Line and Surface Integrals: An Intuitive Understanding

Joseph Breen

Introduction

Multivariable calculus is all about abstracting the ideas of differentiation and integration from the familiar single variable case to that of higher dimensions. Unfortunately, this can result in what may seem like multiple kinds of generalizations. Take the integral, for example. In multivariable calculus, we have double integrals, triple integrals, line integrals, surface integrals — where does it end? As it turns out, *any* kind of integral can be thought about in a unified way. The aim of this discussion is to give some intuition behind these objects, and in particular, line and surface integrals.

Review

First, let's review integration in the single variable case, emphasizing a certain interpretation. In this case, our domain — i.e., the **thing we are integrating over** — is always an interval $I = [a, b] \subseteq \mathbb{R}$. One way to think about an integral

$$\int_I f(x)\, dx$$

of a function over I is that we take I, chop it up into chunks, give each sum a weight, and add everything back together. More concisely, an integral over I is a **weighted sum** of I, where each chunk dx of I is weighted by $f(x)$, the value of the function at that point.

A Probabilistic Interpretation

If you've taken any probability theory, this should sound really familiar. If you haven't, let's start off easy: given a list of numbers $X = \{x_1, x_2, \ldots, x_n\}$, we all know how to calculate its *average value:*

$$X_{avg} = \frac{x_1 + x_2 + \cdots + x_n}{n}$$

We can write this a little differently as:

$$X_{avg} = \frac{1}{n}x_1 + \frac{1}{n}x_2 + \cdots + \frac{1}{n}x_n$$

In other words, we take all the numbers in our set, give them equal weighting ($\frac{1}{n}$) and add them all back together. In probability theory, we abstract this idea to that of an *expected value.* Instead of giving each number equal weighting, we can assign possibly different weightings. Phrased differently, each number x_i can be given a probability $P(x_i)$ of happening. Then the **expected value** of the set X with respect to P is:

$$\mathbb{E}(X) = P(x_1)\, x_1 + P(x_2)\, x_2 + \cdots + P(x_n)\, x_n = \sum_{i=1}^{n} P(x_i)\, x_i$$

Again, we are taking a *weighted sum over our set X.* Note that if $P(x_i) = \frac{1}{n}$ for all i, we get the average value from above.

What if we wanted to the same thing with an infinite set of numbers, like an interval $I = [a, b]$? Informally, we can take our set I and chop it up into infinitely many chunks dx, and look at the probability $p(x)$ of each chunk occuring. Here, each dx is acting like one of the x_i's from above, and our function $p(x)$ is sort of acting like $P(x_i)$. If we calculate the weighted sum of our infinitely many chunks, we get an integral! In this case, the **expected value** with respect to p is:

$$\mathbb{E}(I) = \int_I p(x)\, dx$$

Hopefully it is clear how naturally an integral in general can be interpreted as a sort of weighted sum or expected value. Note that this gives a nice understanding of the fact that:

$$\int_a^b dx = b - a$$

If we take an interval I, chops it up, give everything an equal weighting of 1, then when we add everything together we're just left with the length of I! This interpretation of an integral as a weighted sum of its domain will be the key idea to understanding *every other kind of integral out there.* Let's quickly interpret this in the context of double and triple integrals.

Double and Triple Integrals

Double Integrals

Say I give you a region $E \subseteq \mathbb{R}^2$. Let's chop it up into little chunks of two-dimensional regions dA. For each little chunk of area at position $(x, y) \in E$, we can weight it by some function $f(x, y)$. Adding all the weighted chunks together gives us:

$$\iint_E f(x, y)\, dA$$

And that's a double integral! We're really not doing anything different than we did before — this is just a reinterpretation of the idea of a *weighted sum.* Again, what happens if we give each dA a weight of $f(x, y) = 1$? Well, when we add everything back together, we get the area of E!

