Lecture 13: extrema and critical points

Calculus I, section 10
October 27, 2022

Last time, we saw some first applications of differentiation, including a new concept, related rates. This time we'll introduce another new concept to open up a whole set of applications: using derivatives to find maxima and minima of functions.

Recall from the very first lecture we gave the following problem as an example of a question which is very hard if not impossible without calculus, but supposedly solvable in a reasonable way with calculus: what is the highest point on the graph of $y = f(x) = x^2 - x^4$?

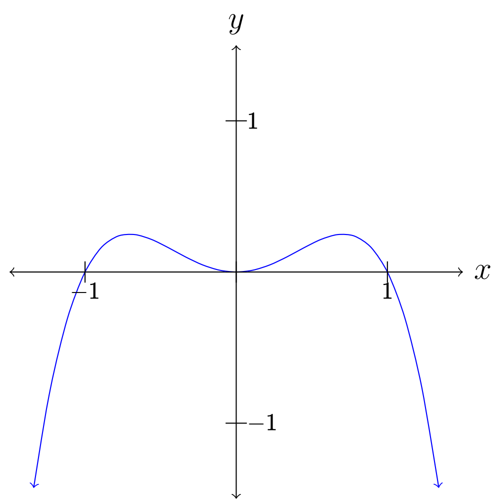

In other words, for which x is $x^2 - x^4$ largest?

From looking at the graph, we might guess that there are actually two such points, which are negatives of each other; let's pick one and call it a. How can we find a?

Well, one observation we can make is about the *derivative* of f near a. Since a is supposed to be a maximum (or at least a *local* maximum, i.e. $f(a)$ is greater than or equal to $f(x)$ for any x near a), if we pick x slightly greater than a then we must have $f(x) \leq f(a)$, so

$$\lim_{x \to a^+} \frac{f(x) - f(a)}{x - a} \leq 0.$$

On the other hand, if x is slightly less than a, we again have $f(x) \leq f(a)$ for the same reason, but now $x - a < 0$ and so

$$\lim_{x \to a^-} \frac{f(x) - f(a)}{x - a} \geq 0.$$

In order for f to be differentiable at a, these must be equal, i.e. $f'(a) = 0$.

This makes sense geometrically, too: for x a little less than a, since f has a maximum at a we expect f to be increasing at x, so $f'(x) > 0$; on the other hand for x a little more than a we expect f to now be decreasing, so $f'(x) < 0$. Since $f'(a)$ is in the middle, the maximum is the point where the derivative "crosses over," i.e. $f'(a) = 0$.

Now that we now this property of a, it's much easier to find it because there are only a few points where $f'(x) = 0$. We have $f'(x) = 2x - 4x^3$ by linearity, and then we can solve like for any polynomial: $2x - 4x^3 = 2x(1 - 2x^2)$, so either $x = 0$ or $1 - 2x^2 = 0$, i.e. $x = 0$ or $x = \pm\frac{1}{\sqrt{2}}$.

What's happening here? We got three solutions to $f'(x) = 0$, but we expect only two maxima. In fact, the same sort of thing happens at another kind of point: there is a local minimum at $x = 0$, and so by the same kind of argument $f'(0) = 0$. Therefore these solutions correspond to local maxima at $\pm\frac{1}{\sqrt{2}}$ and a local minimum at 0.

The $x = 0$ case illustrates the meaning of the word "local." We asked for the highest points on the graph, and got that they must be at $x = \pm\frac{1}{\sqrt{2}}$; we can plug these values into $f(x) = x^2 - x^4$ to get that they are both at $y = \frac{1}{\sqrt{2}^2} - \frac{1}{\sqrt{2}^4} = \frac{1}{2} - \frac{1}{4} = \frac{1}{4}$. However, if we'd asked for the minimum instead of the maximum, even though we find $x = 0$ is a local minimum, it is not a global minimum: as $x \to \pm\infty$, the function goes to $-\infty$, so there is no smallest value! The issue of when global maxima or minima occur is a more complicated one which we'll come back to, but the first step is always to find the local minima by solving $f'(x) = 0$.

