Math 3204: Calculus IV[1] Mikhail Lavrov

Lecture 1: Cylindrical coordinates

August 13, 2024 *Kennesaw State University*

1 Plans for the future

The first two weeks of this semester will seem like they don't quite fit in at first. We will be talking about alternative coordinate systems, and integration in those different coordinate systems. In the first week, we will discuss cylindrical and spherical coordinates. In the second week, we will cover substitution in 2-dimensional and 3-dimensional integrals.

The rest of the semester will pivot to learning about a wider variety of multivariable integrals, especially ones involving vectors and vector fields. Our ultimate goal is Stokes' theorem and its cousins: high-dimensional generalizations of the fundamental theorem of calculus. This seems unrelated. But we will see later on that differentiation and integration in alternative coordinate systems underlies everything we will do—not to mention that describing paths and surfaces in cylindrical and spherical ways will be very useful to us throughout the semester.

2 From polar coordinates to cylindrical coordinates

I will assume you have already seen polar coordinates. These let us represent a 2-dimensional point as a pair (r, θ), where r is a nonnegative distance from the origin, and θ is the angle made with the positive x-axis.

Cylindrical coordinates are one possible generalization of polar coordinates to 3 dimensions. Here, a 3-dimensional point P is represented by a triple (r, θ, z). The z-coordinate means the same thing it usually does: the height of the point P above the xy-plane. The pair (r, θ) is the 2-dimensional polar representation, but not of P: of the point Q that is the "shadow" of P on the xy-plane. This is shown in Figure 1a.

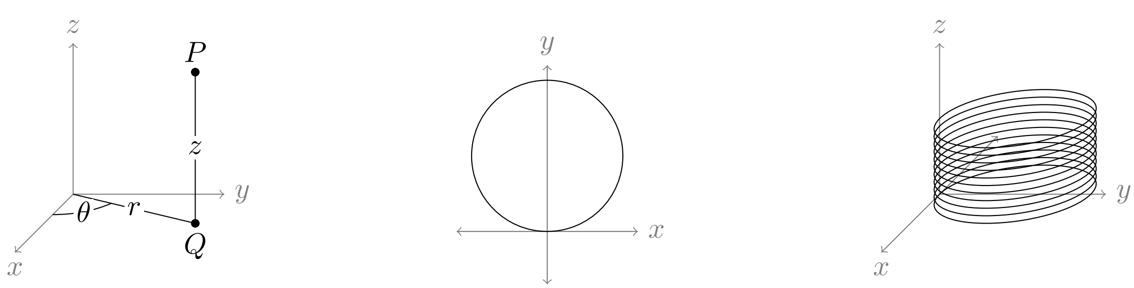

(a) The cylindrical coordinates of point P

(b) The 2-dimensional graph of $r = \sin\theta$

(c) The 3-dimensional graph of $r = \sin\theta$ with $0 \le z \le 1$

Figure 1: The connections between polar and cylindrical coordinates

[1]This document comes from an archive of the Math 3204 course webpage: http://misha.fish/archive/3204-fall-2024

(The "shadow" of P on the xy-plane is more formally known as the **projection** of P onto the xy-plane. We will see this concept again.)

The equations you need to convert between rectangular and cylindrical coordinates are inherited from polar coordinates: the important ones to remember are

$$\begin{cases} x = r\cos\theta \\ y = r\sin\theta \end{cases}$$

Nothing happens to z: it is the same in both coordinate systems.

It is also useful (but not as critical) to keep in mind that $x^2 + y^2$ simplifies to r^2, so $r = \sqrt{x^2 + y^2}$. We can also write θ in terms of x and y, but this involves an inverse tangent we want to be careful about, and we rarely need it.

We don't just want to describe points; we want to describe shapes. For example, in polar coordinates, the equation $r = \sin\theta$ describes a circle of radius $\frac{1}{2}$ centered at the point with rectangular coordinates $(x, y) = (0, \frac{1}{2})$, as seen in Figure 1b.