Here's an important point to note: the abstract notation

$$\iint_E f(x, y)\, dA$$

is conceptually nice, but to actually calculate this value, we have to write it in terms of iterated single-variable integrals:

$$\iint_E f(x, y)\, dA = \int_{I_1} \int_{I_2} f(x, y)\, dx\, dy$$

for some intervals I_1 and I_2. This is important to keep in mind as we press on: *the only thing we know how to actually calculate is a regular single variable integral.*

Triple Integrals

We're on a roll now. If I give you a region $E \subseteq \mathbb{R}^3$, we can chop it up into little chunks of volume dV, weight each chunk by the value of $f(x, y, z)$ at that point, and add everything up to get:

$$\iiint_E f(x, y, z)\, dV$$

This is a triple integral! We make the same observations as before: if $f(x, y, z) = 1$, we're weighting every chunk the same, so we recover the volume of E by calculating:

$$\iiint_E dV$$

Also, to actually calculate these integrals, we have to reduce them to single variable integrals:

$$\iiint_E f(x, y, z)\, dV = \int_{I_1} \int_{I_2} \int_{I_3} f(x, y, z)\, dx\, dy\, dz$$

for some appropriate intervals I_1, I_2, and I_3.

Line Integrals

Time to crank it up a notch. Now that we've had some practice with thinking about integrals as weighted sums, we can talk about integrating over more complicated domains. The integrals we just finished discussing took regions that were, in some sense *flat*: intervals, regions of the plane, regions of space, etc. To make things more interesting, let's say I give you a differentiable curve C in the plane:

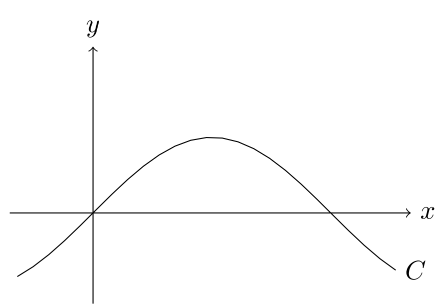

This will be the new "domain" of our integral.

Line Integrals in Scalar Fields

Given a function $f : \mathbb{R}^2 \to \mathbb{R}$, we want to define what it would mean to integrate f over C, i.e., to take a weighted sum of infinitely small chunks of C. Given what we've done so far, this doesn't seem too hard! Let's chop our curve C up into little "chunks of curve" ds, whatever that means, weight each ds by the function value at that point, and add everything back together. That's a **line integral**:

$$\int_C f(x, y)\, ds$$

Conceptually nice? Sure. As far as computation goes, this still seems a little mysterious. After all, we weren't precise about what we meant by ds. To formalize the idea of "a small chunk of curve," we can take advantage of that fact that C is differentiable. Basically, this means that if we pick a point on C and zoom in really closely, it looks like a line! So our ds can be interpreted as an infinitely small line segment, tangent to C. Since C is moving up and down, at any given point, our line segment will have an x-component change and a y-component change. We can visualize this as an infinitely small right triangle:

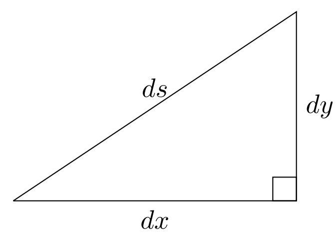

This means, roughly, that $ds = \sqrt{(dx)^2 + (dy)^2}$. Cool.

Now, remember that the *only thing we know how to calculate is single variable integrals.* Our next step is to figure out how to turn $\int_C f(x,y)\sqrt{(dx)^2 + (dy)^2}$ into an integral that we can calculate. In other words, we need some sort of relationship between an interval $I \subseteq \mathbb{R}$ and our curve C. This is exactly what a *parametrization* of a curve is - a mapping from an interval to the graph of the curve. So let's parametrize C in the following way:

$$\mathbf{r} : [a,b] \subseteq \mathbb{R} \to \mathbb{R}^2 \qquad \mathbf{r}(t) = \begin{bmatrix} x(t) \\ y(t) \end{bmatrix}$$

Now we can use the fact that $dx = \frac{dx}{dt} \cdot dt$ (and similarly for dy) to turn our expression into an integral entirely in terms of t:

$$\int_C f(x,y)\, ds = \int_a^b f(x(t), y(t)) \sqrt{\left(\frac{dx}{dt}\right)^2 + \left(\frac{dy}{dt}\right)^2}\, dt$$

Just to reiterate: all we did here was interpret the idea of taking weighted sum in terms of small chunks of our domain. Nothing too different so far! Also, note that if $f(x,y) = 1$ (in other words, every chunk of curve has an equal weighting of 1), we get the length of C!

From here, we could do the same thing for a curve in 3 dimensions, changing only our parametrization ($\mathbf{r}(t)$ will have three component functions) and our ds term (it will have a dz). What's to stop us from going to 4 or 5 dimensions? Nothing! We can just as easily define an line integral of a function $f : \mathbb{R}^n \to \mathbb{R}$ over a curve in $\mathbb{R}^n$ for any natural number n.