Let's look at another example: what is the minimum of $f(x) = x + \frac{1}{x}$ for $x > 0$?

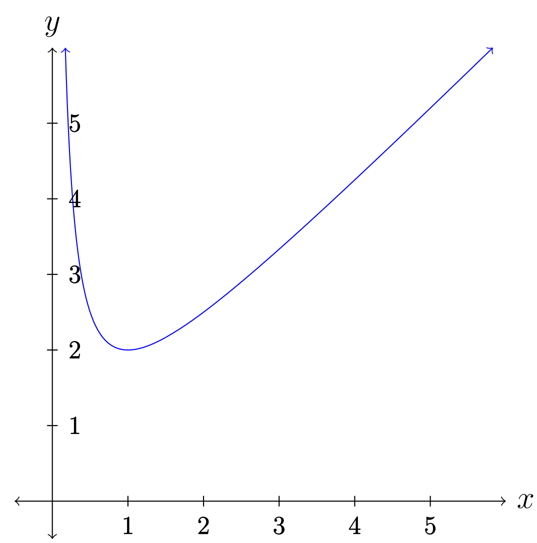

Well, we have

$$\frac{d}{dx}x + \frac{1}{x} = 1 - \frac{1}{x^2},$$

which is equal to 0 at $x = \pm 1$. Since we're restricting to $x > 0$, this means the only solution is at $x = 1$, which matches what we see on the graph. The minimum value is $f(1) = 1 + \frac{1}{1} = 2$.

What about the minimum value of $f(x) = x \cdot 2^x$?

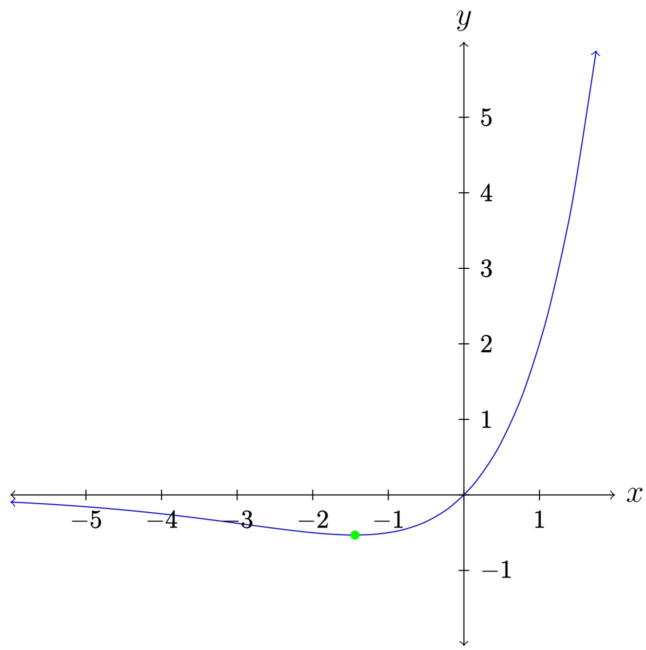

Differentiate: using the product rule,

$$f'(x) = 2^x + x \cdot 2^x \cdot \ln(2) = 2^x \cdot (1 + x \cdot \ln(2)),$$

so solving $f'(x) = 0$ we get $1 + x \cdot \ln(2) = 0$ and so $x = -\frac{1}{\ln(2)}$. The minimum value is then $f(-\frac{1}{\ln(2)}) = -\frac{1}{\ln(2)} \cdot 2^{-\frac{1}{\ln(2)}}$, which by a similar trick to on one of the homework problems is $-\frac{1}{e\ln(2)}$: we have $\frac{1}{\ln(2)} = \log_2(e)$, so $2^{-\frac{1}{\ln(2)}} = 2^{-\log_2(e)} = e^{-1}$.

A somewhat different kind of example is $f(x) = x^3$.

