

Math 290-2: Linear Algebra & Multivariable Calculus

Northwestern University, Lecture Notes

Written by Santiago Cañez

These are notes which provide a basic summary of each lecture for Math 290-2, the second quarter of “MENU: Linear Algebra & Multivariable Calculus”, taught by the author at Northwestern University. The books used as references are the 5th edition of *Linear Algebra with Applications* by Bretscher and the 4th edition of *Vector Calculus* by Colley. Watch out for typos! Comments and suggestions are welcome.

Contents

Lecture 1: Orthonormal Bases	2
Lecture 2: Orthogonal Projections	4
Lecture 3: Gram-Schmidt Process	8
Lecture 4: Orthogonal Matrices	12
Lecture 5: Least Squares	17
Lecture 6: Symmetric Matrices	20
Lecture 7: Quadratic Forms	25
Lecture 8: Curves and Lines	30
Lecture 9: Cross Products	33
Lecture 10: Planes	36
Lecture 11: Polar/Cylindrical Coordinates	39
Lecture 12: Spherical Coordinates	47
Lecture 13: Multivariable Functions	53
Lecture 14: Quadric Surfaces	58
Lecture 15: Limits	64
Lecture 16: Partial Derivatives	68
Lecture 17: Differentiability	72
Lecture 18: Jacobians and Second Derivatives	76
Lecture 19: Chain Rule	81
Lecture 20: Directional Derivatives	86
Lecture 21: Gradients	90
Lecture 22: Taylor Polynomials	95
Lecture 23: Local Extrema	98
Lecture 24: Absolute Extrema	103
Lecture 25: Lagrange Multipliers	108
Lecture 26: More on Lagrange Multipliers	114

Lecture 1: Orthonormal Bases

The beginning of a new quarter, hoorah! Today I gave a brief overview of the course, and mentioned that the problem of classifying extrema points of multivariable functions is one way in which we will see linear algebra pop up in multivariable calculus. Then we started talking about the notion of an orthonormal basis.

Dot products. Recall that the *dot product* of two vectors

$$\vec{u} = \begin{pmatrix} u_1 \\ \vdots \\ u_n \end{pmatrix} \text{ and } \vec{v} = \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix}$$

in \mathbb{R}^n is the number $\vec{u} \cdot \vec{v}$ defined by

$$\vec{u} \cdot \vec{v} = u_1 v_1 + \cdots + u_n v_n.$$

We'll see later that this formula is (surprisingly) the same as

$$\vec{u} \cdot \vec{v} = \|\vec{u}\| \|\vec{v}\| \cos \theta$$

where $\|\vec{u}\|$ and $\|\vec{v}\|$ denote the length of \vec{u} and \vec{v} respectively and where θ is the angle between \vec{u} and \vec{v} . Recall that the *length* of a vector \vec{x} can be written as

$$\|\vec{x}\| = \sqrt{\vec{x} \cdot \vec{x}}.$$

Properties of dot products. The expression for the dot product given above in terms of $\cos \theta$ should make it clear that $\vec{u} \cdot \vec{v} = 0$ if and only if \vec{u} and \vec{v} are *orthogonal*, meaning perpendicular. Indeed, for nonzero \vec{u} and \vec{v} the dot product is zero if and only if $\cos \theta = 0$, which happens if and only if $\theta = 90^\circ$.

But note that even when nonzero, the sign of the dot product still gives useful geometric information: $\vec{u} \cdot \vec{v} > 0$ if and only if the angle between \vec{u} and \vec{v} is less than 90° and $\vec{u} \cdot \vec{v} < 0$ if and only if the angle between \vec{u} and \vec{v} is greater than 90° . We'll use these interpretations later, especially in the spring quarter.

Orthonormal vectors. A collection $\vec{u}_1, \dots, \vec{u}_k$ of vectors in \mathbb{R}^n is said to be *orthonormal* if all vectors are orthogonal to each other and all have length 1. In particular, an *orthonormal basis* of \mathbb{R}^n (or of a subspace of \mathbb{R}^n) is a basis consisting of orthonormal vectors.

Example 1. The standard basis of \mathbb{R}^n is an orthonormal basis, but of course, there can be other orthonormal bases. In particular, the vectors

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

also form an orthonormal basis of \mathbb{R}^2 .

Example 2. The vectors

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

form an orthonormal basis of \mathbb{R}^3 .

Why care about orthonormal bases? Say that $\vec{u}_1, \dots, \vec{u}_n$ is an orthonormal basis of some space. Given another vector \vec{x} , it should be possible to express \vec{x} as a linear combination of these basis vectors: i.e.

$$\vec{x} = c_1\vec{u}_1 + \dots + c_n\vec{u}_n$$

for some scalars c_1, \dots, c_n . Using techniques from last quarter, we could solve this equation for the unknown scalars either by converting it into a system of linear equations and using row operations, or by converting it into a matrix equation and using some kind of inverse. However, both of these methods lead to a lot of extra work, and are completely unnecessary in this case.

Take the above expression and dot both sides with \vec{u}_1 :

$$\vec{x} \cdot \vec{u}_1 = (c_1\vec{u}_1 + \dots + c_n\vec{u}_n) \cdot \vec{u}_1.$$

Dot products are distributive, so the right side breaks up into

$$c_1\vec{u}_1 \cdot \vec{u}_1 + c_1\vec{u}_2 \cdot \vec{u}_1 + \dots + c_n\vec{u}_n \cdot \vec{u}_1.$$

But now the magic happens: since our basis is orthonormal, all of these dot products are zero except for $\vec{u}_1 \cdot \vec{u}_1$, which is 1 since \vec{u}_1 has length 1. Thus we are left with

$$\vec{x} \cdot \vec{u}_1 = c_1.$$

In general, $c_i = \vec{x} \cdot \vec{u}_i$. The point is that we have an easy way of determining the coefficients needed to express \vec{x} in terms of an orthonormal basis: we simply take the dot product of \vec{x} with each basis vector. This is why orthonormal bases will be useful.

Important. Given an orthonormal basis $\vec{u}_1, \dots, \vec{u}_n$ of some space and another vector \vec{x} in that space, we have

$$\vec{x} = (\vec{x} \cdot \vec{u}_1)\vec{u}_1 + \dots + (\vec{x} \cdot \vec{u}_n)\vec{u}_n,$$

so an orthonormal basis gives us an easy way of finding the coefficients in a linear combination expression.

Orthogonal vectors are linearly independent. Note one quick consequence of what we've done above: if $\vec{u}_1, \dots, \vec{u}_n$ are orthogonal, then they are linearly independent. Indeed, starting with

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

a similar technique as above (taking dot products of both sides with some \vec{u}_i) shows that each c_i must be zero, so $\vec{u}_1, \dots, \vec{u}_n$ are linearly independent as claimed.

Back to Example 1. Consider the orthonormal basis

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

of \mathbb{R}^2 from Example 1. Say we want to solve

$$\begin{pmatrix} 3 \\ 4 \end{pmatrix} = c_1 \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} + c_2 \begin{pmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}.$$

The coefficients needed are

$$c_1 = \begin{pmatrix} 3 \\ 4 \end{pmatrix} \cdot \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} = \frac{7}{\sqrt{2}}$$

and

$$c_2 = \begin{pmatrix} 3 \\ 4 \end{pmatrix} \cdot \begin{pmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} = \frac{1}{\sqrt{2}}.$$

You can check on your own that

$$\frac{7}{\sqrt{2}} \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} + \frac{1}{\sqrt{2}} \begin{pmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}$$

does in fact equal $\begin{pmatrix} 3 \\ 4 \end{pmatrix}$.

Back to Example 2. Consider the orthonormal basis

$$\begin{pmatrix} 2/3 \\ 1/3 \\ 2/3 \end{pmatrix}, \begin{pmatrix} -2/3 \\ 2/3 \\ 1/3 \end{pmatrix}, \begin{pmatrix} 1/3 \\ 2/3 \\ -2/3 \end{pmatrix}$$

of \mathbb{R}^3 from Example 2. We can write $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$ as a linear combination of these as

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

where these coefficients come from taking dot product of $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$ with each orthonormal basis vector.

Producing orthonormal bases. Given two nonzero orthogonal vectors \vec{u} and \vec{v} , it is easy to come up with orthonormal vectors which point in the same direction as these: \vec{u} and \vec{v} are already orthogonal, so all we need to do is rescale them to have length 1, which is done by dividing each by their lengths. In other words, if \vec{u} and \vec{v} are nonzero and orthogonal, then

$$\frac{\vec{u}}{\|\vec{u}\|} \text{ and } \frac{\vec{v}}{\|\vec{v}\|}$$

are orthonormal. The same works for a larger collection of nonzero vectors which are orthogonal to begin with.

But, what do we do when the vectors we start with are not orthogonal? How can we use them to produce orthonormal vectors with the same span as the original vectors? This is what the so-called *Gram-Schmidt process* is all about; we'll come back to this on Monday.

Lecture 2: Orthogonal Projections

Today we spoke about orthogonal projections onto arbitrary spaces, not just lines as we saw last quarter. The fact that we can easily compute such projections given an orthonormal basis for our space gives further evidence that orthonormal vectors are good to have around.

Warm-Up 1. Among all unit vectors $\mathbf{u} = \begin{pmatrix} u_1 \\ \vdots \\ u_n \end{pmatrix}$ in \mathbb{R}^n we want to find the one for which $u_1 + \cdots + u_n$ is maximal. The key point is that we can rewrite the expression we want to maximize

as a certain dot product, and then we can express this dot product in terms of the angle between two vectors. In particular, we have

$$u_1 + \cdots + u_n = \begin{pmatrix} u_1 \\ \vdots \\ u_n \end{pmatrix} \cdot \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} = \|\mathbf{u}\| \left\| \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \right\| \cos \theta$$

where θ is the angle between \mathbf{u} and $\begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}$. We are only considering unit vectors \mathbf{u} , so $\|\mathbf{u}\| = 1$, and the length of the vector with all entries equal to 1 is \sqrt{n} . So

$$u_1 + \cdots + u_n = \sqrt{n} \cos \theta.$$

This is maximized when $\cos \theta = 1$, which happens when $\theta = 0$, so we conclude that the vector \mathbf{u} we want should be the unit vector pointing in the same direction as $\begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}$. This unit vector is obtained

by dividing $\begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}$ by its length, so

$$\mathbf{u} = \begin{pmatrix} 1/\sqrt{n} \\ \vdots \\ 1/\sqrt{n} \end{pmatrix}$$

is the unit vector in \mathbb{R}^n which makes $u_1 + \cdots + u_n$ as large as possible.

Warm-Up 2. Let's find a basis for $\text{span} \left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\}^\perp$. Here, for a subspace V of \mathbb{R}^n , V^\perp denotes the *orthogonal complement* of V in \mathbb{R}^n , which is the subspace consisting of all vectors in \mathbb{R}^n which are orthogonal to everything in V . In our case, we're looking at the orthogonal complement of the line spanned by $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$. Geometrically, this orthogonal complement should be a plane.

In order for a vector $\begin{pmatrix} x \\ y \\ z \end{pmatrix}$ to be in $\text{span} \left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\}^\perp$, it must be orthogonal to $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$. Thus a vector in this orthogonal complement satisfies the equation

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

As expected, the orthogonal complement to $\text{span} \left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\}$ is a plane, namely the plane with equation $x + y + z = 0$. We can now find a basis for this plane as we would have done last quarter, and one possible basis turns out to be

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

Now instead say we wanted to find an *orthonormal basis* for this orthogonal complement. The basis we found above doesn't work since those basis vectors aren't orthogonal. But, we can use them to find what we want. Keeping the first basis vector as is, we then want to find another vector \mathbf{v} on the plane $x + y + z = 0$ which is orthogonal to the first basis vector $\begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}$. We know that this

vector \mathbf{v} we want can be written as a linear combination of the basis vectors found above (as any vector on this plane can), so the problem boils down to finding coefficients which make

$$\mathbf{v} = c_1 \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix} + c_2 \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}$$

orthogonal to $\begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}$. This means that $\mathbf{v} \cdot \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}$ should be zero, which gives the equation

$$2c_1 + c_2 = 0.$$

Taking $c_1 = 1$ and $c_2 = -2$ gives one possible set of coefficients, and then

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

This gives

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

as an orthogonal basis for $\text{span} \left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\}^\perp$, and dividing each of these basis vectors by its length then gives an orthonormal basis for this orthogonal complement.

Orthogonal projections. Recall from last quarter that the *orthogonal projection* of a vector \mathbf{x} onto the line spanned by a vector \mathbf{u} is given by

$$\text{proj}_{\mathbf{u}}(\mathbf{x}) = \left(\frac{\mathbf{x} \cdot \mathbf{u}}{\mathbf{u} \cdot \mathbf{u}} \right) \mathbf{u}.$$

We derived this formula in class, but I'll omit that here—you can look it up in the book if interested.

We can now generalize this to the orthogonal projection of a vector onto *any* subspace of \mathbb{R}^n , whether it be a line, a plane, or something higher dimensional. First the definition:

Given a subspace V of \mathbb{R}^n , the *orthogonal projection* of a vector \mathbf{x} in \mathbb{R}^n onto V is the unique vector $\text{proj}_V(\mathbf{x})$ in V which makes $\mathbf{x} - \text{proj}_V(\mathbf{x})$ orthogonal to V .

This orthogonality requirement is why we call this an “orthogonal projection” as opposed to simply a “projection”. It turns out that there is a relatively easy way of computing such orthogonal projections, at least if we have an orthonormal basis of V available to us. The key fact is that given an orthogonal basis of V (what follows would not be true if our basis wasn't orthogonal), the orthogonal projection of \mathbf{x} onto V is obtained by orthogonal projecting \mathbf{x} onto each basis vector separately and adding together all such projections.

Important. If $\mathbf{u}_1, \dots, \mathbf{u}_k$ is an orthonormal basis of V , the orthogonal projection of \vec{x} onto V is given by

$$\begin{aligned} \text{proj}_V(\mathbf{x}) &= \text{proj}_{\mathbf{u}_1}(\mathbf{x}) + \dots + \text{proj}_{\mathbf{u}_k}(\mathbf{x}) \\ &= (\mathbf{x} \cdot \mathbf{u}_1)\mathbf{u}_1 + \dots + (\mathbf{x} \cdot \mathbf{u}_k)\mathbf{u}_k. \end{aligned}$$

If our basis were only orthogonal instead of orthonormal, all that would change is that the formula for the coefficients above would have denominators of the form $\mathbf{u}_i \cdot \mathbf{u}_i$; such denominators all happen to be equal to 1 for an orthonormal basis.

Example. Let's find the orthogonal projection of $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$ onto the subspace

$$V = \text{span} \left\{ \begin{pmatrix} 2 \\ 2 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -2 \\ 2 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ -4 \\ 0 \end{pmatrix} \right\}$$

of \mathbb{R}^4 . The nice thing is that these given vectors are already orthogonal, so to get an orthonormal basis for this span we only need to divide each by its length. This gives

$$\begin{pmatrix} 2/3 \\ 2/3 \\ 1/3 \\ 0 \end{pmatrix}, \begin{pmatrix} -2/3 \\ 2/3 \\ 0 \\ 1/3 \end{pmatrix}, \begin{pmatrix} 1/3\sqrt{2} \\ 1/3\sqrt{2} \\ -4/3\sqrt{2} \\ 0 \end{pmatrix}$$

as an orthonormal basis for the given span.

Denoting these orthonormal basis vectors by $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$ and $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$ by \mathbf{x} , we then have

$$\begin{aligned} \text{proj}_V \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} &= (\mathbf{x} \cdot \mathbf{u}_1) \mathbf{u}_1 + (\mathbf{x} \cdot \mathbf{u}_2) \mathbf{u}_2 + (\mathbf{x} \cdot \mathbf{u}_3) \mathbf{u}_3 \\ &= \frac{5}{3} \mathbf{u}_1 + \frac{1}{3} \mathbf{u}_2 - \frac{2}{3\sqrt{2}} \mathbf{u}_3 \\ &= \frac{5}{3} \begin{pmatrix} 2/3 \\ 2/3 \\ 1/3 \\ 0 \end{pmatrix} + \frac{1}{3} \begin{pmatrix} -2/3 \\ 2/3 \\ 0 \\ 1/3 \end{pmatrix} - \frac{2}{3\sqrt{2}} \begin{pmatrix} 1/3\sqrt{2} \\ 1/3\sqrt{2} \\ -4/3\sqrt{2} \\ 0 \end{pmatrix} \end{aligned}$$

as the explicit orthogonal projection we want. (Don't be scared, it's unlikely that you would ever have to simplify such an expression further. When in doubt, ask!)

Why care about orthogonal projections? All of this is well and good, but it's not yet clear that orthogonal projections are actually useful things. Here is the key fact: among all vectors in V , $\text{proj}_V(\vec{x})$ is the one which is "closest" to \vec{x} in the sense that $\|\mathbf{x} - \mathbf{v}\|$ is minimized for \mathbf{v} in V precisely when $\mathbf{v} = \text{proj}_V(\mathbf{x})$. This is such an important fact that I'll repeat it again, in red! We'll exploit this minimization property of orthogonal projections next week to great effect.

Important. Given a subspace V of \mathbb{R}^n and a vector \mathbf{x} in \mathbb{R}^n , the quantity $\|\mathbf{x} - \mathbf{v}\|$ as \mathbf{v} ranges through all vectors in V is minimized when $\mathbf{v} = \text{proj}_V(\mathbf{x})$.

Lecture 3: Gram-Schmidt Process

Today we spoke about the Gram-Schmidt process, which is a procedure for producing orthonormal bases of any kind of spaces we want. The calculations can be kind of tedious, but the results are well worth it.

Warm-Up 1. For $\mathbf{x} = \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}$, we want to find the minimum possible value of $\|\mathbf{x} - \mathbf{v}\|$ as \mathbf{v} ranges through all vectors in

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

The key fact is that this minimum value is obtained precisely when $v = \text{proj}_V(\mathbf{x})$, so this problem is really just a convoluted way of saying “compute $\|\mathbf{x} - \text{proj}_V(\mathbf{x})\|$ ”. The given basis vectors of V are already orthogonal, so calling them \mathbf{v}_1 and \mathbf{v}_2 respectively we have:

$$\begin{aligned} \text{proj}_V(\mathbf{x}) &= \text{proj}_{\mathbf{v}_1} \mathbf{x} + \text{proj}_{\mathbf{v}_2} \mathbf{x} \\ &= \left(\frac{\mathbf{x} \cdot \mathbf{v}_1}{\mathbf{v}_1 \cdot \mathbf{v}_1} \right) \mathbf{v}_1 + \left(\frac{\mathbf{x} \cdot \mathbf{v}_2}{\mathbf{v}_2 \cdot \mathbf{v}_2} \right) \mathbf{v}_2 \\ &= \frac{2}{6} \begin{pmatrix} 2 \\ 1 \\ 1 \end{pmatrix} - \frac{1}{3} \begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}. \end{aligned}$$

Thus the minimum value of $\|\mathbf{x} - \mathbf{v}\|$ we want is

$$\|\mathbf{x} - \text{proj}_V \mathbf{x}\| = \left\| \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} \right\| = \sqrt{2}.$$

Geometrically, V is a plane in \mathbb{R}^3 and what we have computed is the *distance* from $\begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}$ to this plane. We’ll see more of this later when we do calculus.

The matrix of an orthogonal projection. Suppose that $\mathbf{u}_1, \dots, \mathbf{u}_k$ is an orthonormal basis for a subspace V of \mathbb{R}^n , and let Q be the matrix having these as columns:

$$Q = \begin{pmatrix} | & & | \\ \mathbf{u}_1 & \cdots & \mathbf{u}_k \\ | & & | \end{pmatrix}.$$

Then QQ^T is the matrix of “orthogonal projection onto V ”, that is, the linear transformation which orthogonally projects a vector in \mathbb{R}^n onto the subspace V . Indeed, let’s compute the product $QQ^T \mathbf{x}$ for a vector \mathbf{x} in \mathbb{R}^n :

$$QQ^T \mathbf{x} = \begin{pmatrix} | & & | \\ \mathbf{u}_1 & \cdots & \mathbf{u}_k \\ | & & | \end{pmatrix} \begin{pmatrix} \text{---} & \mathbf{u}_1 & \text{---} \\ & \vdots & \\ \text{---} & \mathbf{u}_k & \text{---} \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$$

$$\begin{aligned}
&= \begin{pmatrix} | & & | \\ \mathbf{u}_1 & \cdots & \mathbf{u}_k \\ | & & | \end{pmatrix} \begin{pmatrix} \mathbf{x} \cdot \mathbf{u}_1 \\ \vdots \\ \mathbf{x} \cdot \mathbf{u}_k \end{pmatrix} \\
&= (\mathbf{x} \cdot \mathbf{u}_1)\mathbf{u}_1 + \cdots + (\mathbf{x} \cdot \mathbf{u}_k)\mathbf{u}_k.
\end{aligned}$$

We've seen already that this resulting expression is the orthogonal projection of \mathbf{x} onto V , so

$$QQ^T \mathbf{x} = \text{proj}_V \mathbf{x}$$

as claimed. The strange looking formula in Chapter 2 of the book for orthogonally projecting onto a line can now be derived using this approach.

Warm-Up 2. Let's compute the matrix of the orthogonal projection onto the plane

$$V = \text{span} \left\{ \begin{pmatrix} 2 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix} \right\}$$

from the first Warm-Up. We need an orthonormal basis for V , which is obtained simply by dividing each of the given orthogonal basis vectors by their lengths:

$$\mathbf{u}_1 = \frac{1}{\sqrt{6}} \begin{pmatrix} 2 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 2/\sqrt{6} \\ 1/\sqrt{6} \\ 1/\sqrt{6} \end{pmatrix}, \quad \mathbf{u}_2 = \frac{1}{\sqrt{3}} \begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} -1/\sqrt{3} \\ 1/\sqrt{3} \\ 1/\sqrt{3} \end{pmatrix}.$$

Using these as the columns of a matrix Q , the matrix for the orthogonal projection onto V is

$$\begin{aligned}
QQ^T &= \begin{pmatrix} 2/\sqrt{6} & -1/\sqrt{3} \\ 1/\sqrt{6} & 1/\sqrt{3} \\ 1/\sqrt{6} & 1/\sqrt{3} \end{pmatrix} \begin{pmatrix} 2/\sqrt{6} & 1/\sqrt{6} & 1/\sqrt{6} \\ -1/\sqrt{3} & 1/\sqrt{3} & 1/\sqrt{3} \end{pmatrix} \\
&= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1/2 & 1/2 \\ 0 & 1/2 & 1/2 \end{pmatrix}.
\end{aligned}$$

As a check, this matrix should have the property that $QQ^T \mathbf{x} = \text{proj}_V \mathbf{x}$, so using the vector $\mathbf{x} = \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}$ from the first Warm-Up we get:

$$QQ^T \mathbf{x} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1/2 & 1/2 \\ 0 & 1/2 & 1/2 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix},$$

which agrees with the answer we found for $\text{proj}_V \mathbf{x}$ in that Warm-Up.

Gram-Schmidt process. The point of the Gram-Schmidt process is to take a collection of linearly independent vectors and to produce a collection of orthonormal vectors with the same span as the original set of vectors. I wrote up some notes on this last year, which you can find on my website at <http://math.northwestern.edu/~scanez/archives/lin-algebra/notes.php>. This describes an approach to the Gram-Schmidt process which is slightly different than how the book does it, but which is I think a bit simpler computationally. Use whichever method works best for you, but be consistent!

The main point of the process is to at each step replace a vector in your original collection by what you get when you subtract from it its orthogonal projection onto all previously constructed vectors. Again, check the notes linked to above for more details. One thing I didn't mention in class is the following: how do we know that the resulting vectors will have the same span as the original vectors? This is because each vector constructed during the Gram-Schmidt process is a *linear combination* of the original vectors, so the span does not change.

One final point: how do we know that the resulting vectors must be orthogonal? For instance, how do we know

$$\mathbf{b}_2 = \mathbf{v}_2 - \left(\frac{\mathbf{v}_2 \cdot \mathbf{b}_1}{\mathbf{b}_1 \cdot \mathbf{b}_1} \right) \mathbf{b}_1$$

is orthogonal to \mathbf{b}_1 ? We simply compute:

$$\begin{aligned} \mathbf{b}_2 \cdot \mathbf{b}_1 &= \left[\mathbf{v}_2 - \left(\frac{\mathbf{v}_2 \cdot \mathbf{b}_1}{\mathbf{b}_1 \cdot \mathbf{b}_1} \right) \mathbf{b}_1 \right] \cdot \mathbf{b}_1 \\ &= \mathbf{v}_2 \cdot \mathbf{b}_1 - \left(\frac{\mathbf{v}_2 \cdot \mathbf{b}_1}{\mathbf{b}_1 \cdot \mathbf{b}_1} \right) (\mathbf{b}_1 \cdot \mathbf{b}_1) \\ &= \mathbf{v}_2 \cdot \mathbf{b}_1 - \mathbf{v}_2 \cdot \mathbf{b}_1 \\ &= 0. \end{aligned}$$

A similar computation works for the other vectors in the construction.

Important. Given linearly independent vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$, the Gram-Schmidt process produces orthonormal vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$ such that $\text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_k\} = \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$.

Example 1. Let's apply the Gram-Schmidt process to

$$\mathbf{v}_1 = \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}, \quad \mathbf{v}_2 = \begin{pmatrix} 1 \\ 3 \\ -1 \end{pmatrix}, \quad \mathbf{v}_3 = \begin{pmatrix} 4 \\ 1 \\ 1 \end{pmatrix}.$$

First we set

$$\mathbf{b}_1 = \mathbf{v}_1 = \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}.$$

Next we compute:

$$\begin{aligned} \mathbf{b}_2 &= \mathbf{v}_2 - \text{proj}_{\mathbf{b}_1} \mathbf{v}_2 \\ &= \mathbf{v}_2 - \left(\frac{\mathbf{v}_2 \cdot \mathbf{b}_1}{\mathbf{b}_1 \cdot \mathbf{b}_1} \right) \mathbf{b}_1 \\ &= \begin{pmatrix} 1 \\ 3 \\ -1 \end{pmatrix} - \frac{-2}{2} \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} 2 \\ 2 \\ -1 \end{pmatrix}. \end{aligned}$$

Note that this is indeed orthogonal to \mathbf{b}_1 . Finally, we compute:

$$\mathbf{b}_3 = \mathbf{v}_3 - \text{proj}_{\mathbf{b}_1} \mathbf{v}_3 - \text{proj}_{\mathbf{b}_2} \mathbf{v}_3$$

$$\begin{aligned}
&= \mathbf{v}_3 - \left(\frac{\mathbf{v}_3 \cdot \mathbf{b}_1}{\mathbf{b}_1 \cdot \mathbf{b}_1} \right) \mathbf{b}_1 - \left(\frac{\mathbf{v}_3 \cdot \mathbf{b}_2}{\mathbf{b}_2 \cdot \mathbf{b}_2} \right) \mathbf{b}_2 \\
&= \begin{pmatrix} 4 \\ 1 \\ 1 \end{pmatrix} - \frac{3}{2} \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix} - \frac{9}{9} \begin{pmatrix} 2 \\ 2 \\ -1 \end{pmatrix} \\
&= \begin{pmatrix} 1/2 \\ 1/2 \\ 2 \end{pmatrix}.
\end{aligned}$$

The final step is to divide each of $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3$ by their lengths, so the vectors resulting from applying the Gram-Schmidt process to $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ are

$$\mathbf{u}_1 = \frac{\mathbf{b}_1}{\|\mathbf{b}_1\|} = \begin{pmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \\ 0 \end{pmatrix}, \quad \mathbf{u}_2 = \frac{\mathbf{b}_2}{\|\mathbf{b}_2\|} = \begin{pmatrix} 2/3 \\ 2/3 \\ -1/3 \end{pmatrix}, \quad \mathbf{u}_3 = \frac{\mathbf{b}_3}{\|\mathbf{b}_3\|} = \frac{1}{\sqrt{9/2}} \begin{pmatrix} 1/2 \\ 1/2 \\ 2 \end{pmatrix} = \begin{pmatrix} 1/3\sqrt{2} \\ 1/3\sqrt{2} \\ 2\sqrt{2}/3 \end{pmatrix}.$$

Example 2. We find an orthonormal basis for the kernel of

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

First we need *any* basis for $\ker A$, which we can find using techniques from last quarter. One possible basis is

$$\mathbf{v}_1 = \begin{pmatrix} -2 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \quad \mathbf{v}_2 = \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \end{pmatrix}, \quad \mathbf{v}_3 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ -1 \end{pmatrix}.$$

Next we apply the Gram-Schmidt process to this basis. We get:

$$\begin{aligned}
\mathbf{b}_1 &= \mathbf{v}_1 = \begin{pmatrix} -2 \\ 1 \\ 0 \\ 0 \end{pmatrix} \\
\mathbf{b}_2 &= \mathbf{v}_2 - \text{proj}_{\mathbf{b}_1} \mathbf{v}_2 = \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \end{pmatrix} - \frac{-2}{5} \begin{pmatrix} -2 \\ 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 1/5 \\ 2/5 \\ 1 \\ 0 \end{pmatrix}.
\end{aligned}$$

At this point, to avoid having to deal with so many fractions, let's scale this vector by 5 and use

$$\mathbf{b}_2 = \begin{pmatrix} 1 \\ 2 \\ 5 \\ 0 \end{pmatrix}$$

instead; this is fine, since this new choice for \mathbf{b}_2 is still orthogonal to \mathbf{b}_1 . Finally:

$$\mathbf{b}_3 = \mathbf{v}_3 - \text{proj}_{\mathbf{b}_1} \mathbf{v}_3 - \text{proj}_{\mathbf{b}_2} \mathbf{v}_3 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ -1 \end{pmatrix} - \frac{-2}{5} \begin{pmatrix} -2 \\ 1 \\ 0 \\ 0 \end{pmatrix} - \frac{1}{30} \begin{pmatrix} 1 \\ 2 \\ 5 \\ 0 \end{pmatrix} = \begin{pmatrix} 1/6 \\ 1/3 \\ -1/6 \\ -1 \end{pmatrix}.$$

Finally we divide by lengths to get

$$\begin{pmatrix} -2/\sqrt{5} \\ 1/\sqrt{5} \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1/\sqrt{30} \\ 2/\sqrt{30} \\ 5/\sqrt{30} \end{pmatrix}, \frac{1}{\sqrt{7/6}} \begin{pmatrix} 1/6 \\ 1/3 \\ -1/6 \\ -1 \end{pmatrix}$$

as an orthonormal basis for $\ker A$.

Lecture 4: Orthogonal Matrices

Today we spoke about orthogonal transformations and matrices, otherwise known as rotations and reflections. Such matrices turn out to have many useful properties, as we'll see.

Warm-Up. We compute the distance from $\mathbf{x} = \begin{pmatrix} -1 \\ 2 \end{pmatrix}$ to each eigenspace of the matrix

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

Recall that such a distance is given by the length of \mathbf{x} minus its orthogonal projection onto the eigenspace we're looking at, and to compute such orthogonal projections we need orthonormal bases for the eigenspaces, which are obtained using the Gram-Schmidt process. This problem touches upon pretty much everything we've looked at so far this quarter, and some things from last quarter.

First we need any bases for the eigenspaces of A . The eigenvalues of A (as you should check) are 2 and 5, and computing bases for the eigenspace corresponding to each using techniques from last quarter gives:

$$\begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix} \text{ as a basis for } E_2 \text{ and } \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \text{ as a basis for } E_5.$$

Let's first work with E_5 . To get an orthonormal basis for E_5 we divide the given basis vector by its length:

$$\mathbf{u} = \begin{pmatrix} 1/\sqrt{3} \\ 1/\sqrt{3} \\ 1/\sqrt{3} \end{pmatrix}.$$

The orthogonal projection of \mathbf{x} onto E_5 is thus

$$\text{proj}_{E_5}(\mathbf{x}) = (\mathbf{x} \cdot \mathbf{u})\mathbf{u} = \frac{3}{\sqrt{3}} \begin{pmatrix} 1/\sqrt{3} \\ 1/\sqrt{3} \\ 1/\sqrt{3} \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}.$$

The distance from \mathbf{x} to E_5 is therefore

$$\|\mathbf{x} - \text{proj}_{E_5}(\mathbf{x})\| = \left\| \begin{pmatrix} -2 \\ 1 \\ 1 \end{pmatrix} \right\| = \sqrt{6}.$$

Now for E_2 . To get an orthonormal basis for E_2 we apply the Gram-Schmidt process to the basis vectors \mathbf{v}_1 and \mathbf{v}_2 we found previously:

$$\begin{aligned}\mathbf{b}_1 &= \mathbf{v}_1 = \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} \\ \mathbf{b}_2 &= \mathbf{v}_2 - \text{proj}_{\mathbf{b}_1} \mathbf{v}_2 = \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix} - \frac{1}{2} \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} -1/2 \\ 1 \\ -1/2 \end{pmatrix} \\ \mathbf{u}_1 &= \frac{\mathbf{b}_1}{\|\mathbf{b}_1\|} = \begin{pmatrix} -1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \end{pmatrix} \\ \mathbf{u}_2 &= \frac{\mathbf{b}_2}{\|\mathbf{b}_2\|} = \frac{1}{\sqrt{3/2}} \begin{pmatrix} -1/2 \\ 1 \\ -1/2 \end{pmatrix} = \begin{pmatrix} -1/\sqrt{6} \\ 2/\sqrt{6} \\ -1/\sqrt{6} \end{pmatrix},\end{aligned}$$

so \mathbf{u}_1 and \mathbf{u}_2 form an orthonormal basis of E_2 . The orthogonal projection of \mathbf{x} onto E_2 is then

$$\text{proj}_{E_2}(\mathbf{x}) = (\mathbf{x} \cdot \mathbf{u}_1)\mathbf{u}_1 + (\mathbf{x} \cdot \mathbf{u}_2)\mathbf{u}_2 = \frac{3}{\sqrt{2}} \begin{pmatrix} -1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \end{pmatrix} + \frac{3}{\sqrt{6}} \begin{pmatrix} -1/\sqrt{6} \\ 2/\sqrt{6} \\ -1/\sqrt{6} \end{pmatrix} = \begin{pmatrix} -2 \\ 1 \\ 1 \end{pmatrix},$$

so the distance from \mathbf{x} to E_2 is

$$\|\mathbf{x} - \text{proj}_{E_2}(\mathbf{x})\| = \left\| \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\| = \sqrt{3}.$$

As you can see, some of the dot product computations get a little messy; don't worry too much about getting all values exactly right, much more important is understanding the thought process which went into figuring out how to compute what we wanted to compute.

Orthogonal Matrices. Suppose that Q is an $n \times n$ matrix with orthonormal columns, say

$$Q = \begin{pmatrix} | & & | \\ \mathbf{u}_1 & \cdots & \mathbf{u}_n \\ | & & | \end{pmatrix}.$$

We call such a matrix an *orthogonal* matrix. For a vector \mathbf{x} in \mathbb{R}^n , let's compute $\|Q\mathbf{x}\|$. We have

$$Q\mathbf{x} = \begin{pmatrix} | & & | \\ \mathbf{u}_1 & \cdots & \mathbf{u}_n \\ | & & | \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = x_1\mathbf{u}_1 + \cdots + x_n\mathbf{u}_n$$

so

$$\begin{aligned}Q\mathbf{x} \cdot Q\mathbf{x} &= (x_1\mathbf{u}_1 + \cdots + x_n\mathbf{u}_n) \cdot (x_1\mathbf{u}_1 + \cdots + x_n\mathbf{u}_n) \\ &= x_1^2(\mathbf{u}_1 \cdot \mathbf{u}_1) + x_2^2(\mathbf{u}_2 \cdot \mathbf{u}_2) + \cdots + x_n^2(\mathbf{u}_n \cdot \mathbf{u}_n),\end{aligned}$$

which we get after distributing dot products and using the fact that $\mathbf{u}_i \cdot \mathbf{u}_j = 0$ for $i \neq j$. But each \mathbf{u}_i has length 1, so $\mathbf{u}_i \cdot \mathbf{u}_i = 1$ for all i and the above becomes

$$Q\mathbf{x} \cdot Q\mathbf{x} = x_1^2 + x_2^2 + \cdots + x_n^2 = \mathbf{x} \cdot \mathbf{x}.$$

Thus

$$\|Q\mathbf{x}\| = \|\mathbf{x}\| \text{ for any } \mathbf{x}.$$

In other words, an orthogonal matrix has the property that it “preserves” lengths.

