LECTURE: METRIC SPACES AND TOPOLOGY

Today: We'll go over some useful notions in topology such as metric spaces, convergence, and open/closed sets.

1. Metric Spaces

Video: Metric Spaces

A metric space is a nonempty set, together with a function called a metric, which measures the distance between any pair of points in the set.

Definition:

Let X be an arbitrary set. A function $d : X \times X \to \mathbb{R}$ is a **metric** on X if the following conditions hold for all $x, y, z \in X$

(1) $d(x, y) \geq 0$

(2) $d(x, y) = 0 \Leftrightarrow x = y$

(3) $d(x, y) = d(y, x)$

(4) $d(x, z) \leq d(x, y) + d(y, z)$ (Triangle Inequality)

We call the pair (X, d) a **metric space**.

(4) is called the triangle inequality and generalizes the notion from geometry that says: "The length of a side of a triangle must be less than or equal to the sum of the lengths of the other two sides."

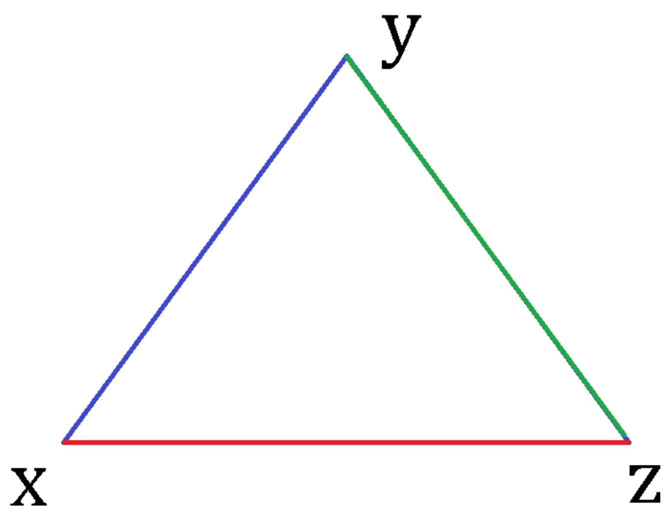

From the triangle inequality, we can derive the

Reverse Triangle Inequality

$$d(x, z) \geq |d(x, y) - d(y, z)|$$

This generalizes that "The length of a side of a triangle must be greater than or equal to the difference of the lengths of the other two sides."

From now on, assume that we are working in a metric space (X, d)

2. Examples

Example 1:

The **Euclidean distance** on $\mathbb{R}^n$ is defined by

$$d(x, y) = |x - y|$$

More precisely, if $x = (x_1, \dots x_n)$ and $y = (y_1, \dots, y_n)$, then

$$d(x, y) = \sqrt{\sum_{k=1}^{n} (x_k - y_k)^2}$$

Example 2:

The **discrete metric** on any set X is defined by

$$d(x, y) = \begin{cases} 0 & \text{if } x = y \\ 1 & \text{otherwise} \end{cases}$$

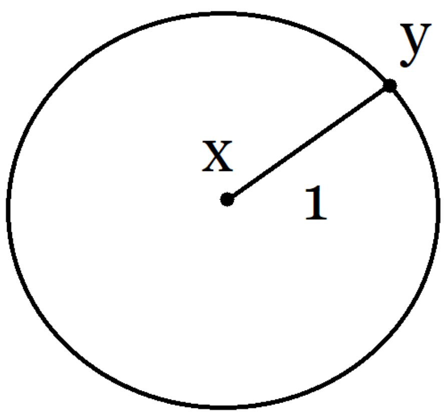

In other words, with the metric d, *all* the points in X are distance 1 apart. Freaky, isn't it? But it's a great source of counterexamples!

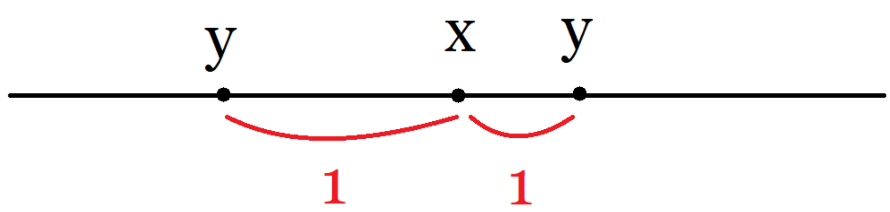

Note: The discrete metric seems weird for $\mathbb{R}$, but is more natural in other examples: If $X = \{1, 2, 3\}$ with the discrete metric. Then X is just an equilateral triangle!