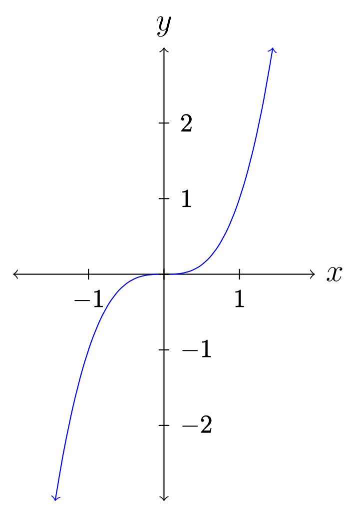

Just from looking at the graph, we can immediately see it doesn't have any maxima or minima: the function gets arbitrarily large in both the positive and negative directions.

Nevertheless, if we differentiate, $f'(x) = 3x^2 = 0$ does have a solution, at $x = 0$. What is happening here?

Nothing too terrible: we've just discovered a third possibility for when $f'(x) = 0$. Either this means that f has a local maximum at x, a local minimum at x, or neither. In this example, it is what's called an inflection point: the graph goes from concave, where the derivative is decreasing, to convex, where the derivative is increasing. In general, we call any point x where $f'(x) = 0$ a *critical point* or *stationary point* (because $f(x)$ is "not changing" at x, since the derivative is zero); local maxima and minima are special kinds of critical points. (Sometimes we'll use the word "extrema" to refer to critical points which are either maxima or minima, without specifying which.)

For example, consider $f(x) = x^4 - 2x^3$.

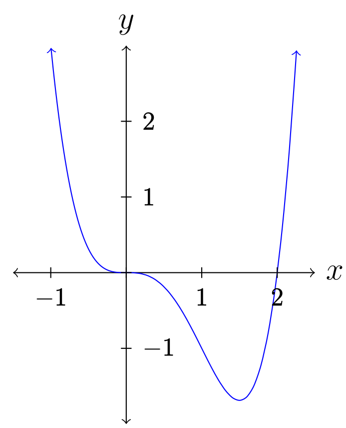

Solving $f'(x) = 4x^3 - 6x^2 = 0$, we see that f has two critical points, at $x = 0$ and $x = \frac{3}{2}$. It looks like the critical point at $x = \frac{3}{2}$ is a local minimum and the one at $x = 0$ is neither a minimum nor a maximum, but how can we know for sure?

One way is just to compare the values at all the critical points. However, this doesn't quite do it: $f(0)$ is greater than $f(\frac{3}{2})$, but although $\frac{3}{2}$ is a minimum this doesn't imply that 0 is a maximum: we'd need to compare the end behavior, too. We'll do this for finding global minima and maxima, but for now we want to avoid this: we're only looking for local minima and maxima, so we shouldn't have to worry about global stuff like end behavior.

Instead, we apply the *second derivative test*. What this means is this: suppose we have a local maximum at a. Our function is supposed to be increasing for $x < a$ and decreasing for $x > a$, so the derivative goes from positive to negative, passing through 0 at $x = a$. This means that the derivative is itself decreasing at a, and therefore $f''(a) \leq 0$!

Similarly, at a local minimum the derivative is going from negative to positive, and so $f''(a) \geq 0$. This gives us a test: given a point a at which f' is zero, to check if it's a local minimum or maximum we compute the second derivative $f''(a)$. If it's greater than zero, it must be a minimum; if it's less than zero, it must be a maximum; and if it's equal to zero, it could be either, or neither!

To resolve the ambiguity, let's think about what it would mean to have e.g. a minimum at a with $f''(x) = 0$. An example is $f(x) = x^4$ with $a = 0$; there's certainly a minimum at

0, but $f''(x) = \frac{d}{dx}4x^3 = 12x^2$ vanishes at $x = 0$. What's happening is that the derivative is negative for $x < 0$ and positive for $x > 0$, passing through 0 at $x = 0$, but since $f''(0) = 0$ the derivative itself has a critical point at $x = 0$, which is neither a maximum nor a minimum, since it goes from negative to positive. This means that not only must $f''(0) = 0$, but we also have to have $f'''(0) = 0$. The same thing applies to maxima: we know that f has a maximum (resp. minimum) at a if $f''(a) < 0$ (resp. $f''(a) > 0$), and if $f''(a) = 0$ it can still have a maximum (or minimum) only if $f'''(a)$ is also 0. (If $f'(a) = f''(a) = f'''(a) = 0$, the situation is again ambiguous, and we proceed to higher and higher derivatives until we get a definitive answer; we can talk about this more later if it comes up.)

