

This is the review I put together for my MATH 311 final. My final was very heavily weighted towards the multivariable calculus (Cal 3) portion of the class, so that's all I covered here, as I had flash cards and messy notes from the previous units. I did okay on this final, not as good as I would have liked, but I believe this was mostly due to not doing enough practice problems. So be warned, concepts alone are not enough to study, you need to build fluency in this information through practice.

Topics Covered

- Basics
 - Cross Product and Dot Product
 - The Del operator
 - Gradient, Divergence, and Curl
- 2D Math
 - Area of a Parallelogram
 - Line Integrals
 - Green's Theorem
- 3D Math
 - Volume of a Parallelepiped
 - Surface Integrals
 - Stokes's Theorem
 - Gauss's Theorem
- Gradient / Conservative Fields

Basics

Vectors

There are two types of vectors, but they are essentially the same. One is just a group of scalars:

$$\vec{a} = \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{bmatrix} = (a_1, a_2, \dots, a_n)$$

While the other is called a vector valued function also known as a vector field:

$$F(x, y, z) = \begin{bmatrix} f_1(x, y, z) \\ f_2(x, y, z) \\ f_3(x, y, z) \end{bmatrix} = (f_1(x, y, z), f_2(x, y, z), f_3(x, y, z))$$

We can think of the first as a set of coordinates in \mathbb{R}^n , whereas the second is a function to give us a set of coordinates. You can technically have vector function to n-dimensions, but that's hard for our brain to hold, and so we usually keep things to 2 or 3 dimensions. Although the vector field has multiple dimensions, each of its components are just scalar (normally defined) functions, which here we will refer to as scalar fields.

Cross Product

Given two 3-dimensional vectors, we can represent the cross product as the the determinant of them, where the top row is the unit vectors \hat{i}, \hat{j} , and \hat{k} , the middle row is the first vector, and the bottom row is the second vector:

$$\vec{a} \times \vec{b} = \begin{vmatrix} \hat{i} & \hat{j} & \hat{k} \\ a_1 & a_2 & a_3 \\ b_1 & b_2 & b_3 \end{vmatrix}$$

Or we can represent it as

$$\|\vec{a} \times \vec{b}\| = \|\vec{a}\| * \|\vec{b}\| \sin(\vec{a}, \vec{b})$$

Dot Product

Given two vectors of n dimensions, the dot product is the sum of the product of the corresponding entries:

$$\vec{a} \cdot \vec{b} = a_1b_1 + a_2b_2 + \dots + a_nb_n$$

Or we can represent it as

$$\|\vec{a} \cdot \vec{b}\| = \|\vec{a}\| * \|\vec{b}\| \cos(\vec{a}, \vec{b})$$

Del Operator

Many of our operations in multiple variables happens in terms of derivatives of different ways. To summarize all of these, we make an operator that is the partial derivatives to an n-th degree:

$$\nabla = \begin{bmatrix} \frac{\partial}{\partial x_1} \\ \frac{\partial}{\partial x_2} \\ \vdots \\ \frac{\partial}{\partial x_n} \end{bmatrix}$$

Generally though, we use this in 2 or 3 dimensions, which looks like

$$\nabla = \begin{bmatrix} \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \\ \frac{\partial}{\partial z} \end{bmatrix}$$

This is an operator, but it's vector valued, so normal vector / matrix operations apply

Gradient

The simplest and easiest thing to comprehend as we start to combine ideas is the gradient. This operation takes in a scalar (normally) defined function and produces a vector-valued function that shows the rate of change in each direction. It can be represented as

$$\text{grad } f = \nabla f = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z} \right)$$

In single-variable calculus, the normal derivative of a function gave the function's rate of change in that single direction, and this is the equivalent when you have multiple variables. Another name for the gradient is the "directional derivative" and turns a scalar field into a vector field. It's defined this way, because the scalar function is treated as a scalar and is multiplied into the vector del operator.

For a physical representation, imagine your vector field as the flow of a fluid, the gradient vector points in the direction of greatest increase.

Divergence

Okay, now lets talk about something that might be slightly harder, but hopefully not terrible. If the gradient turns a scalar field into a vector field, the divergence does the opposite, turning a vector field

into a scalar field. Conceptually, we need to turn a vector into a single scalar expression. The most common way we do that is through the dot product:

$$\operatorname{div} F = \nabla \cdot F = \frac{\partial F_1}{\partial x} + \frac{\partial F_2}{\partial y} + \frac{\partial F_3}{\partial z}$$

If we imagine a vector field as describing the flow of a fluid, we can imagine the divergence as the measure of how fast the fluid (vectors) in the field are moving away from a point in space. Building on our fluid analogy, if we were to throw a ball into the fluid field, the divergence would be the amount of force the ball experiences at that point. Divergence is essentially the flux of an incredibly (infinitesimally) small region, as when the flux grows, there's more lines "pushing".

