Part IV

Multivariable calculus

Before we tackle the very large subject of calculus of functions of several variables, you should know the applications that motivate this topic. Here is a list of some key applications.

1. Totals of quantities spread out over an area.
2. Probabilities of more than one random variable: what is the probability that a pair of random variables (X, Y) is in a certain set of possible values?
3. Marginal cost.
4. Optimization: if I have a limit on how much I can spend on production and advertising in total, and my profit will be some function $f(p, a)$, then how much should I invest in production and how much in advertising?

When dealing with these sorts of questions, the functions and their notation can start to seem difficult and abstract. Geometric understanding of multi-variable functions will help us think straight when doing word problems and algebraic manipulations.

10 Multivariable functions and integrals

10.1 Plots: surface, contour, intensity

To understand functions of several variables, start by recalling the ways in which you understand a function f of one variable.

(i) As a rule, e.g., "double and add 1"

(ii) As an equation, e.g., $f(x) = 2x + 1$

(iii) As a table of values, e.g.,

x	0	1	2	5	20	-95	π
$f(x)$	1	3	5	11	41	-189	$2\pi + 1$

(iv) As a graph, e.g.,

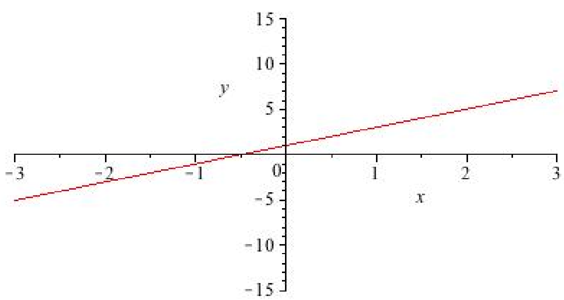

Similarly, a function f of two variables is a way of associating to any pair of values for x and y (two real numbers) a real number $f(x, y)$. The same options apply for understanding f.

(i) We can give the rule if it is easily stated, e.g., "multiply the two inputs."

(ii) We could give an equation, such as $f(x, y) = xy$.

(iii) We could make a table, e.g.,

x	1	1	1	2	2
y	0	1	5	0	π
$f(x, y)$	0	1	5	0	2π

(iv) One might graph f.

You can think of a function of two variables as having two real inputs x and y or as having one input that is a pair (x, y). The second way makes the domain of the function into (some subset of) the xy-plane. For more on how to figure out exactly what subset forms the domain, look at the first few pages of Section 14.1 of the textbook. We won't focus on that, but we will use geometry to understand f via its various visual depictions. There most common way to make a graph of f is to plot the three-dimensional surface $z = f(x, y)$ as in the following figure.

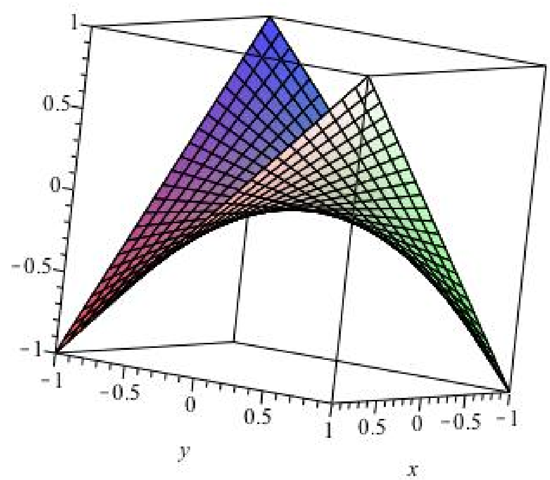

Another way is to plot the **level curves**. To do this, you have to figure out which points (x, y) share the same f-value, say zero, and draw a curve indicating that set. Then, draw the curve indicating another nearby value such as $1/2, 1, -1$, etc. This is shown on the right of the figure above. The book explains this on page 797. The convention when drawing level curves is to pick some fixed increment, such as every $1/2$ or every 100, and draw the level curves corresponding to these regular intervals.

The US Geological Service produces a series of maps drawn this way. These are contour plots of $f(x, y)$, where f is the elevation and x and y are distance east and distance north of the center of a quadrangle.

