

An Intuitive Primer on Calculus

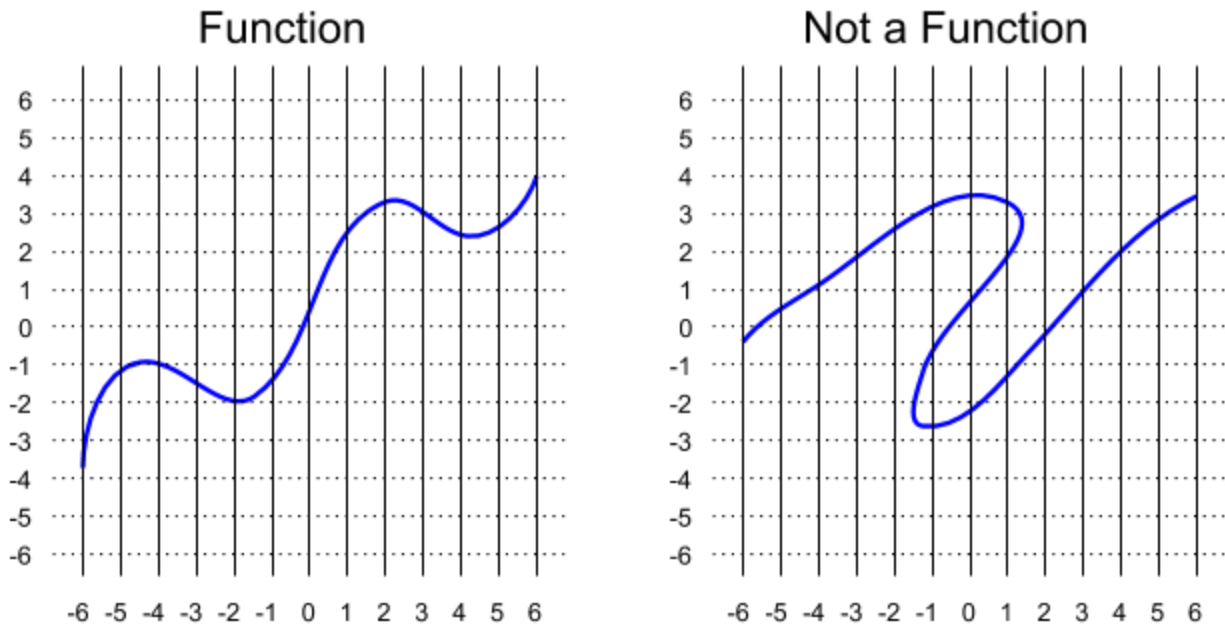
Justin Skycak, 2017

This primer explains the intuition behind the core concepts which tie together all of single-variable calculus. Understanding these core concepts will make it easier for you to understand the details, because you will see where they fit in the big picture of calculus.

Functions	1
Limits	3
Derivatives	7
Integrals	10
Sequences	12
Series	14

Functions

A function is a scribble that crosses each vertical line only once.



The instructions for drawing a function are given by the function's equation. Each vertical line is labeled with a number, and when you plug that number into the x variable in a function, the result tells you how high the function should be when it crosses that line.

For example, to draw the function whose equation is $y = 3x + 1$, we can start by plugging in the numbers -1, 0, and 1.

$$x = -1 \implies y = 3(-1) + 1 = -2$$



$$x = 0 \implies y = 3(0) + 1 = 1$$

$$x = 1 \implies y = 3(1) + 1 = 4$$

Our results mean:

- On the line $x = -1$, the scribble needs to cross at height $y = -2$.
- On the line $x = 0$, the scribble needs to cross at height $y = 1$.
- On the line $x = 1$, the scribble needs to cross at height $y = 4$.

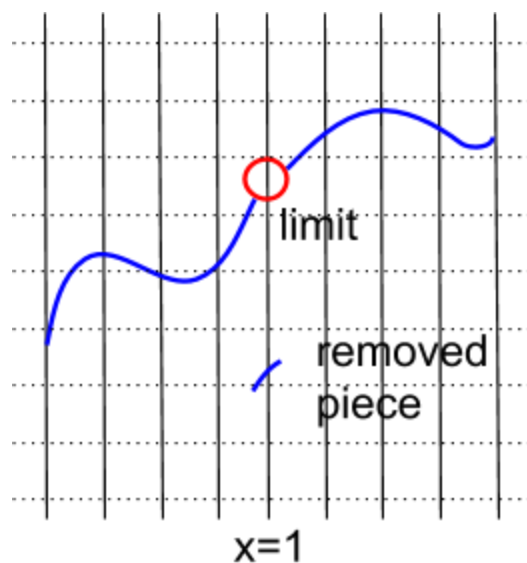
In other words, we need to plot the points $(-1, -2)$, $(0, 1)$, and $(1, 4)$ on the grid. Since this is a *linear* function, we can draw the rest of the function just by connecting the dots.




