Point-Set Topology Minicourse

Zev Chonoles
2011-07-11

My presentation will follow that in Crossley's *Essential Topology*, which I highly recommend.

Motivating questions.

- Is a "continuous" function from $\mathbb{R}$ to $\mathbb{Z}$ necessarily constant?
- Can the range of a "continuous" function from $(0, 1)$ to $\mathbb{R}$ be "missing a point"? That is, if a and b are in the range, must every $x \in (a, b)$ be in the range as well?
- If a sequence "converges", must it converge to only one value?

We'll be able to answer these questions by the end of the talk. As we'll see all of these questions depend on the **topology** of the relevant sets.

What's a naive definition of "continuous" function? "Able to be drawn without lifting pencil from paper."

But consider the function $\frac{1}{x}$. Except for the point $x = 0$, where it isn't defined, this is a continuous function; we can draw *each half of*

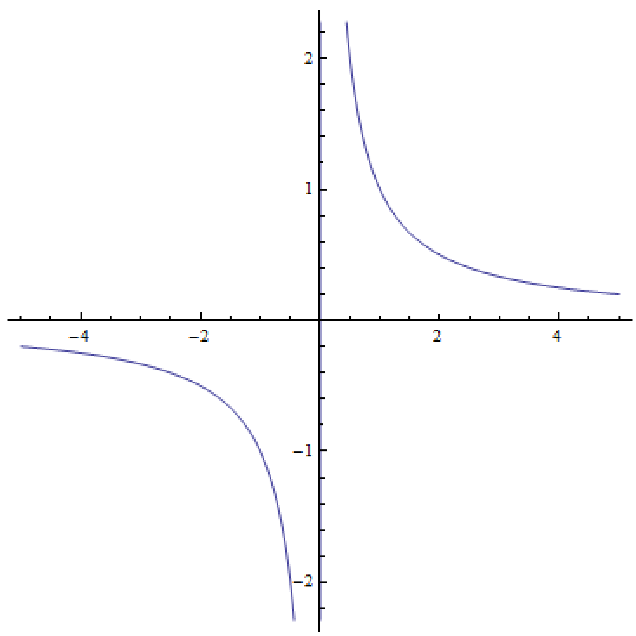

without lifting our pencil. Put another way: around any point *in the domain of* $\frac{1}{x}$, some piece of the graph of $\frac{1}{x}$ is able to be drawn without lifting our pencil. I think $\frac{1}{x}$ deserves to be considered "continuous", as long as we throw away 0.

So, if we could define a notion of what it means for a function to be "continuous at a point a", then we could simply define a function f to be continuous precisely when, for every point a in its domain, it is continuous at a.

What's a naive definition of "f is continuous at a"? "$f(a)$ is the value we expect."

This isn't rigorous though - different people might expect different things! Intuitively, as x approaches a, we want $f(x)$ to approach $f(a)$. But suppose you've gotten into an argument with the Continuity Monster, who wants to show that your function f isn't continuous at a.

> CM: You really think $f(x)$ approaches $f(a)$? Well, I've chosen a **very** small number ϵ, and I bet $f(x)$ doesn't get within ϵ of $f(a)$! Or if it does, only a few times, and not for an entire range of x's.
>
> You: Ha! In fact, I've found a number δ such that for **any** x that's within δ of a, the value of $f(x)$ is within ϵ of $f(a)$!
>
> CM: But what if I make ϵ *even smaller*!?
>
> You: It doesn't matter what ϵ you choose - I can always come up with a δ that works.
>
> CM: Nooooooo!

Let's say that f is continuous at a when we can defeat the Continuity Monster - no matter which $\epsilon > 0$ he chooses, we can find a $\delta > 0$ such that for the entire range of x's with $|x - a| < \delta$, it's true that $|f(x) - f(a)| < \epsilon$. **Note that our choice of δ <u>depends on</u> the value of ϵ.** Using a formal definition:

Definition. A function f is continuous at a when

$$\forall \epsilon > 0, \exists \delta > 0 \text{ such that } |x - a| < \delta \implies |f(x) - f(a)| < \epsilon.$$

Definition. A function f is continuous when it is continuous at a, for every a in the domain of f.

But does this really capture our experience? Consider the function $f : (0, 1) \to \mathbb{R}$ defined by $f(x) = x^2$:

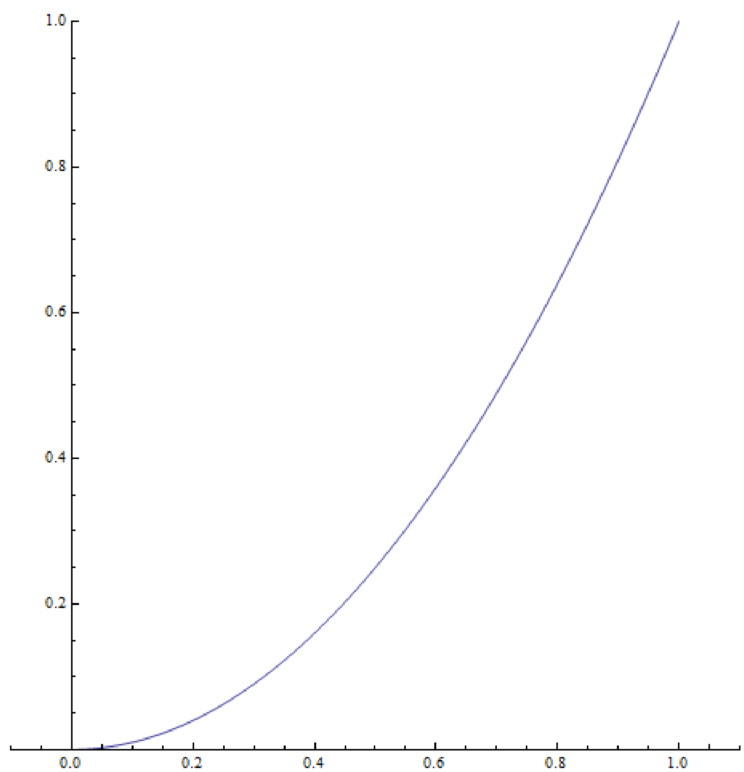

It certainly *looks* continuous to me.

