Lecture 3

3.1 Remarks on writing assignment

I'd like to touch on a few themes that came up in your writing assignments

3.1.1 Continuity of ***objects*** and of ***functions***

Many people thought of continuity in terms of properties of an object—for example, they explored in what sense $\mathbb{R}$ seemed like a continuous object.

As I mentioned in class, we define certain ideas (like topological spaces) to be able to speak of certain functions between them (like continuous functions). So it is indeed a good investment to think about what it means for an object like $\mathbb{R}$ to have certain properties that allow us to speak of continuous functions out of, or to, $\mathbb{R}$.

3.1.2 f is continuous if f is defined and...

Many students learn in calculus that f is continuous at a if three conditions are satisfied:

1. $f(a)$ is defined
2. $\lim_{x \to a} f(x)$ exists, and
3. this limit equals $f(a)$.

However, that first condition is superfluous when you have already declared that $f : \mathbb{R} \to \mathbb{R}$ is a function from $\mathbb{R}$ to $\mathbb{R}$. That the domain is $\mathbb{R}$ (and in particular, *all* of $\mathbb{R}$!) means that f is defined at any $a \in \mathbb{R}$.

The reason that you were taught the above definition is because most calculus educators are also taught things that way, and the reason they are taught things that way is because they are not expected to teach with the precision and sophistication that should be expected of mathematicians (and of math majors, at least). For example, your calculus class probably did not consist only of future math majors, so that level of precision was dropped. If one has a function $f : X \to Y$, in particular, f is defined on every element of X.

Finally, in practice, one sometimes writes a formula, and one needs to verify that f is indeed a function on all of $\mathbb{R}$. For example, if I declare

$$f(x) = \frac{\sin x}{x}$$

it is not obvious at first whether f is defined at $x = 0$. Thus you may not know whether to declare f to be a function on all of $\mathbb{R}$, or only on the set $\mathbb{R} \setminus \{0\}$. So this is at least the beginnings of why so many people are careful about verifying that f is defined somewhere—we often write formulas, but formulas do not always make sense everywhere.

3.1.3 Smooth versus continuous

Some people were confused about the existence of "kinks" or "non-smooth" phenomena, and said that part of their intuition of continuity contained a "smoothness" requirement. I want to dissuade you from thinking about smoothness.

First, "smooth" has a technical meaning in math, as it turns out. A function is smooth if you can take a derivative as many times as you want. And while every function is smooth, not every continuous function is smooth. So keep that in mind.

Next, consider the example of $f(x) = |x|$. The graph of this function is not a "smooth" object, as it clearly has a kink, or a corner, at the origin. But the function is still continuous.

3.1.4 "Approached from either side."

Many people spoke of limits. They said that a limit $\lim_{x \to a} f(x)$ exists if "when approached from either side" the limiting value is equal.

This is a fine intuition for limits in $\mathbb{R}$, but what if you are in $\mathbb{R}^2$? If you chose a point $a \in \mathbb{R}^2$, there are many ways to "approach it"–not just in terms of directions, but also in terms of the shape of the path that you take to get to a. (For example, one could spiral toward a.) You see the situation can be even more complicated in $\mathbb{R}^n$ for high n.

The ϵ-δ definition of continuity, which we gave without discussing the notion of limit, ignores any "choice" of direction or path by which you approach a. It simply says that if you want to guarantee that f attains values close to the value $f(a)$, you simply need to be close to a itself.

3.1.5 Picturing

Many people said they had a hard time picturing the definition of continuity.

That's perfectly normal. Indeed, even the most seasoned mathematicians probably do not imagine the most general and crazy examples of continuity; they may simply imagine something like the function $f(x) = |x|$.

This is also the *power* of abstract definitions; in life we sometimes have to prove or understand things without being able to visualize them.

3.1.6 Understanding

Many people said they do not understand the $\epsilon - \delta$ definition. That is normal.

Let me share a quote from John von Neumann.

"In mathematics you don't understand things. You just get used to them."

I strongly disagree with this quote, but it rests on what you mean by understanding. (Let's not get into that discussion.) I give this quote note as a model, but as comfort; even seasoned mathematicians feel like they do not understand things.

An analogy I dislike, but will use anyway because it is helpful, is the following: You may not understand how a car works, but you can still drive it. Rest assured that most mathematicians do not understand everything they use; and at some point, they have simply had to drive a car without taking apart every component.

