VOLUME COMPUTATIONS USING INTEGRALS

MATH 152, SECTION 55 (VIPUL NAIK)

Corresponding material in the book: Section 6.2, 6.3.

Difficulty level: Hard (degree of hardness depends on your visuo-spatial skills and prior exposure to these ideas).

What students should definitely get: The basic constructive ideas for volume: cylinders with constant and varying cross section, surfaces of revolution (disks and washers). The formula and mechanics for the shell method.

Note: I haven't included pictures, since these are hard to draw. I suggest that you look at pictures in the book, which are pretty well done, and of course, pay attention in class.

Executive summary

Words ...

(1) The cross section method for computing volume is an analogue of the two-dimensional area computation method: our slices are replaced by cross sections by planes parallel to a fixed plane, and the line of integration is a line perpendicular to the planes. One-dimensional slices are replaced by two-dimensional cross sections.
(2) Suppose Ω is a region in the plane. We can construct a right cylinder with base Ω and height h. This is obtained by translating Ω in a direction perpendicular to its plane by a length of h. The cross section of this right cylinder along any plane parallel to the original plane looks like Ω if that plane is within range. The volume is the product of the area of Ω and the height h. This is also called the right cylinder with constant cross section Ω.
(3) We can also construct an oblique cylinder. Here, the direction of translation is not perpendicular to the original plane. The total volume is the product of the area of Ω and the height perpendicular to Ω. Oblique cylinders are to right cylinders what parallelograms are to rectangles.
(4) More generally, the volume of a solid can be computed using the cross section method. Here, we choose a direction. We measure areas of cross sections along planes perpendicular to that direction, and integrate these areas along that direction.
(5) This general approach has another special case that is perhaps as important as right cylinders. These are the *cones* (there are right cones and oblique cones). A cone is obtained by taking a region in a plane and connecting all points in it to a point outside the plane. It is a right cone if that point is directly above the center of the region. The volume of a cone is 1/3 times the product of the base area and the height, i.e., the perpendicular distance from the outside point to the plane. In particular, a cone has one-third the volume of a cylinder of the same base and height.
(6) A solid of revolution is a solid obtained by revolving a region in a plane about a line (called the axis of revolution). The volume of a solid of revolution can be computed by choosing the axis as the axis of integration and using the planes of cross section as planes perpendicular to it. These cross sections are either circular disks or annuli.
(7) The *disk method* is a special case of the above, where the region between revolved is supported on the axis of revolution. For instance, consider the region between the x-axis, the graph of a function f, and the lines $x = a$ and $x = b$. The volume of the corresponding solid of revolution is $\pi \int_a^b [f(x)]^2 \, dx$. This is because the radius of the cross section disk at $x = x_0$ is $|f(x_0)|$.
(8) The *washer method* is the more general case where the region need not adhere to the axis of revolution. For instance, consider two nonnegative functions f, g and suppose $0 \le g \le f$. Consider the region bounded by the graphs of these two functions and the lines $x = a$ and $x = b$. The volume of the corresponding solid of revolution is $\pi \int_a^b ([f(x)]^2 - [g(x)]^2) \, dx$. Note that in the more general case

where the functions cross each other, we may need to split into sub-intervals so that we can apply the washer method on each sub-interval.

(9) The shell method works for situations where we revolve about the y-axis the region made between the graph of a function and the x-axis. The formula here is $2\pi \int_a^b xf(x)\,dx$ for f nonnegative and $0 < a < b$. If f could be positive or negative, we use $2\pi \int_a^b x|f(x)|\,dx$.. More generally, if we are looking at the region between the graphs of f and g (vertically) with $g \le f$, we get $2\pi \int_a^b x[f(x) - g(x)]\,dx$. If we don't know which one is bigger where, we use $2\pi \int_a^b x|f(x) - g(x)|\,dx$.

Actions ...

(1) To compute the volume using cross sections, we first need to set things up so that we know the cross section areas as a function of the position of the plane. For this, it is usually necessary to use either coordinate geometry or basic trigonometry, or a combination.
(2) A solid occurs as a solid of revolution if it has complete rotational symmetry about some axis. In that case, that axis is the axis of revolution and the original region that we need is obtained by taking a cross section in any plane containing the axis of revolution and looking at the part of that cross section that is on one side of the axis of revolution.
(3) For solids of revolution, be particularly wary if the original figure being revolved has parts on both sides of the axis of revolution. If it is symmetric about the axis of revolution, delete one side.
(4) Be careful about the situations where you have to be sign-sensitive and the situations where you do not. In the disk method sensitivity to signs is not important. In the washer method and shell method, it is.
(5) The farther the shape being revolved is from the axis, the greater the volume of the solid of revolution.
(6) The average value point of view is sometimes useful for understanding such situations.

