Contents

1 Finding Local Extrema: Second Derivative Test

Given a function $f(x, y)$ of two variables, to identify local maxima and minima, one uses the following procedure:

> Identify all the critical points of f. Then at each critical point, apply the second derivative test to determine whether it's a local maximum, minimum, or saddle point.

In more detail, the method goes as follows.

1. **To identify critical points**, solve the equation $\nabla f = \mathbf{0}$. Note that this is a vector equation, so it is actually a system of two scalar equations:
$$\frac{\partial f}{\partial x} = 0, \quad \frac{\partial f}{\partial y} = 0.$$
2. **To apply the second derivative test**, for every solution $(x, y) = (a, b)$ found, compute the Hessian determinant
$$D = \det \begin{bmatrix} f_{xx}(a,b) & f_{xy}(a,b) \\ f_{yx}(a,b) & f_{yy}(a,b) \end{bmatrix} = f_{xx}(a,b)f_{yy}(a,b) - (f_{xy}(a,b))^2.$$
(Note that in the equality above we are assuming the function is "nice enough" for Clairaut's theorem to be applicable—which will be the case in problems we consider.)
   (a) If $D > 0$ then we have a local extremum. Note that we necessarily also have $f_{xx}(a,b)f_{yy}(a,b) > 0$, which means that $f_{xx}(a,b)$ and $f_{yy}(a,b)$ have the same sign—they are either both positive or both negative.
      i. If they are both positive, then f has a local minimum at (a, b).
      ii. If they are both negative, then f has a local maximum at (a, b).

   A way to remember this is to think of the analogous situation in single-variable calculus: at a critical point of a function $f(x)$, if $f'' > 0$ (concave up) we have a minimum, and if $f'' < 0$ (concave down) we have a maximum.

   (b) If $D < 0$ then f has a saddle point at (a, b). No further analysis as above is needed. A saddle point is, by definition, any critical point that is not a local minimum or maximum.

   (c) If $D = 0$ then unfortunately the second derivative test is inconclusive, and ad-hoc analysis of the critical point is needed.

Where to look in Stewart: §14.7

2 Finding Global Extrema: Lagrange Multipliers

Suppose we have a function f, either of two or three variables, that we are trying to maximize or minimize, possibly subject to some constraints. In other words, we are trying to find the global extrema of a function $f(x, y)$ on some region D of $\mathbb{R}^2$, or of a function $f(x, y, z)$ on some region D of $\mathbb{R}^3$.

The method of Lagrange multipliers *assumes that the global maximum or minimum we are trying to find actually exists*. If it does exist, then the method will find it.

Remark. A sufficient condition for existence of global extrema is that the region D be closed and bounded. The latter condition just means that D doesn't extend off to infinity. The former is a bit more subtle, but any region defined by $=$, $\leq$, and/or $\geq$ will typically satisfy it ($<$ and $>$ may be problematic, however).

The concept behind the Lagrange multiplier algorithm is the following:

> If we know that the global maximum or minimum that we are after actually exists, then in order to find it, we can make a list of all possible candidate points in D and then compare the value of f at all those points.