Here’s a definition:

A linear transformation T from \mathbb{R}^n to \mathbb{R}^n is an *orthogonal transformation* if it is length-preserving in the sense that $\|T(\mathbf{x})\| = \|\mathbf{x}\|$ for all \mathbf{x} in \mathbb{R}^n .

The upshot is that orthogonal matrices give examples of orthogonal transformations. We will soon see that in fact all orthogonal transformations come about from orthogonal matrices.

Important. To say that a square matrix is orthogonal means that it has orthonormal columns, not just orthogonal columns. We don’t have a specific term for matrices with only orthogonal columns. The term “orthogonal matrix” has been around long enough that we’re stuck with it, even though “orthonormal matrix” might be a more descriptive name.


Rotations and reflections. Geometrically, the only types of linear transformations we’ve seen which preserve lengths are rotations and reflections, and indeed any orthogonal transformation must be one of these. To distinguish between the two, recall the interpretation of the sign of the determinant we saw last quarter: positive determinant means “orientation preserving” while negative determinant means “orientation reversing”. In particular, if $T(\mathbf{x}) = Q\mathbf{x}$ is length-preserving, then considering the interpretation of $|\det Q|$ as an expansion factor gives $|\det Q| = 1$, so $\det Q = \pm 1$; rotations have determinant $+1$ and reflections have determinant -1 .

Properties of orthogonal transformations. Here are two more key properties of orthogonal transformations: they “preserve angles” in the sense that the angle between $T(\mathbf{x})$ and $T(\mathbf{y})$ is the same as the angle between \mathbf{x} and \mathbf{y} , and they “preserve dot products” in the sense that $T(\mathbf{x}) \cdot T(\mathbf{y})$ is the same as $\mathbf{x} \cdot \mathbf{y}$. The book shows that orthogonal transformations preserve right angles at least, and these two general facts are exercises in the book which will show up on Homework 3.

For now, let’s use the fact that orthogonal transformations preserve dot products to justify a property of dot products we saw earlier: the fact that $\mathbf{x} \cdot \mathbf{y} = \|\mathbf{x}\| \|\mathbf{y}\| \cos \theta$ where θ is the angle between \mathbf{x} and \mathbf{y} . First consider the case where \mathbf{y} points along the positive x -axis, so

$$\mathbf{y} = \begin{pmatrix} a \\ 0 \end{pmatrix} \text{ for some } a > 0.$$

Then for $\mathbf{x} = \begin{pmatrix} c \\ d \end{pmatrix}$, we have $\mathbf{x} \cdot \mathbf{y} = ac$. But $a = \|\mathbf{y}\|$ and a correctly-drawn right triangle:



shows that $c = \|\mathbf{x}\| \cos \theta$. Thus $\mathbf{x} \cdot \mathbf{y} = ac = \|\mathbf{y}\| \|\mathbf{x}\| \cos \theta$ is true in this special case. Now, let \mathbf{x} and \mathbf{y} be any nonzero vectors. Take T to be a rotation which rotates \mathbf{y} to the positive x -axis. Then by the special case we just did we have

$$T(\mathbf{x}) \cdot T(\mathbf{y}) = \|T(\mathbf{x})\| \|T(\mathbf{y})\| \cos \theta$$

where θ is the angle between $T(\mathbf{x})$ and $T(\mathbf{y})$. But this is the same as the angle between \mathbf{x} and \mathbf{y} , and since T is orthogonal $\|T(\mathbf{x})\| = \|\mathbf{x}\|$, $\|T(\mathbf{y})\| = \|\mathbf{y}\|$, and $T(\mathbf{x}) \cdot T(\mathbf{y}) = \mathbf{x} \cdot \mathbf{y}$ so the above expression becomes

$$\mathbf{x} \cdot \mathbf{y} = \|\mathbf{x}\| \|\mathbf{y}\| \cos \theta$$

as claimed. The other way I know of deriving this property of dot products is via the so called “law of cosines” in trigonometry, which is nowhere near as enlightening as the way we did it here.

The matrix of an orthogonal transformation. Say that T is an orthogonal transformation. Since the standard basis vectors $\mathbf{e}_1, \dots, \mathbf{e}_n$ are orthonormal and T preserves lengths and angles, $T(\mathbf{e}_1), \dots, T(\mathbf{e}_n)$ are also orthonormal. But these vectors make up the columns of the matrix of T relative to the standard basis, so we are saying that

$$\text{matrix of } T = \begin{pmatrix} | & & | \\ T(\mathbf{e}_1) & \cdots & T(\mathbf{e}_n) \\ | & & | \end{pmatrix}$$

is an orthogonal matrix. So we have come full circle: not only do orthogonal matrices give examples of orthogonal transformations, but the matrix of any orthogonal transformation must actually be an orthogonal matrix.

Examples. The matrix of a 2-dimensional rotation:

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

is an orthogonal matrix, and so is the matrix of a 2-dimensional reflection.

The matrix

$$\begin{pmatrix} 1/\sqrt{2} & 2/3 & 1/3\sqrt{2} \\ -1/\sqrt{2} & 2/3 & 1/3\sqrt{2} \\ 0 & -1/3 & 2\sqrt{2}/3 \end{pmatrix}$$

is a 3×3 orthogonal matrix, meaning that its columns are orthonormal. Being orthogonal, this matrix must represent either a rotation or a reflection. Since it has determinant 1 (as you can check), it describes a 3-dimensional rotation.

The inverse of an orthogonal matrix. Again, suppose that Q is an orthogonal $n \times n$ matrix. We've worked out before that for *any* matrix (not necessarily square) with orthonormal columns, $Q^T Q = I$. But now if Q is square we can say more. The product $Q Q^T$ describes the orthogonal projection onto the space spanned by the columns of Q , which in this case is \mathbb{R}^n since Q has n columns. But "orthogonal projection of \mathbb{R}^n onto \mathbb{R}^n " is the identity transformation since projecting a vector in a given space onto that same space does nothing to it. Thus

$$Q Q^T = I \text{ when } Q \text{ is an orthogonal square matrix.}$$

Since $Q^T Q = I$ and $Q Q^T = I$, $Q^T = Q^{-1}$ so we find that the inverse of an orthogonal matrix is simply its transpose. For instance, the inverse of the 3×3 matrix in the previous example is its transpose. This property gives another characterization of orthogonal matrices.

Important. For a square matrix Q , the following conditions are equivalent:

- The transformation $T(\mathbf{x}) = Q\mathbf{x}$ preserves lengths,
- The transformation $T(\mathbf{x}) = Q\mathbf{x}$ preserves dot products,
- The transformation $T(\mathbf{x}) = Q\mathbf{x}$ describes either a rotation or a reflection,
- The columns of Q are orthonormal (i.e. Q is an orthogonal matrix),
- $Q^T Q = I$ and $Q Q^T = I$, so $Q^{-1} = Q^T$.

Final example. Let's find all orthogonal matrices of the form

$$\begin{pmatrix} 1 & a & d \\ 0 & b & e \\ 0 & c & f \end{pmatrix}.$$

In order for the first two columns to be orthogonal it must be the case that $a = 0$. Then, for the second columns to have length 1 it must be true that $b^2 + c^2 = 1$. Thus the second column must be of the form

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

For the third column to be orthogonal to the first column we again need $d = 0$, and to also be orthogonal to this second column the third column must be of the form

$$\begin{pmatrix} 0 \\ -\sin \theta \\ \cos \theta \end{pmatrix} \text{ or } \begin{pmatrix} 0 \\ \sin \theta \\ -\cos \theta \end{pmatrix}.$$

Thus, we see that all 3×3 orthogonal matrices of the given form must look like

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

The first form has determinant 1 and describes a rotation of the yz -plane around the x -axis while the second has determinant -1 and describes a reflection across a plane containing the x -axis.

Lecture 5: Least Squares

Today we spoke about the method of least squares, which is a beautiful application of orthogonal projections. In the course of working through this, we derived a way of computing orthogonal projections which avoids any mention of Gram-Schmidt or orthonormal bases, and so is usually much quicker to carry out computationally.

Warm-Up. We claim that the rows of an orthogonal matrix Q are in fact orthonormal. The rows of Q are the columns of Q^T , so the claim is that Q^T is also an orthogonal matrix. Here are two ways of seeing this.

First, if Q is orthogonal, the corresponding transformation preserves lengths and hence so does the inverse transformation. Thus Q^{-1} describes an orthogonal transformation, so Q^{-1} is an orthogonal matrix. (Another way to see that Q^{-1} is orthogonal is to note that the inverse of a rotation or reflection is itself a rotation or reflection.) Since $Q^{-1} = Q^T$, Q^T is orthogonal as claimed.


Second, since Q is orthogonal we have $QQ^T = I = Q^TQ$. But $Q = (Q^T)^T$ so

$$[(Q^T)^T]Q^T = I = Q^T[(Q^T)^T].$$

In other words, Q^T times its transpose is the identity, so by one of the equivalent characterizations of orthogonal matrices Q^T is also orthogonal.

Least squares. Say we want to find a function of the form $f(t) = c_0 + c_1t$ whose graph passes through the points $(0, 3), (1, 3), (2, 6)$. Of course here this is not possible since such a function describes a line and these three points do not lie on the same line. Instead, we ask for the function of this type which “best fits” the given points in the following sense.

Given any line, we consider the points on the line with x -coordinates $0, 1, 2$. Each of these points has some vertical distance to the original data points, which we denote by $\epsilon_1, \epsilon_2, \epsilon_3$:



We say that $f(t) = c_0 + c_1t$ “best fits” the given points in the *least squares* sense if it is the line for which the “error” $\epsilon_1^2 + \epsilon_2^2 + \epsilon_3^2$ is minimized. But this expression is minimized when its square root is minimized, so we are looking to minimize

$$\sqrt{\epsilon_1^2 + \epsilon_2^2 + \epsilon_3^2}.$$

The point is that we can rewrite this quantity as the length of a certain vector.

Indeed, if our line were going to pass through the given points exactly it would satisfy

$$\begin{aligned}f(0) &= 3, \text{ so } c_0 = 3 \\f(1) &= 3, \text{ so } c_0 + c_1 = 3 \\f(2) &= 6, \text{ so } c_0 + 2c_1 = 6.\end{aligned}$$

These resulting system of equations can be written as $A\mathbf{x} = \mathbf{b}$ where

$$A = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{pmatrix}, \quad \mathbf{x} = \begin{pmatrix} c_0 \\ c_1 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} 3 \\ 3 \\ 6 \end{pmatrix}.$$

The expression we want to minimize is then precisely $\|\mathbf{b} - A\mathbf{x}\|$. Again, the equation $A\mathbf{x} = \mathbf{b}$ has no solution since our points do not lie on the same line, so our goal is to find \mathbf{x} such that $\|\mathbf{b} - A\mathbf{x}\|$ is as small as possible; this \mathbf{x} is called the *least squares solution* of $A\mathbf{x} = \mathbf{b}$.

Note that the vectors of the form $A\mathbf{x}$ we consider make up the image of A , so we are asking to find the vector in this image which is closest to \mathbf{b} —but this is something we already know how to do! Indeed, we know that the closest such vector is the orthogonal projection of \mathbf{b} onto this image, so our goal is to find \mathbf{x} such that $A\mathbf{x} = \text{proj}_{\text{im } A} \mathbf{b}$.

Example 1. Let's find this \mathbf{x} using previous methods. First we need to find $\text{proj}_{\text{im } A} \mathbf{b}$. Using techniques from last quarter we see that the image of A has basis

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

Applying the Gram-Schmidt process gives

$$\mathbf{u}_1 = \begin{pmatrix} 1/\sqrt{3} \\ 1/\sqrt{3} \\ 1/\sqrt{3} \end{pmatrix}, \quad \mathbf{u}_2 = \begin{pmatrix} -1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \end{pmatrix}$$

as an orthonormal basis of $\text{im } A$. Thus

$$\text{proj}_{\text{im } A} \mathbf{b} = \text{proj}_{\mathbf{u}_1} \mathbf{b} + \text{proj}_{\mathbf{u}_2} \mathbf{b} = \frac{12}{\sqrt{3}} \begin{pmatrix} 1/\sqrt{3} \\ 1/\sqrt{3} \\ 1/\sqrt{3} \end{pmatrix} + \frac{3}{\sqrt{2}} \begin{pmatrix} -1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \end{pmatrix} = \begin{pmatrix} 5/2 \\ 4 \\ 11/2 \end{pmatrix}.$$

Now we need \mathbf{x} such that $A\mathbf{x} = \text{proj}_{\text{im } A} \mathbf{b}$. Solving

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{pmatrix} \mathbf{x} = \begin{pmatrix} 5/2 \\ 4 \\ 11/2 \end{pmatrix}$$

using row operations gives

$$\mathbf{x} = \begin{pmatrix} 5/2 \\ 3/2 \end{pmatrix}.$$

Recall that the entries of \mathbf{x} were the coefficients of the best-fitting line we're looking for, so the line which best fits the given points in the least squares sense is thus $f(t) = \frac{5}{2} + \frac{3}{2}t$.

Normal equation. It turns out there is a quicker way of finding this line without using an orthonormal basis. The key is in the condition that $A\mathbf{x} = \text{proj}_{\text{im } A} \mathbf{b}$. Since $\mathbf{b} - \text{proj}_{\text{im } A} \mathbf{b}$ should be orthogonal to $\text{im } A$ by one of the characterizations of orthogonal projections, this means that $A\mathbf{x}$ should have the property that $\mathbf{b} - A\mathbf{x}$ is orthogonal to $\text{im } A$. But as the book shows, a vector orthogonal to $\text{im } A$ is the *same* as a vector in $\ker A^T$! Thus we're looking for \mathbf{x} such that $\mathbf{b} - A\mathbf{x}$ is in $\ker A^T$, meaning that

$$A^T(\mathbf{b} - A\mathbf{x}) \text{ should be } \mathbf{0}.$$

The equation $A^T(\mathbf{b} - A\mathbf{x}) = \mathbf{0}$ is the same as $A^T A\mathbf{x} = A^T \mathbf{b}$, so the \mathbf{x} which makes $A\mathbf{x} = \text{proj}_{\text{im } A} \mathbf{b}$ is precisely the \mathbf{x} satisfying

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

This equation is called the *normal equation* of $A\mathbf{x} = \mathbf{b}$, and as mentioned before its solution is called the *least square* of $A\mathbf{x} = \mathbf{b}$.

Important. The least squares solution of $A\mathbf{x} = \mathbf{b}$ is the vector \mathbf{x} such that $A\mathbf{x} = \text{proj}_{\text{im } A} \mathbf{b}$, which is precisely the solution of the normal equation $A^T A\mathbf{x} = A^T \mathbf{b}$. Geometrically, this is the vector \mathbf{x} giving the vector $A\mathbf{x}$ in the image of A which is closest to \mathbf{b} .

Back to Example 1. Let's use the newly-derived normal equation to solve our previous example. We wanted to find \mathbf{x} such that $A\mathbf{x} = \text{proj}_{\text{im } A} \mathbf{b}$ where

$$A = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{pmatrix} \text{ and } \mathbf{b} = \begin{pmatrix} 3 \\ 3 \\ 6 \end{pmatrix}.$$

The corresponding normal equation $A^T A\mathbf{x} = A^T \mathbf{b}$ is

$$\begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{pmatrix} \mathbf{x} = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \end{pmatrix} \begin{pmatrix} 3 \\ 3 \\ 6 \end{pmatrix},$$

which becomes

$$\begin{pmatrix} 3 & 3 \\ 3 & 5 \end{pmatrix} \mathbf{x} = \begin{pmatrix} 12 \\ 15 \end{pmatrix}.$$

Solving this using whatever method we want gives $\mathbf{x} = \begin{pmatrix} 5/2 \\ 3/2 \end{pmatrix}$ as we found before. Hopefully you can see why this method is usually much quicker than the "orthonormal basis" method.

Example 2. To drive home the point, let's work out an example we did previously using this new least squares method. In the Warm-Up from last class we wanted to find the distance from $\mathbf{x} = \begin{pmatrix} -1 \\ 2 \end{pmatrix}$ to each eigenspace of

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

We found that the eigenspace corresponding to 2 had basis

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

Let us compute $\text{proj}_{E_2} \mathbf{x}$ using least squares. To be able to do this, we need to express the space we're projecting onto as the image of some matrix, but this is easy: we use the basis vectors for that subspace as the columns of the matrix. In other words, for

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

we have $E_2 = \text{im } B$. So to find $\text{proj}_{\text{im } B} \mathbf{x}$ we first solve the normal equation $B^T B \mathbf{y} = B^T \mathbf{x}$ for \mathbf{y} . This normal equation is

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

or

$$\begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \mathbf{y} = \begin{pmatrix} 3 \\ 3 \end{pmatrix}.$$

Solving gives $\mathbf{y} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$. Since this least-squares solution satisfies $\text{proj}_{\text{im } B} \mathbf{x} = A\mathbf{y}$, this orthogonal projection we want is

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

If you go back to the previous Warm-Up in question you'll see that this is the same answer for $\text{proj}_{E_2} \mathbf{x}$ we found there. Yay!

Lecture 6: Symmetric Matrices

Today we spoke about symmetric matrices and their amazing properties. The culmination of these ideas is the so-called *Spectral Theorem*, which says that symmetric matrices are the same as "orthogonally diagonalizable" ones.

Warm-Up 1. We find the quadratic function $f(t) = c_0 + c_1 t + c_2 t^2$ which best fits the data points $(0, 4)$, $(1, 3)$, $(2, 6)$, and $(-1, 3)$ in the least-squares sense. Recall the idea: there is no function of the specified form which passes through all four given points, so we are looking for the function which comes as close as possible to doing so.

The condition that f pass through the given points gives the following system of equations:

$$\begin{aligned} f(0) = 4 &\implies c_0 = 4 \\ f(1) = 3 &\implies c_0 + c_1 + c_2 = 3 \\ f(2) = 6 &\implies c_0 + 2c_1 + 4c_2 = 6 \\ f(-1) = 3 &\implies c_0 - c_1 + c_2 = 3, \end{aligned}$$

which can be written in matrix form $A\mathbf{x} = \mathbf{b}$ as

$$\begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & -1 & 1 \end{pmatrix} \begin{pmatrix} c_0 \\ c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 4 \\ 3 \\ 6 \\ 3 \end{pmatrix}.$$

(You can check that this system indeed has no solution.) The least squares solution $\mathbf{x} = \begin{pmatrix} c_0 \\ c_1 \\ c_2 \end{pmatrix}$ of $A\mathbf{x} = \mathbf{b}$ is the actual solution of the corresponding normal equation $A^T A\mathbf{x} = A^T \mathbf{b}$, which is:

$$\begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & -1 \\ 0 & 1 & 4 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & -1 & 1 \end{pmatrix} \mathbf{x} = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & -1 \\ 0 & 1 & 4 & 1 \end{pmatrix} \begin{pmatrix} 4 \\ 3 \\ 6 \\ 3 \end{pmatrix},$$

or

$$\begin{pmatrix} 4 & 2 & 6 \\ 2 & 6 & 8 \\ 6 & 8 & 18 \end{pmatrix} \mathbf{x} = \begin{pmatrix} 16 \\ 12 \\ 30 \end{pmatrix}.$$

Solving this gives

$$\mathbf{x} = \begin{pmatrix} 31/10 \\ 3/10 \\ 1/2 \end{pmatrix},$$

so the quadratic function which best fits the given data points is $f(t) = \frac{31}{10} + \frac{3}{10}t + \frac{1}{2}t^2$.

Warm-Up 2. (This second Warm-Up is actually the same as the first Warm-Up in disguise.) Say we want to find the vector in

$$V = \text{span} \left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 2 \\ -1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 4 \\ -1 \end{pmatrix} \right\}$$

which is closest to

$$\mathbf{b} = \begin{pmatrix} 4 \\ 3 \\ 6 \\ 2 \end{pmatrix}.$$

We know that this vector should be $\text{proj}_V \mathbf{b}$, and I claim that we can easily compute this based on what we did in the first Warm-Up. Indeed, noting that the space V can also be described as the image of the matrix

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

we are looking for the orthogonal projection of \mathbf{b} onto $\text{im } A$ and the point is that the least squares solution \mathbf{x} we computed in the first Warm-Up precisely has the property that $A\mathbf{x}$ is this projection! Thus,

$$A\mathbf{x} = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & -1 & 1 \end{pmatrix} \begin{pmatrix} 31/10 \\ 3/10 \\ 1/2 \end{pmatrix} = \begin{pmatrix} 31/10 \\ 39/10 \\ 33/10 \end{pmatrix}$$

is $\text{proj}_V \mathbf{b} = \text{proj}_{\text{im } A} \mathbf{b}$, so this is the vector in V which is closest to \mathbf{b} .

Why are transposes important? Before talking about symmetric matrices, let's be clear about why the transpose of a matrix is a useful thing to look at. The key is the following fact: for any square matrix A and vectors \mathbf{v} and \mathbf{w} ,

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

In other words, when multiplying one vector in a dot product by a matrix A , the resulting dot product is the same as the one we would get when instead multiplying the other vector by A^T . It is this “moving around a dot product expression” property of transposes which accounts for their usefulness.

In particular, if A is symmetric (so that $A = A^T$) it does not matter which vector we multiply by A , the resulting dot products are always the same: $A\mathbf{v} \cdot \mathbf{w} = \mathbf{v} \cdot A\mathbf{w}$ for any \mathbf{v} and \mathbf{w} .

Key properties of symmetric matrices. Suppose that A is symmetric. Then:

- **All eigenvalues of A exist and are real.** Say that $\lambda = a + ib$ is a complex eigenvalue (which always exists) of A with complex eigenvector $\mathbf{v} + i\mathbf{w}$. Then as we saw at the end of last quarter $\bar{\lambda} = a - ib$ is also an eigenvalue of A with complex eigenvector $\mathbf{v} - i\mathbf{w}$. We compute:

$$\begin{aligned} A(\mathbf{v} + i\mathbf{w}) \cdot (\mathbf{v} - i\mathbf{w}) &= \lambda(\mathbf{v} + i\mathbf{w}) \cdot (\mathbf{v} - i\mathbf{w}) \\ &= \lambda(\mathbf{v} \cdot \mathbf{v} - i\mathbf{v} \cdot \mathbf{w} + i\mathbf{w} \cdot \mathbf{v} + \mathbf{w} \cdot \mathbf{w}) \\ &= \lambda(\|\mathbf{v}\|^2 + \|\mathbf{w}\|^2) \\ (\mathbf{v} + i\mathbf{w}) \cdot A(\mathbf{v} - i\mathbf{w}) &= (\mathbf{v} + i\mathbf{w}) \cdot \bar{\lambda}(\mathbf{v} - i\mathbf{w}) \\ &= \bar{\lambda}(\mathbf{v} \cdot \mathbf{v} - i\mathbf{v} \cdot \mathbf{w} + i\mathbf{w} \cdot \mathbf{v} + \mathbf{w} \cdot \mathbf{w}) \\ &= \bar{\lambda}(\|\mathbf{v}\|^2 + \|\mathbf{w}\|^2). \end{aligned}$$

Since A is symmetric $A(\mathbf{v} + i\mathbf{w}) \cdot (\mathbf{v} - i\mathbf{w})$ should equal $(\mathbf{v} + i\mathbf{w}) \cdot A(\mathbf{v} - i\mathbf{w})$, so we get that $\lambda = \bar{\lambda}$; i.e. $a + ib = a - ib$ so $b = 0$ and the eigenvalue $a + ib = a$ is actually real.

- **Eigenvectors of A with different eigenvalues are orthogonal.** Say that \mathbf{v} and \mathbf{w} are eigenvectors of A with eigenvalues $\lambda \neq \mu$ respectively. Then

$$A\mathbf{v} \cdot \mathbf{w} = \lambda(\mathbf{v} \cdot \mathbf{w}) \text{ and } \mathbf{v} \cdot A\mathbf{w} = \mathbf{v} \cdot (\mu\mathbf{w}) = \mu(\mathbf{v} \cdot \mathbf{w}).$$

Since A is symmetric these two expressions are equal, so since $\lambda \neq \mu$ it must be that $\mathbf{v} \cdot \mathbf{w} = 0$.

- **A is diagonalizable.** This is not easy to justify in general, and the book's proof is not very enlightening, but here is an “intuitive” idea why it is true, at least for a 3×3 symmetric matrix. (This same idea generalizes to symmetric matrices of any size.) We know from the first property above that A has at least one real eigenvalue and so at least one eigenvector \mathbf{v}_1 . This eigenvector spans some line in \mathbb{R}^3 , so its orthogonal complement is some 2-dimensional plane in \mathbb{R}^3 . Now, here is the key fact: A *preserves* this orthogonal complement in the sense that if \mathbf{x} is on this orthogonal plane, then $A\mathbf{x}$ remains on this orthogonal plane. So, we can view A as describing a transformation from this orthogonal plane to itself. Applying the first property again now tells us that A has some eigenvector \mathbf{v}_2 on this plane. The space of vectors orthogonal to both \mathbf{v}_1 and \mathbf{v}_2 is now a line, and again A preserves this line, so viewing A as a transformation from this line to itself gives a third eigenvector \mathbf{v}_3 when applying the first property one more time. We end up with three orthogonal (and hence linearly independent) eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$, so A is diagonalizable.

Example 1. Take A to be the matrix

$$A = \begin{pmatrix} 3 & 1 & 1 \\ 1 & 3 & 1 \\ 1 & 1 & 3 \end{pmatrix},$$

which we used in the Warm-Up from January 15th. There we said the eigenvalues of A were 2 and 5, with bases for the eigenspaces being

$$\begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix} \text{ for } E_2 \text{ and } \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \text{ for } E_5.$$

As expected from the key properties listed above, A has all real eigenvalues, is diagonalizable, and eigenvectors corresponding to different eigenvalues are orthogonal.

In the previous Warm-Up with this matrix, using the Gram-Schmidt process we also found orthonormal bases for each eigenspace, which were:

$$\begin{pmatrix} -1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \end{pmatrix}, \begin{pmatrix} -1/\sqrt{6} \\ 2/\sqrt{6} \\ -1/\sqrt{6} \end{pmatrix} \text{ for } E_2 \text{ and } \begin{pmatrix} 1/\sqrt{3} \\ 1/\sqrt{3} \\ 1/\sqrt{3} \end{pmatrix} \text{ for } E_5.$$

Note that again, the first two (orthonormal) eigenvectors are orthogonal to the third. Putting these three together thus gives an orthonormal basis for \mathbb{R}^3 consisting of eigenvectors of A .

Example 2. Let B be the matrix

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

This has eigenvalues 2 and -3 , and possible bases for the eigenspaces are

$$\begin{pmatrix} 1 \\ 0 \\ 2 \end{pmatrix} \text{ for } E_2 \text{ and } \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -2 \\ 1 \\ 1 \end{pmatrix} \text{ for } E_{-3}.$$

Again, B has real eigenvalues, is diagonalizable, and eigenvectors for different eigenvalues are orthogonal, as should be the case since B is symmetric. Applying Gram-Schmidt to each basis separately then gives three orthonormal eigenvectors:

$$\begin{pmatrix} 1/\sqrt{5} \\ 0 \\ 2/\sqrt{5} \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -2/\sqrt{5} \\ 0 \\ 1/\sqrt{5} \end{pmatrix},$$

which as before form an “orthonormal eigenbasis” for \mathbb{R}^3 .

Orthogonal diagonalization. Recall that to diagonalize a matrix A means to write it as $A = SDS^{-1}$ with D diagonal, which amounts to finding a basis for \mathbb{R}^n consisting of eigenvectors of A . (These eigenvectors make up the columns of S .) In the two examples above, using the orthonormal eigenbases we found as the columns of S makes S an *orthogonal* matrix. The nice thing is that, as we’ve seen, the inverse of such a matrix is simply its transpose, so S^{-1} is easy to write down explicitly.

The point is that we can thus actually diagonalize the matrices from the two examples in a particularly nice way, leading to the definition of “orthogonally diagonalizable”:

A square matrix A is *orthogonally diagonalizable* if we can diagonalize it as $A = QDQ^T$ where D is diagonal and Q is an orthogonal matrix, in which case $Q^T = Q^{-1}$.

Note that what makes this work in the two examples above is that since eigenvectors for different eigenvalues were already orthogonal, applying Gram-Schmidt to each eigenspace separately still results in orthogonal vectors *even when* we put eigenvectors with different eigenvalues together into one big list; this is not necessarily going to be true for matrices in general, and in fact we'll see in a second that it's only true for symmetric matrices.

Back to Example 1. Using the orthonormal eigenbasis we found, we can orthogonally diagonalize A as

$$\begin{pmatrix} 3 & 1 & 1 \\ 1 & 3 & 1 \\ 1 & 1 & 3 \end{pmatrix} = \begin{pmatrix} -1/\sqrt{2} & -1/\sqrt{6} & 1/\sqrt{3} \\ 0 & 2/\sqrt{6} & 1/\sqrt{3} \\ 1/\sqrt{2} & -1/\sqrt{6} & 1/\sqrt{3} \end{pmatrix} \begin{pmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 5 \end{pmatrix} \begin{pmatrix} -1/\sqrt{2} & 0 & 1/\sqrt{2} \\ -1/\sqrt{6} & 2/\sqrt{6} & -1/\sqrt{6} \\ 1/\sqrt{3} & 1/\sqrt{3} & 1/\sqrt{3} \end{pmatrix}.$$

We can now do awesome things like (fairly) easily compute arbitrary powers of A .

Back to Example 2. We can orthogonally diagonalize B as

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

As an application of this, say we want to find a matrix C such that $C^3 = B$. Intuitively, we want to take a “cube root” of B , which we can do using this diagonalization simply by taking the cube root of each entry in the diagonal matrix part. Indeed, the matrix

$$C = \begin{pmatrix} 1/\sqrt{5} & 0 & -2/\sqrt{5} \\ 0 & 1 & 0 \\ 2/\sqrt{5} & 0 & 1/\sqrt{5} \end{pmatrix} \begin{pmatrix} \sqrt[3]{2} & 0 & 0 \\ 0 & -\sqrt[3]{3} & 0 \\ 0 & 0 & -\sqrt[3]{3} \end{pmatrix} \begin{pmatrix} 1/\sqrt{5} & 0 & 2/\sqrt{5} \\ 0 & 1 & 0 \\ -2/\sqrt{5} & 0 & 1/\sqrt{5} \end{pmatrix}$$

satisfies $C^3 = B$ as required.

Spectral Theorem. What remains is to determine the matrices for which orthogonal diagonalization is possible. As said before, to be able to carry this out for an $n \times n$ matrix we need to end up with n orthonormal eigenvectors in the end, which can only happen if eigenvectors for different eigenvalues are already orthogonal.

This works for symmetric matrices, and now we can justify that this *only* works for symmetric matrices. Indeed, if $A = QDQ^T$ with D diagonal and Q orthogonal, then

$$A^T = (QDQ^T)^T = (Q^T)^T D^T Q^T = QD^T Q^T = QDQ^T = A,$$

so that A must in fact be symmetric. Thus we get the statement of the famous *Spectral Theorem*: a matrix is orthogonally diagonalizable if and only if it is symmetric. In other words, “symmetric” and “orthogonally diagonalizable” mean the same thing! This is quite surprising, since the definitions of these two terms really seem to be worlds apart. Next time I'll hint at some possible applications of this truly wonderful fact.

Important. To orthogonally diagonalize a symmetric matrix, find eigenvalues and basis eigenvectors as you normally would when diagonalizing, and then apply Gram-Schmidt to the bases you find for each eigenspace separately, using the resulting orthonormal eigenvectors as the columns of Q . The phrases: “ A is symmetric”, “ A is orthogonally diagonalizable”, and “can find an orthonormal eigenbasis of \mathbb{R}^n corresponding to A ” all mean the same thing.

Lecture 7: Quadratic Forms

Today we spoke about quadratic forms, which is a nice application of orthogonally diagonalizing symmetric matrices. Later we will use these ideas to study surfaces and to classify extrema of multivariable functions.

Warm-Up 1. Suppose that A is a 3×3 matrix with eigenvalues $1, 1, -3$ and corresponding eigenvectors

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

We want to orthogonally diagonalize A . Note that according to the Spectral Theorem this should only be possible if A is symmetric, which is not something we are told. However, we can see right away that A is symmetric since both eigenvectors corresponding to 1 are orthogonal to the eigenvector corresponding to -3 , and only symmetric matrices have this property.

To get orthonormal eigenvectors we apply Gram-Schmidt to the vectors we have for each eigenspace separately. We get the orthonormal bases:

$$\begin{pmatrix} 2/3 \\ 1/3 \\ 2/3 \end{pmatrix}, \begin{pmatrix} -2/3 \\ 2/3 \\ 1/3 \end{pmatrix} \text{ for } E_1 \text{ and } \begin{pmatrix} 1/3 \\ 2/3 \\ -2/3 \end{pmatrix} \text{ for } E_{-3}.$$

All three together give an orthonormal eigenbasis for \mathbb{R}^3 , so we can orthogonally diagonalize A as

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

Note that as a consequence, we now see that there is only one matrix having the given eigenvalues and eigenvectors—namely the product above—and that it is symmetric.

Warm-Up 2. With the same setup as above, let's compute $A \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$. Now, using the diagonalization from above we can actually determine A explicitly and then multiply it by $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$, but here we want to do this computation without ever figuring out exactly what A is. The point is that we can easily write $\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$ as a linear combination of the eigenbasis we found above since those basis vectors are orthonormal, and after this we can easily determine what happens when we multiply through by A since A will just scale each eigenvector by its eigenvalue.

We have

$$\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \frac{5}{3} \begin{pmatrix} 2/3 \\ 1/3 \\ 2/3 \end{pmatrix} + \frac{1}{3} \begin{pmatrix} -2/3 \\ 2/3 \\ 1/3 \end{pmatrix} + \frac{1}{3} \begin{pmatrix} 1/3 \\ 2/3 \\ -2/3 \end{pmatrix}$$

using the fact that $\mathbf{x} = (\mathbf{x} \cdot \mathbf{u}_1)\mathbf{u}_1 + \cdots + (\mathbf{x} \cdot \mathbf{u}_n)\mathbf{u}_n$ when $\mathbf{u}_1, \dots, \mathbf{u}_n$ form an orthonormal basis of \mathbb{R}^n . Then multiplying through by A gives:

$$A \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \frac{5}{3} A \begin{pmatrix} 2/3 \\ 1/3 \\ 2/3 \end{pmatrix} + \frac{1}{3} A \begin{pmatrix} -2/3 \\ 2/3 \\ 1/3 \end{pmatrix} + \frac{1}{3} A \begin{pmatrix} 1/3 \\ 2/3 \\ -2/3 \end{pmatrix}$$

$$\begin{aligned}
&= \frac{5}{3} \begin{pmatrix} 2/3 \\ 1/3 \\ 2/3 \end{pmatrix} + \frac{1}{3} \begin{pmatrix} -2/3 \\ 2/3 \\ 1/3 \end{pmatrix} + \frac{1}{3}(-3) \begin{pmatrix} 1/3 \\ 2/3 \\ -2/3 \end{pmatrix} \\
&= \begin{pmatrix} 10/9 \\ 5/9 \\ 10/9 \end{pmatrix} + \begin{pmatrix} -2/9 \\ 2/9 \\ 1/9 \end{pmatrix} + \begin{pmatrix} -3/9 \\ -6/9 \\ 6/9 \end{pmatrix} \\
&= \begin{pmatrix} 5/9 \\ 1/9 \\ 17/9 \end{pmatrix}.
\end{aligned}$$

And we are done, having never found A explicitly.

Can you hear the shape of a drum? (This is purely for your own interest, and is not standard course material.) The idea from the second Warm-Up that computations involving a symmetric matrix can be carried out without knowing it explicitly as long we know eigenvalues and orthonormal eigenvectors is an important one in many applications. Indeed, the study of eigenvalues and eigenvectors of symmetric matrices (and their generalizations) has developed into its own branch of mathematics called *Spectral Theory*. Here is one illustration of these ideas.