1. Motivation: from area to volume

1.1. What are we trying to do?

Our purpose right now is to find formulas for the volumes of various three-dimensional figures. This is a little like our attempts at finding areas of regions, which we successfully did, at least for some regions. Wait, what?

The process that we are going through is something whose broad outlines should be familiar to you. Think back, for instance, to how we dealt with differentiation. We first computed formulas for the derivatives of a few functions. Then, we considered all the ways that new functions can be created from old functions. Finally, we found formulas that tackled each way of creating a new function from an old function. Combined with the knowledge of how to differentiate the basic functions, this allowed us to differentiate any function given to us using a simple set of rules.

Similarly, when trying to figure out general strategies for finding limits, we started out by computing a few basic limits, and looking at rules for computing limits of functions created from simpler functions.

We did a similar thing for integration: we learned rules for finding antiderivatives for some basic functions, and then we learned various processes of combination (something we're not quite done with yet). The overall strategy is:

(1) Find out how to deal with basic situations.
(2) Identify the typical ways that basic situations are combined to create more complicated situations.
(3) For each such process of combining basic situations into more complicated situations, identify a way of reducing the problem for the complicated situation in terms of the basic situations.

We shall consider how to deal with volumes. Our main difficulty in calculating volumes is with step (2) – we don't have an understanding of the systemic processes whereby new three-dimensional figures can be created. Once we do, we can try to find a volume formula for each such process, and use these formulas to calculate the areas of a number of figures.

1.2. A recapitulation of how we handled area computations.

We have so far dealt with two kinds of area computations. The first is computing areas *against the x-axis*. Here, we are measuring the area bounded between a curve and the x-axis, or the area between two curves and two vertical lines.

Let us reflect more carefully on how we can characterize these situations geometrically. In all these situations, the region Ω that we have has the property that the intersection of Ω with any vertical line is either empty or a line segment. Regions of this kind are sometimes called Type I regions. For Type I regions, the general formula for the unsigned area is:

$$\int (\text{Length of the line segment as a function of } x)\, dx$$

This process can be thought of as *vertical slicing*. We are dividing the area that we want to measure into vertical slices, and then integrating the length along the perpendicular axis (which is horizontal).

The other procedure that we saw for integration is integration against the y-axis. This kind of integration works for regions Ω which have the property: the intersection of Ω with any horizontal line is either empty or a line segment. Regions of this type are sometimes called Type II regions. The formula for the area of a Type II region is

$$\int (\text{Length of the line segment as a function of } y)\, dy$$

This process can be thought of as *horizontal slicing*. We are dividing the area that we want to measure into horizontal slices, and the integrating the length along the perpendicular axis (which is vertical).

Thus, we have seen two processes of breaking up an area into slices: vertical slicing (where we integrate the lengths along a horizontal axis) and horizontal slicing (where we integrate the lengths along a vertical axis).

Notice that both these procedures are variants of the same basic procedure: choose two mutually perpendicular directions, such that all lines in one direction have intersection with the region that is either empty or a line segment. Then, integrate the length of the line segment along the perpendicular direction.

Note also that the extreme case of both these occurs in rectangles. Here, whether we use horizontal or vertical slicing, we are integrating a constant function.

1.3. **How does this general idea carry over to three dimensions?** $1 + 1 = 2$, but $1 + 1 \neq 3$. So, the idea of choosing two mutually perpendicular directions, one for the slices and the other as the direction of integration, does not work directly for computing volumes. However, it *is* true that $2 + 1 = 3$. This suggests a slightly different strategy to measure the volume of a three-dimensional region Ω: choose a plane π and a line ℓ perpendicular to π. Now, measure the areas of the intersection of Ω with regions perpendicular to π, and integrate this area along ℓ.

In other words, the *slices* are two-dimensional and parallel to each other, and the direction of the line along which we integrate is perpendicular to those planes.[1]

In the forthcoming section, we look at some systemic processes for creating three-dimensional structures and for slicing them suitably.

