \\ 0 & 0 & 0 & 5 & -5 & 5 \end{array} \right).$$

In usual Gauss-Jordan elimination we would also want to eliminate the 1 above the pivot 2 in the second row, but for now we skip this. To simplify some computations, we next divide the third row by 4 to get

$$\left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ 0 & 0 & 2 & 2 & -5 & 7 \\ 0 & 0 & 0 & 1 & -1 & 2 \\ 0 & 0 & 0 & 5 & -5 & 5 \end{array} \right).$$

We perform the row operation $-5III + IV \rightarrow IV$:

$$\left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ 0 & 0 & 2 & 2 & -5 & 7 \\ 0 & 0 & 0 & 1 & -1 & 2 \\ 0 & 0 & 0 & 0 & 0 & -5 \end{array} \right).$$

Since the last row corresponds to the impossible equation $0 = -5$, the original system has no solutions. Note that we did not have to do a full Gauss-Jordan elimination to determine this.

Important. If you are only interested in determining whether there is a solution, or how many there are, a full Gauss-Jordan elimination is not needed. Only use a full elimination process when trying to actually describe all solutions.

Warm-Up 2. We consider the same system as before, only changing the 2 at the end of the last equation to -2 . If you follow through the same row operations as before, you end up with the augmented matrix:

$$\left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ 0 & 0 & 2 & 2 & -5 & 3 \\ 0 & 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & 0 & 5 & -5 & 5 \end{array}\right) \rightarrow \left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ 0 & 0 & 2 & 2 & -5 & 3 \\ 0 & 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array}\right).$$

We no longer have the issue we had before, so here we will have a solution, and in fact infinitely many. We first do $-2I + II \rightarrow I$ to get rid of the 1 above the pivot 2, which we skipped in the first Warm-Up:

$$\left(\begin{array}{ccccc|c} 1 & 2 & 1 & 4 & -1 & 5 \\ 0 & 0 & 2 & 2 & -5 & 3 \\ 0 & 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array}\right) \rightarrow \left(\begin{array}{ccccc|c} -2 & -4 & 0 & -6 & -3 & -7 \\ 0 & 0 & 2 & 2 & -5 & 3 \\ 0 & 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array}\right).$$

Next we do $6III + I \rightarrow I$ and $-2III + II \rightarrow II$ to get rid of the entries above the pivot 1:

$$\left(\begin{array}{ccccc|c} -2 & -4 & 0 & -6 & -3 & -7 \\ 0 & 0 & 2 & 2 & -5 & 3 \\ 0 & 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array}\right) \rightarrow \left(\begin{array}{ccccc|c} -2 & -4 & 0 & 0 & -9 & -1 \\ 0 & 0 & 2 & 0 & -3 & 1 \\ 0 & 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array}\right).$$

Finally we divide the first row by -2 and the second by 2 to get:

$$\left(\begin{array}{ccccc|c} 1 & 2 & 0 & 0 & 9/2 & 1/2 \\ 0 & 0 & 1 & 0 & -3/2 & 1/2 \\ 0 & 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array}\right).$$

This matrix is now in what's called *row-reduced echelon form* since all pivots are 1, all entries above and below pivots are zero, and each pivot occurs strictly to the right of any pivot above it.

The variables which don't correspond to pivots are the ones we call "free" variables, and when writing down the general form of the solution we express all "pivot" variables in terms of the free ones. The *rank* (i.e. # of pivots in the reduced echelon form) of this matrix, or of the original one we started with, is 3. This final augmented matrix corresponds to the system:

$$\begin{aligned} x_1 + 2x_2 &+ \frac{9}{2}x_5 = \frac{1}{2} \\ x_3 - \frac{3}{2}x_5 &= \frac{1}{2} \\ x_4 - x_5 &= 2 \end{aligned}$$

so we get

$$x_1 = -2x_2 - \frac{9}{2}x_5 + \frac{1}{2}, \quad x_3 = \frac{3}{2}x_5 + \frac{1}{2}, \quad x_4 = x_5 + 1$$

with x_2 and x_5 free. In so-called "vector form", the general solution is

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{pmatrix} = \begin{pmatrix} -2s - \frac{9}{2}t + \frac{1}{2} \\ s \\ \frac{3}{2}t + \frac{1}{2} \\ t + 1 \\ t \end{pmatrix}$$

where s and t are arbitrary numbers.

Fact about reduced echelon form. It is possible to get from one matrix to another using a sequence of row operations precisely when they have the same reduced echelon form. For instance, since the reduced echelon form of

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 10 \end{pmatrix} \text{ and } \begin{pmatrix} 1 & 4 & 3 \\ 9 & 1 & -10 \\ 89 & \pi & 23838 \end{pmatrix}$$

both have

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

as their reduced echelon form, it is possible to get from one to the other by some sequence of row operations.

Relation between rank and number of solutions. Based on the form of the reduced echelon form of a matrix, there is a strong relation between the rank of a matrix and the number of solutions of a system having that matrix as its coefficients. For instance, any system where the rank is $<$ the number of variables cannot possibly have a unique solution. Also, any system where the rank equals the number of variables cannot possibly have an infinite number of solutions. We will explore this further later, but check the book for similar facts.

Vectors. A *vector* is a matrix with one column, and is said to be in \mathbb{R}^n when it has n entries. (\mathbb{R} is a common notation for the set of real numbers.) For instance,

$$\begin{pmatrix} 1 \\ 2 \end{pmatrix} \text{ is in } \mathbb{R}^2, \text{ and } \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} \text{ is in } \mathbb{R}^3.$$

We draw vectors as arrows starting at the “origin” $\begin{pmatrix} 0 \\ 0 \end{pmatrix}$ and ending at the point determined by the vector’s entries. We add vectors simply by adding the corresponding entries together, and multiple vectors by scalars (i.e. numbers) simply by multiplying each entry of the vector by that scalar.

Relation between vectors and linear systems. Consider the system

$$\begin{aligned} 2x + 3y + z &= 0 \\ x - y + z &= 2 \end{aligned}$$

from a previous example. If we imagine each row as the entry of a vector, we can write the entire left-hand side as the vector sum:

$$x \begin{pmatrix} 2 \\ 1 \end{pmatrix} + y \begin{pmatrix} 3 \\ -1 \end{pmatrix} + z \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

The right side is the vector $\begin{pmatrix} 0 \\ 2 \end{pmatrix}$, so the given system can be written as

$$x \begin{pmatrix} 2 \\ 1 \end{pmatrix} + y \begin{pmatrix} 3 \\ -1 \end{pmatrix} + z \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 2 \end{pmatrix}.$$

The upshot is that we have transformed the original system of equations into a single vector equation. This will be a jumping off point into a wide range of new topics, and we will come back to it Wednesday.

Lecture 4: More on Solutions of Systems and Vectors

Today we continued talking about the relation between solutions of linear systems and vectors, and the various ways of representing systems in terms of vectors and matrices.

Warm-Up. Is there a 3×4 matrix A of rank 3 such that the system with augmented matrix

$$\left(\begin{array}{c|c} A & \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} \end{array} \right)$$


has a unique solution? If you think about what the reduced echelon form of A looks like, we know that it should have 3 pivots. However, with 4 columns this means that one column won't have a pivot and so will correspond to a free variable. This means that it is not possible for such a system to have exactly one solution: either it will have no solutions or infinitely many depending on whether the last row in the reduced echelon form corresponds to $0 = 0$ or some impossible equation. The key point is understanding the relation between number of pivots and number of solutions.

As a contrast, we ask if there is a 4×3 matrix A of rank 3 such that the system with augmented matrix


$$(A \mid \vec{0}),$$

where $\vec{0}$ denotes the zero vector in \mathbb{R}^4 , has a unique solution. In this case, in fact for any such matrix this system will have a unique solution. Again the reduced form will have 3 pivots, but now with A having only 3 columns there won't be a column without a pivot and so no free variables. Since we started with the "augmented" piece (i.e. the final column corresponding to the numbers on the right side of equals signs in the corresponding system) consisting of all zeroes, any row operations which transform the A part into reduced form will still result in the final column being all zeroes. Thus there are no contradictions like " $0 = 1$ " and so such a system will always have a unique solution, namely the one where all variables equal 0.

Geometric meaning of vector addition and scalar multiplication. Given a vector \vec{x} and a scalar r , the vector $r\vec{x}$ points in parallel to \vec{x} but its length is scaled by a factor of r ; for negative r the direction is turned around:



Given vectors \vec{x} and \vec{y} , their sum $\vec{x} + \vec{y}$ is the vector which forms the diagonal of the parallelogram with sides \vec{x} and \vec{y} :



Linear combinations. Recall that previously we saw how to write the system

$$2x + 3y + z = 0$$

$$x - y + z = 2$$

as the vector equation

$$x \begin{pmatrix} 2 \\ 1 \end{pmatrix} + y \begin{pmatrix} 3 \\ -1 \end{pmatrix} + z \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 2 \end{pmatrix}.$$

The expression on the left is what is called a *linear combination* of


$$\begin{pmatrix} 2 \\ 1 \end{pmatrix}, \begin{pmatrix} 3 \\ -1 \end{pmatrix}, \text{ and } \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

So, asking whether or not $\begin{pmatrix} 0 \\ 2 \end{pmatrix}$ can be expressed as such a linear combination is the same as asking whether or not the corresponding system has a solution. We already know from previous examples that this system has infinitely many solutions, but let us now understand why this is true geometrically.

Consider first the simpler vector equation given by

$$x \begin{pmatrix} 2 \\ 1 \end{pmatrix} + y \begin{pmatrix} 3 \\ -1 \end{pmatrix} = \begin{pmatrix} 0 \\ 2 \end{pmatrix}. \tag{1}$$


Using the geometric interpretations of vector addition and scalar multiplication, it makes sense that this system has a unique solution since we can eyeball that there are specific scalars x and y we can use to scale $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 3 \\ -1 \end{pmatrix}$ and have the results add up to $\begin{pmatrix} 0 \\ 2 \end{pmatrix}$.



Similarly, the vector equation given by

$$y \begin{pmatrix} 3 \\ -1 \end{pmatrix} + z \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 2 \end{pmatrix} \quad (2)$$

should also have a unique solution, based on the picture:



Now we go back to our original vector equation. The solution for x and y in equation 1 together with $z = 0$ gives a solution of

$$x \begin{pmatrix} 2 \\ 1 \end{pmatrix} + y \begin{pmatrix} 3 \\ -1 \end{pmatrix} + z \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 2 \end{pmatrix}.$$

In the same manner the solution for y and z in equation 2 together with $x = 0$ gives another solution of this same vector equation. Since this vector equation now has at least two solutions, it in fact must have infinitely many since we know that any system has either no, one, or infinitely many solutions. This agrees with what we found previously when solving this system algebraically.

The point is that now that we've rewritten systems in terms of vectors, we have new geometric ideas and techniques available at our disposal when understanding what it means to solve a system of linear equations.

Matrix forms of systems. Continuing on with the same example, the expression

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

is also what we call the result of *multiplying* the matrix $\begin{pmatrix} 2 & 3 & 1 \\ 1 & -1 & 1 \end{pmatrix}$ by the vector $\begin{pmatrix} x \\ y \\ z \end{pmatrix}$:

$$\begin{pmatrix} 2 & 3 & 1 \\ 1 & -1 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = x \begin{pmatrix} 2 \\ 1 \end{pmatrix} + y \begin{pmatrix} 3 \\ -1 \end{pmatrix} + z \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

Thus the system we are considering can also be written in *matrix form* as

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

Solving this “matrix equation” for $\begin{pmatrix} x \\ y \end{pmatrix}$ is the same as solving the original system of equations we considered, which is also the same as solving the vector equation we had before. The idea we will expand on in the coming weeks is that understanding more about matrices and such matrix equations will give us yet another point of view on what it means to solve a system of linear equations.

Important. Systems of linear equations, vector equations involving linear combinations, and matrix equations are all different ways of looking at the same type of problem. These different points of view allow the use of different techniques, and help to make systems more “geometric”.

Lecture 5: Linear Transformations

Today we started talking about what are called “linear transformations” and their relation to matrices.

Warm-Up. Is the vector $\begin{pmatrix} 4 \\ 0 \\ 2 \end{pmatrix}$ a linear combination of the vectors

$$\begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}, \begin{pmatrix} 3 \\ 2 \\ -1 \end{pmatrix}, \text{ and } \begin{pmatrix} -1 \\ -2 \\ -1 \end{pmatrix}?$$

Recall that a linear combination of these three vectors is an expression of the form

$$a \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} + b \begin{pmatrix} 3 \\ 2 \\ -1 \end{pmatrix} + c \begin{pmatrix} -1 \\ -2 \\ -1 \end{pmatrix},$$

so we are asking whether there exist scalars a, b, c such that

$$a \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} + b \begin{pmatrix} 3 \\ 2 \\ -1 \end{pmatrix} + c \begin{pmatrix} -1 \\ -2 \\ -1 \end{pmatrix} = \begin{pmatrix} 4 \\ 0 \\ 2 \end{pmatrix}.$$

After adding together the entire left side this becomes

$$\begin{pmatrix} a + 3b - c \\ 2b - 2c \\ -a - b - c \end{pmatrix} = \begin{pmatrix} 4 \\ 0 \\ 2 \end{pmatrix},$$

so our problem boils down to asking whether this system has a solution. Reducing the augmented matrix gives

$$\left(\begin{array}{ccc|c} 1 & 3 & -1 & 4 \\ 0 & 2 & -2 & 0 \\ -1 & -1 & -1 & 2 \end{array} \right) \rightarrow \left(\begin{array}{ccc|c} 1 & 3 & -1 & 4 \\ 0 & 2 & -2 & 0 \\ 0 & 2 & -2 & 6 \end{array} \right) \rightarrow \left(\begin{array}{ccc|c} 1 & 3 & -1 & 4 \\ 0 & 2 & -2 & 0 \\ 0 & 0 & 0 & 6 \end{array} \right),$$

so we see that there is no solution. Hence the vector $\begin{pmatrix} 4 \\ 0 \\ 2 \end{pmatrix}$ is not a linear combination of the three given vectors.

Instead we now ask if $\begin{pmatrix} 4 \\ 2 \\ -2 \end{pmatrix}$ is a linear combination of the three given vectors. This is asking whether

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

has a solution, which is the same as asking whether the system

$$\begin{aligned} a + 3b - c &= 4 \\ 2b - 2c &= 2 \\ -a - b - c &= -2 \end{aligned}$$

has a solution. Reducing the corresponding augmented matrix gives

$$\left(\begin{array}{ccc|c} 1 & 3 & -1 & 4 \\ 0 & 2 & -2 & 2 \\ -1 & -1 & -1 & -2 \end{array} \right) \rightarrow \left(\begin{array}{ccc|c} 1 & 3 & -1 & 4 \\ 0 & 2 & -2 & 2 \\ 0 & 2 & -2 & 2 \end{array} \right) \rightarrow \left(\begin{array}{ccc|c} 1 & 3 & -1 & 4 \\ 0 & 2 & -2 & 2 \\ 0 & 0 & 0 & 0 \end{array} \right),$$

at which point we know that there will be a solution. Thus $\begin{pmatrix} 4 \\ 2 \\ -2 \end{pmatrix}$ is a linear combination of the three given vectors. To be precise, continuing on and solving this system completely gives $a = -1, b = 2, c = 1$ as one solution (there are infinitely many others), and you can check that

$$- \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} + 2 \begin{pmatrix} 3 \\ 2 \\ -1 \end{pmatrix} + \begin{pmatrix} -1 \\ -2 \\ -1 \end{pmatrix}$$

indeed equals $\begin{pmatrix} 4 \\ 2 \\ -2 \end{pmatrix}$.

Important. Questions about linear combinations often boil down to solving some system of linear equations.

Example 1. Consider the function T from \mathbb{R}^2 to \mathbb{R}^2 defined by

$$T(\vec{x}) = A\vec{x} \text{ where } A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}.$$

To be clear, T is the function which takes a vector \vec{x} in \mathbb{R}^2 as an “input” and “outputs” the vector resulting from multiplying A by \vec{x} . We also call T a “transformation” since it transforms vectors in \mathbb{R}^2 somehow to produce other vectors.

To get a sense for what T is doing “geometrically”, consider what happens T is applied to $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$. Inputting the first vector into T gives

$$T \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

so geometrically T has the effect of transforming $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ into $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$. Similarly,

$$T \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} -1 \\ 0 \end{pmatrix},$$

so T has the effect of transforming $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ into $\begin{pmatrix} -1 \\ 0 \end{pmatrix}$. Note that in this case, $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ is nothing but what you get when rotating $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ counterclockwise by 90° and $\begin{pmatrix} -1 \\ 0 \end{pmatrix}$ is what you get when rotating $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ by the same amount. So we would say that geometrically applying T has the effect of rotating $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ by 90° . In fact, as we will see, it turns out that this is the effect of T on *any* possible input: given a vector \vec{x} in \mathbb{R}^2 , the vector $T(\vec{x})$ obtained after applying T is the vector you get when rotating \vec{x} (visualized as an arrow) by 90° .

The above transformation T has the property that for any input vectors \vec{x} and \vec{y} we have

$$T(\vec{x} + \vec{y}) = T(\vec{x}) + T(\vec{y}),$$

and for any scalar r we have

$$T(r\vec{x}) = rT(\vec{x}).$$

(This just says that $A(\vec{x} + \vec{y}) = A\vec{x} + A\vec{y}$ and $A(r\vec{x}) = rA\vec{x}$, which we will see later are general properties of matrix multiplication.) The first equality says that when taking two inputs vectors, it does matter whether we add them together first and then apply T to the result, or apply T to the two inputs separately and then add, we will always get the same result. The second equality says that scaling an input vector \vec{x} by r and then applying T is the same as applying T to \vec{x} first and then scaling by r . Both of these properties should make sense geometrically since T after all is nothing but a “rotation transformation”. These two properties together give us the first definition of what it means to say a function is a “linear transformation”.

First definition of Linear Transformation. A function T from some space \mathbb{R}^m to some space \mathbb{R}^n is a *linear transformation* if it has the properties that

$$T(\vec{x} + \vec{y}) = T(\vec{x}) + T(\vec{y})$$

for any inputs \vec{x} and \vec{y} in \mathbb{R}^m , and

$$T(r\vec{x}) = rT(\vec{x})$$

for any input \vec{x} in \mathbb{R}^m and any scalar r .

Example 2. Consider the transformation T from \mathbb{R}^2 to \mathbb{R}^2 defined by

$$T \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 2x + 3y + 2 \\ -3x - 4y \end{pmatrix}.$$

We claim that this is not linear. The first requirement says that given two input vectors, adding them together and then applying T should give the same result as applying T to both inputs separately and then adding. Take two arbitrary inputs $\begin{pmatrix} x \\ y \end{pmatrix}$ and $\begin{pmatrix} a \\ b \end{pmatrix}$. Then

$$T \left(\begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} a \\ b \end{pmatrix} \right) = T \begin{pmatrix} x + a \\ y + b \end{pmatrix} = \begin{pmatrix} 2(x + a) + 3(y + b) + 2 \\ -3(x + a) - 4(y + b) \end{pmatrix}.$$

However, this is not the same result as

$$T \begin{pmatrix} x \\ y \end{pmatrix} + T \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} 2x + 3y + 2 \\ -3x - 4y \end{pmatrix} + \begin{pmatrix} 2a + 3b + 2 \\ -3a - 4b \end{pmatrix} = \begin{pmatrix} 2(x + a) + 3(y + b) + 4 \\ -3(x + a) - 4(y + b) \end{pmatrix}.$$

The problem is that in this final expression we have a $+4$ in the first entry but in the one we computed before we only have a $+2$. Thus this transformation does not satisfy

$$T(\vec{x} + \vec{y}) = T(\vec{x}) + T(\vec{y})$$

so T is not a linear transformation.

Even though we already know T is not linear due to the fact that it fails the first requirement in the definition of a linear transformation, for good measure let's note that it also fails the second requirement. Indeed, if r is any scalar we have

$$T\left(r\begin{pmatrix} x \\ y \end{pmatrix}\right) = T\begin{pmatrix} rx \\ ry \end{pmatrix} = \begin{pmatrix} 2rx + 3ry + 2 \\ -3rx - 4yr \end{pmatrix},$$

which is not the same as

$$rT\begin{pmatrix} x \\ y \end{pmatrix} = r\begin{pmatrix} 2x + 3y + 2 \\ -3x - 4y \end{pmatrix} = \begin{pmatrix} 2rx + 3ry + 2r \\ -3rx + 4ry \end{pmatrix}$$

due to the $2r$ term as opposed to simply 2. Thus $T(r\vec{x}) \neq rT(\vec{x})$ so T also fails the second requirement in the definition of linear.

Important. In general it is possible for a function to satisfy neither property in the definition of a linear transformation or to satisfy only one. To be a linear transformation it must satisfy both.

Example 3. Now consider the transformation S from \mathbb{R}^2 to \mathbb{R}^2 defined by

$$S\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 2x + 3y \\ -3x - 4y \end{pmatrix}.$$

This is almost the same as T above only without the $+2$ term in the first entry. If you look back to the computations we did above, without this $+2$ it turns out that

$$S(\vec{x} + \vec{a}) \text{ is the same as } S(\vec{x}) + S(\vec{a})$$

and

$$S(r\vec{x}) \text{ is the same as } rS(\vec{x}).$$

Thus S is a linear transformation. In general the formula for a linear transformation should only involve “linear” (i.e. to the first power) terms in the input variables (so nothing like x^2 , xy , or $\sin y$) and should not have an extra constants added on.

Now, notice that the formula for S here can be rewritten in terms of matrix multiplication as

$$S\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 2x + 3y \\ -3x - 4y \end{pmatrix} = \begin{pmatrix} 2 & 3 \\ -3 & -4 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

So, S is actually a “matrix transformation”, just as the transformation in Example 1 was. It turns out that this is true of *any* linear transformation, giving us our second way to define what it means for a transformation to be linear.

Second definition of Linear Transformation. A function T from \mathbb{R}^m to \mathbb{R}^n is a *linear transformation* if there is some $n \times m$ matrix A with the property that

$$T(\vec{x}) = A\vec{x},$$

that is, applying T to an input vector is the same as multiplying that vector by A . We call A the *matrix of the transformation*.

Example 4. Consider the function L from \mathbb{R}^3 to \mathbb{R}^3 defined by

$$L \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} y \\ z \\ x \end{pmatrix}.$$

This is a linear transformation since the above formula is the same as

$$L \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix}.$$

The matrix of the linear transformation L is thus $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$.

The two definitions of linear are the same. The second definition of linear transformation is the one the book gives as the main definition, and then later on it talks about the first definition. I think it is conceptually better to give the first definition as the main one, and to then realize that such things are represented by matrices. Both definitions are equivalent in the sense that a function satisfying one must satisfy the other. In particular, if T is a function from \mathbb{R}^m to \mathbb{R}^n satisfying the first definition here is how we can see that it will also satisfy the second.

To keep notation simple we only focus on the case of a transformation T from \mathbb{R}^3 to \mathbb{R}^3 . The key point is that for any input vector $\begin{pmatrix} x \\ y \\ z \end{pmatrix}$, we can express it as

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

Then using the properties in the first definition of linear transformation we have:

$$\begin{aligned} T \begin{pmatrix} x \\ y \\ z \end{pmatrix} &= T \left(x \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + y \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} + z \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right) \\ &= T \left(x \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \right) + T \left(y \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \right) + T \left(z \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right) \\ &= xT \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + yT \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} + zT \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \\ &= \left(T \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \quad T \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \quad T \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right) \begin{pmatrix} x \\ y \\ z \end{pmatrix}. \end{aligned}$$

This shows that applying T to a vector is the same as multiplying that vector by the matrix whose first column is the result of applying T to $\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$, second column the result of applying T to $\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$, and third column the result of applying T to $\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$. Thus T is a linear transformation according to the second definition with matrix given by

$$\left(T \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \quad T \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \quad T \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right).$$

A similar reasoning works for T from \mathbb{R}^m to \mathbb{R}^n in general, not just \mathbb{R}^3 to \mathbb{R}^3 .

Important. For any linear transformation, the matrix which represents it is always found in the same way: we determine what T does to input vectors with a single entry equal to 1 and other entries equal to 0, and use these results of these computations as the columns of the matrix.

Lecture 6: Geometric Transformations

Today we looked at some “basic” types of geometric transformations and the matrices which represent them.

Warm-Up 1. Consider the linear transformation T from \mathbb{R}^2 to \mathbb{R}^2 defined by

$$T(\vec{x}) = \begin{pmatrix} 1 & 2 \\ -2 & -4 \end{pmatrix} \vec{x}.$$

Find all vectors \vec{x} such that $T(\vec{x}) = \vec{0}$. Said another way, find all vectors which T sends to the zero vector.

Setting $\vec{x} = \begin{pmatrix} a \\ b \end{pmatrix}$, we want all values of a and b such that

$$T \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} 1 & 2 \\ -2 & -4 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

In other words, we are asked to solve the matrix equation

$$\begin{pmatrix} 1 & 2 \\ -2 & -4 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

which in turn means we have to solve the system

$$\begin{aligned} a + 2b &= 0 \\ -2a - 4b &= 0 \end{aligned}$$

This is straightforward, and the general solution turns out to be

$$\begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} -2t \\ t \end{pmatrix},$$

but the point here is in realizing that the original question just boils down to solving some system. Again, a common phenomenon we will see throughout the course. Writing the general solution above once more by “factoring out t ” as

$$\begin{pmatrix} a \\ b \end{pmatrix} = t \begin{pmatrix} -2 \\ 1 \end{pmatrix},$$

we see that the collection of all such vectors are simply the scalar multiples of $\begin{pmatrix} -2 \\ 1 \end{pmatrix}$, which geometrically in \mathbb{R}^2 looks like the line passing through $\begin{pmatrix} -2 \\ 1 \end{pmatrix}$. Thus all vectors \vec{x} such that $T(\vec{x}) = \vec{0}$ are precisely the vectors on this line.

Warm-Up 2. Suppose that T is a linear transformation from \mathbb{R}^2 to \mathbb{R}^2 such that

$$T \begin{pmatrix} 2 \\ -1 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \text{ and } T \begin{pmatrix} 0 \\ 3 \end{pmatrix} = \begin{pmatrix} 1 \\ -1 \end{pmatrix}.$$

Find the matrix of T .

The point here is that we are only given two pieces of information about T , and yet we will be able to completely determine what T does when applied to *anything*. As usual, the matrix of T has first column equal to the result of applying T to $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and second column equal to the result of applying T to $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$. We must determine these two resulting vectors, and it is the properties spelled out in the first definition of linear transformations we gave last time which come to the rescue. First, since $\begin{pmatrix} 1 \\ 3 \end{pmatrix} = 3 \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ we have

$$\begin{pmatrix} 1 \\ -1 \end{pmatrix} = T \begin{pmatrix} 0 \\ 3 \end{pmatrix} = T \left(3 \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right) = 3T \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

so dividing by 3 gives

$$T \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 1/3 \\ -1/3 \end{pmatrix}.$$

This is thus the second column of the matrix of T . To determine $T \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, we start with the fact that

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

Then rearranging terms we have

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

Again using the properties in the first definition of linear transformation, we compute:

$$T \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \frac{1}{2} \left(T \begin{pmatrix} 2 \\ -1 \end{pmatrix} + T \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right) = \frac{1}{2} \left(\begin{pmatrix} 0 \\ 1 \end{pmatrix} + (1/3 \quad -1/3) \right) = \begin{pmatrix} 1/6 \\ 1/3 \end{pmatrix},$$

which is the first column of the matrix of T . The matrix of T is thus

$$\begin{pmatrix} 1/6 & 1/3 \\ 1/3 & -1/3 \end{pmatrix}$$


and you can double check that multiplying this by each of the original input vectors given in the setup indeed results in the corresponding output. The upshot is that now we can compute the result of applying T to any vector simply by multiplying that vector by this matrix, even though we were only given two pieces of information about T to begin with.

Geometric Transformations. Various geometric transformations in 2 and 3 dimensions are linear and so can be represented by matrices. That these are linear follows from using the geometric interpretations of vector addition and scalar multiplication to convince yourselves that

$$T(\vec{x} + \vec{y}) = T(\vec{x}) + T(\vec{y}) \text{ and } T(r\vec{x}) = rT(\vec{x}).$$

Without knowing that these geometric transformations satisfied these two properties we would have a really hard time guessing that they were represented by matrices. However, now knowing that this is the case, in order to find the corresponding matrices all we have to do is determine the result of applying these transformations to the “standard” vectors $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ in two dimensions, and their analogs in three dimensions.

Example 1. Let T be the transformation from \mathbb{R}^2 to \mathbb{R}^2 which rotates the xy -plane (counterclockwise) by an angle θ . Rotating the standard vectors gives



The x and y -coordinates of the vector obtained by rotating $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ are $\cos \theta$ and $\sin \theta$ respectively, so

$$T \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix}.$$

The x and y -coordinates of the vector obtained by rotating $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ are $-\sin \theta$ (negative since the result is in the negative x -direction) and $\cos \theta$ respectively, so

$$T \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} -\sin \theta \\ \cos \theta \end{pmatrix}.$$

Thus the matrix of T (having these two as columns) is

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

The point is that multiplying any vector by this matrix has the same effect as rotating that vector by θ .

Example 2. Say S is the transformation from \mathbb{R}^3 to \mathbb{R}^3 which rotates by θ around the x -axis when viewed from the positive x -axis. So, if we draw the y and z -axes respectively horizontally and vertically on this page with the positive x -axis coming out at us, this rotations just rotates this page counterclockwise by θ . Under this rotations, the vector $\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$ is left unchanged so

$$S \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}.$$

Rotating $\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$ gives

$$S \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ \cos \theta \\ \sin \theta \end{pmatrix} \text{ and } S \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ -\sin \theta \\ \cos \theta \end{pmatrix},$$

so the matrix of this three dimensional rotation is

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{pmatrix}.$$

Example 3. The matrix of reflection across the line $y = x$ in \mathbb{R}^2 is

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

since reflecting $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ across $y = x$ gives $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ and reflecting $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ gives $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$. As a check:

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

is indeed the result of reflecting the vector $\begin{pmatrix} x \\ y \end{pmatrix}$ across $y = x$.

Example 4. Consider the linear transformation L from \mathbb{R}^2 to \mathbb{R}^2 which first applies the *shear* determined by $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and then scales the result by a factor of 2. (Check the book for a picture of what a “shear” transformation does.) Starting with $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, the shear transformation gives

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

and then scaling by 2 gives $\begin{pmatrix} 2 \\ 2 \end{pmatrix}$. Shearing $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ gives

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

and then scaling by 2 gives $\begin{pmatrix} 2 \\ 2 \end{pmatrix}$. Thus the matrix of this combined transformation, which first shears and then scales, is

$$\begin{pmatrix} 2 & 2 \\ 0 & 2 \end{pmatrix}.$$

Important. You should be familiar with rotations, reflections, shears, and scalings this quarter. Orthogonal projections have a somewhat more complicated formula, and is something we will come back to next quarter. UPDATE: When these notes were first written, it was true that orthogonal projections were not something we focused on in the fall quarter. However, subsequent years DID cover orthogonal projections in the fall, so you should be familiar with these as well. Ask your instructor if it’s unclear what exactly you should know.

Lecture 7: Matrix Multiplication

Today we spoke about what it means to multiply matrices, and how this relates to composing linear transformations.

Warm-Up. Find the matrix of the linear transformation which first rotates \mathbb{R}^2 by $\pi/4$ and then reflects across the line $y = -x$. As usual, we determine the result of applying this transformation to $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and then to $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$. First, rotating $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ gives

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} \rightarrow \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}.$$

Reflecting this across $y = -x$ flips its direction so we get

$$\begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} \rightarrow \begin{pmatrix} -1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix}.$$

Thus overall the transformation in question sends $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ to $\begin{pmatrix} -1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix}$. Second, rotating $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ gives

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} \rightarrow \begin{pmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}$$

and then reflecting this across $y = -x$ does nothing since this vector is on this line. Thus overall $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ is sent to $\begin{pmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}$. The matrix of this combined transformation is thus

$$\begin{pmatrix} -1/\sqrt{2} & -1/\sqrt{2} \\ -1/\sqrt{2} & 1/\sqrt{2} \end{pmatrix}.$$

Example 1. Consider linear transformations T and S from \mathbb{R}^2 to itself represented respectively by matrices

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \text{ and } \begin{pmatrix} m & n \\ p & q \end{pmatrix}.$$

We determine the matrix for the *composed* transformation TS which first applies S and then applies T : $(TS)(\vec{x}) = T(S\vec{x})$. We compute

$$\begin{aligned} (TS) \begin{pmatrix} 1 \\ 0 \end{pmatrix} &= T \left(S \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) \\ &= T \left(\begin{pmatrix} m & n \\ p & q \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) \\ &= T \begin{pmatrix} m \\ p \end{pmatrix} \\ &= \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} m \\ p \end{pmatrix} \\ &= \begin{pmatrix} am + bp \\ cm + dp \end{pmatrix}. \end{aligned}$$

Next we compute

$$\begin{aligned} (TS) \begin{pmatrix} 0 \\ 1 \end{pmatrix} &= T \left(S \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right) \\ &= T \left(\begin{pmatrix} m & n \\ p & q \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right) \\ &= T \begin{pmatrix} n \\ q \end{pmatrix} \\ &= \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} n \\ q \end{pmatrix} \\ &= \begin{pmatrix} an + bq \\ cn + dq \end{pmatrix}. \end{aligned}$$

Thus the matrix for the composition TS is

$$\begin{pmatrix} am + bp & an + bq \\ cm + dp & cn + dq \end{pmatrix}.$$

This is what we define to be the *product* of $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ and $\begin{pmatrix} m & n \\ p & q \end{pmatrix}$.