Theorem. The function $f : (0,1) \to \mathbb{R}$ defined by $f(x) = x^2$ is continuous.

Proof. By definition, f is continuous only if, for every $a \in (0,1)$, f is continuous at a.

Choose some $a \in (0,1)$. For any $x \in (0,1)$,

$$|x^2 - a^2| = |(x+a)(x-a)| = |x+a| \cdot |x-a| < 2|x-a|$$

because $x \in (0,1)$ and $a \in (0,1)$ means that $0 < x < 1$ and $0 < a < 1$, and therefore $0 < x + a < 2$.

For any $\epsilon > 0$,

$$|x-a| < \frac{\epsilon}{2} \implies |f(x) - f(a)| = |x^2 - a^2| < 2|x-a| < 2 \cdot \frac{\epsilon}{2} = \epsilon.$$

Therefore, for any $\epsilon > 0$, there does indeed exist a $\delta > 0$ such that $|x - a| < \delta$ implies $|f(x) - f(a)| < \epsilon$: namely, $\delta = \frac{\epsilon}{2}$. Thus, for any $a \in (0,1)$, f is continuous at a, and so f is continuous. □

OK, I guess we're done with continuity. Except. . . this whole δ and ϵ business is kind of messy! Who really wants to have to find a δ for **every possible** ϵ just to show that a function is continuous? As we'll see later, defining what it means for a function to be "continuous" does not even require a notion of "distance" or "size" (remember, in our definition, we needed to use the absolute value $| \ |$). We can do a lot better!

You should be familiar with the notion of an open interval $(a,b) = \{x \in \mathbb{R} \mid a < x < b\}$ already. But here are two new concepts:

Definition. A subset $S \subseteq \mathbb{R}$ is an **open** set when, for every $p \in S$, there is some open interval (a,b) such that $p \in (a,b)$, and $(a,b) \subseteq S$.

Definition. A subset $S \subseteq \mathbb{R}$ is a **closed** set when its complement, $S^c = \{x \in \mathbb{R} \mid x \notin S\}$, is an open set.

What are some examples of open sets?

- all of $\mathbb{R}$
- the empty set $\varnothing$
- any open interval (a,b)
- any "infinite open interval" (a,∞)

You may be wondering, "Why is the empty set $\varnothing$ open? How a statement that begins 'for every $p \in \varnothing$ . . .' be true?" In fact, *any* statement I make about the elements of $\varnothing$ is automatically **true**, precisely **because** it has no elements. Think of it this way: if I claimed "Every pink elephant in the room is wearing a hat", that would actually be true - because there *aren't any pink elephants in the room to begin with* (hopefully), so *all 0 of them* are wearing hats. This is called "vacuous truth".

Another very important point is that **open is not the opposite of closed** - a subset $S \subseteq \mathbb{R}$ might be open, closed, *both*, or *neither*.

It turns out that the only subsets of $\mathbb{R}$ that are both open and closed are $\mathbb{R}$ and $\varnothing$ - this will come up later, when we talk about connectedness.

Also, here is an example of an $S \subseteq \mathbb{R}$ that is neither open nor closed: $S = [0,1)$. It's not open, because there does not exist any open interval (a,b) that contains 0 such that $(a,b) \subseteq S$; if $0 \in (a,b)$, then $a < 0$, so $\frac{a}{2} < 0$ and $\frac{a}{2} \in (a,b)$, but $\frac{a}{2} \notin S = [0,1)$. The complement of S, namely $S^c = (-\infty,0) \cup [1,\infty)$, is not open, because S^c has the same problem at 1 that S did at 0. Because S^c isn't open, the set S isn't closed, by definition.

A good way to think about open sets is that they have "breathing space" or "wiggle room" around every point - given any p in an open set S, there is some interval (a,b) containing p that we can move around in without leaving S.

Theorem. Arbitrary unions of open sets are open.

Remark. When we say "arbitrary", we mean that we might be looking at literally *any* collection of open sets - there could be an uncountable number of open sets, or maybe just finitely many of them. When we index objects with natural numbers, e.g. U_1, U_2, ..., we usually refer to the collection as $\{U_i\}_{i\in\mathbb{N}}$. However, since we don't know how many objects we'll be working with, the convention is to say only that our indexing set is A, and we refer to the collection as $\{U_\alpha\}_{\alpha\in A}$.

Proof. Let $\{U_\alpha\}_{\alpha\in A}$ be some collection of open sets, and let $V = \bigcup_{\alpha\in A} U_\alpha$ be their union. In order to prove that V is open, we have to prove that for any $p \in V$, there is some open interval (a, b) such that $p \in (a, b)$ and $(a, b) \subseteq V$.

By the definition of union, $p \in V$ if and only if $p \in U_\beta$ for *at least one* $\beta \in A$. By hypothesis, this U_β is open, and because $p \in U_\beta$, there exists an open interval (a, b) such that $p \in (a, b)$ and $(a, b) \subseteq U_\beta$. Then, because (a, b) is contained in this specific U_β, it is also contained in V, the union of all of the U_α's. Thus, we have proven that for any $p \in V$, there is an open interval (a, b) such that $p \in (a, b)$ and $(a, b) \subseteq V$; therefore, V is open. □

Theorem. Finite intersections of open sets are open.