Sidenote: Distinction between a disk and a circle. Henceforth, when I refer to a *circle*, I refer to the *boundary*, i.e., the set of points whose distance from the center equals the radius. When I want to talk of the circle along with the interior region, I will use the term *circular disk*, or, more briefly, *disk*. When I want to simply look at the interior and exclude the boundary, I will use the term *interior of the disk* or *open disk*.

I will, however, switch between a circle and its disk easily, hence when I talk about the center, radius, or diameter of a disk, I am referring to those notions for its boundary circle.

2. Creating three-dimensional structures

2.1. **The general concept of a right cylinder.** What you may have been told is a *cylinder* is more appropriately termed a *right circular cylinder*. The adjectival qualifier *circular* indicates that the base is a circle (more precisely, the boundary is a circle and the base is a circular disk). The term *right cylinder*, in general, means something like a right circular cylinder except that the base need not be circular.

Basically, we take a region Ω in the plane with boundary Λ and then translate Ω along a direction perpendicular to the plane for a fixed length. That fixed length is called the *height* of the right cylinder. This gives the (solid) right cylinder with cross section Ω. The curved surface of the cylinder is the boundary

[1] The reason why we are forced to use $2 + 1 = 3$ rather than $1 + 2 = 3$ is because the only kind of integration that we have explicitly dealt with is integration in one variable, i.e., along a line.

of this, which is obtained by translating Λ in a direction perpendicular to the plane of Ω. The two *caps* are the two copies of Ω located at the two ends.

The term *cross section* here refers to the fact that if we take any plane parallel to the plane of Ω, its intersection with the right cylinder is a copy of Ω if the plane is located in the relevant region; otherwise it is empty.

You may have heard the term *cross section* arising in different contexts. It basically means the intersection with a given plane. For instance, in biology, when studying things ranging from tree trunks to micro-organisms and cells, we take cross-sections in various directions.

The right cylinder has a constant cross section. In this sense, it is similar to a rectangle in two dimensions, which has constant cross sections.

The volume of a right cylinder is given by:

$$\text{Volume of right cylinder} = \text{Area of cross section} \times \text{Height}$$

Some particular cases of interest:

- When the base cross section is a circular disk, we get a *right circular cylinder*.
- When the base cross section is a polygon, we get what is often called a *prism*. In particular, when the base is a rectangle, we get a rectangular prism.

2.2. **Oblique cylinders.** A slight variant on right cylinder is oblique cylinder. Oblique cylinders are to parallelograms what right cylinders are to rectangles. Here is the construction of an oblique cylinder.

Start with a region Ω in a plane π. Now, choose a direction in space that is not parallel to the plane π. Translate Ω by a length l along this direction. The region traced this way is termed an oblique cylinder.

The volume of an oblique cylinder is given by:

$$\text{Volume of oblique cylinder} = \text{Area of cross section} \times \text{Height perpendicular to cross section}$$

Equivalently:

$$\text{Volume of oblique cylinder} = \text{Area of cross section} \times \text{Length } l \times \sin\theta$$

where θ is the angle between the plane π and the direction of translation. In particular, when $\theta = \pi/2$, we get a right cylinder.

If we consider cross sections of oblique cylinder parallel to π, each of these cross sections looks like Ω. However, unlike the right cylinder case, the *location* of the Ω in the cross section plane keeps changing.

2.3. **More oblique than oblique.** In fact, it is possible to get eve nmore oblique than oblique – we translate a shape in a plane along a direction other than the plane, but we keep changing the direction. Thus, each cross section still has the same shape, but its location changes rather unpredictable. We'll see some such situations in a homework/quiz/test.

2.4. **Variable cross sections.** We next consider a situation where the cross sections are variable. This is no longer a right cylinder, but we can use the idea mentioned a little while ago – integrating the area function. Earlier, we multiplied a constant with the height over which that constant was valid. Now, we integrate a variable function over an interval. Remember, integration is like multiplication where the thing you're trying to multiply keeps changing. The important thing is that the area of each cross section should be something we know how to measure. The general formula is:

$$\text{Volume} = \int (\text{Area of cross section perpendicular to } x)\, dx$$

2.5. **Cones.** One case of particular importance, where it is useful to remember a general approach as well as the specific answer, is that of the *cone*. A cone is defined as follows. Suppose Ω is a region in a plane π and P is a pont not in π. The cone corresponding to Ω and P is the union of all the line segments joining P to points in Ω.