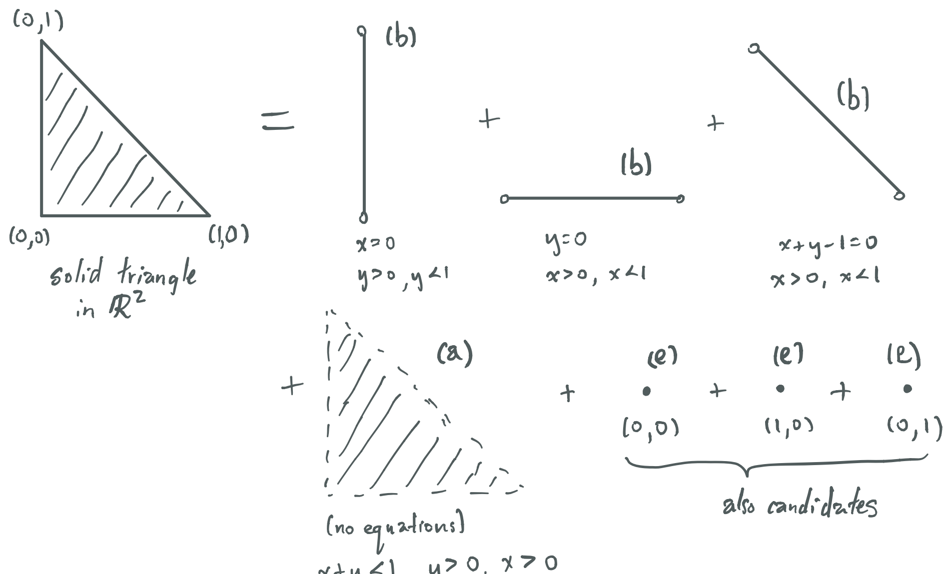

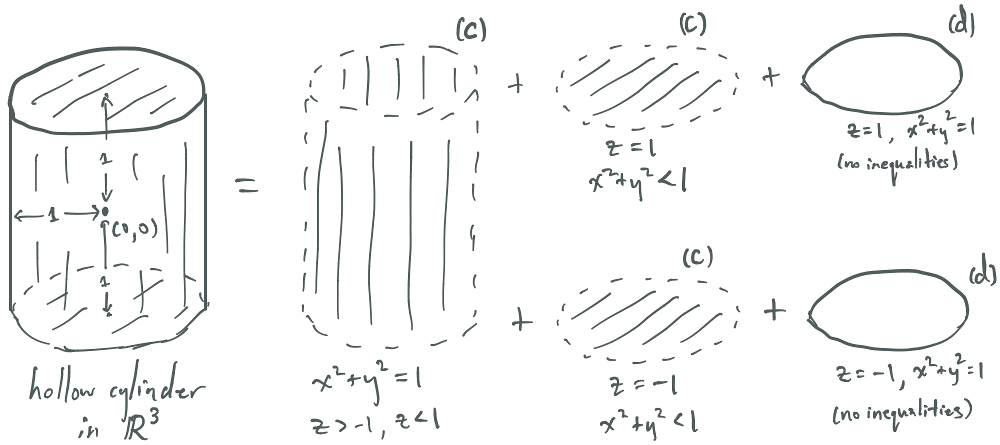

In more detail:

1. **To find the candidates for extrema**, we need to break up our region D into pieces of "homogeneous" dimension, as a different analysis will be needed for each part. A picture of some examples has been provided.

In each case, there is a common algebraic theme. Each component is defined by some combination of *equations* and *(strict) inequalities*. First we look at the *equations*: we will have

$$\text{(dimension of component)} + \text{(number of equations)} = \text{(number of variables)}$$

where the right hand side is 2 if we are in $\mathbb{R}^2$, and 3 if we are in $\mathbb{R}^3$.

Algebraically, Lagrange multipliers introduces sufficiently many new equations and new unknowns so that we end up with an equal number of equations and unknowns:

$$\begin{aligned}&\text{(number of original equations)} + \text{(number of Lagrange multiplier equations)}\\ &= \text{(number of original variables)} + \text{(number of Lagrange multipliers)}\end{aligned} \tag{2.1}$$

Then we have a hope of solving the system. Finally, after we get the solutions to the equations, we make sure that those solutions satisfy the relevant *inequalities*.

Now let's look in detail at how this works for each case.

(a) **A 2D component in $\mathbb{R}^2$ or a 3D component in $\mathbb{R}^3$**. Sometimes the region D will have an *interior*. If our problem is in $\mathbb{R}^2$, this is if the decomposition has a two-dimensional piece. Likewise, if our problem is in $\mathbb{R}^3$, this is if the decomposition has a three-dimensional piece. In the examples sketched above, the triangle has an interior, while the hollow cylinder does not.

The key feature of the interior component (if there is one) is that points in this region are *locally unconstrained*; they have the full 2 or 3 degrees of freedom for movement. And since differential calculus is wholly local, this basically means we get to ignore the constraints entirely.

> For this reason, **the candidates for extrema in the interior are exactly the critical points**. So to find them, solve $\nabla f = \mathbf{0}$.

If we are in $\mathbb{R}^2$, then the situation in regards to (2.1) is $2+0 = 2+0$. If we are in $\mathbb{R}^3$, then it is $3+0 = 3+0$. There are no Lagrange multipliers to be used in this analysis.

Remember to plug in your solutions into any defining inequalities for the interior to make sure that those points are actually in the region of interest!

(b) **A 1D component in $\mathbb{R}^2$**. The geometric idea: at a candidate extrema point, the gradient $\nabla f(x, y)$ ought to be perpendicular to the component (a curve in this case) being analyzed. If the curve is defined by the equation $g(x, y) = 0$, then $\nabla g(x, y)$ is already perpendicular to this curve, so our desired condition is that $\nabla f(x, y)$ and $\nabla g(x, y)$ are parallel.