Linear functions are straight lines, but in general, functions tend to be curvy. There are many different types of functions, each of whose scribble curves in a different way.

Limits

The limit of a function is the height, or y -value, where it looks like the scribble hit a particular vertical line, or x -value -- but sometimes scribbles don't actually hit where they look like they're going to hit. For example, if you have a weird scribble with a broken piece removed, the limit at $x = 1$ is the height where the rest of the function suggested the piece was going to hit, not the "island" segment where the piece actually hits $x = 1$.



On the other hand, the function below actually occupies its limit at $x = 1$, because it hits where it looks like it's going to hit.



Functions which actually hit where they look like they're going to hit, like the one above, are called *continuous* functions. A rule of thumb: a function is continuous if you can draw it without picking up your pen.

A faster way to find a function's limit, without having to draw out the function, is to use the function's equation. If the function is continuous, then all you need to do is plug the x-value of the limit into the function. For example, to find the limit of $x^2 - 1$ at $x = 1$, you can simply plug in 1 for x :

$$\lim_{x \rightarrow 1} (x^2 - 1) = 1^2 - 1 = 0$$

But sometimes you're not able to plug in directly, and you have to manipulate the equation first. For example, if you try plugging in to $\lim_{x \rightarrow 1} \frac{x^2 - 1}{x - 1}$, you get $\frac{0}{0}$, which is undefined because of the 0 in the denominator. But if you manipulate the equation first, you can convert it into a form where you can plug in the limit:

$$\begin{aligned} \lim_{x \rightarrow 1} \frac{x^2 - 1}{x - 1} &= \lim_{x \rightarrow 1} \frac{(x + 1)(x - 1)}{x - 1} \\ &= \lim_{x \rightarrow 1} x + 1 \\ &= 1 + 1 \\ &= 2 \end{aligned}$$

Other times, the limit might not even exist! In $\lim_{x \rightarrow 1} \frac{1}{|x - 1|}$, you cannot plug in

because you'll get $\frac{1}{0}$, which is undefined. However, there is no way to manipulate the function as it is already in its most simplified form. When a limit doesn't exist, it means there is no number the function looks like it should go to at that x-value.

Sometimes this just happens because a function gets infinitely high or infinitely low. To discern this from the equation, it is helpful to think of plugging in x-values that are very close, but not equal to, the limit value. For example, in


$\lim_{x \rightarrow 1} \frac{1}{|x - 1|}$, you might try plugging in 0.999 instead of 1. Then you get

$$\frac{1}{|1 - 0.999|} = \frac{1}{0.001} = 1000$$

. If you try an x-value that is even closer, like 0.999999, then you get $\frac{1}{|1 - 0.999999|} = \frac{1}{0.000001} = 1000000$, which is an even bigger number. The function just keeps getting higher, the closer the x-value gets to 1.

The same thing happens if we approximate from the other side with numbers slightly larger than 1, such as 1.001 and 1.000001. So even though the limit

doesn't exist, we say that the limit is infinity. That's written as $\lim_{x \rightarrow 1} \frac{1}{|1 - x|} = \infty$.



Other times, though, the limit will look like something when we approximate with numbers slightly smaller than x, and something totally different when we approximate with numbers slightly bigger than x. In these cases, the limit just doesn't exist, and there's nothing we can do about it.

For an example, consider $\lim_{x \rightarrow 1} \frac{1}{1 - x}$, which is like the previous example but without the absolute value signs. If we try to approximate the limit using numbers slightly smaller than 1, the function just keeps getting higher and higher, to positive infinity.

$$x = 0.999 \implies \frac{1}{1-x} = 1000$$

$$x = 0.999999 \implies \frac{1}{1-x} = 1000000$$

But if we use numbers slightly bigger than 1, the function just keeps getting lower and lower, to negative infinity.


$$x = 1.001 \implies \frac{1}{1-x} = -1000$$

$$x = 1.000001 \implies \frac{1}{1-x} = -1000000$$


Since we get two different answers depending on whether we approach 1 from above or below, the limit does not exist.