Say there is a drum which we cannot see but which we can hear; the sound which the drum makes when hit depends on the shape of the drum. The question is: from the sound we hear alone, can we determine what shape the drum must have had? Here is how this relates to what we've been studying.

The geometry of the drum can be used to construct a certain matrix called the *Laplacian* of the drum, which encodes some geometric properties of the drum. This turns out to be a huge matrix with an infinite number of rows and columns, but the saving grace is that it is symmetric! Thus, an analog of the Spectral Theorem in this infinite-dimensional setting still says that the Laplacian is orthogonally diagonalizable, so we get eigenvalues and orthonormal eigenvectors. (Probably an infinite number of each.) It turns out that these data are related to the sound waves produced by the drum: the eigenvalues describe the frequencies of the waves and the orthonormal eigenvectors the structure of the wave. (Surprising, no?)

The question is whether we can reverse this process: given the sound, recover the geometry. From the sound we hear we can determine the frequencies of the sound waves and their structure, and hence the eigenvalues and orthonormal eigenvectors of the (at this point) unknown Laplacian. The same idea as in the second Warm-Up (courtesy of the Spectral Theorem) now says that from this we can determine how the Laplacian behaves, without having to write it down explicitly (which would likely be hard since it is a matrix of infinite size). So, now that we have the Laplacian, all that remains is to figure out what geometric shape would have given rise to that Laplacian.

Unfortunately, this is not something we can do exactly, since different shapes can actually have the same Laplacian. BUT, these different shapes can be classified, so that if we know the Laplacian we can at least determine that the shape of the drum must belong to some easy-to-manage (hopefully finite) list. So, we cannot hear the shape of the drum precisely... but we can come pretty close!

Important. A symmetric matrix is completely determined by its eigenvalues and orthonormal eigenvectors. In other words, if two symmetric matrices have the same eigenvalues and associated orthonormal eigenvectors, then they are actually the same matrix.

Quadratic forms. A *quadratic form* (say in two variables x, y for now) is a function which only involves quadratic terms: x^2, xy , and y^2 . (So no linear terms and no additional +constant terms.)

The basic fact is that any such function q can be written as

$$q(x, y) = \mathbf{x} \cdot A\mathbf{x}$$

where $\mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}$ and A is a symmetric matrix called the *matrix of the quadratic form*. (We'll work this out in some examples in a bit.) The point is then the following: by orthogonally diagonalizing the matrix A , we can come up with a new system of coordinates to use which will simplify the description of the form, and which will thus make studying the form simpler.

Example 1. Consider the quadratic form

$$q(x, y) = -7x^2 + 8xy - 13y^2.$$

The matrix of this form is

$$A = \begin{pmatrix} -7 & 4 \\ 4 & -13 \end{pmatrix},$$

whose entries come simply from the coefficients of the various terms in the form. (The 4 comes from half the coefficient 8 of xy .) Indeed, let us compute:

$$\begin{aligned} \mathbf{x} \cdot A\mathbf{x} &= \begin{pmatrix} x \\ y \end{pmatrix} \cdot \begin{pmatrix} -7 & 4 \\ 4 & -13 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \\ &= \begin{pmatrix} x \\ y \end{pmatrix} \cdot \begin{pmatrix} -7x + 4y \\ 4x - 13y \end{pmatrix} \\ &= -7x^2 + 4xy + 4yx - 13y^2 \\ &= q(x, y), \end{aligned}$$

so A is really is the matrix of q . Note that we needed to use 4 in the matrix since there are two terms in the expression we get from $\mathbf{x} \cdot A\mathbf{x}$ which involve xy , so to get coefficient 8 in total we need each of those terms to have coefficient 4. (Also, having coefficient 4 for both the xy piece and yx piece in $\mathbf{x} \cdot A\mathbf{x}$ guarantees that A is symmetric, as we want.)

Now, A has eigenvalues $-5, -15$ with associated orthonormal eigenvectors

$$\begin{pmatrix} 2/\sqrt{5} \\ 1/\sqrt{5} \end{pmatrix}, \begin{pmatrix} -1/\sqrt{5} \\ 2/\sqrt{5} \end{pmatrix}$$

respectively. Let c_1, c_2 be coordinates relative to this basis of \mathbb{R}^2 . Recall from last quarter that this means for a given \mathbf{x} , its coordinates are the values of c_1, c_2 satisfying

$$\mathbf{x} = c_1 \begin{pmatrix} 2/\sqrt{5} \\ 1/\sqrt{5} \end{pmatrix} + c_2 \begin{pmatrix} -1/\sqrt{5} \\ 2/\sqrt{5} \end{pmatrix}.$$

With respect to these coordinates, the equation for the quadratic form becomes

$$-7x^2 + 8xy - 13y^2 = -5c_1^2 - 15c_2^2,$$

and the point is that we've eliminated the mixed term. (We'll see in a second why this is the correct equation.) Note that the resulting coefficients are simply the eigenvalues of A .


To see why this is useful, consider the problem of sketching the curve

$$-7x^2 + 8xy - 13y^2 = -1.$$


Relative to our new coordinates this equation becomes

$$-5c_1^2 - 15c_2^2 = -1, \text{ or } 5c_1^2 + 15c_2^2 = 1,$$

which is the equation of an ellipse! (Note that the number which we set the quadratic form equal to is important; for instance, $-7x^2 + 8xy - 13y^2 = 1$ is not an ellipse, and in fact has no solutions since $-5c_1^2 - 15c_2^2 = 1$ has no solutions.) To draw this ellipse, we draw the set of axes determined by the eigenvectors:



Then the ellipse is



where the intercepts with the c_1 and c_2 -axes are determined as follows: the c_1 intercept occurs when $c_2 = 0$, so when $5c_1^2 = 1$, giving $c_1 = \pm 1/\sqrt{5}$, and the c_2 intercept occurs when $c_1 = 0$ so when $15c_2^2 = 1$, giving $c_2 = \pm 1/\sqrt{15}$.

Why a change of coordinates works. To complete the example above, we should justify the claim that

$$q(x, y) = -7x^2 + 8xy - 13y^2 \text{ becomes } q(c_1, c_2) = -5c_1^2 - 15c_2^2$$

after a change of coordinates. For a general quadratic form $q(\mathbf{x}) = \mathbf{x} \cdot A\mathbf{x}$, we orthogonally diagonalize A as $A = QDQ^T$ with the columns of the orthogonal matrix Q being the orthonormal eigenvectors of A and the diagonal entries of D being the eigenvalues of A . Denote by $\vec{c} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}$ the coordinates of a vector $\mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}$ relative to our orthonormal eigenbasis. In particular, recall from

that quarter that Q is then the “change of basis” matrix satisfying $\mathbf{x} = Q\vec{c}$, meaning Q tells us how to move from new coordinates to standard coordinates.

Plugging in some substitutions, we can now compute:

$$\begin{aligned}\mathbf{x} \cdot A\mathbf{x} &= (Q\vec{c}) \cdot [(QDQ^T)(Q\vec{c})] \\ &= Q\vec{c} \cdot QD\vec{c}, \text{ since } Q^TQ = I \\ &= \vec{c} \cdot D\vec{c}, \text{ since } Q \text{ preserves dot products.}\end{aligned}$$

Thus the quadratic form $\mathbf{x} \cdot A\mathbf{x}$ becomes $\vec{c} \cdot D\vec{c}$ in new coordinates, and working this expression out for $D = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$ gives

$$\vec{c} \cdot D\vec{c} = \lambda_1 c_1^2 + \lambda_2 c_2^2$$

as claimed. Note that the reason why mixed terms are eliminated is because D is diagonal, and so gives coefficient 0 for the mixed terms.

Important. After a change of coordinates, any quadratic form $q(\mathbf{x}) = \mathbf{x} \cdot A\mathbf{x}$ can be written as

$$q(c_1, \dots, c_n) = \lambda_1 c_1^2 + \dots + \lambda_n c_n^2$$

where $\lambda_1, \dots, \lambda_n$ are the eigenvalues of A and c_1, \dots, c_n are coordinates relative to an orthonormal eigenbasis of \mathbb{R}^n corresponding to A .

Example 2. We sketch the curve determined by

$$-3x^2 + 6xy + 5y^2 = 1.$$

The quadratic form $q(x, y) = -3x^2 + 6xy - 5y^2$ has matrix

$$\begin{pmatrix} -3 & 3 \\ 3 & -5 \end{pmatrix},$$


which has eigenvalues $-4, 6$ and associated orthonormal eigenvectors

$$\begin{pmatrix} 3/\sqrt{10} \\ -1/\sqrt{10} \end{pmatrix}, \begin{pmatrix} 1/\sqrt{10} \\ 3/\sqrt{10} \end{pmatrix}.$$

Relative to coordinates c_1, c_2 determined by this basis of \mathbb{R}^2 , the quadratic form q is given by $q(c_1, c_2) = -4c_1^2 + 6c_2^2$, so the given curve becomes

$$-4c_1^2 + 6c_2^2 = 1.$$

This describes a hyperbola which crosses the c_2 -axis. (We can see this since setting $c_2 = 0$ in the equation for the curve gives no solutions for c_1 , meaning that the curve cannot cross the c_1 -axis.) The intercepts on the c_2 -axis are $\pm 1/\sqrt{6}$ (found by setting $c_1 = 0$), so the hyperbola looks like:



Lecture 8: Curves and Lines

Today we started talking about parametric curves, and lines in particular. Parametric equations give us a concrete way to talk about arbitrary curves in 2 and 3-dimensions, where much of the calculus we will eventually talk about takes place.

Warm-Up 1. We sketch the curve in \mathbb{R}^2 given by the equation

$$6x_1^2 + 4x_1x_2 + 3x_2^2 = 1.$$

The left-hand side defines a quadratic form with matrix

$$\begin{pmatrix} 6 & 2 \\ 2 & 3 \end{pmatrix}.$$


This has eigenvalues 2 and 7, with possible corresponding orthonormal eigenvectors respectively:

$$\begin{pmatrix} 1/\sqrt{5} \\ -2/\sqrt{5} \end{pmatrix} \text{ and } \begin{pmatrix} 2/\sqrt{5} \\ 1/\sqrt{5} \end{pmatrix}.$$

In coordinates c_1 and c_2 relative to this basis of \mathbb{R}^2 the equation for the curve becomes

$$2c_1^2 + 7c_2^2 = 1,$$

so the curve is an ellipse. The intercepts with the c_1 -axis occur when $c_2 = 0$, so at $c_1 = \pm 1/\sqrt{2}$, and the intercepts with the c_2 -axis occur when $c_1 = 0$, so at $c_2 = \pm 1/\sqrt{7}$. The ellipse thus looks like:



Note that the part of the c_1 -axis in the fourth quadrant is the *positive* c_1 -axis since the eigenvector spanning that line points in this direction, which is why the intersection point there is labeled with a positive $1/\sqrt{2}$.

Warm-Up 2. Now we determine the point (or points) on the surface $-x_1^2 + 2x_2x_3 = 1$ in \mathbb{R}^3 which is (or are) closest to the origin. The quadratic form $q(x_1, x_2, x_3) = -x_1^2 + 2x_2x_3$ has matrix

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


which has eigenvalues -1 and 1 , with orthonormal eigenvectors

$$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ -1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} \text{ for } -1 \text{ and } \begin{pmatrix} 0 \\ 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} \text{ for } 1.$$

Taking coordinates c_1, c_2, c_3 relative to this orthonormal basis of \mathbb{R}^3 , the equation of the surface becomes

$$-c_1^2 - c_2^2 + c_3^2 = 1.$$

To get a sense for what this surface looks like, we note the following. First, setting $c_3 = 0$ in the given equation gives $-c_1^2 - c_2^2 = 1$, which has no solutions since the left side is never positive. This means that our surface never crosses the plane $c_3 = 0$, which is the c_1c_2 -plane. Second, setting $c_1 = 0$ and $c_2 = 0$ in the given equation determines the points where the surface intersects the c_3 -axis: $c_3^2 = 1$ so $c_3 = \pm 1$. We'll see later how to determine precisely what this surface looks like, but for now I claim that it is what's called a *hyperboloid of two sheets*, which is a 3-dimensional analog of a hyperbola. It looks like:



where we draw the c_1, c_2, c_3 -axes as if they were the ordinary x_1, x_2, x_3 -axes. (With respect to the standard axes, the actual hyperboloid we're looking at would be a rotated version of the one drawn above.)

Thus we see that the points closest to the origin occur when $c_1 = c_2 = 0$ and $c_3 = \pm 1$. In terms of standard coordinates, these are the points

$$0 \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + 0 \begin{pmatrix} 0 \\ -1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} \pm \begin{pmatrix} 0 \\ 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix},$$

so $(0, 1/\sqrt{2}, 1/\sqrt{2})$ and $(0, -1/\sqrt{2}, -1/\sqrt{2})$.

Parametric Curves. We describe arbitrary curves in 2 and 3-dimensions using *parametric equations*, which are equations:

$$\begin{cases} x = x(t) \\ y = y(t) \\ z = z(t) \end{cases}$$

giving the (x, y, z) -coordinates of points along the curve in terms of a parameter t . The idea is that as t varies, we get various points $(x(t), y(t), z(t))$ which trace out the curve we're considering. We can describe only a piece of a curve by restricting the values of the parameter we consider.

For instance, the parametric equations

$$\begin{cases} x = \cos t \\ y = \sin t \end{cases}$$

give a unit circle. Indeed, for these equations one can check that $x^2 + y^2$ does equal 1 (so we are on the unit circle), and that varying gives the entire circle. (Actually, the entire circle is already traced out for $0 \leq t \leq 2\pi$.) In this case, the parameter t simply describes an angle, and the circle is traced out counterclockwise starting at $(1, 0)$ when $t = 0$.

The parametric equations

$$\begin{cases} x = \sin t \\ y = \cos t \end{cases} \quad 0 \leq t \leq \pi$$

describe only the right half of the unit circle due to the restriction on t . In this case, the curve is traced out *clockwise* starting at $(0, 1)$ when $t = 0$.

Finally, the parametric equations

$$\begin{cases} x = \cos t - \sin t \\ y = \cos t + \sin t \end{cases}$$

describe a circle of radius $\sqrt{2}$ centered at the origin. To see this, note that these equations can be written in vector form as

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{pmatrix} \begin{pmatrix} \sqrt{2} \cos t \\ \sqrt{2} \sin t \end{pmatrix}.$$

The parametric equations $x = \sqrt{2} \cos t, y = \sqrt{2} \sin t$ describe the circle of radius 2 centered at the origin, and multiplying this by the given matrix (which is the matrix of a rotation by $\pi/4$) does not change the shape of the circle. (Note, however, that the starting points at $t = 0$ for these two sets of parametric equations of the circle of radius $\sqrt{2}$ are different.)

Lines. Consider the curve with parametric equations

$$\begin{cases} x = 1 + t \\ y = 5 + 2t \\ z = 3t. \end{cases}$$

We claim that this is a line in \mathbb{R}^3 . Indeed, we can rewrite this set of equations in vector form as

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 1 \\ 5 \\ 0 \end{pmatrix} + t \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix},$$

which shows that this is the line parallel to the one spanned by $\begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$ only it is translated by $\begin{pmatrix} 1 \\ 5 \\ 0 \end{pmatrix}$.

Taking $t = 0$ shows that $(1, 5, 0)$ is on this line, and $t = 1$ gives $(2, 7, 3)$ as another point on this line. The vector $(1, 2, 3)$ (from now on we write vectors as either rows or columns depending on which notation is more useful for the task at hand) giving the direction of the line is precisely the vector *from* $(1, 5, 0)$ *to* $(2, 7, 3)$:

$$\text{end point} - \text{start point} = (2, 7, 3) - (1, 5, 0) = (1, 2, 3).$$

So, we can describe this line as the line in \mathbb{R}^3 passing through $(1, 5, 0)$ and $(2, 7, 3)$, or as the line through $(1, 5, 0)$ which is parallel to the vector $(1, 2, 3)$.

Important. Given a point (x_0, y_0, z_0) on a line and a vector (a, b, c) parallel to the line, the equation of the line is given in *vector form* by

$$\mathbf{r}(t) = \mathbf{b} + t\mathbf{a},$$

where $\mathbf{r} = (x, y, z)$, $\mathbf{b} = (x_0, y_0, z_0)$ and $\mathbf{a} = (a, b, c)$. Working out the x, y, z -coordinates of this, the parametric equations for this line are

$$\begin{cases} x = x_0 + at \\ y = y_0 + bt \\ z = z_0 + ct. \end{cases}$$

Similar equations hold in 2-dimensions.

Example. We find parametric equations for the line in \mathbb{R}^2 passing through $(3, -1)$ and perpendicular to the line

$$\begin{cases} x = 1 + 5t \\ y = 2 - 2t. \end{cases}$$

For this we need two things: a point on the line we want (which we are given) and a vector parallel to the line we want. Now, the line whose equations are given is parallel to the vector $(5, -2)$, so in order to be perpendicular to this our line should be parallel to $(2, 5)$. Thus the line we want is given by

$$\mathbf{r}(t) = (3, -1) + (2, 5)t,$$

or by

$$\begin{cases} x = 3 + 2t \\ y = -1 + 5t \end{cases}$$

in parametric form.

Lecture 9: Cross Products

Today we spoke about the *cross product* of vectors, which is a way of combining two vectors to get a third which has really nice geometric properties. Most importantly, the cross product of two vectors is always perpendicular to each of those vectors—a fact which will make certain constructions we look at (both this quarter and next) simpler.

Warm-Up. We find the line L parallel to the line L' given by the parametric equations

$$\begin{cases} x = 1 + t \\ y = 5 + 2t \\ z = 3t \end{cases}$$

and passing through the intersection of L' with the plane $2x + y - z = 4$. First, the vector $\mathbf{a} = (1, 2, 3)$ is parallel to L' and hence to L as well. Second, we need to find the point of intersection of L' with the given plane, which should come from a point on the line which satisfies the equation of the plane. Plugging in the given parametric equations into the plane gives:

$$2(1 + t) + (5 + 2t) - (3t) = 4, \text{ so } t = -3.$$

This says that the intersection point we want occurs when $t = -3$ on the line L' , so at $(-2, -1, -9)$. Possible parametric equations for L are thus

$$\begin{cases} x = -2 + t \\ y = -1 + 2t \\ z = -9 + 3t. \end{cases}$$

Remark. Note that the line L from above is also given by the parametric equations

$$\begin{cases} x = -2 + t^9 \\ y = -1 + 2t^9 \\ z = -9 + 3t^9 \end{cases}$$

since as t ranges through all possible values, so does t^9 . The point is that parametric equations for lines do not necessarily need to have only t to the first power, as long as the “ ta ” part of the equation $\mathbf{r}(t) = \mathbf{b} + t\mathbf{a}$ still gives all possible multiples of \mathbf{a} .

Example. Say we want to find parametric equations for the line L which is perpendicular to both of the lines

$$\begin{cases} x = 1 + 2t \\ y = 2 - t \\ z = 3 + t \end{cases} \quad \text{and} \quad \begin{cases} x = 5 - 3t \\ y = 5 - t \\ z = -2 + 2t \end{cases}$$

and which passes through their point of intersection. First we find this point of intersection. Say that it occurs along the first line when $t = t_1$ and along the second when $t = t_2$. Then we must have

$$\begin{aligned} 1 + 2t_1 &= 5 - 3t_2 \\ 2 - t_1 &= 5 - t_2 \\ 3 + t_1 &= -2 + 2t_2. \end{aligned}$$

Solving this system of equations gives $t_1 = -1$ and $t_2 = 2$, so the intersection point is at $(-1, 3, 2)$. Thus L should pass through $(-1, 3, 2)$.

Now we need a direction vector for L . This vector should be perpendicular to both lines given above, and hence to their respective direction vectors $(2, -1, 1)$ and $(-3, -1, 2)$ as well. Thus all we need to do is find a vector perpendicular to both of these. We can do this using linear algebra,

either by applying Gram-Schmidt to these vectors together with a third linearly independent one, or by solving the system of equations

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

obtained by setting the dot product of our unknown vector with each of these vectors equal to 0. Either of these methods will involve some work (especially the Gram-Schmidt method), but luckily we have a more direct way of finding such a perpendicular vectors using *cross products*.

Cross products. The *cross product* of vectors $\mathbf{u} = (u_1, u_2, u_3)$ and $\mathbf{v} = (v_1, v_2, v_3)$ is the vector defined by

$$\mathbf{u} \times \mathbf{v} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ u_1 & u_2 & u_3 \\ v_1 & v_2 & v_3 \end{vmatrix},$$

where $\mathbf{i}, \mathbf{j}, \mathbf{k}$ denote the standard basis vectors of \mathbb{R}^3 and we compute this determinant as we normally would if $\mathbf{i}, \mathbf{j}, \mathbf{k}$ were actually numbers. Cross products have some nice geometric properties:

- $\mathbf{u} \times \mathbf{v}$ is always perpendicular to both \mathbf{u} and \mathbf{v}
- the direction of $\mathbf{u} \times \mathbf{v}$ is determined by the so-called *right-hand rule*: if you line up the fingers of your right hand in the direction of \mathbf{u} and curl them towards \mathbf{v} , your thumb will point in the direction of $\mathbf{u} \times \mathbf{v}$
- as a consequence of the right hand rule, $\mathbf{w} \times \mathbf{v} = -\mathbf{v} \times \mathbf{w}$
- $\|\mathbf{u} \times \mathbf{v}\| = \|\mathbf{u}\| \|\mathbf{v}\| \sin \theta$ where θ is the angle between \mathbf{u} and \mathbf{v} , and thus $\|\mathbf{u} \times \mathbf{v}\|$ is the area of the parallelogram with sides \mathbf{u} and \mathbf{v}

We'll take a look next time at where these properties come from, and how it is that anyone ever thought of creating the cross product in the first place.

Important. From now on, anytime you need to find a vector perpendicular to two given vectors, use the cross product.

Back to Example. So now we have a direct way of finding a vector perpendicular to both $(2, -1, 1)$ and $(-3, -1, 2)$: their cross product is

$$(2, -1, 1) \times (-3, -1, 2) = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ 2 & -1 & 1 \\ -3 & -1 & 2 \end{vmatrix} = (-2 + 1)\mathbf{i} - (4 + 3)\mathbf{j} + (-2 - 3)\mathbf{j} = -\mathbf{i} - 7\mathbf{j} - 5\mathbf{k},$$

which we can also write as $(-1, -7, -5)$. Note that this is indeed perpendicular to both $(2, -1, 1)$ and $(-3, -1, 2)$. Finishing off the Example, the line L we're looking for should thus pass through $(-1, 3, 2)$ and be parallel to $(-1, -7, -5)$, so it is given by

$$\begin{cases} x = -1 - t \\ y = 3 - 7t \\ z = 2 - 5t. \end{cases}$$

Final example. We find the line passing through $(-3, 1, 5)$ which is perpendicular to the plane $x + y + z = 0$. To find a direction vector of this line, we thus have to find a vector perpendicular to this plane. We do this by finding two vectors on the plane, say $(1, -1, 0)$ and $(0, 1, -1)$, and taking their cross product:

$$(1, -1, 0) \times (0, 1, -1) = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ 1 & -1 & 0 \\ 0 & 1 & -1 \end{vmatrix} = (1, 1, 1).$$

The line we want is then given by

$$\mathbf{r}(t) = (-3 + t, 1 + t, 5 + t),$$

where we have written the parametric equations for the line in vector form.

(Note that the entries of the cross product we found are just the coefficients of the variables in the equation of the plane; this is no accident, and we'll come back to this next time.)

Lecture 10: Planes

Today we spoke about equations of planes, and using cross products to produce vectors perpendicular to planes. In the calculus we'll be doing planes will be the analogs of tangent lines, and will be useful in visualizing what multivariable derivatives mean.

Warm-Up. We find parametric equations for the line perpendicular to the plane $x - 3y + 5z = 15$ and passing through the point where this plane intersects the line with parametric equations $x = 1 - t, y = 2 + t, z = -2 + 2t$. We need a direction vector for the line and a point on the line. The point comes from the intersection of the plane with the given line, which occurs when

$$(1 - t) - 3(2 + t) + 5(-2 + 2t) = 15, \text{ so when } t = 5.$$

Thus the intersection point is $(-4, 7, 8)$.

Now, the direction vector of the line comes from a vector perpendicular to the plane. To find this we start with three points on the plane, say

$$(0, 0, 3), (0, -5, 0), \text{ and } (15, 0, 0)$$

which come from setting two variables equal to zero in the equation of the plane. Then the vectors

$$(0, -5, 0) - (0, 0, 3) = (0, -5, -3) \text{ and } (15, 0, 0) - (0, 0, 3) = (15, 0, -3)$$

are parallel to the plane, so their cross product is perpendicular to the plane. Hence

$$(0, -5, -3) \times (15, 0, -3) = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ 0 & -5 & -3 \\ 15 & 0 & -3 \end{vmatrix} = (15, -45, 75)$$

is a direction vector for the line we want. The perpendicular line thus has parametric equations:

$$\begin{cases} x = -4 + 15t \\ y = 7 - 45t \\ z = 8 + 75t. \end{cases}$$

Note that the direction vector is parallel to the vector $(1, -3, 5)$ obtained by taking the coefficients of x, y, z in the equation of the plane; this is no accident as we'll soon see.

Where do cross products come from? The formula for the cross product seems pretty mysterious and begs the question: how anyone would have ever thought of it in the first place? Let's give another definition for the cross product, which makes some of its properties clearer.

Take vectors \mathbf{u}, \mathbf{v} in \mathbb{R}^3 and define the linear transformation T from \mathbb{R}^3 to \mathbb{R} by

$$T(\mathbf{x}) = \det \begin{pmatrix} | & | & | \\ \mathbf{u} & \mathbf{v} & \mathbf{x} \\ | & | & | \end{pmatrix}.$$

This is linear due to the fact that determinants are linear in each column, as we saw last quarter. So, being linear, there should be a 1×3 matrix A satisfying $T(\mathbf{x}) = A\mathbf{x}$; but a 1×3 matrix is just the transpose of a (column) vector, so there is a vector \mathbf{b} satisfying

$$T(\mathbf{x}) = \mathbf{b}^T \mathbf{x}, \text{ which is the same as } T(\mathbf{x}) = \mathbf{b} \cdot \mathbf{x}.$$

This vector \mathbf{b} is precisely the cross product $\mathbf{u} \times \mathbf{v}$! So, the cross product arises as the 1×3 matrix representing T and thus satisfying:

$$\det \begin{pmatrix} | & | & | \\ \mathbf{u} & \mathbf{v} & \mathbf{x} \\ | & | & | \end{pmatrix} = (\mathbf{u} \times \mathbf{v}) \cdot \mathbf{x}.$$

Now the properties of cross products can be interpreted as follows: taking $\mathbf{x} = \mathbf{u}$ or $\mathbf{x} = \mathbf{v}$ gives zero determinant on the left so the dot product on the right is zero, and hence $\mathbf{u} \times \mathbf{v}$ is perpendicular to both \mathbf{u} and \mathbf{v} ; the right hand rule is related to the geometric interpretation of the sign of the determinant in that positive determinant means "orientation-preserving" (i.e. "right-hand rule"-preserving) and negative determinant means the opposite; and finally the fact that the length of a cross product gives the area of a parallelogram reflects the interpretation of determinants in terms of areas and volumes.

Planes. In order to describe a plane we need to know two things: a point (x_0, y_0, z_0) on the plane a vector \mathbf{n} perpendicular to the plane. (We say that \mathbf{n} is *normal* to the plane.) Given these, the equation of the plane in *vector form* is:

$$\mathbf{n} \cdot (\mathbf{r} - \mathbf{r}_0) = 0$$

where $\mathbf{r}_0 = (x_0, y_0, z_0)$ contains the given point and $\mathbf{r} = (x, y, z)$ encodes any other point on the plane. This equation says the following: for a point (x, y, z) to be on the plane, the vector $\mathbf{r} - \mathbf{r}_0$ between it and the point (x_0, y_0, z_0) on the plane must itself be on the plane and hence should be perpendicular to the given normal vector. If $\mathbf{n} = (a, b, c)$, working out this dot product gives the *scalar equation* of the plane:

$$a(x - x_0) + b(y - y_0) + c(z - z_0) = 0.$$

Note that the coefficients of x, y, z indeed give the entries of the normal vector, as we've alluded to earlier.

Example 1. We find an equation for the plane containing the points $(1, 3, 2)$, $(-4, 2, 1)$, and $(0, 2, 3)$. To find a normal vector to the plane, we use a cross product: the vectors

$$(-4, 2, 1) - (1, 3, 2) = (-5, -1, -1) \text{ and } (0, 2, 3) - (1, 3, 2) = (-1, -1, 1)$$

are on the plane, so their cross product

$$(-5, -1, -1) \times (-1, -1, 1) = (-2, 6, 4)$$

is normal to it. Thus with $\mathbf{n} = (-2, 6, 4)$ and $\mathbf{r}_0 = (1, 3, 2)$ we get

$$(-2, 6, 4) \cdot [(x, y, z) - (1, 3, 2)] = 0$$

or

$$-2(x - 1) + 6(y - 3) + 4(z - 2) = 0$$

as the equation of the plane. After simplifying we can also write this as

$$-2x + 6y + 4z = 24.$$

There is one more way of expressing this, which is in terms of *parametric equations* for the plane. These are equations for the x, y, z coordinates of points on the plane in terms of *two* parameters s, t . (In general, parametric equations with one parameter describe lines, and those with two parameters describe surfaces.) These parametric equations come from thinking of the plane as given by expressions of the form:

$$\mathbf{r}_0 + s\mathbf{u} + t\mathbf{v}$$

where $\mathbf{r}_0 = (x_0, y_0, z_0)$ encodes a point on the plane and \mathbf{u}, \mathbf{v} are vectors parallel to the plane. The point is that the linear combinations $s\mathbf{u} + t\mathbf{v}$ give a plane through the origin which is parallel to our plane, and then adding \mathbf{r}_0 translates this parallel plane onto our plane.

To find \mathbf{u} and \mathbf{v} we consider the plane parallel to ours but passing through the origin, which is given by

$$-2x + 6y + 4z = 0.$$

We find vectors spanning this plane; one possibility is $\mathbf{u} = (-5, -1, -1)$ and $\mathbf{v} = (-1, -1, 1)$. Then our plane can be expressed as

$$\mathbf{r}_0 + s\mathbf{u} + t\mathbf{v} = \begin{pmatrix} 1 \\ 3 \\ 2 \end{pmatrix} + s \begin{pmatrix} -5 \\ -1 \\ -1 \end{pmatrix} + t \begin{pmatrix} -1 \\ -1 \\ 1 \end{pmatrix}.$$

Combining the right side and taking the equations for the x, y, z coordinates gives

$$\begin{cases} x = 1 - 5s - t \\ y = 3 - s - t \\ z = 2 - s + t \end{cases}$$

as parametric equations for this plane.

Important. The plane containing (x_0, y_0, z_0) and normal to $\mathbf{n} = (a, b, c)$ is given by the equation

$$\mathbf{n} \cdot (\mathbf{r} - \mathbf{r}_0) = 0$$

where $\mathbf{r} = (x, y, z)$ and $\mathbf{r}_0 = (x_0, y_0, z_0)$, which is the same as the equation

$$a(x - x_0) + b(y - y_0) + c(z - z_0) = 0.$$

Given vectors \mathbf{u} and \mathbf{v} parallel to this plane, parametric equations for the plane are obtained by taking x, y, z coordinates in the vector sum

$$\mathbf{r}_0 + s\mathbf{u} + t\mathbf{v}$$

describing an arbitrary point on the plane.

Example 2. We find parametric equations for the line of intersection of the planes

$$-3x + 2y - z = 1 \text{ and } 2x - y + 2z = -8.$$

Technically we can already do this using linear algebra: view this as a system of two linear equations and solve for x, y, z using row operations; we will get one free variable, which will play the role of the parameter we need in our parametric equations. But we can also do this using new material as follows.

First we need a point on the line of intersection, so a point on both planes simultaneously. Instead of solving the system given by both planes fully, we can look for an intersection point with $z = 0$, so we only need to find x, y satisfying

$$-3x + 2y = 1 \text{ and } 2x - y = -8.$$

This gives $x = -15, y = -22$, so $(-15, -22, 0)$ is on the line of intersection.

Second we need a direction vector for this line. Since this line is on the first plane, its direction vector should be perpendicular to the normal vector $(-3, 2, -1)$ of that first plane, and similarly the direction vector of the line should also be perpendicular to the normal vector $(2, -1, 2)$ of the second plane. Thus to get a direction vector for the line of intersection we just need a vector perpendicular to both $(-3, 2, -1)$ and $(2, -1, 2)$, so their cross product works! We get

$$(-3, 2, -1) \times (2, -1, 2) = (3, 4, -1)$$

as a direction vector for the line, so the line of intersection has parametric equations

$$\begin{cases} x = -15 + 3t \\ y = -22 + 4t \\ z = -t. \end{cases}$$

Lecture 11: Polar/Cylindrical Coordinates

Today we spoke about polar and cylindrical coordinates, which give us a new (and often simpler) way of describing curves and surfaces in 2 and 3 dimensions.

Warm-Up 1. We first ask if there is a plane parallel to $5x - 3y + 2z = 10$ which contains the line $x = t + 4, y = 3t - 2, z = 5 - 2t$. If so, the direction vector for this line would have to be perpendicular to the normal vector of the plane, which is not true in this case: the direction vector is $(1, 3, -2)$ and the normal vector is $(5, -3, 2)$, and $(1, 3, -2) \cdot (5, -3, 2) \neq 0$. So, no such plane exists.

Instead let us change the z coordinate on the line to $z = 5 + 2t$, and ask for a plane parallel to $5x - 3y + 2z = 10$ containing this line instead. (Here the direction vector is perpendicular to the normal vector, so such a plane does exist.) The plane we want should also have normal vector $(5, -3, 2)$ if we want it to be parallel to $5x - 3y + 2z = 10$, so we only need a point on this new plane. Any point on the line we're looking at will be our plane, so for instance $(4, -2, 5)$ (when $t = 0$ in the line) is on the plane we want. Our plane thus has equation

$$5(x - 4) - 3(y + 2) - 2(z - 5) = 0,$$

which is the same as $5x - 3y - 2z = 16$.

Warm-Up 2. We find the distance between the planes $5x - 3y + 2z = 10$ and $10x - 6y + 4z = 30$. Note that these planes are parallel (since their normal vectors are parallel) so it makes sense to talk about the distance between them. As a nice picture in the book shows (look it up!), this distance can be obtained by orthogonally projecting a vector from a point P on the first plane to a point Q on the second plane onto the normal vector of either plane. Taking $P = (2, 0, 0)$ on the first plane and $Q = (0, -5, 0)$ on the second, we will project

$$\vec{PQ} = (0, -5, 0) - (2, 0, 0) = (-2, -5, 0).$$

With normal vector $\mathbf{n} = (5, -3, 2)$, we compute:


$$\text{proj}_{\mathbf{n}} \vec{PQ} = \frac{5}{38}(5, -3, 2).$$

The distance between the planes is then the length of this, which is:

$$\left\| \text{proj}_{\mathbf{n}} \vec{PQ} \right\| = \frac{5}{38} \sqrt{38} = \frac{5}{\sqrt{38}}.$$

There are other distance formulas in the book which might be good to look at, but notice that they all involve orthogonally projecting something onto something else.

Polar coordinates. The *polar coordinates* (r, θ) of a point (x, y) in \mathbb{R}^2 are defined as in the following picture:



So, r is the distance from the point to the origin and θ is the angle you have to move counterclockwise from the positive x -axis in order to face the point. A negative value of r is interpreted as describing a point in the direction *opposite* to θ ; for instance, $\theta = \frac{\pi}{2}$ points us in the positive y -direction and $r = -1$ then gives the point $(0, -1)$ on the negative y -axis.


By looking at the appropriate right triangles we get the following relation between polar coordinates (r, θ) and *rectangular* (or *Cartesian*) coordinates (x, y) :

$$r^2 = x^2 + y^2, \quad \tan \theta = \frac{y}{x}, \quad x = r \cos \theta, \quad y = r \sin \theta.$$


We'll see again and again that certain curves and regions in \mathbb{R}^2 are much simpler to describe in polar coordinates than in rectangular coordinates.