> If the component is defined by the equation $g(x, y) = 0$ (plus perhaps some inequalities) then candidates for extrema of f on this component are found by solving the system of equations
>
> $$\begin{aligned} g(x,y) &= 0,\\ \frac{\partial f}{\partial x} &= \lambda \frac{\partial g}{\partial x},\\ \frac{\partial f}{\partial y} &= \lambda \frac{\partial g}{\partial y}, \qquad \text{(last two equations can be rewritten as } \nabla f = \lambda \nabla g) \end{aligned}$$

In regards to (2.1), we have $1+2 = 2+1$: we get 2 auxiliary equations and a single Lagrange multiplier λ.

Remember to plug in your solutions into any defining inequalities for this component to make sure that those points are actually in the region of interest!

(c) **A 2D component in** $\mathbb{R}^3$. This is completely analogous to the preceding case. Again, we want to find when the gradient $\nabla f(x,y,z)$ is perpendicular to the component (a surface in this case). If $g(x,y,z) = 0$ defines the surface, then $\nabla g(x,y,z)$ is a normal vector for the surface, and so our desired condition is that $\nabla f = \lambda \nabla g$ for some λ.

> If the component is defined by the equation $g(x,y,z) = 0$ (plus perhaps some inequalities) then candidates for extrema of f on this component are found by solving the system of equations
>
> $$\begin{aligned} g(x,y,z) &= 0, \\ \frac{\partial f}{\partial x} &= \lambda \frac{\partial g}{\partial x}, \\ \frac{\partial f}{\partial y} &= \lambda \frac{\partial g}{\partial y}, \\ \frac{\partial f}{\partial z} &= \lambda \frac{\partial g}{\partial z}, \qquad \text{(last three equations can be rewritten as } \nabla f = \lambda \nabla g\text{)} \end{aligned}$$

In regards to (2.1), we have $1 + 3 = 3 + 1$: we get 3 auxiliary equations and a single Lagrange multiplier λ. Remember to plug in your solutions into any defining inequalities for this component to make sure that those points are actually in the region of interest!

(d) **A 1D component in** $\mathbb{R}^3$. This case is a bit more involved because, in order to define a curve in $\mathbb{R}^3$, one typically needs two equations. The geometric content is once again the same: we want to find when $\nabla f(x,y,z)$ is perpendicular to the curve. But now the collection of vectors perpendicular to the curve at a given point forms a "normal plane" rather than a normal line.

If the curve is defined by two equations $g(x,y,z) = 0$ and $h(x,y,z) = 0$, then ∇g and ∇h are already perpendicular to the curve, so they live in this "normal plane." By a bit of linear algebra, if we know that ∇g and ∇h are *not* parallel, then the condition for ∇f to be in this plane also is that $\nabla f = \lambda \nabla g + \mu \nabla h$ for some scalars λ and μ.

> If the component is defined by the two equations $g(x,y,z) = 0$ and $h(x,y,z) = 0$ (plus perhaps some inequalities) then candidates for the extrema of f on this component are found by solving the system of equations
>
> $$\begin{aligned} g(x,y,z) &= 0, \\ h(x,y,z) &= 0, \\ \frac{\partial f}{\partial x} &= \lambda \frac{\partial g}{\partial x} + \mu \frac{\partial h}{\partial x}, \\ \frac{\partial f}{\partial y} &= \lambda \frac{\partial g}{\partial y} + \mu \frac{\partial h}{\partial y}, \\ \frac{\partial f}{\partial z} &= \lambda \frac{\partial g}{\partial z} + \mu \frac{\partial h}{\partial z}, \qquad \text{(last three equations can be rewritten as } \nabla f = \lambda \nabla g + \mu \nabla h\text{)} \end{aligned}$$

In regards to (2.1), we have $2 + 3 = 3 + 2$: we get 3 auxiliary equations and two Lagrange multipliers λ and μ. That's 5 equations in 5 unknowns, which is kind of terrifying! There is an alternative approach outlined below.

Remember to plug in your solutions into any defining inequalities for this component to make sure that those points are actually in the region of interest!

(d') **A 1D component in** $\mathbb{R}^3$**, alternative approach**. The geometric conclusion was that ∇f should live in the same plane as ∇g and ∇h. But we can translate this to an algebraic condition in a different fashion, using the fact that three vectors are coplanar precisely when their scalar triple product is zero.