Derivatives

The derivative tells the steepness of a function at a given point, kind of like a carpenter's level. For example, the derivative of a flat function is 0 because the bubble doesn't move when you hold up a level to the function:




However, if you hold up a level to a function that inclines upward, the bubble will move right, to a positive number. How big this positive number is, depends on how steeply the function is inclined. If the function is only a slight incline, like a wheelchair ramp, then the derivative will be small. However, if the function is steeply inclined, like a mountainside, then the derivative will be large.



Likewise, if you use a level to measure the derivative of a function that declines downward, the bubble will move left, to a negative number. Functions that are slightly declined will have slightly negative derivative, and functions that are steeply declined will have extremely negative derivative.

Now that we understand the intuition behind the derivative, we can offer a more precise definition: the derivative is the distance upward, per unit of distance


horizontal. For example, if you walk to the right one meter on a hill, you might go up two meters, so the derivative is 2.



The way to compute this mathematically is to start at your initial point, pick a point that's really close to it, compute the change in vertical distance, and then divide by the change in horizontal distance.


For example, in the function $y = x^2$, if we want to compute the derivative at the point $(2, 4)$, we can pick a nearby point $(3, 9)$ and compute the slope:

$$\frac{\text{vertical distance}}{\text{horizontal distance}} = \frac{9 - 4}{3 - 2} = \frac{5}{1} = 5$$



This isn't exactly right, but it's a good ballpark estimate. We can get a better estimate by picking a point that's even closer to the initial point: the slope

between $(2, 4)$ and $(2.1, 4.41)$ is $\frac{4.41 - 4}{2.1 - 2} = \frac{0.41}{0.1} = 4.1$.



And if we move the point even closer, then we get an even better approximation: the slope between $(2, 4)$ and $(2.01, 4.0401)$ is

$\frac{4.0401 - 4}{2.01 - 2} = \frac{0.0401}{0.01} = 4.01$.

Given our increasingly accurate approximations of 5, 4.1, and 4.01, you can probably guess that the exact derivative is 4. That is to say, the limit of the slope is 4, since the closer we move our second point, the closer the slope between the two points gets to 4.

It would be frustrating if we always had to guess the exact derivative by making approximations first. Conveniently, there are rules we can use to calculate the exact derivative quickly. For example, the power rule says that if you have a function x^n for some number n , the derivative is nx^{n-1} . Symbolically, that's


written $\frac{d}{dx}(x^n) = nx^{n-1}$. In the previous example, $\frac{d}{dx}(x^2) = 2x$, meaning that the derivative of x^2 is $2x$. We wanted to find the derivative when $x = 2$, so we plug in 2 to the derivative to get $2(2) = 4$.

There are other rules as well, for different types of functions.


Integrals

Integrals give the area under a portion of a function. For example, for the

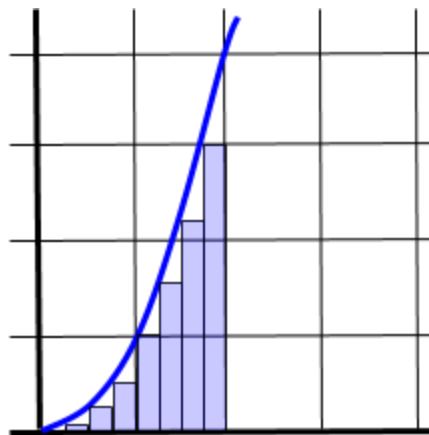

function $y = 2x$, the integral from $x = 0$ to $x = 2$, written $\int_0^2 2x$, is the area of the resulting triangle, $\frac{1}{2}(\text{base})(\text{height}) = \frac{1}{2}(2)(4) = 4$.



That integral was easy to compute, because we already knew the formula for area of a triangle. But what if we wanted to calculate the area of a curvy function like x^2 , where we don't have a formula for the area?



One way to proceed is to approximate the area by breaking it into a bunch of rectangles and adding up the areas of all the rectangles. The more rectangles, the better the approximation.



It turns out that taking the limit of infinitely many rectangles reveals a slick trick to get the exact area. To use this trick on the function x^2 , we need only follow three steps:

1. Work backwards to find the function whose derivative is x^2 (it's $\frac{1}{3}x^3$)
2. Evaluate $\frac{1}{3}x^3$ at $x = 0$ and $x = 2$ (to get 0 and $\frac{8}{3}$)
3. Take the difference (to get $\frac{8}{3}$).