Proof. Let $\{U_1, \dots, U_n\}$ be some finite collection of open sets, and let $V = \bigcap_{i=1}^n U_i$ be their intersection. In order to prove that V is open, we have to prove that for any $p \in V$, there is some open interval (a, b) such that $p \in (a, b)$ and $(a, b) \subseteq V$.

By the definition of intersection, $p \in V$ if and only if $p \in U_i$ for *every* $i \in \{1, \dots, n\}$. By hypothesis, each U_i is open, and because $p \in U_i$, there exists an open interval (a_i, b_i) such that $p \in (a_i, b_i)$ and $(a_i, b_i) \subseteq U_i$. In other words,

$$a_i < p < b_i \text{ for all } i \in \{1, \dots, n\}.$$

Let $a = \max\{a_1, \dots, a_n\}$ and $b = \min\{b_1, \dots, b_n\}$. Because $a \in \{a_1, \dots, a_n\}$ and $b \in \{b_1, \dots, b_n\}$, it is clear that $a < p < b$, and so $p \in (a, b)$. Because $a_i \le a$ and $b \le b_i$ for all $i \in \{1, \dots, n\}$,

$$(a, b) \subseteq (a_i, b_i) \text{ for all } i \in \{1, \dots, n\}.$$

Because $(a_i, b_i) \subseteq U_i$, this implies that

$$(a, b) \subseteq U_i \text{ for all } i \in \{1, \dots, n\},$$

and therefore $(a, b) \subseteq V = \bigcap_{i=1}^n U_i$. Thus, we have proven that for any $p \in V$, there is an open interval (a, b) such that $p \in (a, b)$ and $(a, b) \subseteq V$; therefore, V is open. □

In general, we can't take arbitrary intersections of open sets and expect them to be open. For example, if we take the collection of open sets $\{U_n\}_{n\in\mathbb{N}}$ where $U_n = (-\frac{1}{n}, 1)$, their intersection is

$$\bigcap_{n\in\mathbb{N}} U_n = \{x \in \mathbb{R} \mid -\tfrac{1}{n} < x < 1 \text{ for all } n \in \mathbb{N}\} = \{x \in \mathbb{R} \mid 0 \le x < 1\} = [0, 1)$$

which, as we saw before, is neither open nor closed.

These results make sense, using the "wiggle room" analogy. If a set has wiggle room around each point, adding stuff to the set can't possibly change that. But if we start cutting off parts of the set, we have to be careful that we don't cut out a point's wiggle room, but not the point itself, lest we leave a point having no wiggle room.

Theorem. Arbitrary intersections and finite unions of closed sets are closed.

Proof. This follows directly from the opposite statements about open sets. By definition, $K \subseteq \mathbb{R}$ is closed precisely when $K^c = \{x \in \mathbb{R} \mid x \notin K\}$ is open. Thus, if $\{K_\alpha\}_{\alpha \in A}$ is any collection of closed sets, then every K_α^c is an open set, so that $\bigcup_{\alpha \in A} K_\alpha^c$ is open, and therefore (by de Morgan's laws)

$$\bigcap_{\alpha \in A} K_\alpha = \left(\bigcup_{\alpha \in A} K_\alpha^c \right)^c$$

is the complement of an open set, hence closed. Similary, if $\{K_1, \ldots, K_n\}$ is a finite collection of closed sets, then every K_i^c is an open set, so that $\bigcap_{i=1}^n K_i^c$ is open, and therefore (by de Morgan's laws)

$$\bigcup_{i=1}^{n} K_i = \left(\bigcap_{i=1}^{n} K_i^c \right)^c$$

is the complement of an open set, hence closed. □

Definition. Let X and Y be sets, let $f : X \to Y$ be a function, and let $S \subseteq Y$ be a subset of Y. The **preimage** (a.k.a. **inverse image**) of S is the subset $f^{-1}(S)$ of X consisting of those elements which get sent to S by f. In symbols,

$$f^{-1}(S) = \{x \in X \mid f(x) \in S\}.$$

The preimage $f^{-1}(S)$ is defined for any $S \subseteq Y$, regardless of whether f is a bijection. So, the preimage "f^{-1}" is **not the same concept** as the inverse function "f^{-1}" (which only exists when f is a bijection).

What are some examples of preimages? Let $f : \mathbb{R} \to \mathbb{R}$ be defined by $f(x) = x^2 - 1$, and let

$$Q = (-1, 1) \qquad R = [1, 3] \qquad S = (-5, 1) \qquad T = [5, \infty).$$

Then

$$\begin{aligned} f^{-1}(Q) &= \{x \in \mathbb{R} \mid f(x) \in (-1,1)\} \\ &= \{x \in \mathbb{R} \mid -1 < x^2 - 1 < 1\} \\ &= \{x \in \mathbb{R} \mid 0 < x^2 < 2\} \\ &= \{x \in \mathbb{R} \mid 0 < x < \sqrt{2}\} \\ &= (0, \sqrt{2}) \end{aligned}$$

$$\begin{aligned} f^{-1}(R) &= \{x \in \mathbb{R} \mid f(x) \in [1,3]\} \\ &= \{x \in \mathbb{R} \mid 1 \le x^2 - 1 \le 3\} \\ &= \{x \in \mathbb{R} \mid 2 \le x^2 \le 4\} \\ &= \{x \in \mathbb{R} \mid \sqrt{2} \le x \le 2\} \\ &= (\sqrt{2}, 2) \end{aligned}$$