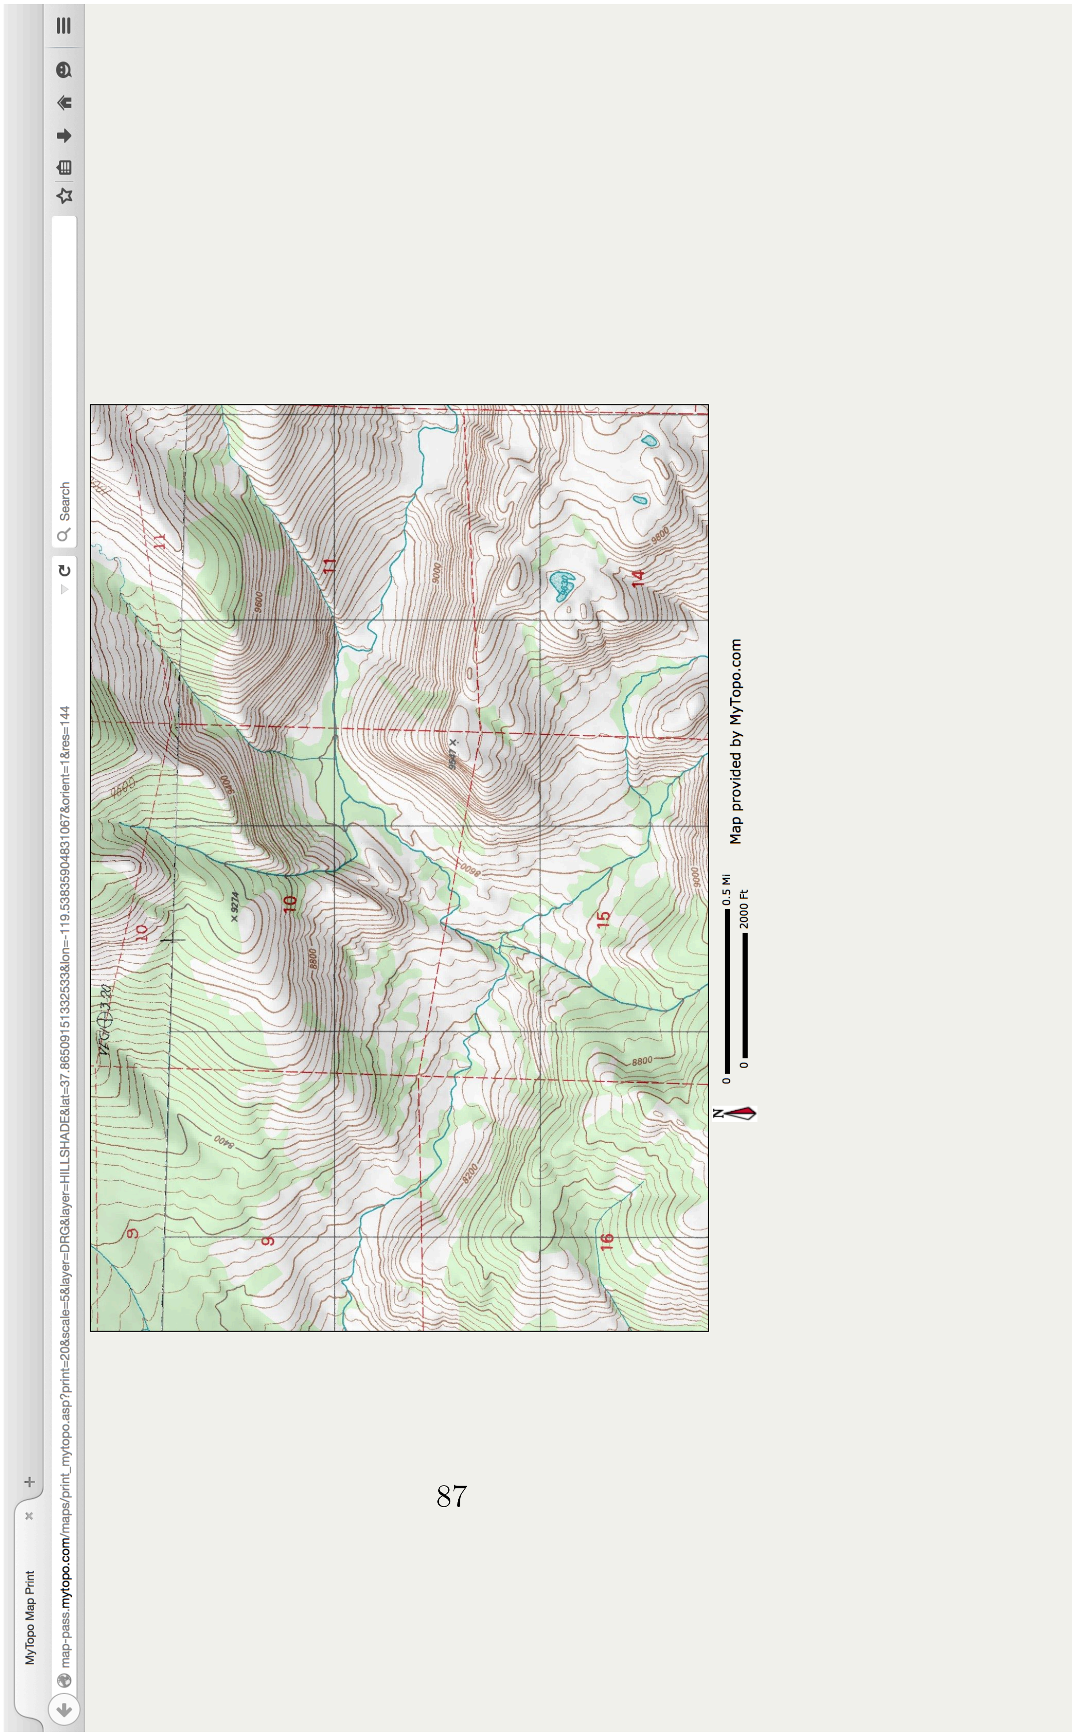
MyTopo Map Print
map-pass.mytopo.com/maps/print_mytopo.asp?print=20&scale=5&layer=DRG&layer=HILLSHADE&lat=37.8650915133253 3&lon=-119.5383590483106 7&orient=1&res=144
Search
Map provided by MyTopo.com
0.5 Mi
2000 Ft
N

