Vectors

This week we will learn about:

- What vectors are,
- How to manipulate vectors, and
- Linear combinations.

Extra reading and watching:

- Section 1.1 in the textbook
- Lecture videos 1, 2, and 3 on YouTube
- Vector at Wikipedia

Extra textbook problems:

⋆ 1.1.1–1.1.3, 1.1.5–1.1.8

⋆⋆ 1.1.9–1.1.12

⋆⋆⋆ 1.1.13(a), 1.1.14, 1.1.15

☠ 1.1.13(b)

Linear algebra is one of the branches of mathematics where everything "just works". Most of the objects that we introduce in this course will seem rather simple and unremarkable at first, but we will be able to do a *lot* with them. Some of the things we will be able to do are motivated very geometrically...

Example. Lengths, angles, and deformations (oh my!)

...while other applications will involve sifting through huge amounts of data:

Example. How does (well, *did*) Google work?

Vectors

A **vector** is an ordered list of numbers like $(3, 1)$. These lists can be as long as we like, but we'll start by considering 2-dimensional vectors in order to establish some intuition for how they work, since we can interpret them geometrically in this case.

Several different notations are used for vectors:

The **coordinates** or **entries** of a vector only tell us how far the vector stretches in the x- and y-directions; **not** where it is located in space.

Example. Coordinates of vectors.

That is, vectors represent...

The order of the coordinates matters: for example, $(2,3) \neq (3,2)$. For this reason, 2D vectors are sometimes called "ordered pairs". In another math class, you might be introduced to objects called "sets", where order does *not* matter.

- Two vectors are equal if and only if...

- The **zero vector** is...

- Recall that the set of all real numbers is denoted by $\mathbb{R}$. Similarly,

Sometimes we want to combine two (or more) vectors to get new ones. For example, we might want to think about what happens if we move along the path of multiple different vectors, one after another. Where do we end up after doing this? The answer is given by **vector addition**.

Example. Vector addition.

Specifically, if $\mathbf{v} = (v_1, v_2)$ and $\mathbf{w} = (w_1, w_2)$, then $\mathbf{v} + \mathbf{w} = (v_1 + w_1, v_2 + w_2)$ is the vector from the tail of $\mathbf{v}$ to the head of $\mathbf{w}$.

Another common way to manipulate vectors is to "scale" them (or "multiply by a scalar"). The idea here is that we want to move in the same direction as a given vector, but we want to change how *far* we move in that direction.

Example. Scalar multiplication.

Specifically, if $\mathbf{v} = (v_1, v_2)$ and c is a real number, then $c\mathbf{v} = (cv_1, cv_2)$ is the vector that points in the same direction as $\mathbf{v}$, but is c times as long (and if $c < 0$ then the direction of the vector is reversed).

Finally, how do you think vector subtraction might be defined? If $\mathbf{v} = (v_1, v_2)$ and $\mathbf{w} = (w_1, w_2)$, then

$$\mathbf{v} - \mathbf{w} =$$

is the vector from the to the .

Example. Vector subtraction.

Example. Suppose that a regular hexagon has its center at the point $(0, 0)$ and one of its corners at the point $(1, 0)$. Find the sum of the 6 vectors that point from its center to its corners.

If $\mathbf{v}$ and $\mathbf{w}$ are the (non-parallel) sides of a parallelogram, then $\mathbf{v}+\mathbf{w}$ and $\mathbf{v}-\mathbf{w}$ appear very naturally in that parallelogram...

Example. The parallelogram rule.

3-Dimensional Vectors

Everything we have learned about vectors so far extends naturally to 3 dimensions.

A vector in 3 dimensions is an **ordered triple** like $(1, 3, 2)$, and the set of all ordered triples is denoted by $\mathbb{R}^3$. These are a bit harder to draw than their 2-dimensional counterparts, but it's still possible...

Example. Drawing 3D vectors.

Everything we have seen in 2D carries over exactly how you would expect in 3D:

- Adding vectors still has the geometric interpretation of "following" both vectors, one after the other.
- Adding vectors has the same formula you might expect:
- Scalar multiplication still has the geometric interpretation of stretching the vector.
- Scalar multiplication has the same formula you might expect:

High-Dimensional Vectors

Everything we have learned about vectors so far extends naturally to 4 (and more!) dimensions.

We'll get a bit more general now and consider an *arbitrary* number of dimensions.

A vector in n dimensions is an **ordered n-tuple** like $(1, 2, 3, \ldots, n)$, and the set of all ordered n-tuples is denoted by $\mathbb{R}^n$. These are a bit harder to draw than their 2- and 3-dimensional counterparts...

Example. Drawing 4D (and 5D, and 6D...) vectors.

However, everything *algebraic* that we have seen for 2D and 3D vectors carries over exactly how you would expect in higher dimensions:

- Adding vectors has the same formula you might expect:

- Scalar multiplication has the same formula you might expect:

Even though we can't draw vectors in $\mathbb{R}^n$, we still want to be able to manipulate them. We have seen that vector addition and scalar multiplication work the "naive" way. The following theorem shows some more properties that are similarly "obvious":

Theorem 1.1 — Properties of Vector Operations

Let $\mathbf{v}, \mathbf{w}, \mathbf{x} \in \mathbb{R}^n$ be vectors and let $c, d \in \mathbb{R}$ be scalars. Then

a) $\mathbf{v} + \mathbf{w} = \mathbf{w} + \mathbf{v}$ (commutativity)

b) $(\mathbf{v} + \mathbf{w}) + \mathbf{x} = \mathbf{v} + (\mathbf{w} + \mathbf{x})$ (associativity)

c) $c(\mathbf{v} + \mathbf{w}) = c\mathbf{v} + c\mathbf{w}$ (distributivity)

d) $(c + d)\mathbf{v} = c\mathbf{v} + d\mathbf{v}$ (distributivity)

e) $\mathbf{v} + \mathbf{0} = \mathbf{v}$

f) $\mathbf{v} + (-\mathbf{v}) = \mathbf{0}$

g) $c(d\mathbf{v}) = (cd)\mathbf{v}$

Proof. We will prove property (a) in class; you can try to prove some of the others on your own (the method is quite similar).

which completes the proof. ■

Why did we even bother with the theorem on the previous page? They're all the type of thing you can just look at and "see" are true, right?

One reason is that we have to make sure that certain combinations of symbols even make sense when we are in new and unfamiliar settings. For example, associativity (property (b)) says that this expression makes sense:

We will soon introduce some operations that do not have these basic properties like commutativity, so we will have to start being very careful.

Example. Simplify $\mathbf{v} + 2(\mathbf{w} - \mathbf{v}) - 3(\mathbf{v} + 2\mathbf{w})$.

Linear Combinations

One common task in linear algebra is to start out with some given collection of vectors $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k$ and then use vector addition and scalar multiplication to construct new vectors out of them. The following definition gives a name to this concept.

Definition 1.1 — Linear Combinations

A **linear combination** of $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k \in \mathbb{R}^n$ is a vector of the form

$$c_1\mathbf{v}_1 + c_2\mathbf{v}_2 + \cdots + c_k\mathbf{v}_k,$$

where $c_1, c_2, \ldots, c_k \in \mathbb{R}$.

Example. Show that $(1, 2, 3)$ is a linear combination of the vectors $(1, 1, 1)$ and $(-1, 0, 1)$.

Example. Show that $(1, 2, 3)$ is *not* a linear combination of $(1, 1, 0)$ and $(2, 1, 0)$.

When working with linear combinations, some particularly important vectors are the ones with a single 1 in one of their entries, and all other entries equal to 0. These are called the **standard basis vectors**:

Example. List and draw all of the standard basis vectors in $\mathbb{R}^2$ and $\mathbb{R}^3$.

For now, the reason for our interest in these standard basis vectors is that every vector $\mathbf{v} = (v_1, v_2, \ldots, v_n) \in \mathbb{R}^n$ can be written as a linear combination of them. In particular,

This idea of writing vectors in terms of the standard basis vectors is one of the most useful tricks that we make use of in linear algebra: in many situations, if we can prove that some property holds for the standard basis vectors, then we can use linear combinations to show that it must hold for *all* vectors.

Example. Compute $3\mathbf{e}_1 - 2\mathbf{e}_2 + \mathbf{e}_3 \in \mathbb{R}^3$.

Example. Write $(3, 5, -2, -1)$ as a linear combination of $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4 \in \mathbb{R}^4$.

Lengths, Angles, and the Dot Product

This week we will learn about:

- The dot product,
- The length of vectors and the angle between them, and
- The Cauchy–Schwarz and triangle inequalities.

Extra reading and watching:

- Section 1.2 in the textbook
- Lecture videos 4, 5, 6, and 7 on YouTube
- Dot product at Wikipedia
- Cauchy–Schwarz inequality at Wikipedia

Extra textbook problems:

$\star$ 1.2.1–1.2.3, 1.2.7, 1.2.8

$\star\star$ 1.2.4–1.2.6, 1.2.9–1.2.11

$\star\star\star$ 1.2.12, 1.2.13, 1.2.17–1.2.21

☠ 1.2.23

The Dot Product

In 2D (and sometimes in 3D), it is fairly intuitive to talk about geometric quantities like lengths or angles. You have used things like similar triangles and the law of cosines for tackling problems like this in the past.

Using vectors, we can now generalize these concepts to arbitrary dimensions (even though we can't picture it)! Our main tool will be...

Definition 2.1 — Dot Product

If $\mathbf{v} = (v_1, v_2, \ldots, v_n) \in \mathbb{R}^n$ and $\mathbf{w} = (w_1, w_2, \ldots, w_n) \in \mathbb{R}^n$ then the **dot product** of $\mathbf{v}$ and $\mathbf{w}$, denoted by $\mathbf{v} \cdot \mathbf{w}$, is the quantity

$$\mathbf{v} \cdot \mathbf{w} \stackrel{\text{def}}{=} v_1 w_1 + v_2 w_2 + \cdots + v_n w_n.$$

Please be wary of what types of objects go into and come out of the dot product:

Intuitively, the dot product $\mathbf{v} \cdot \mathbf{w}$ tells you how much $\mathbf{v}$ points in the direction of $\mathbf{w}$ (or how much $\mathbf{w}$ points in the direction of $\mathbf{v}$).

Example. 2D examples.

Example. Higher-dimensional examples.

We have defined a new mathematical operation, so it's time for another "obvious" theorem telling us what properties it satisfies:

Theorem 2.1 — Properties of the Dot Product

Let $\mathbf{v}, \mathbf{w}, \mathbf{z} \in \mathbb{R}^n$ be vectors and let $c \in \mathbb{R}$ be a scalar. Then

a) $\mathbf{v} \cdot \mathbf{w} = \mathbf{w} \cdot \mathbf{v}$ (commutativity)

b) $\mathbf{v} \cdot (\mathbf{w} + \mathbf{z}) = \mathbf{v} \cdot \mathbf{w} + \mathbf{v} \cdot \mathbf{z}$ (distributivity)

c) $(c\mathbf{v}) \cdot \mathbf{w} = c(\mathbf{v} \cdot \mathbf{w})$

Proof. We will prove property (a). You can try the rest on your own (the method is quite similar).

This completes the proof. ■

Example. Compute $\frac{1}{2}(-1, -3, 2) \cdot (6, -4, 2)$.

Example. Show that $(\mathbf{v} + \mathbf{w}) \cdot (\mathbf{v} + \mathbf{w}) = \mathbf{v} \cdot \mathbf{v} + 2\mathbf{v} \cdot \mathbf{w} + \mathbf{w} \cdot \mathbf{w}$ for all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$.