Let's try applying this test to the example above. The second derivative is $f''(x) = \frac{d}{dx}f'(x) = 12x^2 - 12x = 12x(x-1)$, so at $x = \frac{3}{2}$ we have $f''(\frac{3}{2}) = 9 > 0$ and so this point is indeed a minimum. At $x = 0$, $f''(0) = 0$ and so the second derivative test is indeterminate, contributing to our suspicion that this is neither a maximum nor a minimum.

How could we know for sure? Well, we can try computing the third derivative: $f'''(x) = \frac{d}{dx}f'(x) = 24x - 12$, which at $x = 0$ is -12. Since either a maximum or minimum with $f''(0) = 0$ would have to have $f'''(0) = 0$ as well, this must be neither a maximum nor a minimum.

This is kind of a niche case and I won't make you deal with it too much, but it does sometimes come up in optimization problems, and we'll talk more about it for them.

There's one more kind of points that can be local maxima or minima without being stationary points, i.e. without having $f'(a) = 0$. These are points at which $f'(a)$ doesn't exist at all. For example, $f(x) = |x|$ has no points at which $f'(x) = 0$: the derivative is always ± 1. Nevertheless, it has a minimum at $x = 0$, where the derivative doesn't even exist.

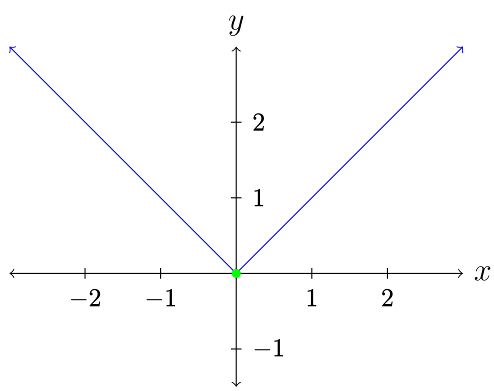

This is because our logic above that the maximum or minimum has to have derivative 0 depended on an implicit assumption that the derivative exists and is continuous: in order to go from negative to positive, or vice versa, we said the derivative had to pass through 0 at some point, which would be the extremum. But if the derivative fails to exist or to be continuous (which, recall, are basically the same thing), then it is possible to go from negative to positive without passing through zero, as in this example.

Thus if we want to find the *global* maxima or minima, there are at least two kinds of points we have to check: the stationary points, where $f'(x) = 0$, and any points where f is not differentiable, i.e. $f'(x)$ doesn't exist. Once we've found all these points, we can simply compute f at all of them and see which is largest or smallest.

In fact, that isn't quite the end of the story. There's another category of points we have to check: the endpoints. Often in applications, in practice our variable x has to be within some range. For example, maybe we're looking for the maximum value of $f(x) = 1 - x$ for $0 \leq x \leq 1$.

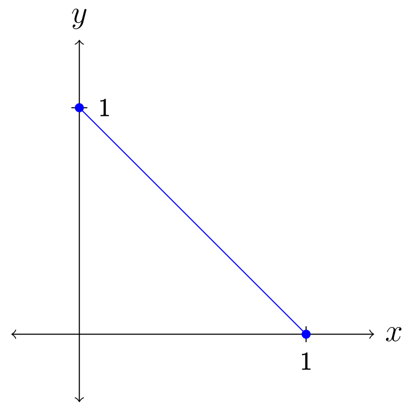

We have $f'(x) = -1$, so there are no stationary points and no points where f is not differentiable. Nevertheless, this function clearly has a maximum, namely at $x = 0$ where $f(0) = 1 - 0 = 1$. The idea is that in addition to the stationary and non-differentiable points, we also need to check the endpoints, if they exist.