Matrix Multiplication. Given two matrices A and B where the number of columns of A is the same as the number of rows of B (the condition needed in order for the product AB to be defined), their product is the matrix AB defined by

$$AB = A(\vec{v}_1 \quad \vec{v}_2 \quad \cdots \quad \vec{v}_t) = (A\vec{v}_1 \quad A\vec{v}_2 \quad \cdots \quad A\vec{v}_t),$$

where $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_t$ denote the columns of B . The columns $A\vec{v}_i$ of the final matrix are just the usual expressions for matrix times vector we've seen before. The resulting matrix AB will have the same number of rows as A does and the same number of columns as B does.

Important. Matrix multiplication is defined in this way precisely so that matrix multiplication corresponds to composing linear transformations: if T is a linear transformation with matrix A and S a linear transformation with matrix B , the matrix of the composed transformation TS is AB . (The order of multiplication matters!)

Example 2. The result of multiplying

$$\begin{pmatrix} 1 & -1 & 0 \\ 3 & 2 & 1 \\ -2 & 0 & -1 \end{pmatrix} \text{ and } \begin{pmatrix} 0 & 3 & -2 \\ -1 & 5 & 1 \\ 2 & -1 & 1 \end{pmatrix},$$

is the 3×3 matrix given by

$$\begin{pmatrix} 1 & -1 & 0 \\ 3 & 2 & 1 \\ -2 & 0 & -1 \end{pmatrix} \begin{pmatrix} 0 & 3 & -2 \\ -1 & 5 & 1 \\ 2 & -1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & -2 & -3 \\ 0 & 18 & -3 \\ -2 & -5 & 4 \end{pmatrix}.$$

Again, the first column of the product is the result of multiplying the first matrix by $\begin{pmatrix} 0 \\ -1 \\ 2 \end{pmatrix}$, the second column is the first matrix times $\begin{pmatrix} 3 \\ 5 \\ -1 \end{pmatrix}$, and the last column is the first matrix times $\begin{pmatrix} -2 \\ 1 \\ 1 \end{pmatrix}$.

Example 3. Consider the matrices

$$A = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \text{ and } B = \begin{pmatrix} \cos \beta & -\sin \beta \\ \sin \beta & \cos \beta \end{pmatrix},$$

which geometrically are rotations by θ and β respectively. On the one hand, we can compute

$$AB = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} \cos \beta & -\sin \beta \\ \sin \beta & \cos \beta \end{pmatrix} = \begin{pmatrix} \cos \theta \cos \beta - \sin \theta \sin \beta & -\cos \theta \sin \beta - \sin \theta \cos \beta \\ \sin \theta \cos \beta + \cos \theta \sin \beta & -\sin \theta \sin \beta + \cos \theta \cos \beta \end{pmatrix}.$$

On the other hand, composing these two rotations is the same as rotating by $\theta + \beta$, so the matrix for this composition should also equal

$$\begin{pmatrix} \cos(\theta + \beta) & -\sin(\theta + \beta) \\ \sin(\theta + \beta) & \cos(\theta + \beta) \end{pmatrix}.$$

These two matrices we computed should be equal, so we must have

$$\begin{pmatrix} \cos(\theta + \beta) & -\sin(\theta + \beta) \\ \sin(\theta + \beta) & \cos(\theta + \beta) \end{pmatrix} = \begin{pmatrix} \cos \theta \cos \beta - \sin \theta \sin \beta & -\cos \theta \sin \beta - \sin \theta \cos \beta \\ \sin \theta \cos \beta + \cos \theta \sin \beta & -\sin \theta \sin \beta + \cos \theta \cos \beta \end{pmatrix}.$$

Comparing corresponding entries on both sides gives the well-known trig identities for $\cos(\theta + \beta)$ and $\sin(\theta + \beta)$. Imagine trying to justify these trig identities without using linear algebra!

Back to Warm-Up. Here is another way to approach the Warm-Up. The rotation part of the composed transformation has matrix

$$\begin{pmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{pmatrix}$$

and the reflection part has matrix

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

(This second one is obtained by reflecting $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ across $y = -x$ and then doing the same for $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$.) Thus, since matrix multiplication corresponds to composition of transformations, the combined transformation of the Warm-Up has matrix equal to the product

$$\begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{pmatrix} = \begin{pmatrix} -1/\sqrt{2} & -1/\sqrt{2} \\ -1/\sqrt{2} & 1/\sqrt{2} \end{pmatrix},$$

agreeing with the matrix we found in the Warm-Up. Again note the order of composition: our combined transformation first applied the rotation and then the reflection, so the matrix for rotation is on the right and reflection on the left. (We always read compositions from right to left.)

Example 3. Let $A = \begin{pmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{pmatrix}$. We want to compute A^{80} , which means A multiplied by itself 80 times. Of course, doing this multiplication by hand 80 times would be crazy. As well, if you start multiplying A by itself a few times you might notice some pattern which would help, but this is still not the most efficient way to approach this. Instead, recognize that A is the matrix for rotation by $\pi/4$, so A^{80} is the matrix for composing this rotation with itself 80 times. The point is that we know that rotating by $\pi/4$ eight times is the same as a rotation by 2π , which geometrically puts a vector back where it started. Thus A^8 should be the matrix for the transformation which leaves a vector untouched; this is called the “identity” transformation and its matrix is the *identity matrix*:

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

Thus $A^8 = I_2$ (the subscript denotes the fact that we are looking at the 2×2 identity matrix), and every further eighth power will again result in I_2 . Since 80 is a multiple of 8, we have $A^{80} = I_2$ without having to explicitly multiply A by itself 80 times.

Similarly, $A^{40} = I_2$ so $A^{43} = A^{40}A^3 = A^3$, which is the matrix for rotation by $3\pi/4$, which is the same as rotating by $\pi/4$ three times in a row. Thus

$$A^{43} = \begin{pmatrix} -1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{pmatrix}.$$

Properties and non-properties of matrix multiplication. Just like usual multiplication of numbers, matrix multiplication is associative:

$$(AB)C = A(BC)$$

for any matrices A, B, C for which all products above make sense; it is distributive:

$$A(B + C) = AB + AC \text{ and } (A + B)C = AC + BC,$$

and it has an *identity*, namely the identity matrix I :

$$AI = A = IA \text{ for any } A.$$

However, matrix multiplication is *unlike* multiplication of numbers in that it is not necessarily commutative:

$$AB \text{ does not necessarily equal } BA,$$

and it is possible to multiply nonzero matrices together to get the zero matrix:

$$AB = 0 \text{ does not necessarily mean } A = 0 \text{ or } B = 0.$$

Here, 0 denotes the zero matrix, which is the matrix with all entries equal to 0 .

Example 4. For matrices A, B , it is not necessarily true that

$$(A + B)(A - B) = A^2 - B^2$$

as is true for numbers. Indeed, the left side expands as

$$A^2 - AB + BA - B^2$$

and so equals $A^2 - B^2$ only when $AB = BA$. We say that A and B *commute* in this case.

Also, there are solutions to $A^2 = I$ apart from $A = \pm I$, as opposed to what happens for numbers. Indeed, this equation can be rewritten as $A^2 - I = 0$ which factors as

$$(A + I)(A - I) = 0.$$

However, this does not mean that one of the factors has to be zero. For example, taking $A = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ gives an example where A is neither $\pm I$ and yet $A^2 = I$.

Important. Be careful when manipulating equations involving matrix multiplication: sometimes things work like they do when multiplying numbers, but not always.

Lecture 8: Invertibility and Inverses

Today we spoke about the notion of a matrix being invertible, and finding the inverse of a matrix which is invertible. We will continue looking at properties of invertible matrices next time.

Warm-Up 1. Let A be the matrix

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

We compute A^{100} . The point is that we should not sit down and multiply A by itself 100 times, but rather we should think about what the linear transformation T corresponding to A is actually doing. We have

$$T \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} y \\ z \\ x \end{pmatrix},$$

so we see that T has the effect of “shifting” the entries of an input vector up by one while moving the first entry down to the end. Thus, T^2 has the effect of doing this twice in a row and T^3 the

effect of shifting three times in a row. However, with only three entries in an input vector, after doing the shift three times in a row we're back where we started so

$$T^3 \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} x \\ y \\ z \end{pmatrix}.$$

In other words, T^3 is the “identity” transformation which does nothing to input vectors. Thus A^3 , which is supposed to be the matrix for the composition T^3 , equals the matrix for the identity transformation, which is the identity matrix:

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

Then $A^6 = I$, $A^9 = I$, and so on every time we take a power that's a multiple of 3 we get the identity, so

$$A^{100} = A^{99}A = IA = A = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}.$$

Of course, you can multiply A by itself three times directly and see that $A^3 = I$, but's important to see why this is so from the point of view of composing linear transformations.

Warm-Up 2. Find all matrices B commuting with $\begin{pmatrix} 2 & -1 \\ 7 & 5 \end{pmatrix}$; that is, find all B such that

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

Note that in order for both of these products to be defined B must be 2×2 . Writing down an arbitrary expression for B :

$$B = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

the problem is to find all values of a, b, c, d such that

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} 2 & -1 \\ 7 & 5 \end{pmatrix} = \begin{pmatrix} 2 & -1 \\ 7 & 5 \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Multiplying out both sides gives the requirement

$$\begin{pmatrix} 2a + 7b & -a + 5b \\ 2c + 7d & -c + 5d \end{pmatrix} = \begin{pmatrix} 2a - c & 2b - d \\ 7a + 5c & 7b + 5d \end{pmatrix}.$$

Equating corresponding entries on both sides gives the requirements

$$\begin{array}{ll} 2a + 7b = 2a - c & -a + 5b = 2b - d \\ 2c + 7d = 7a + 5c & -c + 5d = 7b + 5d, \end{array}$$

so after moving everything to one side of each of these we see that the values of a, b, c, d we are looking for must satisfy the system

$$\begin{array}{rcl} 7b + c & = & 0 \\ -7a - 3c + 7d & = & 0 \\ -a + 3b + d & = & 0 \\ -7b - c & = & 0 \end{array}.$$

Hence our question boils down to solving this system of equations. (Imagine that!)

Row-reducing the augmented matrix gives

$$\left(\begin{array}{cccc|c} 0 & 7 & 1 & 0 & 0 \\ -7 & 0 & -3 & 7 & 0 \\ -1 & 3 & 0 & 1 & 0 \\ 0 & -7 & -1 & 0 & 0 \end{array} \right) \rightarrow \left(\begin{array}{cccc|c} 1 & 0 & 3/7 & -1 & 0 \\ 0 & 1 & 1/7 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right)$$

as the reduced echelon form. Thus we find the solution to be

$$\begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix} = \begin{pmatrix} -\frac{3}{7}c + d \\ -\frac{1}{7}c \\ c \\ d \end{pmatrix},$$

so our conclusion is that the matrices B which commute with $\begin{pmatrix} 2 & -1 \\ 7 & 5 \end{pmatrix}$ are those of the form

$$B = \begin{pmatrix} -\frac{3}{7}c + d & -\frac{1}{7}c \\ c & d \end{pmatrix} \text{ for any numbers } c \text{ and } d.$$

Note that taking $c = 7$ and $d = 5$ gives $\begin{pmatrix} 2 & -1 \\ 7 & 5 \end{pmatrix}$, which makes sense since any matrix indeed commutes with itself.

Inverse Transformations. Say that $A = \begin{pmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{pmatrix}$ is rotation by $\pi/4$. We call rotation by $-\pi/4$ the *inverse* transformation since it “undoes” what the first one does. In other words, applying the linear transformation determined by A and then following it with the inverse transformation always gives you back what you started with. The matrix for this inverse transformation (rotation by $-\pi/4$) is $\begin{pmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ -1/\sqrt{2} & 1/\sqrt{2} \end{pmatrix}$. Similarly, the inverse of the reflection transformation determined by $B = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ is that same reflection since to “undo” what a reflection does you simply apply that same reflection again. We say that a reflection is its own “inverse”.

In general, given a transformation T such that with input \vec{x} you get output \vec{y} :

$$T(\vec{x}) = \vec{y},$$

the inverse transformation (if it exists) is the linear transformation with the property that inputting \vec{y} gives as output \vec{x} .

Invertible Matrices. A (square) matrix A is *invertible* if there is another (square) matrix B (it will necessarily be of the same size as A) with the property that $AB = I$ and $BA = I$. We call this matrix B the *inverse* of A and denote it by A^{-1} . (It turns out that for square matrices the requirement that $AB = I$ automatically implies $BA = I$, but this is not at all obvious.) If A is the matrix for a linear transformations T , then A^{-1} is the matrix for the inverse transformation of T .

Back to previous geometric examples. The geometric examples we just looked at say that

$$\begin{pmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{pmatrix}^{-1} = \begin{pmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ -1/\sqrt{2} & 1/\sqrt{2} \end{pmatrix}$$

and

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

You can verify this by checking that

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

and

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

Inverses of 2×2 matrices. In general, a 2×2 matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ is invertible only when $ad - bc \neq 0$, in which case its inverse is given by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}^{-1} = \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix}.$$

KNOW THIS FORMULA BY HEART. The denominator of the fraction involved is called the *determinant* of $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$; we will come back to determinants later. You can verify that this formula for the inverse of a 2×2 matrix is correct by checking that

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \text{ times } \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix} \text{ equals } \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Also, note that applying this formula to the geometric 2×2 examples above indeed gives what we claimed were the inverses.

Inverses in general. As in the 2×2 case, there is an explicit formula for the inverse of any invertible $n \times n$ matrix. However, this formula gets a lot more complicated even in the 3×3 case and is NOT worth memorizing. Instead, we compute inverses in general using the following method, which will most often be much, much quicker.

To find the inverse of an invertible matrix A , set up a big augmented matrix

$$(A \mid I)$$

with A on the left and the appropriately-sized identity on the right, then start doing row operations to reduce that “ A ” part to the identity while at the same time doing the same operations to the identity part:

$$(A \mid I) \rightarrow (I \mid A^{-1}).$$

The matrix you end up with on the right side is the inverse of A . Check the book for an explanation of why this works.

Note that this process only works if it is indeed possible to reduce A to the identity matrix, giving us our first way to check whether a given matrix is invertible.

Important. A is invertible if and only if the reduced echelon form of A is the identity matrix, which can happen if and only if A has “full” rank, meaning rank equal to the numbers of rows and columns.

Example 1. Let A be the matrix

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

Reducing $(A \mid I)$ gives:

$$\left(\begin{array}{ccc|ccc} 1 & -1 & 0 & 1 & 0 & 0 \\ 3 & 2 & 1 & 0 & 1 & 0 \\ -2 & 0 & -1 & 0 & 0 & 1 \end{array} \right) \rightarrow \left(\begin{array}{ccc|ccc} 1 & -1 & 0 & 1 & 0 & 0 \\ 0 & 5 & 1 & -3 & 1 & 0 \\ 0 & -2 & -1 & 2 & 0 & 1 \end{array} \right) \rightarrow \left(\begin{array}{ccc|ccc} 1 & -1 & 0 & 1 & 0 & 0 \\ 0 & 5 & 1 & -3 & 1 & 0 \\ 0 & 0 & -3 & 4 & 2 & 5 \end{array} \right).$$

Note that at this point we know that A is invertible since we can already see that the reduced echelon form of A will end up with three pivots and will be the 3×3 identity matrix. Continuing on yields

$$\left(\begin{array}{ccc|ccc} 1 & -1 & 0 & 1 & 0 & 0 \\ 0 & 5 & 1 & -3 & 1 & 0 \\ 0 & 0 & -3 & 4 & 2 & 5 \end{array} \right) \rightarrow \left(\begin{array}{ccc|ccc} 5 & 0 & 1 & 2 & 1 & 0 \\ 0 & 5 & 1 & -3 & 1 & 0 \\ 0 & 0 & -3 & 4 & 2 & 5 \end{array} \right) \rightarrow \left(\begin{array}{ccc|ccc} 15 & 0 & 0 & 10 & 5 & 5 \\ 0 & 15 & 0 & -5 & 5 & 5 \\ 0 & 0 & -3 & 4 & 2 & 5 \end{array} \right),$$

Dividing by the appropriate scalars turns the right side into

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

which is the inverse of A . You can check that

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

as required of the inverse of A .

Example 2. For the matrix

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

we have

$$\left(\begin{array}{ccc|ccc} 1 & 2 & 3 & 1 & 0 & 0 \\ 4 & 5 & 6 & 0 & 1 & 0 \\ 7 & 8 & 9 & 0 & 0 & 1 \end{array} \right) \rightarrow \left(\begin{array}{ccc|ccc} 1 & 2 & 3 & 1 & 0 & 0 \\ 0 & -3 & -6 & -4 & 1 & 0 \\ 0 & -6 & -12 & -7 & 0 & 1 \end{array} \right) \rightarrow \left(\begin{array}{ccc|ccc} 1 & 2 & 3 & 1 & 0 & 0 \\ 0 & -3 & -6 & -4 & 1 & 0 \\ 0 & 0 & 0 & 1 & -2 & 1 \end{array} \right).$$

Since we can now see that the left part will not give us a reduced form with three pivots, we can stop here: B is not invertible.

Lecture 9: More on Inverses, the Amazingly Awesome Theorem

Today we continued talking about invertible matrices, focusing on the various characterizations of what it means for a matrix to be invertible. These various characterizations make up what I call the “Amazingly Awesome Theorem” in order to emphasize how important they are.

Warm-Up 1. Suppose A is a square matrix such that $A^5 = I$. We claim that A is invertible. Recall that to be invertible means that there is a matrix B with the property that $AB = I = BA$. If $A^5 = I$, we have

$$AA^4 = A^5 = I \text{ and } A^4A = A^5 = I,$$

so $B = A^4$ satisfies the requirement in the definition of what it means for A to be invertible. So A is invertible with inverse A^4 .

Another approach might be to say that since $A^5 = I$, $\text{rref}(A^5) = I$ and so $\text{rref}(A) = I$ as well. However, it is not immediately obvious that just because the reduced echelon form of A^5 is the identity means that same will be true of A . Indeed, note that in general

$$\text{rref}(AB) \neq \text{rref}(A) \text{rref}(B),$$

so that whether we perform row operations before or after multiplying matrices matters. If A and B are both invertible, then AB is also invertible (with inverse $B^{-1}A^{-1}$) so in this case $\text{rref}(A)$, $\text{rref}(B)$, and $\text{rref}(AB)$ are all the identity and here we do have

$$\text{rref}(AB) = \text{rref}(A) \text{rref}(B).$$

However, for $A = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ and $B = \begin{pmatrix} 1 & 1 \\ -1 & -1 \end{pmatrix}$ we have

$$\text{rref}(A) = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}, \text{rref}(B) = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}, \text{ and } \text{rref}(AB) = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix},$$

so

$$\text{rref}(AB) \neq \text{rref}(A) \text{rref}(B)$$

in this case. That is, $\text{rref}(AB) = \text{rref}(A) \text{rref}(B)$ is only sometimes true.

Warm-Up 2. Suppose that A and B are square matrices such that $AB = I$. We claim that both A and B are then invertible. Comparing with the definition of what it means for either A or B to be invertible, the point is that here we only know that multiplying A and B in one order gives the identity, whereas the definition would require that $BA = I$ as well. It is not at all obvious that just because $AB = I$ it must also be true that $BA = I$, and indeed this is only true for *square* matrices.

A first thought might be to multiply both sides of $AB = I$ on the left by A^{-1} to give

$$B = A^{-1}.$$

Then multiplying by A on the right gives $BA = I$, which is what we want. However, this assumes that A^{-1} already exists! We can't start multiplying by A^{-1} before we know A is actually invertible, so this is no good. Instead, we use the fact that a matrix B is invertible if and only if the only solution to $B\vec{x} = \vec{0}$ is $\vec{x} = \vec{0}$. Indeed, if this is true then the reduced echelon form of B has to be the identity, so B will be invertible.

So, we want to show that only solution of $B\vec{x} = \vec{0}$ is the zero vector. Multiplying on the left by A here gives

$$A(B\vec{x}) = A\vec{0}, \text{ so } (AB)\vec{x} = \vec{0}.$$

But since $AB = I$ this gives $\vec{x} = \vec{0}$ as we wanted. This means that B is invertible, and hence B^{-1} exists. Now we can multiply both sides of $AB = I$ on the right by B^{-1} to give

$$A = B^{-1}, \text{ so } BA = I$$

after multiplying on the left by B . Thus $AB = I$ and $BA = I$ so A is invertible as well with inverse B . Again, note that we said nothing about B^{-1} above until we already knew that B was invertible, and we were able to show that B is invertible using another way of thinking about invertibility. (More on this to come.)

Remark. As said above, for non-square matrices A and B it is not necessarily true that $AB = I$ automatically implies $BA = I$. (In fact, it never does but to understand this we need to know a bit more about what the rank of a matrix really means.) For example, for

$$A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \text{ and } B = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}$$

we have

$$AB = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \text{ but } BA = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

so $AB = I$ does not always imply that $BA = I$.

The Amazingly Awesome Theorem. (The name was chosen to emphasize how important and useful this can be.) The following are equivalent to a square matrix A being invertible:

- The row-reduced echelon form of A is the identity matrix.
- A has full rank, meaning rank equal to the number of rows or columns in A .
- For any vector \vec{b} , $A\vec{x} = \vec{b}$ has a solution, which will in fact be unique.
- The only solution of $A\vec{x} = \vec{0}$ is $\vec{x} = \vec{0}$.
- The only way in which to express $\vec{0}$ as a linear combination of the columns $\vec{v}_1, \dots, \vec{v}_n$ of A is to take all coefficients to be 0; i.e. the only solution of

$$c_1\vec{v}_1 + \dots + c_n\vec{v}_n = \vec{0}$$

is $c_1 = \dots = c_n = 0$. (Later we will see that this is what it means to say that the columns of A are “linearly independent”.)

- Any vector \vec{b} can be written as a linear combination of the columns $\vec{v}_1, \dots, \vec{v}_n$ of A ; i.e. for any \vec{b} the equation

$$c_1\vec{v}_1 + \dots + c_n\vec{v}_n = \vec{b}$$

has a solution for c_1, \dots, c_n . (Later we will see that this is what it means to say that the columns of A “span” all of \mathbb{R}^n .)

Elementary Matrices (Optional). This topic is purely extra and will never be on an exam, quiz, nor homework. It is only meant to introduce an idea which can be useful in certain contexts.

The idea behind elementary matrices is that any row operation can in fact be expressed as a matrix multiplication. To get a feel for this let’s just focus on a 2×2 matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$. First, note that

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} c & d \\ a & b \end{pmatrix},$$

so the operation of swapping the rows of a matrix can be obtained as a result of multiplying by $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Second, scaling a row (say the first) by k is obtained via

$$\begin{pmatrix} k & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} ka & kb \\ c & d \end{pmatrix}.$$

Finally, adding a multiple of one row to another is obtained as

$$\begin{pmatrix} 1 & 0 \\ k & 1 \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} a & b \\ ka + c & kb + d \end{pmatrix}.$$

The matrices

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} k & 0 \\ 1 & 0 \end{pmatrix}, \text{ and } \begin{pmatrix} 1 & 0 \\ k & 1 \end{pmatrix}$$

are examples of what are called *elementary matrices* because they induce elementary row operations. There are analogues for matrices of any size.

As an application, we can fully justify why the process we use for computing inverses actually works. Starting with an invertible A , we know there are row operations which will reduce A to the identity I . Each of these row operations can be obtained via multiplication by an elementary matrix, so there are elementary matrices E_1, E_2, \dots, E_m such that

$$E_m \cdots E_2 E_1 A = I.$$

But then this equation says that $E_m \cdots E_2 E_1$ satisfies the requirement of being the inverse of A , so $A^{-1} = E_m \cdots E_2 E_1$. We can write this as

$$A^{-1} = E_m \cdots E_2 E_1 I,$$

where the right side now means we take the operations which reduced A to I and instead perform them on I ; the result is A^{-1} , which is precisely what our method for the finding inverses says.

Lecture 10: Images and Kernels

Today we spoke about the notions of the “kernel” of a matrix and the “image” of a matrix. This is the start of a whole new perspective on what we’ve been doing so far.

Warm-Up 1. For $A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$, we want to find a vector \vec{b} in \mathbb{R}^3 such that $A\vec{x} = \vec{b}$ has no solution. Rather than just taking some random \vec{b} and seeing whether $A\vec{x} = \vec{b}$ has a solution, here is a more systematic way of finding such a vector. Row-reducing A a few steps gives:

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 2 & 3 \\ 0 & -3 & -6 \\ 0 & -6 & -12 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 2 & 3 \\ 0 & -3 & -6 \\ 0 & 0 & 0 \end{pmatrix}. \quad (3)$$

Given this final form, we know that if we had a vector \vec{b} such that the final “augmented” piece of the corresponding reduced augmented matrix ended up being something like $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$:

$$\begin{pmatrix} 1 & 2 & 3 & | & \\ 4 & 5 & 6 & | & \vec{b} \\ 7 & 8 & 9 & | & \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 2 & 3 & | & 1 \\ 0 & -3 & -6 & | & 1 \\ 0 & 0 & 0 & | & 1 \end{pmatrix},$$

then $A\vec{x} = \vec{b}$ would indeed have no solution. The point is that starting from this reduced form we can work our way *backwards* to find such a \vec{b} .

Going back to the row operations we did in (3), to “undo” the last one we multiply the second row by 2 and add it to the last row:

$$\left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 0 & -3 & -6 & 1 \\ 0 & 0 & 0 & 1 \end{array}\right) \rightarrow \left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 0 & -3 & -6 & 1 \\ 0 & -6 & -12 & 3 \end{array}\right).$$

To undo the first row operations we did in (3), we now do $4I + II \rightarrow II$ and $7I + III \rightarrow III$:

$$\left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 0 & -3 & -6 & 1 \\ 0 & -6 & -12 & 3 \end{array}\right) \rightarrow \left(\begin{array}{ccc|c} 1 & 2 & 3 & 1 \\ 4 & 5 & 6 & 5 \\ 7 & 8 & 9 & 10 \end{array}\right).$$

Thus, $\vec{b} = \begin{pmatrix} 1 \\ 5 \\ 10 \end{pmatrix}$ is an example of a vector \vec{b} such that $A\vec{x} = \vec{b}$ has no solution. We’ll reinterpret this fact in terms of the image of A in a bit.

Warm-Up 2. The *transpose* of an $n \times n$ matrix B is the $n \times n$ matrix B^T obtained by turning the columns of B into the rows of B^T . For instance,

$$\begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix}^T = \begin{pmatrix} a & d & g \\ b & e & h \\ c & f & i \end{pmatrix}.$$

We claim that if B is invertible then so is B^T , and the only fact we need is the following identity:

$$(AB)^T = B^T A^T \text{ for any matrices } A \text{ and } B.$$

(We’ll come back to this identity and work with transposes more later on.) To show that B^T is invertible we use one of the equivalent characterizations of invertibility from the Amazingly Awesome Theorem and verify that the only vector satisfying $B^T \vec{x} = \vec{0}$ is $\vec{x} = \vec{0}$.

So, start with $B^T \vec{x} = \vec{0}$. Since B is invertible B^{-1} exists and it has a transpose $(B^{-1})^T$. Multiplying both sides of our equation by this gives

$$(B^{-1})^T B^T \vec{x} = (B^{-1})^T \vec{0} = \vec{0}.$$

Using the identity for transposes stated above, this left-hand side equals $(BB^{-1})^T$, which equals the transpose of the identity matrix, which is the identity matrix itself. So the above equation becomes

$$(BB^{-1})^T \vec{x} = \vec{0} \implies I\vec{x} = \vec{0},$$

so the only vector satisfying $B^T \vec{x} = \vec{0}$ is $\vec{x} = \vec{0}$. This means that B^T is invertible.

The kernel of a matrix. The *kernel* of a matrix A , denoted by $\ker A$, is the space of all solutions to $A\vec{x} = \vec{0}$. In terms of the linear transformation T determined by A , this is the space of all input vectors which are sent to $\vec{0}$ under T .

Example 1. Looking at the matrix A from Warm-Up 1, to find its kernel we must find all solution of $A\vec{x} = \vec{0}$. Continuing the row operations started before, we have:

$$(A \mid \vec{0}) \rightarrow \left(\begin{array}{ccc|c} 1 & 2 & 3 & 0 \\ 0 & -3 & -6 & 0 \\ 0 & 0 & 0 & 0 \end{array}\right) \rightarrow \left(\begin{array}{ccc|c} 1 & 0 & -1 & 0 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 0 \end{array}\right).$$

Denoting $\vec{x} = \begin{pmatrix} x \\ y \\ z \end{pmatrix}$, we have that z is free and $x = z, y = -2z$. Thus solutions of $A\vec{x} = \vec{0}$, i.e. vectors \vec{x} in the kernel of A , look like:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} t \\ -2t \\ t \end{pmatrix} = t \begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix}.$$

Thus, any vector in $\ker A$ is a multiple of $\begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix}$. The collection of all such multiples is what we call the *span* of $\begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix}$, so we can say that

$$\ker A = \text{span} \left\{ \begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix} \right\}.$$

In words we would also say that $\begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix}$ spans the kernel of A . Geometrically, this kernel is the line containing $\begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix}$, which is what the set of all multiples of this vector looks like.

Definition of Span. The *span* of vectors $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_k$ is the collection of all linear combinations of $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_k$; i.e. all vectors expressible in the form

$$c_1\vec{v}_1 + c_2\vec{v}_2 + \dots + c_k\vec{v}_k$$

for scalars c_1, \dots, c_k .

Example 2. We find vectors which span the kernel of

$$B = \begin{pmatrix} 1 & 2 & -3 \\ 2 & 4 & -6 \\ 3 & 6 & -9 \end{pmatrix}.$$

First, row reducing B gives

$$\begin{pmatrix} 1 & 2 & -3 \\ 2 & 4 & -6 \\ 3 & 6 & -9 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 2 & -3 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

When solving $B\vec{x} = \vec{0}$ we would add on an extra “augmented” column of zeros, but from now on when finding the kernel of a matrix we will skip this additional step. From the reduced echelon form above, we can see that the solutions of $B\vec{x} = \vec{0}$ all look like

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} -2s + 3t \\ s \\ t \end{pmatrix}.$$

To find vectors which span the collection of all vectors which look like this, we “factor out” each free variable:

$$\begin{pmatrix} -2s + 3t \\ s \\ t \end{pmatrix} = s \begin{pmatrix} -2 \\ 1 \\ 0 \end{pmatrix} + t \begin{pmatrix} 3 \\ 0 \\ 1 \end{pmatrix}.$$

Thus, any solution of $B\vec{x} = \vec{0}$ is expressible as a linear combination of $\begin{pmatrix} -2 \\ 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 3 \\ 0 \\ 1 \end{pmatrix}$, so these two vectors span the kernel of B :

$$\ker B = \text{span} \left\{ \begin{pmatrix} -2 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 3 \\ 0 \\ 1 \end{pmatrix} \right\}.$$

Important. To find vectors spanning the kernel of any matrix A , find all solutions of $A\vec{x} = \vec{0}$ and express all variables in terms of free variables. Then “factor out” each free variable to express the solution as a linear combination of some vectors; these vectors span the kernel.