If the component is defined by the two equations $g(x, y, z) = 0$ and $h(x, y, z) = 0$ (plus perhaps some inequalities) then an alternative way for identifying candidates for the extrema of f on this component are found by solving the system of equations

$$\begin{aligned} g(x, y, z) &= 0, \\ h(x, y, z) &= 0, \\ \nabla f \cdot (\nabla g \times \nabla h) &= 0. \end{aligned}$$

Note that the third equation is a scalar equation too, so this is truly 3 equations in 3 unknowns.

So rather than having to deal with $2 + 3 = 3 + 2$, you can just deal with $2 + 1 = 3 + 0$ (i.e. getting one auxiliary equation, and not introducing any actual Lagrange multipliers). But beware: the drawback is that the third equation above may be very complicated!

(e) **"Zero dimensional" components—points**. If your decomposition has points, each of those points is a candidate too. See the triangle example.

2. **Compare the value of f at each candidate**. This part is easy; the absolute maximum of f is whatever largest value is attained among the candidates, and analogously for the absolute minimum.

Some potentially helpful remarks:

- For all of these different cases, it is worth remembering that at the end of the day we really don't care about the values of λ and μ. So if you see an opportunity to eliminate them and reduce to a system of equations just about your original variables, you should do so.
- If you know that both a minimum and maximum exist, then that certainly means you should find *at least two* candidates for extrema! If you only find one, check your work. (It is possible in some cases where e.g. you only have a minimum and not a maximum to find only one candidate via Lagrange multipliers. But if the constraint region is closed and bounded, you're guaranteed the existence of both types of extrema, and so you had better find at least two candidates.)

Where to look in Stewart: §14.8

In the preceding discussion about differential calculus and its applications to optimization, the distinction between strict and non-strict inequalities is a hugely important one, because the analysis is very sensitive to such changes. For instance, a function is guaranteed to have extrema on the closed disk $x^2 + y^2 \le 1$, but not necessarily on the open disk $x^2 + y^2 < 1$.

On the other hand, the distinction between $\le$ and $<$ is completely irrelevant for integral calculus, so don't worry about whether you should be using strict or non-strict inequalities for the rest of this write-up.

3 Writing Multiple Integrals

This section covers how to set up the bounds of integration for a problem which asks you to integrate some function $f(x, y, z)$ over a prescribed region E in $\mathbb{R}^3$. For a double integral over a region D in $\mathbb{R}^2$, the process is similar: just start at step 3 in the procedure below.

To figure out the bounds of integration, it is usually easiest to work from the "inside out." A detailed example is given in the §4; it may be a good idea to see how it fits with the procedure described below.

1. **Suppose we decided to integrate with respect to z first. To find the z bounds**, we ask: for fixed values of x and y, what is the allowed range of z? We can answer this problem either geometrically or algebraically.

   - To answer it geometrically, if we fix values for x and y, then the points (x, y, z) form a line in the z direction. Traverse this line upwards (i.e. in the positive z direction) and look at where you enter and exit the region E; the z coordinates of those points will give you the lower and upper z bounds respectively.
   - To answer it algebraically, look at the inequalities defining the region E and see which ones involve z. Isolate z in each, and then take

     $$\max(\text{all expressions less than } z) < z < \min(\text{all expressions greater than } z).$$

     Here the "expressions" should involve x and y only. Note that if there are multiple expressions on either side then you may need to split your integral into multiple parts (although typically you should pick an integration order that avoids this).

2. **Next, we identify the relevant region D in the xy-plane.** Again we can do this either geometrically or algebraically.

   - Geometrically, for what points (x, y) in the xy-plane did the line you drew in the preceding step actually intersect the region E? This collection of points forms the region D. Alternatively: imagine "projecting" E onto the xy-plane—what you see is the region D.
   - To do this algebraically, from the preceding part you found an inequality (lower z bound) $<$ z $<$ (upper z bound). Forget about the z in the middle and just take the inequality (lower z bound) $<$ (upper z bound), which is only about x and y. Take this inequality together with the original inequalities defining E that only involved x and y.

3. **Suppose we decided to integrate with respect to y next. To find the y bounds (if setting up a double integral $\mathrm{d}y\,\mathrm{d}x$, start here)**, repeat step 1 but now in 2D rather than 3D. Our innermost variable (z in this example) is now out of the picture entirely, and now we are looking at the integration order $\mathrm{d}y\,\mathrm{d}x$. So we ask: for a fixed value of x, what is the allowed range of y? Answering this is completely analogous to step 1:

   - To answer it geometrically, if we fix a value for x, then the points (x, y) form a line in the xy-plane in the y direction. Traverse this line upwards (i.e. in the positive y direction) and look at where you enter and exit the region D; the y coordinates of those points will give you the lower and upper y bounds respectively.