What happens when we take the same equation, and view it in cylindrical coordinates? Well, if we ask for a point P with cylindrical coordinates (r, θ, z) to satisfy $r = \sin\theta$, then we're saying that its "shadow" Q must lie on the circle we graphed earlier. The z-coordinate is free to vary, so we get a "stack" of circles, or a cylinder. By default, it's an infinite cylinder. If we add a condition on z, such as $0 \le z \le 1$, we get the finite cylinder shown in Figure 1c.

Another very useful way of leveraging our 2-dimensional intuition to reason about 3-dimensional cylindrical coordinates is to forget about θ at first, and work only with z and r. Because r is nonnegative, the pair (r, z) does not live in a plane, but in the rz-half-plane.

Suppose we want to understand the surface described in cylindrical coordinates by the equation $z = r^2$, with $0 \le \theta \le 2\pi$. (Remember: the angle θ can only ever range from 0 to 2π, so these bounds on θ are just telling us that θ can do anything it likes.) Then we can start by graphing $z = r^2$ in the rz-half-plane, as shown in Figure 2a.

As θ ranges from 0 to 2π, that half-plane drawing will be rotated around the z-axis. (Each direction away from the z-axis is a particular value of θ, and for that particular value of θ, we see a copy of our 2-dimensional drawing in that rz-half-plane. In the case of $z = r^2$, the result is called a paraboloid iand shown in Figure 2b.

To get another common example, start with the region defined $0 \le z \le 1 - r$ in the rz-half-plane. This is a triangle: it is bounded from below by $z \ge 0$, from the left by the universal rule that $r \ge 0$, and from above by $z \le 1 - r$.

If we have $0 \le z \le 1 - r$ in three dimensions, with $0 \le \theta \le 2\pi$, we get a solid cone. (A shape is called *solid* to emphasize that we include its interior, not just its boundary.) The cone will be oriented like a traffic cone or party hat, with the tip facing up, though it will be shorter and wider—or more "squashed"—than the typical traffic cone or party hat. Try drawing it yourself.

For the rest of class, we'll stick to one specific region as an example: the region R_{sp} defined by $0 \le \theta \le 2\pi$ and $r^2 \le z \le 1$. What is this region? It is a solid paraboloid: a "filled-in" version of Figure 2b. (That's why I'm calling it R_{sp}.)

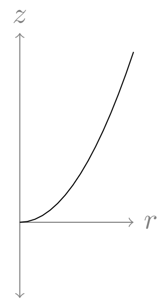

(a) The graph of $z = r^2$ in the rz-half-plane

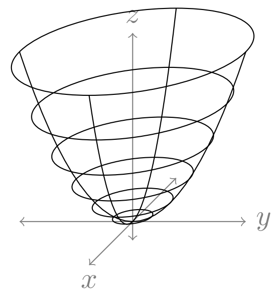

(b) The graph of $z = r^2$ in 3D space

Figure 2: Going from the rz-half-plane to 3-dimensional cylindrical coordinates

This is a class primarily about integration. The most basic integral we can take over R_{sp} is the volume integral: we can find the volume of R_{sp} with the formula

$$\text{Vol}(R_{sp}) = \iiint_{R_{sp}} \mathrm{d}V.$$

But what is this $\mathrm{d}V$?

3 Integration in cylindrical coordinates

You might remember that when we integrate over a region in polar coordinates, we cannot simply replace $\mathrm{d}x\,\mathrm{d}y$ by $\mathrm{d}r\,\mathrm{d}\theta$. We must replace it by $r\,\mathrm{d}r\,\mathrm{d}\theta$. Why is that?

The intuition is that the differential $\mathrm{d}x\,\mathrm{d}y$ stands in for a factor of $\Delta x \Delta y$ in the limit definition of a Riemann integral. That product $\Delta x \Delta y$ stands in for the area of a tiny square.