Example 3. We have

$$\begin{pmatrix} 1 & 2 & 0 & -2 & 3 & 1 \\ -2 & -4 & 4 & 3 & -5 & -1 \\ 3 & 6 & 4 & -7 & 11 & 5 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 2 & 0 & -2 & 0 & -2 \\ 0 & 0 & 1 & -1/4 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}.$$

Thus the kernel of the first matrix consists of things which look like

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{pmatrix} = \begin{pmatrix} -2s + 2t + 2u \\ s \\ \frac{1}{4}t \\ t \\ -u \\ u \end{pmatrix} = s \begin{pmatrix} -2 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} 2 \\ 0 \\ 1/4 \\ 1 \\ 0 \\ 0 \end{pmatrix} + u \begin{pmatrix} 2 \\ 0 \\ 0 \\ 0 \\ -1 \\ 1 \end{pmatrix},$$

and so

$$\ker \begin{pmatrix} 1 & 2 & 0 & -2 & 3 & 1 \\ -2 & -4 & 4 & 3 & -5 & -1 \\ 3 & 6 & 4 & -7 & 11 & 5 \end{pmatrix} = \text{span} \left\{ \begin{pmatrix} -2 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ 0 \\ 1/4 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ 0 \\ 0 \\ 0 \\ -1 \\ 1 \end{pmatrix} \right\}.$$

The image of a matrix. The *image* of a matrix A is the collection $\text{im } A$ of all possible outputs of the linear transformation determined by A . More concretely, any such output has the form $A\vec{x}$, so the image of A consists of all vectors \vec{b} such that $A\vec{x} = \vec{b}$ has a solution. Better yet, the product $A\vec{x}$ is expressible as a linear combination of the columns of A , so \vec{b} is in the image of A if \vec{b} is a linear combination of the columns of A . Thus we can say that

$$\text{im } A = \text{span} \{ \text{columns of } A \}.$$

Back to Example 1. The first Warm-Up shows that $\vec{b} = \begin{pmatrix} 1 \\ 5 \\ 10 \end{pmatrix}$ is not in the image of $A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$ since $A\vec{x} = \vec{b}$ has no solution.

Back to Example 2. The image of B from Example 2 is the span of its columns, so

$$\text{im } B = \text{span} \left\{ \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ 4 \\ 6 \end{pmatrix}, \begin{pmatrix} -3 \\ -6 \\ -9 \end{pmatrix} \right\}.$$

However, we can simplify this description. Note that second and third vectors in this spanning set are multiples of the first. This means that any vector which is expressible as a linear combination of all three is in fact expressible as a linear combination (i.e. multiple) of the first alone. Thus, the above span is the same as the span of the first column alone, so

$$\text{im} \begin{pmatrix} 1 & 2 & -3 \\ 2 & 4 & -6 \\ 3 & 6 & -9 \end{pmatrix} = \text{span} \left\{ \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} \right\}.$$

Geometrically, the span of one vector is the line containing that vector, so in this case the image of B is a line.

Important. If \vec{b} is a linear combination of vectors $\vec{v}_1, \dots, \vec{v}_k$, then

$$\text{span} \{ \vec{v}_1, \dots, \vec{v}_k, \vec{b} \} = \text{span} \{ \vec{v}_1, \dots, \vec{v}_k \}.$$

In other words, if one vector is itself in the span of other vectors, throwing that vector away from our spanning set does not change the overall span.

Final Example. We want to find matrices A and B such that the plane $x - 2y + 3z = 0$ is at the same time the kernel of A and the image of B . First, the equation defining this plane is precisely what it means to say that

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \text{ is in the kernel of } \begin{pmatrix} 1 & -2 & 3 \end{pmatrix},$$

so for the 1×3 matrix $A = \begin{pmatrix} 1 & -2 & 3 \end{pmatrix}$, the plane $x - 2y + 3z = 0$ is the kernel of A . There are tons of other matrices which work; for instance, the kernel of

$$\begin{pmatrix} 1 & -2 & 3 \\ 1 & -2 & 3 \\ 1 & -2 & 3 \end{pmatrix}$$

is also the plane $x - 2y + 3z = 0$.

Now, note that we can find vectors which span the plane as follows. From the equation of the plane we find that $x = 2y - 3z$, so vectors on the plane look like

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 2y - 3z \\ y \\ z \end{pmatrix} = y \begin{pmatrix} 2 \\ 1 \\ 0 \end{pmatrix} + z \begin{pmatrix} -3 \\ 0 \\ 1 \end{pmatrix}.$$

Hence the plane is equal to the span of

$$\begin{pmatrix} 2 \\ 1 \\ 0 \end{pmatrix} \text{ and } \begin{pmatrix} -3 \\ 0 \\ 1 \end{pmatrix},$$

which is equal to the image of the matrix

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

That is, $\text{im } B$ is the plane $x - 2y + 3z = 0$.

Lecture 11: Subspaces of \mathbb{R}^n

Today we started talking about the notion of a *subspace* of \mathbb{R}^n . The definition is kind of abstract, but in the end subspaces are pretty simple to visualize geometrically.

Warm-Up 1. Suppose that A is an invertible $n \times n$ matrix. We determine $\ker A$ and $\text{im } A$. First, $\ker A$ consists of the vectors \vec{x} satisfying $A\vec{x} = \vec{0}$. However, since A is invertible only $\vec{x} = \vec{0}$ satisfies this, so

$$\ker A = \{\vec{0}\}.$$

That is, $\ker A$ is the set only containing the zero vector. Note that conversely, if A is a square matrix such that $\ker A = \{0\}$ then in fact A must be invertible! So, we can add this condition on to our Amazingly Awesome Theorem characterizing the various things equivalent to a matrix being invertible.

The image of A consists of all \vec{b} in \mathbb{R}^n such that $A\vec{x} = \vec{b}$ has a solution. But since A is invertible, this equation *always* has the solution $\vec{x} = A^{-1}\vec{b}$, so any vector in \mathbb{R}^n is in the image of A . Thus

$$\text{im } A = \mathbb{R}^n.$$

Conversely, an $n \times n$ matrix whose image is all of \mathbb{R}^n must be invertible, again adding on to our Amazingly Awesome Theorem.

Amazingly Awesome Theorem, continued. The following are also equivalent to a square $n \times n$ matrix A being invertible:

- $\ker A = \{\vec{0}\}$, i.e. the kernel of A consists of only the zero vector
- $\text{im } A = \mathbb{R}^n$, i.e. the image of A consists of every vector in \mathbb{R}^n

Warm-Up 2. Let A be the matrix

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

We want to find vectors spanning $\ker A$ and $\text{im } A$, and in each case we want to use as few vectors as possible. First we row reduce:

$$\begin{pmatrix} 1 & -1 & 2 & 1 \\ 1 & -1 & 3 & 4 \\ 2 & -2 & 2 & -4 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & -1 & 0 & -5 \\ 0 & 0 & 1 & 3 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

Thus solution of $A\vec{x} = \vec{0}$ are of the form

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = \begin{pmatrix} s + 5t \\ s \\ -3t \\ t \end{pmatrix} = s \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} 5 \\ 0 \\ -3 \\ 1 \end{pmatrix},$$

so

$$\ker A = \text{span} \left\{ \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 5 \\ 0 \\ -3 \\ 1 \end{pmatrix} \right\}.$$

Since neither of these vectors is a linear combination of the other (recall that a linear combination of a single vector is simply a multiple of that vector), throwing one vector away won't give us the same span, so we need both of these in order to span the entire kernel.

Now, $\text{im } A$ is spanned by the columns of A . However, note that the second column is a multiple of the first so that throwing it away gives the same span as all four columns:

$$\text{im } A = \text{span} \left\{ \begin{pmatrix} 1 \\ 1 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ 3 \\ 2 \end{pmatrix}, \begin{pmatrix} 1 \\ 4 \\ -4 \end{pmatrix} \right\}.$$

But we're not done yet! The third vector here is actually a linear combination of the first two:

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

so that the first two vectors have the same span as all three. (Again, to be clear, we are using the fact that if $\vec{v}_3 = a\vec{v}_1 + b\vec{v}_2$, then

$$c_1\vec{v}_1 + c_2\vec{v}_2 + c_3\vec{v}_3 = c_1\vec{v}_1 + c_2\vec{v}_2 + c_3(a\vec{v}_1 + b\vec{v}_2) = (c_1 + c_3a)\vec{v}_1 + (c_2 + c_3b)\vec{v}_2,$$

so a linear combination of $\vec{v}_1, \vec{v}_2, \vec{v}_3$ can be rewritten as a linear combination of \vec{v}_1 and \vec{v}_2 alone.) The first two vectors in the above span are not multiples of each other, so finally we conclude that

$$\text{im } A = \text{span} \left\{ \begin{pmatrix} 1 \\ 1 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ 3 \\ 2 \end{pmatrix} \right\}$$

and that we need both of these vectors to span the entire image.

Important properties of kernels and images. For any matrix A , its kernel and image both have the following properties:

- Both contain the zero vector: $\vec{0}$ is in $\ker A$ since $A\vec{0} = \vec{0}$ and $\vec{0}$ is in $\text{im } A$ since $A\vec{x} = \vec{0}$ has a solution,
- Adding two things in either one gives something still in that same space: if \vec{x} and \vec{y} are in $\ker A$ then $A(\vec{x} + \vec{y}) = A\vec{x} + A\vec{y} = \vec{0} + \vec{0} = \vec{0}$ so $\vec{x} + \vec{y}$ is also in $\ker A$, and if \vec{b}_1 and \vec{b}_2 are in $\text{im } A$, meaning that $A\vec{x} = \vec{b}_1$ has a solution \vec{x}_1 and $A\vec{x} = \vec{b}_2$ has a solution \vec{x}_2 , then $A\vec{x} = \vec{b}_1 + \vec{b}_2$ has solution $\vec{x}_1 + \vec{x}_2$, so $\vec{b}_1 + \vec{b}_2$ is still in $\text{im } A$, and
- Scaling something in either one gives something still in that same space: if \vec{x} is in $\ker A$, then $A(c\vec{x}) = cA\vec{x} = c\vec{0} = \vec{0}$ so $c\vec{x}$ is in $\ker A$, and if $A\vec{x} = \vec{b}$ has solution \vec{x} , then $A(c\vec{x}) = cA\vec{x} = c\vec{b}$ so $c\vec{b}$ is still in $\text{im } A$.

Definition of a subspace. A collection V of vectors in \mathbb{R}^n is a *subspace* of \mathbb{R}^n if it has all of the following properties:

- The zero vector $\vec{0}$ is in V ,
- V is *closed under addition* in the sense that if \vec{u} and \vec{v} are in V , then $\vec{u} + \vec{v}$ is also in V ,

- V is *closed under scalar multiplication* in the sense that if \vec{v} is in V and c is any scalar, then $c\vec{v}$ is also in V .

Back to kernels and images. So, for an $n \times m$ matrix, $\ker A$ is a subspace of \mathbb{R}^m and $\text{im } A$ is a subspace of \mathbb{R}^n .

Example 1. Consider a line in \mathbb{R}^2 which passes through the origin, such as $y = x$. This line consists of all vectors $\begin{pmatrix} x \\ y \end{pmatrix}$ whose coordinates satisfy $y = x$. We claim that this is a subspace of \mathbb{R}^2 . Indeed, $\begin{pmatrix} 0 \\ 0 \end{pmatrix}$ satisfies $y = x$, so this line contains the zero vector. Given two vectors on this line, say $\begin{pmatrix} a \\ a \end{pmatrix}$ and $\begin{pmatrix} b \\ b \end{pmatrix}$, their sum $\begin{pmatrix} a+b \\ a+b \end{pmatrix}$ is also on this line since its x and y coordinates satisfy $y = x$. Finally, given a vector $\begin{pmatrix} a \\ a \end{pmatrix}$ on this line, any multiple $\begin{pmatrix} ca \\ ca \end{pmatrix}$ of it is still on this line. Thus the line $y = x$ is a subspace of \mathbb{R}^2 .

Remark. For similar reasons, any line which passes through the origin is a subspace of \mathbb{R}^2 .

Example 2. Consider a line in \mathbb{R}^2 which does not pass through the origin, such as $y = x + 1$. To be clear, this line consists of all vectors $\begin{pmatrix} x \\ y \end{pmatrix}$ satisfying $y = x + 1$. This is not a subspace of \mathbb{R}^2 . First, it does not contain the zero vector $\begin{pmatrix} 0 \\ 0 \end{pmatrix}$ since these coordinates do not satisfy $y = x + 1$. Second, it is not closed under addition since $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 1 \\ 2 \end{pmatrix}$ are both on this line but their sum $\begin{pmatrix} 1 \\ 3 \end{pmatrix}$ is not. Third, it is not closed under scalar multiplication since $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ is on this line but $2\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 2 \end{pmatrix}$ is not. Failing at least one of these conditions is enough to say that this line is not a subspace of \mathbb{R}^2 , this example just happens to fail all three.

Remark. For similar reasons, any line which does not pass through the origin won't be a subspace of \mathbb{R}^2 either. In particular, any such line does not contain the zero vector.

Example 3. Consider the parabola $y = x^2$. This does contain the zero vector since $\begin{pmatrix} 0 \\ 0 \end{pmatrix}$ satisfies the equation of the parabola. However, this is not closed under addition since $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 2 \\ 4 \end{pmatrix}$ are both on this parabola but their sum $\begin{pmatrix} 3 \\ 5 \end{pmatrix}$ is not. Thus this parabola is not a subspace of \mathbb{R}^2 . It is also not closed under scalar multiplication.

Similarly, any other curve which is not a line through the origin won't be a subspace of \mathbb{R}^2 either.

Subspaces of \mathbb{R}^2 in general. The set consisting of only the zero vector $\{\vec{0}\}$ is a subspace of \mathbb{R}^2 , as is the entire xy -plane, i.e. \mathbb{R}^2 , itself. Apart from these, the only other subspaces of \mathbb{R}^2 as found earlier are lines through the origin.

Subspaces of \mathbb{R}^3 . The only subspaces of \mathbb{R}^3 are the origin $\{\vec{0}\}$, lines through the origin, planes through the origin, and all of \mathbb{R}^3 .

Important. Geometrically, subspaces of \mathbb{R}^n are either the single point consisting of the origin alone, lines through the origin, planes through the origin, higher-dimensional analogues of lines and planes through the origin, and all of \mathbb{R}^n itself. In particular, the kernel or image of any matrix looks like one of these.

Lecture 12: Linear Dependence/Independence and Bases

Today we spoke about the notion of vectors being linearly “dependent” or “independent”, and the idea of a “basis” of a subspace of \mathbb{R}^n . Bases will give us a simple and efficient way to describe any subspace.

Warm-Up 1. Consider the following set of vectors in \mathbb{R}^2 :

$$V = \left\{ \begin{pmatrix} x \\ y \end{pmatrix} : x \geq 0 \text{ and } y \geq 0 \right\}.$$

To be clear, this notation means we are looking at the vectors $\begin{pmatrix} x \\ y \end{pmatrix}$ in \mathbb{R}^2 satisfying the condition that $x \geq 0$ and $y \geq 0$. (Read the colon as “such that”.) In other words, V is the first quadrant of the xy -plane including the nonnegative x and y -axes. We check whether V is a subspace of \mathbb{R}^2 . Based on what we said last time, this answer should be no since V is not $\{\vec{0}\}$, nor a line through the origin, nor all of \mathbb{R}^2 .

First, V does contain the zero vectors since $\begin{pmatrix} 0 \\ 0 \end{pmatrix}$ satisfies the requirement that its x and y coordinates are ≥ 0 . V is also closed under addition: if $\begin{pmatrix} x \\ y \end{pmatrix}$ and $\begin{pmatrix} a \\ b \end{pmatrix}$ are in V (meaning both coordinates of each are ≥ 0), then so is $\begin{pmatrix} x+a \\ y+b \end{pmatrix}$ since

$$x + a \geq 0 \text{ and } y + b \geq 0.$$

However, V is not closed under scalar multiplication: $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is in V but $-2\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is not. Thus V is not a subspace of \mathbb{R}^2 .

Warm-Up 2. We now ask whether the following set of vectors in \mathbb{R}^3 is a subspace of \mathbb{R}^3 :

$$W = \left\{ \begin{pmatrix} x \\ y \\ z \end{pmatrix} : x - y + z = 0 \text{ and } 2x - y - 2z = 0 \right\}.$$

We can check the subspace conditions one at a time again. For instance, if $\begin{pmatrix} x \\ y \\ z \end{pmatrix}$ and $\begin{pmatrix} a \\ b \\ c \end{pmatrix}$ are in W , then each satisfies the equations defining W , so $\begin{pmatrix} x+a \\ y+b \\ z+c \end{pmatrix}$ satisfies

$$(x + a) - (y + b) + (z + c) = (x - y + z) + (a - b + c) = 0 + 0 = 0$$

and

$$2(x + a) - (y + b) - 2(z + c) = (2x - y - 2z) + (2a - b - 2c) = 0 + 0 = 0.$$

Thus $\begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} a \\ b \\ c \end{pmatrix}$ also satisfies the equations defining W , so this is in W and hence W is closed under addition.

However, there is a simpler way to see that W is a subspace of \mathbb{R}^3 . The equations defining W say precisely that

$$\begin{pmatrix} 1 & -1 & 1 \\ 2 & -1 & -2 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

so W is the same as the kernel of $\begin{pmatrix} 1 & -1 & 1 \\ 2 & -1 & -2 \end{pmatrix}$. Since kernels are always subspaces, W is indeed a subspace of \mathbb{R}^3 .

Definition of linearly dependent/independent. A collection of vectors $\vec{v}_1, \dots, \vec{v}_k$ in \mathbb{R}^n is said to be *linearly dependent* if one vector is a linear combination of the rest. We call such a vector a “redundant” vector since eliminating that vector from the collection will not change the span of the collection. If no vector is a linear combination of the others, we say that $\vec{v}_1, \dots, \vec{v}_k$ are *linearly independent*.

Intuitive idea behind independence. Vectors $\vec{v}_1, \dots, \vec{v}_k$ are linearly independent if each vector adds a new “dimension” to their span. (We will give a precise meaning to dimension later.)

Example 1. We check whether the vectors

$$\begin{pmatrix} 1 \\ 2 \\ -2 \\ -3 \end{pmatrix}, \begin{pmatrix} -1 \\ 0 \\ 4 \\ 5 \end{pmatrix}, \begin{pmatrix} 2 \\ 3 \\ -7 \\ -5 \end{pmatrix}, \text{ and } \begin{pmatrix} 0 \\ 3 \\ 4 \\ 1 \end{pmatrix}$$

in \mathbb{R}^4 are linearly dependent or independent. According to the definition, we must see if any vector is a linear combination of the rest. For instance, to check if the first is a linear combination of the other three we ask whether

$$a \begin{pmatrix} -1 \\ 0 \\ 4 \\ 5 \end{pmatrix} + b \begin{pmatrix} 2 \\ 3 \\ -7 \\ -5 \end{pmatrix} + c \begin{pmatrix} 0 \\ 3 \\ 4 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 2 \\ -2 \\ -3 \end{pmatrix}$$

has a solution for a, b, c . Reducing the corresponding augmented matrix gives

$$\left(\begin{array}{ccc|c} -1 & 2 & 0 & 1 \\ 0 & 3 & 3 & 2 \\ 4 & -7 & 4 & -2 \\ 5 & -5 & 1 & -3 \end{array} \right) \rightarrow \left(\begin{array}{ccc|c} -1 & 2 & 0 & 1 \\ 0 & 3 & 3 & 2 \\ 0 & 1 & 4 & 2 \\ 0 & 5 & 1 & 2 \end{array} \right) \rightarrow \left(\begin{array}{ccc|c} -1 & 2 & 0 & 1 \\ 0 & 3 & 3 & 2 \\ 0 & 0 & -9 & -4 \\ 0 & 0 & 0 & 4 \end{array} \right),$$

from which we can see there is no solution. Thus the first vector in our collection is not a linear combination of the others, so it is not redundant.

We move on and ask whether the second vector is a linear combination of the rest; that is, does

$$a \begin{pmatrix} 1 \\ 2 \\ -2 \\ -3 \end{pmatrix} + b \begin{pmatrix} 2 \\ 3 \\ -7 \\ -5 \end{pmatrix} + c \begin{pmatrix} 0 \\ 3 \\ 4 \\ 1 \end{pmatrix} = \begin{pmatrix} -1 \\ 0 \\ 4 \\ 5 \end{pmatrix}$$

have a solution for a, b, c . As before, we can reduce the corresponding augmented matrix:

$$\left(\begin{array}{ccc|c} 1 & 2 & 0 & -1 \\ 2 & 3 & 3 & 0 \\ -2 & -7 & 4 & 4 \\ -3 & -5 & 1 & 5 \end{array} \right)$$

until we see that this won't have a solution either. So, the second vector in our collection is not redundant either.

And so on, we can do the same for the third vector and then the fourth. However, note that this gets pretty tedious, and you can imagine that if we had more vectors in our collection this process would become way to much work. We will need a better way to check for dependence/independence.

Important. Vectors $\vec{v}_1, \dots, \vec{v}_k$ in \mathbb{R}^n are linearly independent if and only if the only solution of

$$c_1 \vec{v}_1 + \dots + c_k \vec{v}_k = \vec{0}$$

is the zero solution $c_1 = \dots = c_k = 0$. So, to test whether some vectors are linearly independent we set a linear combination of them equal to $\vec{0}$ and solve for the corresponding coefficients; if all the coefficients must be 0 the vectors are independent, if at least one is nonzero they are dependent.

Remark. The idea behind this fact is simple: if one vector is a linear combination of the others, say

$$\vec{v}_1 = c_2\vec{v}_2 + \dots + c_k\vec{v}_k,$$

we can rewrite this as

$$-\vec{v}_1 + c_2\vec{v}_2 + \dots + c_k\vec{v}_k = \vec{0}$$

with at least one coefficient, namely the one in front of \vec{v}_1 equal to 0. Similarly, if we have

$$c_1\vec{v}_1 + \dots + c_k\vec{v}_k = \vec{0}$$

with at least one coefficient nonzero, say c_1 , we can use this equation to “solve” for \vec{v}_1 in terms of the other vectors by moving $c_1\vec{v}_1$ to one side and dividing by c_1 (which is why we need a nonzero coefficient).

Back to Example 1. Consider the equation

$$c_1 \begin{pmatrix} 1 \\ 2 \\ -2 \\ -3 \end{pmatrix} + c_2 \begin{pmatrix} -1 \\ 0 \\ 4 \\ 5 \end{pmatrix} + c_3 \begin{pmatrix} 2 \\ 3 \\ -7 \\ -5 \end{pmatrix} + c_4 \begin{pmatrix} 0 \\ 3 \\ 4 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

Reducing the corresponding augmented matrix gives:

$$\left(\begin{array}{cccc|c} 1 & -1 & 2 & 0 & 0 \\ 2 & 0 & 3 & 3 & 0 \\ -2 & 4 & -7 & 3 & 0 \\ -3 & 5 & -5 & 1 & 0 \end{array} \right) \rightarrow \left(\begin{array}{cccc|c} 1 & -1 & 2 & 0 & 0 \\ 0 & 2 & -1 & 3 & 0 \\ 0 & 0 & -2 & 1 & 0 \\ 0 & 0 & 0 & -1 & 0 \end{array} \right),$$

and we see that the only solution to our equation is the zero solution $c_1 = c_2 = c_3 = c_4 = 0$. Hence the vectors in Example 1 are linearly independent.

Definition of basis. Suppose that V is a subspace of \mathbb{R}^n . A collection of vectors $\vec{v}_1, \dots, \vec{v}_k$ in V are said to be a *basis* of V if:

- $\vec{v}_1, \dots, \vec{v}_k$ span all of V , and
- $\vec{v}_1, \dots, \vec{v}_k$ are linearly independent.

The point is the following: the first condition says that the basis vectors are enough to be able to describe any vector in V since anything in V can be written as a linear combination of the basis vectors, and the second condition says that the basis vectors constitute the fewest number of vectors for which this spanning condition is true.

Important. Intuitively, a basis of V is a “minimal” spanning set of V .

Standard basis of \mathbb{R}^n . The vectors

$$\vec{e}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ and } \vec{e}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

form a basis of \mathbb{R}^2 , called the *standard* basis of \mathbb{R}^2 . In general, the *standard* basis of \mathbb{R}^n is the collection of vectors $\vec{e}_1, \dots, \vec{e}_n$ where \vec{e}_i is the vector with 1 in the i -th position and 0's everywhere else. Expressing an arbitrary vector in \mathbb{R}^n as a linear combination of the standard basis of \mathbb{R}^n is easy:

$$\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} = x_1\vec{e}_1 + x_2\vec{e}_2 + \cdots + x_n\vec{e}_n.$$

Remark. Note that a space can (and will) have more than one possible basis. For example, the vectors

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix} \text{ and } \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

also form a basis of \mathbb{R}^2 . Indeed, for any $\begin{pmatrix} x \\ y \end{pmatrix}$ in \mathbb{R}^2 the equation

$$\begin{pmatrix} x \\ y \end{pmatrix} = c_1 \begin{pmatrix} 1 \\ 1 \end{pmatrix} + c_2 \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

has a solution, so these two vectors span all of \mathbb{R}^2 , and these two vectors are linearly independent since neither is a multiple of the other. So, although above we defined what we mean by the “standard” basis of \mathbb{R}^n , it is important to realize that \mathbb{R}^n has many other bases. What is common to them all however is that they will *all* consist of n vectors; this is related to the notion of “dimension”, which we will come back to next time.

Example 2. We determine bases for the kernel and image of

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

There is a standard way of doing this, using the reduced echelon form of A , which is:

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

We first use this to find vectors spanning the kernel of A . Anything in the kernel looks like

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{pmatrix} = \begin{pmatrix} s \\ -2t - 3u \\ t \\ -4u \\ u \end{pmatrix} = s \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} 0 \\ -2 \\ 1 \\ 0 \\ 0 \end{pmatrix} + u \begin{pmatrix} 0 \\ -3 \\ 0 \\ -4 \\ 1 \end{pmatrix},$$

so

$$\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ -2 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \text{ and } \begin{pmatrix} 0 \\ -3 \\ 0 \\ -4 \\ 1 \end{pmatrix}$$

span $\ker A$. Now, note that these vectors are actually linearly independent! Indeed, the first is not a linear combination of the other two since any combination of the other two will have a zero first entry, the second is not a linear combination of the other two since any combination of the first and third will have a zero third entry, and the third is not a linear combination of the first two since any combination of the first two will have a zero fifth entry. (This will *always* happen when using the procedure we've described before for finding vectors which span the kernel of a matrix.) Thus the three vectors above form a basis of $\ker A$.

To find a basis of $\text{im } A$, we also look at the reduced echelon form. In this reduced form, note that the second and fourth columns are the “pivot columns”, meaning columns which contain a pivot. It turns out that the corresponding columns of the *original* matrix form a basis of the image of A ! (We'll see why next time.) In our case then, the second and fourth columns of A :

$$\begin{pmatrix} 1 \\ -2 \\ 3 \end{pmatrix} \text{ and } \begin{pmatrix} 0 \\ 1 \\ -2 \end{pmatrix}$$

form a basis of $\text{im } A$.

Important. To find a basis for the kernel or image of a matrix, find the reduced echelon form of that matrix. Then:

- For the kernel, find vectors spanning the kernel as we have done before using the idea of “factoring” out free variables. The resulting vectors will be a basis for the kernel.
- For the image, take the columns in the original matrix which correspond to the pivot columns in the echelon form. These columns form a basis for the image.

Lecture 13: Bases and Dimension

Today we continued talking about the notion of a basis of a subspace of \mathbb{R}^n , and introduced the idea of the “dimension” of a subspace. The dimension of a subspace matches up with our usual geometric intuition as to what “dimension” should mean: i.e. lines are 1-dimensional, planes are 2-dimensional, and so on.

Warm-Up 1. We find all numbers k such that

$$\begin{pmatrix} 1 \\ -2 \\ -1 \end{pmatrix}, \begin{pmatrix} 0 \\ 3 \\ -2 \end{pmatrix}, \text{ and } \begin{pmatrix} -1 \\ -4 \\ k \end{pmatrix}$$

are linearly independent. These are linearly independent when

$$c_1 \begin{pmatrix} 1 \\ -2 \\ -1 \end{pmatrix} + c_2 \begin{pmatrix} 0 \\ 3 \\ -2 \end{pmatrix} + c_3 \begin{pmatrix} -1 \\ -4 \\ k \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

has only the zero solution. Reducing the augmented matrix of the corresponding system gives:

$$\left(\begin{array}{ccc|c} 1 & 0 & -1 & 0 \\ -2 & 3 & -4 & 0 \\ -1 & -2 & k & 0 \end{array} \right) \rightarrow \left(\begin{array}{ccc|c} 1 & 0 & -1 & 0 \\ 0 & 3 & -6 & 0 \\ 0 & -2 & k-1 & 0 \end{array} \right) \rightarrow \left(\begin{array}{ccc|c} 1 & 0 & -1 & 0 \\ 0 & 3 & -6 & 0 \\ 0 & 0 & k-5 & 0 \end{array} \right).$$

Thus our equation has only the zero solution $c_1 = c_2 = c_3 = 0$ whenever $k \neq 5$, so the three given vectors are linearly independent as long as $k \neq 5$.

For $k = 5$ our vectors are then linearly dependent, and it should be possible to express one as a linear combination of the rest. Indeed, when $k = 5$ we have

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

and this can be used to express any vector in our original collection as a linear combination of the other two.

Warm-Up 2. We find a basis for the span of

$$\begin{pmatrix} 1 \\ -2 \\ 3 \\ -1 \end{pmatrix}, \begin{pmatrix} 2 \\ -4 \\ 6 \\ -2 \end{pmatrix}, \begin{pmatrix} 1 \\ -2 \\ 5 \\ 1 \end{pmatrix}, \begin{pmatrix} 4 \\ -3 \\ 10 \\ -2 \end{pmatrix}, \text{ and } \begin{pmatrix} -1 \\ -3 \\ -4 \\ -4 \end{pmatrix}.$$

Note that for sure the second vector is redundant since it is a multiple of the first, so removing it from our collection won't change the span. Instead of trying to see which other vectors are redundant by inspection, we use a more systematic approach. First, note that the span of these vectors is the same as the image of the matrix

$$A = \begin{pmatrix} 1 & 2 & 1 & 4 & -1 \\ -2 & -4 & -2 & -3 & -3 \\ 3 & 6 & 5 & 10 & -4 \\ -1 & -2 & 1 & -2 & -4 \end{pmatrix}.$$

So, we are really looking for a basis for this image, and we saw last time how to find one. Reducing this matrix a bit gives:

$$\begin{pmatrix} 1 & 2 & 1 & 4 & -1 \\ -2 & -4 & -2 & -3 & -3 \\ 3 & 6 & 5 & 10 & -4 \\ -1 & -2 & 1 & -2 & -4 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 2 & 1 & 4 & -1 \\ 0 & 0 & 2 & -2 & -1 \\ 0 & 0 & 0 & -4 & 4 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

The claim was that the columns in the original matrix which correspond to pivot columns in the reduced echelon form give a basis for the image. Our matrix so far is not yet in reduced echelon form, but we can already tell that the first, third, and fourth columns will be the ones containing pivots in the end. Thus

$$\begin{pmatrix} 1 \\ -2 \\ 3 \\ -1 \end{pmatrix}, \begin{pmatrix} 1 \\ -2 \\ 5 \\ 1 \end{pmatrix}, \begin{pmatrix} 4 \\ -3 \\ 10 \\ -2 \end{pmatrix}$$

forms a basis for the image of A , and hence a basis for the span of the original five vectors.

Let us justify a bit why these vectors indeed form a basis for $\text{im } A$. First, they are linearly independent: the only solution of

$$\begin{pmatrix} 1 & 1 & 4 \\ -2 & -2 & -3 \\ 3 & 5 & 10 \\ -1 & 1 & -2 \end{pmatrix} \vec{x} = \vec{0}$$

is $\vec{x} = \vec{0}$ since row-reducing this matrix ends up giving

$$\begin{pmatrix} 1 & 1 & 4 \\ 0 & 2 & -2 \\ 0 & 0 & -4 \\ 0 & 0 & 0 \end{pmatrix},$$

which are the first, third and fourth columns from the reduced form of A we computed above. Second, any other vector among our original list is a linear combination of these three: for instance if we want to write the fifth vector as

$$\begin{pmatrix} -1 \\ -3 \\ -4 \\ -4 \end{pmatrix} = c_1 \begin{pmatrix} 1 \\ -2 \\ 3 \\ -1 \end{pmatrix} + c_2 \begin{pmatrix} 1 \\ -2 \\ 5 \\ 1 \end{pmatrix} + c_3 \begin{pmatrix} 4 \\ -3 \\ 10 \\ -2 \end{pmatrix},$$

the corresponding augmented matrix reduces to

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

based on the reduced form of A we previously computed. From here we can see that there are scalars c_1, c_2, c_3 which express the fifth vector in our original list as a linear combination of the three we are claiming is a basis. These ideas carry over for any matrix, which is why our method for finding a basis for the image of a matrix always works.

Number of linearly independent and spanning vectors. For a subspace V of \mathbb{R}^n , any linearly independent set of vectors in V always has fewer (or as many) vectors than any spanning set of vectors in V :

$$\# \text{ of any linearly independent vectors} \leq \# \text{ of any spanning vectors.}$$

For instance, say we have four vectors $\vec{v}_1, \vec{v}_2, \vec{v}_3, \vec{v}_4$ in \mathbb{R}^3 . Since the standard basis of \mathbb{R}^3 is a spanning set with three vectors, our four vectors cannot be linearly independent. So, as soon as we have more than three vectors in \mathbb{R}^3 they must be linearly dependent. Similarly, anytime we have fewer than three vectors in \mathbb{R}^3 , they cannot possibly span all of \mathbb{R}^3 since the standard basis is a linearly independent set with three vectors.