Curl

Okay, one more base operation. This one turns a vector field into another vector field, but instead of measuring the rate of change or the flux, this one measures the "spin" of the field. It takes and see's how each component is changing with respect to the other two dimensions, and to accomplish this, we can use the cross product:

$$\operatorname{curl} F = \nabla \times F$$

For my sanity in writing latex, I left off an expanded form, but it wouldn't be terrible to do by hand, it's just the cross product. After doing it, you'll see that the functions for each coordinate become the partial derivatives of each functions coordinate in terms of the other two.

Parametric Equations

Here's something thats a little different. If we have a scalar function in terms of two variables $f(x, y)$ we can do a formula, to change it into a vector valued function $f(t) = \begin{bmatrix} x(t) \\ y(t) \end{bmatrix}$. For any general function, the easiest way to do this is to assign t to one of the variables and then solve for the other.

2-D Math

Area of a Parallelogram

Let's suppose we're given two vectors \vec{a} and \vec{b} , and say we want to find the area of the parallelogram bounded by them. If these were perpendicular, they would form the base and height of the parallelo-

gram, and let's denote them with \vec{b} and \vec{h} respectively. Now suppose we took the magnitude of their cross product and see what happens:

$$\|\vec{b} \times \vec{h}\| = \|\vec{b}\| * \|\vec{h}\| \sin(\vec{b}, \vec{h}) = \|\vec{b}\| * \|\vec{h}\| (1) = \|\vec{b}\| * \|\vec{h}\|$$

How interesting, we just found that the area of the rectangle bounded by two perpendicular vectors is the area of the cross product bounded by them. If we generalize this into any two vectors, through some proofs I'm not qualified to give here, we can find that:

$$A = \|\vec{a} \times \vec{b}\|$$

Line Integrals

Now that we have the basics, let's do some cool stuff. Suppose I gave you a tiny massless particle and said "please move it from point a to point b, and keep track of *any* work done on the particle please, thank you". That would be a very odd request, but if you were to agree, you'd move the ball from a to b, see that since it was massless you had to apply no force so you didn't do any work, and with nothing else acting on it, you report that the work was zero. Well now suppose I have a magic machine that can apply forces to space at given points by a certain scalar function. Now how do you find the work done on the particle?

Well, if the function is the force, we can check that force at every single point that the particle travels to, multiply it by the distance that force is applied, and sum all of our values together. By hand, this would be very very slow, and pretty inaccurate as the force changes very quickly, so you'd need many, many, teeny tiny partitions to get it right. But! We have a very nice tool that can do all of this partitioning mess for us: the integral! So let's think about it, if our force function is $f(x, y)$, and our position function is $r(x, y)$, how could we do our work formula? Well, we could do this with a double integral (you'll find out more about this later), or this is a great time to use parametrization to get everything in terms of a single variable. So do the work, and you'll find that your force function becomes $f(r(t))$ and your position function becomes $r(t)$. Great, now we can calculate the force at t, where t describes our original curve. Now we need to know the distance it's applied for to get the work. Lucky for us, we can leverage the fact that we have a function describing the position in terms of one variable, as we can take the derivative without doing a partial derivative, and see how quickly we're moving through space, represented by $r'(t)$, and if we take the magnitude of it, we'll have a step size. Remember, when you parametrize equations, you end up with a vector function, so the magnitude is normal vector magnitude. So, we have a function representing force at a point t AND we have the distance that force is applied for through $\|r'(t)\|$, now we just have to apply that for every point in our path:

$$\int_x f ds := \int_a^b f(r(t)) * ||r'(t)|| dt$$

What we just found is called a line integral, which makes sense, it's what happens to us as we move down a line. Okay, cool! Now, going back to our original analogy, what happens if I say "Well dammit, I didn't expect you to actually figure that out, so I'm changing my machine, find the work done if this produces a vector valued function instead." Confused why I'm upset for doing exactly what you ask, but pushing on nonetheless, you begin to think. A vector valued function is basically just a parametrized function where you still have multiple variables but have multiple separate functions for each coordinate. Okay... And you still have access to your parametrized $r(t)$ function, which if you'll remember, is also a vector field. Okay... So suppose, since we have two vector fields, but want a single, integrable scalar function, we use the dot product to get these two together. We don't turn the $r'(t)$ into it's magnitude this time because we need to see how it affects each dimension of the vector field. (I think, answers weren't super clear)

$$\int_x F \cdot ds := \int_a^b F(r(t)) \cdot r'(t) dt$$

Well look at that, we solved the problem again. Considering this is the only two types of functions we know about, I just turn off my machine and run away crying because I couldn't stump you, leaving you with the knowledge of line integrals.

