MATH 521, WEEK 15:
Sequences and Series of Functions

There are many applications where we are interested in the limit of some sequence or series, but where the objects are not points in $\mathbb{R}^n$ (or some abstract metric space X) but are *functions* themselves. This is particularly prevalent, for instance, in the study of *differential equations*, where the exact form of the solution may be unknown (or too complicated to work with) but may be known to be the limit of a sequence of *approximate* solutions.

It is also something we have seen previously in Calculus when we studied *Taylor series*. For instance, we were able to convince ourselves that

$$e^x = \sum_{n=0}^{\infty} \frac{x^n}{n!} = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \cdots . \tag{1}$$

But what exactly did this mean? After all, we cannot compute an infinite number of functions in a finite amount of time. Similarly to partial sums, we reasoned that what we were really interested in was the sequence of functions $\{f_n\}$, $n \in \mathbb{N}$, defined by

$$f_n(x) = \sum_{k=0}^{n} \frac{x^k}{k!} = 1 + x + \frac{x^2}{2!} + \cdots + \frac{x^n}{n!}.$$

So what we really meant by (1) was that

$$\lim_{n \to \infty} f_n(x) = e^x.$$

This is still not fully rigorous, but it will serve as a good template for what we want to study now. Namely, we want to consider:

1. What does it mean for a sequence $\{f_n\}$ of *functions* to converge to another function f?
2. Does the function we are converging toward share the same properties (e.g. continuity, differentiability, etc.) as the functions which are converging?

1 Pointwise Convergence

Consider the following sequence of functions $\{f_n\}$, $n \in \mathbb{N}$, where $f_n : [0,1] \mapsto \mathbb{R}$ are defined by:

$$f_n(x) = \begin{cases} 1 + nx, & \text{for } -\frac{1}{n} < x \leq 0 \\ 1 - nx, & \text{for } 0 < x \leq \frac{1}{n} \\ 0, & \text{otherwise} \end{cases}$$

We want (eventually) consider what happens as $n \to \infty$. Before jumping into this, however, we consider what the individual functions $f_n(x)$ look like. We can quickly see that, in the interval $-1/n < x \leq 0$, at the end points we have

$$f(-1/n) = 0$$
$$f(0) = 1.$$

The interval is similar in the region $0 < x \leq 1/n$. That is, the function looks like a series of "tents" where the width is determine inversely by the index $n \in \mathbb{N}$ (see Figure 1).

We now want to ask the question of what happens as $n \to \infty$. That is, we want to make sense out of an expression like $\lim_{n\to\infty} f_n(x)$. The way we will do this will be similar to what we would do with power series. We do not have a general notion of convergence for functions, but for any given value of x, the sequence simplifies to a sequence of *numbers*. We are very good at determining convergence properties of sequences of numbers.

Consider the following definition. (We will restrict to functions on the real numbers, but obvious extensions may be made.)

Definition 1.1. *Suppose $f_n : [a,b] \mapsto \mathbb{R}$ for $n \in \mathbb{N}$, and $f : [a,b] \mapsto \mathbb{R}$. We say that f_n* ***converges pointwise*** *to f on a subset $[a,b]$ if, for every $x \in [a,b]$,*

$$\lim_{n\to\infty} f_n(x) = f(x).$$

Let's reconsider our example. We consider the convergence for each point $x \in [-1,1]$ *independently.* For instance, if we consider the point $x = -1/5$, we generate the sequence

$$\{f_n(-1/5)\} = \left\{\frac{4}{5}, \frac{3}{5}, \frac{2}{5}, \frac{1}{5}, 0, 0, 0, \ldots\right\}$$