$$\begin{aligned} f^{-1}(\varnothing) &= \{x \in \mathbb{R} \mid f(x) \in \varnothing\} \\ &= \varnothing \\ &\left(\begin{matrix} \text{because } \varnothing \\ \text{has no elements} \end{matrix} \right) \end{aligned}$$

$$\begin{aligned} f^{-1}(S) &= \{x \in \mathbb{R} \mid f(x) \in (-5,1)\} \\ &= \{x \in \mathbb{R} \mid -5 < x^2 - 1 < 1\} \\ &= \{x \in \mathbb{R} \mid -4 < x^2 < 2\} \\ &= \{x \in \mathbb{R} \mid 0 < x < \sqrt{2}\} \\ &= (0, \sqrt{2}) \end{aligned}$$

$$\begin{aligned} f^{-1}(T) &= \{x \in \mathbb{R} \mid f(x) \in [5,\infty)\} \\ &= \{x \in \mathbb{R} \mid 5 \le x^2 - 1\} \\ &= \{x \in \mathbb{R} \mid 6 \le x^2\} \\ &= \{x \in \mathbb{R} \mid \sqrt{6} \le x\} \\ &= [\sqrt{6}, \infty) \end{aligned}$$

$$\begin{aligned} f^{-1}(\mathbb{R}) &= \{x \in \mathbb{R} \mid f(x) \in \mathbb{R}\} \\ &= \mathbb{R} \\ &\left(\begin{matrix} \text{because } f \text{ has to send} \\ \text{each } x \in \mathbb{R} \textit{ somewhere} \end{matrix} \right) \end{aligned}$$

Some important things to remember about preimages: given two sets X and Y, and a function $f : X \to Y$, then for any subsets $A \subseteq X$ and $B_1, B_2 \subseteq Y$,

- $f^{-1}(B_1 \cup B_2) = f^{-1}(B_1) \cup f^{-1}(B_2)$
- $f^{-1}(B_1 \cap B_2) = f^{-1}(B_1) \cap f^{-1}(B_2)$
- $f^{-1}(f(A)) \supseteq A$ (**not necessarily equal**)

Theorem. A function $f : \mathbb{R} \to \mathbb{R}$ is continuous if and only if, for every open set $U \subseteq \mathbb{R}$, the preimage $f^{-1}(U)$ is also an open set.

Proof. ($\Longrightarrow$) Suppose $f : \mathbb{R} \to \mathbb{R}$ is a continuous function. By our definition, this means that for every $a \in \mathbb{R}$, f is continuous at a, i.e. for every $\epsilon > 0$, there exists a $\delta > 0$ such that

$$|x - a| < \delta \implies |f(x) - f(a)| < \epsilon,$$

or in other words

$$x \in (a - \delta, a + \delta) \implies f(x) \in (f(a) - \epsilon, f(a) + \epsilon).$$

We want to show that for any open $U \subseteq \mathbb{R}$, the set $f^{-1}(U)$ is also open, i.e., that for any $t \in f^{-1}(U)$, there is an open interval (a, b) for which $t \in (a, b)$ and $(a, b) \subseteq f^{-1}(U)$. By the definition of $f^{-1}(U)$, $f(t) \in U$ for any $t \in f^{-1}(U)$. Because U is open, there is an open interval (c, d) such that $f(t) \in (c, d)$ and $(c, d) \subseteq U$. Letting $\epsilon = \min\{f(t) - c, d - f(t)\}$ (note that $\epsilon > 0$ because $c < f(t) < d$), we have that

$$(f(t) - \epsilon, f(t) + \epsilon) \subseteq (c, d) \subseteq U.$$

Because f is continuous, there exists a $\delta > 0$ such that

$$x \in (t - \delta, t + \delta) \implies f(x) \in (f(t) - \epsilon, f(t) + \epsilon).$$

Because $(f(t) - \epsilon, f(t) + \epsilon) \subseteq U$, we see that $x \in (t - \delta, t + \delta) \implies f(x) \in U$. Therefore, $(t - \delta, t + \delta) \subseteq f^{-1}(U)$. Thus, we have shown that for any $t \in f^{-1}(U)$, there is an open interval $(a, b) = (t - \delta, t + \delta)$ for which $t \in (a, b)$ and $(a, b) \subseteq f^{-1}(U)$. Therefore, $f^{-1}(U)$ is an open set for any open set U.

($\Longleftarrow$) Suppose that for any open set U, $f^{-1}(U)$ is also an open set. We want to show that f is continuous, i.e., that for every $a \in \mathbb{R}$, f is continuous at a.

Fix any $a \in \mathbb{R}$. For any $\epsilon > 0$, the set $V = (f(a) - \epsilon, f(a) + \epsilon)$ is a open set, and therefore $f^{-1}(V)$ is open as well. Note that $a \in f^{-1}(V)$, because $f(a) \in V$. Because $f^{-1}(V)$ is open and $a \in f^{-1}(V)$, there is some open interval (s, t) for which $a \in (s, t)$ and $(s, t) \subseteq f^{-1}(V)$. Letting $\delta = \min\{a - s, t - a\}$ (note that $\delta > 0$ because $s < a < t$), we have that

$$(a - \delta, a + \delta) \subseteq (s, t) \subseteq f^{-1}(V).$$

Thus, $x \in (a - \delta, a + \delta) \implies f(x) \in V = (f(a) - \epsilon, f(a) + \epsilon)$. In other words,

$$|x - a| < \delta \implies |f(x) - f(a)| < \epsilon.$$

Because $a \in \mathbb{R}$ and $\epsilon > 0$ were arbitrary, this proves that f is continuous at a, for all $a \in \mathbb{R}$, and therefore that f is continuous. □

Now we're going to take the experience we gained from working with open sets in $\mathbb{R}$, and generalize it (much like we take our experience with the operations $+$ and $\times$ in $\mathbb{Z}$, and define the concept of "ring").

