INTEGRATION OF FUNCTIONS OF SEVERAL VARIABLES

Contents

1. Integration

Now that the quick review of differential calculus of several variables is finished, let's start with the new material in this class. We are interested in developing a theory of integration for functions of several variables. Let us begin with the case of functions $f(x, y)$ of two real variables. We want to define an integral for this function. How should we proceed?

Notice that there is no obvious candidate for an antiderivative, so the strategy of trying to define an integral as an antiderivative is not going to work. We should therefore emulate the definition of a definite integral for functions of a single variable.

How is a definite integral

$$\int_a^b f(x)\,dx$$

defined? Recall that the definition of this integral is as the value of a limit

$$\lim_{||P||\to 0} \sum f(x_i^*)\Delta x_i$$

where the x_i are a partition of $[a, b]$, with largest Δx_i equal to $||P||$, and x_i^* a number between x_i and x_{i+1}. Each of these sums is called a Riemann sum, and can be thought of as an approximation to the area under the graph of $y = f(x)$ on the interval $[a, b]$. This limit is just a fancy way of saying that we are interested in the limit of better and better approximations to this area, which we obtain by making the rectangles thinner and thinner. Presumably, our intuition tells us that this should make the error in approximation smaller and smaller.

What is the correct analogue of the notion of a Riemann sum for functions of two variables? Such a sum should estimate the volume under the graph of $z = f(x, y)$ on some region. The closest analogue to an interval $[a, b]$ in $\mathbb{R}$ might be a rectangle R in $\mathbb{R}^2$, perhaps of the form $[a, b] \times [c, d]$. We want a sum which estimates the volume of the region under $z = f(x, y)$ over the rectangle R.

Just like how we use rectangles to approximate the area under $y = f(x)$, we can use rectangular prisms to approximate the area under $z = f(x, y)$. In particular, over a rectangle $[x_i, x_{i+1}] \times [y_i, y_{i+1}]$, we can use a rectangular prism with that rectangle as base, and height given by $f(x_i^*, y_i^*)$ as an approximation to the volume over this smaller rectangle. This rectangular prism has volume

$$f(x_i^*, y_i^*)\Delta x_i \Delta y_i.$$

($\Delta x_i = x_{i+1} - x_i$, and similarly for Δy_i.) Therefore, the sum of these rectangular prisms, which is approximating the volume under $z = f(x, y)$ over the rectangle R, is

$$\sum f(x_i^*, y_i^*)\Delta x_i \Delta y_i$$

where the summation is over all small rectangles making up the large rectangle R. This sum is called a *Riemann sum* for $z = f(x, y)$ over the rectangle R. If we make these small rectangles smaller and smaller, then the limit of these sums (if the limit exists) will be called the *definite integral* of $f(x, y)$ over R and will be denoted

$$\iint\limits_R f(x, y)\, dA.$$

The two integral signs remind us that we are integrating over a two-dimensional region; namely, the rectangle R which is written underneath the two integral signs. The differential dA tells us that we are integrating with respect to an *area element*. It is mostly to be thought of as notation further reminding us that we are integrating over a two-dimensional region.

Of course, in practice, we never use this definition of a definite integral to actually evaluate definite integrals like

$$\int_a^b f(x)\, dx.$$

We sometimes see how to evaluate these integrals directly using the limit definition for certain special functions, like $f(x) = x^2$, but these calculations are long and tedious. Instead, we use the fundamental theorem of Calculus, which tells us that to evaluate a definite integral we need only find an antiderivative for $f(x)$, and then evaluate it at a, b.

The Riemann sum definition is still useful from a computational point of view sometimes. For instance, there are many functions which are impossible to integrate exactly in any sort of useful way, but which appear frequently in real life (such as the normal distribution, perhaps the most important function in probability and statistics). Nevertheless, we may want to obtain a numerical estimate for these integrals. Since Riemann sums are meant to be approximations of definite integrals, and the computation of a particular Riemann sum is a finite calculation, we can use Riemann sums to approximate integrals which we may not be able to do not want to evaluate exactly. In general, the more terms we take in a Riemann sum, the more accurate our estimate will be, although this is not always true.

Example. Consider the function $f(x, y) = x^2+y^2$ over the rectangle $R = [0, 4]\times[0, 4]$. Use a Riemann sum to estimate the integral of $f(x, y)$ over R, where the Riemann sum splits R up into four congruent squares. Use the bottom-left points in a square to obtain one estimate, and the midpoints to obtain another.

