Lecture 14: optimization

Calculus I, section 10
November 1, 2022

Last time, we saw how to find maxima and minima (both local and global) of functions using derivatives. Today, we'll apply this tool to some real-life optimization problems. We don't really have a new mathematical concept today; instead, we'll focus on building mathematical models from a given problem so that we can apply our mathematical tools.

Here's a motivating example, which might even be useful in real life: if you want to throw something (or someone) as far as possible, what angle should you throw at?

Let's make some simplifying assumptions, to make the situation easier to model. We'll assume that we're starting at ground level, height $y = 0$; that the initial speed we can throw at is constant, independent of the angle θ we're throwing at; and that there is no air resistance, just gravity. How should we go about maximizing the distance we can throw?

Let's start by drawing a picture:

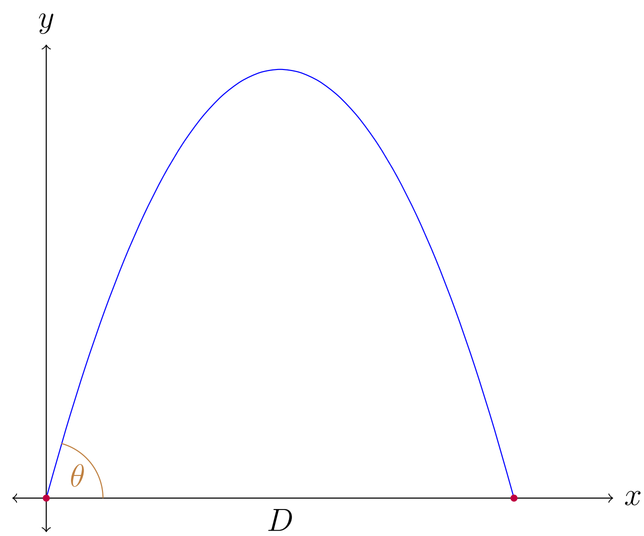

What we want to maximize is this quantity D; the only variable we can control is θ, so we want to think of D as depending on θ. How should we do this?

Well, D is the point at which whatever we're throwing hits the ground, so it actually depends on this extra quantity y, specifically on the point at which $y = 0$. I don't want to get too much into the physics and vectors and such, so let's just take the following as a black box: if we can throw with speed v, then the position of the object after time t will be given by

$$x(t) = v\cos(\theta)t, \qquad y(t) = -16t^2 + v\sin(\theta)t.$$

Thus $y = 0$ at $t = \frac{1}{16}v\sin(\theta)$. (For example, if $\theta = 0$ then the object hits the ground immediately; if $\theta = \frac{\pi}{2}$ then it takes the maximal amount of time, so $D = x(\frac{1}{16}v\sin(\theta)) = \frac{1}{16}v^2\sin(\theta)\cos(\theta)$.

We're now ready to do our optimization: the global maximum for D will be either were $\frac{dD}{d\theta} = 0$, where the derivative does not exist, or at the endpoints. Let's take these in reverse order: first of all, what are our endpoints? If we throw at $\theta = 0$, the object hits the ground immediately, so $D = 0$; this is the worst we can possibly do, and $\theta < 0$ we're throwing down so the same thing will happen. Therefore let's put one endpoint at $\theta = 0$. On the other hand, if we throw straight up, at $\theta = \frac{\pi}{2}$, although the object will be in the air for a relatively long time, it will also come down exactly where we threw it from, i.e. $D = 0$ again, so again this is a poor choice, and for $\theta > \frac{\pi}{2}$ we're actually throwing backwards, which we again discard. We expect the answer to be somewhere between 0 and $\frac{\pi}{2}$, so we'll take these as our endpoints; both give $D = 0$ as we saw, so if we get anything better from local maxima or non-differentiable points that will have to be the right answer.