so that

$$f(-1/5) = \lim_{n\to\infty} f_n(-1/5) = 0.$$

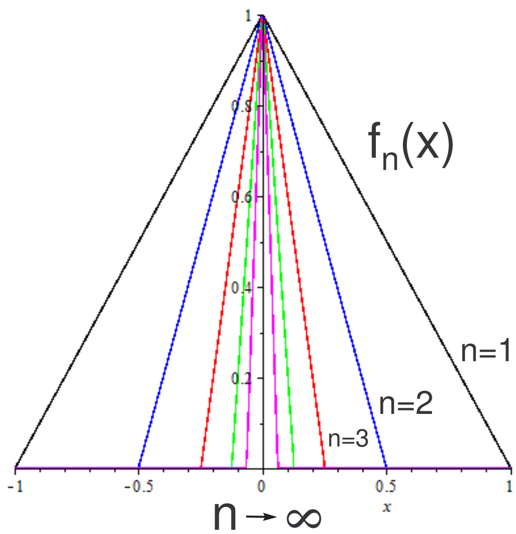

Figure 1: Graphs of the functions $f_n(x)$ for various $n \in \mathbb{N}$. As n increases, the "tents" become more peaked around $x = 0$, with the "tentpole" at the point $f(0) = 1$.

Now consider an arbitrary $x \in [-1, 0)$. We can see that, for any such choice, there is an $N \in \mathbb{N}$ such that, for all $n \geq N$ we have $-1 < x < -1/n < 0$. After this point, we are no longer in the domain $(-1/n, 0]$ and the function returns the value zero. The same holds for all $x \in (0, 1]$. It follows that

$$\lim_{n\to\infty} f_n(x) = 0$$

for all $x \in [-1, 0) \cup (0, 1]$. But what about $x = 0$? At this point, the intuition is different. For every $n \in \mathbb{N}$, $x = 0$ is contained in the domain $(-1/n, 0]$ so that $f_n(0) = 1$ for all $n \in \mathbb{N}$. It follows that the pointwise limit at $x = 0$ is 1. In summation, we have that

$$\lim_{n\to\infty} f_n(x) = f(x)$$

where $f : [-1, 1] \mapsto \mathbb{R}$ is defined by

$$f(x) = \begin{cases} 1, & \text{for } x = 0 \\ 0, & \text{otherwise.} \end{cases}$$

So we can define a limiting process for sequences of functions $\{f_n\}$ by simply considering the convergence independently at individual points in the domain. The previous example should cause some warning bells to sound, however. We have a sequence which converges at every point $x \in [-1, 1]$ but for which the limiting function is of a fundamentally different class of function. Namely, the limiting function is *discontinuous* while the functions in the sequence are all *continuous*.

Another way to think about this is that, for sequences of functions $\{f_n(x)\}$ which converge pointwise to a function $f(x)$, we have, in general, that

$$\lim_{n\to\infty} f_n(a) = \lim_{n\to\infty} \lim_{x\to a} f_n(x) \neq \lim_{x\to a} \lim_{n\to\infty} f_n(x) = \lim_{x\to a} f(x).$$

That is, we may not generally exchange the limiting process over the sequence of functions with the limiting process at a point (which is equivalent to continuity at a).

2 Uniform Convergence

We want to develop a stronger notion of convergence for sequences of functions. In particular, we want to develop a notion which will preserve desirable properties of the sequence functions, particularly continuity. Consider the following.

Definition 2.1. *Suppose $f_n : [a, b] \mapsto \mathbb{R}$ for $n \in \mathbb{N}$, and $f : [a, b] \mapsto \mathbb{R}$. We say that f_n* ***converges uniformly*** *to f on a $[a, b]$ if for every $\epsilon > 0$, there is a $N \in \mathbb{N}$ such that $n > N$ implies $|f_n(x) - f(x)| < \epsilon$ for all $x \in [a, b]$.*

The difference between *pointwise* and *uniform* convergence is similar to the difference between continuity at a point and uniform continuity. We must be able to select our small parameter (in this case, $\epsilon > 0$) independent of where we are on the domain. In fact, we cannot! To see this, consider the points near $x = 0$ (but not $x = 0$ itself).

