Determinants

This document is a supplement to the Linear Algebra class from Canada/USA Mathcamp 2013. You should understand the definitions of vector spaces, linear independence, and spanning, and know the Rank-Nullity Theorem. The material on eigenvectors and eigenvalues from that course aren't necessary for understanding the material presented here.

These notes will go pretty quickly, and exercises are scattered throughout that you should solve if you want to understand what's going on. If you're reading this during camp, find me if you have any questions.

Area

Our motivating question will be to find a way to talk about area and volume in a vector space. Since the only objects we have to work with are vectors, not two- or three-dimensional objects like polygons or balls, we'll study the areas of shapes that can be made out of vectors. We'll start in $\mathbb{R}^2$:

Definition 1. Given two vectors $\vec{v}$ and $\vec{w}$ in $\mathbb{R}^2$, the *parallelogram* defined by $\vec{v}$ and $\vec{w}$ is the set

$$\{a\vec{v}+b\vec{w} : 0 \le a \le 1,\ 0 \le b \le 1\}.$$

(See Figure 1.) The *unit square* is the parallelogram defined by $\vec{e}_1$ and $\vec{e}_2$.

We're going to describe a function D on pairs of vectors in $\mathbb{R}^2$ by saying that $D(\vec{v},\vec{w})$ is the area of the parallelogram defined by $\vec{v}$ and $\vec{w}$ (at least, that's almost what D will be — we'll see later that a small adjustment will need to be made). It would be possible to work out a formula for D by hand using Euclidean geometry, but that would be much harder to generalize to higher-dimensional vector spaces. We'll take a different approach: by applying some simple geometric reasoning, we'll be able to determine a set of conditions that D will have to satisfy that will determine it completely.

What happens to the area of the parallelogram if you multiply one of $\vec{v}$ or $\vec{w}$ by a scalar? This corresponds to scaling the length of one of the edges of the parallelogram, so the area has to multiply by the same factor (See Figure 2):

$$D(a\vec{v},\vec{w}) = D(\vec{v},a\vec{w}) = aD(\vec{v},\vec{w}).$$

Exercise 2. By similar geometric reasoning, prove that for any $\vec{v}$, $\vec{w}$ and $\vec{x}$,

$$D(\vec{v}+\vec{w},\vec{x}) = D(\vec{v},\vec{x})+D(\vec{w},\vec{x}), \text{ and } D(\vec{v},\vec{w}+\vec{x}) = D(\vec{v},\vec{w})+D(\vec{v},\vec{x}).$$

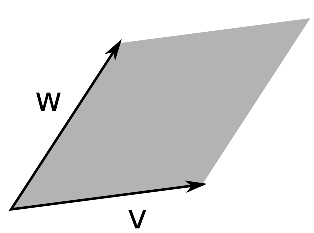

Figure 1: The parallelogram defined by $\vec{v}$ and $\vec{w}$.

Functions of two vectors that satisfy these two properties — scalars and addition pulling out of each component separately — are called *bilinear*. (The name is meant to suggest the fact that, for a fixed $\vec{v}$, $D(\vec{v},\vec{w})$ is a linear function of $\vec{w}$, and vice versa.) There's one more property of D that doesn't form part of the definition of bilinearity in general, but is an important fact about area: if $\vec{v}$ and $\vec{w}$ are equal, then the "parallelogram" defined by them is just a line segment, so $D(\vec{v},\vec{v})$ should be 0. When a bilinear function satisfies this property, we call it *alternating*.

This observation leads us to a somewhat unexpected fact: in order for these properties to all be true, D is going to have to sometimes be negative:

Exercise 3. Prove from the properties of D outlined above that, for any $\vec{v}$ and $\vec{w}$,

$$D(\vec{w},\vec{v}) = -D(\vec{v},\vec{w}).$$

(This partially justifies the name "alternating.") Hint: Do something clever with $\vec{v}+\vec{w}$.