In fact, we can compute

$$\frac{dD}{d\theta} = \frac{d}{d\theta}\frac{1}{16}v^2 \sin(\theta)\cos(\theta) = \frac{1}{16}v^2\left(\cos(\theta)^2 - \sin(\theta)^2\right)$$

by the product rule. This always exists, so there are no non-differentiable points. It is equal to zero if and only if $\cos(\theta)^2 = \sin(\theta)^2$, or equivalently if $\tan(\theta)^2 = 1$. For $0 \le \theta \le \frac{\pi}{2}$, we always have $\tan(\theta) \ge 0$, so this is the same thing as $\tan\theta = 1$, so $\theta = \tan^{-1}(1) = \frac{\pi}{4}$.

So all told, our guess is that the right answer is $\frac{\pi}{4}$, or 45°, so long as this gives an answer of more than zero: and indeed plugging it in to our formula for D gives

$$D = \frac{1}{16}v^2 \cdot \frac{\sqrt{2}}{2} \cdot \frac{\sqrt{2}}{2} = \frac{v^2}{32}.$$

This is not a terribly surprising result: 45° seems like the obvious compromise between $\theta = 0$ and $\theta = \frac{\pi}{2} = 90^\circ$.

What's interesting is we can now go back and use this same calculation to make our model a little more realistic. (Warning: this will get long and a bit messy, but some of the calculations can be skipped.) For example, in real life, you're rarely throwing something from height 0; more likely you're throwing from around the level of your shoulder, or maybe higher if you're elevated. We can fix our model to take this into account: now say you're starting from a height h. Then

$$x(t) = v\cos(\theta)t, \qquad y(t) = -16t^2 + v\sin(\theta)t + h.$$

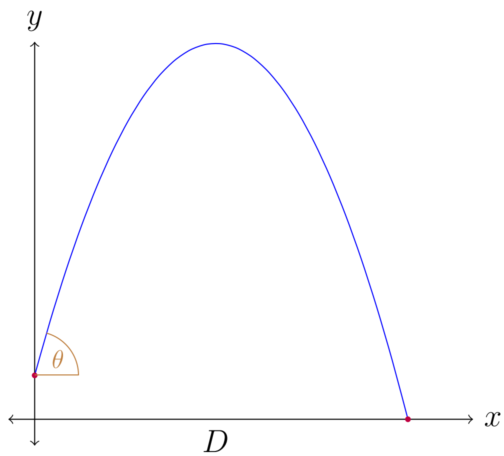

Now to find $y = 0$ we need the quadratic formula. You can imagine this formula is going to get messy once we start plugging things in; instead, let's just write T for the time at which $y(T) = 0$, and figure out how to make it explicit later. (Keep in mind that T also depends on θ.) Then we have

$$D = x(T) = v\cos(\theta)T,$$

so

$$\frac{dD}{d\theta} = v\left(\cos(\theta)\frac{dT}{d\theta} - \sin(\theta)T\right).$$

Again we can take $\theta = 0$ and $\theta = \frac{\pi}{2}$ as endpoints, since throwing down or backwards is never going to be better; in this case D at $\theta = 0$ will be nonzero, since it takes some time for y to drop to 0, but we still expect that the true maximum will be somewhere in the middle.

What we know about T is that for any θ, we have

$$y(T(\theta)) = -16T(\theta)^2 + v\sin(\theta)T(\theta) + h = 0.$$

We can differentiate this relationship with respect to θ:

$$-32T \cdot \frac{dT}{d\theta} + v\cos(\theta)T + v\sin(\theta) \cdot \frac{dT}{d\theta} = 0.$$