Splitting up R into four congruent squares gives the squares $[0,2]\times[0,2]$, $[0,2]\times[2,4]$, $[2,4]\times[0,2]$, $[2,4]\times[2,4]$. The area of each square is 4, so using the left endpoints gives:

$$4\cdot(0+4+4+8)=64.$$

Using the midpoints, which are $(1,1),(1,3),(3,1),(3,3)$ gives

$$4\cdot(2+10+10+18)=160.$$

By way of comparison, the exact answer, which we will shortly see how to compute, is $512/3$. The relatively poor quality of our estimates is due to the fact that we only used four rectangles, and the function $f(x,y)$ is far from approximately constant on each of those rectangles. And of course, in practice, we would probably not need to use a Riemann sum to estimate an integral which we could evaluate exactly.

What we seek now is a convenient way to evaluate double integrals. Even though we should not expect a 'fundamental theorem of Calculus' for double integrals yet, we can still make use of the FTC for functions of a single variable.

2. Double integrals

Recall that we discussed how a definite integral for a function of two variables $f(x,y)$ over a rectangle R is defined, as the limit of Riemann sums which represented sums of volumes of lots of small rectangular prisms. In particular, the definite integral should represent the (signed) volume under $z=f(x,y)$ over the rectangle R.

Examples.

- Evaluate the double integral
$$\iint\limits_R c\,dA$$
in terms of the area of the rectangle R, where c is some constant. This double integral is equal to the volume of a rectangular prism with base of area $A(R)$ and height c, so this double integral should equal $cA(R)$. This is the two-dimensional analogue of the formula
$$\int_a^b c\,dx=c(b-a).$$
- Using the interpretation of a double integral as volume, calculate the definite integral
$$\iint\limits_R \sqrt{1-x^2}\,dA$$
where R is the rectangle $[-1,1]\times[-1,1]$. We begin by making a sketch of the graph of this function over the rectangle R. Notice that $\sqrt{1-x^2}$ does not depend on y, so the cross-sections of the graph of this function taken when we fix y will all look identical. In particular, these cross-sections are of the form $z=\sqrt{1-x^2}$, which is the upper half of a circle of radius 1. So the double integral we are evaluating represents the volume of a region whose

cross-sections by planes of the form $y = C$ are all half-discs of radius 1. In other words, our region is half of a cylinder, whose base is a circle of radius 1 and whose length is 2. The volume of such a solid is given by

$$\frac{1}{2}\pi \cdot 2 = \pi.$$

The double integral also has an interpretation as the average value of a function over the rectangle R. In the single-variable case, the expression

$$\frac{1}{b-a}\int_a^b f(x)\,dx$$

represents the average value of $f(x)$ on the interval $[a, b]$; in the two-variable case, the expression

$$\frac{1}{A(R)}\iint\limits_R f(x, y)\,dA$$

represents the average value of $f(x, y)$ on the rectangle R. For example, if R represents the boundaries of a city, and $f(x, y)$ is the amount of rainfall (that is, the number of inches of rain) the city received at the point (x, y), then this average value represents the average amount of rainfall the city received.

3. Iterated integrals and Fubini's Theorem

The previous example (the half-cylinder) provides the clue to answering the question of how to evaluate double integrals. Notice that we solved that problem by recognizing a cross-section of the solid we wanted to calculate the volume of when we cut it with planes of the form $y = C$. Suppose we want to integrate

$$\iint\limits_R f(x, y)\,dA$$

We might approach this problem by taking cross-sections of the solid under $z = f(x, y)$, over the rectangle R, when we fix either x or y. Concretely, suppose $R = [a, b] \times [c, d]$. Suppose we fix x, so we take cross-sections with planes of the form $x = C$. Then the appropriate cross-section then has area given by the formula

$$A(x) = \int_c^d f(x, y)\,dy$$

where in this integral we treat x as a fixed number. For example, if we are interested in evaluating

$$\iint\limits_R x^2 + y^2\,dA$$

where $R = [0, 2] \times [0, 3]$, then

$$A(x) = \int_0^3 x^2 + y^2\,dy = \left(yx^2 + \frac{y^3}{3}\right)\Big|_{y=0}^{y=3} = 3x^2 + 9.$$

In other words, the cross-sectional area of the solid defined by $z = x^2 + y^2$ over the rectangle $[0, 2] \times [0, 3]$ when we fix x is $3x^2 + 9$. We place $y = 0, y = 3$ in the bounds of integration to remind ourselves that we are evaluating the expression $yx^2 + y^3/3$

at the values $y = 0, 3$ instead of $x = 0, 3$. This procedure, where we calculate an integral by fixing one variable, is very strongly reminiscent of partial differentiation; sometimes this is called *partial integration*.