Some examples of cones are:

(1) A *tetrahedron* is a cone where the base is a triangular region.

(2) A *right circular cone* is a cone where the base is a circular disk.
(3) A *pyramid* is a cone where the base is some polygon.

When we set up the cross section integration for the cone, we see that the shape of any cross section parallel to π is the same as that of Ω, but the size is different. We can use similar triangles to determine the size. If we define:

$$\alpha = \frac{\text{Distance from } P \text{ to cross section}}{\text{Distance from } P \text{ to } \pi}$$

Then the linear measurements for the cross section are α times the corresponding linear measurements for Ω. Since area is two-dimensional, the area of the cross section is α^2 times the area of Ω. We now get that the overall volume is:

$$\int_0^h (x/h)^2 \operatorname{Ar}(\Omega)\, dx$$

Plugging in $x = \alpha h$, we get:

$$\int_0^1 \alpha^2 \operatorname{Ar}(\Omega) h \, d\alpha$$

We pull out the constants, and get:

$$\operatorname{Ar}(\Omega) h \int_0^1 \alpha^2 \, d\alpha$$

The integral now gives 1/3, and we thus get:

$$\text{Volume of cone} = \frac{1}{3} \text{Area of base region} \times \text{Height}$$

Now you understand why you have that 1/3 in the formula for the volume of a cone: $(1/3)(\pi r^2)(h)$.

But not completely. Why 1/3? Well, let's think back to the two-dimensional analogue of this. What's a two-dimensional analogue of a cone? It's just a triangular region. The analogue of the two-dimensional base is a one-dimensional line segment. And we remember that:

$$\text{Area of triangle} = \frac{1}{2} \text{Length of base line segment} \times \text{Height}$$

So why do we get 1/2 in the two-dimensional case and 1/3 in the three-dimensional case? Well, you might guess that we basically get $1/n$ in the n-dimensional case. And then you go back and look at the proof, and see that it essentially works this way:

$$\int_0^1 \alpha^{n-1} \, d\alpha = [\alpha^n/n]_0^1 = 1/n$$

3. Solids of revolution: the disk and washer method

3.1. Definition of solid of revolution. There is another procedure for constructing three-dimensional figures. Three-dimensional figures constructed this way are called *solids of revolution*. This is obtained as follows: we start with a region Ω and a line ℓ. Next, we rotate Ω about the line ℓ in three dimensions. The region obtained in this way is termed the *solid of revolution* of Ω.

For simplicity, we will assume that Ω lies completely to one side of ℓ. We study such surfaces in two steps. First, we study the special case where one boundary of Ω is along ℓ. After that, we study the case where all of Ω could lie on one side of ℓ. The method for the first case is termed the *disk method* and the method for the second case is termed the *washer method*.

Aside: The surface of a solid of revolution includes two capping disks. The remaining part of this surface is the curved surface, and this is often called a *surface of revolution*. Surfaces of revolution turn out to be very important in a variety of natural processes.

3.2. **Disk method.** Consider the area bounded by the graph of the function $y = f(x)$ and the x-axis between $x = a$ and $x = b$ (with $a < b$). Assume, for now, that the graph of f lies completely on the positive side of the x-axis. So, the picture looks something like Figure 6.2.8 (left) of the book. Revolving this about the x-axis gives a solid of revolution as shown in fiure 6.2.8 (right) of the book.

We now consider how to apply the method of parallel cross sections to this volume computation. We consider the axis as the x-axis and the cross sections are thus in the yz-plane. In particular, we see by our construction that all the cross-sections are disks and the disk for a cross section at $x = x_0$ has radius $f(x_0)$ and area $\pi(f(x_0))^2$. The area is thus:

$$\int_a^b \pi(f(x))^2 \, dx$$

We can pull the π out of the integral if we want. This is the general formula for calculating the area.

It turns out that the formula is also valid for a function that crosses the x-axis. In this case, the parts abover the x-axis and the parts below the x-axis are out of phase by π as we revolve them. However, the overall analysis remains the same, with the radius being $|f(x_0)|$ instead of $f(x_0)$. Since we are squaring it anyway, the final answer remains the same.

Here are some particular cases of solids of revolution whose volume can be computed using the disk method:

(1) The right circular cylinder with radius r and height h can be realized as the solid of revolution for the region between the x-axis and the graph of a constant function with value r (bounded by vertical lines) over an interval of length h. The region being rotated is thus a rectangle with dimensions r, h, and h is the fixed side.