The dimension of a subspace. The above fact tells us the following. Say we have two bases $\vec{v}_1, \dots, \vec{v}_k$ and $\vec{w}_1, \dots, \vec{w}_\ell$ for a subspace V of \mathbb{R}^n . Viewing the \vec{v} 's as the linearly independent set and the \vec{w} 's as the spanning set we get that

$$k \leq \ell.$$

Switching roles and viewing \vec{v} 's as spanning and \vec{w} 's as linearly independent gives

$$\ell \leq k.$$

These two inequalities together imply that $k = \ell$, so we come to the conclusion that *any* two bases of V must have the same number of vectors! This common number is what we call the *dimension* of V , and we denote it by $\dim V$.

Important. Any two bases of a subspace V of \mathbb{R}^n have the same number of vectors, and $\dim V$ is equal to the number of vectors in any basis of V .

Simplified basis check. Say $\dim V = n$. Then any n linearly independent vectors must automatically span V . Indeed, if $\vec{v}_1, \dots, \vec{v}_n$ were linearly independent in V and did not span V it would be possible to “extend” this to a basis of V , giving a basis of V with more than $n = \dim V$ vectors. This is not possible, so $\vec{v}_1, \dots, \vec{v}_n$ being linearly independent is enough to guarantee that they actually form a basis of V .

We can also see this computationally, say in the case $V = \mathbb{R}^n$. If $\vec{v}_1, \dots, \vec{v}_n$ are linearly independent vectors in \mathbb{R}^n , the only solution of

$$\begin{pmatrix} | & & | \\ \vec{v}_1 & \cdots & \vec{v}_n \\ | & & | \end{pmatrix} \vec{x} = \vec{0}$$

is the zero vector $\vec{x} = \vec{0}$. (Here, the matrix is the one whose columns are $\vec{v}_1, \dots, \vec{v}_n$.) But, the only way in which this can be possible is for the reduced echelon form of this matrix to be the identity:

$$\begin{pmatrix} | & & | \\ \vec{v}_1 & \cdots & \vec{v}_n \\ | & & | \end{pmatrix} \rightarrow \begin{pmatrix} 1 & & & & 0 \\ & 1 & & & \\ & & \cdots & & \\ 0 & & & & 1 \end{pmatrix}.$$

But with this echelon form, it is true that

$$\begin{pmatrix} | & & | \\ \vec{v}_1 & \cdots & \vec{v}_n \\ | & & | \end{pmatrix} \vec{x} = \vec{b}$$

has a solution no matter what \vec{b} is, meaning that any \vec{b} in \mathbb{R}^n is a linear combination of the columns of this matrix. Thus $\vec{v}_1, \dots, \vec{v}_n$ span all of \mathbb{R}^n and hence form a basis of \mathbb{R}^n .

Similarly, any n vectors in \mathbb{R}^n which span \mathbb{R}^n must automatically be linearly independent and hence form a basis of \mathbb{R}^n .

Important. In an n -dimensional space V , any n linearly independent vectors automatically form a basis, and any n spanning vectors automatically form a basis.

The dimensions of the kernel and image of a matrix. Let us compute the dimension of the kernel and image of

$$A = \begin{pmatrix} 1 & 2 & 1 & 4 & -1 \\ -2 & -4 & -2 & -3 & -3 \\ 3 & 6 & 5 & 10 & -4 \\ -1 & -2 & 1 & -2 & -4 \end{pmatrix}.$$

We previously reduced this in the second Warm-Up to end up with

$$A \rightarrow \begin{pmatrix} 1 & 2 & 1 & 4 & -1 \\ 0 & 0 & 2 & -2 & -1 \\ 0 & 0 & 0 & -4 & 4 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Each pivot column contributes one basis vector for the image, so $\dim \operatorname{im} A = 3$. Notice that this is precisely the same as the rank of A , and we finally have our long-awaited meaning behind the rank of a matrix: it is simply the dimension of its image! Based on the method we saw last time for finding a basis for the kernel of a matrix, each free variable contributes one basis vector for the kernel, so $\dim \ker A$ equals the number of free variables. In this case, $\dim \ker A = 2$.

Important. For any matrix A ,

- $\dim \operatorname{im} A = \operatorname{rank} A = \#$ of pivots in the reduced echelon form of A , and
- $\dim \ker A = \#$ of free variables in the reduced echelon form of A

The rank-nullity theorem. For any matrix A , we now see that

$$\operatorname{rank} A + \dim(\ker A) = \# \text{ of columns of } A,$$

since any column of A either corresponds to a pivot in the echelon form (thus contributing to $\operatorname{rank} A$) or to a free variable in the echelon form (thus contributing to $\dim \ker A$). The dimension of $\ker A$ is often called the *nullity* of A , so this theorem says that $\operatorname{rank} + \text{nullity}$ equals the number of columns, which is where the name of the theorem comes from.

Back to rank. Now that we've seen the true meaning of the rank of a matrix as the dimension of its image, we can justify some properties of rank you may have come across before. For instance, for matrices A and B for which AB is defined, it is true that

$$\operatorname{rank}(AB) \leq \operatorname{rank}(A).$$

Justifying this directly in terms of number of pivots is not easy, but now it is not so bad: a homework problem showed that $\operatorname{im}(AB)$ was always contained inside $\operatorname{im}(A)$, so

$$\dim \operatorname{im}(AB) \leq \dim \operatorname{im}(A)$$

since a subspace of a larger space always has dimension \leq to that of the larger space.

As an application of this, we can justify why the notion of “invertible” only makes sense for square matrices: that is, if A and B aren't square, it is impossible to have both $AB = I$ and $BA = I$. Consider a specific case of this, say A is 5×7 and B is 7×5 . Then BA is 7×7 , so if this were going to equal the 7×7 identity matrix it would have to have rank 7. But

$$\operatorname{rank}(BA) \leq \operatorname{rank} B \leq 5,$$

so it is not possible to have $BA = I$. Note that it can happen that AB equals the 5×5 identity matrix (try to find an example!), but not both $AB = I$ and $BA = I$.

Lecture 14: Coordinates Relative to a Basis

Today we started talking about the idea of “changing your basis” and computing “coordinates” of a vector of \mathbb{R}^n relative to a non-standard basis. We'll continue with this next time.

Warm-Up 1. We find the dimension of the subspace of \mathbb{R}^3 spanned by

$$\begin{pmatrix} 1 \\ -t \\ 2 \end{pmatrix}, \begin{pmatrix} 3 \\ -2t - 2 \\ t + 4 \end{pmatrix}, \text{ and } \begin{pmatrix} 4 \\ 2 - 4t \\ t + 10 \end{pmatrix}.$$

(Note that these are not quite the vectors I used in the actual Warm-Up in class, which is what led to my confusion. I think I dropped the $-t$ in the first vector and just used t by mistake; these vectors above better illustrate what I meant the Warm-Up to show.) We can view this subspace as the image of the matrix

$$\begin{pmatrix} 1 & 3 & 4 \\ -t & -2t - 2 & 2 - 4t \\ 2 & t + 4 & t + 10 \end{pmatrix},$$

so we are really asking about finding the dimension of the image of this matrix, which is equal to the rank of this matrix. We row-reduce:

$$\begin{pmatrix} 1 & 3 & 4 \\ -t & -2t - 2 & 2 - 4t \\ 2 & t + 4 & t + 10 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 3 & 4 \\ 0 & t - 2 & 2 \\ 0 & t - 2 & t + 2 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 3 & 4 \\ 0 & t - 2 & 2 \\ 0 & 0 & t \end{pmatrix}.$$

Now we can see that our answer will depend on what t is: when $t = 0$ or $t = 2$, our matrix has rank 2 and so our original vectors span a 2-dimensional subspace of \mathbb{R}^3 , while if $t \neq 0$ and $t \neq 2$ our matrix has rank 3 and the span of the original vectors is 3-dimensional.

Geometrically, when $t = 0$ or $t = 2$ our vectors span a plane, while for all other values of t our vectors span all of \mathbb{R}^3 . In this case, there are no values of t for which our vectors will span only a line.

Geometric meaning of non-invertibility. Suppose that A is a non-invertible 2×2 matrix and consider the transformation $T(\vec{x}) = A\vec{x}$. Then rank A is either 0 or 1. When rank $A = 0$, the image of T is zero dimensional and so consists of just the origin $\{\vec{0}\}$; thus in this case the transformation T “collapses” all of \mathbb{R}^2 to a single point. When rank $A = 1$, the image of T is a 1-dimensional line, so T in this case “collapses” all of \mathbb{R}^2 to this line. Thus, geometrically a transformation $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ is non-invertible precisely when it “collapses” 2-dimensional things down to a single point or a line.

Analogously, the transformation T corresponding to a non-invertible 3×3 matrix will collapse 3-dimensional things down to a single point, a line, or a plane. A matrix is invertible precisely when the corresponding transformation does not do any such “collapsing”.

Amazingly Awesome Theorem, continued. The following are equivalent to an $n \times n$ matrix A being invertible:

- The columns of A are linearly independent
- The columns of A span all of \mathbb{R}^n
- The columns of A form a basis of \mathbb{R}^n
- The kernel of A is zero dimensional
- The image of A is n -dimensional

The first two conditions were actually included in the original version of the Amazingly Awesome Theorem I gave, only we hadn’t defined “linearly independent” and “span” at that point. I’m listing them here again for emphasis.

Geometric idea behind coordinates. The idea behind coordinates is the following. Usually we represent vectors in terms of the usual x and y -axes, which geometrically are the span of $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and of $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ respectively. Thus usually we are using the standard basis of \mathbb{R}^2 to represent vectors.

However, there is nothing stopping us from using a different set of axes (i.e. a different basis) for \mathbb{R}^2 to represent vectors. Doing so can help to clarify some properties of linear transformations; in particular, we'll see examples where picking the "right" set of axes can make the geometric interpretation of some transformations easier to identify.

Definition of coordinates. Suppose that $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ is a basis of \mathbb{R}^n . We know that we can then express any \vec{x} in \mathbb{R}^n as a linear combination of the basis vectors in \mathcal{B} :

$$\vec{x} = c_1\vec{v}_1 + \dots + c_n\vec{v}_n \text{ for some scalars } c_1, \dots, c_n.$$

We call c_1, \dots, c_n the *coordinates of \vec{x} relative to the basis \mathcal{B}* . The vector

$$[\vec{x}]_{\mathcal{B}} = \begin{pmatrix} c_1 \\ \vdots \\ c_n \end{pmatrix}$$

is called the *coordinate vector of \vec{x} relative to \mathcal{B}* . Thus, the coordinates of a vector relative to a basis give us the coefficients we need in order to express that vector in terms of said basis.

Example 1. The coordinates of a vector \vec{x} relative to the standard basis of \mathbb{R}^n are simply the entries of \vec{x} , i.e. what we normally mean when we refer to the "coordinates" of \vec{x} :

$$\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} = x_1 \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + x_2 \begin{pmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{pmatrix} + \dots + x_n \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix}.$$

This is what I mean when I say that usually when we express vectors, we are doing so in terms of the standard basis.

Example 2. Consider the basis $\mathcal{B} = \left\{ \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right\}$ of \mathbb{R}^2 . This is a basis because these vectors are linearly independent, and any two linearly independent vectors in a 2-dimensional space automatically span that space. We find the coordinates of any vector $\begin{pmatrix} a \\ b \end{pmatrix}$ relative to this basis. We want c_1 and c_2 satisfying

$$\begin{pmatrix} a \\ b \end{pmatrix} = c_1 \begin{pmatrix} 1 \\ 1 \end{pmatrix} + c_2 \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

We can solve this equation using an augmented matrix as usual, or we can note the following: this equation can be written as

$$\begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix},$$

so we can solve for the coordinates we want by multiplying both sides by the inverse of this matrix! This equation says that our "old" coordinates a and b are related to our "new" ones c_1 and c_2 via multiplication by this matrix; because of this we call

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

the *change of basis* matrix from the basis \mathcal{B} to the standard basis. Since

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}^{-1} \begin{pmatrix} a \\ b \end{pmatrix},$$

we would call

$$\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}^{-1} = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$$

the *change of basis* matrix from the standard basis to \mathcal{B} ; in other words, it is this inverse matrix which tells us how to move from “old” (standard) coordinates to “new” coordinates.

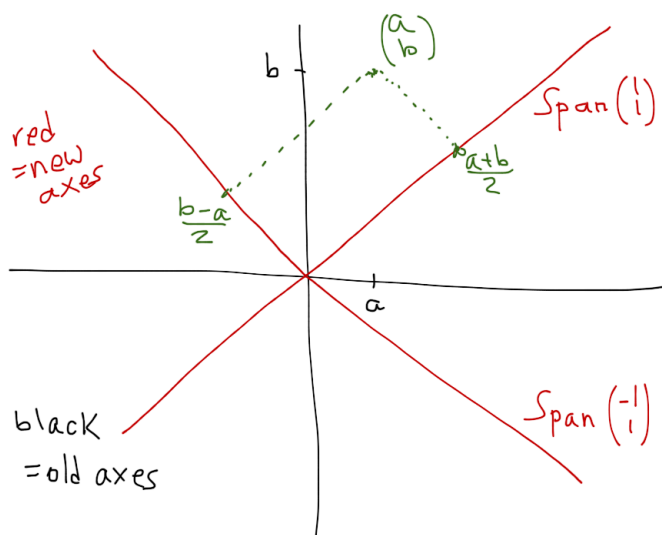
Thus in our case,

$$\left[\begin{pmatrix} a \\ b \end{pmatrix} \right]_{\mathcal{B}} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} \frac{a+b}{2} \\ \frac{b-a}{2} \end{pmatrix},$$

so to express $\begin{pmatrix} a \\ b \end{pmatrix}$ in terms of the basis \mathcal{B} we need to use $\frac{a+b}{2}$ as the coefficient of the first basis vector and $\frac{b-a}{2}$ as the coefficient of the second basis vector:

$$\begin{pmatrix} a \\ b \end{pmatrix} = \frac{a+b}{2} \begin{pmatrix} 1 \\ 1 \end{pmatrix} + \frac{b-a}{2} \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

Geometrically, $\text{span} \left\{ \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right\}$ and $\text{span} \left\{ \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right\}$ give us new axes for \mathbb{R}^2 formed by the lines $y = x$ and $y = -x$ respectively. The coordinates we found tell us how “far along” each of these new axes we have to go in order to reach $\begin{pmatrix} a \\ b \end{pmatrix}$:



Important. Given a basis $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_n\}$ of \mathbb{R}^n , the matrix S whose columns are the basis vectors is called the *change of basis* matrix from \mathcal{B} to the standard basis, and its inverse S^{-1} is the change of basis matrix from the standard basis to \mathcal{B} . (Note that S is invertible according to the Amazingly Awesome Theorem since its columns form a basis of \mathbb{R}^n .) These matrices have the properties that:

$$[\vec{x}]_{\text{old}} = S[\vec{x}]_{\text{new}} \text{ and } [\vec{x}]_{\text{new}} = S^{-1}[\vec{x}]_{\text{old}},$$

where by “new” we mean relative to the basis \mathcal{B} and by “old” we mean relative to the standard basis. The book doesn’t use the term “change of basis” matrix, but I think it is a useful term since it emphasizes the role the matrix S plays in moving between different coordinates.

Example 3. We find the coordinates of $\begin{pmatrix} 3 \\ 5 \end{pmatrix}$ relative to the basis $\mathcal{B} = \left\{ \begin{pmatrix} 3 \\ -5 \end{pmatrix}, \begin{pmatrix} -7 \\ 12 \end{pmatrix} \right\}$ of \mathbb{R}^2 ; that is, we want c_1 and c_2 such that

$$\begin{pmatrix} 3 \\ 5 \end{pmatrix} = c_1 \begin{pmatrix} 3 \\ -5 \end{pmatrix} + c_2 \begin{pmatrix} -7 \\ 12 \end{pmatrix}.$$


The change of basis matrix in this case is

$$S = \begin{pmatrix} 3 & -7 \\ -5 & 12 \end{pmatrix},$$

so the coordinates we want are given by

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = S^{-1} \begin{pmatrix} 3 \\ 5 \end{pmatrix} = \begin{pmatrix} 12 & 7 \\ 5 & 3 \end{pmatrix} \begin{pmatrix} 3 \\ 5 \end{pmatrix} = \begin{pmatrix} 71 \\ 30 \end{pmatrix}.$$

Geometrically, the “axes” determined by our basis vectors are non-perpendicular lines, and these coordinates tells us how far along these axes we must go in order to reach $\begin{pmatrix} 3 \\ 5 \end{pmatrix}$:



Transformations relative to new coordinates. Let T be the linear transformation which reflects \mathbb{R}^2 across the line $y = x$. The matrix of this transformation is

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

We now describe the same transformation only now relative to “new” coordinates determined by the basis $\mathcal{B} = \left\{ \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right\}$. Why would we want to do this? In this case the matrix of T is pretty simple already, and changing coordinates really won’t help to simplify matters much. However, we will see next time examples where we might have a messy-looking transformation and changing coordinates will *definitely* help to simplify the description of what’s going on. Here we work this out in a simple case only to see how the general process works.

Recall that previously found the matrix of T by applying the transformation $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$, and using these as the columns of the matrix. We now do exactly the same thing, only replacing the standard basis with the basis \mathcal{B} ; that is, we apply T to the first basis vector and compute the coordinates of the result relative to \mathcal{B} , and then do the same for the second basis vector. The resulting coordinate vectors form the columns of the *matrix of T relative to \mathcal{B}* :

$$[T]_{\mathcal{B}} = \left(\left[T \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right]_{\mathcal{B}}, \left[T \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right]_{\mathcal{B}} \right).$$

First, T does nothing to $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ since $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is on the line we are reflecting across, so

$$T \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

The coordinates of this relative to \mathcal{B} are $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ since to write $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ as a linear combination of the basis vectors in \mathcal{B} we take

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

After all, $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is the first basis vector already so we don't even need the second basis vector if we want to express the first as a linear combination of the two. The first column of the matrix of T relative to \mathcal{B} is then $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$.

We have

$$T \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$$

since anything perpendicular to the line we reflect across just has its direction flipped around. The coordinates of this relative to \mathcal{B} are

$$\left[T \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right]_{\mathcal{B}} = \left[\begin{pmatrix} 1 \\ -1 \end{pmatrix} \right]_{\mathcal{B}} = \begin{pmatrix} 0 \\ -1 \end{pmatrix}$$

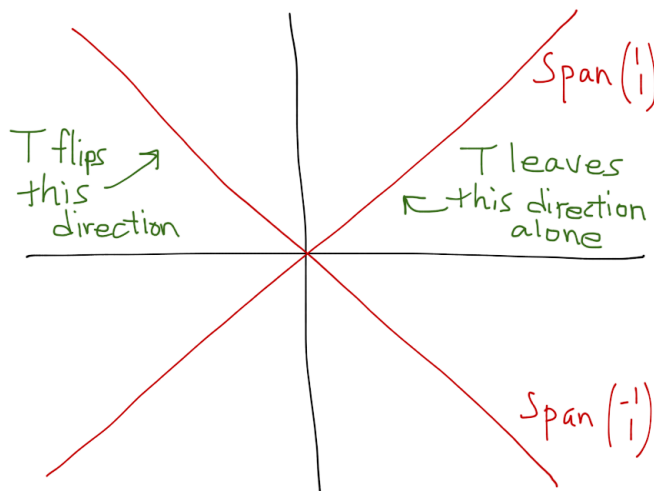
since we don't even need to use the first basis vector in order to express the second as a linear combination of the two. The matrix of T relative to \mathcal{B} is thus

$$[T]_{\mathcal{B}} = \left(\left[T \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right]_{\mathcal{B}} \quad \left[T \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right]_{\mathcal{B}} \right) = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The point of this matrix is the following: say we want to figure out what T does to some vector \vec{x} . We can determine this by taking the coordinate vector of \vec{x} relative \mathcal{B} and multiplying *that* by $[T]_{\mathcal{B}}$; the result will be the coordinate vector of $T(\vec{x})$ relative to \mathcal{B} :

$$[T(\vec{x})]_{\mathcal{B}} = [T]_{\mathcal{B}}[\vec{x}]_{\mathcal{B}}.$$

In this example, the matrix of T relative to \mathcal{B} tells us that geometrically T leaves the span $\left\{ \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right\}$ -coordinate of a vector alone (due to the first column being $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$) but changes the sign of the span $\left\{ \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right\}$ -coordinate (due to the second column being $\begin{pmatrix} 0 \\ -1 \end{pmatrix}$), which is precisely what reflection across $y = x$ should do.



The point is that by switching to a “better” set of axes, we have somewhat simplified the geometric description of T . We’ll see more and better examples of this next time.

Lecture 15: More on Coordinates

Today we continued talking about coordinates and the matrix of a transformation relative to a non-standard basis. We looked at examples which hopefully show why you would want to consider such a thing.


Warm-Up 1. We compute the coordinates of $\begin{pmatrix} 1 \\ 2 \end{pmatrix}$ relative to the basis $\mathcal{B} = \left\{ \begin{pmatrix} 2 \\ 1 \end{pmatrix}, \begin{pmatrix} -3 \\ 1 \end{pmatrix} \right\}$ of \mathbb{R}^2 . We want c_1 and c_2 satisfying

$$\begin{pmatrix} 1 \\ 2 \end{pmatrix} = c_1 \begin{pmatrix} 2 \\ 1 \end{pmatrix} + c_2 \begin{pmatrix} -3 \\ 1 \end{pmatrix}.$$

One way to find these coefficients is by solving the corresponding system of equations, but instead we use the change of basis matrix $\begin{pmatrix} 2 & -3 \\ 1 & 1 \end{pmatrix}$. We have

$$\begin{pmatrix} 1 \\ 2 \end{pmatrix} = \begin{pmatrix} 2 & -3 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}, \text{ so } \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 2 & -3 \\ 1 & 1 \end{pmatrix}^{-1} \begin{pmatrix} 1 \\ 2 \end{pmatrix} = \frac{1}{5} \begin{pmatrix} 1 & 3 \\ -1 & 2 \end{pmatrix} \begin{pmatrix} 1 \\ 2 \end{pmatrix} = \begin{pmatrix} 7/5 \\ 3/5 \end{pmatrix}.$$

Thus the coordinates of $\begin{pmatrix} 1 \\ 2 \end{pmatrix}$ relative to \mathcal{B} are $\frac{7}{5}$ and $\frac{3}{5}$. Viewing $\text{span} \left\{ \begin{pmatrix} 2 \\ 1 \end{pmatrix} \right\}$ and $\text{span} \left\{ \begin{pmatrix} -3 \\ 1 \end{pmatrix} \right\}$ as a new set of axes for \mathbb{R}^2 , these coordinates tell us how far along these new axes we must go in order to reach $\begin{pmatrix} 1 \\ 2 \end{pmatrix}$:



Warm-Up 2. We find the matrix of the reflection across the line spanned by $\begin{pmatrix} 4 \\ 3 \end{pmatrix}$ relative to the basis $\mathcal{B} = \left\{ \begin{pmatrix} 4 \\ 3 \end{pmatrix}, \begin{pmatrix} -3 \\ 4 \end{pmatrix} \right\}$ of \mathbb{R}^2 . Recall that this matrix is given by

$$[T]_{\mathcal{B}} = \left(\left[T \begin{pmatrix} 4 \\ 3 \end{pmatrix} \right]_{\mathcal{B}}, \left[T \begin{pmatrix} -3 \\ 4 \end{pmatrix} \right]_{\mathcal{B}} \right),$$

where we take the coordinate vectors of $T \begin{pmatrix} 4 \\ 3 \end{pmatrix}$ and $T \begin{pmatrix} -3 \\ 4 \end{pmatrix}$ relative to the basis \mathcal{B} . Note that on a previous homework assignment you computed the matrix of this reflection relative to the standard basis, which ended up being

$$\begin{pmatrix} 7/25 & 24/25 \\ 24/25 & -7/25 \end{pmatrix}.$$

The point is that the matrix of T relative to \mathcal{B} is much simpler than this, because the basis \mathcal{B} has nice properties with respect to T .

First,

$$T \begin{pmatrix} 4 \\ 3 \end{pmatrix} = \begin{pmatrix} 4 \\ 3 \end{pmatrix}$$

since $\begin{pmatrix} 4 \\ 3 \end{pmatrix}$ is on the line we are reflecting across. The coordinate vector of this relative to \mathcal{B} is

$$\left[\begin{pmatrix} 4 \\ 3 \end{pmatrix} \right]_{\mathcal{B}} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

since $\begin{pmatrix} 4 \\ 3 \end{pmatrix}$ is itself the first vector in our basis. Next,

$$T \begin{pmatrix} -3 \\ 4 \end{pmatrix} = \begin{pmatrix} 3 \\ -4 \end{pmatrix}$$

since any vector perpendicular to the line of reflection simply has its direction changed by multiplying by -1 . This has coordinate vector

$$\left[\begin{pmatrix} 3 \\ -4 \end{pmatrix} \right]_{\mathcal{B}} = \begin{pmatrix} 0 \\ -1 \end{pmatrix}$$

since $\begin{pmatrix} -3 \\ 4 \end{pmatrix}$ is 0 times the first basis vector plus -1 times the second. The matrix of T relative to \mathcal{B} is then

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

As we know from the geometric description of T , this matrix suggests that T leaves the span $\{\begin{pmatrix} 4 \\ 3 \end{pmatrix}\}$ -direction of a vector alone while flipping the span $\{\begin{pmatrix} -3 \\ 4 \end{pmatrix}\}$ -direction. Also, as we said earlier, this matrix is much simpler than the matrix of T relative to the standard basis.

Similarly, we can consider the transformation R which is orthogonal projection of \mathbb{R}^2 onto the line spanned by $\begin{pmatrix} 4 \\ 3 \end{pmatrix}$. Since this satisfies

$$R \begin{pmatrix} 4 \\ 3 \end{pmatrix} = \begin{pmatrix} 4 \\ 3 \end{pmatrix} \text{ and } R \begin{pmatrix} -3 \\ 4 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

which respectively have coordinate vectors

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ and } \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

relative to \mathcal{B} , the matrix of R relative to \mathcal{B} is

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

This again is much simpler than the matrix of T relative to the standard basis, which is

$$\begin{pmatrix} 31/25 & 24/25 \\ 24/25 & 31/25 \end{pmatrix}.$$

Important. We emphasize that the matrix of T relative to \mathcal{B} satisfies

$$[T(\vec{x})]_{\mathcal{B}} = [T]_{\mathcal{B}}[\vec{x}]_{\mathcal{B}},$$

which says that to determine the result of applying T to a vector \vec{x} , we can take the coordinate vector of \vec{x} relative \mathcal{B} and multiply it by $[T]_{\mathcal{B}}$; the result will be the coordinate vector of $T(\vec{x})$ relative to \mathcal{B} . You should view this equation as analogous to $T(\vec{x}) = A\vec{x}$, only now we are writing everything in terms of a new basis.

Example 1. We want to come up with a geometric description of the transformation defined by $T(\vec{x}) = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix} \vec{x}$. For a first attempt, we compute:

$$T \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 13 \\ -1 \end{pmatrix} \text{ and } T \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} -6 \\ 12 \end{pmatrix}.$$

If we draw these vectors it is not clear at all what kind of geometric properties T has: it doesn't appear to be a rotation nor a reflection, and it's hard to guess whether it might be some kind of shear or something else. The problem is that we're using the standard basis of \mathbb{R}^2 to analyze T .

Instead, let us compute the matrix of T relative to the basis from the first Warm-Up: $\mathcal{B} = \left\{ \begin{pmatrix} 2 \\ 1 \end{pmatrix}, \begin{pmatrix} -3 \\ 1 \end{pmatrix} \right\}$. We have

$$T \begin{pmatrix} 2 \\ 1 \end{pmatrix} = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix} \begin{pmatrix} 2 \\ 1 \end{pmatrix} = \begin{pmatrix} 20 \\ 10 \end{pmatrix},$$

which has coordinates 10 and 0 relative to \mathcal{B} since it is just 10 times the first basis vector. Also,

$$T \begin{pmatrix} -3 \\ 1 \end{pmatrix} = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix} \begin{pmatrix} -3 \\ 1 \end{pmatrix} = \begin{pmatrix} -45 \\ 15 \end{pmatrix},$$

which has coordinates 0 and 15 since it is 15 times the second basis vector. The matrix of T relative to \mathcal{B} is thus

$$[T]_{\mathcal{B}} = \begin{pmatrix} 10 & 0 \\ 0 & 15 \end{pmatrix}.$$

Now, what does this tell us? If we consider the axes corresponding to \mathcal{B} , the form of this matrix tells us that T acts by scaling the span $\left\{ \begin{pmatrix} 2 \\ 1 \end{pmatrix} \right\}$ -direction by a factor of 10 and scales the span $\left\{ \begin{pmatrix} -3 \\ 1 \end{pmatrix} \right\}$ -direction by a factor of 15. Thus we do have a pretty nice description of what T does geometrically, which would have been nearly impossible to determine given the original definition of T .

Now, in the first Warm-Up we computed that

$$\left[\begin{pmatrix} 1 \\ 2 \end{pmatrix} \right]_{\mathcal{B}} = \begin{pmatrix} 7/5 \\ 3/5 \end{pmatrix}.$$

Recall that the matrix of T relative to \mathcal{B} satisfies:

$$[T(\vec{x})]_{\mathcal{B}} = [T]_{\mathcal{B}}[\vec{x}]_{\mathcal{B}},$$

which says that multiplying the coordinate vector of \vec{x} by $[T]_{\mathcal{B}}$ gives the coordinate vector of $T(\vec{x})$. Thus we should have

$$\left[T \begin{pmatrix} 1 \\ 2 \end{pmatrix} \right]_{\mathcal{B}} = \begin{pmatrix} 10 & 0 \\ 0 & 15 \end{pmatrix} \begin{pmatrix} 7/5 \\ 3/5 \end{pmatrix} = \begin{pmatrix} 70/5 \\ 45/5 \end{pmatrix} = \begin{pmatrix} 14 \\ 9 \end{pmatrix}.$$


Indeed, we can directly compute $T \begin{pmatrix} 1 \\ 2 \end{pmatrix}$ as

$$T \begin{pmatrix} 1 \\ 2 \end{pmatrix} = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix} \begin{pmatrix} 1 \\ 2 \end{pmatrix} = \begin{pmatrix} 1 \\ 23 \end{pmatrix},$$

and you can check that the coordinates of this relative to \mathcal{B} are indeed 14 and 9:

$$\begin{pmatrix} 1 \\ 23 \end{pmatrix} = 14 \begin{pmatrix} 2 \\ 1 \end{pmatrix} + 9 \begin{pmatrix} -3 \\ 1 \end{pmatrix}.$$

Geometrically, in terms of our new axes, scaling $7/5$ by 10 and $3/5$ by 15 does look like it should give the coordinates of $\begin{pmatrix} 1 \\ 23 \end{pmatrix}$:



The point again is that T is much simpler to describe now that we've switched to a new basis.

Remark. A fair question to ask at this point is: how did I know that the basis consisting of $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} -3 \\ 1 \end{pmatrix}$ was the right one to use? We'll come back to this later; the answer is related to what are called "eigenvalues" and "eigenvectors".

Example 2. We find the matrix of $T(\vec{x}) = \begin{pmatrix} 4 & 2 & 2 \\ 2 & 4 & 2 \\ 2 & 2 & 4 \end{pmatrix} \vec{x}$ relative to the basis

$$\begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$$

of \mathbb{R}^3 . We have

$$T \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} -2 \\ 0 \\ 2 \end{pmatrix}, T \begin{pmatrix} 0 \\ -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ -2 \\ 2 \end{pmatrix}, \text{ and } T \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 8 \\ 8 \\ 8 \end{pmatrix},$$

which respectively have coordinate vectors

$$\begin{pmatrix} 2 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 2 \\ 0 \end{pmatrix}, \text{ and } \begin{pmatrix} 0 \\ 0 \\ 8 \end{pmatrix}$$

relative to our basis. Thus the matrix of T relative to this basis is

$$\begin{pmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 8 \end{pmatrix}.$$

We can now see that, geometrically, T scales by a factor of 2 in the direction of $\begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}$, it scales by a factor of 2 in the direction of $\begin{pmatrix} 0 \\ -1 \\ 1 \end{pmatrix}$, and it scales by a factor of 8 in the direction of $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$.