3.1.7 What's the use of continuity?

This is a great question.

Let me give some sample applications:

Theorem 3.1.7.1. Let $[a, b]$ be a closed, bounded interval. Fix a function $f : [a, b] \to \mathbb{R}$. If f is continuous, then f attains a maximum and a minimum.

This is a very powerful theorem.

As a non-example, consider $\tan x$, which does not attain a minimum or a maximum—this shows the necessity of the interval of definition being closed.

In this class, we will generalize the above theorem to any continuous function whose domain is compact. Compactness is a useful notion that comes up over and over in mathematics—it is useful because it identifies a large class of spaces that behave well and are easily controlled; it is also a condition that is easy to check in many cases.

3.2 Some facts about metrics

Now let's get back to metric spaces.

Exercise 3.2.0.1. Let (X, d) be a metric space. Prove that for any $x, x' \in X$,

$$d(x, x') \geq 0.$$

Proof. Combining the triangle inequality and property zero, we have

$$d(x, x') + d(x', x) \geq d(x, x) = 0.$$

By symmetry, we have

$$2d(x, x') \geq 0.$$

Dividing by 2, we are done. □

The above exercise shows that the notion of distance in a metric space fits our physical intuition that the distance between any two points ought to be non-negative. (And, by property zero, positive when $x \neq x'$.)

Last time we left off as we were about to study the continuity of the identity function between various metric space structures on $\mathbb{R}^n$.

Of course, there were four different metric spaces. That means there are a total of

$$4 \times 4 = 16$$

different combinations of metric for which we would have to verify continuity. That's a lot. But here's a useful fact that will cut down that number:

Exercise 3.2.0.2. Fix three metric spaces

$$(X, d_X), \qquad (Y, d_Y), \qquad (Z, d_Z)$$

and two functions $f : X \to Y$ and $g : Y \to Z$.

Show that if both f and g are continuous, then so is the composition $g \circ f$.

Proof. We must verify the following:

For all $x \in X$, and for all $\epsilon > 0$, there exists δ such that

$$d_X(x, x') < \delta \implies d_Z(gf(x), gf(x')) < \epsilon.$$

Well, by the continuity of g, we know that for all y (and in particular, for $y = f(x)$) and for all ϵ, there exists some ϵ' so that

$$d_Y(f(x), y') < \epsilon' \implies d_Z(gf(x), g(y')). \tag{3.2.1}$$

And by the continuity of f, we know that for all x, and all ϵ', there exists δ such that

$$d_X(x, x') < \delta \implies d_Y(f(x), f(x')) < \epsilon'. \tag{3.2.2}$$

Thus, given ϵ, choose ϵ' satisfying (3.2.1), then choose δ satisfying (3.2.2). We are finished. $\square$

3.2.1 Our favorite metrics on $\mathbb{R}^n$

Last time we defined

1. The standard metric

$$d_{std}(x, x') = \sqrt{\sum_{i=1}^{n} (x'_i - x_i)^2}$$

2. The discrete metric

$$d_{discrete}(x, x') = \begin{cases} 0 & x = x' \\ 1 & x \neq x' \end{cases}$$

3. The taxicab metric

$$d_{taxi}(x, x') = \sum_{i=1}^{n} |x'_i - x_i|$$

4. The l^∞ metric
$$d_{l^\infty}(x, x') = \max_{i=1,\ldots,n} |x'_i - x_i|.$$

We have the identity function

$$\mathrm{id} : \mathbb{R}^n \to \mathbb{R}^n, \qquad x \mapsto x.$$

We can put any of the metric space structures above on $\mathbb{R}^n$—for which choices is the identity function continuous?

Example 3.2.1.1. Consider

$$\mathrm{id} : (\mathbb{R}^n, d_{std}) \to (\mathbb{R}^n, d_{discrete}).$$

Is this a continuous function?

I want you now to get into groups and investigate.

Upshot: You should find that the above example is the only example for which the identity function is not continuous.

3.2.2 A tip

How do you prove a function is continuous? In practice, it comes down to understanding what the condition

$$\text{If } x' \text{ is such that } d_X(x, x') < \delta, \text{ then } d_Y(f(x), f(x')) < \epsilon$$

means. To understand that such x' "look like," we first try to understand what $f(x')$ might look like.