Length of a Vector

We now start making use of the dot product to talk about things like the length of vectors or the angle between vectors (even in high-dimensional spaces).

Example. Length of vectors in $\mathbb{R}^2$.

Example. Length of vectors in $\mathbb{R}^3$.

In higher dimensions, we *define* the length of a vector so as to continue the pattern that we observed above:

Definition 2.2 — Length of a Vector

The **length** of a vector $\mathbf{v} = (v_1, v_2, \ldots, v_n) \in \mathbb{R}^n$, denoted by $\|\mathbf{v}\|$, is defined by

$$\|\mathbf{v}\| \stackrel{\text{def}}{=} \sqrt{\mathbf{v} \cdot \mathbf{v}} = \sqrt{v_1^2 + v_2^2 + \cdots + v_n^2}.$$

Example. Compute the length of some vectors.

As always, we have defined a new mathematical object, so we want a theorem that tells us what its properties are.

Theorem 2.2 — Properties of Vector Length

Let $\mathbf{v} \in \mathbb{R}^n$ be a vector and let $c \in \mathbb{R}$ be a scalar. Then

a) $\|c\mathbf{v}\| = |c|\|\mathbf{v}\|$

b) $\|\mathbf{v}\| = 0$ if and only if $\mathbf{v} = \mathbf{0}$

Proof. To prove property (a), we just apply the relevant definitions:

To prove property (b), we have to prove two things:

This completes the proof. ∎

A vector with length 1 is called a **unit vector**. Every non-zero vector $\mathbf{v} \in \mathbb{R}^n$ can be divided by its length to get a unit vector:

Scaling $\mathbf{v}$ to have length 1 like this is called **normalizing $\mathbf{v}$** (and this unit vector $\mathbf{w}$ is called the **normalization** of $\mathbf{v}$).

Example. Normalize the vector $(3, 4) \in \mathbb{R}^2$.

Example. Show that the standard basis vectors are unit vectors.

We now start to look at somewhat more interesting properties of the dot product and vector lengths. Our first result in this direction is an inequality that relates the dot product of two vectors to their lengths:

Theorem 2.3 — Cauchy–Schwarz Inequality

Suppose that $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ are vectors. Then $|\mathbf{v} \cdot \mathbf{w}| \leq \|\mathbf{v}\|\|\mathbf{w}\|$.

Proof. Define the vector $\mathbf{x} = \|\mathbf{w}\|\mathbf{v} - \|\mathbf{v}\|\mathbf{w}$ and then expand the quantity $\|\mathbf{x}\|^2$ in terms of the dot product:

$\blacksquare$

The above theorem is our first example of a theorem with a very non-obvious proof: even though we can follow the steps and see that they are individually true, the choice of $\mathbf{x} = \|\mathbf{w}\|\mathbf{v} - \|\mathbf{v}\|\mathbf{w}$ at the start was something like magic. This particular choice of $\mathbf{x}$ was chosen so that the proof would give us what we wanted. Other choices of $\mathbf{x}$ also result in true inequalities, but ones that are less useful than Cauchy–Schwarz.

Example. Do there exist vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^2$ with...

We have two main uses for the Cauchy–Schwarz inequality. The first is that it helps us prove another geometrically "obvious" fact about vector lengths:

Theorem 2.4 — Triangle Inequality

Suppose that $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ are vectors. Then $\|\mathbf{v} + \mathbf{w}\| \leq \|\mathbf{v}\| + \|\mathbf{w}\|$.

Proof. We start by expanding $\|\mathbf{v} + \mathbf{w}\|^2$ in terms of the dot product:

■

Angle Between Vectors

The second immediate use of the Cauchy–Schwarz inequality is that it helps us define angles in $\mathbb{R}^n$. To get an idea of how this works, let's start by thinking about a triangle with sides given by the vectors $\mathbf{v}$, $\mathbf{w}$, and $\mathbf{v} - \mathbf{w}$:

Our reasoning above gave us a formula for the angle between two vectors in $\mathbb{R}^2$ (and in $\mathbb{R}^3$). We now state it as a definition in higher-dimensional spaces.

Definition 2.3 — Angle Between Vectors

The **angle** θ between two non-zero vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ is the quantity

$$\theta = \arccos\left(\frac{\mathbf{v} \cdot \mathbf{w}}{\|\mathbf{v}\|\|\mathbf{w}\|}\right).$$

Example. What is the angle between $\mathbf{v} = (1, 1, 1, 1)$ and $\mathbf{w} = (2, 0, 2, 0)$?

Example. What is the angle between the diagonals of two adjacent faces of a cube?

Recall that $\arccos(x)$ is only defined if $-1 \le x \le 1$. How do we know that $-1 \le \frac{\mathbf{v}\cdot\mathbf{w}}{\|\mathbf{v}\|\|\mathbf{w}\|} \le 1$?

One special case of vector angles that is worth pointing out is the case when $\mathbf{v} \cdot \mathbf{w} = 0$. When this happens...

This special case is important enough that we give it its own name:

Definition 2.4 — Orthogonality

Two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ are called **orthogonal** if $\mathbf{v} \cdot \mathbf{w} = 0$.

Example. Show that the vectors $(1, 1, -2)$ and $(3, 1, 2)$ are orthogonal.

Example. Find a non-zero vector orthogonal to $\mathbf{v} = (v_1, v_2) \in \mathbb{R}^2$

Matrices and Matrix Operations

This week we will learn about:

- Matrices,
- Matrix addition, scalar multiplication, and matrix multiplication,
- The transpose, and
- Matrix powers and adjacency matrices of graphs.

Extra reading and watching:

- Section 1.3 in the textbook
- Lecture videos 8, 9, 10, 11, and 12 on YouTube
- Matrix multiplication at Wikipedia
- Transpose at Wikipedia

Extra textbook problems:

⋆ 1.3.1, 1.3.2, 1.3.4, 1.3.12

⋆⋆ 1.3.3, 1.3.5–1.3.7, 1.3.9, 1.3.11, 1.3.13–1.3.15

⋆⋆⋆ 1.3.8

☠ none this week

Matrices

Previously, we introduced vectors, which can be thought of as 1D lists of numbers. Now we start working with matrices, which are 2D arrays of numbers:

Definition 3.1 — Matrices

A **matrix** is a rectangular array of numbers. Those numbers are called the **entries** or **elements** of the matrix.

Example. Examples of matrices.

The **size** of a matrix is a description of the number of rows and columns that it has. A matrix with m rows and n columns has size $m \times n$.

A $1 \times n$ matrix is called a **row matrix** or **row vector**. An $m \times 1$ matrix is called a **column matrix** or **column vector**. An $n \times n$ matrix is called **square**.

We use double subscripts to specify individual entries of a matrix: the entry of the matrix A in row i and column j is denoted by $a_{i,j}$. For example, if

then $a_{1,3} =$ and $a_{2,2} =$.

Similarly, when we say "the (i,j)-entry of A", we mean $a_{i,j}$. Another notation for this is $[A]_{i,j}$, and we will see some examples shortly where this notation is advantageous.

With this notation in mind, a general $m \times n$ matrix A has the following form:

$$A = \begin{bmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{bmatrix}.$$

Two matrices are **equal** if they have the same size and all of their entries (in the same positions) are equal to each other.

Example. Some (un)equal matrices.

We use $\mathcal{M}_{m,n}$ to denote the set of $m \times n$ matrices, and the shorthand $\mathcal{M}_n$ for the set of $n \times n$ matrices.

Just like we could add vectors or multiply vectors by a scalar, we can also add matrices and multiply matrices by scalars, and their definitions are exactly what you would expect:

Definition 3.2 — Matrix Addition and Scalar Multiplication

Suppose A and B are $m \times n$ matrices, and $c \in \mathbb{R}$ is a scalar. Then their **sum** $A+B$ is the matrix whose (i,j)-entry is $a_{i,j} + b_{i,j}$, and the **scalar multiplication** cA is the matrix whose (i,j)-entry is $ca_{i,j}$.

In other words, these operations are just performed entry-wise, as you might expect. The definition of matrix addition only makes sense when A and B have the same size.

Example. Matrix addition and scalar multiplication.

Matrix subtraction is defined analogously:

Matrix addition, subtraction, and scalar multiplication satisfy all of the "natural" properties you might expect (e.g., $A+B=B+A$). We state these properties as a theorem:

Theorem 3.1 — Properties of Matrix Operations

Let $A, B, C \in \mathcal{M}_{m,n}$ and let $c, d \in \mathbb{R}$ be scalars. Then

a) $A+B=B+A$ (commutativity)

b) $(A+B)+C=A+(B+C)$ (associativity)

c) $c(A+B)=cA+cB$ (distributivity)

d) $(c+d)A=cA+dA$ (distributivity)

e) $c(dA)=(cd)A$

Proof. We will only prove part (c) of the theorem. The remaining parts of the theorem can be proved similarly: just use the definition of matrix addition and use the fact that all of these properties hold for addition of real numbers.

This completes the proof. ■

Matrix Multiplication

What about matrix multiplication? Recall that multiplication was a bit tricky with vectors: we only saw the dot product, which "multiplied" two vectors to give us a number. Matrix multiplication is a bit different than this, and looks quite messy and ugly at first glance. So hold onto your hats...

Definition 3.3 — Matrix Multiplication

If A is an $m \times n$ matrix and B is an $n \times p$ matrix, then their **product** AB is the $m \times p$ matrix whose (i, j)-entry is:

$$[AB]_{i,j} \stackrel{\text{def}}{=} a_{i,1}b_{1,j} + a_{i,2}b_{2,j} + \cdots + a_{i,n}b_{n,j}.$$

In other words, the product AB is the matrix whose entries are all of the possible dot products of the rows of A with the columns of B.

We emphasize that the matrix product AB only makes sense if A has the same number of *columns* as B has *rows*. For example, it does not make sense to multiply a 2×3 matrix by another 2×3 matrix, but it does make sense to multiply a 2×3 matrix by a 3×7 matrix.

Example. Compute the product of two matrices.

When performing matrix multiplication, double-check that the sizes of your matrices actually make sense. In particular, the inner dimensions of the matrices must be equal, and the outer dimensions of the matrices will be the dimensions of the matrix product:

As always, we have defined a new operation (matrix multiplication), so we want to know what properties it satisfies.

Theorem 3.2 — Properties of Matrix Multiplication

Let A, B, and C be matrices (with sizes such that all of the multiplications below make sense) and let $c \in \mathbb{R}$ be a scalar. Then

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

b) $A(B + C) = AB + AC$ (left distributivity)

c) $(A + B)C = AC + BC$ (right distributivity)

d) $c(AB) = (cA)B = A(cB)$

Proof. We will only prove part (b) of the theorem. The remaining parts of the theorem can be proved similarly: just use the definition of matrix multiplication.

■

Notice that we did *not* say anything about commutativity (i.e., we did not claim that $AB = BA$). Why not?

Example. Commutativity of matrix multiplication?

Example. FOILing matrices.

One particularly important square matrix is the one that consists entirely of 0 entries, except with 1s on its diagonal. This is called the **identity matrix**, and it is denoted by I (or sometimes I_n if we want to emphasize it is $n \times n$).

Similarly, the **zero matrix** is the one with all entries equal to 0. We denote it by O (or $O_{m,n}$ if we care that it is $m \times n$).

Example. The identity matrix and zero matrix.