The elevation example is very important even if you don't care about hiking. This is because the traditional way to plot f is to plot the surface $z = f(x, y)$, which means that our brains are primed to accept $f(x, y)$ as an elevation at the point with coordinates (x, y). However, this is far from the only use of contour plotting. The most important application of this is when $f(x, y)$ is profit or some other kind of a utility function (e.g., the level of satisfaction when you have x dollars in the bank and a car that costs y dollars). The contour plot of f shows the **indifference curves**. Later we can use this interpretation of contour plots along with some calculus to compute optimal allocations. The next figure shows the contour plot for $f(x, y) = xy$ along with the height plot $z = f(x, y)$ that you already saw for this function.

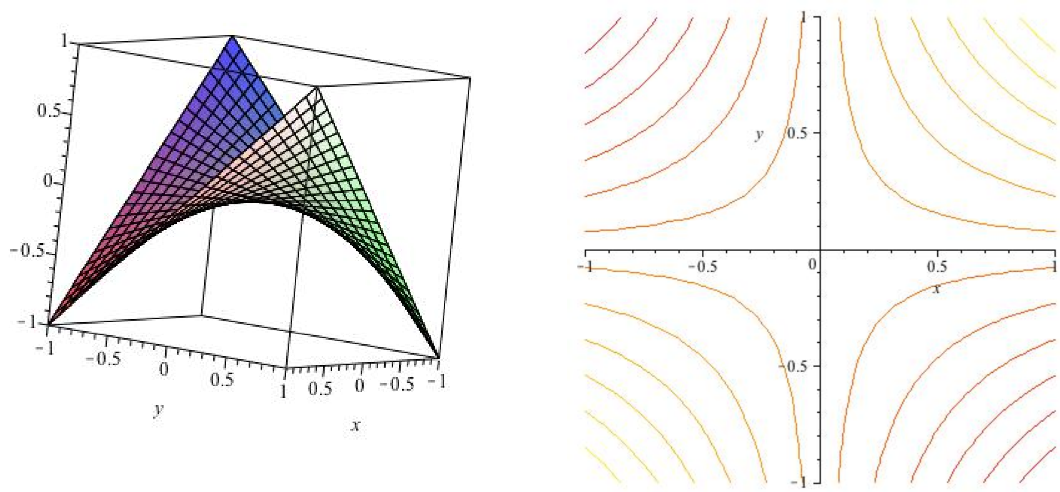

All we are doing in this first section is getting used to functions of more than one variable and their visual depictions. We're almost done, except that we haven't talked about functions of three or more variables. We don't have four dimensions handy, so we can't graph $z = f(x_1, x_2, x_3)$. We can still think of f as a function mapping points in an abstract n-dimensional space to the real numbers, and in the case of exactly three variables, we can make a contour plot which now has contour surfaces in three dimensions; see Figure 14.8 in the book. For now, it suffices to practice going back and forth between the equation for a function of two variables and its visual representations.