What does $d_Y(f(x), f(x')) < \epsilon$ mean? That is, what does it mean for $f(x')$ to be distance less than ϵ apart from $f(x)$? It means—tautologically—that $f(x')$ is in the following set:

$$y' \text{ such that } d_Y(f(x), y') < \epsilon.$$

This is what one calls the *open ball of radius ϵ centered at $f(x)$* (with respect to the metric d_Y).

Example 3.2.2.1. We have drawn examples of open balls of radius 1 for our various metrics on $\mathbb{R}^n$.

A common strategy to prove the continuity of a function $f : X \to Y$ is to understand the set of those x' that end up in the open ball of radius ϵ centered at $f(x)$. That is, under what circumstances do we have that

$$f(x') \in \{\ y' \text{ such that } d_Y(f(x), y') < \epsilon.\ \}?$$

This comes down to understanding the following:

$$\text{What } x' \text{ are contained in the preimage } f^{-1}(\{\ y' \text{ such that } d_Y(f(x), y') < \epsilon.\ \})?$$

In other words,

$$\text{What is} f^{-1}(\{\ y' \text{ such that } d_Y(f(x), y') < \epsilon.\ \})? \tag{3.2.3}$$

Note that x is always contained in this set.

Once you understand this set, you can ask the following: Is there an open ball centered at x (of some radius $\delta > 0$) contained in this set?

If you can find such a δ, and if you can do this for every $\epsilon > 0$ and every $x \in X$, you have proven the continuity of $f : X \to Y$.

Let me summarize. The following statements are more or less equivalent descriptions of the process:

1. To show $f : X \to Y$ is continuous, you must answer in the affirmative: "Is there an open ball of radius $\delta > 0$ centered at x contained in the preimage of the open ball of radius ϵ centered at $f(x)$?" for every choice of $x \in X$ and for every choice of $\epsilon > 0$.

2. Continuity comes down to verifying that you can always find open balls centered at x contained in the preimage of open balls centered at $f(x)$.

We got back together and discussed. We discussed the following three examples:

- $\mathrm{id} : (\mathbb{R}^n, d_{std}) \to (\mathbb{R}^n, d_{std})$.

- $(\mathbb{R}^n, d_{std}) \to (\mathbb{R}^n, d_{taxi})$.

- $(\mathbb{R}^n, d_{std}) \to (\mathbb{R}^n, d_{discrete})$.

3.2.3 The identity function from a metric space to itself

Proposition 3.2.3.1. $\mathrm{id} : (\mathbb{R}^n, d_{std}) \to (\mathbb{R}^n, d_{std})$ is continuous.

Remark 3.2.3.2. We are verifying that f is continuous for the case $X = Y = \mathbb{R}^n$, $d_X = d_Y = d_{std}$, and $f = \mathrm{id}$. The proof will thus use the notations X, Y, d_X, d_Y to make clear when I am speaking of the domain, or of the codomain.

For the proof, we follow the tip:

Proof of Proposition 3.2.3.1. Note that for any $x \in \mathbb{R}^n$, $f(x) = x$ because we have chosen $f = \mathrm{id}$.[1] Note also the following:

$$\begin{aligned} d_X(x, x') &= d_{std}(x, x') \\ &= d_{std}(f(x), f(x')) \\ &= d_Y(f(x), f(x')). \end{aligned} \tag{3.2.4}$$

That is,

$$d_X(x, x') = d_Y(f(x), f(x')). \tag{3.2.5}$$

As such, given any $\epsilon > 0$, let us simply set δ to be any positive real number less than or equal to ϵ. Then if $d_X(x, x') < \delta$, we have that $d_Y(f(x), f(x')) < \epsilon$. This shows f is continuous. $\square$

Note that we did not use anything about $\mathbb{R}^n$ or d_{std} in the proof—indeed, the string of equalities uses the notation d_{std}, but the equalities hold for any choice of metric so long as $d_X = d_Y$. We conclude:

Proposition 3.2.3.3. Let $X = Y$ and $d_X = d_Y$, and let $f : X \to Y$ be the identity function. Then f is continous.

Proof. Note that (3.2.5) is true when $f = \mathrm{id}$ and $d_X = d_Y$. Then follow the rest of the proof of Proposition 3.2.3.1. $\square$

[1] That is, we have chosen f to be the identity function.

3.2.4 Isometric embeddings and isometries

In fact, we have done something even better. We do not need X to equal Y, nor for d_X to equal d_Y. The proof of Proposition 3.2.3.1 relied only on the equality (3.2.5). This is a useful condition, so let's give it a name.