The toughest part of the process is usually the first step, which involves working backwards to find the opposite of the derivative, called the antiderivative. However, there are rules we can use to complete this step quickly -- for example, the power rule says that the integral of any function x^n , where n is

some number, is $\frac{1}{n+1}x^{n+1}$. You can verify that this rule is true because if you differentiate the function $\frac{1}{n+1}x^{n+1}$, you get $\frac{1}{n+1}(n+1)x^n$, which comes out to just x^n .

There are other rules as well, for different types of functions.

Sequences

A sequence is a list of numbers that has some pattern. For example, $2, 4, 6, 8, \dots$ is the sequence of even numbers, and $1, 3, 5, 7, \dots$ is the sequence of odd numbers.

We can interpret sequences as functions evaluated at whole numbers. For example, the sequence of even numbers is like a function $y = 2x$, where x is evaluated at $1, 2, 3, 4, \dots$. The result is $2(1), 2(2), 2(3), 2(4), \dots$ which simplifies to $2, 4, 6, 8, \dots$.

Because we're only plugging in whole numbers for x , we usually use the term n instead of x , and a_n instead of y . So the sequence of even numbers is given by the rule $a_n = 2n$, as follows:

$$a_1 = 2(1) = 2$$

$$a_2 = 2(2) = 4$$

$$a_3 = 2(3) = 6$$

$$a_4 = 2(4) = 8$$

\vdots

To check your understanding, observe that the sequence of odd numbers is given by $a_n = 2n - 1$:

$$a_1 = 2(1) - 1 = 2 - 1 = 1$$

$$a_2 = 2(2) - 1 = 4 - 1 = 3$$


$$a_3 = 2(3) - 1 = 6 - 1 = 5$$

$$a_4 = 2(4) - 1 = 8 - 1 = 7$$


\vdots

The sequences of odd and even numbers keep getting bigger as they go on, but not all sequences are like that. For example, consider the sequence $a_n = \frac{1}{n}$. This

sequence, which starts off $1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \dots$ is a sequence of fractions that not only keeps on getting smaller. Not only does this sequence get smaller, but it also gets closer and closer to a single number as the sequence goes on. That number is zero -- we say that the sequence converges to zero.



Likewise, consider the sequence $a_n = \frac{n-1}{n}$. This sequence is an increasing sequence that gets closer and closer to 1, so we say it converges to 1.



You can usually tell whether a sequence converges by thinking about what happens when you plug in very large numbers for n . For example, the sequence

$a_n = \frac{1}{n}$ converges to 0, and plugging in a million for n gives

$a_{1000000} = \frac{1}{1000000} = 0.000001$, which is really close to 0. Likewise, the sequence

$a_n = \frac{n-1}{n}$ converges to 1, and plugging in a million for n gives

$a_{1000000} = \frac{999999}{1000000} = 0.999999$, which is really close to 1.


Series

A series is the sum of a sequence. For example, for the sequence of even numbers from 0 to 10, the series is given by $0 + 2 + 4 + 6 + 8 + 10 = 30$. It would make no sense if we tried to add up infinitely many of the even numbers, but sometimes it is actually possible to add up infinitely many numbers and still get a sensible answer.

For example, consider the series $\frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \dots$, where the next term is always half the preceding term.

$$\begin{aligned}\frac{1}{2} &= \frac{1}{2} \\ \frac{1}{2} + \frac{1}{4} &= \frac{3}{4} \\ \frac{1}{2} + \frac{1}{4} + \frac{1}{8} &= \frac{7}{8} \\ \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} &= \frac{15}{16}\end{aligned}$$

Each time we add the next term in the sequence, it's like shading in half of what's left in the whole circle. The more numbers we add, the closer we get to filling in the whole circle.



Since we get closer and closer to filling in the whole circle as the series goes on, we say that the series converges to 1. The takeaway of this example is that when terms in a series get small enough fast enough, their sum might not be infinite.

It's easy to see that series only make sense when the terms converge to zero (otherwise, you could surpass any number by summing enough terms).

However, just because a series' terms converge to zero, does not necessarily mean that you can add them all up and get a finite answer. For example, the

series $\frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \frac{1}{5} + \dots$ sums to infinity, or *diverges*.

It is often difficult to tell whether a series converges or diverges just by looking at it. However, there are some tests you can apply to a series to tell whether it converges or diverges. Many of these tests involve computing limits or integrals.