Example 1 (revisited): For the previous example sequence $f_n : [0, 1] \mapsto \mathbb{R}$ where

$$f_n(x) = \begin{cases} 1 + nx, & \text{for } -\frac{1}{n} < x \leq 0 \\ 1 - nx, & \text{for } 0 < x \leq \frac{1}{n} \\ 0, & \text{otherwise} \end{cases}$$

we saw that we had $\lim_{n\to\infty} f_n(x) = f(x)$ where

$$f(x) = \begin{cases} 1, & \text{for } x = 0 \\ 0, & \text{otherwise} \end{cases}$$

and the limit was obtained pointwise.

In order to ascertain whether the convergence is uniform on the domain $[-1, 1]$ we take $\epsilon > 0$ and check whether there is an N such that $n \geq N$ implies $|f_n(x) - f(x)| < \epsilon$ for the whole domain. In fact, we can quickly determine that this is *not* the case! To see this, consider the points near $x = 0$ (but not $x = 0$ itself). We know that, for every $x \in [-1, 0)$,

$$\lim_{n\to\infty} f_n(x) = 0.$$

However, for any $0 < \epsilon < 1$ and any $n \in \mathbb{N}$, we can pick an $x \in [-1, 0)$ so that $f_n(x) > \epsilon$. Explicitly, we have

$$1 + nx > \epsilon \quad \Longrightarrow \quad \frac{\epsilon - 1}{n} < x < 0.$$

Intuitively, this corresponds to the following observation: for any given function $f_n(x)$, near $x = 0$ the function $f_n(x)$ must pass through all values between 0 and 1 to preserve continuity, so there must always be points which are between 0 and 1 from the eventual limit of zero. This is enough to destroy as possibility of uniform continuity. □

Example 2: Prove that the sequence of functions $\{f_n\}$, $n \in \mathbb{N}$, given by

$$f_n(x) = \frac{x}{1 + nx^2}$$

converges uniformly to $f(x) = 0$. (See Figure 2.)

Solution: Take $\epsilon > 0$. We need to bound $|f_n(x) - f(x)|$ by ϵ over the whole domain $\mathbb{R}$. In this case, it is sufficient to bound $|f_n(x)|$, since $f(x) = 0$ for all $x \in \mathbb{R}$. Since the function is differentiable, we will perform a little standard calculus work. First of all, we have that

$$f_n'(x) = \frac{1 - nx^2}{(1 + nx^2)^2}$$

so that $f_n'(x) = 0$ gives $x = \pm\frac{1}{\sqrt{n}}$. We can also quickly check that

$$\begin{aligned} f_n'(x) > 0, &\qquad \text{for } -\tfrac{1}{\sqrt{n}} < x < \tfrac{1}{\sqrt{n}} \\ f_n'(x) < 0, &\qquad \text{for } x < -\tfrac{1}{\sqrt{n}} \text{ and } x > \tfrac{1}{\sqrt{n}}. \end{aligned}$$

Since

$$\lim_{x\to\pm\infty} f_n(x) = 0$$

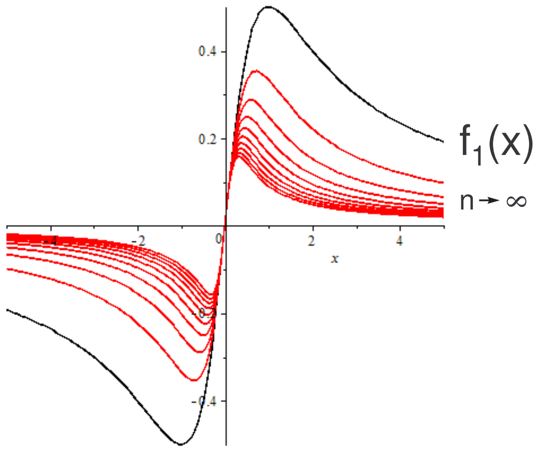

Figure 2: Sequence $f_n(x) = \frac{x}{1+nx^2}$ which converges uniformly to $f(x) = 0$.

we can conclude that $x = -\frac{1}{\sqrt{n}}$ is a global minimum and $x = \frac{1}{\sqrt{n}}$ is a global maximum. We can quickly evaluate that

$$f_n\left(-\frac{1}{\sqrt{n}}\right) = -\frac{1}{2\sqrt{n}}$$

and

$$f_n\left(\frac{1}{\sqrt{n}}\right) = \frac{1}{2\sqrt{n}}$$

so that $f_n(x) \in \left[-\frac{1}{2\sqrt{n}}, \frac{1}{2\sqrt{n}}\right]$ for all $x \in \mathbb{R}$.