10.2 Multivariate integration: rectangular regions

This section is a bit heftier than the previous one because multiple integrals are really, really important. This is a tricky topic for two reasons. First, students often confuse the *definition* of a double integral with the *computation* of a double integral. I will try to help you keep these straight. Secondly, non-rectangular regions of integration (which are the topic of Section 10.2) require deeper understanding of free and bound variables than you have needed for the calculus you've done so far. Please come to class having read Section 15.1 of the textbook!

(i) **Meaning**

Let R be a region in the xy-plane and let $f(x, y)$ be a function. The notation $\int_R f(x, y)\, dA$ is read as "the double integral of f over the region R" and defined as follows (I am parphrasing what is on page 883 of your textbook).

> Divide R into small rectangular regions (ignore for now the fact that these don't quite cover R or sometimes extend a litle beyond R). Multiply the are of each rectangle by f evaluated at some point in that rectange, and add up all of these products. The integral is defined to be the limit of this sum of products as the rectangles get small.

What does this compute? In general it computes the total amount of stuff when f is a density of stuff per unit area. For example, suppose the density of iron ore over a patch of ground is a function $f(x, y)$ that varies due to proximity to some pre-historic lava flow. Then $\int_R f(x, y)\, dA$ will be the total amount of iron ore in the region R. Do you see why? You can get the total by adding up the amount in regions small enough that f doesn't vary significantly; then the amount of ore in the region is roughly the area times f evaluated at any point in the region, so we should expect that adding up these products approximates the total; in the limit, it *is* the total.

Time for a bunch of conceptual remarks!

1. Notice there is now a quantity dA rather than dx or dy. This means, literally, "the teeny amount of area". Starting now, it will be very important to keep track of the infinitesimals.

2. The units of $\int_R f(x,y)\,dA$ are units of f times units of A. The units of A can be area, but more generally, they are whatever unit x represents times whatever unit y represents.

3. Try to see how this is analogous to integrals in one variable. In each case you break up the (interval / region), then in each small part you evaluate f somewhere, use this as a proxy for f throughout the small part, multiply by the (length / area) of the small part, sum and take the limit.

4. You can integrate in three variables. Just chop a 3-D region into subregions, sum their volumes times the value of $f(x,y,z$ somewhere in the region, and take a limit. In fact, you can do this in any number of variables even though we can't visualize space in dimensions higher than three. In Math 110, we'll stick to two variables.

Here are some more meanings for a double integral.

Volume. If $f(x,y)$ is the height of a surface at the point (x,y), then $\int f(x,y)\,dA$ gives the volume underneath the surface but above the xy-plane. That's because the summands (namely the area of a little region times $f(x,y)$ evaluated at a point in the region) is the volume of a tall skinny rectangular shard, many of which together physically approximate the region. If you can't picture this, you have to have a look at Figure 15.3. Notice the units work: f is height (units of length) and $\int_R f(x,y)\,dA$ is volume, which does indeed have units of length times area.

Area. A special case is when $f(x,y)$ is the constant function 1. Who would have thought that integrating 1 could be at all important? But it is. If you build a surface of height 1 over a region R, then the volume of each shard is the area at the base of the shard and the integral is just the limitin sum of these, namely the total area. Notice the units work: in the example f is unitless, and $\int_R f(x,y)\,dA$ is the area of R, which has units of area.

Averages. By definition, the average of a varying quantity $f(x,y)$ over a region R is the total of f divided by the area of the region:

$$\text{Average of } f \text{ over } R = \frac{\int_R f(x,y)\,dA}{\text{Area of } R}\,.$$

Probability. This application will get its own treatment in Section 10.4.

(ii) **Computing the iterated integral: rectangular regions**

Remember how it worked when you learned integration in one variable? It was defined as the limit of Riemann sums, which intuitively captures the notion of area under a curve. Then there's a theorem saying you can figure out the value of the integral over an interval by computing an antiderivative and subtracting its values at the two endpoints. Similarly, we have already defined the integral conceptually, now we need to say something about using calculus to compute it. A lucky fact: we don't need anything as difficult as the Fundamental Theorem of Calculus like we did for one variable integrals. That's because we assume you already know how to compute single variable integrals and that can be harnesses to compute the double integral. Remember, for now we're sticking to to the case where R is a rectangle.