Recall that $\mathcal{P}(X)$, the power set of X, is defined to be the set whose elements are the subsets of X. For example, $\mathcal{P}(\{a, b\}) = \{\varnothing, \{a\}, \{b\}, \{a, b\}\}$.

Definition. A **topological space** is a set X, together with a set $T \subseteq \mathcal{P}(X)$ satisfying

- $\varnothing, X \in T$
- T is closed under arbitrary unions
- T is closed under finite intersections

In other words, if $\{U_\alpha\}_{\alpha \in A}$ is any collection of elements of T, then $\bigcup_{\alpha \in A} U_\alpha \in T$, and if $\{U_1, \ldots, U_n\}$ is any finite collection of elements of T, then $\bigcap_{\alpha \in A} U_\alpha \in T$.

T is called the **topology** on X. Subsets of X that are in T are called **open** subsets of X.

It is very important to understand that we are now *defining* which subsets of X are "open" using the topology T, which we can choose to be whatever we want. Given a set X with a topology T_1, and the same set X but with a different topology T_2, these are *different* topological spaces. So, a statement like "the subset $U \subseteq X$ is open" is technically ambiguous without specifying "in the topology T". However, if one particular topology on the set X has been agreed on beforehand, and there is no concern of confusion, it is quite common to simply say a set is "open in X" without mentioning in which topology we mean.

What are some examples of topological spaces?

- The real numbers $\mathbb{R}$, with the topology we found earlier:

$$T = \left\{ S \subseteq \mathbb{R} \;\middle|\; \begin{array}{c} \text{for all } p \in S, \text{ there is an open interval } (a,b) \\ \text{such that } p \in (a,b) \text{ and } (a,b) \subseteq S \end{array} \right\}.$$

  This is (surprise!) called the **usual topology** on $\mathbb{R}$. If someone says "$U \subseteq \mathbb{R}$ is open" without specifying in what topology, they implicitly mean the usual topology.

- Let X be any set. Let $T = \{\varnothing, X\}$ (i.e. the only subsets of X that are open are $\varnothing$ and X). This is called the **trivial topology** on X.

- Let X be any set. Let $T = \mathcal{P}(X)$ (i.e., *every* subset of X will be labeled "open"). This is called the **discrete topology** on X. A topological space is **discrete** when its topology is the discrete topology.

- Let X_1 and X_2 be topological spaces, with T_1 the topology on X_1 and T_2 the topology on X_2. Then the Cartesian product

$$P = X_1 \times X_2 = \{(a,b) \mid a \in X_1, b \in X_2\}$$

  can be made into a topological space by giving it the **product topology**

$$T = \left\{ S \subseteq P \;\middle|\; \begin{array}{c} S = \text{ a union of sets of the form } U \times V \\ \text{where } U \in T_1 \text{ and } V \in T_2 \end{array} \right\}.$$

  The notation $U \times V$ means the subset of P

$$U \times V = \{(a,b) \in P \mid a \in U, b \in V\}$$

  Using this construction, we can make $\mathbb{R}^n = \overbrace{\mathbb{R} \times \cdots \times \mathbb{R}}^{n \text{ times}}$, the set of n-tuples of real numbers, into a topological space. In general, taking products of topological spaces is an important way in which we make new topological spaces.

- Let X be a topological space, having topology T. For any subset $Y \subseteq X$, we can make Y into a topological space by giving it the **subspace topology**

$$T' = \{U \subseteq Y \mid U = V \cap Y \text{ for some open set } V \subseteq X\}.$$

In other words, a subset U of Y is open in the subspace topology precisely when it is the intersection of Y with some open subset V of X.

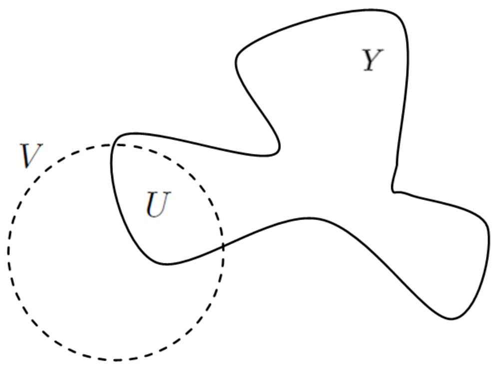

The subspace topology is another extremely important way in which we can make topological spaces. For example, $(0,1)$, $[0,1]$, and $\mathbb{Z}$ are all important subsets of $\mathbb{R}$, and we are often interested in giving them the subspace topology from $\mathbb{R}$.

The subspace topology that $\mathbb{Z}$ gets from $\mathbb{R}$ is actually the same as the discrete topology on $\mathbb{Z}$. This is because, for any subset $S \subseteq \mathbb{Z}$,

$$S = \mathbb{Z} \cap \left(\bigcup_{n \in S} (n - \tfrac{1}{2}, n + \tfrac{1}{2}) \right)$$

That is, for any subset $S \subseteq \mathbb{Z}$, there is an open set of $\mathbb{R}$ whose intersection with $\mathbb{Z}$ is S, and so every subset of $\mathbb{Z}$ is open in the subspace topology; therefore, the subspace topology is the same as the discrete topology. Here is an illustration of an example when $S = \{0, 2, 4, 5\}$:

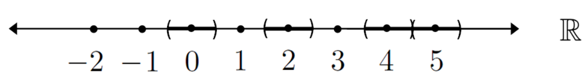

The subspace topology that $[0,1]$ gets from $\mathbb{R}$ is also worth commenting on. Note that, because $(-\frac{1}{2}, 2)$ is an open subset of $\mathbb{R}$, and its intersection with $[0,1]$ is the set $(-\frac{1}{2}, 1]$, the set $(-\frac{1}{2}, 1]$ is **open** in the subspace topology of $[0,1]$, even though it was neither open nor closed as a subset of $\mathbb{R}$. This is yet another illustration of the fact that the question of which sets are "open" is determined by which topology we choose, and is **not** somehow intrinsic to the sets.