Example 1. We sketch the curve with polar equation $r = \sin \theta$, meaning the curve consisting of all points in \mathbb{R}^2 whose polar coordinates satisfy $r = \sin \theta$. For instance, when $\theta = 0$ we get $r = 0$ so the origin (the only point with r value 0) is on this curve; when $\theta = \frac{\pi}{2}$ we get $r = 1$ so the point $(0, 1)$ is on the this curve.


To visualize the entire curve we start with a graph of r in terms of θ :



First we focus on θ between 0 and $\frac{\pi}{2}$, which corresponds to the first quadrant. As we move counterclockwise in the first quadrant, r increases from 0 to 1, so we get a piece of the curve which looks like:

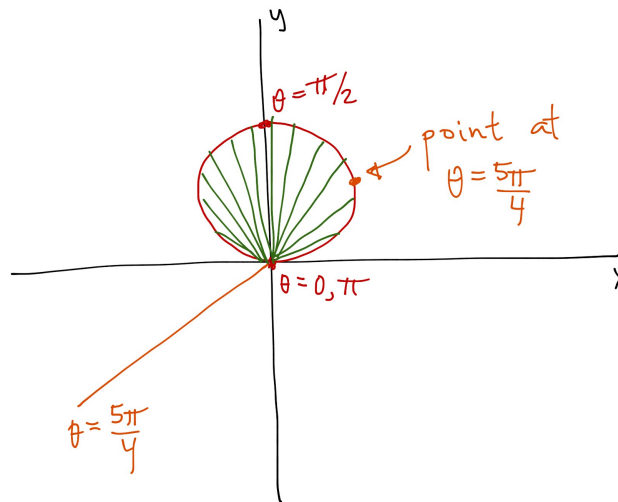


The curve is in red, and the green lines indicate the increasing value of r (distance to the origin) from 0 to 1 as θ moves counterclockwise. Now, for $\frac{\pi}{2} \leq \theta \leq \pi$ (in the second quadrant) the value of r should decrease from 1 to 0, so we get:



Again the green lines in the second quadrant indicate the decreasing value of r from 1 to 0 as θ moves from the positive y -axis to the negative x -axis. At $\theta = \pi$, $r = 0$ we are back at the origin.

Now, notice that for $\pi \leq \theta \leq \frac{3\pi}{2}$ the value of r is negative, moving from 0 to -1 . These values of θ occur in the third quadrant, but the negative value of r means that the corresponding points are actually drawn in the *opposite* (i.e. first) quadrant. For instance, at $\theta = \frac{5\pi}{4}$ we get the point labeled below:



At $\theta = \frac{3\pi}{2}$ we get the point $(0, 1)$ on the positive y -axis again, so for $\pi \leq \theta \leq \frac{3\pi}{2}$ we simply trace out the same piece of the curve we did for $0 \leq \theta \leq \frac{\pi}{2}$. For $\frac{3\pi}{2} \leq \theta \leq 2\pi$ a similar thing happens and we trace out the same piece of the curve we did for $\frac{\pi}{2} \leq \theta \leq \pi$. Thus the full curve is the red circle above the x -axis in the previous picture.

Note that we can also see this by finding the Cartesian equation of the curve. Multiplying the polar equation by r gives

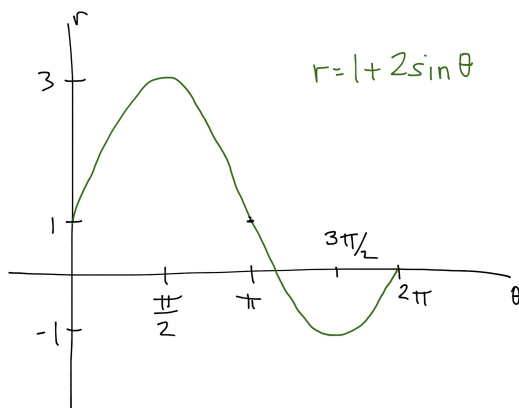
$$r^2 = r \sin \theta$$

and converting to rectangular coordinates gives


$$x^2 + y^2 = y.$$

After completing the square, this becomes $x^2 + (y - \frac{1}{2})^2 = \frac{1}{4}$, which describes a circle of radius $\frac{1}{2}$ centered at $(0, \frac{1}{2})$, which is precisely what our picture suggests.


Example 2. Now we sketch the curve with polar equation $r = 1 + 2 \sin \theta$. Again, we start with a graph of r vs θ to see how r changes as θ increases:




As θ moves from 0 to $\frac{\pi}{2}$, r increases from 1 to 3 so we get:




Then r decreases from 3 back to 1 as θ goes from $\frac{\pi}{2}$ to π in the second quadrant:




Now, at some angle between π and $\frac{3\pi}{2}$, r is zero at which point we're back at the origin; up until this angle the value of r decreases from 1 to 0 so we get:



As θ moves from this angle to $\frac{3\pi}{2}$, r is negative so we actually get a piece of the curve in the first quadrant, ending up at $(0, 1)$ when $\theta = \frac{3\pi}{2}$ and $r = -1$:



At some angle between $\frac{3\pi}{2}$ and 2π the value of r is zero again, so we're at the origin, and up until this angle we have negative values of r so we get a piece of the curve in the second quadrant; after this point we're back in the fourth quadrant with positive values of r increasing from 0 to 1 until θ hits 2π :




Going beyond $\theta = 2\pi$ simply retraces the same curve.

Multiplying the polar equation of the curve through by r gives $r^2 = r + 2r \sin \theta$, which in rectangular coordinates becomes

$$x^2 + y^2 = \sqrt{x^2 + y^2} + 2y, \text{ or } (x^2 + y^2 - 2y)^2 = x^2 + y^2,$$

and hopefully it is clear that it would be impossible to sketch this curve without magical powers given only this Cartesian equation.


Cylindrical coordinates. *Cylindrical coordinates* are 3-dimensional analogs of polar coordinates, and are defined as in the following picture:



So, r and θ are the polar coordinates of the point in the xy -plane lying below (x, y, z) , and z is the usual z -coordinate which gives height. Note that another way to interpret r is as the distance from (x, y, z) to the z -axis. As the term *cylindrical* suggests, these coordinates will be especially useful for describing surfaces “cylindrical” in shape such as cylinders or cones.


Since r and θ are the same as they were for polar coordinates, we have the same conversions between rectangular and cylindrical coordinates as we did for polar coordinates.

Example 3. The cylindrical equation $r = 1$ describes the surface consisting of all points with r value 1. Since r is distance to the z -axis, this is just a cylinder of radius 1 around the z -axis:



Here is another way to see this. In the xy -plane the polar equation $r = 1$ describes a circle of radius 1. Now, in cylindrical coordinates $r = 1$ places no restriction on z , so taking this circle and moving it up and down along the z -axis still gives points satisfying the cylindrical equation $r = 1$. This traces out the cylinder described earlier, which has Cartesian equation $x^2 + y^2 = 1$.


Example 4. Now we identify the surface with cylindrical equation $z = r$. At $z = 0$ we get $r = 0$, which describes the origin in \mathbb{R}^3 . At a height $z = 1$ we get $r = 1$ which is a circle of radius 1 in the $z = 1$ plane. Similarly, at $z = 2$, $r = 2$ so we get a circle of radius 2, and so on: as z increases our surface is traced out by circles of increasing radii, which all together give a cone:



This makes sense since $z = r$ says that the height of a point should be the same as its distance to the z -axis, which is what happens along this cone.

We can also see this as follows. First focus on what's happening at $\theta = \frac{\pi}{2}$, so on the yz -plane in the positive y -direction. On this plane, $x = 0$ so the value of r is the same as y . Thus $z = r$ gives the line $z = y$ on the yz -plane. Now, the cylindrical equation $z = r$ places no restriction on θ , so taking this line and swinging it around the z -axis for all values of θ will sweep out the surface we want, which indeed gives a cone.

For negative values of z we get negative values of r , which, using the same interpretation for negative r as we had for polar coordinates, still give circles. Thus negative values of z give a cone opening downward:



This entire surface is called a *double cone*. It should be made clear from context whether we want to allow negative values of r or not.

Squaring the cylindrical equation $z = r$ gives $z^2 = r^2$, which in rectangular coordinates becomes

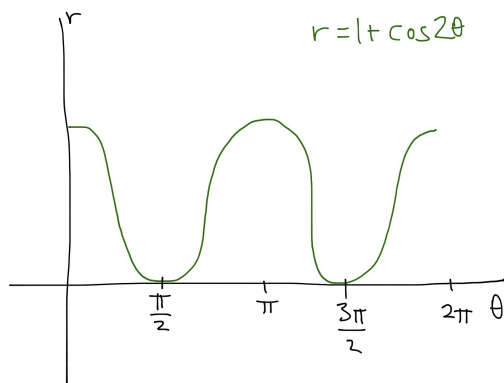
$$z^2 = x^2 + y^2.$$

This is the Cartesian equation of the double cone; taking square roots gives $z = \sqrt{x^2 + y^2}$, which is the top half of the double cone, or $z = -\sqrt{x^2 + y^2}$, which is the bottom half.


Lecture 12: Spherical Coordinates

Today we spoke about spherical coordinates, yet another new type of coordinate system. As with cylindrical coordinates, spherical coordinates give us a simpler way of describing certain surfaces.


Warm-Up 1. We sketch the curve with polar equation $r = 1 + \cos 2\theta$. First, r vs θ looks like




As θ moves from 0 to $\frac{\pi}{2}$ counterclockwise, r decreases from 1 down to 0 so we start at $(1, 0)$ and moves towards the origin:



Now, for θ between $\frac{\pi}{2}$ and π the value of r is positive, going from 0 back to 1 ; this gives a piece of the curve in the second quadrant which looks just like the piece in the first quadrant:



For θ between π and $\frac{3\pi}{2}$, r is positive and decreasing from 1 to 0 giving a piece in the third quadrant, and finally r goes from 0 back to 1 as θ moves from $\frac{3\pi}{2}$ to 2π , giving a piece of the curve in the fourth quadrant:



Using the trig identity $\cos 2\theta = \cos^2 \theta - \sin^2 \theta$, we can write the polar equation of the curve as

$$r = 1 + \cos^2 \theta - \sin^2 \theta.$$

In rectangular coordinates this becomes


$$\sqrt{x^2 + y^2} = 1 + \left(\frac{x}{r}\right)^2 - \left(\frac{y}{r}\right)^2 = 1 + \frac{x^2 - y^2}{x^2 + y^2},$$

with the point being that this does us no good in helping to visualize the curve.

Viewed as a cylindrical equation in \mathbb{R}^3 , $r = 1 + \cos 2\theta$ describes the surface traced out by taking the curve above and sliding it up and down along the z -axis, giving two cylinders which meet along the z -axis.


Warm-Up 2. Now we sketch the intersection of the surface $z = \sin \theta$ with the surface $r = 1$, which is the set of points whose cylindrical coordinates satisfy both equations simultaneously. We've seen before that the second surface is a cylinder, so our curve will lie on this cylinder.

At $\theta = 0$ (positive x -axis direction) we have $z = 0$ (on the xy -plane) and $r = 1$, so we are at $(1, 0, 0)$; at $\theta = \frac{\pi}{2}$ (positive y -axis direction) we have $z = 1$ and $r = 1$, placing us at $(0, 1, 1)$. As θ moves from 0 to $\frac{\pi}{2}$ from the positive x -axis towards the positive y -axis, we get a piece of the curve moving upward and counterclockwise from $(1, 0, 0)$ to $(0, 1, 1)$:

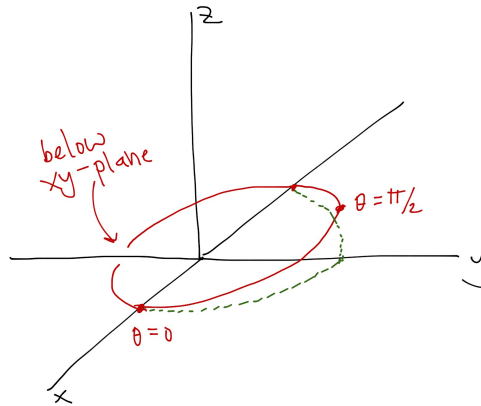


Note that every point along this curve has distance $r = 1$ to the z -axis, so this curve lies directly above the unit circle in the xy -plane, which is the dotted green curve in this picture.

As θ moves from $\frac{\pi}{2}$ to π in the direction of positive y -axis towards the negative x -axis, our curve moves from a height of $z = 1$ back down to a height of $z = 0$, so we get:



which again lies directly above the unit circle in the xy -plane. A similar thing happens for θ moving from π to 2π , only that we have negative values of $z = \sin \theta$, meaning that our curve moves below the xy -plane:

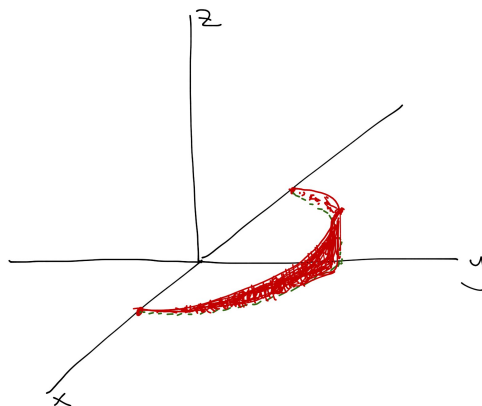


Overall, the intersection curve we're looking at looks like a tilted circle (actually more like a tilted ellipse) with the right half above the xy -plane and the left half below.

Now we consider the points satisfying


$$0 \leq z \leq \sin \theta \text{ and } r = 1,$$

so we no longer require that z be exactly equal to $\sin \theta$, only that it be smaller than that and nonnegative. This rules out anything below the xy -plane, so let's forget the left half of the curve above. Now, we are looking at points (still distance $r = 1$ to the z -axis) which can start on the xy -plane at $z = 0$ and move up to the curve $z = \sin \theta$; this thus includes all points below this curve and above the xy -plane, so we get the following surface shaded in red:



which is on the cylinder $r = 1$. Being able to describe surfaces using such inequalities will be crucial when we talk about integration next quarter.

Spherical coordinates. The *spherical coordinates* of a point (x, y, z) are defined as in the following picture:



So, ρ (pronounced “ro”) is the distance to the origin, ϕ (pronounced “fee” or “fie”) is the angle which you have to move down from the positive z -axis in order to reach the point, and θ is the same θ as in cylindrical coordinates. By considering various right triangles in this picture, we get the following conversions between spherical and rectangular coordinates:


$$\rho^2 = x^2 + y^2 + z^2, \quad x = \rho \sin \phi \cos \theta, \quad y = \rho \sin \phi \sin \theta, \quad z = \rho \cos \phi.$$

In particular, note that r in cylindrical coordinates (distance to z -axis) is $\rho \sin \phi$ in spherical coordinates, which explains the equations for x and y in spherical coordinates: they are simply $x = r \cos \theta$ and $y = r \sin \theta$ with $r = \rho \sin \phi$. A negative value of ρ is interpreted in a similar way to negative values of r in polar or cylindrical coordinates.

As the term “spherical” suggests, spherical coordinates are useful for describing surfaces and regions related to spheres (and cones). In particular, the surface $\rho = 1$ is simply the unit sphere.

Example 1. The set of all points whose spherical coordinates satisfy $\phi = \frac{\pi}{2}$ is the xy -plane. Indeed, we move down $\frac{\pi}{2}$ from the positive z -axis, putting us on the xy -plane, and swing around for all values of θ since $\phi = \frac{\pi}{2}$ places no restriction on θ .


To surface $\phi = \frac{\pi}{4}$ (assuming $\rho \geq 0$, which we will often do) is a cone opening upward:




Indeed, in the yz -plane swinging down $\frac{\pi}{4}$ from the positive z -axis puts us on the line $y = z$, and then we swing this line around for all values of θ . Similarly, $\phi = \frac{3\pi}{4}$ describes a cone opening downward since an angle $\frac{3\pi}{4}$ from the positive z -axis places us below the xy -plane. The equation $\phi = \pi$ describes the negative z -axis.

What about the equation $\phi = \frac{5\pi}{4}$? This would move us beyond the negative z -axis and in fact will give the same cone as $\phi = \frac{3\pi}{4}$. This illustrates something important: any point with an angle $\phi > \pi$ can be described in another way by an angle $\phi < \pi$ by changing the value of θ . For instance, $\rho = 1, \theta = \frac{\pi}{2}, \phi = \frac{3\pi}{2}$ would put us at $(0, -1, 0)$, but so would $\rho = 1, \theta = \frac{3\pi}{2}, \phi = \frac{\pi}{2}$. Thus, we will usually always restrict values of ϕ in spherical coordinates to ones between 0 and π .


All together, the above equations thus describe the following:



Example 2. We determine the surface with spherical equation $\rho = \sin \phi$. Let us first focus on the piece of this surface in the yz -plane. As ϕ moves from 0 in the direction of the positive z -axis to $\frac{\pi}{2}$ in the direction of the positive y -axis, ρ increases from 0 (at the origin) to 1:

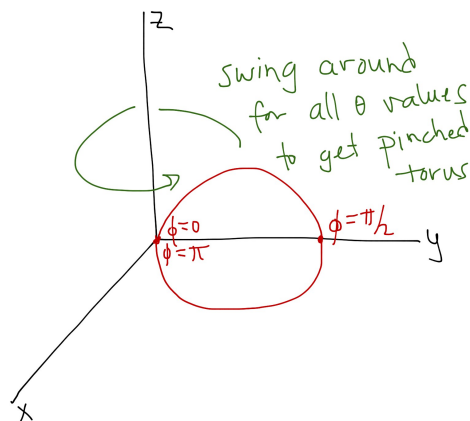


The green lines indicate the increasing radius from 0 to 1 as ϕ swings down. For $0 \leq \phi \leq \pi$ we get the bottom half with ρ decreasing from 1 to 0:



We've drawn this so far as a circle, and we can indeed verify that this makes sense: the spherical equation $\rho = \sin \phi$ becomes $\rho^2 = \rho \sin \phi$ after multiplying through by ρ , which in rectangular coordinates is $x^2 + y^2 + z^2 = \sqrt{x^2 + y^2}$; on the yz -plane $x = 0$ so this becomes $y^2 + z^2 = y$, which is a circle of radius $\frac{1}{2}$ centered at $(0, \frac{1}{2}, 0)$ after completing the square.

Now, to get the entire surface from this, note that $\rho = \sin \phi$ places no restriction on θ , so taking the circle in the yz -plane and swinging it around for all values of θ gives our surface, which is a *pinched torus*:



This will look like the surface of a donut (i.e. a “torus”) only without a complete hole in the middle but rather a “pinching” along the z -axis towards the origin.

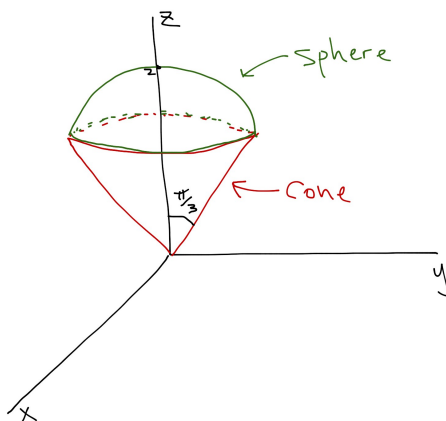
Lecture 13: Multivariable Functions

Today we started talking about functions of several variables, which will be our main object of study the rest of this quarter and next. For now, the main goal is to develop a way of visualizing the graphs of such functions, since this will make it geometrically clear what derivatives in this multivariable setting mean.

Warm-Up. We describe the region above the surface $3z^2 = x^2 + y^2$ and below the surface $x^2 + y^2 + z^2 + 4$ using cylindrical and spherical coordinates. First we should have an idea as to what this region looks like. In cylindrical coordinates, $3z^2 = x^2 + y^2$ becomes $3z^2 = r^2$, or $\sqrt{3}z = r$. This is a cone around the z -axis; the difference between this cone and the cone $z = r$ we saw last time is the angle ϕ which it makes with the positive z -axis. In this case, ϕ (note that r is “opposite” ϕ and z “adjacent”) satisfies

$$\tan \phi = \frac{r}{z} = \frac{\sqrt{3}z}{z} = \sqrt{3},$$

so $\phi = \frac{\pi}{3}$. Thus this cone is wider than the cone $z = r$, which makes an angle $\frac{\pi}{4}$ with the z -axis. Now, $x^2 + y^2 + z^2 = 4$ describes a sphere of radius 2, so the region we want should be above the cone and below this sphere, so it looks like:



(Think of this as an ice cream cone with a scoop of ice cream on top!)

First let’s describe this region in spherical coordinates. To get all points in this region ϕ should move from 0 along the positive z -axis down to $\frac{\pi}{3}$ along the cone, so we need

$$0 \leq \phi \leq \frac{\pi}{3}.$$

This region includes the origin, so we need ρ to start at 0; then moving out in any direction we need ρ to increase all the way until we hit the sphere at $\rho = 2$, giving

$$0 \leq \rho \leq 2.$$

Finally, to get this entire region we need to consider all possible values of θ , so

$$0 \leq \theta \leq 2\pi.$$

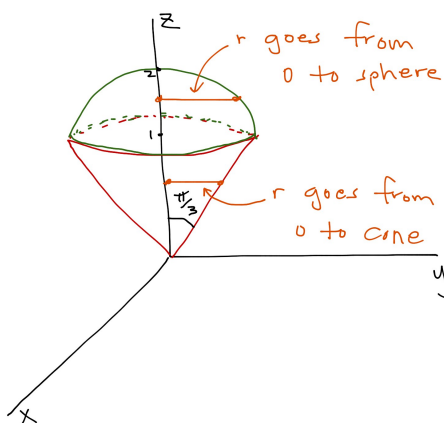
All together then, the given region in spherical coordinates is

$$0 \leq \phi \leq \frac{\pi}{3}, 0 \leq \rho \leq 2, 0 \leq \theta \leq 2\pi.$$

Now we use cylindrical coordinates. Note that the equation of the sphere in cylindrical coordinates is

$$r^2 + z^2 = 4.$$

Since the cylindrical coordinate θ is the same as the one in spherical coordinates, we again need $0 \leq \theta \leq 2\pi$. Now, the smallest which r has among point in this region is 0 for those points on the z -axis, so we need $0 \leq r$. To determine an upper bound on r we must determine how far out horizontally points in this region can move away from the z -axis, but notice that this depends on where in the region we actually are: for certain points (those in the “cone” part of the ice cream cone) it is the cone which is the farthest r can go out to but for other points (those in the “scoop” part) it is the sphere:



Thus the upper bound on r is different for the “cone” piece and the “scoop” piece, so we need different inequalities for each. To be precise, the transition from cone to sphere occurs at $z = 1$, which we find by plugging the equation of the cone $3z^2 = r^2$ into the equation of the sphere $r^2 + z^2 = 4$ and solving for z . Thus for $0 \leq z \leq 1$, the value of r goes out to the cone $r = \sqrt{3}z$, while for $1 \leq z \leq 2$, the value of r goes out to the sphere $r = \sqrt{4 - z^2}$. All together then, the region we’re considering is described in cylindrical coordinates by

$$\begin{cases} 0 \leq \theta \leq 2\pi \\ 0 \leq z \leq 1 \\ 0 \leq r \leq \sqrt{3}z \end{cases} \quad \text{and} \quad \begin{cases} 0 \leq \theta \leq 2\pi \\ 1 \leq z \leq 2 \\ 0 \leq r \leq \sqrt{4 - z^2}. \end{cases}$$

As a consequence, this region is simpler to describe in cylindrical coordinates.

Functions of several variables. We consider functions $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ as we did in linear algebra, only now there is no requirement that these be linear transformations. We will mainly be interested in functions $\mathbb{R}^2 \rightarrow \mathbb{R}$ and $\mathbb{R}^3 \rightarrow \mathbb{R}$, which are of two and three variables respectively. But, we’ve actually already encountered functions of the form


$$\mathbb{R} \rightarrow \mathbb{R}^2, \mathbb{R} \rightarrow \mathbb{R}^3, \mathbb{R}^2 \rightarrow \mathbb{R}^3$$

previously; indeed, such functions are precisely what we use when describing curves and planes using parametric equations.

For instance, consider the function $\gamma : \mathbb{R} \rightarrow \mathbb{R}^3$ defined by

$$\gamma(t) = (\cos t, t, \sin t).$$

The *image* of this function (the collection of all point you get as outputs) is the curve with parametric equations $x = \cos t, y = t, z = \sin t$. To see what this looks like, focus first on the x and z -coordinates, which trace out a unit circle in the xz plane. Then, the value of y increases as this circle is being traced out for increasing t , so as we move around the circle in the x, z -directions at the same time we move to the right along the positive y -axis; this gives a spiral:




better known as a *helix*.

For another example, take the function $\gamma : \mathbb{R} \rightarrow \mathbb{R}^3$ defined by

$$\gamma(t) = (t \cos t, t, t \sin t),$$

whose image is the curve with parametric equations $x = t \cos t, y = t, z = t \sin t$. This curve is similar to the one above, only that as we move in a circular direction in terms of x and z the “radius” should get larger because of the coefficient t in the equations for x and z . Thus we get another spiral-shaped curve, only getting wider as we move along the positive y -axis:



For a function $\mathbb{R}^2 \rightarrow \mathbb{R}^3$, we need something of the form

$$(s, t) \mapsto (x(s, t), y(s, t), z(s, t))$$


where x, y, z coordinates all depend on two input parameters s and t , which are precisely the types of things we saw when looking at parametric equations for planes. So, the upshot is that parametric equations already give us examples of more general functions than just ones from \mathbb{R} to \mathbb{R} .

Level curves. Focusing now on functions $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ of two variables, we wish to be able to visualize what their graphs look like since this is what will geometric meaning to the calculus we will be doing. The graph of f is the surface of points (x, y, z) in \mathbb{R}^3 whose z -coordinate is $f(x, y)$, so it is the surface with equation $z = f(x, y)$. To visualize this, we consider *level curves*, which are the curves obtained when holding $z = k$ fixed at some value. Geometrically, the level curve at $z = k$ is the intersection of the graph of f with the horizontal plane $z = k$ at height k .

For instance, consider the function $f(x, y) = 6x - 3y$ with graph $z = 6x - 3y$. This should be the equation of a plane. Some level curves of this are:

$$\begin{aligned} 6x - 3y &= 0, \text{ or } y = 2x \text{ at } z = 0 \\ 6x - 3y &= 1, \text{ or } y = 2x - \frac{1}{3} \text{ at } z = 1 \\ 6x - 3y &= 2, \text{ or } y = 2x - \frac{2}{3} \text{ at } z = 2 \\ 6x - 3y &= -1, \text{ or } y = 2x + \frac{1}{3} \text{ at } z = -1. \end{aligned}$$

These are all lines, and indeed the general level curve at $z = k$ is the line $k = 6x - 3y$. We draw these level curves on the xy -plane:




To visualize the graph of f , imagine these as lines occurring at the labeled heights: we start on the xy -plane with $6x - 3y = 0$, then as we move up for positive values of z we get the lines drawn with negative y -intercept below, and as we move down for negative z we get lines with positive y -intercepts. Moving the line at $z = 0$ up and down in this manner traces out the graph of f , which is a plane as claimed.


Important. The *level curves* of a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ are the intersections of the graph of f with horizontal planes at various heights. Visualizing these curves as moving up or down depending on the z value give a way of visualizing the graph of f .

More examples. Now we sketch the function $f(x, y) = x^2 + y^2$. The level curve at $z = 0$ is $x^2 + y^2 = 0$, which describes only the origin. For positive $z = k$, the level curves are circles

$x^2 + y^2 = k$ of increasing radii \sqrt{k} as $z = k$ increases. Finally, there are no level curves for negative $z = k$ since $x^2 + y^2$ can never be negative, meaning that the graph of f is never below the xy -plane. These level curves thus look like:




To visualize the graph, imagine starting at the origin and then tracing out circles of increasing radii as you move up:




This surface is a *paraboloid*, which is a 3-dimensional analog of a parabola.

Compare this with the graph of the function $g(x, y) = \sqrt{x^2 + y^2}$, whose level curves for positive $z = k$ are also circles:



The green curves are the level curves of g while the red ones are the level curves of f . The difference is that here, the level curve at $z = k$ is a circle of radius k , whereas in the paraboloid it was a circle of radius \sqrt{k} . This has the consequence that in this case increasing the radius by 1 requires an increase in height of 1 as well, whereas before increasing the radius by 1 caused a larger change in height. So, we get a surface overall that does not slope up as steeply as the paraboloid, and indeed has in a sense “constant slope”; this gives a cone:



The point is that it is the shape of the level curves *together* with the height they occur at that contribute to seeing what the graph looks like.

One final thing to note: the cone as a “sharp” point at the origin, whereas the paraboloid is “smooth” at the origin. This is reflected by the fact, which we will see later, that $f(x, y) = x^2 + y^2$ is *differentiable* at the origin but $g(x, y) = \sqrt{x^2 + y^2}$ is not.


Lecture 14: Quadric Surfaces

Today we continued talking about visualizing surfaces in 3-dimensions, and looked at some examples of *quadric surfaces*. These are not the graphs of functions, but the same ideas we used to visualize graphs works just as well.


Warm-Up 1. We describe the graphs of the functions $f(x, y) = xy$ and $g(x, y) = \frac{x}{y}$ using level curves. The level curves of f are of the form

$$xy = k,$$

which are hyperbolas for nonzero k and the x and y -axis for $k = 0$. For positive $z = k$, we get hyperbolas in the first and third quadrants moving further away from the origin as k increases, while for negative $z = k$ we get hyperbolas in the second and fourth quadrants again moving further away from the origin as k decreases:



So, intersecting the graph of f (which has equation $z = xy$) with the xy -plane gives the two axes, and as we move up or down we get hyperbolas (oriented differently) which are moving further away from the origin; this sweeps out a *hyperbolic paraboloid*, which looks like the surface of a saddle:




Note that indeed, intersecting with horizontal planes above the xy -plane give hyperbolas in one direction, and intersecting it with planes below the xy -plane give hyperbolas in another direction.

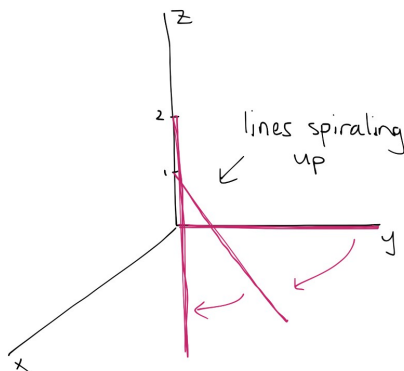
Now, the level curves of g have equations

$$\frac{x}{y} = k.$$

(Note that y can't be zero, so the graph of g never crosses the xz -plane.) These level curves are all lines $x = yk$ of varying slope:



At $z = k = 0$ we have the y -axis, and as $z = k$ increases we get lines of smaller and smaller slope which approach (but never reach since y can't be zero) the x -axis. For negative $z = k$ we also get lines which approach the x -axis, but to make the graph simpler to visualize only the level curves for $z = k \geq 0$ are drawn above. So what does the graph of g with equation $z = \frac{x}{y}$ look like? Starting at $z = 0$ we have the y -axis, and then as z increases this line swings clockwise towards the x -axis, tracing out a kind of “spiraled” surface which tilts more and more steeply as it moves clockwise:




A similar picture works for negative values of z .

Warm-Up 2. Now we describe the *level surfaces* of $h(x, y, z) = x^2 + y^2 + z^2$, which are the analogs of level curves for functions of three variables. Since h is a three variable function, to fully visualize its graph $w = x^2 + y^2 + z^2$ we would need to work in 4-dimensional space, with one axis for each input variable and a fourth for the output variable. Of course, we cannot do this in our 3-dimensional world, but analogously to using level curves to visualize the graph of a two-variable function, level surfaces give a way to “visualize” the graph of h .

At $w = 0$ we get $0 = x^2 + y^2 + z^2$, which only the origin satisfies. There are no level surfaces for negative values of w , and for positive $w = k$ we get spheres:

$$k = x^2 + y^2 + z^2$$

of increasing radii:



Recall that for level curves we visualized the graph by imaging the curves moving up or down as z changed—a similar idea applies here: the graph of h is “traced out” by these spheres of increasing radii. Again, this does not give us a complete picture of the graph of h , but it is enough for our

purposes. In particular, we can tell whether h is increasing or decreasing based on the size of the spheres: h increases in the direction of larger spheres and decreases in the direction of smaller spheres.

Quadric surfaces. Now we consider more general surfaces which are not necessarily the graphs of functions. In particular, we look at *quadric surfaces*, which are the analogs of “conic sections” in 3-dimensions; essentially, a quadric surface is one whose equation involves variables squared. To visualize these, we can look at intersections of the surface with planes $z = k$, which we call *sections* of the surface. (We don’t call these level curves since that term is usually reserved for graphs of functions.)

But the point is that we can do the same with the other variables: we can look at intersections of the surface with *vertical* planes $x = k$ or $y = k$. By looking at all these various sections from different directions, we can piece together what the surface must look like.


The book lists all types of quadric surfaces—ellipsoids, paraboloids, cones, hyperboloids, and hyperbolic paraboloids—at the end of Section 2.1 together with the equations which give rise to them. To get similar surfaces only centered around the x or y -axes (as opposed to the z -axis which is used in the book) we just have to exchange the roles of x, y, z in the equations. I would suggest that it’d be nice to be able to just look at an equation and have a sense for what the corresponding surface looks like, but that memorizing all types of quadric surfaces is probably not necessary; more important is that you can use sections to reconstruct what the surface looks like by hand.

Important. The *sections* of a surface at planes $x = k, y = k$, or $z = k$ are the intersections of the surface with these planes. By imagining how these intersections change as x, y, z increase or decrease we can see how the surface is being traced out.


Example 1. We sketch the quadric surface $x^2 + y^2 - z^2 = 1$. First we consider sections at planes $z = k$, which are given by the curves

$$x^2 + y^2 - k^2 = 1, \text{ or } x^2 + y^2 = 1 + k^2.$$

These curves are all circles since the right side is always positive. The smallest circle of radius 1 occurs on the xy -plane at $k = 0$. Then, as $z = k$ increases or decreases we get circles of larger and larger radii:




This is already enough to determine what the entire surface looks like: take the unit circle on the xy -plane, and as you move up make it larger, and as you move down make it larger. We get a tubular type of surface which is thinnest in the middle, called a *hyperboloid of one sheet*:




For good measure, let's determine the sections at $x = k$ as well. Here we get curves in the yz -plane given by

$$k^2 + y^2 - z^2 = 1, \text{ or } y^2 - z^2 = 1 - k^2.$$

These are (almost all) hyperbolas, changing orientation depending on k . Indeed, for $-1 < k < 1$ the right side is positive so we have hyperbolas crossing the y -axis (since z can be zero), while for $k < -1$ or $k > 1$ the right side is negative, so we get hyperbolas crossing the z -axis (since y can be zero). The only exceptions occur at $k = \pm 1$, where we get $y^2 - z^2 = 0$, or $y = \pm z$:



The point is that you can see these sections on the 3-dimensional hyperboloid itself; for instance, intersecting the hyperboloid with the plane $x = 2$ gives:




which is a hyperbola in the z -direction, just as the section at $x = 2$ in the yz -plane drawing above suggested. You can see visually that intersecting the hyperboloid with vertical planes at $x = 1$ or

$x = 0$ results in something different, which accounts for the difference in the sections we found at these values.

Example 2. Now we sketch the quadric surface $-x^2 + y^2 - z^2 = 1$. The sections at $z = k$ have equations

$$-x^2 + y^2 - k^2 = 1, \text{ or } -x^2 + y^2 = 1 + k^2.$$


The right side is always positive, so these are all hyperbolas crossing the y -axis:




Using these alone it might be kind of tough to see what the surface looks like, so let's look at sections at $y = k$ as well. These have equations

$$-x^2 + k^2 - z^2 = 1, \text{ or } x^2 + z^2 = k^2 - 1.$$

For $k = \pm 1$ we just get the origin, while for $k > 1$ or $k < -1$ we get circles:




Note that for $-1 < k < 1$, $k^2 - 1 < 0$ so there are no points satisfying $x^2 + z^2 = k^2 - 1$, meaning that no piece of our surface lies between $y = -1$ and $y = 1$. Thus, starting with a point at $y = 1$ our surface is traced out by larger and larger circles as we move to the right, and starting with a point at $y = -1$ by larger and larger circles as we move to the left; this gives a *hyperboloid of two sheets*



As expected, there is no piece of this hyperboloid in the region where $-1 < y < 1$.