As the textbook does, we start by assuming R is a rectangle $a \le x \le b$ and $c \le y \le d$, chopped up so that each little square has width Δx and length Δy. We then add up the little bits in an organized way. First add all the tall skinny rectangles over a given x interval as y varies. In the volume interpretation this gives the volume of the slice of the solid that has width Δx. There is a slice for each x-value in the grid.

Here's the thing. If you fix a value $x = M$, then **you're just computing Δx times the area under the one-variable function** $f(M, y)$. You know how to do that: you integrate $\int_c^d f(M, y) dy$ and multiply by Δx. This integral of course depends on M. Call it $g(M)$. Summing all the slice volumes is the same as integrating $g(M)$ from a to b. We don't have to use the variable M, we can just call it x. So the answer is:

$$\int_R f(x, y)\, dA = \int_a^b g(x)\, dx, \quad \text{where } \; g(M) = \int_c^d f(M, y)\, dy\,.$$

This is Fubini's Theorem (first form) on page 885 which you practiced computing in the MML problems from Section 15.1. I prefer to put parentheses into the equation given in the book:

$$\int_R f(x, y)\, dA = \int_a^b \left[\int_c^d f(x, y)\, dy \right] dx = \int_c^d \left[\int_a^b f(x, y)\, dx \right] dy\,. \tag{10.1}$$

At this point it would be a good idea to read Examples 1 and 2 in Section 15.1. Also, you should pay attention to free and bound variables. In the so-called inner integral $\int_c^d f(x, y)\, dy$, the variable y is bound, but x is free. In other words, this integral represents a quantity that depends on x (but not y). That's why we can integrate it against dx in the outer integral, to finally get a number.

In Example 1 of Chapter 15 of the textbook, they do the integral two ways (x-direction first versus y-direction first) to show that you get the same answer (that's part of Fubini's Theorem). Sometimes you need to use this to evaluate an integral that appears difficult: write it in the other order and see if it is easier; one of the homework problems is on this technique.

Magic product formula

Suppose your function $f(x,y)$ is of the form $g(x) \cdot h(y)$ and your region of integration is a rectangle $[a,b] \times [c,d]$. Then

$$\int_R f(x,y)\, dA = \left(\int_a^b g(x)\, dx \right) \times \left(\int_c^d h(y)\, dy \right) .$$

Can you see why? It's due to the distributive law. The Riemann sum for the double integral actually factors into the product of two Riemann sums. I'll do this on the blackboard for you because, when written without narration, it just looks like a mess.

One parting word: circling back to the issue of distinguishing the definition from the computation, the left-hand side of (10.1) refers to the definition – a limit of Riemann sums; the two expressions after the equal signs are single variable integrals, computable as antiderivatives. The theorem is asserting that they are all equal.

10.3 Multivariate integration: general regions

The trickiest thing about learning double integration is when R is not a rectangle. Then, when you cut into slices, the limits of integration will change with each slice. That's OK as long as you can write them as a function of the variable you are not integrating and evaluate properly. There are four examples in the book (Section 15.2), plus I'll give you one more here. But before diving into these, we should review how to write sets of points in the plane.

Writing sets of points in the plane

The notation $\{(x, y) : \text{ blah blah blah}\}$ denotes the set of points in the plane satisfying the condition I have called "blah blah blah". For example, $\{(x, y) : x^2 + y^2 \leq 1\}$ is the unit disk. You will need to become an expert at writing sets of points in a very specific manner: the set of points where x is in some interval $[a, b]$ and y lies between two functions of x, call them g and h. It looks like

$$\{(x, y) : a \leq x \leq b,\ g(x) \leq y \leq h(x)\}.$$

EXAMPLE: can you write the unit disk in this format? For a and b you need the least and greatest x values that appear anywhere in the region. For the unit disk, that's -1 and 1. Then, for each x, you need to figure out the least and greatest y values that can be associated with that x. For the unit disk, the least value is $-\sqrt{1 - x^2}$ and the greatest is $+\sqrt{1 - x^2}$.