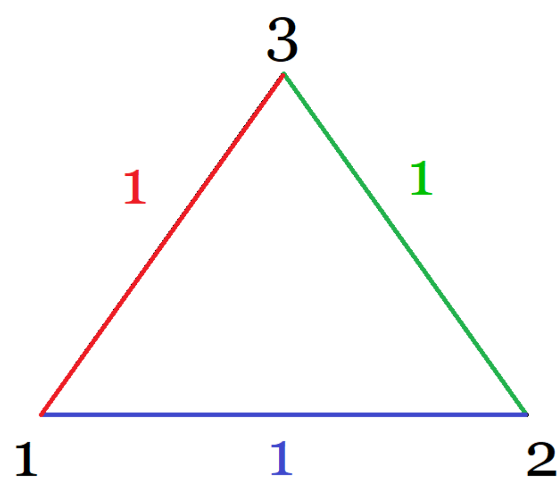

Example 3: Three metrics on $\mathbb{R}^2$

If $x = (x_1, x_2)$ and $y = (y_1, y_2)$ then

(1) Manhattan distance / taxicab metric / ℓ^1 metric:

$$d_1(x, y) = |x_1 - y_1| + |x_2 - y_2|$$

(2) Euclidean distance / ℓ^2 metric:

$$d_2(x, y) = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2}$$

(3) Maximum distance / chessboard metric / ℓ^∞ metric:

$$d_\infty(x, y) = \max\{|x_1 - y_1|, |x_2 - y_2|\}$$

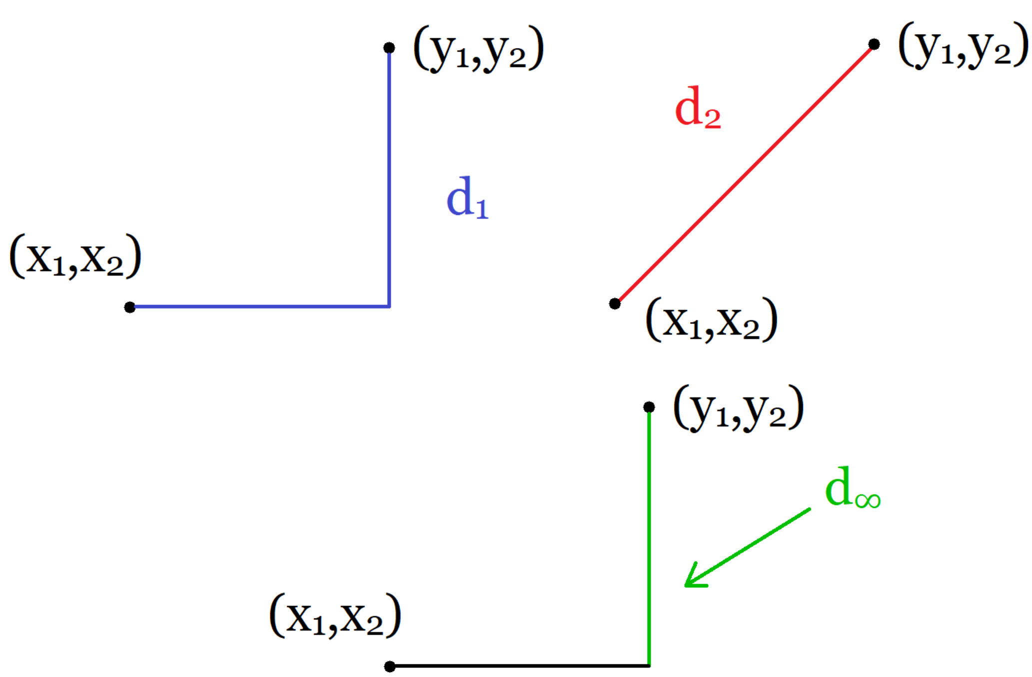

Note: d_1 s called the taxicab metric because taxicabs in New York can't just go diagonally from (x_1, x_2) to (y_1, y_2) without crashing into buildings, they have to go right, and then up.