Line Integrals in Vector Fields

As I mentioned in the beginning of our discussion, the abstraction of mathematical objects in one variable to higher dimensions can result in multiple kinds of generalizations. We run into this issue when we abstract the idea of a function: a function defined on a higher-dimensional domain can spit out a number[1], but it can also spit out a vector. We disposed of the former case in the last section, and now we'll tackle the latter.

Recall that a function on a multidimensional domain that spits out a vector in the same space is called a **vector field**. Formally, a vector field is a function $\mathbf{F} : \mathbb{R}^n \to \mathbb{R}^n$. As before, we'll start in the two-dimensional case ($n = 2$) and easily generalize from there. So let C be a curve in the $x - y$ plane, and let $\mathbf{F} : \mathbb{R}^2 \to \mathbb{R}^2$ be a vector field. We need to figure out a way to integrate $\mathbf{F}$ over C. As always, our key intuition will be to think about taking a weighted sum over C. Before, our weights were given by the numbers spit out by our scalar field f. Now, we have a function $\mathbf{F}$ that spits out a *vector,* so we need to reformulate what we mean by "weighting a chunk of curve." After all, we can't do what we did before and just multiply ds by $\mathbf{F}$, because we would get a vector, and we want our integral to be a number. In order to remedy this, we turn to the *dot product.*

> **Conceptualizing the Dot Product**: Given vectors $\mathbf{a}$ and $\mathbf{b}$, one way to think about the dot product $\mathbf{a} \cdot \mathbf{b}$ is that it is a measurement of how similar $\mathbf{a}$ and $\mathbf{b}$ are; i.e., how well they travel together. So if $\mathbf{a}$ and $\mathbf{b}$ are orthogonal (i.e. as "different" as possible), $\mathbf{a} \cdot \mathbf{b} = 0$, whereas if $\mathbf{a}$ and $\mathbf{b}$ are parallel, the magnitude of $\mathbf{a} \cdot \mathbf{b}$ will be as big as possible. So the dot product is a way to compare two vectors — or, dare I say, *weight one vector by another.* We'll keep this in mind as we continue.

The first thing we'll do is parametrize C. We're going to have to do that eventually (remember, to compute stuff we have reduce our problem to a single variable integral), and it will make formulating the integral easier. So from now on, C is the image of the following function:

$$\mathbf{r} : [a, b] \subseteq \mathbb{R} \to \mathbb{R}^2 \qquad \mathbf{r}(t) = \begin{bmatrix} x(t) \\ y(t) \end{bmatrix}$$

Before, when we were looking at ds, this was like a mini tangent line to the curve. In terms of our parametrization $\mathbf{r}$, instead of calculating a mini tangent line, we can just calculate the tangent vector! So for any point $t \in [a, b]$, our analogue to ds will be $d\mathbf{r}$:

$$d\mathbf{r} = \frac{d\mathbf{r}}{dt} \cdot dt = \mathbf{r}'(t)\, dt = \begin{bmatrix} x'(t) \\ y'(t) \end{bmatrix} dt$$

Now, at any point $t \in [a, b]$, we can calculate the value of our vector field $\mathbf{F}$ at that point on the curve by inputting our position on the curve into the function: $\mathbf{F}(\mathbf{r}(t)) = \mathbf{F}(x(t), y(t))$. Remember that this is a vector. Our goal is to figure out how to give $d\mathbf{r}$ a weight. So what we can do is take the dot product of the vector spit out by the vector field at that point with the tangent vector to the curve:

$$\mathbf{F}(x(t), y(t)) \cdot \begin{bmatrix} x'(t) \\ y'(t) \end{bmatrix} dt$$

[1] These types of functions $f : \mathbb{R}^n \to \mathbb{R}$ are called *real-valued functions,* or sometimes *scalar fields,* hence the name of the previous subsection.