Back to Example 1. Let $A = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix}$ be the matrix of Example 1; we want to now compute A^{100} . It will be of no use to try to multiply A by itself 100 times, or to multiply it by itself enough times until we notice some kind of pattern. We need a better way to do this.

The key comes from the following equation:

$$\begin{pmatrix} 13 & 6 \\ -1 & 12 \end{pmatrix} = \begin{pmatrix} 2 & -3 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 10 & 0 \\ 0 & 15 \end{pmatrix} \begin{pmatrix} 2 & -3 \\ 1 & 1 \end{pmatrix}^{-1}.$$

You can certainly multiply out the right hand side to see why this is true, but I claim that we know it has to be true without doing any further computation by thinking about each of these matrices are supposed to represent. Recall that

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

is the change of basis matrix from the basis $\mathcal{B} = \left\{ \begin{pmatrix} 2 \\ 1 \end{pmatrix}, \begin{pmatrix} -3 \\ 1 \end{pmatrix} \right\}$ of \mathbb{R}^2 to the standard basis and that

$$B = \begin{pmatrix} 10 & 0 \\ 0 & 15 \end{pmatrix}$$

is the matrix of the transformation $T(\vec{x}) = A\vec{x}$ relative to \mathcal{B} . Take a vector \vec{x} and consider

$$SBS^{-1}\vec{x}.$$

First, S^{-1} is the change of basis matrix from the standard basis to \mathcal{B} , so

$$S^{-1}\vec{x}$$

gives the coordinate vector of \vec{x} relative to \mathcal{B} . Now, since B tells us what T does relative to our new basis, multiplying $S^{-1}\vec{x}$ by B gives us the coordinate vector of $T(\vec{x})$ relative to \mathcal{B} . Finally, multiplying by S takes this coordinate vector and expresses it back in terms of coordinates relative to the standard basis. The end result is that

$$SBS^{-1}\vec{x}$$

gives $T(\vec{x})$ expressed in terms of the standard basis, but this is precisely what $A\vec{x}$ is supposed to be! In other words, the transformation corresponding to the product SBS^{-1} does the same thing to \vec{x} as does A , so

$$A = SBS^{-1}$$

as claimed.

Definition of similar matrices. Two matrices A and B are said to be *similar* if there is an invertible matrix S satisfying $A = SBS^{-1}$.

Important. Similar matrices represent the “same” linear transformation only with respect to different bases. In particular, for an $n \times n$ matrix A and a basis \mathcal{B} of \mathbb{R}^n , A is similar to the matrix B of the transformation $T(\vec{x}) = A\vec{x}$ relative to \mathcal{B} ; i.e.

$$A = SBS^{-1}$$

where S is the change of basis matrix from \mathcal{B} to the standard basis.

Finishing up our last computation. Going back to where we previously left off, we now know that

$$A = SBS^{-1}.$$

Note that this gives

$$A^2 = AA = (SBS^{-1})(SBS^{-1}) = SB^2S^{-1}$$

since the $S^{-1}S$ in the middle is the identity. Similarly, $A^3 = SB^3S^{-1}$ and in general

$$A^k = SB^kS^{-1}.$$

Thus,

$$\begin{pmatrix} 13 & 6 \\ -1 & 12 \end{pmatrix}^{100} = \begin{pmatrix} 2 & -3 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 10 & 0 \\ 0 & 15 \end{pmatrix}^{100} \begin{pmatrix} 2 & -3 \\ 1 & 1 \end{pmatrix}^{-1},$$

which is simple to compute now since powers of a diagonal matrix are easy to compute:

$$\begin{pmatrix} 10 & 0 \\ 0 & 15 \end{pmatrix}^{100} = \begin{pmatrix} 10^{100} & 0 \\ 0 & 15^{100} \end{pmatrix},$$

so

$$\begin{aligned} A^{100} &= \begin{pmatrix} 2 & -3 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 10^{100} & 0 \\ 0 & 15^{100} \end{pmatrix} \begin{pmatrix} 2 & -3 \\ 1 & 1 \end{pmatrix}^{-1} \\ &= \begin{pmatrix} 2 \cdot 10^{100} & -3 \cdot 15^{100} \\ 10^{100} & 15^{100} \end{pmatrix} \frac{1}{5} \begin{pmatrix} 1 & 3 \\ -1 & 2 \end{pmatrix} \\ &= \frac{1}{5} \begin{pmatrix} 2 \cdot 10^{100} + 3 \cdot 15^{100} & 6 \cdot 10^{100} - 6 \cdot 15^{100} \\ 10^{100} - 15^{100} & 3 \cdot 10^{100} + 2 \cdot 15^{100} \end{pmatrix}. \end{aligned}$$

as desired. Note how useful it was to compute the matrix of the transformation corresponding to A in terms of another well-chosen basis!

Lecture 16: Determinants

Today we started talking about determinants, which we will continue with all this week. For now all we are interested in is computing determinants; we will talk about just what exactly determinants mean over the next few lectures.

Warm-Up 1. Suppose that A is the 2×2 matrix of a rotation by an angle $0 < \theta < 180$. We claim that A is not similar to a diagonal matrix. If it was, there would be a basis $\mathcal{B} = \{\vec{v}_1, \vec{v}_2\}$ of \mathbb{R}^2 relative to which the matrix of $T(\vec{x}) = A\vec{x}$ had the form

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

Now, let us think about what it would take to get a matrix of this form. This says that the coordinate vectors of $T(\vec{v}_1)$ and $T(\vec{v}_2)$ relative to \mathcal{B} would respectively be $\begin{pmatrix} a \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ b \end{pmatrix}$. So, in order to write $T(\vec{v}_1)$ as a linear combination of the basis vectors in \mathcal{B} we would use

$$T(\vec{v}_1) = a\vec{v}_1 + 0\vec{v}_2 = a\vec{v}_1.$$

Thus the first basis vector \vec{v}_1 would need to have the property that applying T to it gave a multiple of it; however, rotating a nonzero vector by an angle strictly between 0 and 180 can never produce a multiple of that vector! Similarly, the second basis vector would need to satisfy

$$T(\vec{v}_2) = b\vec{v}_2,$$

and again such an equation cannot possibly hold for the type of rotation we are considering.

This shows that there will never be a basis of \mathbb{R}^2 relative to which the matrix of T has the form

$$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix},$$

so A is not similar to a diagonal matrix. Again, the key point is in realizing what having a diagonal matrix as the matrix of T relative to a basis would mean about what happens when we apply T to those basis vectors.

Warm-Up 2. We claim that any 2×2 reflection matrix A is similar to $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. This would require a basis $\{\vec{v}_1, \vec{v}_2\}$ of \mathbb{R}^2 relative to which the matrix of $T(\vec{x}) = A\vec{x}$ is $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. Again by considering the columns of this matrix as the coordinate vectors of $T(\vec{v}_1)$ and $T(\vec{v}_2)$, this means our basis must satisfy

$$T(\vec{v}_1) = \vec{v}_1 \text{ and } T(\vec{v}_2) = -\vec{v}_2.$$

In this case we can always find such a basis: take \vec{v}_1 to be any nonzero vector on the line we are reflecting across and \vec{v}_2 to be any nonzero vector perpendicular to the line we are reflecting across. Such vectors indeed satisfy $T(\vec{v}_1) = \vec{v}_1$ and $T(\vec{v}_2) = -\vec{v}_2$, so the matrix of T relative to the basis $\{\vec{v}_1, \vec{v}_2\}$ is $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ as desired. The upshot is that all reflections “look similar” since they all “look” like the matrix $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ after picking the right basis. This justifies the use of the term “similar” to describe such matrices.

We can also say that any 2×2 reflection matrix is also similar to $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$. Indeed, take the basis $\{\vec{v}_1, \vec{v}_2\}$ from before but switch their order and consider the basis $\{\vec{v}_2, \vec{v}_1\}$. Doing so has the effect of switching the columns in the corresponding matrix of T .

Warm-Up 3. Finally, we find the matrix A of the orthogonal projection T of \mathbb{R}^3 onto the line $\text{span}\left\{\begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}\right\}$ relative to the standard basis. Using previous techniques we would have to compute

$$T \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, T \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \text{ and } T \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix},$$

which give the columns of the matrix we want. This is not so hard in this case, but could be more complicated if we were projecting onto a different line. Here’s another way to answer this which is other situations is much simpler.

We know that the matrix A we’re looking for will be similar to the matrix B of T relative to any basis of \mathbb{R}^3 :

$$A = SBS^{-1}$$

where S is the change of basis matrix. Thus we can find A by computing SBS^{-1} . In order to make this worthwhile, we should find a basis of \mathbb{R}^3 relative to which B will be simpler; thinking about orthogonal projections geometrically we can always find such a basis: take \vec{v}_1 to be a nonzero

vector on the line we're projecting onto and \vec{v}_2 and \vec{v}_3 to be nonzero vectors perpendicular to that line. For instance,

$$\vec{v}_1 = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}, \quad \vec{v}_2 = \begin{pmatrix} 0 \\ -1 \\ 1 \end{pmatrix}, \quad \text{and} \quad \vec{v}_3 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$$

work. For these vectors, we have

$$T(\vec{v}_1) = \vec{v}_1, \quad T(\vec{v}_2) = \vec{0}, \quad \text{and} \quad T(\vec{v}_3) = \vec{0}$$

since projecting a vector already on a line onto that line leaves that vector alone and projecting a vector perpendicular to a line onto that line gives the zero vector. The coordinates of the above vectors relative to the basis $\{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ of \mathbb{R}^3 are respectively

$$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

With the change of basis matrix S given by

$$S = \begin{pmatrix} 0 & 0 & 1 \\ 1 & -1 & 0 \\ 1 & 1 & 0 \end{pmatrix},$$

we then have

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

Using our technique for computing inverses we find that

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

Thus

$$\begin{aligned} A &= \begin{pmatrix} 0 & 0 & 1 \\ 1 & -1 & 0 \\ 1 & 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 0 & 1 \\ 1 & -1 & 0 \\ 1 & 1 & 0 \end{pmatrix}^{-1} \\ &= \begin{pmatrix} 0 & 0 & 1 \\ 1 & -1 & 0 \\ 1 & 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 1/2 & 1/2 \\ 0 & -1/2 & 1/2 \\ 1 & 0 & 0 \end{pmatrix} \\ &= \begin{pmatrix} 0 & 0 & 1 \\ 1 & -1 & 0 \\ 1 & 1 & 0 \end{pmatrix} \begin{pmatrix} 0 & 1/2 & 1/2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \\ &= \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1/2 & 1/2 \\ 0 & 1/2 & 1/2 \end{pmatrix} \end{aligned}$$

is the matrix of T relative to the standard basis. Again, this would not have been hard to find using earlier methods, but this new method might be easier to apply in other situations, say when projecting on the line spanned by $\begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$ for instance.

Determinants. The *determinant* of a (square) matrix is a certain number we compute from that matrix. The determinant of A is denoted by $|A|$ or by $\det A$. This one number will encode much information about A : in particular, it will determine completely whether or not A is invertible. More importantly, it has an important geometric interpretation, which we will come to over the next few lectures.

To start with, the determinant of a 2×2 matrix is a number we've seen before in the formula for the inverse of a 2×2 matrix:

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

In this case, it is true that a 2×2 matrix is invertible if and only if its determinant is nonzero; this will be true in general.

Remark on Section 6.1. The definition of determinants and method for computing them giving in Section 6.1 is ridiculous: technically it is correct, but makes the idea of a determinant way too complicated. In particular, ignore anything having to do with “patterns” and “inversions”. Instead, you should use the method of *cofactor* or *Laplace expansion* described towards the end of Section 6.2 when computing determinants. The rest of 6.1 contains useful facts that we'll come to, but seriously, I have no idea why the author chose to define determinants in the way he does.

Example 1. We illustrate the method of doing a *cofactor expansion* along the first row when computing the following determinant:

$$\begin{vmatrix} 1 & 2 & 3 \\ 4 & 5 & -1 \\ -2 & 3 & 1 \end{vmatrix}.$$

To expand along the first row means we take each entry of the first row and multiply it by the determinant of the matrix leftover when we cross out the row and column that entry is in. So, for the entry 1, crossing out the row and column it is in (first row and first column in this case) leaves us with $\begin{pmatrix} 5 & -1 \\ 3 & 1 \end{pmatrix}$. So we take

$$1 \begin{vmatrix} 5 & -1 \\ 3 & 1 \end{vmatrix}$$

as part of our cofactor expansion. We do the same with the 2 and 3 in the first row, giving the terms

$$2 \begin{vmatrix} 4 & -1 \\ -2 & 1 \end{vmatrix} \text{ and } 3 \begin{vmatrix} 4 & 5 \\ -2 & 3 \end{vmatrix}$$

in our expansion. The last thing to determine is what to do with the terms we've found: starting with a + sign associated to the upper-left corner entry of our matrix, we alternate between assigning +'s and -'s to all other entries, so in the 3×3 case we would have

$$\begin{pmatrix} + & - & + \\ - & + & - \\ + & - & + \end{pmatrix}.$$

These signs tell us what to do with the corresponding terms in the cofactor expansion, so our cofactor expansion along the first row becomes

$$1 \begin{vmatrix} 5 & -1 \\ 3 & 1 \end{vmatrix} - 2 \begin{vmatrix} 4 & -1 \\ -2 & 1 \end{vmatrix} + 3 \begin{vmatrix} 4 & 5 \\ -2 & 3 \end{vmatrix}.$$

We are now down to computing these 2×2 determinants, which we know how to do, so putting it all together we have

$$\begin{aligned} \begin{vmatrix} 1 & 2 & 3 \\ 4 & 5 & -1 \\ -2 & 3 & 1 \end{vmatrix} &= 1 \begin{vmatrix} 5 & -1 \\ 3 & 1 \end{vmatrix} - 2 \begin{vmatrix} 4 & -1 \\ -2 & 1 \end{vmatrix} + 3 \begin{vmatrix} 4 & 5 \\ -2 & 3 \end{vmatrix} \\ &= 1(5 + 3) - 2(4 - 2) + 3(12 + 10) \\ &= 70. \end{aligned}$$

Let us now compute the same determinant, only doing an expansion along the second column. So we look at the terms

$$2 \begin{vmatrix} 4 & -1 \\ -2 & 1 \end{vmatrix}, 5 \begin{vmatrix} 1 & 3 \\ -2 & 1 \end{vmatrix}, \text{ and } 3 \begin{vmatrix} 1 & 3 \\ 4 & -1 \end{vmatrix},$$

which we get the same way as before: moving down the second column and multiplying each entry by the determinant of what's left after crossing out the row and column that entry is in. In this case, the first entry in the second column comes with a $-$ sign, and alternating signs down gives us

$$-2 \begin{vmatrix} 4 & -1 \\ -2 & 1 \end{vmatrix} + 5 \begin{vmatrix} 1 & 3 \\ -2 & 1 \end{vmatrix} - 3 \begin{vmatrix} 1 & 3 \\ 4 & -1 \end{vmatrix}.$$

Thus

$$\begin{aligned} \begin{vmatrix} 1 & 2 & 3 \\ 4 & 5 & -1 \\ -2 & 3 & 1 \end{vmatrix} &= -2 \begin{vmatrix} 4 & -1 \\ -2 & 1 \end{vmatrix} + 5 \begin{vmatrix} 1 & 3 \\ -2 & 1 \end{vmatrix} - 3 \begin{vmatrix} 1 & 3 \\ 4 & -1 \end{vmatrix} \\ &= -2(4 - 2) + 5(1 + 6) - 3(-1 - 12) \\ &= 70, \end{aligned}$$

agreeing with our answer from when we expanded along the first row.

Important. Performing a cofactor expansion along any row or any column of a matrix will *always* give the same value. Choose the row or column which makes computations as simple as possible, which usually means choose the row or column with the most zeroes.

Example 2. We compute the determinant of

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

using a cofactor expansion along the third row, since this has two zeroes in it. The first and last term we get will automatically be zero, so we only get:

$$\begin{vmatrix} 3 & 4 & -1 & 2 \\ 3 & 0 & 1 & 5 \\ 0 & -2 & 1 & 0 \\ -1 & -3 & 2 & 1 \end{vmatrix} = -(-2) \begin{vmatrix} 3 & -1 & 2 \\ 3 & 1 & 5 \\ -1 & 2 & 1 \end{vmatrix} + 1 \begin{vmatrix} 3 & 4 & 2 \\ 3 & 0 & 5 \\ -1 & -3 & 1 \end{vmatrix}.$$

Note that the signs follow the same pattern as before. Now we must compute each of these 3×3 determinants, and we do so by again using a cofactor expansion on each. Expanding along the first row in the first and second row in the second, we get

$$\begin{aligned} \begin{vmatrix} 3 & -1 & 2 \\ 3 & 1 & 5 \\ -1 & 2 & 1 \end{vmatrix} &= 3 \begin{vmatrix} 1 & 5 \\ 2 & 1 \end{vmatrix} - (-1) \begin{vmatrix} 3 & 5 \\ -1 & 1 \end{vmatrix} + 2 \begin{vmatrix} 3 & 1 \\ -1 & 2 \end{vmatrix} \\ &= 3(-9) + 1(8) + 2(7) \\ &= -5 \end{aligned}$$

and

$$\begin{aligned} \begin{vmatrix} 3 & 4 & 2 \\ 3 & 0 & 5 \\ -1 & -3 & 1 \end{vmatrix} &= -3 \begin{vmatrix} 4 & 2 \\ -3 & 1 \end{vmatrix} - 5 \begin{vmatrix} 3 & 4 \\ -1 & -3 \end{vmatrix} \\ &= -3(10) - 5(-5) \\ &= -5. \end{aligned}$$

Putting it all together gives

$$\begin{vmatrix} 3 & 4 & -1 & 2 \\ 3 & 0 & 1 & 5 \\ 0 & -2 & 1 & 0 \\ -1 & 3 & 2 & 1 \end{vmatrix} = -(-2) \begin{vmatrix} 3 & -1 & 2 \\ 3 & 1 & 5 \\ -1 & 2 & 1 \end{vmatrix} + 1 \begin{vmatrix} 3 & 4 & 2 \\ 3 & 0 & 5 \\ -1 & 3 & 1 \end{vmatrix} = 2(-5) + 1(-5) = -15.$$

Amazingly Awesome Theorem, continued. A square matrix is invertible if and only if its determinant is not zero. (We will come back to why this is true later.)

Formula for inverses. As in the 2×2 case, there is actually a concrete formula for the inverse of any square matrix. This formula looks like

$$A^{-1} = \frac{1}{\det A} \begin{pmatrix} \text{something pretty} \\ \text{complicated} \end{pmatrix},$$

which helps to explain why we need $\det A \neq 0$ in order for A to be invertible. The fact that such a concrete formula exists is nice for certain theoretical reasons, but is not very practical due to the complicated nature of the matrix involved. Indeed, even for 3×3 matrices it will be faster to compute inverses using the technique we've previously described. So, we won't say much more about this explicit formula.

Lecture 17: Properties of Determinants

Today we continued talking about determinants, focusing on some of their important properties. In particular, the way in which determinants behave under row operations can give a pretty useful way to compute determinants in many cases.

Warm-Up 1. We find all values of λ such that

$$\begin{pmatrix} \lambda - 1 & -3 & -3 \\ 3 & \lambda + 5 & 3 \\ -3 & -3 & \lambda - 1 \end{pmatrix}$$

is not invertible. Trying to do this using previous techniques we would have to perform row operations until we can determine what the rank of this matrix will be. The trouble is that with all that λ 's floating around, these row operations will get a little tedious. Instead, we can simply figure out when this matrix will have zero determinant.

Doing a cofactor expansion along the first row, we have:

$$\begin{aligned} \begin{vmatrix} \lambda - 1 & -3 & -3 \\ 3 & \lambda + 5 & 3 \\ -3 & -3 & \lambda - 1 \end{vmatrix} &= (\lambda - 1) \begin{vmatrix} \lambda + 5 & 3 \\ -3 & \lambda - 1 \end{vmatrix} - (-3) \begin{vmatrix} 3 & 3 \\ -3 & \lambda - 1 \end{vmatrix} + (-3) \begin{vmatrix} 3 & \lambda + 5 \\ -3 & -3 \end{vmatrix} \\ &= (\lambda - 1)[(\lambda + 5)(\lambda - 1) + 9] + 3[3(\lambda - 1) + 9] - 3[-9 + 3(\lambda + 5)] \\ &= (\lambda - 1)(\lambda^2 + 4\lambda + 4) + 3(3\lambda + 6) - 3(3\lambda + 6) \\ &= (\lambda - 1)(\lambda + 2)^2. \end{aligned}$$

This, the given matrix has determinant equal to 0 only for $\lambda = 1$ and $\lambda = -2$, so these are the only two values of λ for which the matrix is not invertible.

Warm-Up 2. Recall that the transpose of a matrix A is the matrix A^T obtained by turning the rows of A into the columns of A^T . We claim that for any square matrix A , $\det A^T = \det A$. This is actually quite simple: doing a cofactor expansion along row i of A^T is the same as doing a cofactor expansion along column i of A , so both expansions will give the same value.

Determinants are linear in columns and rows. To say that determinants are linear in the columns of a matrix means the following; to simplify matters, let's focus on a 3×3 matrix, but the general case is similar. Suppose that the second column of a 2×2 matrix is written as the sum of two vectors \vec{a} and \vec{b} :

$$\begin{pmatrix} | & | & | \\ \vec{v}_1 & \vec{a} + \vec{b} & \vec{v}_3 \\ | & | & | \end{pmatrix}.$$

Then

$$\det \begin{pmatrix} | & | & | \\ \vec{v}_1 & \vec{a} + \vec{b} & \vec{v}_3 \\ | & | & | \end{pmatrix} = \det \begin{pmatrix} | & | & | \\ \vec{v}_1 & \vec{a} & \vec{v}_3 \\ | & | & | \end{pmatrix} + \det \begin{pmatrix} | & | & | \\ \vec{v}_1 & \vec{b} & \vec{v}_3 \\ | & | & | \end{pmatrix}$$

so that the determinant “breaks up” when splitting the column $\vec{a} + \vec{b}$ into two pieces. What does this have to do with linearity? We can define a linear transformation T from \mathbb{R}^3 to \mathbb{R} by setting

$$T(\vec{x}) = \det \begin{pmatrix} | & | & | \\ \vec{v}_1 & \vec{x} & \vec{v}_3 \\ | & | & | \end{pmatrix}.$$

Then this above property says that $T(\vec{a} + \vec{b}) = T(\vec{a}) + T(\vec{b})$, which is the first property required in order to say that T is a linear transformation. The second property, $T(c\vec{a}) = cT(\vec{a})$, is the second linearity property of determinants:

$$\det \begin{pmatrix} | & | & | \\ \vec{v}_1 & c\vec{a} & \vec{v}_3 \\ | & | & | \end{pmatrix} = c \det \begin{pmatrix} | & | & | \\ \vec{v}_1 & \vec{a} & \vec{v}_3 \\ | & | & | \end{pmatrix},$$

which says that scalars “pull out” when multiplied by a single column. Note that if two columns were multiplied by c , then c would “pull out” twice and we would get c^2 in front.

The same is true no matter which column is written as the sum of two vectors or no matter which column is scaled by a number, and the same is true if we do this all with rows instead.

Example 1. We find the matrix of the linear transformation T from \mathbb{R}^2 to \mathbb{R} defined by

$$T \begin{pmatrix} a \\ b \end{pmatrix} = \det \begin{pmatrix} 3 & a \\ 4 & b \end{pmatrix}.$$

We have

$$T \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \det \begin{pmatrix} 3 & 1 \\ 4 & 0 \end{pmatrix} = -4 \text{ and } T \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \det \begin{pmatrix} 3 & 0 \\ 4 & 1 \end{pmatrix} = 3,$$

so the matrix of T is $A = (-4 \ 3)$. Indeed, let's check that

$$A \begin{pmatrix} a \\ b \end{pmatrix} = (-4 \ 3) \begin{pmatrix} a \\ b \end{pmatrix} = -4a + 3b, \text{ which is the same as } T \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} 3 & a \\ 4 & b \end{pmatrix} = 3b - 4a.$$

Row (and column) operations and determinants. Determinants behave in pretty simple ways when performing row operations:

- swapping rows changes the sign of a determinant,
- scaling a row by a nonzero number multiplies a determinant by that same number, and
- adding a multiple of one row to another does nothing to a determinant.

This last property is the one which makes these observations actually useful, and gives us a new way to compute determinants. Note that in this last property that it is crucial that the row we are replacing is not the row we are scaling; if instead we had scaled the row we replaced the determinant *would* change, it would be scaled by that same amount.

Example 2. We compute the determinant of the matrix

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

using row operations. We did this last time using a cofactor expansion, and it was a little tedious. Row operations give us a bit of a smoother computation. We will reduce A , keeping track of how the operations we do at each step affect the determinant.

First, we swap the first and fourth rows to get the -1 in the upper left corner. Note that if instead we used something like $3IV + I \rightarrow IV$, this does change the determinant we're after since we scaled the row we replaced; this is why we're swapping rows first, so that the row additions we do later do not affect the determinant. After the row swap, the determinant of the resulting matrix is the negative of the one before:

$$\begin{pmatrix} 3 & 4 & -1 & 2 \\ 3 & 0 & 1 & 5 \\ 0 & -2 & 1 & 0 \\ -1 & -3 & 2 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} -1 & -3 & 2 & 1 \\ 3 & 0 & 1 & 5 \\ 0 & -2 & 1 & 0 \\ 3 & 4 & -1 & 2 \end{pmatrix}, \det A \rightarrow -\det A.$$

Now we do $3I + II \rightarrow II$ and $3I + IV \rightarrow IV$, neither of which affect the determinant:

$$\begin{pmatrix} -1 & -3 & 2 & 1 \\ 3 & 0 & 1 & 5 \\ 0 & -2 & 1 & 0 \\ 3 & 4 & -1 & 2 \end{pmatrix} \rightarrow \begin{pmatrix} -1 & -3 & 2 & 1 \\ 0 & -9 & 7 & 8 \\ 0 & -2 & 1 & 0 \\ 0 & -5 & 5 & 5 \end{pmatrix}, \quad -\det A \rightarrow -\det A.$$

Next let's multiply the the last row by $\frac{1}{5}$, which scales the determinant by the same amount:

$$\begin{pmatrix} -1 & -3 & 2 & 1 \\ 0 & -9 & 7 & 8 \\ 0 & -2 & 1 & 0 \\ 0 & -5 & 5 & 5 \end{pmatrix} \rightarrow \begin{pmatrix} -1 & -3 & 2 & 1 \\ 0 & -9 & 7 & 8 \\ 0 & -2 & 1 & 0 \\ 0 & -1 & 1 & 1 \end{pmatrix}, \quad -\det A \rightarrow -\frac{1}{5}\det A.$$

We swap the second and fourth rows to get the -1 where the -9 is:

$$\begin{pmatrix} -1 & -3 & 2 & 1 \\ 0 & -9 & 7 & 8 \\ 0 & -2 & 1 & 0 \\ 0 & -1 & 1 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} -1 & -3 & 2 & 1 \\ 0 & -1 & 1 & 1 \\ 0 & -2 & 1 & 0 \\ 0 & -9 & 7 & 8 \end{pmatrix}, \quad -\frac{1}{5}\det A \rightarrow \frac{1}{5}\det A.$$

The row operations $-2II + III \rightarrow III$ and $-9II + IV \rightarrow IV$ do not affect the determinant:

$$\begin{pmatrix} -1 & -3 & 2 & 1 \\ 0 & -1 & 1 & 1 \\ 0 & -2 & 1 & 0 \\ 0 & -9 & 7 & 8 \end{pmatrix} \rightarrow \begin{pmatrix} -1 & -3 & 2 & 1 \\ 0 & -1 & 1 & 1 \\ 0 & 0 & -1 & -2 \\ 0 & 0 & -2 & -1 \end{pmatrix}, \quad \frac{1}{5}\det A \rightarrow \frac{1}{5}\det A.$$

Finally, we do $-2III + IV \rightarrow IV$, which again does not affect the determinant:

$$\begin{pmatrix} -1 & -3 & 2 & 1 \\ 0 & -1 & 1 & 1 \\ 0 & 0 & -1 & -2 \\ 0 & 0 & -2 & -1 \end{pmatrix} \rightarrow \begin{pmatrix} -1 & -3 & 2 & 1 \\ 0 & -1 & 1 & 1 \\ 0 & 0 & -1 & -2 \\ 0 & 0 & 0 & 3 \end{pmatrix}, \quad \frac{1}{5}\det A \rightarrow \frac{1}{5}\det A.$$

Now, the whole point is that the determinant of this final matrix is super-easy to compute: this matrix is *upper-triangular*, and the determinant of such a matrix is simply the product of its diagonal entries:

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

But also, we found above that the determinant of this final matrix is $\frac{1}{5}\det A$, so we get

$$\frac{1}{5}\det A = -3.$$

Thus $\det A = -15$, as we computed using a cofactor expansion last time. So, now we have a new method for computing determinants. Either this way or using a cofactor expansion will always work. Actually, once you get used to using row operations, this method will almost always be faster, but feel free to compute determinants in whatever way you'd like.

Important. To emphasize: row swaps multiply determinants by -1 , scaling rows multiplies determinants by the same amount, and adding multiples of rows to other rows does not change determinants as long as long as the rows replaced are not the rows which were scaled.

Justifying some of these properties. Again, let's just look at the 3×3 case. Why is it that swapping rows changes the sign of a determinant? For instance, why is

$$\det \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix} = -\det \begin{pmatrix} d & e & f \\ a & b & c \\ g & h & i \end{pmatrix}?$$

A cofactor expansion along the first row of the first matrix looks like

$$a |\text{something}| - b |\text{something}| + c |\text{something}|.$$

Notice that you get the same type of expression when taking a cofactor expansion of the second matrix along the second row, except that all the signs change:

$$-a |\text{something}| + b |\text{something}| - c |\text{something}|.$$

The 2×2 determinants in this and the previous expression are exactly the same, so this last expression is negative the first one, justifying the fact that swapping the first two rows changes the sign the determinant.

Notice that if we swapped the first and third row instead, at first glance it seems that doing a cofactor expansion along the first row in the original matrix and a cofactor expansion along the third row of the result give the same expression since both look like

$$a |\text{something}| - b |\text{something}| + c |\text{something}|.$$

However, if you look at the 2×2 determinants you get in this case it turns out that this is where the extra negative signs show up, so that this type of row swap still changes the sign of the determinant.

As for the third type of row operation, consider something like

$$\begin{pmatrix} - & \vec{r}_1 & - \\ - & \vec{r}_2 & - \\ - & \vec{r}_3 & - \end{pmatrix} \text{ versus } \begin{pmatrix} - & \vec{r}_1 & - \\ - & 3\vec{r}_1 + \vec{r}_2 & - \\ - & \vec{r}_3 & - \end{pmatrix},$$

where the second matrix is the result of doing the row operation $3I + II \rightarrow II$ on the first. The linearity property (in the second row in this case) of determinants tells us that:

$$\det \begin{pmatrix} - & \vec{r}_1 & - \\ - & 3\vec{r}_1 + \vec{r}_2 & - \\ - & \vec{r}_3 & - \end{pmatrix} = 3 \det \begin{pmatrix} - & \vec{r}_1 & - \\ - & \vec{r}_1 & - \\ - & \vec{r}_3 & - \end{pmatrix} + \det \begin{pmatrix} - & \vec{r}_1 & - \\ - & \vec{r}_2 & - \\ - & \vec{r}_3 & - \end{pmatrix}.$$

The first matrix on the right is not invertible since it has linearly dependent rows, so its determinant is zero. Hence we're left with

$$\det \begin{pmatrix} - & \vec{r}_1 & - \\ - & 3\vec{r}_1 + \vec{r}_2 & - \\ - & \vec{r}_3 & - \end{pmatrix} = \det \begin{pmatrix} - & \vec{r}_1 & - \\ - & \vec{r}_2 & - \\ - & \vec{r}_3 & - \end{pmatrix},$$

saying that this type of row operation does not change determinants.

Invertibility and nonzero determinants. Now we can justify the fact that a matrix is invertible if and only if its determinant is not zero. Let A be a square matrix. There is some sequence of row operations getting us from A to its reduced echelon form:

$$A \rightarrow \cdots \rightarrow \text{rref}(A).$$

Now, each row operation either changes the sign of the determinant, multiplies it by some nonzero scalar, or does nothing, so the determinant of the final matrix ends up being related to the original determinant by something like:

$$\det(\text{rref } A) = k_n \cdots k_2 k_1 (-1)^m \det A,$$

where m is the number of row swaps we do and k_1, \dots, k_n are the numbers we scale rows by throughout. None of these are zero, so

$$\det A = 0 \text{ if and only if } \det(\text{rref } A) = 0.$$

But $\det(\text{rref } A) = 0$ only if some diagonal entry of $\text{rref } A$ is zero, which happens if and only if $\text{rref } A$ is not the identity, in which case A is not invertible. Thus A is invertible if and only if $\det A \neq 0$.