Example 4:

Let $C([a, b])$ be the space of cont. $\mathbb{R}$-valued functions on $[a, b]$

Then the sup metric on $C([a, b])$ is defined by

$$d(f, g) = \max_{x \in [a,b]} |f(x) - g(x)|$$

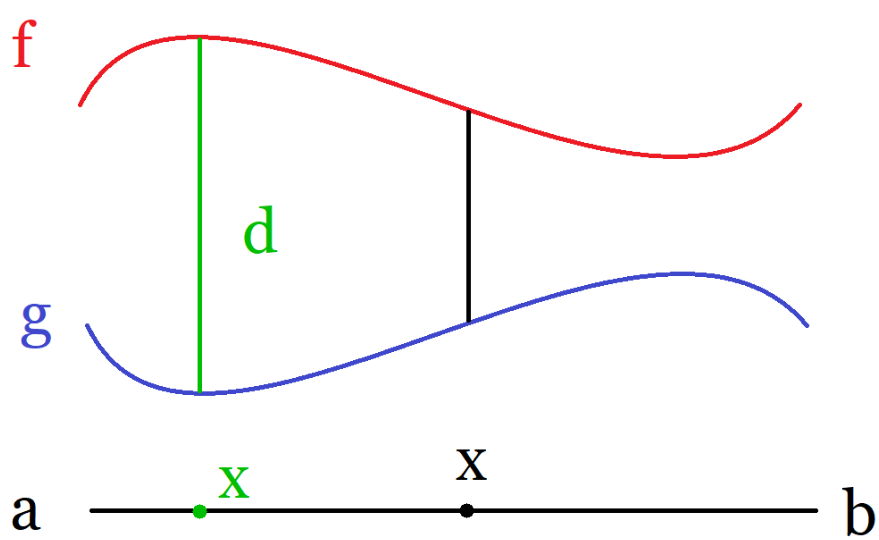

Note: Since a continuous function attains its maximum and minimum on a closed interval (extreme value theorem), this is well-defined.

Example 5: Integral Metrics

Let $X = C[a,b]$ then the following are metrics on X

$$d(f,g) = \int_a^b |f(x) - g(x)| \, dx$$

$$d(f,g) = \sqrt{\int_a^b |f(x) - g(x)|^2 \, dx}$$

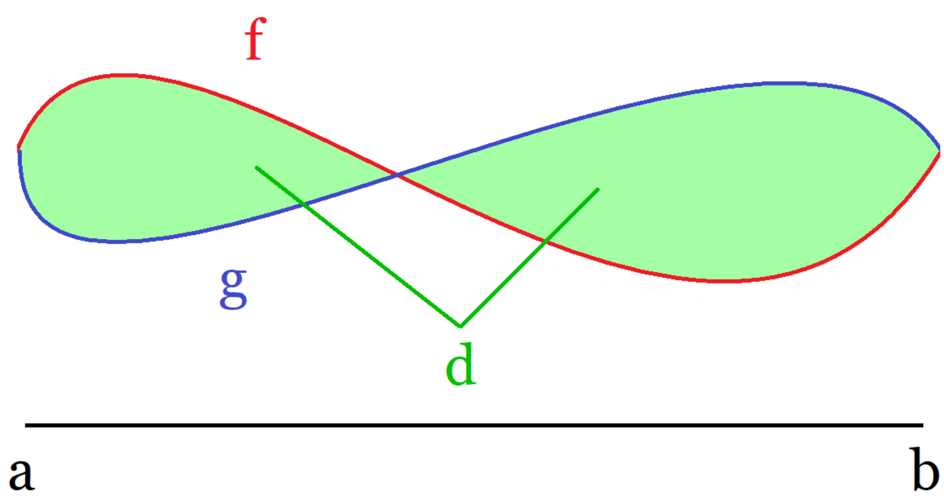

Those are very natural on X if you remember that an integral is just a sum. The second one is nice because X becomes a Hilbert space.