Note that now that we have this picture, the sections at $z = k$ we found earlier make sense since intersecting this two-sheeted hyperboloid with horizontal planes always gives hyperbolas:



Example 3. Returning to the surface $z = xy$ we had in the first Warm-Up, we have another way of seeing what this looks like. The point is that the right-hand side is actually a quadratic form in terms of x and y , so by diagonalizing this form we can find a simpler equation for this surface with respect to some rotated axes. The quadratic form $q(x, y) = xy$ has matrix

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

After orthogonally diagonalizing and taking coordinates with respect to an orthonormal eigenbasis, this quadratic form becomes $q(x, y) = \frac{1}{2}c_1^2 - \frac{1}{2}c_2^2$, so our surface becomes

$$z = \frac{1}{2}c_1^2 - \frac{1}{2}c_2^2.$$

This equation looks similar to the quadric surface $z = x^2 - y^2$, which as the book shows is a hyperbolic paraboloid (i.e. saddle). So, our surface looks like this (as we saw in the Warm-Up), only with the x and y -axes (but not the z -axis!) rotated by some amount.

Lecture 15: Limits

Today we spoke about limits of multivariable functions. The intuitive idea is similar to what you would have seen before for single variable functions, with the new twist being that you can approach a point from multiple directions.

Warm-Up 1. We describe the surface given by the equation

$$6x^2 + 2y^2 + 3z^2 + 4xz = 1.$$

The left-hand side is a quadratic form with matrix

$$\begin{pmatrix} 6 & 0 & 2 \\ 0 & 2 & 0 \\ 2 & 0 & 3 \end{pmatrix},$$

which has eigenvalues 2, 2, 7. Taking coordinates c_1, c_2, c_3 relative to an orthonormal eigenbasis, the equation for our surface becomes

$$2c_1^2 + 2c_2^2 + 7c_3^2 = 1.$$

This is an ellipsoid with respect to the c_1, c_2, c_3 -axes, so a rotated version of a usual ellipsoid.

Warm-Up 2. Now we describe the surface given by

$$z = -x^2 + 2x - 2y^2 + 8y - 8.$$

First we complete squares on the right-hand side:

$$\begin{aligned} z &= -(x^2 - 2x) - 2(y^2 - 4y) - 8 \\ &= -(x - 1)^2 + 1 - 2(y - 2)^2 + 8 - 8 \\ &= -(x - 1)^2 - 2(y - 2)^2 + 1. \end{aligned}$$

Compare this with the surface $z = -x^2 - 2y^2 + 1$, which is a paraboloid opening downward with highest point at $(0, 0, 1)$. Our surface is a translation of this, so a paraboloid opening downward with highest point at $(1, 2, 1)$.

Limits. Suppose that f is a function of two variables. Intuitively, the *limit of f as (x, y) approaches (a, b)* , denoted $\lim_{(x,y) \rightarrow (a,b)} f(x, y)$, is the value (if any) which the height $z = f(x, y)$ approaches on the graph of f as we consider points (x, y) getting closer to (a, b) . The precise definition is in terms of so-called “ ϵ ’s and δ ’s”, and is left for a course in real analysis. For us the intuitive idea behind limits will be enough. Limits of function with more variables have a similar interpretation.

Example 1. Consider the function $f(x, y) = x^2y + e^{xy}$. This is continuous (intuitively, its graph has no “jumps”) since it is made up by multiplying and adding continuous things. Thus the limit of f as we approach a point is simply the value of f at that point; for instance

$$\lim_{(x,y) \rightarrow (1,2)} (x^2y + e^{xy}) = 1^2 \cdot 2 + e^{1 \cdot 2} = 2 + e^2.$$

This property, that the limit is just the value of the function at the point we’re approaching, is one way of defining what it means for a function to be continuous:

A function f of two variables is *continuous* at (a, b) if $\lim_{(x,y) \rightarrow (a,b)} f(x, y) = f(a, b)$.

So, for continuous functions, limits are easy to compute. An analogous definition holds for functions of any number of variables.

Example 2. Now we consider

$$\lim_{(x,y) \rightarrow (0,0)} \frac{x+y}{2x+y}.$$

This limit does not actually exist, and the key to seeing why lies in the fact that in order for a limit to exist, it must exist and be the same *no matter* how we choose to approach the point we're approaching. In this case, note that approaching $(0,0)$ along the x -axis or y -axis gives different answers:

$$\begin{aligned} \text{along } y = 0 \text{ we have: } & \lim_{(x,0) \rightarrow (0,0)} \frac{x+0}{2x+0} = \lim_{(x,0) \rightarrow (0,0)} \frac{1}{2} = \frac{1}{2} \\ \text{along } x = 0 \text{ we have: } & \lim_{(0,y) \rightarrow (0,0)} \frac{0+y}{0+y} = \lim_{(0,y) \rightarrow (0,0)} 1 = 1. \end{aligned}$$

Since approaching $(0,0)$ in different ways can give different values, the limit in question does not exist.

Important. In order for $\lim_{(x,y) \rightarrow (a,b)} f(x,y)$ to exist, we should obtain the same value no matter which curve we choose to approach (a,b) along. Thus, if two different curves give different values for the limit, the limit does not exist.

Example 3. Now we look at

$$\lim_{(x,y) \rightarrow (0,0)} \frac{xy}{x^2 + y^2}.$$

In this case, approaching along the x and y -axes gives the same value:

$$\begin{aligned} \text{along } y = 0: & \lim_{(x,0) \rightarrow (0,0)} \frac{0}{x^2} = 0 \\ \text{along } x = 0: & \lim_{(0,y) \rightarrow (0,0)} \frac{0}{y^2} = 0. \end{aligned}$$

Of course, this is not enough to say that the limit in question exists, since there are many other possible ways in which we can approach $(0,0)$. In particular, along $y = x$ we have:

$$\lim_{(x,x) \rightarrow (0,0)} \frac{x^2}{x^2 + x^2} = \frac{1}{2}.$$

Thus again, since approaching $(0,0)$ along different curves can give different values for the limit, the overall limit does not exist.

Example 4. Finally let's look at

$$\lim_{(x,y) \rightarrow (0,0)} \frac{x^4 y^4}{(x^2 + y^4)^3}.$$

Approaching the origin along the x or y -axis gives the value 0, and indeed approaching along *any* line $y = mx$ also gives 0;

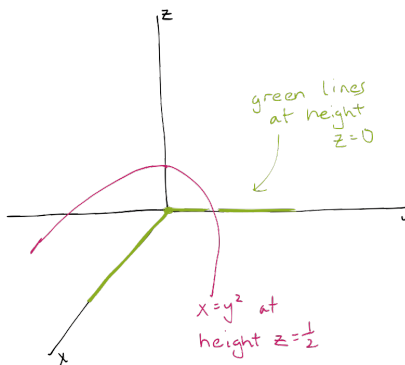
$$\lim_{(x,mx) \rightarrow (0,0)} \frac{mx^8}{(x^2 + m^4 x^4)^3} = \lim_{(x,mx) \rightarrow (0,0)} \frac{mx^2}{(1 + m^4 x^2)^3} = \frac{0}{1} = 0,$$

where in the second step we factored x^6 out of the denominator.

But the point is that we don't necessarily have to approach $(0,0)$ along lines. Approaching along $x = y^2$ (which does pass through the origin) gives:

$$\lim_{(y^2, y) \rightarrow (0,0)} \frac{y^{12}}{(y^4 + y^4)^3} = \lim_{(y^2, y) \rightarrow (0,0)} \frac{y^{12}}{8y^{12}} = \frac{1}{8}.$$

Thus again the limit in question does not exist. If you look at the graph of this function on a computer you can see the limiting behavior changing depending on how you approach the origin; here is a picture giving the rough idea:



Limits in other coordinates. Limits can also be computed by switching to either polar coordinates in the two-variable case or cylindrical/spherical coordinates in the three-variable case. The only thing to keep in mind is that we have to describe the point being approached in new coordinates when describing the resulting limit.

Example 5. Consider

$$\lim_{(x,y) \rightarrow (0,0)} \frac{x+y}{\sqrt{x^2+y^2}}.$$

The form of the denominator suggests that converting to polar coordinates might be useful. In polar coordinates, the given function becomes

$$\frac{x+y}{\sqrt{x^2+y^2}} = \frac{r \cos \theta + r \sin \theta}{r} = \cos \theta + \sin \theta,$$

so our limit becomes

$$\lim_{(x,y) \rightarrow (0,0)} \frac{x+y}{\sqrt{x^2+y^2}} = \lim_{(r,\theta) \rightarrow (0,\theta_0)} (\cos \theta + \sin \theta),$$

where we approach $(0, \theta_0)$ since the origin has $r = 0$ but any possible value of θ . However, now in this form we see that the limit does not exist, since the value depends on which θ -direction we choose to approach the origin in: along $\theta_0 = 0$ gives $\cos 0 + \sin 0 = 1$ but along $\theta_0 = \frac{\pi}{2}$ gives $\cos \frac{\pi}{2} + \sin \pi = 1$.

Instead if we had

$$\lim_{(x,y) \rightarrow (0,0)} \frac{x^2+y^2}{\sqrt{x^2+y^2}},$$

then in polar coordinates this would become

$$\lim_{(r,\theta) \rightarrow (0,\theta_0)} \frac{r^2(\cos \theta + \sin \theta)}{r} = \lim_{(r,\theta) \rightarrow (0,\theta_0)} r(\cos \theta + \sin \theta).$$

In this case, the specific way in which we choose to approach the origin does not matter since the r in front forces the entire expression to go to 0 as $r \rightarrow 0$. Thus the limit in question exists and equals 0.

Example 6. The limit

$$\lim_{(x,y,z) \rightarrow (0,0,0)} \frac{xz}{x^2 + y^2 + z^2}$$

does not exist. In spherical coordinates this becomes

$$\lim_{(\rho,\phi,\theta) \rightarrow (0,\phi_0,\theta_0)} \frac{\rho \sin \phi \cos \theta \rho \cos \phi}{\rho^2},$$

where again ϕ and θ can approach any fixed values ϕ_0 and θ_0 since ρ approaching 0 guarantees that we approach the origin. This simplifies to

$$\lim_{(\rho,\phi,\theta) \rightarrow (0,\phi_0,\theta_0)} \sin \phi \cos \theta \cos \phi,$$

which does not exist since the value depends on exactly what ϕ_0 and θ_0 we use.

Lecture 16: Partial Derivatives

Today we starting talking about partial derivatives of multivariable functions. These are relatively straightforward to compute using the differentiation techniques we all know and love, and have a nice geometric interpretation as well.

Warm-Up 1. We want to know whether the function f defined by

$$f(x, y) = \begin{cases} \frac{x \ln y}{y-x-1} & \text{if } (x, y) \text{ is not on } y = x + 1 \\ 0 & \text{otherwise} \end{cases}$$

is continuous at $(0, 1)$. Recall that to be continuous at a point means that the limit as you approach that point should be the value of the function there, so we want to know whether

$$\lim_{(x,y) \rightarrow (0,1)} f(x, y) = f(0, 1) = 0.$$

However, this limit does not exist. Indeed, approaching $(0, 1)$ along the y -axis gives:

$$\lim_{(0,y) \rightarrow (0,1)} f(0, y) = \lim_{(0,y) \rightarrow (0,1)} \frac{0}{y-1} = 0$$

where we use the fact that point on the y -axis are not on $y = x + 1$ in order to say that $f(0, y) = \frac{0 \ln y}{y-0-1}$ for these points. On the other hand, approaching along the curve $y = e^x$ (which indeed passes through $(0, 1)$) gives:

$$\lim_{(x,e^x) \rightarrow (0,1)} \frac{x \ln e^x}{e^x - x - 1} = \lim_{(x,e^x) \rightarrow (0,1)} \frac{x^2}{e^x - x - 1}.$$

After substituting $y = e^x$ we are left with a single-variable limit, so any technique we know for such limits applies. In particular, since the numerator and denominator here both go to 0 as $x \rightarrow 0$, L'Hopital's rule says that:

$$\lim_{x \rightarrow 0} \frac{x^2}{e^x - x - 1} = \lim_{x \rightarrow 0} \frac{2x}{e^x - 1}.$$

The numerator and denominator here still go to 0 as $x \rightarrow 0$ so applying L'Hopital's rule once more gives:

$$\lim_{x \rightarrow 0} \frac{2x}{e^x - 1} = \lim_{x \rightarrow 0} \frac{2}{e^x} = 2.$$

Since approaching $(0, 1)$ along $y = e^x$ gave a limit value of 2 but approaching along the y -axis gave a value of 0, the limit in question does not exist. Hence f is not continuous at $(0, 1)$.

Warm-Up 2. We determine the value of c which makes

$$f(x, y, z) = \begin{cases} \frac{x^4 - y^4}{x^2 + y^2 + z^2} & \text{if } (x, y, z) \neq (0, 0, 0) \\ c & \text{if } (x, y, z) = (0, 0, 0) \end{cases}$$

continuous at the origin. This value should be

$$c = f(0, 0, 0) = \lim_{(x, y, z) \rightarrow (0, 0, 0)} f(x, y, z).$$

To compute this limit, we switch to spherical coordinates:

$$\begin{aligned} \lim_{(x, y, z) \rightarrow (0, 0, 0)} f(x, y, z) &= \lim_{(\rho, \phi, \theta) \rightarrow (0, \phi_0, \theta_0)} \frac{\rho^4 \sin^4 \phi (\cos^4 \theta - \sin^4 \theta)}{\rho^2} \\ &= \lim_{(\rho, \phi, \theta) \rightarrow (0, \phi_0, \theta_0)} \rho^2 \sin^4 \phi (\cos^4 \theta - \sin^4 \theta) \\ &= 0. \end{aligned}$$

Note that this is zero since the ρ^2 terms causes the entire expression to go to 0 regardless of what's happening to ϕ and θ . Thus $c = 0$ will make f continuous at $(0, 0, 0)$.

Partial derivatives. Given a function of two variables f , the *partial derivative of f with respect to x* is computed by thinking of the variable y as a constant and differentiating with respect to x as normal. Notations for the partial derivative of f with respect to x are:

$$\frac{\partial f}{\partial x} \text{ or } f_x.$$

For instance, if $f(x, y) = x^2y + e^x \sin xy$, then

$$\frac{\partial f}{\partial x} = 2xy + e^x \sin xy + ye^x \cos xy.$$

Similarly, the *partial derivative of f with respect to y* (denoted $\frac{\partial f}{\partial y}$ or f_y) is computed by thinking of x as constant and differentiating with respect to y as normal. For the function $f(x, y) = x^2y + e^x \sin xy$, we have

$$\frac{\partial f}{\partial y} = x^2 + xe^x \cos xy.$$

We can evaluate either of these partial derivatives at a point (a, b) , which we denote using one of


$$\frac{\partial f}{\partial x}(a, b), f_x(a, b), \frac{\partial f}{\partial y}(a, b), f_y(a, b).$$

Geometric interpretation. Recalling the definition of a single-variable derivative as a limit, partial derivatives are concretely defined by:

$$\frac{\partial f}{\partial x}(a, b) = \lim_{x \rightarrow a} \frac{f(x, b) - f(a, b)}{x - a} \text{ and } \frac{\partial f}{\partial y}(a, b) = \lim_{y \rightarrow b} \frac{f(a, y) - f(a, b)}{y - b}.$$

Note that in the first expression we hold y constant at $y = b$ and in the second we hold x constant at $x = a$, which is why the method we described earlier for computing partial derivatives (thinking of one variable as constant) works.

But now, recall that these limit definitions in the single-variable case show that derivatives are certain slopes, and indeed the same is true here: $f_x(a, b)$ is the “slope” (or rate of change) of $z = f(x, y)$ in the x -direction at (a, b) , and $f_y(a, b)$ is the “slope” (or rate of change) of $z = f(x, y)$ in the y -direction. To make this clear, when holding $y = b$ constant and varying the value of x we trace out a curve on the graph of f in the x -direction passing through $(a, b, f(a, b))$, and $f_x(a, b)$ is the slope of this curve; a similar explanation works for $f_y(a, b)$:




Here, the green and blue lines through (a, b) in the xy -plane indicate the direction you get when holding one variable fixed and changing the other, and the green and blue curves above the xy -plane are the corresponding curves traced out on the graph of f .

Example. Let $f(x, y) = x^2 + y^2$. Then $f_x = 2x$ and $f_y = 2y$, so

$$f_x(0, 1) = 0 \text{ and } f_y(0, 1) = 2.$$

This makes sense geometrically: the graph of f is a paraboloid and the point $(0, 1, 1)$ is on the right edge of it; moving a bit either way in the x -direction at this point traces out a parabola (in green below) and moving a bit either way in the y -direction also traces out a piece of a parabola (in blue):



The difference is that $(0, 1, 1)$ is at the minimum of the parabola in the x -direction but the parabola in the y -direction is increasing through $(0, 1, 1)$, so $f_x(0, 1)$ should be zero and $f_y(0, 1)$ should be positive.

Important. Partial derivatives gives slopes or rates of change in the direction of one of the variables. To compute partial derivatives, think of the other variables as constant and differentiate as you normally would.

Tangent planes. With partial derivatives at our disposal we can now compute tangent lines. The *tangent line in the x -direction* to the graph of $f(x, y)$ at (a, b) should only come out in the x -direction, so should have direction vector with y -component equal to 0, say $(x_0, 0, z_0)$. The slope $\frac{z_0}{x_0}$ of this tangent line should be $f_x(a, b)$, so we get that

$$(1, 0, f_x(a, b))$$

is a possible direction vector for the tangent line in the x -direction. Thus this tangent line is

$$\mathbf{r}(t) = (a, b, f(a, b)) + t(1, 0, f_x(a, b)).$$

Similarly, the *tangent line in the y -direction* has

$$(0, 1, f_y(a, b))$$

as a possible direction vector, so the tangent line in the y -direction is

$$\mathbf{r}(t) = (a, b, f(a, b)) + t(0, 1, f_y(a, b)).$$

Now, these two tangent lines should lie on the *tangent plane* to the graph of f at (a, b) , so the cross product of the direction vectors of these lines gives a normal vector to the tangent plane:

$$\mathbf{n} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ 1 & 0 & f_x(a, b) \\ 0 & 1 & f_y(a, b) \end{vmatrix} = (-f_x(a, b), -f_y(a, b), 1).$$

Thus, the tangent plane is given by

$$-f_x(a, b)(x - a) - f_y(a, b)(y - b) + (z - f(a, b)),$$

which after rearranging terms becomes

$$z = f(a, b) + f_x(a, b)(x - a) + f_y(a, b)(y - b).$$

You should view this is analogous to the expression for the tangent line to the graph of a function of one variable.

Back to previous example. The paraboloid $z = x^2 + y^2$ has tangent plane at $(0, 1)$ given by

$$z = 1 + 0(x - 0) + 2(y - 1) = -1 + 2y,$$

using the partial derivatives computed previously. This makes sense: this tangent plane is obtained by taking the line $z = -1 + 2y$ in the yz -plane, which looks like it should be tangent to the right edge of the paraboloid, and sliding it out in the x -direction.

Another example. Say that $g(x, y) = xy$. Then $g_x = y$ and $g_y = x$, so

$$g_x(2, 3) = 3 \text{ and } g_y(2, 3) = 2.$$

Thus the tangent plane to the graph of g at $(2, 3)$ is

$$z = 6 + 3(x - 2) + 2(y - 3) = 3x + 2y - 6.$$

Important. The tangent plane to the graph of a two-variable function f at (a, b) is


$$z = f(a, b) + f_x(a, b)(x - a) + f_y(a, b)(y - b).$$

This tangent plane is supposed to provide the best “linear approximation” to f at (a, b) .

Lecture 17: Differentiability

Today we spoke about what it means for a function of two variables to be *differentiable*, which requires more than simply saying that partial derivatives exist. This is a distinction which does not show up with single-variable functions, but is important in many applications of multivariable calculus.

Warm-Up 1. Consider a function f with some level curves drawn below:



We determine the signs of the partial derivatives of f at P and at Q . First let's look at the point P . In this case, when moving horizontally through P the height z changes from 2, to 1, to 0, so $z = f(x, y)$ is decreasing in the x -direction at P :

$$\frac{\partial f}{\partial x}(P) < 0.$$

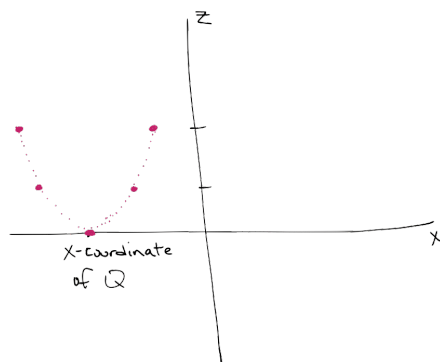
When moving vertically through P , z increases from 0, to 1, to 2, so

$$\frac{\partial f}{\partial y}(P) > 0.$$

Now we move to the point Q . First, note that Q is on the level curve at $z = 0$. As we move only in the y -direction through Q (indicated by the vertical blue line), the height z remains unchanged at $z = 0$, so the rate of change of f in the y -direction at Q is zero:

$$\frac{\partial f}{\partial y}(Q) = 0.$$

Now, imagine moving in the x -direction (indicated by the horizontal green line) through Q . To the left of Q the value of z decreases from 2, to 1, to 0, while to the right of Q the value of z increases from 0, to 1, to 2. So, plotting the change in z with respect to x we get something like:



Note that in the first picture, the horizontal distance to move from $z = 2$ to $z = 1$ to the left of Q looks to be about the same as the horizontal distance to move from $z = 1$ to $z = 2$ to the right of Q , and similarly the distance to move from $z = 1$ to $z = 0$ to the left of Q is the same as the distance to move from $z = 0$ to $z = 1$ to the right of Q , which is why we're getting a parabola-like curve in the second picture. The point Q sits at the bottom of this parabola, so the rate of change in the value of $z = f(x, y)$ in the x -direction at Q is zero! That is, $\frac{\partial f}{\partial x}$ switches from being negative to being positive at Q itself, so the rate of change at Q itself is zero:

$$\frac{\partial f}{\partial x}(Q) = 0.$$

So, both partial derivatives at Q are zero, but for different reasons.

Warm-Up 2. Let $f(x, y) = \sin(x^2y)$. We find an equation for the line through $(\frac{\sqrt{\pi}}{2}, 1, \frac{\sqrt{2}}{2})$ which is perpendicular to the tangent plane to the graph of f at that point. For this, we need a direction vector for the line, which is given by the normal vector to the tangent plane. We have

$$f_x = 2xy \cos(x^2y) \text{ and } f_y = x^2 \cos(x^2y),$$

so

$$f_x\left(\frac{\sqrt{\pi}}{2}, 1\right) = \frac{\sqrt{2\pi}}{2} \text{ and } f_y\left(\frac{\sqrt{\pi}}{2}, 1\right) = \frac{\pi\sqrt{2}}{8}.$$

Hence the tangent plane to the graph of f at the given point is

$$\begin{aligned} z &= f(\sqrt{\pi}/2, 1) + f_x(\sqrt{\pi}/2, 1)(x - \sqrt{\pi}/2) + f_y(\sqrt{\pi}/2, 1)(y - 1) \\ &= \frac{\sqrt{2}}{2} + \frac{\sqrt{2\pi}}{2} \left(x - \frac{\sqrt{\pi}}{2}\right) + \frac{\pi\sqrt{2}}{8}(y - 1). \end{aligned}$$

Writing this in the form $ax + by + cz = d$ by moving z to the right side, we see that

$$\mathbf{n} = \left(\frac{\sqrt{2\pi}}{2}, \frac{\pi\sqrt{2}}{8}, -1\right)$$

is normal to the tangent plane. Thus the line through the given point which is perpendicular to the tangent plane at that point is

$$\mathbf{r}(t) = \left(\frac{\sqrt{\pi}}{2}, 1, \frac{\sqrt{2}}{2}\right) + t \left(\frac{\sqrt{2\pi}}{2}, \frac{\pi\sqrt{2}}{8}, -1\right).$$

Example. Let f be the function defined by

$$f(x, y) = \begin{cases} x & \text{if } |y| < |x| \\ -x & \text{otherwise.} \end{cases}$$

We claim that the partial derivatives of f at the origin both exist. The partial with respect to x is given by:

$$\frac{\partial f}{\partial x}(0, 0) = \lim_{x \rightarrow 0} \frac{f(x, 0) - f(0, 0)}{x - 0} = \lim_{x \rightarrow 0} \frac{x - 0}{x - 0} = 1,$$

where $f(x, 0) = x$ since points $(x, 0)$ along the x -axis are in the region where $|y| < |x|$. The partial with respect to y is:

$$\frac{\partial f}{\partial y}(0, 0) = \lim_{y \rightarrow 0} \frac{f(0, y) - f(0, 0)}{y - 0} = \lim_{y \rightarrow 0} \frac{0 - 0}{y - 0} = 0.$$

Here, points $(0, y)$ on the y -axis are in the region where $f(x, y) = -x$, so $f(0, y) = -0 = 0$.

With these two partial derivatives, we can write down the (candidate) tangent plane at the origin:

$$z = f(0, 0) + f_x(0, 0)(x - 0) + f_y(0, 0)(y - 0) = x.$$

However, it turns out that this is actually *not* the tangent plane to the graph of f at the origin, and indeed no such tangent plane actually exists. The problem is that this f is not *differentiable* at $(0, 0)$.

Differentiability. As a motivation, recall the definition of a single-variable derivative:

$$f'(a) = \lim_{x \rightarrow a} \frac{f(x) - f(a)}{x - a}.$$

This can be rewritten as

$$\lim_{x \rightarrow a} \frac{f(x) - f(a) - f'(a)(x - a)}{x - a} = 0.$$

But $y = f(a) + f'(a)(x - a)$ is the equation of the tangent line to the graph of f at $x = a$, so this limit says that

$$\lim_{x \rightarrow a} \frac{f(x) - (\text{tangent line at } a)}{x - a} = 0,$$

which intuitively says that the numerator approaches 0 faster than the denominator, so the tangent line “approximates” f better and better the closer you get to $x = a$. Thus, the tangent line really is the tangent line. Note that the denominator above is simply the distance between x and the point we’re approaching.

Now we state the same thing in the two-variable setting. Given a function f of two variables whose partial derivatives at (a, b) both exist, we can construct a candidate tangent plane using the equation we’ve seen previously. Then we say that f is *differentiable* at (a, b) if

$$\lim_{(x,y) \rightarrow (a,b)} \frac{f(x, y) - (\text{tangent plane at } (a, b))}{\|(x, y) - (a, b)\|} = 0.$$

As in the single-variable setting, intuitively this says that the candidate tangent plane provides a better and better approximation to f as we approach (a, b) , so that the candidate tangent plane really is the tangent plane.

Back to example. We claimed that the function f from the previous example is not differentiable at the origin. We found the candidate tangent plane to be $z = x$, so in order to be differentiable we would need

$$\lim_{(x,y) \rightarrow (0,0)} \frac{f(x, y) - x}{\sqrt{x^2 + y^2}} = 0.$$

However, approaching the origin along $y = \frac{1}{2}x$ (where $f(x, y) = x$) gives

$$\lim_{(x, \frac{1}{2}x) \rightarrow (0,0)} \frac{x - x}{\sqrt{x^2 + (\frac{1}{2}x)^2}} = 0$$

but approaching along $y = 2x$ (where $f(x, y) = -x$) gives

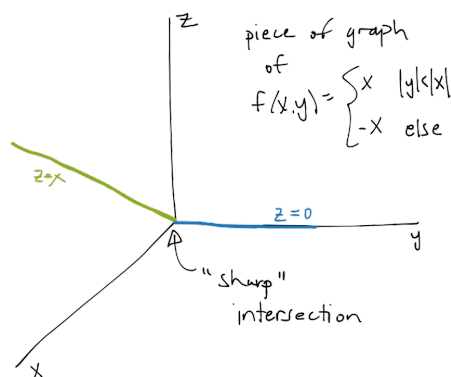
$$\lim_{(x, 2x) \rightarrow (0, 0)} \frac{-x - x}{\sqrt{x^2 + 4x^2}} = \lim_{(x, 2x) \rightarrow (0, 0)} \frac{-2x}{5x} = -\frac{2}{5}.$$

Hence

$$\lim_{(x, y) \rightarrow (0, 0)} \frac{f(x, y) - x}{\sqrt{x^2 + y^2}}$$

does not exist, so it certainly does not equal 0, and thus f is not differentiable at $(0, 0)$. In other words, the candidate tangent plane $z = x$ we found really isn't a tangent plane after all.

We can see this geometrically, which hints at what differentiability really means. The graph of f contains the following two lines:



The green line $z = x$ is the piece of the graph lying above the x -axis and the blue line $z = 0$ is the piece of the graph lying on the y -axis. Note that these two intersect at the origin in a "corner" instead of at a smoothed-out point: this "sharp" point on the graph of f is what prevents f from being differentiable at the origin, and is what prevents there from actually being a tangent plane.

Important. A function $f(x, y)$ is differentiable at (a, b) if

$$\lim_{(x, y) \rightarrow (a, b)} \frac{f(x, y) - [f(a, b) + f_x(a, b)(x - a) + f_y(a, b)(y - b)]}{\sqrt{(x - a)^2 + (y - b)^2}} = 0.$$

Geometrically, this means that the graph of f has a "smooth" instead of a "sharp" point at (a, b) , which means that the candidate tangent plane $z = f(a, b) + f_x(a, b)(x - a) + f_y(a, b)(y - b)$ is a tangent plane after all.

Other examples. The functions

$$f(x, y) = \begin{cases} \frac{-2x^3 + 3y^4}{x^2 + y^2} & \text{if } (x, y) \neq (0, 0) \\ 0 & \text{if } (x, y) = (0, 0) \end{cases} \quad \text{and} \quad g(x, y) = ||x| - |y|| - |x| - |y|$$

are also not differentiable at the origin even though both have partial derivatives at the origin. The graph of the second function is in the book, where you can see the sharpness of the graph at $(0, 0)$, which is geometrically what non-differentiable means.

Our saving grace. Clearly, checking this definition of differentiable every single time we work with a function is tedious, and in particular the limit we need to compute might be quite difficult. However, the following fact is what tells us that for most purposes, we never actually have to worry about it:

If f has *continuous* partial derivatives at (a, b) , then f is differentiable at (a, b) .

So, as long as our function has continuous partial derivatives (which will be the case for the vast majority of functions we consider in this class), we don't have to check differentiability separately.

Final example. Take our standard paraboloid $f(x, y) = x^2 + y^2$. The partial derivatives of f :

$$f_x = 2x \text{ and } f_y = 2y$$

are continuous everywhere, so f is automatically differentiable everywhere. Still, let's work out the limit definition of differentiable anyway, just to see how it would work.

The candidate tangent plane at $(0, 1)$ is $z = -1 + 2y$. We compute:

$$\begin{aligned} \lim_{(x,y) \rightarrow (0,1)} \frac{f(x,y) - [\text{tangent plane at } (0,1)]}{\|(x,y) - (0,1)\|} &= \lim_{(x,y) \rightarrow (0,1)} \frac{x^2 + y^2 - [-1 + 2y]}{\sqrt{x^2 + (y-1)^2}} \\ &= \lim_{(x,y) \rightarrow (0,1)} \frac{x^2 + y^2 - 2y + 1}{\sqrt{x^2 + (y-1)^2}} \\ &= \lim_{(x,y) \rightarrow (0,1)} \frac{x^2 + (y-1)^2}{\sqrt{x^2 + (y-1)^2}} \\ &= \lim_{(x,y) \rightarrow (0,1)} \sqrt{x^2 + (y-1)^2} \\ &= 0. \end{aligned}$$

Thus, as expected since the partial derivatives are continuous at the origin, f is differentiable at the origin.

Lecture 18: Jacobians and Second Derivatives

Today we spoke about derivatives of functions between higher-dimensional spaces (which are expressed in terms so-called Jacobian matrices), and about second derivatives. Both of these topics extend what you already know about single-variable functions to the multivariable setting, but as usual there are some twists which weren't apparent before.

Warm-Up. We show that the function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined by

$$f(x, y) = \begin{cases} \frac{x^3 + y^3}{\sqrt{x^2 + y^2}} & \text{if } (x, y) \neq (0, 0) \\ 0 & \text{if } (x, y) = (0, 0) \end{cases}$$

is differentiable at the origin. For this we first need a candidate for the tangent plane at the origin, so we need the partial derivatives of f at the origin. We have

$$\frac{\partial f}{\partial x}(0, 0) = \lim_{x \rightarrow 0} \frac{f(x, 0) - f(0, 0)}{x - 0} = \lim_{x \rightarrow 0} \frac{\frac{x^3}{\sqrt{x^2}} - 0}{x - 0} = \lim_{x \rightarrow 0} \frac{x^2}{|x|} = 0$$

and

$$\frac{\partial f}{\partial y}(0, 0) = \lim_{y \rightarrow 0} \frac{f(0, y) - f(0, 0)}{y - 0} = \lim_{y \rightarrow 0} \frac{\frac{y^3}{\sqrt{y^2}} - 0}{y - 0} = \lim_{y \rightarrow 0} \frac{y^2}{|y|} = 0.$$

Note that the absolute values came from the fact that $\sqrt{x^2} = |x|$ and similarly for y . (Or, to avoid using absolute values, you can instead look at what happens if you approach 0 from the left or right separately.)

Thus the candidate tangent plane to the graph of f at the origin is

$$z = f(0, 0) + f_x(0, 0)(x - 0) + f_y(0, 0)(y - 0) = 0,$$

or in other words the xy -plane. In order for f to be differentiable at the origin we need

$$\lim_{(x,y) \rightarrow (0,0)} \frac{f(x, y) - [\text{tangent plane}]}{\sqrt{x^2 + y^2}} = 0.$$

In our case we get:

$$\lim_{(x,y) \rightarrow (0,0)} \frac{\frac{x^3+y^3}{\sqrt{x^2+y^2}} - 0}{\sqrt{x^2 + y^2}} = \lim_{(x,y) \rightarrow (0,0)} \frac{x^3 + y^3}{x^2 + y^2}.$$

After converting to polar coordinates, this becomes

$$\lim_{(r,\theta) \rightarrow (0,\theta_0)} \frac{r^3 \cos^3 \theta + r^3 \sin^3 \theta}{r^2} = \lim_{(r,\theta) \rightarrow (0,\theta_0)} r(\cos^3 \theta + \sin^3 \theta) = 0,$$

so f is indeed differentiable at $(0, 0)$.

If you plot the graph of f on a computer you should be able to see that the graph is smooth at the origin instead of “sharp” there, which is geometrically what it means to be differentiable.