What should we make of this? We'll interpret D as a "signed area," that is, an area that might carry a negative sign. The actual area of the parallelogram we mentioned is actually $|D(\vec{v},\vec{w})|$.

Exercise/Definition 4. The sign carries information about what we call the *orientation* of the vectors $\vec{v}$ and $\vec{w}$. Make this precise; what should we mean by "positively oriented" and "negatively oriented" in order for it to line up with the sign of D? Hint: Draw lots of pictures.

There's one final property we need to assert in order to nail down what D has to be:

Exercise 5. If D is a function satisfying all the properties we have laid out so far, show that $E(\vec{v},\vec{w}) = kD(\vec{v},\vec{w})$ also satisfies those properties for any real number k, possibly including 0.

In order to remove this ambiguity, we need to nail down the area of some specific parallelogram in order to use it to measure the others. The unit square is a natural choice: let's declare that it should have area 1.

All together, then, we have made an argument for each of the following properties of D:

- For any vectors $\vec{v}$ and $\vec{w}$ and any scalar a, $D(a\vec{v},\vec{w}) = D(\vec{v},a\vec{w}) = aD(\vec{v},\vec{w})$.
- For any vectors $\vec{v}$, $\vec{w}$, and $\vec{x}$, $D(\vec{v}+\vec{w},\vec{x}) = D(\vec{v},\vec{x}) + D(\vec{w},\vec{x})$ and $D(\vec{v},\vec{w}+\vec{x}) = D(\vec{v},\vec{w}) + D(\vec{v},\vec{x})$.
- For any $\vec{v}$, $D(\vec{v},\vec{v}) = 0$.
- $D(\vec{e}_1,\vec{e}_2) = 1$.

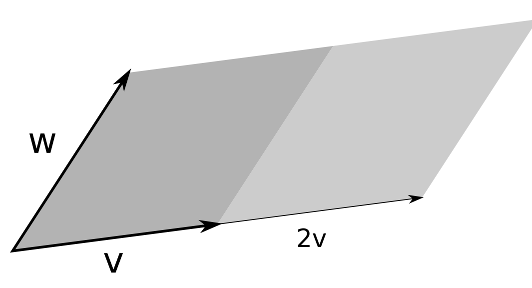

Figure 2: The behavior of D under scalar multiplication.

Now that we have the properties we care about, we can make our goal more formal: we want to find a function D that satisfies those four properties, and we want it to be unique:

Exercise 6. Show that these properties completely characterize the function D and that such a D exists. That is, any function f that has these four properties must be equal to D, and there is such a function. Hint: if $\vec{v}=(a,b)$ and $\vec{w}=(c,d)$, figure out a formula for what $D(\vec{v},\vec{w})$ has to be using only these properties, and then show that that formula indeed gives a function that satisfies the properties.

We mentioned above that plugging in the same vector as both of the inputs to D should give 0 because the corresponding parallelogram has no area. In fact, these axioms show that something slightly more general is true. D allows us to measure linear dependence:

Exercise 7. Prove that $\vec{v}$ and $\vec{w}$ are linearly independent if and only if $D(\vec{v},\vec{w})\neq 0$.

Generalizing to More Dimensions

We can generalize all of these ideas to more dimensions than two. In three dimensions, of course, we'll want to be measuring volume instead of area, and we'll use three vectors instead of two. The shape that plays the same role here as the parallelogram is called a parallelepiped:

Definition 8. Given three vectors $\vec{v}$, $\vec{w}$, and $\vec{x}$ in $\mathbb{R}^3$, the *parallelepiped* defined by $\vec{v}$, $\vec{w}$, and $\vec{x}$ is the set

$$\{a\vec{v}+b\vec{w}+c\vec{x}: 0\leq a\leq 1,\ 0\leq b\leq 1,\ 0\leq c\leq 1\}.$$

In more than three dimensions, we do more or less the same thing:

Definition 9. Given n vectors $\vec{v}_1,\ldots,\vec{v}_n$ in $\mathbb{R}^n$, the *parallelotope* defined by them is the set

$$\left\{\sum_{i=1}^{n} a_i\vec{v}_i : 0\leq a_i\leq 1 \text{ for each } i\right\}.$$

So a parallelogram is a two-dimensional parallelotope, and a parallelepiped is a three-dimensional parallelotope.