Example 6: Distance between sets

If A and B are two subsets of $\mathbb{R}$ (or of any metric space), then

$$d(A,B) = \inf\{|a-b| \mid a \in A, b \in B\}$$

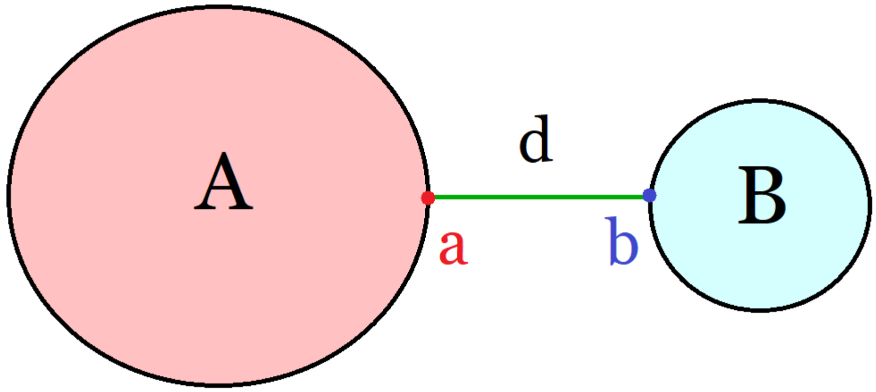

Example 7: Information Theory

(1) **Hamming Distance:** Distance between two strings of identical length: number of positions at which the two strings are different.

$$d(0100, 0010) = 2$$

(2) **Levenshtein Distance:** Distance between two arbitrary strings: minimum number of single-character edits (insertions, deletions or substitutions) required to change one string into the other.

$$d(0100, 0010) = 2$$

(3) **Damerau-Levenshtein distance:** Same as Levenshtein distance, except swaps of adjacent characters are also allowed operations.

$$d(0100, 0010) = 1$$

Note: "More than 80% of all human misspellings can be expressed by a single error of one of the four types" (Damerau, 1964).

Example 8: Geodesic Distance on Connected Graphs

The Number of edges in a shortest path connecting two vertices.

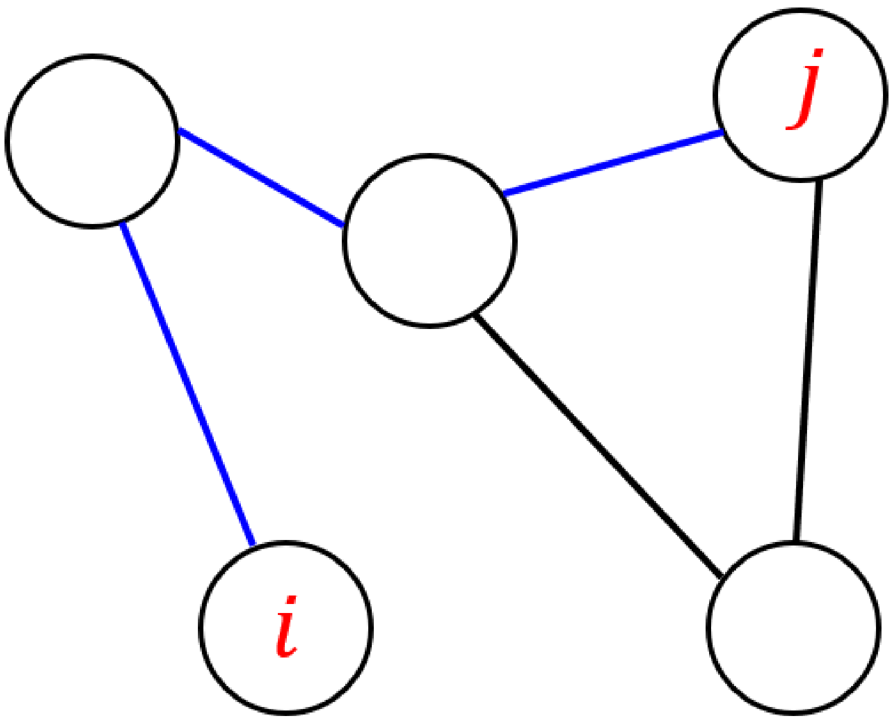

Here $d(i, j) = 3$

Note: There may not be a unique shortest path, but the geodesic distance is unique. The graph needs to be connected for this to work.

Take-away: *Everything* we're going to say about metric spaces holds **ALL** the examples at once, so we're really killing multiple birds with one stone! **THIS** is the power of abstract mathematics!

3. Product metrics

Here is a fun way of constructing new metrics from old ones

Definition:

If X and Y are metric spaces, then the **Cartesian Product** $X \times Y$ is simply

$$X \times Y = \{(x, y) \mid x \in X, y \in Y\}$$

Similarly if $X_1, \cdots, X_n$ are metric spaces then

$$\prod_{k=1}^{n} X_k = X_1 \times \cdots \times X_n := \{(x_1, \ldots, x_n) : x_k \in X_k\}.$$

Take a finite sequence of metric spaces $(X_1, d_1), \ldots, (X_n, d_n)$, potentially equipped with different metrics. Then any of the following are metrics on $X_1 \times \cdots \times X_n$: We call these product metrics.