The previous example suggests the following general result, which is indeed true:

Theorem 3.3 — Multiplication by Identity or Zero

If $A \in \mathcal{M}_{m,n}$ then $AI_n = A = I_m A$ and $AO_n = O_{m,n} = O_m A$.

We won't prove the above theorem, but hopefully it seems believable enough.

Example. Diagonal matrices.

In general, the product of two diagonal matrices is just the entry-wise product of the two matrices:

The Transpose of a Matrix

We now introduce an operation on matrices that changes the shape of a matrix, but not its contents. Specifically, it swaps the role of the rows and columns of a matrix:

Definition 3.4 — The Transpose

Suppose $A \in \mathcal{M}_{m,n}$ is an $m \times n$ matrix. Then its **transpose**, which we denote by A^T, is the $n \times m$ matrix whose (i, j)-entry is $a_{j,i}$.

Intuitively, the transpose of a matrix is obtained by mirroring it across its main diagonal.

Example. Let's compute a transpose or two.

Let's now think about some basic properties that the transpose satisfies:

Theorem 3.4 — Properties of the Transpose

Let A and B be matrices (with sizes such that the operations below make sense) and let $c \in \mathbb{R}$ be a scalar. Then

a) $(A^T)^T = A$

b) $(A + B)^T = A^T + B^T$

c) $(AB)^T = B^T A^T$

d) $(cA)^T = cA^T$

Proof. Parts (a), (b), and (d) of the theorem are intuitive enough, so we will only prove part (c):

■

As a bit of a side note: would you have initially guessed that $(AB)^T = A^T B^T$? Situations like this are why we prove things rather than just guessing based on what "looks believable".

Example. Let's compute some more transposes.

Example. Transpose of the product of many matrices.

The transpose has the useful property that it converts a column vector into the corresponding row vector, and vice-versa. Furthermore, if $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ are column vectors, then we can use our usual matrix multiplication rule to see that

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

In other words, we can use matrix multiplication to recover the dot product.

Matrix Powers

Matrix multiplication also lets us define *powers* of square matrices. For an integer $k \geq 1$, we define

and we also define $A^0 = I$ (analogously to how we define $a^0 = 1$ whenever a is a non-zero real number). The next theorem follows almost immediately from this definition:

Theorem 3.5 — Properties of Matrix Powers

If A is square and k and r are nonnegative integers, then

- $A^k A^r = A^{k+r}$, and
- $(A^k)^r = A^{kr}$.

Example. Compute some matrix powers.

Later in the course, we will see how to define things like and .

Block Matrices

Oftentimes, there are clear "patterns" in the entries of a large matrix, and it might be useful to break that matrix down into smaller chunks based on some partition of its rows and columns. For example:

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

When A and B are written in this way, as matrices whose elements are themselves matrices, they are called **block matrices**. Viewing matrices in this way often simplifies calculations and reveals structure, especially when the matrix has a lot of zeroes.

Remarkably, multiplication of block matrices works exactly as it does for regular matrices:

And indeed, this is the exact same answer we would have gotten if we computed AB the "long way".

We have to be careful when performing block matrix multiplication: it is only valid if we choose the sizes of the blocks so that each and every matrix multiplication being performed makes sense.

Example. Suppose

Which of the following block matrix multiplications make sense?

$$\begin{bmatrix} A & B \\ B & A \end{bmatrix}^2$$

$$\begin{bmatrix} A & B \\ O & I_3 \end{bmatrix} \begin{bmatrix} A & A \\ O & A \\ I_2 & O \end{bmatrix}$$

$$\begin{bmatrix} A & B \\ O & I_3 \end{bmatrix} \begin{bmatrix} A & A \\ O & A \end{bmatrix}$$

$$\begin{bmatrix} A & B \\ O & I_3 \end{bmatrix} \begin{bmatrix} B & O \\ I_3 & I_3 \end{bmatrix}$$

Partitioning matrices in different ways can lead to new insights about how matrix multiplication works.

Theorem 3.6 — Matrix–Vector Multiplication

Suppose $A \in \mathcal{M}_{m,n}$ has columns $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n$ and $\mathbf{v} \in \mathbb{R}^n$ is a column vector. Then

$$A\mathbf{v} = v_1\mathbf{a}_1 + v_2\mathbf{a}_2 + \cdots + v_n\mathbf{a}_n.$$

Proof. We simply perform block matrix multiplication:

■

We of course can compute $A\mathbf{v}$ directly from the definition, but it's nice to have multiple ways to think about things.

Theorem 3.7 — Matrix Multiplication is Column-Wise

Suppose $A \in \mathcal{M}_{m,n}$ and $B \in \mathcal{M}_{n,p}$ are matrices. If $\mathbf{b}_j$ is the j-th column of B, then

$$AB = \left[\, A\mathbf{b}_1 \mid A\mathbf{b}_2 \mid \cdots \mid A\mathbf{b}_p \,\right].$$

Proof. Again, we perform block matrix multiplication:

■

Linear Algebra – Week 4

Linear Transformations

This week we will learn about:

- Understanding linear transformations geometrically,
- The standard matrix of a linear transformation, and
- Composition of linear transformations.

Extra reading and watching:

- Section 1.4 in the textbook
- Lecture videos 13, 14, 15, and 16 on YouTube
- Linear map at Wikipedia

Extra textbook problems:

⋆ 1.4.1, 1.4.4, 1.4.5(a,b,e,f)

⋆⋆ 1.4.2, 1.4.3, 1.4.6, 1.4.7(a,b), 1.4.8, 1.4.14–1.4.16

⋆⋆⋆ 1.4.18, 1.4.22, 1.4.23

☠ 1.4.19, 1.4.20

Linear Transformations

The final main ingredient of linear algebra, after vectors and matrices, are linear transformations: functions that act on vectors and that do not "mess up" vector addition and scalar multiplication:

Definition 4.1 — Linear Transformations

A **linear transformation** is a function $T : \mathbb{R}^n \to \mathbb{R}^m$ that satisfies the following two properties:

a) $T(\mathbf{v} + \mathbf{w}) = T(\mathbf{v}) + T(\mathbf{w})$ for all vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$, and

b) $T(c\mathbf{v}) = cT(\mathbf{v})$ for all vectors $\mathbf{v} \in \mathbb{R}^n$ and all scalars $c \in \mathbb{R}$.

Before looking at specific examples of linear transformations, let's think geometrically about what they do to $\mathbb{R}^n$:

Another way of thinking about this: linear transformations are exactly the functions that preserve linear combinations:

Example. Which of the following functions are linear transformations?

Recall that every vector $\mathbf{v} = (v_1, v_2, \ldots, v_n) \in \mathbb{R}^n$ can be written in the form

By using the fact that linear transformations preserve linear combinations, we see that

But this is exactly what we said before: if $\mathbf{v} \in \mathbb{R}^2$ extends a distance of v_1 in the direction of $\mathbf{e}_1$ and a distance of v_2 in the direction of $\mathbf{e}_2$, then $T(\mathbf{v})$ extends the same amounts in the directions of $T(\mathbf{e}_1)$ and $T(\mathbf{e}_2)$, respectively.

This also tells us one of the most important facts to know about linear transformations:

Example. Suppose $T : \mathbb{R}^2 \to \mathbb{R}^2$ is a linear transformation for which $T(\mathbf{e}_1) = (1, 1)$ and $T(\mathbf{e}_2) = (-1, 1)$. Compute $T(2, 3)$ and then find a general formula for $T(v_1, v_2)$

One of the earlier examples showed that if $A \in \mathcal{M}_{m,n}$ is a matrix, then the function $T : \mathbb{R}^m \to \mathbb{R}^n$ defined by $T(\mathbf{v}) = A\mathbf{v}$ is a linear transformation. Amazingly, the converse is also true: *every* linear transformation can be written as matrix multiplication.

Theorem 4.1 — Standard Matrix of a Linear Transformation

A function $T : \mathbb{R}^n \to \mathbb{R}^m$ is a linear transformation if and only if there exists a matrix $[T] \in \mathcal{M}_{m,n}$ such that

$$T(\mathbf{v}) = [T]\mathbf{v} \quad \text{for all} \quad \mathbf{v} \in \mathbb{R}^n.$$

Furthermore, the unique matrix $[T]$ with this property is called the **standard matrix** of T, and it is

Proof. We already proved the "if" direction, so we just need to prove the "only if" direction. That is, we want to prove that if T is a linear transformation, then $T(\mathbf{v}) = [T]\mathbf{v}$, where the matrix $[T]$ is as defined in the theorem.

∎

Example. Find the standard matrix of the following linear transformations:

A Catalog of Linear Transformations

To get more comfortable with the relationship between linear transformations and matrices, let's find the standard matrices of a few linear transformations that come up fairly frequently.

Example. The zero and identity transformations.

Example. Diagonal transformations/matrices.

Example. Projection onto the x-axis.

Example. Projection onto a line, $P_{\mathbf{u}} : \mathbb{R}^n \to \mathbb{R}^n$.

Example. Find the standard matrix of the linear transformation that projects $\mathbb{R}^3$ onto the line in the direction of the vector...

Example. Rotation counter-clockwise around the origin by 90° ($\pi/2$ radians).

Example. Rotation $R^\theta : \mathbb{R}^2 \to \mathbb{R}^2$ counter-clockwise around the origin by an angle of θ.

Example. What vector is obtained if we rotate $\mathbf{v} = (1, 3)$ by $\pi/4$ radians counter-clockwise?

Composing Linear Transformations

If $T : \mathbb{R}^n \to \mathbb{R}^m$ and $S : \mathbb{R}^m \to \mathbb{R}^p$ are linear transformations, then we can consider the function defined by first applying T to a vector, and then applying S. This function is called the **composition** of T and S, and is denoted by $S \circ T$.

Formally, the composition $S \circ T$ is defined by $(S \circ T)(\mathbf{v}) = S(T(\mathbf{v}))$ for all vectors $\mathbf{v} \in \mathbb{R}^n$. It turns out that $S \circ T$ is a linear transformation whenever S and T are linear transformations themselves, as shown by the next theorem.

Theorem 4.2 — Composition of Linear Transformations

Suppose $T : \mathbb{R}^n \to \mathbb{R}^m$ and $S : \mathbb{R}^m \to \mathbb{R}^p$ are linear transformations with standard matrices $[T] \in \mathcal{M}_{m,n}$ and $[S] \in \mathcal{M}_{p,m}$, respectively. Then $S \circ T : \mathbb{R}^n \to \mathbb{R}^p$ is a linear transformation, and its standard matrix is $[S \circ T] = [S][T]$.

Proof. Let $\mathbf{v} \in \mathbb{R}^n$ and compute $(S \circ T)(\mathbf{v})$:

■

The previous theorem shows us that matrix multiplication tells us how the composition of linear transformations behaves. In fact, this is exactly why matrix multiplication is defined the way it is.

Example. What vector is obtained if we rotate $\mathbf{v} = (4, 2)$ 45° counter-clockwise around the origin and then project it onto the line $y = 2x$?

Example. Find the standard matrix of the linear transformation T that projects $\mathbb{R}^2$ onto the line $y = (4/3)x$ and then stretches it in the x-direction by a factor of 2 and in the y-direction by a factor of 3.

Example. Derive the angle-sum formulas for sin and cos.

Systems of Linear Equations

This week we will learn about:

- Systems of linear equations,
- Elementary row operations and Gaussian elimination, and
- The (reduced) row echelon form of a matrix.