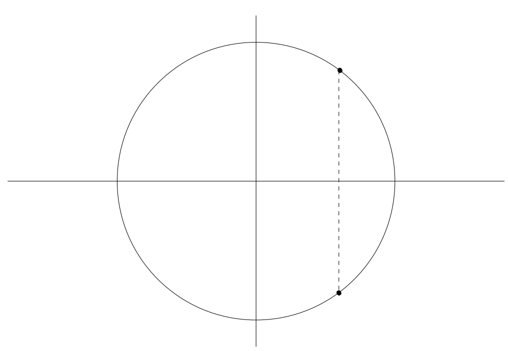

The y-value goes from $-\sqrt{1 - x^2}$ to $+\sqrt{1 - x^2}$

So in the end, the unit disk $\{(x, y) : x^2 + y^2 \leq 1\}$ can be written in our standard form as

$$\{(x, y) : -1 \leq x \leq 1,\ -\sqrt{1 - x^2} \leq y \leq \sqrt{1 - x^2}\}\,.$$

This way of writing it naturally breaks the unit disk into vertical strips where x is held constant and y varies from some least to some greatest value depending on x. I should have said this is "a standard form" not "the standard form" because it is equally useful to break into horizontal strips. These correspond to the format

$$\{(x, y) : c \leq y \leq d,\ g(y) \leq x \leq h(y)\}$$

where for each fixed y, the x values range from some minimum to some maximum value depending on y. You will be practicing a lot with these two formats!

Limits of integration for non-rectangular regions

What I am explaining here is Theorem 2 on page 889 of the textbook. When computing $\int_R f(x, y)\, dA$, if you can write R as a region in the form above.

There are three steps. First, specify the region of integration in terms of varying limits of integration. Second, use these as limits of integration. If x goes from z to b while y goes from $g(x)$ to $h(x)$ then the integral will look like $\int_z^b \int_{g(x)}^{h(x)} f(x, y)\, dy\, dx$. Third, carry out the integration with these limits.

EXAMPLE: Let R be the unit disk and let $f(x, y) = 1$. The possible x-values in R range from -1 to 1. So we put this on the outer integral: $\int_{-1}^1 [\,\cdots\,]$. Now fix a value of x and figure out what the limits are on y. As we have seen, y goes from $-\sqrt{1 - x^2}$ to $\sqrt{1 - x^2}$. So now we can write the whole integral as

$$\int_{-1}^1 \left[\int_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} 1\, dy \right] dx\,.$$

When we do the inner integral we get the antiderivatite y, which we evaluate at the upper and lower limits: $y\ |_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} = 2\sqrt{1 - x^2}$. Finally, we evaluate the outer integral, obtaining $\int_{-1}^1 2\sqrt{1 - x^2}\, dx$. This is a tough integral if you do it honestly: integrating by parts and using #18 in the integral table will give you

$$\left(x\sqrt{1 - x^2} + \arcsin x \right)\Big|_{-1}^{1}\,.$$

The value of $x\sqrt{1-x^2}$ is zero at both endpoints, so this evaluates to $\arcsin(1) - \arcsin(-1) = \pi/2 - (-\pi/2) = \pi$.

Here are an FAQ about what we just did.

1. When we took the anti-derivative of the constant function 1 why did we get y and not x? Answer: we were integrating in the y-variable at that time.

2. How can you know whether the limits on the inner integral will be functions of y or functions of x? Answer: if you choose vertical strips then the inner integral is dy, the outer integral is dx and the limits on the inner integral can be functions of x but not y.

3. Is it a coincidence that after a complicated computation, the integral came out to be a very simple expression? Answer: No! It's because the integral of 1 over a region gives the area, and the area of a circle is a very simple expression. In fact, if you were asked to do this integral on a test or homework, you should probably not do any calculation and just say it's the area of a circular region with radius 1 and is therefore equal to π.

Switching the order of integration

You have seen how to take a region R and write it in either standard form: horizontal or vertical strips. Sometimes, in order to make an integral do-able, you will want to switch from horizontal strips to vertical strips or vice versa. Starting with one standard form, you convert to a geometric region R, then write that in the other standard form. This allows you to switch between an iterated integral with x in the inside and one with y on the inside.