Definition 3.2.4.1. Let (X, d_X) and (Y, d_Y) be metric spaces. Fix a function $f : X \to Y$. We say that f is an *isometric embedding* if f preserves distances. That is,

$$d_X(x, x') = d_Y(f(x), f(x'))$$

for all $x, x' \in X$.

If f is further a bijection, we say that f is an *isometry*.

Proposition 3.2.4.2. Let (X, d_X) and (Y, d_Y) be metric spaces. Fix a function $f : X \to Y$. If f is an isometric embedding, then f is continuous. In particular, any isometry is continuous.

Proof. Follow the proof of Proposition 3.2.3.1 beginning with (3.2.5). □

3.2.5 From standard to taxi

Then we verified that the identity function from $\mathbb{R}^n$ with the standard metric to $\mathbb{R}^n$ with the taxicab metric is continuous:

Proposition 3.2.5.1. Let $(X, d_X) = (\mathbb{R}^n, d_{std})$ and $(Y, d_Y) = (\mathbb{R}^n, d_{taxi})$. Then the identity function

$$f = \mathrm{id} : X \to Y$$

is continuous.

Proof. Fix $x \in X$. We note that the open ball of radius ϵ centered at $f(x)$[2] is a diamond centered at $f(x)$, whose distance from $f(x)$ to any of the corners of the diamond is ϵ. Because $f = \mathrm{id}$, the preimage of this diamond is the diamond itself (now considered as a subset of X).

In $\mathbb{R}^2$, we drew a picture to see that any diamond with $\epsilon > 0$ "clearly" contains an open ball of radius $\delta > 0$ centered at x, so long as δ is small enough. For a picture and precise formula, see the scanned notes.

To prove for $\mathbb{R}^n$ for general n, one employs the formula from the scanned notes to verify indeed that one can find $\delta > 0$ small enough so that the continuity condition holds. □

[2] with respect to $d_Y = d_{taxi}$!)

3.2.6 From standard to discrete

Proposition 3.2.6.1. The function

$$\mathrm{id} : (\mathbb{R}^n, d_{std}) \to (\mathbb{R}^n, d_{discrete})$$

is not continuous.

Proof. It suffices to exhibit an $x \in X$ and $\epsilon > 0$ such that for any $\delta > 0$, there exists some x' such that

$$d_X(x, x') < \delta \text{ and } d_Y(f(x), f(x')) \geq \epsilon.$$

For this, choose ϵ to be any positive real number less than or equal to 1. We saw last time that the open ball of radius ϵ (with respect to $d_{discrete}$) centered at $f(x)$ is then a single point, given by $f(x)$ itself.

Because $f = \mathrm{id}$, the preimage of this set is the same set—that is, the preimage if the set $\{x\}$ consisting of a single point, called x.

Of course, for any $\delta > 0$, the open ball of radius δ centered at x can *not* be contained in this singleton set. $\square$

Summary: We've seen four metrics on $\mathbb{R}^n$!

Notation

$$d_{std}(x,x') = \sqrt{\sum_{i=1}^{n} (x_i' - x_i)^2}$$

(standard metric)

$$d_{taxi}(x,x') = \sum_{i=1}^{n} |x_i' - x_i|$$

(taxicab metric)

$$d_{\ell^\infty}(x,x') = \max_{i=1,\dots,n} |x_i' - x_i|$$

(ℓ^∞ metric)

$$d_{discrete}(x,x') = \begin{cases} 0 & x = x' \\ 1 & x \neq x' \end{cases}$$

(discrete metric)

Project for today

Let $X = Y = \mathbb{R}^2$, and let

$$f: X \longrightarrow Y$$

be the identity function. (That is, $f(x) = x$.)

For which metrics d_X and d_Y on X and Y (from the four metrics we've discussed) is f continuous?

⚠ When you do NOT yet have an intuition, definitions are all you've got to go on.

(Today, you probably have very little intuition for these new metrics!)

· $d_X = d_{std}$, $d_Y = d_{taxi}$.

$\forall \varepsilon$, need to find δ so the (usual) ball of radius δ

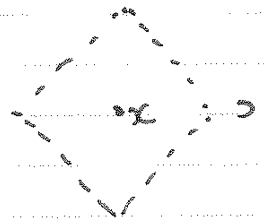

is contained in the cube whose corners are ε away from the center.