(Note: we can't just use the chain rule, because y depends both on $T(\theta)$ and directly on θ—this tripped me up for a minute while writing these notes!) Solving for $\frac{dT}{d\theta}$, we get

$$\frac{dT}{d\theta} = \frac{v\cos(\theta)T}{32T - v\sin(\theta)}.$$

Therefore all together we have $\frac{dD}{d\theta} = 0$ when

$$\frac{v\cos(\theta)^2T}{32T - v\sin(\theta)} = \sin(\theta)T,$$

which we can simplify to

$$32T\sin(\theta) = v(\cos(\theta)^2 + \sin(\theta)^2) = v$$

and so

$$\sin(\theta) = \frac{v}{32T}.$$

(Recall that when $h = 0$, we had $T = \frac{1}{16}v\sin(\theta)$, so in this case this formula would give $\sin(\theta) = \frac{v}{\frac{32}{16}v\sin(\theta)} = \frac{1}{2\sin\theta}$ and so $\sin(\theta)^2 = \frac{1}{2}$, which gives $\theta = \frac{\pi}{4}$ as before.)

Now we can solve for T and only have to use this big expression once: since $y(T) = -16T^2 + v\sin(\theta)T + h = 0$, by the quadratic formula (and since $T > 0$) we have

$$T = \frac{v\sin(\theta) + \sqrt{v^2\sin(\theta)^2 + 64h}}{32}.$$

All in all we get

$$\sin(\theta) = \frac{v}{v\sin(\theta) + \sqrt{v^2\sin(\theta)^2 + 64h}} = \frac{1}{\sin(\theta) + \sqrt{\sin(\theta)^2 + 64h/v^2}}.$$

Solving for $\sin(\theta)$ gives

$$\sin(\theta) = \frac{1}{\sqrt{2 + 64h/v^2}}$$

and so

$$\theta = \sin^{-1}\left(\frac{1}{\sqrt{2 + 64h/v^2}}\right).$$

In particular this makes clear that at $h = 0$ we recover $\sin^{-1}(\frac{1}{\sqrt{2}})$.

We can also get some nice concrete numbers out of this formula. For example, suppose you're throwing from a height of 5 feet at 40 feet per second. Then the optimal angle would be

$$\theta = \sin^{-1}\left(\frac{1}{\sqrt{2 + 64 \cdot 5/40^2}}\right) = \sin^{-1}\left(\frac{1}{\sqrt{2.2}}\right) \approx 0.73988 = 42.392^\circ.$$

Generally this also tells you that the higher you are, the closer to a flat angle you should throw at: as $h \to \infty$, we see that $\theta \to \sin^{-1}(0) = 0$. Nevertheless no matter how high you are you should always be aiming a little bit up, and the lower you are the closer to 45°. (Of course, we could keep adding complicating factors, such as variable throwing power or air resistance.)

Let's do another practical example, though less so in New York: driving. Let's say we're driving from point A to point B, which are some distance D miles apart by the most direct route, along which we can drive at 30 miles per hour. Alternatively, there's also a highway which goes to point B, along which we could drive at 60 miles per hour for a distance ℓ, but it doesn't pass through point A; instead it's r miles away, which we'd again have to drive at 30 miles per hour.

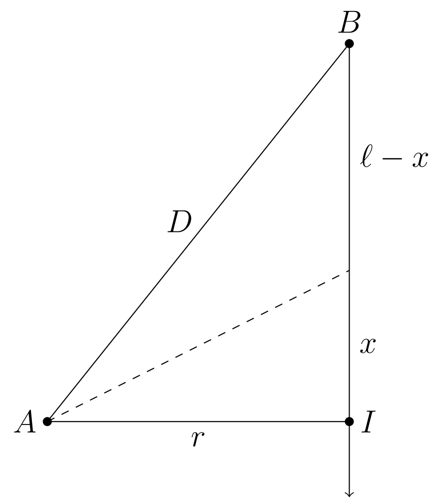

At what point x should we get onto the highway, if at all?