Lecture 18: Geometric Interpretation of Determinants

Today we spoke about the long-awaited geometric interpretation of determinants, and used it to justify some properties of determinants we've already seen. The key is the interpretation of a determinant as an "expansion factor".

Warm-Up 1. We compute the determinant of

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

using row operations. First we multiply the last row by 2, which multiplies the determinant by the same amount:

$$\begin{pmatrix} 2 & 1 & -1 \\ 2 & 0 & 3 \\ 3 & 1 & 2 \end{pmatrix} \rightarrow \begin{pmatrix} 2 & 1 & -1 \\ 2 & 0 & 3 \\ 6 & 2 & 4 \end{pmatrix}, \det A \rightarrow 2 \det A.$$

Taking $-I + II \rightarrow II$ and $-3I + III \rightarrow III$ does not affect the determinant:

$$\begin{pmatrix} 2 & 1 & -1 \\ 2 & 0 & 3 \\ 6 & 2 & 4 \end{pmatrix} \rightarrow \begin{pmatrix} 2 & 1 & -1 \\ 0 & -1 & 4 \\ 0 & -1 & 7 \end{pmatrix}, 2 \det A \rightarrow 2 \det A.$$

Finally, $-II + III \rightarrow III$ does not affect the determinant either:

$$\begin{pmatrix} 2 & 1 & -1 \\ 0 & -1 & 4 \\ 0 & -1 & 7 \end{pmatrix} \rightarrow \begin{pmatrix} 2 & 1 & -1 \\ 0 & -1 & 4 \\ 0 & 0 & 3 \end{pmatrix}, 2 \det A \rightarrow 2 \det A.$$

The determinant of this final matrix is $2(-1)3 = -6$, which should also equal $2 \det A$, so $\det A = -3$.

Remark. Determinants also have the property that $\det(AB) = (\det A)(\det B)$, which we will justify soon. For now, let's use this to prove some more basic formulas. First, if A is invertible, then

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

which follows by taking the determinant of both sides of

$$AA^{-1} = I$$

and using the fact that $\det(AA^{-1}) = (\det A)(\det A^{-1})$. (Note that the fraction $1/\det A$ is defined since $\det A \neq 0$ for an invertible matrix.)

Second, if A and B are similar, so that $A = SBS^{-1}$ for some invertible matrix S , we have

$$\det A = \det(SBS^{-1}) = (\det S)(\det B)(\det S^{-1}) = \det B$$

since $\det S$ and $\det S^{-1}$ cancel out according to the previous fact. Hence similar matrices always have the same determinant.

Warm-Up 2. Say that A is a 4×4 matrix satisfying $A^3 = A^5$. We claim that $\det A$ must be 0 or ± 1 . Indeed, taking determinants of both sides of $A^3 = A^5$ and using the property from the previous remark gives

$$(\det A)^3 = (\det A)^5,$$

and 0, -1 , and 1 are the only numbers satisfying this equality.

To be complete, we give examples showing that each of these determinants is possible. First, the zero matrix A satisfies $A^3 = A^5$ and $\det A = 0$. Second, the identity matrix satisfies $I^3 = I^5$ and $\det I = 1$. For the final possibility, the matrix

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

satisfies $B^3 = B^5$ and $\det B = -1$. Note that negative the identity matrix doesn't work for this last example because although $(-I)^3 = (-I)^5$, the 4×4 identity has determinant $+1$.


Determinants as areas and volumes. In the 2×2 case, the absolute value of

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

is the *area* of the parallelogram with sides formed by $\begin{pmatrix} a \\ c \end{pmatrix}$ and $\begin{pmatrix} b \\ d \end{pmatrix}$. Similarly, in the 3×3 case, $|\det A|$ is the *volume* of the parallelepiped with sides formed by the columns of A .

Note that in both of these case it is the absolute value $|\det A|$ and not just $\det A$ itself which is interpreted as an area or volume. It makes sense that this should be true: $\det A$ itself could be negative, but areas and volumes cannot be negative.

Expansion factors. Say that T is a linear transformation from \mathbb{R}^2 to \mathbb{R}^2 given by $T(\vec{x}) = A\vec{x}$ and take some region D in the xy -plane. After applying T to the points of D we obtain a region $T(D)$ of the xy -plane called the *image* of D under T :



We want to compare the area of $T(D)$ with that of D , and $|\det A|$ is precisely what allows us to do this: the area of $T(D)$ is $|\det A|$ times the area of D . So, $|\det A|$ is the factor by which areas are altered by after applying T . Note that areas are indeed expanded when $|\det A| > 1$ but are actually “contracted” (or shrunk) when $|\det A| < 1$. Regardless, we will always refer to $|\det A|$ as an *expansion* factor.


Important. For $T(\vec{x}) = A\vec{x}$ and a region D , we have

$$\text{area or volume of } T(D) = |\det A|(\text{area or volume of } D)$$


where we use “area” in the 2-dimensional setting and “volume” in the 3-dimensional setting. This is the crucial geometric interpretation of determinants.

The sign of the determinant. Before looking at examples, if the absolute value of a determinant is giving us the “expansion factor” for the corresponding transformation, it is natural to wonder what the *sign* of a determinant tells us. The sign of a determinant also has a nice geometric interpretation in terms of what’s called an “orientation”, and is relatively simple to state in the 2×2 case.


Suppose we have two vectors \vec{v}_1 and \vec{v}_2 where \vec{v}_2 occurs “counterclockwise” to the left of \vec{v}_1 , meaning you have to move counterclockwise to get from \vec{v}_1 to \vec{v}_2 ; for instance,



are both examples of when this happens. After we apply a transformation $T(\vec{x}) = A\vec{x}$ we get two new vectors $T(\vec{v}_1)$ and $T(\vec{v}_2)$, and we can ask whether $T(\vec{v}_2)$ occurs “counterclockwise” to the left of $T(\vec{v}_1)$. The matrix A has positive determinant when this is true, and negative determinant when it isn’t. So, something like



would correspond to a matrix with $\det A > 0$ while something like



corresponds to a matrix with $\det A < 0$. The technical explanation is that matrices with positive determinant *preserve* orientation while those with negative determinant *reverse* orientation. A similar explanation works for 3×3 matrices, although it gets a little trickier to talk about what “orientation” means in higher dimensions; we’ll come back to this later when we do “vector calculus” in the spring.

Remark. Recall that similar matrices have the same determinant. Now this makes sense: similar matrices represent the same linear transformation only with respect to different bases, and thus the expansion factor for each should be the same and so should the property of orientation preserving or reversing.

Example 1. Consider the transformation T given by

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

Taking D to be the unit square with sides the standard basis vectors \vec{e}_1 and \vec{e}_2 , we have

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

The image $T(D)$ of the unit square is then the parallelogram with sides $\begin{pmatrix} a \\ c \end{pmatrix}$ and $\begin{pmatrix} b \\ d \end{pmatrix}$, so the area of this parallelogram is

$$\text{area } T(D) = |\det A|(\text{area } D) = |\det A|.$$

Hence $|\det A|$ is the area of the parallelogram with sides formed by the columns of A , recovering the first geometric interpretation of determinants we gave above.

Example 2. For a transformation T which is either a rotation or a reflection, we have

$$\text{area } T(D) = \text{area } D$$

since rotations and reflections preserve lengths and angles. Thus the expansion factor for a rotation or a reflection is 1, so the corresponding matrix A has


$$|\det A| = 1, \text{ and thus } \det A = \pm 1.$$

Using the interpretation of the sign of a determinant, rotations have $\det A = 1$ while reflections have $\det A = -1$. Note that if A and B both describe reflections, then


$$\det(AB) = (\det A)(\det B) = (-1)(-1) = 1,$$

meaning that AB actually describes a rotation! There was a homework problem a while back showing this was true in a particular example, where the composition of two reflections turned out to be a rotation.

Example 3. We ask whether there can be a linear transformation $T(\vec{x}) = A\vec{x}$ which sends



and



In the first case the area of the region gets larger so we would need $|\det A| > 1$ while in the second the area gets smaller so we would need $|\det A| < 1$. We cannot have a matrix satisfying both, so there is no such transformation.

Example 4. Suppose that T is a linear transformation from \mathbb{R}^3 to \mathbb{R}^3 with matrix A , and that T sends a cube to a plane. We can ask whether T can then send a sphere of radius 5 to the unit sphere. We have

$$\text{volume of plane} = |\det A|(\text{volume of cube}).$$

But a cube has nonzero volume while a plane has zero volume, so this means that $|\det A| = 0$ and thus $\det A = 0$. Thus T cannot send a sphere of positive volume to another with positive volume.

Note that A is not invertible, which makes sense since T sends a 3-dimensional cube to a 2-dimensional plane, so T “collapses” dimension. This means that $\text{rank } A < 3$, and we see that the geometric interpretation of a determinant as an expansion factor gives us another way to see that matrices of non-full rank (i.e. noninvertible matrices) must “collapse” dimension.

Justifying $\det(AB) = (\det A)(\det B)$. Suppose that A and B are 2×2 matrices and take some region D in \mathbb{R}^2 . Applying the transformation B gives a region $B(D)$ with

$$\text{area of } B(D) = |\det B|(\text{area of } D).$$

Taking the resulting region $B(D)$ and applying the transformation A gives a region $A(B(D))$ with

$$\text{area of } A(B(D)) = |\det A|(\text{area of } B(D)) = |\det A||\det B|(\text{area of } D).$$

Thus the composed transformation has expansion factor $|\det A||\det B|$.

However, the matrix of this composition is equal to the product AB , so the expansion factor is also $|\det(AB)|$ and hence

$$|\det(AB)| = |\det A||\det B|.$$

By considering the cases where each of these determinants are positive or negative, we get that

$$\det(AB) = (\det A)(\det B).$$

For instance, suppose that $\det A$ and $\det B$ are both negative. We want to show that then $(\det A)(\det B)$ is positive, so $\det(AB)$ should be positive. But this makes sense: if the transformation B “reverses” orientation and A “reverses” it right back, the transformation AB will preserve orientation and so will have positive determinant. Thus

$$\det(AB) = (\det A)(\det B)$$

is true when $\det A$ and $\det B$ are both negative, and the other possibilities are similar to check.

Important. For square matrices A and B of the same size, $\det(AB) = (\det A)(\det B)$.

Lecture 19: Eigenvalues

Today we started talking about eigenvalues of eigenvectors, which make up our last topic for the quarter, and probably the most important as well. Eigenvalues and eigenvectors are key to understanding many properties of matrices, which often lead to their diverse applications.

Warm-Up. Suppose that the matrix of a linear transformation $T(\vec{x}) = A\vec{x}$ relative to the basis $\left\{\begin{pmatrix} 1 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \end{pmatrix}\right\}$ of \mathbb{R}^2 is $\begin{pmatrix} 2 & 1 \\ 0 & -3 \end{pmatrix}$. We want to find the area of $T(D)$ where D is the parallelogram with sides $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 1 \\ 5 \end{pmatrix}$. Using the interpretation of $|\det A|$ as an expansion factor, we know that

$$\text{area of } T(D) = |\det A|(\text{area of } D),$$

so our task is to find the two numbers on the right.

First, the matrix A is similar to any matrix which represents the same transformation T relative to a non-standard basis, so A is similar to $\begin{pmatrix} 2 & 1 \\ 0 & -3 \end{pmatrix}$. Last time we saw that similar matrices have the same determinant, so $\det A = -6$ and hence the expansion factor for T is 6. Next, using another geometric interpretation of determinants from last time, we know that the area of the parallelogram D is

$$\left| \det \begin{pmatrix} 1 & 1 \\ 1 & 5 \end{pmatrix} \right| = 4.$$

Thus, putting it all together we have

$$\text{area of } T(D) = |\det A|(\text{area of } D) = 6(4) = 24.$$

Note that the specific basis $\left\{\begin{pmatrix} 1 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \end{pmatrix}\right\}$ relative to which the matrix of T was $\begin{pmatrix} 2 & 1 \\ 0 & -3 \end{pmatrix}$ was irrelevant; the only thing we needed was that A is similar to this matrix.

Motivation for eigenvalues and eigenvectors. Recall a previous example we did when talking about coordinates, where we asked for a geometric interpretation of the linear transformation $T(\vec{x}) = A\vec{x}$ where

$$A = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix}.$$

We saw that if we use the basis vectors $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} -3 \\ 1 \end{pmatrix}$ for \mathbb{R}^2 , the matrix of T relative to this basis is

$$\begin{pmatrix} 10 & 0 \\ 0 & 15 \end{pmatrix},$$

meaning that T scales the axis span $\{\begin{pmatrix} 2 \\ 1 \end{pmatrix}\}$ by a factor of 10 and the axis span $\{\begin{pmatrix} -3 \\ 1 \end{pmatrix}\}$ by a factor of 15. The lingering question is: why is this the “right” basis to use?

The key fact which made everything work is that these basis vectors satisfy

$$A \begin{pmatrix} 2 \\ 1 \end{pmatrix} = 10 \begin{pmatrix} 2 \\ 1 \end{pmatrix} \quad \text{and} \quad A \begin{pmatrix} -3 \\ 1 \end{pmatrix} = 15 \begin{pmatrix} -3 \\ 1 \end{pmatrix}.$$

The first of these equations says that 10 is an *eigenvalue* of A with *eigenvector* $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ and the second says that 15 is an eigenvalue of A with eigenvector $\begin{pmatrix} -3 \\ 1 \end{pmatrix}$. Thus, by finding the eigenvalues and eigenvectors of A is how we could determine the right basis to use above.

Definition of eigenvalues and eigenvectors. Say that A is a square matrix. We say that a scalar λ is an *eigenvalue* of A if there is a nonzero vector \vec{x} satisfying $A\vec{x} = \lambda\vec{x}$; in other words, λ is an eigenvalue of A if

$$A\vec{x} = \lambda\vec{x} \text{ has a nonzero solution for } \vec{x}.$$

For such an eigenvalue λ , we call a nonzero vector \vec{x} satisfying this equation an *eigenvector* of A corresponding to the eigenvalue λ .

Important. Geometrically, the eigenvectors of a matrix A are those nonzero vectors with the property that applying the transformation corresponding to A to them results in a multiple of that vector; i.e. eigenvectors are “scaled” by the matrix A . In terms of axes, eigenvectors describe the axes upon which A acts as a scaling, and the eigenvalues of A are the possible scalars describing these scalings.

Remark. Let’s be clear about why we require that the eigenvalue/eigenvector equation $A\vec{x} = \lambda\vec{x}$ have a *nonzero* solution. The point is that for *any* scalar λ the equation $A\vec{x} = \lambda\vec{x}$ always has at least one solution: $\vec{x} = \vec{0}$, so without this nonzero requirement any scalar would satisfy the eigenvalue definition and the notion of an eigenvalue would not be very interesting. The key is that $A\vec{x} = \lambda\vec{x}$ should have a solution apart from $\vec{x} = \vec{0}$.

Linguistic remark. The term “eigen” comes from the German word for “proper” or “characteristic”, and in older books you might see the phrase “proper value” or “characteristic value” instead of “eigenvalue”. Nowadays the terms eigenvalue and eigenvector are much more standard, but the old phrase is what suggests that eigenvalues capture something “characteristic” about a matrix.

Example 1. Suppose that A is the 2×2 matrix of a reflection across some line L in \mathbb{R}^2 . We determine the eigenvalues of A using only geometry. First, note that we cannot possibly have a vector satisfying something like

$$A\vec{x} = 2\vec{x}$$

since a reflection will never make a vector twice as long. In fact, since reflections preserve lengths, the only way in which reflecting a vector \vec{x} could result in a multiple of that vector is when that vector satisfied either

$$A\vec{x} = \vec{x} \text{ or } A\vec{x} = -\vec{x}.$$

The first equation is satisfied for any nonzero vector the line L of reflection and the second for any nonzero vector perpendicular to L . Thus, 1 is an eigenvalue of A and any nonzero \vec{x} on L is an eigenvector for 1, and -1 is also eigenvalue of A where any nonzero \vec{x} perpendicular to L is an eigenvector for -1 .

Example 2. Consider the 2×2 matrix of a rotation by an angle $0 < \theta < 180$. As in the case of a reflection, there is no way that rotating a vector by such an angle could result in a longer vector, so only 1 or -1 might be possible eigenvalues. However, when rotating by an angle $0 < \theta < 180$, no nonzero vector can be left as is, so no nonzero vector satisfies $A\vec{x} = \vec{x}$, and no nonzero vector will be flipped completely around, so no nonzero vector satisfies $A\vec{x} = -\vec{x}$. Thus neither 1 nor -1 are actually eigenvalues of A , and we conclude that A has no eigenvalues.

(Actually, we can only conclude that A has no *real* eigenvalues. After we talk about complex numbers we'll see that A actually does have eigenvalues, they just happen to both be complex. This hints at a deep relation between rotations and complex numbers.)

Having 0 as an eigenvalue. Let us note what it would mean for a matrix to have 0 as an eigenvalue. This requires that there be some nonzero \vec{x} satisfying

$$A\vec{x} = 0 \cdot \vec{x} = \vec{0}.$$

However, there can be such a nonzero vector only when A is not invertible, since this equation would say that \vec{x} is in the kernel of A . So, saying that a matrix has 0 as an eigenvalue is the same as saying that it is not invertible, giving us yet another addition to the Amazingly Awesome Theorem; this is probably the last thing we'll add to this theorem, and hopefully we can now all see exactly why it is “amazingly awesome”.

Amazingly Awesome Theorem, continued. A square matrix A is invertible if and only if 0 is not an eigenvalue of A .

Finding eigenvalues. The question still remains as to how we find eigenvalues of a matrix in general, where a simple geometric interpretation might not be readily available as in the previous examples. At first glance it might seem as if finding eigenvalues of a general matrix might be tough since there are essentially *two* unknowns in the equation

$$A\vec{x} = \lambda\vec{x}$$

we must consider: namely, λ and \vec{x} are both “unknown”. In particular, it seems as though knowing whether or not λ is an eigenvalue depends on knowing its eigenvectors ahead of time, but knowing whether or not \vec{x} is an eigenvector depends on knowing its eigenvalue ahead of time. However, it turns out that we can completely determine the eigenvalues first without knowing anything about its eigenvectors. Here's why.

To say that λ is an eigenvalue of A means that $A\vec{x} = \lambda\vec{x}$ should have a nonzero solution. But this equation can be rewritten as

$$A\vec{x} - \lambda\vec{x} = \vec{0}, \text{ or } (A - \lambda I)\vec{x} = \vec{0}$$

after “factoring” out \vec{x} . (Note that we can’t factor out \vec{x} to get $(A - \lambda)\vec{x} = \vec{0}$ since it does not make sense to subtract a scalar from a matrix. But this is easy to get around: we write $A\vec{x} - \lambda\vec{x}$ as $A\vec{x} - \lambda I\vec{x}$ and then we factor out \vec{x} as desired.) So, to say that λ is an eigenvalue of A means that

$$(A - \lambda I)\vec{x} = \vec{0}$$

should have a nonzero solution. But this is only possible precisely when the matrix $A - \lambda I$ is not invertible! (A solution \vec{x} of this equation will then be in the kernel of $A - \lambda I$.) And finally, $A - \lambda I$ is not invertible precisely when its determinant is zero, so we get that

$$\lambda \text{ is an eigenvalue of } A \iff \det(A - \lambda I) = 0.$$

So, to find the eigenvalues of A we must solve the equation $\det(A - \lambda I) = 0$, and this does not involve eigenvectors at all.

Definition. We call $\det(A - \lambda I)$ the *characteristic polynomial* of A . Thus, the eigenvalues of A are the roots of its characteristic polynomial.

Important. To find the eigenvalues of a matrix A , write down $A - \lambda I$ and then compute $\det(A - \lambda I)$. Setting this equal to 0 and solving for λ gives the eigenvalues of A .

Example 3. Recall the matrix $A = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix}$ from our motivating example. We have

$$A - \lambda I = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix} - \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix} = \begin{pmatrix} 13 - \lambda & -6 \\ -1 & 12 - \lambda \end{pmatrix}.$$

Hence

$$\det(A - \lambda I) = (13 - \lambda)(12 - \lambda) - 6 = \lambda^2 - 25\lambda + 150$$

is the characteristic polynomial of A . (Hopefully this example makes it clear why we call refer to this as a “polynomial”.) Since this factors as $(\lambda - 10)(\lambda - 15)$, the roots of the characteristic polynomial of A are 10 and 15, so 10 and 15 are the eigenvalues of A , precisely as we said in our motivating example.

Example 4. Let $B = \begin{pmatrix} 1 & 3 \\ 1 & 2 \end{pmatrix}$. Then

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

According to the quadratic formula, the roots of this are

$$\lambda = \frac{3 \pm \sqrt{9 + 4}}{2}, \text{ so } \frac{3 + \sqrt{13}}{2} \text{ and } \frac{3 - \sqrt{13}}{2}$$

are the eigenvalues of B . In particular, this means that there should be a nonzero vector \vec{x} satisfying

$$\begin{pmatrix} 1 & 3 \\ 1 & 2 \end{pmatrix} \vec{x} = \left(\frac{3 + \sqrt{13}}{2} \right) \vec{x},$$

and such a vector is an eigenvector with eigenvalue $\frac{3 + \sqrt{13}}{2}$. We will talk about how to precisely find eigenvectors next time.

Remark. The 2×2 examples above illustrate a general fact about 2×2 matrices: in general, the characteristic polynomial of $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ is

$$\lambda^2 - (a + d)\lambda + (ad - bc).$$

The constant term is $\det A$, and the sum $a + d$ is called the *trace* of A and is denoted by $\operatorname{tr} A$. (The trace of any square matrix is the sum of its diagonal entries.) Thus we can rewrite this above characteristic polynomial as

$$\lambda^2 - (\operatorname{tr} A)\lambda + \det A,$$

a nice formula which may help to simplify finding eigenvalues of 2×2 matrices. However, note that there is no nice analog of this for larger matrices, so for larger matrices we have to work out $\det(A - \lambda I)$ by hand.

Example 5. Let C be the matrix

$$\begin{pmatrix} -6 & 1 & 4 \\ -9 & 0 & 3 \\ 0 & 0 & 1 \end{pmatrix}.$$

We have (using a cofactor expansion along the third row)

$$\det(C - \lambda I) = \det \begin{pmatrix} -6 - \lambda & 1 & 4 \\ -9 & -\lambda & 3 \\ 0 & 0 & 1 - \lambda \end{pmatrix} = (1 - \lambda) \begin{vmatrix} -6 - \lambda & 1 \\ -9 & -\lambda \end{vmatrix} = (1 - \lambda)(\lambda + 3)^2.$$

Thus the eigenvalues of C are 1 and -3 . We will talk later about what it means for a 3×3 matrix to only have two (real) eigenvalues.

Lecture 20: Eigenvectors

Today we continued talking about eigenvalues and eigenvectors, focusing on *finding* eigenvectors of a matrix once its eigenvalues are known.

Warm-Up. We find the eigenvalues of

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

The characteristic polynomial of A is:

$$\begin{aligned} \det(A - \lambda I) &= \det \begin{pmatrix} 4 - \lambda & 2 & 2 \\ 2 & 4 - \lambda & 2 \\ 2 & 2 & 4 - \lambda \end{pmatrix} \\ &= (4 - \lambda) \begin{vmatrix} 4 - \lambda & 2 \\ 2 & 4 - \lambda \end{vmatrix} - 2 \begin{vmatrix} 2 & 2 \\ 2 & 4 - \lambda \end{vmatrix} + 2 \begin{vmatrix} 2 & 4 - \lambda \\ 2 & 2 \end{vmatrix} \\ &= (4 - \lambda) \begin{vmatrix} 4 - \lambda & 2 \\ 2 & 4 - \lambda \end{vmatrix} + 4 \begin{vmatrix} 2 & 4 - \lambda \\ 2 & 2 \end{vmatrix} \\ &= (4 - \lambda)(16 - 8\lambda + \lambda^2 - 4) + 4(4 - 8 + 2\lambda) \\ &= (4 - \lambda)(\lambda - 2)(\lambda - 6) + 8(\lambda - 2) \end{aligned}$$

$$\begin{aligned}
&= (\lambda - 2)[(4 - \lambda)(\lambda - 6) + 8] \\
&= (\lambda - 2)(-\lambda^2 + 10\lambda - 16) \\
&= -(\lambda - 2)^2(\lambda - 8).
\end{aligned}$$

Thus the eigenvalues of A , which are the roots of the characteristic polynomial, are 2 and 8. Since $\lambda - 2$ appears twice in the characteristic polynomial and $\lambda - 8$ appears once, we say that the eigenvalue 2 has *algebraic multiplicity 2* and the eigenvalue 8 has algebraic multiplicity 1.

Remark. For the matrix A above, you can check that $\det A = 32$. Note also that this is what you get when you multiply the eigenvalues of A together, using 2 twice since it has multiplicity 2:

$$\det A = 2 \cdot 2 \cdot 8.$$

This is true in general: for any square matrix A , $\det A$ equals the product of the eigenvalues of A taking into account multiplicities and possibly having to use complex eigenvalues, which we'll talk about later. We can see this using the fact that the eigenvalues of A are the roots of its characteristic polynomial: if the eigenvalues are $\lambda_1, \dots, \lambda_n$, the characteristic polynomial factors as

$$\det(A - \lambda I) = (\lambda_1 - \lambda)(\lambda_2 - \lambda) \cdots (\lambda_n - \lambda)$$

and setting $\lambda = 0$ in this expression gives $\det A = \lambda_1 \lambda_2 \cdots \lambda_n$.

This makes sense geometrically: each eigenvalue tells us the amount by which A scales a certain direction (the direction corresponding to an eigenvector), and so the overall “expansion factor” corresponding to A is the product of these individual scaling factors.

Finding eigenvectors. Now that we know how to find the eigenvalues of a matrix, the next step is to find its eigenvectors, which geometrically are the vectors on which your matrix acts as a scaling. But we essentially worked out how to do this last time: recall that we derived the condition that $\det(A - \lambda I) = 0$ for an eigenvalue λ using the fact that

$$A\vec{x} = \lambda\vec{x} \text{ is the same as } (A - \lambda I)\vec{x} = \vec{0}.$$

A nonzero vector satisfying the first equation is an eigenvector with eigenvalue λ , so this should be the same as a nonzero vector satisfying the second equation. Thus the eigenvectors corresponding to λ are precisely the nonzero vectors in the kernel of $A - \lambda I$!

Important. For a square matrix A with eigenvalue λ , the eigenvectors of A corresponding to λ are the nonzero vectors in $\ker(A - \lambda I)$. We call this kernel the *eigenspace* of A corresponding to λ .

Eigenspaces are subspaces. Since for an $n \times n$ matrix A , $\ker(A - \lambda I)$ is a subspace of A , it must be true that adding eigenvectors corresponding to the same eigenvalue should give an eigenvector for that same eigenvalue, and similarly scaling an eigenvector should give an eigenvector with the same eigenvalue. We can actually see this directly without having to resort to looking at kernels: if \vec{v}_1 and \vec{v}_2 are both eigenvectors of A with eigenvalue λ , then

$$A(\vec{v}_1 + \vec{v}_2) = A\vec{v}_1 + A\vec{v}_2 = \lambda\vec{v}_1 + \lambda\vec{v}_2 = \lambda(\vec{v}_1 + \vec{v}_2)$$

so $\vec{v}_1 + \vec{v}_2$ is also an eigenvector of a with eigenvalue λ , and for any scalar r we have

$$A(r\vec{v}_1) = r(A\vec{v}_1) = r(\lambda\vec{v}_1) = \lambda(r\vec{v}_1),$$

so $r\vec{v}_1$ is also an eigenvector of A with eigenvalue λ . Thus, as expected, eigenspaces are closed under addition and scalar multiplication.

Since an eigenspace for an $n \times n$ matrix is a subspace of \mathbb{R}^n , it makes sense to ask about possible bases for it. Expressing an eigenspace as $\ker(A - \lambda I)$ tells us how to find such a basis: we simply find a basis for the kernel of $A - \lambda I$ as we would for the kernel of any matrix.

Example 1. We find bases for each eigenspace of $A = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix}$. We found the eigenvalues of A last time to be 10 and 15. The eigenvectors corresponding to 10 are the nonzero vectors in the kernel of

$$A - 10I = \begin{pmatrix} 3 & -6 \\ -1 & 2 \end{pmatrix}.$$

Note that this matrix is not invertible, which makes sense since we found the eigenvalue 10 by determining that this would make $\det(A - 10I) = 0$. This matrix reduces to

$$\begin{pmatrix} 3 & -6 \\ -1 & 2 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & -2 \\ 0 & 0 \end{pmatrix},$$

so a possible basis for $\ker(A - 10I)$ is given by $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$. That is,

$$\text{eigenspace of } A \text{ corresponding to } 10 = \text{span} \left\{ \begin{pmatrix} 2 \\ 1 \end{pmatrix} \right\},$$

which geometrically is the line passing through $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ and the origin. The claim is that any nonzero vector on this line is an eigenvector of A with eigenvalue 10, which geometrically means that anything on this line is scaled by a factor of 10 after under the transformation corresponding to A . For good measure, note that

$$\begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix} \begin{pmatrix} 2 \\ 1 \end{pmatrix} = \begin{pmatrix} 20 \\ 10 \end{pmatrix} = 10 \begin{pmatrix} 2 \\ 1 \end{pmatrix},$$

so $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ is indeed an eigenvector of A with eigenvalue 10.

The eigenvectors corresponding to 15 are the nonzero vectors in the kernel of

$$A - 15I = \begin{pmatrix} -2 & -6 \\ -1 & -3 \end{pmatrix}.$$

This matrix reduces to

$$\begin{pmatrix} -2 & -6 \\ -1 & -3 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 3 \\ 0 & 0 \end{pmatrix},$$

so $\begin{pmatrix} -3 \\ 1 \end{pmatrix}$ alone forms a basis for $\ker(A - 15I)$:

$$\text{eigenspace of } A \text{ corresponding to } 15 = \text{span} \left\{ \begin{pmatrix} -3 \\ 1 \end{pmatrix} \right\}.$$

Again, note that

$$\begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix} \begin{pmatrix} -3 \\ 1 \end{pmatrix} = \begin{pmatrix} -45 \\ 15 \end{pmatrix} = 15 \begin{pmatrix} -3 \\ 1 \end{pmatrix},$$

and similarly anything on the line spanned by $\begin{pmatrix} -3 \\ 1 \end{pmatrix}$ is scaled by 15 under the transformation corresponding to A .

Remark. This matrix above is the one we looked at last time when motivating eigenvalues and eigenvectors, and came from a previous example dealing with coordinates. If you go back to that coordinate example (from November 1st on the Week 6 Lecture Notes), it should now be clear why we used the basis we did in that example: those basis vectors are precisely the eigenvectors we found above, and give us the directions along which A acts as a scaling!

Example 2. We find bases for the eigenspaces of $B = \begin{pmatrix} 7 & 2 \\ -4 & 1 \end{pmatrix}$. First, the characteristic polynomial of B is

$$\det(B - \lambda I) = \begin{vmatrix} 7 - \lambda & 2 \\ -4 & 1 - \lambda \end{vmatrix} = \lambda^2 - 8\lambda + 15 = (\lambda - 5)(\lambda - 3).$$

Thus the eigenvalues of B are 5 and 3. For the eigenvalue 5 we have

$$B - 5I = \begin{pmatrix} 2 & 2 \\ -4 & -4 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}$$

so $\begin{pmatrix} -1 \\ 1 \end{pmatrix}$ forms a basis for the eigenspace of B corresponding to 5. For the eigenvalue 3 we have

$$B - 3I = \begin{pmatrix} 4 & 2 \\ -4 & -2 \end{pmatrix} \rightarrow \begin{pmatrix} 2 & 1 \\ 0 & 0 \end{pmatrix},$$

so $\begin{pmatrix} 1 \\ -2 \end{pmatrix}$ forms a basis for the eigenspace of B corresponding to 3. As a check, note that

$$\begin{pmatrix} 7 & 2 \\ -4 & 1 \end{pmatrix} \begin{pmatrix} -1 \\ 1 \end{pmatrix} = 5 \begin{pmatrix} -1 \\ 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 7 & 2 \\ -4 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ -2 \end{pmatrix} = 3 \begin{pmatrix} 1 \\ -2 \end{pmatrix},$$

so our proposed basis eigenvectors are indeed eigenvectors with the claimed eigenvalues.

(Of course, the matrix we got above when reducing $B - 3I$ is not in reduced echelon form; if you had put it into reduced form you might have gotten $\begin{pmatrix} -1/2 \\ 1 \end{pmatrix}$ as a basis vector. But of course, this vector and $\begin{pmatrix} 1 \\ -2 \end{pmatrix}$ have the same span, so they are both bases for the same eigenspace. I used the vector I did to avoid fractions, which is something you should be on the lookout for as well.)