So our generalization of D should measure the oriented volume of parallelepipeds in $\mathbb{R}^3$, or the oriented *hypervolume*, the n-dimensional analogue of volume, in $\mathbb{R}^n$. Just as our two-dimensional D was bilinear, alternating, and took the unit square to 1, our n-dimensional version should...

- ...be *multilinear*, that is, linear in each component separately: for any choice of $n-1$ vectors $\vec{v}_1,\vec{v}_2,\ldots,\vec{v}_{i-1},\vec{v}_{i+1},\ldots,\vec{v}_n$, the function $\vec{w}\mapsto D(\vec{v}_1,\ldots,\vec{v}_{i-1},\vec{w},\vec{v}_{i+1},\ldots,\vec{v}_n)$ should be a linear function of $\vec{w}$.
- ...be alternating: if any two of the inputs to D are the same, the result should be 0.
- ...take the unit hypercube to 1: $D(\vec{e}_1,\vec{e}_2,\ldots,\vec{e}_n)=1$.

Exercise 10. Just as in the two-dimensional case, show for any function D satisfying these properties, that switching any two of the inputs to D puts a minus sign on the result.

Exercise 11. Also as in the two-dimensional case, show that these conditions completely determine the values of D for all collections of vectors. Work out a formula for D in terms of the coordinates of the vectors when $n=3$.

Note that we haven't shown that such a function D actually exists, only that, if it did exist, it would be unique. We're going to work somewhat harder to find a formula for D that works in general, but at the end we're going to have both the existence of D and a formula for it. To do so, we'll have to develop a little bit of the theory of permutations:

Exercise/Definition 12. A *permutation of n* is a way of rearranging the numbers $(1,2,\dots,n)$ so that each occurs exactly once, that is, a bijective function from the set $\{1,2,\dots,n\}$ to itself. For example, $\pi=(4,2,1,3)$ and $\sigma=(2,1,3,4)$ are permutations of 4. (This is called the *one-line representation* of a permutation.) Thinking of them as functions, π, for example, is the function defined by

$$\pi(1)=4 \quad ,\pi(2)=2, \quad \pi(3)=1, \quad \pi(4)=3.$$

The set of all permutations of n is called the *symmetric group on n letters*, written S_n.

An *inversion* in a permutation s is a pair of numbers i,j with $i<j$ but $s(i)>s(j)$, that is, a pair of numbers that appears out of order in the one-line representation. For example, there are 4 inversions in π above: positions $(1,2)$, $(1,3)$, $(1,4)$, and $(2,3)$ are out of order. The permutation σ has only one inversion, in the first two positions. If a permutation has an odd number of inversions, it's called *odd*, otherwise it's *even*.

Prove that if you take an odd permutation and switch two of the numbers in it, you get an even permutation, and vice versa.

Exercise/Definition 13. (Trickier) The *sign* of a permutation π, written $\operatorname{sgn}\pi$, is 1 if π is even and -1 if π is odd. Suppose we have a collection of vectors $\vec{v}_1,\dots,\vec{v}_n\in\mathbb{R}^n$, with $\vec{v}_j=(a_{1j},a_{2j},\dots,a_{nj})$. Prove that the formula

$$D(\vec{v}_1,\dots,\vec{v}_n)=\sum_{\pi\in S_n}\operatorname{sgn}\pi\prod_{j=1}^{n}a_{\pi(j),j}$$

satisfies all the properties we asked of D, and therefore, by the uniqueness proved above, is the only such function. (Hint: What happens if you take $D(\vec{v}_1,\dots,\vec{v}_n)$, expand it in terms of the $\vec{e}$'s, and try to expand using multilinearity?)