This was a little bit of work that very likely seems disconnected from original problem; however, we can now see that, for every $n \in \mathbb{N}$,

$$|f_n(x) - f(x)| = \left|\frac{x}{1+nx^2}\right| \leq \frac{1}{2\sqrt{n}}.$$

Importantly, we can bound this by ϵ so long as we choose a sufficiently large n! Explicitly, if we choose $N > \frac{1}{4\epsilon^2}$ then, for all $n > N$ and all $x \in \mathbb{R}$, we have

$$|f_n(x) - f(x)| \leq \frac{1}{2\sqrt{n}} < \epsilon.$$

That is, the sequence $\{f_n(x)\}$ converges uniformly to $f(x)$ on $\mathbb{R}$, and we are done. □

Naturally, in practice we may not always know what the function $f(x)$ looks like or have an explicit form for it. Nevertheless, we would like to be able to say that the sequence of functions converges to something. It is necessarily, therefore, to introduce the following **Cauchy criterion** for sequence convergence.

Theorem 2.1 (7.8 in Rudin)**.** *The sequence of functions* $\{f_n\}$ *where* $f_n : [a,b] \mapsto \mathbb{R}$ *for all* $n \in \mathbb{N}$ *converges uniformly to some function* $f : [a,b] \mapsto \mathbb{R}$ *if and only if, for every* $\epsilon > 0$*, there is an* $N \in \mathbb{N}$ *such that, if* $n, m \geq N$ *then* $|f_n(x) - f_m(x)| < \epsilon$ *for every* $x \in [a,b]$*.*

Proof. Suppose that $f_n(x)$ converges to some function $f(x)$. Take $\epsilon > 0$. It follows that from the convergence of $f_n(x)$ to $f(x)$ that there is an $N \in \mathbb{N}$ such that, if $n, m \geq N$, then

$$\begin{aligned} |f_n(x) - f(x)| &< \frac{\epsilon}{2} \\ |f_m(x) - f(x)| &< \frac{\epsilon}{2} \end{aligned}$$

for all $x \in [a,b]$. It follows immediately that

$$|f_n(x) - f_m(x)| \leq |f_n(x) - f(x)| - |f(x) - f_m(x)| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

That is, the sequence $\{f_n\}$ satisfies the Cauchy criterion.

Now suppose that $\{f_n\}$ satisfies the Cauchy criterion. Take $\epsilon > 0$. We have that, there is an $N \in \mathbb{N}$ such that $n, m \geq N$ implies $|f_n(x) - f_m(x)| < \epsilon$ for every $x \in [a,b]$. That is, $\{f_n\}$ is a Cauchy sequence in $\mathbb{R}$, and therefore converges. Call this convergence point $f(x)$ and note that, after taking the limit $m \to \infty$, we have that $|f_n(x) - f(x)| < \epsilon$ for all $x \in [a,b]$ so that the convergence is uniform. □

Note: Another obvious but important characterization of uniform continuity is that

$$\sup_{x \in [a,b]} |f_n(x) - f(x)|$$

must go to zero as $n \to \infty$. This may seem like mundane point, but it will actually be a key observation in extending the notion of a metric space to functional spaces.

3 Uniform Convergence and Continuity

Of course, we notice something different about convergence in Example 2 compared to Example 1. Namely, the limit of the uniformly convergent sequence of functions in Example 2 converged to a *continuous* function ($f(x) = 0$) while the pointwise convergence sequence of functions in Example 1 did not ($f(x)$ has a discontinuity at $x = 0$).