P.S. That generic term, before we do the parametrizing is worth understanding. The F term is the function / vector field, while the ds term stands for the change in the path of the line. In our parametrizations, we do the work to get everything in terms of one variable t, but you just have to set up your integral to where the integrand is with respect to one variable.

P.P.S. These line integrals happen for one, continuous curve. If you are given some crazy complex shape, if you can break it down into a sum of pieces, you can still do a line integral. For example, a square becomes 4 different integrals summed together, one for each side of the square.

Greens Theorem

Many people, as they study engineering or math, come to the last unit of Calculus 3, cry their way through it as they try to learn all of the random, strange, conceptual things, get to the end and celebrate that they will never have to see the name of Green again. Then they go to differential equations, confirming that that stuff at the end of calculus 3 was just a long, but over, nightmare. And then they get to this course. And they realize to their great dismay, no, the horrors did not end, they are coming back.

Well, lets play with it, see what we can find.

To begin our study, let's define some things to make our lives easier. For all of this discussion, suppose we have a 2-dimensional vector field $F = \begin{bmatrix} P(x, y) \\ Q(x, y) \end{bmatrix}$ and a vector field describing our path $r(x, y) = \begin{bmatrix} x \\ y \end{bmatrix}$ and we're working on the path of the line from a to b.

Now that we have some definitions, lets do some interesting things. In these terms, lets take the line integral from a to b, but instead of parametrizing the curve, lets keep it in terms of x and y. When we do this, our ds term, which represents the path, is just the derivative of our $r(x, y)$ vector:

$$\int_a^b F(x, y) \cdot r'(x, y) = \int_a^b \begin{bmatrix} P(x, y) \\ Q(x, y) \end{bmatrix} \cdot \begin{bmatrix} dx \\ dy \end{bmatrix} = \int_a^b P(x, y)dx + Q(x, y)dy$$

Interesting. Okay. So now we have a term that is the sum of two things. It's very very very important and vital to remember, however, that dx and dy can *not* be treated as the term to integrate with respect to, because that's not what they are, its the change in x of the path, you can't integrate with respect to them because the two other functions are in terms of both variables. Anyway, if you go through and do your parametrization, you can see how you end up in this form when you put everything in terms of t, try it if you don't believe me.

So now that we've seen that form of a line integral, lets play with it. That thing can be... very gross to integrate for anything that's not super simple. But lets think big picture and come up with something better, but, an important note is that from here on out, we're going to be talking about **Closed** lines, as in, the line loops around somehow, someway, and we end up back where we started, wether that's by a square, or a circle, or an ellipse. It doesn't matter *what* the shape is, only that it *is* a shape. We will denote this special kind of surface with the symbol:



Notice the *closed* circle in the middle? But anyway, pressing forward. It is a convention in math (and a lot of other fields) that moving counter-clockwise is deemed "positive". When you're doing math yourself, you'll see this naturally happens. I'm not sure the exact reason why, but I blame the unit circle and the trig functions defining the angle as moving counterclockwise.

Building on this, and using our typical analogy of a vector field representing fluid flow, we know that the line integral measures the total force applied as you move from a to b in a line integral. Along this idea, it's (hopefully) not too much of a stretch to think a positive result means the field was helping push you along your path, while a negative result means the field was pushing against you as you moved. Neat! Now lets take this into a closed loop. If we move counter clockwise around a loop and

end up with a positive value, then the vector field must also have some sort of spin to it in the counter clockwise direction because it overall helped push us. If we end up with a negative value, then the vector field will still be spinning, but it has to be spinning in the clockwise direction, because it pushed against us as we moved. What we just measured here is called the **Circulation** of the field, and it follows our convention of counter clockwise being positive.

What we measured in our line integral, however, is the big picture circulation, it's what overall happens as you move along that path. What Green found, is that if you take your area enclosed by that curve and split it up into a bunch of super-duper tiny regions, to the point where your basically looking at the circulation of a specific point, and then summed them all together you'd end up with the *same* result! This is because, all of the boundaries would be part of 2 calculations, one going clockwise and one going counterclockwise, *except* for the outermost boundary of the original line, so everything would cancel except for the original line.