What about a tiny sliver of area in polar coordinates (Figure 3a), corresponding to a change of $\Delta\theta$ in θ and a change of Δr in r? Its area is not just $\Delta\theta\Delta r$. To a first approximation (which becomes exact in the limit), that sliver of area is a rectangle. One side of that rectangle is Δr, but the other side is $r\Delta\theta$: it is a $\frac{\Delta\theta}{2\pi}$ fraction of a circle with perimeter $2\pi r$. So we get an area of $r\Delta r\Delta\theta$, which means that we have $r\,\mathrm{d}r\,\mathrm{d}\theta$ in every polar integral.

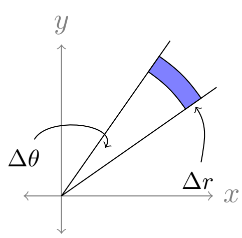

(a) A tiny area in polar coordinates

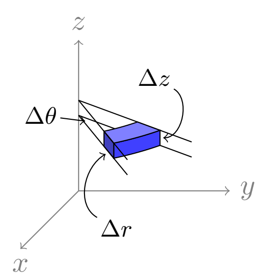

(b) A tiny volume in cylindrical coordinates

Figure 3: Where the factor of r comes from in polar and cylindrical integrals

For example, suppose that we want to know the area inside the circle in Figure 1b. That circle is in the top half of the xy-plane, so we have $0 \le \theta \le \pi$. The boundary of the circle is $r = \sin\theta$, so the interior of the circle is $0 \le r \le \sin\theta$. Therefore the area of the circle is

$$\int_{\theta=0}^{\pi} \int_{r=0}^{\sin\theta} r \,\mathrm{d}r \,\mathrm{d}\theta.$$

Let's actually do this integral, because it's good practice. The inside integral is

$$\int_{r=0}^{\sin\theta} r \,\mathrm{d}r = \left.\frac{r^2}{2}\right|_{r=0}^{\sin\theta} = \frac{\sin^2\theta}{2} - \frac{0^2}{2} = \frac{1}{2}\sin^2\theta.$$

For the outside integral, we have to integrate $\frac{1}{2}\sin^2\theta$, which you might not remember how to do. The tool we use is the double-angle formula for cosine: $\cos 2\theta = 2\cos^2\theta - 1 = 1 - 2\sin^2\theta$. We can, instead, solve for $\sin^2\theta$ in terms of $\cos 2\theta$, getting $\sin^2\theta = \frac{1-\cos 2\theta}{2}$.

Therefore the area of our circle is

$$\int_{\theta=0}^{\pi} \frac{1}{2}\sin^2\theta \,\mathrm{d}\theta = \int_{\theta=0}^{\pi} \frac{1-\cos 2\theta}{4} \,\mathrm{d}\theta = \left.\frac{\theta}{4} - \frac{\sin 2\theta}{8}\right|_{\theta=0}^{\pi} = \frac{\pi}{4}.$$

Of course, we already know that a circle with radius $\frac{1}{2}$ has area $\pi(\frac{1}{2})^2 = \frac{\pi}{4}$. Don't worry, we'll see more surprising integrals in the future.

Similarly, when integrating with respect to cylindrical coordinates, $\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z$ must be replaced by $r\,\mathrm{d}z\,\mathrm{d}r\,\mathrm{d}\theta$. The logic is the same: if we take a tiny volume corresponding to a change of Δz in z, a change of Δr in r, and a change of $\Delta\theta$ in θ, as shown in Figure 3b, then its volume will be $r\Delta z\Delta r\Delta\theta$, because it is approximately a $\Delta z \times \Delta r \times (r\Delta\theta)$ cuboid.

For example, the volume of the solid paraboloid R_{sp} is given by the iterated integral

$$\int_{\theta=0}^{2\pi} \int_{r=0}^{1} \int_{z=r^2}^{1} r \,\mathrm{d}z \,\mathrm{d}r \,\mathrm{d}\theta.$$