Here, what we want to minimize is time traveled. If we travel along the dashed path to get on the highway at distance x from the intersection I, this means traveling the distance of the dashed path, $\sqrt{r^2 + x^2}$, at 30 miles per hour, and the remaining distance $\ell - x$ at 60 miles per hour. Thus the total time is

$$t = \frac{1}{30}\sqrt{r^2 + x^2} + \frac{1}{60}(\ell - x).$$

First, let's think about endpoints. It doesn't make sense to travel away from B, since this makes both the distance to the highway and on the highway longer, so $x \geq 0$; and it doesn't make sense to overshoot B, so $x \leq \ell$. Let's first compute the total times at these endpoints: at $x = 0$, we drive straight to the highway and then straight along the highway to point B for total time

$$t = \frac{r}{30} + \frac{\ell}{60} = \frac{2r + \ell}{60}.$$

At $x = \ell$, we never get on the highway, we just drive straight to B at 30 miles per hour, and so

$$t = \frac{D}{30} = \frac{\sqrt{r^2 + \ell^2}}{30}$$

(since $r^2 + \ell^2 = D^2$ by the Pythagorean theorem). Which of these is faster depends on the particular values of r and ℓ (and so D). For example, if $r = \ell = 10$ miles (so $D = \sqrt{10^2 + 10^2} = 10\sqrt{2} \approx 14.14$ miles), then the $x = 0$ route gives $t = \frac{2\cdot 10 + 10}{60} = \frac{1}{2}$ hours, i.e. 30 minutes, while the $x = \ell$ route gives $t = \frac{10\sqrt{2}}{30} = \frac{\sqrt{2}}{3} \approx 0.47$ hours, or about 28 minutes, so in this case driving directly there is faster. If $r = 5$ miles and $\ell = 12$ miles (so $D = \sqrt{5^2 + 12^2} = 13$ miles), then the $x = 0$ route gives $t = \frac{2\cdot 5 + 12}{60} = \frac{11}{30}$ hours, or 22 minutes, while the $x = \ell$ route gives $t = \frac{13}{30}$ hours, or 26 minutes, so here the $x = 0$ route is faster.

What about in between? We have

$$\frac{dt}{dx} = \frac{2x}{60\sqrt{r^2 + x^2}} - \frac{1}{60},$$

which always exists assuming $r > 0$. It is equal to 0 when $2x = \sqrt{r^2 + x^2}$, i.e. when $4x^2 = r^2 + x^2$, so $3x^2 = r^2$. This gives $x = \frac{r}{\sqrt{3}}$ (we discard the negative solution since we assume $x \geq 0$, and in any case it wouldn't actually give a true solution to the original equation but instead is introduced by squaring).

There is something surprising here: this solution does not depend on the distance ℓ at all! This is a warning sign that for some values of ℓ, we should guess that the minimum is not at this local minimum $r/\sqrt{3}$ at all, but at one of the endpoints.

Let's find the value at this point to compare: first we travel a distance of $\sqrt{r^2 + r^2/3} = \frac{2r}{\sqrt{3}}$ at 30 miles per hour, and then a distance of $\ell - \frac{r}{\sqrt{3}}$ at 60 miles per hour for a total of

$$\frac{2r}{30\sqrt{3}} + \frac{\ell - \frac{r}{\sqrt{3}}}{60} = \frac{3r + \ell}{60\sqrt{3}}.$$

For given values of r and ℓ, we can try and figure out which out of this point or the endpoints is fastest. For example, in the case above with $r = 5$ and $\ell = 12$, we found that driving straight to the highway was faster than driving directly to point B, at 22 minutes instead of 26; for this middle point, we in this case get $\frac{3 \cdot 5 + 12}{60\sqrt{3}} \approx 0.26$ hours, or about 15.6 minutes, so it's faster than either of the extreme strategies.

On the other hand, what if we switch those parameters and take $r = 12$ miles and $\ell = 5$ miles? Then $x = \frac{12}{\sqrt{3}} \approx 6.928$ is greater than ℓ, and so isn't in the domain! Therefore the minimum distance must be at one of the endpoints: we can work out that the direct route will be faster, at $t = \frac{\sqrt{12^2+5^2}}{30} = \frac{13}{30}$ hours, or 26 minutes, as opposed to $t = \frac{2 \cdot 12 + 5}{60}$ hours, i.e. 29 minutes.