Visually, at every point on the curve, we have the following picture:

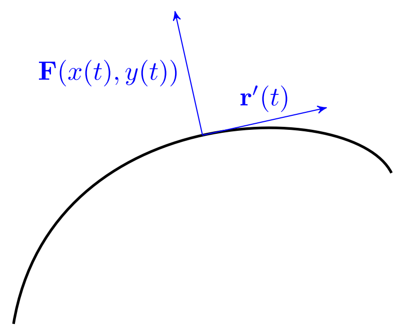

This represents one weighted chunk, so we add up all the chunks on the curve C to give us the line integral of the vector field! Notationally, we write:

$$\int_C \mathbf{F} \cdot d\mathbf{r} = \int_a^b \mathbf{F}(x(t), y(t)) \cdot \begin{bmatrix} x'(t) \\ y'(t) \end{bmatrix} dt$$

Note that by our interpretation of the dot product, this quantity gives a sort of measurement as to *how well the curve flows with the vector field.* If all of the tangent vectors are close in direction to the vector field vectors, our dot product will be larger and therefore the sum will be larger[2]. If the tangent vectors are close to orthogonal to the vector field (so the curve is not following the flow), our dot products will be close to 0 and the line integral will be very small.

As before, we can easily generalize this for a curve C which is the image of some function $\mathbf{r} : [a, b] \subseteq \mathbb{R} \to \mathbb{R}^n$ and a vector field $\mathbf{F} : \mathbb{R}^n \to \mathbb{R}^n$ by just increasing the number of component functions each vector-valued function has.

Surface Integrals

We turn our attention now to integrating over **surfaces**. As before, our formulation will be grounded in the intuitive idea of taking our surface, chopping it up, and calculating a weighted sum of all the pieces. As before, we have to be precise about a couple things: what we mean by a "chunk of surface", and what it means to "weight" a chunk.

Surface Integrals in Scalar Fields

We begin by considering the case when our function spits out numbers, and we'll take care of the vector-valued case afterwards. So let $f : \mathbb{R}^3 \to \mathbb{R}$ be a scalar field, and let M be some surface sitting in $\mathbb{R}^3$.

[2]If the tangents vectors are moving in the same direction *in the opposite direction,* we would get a very large *negative* number.

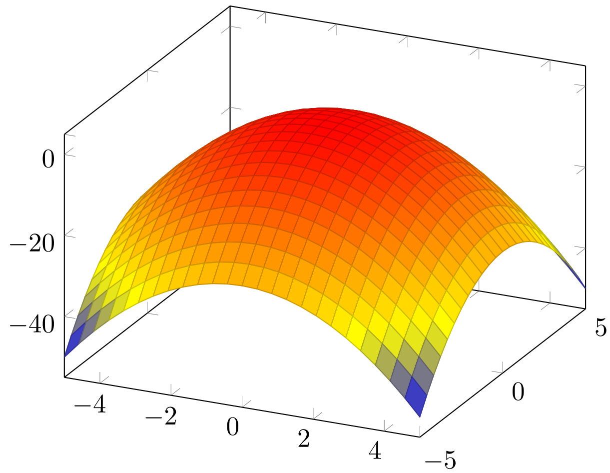

We may as well parametrize our surface, because like before, we're going to have to do that eventually. Note that our parametrization will now pick up a *rectangle* from the plane, as opposed to picking up an interval:

$$\mathbf{r} : [a, b] \times [c, d] \subseteq \mathbb{R}^2 \to \mathbb{R}^3 \qquad \mathbf{r}(s, t) = \begin{bmatrix} x(s, t) \\ y(s, t) \\ z(s, t) \end{bmatrix}$$

So our surface M is the image of $\mathbf{r}$.

Our conceptual development of the surface integral of f over M will mimic what we've done before: chop the surface up into chunks of surface dS, weight each chunk by the value of f at that point, and add everything up:

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

Cool! Next up is to figure out what dS means — if we do that, we're basically done. Recall that with line integrals, we used the tangent vector to encapsulate the information needed for our small chunks of curve. We can try to do the same thing with a surface, but we have an issue: at any given point on M, there are tons of tangent vectors! So for a surface, we want to think of our chunks as little *tangent planes*. Fortunately, tangent plans and tangent vectors are very strongly related, because we can describe a tangent plane using two important tangent vectors: the tangent vector in the "s-direction," and the tangent vector in the "t-direction":

$$\frac{\partial \mathbf{r}}{\partial s} \qquad \text{and} \qquad \frac{\partial \mathbf{r}}{\partial t}$$

To formalize this, we turn to the *cross product*.