  - To answer it algebraically, look at the inequalities defining the region D and see which ones involve y. Isolate y in each, and then take

$$\max(\text{all expressions less than } y) < y < \min(\text{all expressions greater than } y).$$

  Here the "expressions" should involve only x. Note that if there are multiple expressions on either side then you may need to split your integral into multiple parts (although typically you should pick an integration order that avoids this).

4. **To identify the relevant interval(s) in the x-axis, thus finding the x bounds**, repeat step 2 but now a dimension lower:

  - Geometrically, for which values of x does the line you drew in the preceding step actually intersect the region D? The collection of all such x gives you the desired interval(s).
  - Algebraically, look at (lower y bound) < (upper y bound) together with the inequalities describing D that depend only on x. This gives you a system of inequalities solely about x; together they define the relevant interval(s) in the x-axis.

Important things to remember:

- The bounds for a given variable can depend on variables further out, but NEVER on variables deeper inside the integral. So if your integration order is $\mathrm{d}z\,\mathrm{d}y\,\mathrm{d}x$ for example, the y bounds may depend on x but not on z.
- In particular, the outer-most bounds cannot depend on any variables! They should be constants.

Where to look in Stewart: §15.2 for double integrals, §15.6 for triple integrals. See in particular Exercises 15.6.9-18, and also 15.6.29-32.

4 Rewriting Multiple Integrals: Changing Order

Throughout this whole section we will look at a specific example:

$$\int_{-1}^{1}\int_{0}^{4}\int_{-2\sqrt{1-z^2}}^{2\sqrt{1-z^2}} f(x,y,z)\,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z.$$

Suppose we were asked to switch the integration order to $\mathrm{d}z\,\mathrm{d}y\,\mathrm{d}x$ instead. We can easily read off an algebraic description of the region E over which we are trying to integrate:

$$-1 < z < 1,$$
$$0 < y < 4,$$
$$-2\sqrt{1-z^2} < x < 2\sqrt{1-z^2}.$$

Now there are basically two ways to proceed.

4.1 Using an intermediate geometric step

We can convert the original algebraic bounds to a picture, and then to convert the picture into new algebraic bounds:

$$\text{old integration bounds} \rightsquigarrow \text{geometric picture} \rightsquigarrow \text{new integration bounds}$$

Although the particular example we're looking at is a triple integral, the same method works for a double integral (indeed, it is only simpler!).

To build up a picture of the region E from the bounds, it's perhaps easier to work from the "outside in." The outermost bound says that the relevant interval on the z-axis to consider is from -1 to 1. Then the y bounds say that, for every z value from that interval, the y bounds are from 0 to 4, giving us a rectangle in the yz-plane. Finally, we look at the innermost bounds, which say that for every point (y, z) in that rectangle, the x bounds go from $-2\sqrt{1-z^2}$ to $2\sqrt{1-z^2}$. This process is illustrated in the top half of the picture.