Extra reading and watching:

- Section 2.1 in the textbook
- Lecture videos 17, 18, 19, 20, and 21 on YouTube
- System of linear equations at Wikipedia
- Gaussian elimination at Wikipedia

Extra textbook problems:

⋆ 2.1.1, 2.1.2, 2.1.4, 2.1.5

⋆⋆ 2.1.7–2.1.9, 2.1.11, 2.1.15–2.1.17, 2.1.25, 2.1.26

⋆⋆⋆ 2.1.18, 2.1.23, 2.1.24, 2.1.27–2.1.29

☠ none this week

(Systems of) Linear Equations

Much of linear algebra is about solving and manipulating the simplest types of equations that exist—linear equations:

Definition 5.1 — Linear Equations

A **linear equation** in n variables $x_1, x_2, \ldots, x_n$ is an equation that can be written in the form

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n = b,$$

where $a_1, a_2, \ldots, a_n$ and b are constants.

Example. Examples of linear and non-linear equations.

The point is that an equation is linear if each variable is only multiplied by a constant: variables cannot be multiplied by other variables, they can only be raised to the first power, and they cannot have other functions applied to them.

You (hopefully) learned how to manipulate linear equations quite some time ago, and then you "ramped up" to non-linear equations (like $x^2 = 2$ or $2^x = 8$). In this course, we instead "ramp up" in a different direction: we work with multiple linear equations simultaneously.

Definition 5.2 — Systems of Linear Equations

A **system of linear equations** (or a **linear system**) is a finite set of linear equations, each with the same variables $x_1, x_2, \ldots, x_n$.

Some more terminology:

- A **solution** of a system of linear equations is a vector $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ whose entries satisfy *all* of the linear equations in the system.
- The **solution set** of a system of linear equations is the set of *all* solutions of the system.

Example. Solving a linear system geometrically.

Example. Two more (weirder!) systems of linear equations.

The above examples show that systems of linear equations can have no solutions, exactly one solution, or infinitely many solutions. We will show shortly that these are the only possibilities.

Note that we can also visualize systems of linear equations with 3 variables in 3 dimensions, but it's a bit tougher:

Matrix Equations

One of the primary uses of matrices is that they give us a way of working with linear systems much more compactly and cleanly. In particular, any system of linear equations...

can be rewritten as a single matrix equation:

Example. Write the following system of linear equations as a single matrix equation:

The advantage of writing linear systems in this way (beyond the fact that it requires less writing) is that we can now make use of the various properties of matrices and matrix multiplication that we already know to help us understand linear systems a bit better. For example, we can now prove the observation that we made earlier: every linear system has either 0, 1, or infinitely many solutions.

Theorem 5.1 — Trichotomy for Linear Systems

Every system of linear equations has either

a) no solutions,

b) exactly one solution, or

c) infinitely many solutions.

Proof. We just need to show that if a linear system has at least two different solutions, then it has infinitely many solutions.

■

When a system of linear equations has at least one solution (i.e., in cases (b) and (c) of the theorem), it is called **consistent**. If it has no solutions (i.e., in case (a) of the theorem), it is called **inconsistent**.

Solving Linear Systems

Let's now discuss how we might find the solutions of a system of linear equations. If the linear system has a certain special form, then solving it is fairly intuitive.

Example. Solve the following system of linear equations:

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

The procedure that we used to solve the previous example is called **back substitution**, and it worked because of the "triangular" nature of the equations. We were able to easily solve for z, which we then could plug into the second equation and easily solve for y, which we could plug into the first equation and easily solve for x.

So let's try to put *every* system of equations into this triangular form! We start by

To reduce the amount of writing we have to do when solving the linear system $A\mathbf{x} = \mathbf{b}$, we typically use the block matrix $[\,A \mid \mathbf{b}\,]$.

Example. Solve the following (much uglier) system of linear equations:

$$\begin{aligned} x + 3y - 2z &= 5 \\ 3x + 5y + 6z &= 7 \\ 2x + 4y + 3z &= 8 \end{aligned}$$

There were three basic types of operations that we performed on the matrix when solving the previous system of linear equations. These are called the **elementary row operations**:

a) Adding a multiple of a row to another row ($R_i + cR_j$).

b) Multiplying a row by a non-zero constant (cR_i).

c) Interchanging rows ($R_i \leftrightarrow R_j$).

These are the only operations we will ever need to solve a linear system!

As mentioned before, our goal when solving these systems of equations is to first make the matrix "triangular." We now make this a bit more precise.

Definition 5.3 — (Reduced) Row Echelon Form

A matrix is in **row echelon form** if it satisfies both of these properties:

a) All rows consisting entirely of zeros are below the non-zero rows.

b) In each non-zero row, the first non-zero entry (called the **leading entry**) is to the left of any leading entries below it.

If the matrix also satisfies the following additional constraints, then it is in **reduced row echelon form** (**RREF**):

c) The leading entry in each non-zero row is 1.

d) Each leading 1 is the only non-zero entry in its column.

Example. Some matrices that are and are not in (reduced) row echelon form.

To solve a system of linear equations, we use elementary row operations to bring it into row echelon form. Once it is in this form, we can easily solve it via back-substitution.

Alternatively, we can use elementary row operations to bring a matrix all the way into reduced row echelon form. Once an augmented matrix is in this form, the solutions of the associated linear system can be read directly from the entries of the matrix.

Example. Find the solutions of the systems of equations represented by the following augmented matrices:

The process of using elementary row operations to bring a matrix into a row echelon form is called **row reduction**. The process of using row reduction to find a row echelon form, and then back substitution to solve the system of linear equations, is called **Gaussian elimination**.

Example. Use Gaussian elimination to solve the following system of linear equations:

$$\begin{aligned} x + 2y - 4z &= -4 \\ 2x + 4y \phantom{{}-4z} &= 0 \\ -x + y + 3z &= 6 \end{aligned}$$

Some notes about row echelon form and elementary row operations are in order:

- The elementary row operations are reversible: if there is an elementary row operation that transforms A into B, then there is an elementary row operation that transforms B into A.
- Is the row echelon form of a matrix **unique** or **not unique**?
- Two matrices are called **row equivalent** if one can be converted to the other via elementary row operations.

The process of using row reduction to find a *reduced* row echelon form, and hence solve the system of linear equations, is called **Gauss–Jordan elimination**.

Example. Use Gauss–Jordan elimination to solve the following linear system:

$$\begin{aligned} x + 2y - 4z &= -4 \\ 2x + 4y \phantom{{}- 4z} &= 0 \\ -x + y + 3z &= 6 \end{aligned}$$

Some notes about reduced row echelon form and Gauss–Jordan elimination are in order:

- Neither Gaussian elimination nor Gauss–Jordan elimination is a "better" method than the other. Which one you use is typically just based on personal preference.
- Is the reduced row echelon form of a matrix **unique** or **not unique**?
- To check if two matrices are row equivalent, check whether or not they have the same reduced row echelon form.

Free Variables and Systems Without Unique Solutions

Recall that systems of linear equations do not always have a unique solution: they might have no solutions or infinitely many solutions. Identifying systems with no solutions is intuitive enough...

Example. Solve the following system of linear equations:

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

The behaviour in the previous example is what happens in general: a linear system has no solutions if and only if the row echelon forms of its augmented matrix $[\,A \mid \mathbf{b}\,]$ have a row consisting of zeros in the left (A) block and a non-zero entry in the right ($\mathbf{b}$) block.

Things are somewhat more complicated when a system of equations has infinitely many solutions, though. After all, how can we even *describe* all of the solutions in this case? We illustrate the method with a couple more examples:

Example. Solve the following system of linear equations:

$$\begin{aligned} v - 2w \qquad\qquad + 2z &= 3 \\ x \qquad - 3z &= 7 \\ y + \;\; z &= 4 \end{aligned}$$

Example. Solve the following system of linear equations:

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

Again, the behaviour in the previous example is completely general: variables corresponding to columns that have a leading entry in the row echelon form are called **leading variables**, and we write these variables in terms of the non-leading variables (called **free variables**).

Each free variable corresponds to one "dimension" or "degree of freedom" in the solution set. For example, if there is one free variable then the solution set is a line, if there are two then it is a plane, and so on.

Elementary Matrices and Inverses

This week we will learn about:

- Elementary matrices,
- The inverse of a matrix, and
- How awesome inverses are.

Extra reading and watching:

- Section 2.2 in the textbook
- Lecture videos 22, 23, 24, and 25 on YouTube
- Elementary matrix at Wikipedia
- Invertible matrix at Wikipedia

Extra textbook problems:

⋆ 2.2.1, 2.2.2

⋆⋆ 2.2.4–2.2.6, 2.2.8, 2.2.9, 2.2.13, 2.2.15, 2.2.20

⋆⋆⋆ 2.2.7, 2.2.10, 2.2.11, 2.2.21, 2.2.22

☠ 2.2.23

Elementary Matrices

Last week, we learned how to solve systems of linear equations by repeatedly applying one of three row operations to the augmented matrix associated with that linear system. Remarkably, all three of those row operations can be carried out by matrix multiplication (on the left) by carefully-chosen matrices.

For example, if we wanted to swap the first and second rows of the matrix

we could multiply it on the left by the matrix

Similarly, to perform the row operations and
we could multiply on the left by the matrices

Matrices that implement one of these three row operations in this way have a name:

Definition 6.1 — Elementary Matrices

A square matrix $A \in \mathcal{M}_n$ is called an **elementary matrix** if it can be obtained from the identity matrix via a single row operation.

For example, the elementary matrix corresponding to the "Swap" row operation $R_i \leftrightarrow R_j$ looks like

Similarly, the elementary matrices corresponding to the "Addition" row operation $R_i + cR_j$ and the "Multiplication" row operation cR_i look like

Notice that if the elementary matrices $E_1, E_2, \ldots, E_k$ are used to row reduce a matrix A to its reduced row echelon form R, then

In particular, $E_1, E_2, \ldots, E_k$ act as a log that keeps track of which row operations should be performs to put A into RREF. Furthermore, if we define $E = E_k \cdots E_2 E_1$, then $R = EA$, so E acts as a condensed version of that log. Let's now do an example to see how to construct this matrix E.

Example. Let $A =$

Find a matrix E such that $EA = R$, where R is the RREF of A.

The fact that the method of the previous example works in general can be seen by combining some block matrix multiplication trickery with the fact that multiplication on the left by an elementary matrix is equivalent to performing the corresponding row operation. In particular, if row reducing $[\,A \mid I\,]$ to some other matrix $[\,R \mid E\,]$ makes use of the row operations corresponding to elementary matrices $E_1, E_2, \ldots, E_k$, then

This means (by looking at the right half of the above block matrix) that $E = E_k \cdots E_2 E_1$, which then implies (by looking at the left half of the block matrix) that $R = EA$. We state this observation as a theorem:

Theorem 6.1 — Row Reduction is Multiplication on the Left

If the block matrix $[\,A \mid I\,]$ can be row reduced to $[\,R \mid E\,]$ then...

This theorem says that, not only is performing a single row operation equivalent to multiplication on the left by an elementary matrix, but performing a *sequence* of row operations is also equivalent to multiplication on the left (by some potentially non-elementary matrix).

The Inverse of a Matrix

When working with (non-zero) real numbers, we have an operation called "division," which acts as an inverse of multiplication. In particular, $a(1/a) = 1$ for all $a \neq 0$. It turns out that we can (usually) do something very similar for matrix multiplication:

Definition 6.2 — Inverse of a Matrix

If A is a square matrix, the **inverse** of A, denoted by A^{-1}, is a matrix (of the same size as A) with the property that

If such a matrix A^{-1} exists, then A is called **invertible**.