This is, in fact, a general result.

Theorem 3.1 (7.11 and 7.12 in Rudin)**.** *Suppose* $\{f_n\}$ *where* $f_n : [a,b] \mapsto \mathbb{R}$*,* $n \in \mathbb{N}$*, is a sequence of continuous functions which converges uniformly to* $f : [a,b] \mapsto \mathbb{R}$*. Then* $f(x)$ *is continuous on* $[a,b]$*.*

Proof. Take $\epsilon > 0$. We need to show $f(x)$ is continuous at every $c \in [a,b]$.

Now expand our given information. We have that the sequence $\{f_n\}$ converges uniformly to f on $[a,b]$. It follows that there is an $n \in \mathbb{N}$ such that, for all $x \in [a,b]$, we have

$$|f_n(x) - f(x)| < \frac{\epsilon}{3}.$$

We also have that $f_n(x)$ is continuous at $c \in [a,b]$ for the chosen $n \in \mathbb{N}$. It follows that there is a $\delta > 0$ such that $|x - c| < \delta$ implies

$$|f_n(x) - f_n(c)| < \frac{\epsilon}{3}.$$

We now piece everything together. We can choose $n \in \mathbb{N}$ and $\delta > 0$ as above to show that, if $|x - c| < \delta$, then

$$\begin{aligned} |f(x) - f(c)| &\leq |f(x) - f_n(x)| + |f_n(x) - f_n(c)| + |f_n(c) - f(c)| \\ &< \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon. \end{aligned}$$

That is, $f(x)$ is continuous at $c \in [a,b]$, and we are done. □

It is worth noting in this proof that we only need to select a single $n \in \mathbb{N}$ where the bound works, rather than having it hold for all $n \geq N$. That is, we do not need the full generality afforded to us by uniform convergence.

4 Metric Spaces of Continuous Functions

We may naively look at the title of this section, and think we are repeating our earlier work. After all, we know we can define a function $f : X \mapsto Y$

where (X, d_X) and (Y, d_Y) are two distinct metric space. We also know that we can understand continuity in this context by considering the metrics separately in X and Y.

But what we want to consider now is a metric space where the *points* in our space are continuous functions themselves. This shift of focus should not come as a complete surprise. After all, why bother with the generality of metric spaces as abstract sets if you are only ever going to deal with vectors?

In fact, we have already briefly seen that there is a way to measure the "distance" between two objects which are a little outside of our typical understand. Namely, we saw that *Hausdorff metric* was a metric on the space of closed intervals in $\mathbb{R}$. That is, we were able to come up with a way to quantify how close or far apart two closed intervals in $\mathbb{R}$ were. So how do we measure how close or far apart two *continuous functions* are?

We start by formally define the set in which we will be interested.

Definition 4.1. *Define $\mathscr{C}[a,b]$ to be the set all functions $f : [a,b] \mapsto \mathbb{R}$ which are continuous on $[a,b]$.*

There is no significant subtlety here. On the interval $[0,1]$, $f(x) = e^x$ and $g(x) = \sin(x)$ are examples of functions which are continuous on $[0,1]$ and therefore $f, g \in \mathscr{C}[0,1]$. The function $h(x) = 1/x$ is not continuous on $[0,1]$ and therefore $h \notin \mathscr{C}[0,1]$ (but $h \in \mathscr{C}[1,2]$, for instance).

The point of this whole process is that we have defined a *set* of functions. This is the first half of our metric space. The remaining piece is to determine our metric—i.e. to assign a "distance" between each pair of functions in the set $\mathscr{C}[a,b]$. Consider the following.

Theorem 4.1. *The pairing $(\mathscr{C}[a,b], d_\infty)$ where*

$$d_\infty(f,g) = \sup_{x \in [a,b]} |f(x) - g(x)|$$

is a metric space.