(2) The right circular cone with radius r and height h can be realized as the solid of revolution for the region between the x-axis and the graph of the function $y = rx/h$ on the interval $[0, h]$ (bounded by a vertical line at $x = h$). The region being rotated is thus a right triangle with legs r and h and h is the fixed side.

We can verify that we get the same answer as usual when we apply the disk method.

3.3. **Solids of revolution: the washer method.** What if the region being rotated is completely on one side of the axis of rotation? For instance, imagine a disk far away from the x-axis being revolved about the x-axis. The corresponding solid is sometimes called a *filled torus* or *solid torus* (the boundary of this, which is a surface obtained by revolving the boundary circle, is usually simply called a *torus*).

The washer method is a method that allows us to compute the areas of such solids. Again, the idea is to use parallel cross sections. In this case, the cross sections are not disks, but regions called *annuli*. Given a point P and two concentric circles centered at P (in the same plane) with radii $r < R$, the annulus for these two radii is the set of points in the bigger disk that are not there in the interior of the smaller disk. Thus, it is the region between the circles of radii r and R, along with the two boundary circles.

The area of such an annulus is given by $\pi(R^2 - r^2)$.

The upshot of this is that the volume of the solid of revolution obtained by revolving the region between $y = g(x)$ and $y = f(x)$, with $0 \le g(x) \le f(x)$, on $[a, b]$, is:

$$\int_a^b \pi[(f(x))^2 - (g(x))^2] \, dx$$

If the two functions cross each other, then if we are interested in the unsigned volume, we need to split into intervals based on which one is bigger where, calculate the volumes of the solids of revolution corresponding to each interval, and add up. In other words, we need to compute:

$$\int_a^b \pi|(f(x))^2 - (g(x))^2| \, dx$$

3.4. **Solids of revolution: the tale of the receding axis.** The first thing worth noticing about the volumes of solids of revolution is that the volume is *not determined* by the area of th region being rotated. It *also depends* on the choice of axis. As a general rule, the farther the axis from the region being rotated, the bigger the volume.

To understand this, consider the question: given a fixed number $h > 0$, what can we say about the area of the annulus of thickness h, i.e., where the outer radius is h more than the inner radius? For fixed h, this number increases as we increase the two radii. This is because the area is $\pi[(r+h)^2 - r^2] = \pi(2r+h)h$. The $2r + h$ term increases as r increases.

To give you some intuition about this, here is something that might strike you as visually counterintuitive: the area of the annulus with inner radius 4 and outer radius 5 equals the area of the disk of radius 3 (since $5^2 - 4^2 = 3^2$) even though the former has a much smaller thickness. The smaller thickness is compensated for (roughly) by the larger circumference.

The calculations that we did for the annulus show that as we move our axis farther and farther from Ω, the solid of revolution becomes larger and larger in volume. Remember: to calculate the volume of the solid of revolution, we create slices perpendicular to the axis of revolution, but we are not integrating the length of these slices; we are integrating the differences of *squares* of the endpoints of the slices. And this difference of squares increases as both numbers get bigger, even when the actual difference between them is constant.

3.5. **Solids of revolution: when the axis straddles the region.** So far, we have considered a situation where the region being revolved is completely on one side of the axis of revolution.

In the case that the region being revolved is partly on one side and partly on the other side of the solid of revolution, we must keep the following things in mind:

(1) If the region has *mirror symmetry* about the axis of revolution, then we can simply delete the half of the region on one side and consider the solid of revolution for the other half.
(2) Otherwise, in general, we must *fold* the region being rotated along the axis, i.e., reflect all the stuff on one side to the other, while keeping the stuff on the other side unchanged. Note that in the case of mirror symmetry, the reflected material overlaps with the original material. In some cases, such as the graph of a function about the x-axis, the reflected portion has no area of overlap with the stuff already there. In yet other cases, part of the reflected region overlaps, and the rest doesn't.

4. The shell method

4.1. **Formula.** The shell method applies to situations where we revolve about the y-axis the region made between a graph $y = f(x)$ and the x-axis. As before, we work with a nonnegative continuous function f on a closed interval $[a, b]$ with $0 < a < b$. Consider the region bounded by the graph of f, the x-axis, and the vertical lines $x = a$ and $x = b$. Now, consider the solid of revolution obtained by revolving this region about the y-axis. The volume of this solid of revolution is given by the formula

$$\int_a^b 2\pi x f(x)\,dx = 2\pi \int_a^b x f(x)\,dx$$

In the case that the function is not nonnegative throughout, we can use the more general formula:

$$\int_a^b 2\pi x |f(x)|\,dx$$

The best way of doing this is to partition the interval according to the sign of f.