The Determinant

Our volume function D leads to the definition of one of the most important concepts in linear algebra, if not all of mathematics:

Definition 14. Suppose T is a linear transformation from $\mathbb{R}^n$ to itself. The *determinant* of T, written $\det T$, is the oriented hypervolume of the images of the standard basis vectors, that is,

$$\det T = D(T(\vec{e}_1), T(\vec{e}_2),\dots,T(\vec{e}_n)).$$

It's important to remark that the notation "det" for determinant is standard, while the notation "D" was invented for these notes.

The determinant measures the amount by which T affects hypervolumes (or areas or volumes): we're measuring the hypervolume of the parallelotope defined by $T(\vec{e}_1),\dots,T(\vec{e}_n)$; this is just where T sends the unit hypercube. The sign of $\det T$ measures whether it takes the unit hypercube to something of the same or the opposite orientation. If $\det T$ is positive, we say T is *orientation-preserving*, and if it's negative we say T is *orientation-reversing*.

Exercise 15. Show that, for any collection of vectors $\vec{v}_1, \ldots, \vec{v}_n$, we have

$$D(T\vec{v}_1, \ldots, T\vec{v}_n) = (\det T) D(\vec{v}_1, \ldots, \vec{v}_n).$$

(Hint: Show that the function

$$f(\vec{v}_1, \ldots, \vec{v}_n) = \frac{D(T\vec{v}_1, \ldots, T\vec{v}_n)}{\det T}$$

satisfies the axioms that determine D.)

While we defined the determinant in terms of what it does to the unit hypercube, we've shown that applying T multiplies the hypervolume of *any* parallelotope by $\det T$. This gives an even stronger geometric interpretation of the determinant: it's the factor by which hypervolumes are multiplied when you apply T, with a sign that keeps track of orientation.

Exercise 16. Prove that, for any maps T and U from $\mathbb{R}^n$ to itself, $\det(TU) = (\det T)(\det U)$. Conclude that if T is invertible, then $\det T \neq 0$.

Exercise 17. Generalize Exercise 7: prove that for any collection of vectors $\vec{v}_1, \ldots, \vec{v}_n$, they form a basis if and only if $D(\vec{v}_1, \ldots, \vec{v}_n) \neq 0$. (Hint: If they form a basis, consider the map that takes the standard basis to the $\vec{v}$'s.) Conclude that if T is not invertible, $\det T = 0$.

Exercise 18. Suppose $\vec{v}_1, \ldots, \vec{v}_n$ is a basis for $\mathbb{R}^n$, and let A be the map from $\mathbb{R}^n$ to itself that takes the standard basis to this one. If T is any linear map on $\mathbb{R}^n$, convince yourself that ATA^{-1} is the matrix of T in terms of the $\vec{v}$ basis, that is, the columns of ATA^{-1} are the coefficients of each $T(\vec{v}_j)$ when it's written as a linear combination of the $\vec{v}$'s. Prove that $\det(ATA^{-1}) = \det T$, that is, the determinant of a linear map doesn't depend on which basis you use to write the matrix.

Finally, it's worth pointing out that exercise 13 gives us a formula for the determinant in terms of the entries of the matrix: if a_{ij} is the entry in row i, column j, of the matrix for T then

$$\det T = \sum_{\pi \in S_n} \operatorname{sgn} \pi \prod_j a_{\pi(j),j}.$$

There are many more things to say about determinants than we will be able to here: there is a recursive formula for the determinant of a matrix in terms of determinants of smaller matrices, Cramer's rule for determining the inverse of a matrix, and innumerable connections to differential equations, combinatorics, real and functional analysis, commutative algebra, algebraic geometry, and a multitude of other fields. If you're reading this at Mathcamp, feel free to ask me or anyone else about any of these.