Inverses (when they exist) are unique (i.e., every matrix has at most one inverse). To see this...

Example. Show that $\begin{bmatrix} 3 & -5 \\ -1 & 2 \end{bmatrix}$ is the inverse of $\begin{bmatrix} 2 & 5 \\ 1 & 3 \end{bmatrix}$.

So if we are given a particular pair of matrices, it is easy to check whether or not they are inverses of each other. But how could we *find* the inverse of a matrix in the first place? We'll see how soon!

As always, let's think about what properties our new mathematical operation (matrix inversion) has.

Theorem 6.2 — Properties of Matrix Inverses

Let A and B be invertible matrices of the same size, and let c be a non-zero real number. Then

a) A^{-1} is invertible and $(A^{-1})^{-1} = A$

b) cA is invertible and $(cA)^{-1} = \frac{1}{c}A^{-1}$

c) A^T is invertible and $(A^T)^{-1} = (A^{-1})^T$

d) AB is invertible and $(AB)^{-1} = B^{-1}A^{-1}$

Proof. Most parts of this theorem are intuitive enough, so we just prove part (d) (you can prove parts (a), (b) and (c) on your own: they're similar). To this end...

∎

The fact that $(AB)^{-1} = B^{-1}A^{-1}$ (as opposed to the incorrect $(AB)^{-1} = A^{-1}B^{-1}$) is actually intuitive enough: you put on your socks before your shoes, but when reversing that operation, you take off your shoes before your socks.

Not every matrix is invertible. For example,

However, there are even more exotic examples of non-invertible matrices. For example, recall that if $\mathbf{u}$ is a unit vector then the matrix $A = \mathbf{u}\mathbf{u}^T$...

In order to come up with a general method for determining whether or not a matrix is invertible (and constructing its inverse if it exists), we first notice that if A has reduced row echelon form equal to I, then Theorem 6.1 tells us that

It thus seems like A being invertible is closely related to whether or not it can be row reduced to the identity matrix. The following theorem shows that this is indeed the case (along with a whole lot more):

Theorem 6.3 — Characterization of Invertible Matrices

Let $A \in \mathcal{M}_n$. The following are equivalent:

a) A is invertible.

b) The reduced row echelon form of A is I (the identity matrix).

c) There exist elementary matrices $E_1, E_2, \ldots, E_k$ such that $A = E_1 E_2 \cdots E_k$.

d) The linear system $A\mathbf{x} = \mathbf{b}$ has a solution for all $\mathbf{b} \in \mathbb{R}^n$.

e) The linear system $A\mathbf{x} = \mathbf{b}$ has a unique solution for all $\mathbf{b} \in \mathbb{R}^n$.

f) The linear system $A\mathbf{x} = \mathbf{0}$ has a unique solution.

Example. Determine whether or not the matrix $A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix}$ is invertible.

We won't rigorously prove the above theorem, but we'll try to give a rough idea for why some of its equivalences hold. First, notice that every elementary matrix is invertible:

Since the product of invertible matrices is still invertible, it follows that any matrix of the form $A = E_1E_2 \cdots E_k$ (where $E_1, E_2, \ldots, E_k$ are elementary) is invertible, which shows why (c) $\implies$ (a).

The connection between invertibility and linear systems can be clarified by noting that if A is invertible, then we can rearrange the linear system

Thus (a) $\implies$ (e), which implies each of (d) and (f).

When we combine our previous two theorems, we get a method for not only determining whether or not a matrix is invertible, but also for computing its inverse if it exists:

Theorem 6.4 — How to Compute Inverses

A matrix $A \in \mathcal{M}_n$ is invertible if and only if the RREF of $[\, A \mid I \,]$ has the form $[\, I \mid B \,]$ for some $B \in \mathcal{M}_n$. If the RREF has this form then $A^{-1} = B$.

Example. Determine whether or not the matrix $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$ is invertible, and find its inverse if it exists.

Example. Solve the linear system $x + 2y = 3$, $3x + 4y = 5$.

Example. Find the inverse of $\begin{bmatrix} 1 & 2 & -1 \\ 2 & 2 & 4 \\ 1 & 3 & -3 \end{bmatrix}$ if it exists.

Example. Find the inverse of $\begin{bmatrix} 2 & 1 & -4 \\ -4 & -1 & 6 \\ -2 & 2 & -2 \end{bmatrix}$ if it exists.

Using our characterization of invertible matrices, we can prove all sorts of nice properties of them. For example, even though the definition of invertibility required that both $AA^{-1} = I$ and $A^{-1}A = I$, the following theorem shows that it is enough to just multiply on the left *or* the right: you don't need to check both.

Theorem 6.5 — One-Sided Matrix Inverses

Let $A \in \mathcal{M}_n$ be a square matrix. If $B \in \mathcal{M}_n$ is a matrix such that either $AB = I$ or $BA = I$, then A is invertible and $A^{-1} = B$.

Proof. Suppose $BA = I$, and consider the equation $A\mathbf{x} = \mathbf{0}$.

This completes the proof of the $BA = I$ case. Try to prove the case when $AB = I$ on your own. ■

Similarly, we can even come up with an explicit formula for the inverse of matrices in certain small cases. For example, for 2×2 matrices, we have the following formula:

Theorem 6.6 — Inverse of a 2×2 Matrix

Suppose A is the 2×2 matrix

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

Then A is invertible if and only if $ad - bc \neq 0$, and if it is invertible then

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

Proof. If $ad - bc \neq 0$ then we can show that the inverse of A is as claimed just by multiplying it by A:

On the other hand, if $ad - bc = 0$ then $ad = bc$.

■

Example. Compute the inverse (or show that none exists) of the following matrices:

Keep in mind that you can always use the general method of computing inverses (row reduce $[\, A \mid I \,]$ to $[\, I \mid A^{-1} \,]$) if you forget this formula for the 2×2 case.

Subspaces, Spans, and Linear Independence

This week we will learn about:

- Subspaces,
- The span of a set of vectors, and
- Linear (in)dependence.

Extra reading and watching:

- Section 2.3 in the textbook
- Lecture videos 26, 27, 28, and 29 on YouTube
- Linear subspace at Wikipedia
- Linear independence at Wikipedia

Extra textbook problems:

⋆ 2.3.1, 2.3.2, 2.3.4

⋆⋆ 2.3.3, 2.3.5, 2.3.6, 2.3.9–2.3.11, 2.3.18, 2.3.19

⋆⋆⋆ 2.3.12, 2.3.14, 2.3.16, 2.3.22

☠ 2.3.27

Subspaces

Recall that linear systems can be interpreted geometrically as asking for the point(s) of intersection of a collection of lines or planes (depending on the number of variables involved). The following definition introduces "subspaces", which can be thought of as any-dimensional analogues of lines and planes.

Definition 7.1 — Subspaces

A **subspace** of $\mathbb{R}^n$ is a non-empty set $\mathcal{S}$ of vectors in $\mathbb{R}^n$ such that:

a) If $\mathbf{v}$ and $\mathbf{w}$ are in $\mathcal{S}$ then $\mathbf{v} + \mathbf{w}$ is in $\mathcal{S}$.

b) If $\mathbf{v}$ is in $\mathcal{S}$ and c is a scalar, then $c\mathbf{v}$ is in $\mathcal{S}$.

Properties (a) and (b) above together are equivalent to requiring that $\mathcal{S}$ is closed under linear combinations:

Example. Is the set of vectors (x, y) satisfying $y = x^2$ a subspace of $\mathbb{R}^2$?

Example. Is the set of vectors (x, y, z) satisfying $x = 3y$ and $z = -2y$ a subspace of $\mathbb{R}^3$?

Example. Is the set of vectors (x, y, z) satisfying $x = 3y + 1$ and $z = -2y$ a subspace of $\mathbb{R}^3$?

In $\mathbb{R}^3$, lines and planes through the origin are subspaces (this is hopefully not difficult to see for lines, and it can be seen for planes by using the parallelogram law):

Even though we can't visualize subspaces in higher dimensions, you should keep the line/plane intuition in mind: a subspace of $\mathbb{R}^n$ looks like a copy of $\mathbb{R}^m$ (for some $m < n$) going through the origin.

Subspaces Associated with Matrices

Let's now look at some other natural examples of subspaces that appear frequently when working with matrices.

Definition 7.2 — Matrix Subspaces

Let $A \in \mathcal{M}_{m,n}$ be an $m \times n$ matrix.

a) The **range** of A is the subspace of $\mathbb{R}^m$, denoted by $\text{range}(A)$, that consists of all vectors of the form $A\mathbf{x}$.

b) The **null space** of A is the subspace of $\mathbb{R}^n$, denoted by $\text{null}(A)$, that consists of all solutions $\mathbf{x}$ of the linear system $A\mathbf{x} = \mathbf{0}$.

Some remarks about these matrix subspaces are in order:

- $\text{null}(A)$ is a subspace. Why?

- $\text{range}(A)$ is a subspace. Why?

- The term "range" is being used here in the exact same sense as in previous courses.

Example. Describe the range and null space of the 2×3 matrix

The Span of a Set of Vectors

One way to turn a set that is *not* a subspace into a subspace is to add linear combinations to it. For example, the set containing only the vector $(2, 1)$ is not a subspace of $\mathbb{R}^2$ because

To fix this problem, we could

In general, if our starting set contains more than just one vector, we might also have to add general linear combinations of those vectors (not just their scalar multiples) in order to create a subspace. This idea of enlarging a set so as to create a subspace is an important one that we now give a name and explore.

Definition 7.3 — Span

If $B = \{\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k\}$ is a set of vectors in $\mathbb{R}^n$, then the set of all linear combinations of those vectors is called their **span**, and is denoted by $\mathrm{span}(B)$ or $\mathrm{span}(\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k)$.

For example, $\mathrm{span}\big((2, 1)\big)$ is the line through the origin and the point $(2, 1)$, as we discussed earlier.

Example. Show that $\operatorname{span}(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3) = \mathbb{R}^3$.

The natural generalization of this fact holds in all dimensions:

Example. What is $\operatorname{span}\big((1, 0, 3), (-1, 1, -3)\big)$ – a line, a plane, or something else?

We motivated the span of a set of vectors as a way of turning that set into a subspace. We now state (but for the sake of time, do not prove) a theorem that says the span of a set of vectors is indeed always a subspace, as we would hope.

Theorem 7.1 — Spans are Subspaces

Suppose $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k \in \mathbb{R}^n$. Then $\operatorname{span}(\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k)$ is a subspace of $\mathbb{R}^n$.

In fact, you can think of the span of a set of vectors as the *smallest* subspace containing those vectors.

The range of a matrix can be expressed very conveniently as the span of a set of vectors in a way that requires no calculation whatsoever:

Theorem 7.2 — Range Equals the Span of Columns

If $A \in \mathcal{M}_{m,n}$ has columns $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n$ then $\mathrm{range}(A) = \mathrm{span}(\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n)$.

This theorem follows immediately from

For example, if we return to the 2×3 matrix from earlier, we see that its range is...

We close this section by introducing a connection between the range of a matrix and invertible matrices.

Theorem 7.3 — Spanning Sets and Invertible Matrices

Let $A \in \mathcal{M}_n$. The following are equivalent:

a) A is invertible.

b) $\mathrm{range}(A) = \mathbb{R}^n$.

c) The columns of A span $\mathbb{R}^n$.

d) The rows of A span $\mathbb{R}^n$.