Definition:

(1) $d_1(x,y) = \sum_{k=1}^{n} d(x_k, y_k)$

(2) $d_2(x,y) = \sqrt{\sum_{k=1}^{n} d(x_k, y_k)^2}$

(3) $d_\infty(x,y) = \max_{k=1,\ldots,n} d(x_k, y_k)$

The same thing works for **countable** products of metric spaces:

$$\prod_{k=1}^{\infty} X_k := \Big\{ \{x_k\} : x_k \in X_k \Big\}.$$

This is not as easy as before, since we are now dealing with infinite sums and need to worry about points being a finite distance from each other. Luckily we can use the following trick, which allows us to construct metrics from other metrics:

Fact:

If $f : [0,\infty) \to [0,\infty)$ is an increasing concave function such that

$$f(x) = 0 \Leftrightarrow x = 0$$

Then $f(d(x,y))$ is also a metric.

Note that f does not have to be smooth. Important examples are:

$$f(x) = \min\{x, 1\} \text{ or } f(x) = \frac{x}{x+1}$$

Both of these functions "cut off" the metric at 1, so that the greatest possible distance between any two points is 1.

We can use the second of these "cutoff" functions to define a metric for a countable product of metric spaces $\{(X_k, d_k)\}$:

Example:

$$d(x, y) = \sum_{k=1}^{\infty} \frac{1}{2^k} \left(\frac{d_k(x_k, y_k)}{1 + d_k(x_k, y_k)} \right)$$

Note: There are many variations of this: We could alternatively use the "cutoff" metric $\min\{d_k(x_k, y_k), 1\}$. And we can replace $1/2^k$ with any convergent series of positive terms.

4. Convergence

Video: Convergence in $\mathbb{R}^n$

The neat thing about metric spaces is that it's really easy to generalize the notion of convergence to those spaces.

Definition:

If (x_n) is a sequence in X, then $x_n \to L$ if for all $\epsilon > 0$ there is N such that if $n > N$, then $d(x_n, L) < \epsilon$

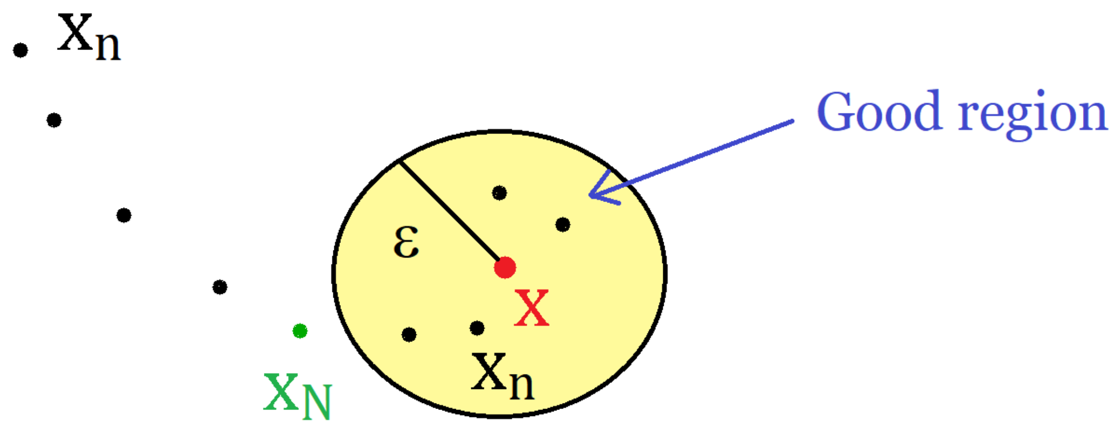

Related to this is the notion of a Cauchy sequence. Intuitively, a sequence $\{x_n\}$ is a Cauchy sequence if its elements get arbitrarily close to *each other* (rather than approach a limit).