Well, we have a way to check how much a field is spinning at a point, that's the curl of the field! Since we're in two dimensions, we're going to only deal with what's called the "scalar curl", which is defined as:

$$(\nabla \times F) \cdot \vec{k} = \frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y}$$

Okay, so we know how to find the spinning at a specific point, now we need to do that for every point in the region. That's not hard, it's just a double integral!

$$\iint_D \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} \right) dA$$

How neat!! We just found a way to represent a line integral as the integral over an area, and that's what Greens theorem says:

$$\int_C F \cdot ds = \int_C Pdx + Qdy = \iint_D \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} \right) dA = \iint_D (\nabla \times F) \cdot \vec{k} dA$$

If you understand what Greens theorem represents, it's actually a super cool theorem, and hopefully a bit easier to remember.

3-D Math

This is where things can get... weird. A lot of this math is taking what we did in 2-D and generalizing it to the third.

Volume of a Parallelepiped

Parallelepiped is a big scary word, it's literally just the 3-d equivalent of a parallelogram, where it's bounded by flat surfaces, but the angles and ratios aren't set like in the case of a cube. Because that word is big (and quite frankly hard to spell) I'll probably just call it a shape from now on.

Suppose we're given a shape bounded by straight lines that exists in 3 dimensions. That is, it has a width, height, and depth. Now if we were to represent each of these dimensions with vectors, how could we find the volume of that shape? Well, as we saw in the 2-d section, the area of a parallelogram is the cross product of the two vectors. Building on that, if we have the area of the base, then we only need to multiply it by the height of the shape to get our volume, a concept we learned in geometry. So we need to find a way to express: $A_{base} * height$

Well, suppose we have 3 vectors: $\vec{a}, \vec{b}, \vec{c}$. We have the area of the base, it's $\|\vec{a} \times \vec{b}\|$. But that cross product there also produces a vector perpendicular to the base. And if we were finding the area of a cube, where \vec{c} would be perpendicular to both \vec{a} and \vec{b} then our volume would be:

$$V = \|\vec{a} \times \vec{b}\| * \|\vec{c}\|$$

Hey, that looks like a formula we know, it's the magnitude of the dot product, where the angle between \vec{c} and the vector resulting from the cross of \vec{a} and \vec{b} is 0° !

$$V = \|\vec{a} \times \vec{b}\| * \|\vec{c}\| \cos(\vec{a} \times \vec{b}, \vec{c})$$

We can generalize this into the formula we want:

$$V = (\vec{a} \times \vec{b}) \cdot \vec{c}$$

Surface Integrals

Okay. This is basically equivalent to a line integral in 3-D space. It's not the same, but the ideas are very similar. Suppose we have some three dimensional function $f(x, y, z)$ and we want to integrate across the surface area of it. The general form of a surface integral looks like this:

$$\iint_S f(x, y, z) dS$$

Well... That's a thing. I sure don't know how to solve it though. What does the dS term mean? In line integrals, the lowercase ds meant the differential with respect to each step of the line, and its the same

here, but across the surface. dS means all of the little bits on the surface summed together. You'll notice very quickly that this also looks similar to the area of a normal double integral ($\iint f(x, y) dA$) and that's because it *is* very similar to a double integral, we just now have the potential for a curved surface, and are working in three dimensions. Building on this idea, we're going to figure out how to represent our 3d shape as a standard double integral.

The reason a surface integral get's tricky is that we're trying to calculate a 2-d value (area) from a 3-d object, but if we can represent our shape in only 2 dimensions, then we can work more easily. So... suppose we did something of a projection onto one of the planes, say the x-y plane, so that we only saw the 2-d shape. That'd be cool, we'd have a 2-d shape, do the double integral, get the area, walk away right?

Wrong.