Finally, note that most of the shapes we are familiar with (a circle, a cylinder, a sphere, etc.) can be defined as being a subset of $\mathbb{R}^n$ (for some choice of dimension n), and then these become topological spaces when we give them the subspace topology from $\mathbb{R}^n$. For example, $\mathbb{S}^1$, the 1-dimensional sphere (a.k.a. the circle) is the set

$$\mathbb{S}^1 = \{(x,y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$$

given the subspace topology from $\mathbb{R}^2$.

Definition. Let X and Y be topological spaces. A function $f : X \to Y$ is **continuous** when, for any open set $U \subseteq Y$, its preimage $f^{-1}(U)$ is an open set of X.

What are some examples of continuous functions?

- We already saw that a function $f : \mathbb{R} \to \mathbb{R}$ is continuous in the usual topology if and only if it satisfies the standard δ and ϵ definition, so all of the usual continuous functions we know and love (e.g. x^2, $\sin(x)$, e^x, $|x|$, etc.) are in fact continuous in the sense of topological spaces.
- For any topological space X, the identity function $f : X \to X$ (i.e., $f(x) = x$ for all $x \in X$) is continuous: for any open subset $U \subseteq X$, we obviously have that $f^{-1}(U) = U$ is an open subset of X.
- Let X be a set given the discrete topology (recall that this means *every* subset of X will be open), and let Y be *any* topological space. Then *any* function $f : X \to Y$ is continuous because, for any open subset $U \subseteq Y$, $f^{-1}(U)$ is automatically an open subset of X, because *every* subset of X is open.
- Let Y be a set given the trivial topology (recall that this means the *only* open subsets of Y will be $\varnothing$ and Y), and let X be *any* topological space. Then *any* function $f : X \to Y$ is continuous, because $f^{-1}(\varnothing) = \varnothing$ and $f^{-1}(Y) = X$, both of which are necessarily open subsets of X.

Theorem. Compositions of continuous functions are continuous.

Proof. Let X, Y, and Z be topological spaces, and suppose $f : X \to Y$ and $g : Y \to Z$ are continuous functions. For any open subset $U \subseteq Z$, $g^{-1}(U)$ is an open subset of Y, because g is continuous. Because $g^{-1}(U)$ is an open subset of Y, $f^{-1}(g^{-1}(U))$ is an open subset of X, because f is continuous. Thus

$$(g \circ f)^{-1}(U) = f^{-1}(g^{-1}(U))$$

is an open subset of X for any open subset $U \subseteq Z$, so that $g \circ f$ is continuous. $\square$

Excellent! We've taken our experience with $\mathbb{R}$, and seen how the definition of topological space naturally comes up. We also have a precise definition of what it means for a function to be continuous, and this definition is particularly valuable because it works in far greater generality than only the δ and ϵ definition for $\mathbb{R}$. Now we can begin to answer our motivating questions from the start of the talk.

Definition. A topological space X is called **connected** when there **do not** exist disjoint non-empty open sets $U, V \subseteq X$ whose union is X. In symbols: $\nexists$ open sets $U, V \subseteq X$ such that

- $U, V \neq \varnothing$ (non-empty)
- $U \cap V = \varnothing$ (disjoint)
- $U \cup V = X$

If X is not connected (i.e., there *do* exist such U and V), we call X **disconnected**.

Theorem. Let X be a topological space. Then X is connected $\iff$ the only subsets of X which are both closed and open are $\varnothing$ and X.

Proof. ($\Longrightarrow$) Suppose X is connected. If S is a closed subset of X, then by definition $S^c = \{x \in X \mid x \notin S\}$ is open. Note that $S \cap S^c = \varnothing$ and that $S \cup S^c = X$. Therefore, if S is a subset of X that is both closed and open, then S and S^c are disjoint open sets whose union is all of X. But because X is connected, this is impossible unless either $S = \varnothing$ or $S^c = \varnothing$, i.e. either $S = \varnothing$ or $S = X$.

($\Longleftarrow$) Suppose the only subsets of X which are both closed and open are $\varnothing$ and X. If $U, V \subseteq X$ are disjoint open sets whose union is all of X, then because $U \cap V = \varnothing$ and $U \cup V = X$, we have that $U = V^c$ and $V = U^c$, so that by definition, both U and V are also closed. Thus, U is either $\varnothing$ or X, and V is either $\varnothing$ or X. Because $U \cap V = \varnothing$ and $U \cup V = X$, either $U = \varnothing$ and $V = X$, or $U = X$ and $V = \varnothing$. Thus, the sets U and V are not both non-empty. Thus, there are no disjoint non-empty open sets U and V for which $U \cup V = X$, so X is connected. $\square$

Theorem. $\mathbb{R}$ is connected.

Proof. Suppose $U, V \subseteq \mathbb{R}$ are open subsets of $\mathbb{R}$, such that $U, V \neq \varnothing$, and $U \cap V = \varnothing$, and $U \cup V = \mathbb{R}$. Since U and V are non-empty, there exists some $x \in U$, $y \in V$. Because $U \cap V = \varnothing$, we cannot have $x = y$. WLOG, let's suppose that $x < y$. Let $I_0 = [x, y]$.