Definition:

$\{x_n\}$ is a **Cauchy sequence** if for all $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that $d(x_n, x_m) < \epsilon$ for all $n, m \geq N$

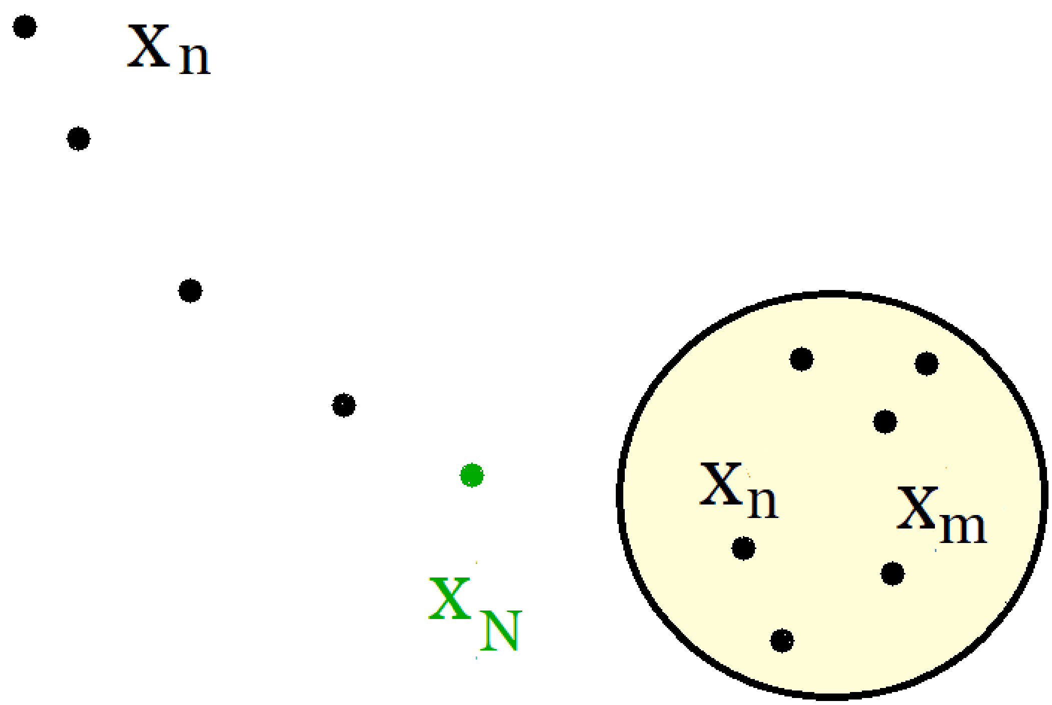

We sometimes abbreviate this as $d(x_m, x_n) \to 0$. Every convergent sequence is a Cauchy sequence, but not every Cauchy sequence is convergent

Non-Example:

Take $\mathbb{Q}$ with $d(x,y) = |x-y|$, and consider the sequence $\{x_n\}$, with $x_1 = 1$ and

$$x_{n+1} = \frac{x_n}{2} + \frac{1}{x_n}$$

This is a Cauchy sequence, but its limit is $\sqrt{2}$, which is not in $\mathbb{Q}$, so it doesn't converge in $\mathbb{Q}$

Analogy: Just because you see a large crowd (Cauchy), it doesn't mean that the crowd is going somewhere (convergent)

Definition:

A metric space is **complete** if every Cauchy sequence converges

Hence $\mathbb{Q}$ with the above metric is not complete.

An example of a complete metric space is $\mathbb{R}$ (the completeness of $\mathbb{R}$ follows from its construction). Using the product metric on $\mathbb{R}^n$ (the maximum version of the metric is easiest here) and the completeness of $\mathbb{R}$, it follows that $\mathbb{R}^n$ is complete.

Note: Every metric space can be completed

Convergence Tests

Two ways to show a sequence $\{x_n\}$ converges.

(1) Use the definition of convergence to show that $x_n \to L$. This means that we need a guess for what L.

(2) Work in a complete metric space, and show $\{x_n\}$ is a Cauchy sequence. This is often easier, since we do not need a guess for the limit, but it has the drawback of not giving us the actual limit.

5. Open Sets

Video: Open Sets

For this, we first need to define what an open ball is.

Definition:

The open ball centered at x and radius r is:

$$B(x,r) = \{y \in S \mid d(x,y) < r\}$$

That is, the set of points that are a distance of at most r away from x.