What if they don't exist? For example, consider

$$f(x) = \tan^{-1}(x) + \frac{2}{x^2+1}.$$

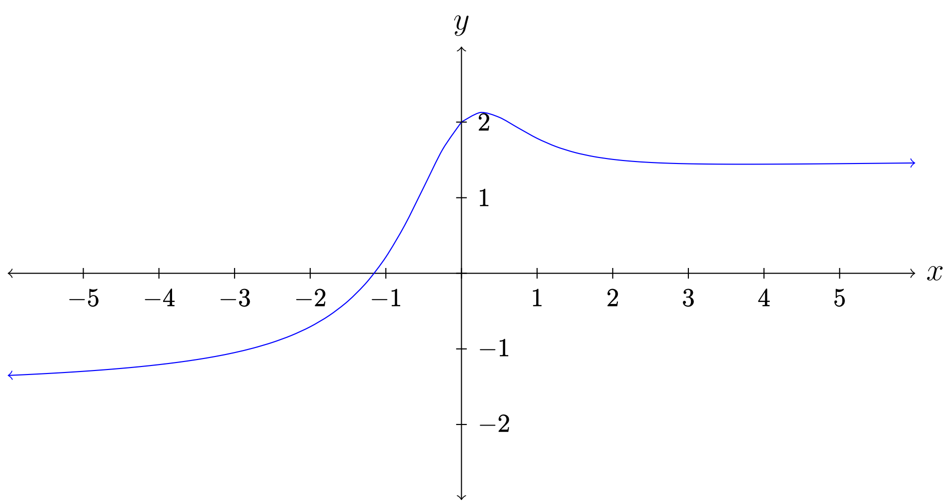

We can compute

$$f'(x) = \frac{1}{x^2+1} - \frac{4x}{(x^2+1)^2} = \frac{x^2-4x+1}{(x^2+1)^2},$$

which is 0 at the zeros of $x^2 - 4x + 1$, i.e. at $x = 2 \pm \sqrt{3}$. We can compute

$$f(2+\sqrt{3}) \approx 1.443, \qquad f(2-\sqrt{3}) \approx 2.128,$$

so we might be tempted to say that this function has a global maximum at $x = 2 - \sqrt{3}$ and a global minimum at $x = 2 + \sqrt{3}$. (We could also use the second derivative to check that these are actually local maxima and minima respectively.)

However, we can see from the graph that while a maximum at $x = 2 - \sqrt{3}$ looks possible, the value of 1.443 at $x = 2 + \sqrt{3}$ is clearly not the lowest value this function takes—when $x \to -\infty$, it becomes negative. This is what we should think of as our endpoint: instead of imposing one like 0 or 1 as above, here our endpoints are $\pm\infty$, and to check the values at these endpoints we take the limit:

$$\lim_{x\to+\infty} f(x) = \frac{\pi}{2}, \qquad \lim_{x\to-\infty} f(x) = -\frac{\pi}{2}.$$

Therefore $x = 2 - \sqrt{3}$ is a true global maximum, but $x = 2 + \sqrt{3}$ is not a global minimum since we get a lower value at the endpoint $-\infty$. Since $-\infty$ is not in the domain of the function, this means that *there is no global minimum*: for whatever number x we choose, we can always find another number y with $f(y) < f(x)$.

The case above with $f(x) = x + \frac{1}{x}$ for $x > 0$ is actually similar. We found a local minimum at $x = 1$, which is the global minimum (for this restriction), but there is no global maximum: the endpoints are 0 and $+\infty$, and $\lim_{x\to 0} f(x) = \lim_{x\to+\infty} f(x) = +\infty$, so no matter what x we choose we can always find some y with $f(y) > f(x)$.

Next time we'll see how to apply this technique to optimization problems, and talk more about some of these weirder cases.