Proof. The fact that properties (a) and (c) are equivalent follows from combining...

The equivalence of properties (c) and (d) follows from the fact that

Finally, the equivalence of properties (b) and (c) follows immediately from

■

The geometric interpretation of the equivalence of properties (a) and (b) in the above theorem is

Linear Dependence and Independence

Recall from earlier that a row echelon form of a matrix can have entire rows of zeros at the end of it. For example, the reduced row echelon form of

$$\left[\begin{array}{cc|c} 1 & -1 & 2 \\ -1 & 1 & -2 \end{array}\right] \quad \text{is}$$

This happens when there is some linear combination of the rows of the matrix that equals the zero row, and we interpret this roughly as saying that one row the rows of the matrix (i.e., one of the equations in the associated linear system) does not "contribute anything new." In the example above, The following definition captures this idea that a redundancy among vectors or linear equations can be identified by whether or not some linear combination of them equals zero.

Definition 7.4 — Linear Dependence and Independence

A set of vectors $B = \{\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k\}$ is **linearly dependent** if there exist scalars $c_1, c_2, \ldots, c_k \in \mathbb{R}$, at least one of which is not zero, such that

$$c_1\mathbf{v}_1 + c_2\mathbf{v}_2 + \cdots + c_k\mathbf{v}_k = \mathbf{0}.$$

If a set of vectors is not linearly dependent, it is called **linearly independent**.

For example, the set of vectors $\{(2,3),(1,0),(0,1)\}$ is linearly...

On the other hand, the set of vectors $\{(1,0,0),(0,1,0),(0,0,1)\}$ is linearly...

In general, to check whether or not a set of vectors $\{\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k\}$ is linearly independent, you should set

and then try to solve for the scalars $c_1, c_2, \ldots, c_k$. If they must all equal 0, then the set is linearly independent, and otherwise it is linearly dependent.

Example. Are these vectors linearly independent?

We saw in the previous example that we can check linear (in)dependence of a set of vectors by placing those vectors as columns in a matrix and augmenting with a $\mathbf{0}$ right-hand side. This is true in general:

Theorem 7.4 — Checking Linear Dependence

Let $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n \in \mathbb{R}^m$ be vectors and let A be the $m \times n$ matrix with these vectors as its columns. The following are equivalent:

a) $\{\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n\}$ is a linearly dependent set.

b) The linear system $A\mathbf{x} = \mathbf{0}$ has a non-zero solution.

Some notes about linear (in)dependence are in order:

- A set of vectors is linearly dependent if and only if at least one of the vectors can be written as a linear combination of the others.

- Every set of vectors containing the zero vector is linearly...

- Geometrically, linear dependence means that...

- For a set of just 2 vectors, linear dependence means that...

Example. Is this set linearly independent?

We close this section by introducing a connection between linear independence and invertible matrices, which we unfortunately have to state without proof due to time constraints.

Theorem 7.5 — Independence and Invertible Matrices

Let $A \in \mathcal{M}_n$. The following are equivalent:

a) A is invertible.

b) The columns of A form a linearly independent set.

c) The rows of A form a linearly independent set.

Bases of Subspaces and the Rank of a Matrix

This week we will learn about:

- The dimension of a subspace,
- Bases of subspaces,
- The rank of a matrix, and
- The rank–nullity theorem.

Extra reading and watching:

- Sections 2.4 in the textbook
- Lecture videos 30, 31, 32, and 33 on YouTube
- Basis (linear algebra) at Wikipedia
- Rank (linear algebra) at Wikipedia

Extra textbook problems:

⋆ 2.4.1, 2.4.2, 2.4.8

⋆⋆ 2.4.5, 2.4.6, 2.4.9, 2.4.10

⋆⋆⋆ 2.4.11, 2.4.12, 2.4.13, 2.4.25, 2.4.30

☠ 2.4.27

Bases of Subspaces

A plane in $\mathbb{R}^3$ is spanned by any two vectors that are parallel to the plane, but not parallel to each other (i.e., are linearly independent). More than two vectors could be used to span the plane, but they would necessarily be linearly dependent. On the other hand, there is no way to use *fewer* than two vectors to span a plane (the span of just one vector is just a line). This leads to the idea of a *basis* of a subspace:

Definition 8.1 — Bases

A **basis** of a subspace $\mathcal{S} \subseteq \mathbb{R}^n$ is a set of vectors in $\mathcal{S}$ that

a) spans $\mathcal{S}$, and

b) is linearly independent.

The idea of a basis is that it is a set that is "big enough" to span the subspace, but it is not "so big" that it contains redundancies. That is, it is "just" big enough to span the subspace.

Example. The standard basis of $\mathbb{R}^n$.

Example. Show that the set $\{(2,1),(1,3)\}$ is a basis of $\mathbb{R}^2$.

The above example demonstrates that the same subspace can (and will!) have more than one basis:

However, the *number* of vectors in a basis of a given subspace is always the same, which we now state as a theorem.

Theorem 8.1 — Uniqueness of Size of Bases

Let $\mathcal{S}$ be a subspace of $\mathbb{R}^n$. Then every basis of $\mathcal{S}$ has the same number of vectors.

We don't prove the above theorem (it is a fairly long and ugly mess), but we can use it to pin down something we have been hand-wavey about up until now: we have never actually defined exactly what we mean by the "dimension" of a subspace of $\mathbb{R}^n$. We now fill in this gap:

Definition 8.2 — Dimension of a Subspace

Let $\mathcal{S}$ be a subspace of $\mathbb{R}^n$. The number of vectors in a basis of $\mathcal{S}$ is called the **dimension** of $\mathcal{S}$.

As one minor technicality, we notice that the set $\mathcal{S} = \{\mathbf{0}\}$ is a subspace of $\mathbb{R}^n$. However, the only basis of this subspace is the empty set $\{\}$ (why?), so its dimension is 0.

Example. What is the dimension of $\mathbb{R}^n$?

Example. Find a basis for $\mathcal{S} = \text{span}(\mathbf{v}, \mathbf{w}, \mathbf{x})$, where $\mathbf{v} = (1, 2, 3)$, $\mathbf{w} = (3, 2, 1)$, $\mathbf{x} = (1, 1, 1)$. What is the dimension of this subspace?

We will now show how to find bases of the fundamental subspaces associated with a matrix: range(A) and null(A).

Example. Find bases for the range and null space of the following matrix and thus compute their dimensions:

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

The quantities $\dim(\text{range}(A))$ and $\dim(\text{null}(A))$ that we computed in the previous example highlight a lot of the structure of the matrix A, so let's have a closer look at them now.

The Rank of a Matrix

With many of the technical details of this course out of the way, we are now in a position to introduce one of the most important properties of a matrix: its rank.

Definition 8.3 — Rank of a Matrix

Let $A \in \mathcal{M}_{m,n}$ be a matrix. Then its **rank**, denoted by $\text{rank}(A)$, is the dimension of its range.

Rank can be thought of as a measure of how degenerate a matrix is, as it describes how much of the output space can actually be reached by A.

Example. Suppose $A \in \mathcal{M}_n$ is the standard matrix of a projection onto a line. What is $\text{rank}(A)$?

Example. Suppose $C \in \mathcal{M}_2$ is the standard matrix of a rotation. What is $\text{rank}(C)$?

One of the reasons why the rank of a matrix is so useful is that it can be interpreted in so many different ways. While it equals the dimension of the range, it also equals some other quantities that we have already seen as well:

Theorem 8.2 — Characterization of Rank

Let $A \in \mathcal{M}_{m,n}$ be a matrix. Then the following quantities are all equal to each other:

a) $\text{rank}(A)$

b) $\text{rank}(A^T)$.

c) The number of non-zero rows in any row echelon form of A.

d) The number of leading columns in any row echelon form of A.

Proof. To see the equivalence of (c) and (d)...

To see the equivalence of (a) and (d)...

The equivalence of (b) and (c) is similar:

■

Example. Find the rank of the matrix $A = \begin{bmatrix} 0 & 0 & -2 & 2 & -2 \\ 2 & -2 & -1 & 3 & 3 \\ -1 & 1 & -1 & 0 & -3 \end{bmatrix}$

Similarly, the **nullity** of a matrix, denoted by nullity(A), is the dimension of its null space (i.e., the dimension of the solution set of the linear system $A\mathbf{x} = \mathbf{0}$). The following theorem demonstrates the close connection between the rank and nullity of a matrix:

Theorem 8.3 — Rank–Nullity

Let $A \in \mathcal{M}_{m,n}$ be a matrix. Then $\text{rank}(A) + \text{nullity}(A) = n$.

Proof. We use the equivalence of the quantities (a) and (d) from the previous theorem:

■

Example. Find the nullity of the matrix $A = \begin{bmatrix} 0 & 0 & -2 & 2 & -2 \\ 2 & -2 & -1 & 3 & 3 \\ -1 & 1 & -1 & 0 & -3 \end{bmatrix}$

The previous theorem makes some geometric sense—there are n dimensions that go into A. $\mathrm{rank}(A)$ of them are sent to the output space, and the other $\mathrm{nullity}(A)$ of them are "squashed away" by A. This observation leads immediately to *yet another* characterization of invertibility:

Theorem 8.4 — Rank and Invertible Matrices

Let $A \in \mathcal{M}_n$. The following are equivalent:

a) A is invertible.

b) $\mathrm{rank}(A) = n$

c) $\mathrm{nullity}(A) = 0$

Linear Algebra – Week 9

Determinants

This week we will learn about:

- Determinants of matrices, and
- That's it. Determinants, determinants, determinants.

Extra reading and watching:

- Section 3.2 in the textbook
- Lecture videos 34, 35, and 36 on YouTube
- Determinant at Wikipedia

Extra textbook problems:

⋆ 3.2.1, 3.2.3, 3.2.4, 3.2.9

⋆⋆ 3.2.5–3.2.8, 3.2.10, 3.2.12, 3.2.17

⋆⋆⋆ 3.2.14, 3.2.16, 3.2.18

☠ 3.2.19–3.2.21

We now introduce one of the most important properties of a matrix: its **determinant**, which roughly is a measure of how "large" the matrix is. More specifically, recall that...

The determinant of A, which we denote by $\det(A)$, is the area (or volume) of this image of the unit hypercube. In other words, it measures how much A expands space when acting as a linear transformation.

Let's now start looking at some of the properties of determinants, so that we can (eventually!) learn how to compute it.

Definition and Basic Properties

Before we even properly define the determinant, let's think about some properties that it should have. The first important property is that, since the identity matrix does not stretch or shrink $\mathbb{R}^n$ at all...

Next, since every $A \in \mathcal{M}_n$ expands space by a factor of $\det(A)$, and similarly each $B \in \mathcal{M}_n$ expands space by a factor of $\det(B)$...

We will also need one more property of determinants, which is a bit more difficult to see. What happens to $\det(A)$ if we multiply one of the columns of A by a scalar $c \in \mathbb{R}$?

Similarly, if we add a vector to one of the columns of a matrix, then...