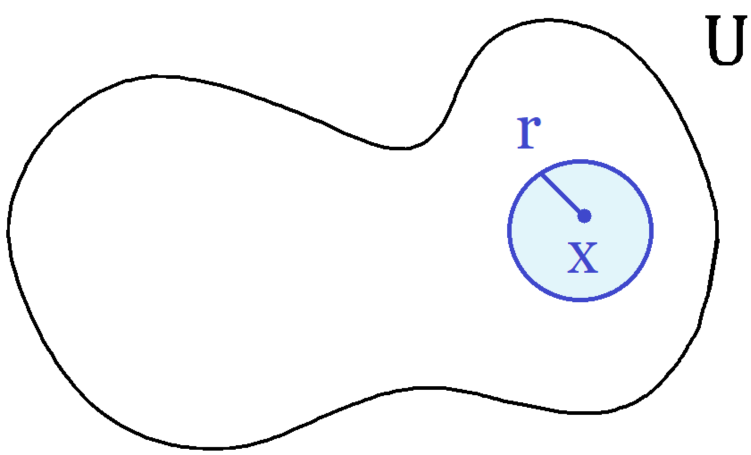

Note: You may see this written as $B_r(x)$ or $U_r(x)$

Using this, we can define the concept of an open set:

Definition:

A subset $U \subseteq X$ is **open** if for all $x \in U$ there is $r > 0$ such that $B(x,r) \subseteq U$.

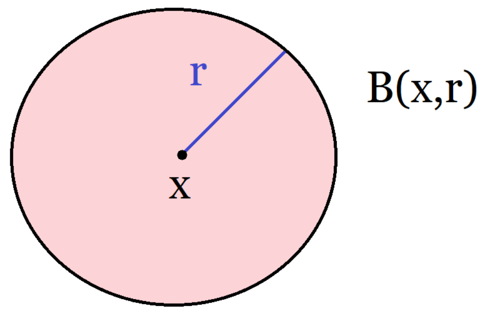

In other words, for every point in U there is some tiny ball that is contained in U.

Interpretation: For every point x in E, you can move around x a little bit and still be in your set. So there is some wiggle room/breathing room around every point.

This open ball property is so useful that we can give it a name

Definition:

$x \in U$ is an **interior point** of U if $B(x, r) \subseteq U$ for some $r > 0$.

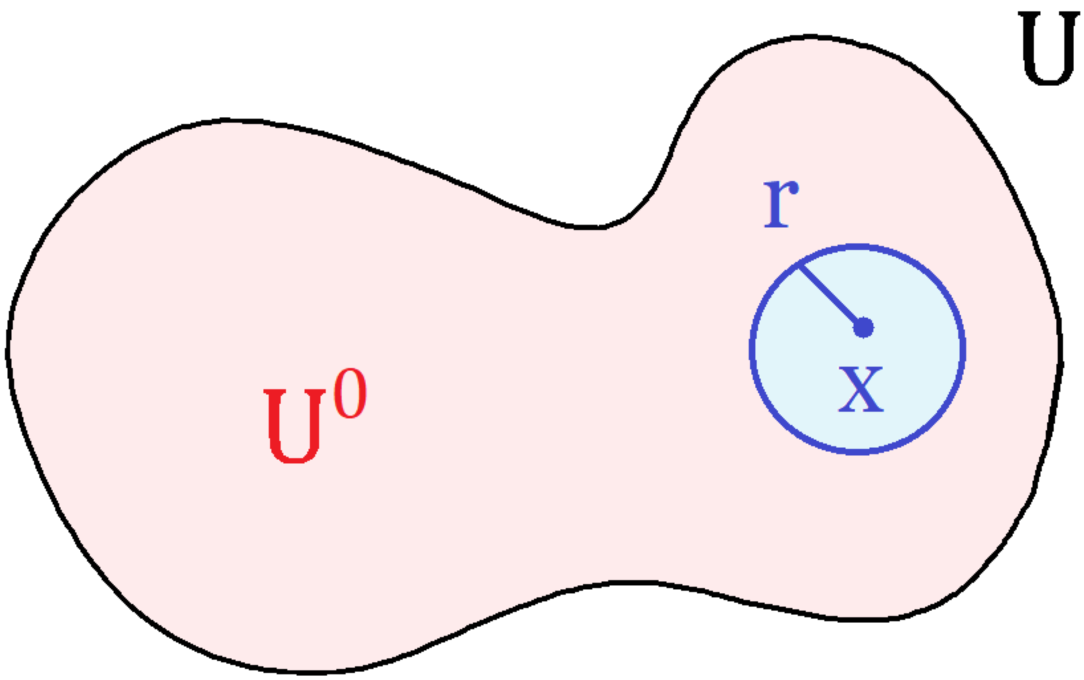

Note: It's similar to the definition of open set except here we're *fixing* a point x. Before, this was true *for all* x.