Example 3. We find bases for the eigenspaces of the matrix A from the Warm-Up. The eigenvalues were 2 with algebraic multiplicity 2 and 8 with multiplicity 1. For 2 we have:

$$A - 2I = \begin{pmatrix} 2 & 2 & 2 \\ 2 & 2 & 2 \\ 2 & 2 & 2 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

giving

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

as a basis for the eigenspace corresponding to 2. You can check that multiplying A by either of these does give 2 times that same vector, as should happen if these are eigenvectors with eigenvalue 2. For the eigenvalues 8 we have:

$$A - 8I = \begin{pmatrix} -4 & 2 & 2 \\ 2 & -4 & 2 \\ 2 & 2 & -4 \end{pmatrix} \rightarrow \begin{pmatrix} -4 & 2 & 2 \\ 0 & -6 & 6 \\ 0 & 0 & 0 \end{pmatrix},$$

so

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

is a basis for the eigenspace corresponding to 8. (I jumped some steps here which you might want to fill in. In fact, since I know that this eigenspace will only be 1-dimensional, if I want to get a basis for this eigenspace all I need to do is find *one* nonzero vector in $\ker(A - 8I)$, and the vector I gave is such a vector. How did I know that this eigenspace would only be 1-dimensional? More on that in a bit.)

This matrix was also one we looked at when dealing with coordinates (Example 2 from November 1st), and low-and-behold the basis eigenvectors we found here were precisely the basis vectors we used in that example. Again, this was no accident ;)

Eigenvectors for different eigenvalues are linearly independent. Note in these three examples that in each case the eigenvectors we found for different eigenvalues turned out to be linearly independent. In fact this is always true: for a square matrix A , if $\vec{v}_1, \dots, \vec{v}_k$ are eigenvectors of A corresponding to distinct eigenvalues $\lambda_1, \dots, \lambda_k$, then $\vec{v}_1, \dots, \vec{v}_k$ must be linearly independent.

The book has a full justification for this, but to get a feel for it let's just work it out when $k = 3$, so we have eigenvectors $\vec{v}_1, \vec{v}_2, \vec{v}_3$ of A corresponding to the different eigenvalues $\lambda_1, \lambda_2, \lambda_3$. To show that $\vec{v}_1, \vec{v}_2, \vec{v}_3$ are linearly independent, we setup the equation

$$c_1\vec{v}_1 + c_2\vec{v}_2 + c_3\vec{v}_3 = \vec{0} \tag{4}$$

and show that for this to be true all coefficients must be zero. Multiplying this through by A and using the fact that

$$A\vec{v}_1 = \lambda_1\vec{v}_1, \quad A\vec{v}_2 = \lambda_2\vec{v}_2, \quad \text{and} \quad A\vec{v}_3 = \lambda_3\vec{v}_3,$$

we get

$$c_1\lambda_1\vec{v}_1 + c_2\lambda_2\vec{v}_2 + c_3\lambda_3\vec{v}_3 = \vec{0}.$$

Now, multiplying equation (1) through by λ_1 gives

$$c_1\lambda_1\vec{v}_1 + c_2\lambda_1\vec{v}_2 + c_3\lambda_1\vec{v}_3 = \vec{0}.$$

Subtracting this from the previous equation gets rid of $c_1\lambda_1\vec{v}_1$, giving

$$c_2(\lambda_2 - \lambda_1)\vec{v}_2 + c_3(\lambda_3 - \lambda_1)\vec{v}_3 = \vec{0}. \tag{5}$$

Multiplying this through by A gives

$$c_2(\lambda_2 - \lambda_1)\lambda_2\vec{v}_2 + c_3(\lambda_3 - \lambda_1)\lambda_3\vec{v}_3 = \vec{0},$$

and multiplying equation (2) through by λ_2 gives

$$c_2(\lambda_2 - \lambda_1)\lambda_2\vec{v}_2 + c_3(\lambda_3 - \lambda_1)\lambda_2\vec{v}_3 = \vec{0}.$$

Subtracting these gets rid of the first term, and we're left with

$$c_3(\lambda_3 - \lambda_1)(\lambda_3 - \lambda_2)\vec{v}_3 = \vec{0}.$$

Since $\vec{v}_3 \neq 0$ and $\lambda_3 - \lambda_1, \lambda_3 - \lambda_2$ are nonzero (since eigenvalues are distinct), this means that c_3 must be zero. Equation (2) then becomes

$$c_2(\lambda_2 - \lambda_1)\vec{v}_2 = \vec{0},$$

so $c_2 = 0$, and equation (1) becomes

$$c_1 \vec{v}_1 = \vec{0},$$

so $c_1 = 0$. Hence $\vec{v}_1, \vec{v}_2, \vec{v}_3$ are linearly independent, and a similar idea works no matter how many eigenvectors corresponding to different eigenvalues we have.

Important. Eigenvectors of a matrix corresponding to distinct eigenvalues are linearly independent. Thus, the collection of vectors formed by putting the bases of all eigenspaces together in one big list is always linearly independent: indeed, basis eigenvectors from the same eigenspace are linearly independent since they came from the same basis, and basis eigenvectors for different eigenspaces are linearly independent by the fact above.

Example 4. One more time, let's find bases for the eigenspaces of

$$C = \begin{pmatrix} 4 & 1 & 0 \\ 0 & 4 & 1 \\ 0 & 0 & 4 \end{pmatrix}.$$

Here, the only eigenvalue is 4, since the eigenvalues of any upper-triangular (or lower-triangular) matrix are simply the entries on its diagonal; indeed, the characteristic polynomial of C is $(4 - \lambda)^3$. We have

$$C - 4I = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix},$$

and basis for $\ker(C - 4I)$, and hence a basis for the eigenspace of C corresponding to 4, is given by

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

Geometric multiplicities and eigenbases. Note that something happened in the above example which had not happened in previous examples: even though the eigenvalue 4 has algebraic multiplicity 3, the dimension of the eigenspace corresponding to 4 is only 1. In previous examples, it was always true that the dimension of each eigenspace was *equal* to the multiplicity of the corresponding eigenvalue.

In general, all we can say is that

$$\dim(\text{eigenspace corresponding to } \lambda) \leq \text{algebraic multiplicity of } \lambda,$$

but these two numbers are *not* necessarily the same. (This is how I knew in Example 3 that the eigenspace corresponding to 8 was only going to be 1-dimensional.) We call $\dim \ker(A - \lambda I)$ the *geometric multiplicity* of the eigenvalue λ , since it describes the dimension of the eigenspace when pictured geometrically.

That leaves the question: when is geometric multiplicity = algebraic multiplicity? Going back to Example 3, note that if we put together all the basis eigenvectors we found into one big list:

$$\begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$$

we get a *basis* of \mathbb{R}^3 . Indeed, these are linearly independent (the first two are linearly independent since they came from a basis for the same eigenspace, and they are linearly independent from the third since they and the third correspond to different eigenvalues), and any three linearly independent vectors in \mathbb{R}^3 is automatically a basis. Such a basis of \mathbb{R}^n , where each basis vector is an eigenvector of a given matrix, is called an *eigenbasis*. In Example 4, we only got one linearly independent eigenvector overall, and that one vector is not enough to give an eigenbasis of \mathbb{R}^3 . Indeed, finding an eigenbasis for \mathbb{R}^n from a given matrix is only possible when the geometric and algebraic multiplicities of each eigenvalue are the same!

The 2×2 matrices in Examples 1 and 2 both gave eigenbases for \mathbb{R}^2 (the basis of \mathbb{R}^2 formed by putting together all basis eigenvectors we found), since for each of those the geometric and algebraic multiplicities of each eigenvalue was 1. We'll come back to the notion of an eigenbasis next week, and see what it *really* means for such a basis to exist.

Important. For any eigenvalue λ of a square matrix A , we have

$$\text{geometric multiplicity of } \lambda \leq \text{algebraic multiplicity of } \lambda.$$

The matrix A gives rise to an *eigenbasis* of \mathbb{R}^n precisely when the geometric multiplicities of all eigenvalues agree with their algebraic multiplicities.

Lecture 21: Applications of Eigenvectors

Today we looked at some applications of eigenvectors outside the scope of this course. This was purely done to convince you that the types of things we are looking at really do show up elsewhere, but none of this (**apart from the Warm-Up**) is part of the standard course material and so will never be on an exam or anything else. Think of today as a break from actual course material, which we will come back to after the midterm.

Warm-Up. We find bases for each eigenspace of

$$A = \begin{pmatrix} 5 & 8 & 1 \\ 0 & 0 & 7 \\ 0 & 0 & -2 \end{pmatrix}, B = \begin{pmatrix} -1 & 0 & 25 \\ 3 & 2 & -15 \\ -1 & 0 & 9 \end{pmatrix}, \text{ and } C = \begin{pmatrix} 2 & 3 & 0 \\ 4 & 3 & 0 \\ 0 & 0 & 6 \end{pmatrix}.$$

First, A is upper-triangular so its eigenvalues are its diagonal entries: 5, 0, 2. The algebraic multiplicity of each is 1, so the geometric multiplicity of each is also 1. Thus each eigenspace should only have 1 basis eigenvector. We have:

$$\begin{aligned} A - 5I &= \begin{pmatrix} 0 & 8 & 1 \\ 0 & -5 & 7 \\ 0 & 0 & -7 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 8 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for this eigenspace is } \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \\ A - 0I &= \begin{pmatrix} 5 & 8 & 1 \\ 0 & 0 & 7 \\ 0 & 0 & -2 \end{pmatrix} \rightarrow \begin{pmatrix} 5 & 8 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for this eigenspace is } \begin{pmatrix} -8 \\ 5 \\ 0 \end{pmatrix}, \\ A - (-2)I &= \begin{pmatrix} 7 & 8 & 1 \\ 0 & 2 & 7 \\ 0 & 0 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 8/7 & 1/7 \\ 0 & 1 & 7/2 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for this eigenspace is } \begin{pmatrix} 54 \\ -49 \\ 14 \end{pmatrix}. \end{aligned}$$

Note that in the last one if you use the given reduced form to get a basis for $\ker(A + 2I)$ you might end up with

$$\begin{pmatrix} 1/7 - 56/14 \\ -7/2 \\ 1 \end{pmatrix},$$

which is fine, but to get a “cleaner” eigenvector I set the free variable equal to 14 to end up with only integer entries. Note that

$$\begin{pmatrix} 5 & 8 & 1 \\ 0 & 0 & 7 \\ 0 & 0 & -2 \end{pmatrix} \begin{pmatrix} 54 \\ -49 \\ 14 \end{pmatrix} = -2 \begin{pmatrix} 54 \\ -49 \\ 14 \end{pmatrix}$$

so the vector I used is indeed an eigenvector of A with eigenvalue -2 . In this case putting all three eigenvectors we found into one list gives an eigenbasis of \mathbb{R}^3 .

The characteristic polynomial of B is

$$\det(B - \lambda I) = -(\lambda - 2)(\lambda - 4)^2,$$

so 2 is an eigenvalue of algebraic multiplicity 1 and 4 an eigenvalue with algebraic multiplicity 2. Thus we know that the eigenspace corresponding to 2 is 1-dimensional and the eigenspace corresponding to 4 is either 1 or 2-dimensional. We have:

$$B - 2I = \begin{pmatrix} -3 & 0 & 25 \\ 3 & 0 & -15 \\ -1 & 0 & 7 \end{pmatrix} \rightarrow \begin{pmatrix} -3 & 0 & 25 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for this eigenspace is } \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix},$$

$$B - 4I = \begin{pmatrix} -5 & 0 & 25 \\ 3 & -2 & -15 \\ -1 & 0 & 5 \end{pmatrix} \rightarrow \begin{pmatrix} -5 & 0 & 25 \\ 0 & -10 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for this eigenspace is } \begin{pmatrix} 5 \\ 0 \\ 1 \end{pmatrix}.$$

In this case, there are only two linearly independent eigenvectors, so there can be no eigenbasis for \mathbb{R}^3 consisting of eigenvectors of B .

Finally, the characteristic polynomial of C is

$$\det(C - \lambda I) = -(\lambda + 1)(\lambda - 6)^2,$$

so the eigenvalues are -1 , with algebraic multiplicity 1, and 6, with algebraic multiplicity 2. We have:

$$C - (-1)I = \begin{pmatrix} 3 & 3 & 0 \\ 4 & 4 & 0 \\ 0 & 0 & 7 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for this eigenspace is } \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix},$$


$$C - 6I = \begin{pmatrix} -4 & 3 & 0 \\ 4 & -3 & 0 \\ 0 & 0 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} -4 & 3 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for this eigenspace is } \left\{ \begin{pmatrix} 3 \\ 4 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right\}.$$

In this case, the geometric multiplicity of each eigenvalue is the same as its algebraic multiplicity, and the three linearly independent eigenvectors we found gives an eigenbasis of \mathbb{R}^3 .

Web search rankings. When you search for something on the internet, whatever search engine you use takes your search terms and goes through its catalog of all possible web pages, picking out

the ones which might in some way be relevant to your search. What then determines the order in which these resulting pages are presented, with whatever the engine thinks is most relevant being listed first? The answer heavily depends on the theory of eigenvectors!

To see how eigenvectors naturally come up in such a problem, let's consider a simplified version of the internet with only three web pages:



where the arrows indicate links from one page to another. The basic assumption is that sites with links from “relevant” pages should themselves be relevant, and the more links from relevant pages it has the more relevant it is. Let's denote the “relevance” of page k by x_k . The goal is to find values for these, which then determine the order in which our three pages are listed after a search, with the one with largest x_k value appearing first.

The assumption that a page's relevance depends on links coming to it from relevant pages turns into a relation among x_1, x_2, x_3 . For instance, page 1 has links from page 2 and from page 3, so x_1 should depend on x_2 and x_3 . Since page 2 only has one link coming out of it, its entire “relevance” contributes to the relevance of page 1, while since page 3 has two links coming out of it, only half of its “relevance” contributes to that of page 1 with the other half contributing to the relevance of page 2. This gives the relation

$$x_1 = x_2 + \frac{x_3}{2}.$$

Similarly, page 2's relevance comes from that of pages 1 and 3, with half of x_1 contributing to x_2 and half of x_3 contributing to x_2 since each of pages 1 and 3 have two links coming out of them; this gives

$$x_2 = \frac{x_1}{2} + \frac{x_3}{2}.$$

Only page 1 links to page 3, so $x_3 = \frac{x_1}{2}$ since page 1 has two links coming out of it. The resulting system of equations

$$\begin{aligned} x_1 &= x_2 + \frac{x_3}{2} \\ x_2 &= \frac{x_1}{2} + \frac{x_3}{2} \\ x_3 &= \frac{x_1}{2} \end{aligned}$$

can be written in matrix form as

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 1/2 \\ 1/2 & 0 & 1/2 \\ 1/2 & 0 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix},$$

which says that the “relevance” vector we’re looking for should be an eigenvector with eigenvalue 1 of $\begin{pmatrix} 0 & 1 & 1/2 \\ 1/2 & 0 & 1/2 \\ 1/2 & 0 & 0 \end{pmatrix}$. Thus, we find the relevance of each page by finding eigenvectors of this matrix!

This was a really simplified example, but the basic idea works for real-life web searches: ranking orders are determined by looking at eigenvectors of some matrices whose entries have something to do with links between pages. All modern search engines somehow use this idea, with various tweaks. In particular, Google’s ranking algorithm, known as “PageRank”, works as follows. Given some search terms, take all pages that might have some relevance; this likely gives over a million pages. Define the matrix A by saying that its ij -th entry is 1 if page i links to page j , and 0 otherwise. As with the example above, now some modifications are done to “weight” the relevances we want, with more links to a page give that page a higher weight; in its most basic form this amounts to replacing A by something of the form

$$D + A$$

where D is some type of “weighting” matrix. (The exact nature of D is one of Google’s trade secrets, as well as any additional modifications which are done to A .) The claim is that the rankings determined by Google’s search engine come from *dominant* eigenvectors of $D + A$, which are eigenvectors corresponding to the largest eigenvalue. In practice, such eigenvectors are almost impossible to find directly, even for a computer, since $D + A$ will be some huge (over 1000000×1000000 in size) matrix, but fortunately there exist good algorithms for approximating the entries of these dominant eigenvectors.

The internet would surely be a much different place if it weren’t for the existence of eigenvectors and eigenvalues!

Population models. Suppose we have populations of deer and wolves in some forest, with $x_1(t)$ denoting the population of deer at time t and $x_2(t)$ the population of wolves at time t . We are interested in understanding the long-term behavior of these two. The basic assumption is that the rate at which these changes (i.e. the values of their derivatives) depend on the values of both at any specific time.

For instance, since wolves feed on deer, the rate of change in the population of deer should obey something like

$$x_1'(t) = (\textit{positive})x_1(t) + (\textit{negative})x_2(t)$$

where the first term comes from deer reproducing (so having a positive effect on population) and the second from deer lost due to the population of wolves (so having a negative effect on deer population). Similarly, the rate of change in the population of wolves might be something like

$$x_2'(t) = (\textit{positive})x_1(t) + (\textit{positive})x_2(t)$$

since the more wolves there are the more wolves there will be, and the more food there is the more wolves there will be. In reality, these are somewhat naive assumptions since there are many other factors contributing to these populations, and maybe in fact since a forest can only support so many wolves, maybe $x_2'(t)$ should actually depend *negatively* on $x_2(t)$. Regardless, we’ll just use this basic model.

Suppose that our populations are modeled by

$$\begin{aligned} x_1'(t) &= 13x_1(t) - 6x_2(t) \\ x_2'(t) &= -x_1(t) + 12x_2(t), \end{aligned}$$

which again probably isn't very realistic, but whatever. This is what is known as a *system of linear differential equations*, and we are interested in determining what functions $x_1(t), x_2(t)$ satisfy these equations. A key observation is that this system can be written as

$$\begin{pmatrix} x_1'(t) \\ x_2'(t) \end{pmatrix} = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix} \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix}.$$

Suppose that the functions we want have the form

$$x_1(t) = r_1 e^{\lambda t} \text{ and } x_2(t) = r_2 e^{\lambda t}.$$

Plugging this into the rewritten system gives

$$\begin{pmatrix} r_1 \lambda e^{\lambda t} \\ r_2 \lambda e^{\lambda t} \end{pmatrix} = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix} \begin{pmatrix} r_1 e^{\lambda t} \\ r_2 e^{\lambda t} \end{pmatrix}.$$

Since $e^{\lambda t}$ is never zero, we can divide both sides through by it to get

$$\lambda \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} = \begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix} \begin{pmatrix} r_1 \\ r_2 \end{pmatrix},$$

which says that the unknowns λ, r_1, r_2 in our expressions for $x_1(t)$ and $x_2(t)$ come from eigenvalues and eigenvectors of the matrix $\begin{pmatrix} 13 & -6 \\ -1 & 12 \end{pmatrix}$! So, finding these eigenvectors is how we are able to find solutions of our population model.

This is a matrix we've seen before, where we computed that its eigenvalues were 10 and 15 with corresponding eigenvectors

$$\begin{pmatrix} 2 \\ 1 \end{pmatrix} \text{ and } \begin{pmatrix} -3 \\ 1 \end{pmatrix}.$$


Thus we get as solutions of our model:

$$\begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} = \begin{pmatrix} 2e^{10t} \\ e^{10t} \end{pmatrix} \text{ and } \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} = \begin{pmatrix} -3e^{15t} \\ e^{15t} \end{pmatrix}.$$

As part of the general theory you would learn about for systems of differential equations in a more advanced differential equations course, it turns out that the general solution of the system above looks like

$$\begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} = c_1 \begin{pmatrix} 2e^{10t} \\ e^{10t} \end{pmatrix} + c_2 \begin{pmatrix} -3e^{15t} \\ e^{15t} \end{pmatrix}.$$

This helps us to visualize our solutions and determine the long-term behavior of our system. Plotting such solutions (for varying c_1 and c_2) on the $x_1 x_2$ -axes gives something which looks like:



The red lines are determined by the directions corresponding to the eigenvectors we found, and each green curve represents a solution for some specific c_1 and c_2 . The observation is that no matter which solution we're on (say we're at the orange dot at time $t = 0$), we will also move "towards" the line determined by the eigenvector with eigenvalue 15 as $t \rightarrow \infty$, essentially because this is the larger eigenvalue. So, long-term, no matter what the initial population of deer and wolves are, the populations will always approach some "ideal" populations determined by the eigenvalue 15.

Again, this is all something you would learn more about in any course which heavily uses differential equations. This was based on a population model, but the same types of models show up in economics and finance, chemistry, engineering, and pretty much everywhere. We'd be lost in all these applications were it not for eigenvalues and eigenvectors!

Eigenfunctions. The previous type of application suggests a strong relation between derivatives and matrices, which indeed we will come to next quarter when we do multivariable calculus. But here's another fun realization.

In this class, all spaces we've dealt with are either \mathbb{R}^n or subspaces of \mathbb{R}^n . However, linear algebra works in more general types of settings; in particular, in other contexts we can consider spaces of "functions" whose elements are themselves functions. Say that V is such a "space", containing say the function $f(x) = x^2$, or $f(x) = e^x$, or $f(x) = \sin x$, etc. We can define an operation D from V to V by

$$D(f) = f'.$$

In other words, D is the "transformation" which takes as input a function and spits out its derivative. The well-known properties of derivatives which say:

$$(f + g)' = f' + g' \text{ and } (cf)' = cf' \text{ for a scalar } c$$

then become the statements that

$$D(f + g) = D(f) + D(g) \text{ and } D(cf) = cD(f),$$

so that D is actually a linear transformation in this more general context!

Now, observe that since the derivative of e^x is e^x we have

$$D(e^x) = e^x,$$

so e^x is an "eigenvector" of D with eigenvalue 1! Also, the derivative of e^{2x} is $2e^{2x}$:

$$D(e^{2x}) = 2e^{2x},$$

so e^{2x} is an eigenvector of D with eigenvalue 2. Such functions are called *eigenfunctions* of D , and D is called a *differential operator*. The study of differential operators and their eigenfunctions has led to deep advancements in physics, chemistry, economics, and pretty much anywhere differential equations show up. Again, things you would learn about in more advanced courses.

Remark. Again, today's lecture was outside the scope of this course, and is purely meant to illustrate how eigenvectors and eigenvalues show up in various contexts. Hopefully you can now somewhat better appreciate why we spend time learning about these things!

Lecture 22: Diagonalization

Today we started talking about what it means for a matrix to be diagonalizable, which we've actually been secretly talking about for a while now. This will be our final topic, apart from the notion of complex eigenvalues, and really brings together many different concepts. It's not an exaggeration to say that pretty much every single thing we've covered this entire quarter plays some kind of role in questions dealing with diagonalizability.

Warm-Up 1. We claim that similar matrices always have the same eigenvalues, with the same algebraic multiplicities. Indeed, suppose that A and B are similar, so that $A = SBS^{-1}$ for some invertible S . Note that then

$$\det(A - \lambda I) = \det(SBS^{-1} - \lambda I).$$

Now, we can write the identity I as $I = SIS^{-1}$ since $SS^{-1} = I$, and making this substitution above gives

$$\det(A - \lambda I) = \det(SBS^{-1} - S\lambda IS^{-1}).$$

The point is that now we can factor S out from the left and S^{-1} out from the right on both sides of the matrix we're taking the determinant above, giving

$$\det(A - \lambda I) = \det(S[B - \lambda I]S^{-1}).$$

The right-side breaks up into $(\det S)(\det[B - \lambda I])(\det S^{-1})$, so we get

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

Hence A and B have the same characteristic polynomial, and so have the same eigenvalues with the same multiplicities as well.

Remark. Let's be careful with what the above is saying. It is NOT true that two matrices with the same eigenvalues must be similar: for instance,

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

have the same eigenvalues with the same algebraic multiplicities, but are not similar, as we will soon see. It is also NOT true that similar matrices must have the same eigenvectors: for instance,

$$\begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix} \text{ and } \begin{pmatrix} 1 & 1 \\ 0 & 2 \end{pmatrix}$$

are similar, as we'll see, but $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ is an eigenvector of the first which is not an eigenvector of the second. On the other hand, it is true that two matrices with *different* eigenvalues or multiplicities cannot be similar.

Warm-Up 2. We find bases for the eigenspaces of

$$A = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & -3 \end{pmatrix} \text{ and } B = \begin{pmatrix} 2 & -5 & 5 \\ 0 & 3 & -1 \\ 0 & -1 & 3 \end{pmatrix}.$$

Since A is upper-triangular, its eigenvalues are 2 with algebraic multiplicity 2 and -3 with multiplicity 1. Thus the eigenspace corresponding to 2 is 1 or 2-dimensional, while the eigenspace corresponding to -3 is 1-dimensional. We have:

$$A - 2I = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -5 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for } E_2 \text{ is } \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix},$$

$$A + 3I = \begin{pmatrix} 5 & 1 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for } E_{-3} \text{ is } \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

(Note that $E_\lambda = \ker(A - \lambda I)$ is just notation for the eigenspace corresponding to λ .) Putting these basis vectors together only gives two linearly independent eigenvectors, so there does not exist an eigenbasis of \mathbb{R}^3 associated to A .

Using a cofactor expansion along the first column, the characteristic polynomial of B is

$$\begin{aligned} \det(B - \lambda I) &= \begin{vmatrix} 2 - \lambda & -5 & 5 \\ 0 & 3 - \lambda & -1 \\ 0 & -1 & 3 - \lambda \end{vmatrix} \\ &= (2 - \lambda) \begin{vmatrix} 3 - \lambda & -1 \\ -1 & 3 - \lambda \end{vmatrix} \\ &= (2 - \lambda)(\lambda^2 - 6\lambda + 8) \\ &= -(\lambda - 2)^2(\lambda - 4). \end{aligned}$$

Thus the eigenvalues of B are 2 with algebraic multiplicity 2 and 4 with multiplicity 1. We have:

$$B - 2I = \begin{pmatrix} 0 & -5 & 5 \\ 0 & 1 & -1 \\ 0 & -1 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 1 & -1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for } E_2 \text{ is } \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} \right\}$$

$$B - 4I = \begin{pmatrix} -2 & -5 & 5 \\ 0 & -1 & -1 \\ 0 & -1 & -1 \end{pmatrix} \rightarrow \begin{pmatrix} -2 & -5 & 5 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \text{ so a basis for } E_4 \text{ is } \begin{pmatrix} 5 \\ -1 \\ 1 \end{pmatrix}.$$

Putting these basis vectors together gives an eigenbasis of \mathbb{R}^3 associated to B , meaning that

$$\left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 5 \\ -1 \\ 1 \end{pmatrix} \right\}$$

is a basis of \mathbb{R}^3 consisting of eigenvectors of B . Note that here the geometric multiplicity of each eigenvalue is the same as its algebraic multiplicity.

Eigenbases are good. Finally we come to the question: why do we care about eigenbases, and whether or not a matrix gives rise to one? The answer is one we've been hinting at for a while now. Consider the matrix B from the Warm-Up, and the associated transformation $T(\vec{x}) = B\vec{x}$. The matrix of T relative to the eigenbasis \mathcal{B} we found turns out to be

$$[T]_{\mathcal{B}} = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 4 \end{pmatrix},$$

precisely because each basis vector was an eigenvector of B ! Indeed, the fact that $T(\vec{v}_i) = \lambda_i \vec{v}_i$ for each of these basis vectors tells us that the coordinate vector of $T(\vec{v}_i)$ simply has λ_i in the i -th position and zeroes elsewhere, which is why the matrix of T turns out to be diagonal. This is good, since it says that geometrically T scales the axes corresponding to these specific basis eigenvectors by an amount equal to the corresponding eigenvalue.

In general, given some transformation $T(\vec{x}) = A\vec{x}$, the only possible bases relative to which the matrix of T is diagonal are ones where each basis vector is an eigenvector of A , since having the i -th column in the matrix of T relative to this basis be of the form

$$i\text{-th column of } [T]_{\mathcal{B}} = \begin{pmatrix} 0 \\ \vdots \\ \lambda_i \\ \vdots \\ 0 \end{pmatrix},$$

with λ_i in the i -th position, requires that the i -th basis vector v_i satisfy $Av_i = \lambda_i v_i$. In other words, such a basis must be an eigenbasis corresponding to A !

Definition. A square matrix A *diagonalizable* if it is similar to a diagonal matrix; i.e. if there exists an invertible matrix S and a diagonal matrix D satisfying $A = SDS^{-1}$. To *diagonalize* a matrix A means to find such an S and D and to express A as $A = SDS^{-1}$. Geometrically, diagonalizable matrices are the ones for which there exists a complete set of “axes” for \mathbb{R}^n upon which the corresponding transformations acts via scalings.

Important. An $n \times n$ matrix A is diagonalizable precisely when it gives rise to an eigenbasis of \mathbb{R}^n . Thus, to diagonalize A (if possible):

- (i) find all eigenvalues of A ,
- (ii) find a basis for each eigenspace of A , and
- (iii) count the *total* number of basis eigenvectors you find and see if you have n of them.

If so, A is diagonalizable and $A = SDS^{-1}$ with S being the matrix having the eigenvectors you found as columns and D being the diagonal matrix with the corresponding eigenvalues down the diagonal. If you end up with fewer than n basis eigenvectors, A is not diagonalizable.

Remark. If you only want to determine diagonalizability without explicitly finding an eigenbasis (i.e. without explicitly diagonalizing it), it is enough to check whether the geometric multiplicity of each eigenvalue is the same as its algebraic multiplicity: these have to be equal in order for a matrix to be diagonalizable.

Example 1. Let $A = \begin{pmatrix} 1 & 3 \\ 4 & -3 \end{pmatrix}$. The eigenvalues of A are -5 and 3 , so already we know that A is diagonalizable: the algebraic multiplicity of each eigenvalue is 1 and so the geometric multiplicity of each eigenvalue must also be 1. Finding a basis for each of $\ker(A + 5I)$ and $\ker(A - 3I)$ gives

$$E_{-5} = \text{span} \left\{ \begin{pmatrix} -1 \\ 2 \end{pmatrix} \right\} \text{ and } E_3 = \text{span} \left\{ \begin{pmatrix} 3 \\ 2 \end{pmatrix} \right\}.$$

Thus $\left\{\begin{pmatrix} -1 \\ 2 \end{pmatrix}, \begin{pmatrix} 3 \\ 2 \end{pmatrix}\right\}$ forms an eigenbasis for \mathbb{R}^2 and we can diagonalize A as

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

Note that the order in which we write the eigenvalues in the diagonal matrix D matters: they should correspond to the order in which we write the eigenvectors as columns of S .

Of course, a diagonalization of A is not unique. For one thing, we can change the order of the columns of S and the order of the eigenvalues in D :

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

or we can use different eigenvectors altogether; for instance, $\begin{pmatrix} 2 \\ -4 \end{pmatrix}$ is also an eigenvector of A with eigenvalue -5 and $\begin{pmatrix} 9 \\ 6 \end{pmatrix}$ is another eigenvector with eigenvalue 3 so

$$\begin{pmatrix} 1 & 3 \\ 4 & -3 \end{pmatrix} = \begin{pmatrix} 2 & 9 \\ -4 & 6 \end{pmatrix} \begin{pmatrix} -5 & 0 \\ 0 & 3 \end{pmatrix} \begin{pmatrix} 2 & 9 \\ -4 & 6 \end{pmatrix}^{-1}$$

as well. There is no preference for one diagonalization over another, except that trying to avoid fractions might be a good idea.

Example 2. Let $B = \begin{pmatrix} 4 & -1 \\ 1 & 2 \end{pmatrix}$. This only has one eigenvalue, namely 3 . Since $B - 3I = \begin{pmatrix} 1 & -1 \\ 1 & -1 \end{pmatrix}$, we only come up with one basis eigenvector for the eigenspace corresponding to 3 . With only one eigenvalue, there are no other eigenspaces which could produce basis eigenvectors, so B is not diagonalizable.

Actually, there is a way to see this is true only knowing that B has one eigenvalue. In general, if A is an $n \times n$ diagonalizable matrix with only one eigenvalue λ , then A must equal λI . Indeed, A diagonalizable gives $A = SDS^{-1}$ with D diagonal, but if λ is the only eigenvalue of A then D must be $D = \lambda I$. But then

$$A = SDS^{-1} = S(\lambda I)S^{-1} = \lambda(SIS^{-1}) = \lambda I,$$

so A must have been $A = \lambda I$ to start with. In the example above, since the only eigenvalue of B is 3 , if B was going to be diagonalizable it must have been equal to $3I$, which it is not.