In other words, the determinant is linear in the columns of a matrix (sometimes called **multilinearity**). We now *define* the determinant to be the function that satisfies this multilinearity property, as well as the other two properties that we demonstrated earlier:

Definition 9.1 — Determinant

The **determinant** is the (unique!) function $\det : \mathcal{M}_n \to \mathbb{R}$ that satisfies the following three properties:

a) $\det(I) = 1$,

b) $\det(AB) = \det(A)\det(B)$ for all $A, B \in \mathcal{M}_n$, and

c) for all $c \in \mathbb{R}$ and all $\mathbf{v}, \mathbf{w}, \mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n \in \mathbb{R}^n$, it is the case that

$$\det\big([\ \mathbf{a}_1 \mid \cdots \mid \mathbf{v} + c\mathbf{w} \mid \cdots \mid \mathbf{a}_n\ ]\big) = \det\big([\ \mathbf{a}_1 \mid \cdots \mid \mathbf{v} \mid \cdots \mid \mathbf{a}_n\ ]\big) + c \cdot \det\big([\ \mathbf{a}_1 \mid \cdots \mid \mathbf{w} \mid \cdots \mid \mathbf{a}_n\ ]\big).$$

Example. Compute the determinant of the matrix $A = \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix}$.

Let's start looking at some of the basic properties of the determinant. First, if $A \in \mathcal{M}_n$ is invertible then properties (a) and (b) tell us that

This makes sense geometrically, since if A expands space by a factor of $\det(A)$ then

On the other hand, if A is not invertible, then

We summarize our observations about the determinant of invertible and non-invertible matrices in the following theorem:

Theorem 9.1 — Determinants and Invertibility

Suppose $A \in \mathcal{M}_n$. Then A is invertible if and only if $\det(A) \neq 0$, and if it is invertible then $\det(A^{-1}) = 1/\det(A)$.

There are also a few other basic properties of determinants that are useful to know, so we state them here (but for time reasons we do not explicitly prove them):

Theorem 9.2 — Other Properties of the Determinant

Suppose $A \in \mathcal{M}_n$ and $c \in \mathbb{R}$. Then

a) $\det(cA) = c^n \det(A)$, and

b) $\det(A^T) = \det(A)$.

Example. Suppose $A, B \in \mathcal{M}_3$ are matrices with $\det(A) = 2$ and $\det(B) = 5$. Compute...

Computation

In order to come up with a general method of computing the determinant, we start by computing it on elementary matrices.

The elementary matrix corresponding to the row operation cR_i has the form

This matrix has determinant equal to...

The elementary matrix corresponding to the row operation $R_i + cR_j$ has the form

This matrix has determinant equal to...

The elementary matrix corresponding to the row operation $R_i \leftrightarrow R_j$ has the form

This matrix has determinant equal to...

Wait, so the determinant of a matrix can be ***negative***? But it measures area/volume!

Since multiplication on the left by an elementary matrix corresponds to performing a row operation, we can rephrase our above calculations as the following theorem:

Theorem 9.3 — Computing Determinants via Row Operations

Suppose $A, B \in \mathcal{M}_n$. If B is obtained from A via a single row operation, then their determinants are related as follows:

$$cR_i\textbf{:}\ \det(B) = c \cdot \det(A),$$

$$R_i + cR_j\textbf{:}\ \det(B) = \det(A), \text{ and}$$

$$R_i \leftrightarrow R_j\textbf{:}\ \det(B) = -\det(A).$$

The above theorem gives us everything we need to know to be able to compute determinants in general – row reduce A to I, keeping track of the row operations that we performed along the way, and use the fact that $\det(I) = 1$. If we cannot row reduce to I, then A is not invertible, so $\det(A) = 0$.

Example. Compute the determinant of $A = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{bmatrix}$.

In the previous example, the determinant of the row echelon form ended up being the product of its diagonal entries. We now state this observation as a theorem:

Theorem 9.4 — Determinant of a Triangular Matrix

Let $A \in \mathcal{M}_n$ be a triangular matrix. Then $\det(A)$ is the product of its diagonal entries:

$$\det(A) = a_{1,1} a_{2,2} \cdots a_{n,n}.$$

Proof. The idea is that a triangular matrix can be row-reduced to I just by operations of the form $R_i + cR_j$ (which do not affect the determinant) and $(1/a_{1,1})R_1$, …, $(1/a_{n,n})R_n$:

∎

By using this fact, we can compute determinants a bit more quickly, by just row-reducing to row echelon form (instead of *reduced* row echelon form). This method is best illustrated with another example.

Example. Compute the determinant of $A = \begin{bmatrix} 1 & 2 & 3 & 0 \\ 1 & 1 & 2 & 3 \\ 0 & -1 & 1 & 0 \\ 0 & -1 & 3 & 0 \end{bmatrix}$.

Explicit Formulas and Cofactor Expansions

Remarkably, the determinant can be computed via an explicit formula just in terms of multiplication and addition of the entries of the matrix. Before presenting the general formula for $n \times n$ matrices, let's start with what it looks like for 2×2 matrices.

Theorem 9.5 — Determinant of 2×2 Matrices

The determinant of a 2×2 matrix is given by

$$\det\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) = ad - bc.$$

Proof. We prove this theorem by making use of multilinearity (i.e., defining property (c) of the determinant):

Well,

Adding these two quantities together gives the desired formula. ■

The above theorem is perhaps best remembered in terms of diagonals of the matrix – the determinant of a 2×2 matrix is the product of its forward diagonal minus the product of its backward diagonal.

Example. Compute the determinant of $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$.

The formula for the determinant of a 3×3 matrix is somewhat more complicated:

Theorem 9.6 — Determinant of 3×3 Matrices

The determinant of a 3×3 matrix is given by

$$\det\left(\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}\right) = aei + bfg + cdh - afh - bdi - ceg.$$

Proof. Again, we make use of multilinearity (i.e., defining property (c) of the determinant) to write

Let's compute the first of the three determinants on the right by using a similar trick on its second column:

Well, these two determinants are

The computation of the remaining terms in the determinant is similar. ■

We can also think of the formula for determinants of 3×3 matrices in terms of diagonals of the matrix – it is the sum of the products of its forward diagonals minus the sum of the products of its backward diagonals, with the understanding that the diagonals "loop around" the matrix:

Example. Compute the determinant of $A = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{bmatrix}$.

The following theorem tells us how to come up with these formulas in general, and it is just a direct generalization of the 2×2 and 3×3 formulas that we already saw.

Theorem 9.7 — Cofactor Expansion

Let $A \in \mathcal{M}_n$. For each $1 \leq i, j \leq n$, define $c_{i,j} = (-1)^{i+j} \det(\overline{A_{i,j}})$, where $\overline{A_{i,j}}$ is the matrix obtained by removing the i-th row and j-th column of A. Then the determinant of A can be computed via

$$\det(A) = a_{i,1}c_{i,1} + a_{i,2}c_{i,2} + \cdots + a_{i,n}c_{i,n} \quad \text{for all } 1 \leq i \leq n, \quad \text{and}$$
$$\det(A) = a_{1,j}c_{1,j} + a_{2,j}c_{2,j} + \cdots + a_{n,j}c_{n,j} \quad \text{for all } 1 \leq j \leq n.$$

That theorem is a mouthful! Several remarks are in order:

- If we use this theorem to compute the determinant of a 2×2 or 3×3 matrix,

- The above method of computing the determinant is called a "cofactor expansion," since the number $c_{i,j}$ is called the "(i,j)-cofactor of A."

- The theorem gives *multiple* formulas for $\det(A)$:

Example. Compute the determinant of $A = \begin{bmatrix} 2 & 1 & -1 & 0 \\ 0 & -2 & 1 & 3 \\ 0 & 0 & 1 & 0 \\ -1 & 3 & 0 & 2 \end{bmatrix}$.

Example. Compute the determinant of $A = \begin{bmatrix} 0 & -1 & 2 & 1 & 3 \\ 0 & 0 & 0 & 2 & 0 \\ -2 & 1 & 1 & -1 & 0 \\ 1 & 0 & -3 & 1 & 0 \\ 2 & 1 & -1 & 0 & 0 \end{bmatrix}$.

In general, computing determinants via cofactor expansions is extremely inefficient. It's not too bad for 2×2, 3×3, or maybe 4×4 matrices. But for an $n \times n$ matrix A, a cofactor expansion contains $n!$ terms being added up, and each of those terms is the product of n entries of A. For example,

$$\det\left(\begin{bmatrix} a & b & c & d \\ e & f & g & h \\ i & j & k & \ell \\ m & n & o & p \end{bmatrix}\right) = afkp - af\ell o - agjp + ag\ell n + ahjo - ahkn$$
$$- bekp + be\ell o + bgip - bg\ell m - bhio + bhkm$$
$$+ cejp - ce\ell n - cfip + cf\ell m + chin - chjm$$
$$- dejo + dekn + dfio - dfkm - dgin + dgjm.$$

Ugh! So for large matrices, use the Gaussian elminiation method instead. Nonetheless, cofactor expansions will be useful for us for theoretical reasons next week.

Eigenvalues and Eigenvectors

This week we will learn about:

- Complex numbers,
- Eigenvalues, eigenvectors, and eigenspaces,
- The characteristic polynomial of a matrix, and
- Algebraic and geometric multiplicity.

Extra reading:

- Section 3.3 in the textbook
- Lecture videos 37, 38, and 39 on YouTube
- Complex number at Wikipedia
- Eigenvalues and eigenvectors at Wikipedia

Extra textbook problems:

⋆ 3.3.1, 3.3.2

⋆⋆ 3.3.3, 3.3.5, 3.3.7, 3.3.9, 3.3.16, 3.3.20

⋆⋆⋆ 3.3.6, 3.3.11–3.3.14

☠ 3.3.19, 3.3.23, 3.3.24

Eigenvalues and Eigenvectors

Some linear transformations behave very well when they act on certain specific vectors. For example, diagonal matrices behave very well on the standard basis vectors:

In the above example, we saw that there are vectors such that matrix multiplication behaved just like scalar multiplication: $A\mathbf{v} = \lambda\mathbf{v}$. This is extremely desirable in many situations: we often want matrix multiplication to behave like scalar multiplication, and we often want general matrices to behave like diagonal matrices. This leads to the following definition.

Definition 10.1 — Eigenvalues and Eigenvectors

Let A be a square matrix. A scalar λ is called an **eigenvalue** of A if there is a non-zero vector $\mathbf{v}$ such that $A\mathbf{v} = \lambda\mathbf{v}$. Such a vector $\mathbf{v}$ is called an **eigenvector** of A corresponding to λ.

Example. Show that $\mathbf{v} = (1, 1)$ is an eigenvector of $A = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$, and find the corresponding eigenvalue.

OK, how do we go about actually *finding* eigenvalues and eigenvectors? It's easy enough when the eigenvector is given to us, but the real world isn't that nice.

Well, we find them via a two-step process: first, we find the eigenvalues, then we find the eigenvectors.

Step 1: Find the eigenvalues. Recall that λ is an eigenvalue of A if and only if there is a non-zero vector $\mathbf{v}$ such that $A\mathbf{v} = \lambda\mathbf{v}$. This is equivalent to...

In other words, λ is an eigenvalue of A if and only if the matrix $A - \lambda I$ has non-zero null space. How can we find when a matrix has a non-zero null space? Well...

- $\dim(\text{null}(A - \lambda I)) > 0$ if and only if...

- ...if and only if...

A-ha! This is the type of equation we can actually solve! So to find the eigenvalues of A, we find all numbers λ such that $\det(A - \lambda I) = 0$.

Example. Find all eigenvalues of $A = \begin{bmatrix} 1 & 2 \\ 5 & 4 \end{bmatrix}$.

Step 2: Find the eigenvectors. Once you know the eigenvalues (from step 1), the associated eigenvectors are the vectors $\mathbf{v}$ satisfying $A\mathbf{v} = \lambda\mathbf{v}$. But this equation holds if and only if...