Gradients and Jacobians. The equation for a tangent plane (with $\mathbf{a} = (a, b)$ and $\mathbf{x} = (x, y)$) can be written as

$$z = f(\mathbf{a}) + f_x(\mathbf{a})(x - a) + f_y(\mathbf{a})(y - b) = f(\mathbf{a}) + (f_x(\mathbf{a}), f_y(\mathbf{a})) \cdot (x - a, y - b).$$

The vector $(f_x(\mathbf{a}), f_y(\mathbf{a}))$ which shows up here is important enough that we give it its own name: it is called the *gradient* of f at the point $\mathbf{a} = (a, b)$, and we denote it by $\nabla f(a, b)$:

$$\nabla f(a, b) = (f_x(a, b), f_y(a, b)).$$

(We’ll see next week the amazing properties this vector has.) So, with this notation, the equation for a tangent plane is

$$z = f(\mathbf{a}) + \nabla f(\mathbf{a}) \cdot (\mathbf{x} - \mathbf{a}).$$

Compare this to the equation for a tangent line in single-variable calculus:

$$y = f(a) + f'(a)(x - a),$$

and note that the gradient $\nabla f(\mathbf{a})$ in the tangent plane equation plays the same role as the usual derivative $f'(a)$ does in the tangent line equation. The condition for f to be differentiable at \mathbf{a} is thus:

$$\lim_{\mathbf{x} \rightarrow \mathbf{a}} \frac{f(\mathbf{x}) - [f(\mathbf{a}) + \nabla f(\mathbf{a}) \cdot (\mathbf{x} - \mathbf{a})]}{\|\mathbf{x} - \mathbf{a}\|} = 0.$$

Now we see how to generalize all this to functions between spaces of arbitrary dimension, say $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$. The gradient ∇f is replaced by the *Jacobian* of f , denoted Df , which is the *matrix of partial derivatives* of f ; that is, if f looks like

$$f(x_1, \dots, x_m) = (f_1(\mathbf{x}), \dots, f_n(\mathbf{x})),$$

then the Jacobian is the $n \times m$ matrix encoding the partial derivatives of all the component functions f_1, \dots, f_n :

$$Df = \begin{pmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_m} \\ \frac{\partial f_2}{\partial x_1} & \cdots & \frac{\partial f_2}{\partial x_m} \\ \vdots & & \vdots \\ \frac{\partial f_m}{\partial x_1} & \cdots & \frac{\partial f_m}{\partial x_m} \end{pmatrix}.$$

The pattern is that as we move along a row we change the variable we differentiate with respect to, and as we move down a column we change which component function we differentiate. Note that for a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, the Jacobian is the 1×2 matrix $(f_x \ f_y)$, which is precisely the gradient ∇f . Then, the expression

$$f(\mathbf{a}) + \nabla f(\mathbf{a}) \cdot (\mathbf{x} - \mathbf{a})$$

showing up in the equation of a tangent plane when f goes from \mathbb{R}^2 to \mathbb{R} gets replaced by

$$f(\mathbf{a}) + Df(\mathbf{a})(\mathbf{x} - \mathbf{a}),$$

where $Df(\mathbf{a})$ is the Jacobian evaluated at \mathbf{a} and $Df(\mathbf{a})(\mathbf{x} - \mathbf{a})$ is the usual matrix multiplication of the matrix $Df(\mathbf{a})$ with the vector (written as a column) $\mathbf{x} - \mathbf{a}$.

So, $Df(\mathbf{a})$ plays the same role which the usual derivative does in the tangent line equation, and in this sense the Jacobian of f should really be thought of as *the* derivative of f . In the single-variable case, the tangent line provides the best “linear approximation” to f at $x = a$, in the $\mathbb{R}^2 \rightarrow \mathbb{R}$ case the tangent plane provides the best “linear approximation” to f at $(x, y) = (a, b)$, and in general the function

$$g(\mathbf{x}) = f(\mathbf{a}) + Df(\mathbf{a})(\mathbf{x} - \mathbf{a})$$

provides the best linear approximation to f at $\mathbf{x} = \mathbf{a}$, where g is “linear“ in the sense that the formula for g only involves linear terms, as we’ll see in an explicit example in a bit. Finally, to say that $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is *differentiable* at \mathbf{a} means that a certain limit should be zero, namely the limit you get when you take the definition of differentiable for a two-variable function expressed using a gradient and replace the gradient with a Jacobian. Now, this general definition of differentiable is not something we’ll ever actually work with since pretty much all functions we consider will have continuous partial derivatives, in which case the differentiability condition is automatically satisfied.

Important. The *Jacobian* of a function $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is the $n \times m$ matrix Df of partial derivatives of f . In the case of a function $f : \mathbb{R}^m \rightarrow \mathbb{R}$, the Jacobian (which is a row vector) is called the *gradient* of f and is denoted ∇f . Given a point \mathbf{a} , the function

$$g(\mathbf{x}) = f(\mathbf{a}) + Df(\mathbf{a})(\mathbf{x} - \mathbf{a})$$

provides the best linear approximation to f near \mathbf{a} . In the case of a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, this linear approximation g is just the tangent plane to the graph of f at \mathbf{a} .

Example 1. This seems like a lot of new material, but let’s work out an example to see that it all essentially boils down to computing partial derivatives. Let $f : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ be the function

$$f(x, y, z) = (xy^2z + ze^y, z + \sin(xyz)).$$

We want to determine the best “linear approximation” to f at the point $(1, 1, 1)$. First we need the Jacobian of f . The first row of the Jacobian is what you get when you differentiate the first component function of f , which is $f_1(x, y, z) = xy^2z + ze^y$, with respect to x, y, z :

$$\frac{\partial f_1}{\partial x} = y^2z, \quad \frac{\partial f_1}{\partial y} = 2xyz + ze^y, \quad \frac{\partial f_1}{\partial z} = xy^2 + e^y.$$

The second row of the Jacobian is doing the same thing with the second component of f , which is $f_2(x, y, z) = z + \sin(xyz)$, so in the end the Jacobian of f is

$$Df = \begin{pmatrix} y^2z & 2xyz + ze^y & xy^2 + e^y \\ yz \cos(xyz) & xz \cos(xyz) & 1 + xy \cos(xyz) \end{pmatrix}.$$

To get the linear approximation we need we evaluate this Jacobian at the given point $(1, 1, 1)$:

$$Df(1, 1, 1) = \begin{pmatrix} 1 & 2 + e & 1 + e \\ \cos 1 & \cos 1 & 1 + \cos 1 \end{pmatrix}.$$

Then the linear approximation we want is given by the function (where $\mathbf{x} = (x, y, z)$, $\mathbf{a} = (1, 1, 1)$, and we now write vectors as columns):

$$\begin{aligned} g(\mathbf{x}) &= f(\mathbf{a}) + Df(\mathbf{a})(\mathbf{x} - \mathbf{a}) \\ &= \begin{pmatrix} 1 + e \\ 1 + \sin 1 \end{pmatrix} + \begin{pmatrix} 1 & 2 + e & 1 + e \\ \cos 1 & \cos 1 & 1 + \cos 1 \end{pmatrix} \begin{pmatrix} x - 1 \\ y - 1 \\ z - 1 \end{pmatrix} \\ &= \begin{pmatrix} 1 + e \\ 1 + \sin 1 \end{pmatrix} + \begin{pmatrix} (x - 1) + (2 + e)(y - 1) + (1 + e)(z - 1) \\ (\cos 1)(x - 1) + (\cos 1)(y - 1) + (1 + \cos 1)(z - 1) \end{pmatrix} \\ &= \begin{pmatrix} 1 + e + (x - 1) + (2 + e)(y - 1) + (1 + e)(z - 1) \\ 1 + \sin 1 + (\cos 1)(x - 1) + (\cos 1)(y - 1) + (1 + \cos 1)(z - 1) \end{pmatrix}. \end{aligned}$$

As mentioned earlier, this function $g : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ is linear since it only involves x, y, z appearing to the first power, and the point is that it is supposed to be the linear function which best approximates f close to $(1, 1, 1)$ among all possible linear functions. Again, this function g is the analog of a tangent line or a tangent plane now in this higher-dimensional setting.

Second derivatives. Say that we have a two-variable function $f(x, y)$. Then there are two first partial derivatives f_x and f_y . Now, we can differentiate *each* of these also in two ways, either with respect to x or y , and the four resulting expressions are the *second partial derivatives* of f . Notation-wise, f_{xx} denotes the result of differentiating with respect to x and then with respect to x again, f_{xy} the result of differentiating with respect to x and then with respect to y , and so on. Alternate notations, analogous to $\frac{\partial f}{\partial x}$ for first partial derivatives, include:

$$\frac{\partial^2 f}{\partial x^2} \text{ for } f_{xx}, \quad \frac{\partial^2 f}{\partial y \partial x} \text{ for } f_{xy}, \text{ and so on.}$$

Note that in this alternate notation, the order in which you write the variables you are differentiating with respect to is different than in the previous notation; this comes from thinking about a second partial derivative like f_{xy} as “differentiate with respect to y the function f_x ”, which symbolically is:

$$\frac{\partial}{\partial y} \left(\frac{\partial f}{\partial x} \right),$$

and is why we write f_{xy} as $\frac{\partial^2 f}{\partial y \partial x}$. A function of three variables would have nine second partial derivatives, although as we'll see in a second a lot of these second partials turn out to be the same.

Example 2. We compute the second partial derivatives of $f(x, y) = xy^2 + e^{xy}$. We have:

$$\frac{\partial f}{\partial x} = f_x = y^2 + ye^{xy} \quad \text{and} \quad \frac{\partial f}{\partial y} = f_y = 2xy + xe^{xy}.$$

The second partials are thus:

$$\begin{aligned} \frac{\partial^2 f}{\partial x^2} &= \frac{\partial}{\partial x} \left(\frac{\partial f}{\partial x} \right) = f_{xx} = y^2 e^{xy} \\ \frac{\partial^2 f}{\partial y \partial x} &= \frac{\partial}{\partial y} \left(\frac{\partial f}{\partial x} \right) = f_{xy} = 2y + xy e^{xy} \\ \frac{\partial^2 f}{\partial x \partial y} &= \frac{\partial}{\partial x} \left(\frac{\partial f}{\partial y} \right) = f_{yx} = 2y + xy e^{xy} \\ \frac{\partial^2 f}{\partial y^2} &= \frac{\partial}{\partial y} \left(\frac{\partial f}{\partial y} \right) = f_{yy} = 2x + x^2 e^{xy}. \end{aligned}$$

It should immediately jump out at you that the “mixed partials” f_{xy} and f_{yx} are the same in this case; this is no accident:

Clairaut's Theorem. If a multivariable function $f(x_1, x_2, \dots, x_n)$ has continuous first and second-order partial derivatives, then the order in which we differentiate with respect to two variables does not matter: $f_{x_i x_j} = f_{x_j x_i}$.

Example 3. Let $g(x, y, z) = xy^2 z^3 + e^z \cos(xz)$. The first-order partials with respect to x and z are:

$$g_x = y^2 z^3 - ze^z \sin(xz) \quad \text{and} \quad g_z = 3xy^2 z^2 + e^z \cos(xz) - xe^z \sin(xz).$$

Then

$$g_{xz} = 3y^2 z^2 - (e^z + ze^z) \sin(xz) - xze^z \cos(xz)$$

and

$$g_{zx} = 3y^2 z^2 - ze^x \sin(xz) - e^z \sin(xz) - xze^z \cos(xz),$$

so $g_{xz} = g_{zx}$ as predicted by Clairaut's Theorem.

Hessians. Finally, if we should think of the Jacobian of f as *the* derivative of f , what should we think of as *the* second derivative of f ? The answer is given by the *Hessian* of f , which is a square matrix Hf encoding all second-order partial derivatives of f . For a function of two variables $f(x, y)$, the Hessian is

$$Hf = \begin{pmatrix} \frac{\partial^2 f}{\partial x^2} & \frac{\partial^2 f}{\partial y \partial x} \\ \frac{\partial^2 f}{\partial x \partial y} & \frac{\partial^2 f}{\partial y^2} \end{pmatrix}.$$

For instance, the function from Example 2 has Hessian:

$$Hf = \begin{pmatrix} y^2 e^{xy} & 2y + xy e^{xy} \\ 2y + xy e^{xy} & 2x + x^2 e^{xy} \end{pmatrix}.$$

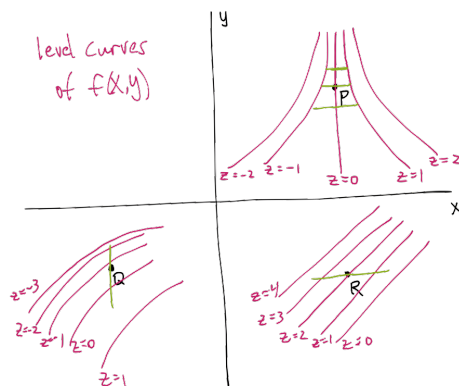
A function of three variables will have a 3×3 Hessian.

As a consequence of Clairaut's Theorem, the Hessian of a function is symmetric(!), so all the great things we learned about symmetric matrices will apply, as yet another example of how linear algebra shows up in multivariable calculus. We'll return to this later.

Lecture 19: Chain Rule

Today we spoke about the chain rule for multivariable derivatives, which applies whenever we have a function depending on variables which themselves depend on additional parameters. There are two ways of expressing this chain rule, and the version in terms of Jacobians is the most versatile.

Warm-Up 1. Suppose we are given level curves of a function f as follows:



We want to determine the signs of some second partial derivatives. First we consider $f_{xx}(R)$, which is

$$\frac{\partial}{\partial x} \left(\frac{\partial f}{\partial x} \right) (R).$$

This gives the rate of change in the x -direction of $\frac{\partial f}{\partial x}$, so in other words the rate of change in the x -direction of the slope in the x -direction. Imagine moving horizontally through the point R . The slope in the x -direction at R is negative since z decreases moving horizontally through R , and the same is true a bit to the left of R and a bit to the right. Now, the equal spacing between the level curves tells us that the negative slope at the point R is the same as the negative slope a bit to the left and the same as the negative slope a bit to the right, so the slope in the x -direction $\frac{\partial f}{\partial x}$ stays constant as we move through R in the x -direction. Thus

$$f_{xx}(R) = \frac{\partial^2}{\partial x^2}(R) = \frac{\partial}{\partial x} \left(\frac{\partial f}{\partial x} \right) (R) = 0$$

since $\frac{\partial f}{\partial x}$ does not change with respect to x . Geometrically, the graph of f in the x -direction at R looks like a straight line, so it has zero concavity.

Next we look at $f_{yy}(Q)$, which is the rate of change in the y -direction of the slope in the y -direction. At Q the slope in the y -direction is negative since z is decreasing vertically through Q , and the slope in the y -direction is also negative a bit below Q as well as a bit above Q . However, the level curves here are not equally spaced: below Q it takes a longer distance to decrease by a height of 1 than it does at Q , so the slope in the y -direction below Q is a little less negative than it is at Q itself. Similarly, above Q the slope in the y -direction is even more negative than it is at Q since it takes a shorter distance to decrease by a height of 1. Thus moving vertically through Q , the slope in the y -direction gets more and more negative, so $\frac{\partial f}{\partial y}$ is decreasing with respect to y at Q , meaning that

$$f_{yy}(Q) = \frac{\partial^2}{\partial y^2}(Q) = \frac{\partial}{\partial y} \left(\frac{\partial f}{\partial y} \right) (Q) < 0.$$

Geometrically, the graph of f at Q in the y -direction is concave down since the downward slope gets steeper and steeper.

Finally we look at $f_{xy}(P)$, which is the rate of change in the y -direction of the slope in the x -direction. At P the slope in the x -direction is positive since z increases when moving horizontally through P . Now, a bit below P the slope in the x -direction is also positive but not as positive as it is at P since it takes a longer distance to increase the height than it does at P . A bit above P it takes an even shorter distance to increase the height in the x -direction, so $\frac{\partial f}{\partial x}$ is larger above P than it is at P . Hence the slope $\frac{\partial f}{\partial x}$ in the x -direction is increasing (getting more and more positive) as you move through P in the y -direction, so

$$f_{xy}(P) = \frac{\partial^2}{\partial y \partial x}(P) = \frac{\partial}{\partial y} \left(\frac{\partial f}{\partial x} \right) (P) > 0.$$

Then by Clairaut's Theorem, $f_{yx}(P)$ is also positive.

Important. For a function f of two variables, f_{xx} measures the concavity of the graph of f in the x -direction (concave up if positive, concave down if negative), f_{yy} measures the concavity of the graph in the y -direction, and $f_{xy} = f_{yx}$ measures the rate of change of the slope in the x -direction as you move in the y -direction, or equivalently the rate of change of the slope in the y -direction as you move in the x -direction.

Warm-Up 2. We use a linear approximation to approximate the (x, y) -coordinates of the point with polar coordinates $r = 2.1$ and $\theta = \frac{\pi}{3} - 0.1$. Consider the function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ which sends the polar coordinates of a point to the corresponding rectangular coordinates:

$$f(r, \theta) = (r \cos \theta, r \sin \theta).$$

The Jacobian of f (with $x = r \cos \theta$ and $y = r \sin \theta$) is

$$Df = \begin{pmatrix} \frac{\partial x}{\partial r} & \frac{\partial x}{\partial \theta} \\ \frac{\partial y}{\partial r} & \frac{\partial y}{\partial \theta} \end{pmatrix} = \begin{pmatrix} \cos \theta & -r \sin \theta \\ \sin \theta & r \cos \theta \end{pmatrix}.$$

Now consider the point with polar coordinates $\mathbf{a} = (2, \frac{\pi}{3})$. The rectangular coordinates of this point are

$$f(\mathbf{a}) = f(2, \frac{\pi}{3}) = (1, \sqrt{3}).$$

The Jacobian of f at \mathbf{a} is

$$Df(\mathbf{a}) = \begin{pmatrix} 1/2 & -\sqrt{3} \\ \sqrt{3}/2 & 1 \end{pmatrix}.$$

The linear approximation to f at \mathbf{a} is given by the function

$$g(\mathbf{r}) = f(\mathbf{a}) + Df(\mathbf{a})(\mathbf{r} - \mathbf{a}) = \begin{pmatrix} 1 \\ \sqrt{3} \end{pmatrix} + \begin{pmatrix} 1/2 & -\sqrt{3} \\ \sqrt{3}/2 & 1 \end{pmatrix} \begin{pmatrix} r - 2 \\ \theta - \frac{\pi}{3} \end{pmatrix},$$

where $\mathbf{r} = \begin{pmatrix} r \\ \theta \end{pmatrix}$ and we write vectors as columns. The approximation we want is the value when $\mathbf{r} = \begin{pmatrix} 2.1 \\ \pi/3 - 0.1 \end{pmatrix}$, which is

$$g \left(\begin{pmatrix} 2.1 \\ \pi/3 - 0.1 \end{pmatrix} \right) = \begin{pmatrix} 1 \\ \sqrt{3} \end{pmatrix} + \begin{pmatrix} 1/2 & -\sqrt{3} \\ \sqrt{3}/2 & 1 \end{pmatrix} \begin{pmatrix} 0.1 \\ -0.1 \end{pmatrix} = \begin{pmatrix} 1 + \frac{1}{20} + \frac{\sqrt{3}}{10} \\ \sqrt{3} + \frac{\sqrt{3}}{20} - \frac{1}{10} \end{pmatrix}.$$

These values are approximately $(\frac{1.22}{1.72})$, which are indeed pretty close to the actual values of the x and y coordinates of the point with polar coordinates $(2.1, \frac{\pi}{3} - 0.1)$.

Example 1. Suppose that $f(x, y) = x^2y$, and that x and y themselves depend on variables s and t as follows:

$$x = st \quad y = e^{st}.$$

Then we can express f in terms of s and t as

$$f(s, t) = (st)^2 e^{st}.$$

Using this we can compute the partial derivatives of f with respect to s and t . However, there is a way to do this without substituting in for x and y , and this is the chain rule.


Chain rule, Version I. For a function $f(x, y)$ of variables x and y , with $x = x(s, t)$ and $y = y(s, t)$ themselves depending on s and t , we have:

$$\begin{aligned} \frac{\partial f}{\partial s} &= \frac{\partial f}{\partial x} \frac{\partial x}{\partial s} + \frac{\partial f}{\partial y} \frac{\partial y}{\partial s} \\ \frac{\partial f}{\partial t} &= \frac{\partial f}{\partial x} \frac{\partial x}{\partial t} + \frac{\partial f}{\partial y} \frac{\partial y}{\partial t}. \end{aligned}$$

To get a feel for the expression for $\frac{\partial f}{\partial s}$ which the chain rule gives us, compare it to the usual single-variable chain rule. In that case we have a function $g(y)$ of a single variable $y = f(x)$ which in turn depends on a single variable x . Then the chain rule says that the derivative of $g(f(x))$ with respect to x is

$$\frac{dg}{dx} = \frac{dg}{dy} \frac{dy}{dx} = g'(y)f'(x) = g'(f(x))f'(x).$$

The multivariable chain rule gives a sum of the same type of terms, with one such term for each variable our original function depends on. We can visualize this dependence via a “tree” diagram:



where each thing depends on the things below it. In this case f depends on x and y and x, y themselves each depend on s and t . To get the expression for $\frac{\partial f}{\partial s}$ we look at all ways of getting from f down to s : doing this through x gives the $\frac{\partial f}{\partial x} \frac{\partial x}{\partial s}$ term in the chain rule and doing this through y gives the $\frac{\partial f}{\partial y} \frac{\partial y}{\partial s}$ term.

Back to Example 1. In this example we thus get:

$$\begin{aligned} \frac{\partial f}{\partial s} &= 2xyt + x^2te^{st} = 2(st)(e^{st})t + (st)^2te^{st} \\ \frac{\partial f}{\partial t} &= 2xys + x^2se^{st} = 2(st)(e^{st})t + (st)^2se^{st}, \end{aligned}$$

which are the same things we would get by differentiating the expression

$$f(s, t) = (st)^2 e^{st}$$

found above.

Example 2. Consider a cylinder made out of wax, melting at a rate of 3 in³/sec. (The units aren't important.) Suppose that at some instant the radius and height of the cylinder both happen to be 1 in and that at this instant the radius is decreasing at a rate of 1 in/sec. How quickly is the height changing at this instant?

Here, the volume V depends on the radius r and height h of the cylinder, each of which depend on time t . Since the volume is $V = \pi r^2 h$, the chain rule gives

$$\frac{\partial V}{\partial t} = \frac{\partial V}{\partial r} \frac{\partial r}{\partial t} + \frac{\partial V}{\partial h} \frac{\partial h}{\partial t} = 2\pi r h \frac{\partial r}{\partial t} + \pi r^2 \frac{\partial h}{\partial t}.$$

At the instant we're interested in, $\frac{\partial V}{\partial t} = -3$ (since the volume is decreasing), $\frac{\partial r}{\partial t} = -1$, and $r = h = 1$, giving

$$-3 = 2\pi(1)(1)(-1) + \pi(1)^2 \frac{\partial h}{\partial t},$$

so

$$\frac{\partial h}{\partial t} = \frac{-3 + 2\pi}{\pi}.$$


Thus the height is changing a rate of $\frac{2\pi-3}{\pi}$ in/sec.

(As pointed out by one of your fellow students, this is not very realistic since it says that the height increases as the cylinder melts, whoops! The problem is that I did not choose realistic values for the various things involved. The math is still sound however!)

Example 3. Suppose that $g(x, y, z) = xy^2z^3$ and that $x = st, y = s + t$, and $z = s^2t$. In this case the chain rule gives three terms making up $\frac{\partial g}{\partial s}$, since g depends on three things:

$$\frac{\partial g}{\partial s} = \frac{\partial g}{\partial x} \frac{\partial x}{\partial s} + \frac{\partial g}{\partial y} \frac{\partial y}{\partial s} + \frac{\partial g}{\partial z} \frac{\partial z}{\partial s}$$

where we used the tree diagram:




Similarly, $\frac{\partial g}{\partial t}$ will be made up of three terms as well. Concretely, we get:

$$\frac{\partial g}{\partial s} = y^2 z^3 t + 2xyz^3 + 3xy^2 z^2 2st,$$

which we can then write solely in terms of s and t by substituting in for x, y, z .

Suppose further that $s = s(u)$ and $t(u)$ themselves depend on an additional parameter u . Then g depends on x, y, z , which depend on s, t , which depend on u , so working out the chain rule using the corresponding tree diagram



gives:

$$\frac{\partial g}{\partial u} = \frac{\partial g}{\partial x} \frac{\partial x}{\partial s} \frac{\partial s}{\partial u} + \frac{\partial g}{\partial x} \frac{\partial x}{\partial t} \frac{\partial t}{\partial u} + \text{similar terms involving } \frac{\partial g}{\partial y} \text{ and } \frac{\partial g}{\partial z},$$

which has six terms in total since there are six ways to get from g down to u in the tree.

Chain rule, Version II. Going back to the type of setup in Example 1, where we have $f(x, y)$ depending on x, y and $x = x(s, t), y = y(s, t)$ depending on s, t , what is really going on is the following. First, f is a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$. We can package the expression for x, y in terms of s, t as a function $g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by

$$g(s, t) = (x(s, t), y(s, t)).$$

Then thinking of f in terms of s and t as $f(x(s, t), y(s, t))$ means we are considering the composition $f \circ g$. The chain rule gives us a way of computing the derivatives of this composition, but there is a better way of thinking about this.

As a motivation, consider again the single variable chain rule. In that case the function $f(g(x))$ is a composition of $g : \mathbb{R} \rightarrow \mathbb{R}$ followed by $f : \mathbb{R} \rightarrow \mathbb{R}$, and the chain rule says that

$$\text{derivative of } f \circ g = (\text{derivative of } f)(\text{derivative of } g).$$

But now, notice that we can make sense of this kind of statement even when f, g are functions between higher-dimensional spaces as long as we interpret “derivative” as “Jacobian”! In this version, the chain rule says that

$$\text{Jacobian of } f \circ g = (\text{Jacobian of } f)(\text{Jacobian of } g)$$

where the multiplication on the right is just matrix multiplication.

Back to Example 1 again. As stated above, the setup in Example 1 is really about the composition $f \circ g$ where $g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $f : \mathbb{R}^2 \rightarrow \mathbb{R}$. The version of the chain rule stated in terms of Jacobians says that

$$D(f \circ g) = Df \cdot Dg.$$

We have

$$Df = \begin{pmatrix} \frac{\partial f}{\partial x} & \frac{\partial f}{\partial y} \end{pmatrix} \text{ and } Dg = \begin{pmatrix} \frac{\partial x}{\partial s} & \frac{\partial x}{\partial t} \\ \frac{\partial y}{\partial s} & \frac{\partial y}{\partial t} \end{pmatrix},$$

so

$$D(f \circ g) = \begin{pmatrix} \frac{\partial f}{\partial x} & \frac{\partial f}{\partial y} \end{pmatrix} \begin{pmatrix} \frac{\partial x}{\partial s} & \frac{\partial x}{\partial t} \\ \frac{\partial y}{\partial s} & \frac{\partial y}{\partial t} \end{pmatrix} = \left(\frac{\partial f}{\partial x} \frac{\partial x}{\partial s} + \frac{\partial f}{\partial y} \frac{\partial y}{\partial s} \quad \frac{\partial f}{\partial x} \frac{\partial x}{\partial t} + \frac{\partial f}{\partial y} \frac{\partial y}{\partial t} \right).$$

On the other hand, with $f \circ g$ written as $f(x(s, t), y(s, t))$ we have

$$D(f \circ g) = \begin{pmatrix} \frac{\partial f}{\partial s} & \frac{\partial f}{\partial t} \end{pmatrix},$$

and so comparing entries in these two expressions for $D(f \circ g)$ gives

$$\begin{aligned}\frac{\partial f}{\partial s} &= \frac{\partial f}{\partial x} \frac{\partial x}{\partial s} + \frac{\partial f}{\partial y} \frac{\partial y}{\partial s} \\ \frac{\partial f}{\partial t} &= \frac{\partial f}{\partial x} \frac{\partial x}{\partial t} + \frac{\partial f}{\partial y} \frac{\partial y}{\partial t}.\end{aligned}$$

These are precisely the expressions given for $\frac{\partial f}{\partial s}$ and $\frac{\partial f}{\partial t}$ in the first version of the chain rule, and the point is that this is just a special case of the more general version of the chain rule stated in terms of Jacobians.

Important. Given functions $g : \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $f : \mathbb{R}^n \rightarrow \mathbb{R}^k$, the Jacobian of $f \circ g : \mathbb{R}^m \rightarrow \mathbb{R}^k$ is the product of the Jacobian of f with the Jacobian of g :

$$D(f \circ g) = Df \cdot Dg.$$

So, the chain rule is nothing but an instance of matrix multiplication, a fact you'll only see in MENU ;).

Lecture 20: Directional Derivatives

Today we spoke about directional derivatives, which give the rate of change of a function (or the slope of its graph) in arbitrary directions. In particular, here we start to see some of the amazing and unexpected properties which gradients have.

Warm-Up 1. Suppose we have a rectangle moving through space on a rocket, with length ℓ and width w changing with respect to time t according to


$$\ell = 1 + 3t \text{ and } w = e^{2t}.$$

But now an added twist: according to Einstein's theory of special relativity, time is not absolute but itself depends on the speed v with which something is traveling. (Don't worry about the physics involved here, the application of the chain rule we'll see is what's important. But, yes, if you've never seen this before it is indeed true that objects moving at different velocities experience time differently, strange as it may seem.) To be precise suppose that time t depends on speed v according to

$$t = \frac{1}{\sqrt{1 - \frac{v^2}{c^2}}}$$

where c is the speed of light.

We determine the rate of change of the area A of the rectangle with respect to speed v . We have the dependence diagram:



Following the ways to get from A to v gives according to the chain rule:

$$\frac{\partial A}{\partial v} = \frac{\partial A}{\partial \ell} \frac{\partial \ell}{\partial t} \frac{\partial t}{\partial v} + \frac{\partial A}{\partial w} \frac{\partial w}{\partial t} \frac{\partial t}{\partial v}.$$

Since area $A = \ell w$, we get

$$\begin{aligned} \frac{\partial A}{\partial v} &= w(3)\left(-\frac{1}{2}\left(1 - \frac{v^2}{c^2}\right)^{-3/2}\right)\left(-\frac{2v}{c^2}\right) + \ell(2e^{2t})\left(-\frac{1}{2}\left(1 - \frac{v^2}{c^2}\right)^{-3/2}\right)\left(-\frac{2v}{c^2}\right) \\ &= \frac{3vw + 2\ell ve^{2t}}{c^2} \frac{1}{\left(1 - \frac{v^2}{c^2}\right)^{3/2}} \end{aligned}$$

If desired, you could now substitute in for ℓ , w , and t in terms of v to express everything in terms of v .

Warm-Up 2. Suppose that $h(x, y, z) = (xyz, x + y)$ and $x = st, y = s + t, z = s^2 - t^2$. Substituting these expressions in for x, y, z in h will give an expression for h in terms of s, t ; we want to compute the Jacobian of this. To be clear, viewing the expressions for x, y, z in terms of s, t as defining a function $f: \mathbb{R}^2 \rightarrow \mathbb{R}^3$:

$$f(s, t) = (st, s + t, s^2 - t^2),$$

we are looking for the Jacobian of $h \circ f$. According to the Jacobian version of the chain rule this is:

$$D(h \circ f) = Dh \cdot Df = \begin{pmatrix} yz & xz & xy \\ 1 & 1 & 0 \end{pmatrix} \begin{pmatrix} t & s \\ 1 & 1 \\ 2s & -2t \end{pmatrix} = \begin{pmatrix} yzt + xz + 2sxy & syz + xz - 2xyt \\ t + 1 & s + t \end{pmatrix}.$$

To find the Jacobian specially at $(s, t) = (3, 3)$, we could then find the values of x, y, z at $s = 3, t = 3$ and substitute these all into this matrix.

But the chain rule already gives us a way to evaluate the Jacobian of $h \circ f$ at a specified point, using:

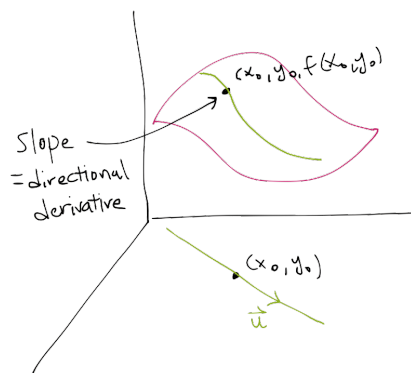
$$D(h \circ f)(\mathbf{a}) = Dh(f(\mathbf{a})) \cdot Df(\mathbf{a}).$$

Note that it is not \mathbf{a} itself we plug into the Jacobian of h but rather $f(\mathbf{a})$, analogous to the fact that the single-variable derivative of $g(f(x))$ is not $g'(x)f'(x)$ but rather $g'(f(x))f'(x)$. In our case, we have

$$\begin{aligned} D(h \circ f)(3, 3) &= Dh(f(3, 3))Df(3, 3) \\ &= Dh(9, 6, 0)Df(3, 3) \\ &= \begin{pmatrix} 0 & 0 & 54 \\ 1 & 1 & 0 \end{pmatrix} \begin{pmatrix} 3 & 3 \\ 1 & 1 \\ 6 & -6 \end{pmatrix} \\ &= \begin{pmatrix} 324 & -324 \\ 4 & 4 \end{pmatrix}. \end{aligned}$$

Directional derivatives. Given a function f of two variables, we know already that the partial derivatives at a point give the rate of the change of f in the x and y -directions, or geometrically the slope of the graph of f in the x and y -directions. But what about the rate of change or slope in other directions? This is what directional derivatives give us.

The setup is as follows. Take a point (x_0, y_0) in the xy -plane and a line going through it with direction given by a vector \mathbf{u} . Plugging points along this line into f gives a curve up on the graph of f passing through $(x_0, y_0, f(x_0, y_0))$, and the *directional derivative* of f at (x_0, y_0) in the direction of the vector \mathbf{u} is the slope of this curve:



The notation we use for this directional derivative is

$$D_{\mathbf{u}}f(x_0, y_0),$$

and we also call it the rate of change of f at (x_0, y_0) in the direction of \mathbf{u} . In fact, partial derivatives are special cases of this: taking $\mathbf{u} = \mathbf{i}$ gives the directional derivative in the direction of \mathbf{i} , which is the same as the derivative in the x -direction

$$D_{\mathbf{i}}f(x_0, y_0) = \frac{\partial f}{\partial x}(x_0, y_0),$$

and similarly the directional derivative in the direction of \mathbf{j} is the rate of change in the y -direction:

$$D_{\mathbf{j}}f(x_0, y_0) = \frac{\partial f}{\partial y}(x_0, y_0).$$

To compute a general directional derivative, we proceed as follows. The line through (x_0, y_0) in the direction of $\mathbf{u} = (a, b)$ has parametric equations

$$x(t) = x_0 + at \text{ and } y(t) = y_0 + bt.$$

The values of f for points along this line are given by $f(x(t), y(t))$, and we want the derivative of this as we move along the line, so the derivative of f with respect to t . The chain rule tells us that this equals:

$$\frac{\partial f}{\partial t} = \frac{\partial f}{\partial x} \frac{\partial x}{\partial t} + \frac{\partial f}{\partial y} \frac{\partial y}{\partial t} = \frac{\partial f}{\partial x} \cdot a + \frac{\partial f}{\partial y} \cdot b.$$

And now the point is that this can be rewritten as the gradient of f dot the direction vector \mathbf{u} :

$$D_{\mathbf{u}}f(x_0, y_0) = \frac{\partial f}{\partial x}(x_0, y_0)a + \frac{\partial f}{\partial y}(x_0, y_0)b = \nabla f(x_0, y_0) \cdot \mathbf{u},$$

which gives us our final formula. Thus to compute a directional derivative at a point we just need to take the dot product of the gradient at that point with the vector given us the direction we're interested in.

One final thing to note: the value of the directional derivative should not depend on the length of the vector we are using to specify the direction, but the formula we derived above does depend on whether we use \mathbf{u} , or $2\mathbf{u}$, or $3\mathbf{u}$, etc. For this reason we always take *unit vectors* when specifying directions.

Important. The *directional derivative* of f at (x_0, y_0) in the direction of the unit vector \mathbf{u} is $D_{\mathbf{u}}f(x_0, y_0) = \nabla f(x_0, y_0) \cdot \mathbf{u}$.

Example. Take the function $f(x, y) = xye^y$. We want the directional derivative of f at $(3, 1)$ in the direction of $(1, 2)$. Geometrically, standing at the point $(3, 1, 3e)$ on the graph of f corresponding to $(3, 1)$ and facing in the direction of the vector $(1, 2)$, this directional derivative gives us the slope of the graph in the direction we're facing.

We first take a unit vector in the direction we want: $\mathbf{u} = (\frac{1}{\sqrt{5}}, \frac{2}{\sqrt{5}})$. The gradient of f in general is

$$\nabla f = (ye^y, xe^y + xye^y),$$

so at $(3, 1)$ it is $\nabla f(3, 1) = (e, 6e)$. The directional derivative we want is thus

$$D_{\mathbf{u}}f(3, 1) = \nabla f(3, 1) \cdot \mathbf{u} = (e, 6e) \cdot (\frac{1}{\sqrt{5}}, \frac{2}{\sqrt{5}}) = \frac{13e}{\sqrt{5}}.$$

This is positive, so the graph of f slopes upward at the point $(3, 1, 3e)$ in the direction of $(1, 2)$.