This is not too bad, as integrals go! As a general rule: if we're integrating with respect to a variable that doesn't appear in the expression we integrate, we just multiply the expression by the length of the integrating interval. For example, nothing inside the integral with respect to θ depends on θ, so we can simplify $\mathrm{Vol}(R_{sp})$ to

$$2\pi \int_{r=0}^{1} \int_{z=r^2}^{1} r \,\mathrm{d}z \,\mathrm{d}r.$$

Similarly, we're integrating r with respect to z over an interval of length $1 - r^2$, so we can further simplify

$$\mathrm{Vol}(R_{sp}) = 2\pi \int_{r=0}^{1} r(1-r^2) \,\mathrm{d}z.$$

Since $r(1-r^2) = r - r^3$, its antiderivative is $\frac{r^2}{2} - \frac{r^4}{4}$, and $\left.\frac{r^2}{2} - \frac{r^4}{4}\right|_{r=0}^{1} = (\frac{1}{2} - \frac{1}{4}) - (0 - 0) = \frac{1}{4}$. Multiplying by 2π, we get $\mathrm{Vol}(R_{sp}) = \frac{1}{2}\pi$.

4 Using integrals to find an average

Let's use these integrals to do something different. Where is the center of R_{sp}?

To find the center, we should develop a general theory of averaging over a region.[2] My claim is the following:

Claim 4.1. *Let R be a solid region in $\mathbb{R}^3$, and let $f\colon \mathbb{R}^3 \to \mathbb{R}$ be any function. Then the average value of f on R is given by*

$$\operatorname{Avg}(f) = \frac{1}{\operatorname{Vol}(R)} \iiint_R f(x, y, z)\,\mathrm{d}V = \frac{\iiint_R f(x, y, z)\,\mathrm{d}V}{\iiint_R \mathrm{d}V}.$$

(Better notation would be $\operatorname{Avg}_R(f)$, to emphasize that it also depends on R, but usually in our examples, there will only be one R to average over, and we can leave it implied.)

To get some intuition about this claim, think about what you would do if you wanted to find the average value of f on a set of finitely many points. You would simply add up the values of f on all the points, and divide by the number of points you have. When we have a continuous region R, all that changes is that the sum becomes an integral—and the number of points becomes the volume of the region.

These averages let us compute one important notion of what a "center" of a region is: the **centroid** of that region. The centroid of a region R is the point

$$(\overline{x}, \overline{y}, \overline{z}) = (\operatorname{Avg}(x), \operatorname{Avg}(y), \operatorname{Avg}(z)).$$

Returning to our solid paraboloid R_{sp}, we can skip computing $\overline{x}$ and $\overline{y}$ due to symmetry: the center of the paraboloid should be *somewhere* on the z-axis. (If you did take the integrals anyway, you would find yourself integrating an odd function, and end up with 0.) So the centroid is at the point $(0, 0, \overline{z})$, and we just have to compute $\overline{z} = \operatorname{Avg}(z)$. For this, let's use cylindrical coordinates.

We've already computed $\operatorname{Vol}(R_{sp}) = \frac{1}{2}\pi$. In the numerator of $\overline{z}$, we'll have to do a bit more work than we did then, because we're integrating rz, which *does* depend on z. The inmost integral becomes:

$$\int_{z=r^2}^{1} rz\,\mathrm{d}z = \frac{rz^2}{2}\bigg|_{z=r^2}^{1} = \frac{r(1)^2}{2} - \frac{r(r^2)^2}{2} = \frac{r}{2} - \frac{r^5}{2}.$$

Next, we must integrate with respect to r:

$$\int_{r=0}^{1} \left(\frac{r}{2} - \frac{r^5}{2}\right) \mathrm{d}r = \left(\frac{r^2}{4} - \frac{r^6}{12}\right)\bigg|_{r=0}^{1} = \frac{1}{4} - \frac{1}{12} = \frac{1}{6}.$$

Finally, integrating with respect to θ multiplies the result by 2π, and we get $\frac{1}{3}\pi$ as the answer.

We conclude that $\overline{z} = \frac{\pi/3}{\pi/2} = \frac{2}{3}$. Unsurprisingly, the paraboloid R_{sp} is top-heavy: its centroid $(0, 0, \frac{2}{3})$ is two-thirds of the way from the vertex at the bottom to the flat circle at the top!