Note: All we are saying here is that it is sufficient to consider the largest variance between $f(x)$ and $g(x)$ on the domain $[a,b]$. This is an admittedly crude, but still somewhat sensible measure of how "close" two functions are. After all, for a small value ϵ, $f(x)$ and $g(x)$ can be no farther apart than ϵ on the entire domain $[a,b]$. So they are definitely "close". At the same time, however, we may wish to distinguish between two functions which differ by ϵ on only a very small set, and those which vary by ϵ, say, *everywhere*. This notion does not allow us to do that. Nevertheless, we can quickly verify that it is a metric space.

Proof. We just need to verify the metric space axioms. (1) and (2) are trivial, so we will exclude them. We also notice that, for any $x \in [a, b]$, we have

$$|f(x) + g(x)| \leq |f(x)| + |g(x)|$$

by the triangle inequality for $\mathbb{R}$. It follows that

$$\sup_{x\in[a,b]} |f(x) + g(x)| \leq \sup_{x\in[a,b]} |f(x)| + \sup_{x\in[a,b]} |g(x)|.$$

Making a slight switch in interpretation (as we have done before when consider metric spaces, in particular with Euclidean metrics), we have that

$$\begin{aligned} d_\infty(f, g) &= \sup_{x\in[a,b]} |f(x) - g(x)| \\ &= \sup_{x\in[a,b]} |f(x) - h(x) + h(x) - g(x)| \\ &\leq \sup_{x\in[a,b]} |f(x) - h(x)| + \sup_{x\in[a,b]} |h(x) - g(x)| \\ &= d_\infty(f, h) + d_\infty(h, g). \end{aligned}$$

That is, the distance from any function f to any function g, cannot be greater than the distance through any intermediate function h. Although we would be hard pressed to correspond this to any geometric picture of the triangle inequality, this is exactly what we have just proved for the metric space $(\mathscr{C}[a, b], d_\infty)$, so we are done. □

This result turns out to only be the tip of the iceberg. Recall that a metric space is **complete** if every Cauchy sequence converges to an element in the metric space itself. We could not generally take this for granted! After all, a sequence being Cauchy in some metric space only meant that the element eventually got close together. There could, however, be a "gap" in the space which allows something to slip through. This was easy to see when comparing the metric spaces $(\mathbb{Q}, d_2)$ and $(\mathbb{R}, d_2)$, since we immediately understand that the rational numbers have "gaps" in them. It is less clear what we mean geometrically by "gaps" in $\mathscr{C}[a, b]$!

In fact, however, the following result tells us that there are not gaps (at least which with respect to the metric d_∞). This will be the final result in this course.

Theorem 4.2. *The metric space* $(\mathscr{C}[a, b], d_\infty)$ *is complete.*

Proof. We clarify first of all the interpretation of the result. We need to show that every Cauchy sequence $\{f_n\}$ in $(\mathscr{C}[a,b], d_\infty)$ converges to an element in $\mathscr{C}[a,b]$. That is, we want to be able to say that if the $d_\infty(f_n, f_m)$ becomes small, then there is a continuous function $f : [a,b] \mapsto \mathbb{R}$ such that

$$\lim_{n\to\infty} f_n(x) = f(x).$$

We already have all of the machinery for this result! Take $\epsilon > 0$ and let $\{f_n\}$ be a Cauchy sequence in $(\mathscr{C}, d_\infty)$. To have a Cauchy sequence with respect to the metric d_∞, it follows that there is an $N \in \mathbb{N}$ such that $n, m \geq N$ implies

$$d_\infty(f_n, f_m) = \sup_{x\in[a,b]} |f_n(x) - f_m(x)| < \epsilon.$$

Of course, the only way this can be true is if $|f_n(x) - f_m(x)| < \epsilon$ for all $x \in [a,b]$. It then follows from Theorem 2.1 that $\{f_n\}$ is uniformly convergent to some function $f(x) : [a,b] \mapsto \mathbb{R}$. However, from Theorem 3.1, we have that every uniformly convergence sequence of continuous functions converges to a *continuous* function. It follows that $f \in \mathscr{C}[a,b]$, and we are done. $\square$