EXAMPLE: Compute $\displaystyle\int_0^1 \int_y^1 \frac{\sin x}{x}\,dx\,dy$. Unfortunately you can't integrate $\sin x/x$. But wait! The region $\{0 \le y \le 1,\ y \le x \le 1\}$ is triangular and can also be written in vertical strips: $\{0 \le x \le 1,\ 0 \le y \le x\}$. The integral is therefore equal to $\displaystyle\int_0^1 \int_0^x \frac{\sin x}{x}\,dy\,dx$. We can now see that this is equal to

$$\int_0^1 \left(y\frac{\sin x}{x} \right)\Bigg|_0^x dx = \int_0^1 \sin x\,dx = 1 - \cos(1)\,.$$

10.4 Applications: spatial totals, averages, probabilities

No new math in this section, just some applications. Two of them are pretty straightforward: integrals to yield total amounts and integrals to compute averages. The third, probability densities in two variables, will involve a couple of new concepts.

Integrals to compute totals

This is essentially just a reminder that the integral of stuff per unit area over an area yields total stuff.

EXAMPLE: The population density east of a river running north-south is $f(x, y) = 6000e^{-x^2}$ people per square mile. The county is divided into quadrants as shown in the figure. Roughly how many people are there in the east quadrant?

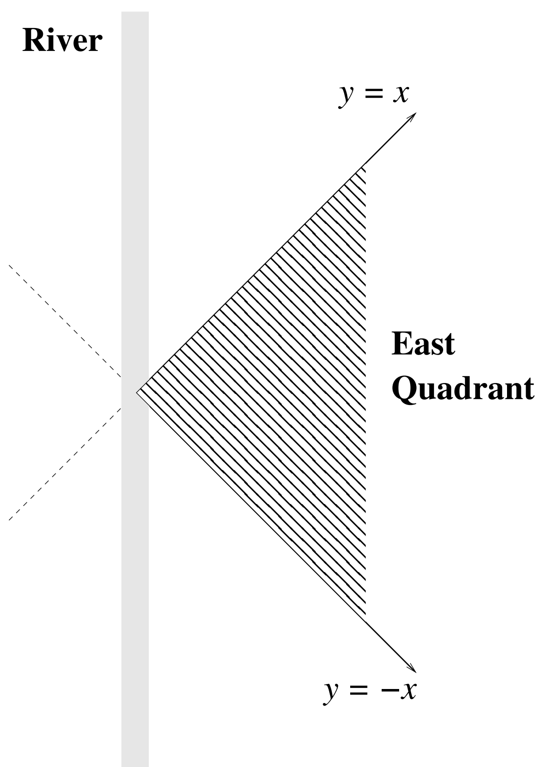

SOLUTION: Make coordinates in which the quadrant is the region represented in standard form by the set

$$\{(x, y) : 0 \leq x < \infty, -x \leq y \leq x\}\,.$$

We're not going to re-do the theory of improper integrals in two variables, we'll only deal with cases where you can just plug in ∞ and get the right answer. The region is in standard form, so the total population is given by

$$\int_0^\infty \int_{-x}^{x} 6000e^{-x^2}\, dy\, dx\,.$$

The inner integral might look tough but it's not (look carefully at which is the variable of integration):

$$\int_{-x}^{x} 6000\, e^{-x^2}]\, dy = 6000\, ye^{-x^2}\Big|_{-x}^{x} = 12000\, xe^{-x^2}\,.$$

The outer integral can then be done by the substitution $u = e^{-x^2}$, leading to

$$\int_0^\infty 12000\, xe^{-x^2}\, dx = -6000\, e^{-x^2}\Big|_0^\infty = 0 - (-6000) = 6000\,.$$

This is a good example of an integral which is not too hard one way but impossible the other. Try to integrate e^{-x^2} against dx rather than dy and you will be stuck at the first step! If you come across this, you will always want to switch the order of the integrals.

Averages

The average of a quantity over a region is just the total of the quantity divided by the size of the region.

EXAMPLE: What is the average of e^x over the triangular region where $0 \le x \le y \le 1$?

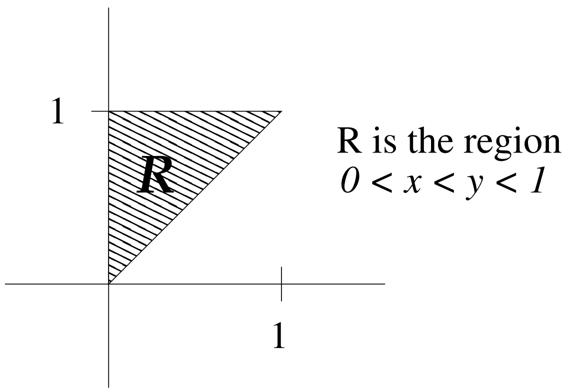

SOLUTION:

$$\int_0^1 \int_0^y e^x \, dx \, dy = \int_0^1 \left(e^x \big|_0^y \right) \, dy = \int_0^1 \left(e^y - 1 \right) \, dy = e - 2 \, .$$

The average value is the total, $e - 2$, divided by the area. The area is $1/2$ therefore the average value is $2(e - 2)$.