[2] This may be new to some of you, and not to others; I think it's important to talk about, because it will provide intuition for many of the things we do later in the semester.

5 A triangular prism in cylindrical coordinates (optional)

So far, we've seen only simple cases of cylindrically-described regions: cases where z is bounded separately from (r, θ), and cases where θ is bounded separately from (r, z). These are very common, because they leverage the advantages of a cylindrical representation. This is not true of all shapes! Conceivably, we might want to suffer through describing a shape in cylindrical coordinates even when it doesn't look at all cylindrical. If you'd like to see a much trickier example of describing a shape in cylindrical coordinates, consider the triangular prism in Figure 4.

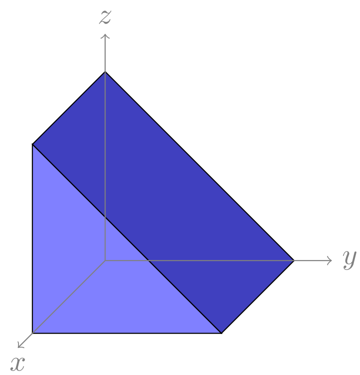

Figure 4: The triangular prism bounded by $x \geq 0$, $y \geq 0$, $z \geq 0$, $x \leq 1$, and $y + z \leq 1$.

Ultimately, we want to describe regions like this one so we can integrate over them. So the best way to describe it depends on the order in which we'll integrate. A common way to do a cylindrical integral is with z on the inside, r in the middle, and θ on the outside. This means we will want constant bounds on θ, bounds on r that depend only on θ, bounds on z that depend only on θ and r. We don't have to stick to this order, but it is often the easiest to work with.

The bounds $x \geq 0$ and $y \geq 0$ tell us that we're in the first quadrant, which corresponds to $0 \leq \theta \leq \frac{\pi}{2}$. This tells us our bounds on θ.

To determine our bounds on r, we want to ask: for a fixed direction θ, what is the *furthest* distance r we can go from the z-axis? Intuitively, if you want to stay inside the triangular prism in Figure 4, but get as far from the z-axis as possible, you want to do be on the xy-plane.

Here, the bounds on x and y are $0 \leq x \leq 1$ and $0 \leq y \leq 1$: the intersection of the prism with the xy-plane is a square. This means that $0 \leq r\cos\theta \leq 1$ and $0 \leq r\sin\theta \leq 1$, or in other words $0 \leq r \leq \frac{1}{\cos\theta}$ and $0 \leq r \leq \frac{1}{\sin\theta}$. That's two conditions on r: which do we use?

Trick question: we use both of them! We want both $0 \leq x \leq 1$ and $0 \leq y \leq 1$ to hold: otherwise, we're not inside the prism. There are two ways to say that we want both bounds on r to hold:

- Combine them into $0 \leq r \leq \min\{\frac{1}{\cos\theta}, \frac{1}{\sin\theta}\}$. The min will determine which bound is the one that restricts r more, and apply that one.
- Determine when each bound applies. For $0 \leq \theta \leq \frac{\pi}{4}$, the bound $0 \leq r \leq \frac{1}{\cos\theta}$ is the more restrictive one, and we use that. For $\frac{\pi}{4} \leq \theta \leq \frac{\pi}{2}$, the bound $0 \leq r \leq \frac{1}{\sin\theta}$ is the more restrictive one, and we use that.

The only constraints we haven't used yet are $z \geq 0$ and $y + z \leq 1$, or in other words, $0 \leq z \leq 1 - y$. Since we are free to use both r and θ when bounding z, this is quick: we can just write $0 \leq z \leq 1 - r\sin\theta$, because $y = r\sin\theta$.

If we wanted to find the volume of this triangular prism, and for some reason we wanted to use cylindrical coordinates to do so, we'd write the integral as

$$\int_{\theta=0}^{\pi/2} \int_{r=0}^{\min\{\frac{1}{\cos\theta}, \frac{1}{\sin\theta}\}} \int_{z=0}^{1-r\sin\theta} r \, \mathrm{d}z \, \mathrm{d}r \, \mathrm{d}\theta.$$