> **Conceptualizing the Cross Product**: Given two vectors $\mathbf{a}$ and $\mathbf{b}$, we can calculate their cross product $\mathbf{a} \times \mathbf{b}$. This gives us *another vector*. In particular, it's a vector *orthogonal to both* $\mathbf{a}$ *and* $\mathbf{b}$. *The other cool thing about the vector* $\mathbf{a} \times \mathbf{b}$ *is that its length* $\|\mathbf{a} \times \mathbf{b}\|$ *represents the area of the parallelogram defined by* $\mathbf{a}$ *and* $\mathbf{b}$.

So, at each point on M, we have two important tangent vectors, $\frac{\partial \mathbf{r}}{\partial s}$ and $\frac{\partial \mathbf{r}}{\partial t}$. These define a parallelogram, and this parallelogram will essentially be our mini-tangent plane!

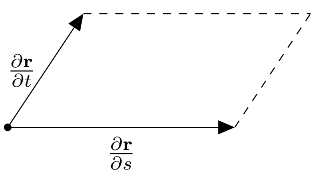

So we can make dS be the area of this little tangent plane, which we calculate using the length of the cross product:

$$dS = \left\| \frac{\partial \mathbf{r}}{\partial s} \times \frac{\partial \mathbf{r}}{\partial t} \right\| ds\, dt$$

From here, we can convert our absract surface integral into a double integral over the $s - t$ region:

$$\iint_M f(x, y, z)\, dS = \int_a^b \int_c^d f(x(s,t), y(s,t), z(s,t)) \left\| \frac{\partial \mathbf{r}}{\partial s}(s,t) \times \frac{\partial \mathbf{r}}{\partial t}(s,t) \right\| dt\, ds$$

Not the most beautiful expression in the world, but just keep in mind that all we're doing is taking a weighted sum of the chunks of the surface.

Surface Integrals in Vector Fields

As with line integrals, the story is (seemingly) different when we have a vector field $\mathbf{F} : \mathbb{R}^n \to \mathbb{R}^n$ instead of a scalar field f. The issue is that we can't weight our chunks of surface by multiplication, because the vector field is spitting out a vector. But, as before, we can remedy this by turning to the dot product. However, instead of calculating the dot product of our vector field with one of the tangent vectors, we will use the *normal vector* to our surface. Recall that we can calculate the normal vector $\mathbf{N}$ to a surface by computing the cross product of the tangent vectors, since this gives us a vector perpendicular to the tangent plane. So our little chunk of surface $d\mathbf{S}$ will be calculated in the following way:

$$d\mathbf{S} = \mathbf{N}\, dS = \frac{\partial \mathbf{r}}{\partial s} \times \frac{\partial \mathbf{r}}{\partial t}\, dt\, ds$$

Note that this is the same cross product that represented the area of our tangent plane in the scalar-field version of our line integral. So given our normal vector $\mathbf{N}$, we can weight the chunk of surface represented by this vector by computing $\mathbf{F} \cdot \mathbf{N}$. Then, as before, add everything up:

$$\iint_M \mathbf{F} \cdot d\mathbf{S} = \iint_M \mathbf{F} \cdot \mathbf{N}\, dS = \int_a^b \int_c^d \mathbf{F}(x(s,t), y(s,t), z(s,t)) \cdot \left(\frac{\partial \mathbf{r}}{\partial s}(s,t) \times \frac{\partial \mathbf{r}}{\partial t}(s,t) \right) dt\, ds$$

And we're done!

Let's think about what we're calculating here. Remember that the dot product measures how similar two vectors are, or, how well they travel in the same direction. Since we're calculating the sum of the dot product of the vector field with the *normal* vector to our surface, the surface integral gives an indication of how much the vector field is moving directly through the surface. In other words, if the surface flowed tangentially to the vectors in the vector field, the dot products would be close to 0, so the surface integral would be close to 0. One way to picture this is by thinking of the surface as a sail, and the vector field as the wind: if the wind is blowing directly into the sail and making it very taut, then the surface integral will be very large (or very negatively large if it's blowing in the opposite direction); if the wind is either not blowing or blowing along side the sail, the surface integral will be very close to 0.

Review

Time to review everything that we've done, just to drive the point home. At this point in your life, you've seen many kinds of integrals: single variable integrals, double integrals, triple integrals, line integrals, and surface integrals. *Every kind of integral can be formulated and interpreted as a weighted sum over its appropriate domain.* But to actually calculate these things, we have to turn them into single-variable integrals, because that's the only thing we know how to calculate! The main thing to take away from this is to not be intimidated when you see any sort of integral — they all represent more or less the same thing.