Geometric interpretation of gradients. We can also express a directional derivative as:

$$D_{\mathbf{u}}f(x_0, y_0) = \nabla f(x_0, y_0) \cdot \mathbf{u} = \|\nabla f(x_0, y_0)\| \|\mathbf{u}\| \cos \theta$$

where θ is the angle between $\nabla f(x_0, y_0)$ and the direction vector \mathbf{u} . Since \mathbf{u} is taken to be a unit vector, this just becomes

$$D_{\mathbf{u}}f(x_0, y_0) = \|\nabla f(x_0, y_0)\| \cos \theta,$$

so that as expected the directional derivative only depends on the function f and the direction we're facing (determined by θ) and not on which specific vector we use in a given direction.

But now note the following: $\|\nabla f(x_0, y_0)\| \cos \theta$ is maximized when $\cos \theta = 1$, so when $\theta = 0$. In other words, the direction in which the directional derivative at a point is as large as possible is precisely the direction given by the gradient at that point! Then in this direction, when $\cos \theta = 1$, we get that the value of the directional derivative itself is $\|\nabla f(x_0, y_0)\|$. Thus the gradient at a point has an important geometric interpretation: it points in the direction where the graph of f has the steepest positive slope (i.e. the directional of maximum rate of change of f , or the maximum directional derivative), and its length *is* the value of the slope in that steepest direction. This is kind of unexpected given the definition of a gradient as a vector made out of partial derivatives, but it key to understanding what gradients mean on a deeper level.

Back to example. Going back to the function $f(x, y) = xye^{xy}$, we want the direction of the maximum rate of change of f at the point $(3, 1)$, which geometrically is the direction in which the graph of f has largest positive slope. As determined above, this is the direction determined by the gradient of f at $(3, 1)$:

$$\nabla f(3, 1) = (e, 6e).$$

(Of course, the same direction is also given by $(1, 6)$, or by any positive multiple of this.) The rate of change of f in this direction (i.e. the maximum rate of change of f at $(3, 1)$) is

$$\|\nabla f(3, 1)\| = \|(e, 6e)\| = \sqrt{37e^2}.$$

Geometrically this is the slope in the direction where the upward slope is the steepest it can be.

Say instead we want to the direction of the minimum rate of change of f , or in other words the direction of steepest *downward* slope. This is obtained when

$$D_{\mathbf{u}}f(3, 1) = \nabla f(3, 1) \cdot \mathbf{u} = \|\nabla f(3, 1)\| \cos \theta$$

is as small as possible, so when $\cos \theta = -1$ and $\theta = \pi$. Thus the direction *opposite* the gradient points in the direction of steepest negative slope:

$$-\nabla f(3, 1) = -(e, 6e) = (-e, -6e).$$

The steepest negative slope at $(3, 1)$ itself is then $-\|\nabla f(3, 1)\| = -\sqrt{37e^2}$.

Finally, there are two directions in which f has no rate of change at all at the point $(3, 1)$; these are the directions which make the directional derivative

$$D_{\mathbf{u}}f(3, 1) = \nabla f(3, 1) \cdot \mathbf{u}$$

zero, so the directions perpendicular to the gradient:

$$(-6e, e) \text{ and } (6e, -e).$$

Geometrically, the graph of f at the point $(3, 1, 3e)$ has zero slope in these directions.

Important. At a point (x_0, y_0) , the gradient $\nabla f(x_0, y_0)$ points in the direction where f is increasing most rapidly and this rate of most rapid increase is $\|\nabla f(x_0, y_0)\|$. The direction in which f decreases most rapidly is $-\nabla f(x_0, y_0)$, and this rate of most rapid decrease is $-\|\nabla f(x_0, y_0)\|$. The rate of change of f at (x_0, y_0) is zero in directions perpendicular to $\nabla f(x_0, y_0)$.

Lecture 21: Gradients

Today we spoke more about properties of gradients, in particular the fact they are always perpendicular to level sets. This gives a nice way of finding tangent planes to arbitrary surfaces, not just ones which arise as the graphs of functions.

Warm-Up 1. Let f be the function $f(x, y) = xy$. We want to determine the direction in which the rate of increase of f at the point $(1, 1)$ is half the value of the largest rate of increase at this same point; in other words, if at the point $(1, 1)$ the largest rate of increase of f is some number M , we want the direction in which the rate of increase at $(1, 1)$ is $\frac{M}{2}$.

Before doing this, let's look at a related question. The gradient of f at a point is $\nabla f = (y, x)$, so the gradient at $(1, 1)$ is

$$\nabla f(1, 1) = (1, 1).$$

Thus if we are standing on the graph of f (which is a hyperbolic paraboloid) at the point $(1, 1, 1)$, then the largest positive slope of the graph occurs in the direction of the vector $(1, 1)$. The value of this largest positive slope itself is

$$\|\nabla f(1, 1)\| = \sqrt{2}.$$

To be sure, if we look at the directional derivative in some other direction, say $(3, 1)$, we should get a value smaller than $\sqrt{2}$. The direction derivative of f in the direction of $(3, 1)$ at $(1, 1)$ is (using $\mathbf{u} = (\frac{3}{10}, \frac{1}{10})$ as a unit vector in the direction we want):

$$D_{\mathbf{u}}f(1, 1) = \nabla f(1, 1) \cdot \mathbf{u} = (1, 1) \cdot \left(\frac{3}{10}, \frac{1}{10}\right) = \frac{4}{10},$$

which is indeed less than $\sqrt{2}$.

Now back to our actual question. We want the direction in which the rate of increase of f at $(1, 1)$ is $\frac{1}{2} \|\nabla f(1, 1)\| = \frac{\sqrt{2}}{2}$. Using

$$D_{\mathbf{u}}f(1, 1) = \nabla f(1, 1) \cdot \mathbf{u} = \|\nabla f(1, 1)\| \cos \theta$$


where θ is the angle between \mathbf{u} and $\nabla f(1,1)$, we are thus looking for directions such that $\cos \theta = \frac{1}{2}$, so direction vectors making an angle $\frac{\pi}{3}$ with $\nabla f(1,1) = (1,1)$. These are obtained by rotating $(1,1)$ by $\frac{\pi}{3}$ either counterclockwise or clockwise, and so (recalling the formulas from last quarter for rotation matrices) are

$$\begin{pmatrix} 1/2 & -\sqrt{3}/2 \\ \sqrt{3}/2 & 1/2 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} \frac{1-\sqrt{3}}{2} \\ \frac{1+\sqrt{3}}{2} \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1/2 & \sqrt{3}/2 \\ -\sqrt{3}/2 & 1/2 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} \frac{1+\sqrt{3}}{2} \\ \frac{1-\sqrt{3}}{2} \end{pmatrix}.$$

Hence, at the point $(1,1)$, the slope of the graph of f is half the value of its largest possible slope at this point in the directions of $([1 - \sqrt{3}]/2, [1 + \sqrt{3}]/2)$ and $([1 + \sqrt{3}]/2, [1 - \sqrt{3}]/2)$.

Note that there is nothing here which is specific to this particular function and this particular point: for *any* function f and at *any* point (x,y) , the directional derivative is half its largest possible value in directions making an angle $\pi/3$ with $\nabla f(x,y)$, just like the directional derivative is 0 in directions perpendicular to $\nabla f(x,y)$ and is as small as possible in directions opposite that of $\nabla f(x,y)$. So, the gradient $\nabla f(x,y)$ really does control everything that's happening.


Warm-Up 2. Say we have a function $f(x,y)$ with level curves looking like:



We want to sketch the gradients of f at the given points P and Q .

First we consider $\nabla f(P)$. This should point in the direction at P where f increases most rapidly. In particular, it should point in a direction where f actually increases which rules out the possibility that it can point in a direction to the left of the drawn level curve, or in the same direction as the level curve itself. So, $\nabla f(P)$ must point to the right side of the given level curve. The precise direction is determined by looking at the shortest way to get from the level curve at $z = 2$ to the one at $z = 3$.

Similarly, $\nabla f(Q)$ should point to the left side of the level curve passing through Q since this is the direction in which f increases. Again, the precise direction is determined by the shortest way to get from the level curve at $z = -1$ to the one at $z = 0$. These two gradients thus look like:



The final thing to make sure is that these gradient vectors have appropriate lengths. Recall that the length of the gradient at a point gives the value of the maximum positive slope at that point, and in this case the maximum slope at Q is larger than that at P since it takes a shorter distance at Q to increase by a height of 1 than it does at P . Thus $\nabla f(Q)$ should be longer than $\nabla f(P)$, which is true in the drawing above.

Gradients are perpendicular to level sets. The gradients drawn in the second Warm-Up seem to have another interesting property: they look to be perpendicular to the level curves themselves. In fact, this is true in general, for the following reason. The level curve of a function $f(x, y)$ at $z = k$ has equation

$$f(x, y) = k.$$

Imagine that we have some parametric equations $\mathbf{r}(t) = (x(t), y(t))$ for this curve, so that these equations satisfy

$$f(x(t), y(t)) = k.$$

Differentiating both sides with respect to t gives

$$\frac{\partial f}{\partial x} \frac{dx}{dt} + \frac{\partial f}{\partial y} \frac{dy}{dt} = 0,$$

where the left-hand side comes from the chain rule. This can be rewritten as

$$\left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right) \cdot (x'(t), y'(t)) = 0, \text{ or } \nabla f(x, y) \cdot \mathbf{r}'(t) = 0$$

where $\mathbf{r}'(t) = (x'(t), y'(t))$. Geometrically $\mathbf{r}'(t)$ describes the *tangent vector* to the curve at a given point, so we find that the gradient $\nabla f(x, y)$ at a point is perpendicular to this tangent vector, meaning that it is perpendicular to the level curve containing (x, y) itself. So, it is no accident that the gradients we drew in the second Warm-Up look like they are perpendicular to the given level sets. Together with the fact that gradients should point in the direction of maximum increase, this gives us a complete way of visualizing what gradient vectors look like in general.

An analogous thing is true in 3-dimensions, where the statement is that the gradient of a three variable function at a point is perpendicular to the level *surface* of that function containing the given point.

Important. For a function f of two variables, $\nabla f(x_0, y_0)$ is perpendicular to the level curve of f containing (x_0, y_0) . For a function g of three variables, $\nabla g(x_0, y_0, z_0)$ is perpendicular to the level surface of g containing (x_0, y_0, z_0) .

Example 1. Consider the curve in the xy -plane with equation

$$xy - y^2 e^{xy} = 2 - e^2.$$

We want to find parametric equations for the tangent line to this curve at the point $(2, 1)$. (Note that this point indeed satisfies the equation of the curve, which is why I needed to use $2 - e^2$ on the right side.) We need two things to describe this tangent line: a point on it, which we have, and a vector giving its direction. To find this direction we proceed as follows.

Let f be the function $f(x, y) = xy - y^2 e^{xy}$. Then the curve in question is precisely the level curve of this function at $z = 2 - e^2$. Hence $\nabla f(2, 1)$ should be a vector perpendicular to the given curve at the point $(2, 1)$. We have

$$\nabla f = (y - y^3 e^{xy}, x - 2ye^{xy} - xy^2 e^{xy}),$$

so $\nabla f(2, 1) = (1 - e^2, 2 - 4e^2)$. The tangent line we want is perpendicular to this vector, so a possible direction vector for the line is

$$\mathbf{v} = (2 - 4e^2, -1 + e^2),$$

or any other vector perpendicular to $(1 - e^2, 2 - 4e^2)$. Hence the tangent line has equation

$$\mathbf{x}_0 + t\mathbf{v} = (2, 1) + t(2 - 4e^2, -1 + e^2),$$

and so has parametric equations $x = 2 + t(2 - 4e^2)$ and $y = 1 + t(e^2 - 1)$.

Example 2. We find an equation for the tangent plane to the unit sphere $x^2 + y^2 + z^2 = 1$ at the point $(\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}})$. Now, we can do this using older material by using the equation

$$z = f(a, b) + f_x(a, b)(x - a) + f_y(a, b)(y - b)$$

for the tangent plane to the graph of a function $f(x, y)$ at (a, b) . However, to do this we need to express our surface as the graph of a function, and there is no way to do this for the entire sphere at once. What we can do is solve for z in the equation of the sphere:

$$z = \pm \sqrt{1 - x^2 - y^2}$$

to obtain equations for the top and bottom half separately. Our points lies on the top half, so we would use

$$f(x, y) = \sqrt{1 - x^2 - y^2}$$

as the function in the equation of the tangent plane given above. This is doable, but notice that things get a little messy when working out

$$\frac{\partial f}{\partial x}\left(\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}\right) \text{ and } \frac{\partial f}{\partial y}\left(\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}\right)$$

since we end up with square roots in the denominator, which are things we usually want to avoid if possible.

But there is another way to approach this which avoids having to identify our surface as the graph of a function of two variables, where instead we think of it as a level surface of a three-variable function! In particular, the unit sphere is the level surface at $g(x, y, z) = 1$ of

$$g(x, y, z) = x^2 + y^2 + z^2.$$

Then $\nabla g(x, y, z) = (2x, 2y, 2z)$ is perpendicular to the level surface at any point, so in particular

$$\nabla g\left(\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}\right) = \left(\frac{2}{\sqrt{3}}, \frac{2}{\sqrt{3}}, \frac{2}{\sqrt{3}}\right)$$

is perpendicular to the unit sphere at $(\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}})$, and so provides a normal vector to tangent plane we're looking for. Using this normal vector and the point we're given, we get

$$\frac{2}{\sqrt{3}}\left(x - \frac{1}{\sqrt{3}}\right) + \frac{2}{\sqrt{3}}\left(y - \frac{1}{\sqrt{3}}\right) + \frac{2}{\sqrt{3}}\left(z - \frac{1}{\sqrt{3}}\right) = 0$$

as the equation of the desired tangent plane. Note that this is less work than the method described previously where we think of the top half of the sphere as the graph of the function $f(x, y) = \sqrt{1 - x^2 - y^2}$.

Example 3. Consider the surface $xyz = 8$. We determine the points at which the tangent plane to this surface is parallel to the plane $x + 2y + 4z = 100$. Two planes are parallel when their normal vectors are parallel, so we're looking for points at which the normal vector to the tangent plane is parallel to $(1, 2, 4)$, and for this we need a description of these normal vectors. One way to do this is to express the given surface as the graph of the function

$$f(x, y) = \frac{8}{xy},$$

which we get after solving for z , and then using the old tangent plane equation. This however will involve taking derivatives of expressions with variables in the denominator, which we want to avoid.

Instead, we can view the given surface as the level surface at 8 of the three variable function

$$g(x, y, z) = xyz.$$

Then the gradient of g at a point is perpendicular to such a level surface, and so gives a normal vector to the tangent plane. We have

$$\nabla g(x, y, z) = (yz, xz, xy),$$

so in the end we're looking for points where this normal vector is parallel to $(1, 2, 4)$; that is, points where

$$(yz, xz, xy) = \lambda(1, 2, 4)$$

for some scalar λ . Comparing components on both sides gives the equations

$$\begin{aligned} yz &= \lambda \\ xz &= 2\lambda \\ xy &= 4\lambda \end{aligned}$$

which points we're looking for have to satisfy, in addition to the equation $xyz = 8$ of the surface we're looking at. Note that none of x, y, z can be zero since then they wouldn't satisfy the equation of the surface. From the first two equations above we get

$$xz = 2\lambda = 2yz, \text{ so } x = 2y,$$

and from the first and third we get

$$xy = 4\lambda = 4yz, \text{ so } x = 4z.$$

(Note that here we're using the fact that all of x, y, z are nonzero, which is why we cancel these terms out of these equations.)

Thus the points we're looking for must have $y = \frac{x}{2}$ and $z = \frac{x}{4}$, so the equation of the surface gives

$$xyz = x \left(\frac{x}{2}\right) \left(\frac{x}{4}\right) = 8.$$

Thus $x^3 = 64$, so $x = 4$ in which case $y = \frac{x}{2} = 2$ and $z = \frac{x}{4} = 1$. Hence $(4, 2, 1)$ is the *only* point on the surface $xyz = 8$ at which the tangent plane is parallel to $x + 2y + 4z = 100$.

Lecture 22: Taylor Polynomials

Today we spoke about Taylor polynomials of multivariable functions. These are incredibly useful in approximation problems, but for us the main point is they will lead to a characterization of local extrema points in optimization problems.

Warm-Up. We find the tangent plane to the surface $6x^2 + 2y^2 + 3z^2 + 4xz = 1$ at $(1/\sqrt{12}, 1/2, 0)$. Taking

$$f(x, y, z) = 6x^2 + 2y^2 + 3z^2 + 4xz,$$

our surface is the level surface of f at 1. Since gradients are perpendicular to level sets, the gradient of f at the given point will be perpendicular to our surface at that point, and hence gives a normal vector for the tangent plane. We have

$$\nabla f = (12x + 4z, 4y, 6z + 4x),$$

so $\nabla f(1/\sqrt{12}, 1/2, 0) = (\sqrt{12}, 2, 4/\sqrt{12})$.

With this as a normal vector and $(1/\sqrt{12}, 1/2, 0)$ as a point on the tangent plane, the tangent plane to $6x^2 + 2y^2 + 3z^2 + 4xz = 1$ at $(1/\sqrt{12}, 1/2, 0)$ is

$$\frac{1}{\sqrt{12}} \left(x - \frac{1}{\sqrt{12}}\right) + 2 \left(y - \frac{1}{2}\right) + \frac{4}{\sqrt{12}}z = 0.$$

Taylor polynomials. The Taylor polynomials of a function are polynomials which provide the best polynomial approximations to that function. The question as to why this is true and in what sense we mean by "best" are somewhat outside the scope of this course, and will be left to a course in *real analysis*. (Math 320 for the win!)

Say that f is a function of two variables. The *first-order Taylor polynomial of f at (a, b)* is:

$$f(a, b) + f_x(a, b)(x - a) + f_y(a, b)(y - b),$$

which is just the polynomial describing the tangent plane at (a, b) . The *second-order Taylor polynomial of f at (a, b)* is:

$$f(a, b) + f_x(a, b)(x - a) + f_y(a, b)(y - b) + \frac{1}{2} [f_{xx}(a, b)(x - a)^2 + 2f_{xy}(a, b)(x - a)(y - b) + f_{yy}(a, b)(y - b)^2],$$

which is the tangent plane together with some additional quadratic terms. Apart from the factor of $\frac{1}{2}$ (the reason why this is there is again left to a later course), the coefficients of the quadratic terms just come from second-derivatives where the variables used to differentiate with respect to correspond to the variables used in forming that quadratic piece. The $(x - a)(y - b)$ term has an extra 2 in front since this actually accounts for two terms:

$$f_{xy}(a, b)(x - a)(y - b) \text{ and } f_{yx}(a, b)(y - b)(x - a),$$

which are of course equal since $f_{xy} = f_{yx}$.

Later we will see a much more compact way of writing this. A similar expression holds for functions of three variables, where we just add on more quadratic terms involving z as well with coefficients given by second-derivatives involving differentiation with respect to z .

Example 1. We find the first and second-order Taylor polynomials of $f(x, y) = 3x - 2y + 1$ at $(1, 1)$. We have

$$f_x = 3 \text{ and } f_y = -2,$$

so the first-order Taylor polynomial is

$$2 + 3(x - 1) - 2(y - 1).$$

Since

$$f_{xx} = 0, f_{xy} = 0, f_{yy} = 0,$$

The second-order Taylor polynomial is the same as the first-order Taylor polynomial.

This makes sense! Since f is already *linear* (it's graph is a plane), it already provides the best linear approximation and quadratic approximations to itself. In fact, note that the expression for the first or second-order Taylor polynomials can be written as

$$2 + 3(x - 1) - 2(y - 1) = 3x - 2y + 1,$$

which is just f ; in other words, the tangent plane to a plane is that plane itself.

Example 2. Now let's work out the first and second-order Taylor polynomials of $f(x, y) = e^{2x} \cos 3y$ at $(0, \pi)$. We have

$$f_x = 2e^{2x} \cos 3y \text{ and } f_y = -3e^{2x} \sin 3y.$$

Hence the first-order Taylor polynomial is

$$-1 - 2(x - 0) + 0(y - \pi) = -1 - 2x.$$

Next we have:

$$f_{xx} = 4e^{2x} \cos 3y, f_{xy} = -6e^{2x} \sin 3y = f_{yx}, f_{yy} = -9e^{2x} \cos 3y.$$

So the second-order Taylor polynomial is

$$-1 - 2x + \frac{1}{2} [-4(x - 0)^2 + 2(0)(x - 0)(y - \pi) + 9(y - \pi)^2] = -1 - 2x - 2x^2 + \frac{9}{2}(y - \pi)^2.$$

Example 3. Finally, we find the first and second-order Taylor polynomials of the three variable function $f(x, y, z) = ye^{3x} + ze^{2y}$ at $(0, 0, 2)$. First:

$$f_x = 3ye^{3x}, f_y = e^{3x} + 2ze^{2y}, f_z = e^{2y},$$

so the first-order Taylor polynomial is

$$2 + 0(x - 0) + 5(y - 0) + 1(z - 2) = 2 + 5y + (z - 2).$$

Next:

$$\begin{aligned} f_{xx} &= 9ye^{3x}, & f_{yy} &= 4ze^{2y}, & f_{zz} &= 0 \\ f_{xy} &= 3e^{3x} = f_{yx}, & f_{xz} &= 0 = f_{zx}, & f_{yz} &= 2e^{2y} = f_{zy}, \end{aligned}$$

so the second-order Taylor polynomial is

$$\begin{aligned} &2 + 5y + (z - 2) \\ &+ \frac{1}{2} \left[0(x - 0)^2 + 8(y - 0)^2 + 0(z - 2)^2 + 2(3)(x - 0)(y - 0) \right. \\ &\quad \left. + 2(0)(x - 0)(z - 2) + 2(2)(y - 0)(z - 2) \right] \\ &= 2 + 5y + (z - 2) + 4y^2 + 3xy + 2y(z - 2). \end{aligned}$$

Note again that the coefficients of the quadratic terms come from second derivatives, so for example the coefficient of $(x - a)(z - c)$ comes from $f_{xz}(a, b, c)$ times an extra 2 since the same term arises from $(z - c)(x - a)$ with coefficient $f_{zx}(a, b, c)$.

An alternate expression. Here's a much better way of expressing the second-order Taylor polynomial of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. First, recall that we can express the first-order Taylor polynomial (i.e. the "linear approximation" of f) at \mathbf{a} as

$$f(\mathbf{a}) + Df(\mathbf{a})(\mathbf{x} - \mathbf{a})$$

where $Df(\mathbf{a})$ is the Jacobian evaluated at \mathbf{a} and vectors are written as columns so that the matrix multiplication $Df(\mathbf{a})(\mathbf{x} - \mathbf{a})$ makes sense. Then the second-order Taylor polynomial is:

$$f(\mathbf{a}) + Df(\mathbf{a})(\mathbf{x} - \mathbf{a}) + \frac{1}{2}(\mathbf{x} - \mathbf{a}) \cdot Hf(\mathbf{a})(\mathbf{x} - \mathbf{a})$$

where $Hf(\mathbf{a})$ is the Hessian of f at \mathbf{a} and \cdot denotes dot product.

The point is that

$$\frac{1}{2}(\mathbf{x} - \mathbf{a}) \cdot Hf(\mathbf{a})(\mathbf{x} - \mathbf{a})$$

is a quadratic form (!!!) which reproduces all the second-order terms in the Taylor polynomial. For instance, in the two variable $f(x, y)$ case with $\mathbf{a} = (a, b)$ and $\mathbf{x} = (x, y)$, this Hessian term turns out to be:

$$\begin{aligned} \frac{1}{2}(\mathbf{x} - \mathbf{a}) \cdot Hf(\mathbf{a})(\mathbf{x} - \mathbf{a}) &= \frac{1}{2} \begin{pmatrix} x - a \\ y - b \end{pmatrix} \cdot \begin{pmatrix} f_{xx}(a, b) & f_{xy}(a, b) \\ f_{yx}(a, b) & f_{yy}(a, b) \end{pmatrix} \begin{pmatrix} x - a \\ y - b \end{pmatrix} \\ &= \frac{1}{2} \begin{pmatrix} x - a \\ y - b \end{pmatrix} \cdot \begin{pmatrix} f_{xx}(a, b)(x - a) + f_{xy}(a, b)(y - b) \\ f_{yx}(a, b)(x - a) + f_{yy}(a, b)(y - b) \end{pmatrix} \\ &= \frac{1}{2} [f_{xx}(a, b)(x - a)^2 + f_{xy}(a, b)(x - a)(y - b) + \text{and so on.}] \end{aligned}$$

Tongue-in-cheek comment: I wonder if all the things we learned about quadratic forms will be somehow applicable to the quadratic form $\frac{1}{2}(\mathbf{x} - \mathbf{a}) \cdot Hf(\mathbf{a})(\mathbf{x} - \mathbf{a})$? The answer is of course, as we'll see next time.

Important. For a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the second-order Taylor polynomial of f at \mathbf{a} is given by

$$f(\mathbf{a}) + Df(\mathbf{a})(\mathbf{x} - \mathbf{a}) + \frac{1}{2}(\mathbf{x} - \mathbf{a}) \cdot Hf(\mathbf{a})(\mathbf{x} - \mathbf{a}).$$

This provides the best *quadratic* approximation to f near \mathbf{a} .

Lecture 23: Local Extrema

Today we spoke about finding and classifying local extrema of multivariable functions, where we see that the Hessian plays a big role in describing such points.

Warm-Up. We want to approximate the value of $\cos(\pi/4 - 0.1) \sin(\pi/3 + 0.1)$. Of course, you can just plug this into a calculator and see that the value is

$$0.7057427798\dots$$

or at least that's how many decimal places my computer gave for this value. But, we can come up with what this value should approximately be by hand using Taylor polynomials. The point is that Taylor polynomials are what calculators and computers use to come up with such values in the first place: your calculator has no idea what "sin" or "cos" mean, all it knows are the Taylor polynomials approximating them which are stored into its memory.

We find the second-order Taylor polynomial of $f(x, y) = \cos x \sin y$ at $(\pi/4, \pi/3)$. We have

$$f_x = -\sin x \sin y, \quad f_y = \cos x \cos y$$

and then

$$f_{xx} = -\cos x \sin y, \quad f_{xy} = -\sin x \cos y = f_{yx}, \quad f_{yy} = -\cos x \sin y.$$

Thus the Jacobian and Hessian at $\mathbf{a} = (\pi/4, \pi/3)$ respectively are:

$$Df(\mathbf{a}) = \begin{pmatrix} -\frac{\sqrt{6}}{4} & \frac{\sqrt{2}}{4} \end{pmatrix} \quad \text{and} \quad Hf(\mathbf{a}) = \begin{pmatrix} -\frac{\sqrt{6}}{4} & -\frac{\sqrt{2}}{4} \\ -\frac{\sqrt{2}}{4} & -\frac{\sqrt{6}}{4} \end{pmatrix}.$$

Hence the second-order Taylor polynomial (in the more compact notation) is:

$$\begin{aligned} f(\mathbf{a}) + Df(\mathbf{a})(\mathbf{x} - \mathbf{a}) + \frac{1}{2}(\mathbf{x} - \mathbf{a}) \cdot Hf(\mathbf{a})(\mathbf{x} - \mathbf{a}) \\ = \frac{\sqrt{6}}{4} + \begin{pmatrix} -\frac{\sqrt{6}}{4} & \frac{\sqrt{2}}{4} \end{pmatrix} \begin{pmatrix} x - \frac{\pi}{4} \\ y - \frac{\pi}{3} \end{pmatrix} + \frac{1}{2} \begin{pmatrix} x - \frac{\pi}{4} \\ y - \frac{\pi}{3} \end{pmatrix} \cdot \begin{pmatrix} -\frac{\sqrt{6}}{4} & -\frac{\sqrt{2}}{4} \\ -\frac{\sqrt{2}}{4} & -\frac{\sqrt{6}}{4} \end{pmatrix} \begin{pmatrix} x - \frac{\pi}{4} \\ y - \frac{\pi}{3} \end{pmatrix}. \end{aligned}$$


This should be a good approximation to f near $(\pi/4, \pi/3)$, so plugging in $x = \pi/4 - 0.1$ and $y = \pi/3 + 0.1$ should give the approximate value of $f(\pi/4 - 0.1, \pi/3 + 0.1) = \cos(\pi/4 - 0.1) \sin(\pi/3 + 0.1)$. We thus have:

$$\begin{aligned} f(\pi/4 - 0.1, \pi/3 + 0.1) &\approx \frac{\sqrt{6}}{4} + \begin{pmatrix} -\frac{\sqrt{6}}{4} & \frac{\sqrt{2}}{4} \end{pmatrix} \begin{pmatrix} -0.1 \\ 0.1 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} -0.1 \\ 0.1 \end{pmatrix} \cdot \begin{pmatrix} -\frac{\sqrt{6}}{4} & -\frac{\sqrt{2}}{4} \\ -\frac{\sqrt{2}}{4} & -\frac{\sqrt{6}}{4} \end{pmatrix} \begin{pmatrix} -0.1 \\ 0.1 \end{pmatrix} \\ &= \frac{\sqrt{6}}{4} + \frac{\sqrt{6} + \sqrt{2}}{40} + \frac{1}{2} \left(\frac{-\sqrt{6} + \sqrt{2}}{400} + \frac{\sqrt{2} - \sqrt{6}}{400} \right). \end{aligned}$$

For comparison, this value is 0.7063768279... which is indeed pretty close to the actual value. The first-order Taylor polynomial (stopping at the Jacobian term) would give 0.7089650184... as the approximate value, so the second-order Taylor polynomial gives a better approximation. In general, higher-order Taylor polynomials (involving higher-order derivatives) give better and better approximations.

Critical points and extrema. The *local extrema* of a function, as the name suggests, are points where the function has some sort of "extreme" behavior. In our case, we are interested in points

where the function has a local maximum, a local minimum, or a *saddle point*. Saddle points have no analogs in single-variable calculus, and are points which are local maximums in one direction but local minimums in another, such as what happens on the surface of a saddle:



Note that in all of these cases, all partial derivatives at such points are zero so the gradient of the function at such points is the zero vector. Points where $\nabla f(P) = \mathbf{0}$ are called *critical points* of f and are the candidate points as to where a local maximum, local minimum, or saddle point can occur.

Important. To find extrema points of a function f , first find the critical points where $\nabla f = \mathbf{0}$, and then determine whether these points give maximums, minimums, or saddle points.

Example 1. We find all local extrema of $f(x, y) = 4x + 6y - 12 - x^2 - y^2$. We have

$$\nabla f = (4 - 2x, 6 - 2y),$$

which is zero only when $x = 2$ and $y = 3$. Hence f has one critical point at $(2, 3)$.

Now, to determine what type of critical point this is note that in this case we can rewrite f as

$$f(x, y) = -(x - 2)^2 - (y - 3)^2 + 1$$

after completing the square. The graph of this is a paraboloid opening downward with topmost point at $(2, 3, 1)$, so $(2, 3)$ gives a maximum of f . However, this only worked because of the specific form of the function f (namely that it is quadratic with no mixed xy terms) and is not a technique which will generalize.

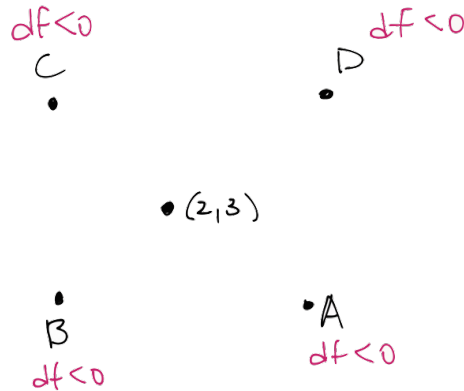
Instead, here's a way to determine the nature of the critical point at $(2, 3)$ which will generalize. We use what's called the *differential* df of f :

$$df = \frac{\partial f}{\partial x} dx + \frac{\partial f}{\partial y} dy$$

which we interpret as giving the "infinitesimal" change df in f given some infinitesimal change dx in x and dy in y . In our case, the differential of f is

$$df = (4 - 2x)dx + (6 - 2y)dy.$$

Consider points A, B, C, D near the critical point $(2, 3)$ as follows:



The change in x at A (as measured from $(2,3)$) is positive since A has larger x -coordinate than $(2,3)$, so $dx > 0$, and since A has x -coordinate larger than 2 the coefficient $4 - 2x$ of dx is negative. The change in y at A is negative, so $dy < 0$ and the coefficient $6 - 2y$ of dy is positive since A has y -coordinate smaller than 3. Thus the change in f at A is

$$df = (-)(+) + (+)(-) < 0,$$

so the value of f at A is smaller than that at $(2,3)$.

At B , $dx < 0$ and $4 - 2x > 0$, and $dy < 0$ with $6 - 2y > 0$, so the change in f at B is

$$df = (+)(-) + (+)(-) < 0,$$

which again means that f has a smaller value at B than at $(2,3)$. Continuing on, at C we have

$$df = (+)(-) + (-)(+) < 0$$

and at D we have

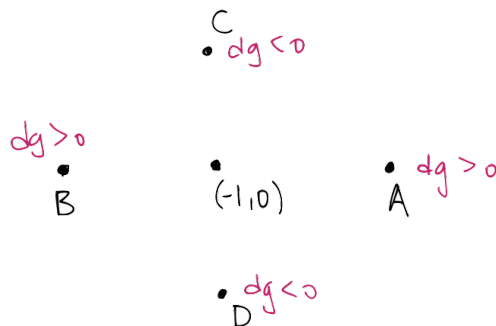
$$df = (-)(+) + (-)(+) < 0.$$

Thus no matter how we move away from $(2,3)$ the value of f will always decrease, so $(2,3)$ is indeed a local minimum of f as claimed.

Example 2. Next we classify the local extrema of $g(x,y) = x^2 - 2y^2 + 2x + 3$. We have

$$\nabla g = (2x + 2, -4y),$$

which is $\mathbf{0}$ only at $(-1,0)$, so this is the only critical point. Consider points A, B, C, D as follows:



At A we have $dy = 0$ since A has the same y -coordinate as $(-1, 0)$, so

$$dg = (2x + 2)dx - 4y dy = (+)(+) + 0 > 0.$$

At B we also have $dy = 0$ but $dx < 0$ so

$$dg = (-)(-) + 0 > 0.$$

Thus the change in g at either A or B as measured from $(-1, 0)$ is positive, so g has larger value at A and B than it does at $(-1, 0)$. This suggests that $(-1, 0)$ is sitting at a minimum in the horizontal direction.

At C and D we have $dx = 0$ since both C and D have the same x -coordinate as does $(-1, 0)$. Then at C we have

$$dg = 0 - (+)(+) < 0$$

and at D we have

$$dg = 0 - (-)(-) < 0,$$

so g has a smaller value at C and D than it does at $(-1, 0)$. Hence $(-1, 0)$ is a maximum in the vertical direction, so overall $(-1, 0)$ is a saddle point of g .

Hessian criteria. The differential method is nice when it works, but most times it may be difficult to determine how the function is changing as we move away from a critical point. For instance, it's hard to determine what dg is in the previous example for a point to the lower-right of $(-1, 0)$, or for a point in any diagonal direction away from $(-1, 0)$. Fortunately we have an even better way of determining the nature of critical points.

Consider the second-order Taylor polynomial of a function f at a point \mathbf{a} :

$$f(\mathbf{a}) + Df(\mathbf{a})(\mathbf{x} - \mathbf{a}) + \frac{1}{2}(\mathbf{x} - \mathbf{a}) \cdot Hf(\mathbf{a})(\mathbf{x} - \mathbf{a}),$$

which we've said before gives a good approximation to the behavior of f near \mathbf{a} . If \mathbf{a} is actually a critical point of f , the Jacobian term is $\mathbf{0}$ so the Taylor polynomial becomes

$$f(\mathbf{a}) + \frac{1}{2}(\mathbf{x} - \mathbf{a}) \cdot Hf(\mathbf{a})(\mathbf{x} - \mathbf{a}).$$

Thus near \mathbf{a} the function f behaves in essentially the same way as the *quadratic form*

$$\frac{1}{2}(\mathbf{x} - \mathbf{a}) \cdot Hf(\mathbf{a})(\mathbf{x} - \mathbf{a}).$$

After picking coordinates c_1, c_2 relative to an orthonormal eigenbasis, this quadratic form becomes

$$\lambda_1(c_1 - a)^2 + \lambda_2(c_2 - b)^2$$

where λ_1, λ_2 are the eigenvalues of the Hessian $Hf(\mathbf{a})$. But the graphs of such functions are easy to describe: we get a paraboloid opening upward when both eigenvalues are positive, a paraboloid opening downward when both eigenvalues are negative, and a hyperbolic paraboloid (saddle) when we have one positive and one negative eigenvalue.