We could try to determine in general which of these three points will be better for which values of r and ℓ: $x = 0$ gives $\frac{2r+\ell}{60}$, $x = \ell$ gives $\frac{D}{30} = \frac{\sqrt{r^2+\ell^2}}{30}$, and the local minimum $x = \frac{r}{\sqrt{3}}$ gives $\frac{3r+\ell}{60\sqrt{3}}$. I claim that we always have $\frac{3r+\ell}{60\sqrt{3}} < \frac{2r+\ell}{60}$, i.e. we can throw out the $x = 0$ possibility. Indeed, we could write each side as a linear function of the positive numbers r and ℓ: $\frac{3}{60\sqrt{3}}r + \frac{1}{60\sqrt{3}}\ell$ and $\frac{2}{60}r + \frac{1}{60}\ell$. Since the coefficients on the left are smaller than on the right, the left-hand side must always be smaller.

This leaves two possibilities, and we've already seen that either can be the right minimum. If $\ell < r/\sqrt{3}$, then we must take the direct route $x = \ell$, since $x = r/\sqrt{3}$ is not within the allowed domain and will have to be slower. (We can also check that in this case, the $x = \ell$ case is still faster than than $x = 0$; this is left as an exercise if you feel like it.) If $\ell \geq r/\sqrt{3}$, then we might guess that the local minimum $x = r/\sqrt{3}$ is always fastest; with some algebra you can work out that this is in fact true.

Let's give one more quick example, our first business/economics application. Suppose we're a company making some kind of product; it costs us $C(x)$ to produce x products. In a very simple situation, this might just look like $C(x) = cx$ for some constant c, but usually there are some complicating factors; e.g. maybe there's some initial cost A to get off the ground, plus a per-product cost c, so $C(x) = A + cx$, or maybe there's some bulk savings which decreases $C(x)$ for x large, and so on. On the other hand we have income $I(x)$ from producing x products: again this might be linear in simple situations, $I(x) = px$ for p the price of each product, but there might also be complicating factors such as diminishing returns, advertising and sales relationships, etc. Overall what we care about is profit $P(x) = I(x) - C(x)$: how many products should we make to maximize profit?

Well, there are two obvious endpoints, $x = 0$ (do nothing, i.e. shut down; hopefully this isn't the best we can do) or $x = +\infty$ (make as many products as we can; this would be nice, but usually doesn't happen either since the market isn't infinite). Other than that, we look at the critical points where $\frac{dP}{dx} = \frac{dI}{dx} - \frac{dC}{dx} = 0$, i.e. $\frac{dI}{dx} = \frac{dC}{dx}$: where the marginal cost is equal to the marginal income.

For example, say we're making books, which has a startup cost of \$1000, a per-book cost of \$2, and a bulk discount such that our total cost function is $C(x) = 1000 + 2x - \frac{1}{10000}x^2$ in dollars for up to 5000 books (our publisher won't let us print them for free, after all). We can sell each book at \$10, but again we have diminishing returns so that $I(x) = 10x - \frac{1}{1000}x^2$. How many books should we produce?

Well, at $x = 0$ we've lost money: $I(0) - C(0) = -1000$, so this isn't a good choice; our limit is bounded at $x = 5000$, where we have $I(5000) - C(5000) = 16500$ in profit. Next we look for when the marginal cost and income are equal:

$$\frac{dI}{dx} = 10 - \frac{1}{500}x = \frac{dC}{dx} = 2 - \frac{1}{5000}x,$$

which occurs at $x = \frac{40000}{9} \approx 4444$. Here the profit is $P = 16777.78$, so it's best to stop at 4444 rather than print as many books as possible, in this case 5000.

For more complicated cost and income functions, you could imagine there could be multiple local maxima or minima, and we'd have to start worrying about which is which. Next time we'll see a few such more complicated examples before moving on to talk about some more applications of differentiation, this time to limits.