Because $U \cup V = \mathbb{R}$, every real number is in either U or V. In particular, $\frac{x+y}{2}$ is in either U or V. If $\frac{x+y}{2} \in V$, then the interval $[x, \frac{x+y}{2}]$ is not contained entirely U, nor contained entirely in V (because $x \in U$ and $\frac{x+y}{2} \in V$); in symbols, we say that $[x, \frac{x+y}{2}] \not\subseteq U$ and $[x, \frac{x+y}{2}] \not\subseteq V$. On the other hand, if $\frac{x+y}{2} \in U$, then the same argument applies to the interval $[\frac{x+y}{2}, y]$, and we would have $[\frac{x+y}{2}, y] \not\subseteq U$ and $[\frac{x+y}{2}, y] \not\subseteq V$.

At any rate, at least one of the intervals $[x, \frac{x+y}{2}]$ and $[\frac{x+y}{2}, y]$ is not fully contained in either U or V. Let's choose one, and call it I_1. In fact, by induction, we can choose a nested sequence of closed intervals I_n satisfying, for all n,

- $I_n \subset I_{n-1}$
- length of $I_n = \dfrac{y-x}{2^n}$
- $I_n \not\subseteq U$ and $I_n \not\subseteq V$

The first two properties actually imply that **the intersection $\bigcap_{n=0}^{\infty} I_n$ consists of a single point** z! Intuitively, this is because the intervals form a nested chain $I_0 \supset I_1 \supset I_2 \cdots$, and the sizes of the intervals I_n are decreasing geometrically; but in order to prove this rigorously, we would technically require something called the Cantor Intersection Theorem, that will not be explained here. However, it is not too hard to show that $\bigcap_{n=0}^{\infty} I_n$ has **at most** one element (the key fact is that the intersection is not actually empty).

Because $U \cup V = \mathbb{R}$ and $U \cap V = \varnothing$, either $z \in U$ or $z \in V$, but not both. If $z \in U$, then because U is open, there is some open interval (a, b) such that $z \in (a, b)$ and $(a, b) \subseteq U$. Let $\delta = \min\{z - a, b - z\}$ (note that $\delta > 0$ because $a < z < b$). Then $z \in (z - \delta, z + \delta)$, and $(z - \delta, z + \delta) \subseteq (a, b) \subseteq U$. But for any

$$n > \log_2\left(\frac{y-x}{\delta}\right),$$

the length of the interval I_n is

$$\frac{y-x}{2^n} < \frac{y-x}{2^{\log_2\left(\frac{y-x}{\delta}\right)}} = \frac{y-x}{\left(\frac{y-x}{\delta}\right)} = \delta.$$

Because $\{z\} = \bigcap_{n=0}^{\infty} I_n$, we have that $z \in I_n$. Any interval of length $< \delta$ containing z is contained in $(z - \delta, z + \delta)$, so $I_n \subseteq (z - \delta, z + \delta) \subseteq U$. But we constructed the I_n's in such a way that we would never have $I_n \subseteq U$; we have reached a contradiction. The same problem occurs if $z \in V$. Thus, the open sets U and V cannot exist; in other words, $\mathbb{R}$ is connected. $\square$

Some other examples of connected topological spaces are $(0, 1)$, $[0, 1]$, $\mathbb{R}^n$, and $\mathbb{S}^1$.

Some examples of disconnected topological spaces are

- $\{0, 1\}$ with the discrete topology, because $\{0\}$ and $\{1\}$ are disjoint non-empty open sets whose union is the entire space, $\{0, 1\}$.
- $(0, 1) \cup (2, 3)$ with the subspace topology from $\mathbb{R}$, because $(0, 1)$ and $(2, 3)$ are disjoint non-empty open sets whose union is the entire space, $(0, 1) \cup (2, 3)$.

Theorem. Let X and Y be topological spaces, and let $f : X \to Y$ be a continuous function. If X is connected, then the image $f(X) \subseteq Y$ is connected.

Proof. First, we give $f(X)$ the subspace topology, so that $A \subseteq f(X)$ is open precisely when there is an open set $B \subseteq Y$ such that $A = f(X) \cap B$.

Suppose the subsets $U, V \subseteq f(X)$ are disjoint non-empty open subsets of $f(X)$ such that $U \cup V = f(X)$. By the definition of the subspace topology, there are open sets $C, D \subseteq Y$ such that $U = f(X) \cap C$ and $V = f(X) \cap D$. Because f is continuous, $f^{-1}(C)$ and $f^{-1}(D)$ are open subsets of X. Thus, the sets

$$f^{-1}(U) = f^{-1}(f(X) \cap C) = f^{-1}(f(X)) \cap f^{-1}(C) = X \cap f^{-1}(C) = f^{-1}(C)$$

$$f^{-1}(V) = f^{-1}(f(X) \cap D) = f^{-1}(f(X)) \cap f^{-1}(D) = X \cap f^{-1}(D) = f^{-1}(D)$$

are open in X.

Because $U \cap V = \varnothing$, we have that $f^{-1}(U) \cap f^{-1}(V) = f^{-1}(U \cap V) = f^{-1}(\varnothing) = \varnothing$.

Because $U \cup V = f(X)$, we have that $f^{-1}(U) \cup f^{-1}(V) = f^{-1}(U \cup V) = f^{-1}(f(X)) = X$.

Because $U, V \neq \varnothing$ and $U, V \subseteq f(X)$, we also have $f^{-1}(U), f^{-1}(V) \neq \varnothing$.