If you were to do this with say a sphere, the only thing you'd project onto that plane is the circle at the diameter of the sphere, losing all of the other data, so we can't stop with just a projection. To solve this, suppose we got our third variable, in this case z , as a function of the other two. Then, still using only two variables, we'd have a way to find the area of the whole surface! Doing this with a scalar function produces:

$$\iint_S f(x, y, z) ds = \iint_D f(x, y, g(x, y)) * \|\nabla g(x, y)\| dA = \iint_D f(x, y, g(x, y)) * \sqrt{\left(\frac{\partial g}{\partial x}\right)^2 + \left(\frac{\partial g}{\partial y}\right)^2 + 1}$$

If you got tripped up on that second conversion like I did, don't forget $z = g(x, y)$ so the partial with respect to it is just 1. If we remember that the gradient is the multi-variable equivalent to the single-variable, full, derivative, our surface integral formula starts to look reeeally similar to our line integral which looked like:

$$\int_D f(r(t)) * \|r'(t)\| dt$$

Bringing up the topic of parametrization, we need to figure out how to do that to continue the topic of surface integrals. Notice here, that in the end, we get everything in terms of only two variables, x and y . What if we generalize this and actually just parametrize the surface to find some vector valued function

$v(s, t) = \begin{bmatrix} v_1(s, t) \\ v_2(s, t) \\ v_3(s, t) \end{bmatrix}$ from here, we could substitute that into our equation and find (for really complex reasons) that:

$$\iint_S f(v(s, t)) * \left\| \frac{\partial v}{\partial s} \times \frac{\partial v}{\partial t} \right\| ds dt$$

Parametrization is weird, that's something you'll need to study, but this is what a surface integral is... yay.

Repeating for a vector function, it's not too dissimilar and is pretty similar to the form of a line integral over a vector field, so I'll just give it:

$$\iint_S F(x, y, z) dS = \iint_S F(v(s, t)) \cdot \left(\frac{\partial v}{\partial s} \times \frac{\partial v}{\partial t} \right) ds dt$$

This is also known as the flux integral, because if you think about it, it's how much of that "stuff" is going through the surface

Stokes's Theorem

So, now that we know what a surface integral is, we're going to take Green's theorem and generalize it to the *third* dimension. OOoooooh Scary. Just kidding it's not. If you understand green's theorem, this one's not hard.

Remember in Green's theorem, when we took the line integral around a loop, what that ended up measuring was the "circulation" of the vector field, and we found that we could get the big level line integral by finding all of the tiny, circulations? Again, remember how to do this across the area, we used the "scalar curl" which was the curl with respect to just the 2 dimensions we cared about. Well we care about all 3 now, so no more of that silliness. Once we hit 3 dimensions, we get:

$$\int_C F ds = \iint_S \text{curl } F dS = \iint_S \nabla \times F dS$$

Gauss's Theorem

This thing is also called the divergence theorem. I'm tired of writing, so I'm just giving it to you:

$$\iint_S F dS = \iiint_E \text{div } F dV = \iiint_E \nabla \cdot F dV$$

What's cool is that this turns a surface integral into a easy, simple, triple integral. It probably has a cool derivation, I just don't care at the moment.

Gradient / Conservative Fields

Here's a section I couldn't decide where to put because it applies to both two and three dimensions, so I put it here in its own thing. We'll talk about a few things and then build towards a nice new theorem.

Suppose I give you a scalar function $f(x, y)$ and then took the gradient of it, and ended up with:

$$\nabla f(x, y) = \left[\frac{\partial f}{\partial x} \quad \frac{\partial f}{\partial y} \right]$$

Well, like we said when we defined the gradient, we got a vector field from a scalar function. But what if we wanted to go the other direction? Produce a scalar function from a vector field? Well, vector fields where we can do this are a special type of fields known as gradient fields, and gradient fields are what's known as "conservative".

What's cool about these is that in a conservative field, line integrals are path independent! So given any two points, we don't have to worry about the specific path we're taking to get there, whether it's a straight line or some kind of crazy curve, the line integral from point A to point B will give the same value. We summarize this in the **Fundamental Theorem for Line Integrals**:

$$\int_a^b F ds = \int_a^b (\nabla f) ds = f(b) - f(a)$$

So this is super cool! Now, instead of having to do a bunch of crazy math, it just comes down to evaluating that function, known as the potential function, at the two points. But, importantly, this *only* works for conservative fields because we need their potential function, so how do we know, given some arbitrary vector field, if it is conservative? A vector field is said to be conservative if two conditions are met:

1. It is a simply connected region (There's no holes)
2. It has a curl of 0 (It's not doing a spinny spin)

The first is based off the fact that when a field is conservative, the line integral is the same no matter what, anywhere inside the boundary, but if you evaluate a line integral where the line between them goes through that hole, it won't be the same as connecting two other points. The second is based on the property that the second (and third) partial derivatives of functions are equal. To test this, we use the curl. I'm going to show two dimensions, but 3 is the same:

$$F = \nabla f$$

$$\text{curl } F = \nabla \times F$$