<u>$n=1$</u>: $\delta = \varepsilon$ will do.

<u>$n=2$</u>:

By symmetry, $\{x_1 + x_2 = \varepsilon\}$ has point closest to origin @ $x_1 = x_2 = \varepsilon/2$.
Hence take $\delta = \sqrt{(\varepsilon/2)^2 + (\varepsilon/2)^2} = \varepsilon \frac{\sqrt{2}}{2}$.

For general n, again by symmetry, $x_1 = x_2 = \cdots = x_n = \varepsilon/n$ is closest point to origin (on $\{x_1 + \cdots + x_n = \varepsilon\}$ plane). So take

$$\delta = \varepsilon \frac{\sqrt{n}}{n}.$$

· $d_X = d_{taxi}$ $\quad$ $d_Y = d_{std}$.

Given ε, need δ so cube whose corners are length δ from center fits in ball of radius ε.

Can take $\delta = \varepsilon$

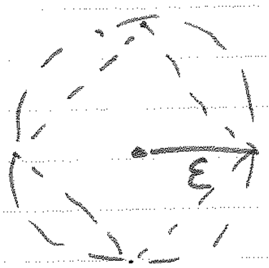

For:

$$\sum |x_i' - x_i| < \delta$$

$$\Longrightarrow \left(\sum |x_i' - x_i|\right)^2 < \delta^2 = \varepsilon^2$$

$$\Longrightarrow \sum_{i=1,\dots,n} |x_i' - x_i|^2 + \underbrace{2\sum_{i \neq j} |x_i' - x_i||x_j' - x_j|}_{\text{positive}} < \varepsilon^2$$

$$\Longrightarrow \sum |x_i' - x_i|^2 < \varepsilon^2$$

$$\Longrightarrow \sqrt{\sum (x_i' - x_i)^2} < \varepsilon$$

$$\Longrightarrow d_{std}(x, x') < \varepsilon.$$

$d_X = d_{std}$, $d_Y = d_{\ell^\infty}$.

Need δ so that, given a box/cube of width 2ε, a ball of radius δ fits inside.

Can take $\delta = \varepsilon$, for:

If $d_{std}(x, x') < \delta$, have

$$\sqrt{\sum_{i=1}^{n} (x_i' - x_i)^2} < \delta$$

$$\Rightarrow \sum_{i=1}^{n} (x_i' - x_i)^2 < \delta^2$$

$\Rightarrow$ for each i, $(x_i' - x_i)^2 < \delta^2$

$$\Rightarrow \forall i, \quad |x_i' - x_i| < \delta$$

$$\Rightarrow \max_{i=1,\dots,n} |x_i' - x_i| < \delta$$

$$\Rightarrow d_{\ell^\infty}(x', x) < \delta = \varepsilon. \quad //$$

$d_X = d_{\ell^\infty}$, $\quad d_Y = d_{std}$

Suppose $d_{\ell^\infty}(x', x) < \delta$, so $\max\limits_{i=1,\dots,n} |x_i' - x_i| < \delta$.

Then $\forall\, i$, $(x_i' - x_i)^2 < \delta^2$. Hence

$$d_{std}(x', x) = \sqrt{\sum (x_i' - x_i)^2}$$

$$< \sqrt{\sum \delta^2}$$

$$= \sqrt{n\delta^2}$$

$$= \delta\sqrt{n}$$

Hence if $\delta \leq \varepsilon \frac{\sqrt{n}}{n}$, we have

$$d_{\ell^\infty}(x', x) < \delta \implies d_{std}(x', x) < \varepsilon.$$

//

$d_X = d_{discrete}$, $d_Y = d_{std}$.

For any ε, just set $\delta = \frac{1}{2}$ (or anything between 0 and 1). For we see

$$\{x' \text{ s.t. } d_{discrete}(x,x') < \tfrac{1}{2}\} = \{x\}.$$

(That is, x is the <u>only</u> point w/in $\frac{1}{2}$ distance of x).

So

$$\begin{aligned} d_{discrete}(x,x') < \tfrac{1}{2} &\Longrightarrow x = x' \\ &\Longrightarrow d_{std}(x,x') = 0 \\ &\Longrightarrow d_{std}(x,x') < \varepsilon. \end{aligned}$$

Strangely, δ does NOT depend on ε.

(As though f behaves like a constant function.)