Therefore, it is plausible that the volume of the solid we are interested in can be obtained by integrating the cross-sectional area $A(x)$ with respect to the remaining variable x. That is, we should expect

$$\iint\limits_R f(x,y)\,dA = \int_a^b A(x)\,dx = \int_a^b \left(\int_c^d f(x,y)\,dy \right) dx.$$

From now on, we will drop the parentheses which enclose the inner integral. We call such an expression for a double integral, where we take an integral with respect to one variable first, and then with respect to the remaining variable, an *iterated integral*. In the example we were looking at, then, we should have

$$\iint\limits_R x^2 + y^2\,dA = \int_0^2 3x^2 + 9\,dx = x^3 + 9x\Big|_0^2 = 8 + 18 = 26.$$

There is nothing special about taking cross-sections by fixing x. We could just as well have taken cross-sections by fixing y, and obtained a formula for the cross-sectional area at y:

$$A(y) = \int_a^b f(x,y)\,dx$$

and then integrated this formula over y:

$$\int_c^d A(y)\,dy = \int_c^d \int_a^b f(x,y)\,dx\,dy.$$

One probably expects that the order of integration should not change the final answer. While this is not always true (see Problem 37 in Chapter 16.2), the following theorem tells us that this is true in virtually every situation we will encounter:

Theorem. (Fubini's Theorem) If $f(x,y)$ is continuous on a rectangle $R = [a,b] \times [c,d]$, then

$$\iint\limits_R f(x,y)\,dA = \int_a^b \int_c^d f(x,y)\,dy\,dx = \int_c^d \int_a^b f(x,y)\,dx\,dy$$

More generally, this is true whenever $f(x,y)$ is bounded on R, continuous on R except possibly at a finite number of smooth curves, and both iterated integrals exist.

You can check that the hypotheses of Fubini's Theorem hold true for basically every example we will look at, and if we encounter any situations where Fubini's Theorem does not hold (which is rather unlikely) this will be explicitly noted.

Notice that once we write an iterated integral down, the ordering of $dx\,dy$ or $dy\,dx$ determines the order in which we should perform the partial integrations. Pay very close attention to the ordering of the differentials, and remember that you begin by integrating with respect to the left-most variable and then work your way to the right. (While we only have two differentials right now, we will be performing triple integrals in the future.)

Examples.

- Let $R = [0,2] \times [0,1]$. Evaluate the double integral
$$\iint\limits_R xe^y \, dA$$
using both possible orders of integration. Since Fubini's Theorem is obviously valid, we can write this double integral as either of the iterated integrals
$$\int_0^2 \int_0^1 xe^y dy \, dx, \int_0^1 \int_0^2 xe^y dx \, dy.$$
The former integral is equal to
$$\int_0^2 xe^y \Big|_{y=0}^{y=1} dx = \int_0^2 x(e-1) \, dx = \frac{(e-1)x^2}{2} \Big|_0^2 = 2(e-1).$$
The latter integral is equal to
$$\int_0^1 e^y \frac{x^2}{2} \Big|_{x=0}^{x=2} dy = \int_0^1 2e^y \, dy = 2e^y \Big|_0^1 = 2(e-1).$$
As expected, these two iterated integrals are equal to each other.
- Sometimes it is easier to integrate with respect to one variable first instead of the other variable. For example, let $R = [0, \pi] \times [0,1]$, and evaluate the double integral
$$\iint\limits_R x\cos(xy) \, dA.$$
Which variable is it easier to integrate with respect to first? If we want to integrate with respect to x, we will need to perform an integration by parts. However, if we integrate with respect to y, we need only use a quick u-substitution, $u = xy$. Then $du = x\,dy$, and we get
$$\iint\limits_R x\cos(xy)\, dA = \int_0^\pi \int_0^1 x\cos(xy)\, dy\, dx = \int_0^\pi \left(\sin(xy) \Big|_{y=0}^{y=1} \right) dx = \int_0^\pi \sin x \, dx = -\cos x \Big|_0^\pi = 2.$$
While in principle it does not matter which variable you integrate with respect to first, in practice it can be computationally easier to integrate with respect to one variable first instead of using the other variable.
- Of course, there is nothing special about the variables x, y. For example, we can evaluate an iterated integral
$$\int_0^2 \int_{-1}^1 (u+v)^2 \, dv \, du.$$
The first integration gives
$$\int_0^2 \left(\frac{(u+v)^3}{3} \Big|_{v=-1}^{v=1} \right) du = \int_0^2 \frac{(u+1)^3}{3} - \frac{(u-1)^3}{3} \, du.$$
This integral is equal to
$$\frac{(u+1)^4}{12} - \frac{(u+1)^4}{12} \Big|_0^2 = \frac{3^4}{12} - \frac{1^4}{12} - \left(\frac{1^4}{12} - \frac{(-1)^4}{12} \right) = \frac{80}{12} = \frac{20}{3}$$

If you want, you can check that this is equal to the answer you would have found had you integrated with respect to u first instead of v.