4.2. **Slight generalization of this formula.** Consider now a slightly more general situation: we are looking at the region between the graphs of the functions f and g between $x = a$ and $x = b$. We consider the solid of revolution obtained by revolving this region about the y-axis. If $g \leq f$ on $[a, b]$, then the volume is given by:

$$\int_a^b 2\pi x[f(x) - g(x)]\,dx = 2\pi \int_a^b x[f(x) - g(x)]\,dx$$

(Note: We don't need any conditions on the nonnegativity of f and g here).

If f and g cross each other, we can use the general formula:

$$\int_a^b 2\pi x|f(x) - g(x)|\,dx$$

This is best handled by partitioning the interval according to where f is greater and where g is greater.

5. Average value point of view

5.1. **Overview.** For the various approaches we have seen so far for volume computation, there is an *average value point of view*. This can be thought of as a process whereby we compare our actual imperfect solid to a more perfect solid which is more uniform, and where the volume is given as a simple product. Let's illustrate this by beginning with our interpretation of volume as the integral of a variable cross section area.

Here, our *ideal* figure is a right cylinder, where the cross section area does not change for the cross sections (more generally, this is also true for oblique cylinders). In these ideal figures, the volume is the product of the constant cross section area and the height.

The volume in general can be thought of as the product of the *average* cross section area and the height. Here, the *average* cross section area is *defined* the way we calculate the average value for a function: we integrate it over the entire interval, and then divide by the length of the interval. In other words, the average cross section area is defined so that a right cylinder with that cross section and the height of our current figure has the same volume.

How does the average value point of view help? Computationally, it doesn't, but it gives us some intuition as to what kind of answers to expect. This is because, looking at the figure, we have some ideas about the average value: it must be somewhere between the minimum and the maximum value, for instance. This provides a reality check on the computations that we do.

5.2. **Average value for shell method.** Here, the ideal function is a constant function f on $[a, b]$ with constant value C. Revolving it about the y-axis yields a cylindrical shell with inner radius a, outer radius b, and height C. The volume is $\pi C(b^2 - a^2) = \pi C(b + a)(b - a)$. The value $\pi C(b + a) = 2\pi C(b + a)/2$ is the curved surface area of the cylinder whose radius is $(b + a)/2$, which is the cylinder whose radius is halfway between the inner and outer radius. We thus see that:

Volume of cylindrical shell = Curved surface area of mid-value cylinder × Difference of outer and inner radii

This is the ideal situation. In the real situation, we define the *average curved surface area* as:

$$\text{Average curved surface area} = \frac{\text{Volume of solid of revolution}}{\text{Difference of upper and lower limits}}$$

In our notation, the denominator is $b - a$. Thus, we obtain that the volume of the solid of revolution is $b - a$ times the *average curved surface area.* As before, this is not really computationally useful, but it might give us some intuition.

5.3. **Average value for disk method: different notions of average!** The average value point can also be used to understand the disk method.

Recall that the volume of a solid of revolution obtained by revolving about the x-axis the region between the x-axis and the graph of f from $x = a$ to $x = b$ is given by $\pi \int_a^b (f(x))^2\,dx$. Recall that we proved this formula by taking cross sections perpendicular to the x-axis. The area of a cross section at the value x is $\pi(f(x))^2$, because the cross section is a disk of radius $|f(x)|$.

Under the average value point of view, we are interested in the average value of this cross section area. There's a little subtlety in this.

To find the *area* between the graph of f and the x-axis from $x = a$ to $x = b$, we perform a simple integration $\int_a^b f(x)\,dx$ (or $\int_a^b |f(x)|\,dx$). On the other hand, to find the volume of the solid of revolution, we perform the integration $\int_a^b (f(x))^2\,dx$.

In other words, when finding the volume of the solid of revolution, we give a lot more weight to larger radii – because the radius is being squared. Remember the discussion from last time where we saw that an annulus with inner and outer radii 4 and 5 has the same area as the disk of radius 3. This is because the square of a number grows much faster than the number itself.