EXAMPLE: The cost of providing fiber optic service to a resident is proportional to the distance to the nearest hub, with constant of proportionality 5 dollars per meter. If a township is a square, two kilometers on a side, and there is a single hub in the center, what is the average of the service cost over this area?

SOLUTION: Make coordinates with the hub in the center. The township is the square $[-1000, 1000] \times [-1000, 1000]$, with x and y representing East-West displacement and North-South displacement in meters. The cost of providing service to the point (x, y) is given by $f(x, y) = 5\sqrt{x^2 + y^2}$. The average is therefore given by

$$\text{Ave.} = \frac{1}{1000^2} \int_R 5\sqrt{x^2 + y^2} \, dA = \frac{1}{1000^2} \int_{-1000}^{1000} \int_{-1000}^{1000} 5\sqrt{x^2 + y^2} \, dy \, dx \, .$$

If you can do this integral, you are smarter than I am. I tried it numerically with a 5×5 grid (I used midpoints and I used symmetry to restrict to the quadrant $[0, 1000] \times [0, 1000]$ in order to make my grid squares smaller) and got roughly \$3812 which is pretty close to what my computer tells me is the correct numeric value of \$3826.

Two-variable probability densities

It is often useful to consider a random pair of real numbers, that is, a random point in the plane. A probability density on the plane[4] is a nonnegative function $f(x, y)$ such that $\int f(x, y) dA = 1$. As before, the mean of the X variable is $\int x f(x, y)\, dA$ and the mean of the Y variable is $\int y f(x, y)\, dA$. Here are a couple of special cases.

EXAMPLE: UNIFORM DENSITY ON A REGION. Let R be a finite region and let $f(x, y) = C$ on R and zero elsewhere. For this to be a probaiblity denstiy, the normalizing constant C must be the reciprocal of the area of R (that's because the integral of $1\, dA$ over R is just the area of R). For example, if R is the interior of the unit circle then C would be $1/\pi$. If R is the recangle $[a, b] \times [c, d]$ then $C = 1/((b - a)(d - c))$.

EXAMPLE: PLANAR STANDARD NORMAL DISTRIBUTION. Let $f(x, y) = \frac{1}{2\pi} e^{-(x^2+y^2)/2}$. This has integral equal to 1 because it is the product of $(1/\sqrt{2\pi})e^{x^2/2}$ and $(1/\sqrt{2\pi})e^{y^2/2}$, which we already know integrate to 1 over the whole plane $(-\infty, \infty) \times (-\infty, \infty)$ because each one is just the one-variable standard normal density. This uses the "magic product formula".

A two-variable probability density corresponds to picking simultaneously two numbers X and Y such that the probability of finding the pair (X, Y) in some region A is equal to the integral of the density over the region A.

EXAMPLE: A probability density on the rectangle $[0, 3] \times [0, 2]$ is given by Ce^{-x}. What is C, and what is the probability of finding the pair (X, Y) in the unit square $[0, 1] \times [0, 1]$?

[4] The integral, if it is over the whole plane, is technically an improper integral, but we won't worry about that; in all our examples either the density will be nonzero on just a finite region or it will be obvious that there is a limit as the region becomes infinite.

Solution: We are integrating over a rectangle and Ce^{-x} is a product of $g(x) = Ce^{-x}$ and $h(y) = 1$. By the magic product formula, the integral is

$$\left(\int_0^3 Ce^{-x}\, dx \right) \times \left(\int_0^2 dy \right) = 2C \cdot \left(1 - e^{-3}\right) .$$

Therefore $C = \dfrac{1}{2(1 - e^{-3})}$ which is just a shade over 1/2. Now using the product formula again to integrate over the unit square gives a probability of

$$\int_{[0,1]\times[0,1]} Ce^{-x}\, dA = C(1 - e^{-1}) = \frac{1 - e^{-1}}{2(1 - e^{-3})} \approx 0.3326.$$