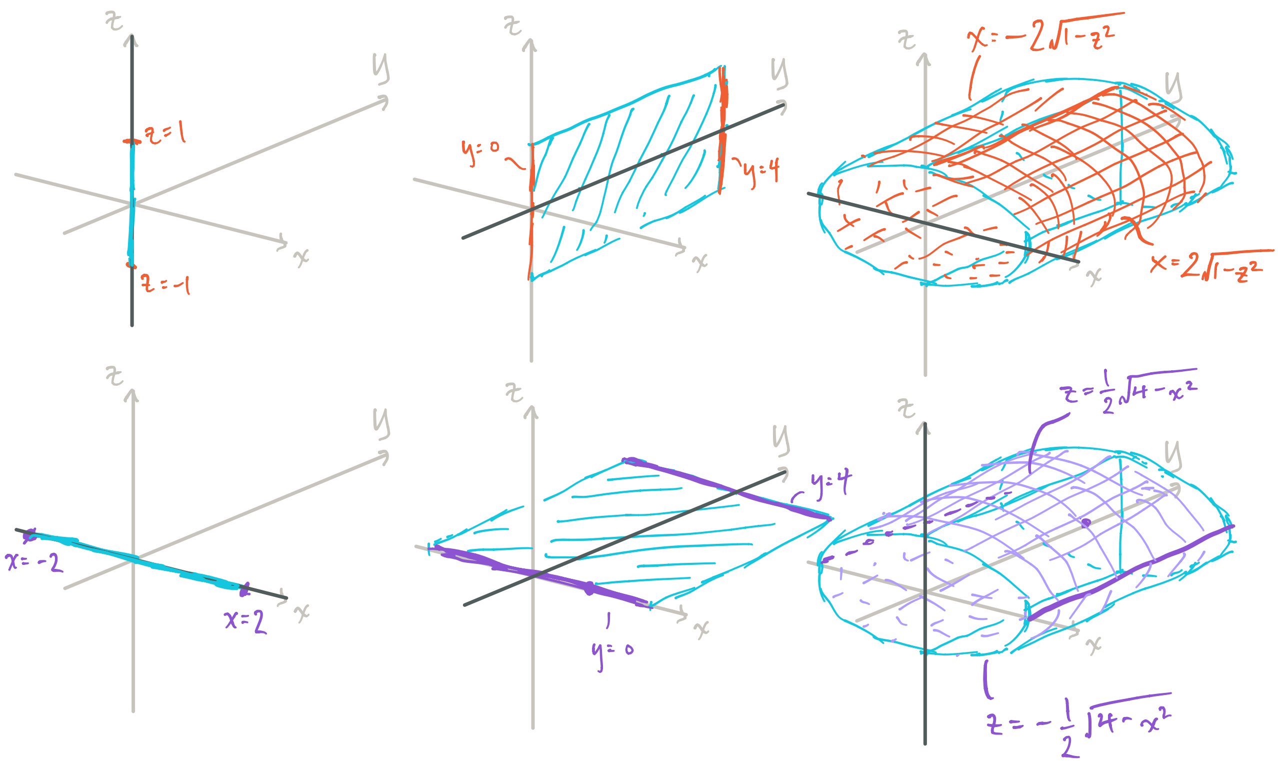

Now we employ the "geometric" half of the procedure in §3 to break this picture down into our new desired bounds, which is illustrated in the bottom half of the picture, from right to left. As outlined in §3 we do this from the "inside out."

1. If we fix x, y, then from the picture we see that the z bounds depend only on x. The cross-sections parallel to the xz-plane are given by $x^2 + 4z^2 < 4$, which is a filled-in ellipse. So for a fixed value of x (and y), the z bounds are from $-\frac{1}{2}\sqrt{4-x^2}$ to $\frac{1}{2}\sqrt{4-x^2}$.

2. Lines parallel to the z-axis only intersect E when they stems from the rectangle depicted in the xy-plane with vertices $(-2, 0), (-2, 4), (2, 4), (2, 0)$. Alternatively, if you project this cylinder onto the xy-plane (as in, view it from a camera placed far away on the z-axis), you see that rectangle.

3. Setting up the y bounds for this rectangle are easy: they go from 0 to 4 independently of x.

4. Likewise, the relevant interval in the x-axis is straightforward to identify: it's just the interval from -2 to 2.

Altogether the integral is

$$\int_{-2}^{2}\int_{0}^{4}\int_{-\frac{1}{2}\sqrt{4-x^2}}^{\frac{1}{2}\sqrt{4-x^2}} f(x, y, z)\,dz\,dy\,dx.$$

4.2 Going directly via algebra

Alternatively one can skip the geometric portion entirely: the original bounds already give us an algebraic description of the region of interest E, so we can just immediately apply the "algebraic" half of the procedure in §3. This has the potential to be faster, but perhaps less enlightening. Ideally one should be comfortable with both (knowledge of one approach helps understanding of the other).

Our region E is described by the inequalities

$$-1 < z < 1,$$
$$0 < y < 4,$$
$$-2\sqrt{1-z^2} < x < 2\sqrt{1-z^2}.$$

The desired new integration order is $\mathrm{d}z\,\mathrm{d}y\,\mathrm{d}x$. We apply the procedure from §3:

1. The inequalities involving z are:

$$-1 < z, \quad z < 1, \quad -2\sqrt{1-z^2} < x, \quad x < 2\sqrt{1-z^2}.$$

   The last two inequalities are the same as saying $x^2 < 4 - 4z^2$, or $z^2 < (4-x^2)/4$, so we can isolate z as

$$-1 < z, \quad z < 1, \quad -\frac{1}{2}\sqrt{4-x^2} < z, \quad z < \frac{1}{2}\sqrt{4-x^2}.$$

   Hence our z bounds are

$$-\frac{1}{2}\sqrt{4-x^2} = \max\left(-1, -\frac{1}{2}\sqrt{4-x^2}\right) < z < \min\left(1, \frac{1}{2}\sqrt{4-x^2}\right) = \frac{1}{2}\sqrt{4-x^2}$$

   because $\frac{1}{2}\sqrt{4-x^2} \le 1$.