Our averaging process is also correspondingly biased. When we are calculating the average value in the ordinary sense, we do $\int_a^b f(x)\,dx/(b - a)$. However, when calculating the average of the areas of the disks, we are doing $\pi \int_a^b (f(x))^2\,dx/(b-a)$. The latter average value is *usually not the same* as the area of the disk whose radius is the average radius. Rather, it is usually larger, because taking the squares assigns greater weight to the bigger radii.[2]

6. Stock-taking

We have now seen some formulas and general approaches that use the ideas of integration to compute areas and volumes. Later in the course and/or in later life, you will encounter formulas to do a lot of the other things you've always wanted to do, such as formulas for arc lengths and surface areas. We are not getting into those formulas right now for two reasons: (i) they require more conceptual apparatus to understand, (ii) the kind of expressions that you typically get to integrate are expressions that you do not know how to deal with.

This brings us to one of the things that differentiates (pun!) differentiation from integration. Differentiation was based on a set of rules that we could apply blindly, because for every way of combining and composing existing functions, we had a corresponding way of breaking down the differentiation problem. With integration, however, we are in more wild territory, since there are no easy hard-and-fast rules and a lot depends on creativity and spotting persuasive patterns. This makes integration more fascinating, but it also means that ever so often, we come across a situation from the real world that boils down to computing an integral, and we don't really have an idea how to go about it.

Nonetheless, reducing a geometric problem of volume computation into a purely arithmetic/algebraic problem of evaluating a definite integral should be seen as a major step forward. Even if we have no clue about what an antiderivative might be, we can still use the upper sum/lower sum method to approximate this integral.

7. More computational intuition

7.1. **Stretching, shrinking, and scaling.**

We can use the "solid of revolution" idea to compute the volume of a sphere. A solid sphere of radius r is obtained by revolving a semicircular region of radius r about its diameter. The volume formula is thus:

$$\pi \int_{-r}^{r} (r^2 - x^2)\,dx$$

This gives the familiar formula $(4\pi/3)r^3$.

Note that a sphere is also the solid of revolution of a circular disk about its diameter. As noted earlier, since a circular disk has mirror symmetry about its diameter, so we can delete one of the semicircular pieces and still get the same solid of revolution.

Now, let's think about what happens if, instead of revolving a circular (or semicircular) disk, we revolve the region enclosed by an ellipse about its major or minor axis. An ellipse oriented along the axes and centered at the origin is a curve given by the equation:

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$$

with a, b positive.

If $a > b$ then the x-axis is the major axis and the y-axis is the minor axis. Note that both the x-axis and y-axis are axes of mirror symmetry; however, unlike the case of the circle, it is no longer true that every line through the origin is an axis of mirror symmetry.

Now, we could do the calculations pretty easily to compute the volume of the solid of revolution about either axis, but let's give an intuitive explanation that allows us to get at the answer. If we start with a circle centered at the origin and of radius b and stretch it by a factor of a/b in the x-direction, we get an ellipse. Clearly, the *area* of the ellipse is therefore a/b times the area of the circle, hence it is πab. What about the volume? We note that the axis along which we integrate gets stretched by a factor of a/b. A little thought now tells us that the answer will be $(4\pi/3)ab^2$.

[2]It turns out that the two averages are equal only for a constant function. The inequality being alluded to here indirectly is known as the arithmeti mean-quadratic mean (AM-QM) inequality or the arithmetic mean-root mean square (AM-RMS) inequality.

More generally, we see that:

(1) If the region being revolved is stretched by a factor of λ *along* the axis of revolution, the volume is multiplied by a factor of λ.
(2) If the region being revolved is stretch by a factor of μ *along* the axis of revolution, the volume is multiplied by a factor of μ^2. The square happens because when we revolve, the area contribution in each slice is proportional to the square of the radius or difference of squares of radius.

Thus, if the same ellipse were revolved about its minor axis, we'd get a volume of $(4\pi/3)a^2b$.

7.2. **Brief mention: Pappus' theorem.** Pappus' theorem is in a later section of the chapter that we're not including in this course, but it's a theorem worth taking a look at and understanding at least temporarily. For Exercise 6.3.44 (featuring in Homework 8), Pappus' theorem gives an alternative solution approach that is much shorter than the disk and shell methods that we will use to solve the problem. It also tells us what the answer will be – in this case $2\pi^2a^3$. The reason it is easier is because for the case of the circle, we know exactly where the center (centroid) is.