But this is supposed to also be what the graph of f essentially looks like near \mathbf{a} , so we find that the critical point \mathbf{a} of f is: a local minimum when $Hf(\mathbf{a})$ has all positive eigenvalues (i.e. is positive-definite), a local maximum when $Hf(\mathbf{a})$ has all negative eigenvalues (i.e. is negative-definite), and a saddle point when $Hf(\mathbf{a})$ has eigenvalues of opposite signs (i.e. is indefinite). Thus

the eigenvalues of the Hessian determine the nature of a critical point, **as long as zero is not an eigenvalue!** (When zero is an eigenvalue, we have to fall back to the method using differentials or something else.) All this works the same for functions of more than two variables.

Important. Say that \mathbf{a} is a critical point of f . If $Hf(\mathbf{a})$ does not have zero as an eigenvalue (i.e. if $Hf(\mathbf{a})$ is invertible), then \mathbf{a} is a:

- local maximum if $Hf(\mathbf{a})$ is negative definite,
- local minimum if $Hf(\mathbf{a})$ is positive definite,
- saddle point if $Hf(\mathbf{a})$ is indefinite.

More precisely, each positive eigenvalue gives axes (determined by the corresponding orthonormal eigenvectors) along which f increases and each negative eigenvalue gives axes (determined by the corresponding orthonormal eigenvectors) along which f decreases.

Back to examples. The function f from Example 1 has Hessian:

$$Hf = \begin{pmatrix} -2 & 0 \\ 0 & -2 \end{pmatrix},$$

which is negative definite at the critical point $(2, 3)$. Hence this also shows that $(2, 3)$ is a local maximum of f . The function g from Example 2 has Hessian:

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

which is indefinite at the critical point $(-1, 0)$, which is another way of showing that $(-1, 0)$ is a saddle point of g .

Example 3. Finally, we find and classify all local extrema of the three-variable function $f(x, y, z) = xy + xz + 2yz + \frac{1}{x}$. The gradient is

$$\nabla f = \left(y + z - \frac{1}{x^2}, x + 2z, x + 2y \right).$$

Hence the critical points come from solutions of

$$y + z - \frac{1}{x^2} = 0, \quad x + 2z = 0, \quad x + 2y = 0.$$

The second and third equations give

$$y = z = -\frac{x}{2},$$

and then substituting into the first equation gives

$$-\frac{x}{2} - \frac{x}{2} - \frac{1}{x^2} = 0, \quad \text{so } x^3 = 1.$$

Hence there is one critical point with $x = 1$ and $y = z = -\frac{1}{2}$.

The Hessian of f is

$$Hf = \begin{pmatrix} \frac{2}{x^3} & 1 & 1 \\ 1 & 0 & 2 \\ 1 & 2 & 0 \end{pmatrix},$$

and evaluated at the critical point this becomes

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

This has eigenvalues $-2, -2, 2$, so this is indefinite and hence $(1, -\frac{1}{2}, -\frac{1}{2})$ is a saddle point of f . Note that in this case, since we're working with a function of three variables, using the differential f to determine this will likely take a long time since we would have to essentially figure out what's happening at 8 different points away from the critical point. The point is that the Hessian method should always be your first attempt, and most times it will work.

If we went ahead and found orthonormal eigenvectors, the claim more precisely is that f increases in the direction of the axis corresponding to the eigenvalue 2 and decreases in the directions of the two axes corresponding to -2 .

Lecture 24: Absolute Extrema

Today we continued talking about extrema of multivariable functions, focusing on finding the *absolute* maximums and minimums of a function restricted to some specific region.

Warm-Up. We classify the local extrema of $f(x, y) = x^2 - y^3 - x^2y + y$. The gradient of f is

$$\nabla f = (2x - 2xy, -3y^2 - x^2 + 1).$$

The first component is zero when $2x - 2xy = 2x(1 - y) = 0$, so when $x = 0$ or $y = 1$. If $x = 0$, the second component of ∇f becomes $-3y^2 + 1$, which is zero when $y = \pm\frac{1}{\sqrt{3}}$. Thus we get $(0, \pm\frac{1}{\sqrt{3}})$ as critical points. If $y = 1$, the second component of ∇f becomes $-3 - x^2 + 1$, which can never be zero. Hence $y = 1$ gives no additional critical points.

The Hessian of f is

$$Hf = \begin{pmatrix} 2 - 2y & -2x \\ -2x & -6y \end{pmatrix}.$$

At the critical points, we get

$$Hf\left(0, \frac{1}{\sqrt{3}}\right) = \begin{pmatrix} 2 - \frac{2}{\sqrt{3}} & 0 \\ 0 & -\frac{6}{\sqrt{3}} \end{pmatrix} \text{ and } Hf\left(0, -\frac{1}{\sqrt{3}}\right) = \begin{pmatrix} 2 + \frac{2}{\sqrt{3}} & 0 \\ 0 & \frac{6}{\sqrt{3}} \end{pmatrix}.$$

The first Hessian is indefinite, so $(0, 1/\sqrt{3})$ is a saddle point, and the second is positive definite, so $(0, -1/\sqrt{3})$ is a local minimum.

When the Hessian doesn't work. Consider $f(x, y) = x^4 + 2y^4$. The only critical point is $(0, 0)$, and the Hessian at $(0, 0)$ is

$$Hf(0, 0) = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

This is not invertible so here the Hessian tells us nothing about what type of critical point $(0, 0)$ is. In cases like this we have to think of something else. However, it should be clear from the function itself that f is positive for any non-origin point and 0 at the origin, so $(0, 0)$ is a minimum of f .

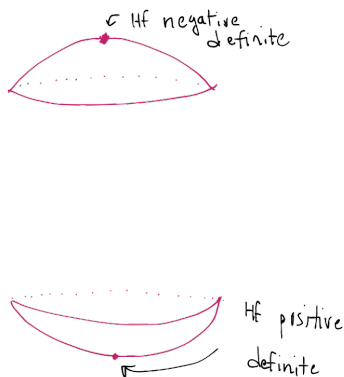
Or, in case determining this just by looking at the function is not easy, we can always use differentials. The differential of f in this case is

$$df = 4x^3 dx + 8y^3 dy,$$


and by consider points close to $(0, 0)$ in any diagonal direction we should also be able to determine that $df > 0$ everywhere away from the origin, so $(0, 0)$ is indeed a minimum.

Morse theory (going off on a tangent here). Let me describe one nice application of some of this stuff, going beyond standard course material, so not something which you'd be expected to know about. *Morse theory* is concerned with the following question. Say we have some unknown surface (or possible higher-dimensional geometric object) X and we have the data of some function $f : X \rightarrow \mathbb{R}$. Then, using only information about the critical points of f , is it possible to determine what X must actually look like? Such questions turn out to arise in many applications, from computer graphics to biological models involve vision, and provide an interesting application of multivariable extrema.


Morse theory tells us that using only information about critical points of some function on X it is indeed possible to determine the basic shape which X must take. For instance, say we have a function $f : X \rightarrow \mathbb{R}$ which we know has exactly two critical points: one where the Hessian is negative definite and the other where it is positive definite. Then near the negative definite point X must look like a downward paraboloid and near the positive definite one it must look like an upward paraboloid:



Now, we cannot have a piece of the surface between these looking like




since this would require more critical points. Thus, the only possibility is for the surface to look like




so in other words X must be an ellipsoid.

As another example, suppose that now we have a function $f : X \rightarrow \mathbb{R}$ which has exactly four critical points: one where the Hessian is negative definite, one where it is positive definite, and two where it is indefinite:



It's a little harder to visualize this, but it turns out the only way this could be possible is if X was actually a torus (!), which looks like the surface of a donut:




More amazingly, the same kind of thing works in higher dimensions, which is where Morse theory really shines.

Absolute extrema. Now we move away from the problem of finding the local extrema of a function f to that of finding its *absolute* or *global* extrema, which are the largest and/or smallest

values a function can have overall. To make matters more interesting, we are interested in finding such values only over a restricted region D , meaning we ask for the absolute max/min values of f among points of D .


We start the same way as before by finding points where f possibly has a local max or min (we don't care about saddle points here), which means finding the critical points of f . After finding these critical points we can simply plug them into f to see which gives the largest value and which gives the smallest. **However**, this method does not account for the fact that the absolute max/min of f might actually occur along the *boundary* of D , since it is possible that a point on the boundary might give the largest or smallest value overall and yet not be a critical point. For instance, for a function and region looking like



we see that the maximum of f over D occurs on the boundary of D and not at the local maximum in the interior of D ; in this case the partial derivatives of f at the boundary point are not zero, so the boundary point is not a critical point of f .

So, after finding critical points of f we still have to check for any possible maximums/minimums on the boundary. Usually this means that we use the equation(s) of the boundary to come up with a simplified version of f along the boundary, and optimize that simplified function instead. The following examples show how this all works.

Example 1. We find the absolute extrema of $f(x, y) = x^2 + xy + y^2 - 6y$ over the rectangle described by $-3 \leq x \leq 3$ and $0 \leq y \leq 5$, which looks like:



(Note: I botched this example in class, and even when I tried to correct it later I still made a mistake. This version is (finally) correct!) First we find critical points. We have

$$\nabla f = (2x + y, x + 2y - 6),$$

so critical points satisfy

$$2x + y = 0 \text{ and } x + 2y - 6 = 0.$$

The first equation gives $y = -2x$ and substituting into the second gives

$$x + 2(-2x) - 6 = 0, \text{ so } -3x - 6 = 0.$$

Thus $x = -2$ and as a result $y = -2x = 4$, so $(-2, 4)$ is the only critical point. Note that at this point we can use the Hessian of f to determine that $(-2, 4)$ is a local minimum, but this is not necessary since at the end we'll just test all points we find anyway to determine which give the absolute max and min.

Now we check the boundary of the rectangle, which consists of four different line segments. The bottom has equation $y = 0$, so the function f along the bottom edge becomes

$$f(x, 0) = x^2.$$

This is now just a function of one variable, which we optimize using techniques from single variable calculus. In this case the only (single-variable) critical point is at $x = 0$, giving $(0, 0)$ as a candidate point for the absolute max and min overall. The right edge has equation $x = 3$, so the function becomes

$$f(3, y) = 9 + 3y + y^2 - 6y = y^2 - 3y + 9.$$

Then $f_y = 2y - 3$ along the right edge, so $(3, \frac{3}{2})$ is a candidate max/min point along the right edge. The top edge is $y = 5$ so f becomes

$$f(x, 5) = x^2 + 5x - 5.$$

Then $f_x = 2x + 5$ along the top, so $(-\frac{5}{2}, 5)$ is another candidate max/min. Finally, the left edge is $x = -3$, so f becomes

$$f(-3, y) = y^2 - 9y + 9,$$

which gives $(-3, \frac{9}{2})$ as another candidate.

To recap, so far we have

$$(-2, 4), (0, 0), (3, \frac{3}{2}), (-\frac{5}{2}, 5), (-3, \frac{9}{2})$$

as possible points where the absolute maximum and minimum occur. But these aren't the only possible points since checking each boundary edge does not take into account what happens at the corners of the rectangle! For instance, along the right edge we had

$$f(3, y) = y^2 - 3y + 9,$$

which has its maximum value along the right edge at the corner $(3, 5)$, and yet this point is not a critical point of the function f restricted to the right edge. In other words, for the same reason why finding critical points of $f(x, y)$ does not necessarily give candidate max/min point along the boundary, finding critical points of f restricted to each boundary piece does not necessarily give the candidate max/min points which occur at the corners of each boundary piece. So, we have to include the four corners

$$(3, 0), (3, 5), (-3, 5), (-3, 0)$$

among the candidate points for an absolute max/min.

In total then we have nine points to test: the one critical point, the four points we found along the boundary pieces, and the four corner points. Plugging all of these into the function gives:

$$\begin{array}{lll} f(-2, 4) = -12 & f(0, 0) = 0 & f(3, 3/2) = 6.75 \\ f(-5/2, 5) = -11.25 & f(-3, 9/2) = -11.25 & f(3, 0) = 9 \\ f(3, 5) = 19 & f(-3, 5) = -11 & f(-3, 0) = 9, \end{array}$$

so the absolute maximum value of f is 19, which is attained at $(3, 5)$, while the absolute minimum value of f is -12 , which is attained at $(-2, 4)$.

Example 2. We find the absolute extrema of the function $f(x, y) = x^2y$ over the region described by $3x^2 + 4y^2 \leq 12$, which is just the region enclosed by the ellipse $3x^2 + 4y^2 = 12$. First,

$$\nabla f = (2xy, x^2),$$

which is $\mathbf{0}$ only when $x = 0$. Thus points on the y -axis are the critical points of f . Now, the points on the boundary satisfy $3x^2 + 4y^2 = 12$, so $x^2 = 4 - \frac{4}{3}y^2$. Hence along the boundary the function f becomes

$$f(y) = \left(4 - \frac{4}{3}y^2\right)y = -\frac{4}{3}y^3 + 4y.$$

This has derivative $-4y^2 + 4$, so only $y = \pm 1$ gives critical points. Then $x^2 = 4 - \frac{4}{3}y^2 = 4 - \frac{4}{3}$, so $x = \pm\sqrt{\frac{8}{3}}$. Hence the candidate max/min points along the boundary ellipse are

$$\left(\sqrt{\frac{8}{3}}, 1\right), \left(-\sqrt{\frac{8}{3}}, 1\right), \left(-\sqrt{\frac{8}{3}}, -1\right), \left(\sqrt{\frac{8}{3}}, -1\right).$$

Note that there are no corner points to test in this case.

Plugging in these points together with the critical points on the y -axis, we find that the absolute maximum value of f is $\frac{8}{3}$, which is attained at $(\sqrt{8/3}, 1)$ and $(-\sqrt{8/3}, 1)$, and the absolute minimum value is $-\frac{8}{3}$, which is attained at $(-\sqrt{8/3}, -1)$ and $(\sqrt{8/3}, -1)$.

Important. To find the absolute extrema of a function f restricted to a region D :

- Find any critical points of f which lie in D ,
- Find any candidate extrema points along the boundary of D , which usually means to use the equations describing the boundary to replace f by a function of one-variable along that boundary, or to use Lagrange multipliers (next lecture!) on the boundary,
- Plug all points you found including any corner points into f to determine which give the largest value and which give the smallest value.

Lecture 25: Lagrange Multipliers

Today we starting talking about the method of Lagrange multipliers, which gives a nice way of optimizing a function subject to some constraint. Such optimization problems are ubiquitous in applications, and indeed the method of Lagrange multipliers shows up all over the place in other subject areas.

Warm-Up. We find the absolute extrema of $f(x, y) = xy$ over the region $x^2 + y^2 \leq 1$, which is the region enclosed by the unit circle. First, $\nabla f = (y, x)$ is $\mathbf{0}$ only when $x = y = 0$, so $(0, 0)$ is the only critical point of f .

Now we check for possible maximums and minimums along the boundary circle. There are a few ways to do this. We can use the equation of the circle to replace y in the function by $y = \pm\sqrt{1 - x^2}$, and then find the maximums and minimums of the resulting single-variable function. This is not impossible, but is a lot of work since we have to test the top and bottom halves of the boundary separately (corresponding to taking the positive or negative square root in the expression for y), and because we end up having to take derivatives of square roots, which is not so nice. Or, we can argue that xy is at a maximum precisely when $(xy)^2$ is, so the maximum of f occurs at the same point as the maximum of $g(x, y) = x^2y^2$, and then do something similar for the minimum. This is better since now substituting $y^2 = 1 - x^2$ in for y^2 in $g(x, y) = x^2y^2$ will avoid having to use square roots, but this approach will still require some work.

Instead, we can check the boundary more easily by converting to polar coordinates, as was suggested by one of your fellow classmates. The function f in polar coordinates is

$$f(r, \theta) = r^2 \cos \theta \sin \theta,$$

so on the boundary $r = 1$ this becomes

$$f(1, \theta) = \cos \theta \sin \theta.$$

Thus on the boundary we have

$$f_\theta = -\cos^2 \theta + \sin^2 \theta,$$

which is zero when $\cos \theta = \pm \sin \theta$. Hence we get candidate max/min points for $\theta = \frac{\pi}{4}, \frac{3\pi}{4}, \frac{5\pi}{4}, \frac{7\pi}{4}$, which give the points

$$\left(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}\right), \left(-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}\right), \left(-\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}}\right), \left(\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}}\right).$$


Finally, testing the critical point $(0, 0)$ and the four points above shows that the absolute maximum of f is $\frac{1}{2}$, which occurs at $(1/\sqrt{2}, 1/\sqrt{2})$ and $(-1/\sqrt{2}, -1/\sqrt{2})$, and the absolute minimum value is $-\frac{1}{2}$, which occurs at $(-1/\sqrt{2}, 1/\sqrt{2})$ and $(1/\sqrt{2}, -1/\sqrt{2})$.

Lagrange Multipliers. The goal of the method of *Lagrange multipliers* is to optimize (meaning maximize or minimize) a function subject to a constraint, meaning that we want to optimize the function only among points satisfying the given constraint. For instance, in the two-variable case we have a function $f(x, y)$ we want to optimize and the constraint is described by an equation of the form


$$g(x, y) = k.$$

In the three-variable case, we have a three-variable function to optimize and the constraint will be described by a three-function as well, and so on.


Here is the key geometric picture to have in mind, at least in the two-variable case. Say that the level curves of f look like



with the maximum of f among points satisfying the constraint occurring at the point P . The question is: what does the constraint curve have to look like in relation to these level curves? It should certainly pass through P if we are assuming P satisfies the constraint, but we can say more. The constraint curve cannot look like



since this would lead to points satisfying the constraint curve which give a *larger* value for f than P does, which is not possible if we are saying that P is where the maximum occurs. Thus, the constraint curve can only look like



with the point being that at a maximum the constraint curve and level curve must be *tangent* to each other. A similar reasoning shows that the same is true at a minimum.

Now, $\nabla f(P)$ is perpendicular to the level curve of f containing P and $\nabla g(P)$ is perpendicular to the constraint curve at P , so since these two curves are tangent to each other, these two gradients must be parallel to each other. Hence the conclusion is:

At a point which gives the maximum or minimum value of f subject to the constraint determined by a function g , $\nabla f = \lambda \nabla g$ for some scalar λ .

Thus, solving $\nabla f = \lambda \nabla g$ gives us the candidate points for the maximums/minimums of f subject to the constraint $g = k$. All this works for three-variable optimization problems as well.

Important. To optimize a function $f(x, y)$ subject to the constraint $g(x, y) = k$, first find the points satisfying

$$\nabla f(x, y) = \lambda \nabla g(x, y) \text{ for some } \lambda,$$

and then determine whether those points give maximums or minimums. The same applies to three-variable functions with three-variable constraints.

Back to Warm-Up. Going back to the Warm-Up, after finding the critical points of $f(x, y) = xy$ inside the disk we were left with checking the boundary $x^2 + y^2 = 1$. Now we can do this part using Lagrange multipliers. The constraint is given by

$$g(x, y) = x^2 + y^2 = 1,$$

so the Lagrange multiplier equation $\nabla f = \lambda \nabla g$ becomes

$$(y, x) = \lambda(2x, 2y).$$

Comparing components on both sides gives $y = \lambda 2x$ and $x = \lambda 2y$, so together with the constraint we get the three equations:

$$\begin{aligned} y &= 2\lambda x \\ x &= 2\lambda y \\ x^2 + y^2 &= 1 \end{aligned}$$

which must be satisfied by any point giving the max/min value of f along the circle $x^2 + y^2 = 1$.

To solve these, note that we can assume x and y are nonzero since if one of them is zero we get the value 0 for f , and there are definitely points on the circle which give positive values for f and points which give negative values, so 0 will be neither the max nor the min. We can then also assume $\lambda \neq 0$ since otherwise the first and second equations would give $y = x = 0$. Then dividing first equation by the second gives

$$\frac{y}{x} = \frac{x}{y}, \text{ so } y^2 = x^2 \text{ and hence } y = \pm x.$$

Substituting into the constraint gives

$$x^2 + x^2 = 1, \text{ so } x = \pm \frac{1}{\sqrt{2}}.$$

Then finding the corresponding y -values gives the same points $(\pm \frac{1}{\sqrt{2}}, \pm \frac{1}{\sqrt{2}})$ we found in the Warm-Up for where the maximum and minimum of f can occur along the boundary.

Example 1. We want to find the largest possible product of three positive numbers x, y, z whose sum is 100. That is, we want to maximize the function $f(x, y, z) = xyz$ subject to the constraint that $g(x, y, z) = x + y + z = 100$. The Lagrange multiplier equation $\nabla f = \lambda \nabla g$ is

$$(yz, xz, xy) = \lambda(1, 1, 1).$$

Equating components gives—together with the constraint—the equations

$$\begin{aligned}yz &= \lambda \\xz &= \lambda \\xy &= \lambda \\x + y + z &= 100.\end{aligned}$$

All of our numbers are positive so λ cannot be zero and we can thus divide any equation by any other. (Even if we had assumed our numbers were only nonnegative, if one were zero then xyz would be 0, which is not going to be the maximum value we're looking for.) Dividing the first equation by the second shows that $y = x$ and dividing the first by the third shows that $x = z$. Then the constraint gives

$$x + x + x = 100, \text{ so } x = \frac{100}{3}.$$

Thus the maximum value of f subject to the given constraint occurs when $x = y = z = \frac{100}{3}$ and is the value

$$f\left(\frac{100}{3}, \frac{100}{3}, \frac{100}{3}\right) = \frac{100^3}{27}.$$

To be thorough, we should give a reason why the value we found is indeed a maximum value and not a minimum value. After all, the method of Lagrange multipliers only gives us points where we have either a max or min, and without having a second candidate point to compare our value to we can't say with certainty yet that we've actually found the maximum value. However, since 1, 1, and 98 also satisfy the constraint and

$$f(1, 1, 98) = 98 < \frac{100^3}{27},$$

the value we found cannot be a minimum and so must be the maximum as claimed.

Example 2. Consider an open rectangular box without a lid. We want to determine the dimension of the box which result in the maximum possible volume among those boxes with surface area 100. Denoting the dimensions by x, y, z (z is height) we thus want to maximize the volume function $f(x, y, z) = xyz$ subject to the constraint

$$g(x, y, z) = xy + 2yz + 2xz = 100,$$

which comes from figuring out the surface area of the box. (The xy term has no 2 in front since the box has no lid.) Then $\nabla f = \lambda \nabla g$ becomes

$$(yz, xz, xy) = \lambda(y + 2z, x + 2z, 2x + 2y).$$

Equation components and including the constraint gives the equations

$$\begin{aligned}yz &= \lambda(y + 2z) \\xz &= \lambda(x + 2z) \\xy &= \lambda(2x + 2y) \\xy + 2yz + 2xz &= 100.\end{aligned}$$

To solve these, note that the left sides of the first three equations are pretty similar and become equal after multiplying the first equation through by x , the second by y , and the third by z :

$$\begin{aligned}xyz &= \lambda(xy + 2xz) \\xyz &= \lambda(xy + 2yz) \\xyz &= \lambda(2xz + 2yz).\end{aligned}$$

Then subtracting the first two equations gives

$$0 = 2\lambda z(x - y).$$

None of the dimensions can be zero since this certainly wouldn't give a maximum volume (we wouldn't even really have a box at all), and λ can't be zero since this would imply that some of the dimensions were zero. Thus we must have

$$x - y = 0, \text{ so } x = y.$$

Subtracting the first and third equations from before gives

$$0 = \lambda x(y - 2z) + 2\lambda z(x - y) = \lambda x(y - 2z)$$

since we already know that $x = y$. Again, λ and x are not zero so $y - 2z = 0$ and hence $y = 2z$. Thus so far we know that the dimensions of the box we're looking for will result in the same length and width with the height being half the length.

Now we find the exact values of $x.y.z$. Substituting $x = y$ and $z = \frac{y}{2}$ into the constraint gives

$$y^2 + y^2 + y^2 = 100, \text{ so } y = \frac{10}{\sqrt{3}}.$$

(We ignore the negative square root since y should be a positive width.) Hence we have

$$x = \frac{10}{\sqrt{3}}, \quad y = \frac{10}{\sqrt{3}}, \quad z = \frac{5}{\sqrt{3}}.$$

To show that these dimensions indeed give a maximum volume and not a minimum volume, we argue as follows. Consider shrinking the height and width of the box but at the same time increasing the length so that the surface area stays fixed at 100. Then the volume, because the height and width are approaching 0, will approach zero as well. Since we can make the volume arbitrarily small while keeping the surface area at 100, there is no minimum volume so the dimensions we found must give a maximum volume.

(If we want to be really explicit, consider an open box with $y = z$ and

$$x = \frac{100 - 2z^2}{3z}.$$

Then this box always has surface area 100 and the volume

$$xyz = \left(\frac{100 - 2z^2}{3z}\right) z^2 = \left(\frac{100 - 2z^2}{3}\right) z$$

approaches 0 as z approaches 0. Note that in turn the length x gets larger and larger.)

Lecture 26: More on Lagrange Multipliers

Last day of class! Today we continued talking about the method of Lagrange multipliers, finishing the quarter by looking at some more examples. We also spoke about to generalize all this to multiple constraint scenarios.

Warm-Up. We find the maximum value of $f(x, y) = x^2 + y^3 - 3y + 10$ over the region $x^2 + (y - 1)^2 \leq 1$, which is the region enclosed by the circle $x^2 + (y - 1)^2 = 1$. First:

$$\nabla f = (2x, 3y^2 - 3)$$

is $\mathbf{0}$ only when $x = 0$ and $y = \pm 1$, so $(0, 1)$ and $(0, -1)$ are the only critical points of f . Of these, only $(0, 1)$ is the region we're considering, so we forget about the other critical point.

Now, to test for possible maximums along the boundary of our region we use the method of Lagrange multipliers; that is, we maximize f subject to the constraint $g(x, y) = x^2 + (y - 1)^2 = 1$. The equation $\nabla f = \lambda \nabla g$ is

$$(2x, 3y^2 - 3) = \lambda(2x, 2y - 2),$$

so the maximum point must satisfy the equations

$$\begin{aligned} 2x &= 2\lambda x \\ 3y^2 - 3 &= \lambda(2y - 2) \\ x^2 + (y - 1)^2 &= 1. \end{aligned}$$

The first equation can be written as $2x(\lambda - 1) = 0$, so $x = 0$ or $\lambda = 1$, and we consider these possibilities separately. When $x = 0$ then the constraint becomes

$$(y - 1)^2 = 1, \text{ so } y = 0 \text{ or } 2.$$

Thus $(0, 0)$ and $(0, 2)$ are candidate points for the maximum. When $\lambda = 1$, the second equation above becomes

$$3y^2 - 3 = 2y - 2, \text{ so } 3y^2 - 2y - 1 = (3y + 1)(y - 1) = 0,$$

so $y = -\frac{1}{3}$ or 1 ; $y = -\frac{1}{3}$ gives a point which is not in our region so we ignore it, and when $y = 1$ the constraint gives $x = \pm 1$, so $(1, 1)$ and $(-1, 1)$ are also candidate for the maximum point.

Thus, all together, the maximum value of f over the given region must occur at one of

$$(0, 1), (0, 0), (0, 2), (1, 1), (-1, 1).$$

Plugging these all into f gives $f(0, 2) = 12$ as the maximum value over the given region.

Example 1. Suppose we are constructing an open (i.e. no lid) can in the shape of a cylinder, where the material for the base costs $\$5/\text{cm}^2$ and the material for the upright side costs $\$2/\text{cm}^2$. We determine the dimensions which minimize the cost of constructing the can if we want the volume to be $40\pi \text{ cm}^3$.

Letting r, h denote the radius and height, the total cost of making the can is

$$f(r, h) = 5\pi r^2 + 4\pi r h,$$

which is obtained by multiplying the area of the base and side by the corresponding cost per unit area. Thus we want to minimize f subject to the constraint $g(r, h) = \pi r^2 h = 40\pi$. Lagrange multipliers gives the equation

$$(10\pi r + 4\pi h, 4\pi r) = \lambda(2\pi r h, \pi r^2),$$

so the dimensions we want must satisfy

$$\begin{aligned}10\pi r + 4\pi h &= 2\lambda\pi r h \\4\pi r &= \lambda\pi r^2 \\ \pi r^2 h &= 40\pi.\end{aligned}$$

We can assume r and h are nonzero since otherwise the volume could not be 40π , and hence we can also assume $\lambda \neq 0$ since otherwise the second equation above would give $r = 0$. The second equation then gives

$$r = \frac{4}{\lambda}.$$

Substituting into the first equation gives

$$10\pi \left(\frac{4}{\lambda}\right) + 4\pi h = 2\lambda\pi \left(\frac{4}{\lambda}\right) h,$$

which simplifies to $h = \frac{10}{\lambda}$. Comparing $r = \frac{4}{\lambda}$ and $h = \frac{10}{\lambda}$ gives

$$h = \frac{2}{5}r,$$

and plugging this into the constraint gives

$$\pi r^2 \left(\frac{2}{5}r\right) = 40\pi, \text{ so } r^3 = 16.$$

Hence $r = \sqrt[3]{16}$ and $h = \frac{2}{5}\sqrt[3]{16}$ are the dimensions which minimize cost.

To be sure that this gives a minimum and not a maximum, note that we can increase r and decrease $h = \frac{40}{r^2}$ accordingly to keep the volume at 40π , and this will lead to larger and larger costs since increasing the area of the base has a greater effect on cost than decreasing the height. Thus there can be no maximum cost, so the dimensions we found indeed give minimum cost.

Example 2. Suppose a business sells three products, with product i costing p_i dollars per unit to produce. If x_i denotes the amount of product i produced, then the total cost of production is

$$C(x_1, x_2, x_3) = x_1p_1 + x_2p_2 + x_3p_3.$$

Say that the “utility” (meaning monetary utility or some other measure of utility) derived from producing such amounts is given by the *utility function*

$$U(x_1, x_2, x_3) = x_1x_2^2x_3^3.$$

We want to determine how to allocate our fixed amount of funds, say D dollars, to producing these three products in order to maximize utility. Note that $x_i p_i$ is the total cost of producing product i , so we want to determine how much of D each of $x_1 p_1$, $x_2 p_2$, and $x_3 p_3$ should take up so that U is maximized subject to the constraint $C(x_1, x_2, x_3) = D$.

The Lagrange multipliers equation $\nabla U = \lambda \nabla C$ is

$$(x_2^2 x_3^3, 2x_1 x_2 x_3^3, 3x_1 x_2^2 x_3^2) = \lambda(p_1, p_2, p_3).$$

Together with the constraint this gives the equations

$$x_2^2 x_3^3 = \lambda p_1$$

$$\begin{aligned}2x_1x_2x_3^3 &= \lambda p_2 \\3x_1x_2^2x_3^3 &= \lambda p_3 \\x_1p_1 + x_2p_2 + x_3p_3 &= D.\end{aligned}$$

Note that multiplying the first equation through by $6x_1$, the second by $3x_2$, and the third by $2x_3$ gives

$$\begin{aligned}6x_1x_2^2x_3^3 &= \lambda 6x_1p_1 \\6x_1x_2^2x_3^3 &= \lambda 3x_2p_2 \\6x_1x_2^2x_3^3 &= \lambda 2x_3p_3\end{aligned}$$

with all left-hand sides being the same. Thus all right-hand sides are the same, and since we can assume $\lambda \neq 0$ (since otherwise we'd have some of the amounts x_i being zero, in which case utility would be 0), we have

$$6x_1p_1 = 3x_2p_2 = 2x_3p_3.$$

Thus $x_2p_2 = 2x_1p_1$ and $x_3p_3 = 3x_1p_1$, so the constraint becomes

$$6x_1p_1 = D, \text{ so } x_1p_1 = \frac{D}{6}.$$

Thus $x_2p_2 = \frac{D}{3}$ and $x_3p_3 = \frac{D}{2}$, meaning that to maximize utility we should devote half of our funds to product 3, a third to product 2, and a sixth to product 1. Note that it makes sense that product 3 should have the most funds dedicated to it since increasing x_3 has a greater affect on the utility function than increasing x_1 or x_2 do.

What is λ ? Now we can give a partial to question: what does λ in the equation $\nabla f = \lambda \nabla g$ mean? Remember that this equation came about from wanting ∇f and ∇g to be parallel, so at first glance it seems that λ is just some random scalar, but it turns out that in many applications λ has a real interpretation.

In general, λ is called a *Lagrange multiplier* of the optimization product. In the example above λ actually describes what's called *marginal utility*, and similarly in many other economics or financial applications the multipliers often related to *marginal cost*, *marginal price*, and so on. In general, the multipliers in a sense describe the "change in f with respect to the constraint variables", whatever that means. (Left for future courses.)

Two constraint example. Finally, let us finish with an example of a Lagrange multiplier setup with two constraints. Say we want to find the point on the line of intersection of the planes $x - 2y + 3z = 8$ and $2z - y = 3$ which is closest to the point $(2, 5, -1)$. In other words, we want to maximize the function

$$f(x, y, z) = \sqrt{(x - 2)^2 + (y - 5)^2 + (z + 1)^2}$$

describing the distance from a point to $(2, 5, -1)$ subject to the constraints that (x, y, z) should satisfy

$$x - 2y + 3z = 8 \text{ and } 2z - y = 3.$$

As a simplification, note that distance is minimized precisely when the term under the square root is minimized, so instead we will minimize

$$f(x, y, z) = (x - 2)^2 + (y - 5)^2 + (z + 1)^2,$$

which avoids having to work with square roots. With the constraint functions

$$g_1(x, y, z) = x - 2y + 3z \text{ and } g_2(x, y, z) = 2z - y,$$

the method of Lagrange multipliers with *two constraints* says that the extrema occur at points where

$$\nabla f = \lambda_1 \nabla g_1 + \lambda_2 \nabla g_2,$$

which is similar to the single-constraint Lagrange multiplier equation only with an additional gradient term added in. In general, more constraints would lead to more gradient terms.

In our case, this equation is

$$(2x - 4, 2y - 10, 2z + 2) = \lambda_1(1, -2, 3) + \lambda_2(0, -1, 2),$$

which together with the two constraints gives the equations:

$$\begin{aligned} 2x - 4 &= \lambda_1 \\ 2y - 10 &= -2\lambda_1 - \lambda_2 \\ 2z + 2 &= 3\lambda_1 + 2\lambda_2 \\ x - 2y + 3z &= 8 \\ -y + 2z &= 3 \end{aligned}$$

which must be satisfied by the point we're looking for. One way to solve these is to rewrite them as a system of five linear equations and then use row operations. Instead, from the first three equations we have

$$x = \frac{1}{2}(\lambda_1 + 4), \quad y = \frac{1}{2}(10 - 2\lambda_1 - \lambda_2), \quad z = \frac{1}{2}(-2 + 3\lambda_1 + 2\lambda_2),$$

and plugging these into the constraints results in

$$\begin{aligned} 14\lambda_1 + 8\lambda_2 &= 38 \\ 8\lambda_1 + 5\lambda_2 &= 20. \end{aligned}$$

Solving these two linear equations gives $\lambda_1 = 5$ and $\lambda_2 = -4$, and plugging these back into the expressions for x, y, z above gives $(\frac{9}{2}, 2, \frac{5}{2})$ as the point on the line of intersection of the two planes which is closest to $(2, 5, -1)$. To guarantee that is a "closest" point and not a furthest point, note that the line of intersection extends nonstop in either of its directions, and moving along on such nonstop direction gives points further and further away from $(2, 5, -1)$, so the point we found must give a minimum distance to this point.