Back to Warm-Up. The matrix A from the second Warm-Up is not diagonalizable, while the matrix B is. Using the eigenbasis we found associated to B , we can diagonalize B as

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

Example 3. Suppose that A is a diagonalizable 3×3 matrix with eigenvectors

$$\begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \text{ and } \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$$

corresponding to the eigenvalues $2, 2,$ and 3 respectively. We want to compute $A \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix}$. Now, we are not given A explicitly, but we can actually find A using the given information: note that the

two given eigenvectors with eigenvalue 2 are linearly independent, so all three given eigenvectors together form an eigenbasis for \mathbb{R}^3 . Hence we can diagonalize A as

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

Thus by computing out the right-hand side, we can figure out what A actually is.

BUT, this is totally unnecessary! To save this extra work, I claim that we can actually find $A \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix}$ *without* knowing A explicitly. Think about what is going on: we have a basis for \mathbb{R}^3 made up of eigenvectors of A , and we know that A acts as a scaling on each of those basis vectors. Thus, if we know how to express the vector $\begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix}$ in terms of that basis, then we can use linearity properties to easily determine $A \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix}$.

To be clear, since the given vectors form a basis of \mathbb{R}^3 there are coefficients satisfying.

$$\begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix} = c_1 \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + c_2 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + c_3 \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

Multiplying by A gives

$$\begin{aligned} A \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix} &= A \left(c_1 \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + c_2 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + c_3 \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right) \\ &= c_1 A \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + c_2 A \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + c_3 A \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \\ &= 2c_1 \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + 2c_2 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + 3c_3 \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \end{aligned}$$

where in the last step we use the fact that the basis vectors are eigenvectors of A , so we know what A times each of them is. Thus all we need to know to be able to compute $A \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix}$ are the coefficients c_1, c_2, c_3 . Solving

$$\begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix} = c_1 \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + c_2 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + c_3 \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$$

gives $c_1 = 2, c_2 = -1, c_3 = 1$, so

$$A \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix} = 2(2) \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + 2(-1) \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + 3(1) \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 7 \\ 3 \\ 5 \end{pmatrix},$$

as desired. Again, note that we still don't even know what A actually is, and yet using the same method as above we can in fact compute $A\vec{x}$ for any possible \vec{x} !

Remark. This idea, that we can determine how a (diagonalizable) linear transformation acts without explicitly knowing that transformation, lies at the core of many important applications

of eigenvalues and eigenvectors. In practice, it is often the case that you have enough data to determine enough eigenvectors and eigenvalues of some transformation without explicitly knowing what that transformation is, and if you're lucky this is enough information to do what you want. In particular, most computations in quantum physics are based on this idea ;)

Lecture 23: More on Diagonalization

Today we continued talking about diagonalization, and looked at some interesting applications which were based on the fact that diagonalizing a matrix makes its powers relatively easy to compute. As I said in class, these applications were purely meant to show why this is something we might want to know how to do, but they will NOT be on the final.

Warm-Up 1. We determine the values of k for which

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

is diagonalizable. First, the characteristic polynomial of A is

$$\det(A - \lambda I) = \begin{vmatrix} -\lambda & 0 & 3 \\ 3 & k - \lambda & 3 \\ 1 & 0 & -2 - \lambda \end{vmatrix} = (k - \lambda) \begin{vmatrix} -\lambda & 3 \\ 1 & -2 - \lambda \end{vmatrix} = (k - \lambda)(\lambda - 1)(\lambda + 3),$$

where we used a cofactor expansion along the second column. Hence the eigenvalues of A are $k, 1, -3$. So, there are either two or three distinct eigenvalues depending on what k is.

If $k \neq 1, -3$, there are three distinct eigenvalues and so in this case A is for sure diagonalizable: with three distinct eigenvalues each eigenspace is 1-dimensional and finding a basis vector for each gives 3 linearly independent eigenvectors overall.

If $k = 1$, then there are only two eigenvalues: 1 with algebraic multiplicity 2 and -3 with algebraic multiplicity 1. We will get one basis eigenvector corresponding to -3 , so what determines whether or not A is diagonalizable is how many basis eigenvectors we get for the eigenvalue 1. We have (keeping in mind that $k = 1$):

$$A - I = \begin{pmatrix} -1 & 0 & 3 \\ 3 & 0 & 3 \\ 1 & 0 & -3 \end{pmatrix} \rightarrow \begin{pmatrix} -1 & 0 & 3 \\ 0 & 0 & 12 \\ 0 & 0 & 0 \end{pmatrix},$$

so E_1 is 1-dimensional. Hence we only get one basis eigenvector for $\lambda = 1$, and together with the basis eigenvector for -3 we only get two overall, so A is not diagonalizable.

If $k = -3$, then again there are two eigenvalues, but now 1 has algebraic multiplicity 1 and -3 has algebraic multiplicity 2. We will get one basis eigenvector corresponding to 1, and since (keeping in mind that $k = -3$)

$$A + 3I = \begin{pmatrix} 3 & 0 & 3 \\ 3 & 0 & 3 \\ 1 & 0 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} 3 & 0 & 3 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

has a 2-dimensional kernel, E_{-3} is two dimensional so we get two basis eigenvectors. These together with the basis eigenvector for 1 gives three in total, so A is diagonalizable.

To summarize, A is diagonalizable for all $k \neq 1$. Note however that the reasons differ for $k \neq -3$ and $k = -3$: in the former case there are three distinct eigenvalues, while in the latter there are only two but the geometric multiplicity of each eigenvalue agrees with its algebraic multiplicity.

Warm-Up 2. Suppose that A is diagonalizable and that A is similar to B . We claim that B must also be diagonalizable. Indeed, A diagonalizable gives

$$A = SDS^{-1}$$

for some invertible S and diagonal D , while A similar to B gives

$$A = PBP^{-1}$$

for some invertible P . Then

$$SDS^{-1} = PBP^{-1}, \text{ so } B = P^{-1}SDS^{-1}P = (P^{-1}S)D(P^{-1}S)^{-1},$$

so B is similar to the diagonal matrix D and is hence diagonalizable.

Computing powers. Why do we care about diagonalizable matrices? Here is perhaps the main practical reason: if $A = SDS^{-1}$, then

$$A^k = SD^kS^{-1}.$$

Indeed, if you write out A^k as $(SDS^{-1})^k$:

$$(SDS^{-1})(SDS^{-1})(SDS^{-1}) \cdots (SDS^{-1}),$$

note that all the $S^{-1}S$ terms cancel out so we're left with the first S , then a bunch of D 's, and the final S^{-1} . In addition, if D is diagonal, its power are easy to compute:

$$\begin{pmatrix} \lambda_1 & \cdots & 0 \\ & \ddots & \\ 0 & \cdots & \lambda_n \end{pmatrix}^k = \begin{pmatrix} \lambda_1^k & \cdots & 0 \\ & \ddots & \\ 0 & \cdots & \lambda_n^k \end{pmatrix},$$

that is, D^k is the diagonal matrix whose entries are the k -th powers of the diagonal entries of D . Putting this all together gives a relatively easy way to find A^k when A is diagonalizable.

Important. If A is diagonalizable and $A = SDS^{-1}$ with D diagonal, then $A^k = SD^kS^{-1}$ where D^k is still diagonal with diagonal entries equal to the k -th powers of the diagonal entries of D .

Example 1. Last time we diagonalized $\begin{pmatrix} 1 & 3 \\ 4 & -3 \end{pmatrix}$ as

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

Then

$$\begin{pmatrix} 1 & 3 \\ 4 & -3 \end{pmatrix}^k = \begin{pmatrix} -1 & 3 \\ 2 & 2 \end{pmatrix} \begin{pmatrix} -5 & 0 \\ 0 & 3 \end{pmatrix}^k \begin{pmatrix} -1 & 3 \\ 2 & 2 \end{pmatrix}^{-1} = \begin{pmatrix} -1 & 3 \\ 2 & 2 \end{pmatrix} \begin{pmatrix} (-5)^k & 0 \\ 0 & 3^k \end{pmatrix} \begin{pmatrix} -1 & 3 \\ 2 & 2 \end{pmatrix}^{-1}.$$

The right side is now pretty straightforward to compute, and so we get a concrete description of $\begin{pmatrix} 1 & 3 \\ 4 & -3 \end{pmatrix}^k$ for any $k > 0$.

Example 2. Say we want to solve

$$\begin{pmatrix} -1 & 1 & -1 \\ -2 & 2 & -1 \\ -2 & 2 & -1 \end{pmatrix}^{100} \vec{x} = \begin{pmatrix} 4 \\ 1 \\ 2 \end{pmatrix}.$$

The wrong way to go about is to try to actually this 100th power directly. Instead, the matrix

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

has three distinct eigenvalues 0, 1, and -1 , so it is diagonalizable. Finding a basis for each eigenspace gives one possible diagonalization as:

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

From this we have

$$\begin{aligned} \begin{pmatrix} -1 & 1 & -1 \\ -2 & 2 & -1 \\ -2 & 2 & -1 \end{pmatrix}^{100} &= \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1^{100} & 0 \\ 0 & 0 & (-1)^{100} \end{pmatrix} \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}^{-1} \\ &= \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}^{-1}. \end{aligned}$$

The inverse on the right is

$$\begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 1 & 0 \\ 1 & -1 & 1 \end{pmatrix},$$

so

$$\begin{aligned} \begin{pmatrix} -1 & 1 & -1 \\ -2 & 2 & -1 \\ -2 & 2 & -1 \end{pmatrix}^{100} &= \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}^{-1} \\ &= \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 & -1 \\ -1 & 1 & 0 \\ 1 & -1 & 1 \end{pmatrix} \\ &= \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ -1 & -1 & 0 \\ 1 & -1 & 1 \end{pmatrix} \\ &= \begin{pmatrix} 1 & -1 & 1 \\ 0 & -2 & 1 \\ 0 & -2 & 1 \end{pmatrix}. \end{aligned}$$

Now we can solve

$$\begin{pmatrix} -1 & 1 & -1 \\ -2 & 2 & -1 \\ -2 & 2 & -1 \end{pmatrix}^{100} \vec{x} = \begin{pmatrix} 1 & -1 & 1 \\ 0 & -2 & 1 \\ 0 & -2 & 1 \end{pmatrix} \vec{x} = \begin{pmatrix} 4 \\ 1 \\ 2 \end{pmatrix}$$

fairly straightforwardly, and it turns out that there are no solutions. The point is that diagonalizing

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

gave us a direct way to compute its powers.

Remark. The two examples which follow are only meant to illustrate how computing powers of a matrix might come up in applications, but as said in the intro you will not be expected to know how to do these types of examples on the final.

Fibonacci numbers. The *Fibonacci numbers* are the numbers defined as follows: start with 1, 1, and take as the next term the sum of the previous terms. So, the first few fibonacci numbers are

$$1, 1, 2, 3, 5, 8, 13, 21, 35, \dots$$

Our goal is to find an explicit expression for the n -th Fibonacci number F_n . The key is the equation:

$$F_{n+2} = F_{n+1} + F_n,$$

which says precisely that each term is the sum of the two previous terms. Take this equation and throw in the silly-looking equation $F_{n+1} = F_{n+1}$ to get the system

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

Now, this can be written in matrix form as

$$\begin{pmatrix} F_{n+2} \\ F_{n+1} \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} F_{n+1} \\ F_n \end{pmatrix},$$

and thus the matrix $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ tells us how to move from F_n and F_{n+1} to F_{n+1} and F_{n+2} . Similarly,

$$\begin{pmatrix} F_{n+1} \\ F_n \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} F_n \\ F_{n-1} \end{pmatrix},$$

and combining this with our previous equation gives

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

Continuing in this manner, the end result is that

$$\begin{pmatrix} F_{n+2} \\ F_{n+1} \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}^n \begin{pmatrix} F_2 \\ F_1 \end{pmatrix}.$$

We know that $F_2 = F_1 = 1$, so if can find an expression for $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}^n$ we will be able to explicitly compute F_{n+2} and F_{n+1} .

Luckily the matrix $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ is diagonalizable! Its characteristic polynomial is $\lambda^2 - \lambda - 1$, so its eigenvalues (using the quadratic formula) are $\frac{1 \pm \sqrt{5}}{2}$. Let's denote these by

$$\lambda_+ = \frac{1 + \sqrt{5}}{2} \text{ and } \lambda_- = \frac{1 - \sqrt{5}}{2}.$$

We have

$$A - \lambda_{\pm} I = \begin{pmatrix} 1 - \lambda_{\pm} & 1 \\ 1 & -\lambda_{\pm} \end{pmatrix} \rightarrow \begin{pmatrix} 1 & -\lambda_{\pm} \\ 1 - \lambda_{\pm} & 1 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & -\lambda_{\pm} \\ 0 & 0 \end{pmatrix},$$

so possible eigenvectors corresponding to each of these respectively are

$$\begin{pmatrix} \lambda_+ \\ 1 \end{pmatrix} \text{ and } \begin{pmatrix} \lambda_- \\ 1 \end{pmatrix}.$$

Thus $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ diagonalizes as

$$\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} \lambda_+ & \lambda_- \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \lambda_+ & 0 \\ 0 & \lambda_- \end{pmatrix} \begin{pmatrix} \lambda_+ & \lambda_- \\ 1 & 1 \end{pmatrix}^{-1},$$

and so

$$\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}^n = \begin{pmatrix} \lambda_+ & \lambda_- \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \lambda_+^n & 0 \\ 0 & \lambda_-^n \end{pmatrix} \begin{pmatrix} \lambda_+ & \lambda_- \\ 1 & 1 \end{pmatrix}^{-1}.$$

The inverse on the right is

$$\begin{pmatrix} \lambda_+ & \lambda_- \\ 1 & 1 \end{pmatrix}^{-1} = \frac{1}{\sqrt{5}} \begin{pmatrix} 1 & -\lambda_- \\ -1 & \lambda_+ \end{pmatrix},$$

so

$$\begin{aligned} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}^n &= \frac{1}{\sqrt{5}} \begin{pmatrix} \lambda_+ & \lambda_- \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \lambda_+^n & 0 \\ 0 & \lambda_-^n \end{pmatrix} \begin{pmatrix} 1 & -\lambda_- \\ -1 & \lambda_+ \end{pmatrix} \\ &= \frac{1}{\sqrt{5}} \begin{pmatrix} \lambda_+ & \lambda_- \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \lambda_+^n & -\lambda_- \lambda_+^n \\ -\lambda_-^n & \lambda_+ \lambda_-^n \end{pmatrix} \\ &= \frac{1}{\sqrt{5}} \begin{pmatrix} \lambda_+^{n+1} - \lambda_-^{n+1} & -\lambda_- \lambda_+^{n+1} + \lambda_+ \lambda_-^{n+1} \\ \lambda_+^n - \lambda_-^n & -\lambda_- \lambda_+^n + \lambda_+ \lambda_-^n \end{pmatrix}. \end{aligned}$$

Putting it all together gives

$$\begin{pmatrix} F_{n+2} \\ F_{n+1} \end{pmatrix} = \frac{1}{\sqrt{5}} \begin{pmatrix} \lambda_+^{n+1} - \lambda_-^{n+1} & -\lambda_- \lambda_+^{n+1} + \lambda_+ \lambda_-^{n+1} \\ \lambda_+^n - \lambda_-^n & -\lambda_- \lambda_+^n + \lambda_+ \lambda_-^n \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix},$$

and multiplying out the right-hand side gives an explicit expression for F_{n+1} ; after readjusting n we get an explicit expression for F_n . Note that this expression can still be simplified further, since for instance

$$\lambda_- \lambda_+ = \left(\frac{1 - \sqrt{5}}{2} \right) \left(\frac{1 + \sqrt{5}}{2} \right) = -1.$$

But, you get the idea.

Remark. Again, this is not a computation you'd be expected to be able to on the final since, as you can see, although fairly straightforward it does get a little messy. The system above where we

start with a vector and repeatedly multiply by the same matrix in order to generate new vectors is an example of what's called a *discrete dynamical system*. Diagonalization plays a big role in the study of such systems, and in a related concept known as a *Markov chain*. These are topics you would perhaps come across in some later course, in other departments as well.

Matrix exponentials. Take $A = \begin{pmatrix} 1 & 3 \\ 4 & -3 \end{pmatrix}$ to be the matrix from Example 1. We want to compute the *matrix exponential* e^A . The first question is: what on Earth is meant by taking e to the power of a matrix? The answer comes from recalling some calculus and looking at Taylor series for e^x :

$$e^x = \sum_{n=0}^{\infty} \frac{x^n}{n!} = 1 + x + \frac{x^2}{2} + \frac{x^3}{3!} + \dots$$

(If you didn't see Taylor series in your calculus course, no worries, this again is just to illustrate how diagonalization can be useful.) The point is that this series makes sense when we substitute a matrix in place of x , so we define e^A to be the matrix

$$e^A = \sum_{n=0}^{\infty} \frac{A^n}{n!} = I + A + \frac{1}{2}A^2 + \frac{1}{3!}A^3 + \dots$$

Since A is diagonalizable, this series is actually something we can directly compute.

First, diagonalizing $A = SDS^{-1}$ as we did in Example 1, we have $A^n = SD^nS^{-1}$ so

$$e^A = I + SDS^{-1} + \frac{1}{2}SD^2S^{-1} + \frac{1}{3!}SD^3S^{-1} + \dots$$

Now, if we rewrite I as SIS^{-1} , then we can factor S out on the left and S^{-1} out on the right of the resulting expression, so

$$\begin{aligned} e^A &= SIS^{-1} + SDS^{-1} + \frac{1}{2}SD^2S^{-1} + \frac{1}{3!}SD^3S^{-1} + \dots \\ &= S \left(I + D + \frac{1}{2}D^2 + \frac{1}{3!}D^3 + \dots \right) S^{-1}. \end{aligned}$$

But the infinite sum in the middle is the definition of e^D , so we get that

$$e^A = Se^D S^{-1}.$$

Now all that's left is to compute e^D for $D = \begin{pmatrix} -5 & 0 \\ 0 & 3 \end{pmatrix}$. This is (fairly) straightforward:

$$\begin{aligned} e^D &= I + D + \frac{1}{2}D^2 + \frac{1}{3!}D^3 + \dots \\ &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \begin{pmatrix} -5 & 0 \\ 0 & 3 \end{pmatrix} + \begin{pmatrix} \frac{(-5)^2}{2} & 0 \\ 0 & \frac{3^2}{2} \end{pmatrix} + \begin{pmatrix} \frac{(-5)^3}{3!} & 0 \\ 0 & \frac{3^3}{3!} \end{pmatrix} + \dots \\ &= \begin{pmatrix} 1 + (-5) + \frac{1}{2}(-5)^2 + \frac{1}{3!}(-5)^3 + \dots & 0 \\ 0 & 1 + 3 + \frac{1}{2}3^2 + \frac{1}{3!}3^3 + \dots \end{pmatrix} \end{aligned}$$

and we can now recognize these diagonal terms as the series expressions for e^{-5} and e^3 respectively. Thus

$$e^D = \begin{pmatrix} e^{-5} & 0 \\ 0 & e^3 \end{pmatrix},$$

so

$$e^A = \begin{pmatrix} -1 & 3 \\ 2 & 2 \end{pmatrix} \begin{pmatrix} e^{-5} & 0 \\ 0 & e^3 \end{pmatrix} \begin{pmatrix} -1 & 3 \\ 2 & 2 \end{pmatrix}^{-1}$$

and multiplying this out gives an explicit expression for e^A .

Matrix exponentials show up in many applications, such as in any advanced study of differential equations or in applications of differential equations. The upshot is that matrix exponentials are pretty easy to compute for diagonalizable matrices, which is what makes them manageable. Note that you can now do all sorts of crazy looking computations, such as

$$\sin A \text{ or } \cos A$$

where A is a matrix: you just take the Taylor series definitions for $\sin x$ and $\cos x$ and replace x in those series by the matrix A .

Non-diagonalizable matrices aren't so bad. Again, this final remark is not something you would have to know for the final. The question remains: what happens when you try to compute powers of a non-diagonalizable matrix? It turns out that this is still somewhat manageable to do (and thus so is computing e to the power of a non-diagonalizable matrix), since even matrices which aren't diagonalizable are still "almost diagonalizable" in the sense that any matrix whatsoever will always (as long as we allow complex eigenvalues and complex eigenvectors) be similar to one of the form

$$\begin{pmatrix} \lambda_1 & * & & & \\ & \lambda_2 & * & & \\ & & \ddots & \ddots & \\ & & & \lambda_{n-1} & * \\ & & & & \lambda_n \end{pmatrix}$$

where all non-diagonal entries are zero *except* for possibly some 1's in the starred locations right above the diagonal. Such a matrix is called a *Jordan matrix* and if A is similar to this, this is called the *Jordan form* of A . The point is that for such matrices, powers are still somewhat straightforward to compute in a way which is good enough for most applications.

Jordan forms are something you would learn about in a later linear algebra course, such as Math 334. They are related to what are called *generalized eigenvectors*, which as the name suggests are generalizations of eigenvectors. As one final fact, we can now answer a questions which I'm sure has been on all of your minds: when are two matrices similar? The answer: two square matrices are similar if and only if they have the same Jordan form.

Lecture 24: Complex Eigenvalues

Today we spoke about complex eigenvalues and eigenvectors. The point is that this all works the same way as real eigenvalues and eigenvectors do, only that now we work with complex numbers.

Warm-Up. Say A is a 4×4 matrix with eigenvectors

$$\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ 3 \\ -1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 2 \\ 3 \\ -1 \end{pmatrix}, \begin{pmatrix} 1 \\ 2 \\ 0 \\ 0 \end{pmatrix}$$

corresponding to the eigenvalues 2, 2, 1, 1 respectively. We want to find A^{-1} and justify along the way that A is indeed invertible.

First note that the given eigenvectors are linearly independent, which is maybe easier to see by writing them in the order

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

The matrix with these as columns has full rank, so those columns are linearly independent. Hence we know that the eigenspaces E_1 and E_2 are each *at least* 2-dimensional, since we have at least two linearly independent eigenvectors in each. Thus the algebraic multiplicities of 1 and 2 are at least 2, so the characteristic polynomial of A looks like

$$\det(A - \lambda I) = (\lambda - 1)^m(\lambda - 2)^n(\text{other stuff}), \text{ where } m, n \geq 2.$$

But A is only 4×4 , so since the degree of this polynomial has to be 4 there is no choice but to have

$$\det(A - \lambda I) = (\lambda - 1)^2(\lambda - 2)^2.$$

Thus 0 is not an eigenvalue of A so A is invertible, and since now we see that the geometric multiplicity of each eigenvalue equals its algebraic multiplicity, A is diagonalizable.

Using the given eigenvectors, we can diagonalize A as

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

Now, we could multiply this out to find A , and then use that to find A^{-1} . However, we save some time as follows. Recall that in general the inverse of a product of matrices is the product of the individual inverses but in *reverse* order. So, for instance, $(BCD)^{-1} = D^{-1}B^{-1}C^{-1}$. In our case, this means that

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

But $(S^{-1})^{-1} = S$, so

$$\begin{aligned} A^{-1} &= \begin{pmatrix} 1 & 1 & 2 & 1 \\ 0 & 2 & 3 & 2 \\ 0 & 0 & -1 & 3 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}^{-1} \begin{pmatrix} 1 & 1 & 2 & 1 \\ 0 & 2 & 3 & 2 \\ 0 & 0 & -1 & 3 \\ 0 & 0 & 0 & 1 \end{pmatrix}^{-1} \\ &= \begin{pmatrix} 1 & 1 & 2 & 1 \\ 0 & 2 & 3 & 2 \\ 0 & 0 & -1 & 3 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1/2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 & 2 & 1 \\ 0 & 2 & 3 & 2 \\ 0 & 0 & -1 & 3 \\ 0 & 0 & 0 & 1 \end{pmatrix}^{-1} \end{aligned}$$

(Note that if $A = SDS^{-1}$, then $A^{-1} = SD^{-1}S^{-1}$ so the formula $A^k = SD^kS^{-1}$ for the powers of a diagonalizable matrix we saw last time works even for *negative* powers.) The inverse on the right is

$$\begin{pmatrix} 1 & 1 & 2 & 1 \\ 0 & 2 & 3 & 2 \\ 0 & 0 & -1 & 3 \\ 0 & 0 & 0 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} 1 & -1/2 & 1/2 & -3/2 \\ 0 & 1/2 & 3/2 & -11/2 \\ 0 & 0 & -1 & 3 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

so multiplying out

$$A^{-1} = \begin{pmatrix} 1 & 1 & 2 & 1 \\ 0 & 2 & 3 & 2 \\ 0 & 0 & -1 & 3 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1/2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & -1/2 & 1/2 & -3/2 \\ 0 & 1/2 & 3/2 & -11/2 \\ 0 & 0 & -1 & 3 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

will give us A^{-1} . As opposed to finding A first and then A^{-1} , this method only requires us to find one inverse explicitly (namely S^{-1}) using row operations instead of two: S^{-1} and A^{-1} .

Example 1. Consider the matrix $A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$. Its characteristic polynomial is

$$\det(A - \lambda I) = \lambda^2 + 1$$

so A has no real eigenvalues. (This makes sense, since a rotation by 90° will turn no vector into a scalar multiple of itself.) However, A does have two complex eigenvalues: i and $-i$. We can find eigenvectors for each of these using the same method as for real eigenvalues.

For $\lambda = i$, we have

$$A - iI = \begin{pmatrix} -i & -1 \\ 1 & -i \end{pmatrix}.$$

Now, we can reduce this by multiplying the first row by $-i$ and adding to the second row, but instead since we know this matrix will not be invertible, we know that the second row will have to become all zeroes after reducing, so

$$\begin{pmatrix} -i & -1 \\ 1 & -i \end{pmatrix} \rightarrow \begin{pmatrix} -i & -1 \\ 0 & 0 \end{pmatrix}.$$

Now, for a matrix of this form, finding a nonzero vector in the kernel is easy: we just switch the two entries in the top row and multiply one by a negative. Thus,

$$\begin{pmatrix} 1 \\ -i \end{pmatrix}$$

is in $\ker(A - iI)$, so this is an eigenvector of A with eigenvalue i . As a check:

$$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ -i \end{pmatrix} = \begin{pmatrix} i \\ 1 \end{pmatrix} = i \begin{pmatrix} 1 \\ -i \end{pmatrix},$$

so $\begin{pmatrix} 1 \\ -i \end{pmatrix}$ is indeed an eigenvector with eigenvalue i .

For $\lambda = -i$, we have

$$A + iI = \begin{pmatrix} i & -1 \\ 1 & i \end{pmatrix} \rightarrow \begin{pmatrix} i & -1 \\ 0 & 0 \end{pmatrix}$$

so $\begin{pmatrix} 1 \\ i \end{pmatrix}$ is an eigenvector of A with eigenvalue $-i$. Thus, even though A is not diagonalizable over \mathbb{R} , we would say that it is *diagonalizable over \mathbb{C}* (the set of complex numbers) as:

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

Finally, this inverse is obtained using the same formula for the inverse of a 2×2 matrix with real entries:

$$\begin{pmatrix} 1 & 1 \\ -i & i \end{pmatrix}^{-1} = \frac{1}{2i} \begin{pmatrix} i & -1 \\ i & 1 \end{pmatrix}.$$

Remark. Note something special that happened above: the eigenvalues of A were complex conjugates of each other, and the associated eigenvectors were also complex conjugates of each other in the sense that one is obtained by taking the complex conjugate of each entry of the other. This is actually true for *any* matrix with real entries which has complex eigenvalues, and cuts down on many computations which come up when dealing with complex eigenvalues.

Important. Suppose that A is a square matrix with real entries. If $a + ib$ is an eigenvalue of A , then $a - ib$ is also an eigenvalue of A . Moreover, if \vec{v} is a complex eigenvector of A with eigenvalue $a + ib$, then $\overline{\vec{v}}$ is an eigenvector of A with eigenvalue $a - ib$ where $\overline{\vec{v}}$ means the vector obtained by taking the complex conjugate of each entry of \vec{v} .

Example 2. Let $A = \begin{pmatrix} 6 & 1 \\ -17 & 2 \end{pmatrix}$. The characteristic polynomial of A is

$$\det(A - \lambda I) = \lambda^2 - 4\lambda + 5.$$

Using the quadratic formula, the eigenvalues of A are then $2 \pm i$. To find eigenvectors, all we need to do is find an eigenvector for $2 + i$ and then take its conjugate to get one for $2 - i$. We have:

$$A - (2 + i)I = \begin{pmatrix} 4 - i & 1 \\ -17 & -4 - i \end{pmatrix} \rightarrow \begin{pmatrix} 4 - i & 1 \\ 0 & 0 \end{pmatrix},$$

so $\begin{pmatrix} 1 \\ -4+i \end{pmatrix}$ is an eigenvector of A with eigenvalue $2 + i$. Hence $\begin{pmatrix} 1 \\ -4-i \end{pmatrix}$ is an eigenvector of A with eigenvalue $2 - i$, so we can diagonalize A over \mathbb{C} as

$$\begin{pmatrix} 6 & 1 \\ -17 & 2 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ -4+i & -4-i \end{pmatrix} \begin{pmatrix} 2+i & 0 \\ 0 & 2-i \end{pmatrix} \begin{pmatrix} 1 & 1 \\ -4+i & -4-i \end{pmatrix}^{-1}.$$

Now, notice that the matrix $B = \begin{pmatrix} 2 & -1 \\ 1 & 2 \end{pmatrix}$ has the same eigenvalues as A , and after finding some eigenvectors we see that we can diagonalize B as

$$B = \begin{pmatrix} 1 & 1 \\ -i & i \end{pmatrix} \begin{pmatrix} 2+i & 0 \\ 0 & 2-i \end{pmatrix} \begin{pmatrix} 1 & 1 \\ -i & i \end{pmatrix}^{-1}.$$

Since A and B are both similar to $\begin{pmatrix} 2+i & 0 \\ 0 & 2-i \end{pmatrix}$, they are similar to each other! The matrix $\begin{pmatrix} 2 & -1 \\ 1 & 2 \end{pmatrix}$ geometrically represents a rotation combined with some scalings—compare to the rotation matrix $\begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$ —so we conclude that $A = \begin{pmatrix} 6 & 1 \\ -17 & 2 \end{pmatrix}$ also represents a rotation combined with scalings.

Remark. In general, a 2×2 matrix with complex eigenvalues $a \pm ib$ will be similar to

$$\begin{pmatrix} a & -b \\ b & a \end{pmatrix},$$

and so geometrically represents a rotation combined with scalings. This is further evidence that there is a deep relation between complex numbers and rotations, which you would elaborate more on in a complex analysis course.

Example 3. Let A be the matrix

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

which has characteristic polynomial

$$(1 - \lambda)(\lambda^2 - 2\lambda + 10).$$

Hence the eigenvalues of A are 1 and $1 \pm 3i$. For $\lambda = 1$ we get $\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$ as an eigenvector. For $\lambda = 1 + 3i$ we have

$$A - (1 + 3i)I = \begin{pmatrix} -3i & 2 & -1 \\ 0 & -3i & -3 \\ 0 & 3 & 3i \end{pmatrix} \rightarrow \begin{pmatrix} -3i & 2 & -1 \\ 0 & -3i & -3 \\ 0 & 0 & 0 \end{pmatrix}.$$

Setting the free variable equal to $3i$, we find that one possible eigenvector is

$$\begin{pmatrix} -1 + 2i \\ -3 \\ 3i \end{pmatrix}.$$

Hence

$$\begin{pmatrix} -1 - 2i \\ -3 \\ -3i \end{pmatrix}$$

is an eigenvector of A for the eigenvalue $1 - 3i$, so we can diagonalize A as

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

Note that things get tougher once you move past 2×2 matrices with complex eigenvalues!

3-dimensional rotations have axes of rotations. And now, after a full quarter, we can finally justify something I claimed on the very first day of class, and which I included as part of the introduction to the class on the syllabus: any 3-dimensional rotation has an axis of rotation. Note how much we had to develop in order to get to this point!

Say that A is a 3×3 rotation matrix. First, we know that A must have at least one real eigenvalue, since complex eigenvalues come in (conjugate) pairs and a 3×3 matrix will have 3 eigenvalues counted with multiplicity. Now, we also know that since A describes a rotation, only -1 and 1 can be real eigenvalues. We claim that 1 must be an eigenvalue. There are two possibilities: either A has 3 real eigenvalues or it has 1 real eigenvalue.

If A has 3 real eigenvalues and they are all -1 , then $\det A$, which is the product of the eigenvalues of A , would be -1 , but we've seen that a rotation must have positive determinant. Hence if A has 3 real eigenvalues at least one of them must be 1 .

If A has only 1 real eigenvalue, it has two other complex eigenvalue $a \pm ib$. If -1 is the one real eigenvalue, then

$$\det A = (-1)(a - ib)(a + ib) = -(a^2 + b^2)$$

is negative, which again is not possible. Hence the one real eigenvalue of A must be 1.

Thus either way, 1 is an eigenvalue of A ; take \vec{x} to be a corresponding eigenvector. Then the line spanned by \vec{x} is an axis of rotation for A . Tada!