Thus, $f^{-1}(U)$ and $f^{-1}(V)$ are disjoint non-empty open subsets of X such that $f^{-1}(U) \cup f^{-1}(V) = X$, which contradicts the assumption that X is connected. Therefore, there do not exist such U and V; therefore, $f(X)$ is connected. □

Theorem. Suppose X is a connected topological space, and Y is a discrete topological space. Then any continuous function $f : X \to Y$ is constant.

Proof. A function $f : X \to Y$ is constant precisely when there is only one element in its image $f(X)$. Because Y is discrete, *every* subset of Y is open. When we give $f(X)$ the subspace topology, $f(X)$ is itself discrete, because any $S \subseteq f(X)$ is open in Y, so that $S = S \cap f(X)$ is the intersection of an open set of Y with $f(X)$, and hence S is open in $f(X)$ by the definition of the subspace topology.

Thus, if the image $f(X)$ contained at least two distinct elements $a, b \in f(X)$, then $\{a\}$ and $\{a\}^c = \{y \in f(X) \mid y \neq a\}$ would be disjoint, non-empty (because $b \in \{a\}^c$) open subsets of $f(X)$ such that $\{a\} \cup \{a\}^c = f(X)$; therefore, $f(X)$ is disconnected, contradicting our previous theorem. Therefore, the function f cannot have more than 2 elements in its image; and thus, it must have exactly one, i.e. f must be constant. □

Corollary. The only continuous functions from $\mathbb{R}$ to $\mathbb{Z}$ are constant.

Proof. $\mathbb{R}$ is connected and $\mathbb{Z}$ is discrete, so the previous theorem shows that any continuous function from $\mathbb{R}$ to $\mathbb{Z}$ must be constant. □

Corollary. Let X be a connected topological space. For any continuous function $f : X \to \mathbb{R}$, if a and b are in the image of f, then so is all of $[a, b]$.

Proof. Let X be connected, and let $f : X \to \mathbb{R}$ be a continuous function. Let $a = f(x)$, $b = f(y)$ for some $x, y \in X$. By our theorem, $f(X)$ is connected.

Suppose $c \in (a, b)$ is not in $f(X)$, i.e. there is no $z \in X$ such that $c = f(z)$. Because $(-\infty, c)$ and (c, ∞) are open sets of $\mathbb{R}$, the sets $U = f(X) \cap (-\infty, c)$ and $V = f(X) \cap (c, \infty)$ are open sets in the subspace topology of $f(X)$. Because $c \notin f(X)$, we have that $U \cup V = f(X)$. Because $(-\infty, c) \cap (c, \infty) = \varnothing$, we have that $U \cap V = \varnothing$. Because $c \in (a, b)$, we have that $a < c < b$, and therefore $a \in U$ and $b \in V$; thus, $U, V \neq \varnothing$.

Thus, U and V are disjoint non-empty open subsets of $f(X)$ whose union is $f(X)$, so that $f(X)$ is disconnected, which is a contradiction. □

The above result is commonly known as the Intermediate Value Theorem when the connected set X we started with is an interval $[c, d]$ or (c, d) in $\mathbb{R}$.

Let X be a topological space, and let $\{x_n\}_{n \in \mathbb{N}}$ be a sequence of elements of X.

Definition. We say $\{x_n\}$ **converges to** $a \in X$ if, for all open sets $U \subseteq X$ such that $a \in U$, there exists an $m \in \mathbb{N}$ such that $x_n \in U$ for all $n > m$. In symbols:

$$\forall \text{ open } U \subseteq X \text{ such that } a \in U, \exists m \in \mathbb{N} \text{ such that } n > m \implies x_n \in U.$$

Definition. A topological space X is **Hausdorff** if, for any two distinct $a, b \in X$, there exist open sets $U, V \subseteq X$ such that $a \in U$, $b \in V$, and $U \cap V = \varnothing$.

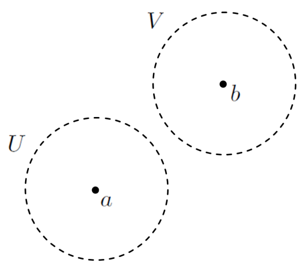

One (slightly silly) way of remembering the definition of the Hausdorff property is that distinct points can be “housed off” from each other using open sets.

An example of a Hausdorff topological space is $\mathbb{R}$, with its usual topology: given any distinct $a, b \in \mathbb{R}$, we can choose $U = (a - \delta, a + \delta)$ and $V = (b - \delta, b + \delta)$ where $\delta = \frac{|b-a|}{2}$ is half the distance between a and b.

A set X with more than 1 element, with the trivial topology, is an example of a non-Hausdorff topological space. Because X has more than 1 element, we can choose distinct $a, b \in X$; but the only open sets are $\varnothing$ and X, so there are no disjoint non-empty subsets U and V, regardless of whether $a \in U$ and $b \in V$.

Theorem. Let X be a Hausdorff topological space. Then any sequence $\{x_n\}_{n \in \mathbb{N}}$ in X converges to at most one element of X.

Proof. Suppose $\{x_n\}$ converges to $a \in X$, and also to $b \in X$, where $a \neq b$. Because X is Hausdorff, there exist open sets $U, V \subseteq X$ such that $a \in U$, $b \in V$, and $U \cap V = \varnothing$. By hypothesis, $\{x_n\}$ converges to a and b, so by definition, there exists an $m_U \in \mathbb{N}$ such that $n > m_U \implies x_n \in U$, and also there exists an $m_V \in \mathbb{N}$ such that $n > m_V \implies x_n \in V$. But then, when $n > \max\{m_U, m_V\}$, we have that $x_n \in U \cap V = \varnothing$, which is a contradiction. □