In other words, to find all eigenvectors of A associated with the eigenvalue λ, we compute $\text{null}(A - \lambda I)$.

Example. Find all eigenvectors of $A = \begin{bmatrix} 1 & 2 \\ 5 & 4 \end{bmatrix}$.

The eigenvalues and eigenvectors can help us understand what a linear transformation "looks like."

Since the set of all eigenvectors of A corresponding to λ is the nullspace of $A-\lambda I$, the set of eigenvectors forms a subspace (the nullspace is always a subspace). We give this subspace a name:

Definition 10.2 — Eigenspace

Let A be a square matrix and let λ be an eigenvalue of A. The set of all eigenvectors of A corresponding to λ, together with the zero vector, is called the **eigenspace** of λ.

Example. Find all eigenvalues, and bases of their corresponding eigenspaces, for the matrix $A = \begin{bmatrix} 2 & -3 & 3 \\ 0 & -1 & 3 \\ 0 & -2 & 4 \end{bmatrix}$.

There are some matrices that do not have any (real) eigenvalues or eigenvectors. For example...

Example. Find all eigenvalues and eigenvectors of $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$.

Complex Numbers

There are a few operations in your mathematical career that you have been told you cannot do:

We now introduce something called **complex numbers** that let us "fix" one of these "problems": they let us work with square roots of negative numbers algebraically.

Remarkably, you can do arithmetic with i just like you're used to with real numbers, and things have a way of just working out. But first, let's get some terminology out of the way:

- An **imaginary number** is a number of the form

- A **complex number** is a number of the form

Arithmetic with complex numbers works just like it does with real numbers, so nothing surprising happens when you add or multiply them.

Example. Add and multiply some complex numbers.

Slightly more generally,

$(a + bi) + (c + di) =$

$(a + bi)(c + di) =$

However, division of complex numbers requires one minor "trick" to get our hands on.

Example. Divide some complex numbers.

The number that we multiplied the top and bottom by in the above example was called the **complex conjugate** of the bottom (denominator). That is,

With just these basic tools under our belt, we can now find roots of quadratics that don't have real roots! We just use the quadratic formula like usual.

Example. Find the (potentially complex) solutions of the equation $x^2 - 2x + 2 = 0$.

The previous example hints at the following observation, which is indeed true:

Just like we think of $\mathbb{R}$ as a line, we can think of $\mathbb{C}$ as a plane, and the number $a + bi$ has coordinates (a, b) on that plane.

Example. Find all eigenvalues and eigenvectors of $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$ (again).

Back to Eigenvalues and Eigenvectors

Recall that the eigenvalues of a matrix A are the solutions λ to the equation $\det(A - \lambda I) = 0$. This is a polynomial in λ, and we give it a special name:

Definition 10.3 — Characteristic Polynomial

Let A be a square matrix. Then $\det(A - \lambda I)$ is called the **characteristic polynomial** of A, and $\det(A - \lambda I) = 0$ is called the **characteristic equation** of A.

The characteristic polynomial of an $n \times n$ matrix is always of degree n. Since every degree-n polynomial has at most n distinct roots, this immediately tells us that

Example. Find the characteristic polynomial, eigenvalues, and bases of the corresponding eigenspaces, of $A = \begin{bmatrix} 1 & -1 & -1 \\ 2 & 2 & 1 \\ -1 & 1 & 2 \end{bmatrix}$.

In the previous example, we had a 3×3 matrix with only 2 distinct eigenvalues. However, the matrix has 3 eigenvalues if we count the multiplicities of the roots of the characteristic polynomial: the eigenvalue $\lambda = 1$ once and the eigenvalue $\lambda = 2$ twice.

There is actually another notion of multiplicity of an eigenvalue that is also important: the dimension of the corresponding eigenspace. These ideas lead to the following definition:

Definition 10.4 — Multiplicity

Let A be a square matrix with eigenvalue λ.

- The **algebraic multiplicity** of λ is the multiplicity of λ as a root of the characteristic polynomial of A.
- The **geometric multiplicity** of λ is the dimension of its eigenspace.

In the previous example...

The fact that the geometric multiplicity of each eigenvalue was $\leq$ the algebraic multiplicity was not a coincidence: it is our next theorem.

Theorem 10.1 — Geo. Mult. $\leq$ Alg. Mult.

Let A be a square matrix. Then the geometric multiplicity of each eigenvalue is less than or equal to its algebraic multiplicity.

A remarkable fact called the Fundamental Theorem of Algebra says that every polynomial of degree n has *exactly* n roots, counted according to multiplicity. This immediately tells us that...

Example. Compute the algebraic and geometric multiplicities of the eigenvalues of all matrices that we considered this week.

Just like with determinants, our eigenvalue life becomes much easier when dealing with triangular matrices.

Example. Compute the eigenvalues of the matrix $\begin{bmatrix} 1 & 2 & 3 \\ 0 & 4 & 5 \\ 0 & 0 & 6 \end{bmatrix}$.

In general, because the determinant of a triangular matrix is just the product of its diagonal entries, the eigenvalues of a triangular matrix are exactly its diagonal entries:

Theorem 10.2 — Eigenvalues of Triangular Matrices

Let A be a triangular matrix. Its eigenvalues are exactly the entries on its main diagonal (i.e., $a_{1,1}, a_{2,2}, \ldots, a_{n,n}$).

Diagonalization

This week we will learn about:

- Diagonalization of matrices,
- Matrix functions, and
- Why diagonalization is amazing.

Extra reading:

- Section 3.4 in the textbook
- Lecture videos 40, 41, 42, 43, and 44 on YouTube
- Diagonalizable matrix at Wikipedia
- Matrix exponential at Wikipedia

Extra textbook problems:

⋆ 3.4.1

⋆⋆ 3.4.2, 3.4.4, 3.4.6, 3.4.7, 3.4.22

⋆⋆⋆ 3.4.8–3.4.12, 3.4.21

☠ 3.4.23

Diagonalization

One of the primary uses of eigenvalues and eigenvectors is that they let us put (most) matrices into a form that makes them almost as easy to work with as diagonal matrices.

Definition 11.1 — Diagonalizable Matrices

A square matrix A is called **diagonalizable** if there is a diagonal matrix D and an invertible matrix P such that $A = PDP^{-1}$.

To get an idea of why diagonalizability is useful, consider the problem of computing a large power of a matrix, like A^{500}.

- If A is a general matrix...

- If A is a diagonal matrix...

- If A is diagonalizable...

But how could we ever hope to determine whether or not a matrix is diagonalizable? It turns out that eigenvalues and eigenvectors give us the answer:

Theorem 11.1 — Diagonalizability

Let A be an $n \times n$ matrix. Then the following are equivalent:

a) A is diagonalizable.

b) A has a set of n linearly independent eigenvectors.

Furthermore, if A is diagonalizable then $A = PDP^{-1}$, where P is the matrix whose columns are the n linearly independent eigenvectors, and D is the diagonal matrix whose diagonal entries are the eigenvalues corresponding to the eigenvectors in P in the same order.

Proof. Start by noticing that the equation $A = PDP^{-1}$ is equivalent to...

■

Example. Diagonalize the matrix $A = \begin{bmatrix} 1 & 2 \\ 5 & 4 \end{bmatrix}$.

OK, now that we've diagonalized this matrix... so what? How does it make our lives easier? Well...

Example. Find a formula for A^n when $A = \begin{bmatrix} 1 & 2 \\ 5 & 4 \end{bmatrix}$.

Example. Find a formula for A^n when $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$.

The previous theorem is nice since it completely characterizes when a matrix is diagonalizable. However, there is one special case that is worth pointing out where it is actually much easier to prove that a matrix is diagonalizable.

Theorem 11.2 — Matrices with Distinct Eigenvalues

Let A be an $n \times n$ matrix with **distinct** eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_n$. Then A is diagonalizable.

Proof. We just need to prove that eigenvectors corresponding to different eigenvalues are linearly independent. To show this...

■

Example. Show that the matrix $A = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$ is diagonalizable.

However, to actually perform the diagonalization itself, we still need to know the eigenvectors.

The Fibonacci Sequence

As an example to demonstrate the usefulness of diagonalization, let's investigate the Fibonacci sequence, which is the sequence of integers that starts as follows:

A bit more properly, it is defined by

My question to you is: can we find a simple closed-form formula for the n-th term of this sequence, without computing all of the previous terms first?

What we will do is represent the Fibonacci sequence via matrix multiplication. Notice that

Well, if we iterate this line of thinking, we get

Matrix multiplication is kinda nasty, but fortunately we already saw that this matrix is diagonalizable, so we can compute large powers of it easily! So let's find its eigenvectors (we already computed its eigenvalues):

Thus its diagonalization is:

Finally, we can use this diagonalization to compute arbitrary powers of the matrix:

Thus we obtain the following simple formula for the n-th Fibonacci number:

The idea used throughout this example applies in a lot of generality: if we can represent something by matrix multiplication, then there's a good chance that diagonalization (via eigenvalues/eigenvectors) can shed light on the problem.

Arbitrary Matrix Powers

Once we have diagonalized a matrix, performing an operation on it is almost as easy as performing that operation on a number. We already saw this with computing large powers of a matrix: our procedure was...

a) First,

b) Next,

c) Finally,

This same basic idea works in lots of generality, and helps us talk about things we wouldn't even know how to define otherwise. For example...

What is a square root B of the matrix A?

OK, how could we find a square root of A?

Example. Find a square root of the matrix $A = \begin{bmatrix} 0 & 4 \\ -1 & 5 \end{bmatrix}$.

Similarly, we can use this method to define the r-th power of a diagonalizable matrix for any real number r (i.e., r doesn't need to be an integer):

Definition 11.2 — Matrix Powers

Let r be a real number and let A be a diagonalizable matrix (i.e., $A = PDP^{-1}$ for some invertible P and diagonal D). Then $A^r = PD^rP^{-1}$, where D^r is obtained by raising each of its diagonal entries to the r-th power.

Example. Compute A^π when $A = \begin{bmatrix} 0 & 4 \\ -1 & 5 \end{bmatrix}$.

Question to ponder: What happens when $r = -1$ in the above definition?

Arbitrary Matrix Functions

The previous section just touched the tip of the iceberg: we can also extend any function with a Taylor series to matrices now. For example, let's consider the function e^x. Recall that

$$e^x = 1 + x +$$

With that in mind, we define e^A, where A is a square matrix, as follows:

That seems like it might be nasty to calculate in general. Fortunately, we can just do what we usually do: diagonalize, apply the function e^x to each diagonal entry, and then un-diagonalize.

Example. Compute e^A, where $A = \begin{bmatrix} 0 & 4 \\ -1 & 5 \end{bmatrix}$.

Why does this work?

The following properties of the matrix exponential are straightforward to check:

- $e^O = I$
- $e^A e^{-A} = I$

There wasn't really anything too special about the function e^x here that let us define it for matrices: we can do the same thing for any function that equals its Taylor series, and the idea is exactly the same: just apply the function to the diagonal entries in the diagonalization of the matrix.

Example. Compute $\sin(A)$, where $A = \begin{bmatrix} 0 & 4 \\ -1 & 5 \end{bmatrix}$.

So go forth, and compute the sin, arctan, and log of matrices to your heart's content!

(Unless the matrices aren't diagonalizable... in that case, take Advanced Linear Algebra (MATH 3221) to learn what to do.)