2. To identify the relevant region in the xy-plane, we have the inequality

$$-\frac{1}{2}\sqrt{4-x^2} < \frac{1}{2}\sqrt{4-x^2}$$

   from the preceding part, together with the inequality $0 < y < 4$ (this is the only original inequality not involving z).

   The first inequality is just saying that $\sqrt{4-x^2} > 0$, so it's equivalent to requiring $-2 < x < 2$. Hence the relevant region in the xy-plane is a rectangle:

$$-2 < x < 2,$$
$$0 < y < 4.$$

3. At this point it's pretty clear what the x and y bounds are. But to be wholly systematic, we'll continue with the procedure even if it seems a little silly: to find the y bounds, we look at the inequalities involving y. These are just

$$0 < y, \quad y < 4.$$

   Because y is already isolated in these, there is nothing to do, so the y bounds are $0 < y < 4$ (though that was clear already).

4. Finally we go to the x-axis. From the preceding part, we have the inequality $0 < 4$, which is obviously always true regardless of x. Then we also have the inequalities $-2 < x < 2$. So the relevant interval in the x axis is from -2 to 2.

Hence the new integral is

$$\int_{-2}^{2}\int_{0}^{4}\int_{-\frac{1}{2}\sqrt{4-x^2}}^{\frac{1}{2}\sqrt{4-x^2}} f(x,y,z)\,\mathrm{d}z\,\mathrm{d}y\,\mathrm{d}x.$$

Where to look in Stewart: §15.2 for double integrals; Exercises 15.2.45-56. §15.6 for triple integrals; Exercises 15.6.33-36.

5 Rewriting Multiple Integrals: Change of Variables

The multivariate change of variables formula for double integrals states that if

$$\begin{cases} x = g(u, v) \\ y = h(u, v) \end{cases}$$

is a (differentiable) one-to-one and onto map from a region S in the uv-plane to a region R in the xy-plane, then

$$\boxed{\iint_S f(g(u,v), h(u,v)) \left| \frac{\partial(x,y)}{\partial(u,v)} \right| \mathrm{d}u \, \mathrm{d}v = \iint_R f(x,y) \, \mathrm{d}x \, \mathrm{d}y.}$$

Using the change of variables formula is not as involved as the preceding topics in this write-up—provided you can remember everything that goes into the formula! So let's break this down into easily digestible pieces.

In practice, you will be given the integral on the right and asked to produce the integral on the left. There are two essentially disjoint subproblems that need to be solved.

1. **What happens to the integrand (the thing inside the integral)?** This is purely computational. First of all, we need to rewrite the existing integrand $f(x, y)$ in terms of u and v instead, but this just amounts to **substituting the given expressions $g(u, v)$ and $h(u, v)$ in place of x and y.** So that explains the $f(g(u, v), h(u, v))$ portion of the left hand side's integrand.

   However, that alone is not enough, because our map from the uv-plane to the xy-plane could *distort area measurements*. **A corrective factor is needed: the absolute value of the Jacobian determinant**. You should think of this term as measuring the area distortion factor:

   $$\left| \frac{\partial(x,y)}{\partial(u,v)} \right| = \left| \det \begin{bmatrix} \partial x/\partial u & \partial x/\partial v \\ \partial y/\partial u & \partial y/\partial v \end{bmatrix} \right|$$

   which is equal to $|g_u h_v - g_v h_u|$ if $x = g(u, v)$ and $y = h(u, v)$.

   Don't forget the absolute values! They're not present in the single-variable change of variables formula, but they are necessary in the multivariate setting.

2. **What happens to the bounds?** Although there is perhaps less computation in this part, this question is usually conceptually more challenging. The strategy is as follows:

   (a) Write down an algebraic description of the region R in the xy-plane, using inequalities involving x and y.

   (b) Then substitute $x = g(u, v)$ and $y = h(u, v)$ into those inequalities. Now you have an algebraic description of some region in the uv-plane.

   (c) Typically that region will suffice as S. But there is something we need to make sure of: for every (x, y) in R, does the system of equations

   $$x = g(u, v), \quad y = h(u, v)$$

   have *exactly one solution* in the variables u and v? If the answer is "yes," then we are fine, and the region identified in (b) will suffice as S (move on to the next step). But if not:

- Is it because there are values of (x, y) in R for which there are multiple solutions for u and v? If this is the case, we will need to make sure that S is small enough so that there is always only one solution inside of S.
- Is it because there are values of (x, y) in R for which there are *no solutions* for u and v? If this is the case... we are doomed. The transformation cannot be used to rewrite the integral!

(d) Set up the u, v integral bounds for the region S, e.g. using the algorithm outlined in §3 (start at step 3 because this is a double integral, and of course replace "y" and "x" by the variables u, v).

3. Put the answers from the two subproblems together to get the desired u, v integral.

Example 1 (Polar integration)**.** As an example, we can give a purely algebraic treatment of polar integration. This is primarily to highlight the various parts of the above algorithm—in practice, a geometric standpoint towards polar integration is often preferable to the below, which is rather contrived by comparison.

The point is that we have much weaker geometric intuition towards arbitary changes of variables. If you understand precisely how it works for a transformation that we can visualize well, perhaps that will give you some faith in how it works, and then you can use it to more comfortably study arbitrary transformations algebraically.

Let R denote the disk of radius 1 centered at $(0, 1)$ in the xy-plane, and suppose we want to evaluate the integral

$$\iint_R y^2 \, \mathrm{d}x \, \mathrm{d}y$$

by applying the transformation $x = r\cos\theta$, $y = r\sin\theta$ (so r, θ play the roles of u, v in the above algorithm).

1. The integrand becomes:

$$\begin{aligned}&(\text{original integrand with subsitution applied})(\text{corrective factor})\\ &= r^2(\sin\theta)^2 \left|\det\begin{bmatrix}\cos\theta & -r\sin\theta \\ \sin\theta & r\cos\theta\end{bmatrix}\right| = r^2(\sin\theta)^2|r|.\end{aligned}$$

2. To identify the corresponding region S in the $r\theta$-plane:

(a) We write down an algebraic description of the original region R in the xy-plane:

$$x^2 + (y-1)^2 < 1.$$

(b) Then we substitute to get

$$\begin{aligned} r^2(\cos\theta)^2 + r^2(\sin\theta)^2 - 2r\sin\theta + 1 &< 1 \\ r^2 - 2r\sin\theta &< 0 \\ r^2 &< 2r\sin\theta. \end{aligned}$$

(c) However, recall that any given point (x, y) has infinitely many polar representations! That is to say, there are infinitely many solutions to $x = r\cos\theta$, $y = r\sin\theta$ for r and θ. We will need to make sure that we pick a region in which this is not an issue. The standard way of doing this is to require that $r \geq 0$ and that $0 \leq \theta < 2\pi$. With these additional stipulations, polar representation becomes unique[1].

As we are now assuming $r > 0$ (the origin isn't in $x^2 + (y-1)^2 < 1$, and even if it were, we could throw it out because it's only one point) we can divide by r and get $r < 2\sin\theta$.

[1]... well, except for the origin, but that's not an issue because it's only one point.

(d) Our region is described by $r < 2\sin\theta$, $r > 0$, $0 \leq \theta < 2\pi$.

If we use the integration order $\mathrm{d}r\,\mathrm{d}\theta$, then we see that the r bounds are $0 < r < 2\sin\theta$.

To get the θ bounds, we have the inequality $0 < 2\sin\theta$ from the r bound, in addition to the original inequality $0 \leq \theta < 2\pi$. So the θ bounds are $0 < \theta < \pi$ (see §3 for more discussion about this process).

3. Putting everything together, we get

$$\int_0^{\pi} \int_0^{2\sin\theta} r^2(\sin\theta)^2 |r| \,\mathrm{d}r\,\mathrm{d}\theta.$$

But since r is always positive in our region of integration, we can drop the absolute values:

$$\int_0^{\pi} \int_0^{2\sin\theta} r^3(\sin\theta)^2 \,\mathrm{d}r\,\mathrm{d}\theta.$$

The rest is just an integral computation.

Where to look in Stewart: §15.9. Exercises 15.9.23-27 may be good practice in particular. Note that you will want to start by making a guess for what u and v should be in terms of x and y; to apply the change of variables formula you will first need to solve that system of equations to get x and y in terms of u and v.