Then a set U is open iff it consists entirely of interior points

6. Closed Sets

Video: Closed Sets

On the other side of the spectrum comes the notion of a closed set, which has to do with limits of sequences.

Definition:

$K \subseteq X$ is **closed** if, whenever (x_n) is a sequence in K that converges to x, then $x \in K$

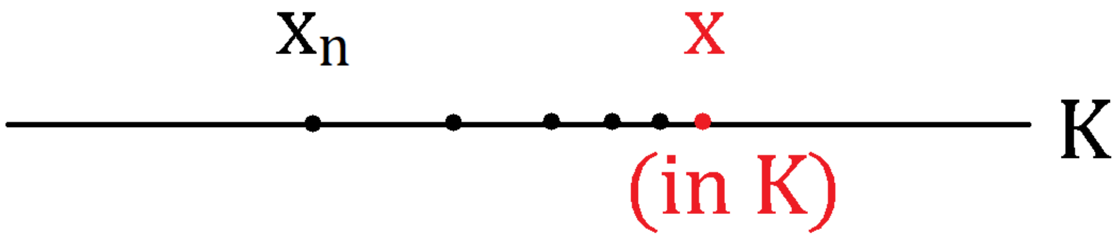

In other words, K must contain all the limits of all the sequences in it.

Non-Example:

$(0, 1]$ is not closed

For instance, $x_n = \frac{1}{2^n}$ is a sequence in $(0, 1]$ that converges to 0, but $0 \notin (0, 1]$

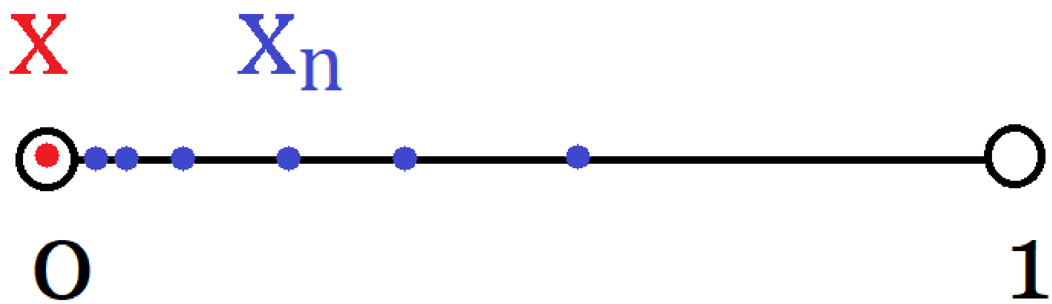

In some sense, you can escape $(0, 1]$ by taking limits, like a prisoner getting out of prison.

Sometimes you will see closed sets defined in terms of limit points

Definition:

We say x is a **limit point** of K if there is a sequence (x_n) in K that converges to x.

Alternatively x is a **limit point** if for every $\epsilon > 0$ $B(x, \epsilon)$ contains a point in E (note that x may or may not be in E).

Non-Example:

0 is a limit point of $(0, 1]$ because for all $\epsilon > 0$, $B(0, \epsilon) = (-\epsilon, \epsilon)$ contains points of $(0, 1]$. Notice $0 \notin (0, 1]$

Definition:

We a subset $K \subseteq X$ is **closed** if it contains all of its limit points.

Basic Properties:

(1) K is closed if and only if $X \backslash K$ is open

(2) The union of *any* collection of open sets is open

(3) The intersection of finitely many open sets is also open

Warning: The intersection of *infinitely* open sets isn't necessarily open:

Non-Example:

Consider $U_n = \left(-\frac{1}{n}, \frac{1}{n}\right)$ in $\mathbb{R}$

Then each U_n is open, but the intersection of all U_n is $\{0\}$, which is not open.

It follows from $(1)-(3)$ that arbitrary intersections of closed sets if closed and finite unions of closed sets are closed.

Note: Topologists actually use those properties to define open sets! More precisely

Definition:

A **topology** on a set X is a family $\mathcal{T}$ of subsets of X which contains $\emptyset$ and X, is closed under arbitrary unions, and under finite intersections. The sets in $\mathcal{T}$ is called the open sets.

7. Closure, Interior, and Boundary

Definition:

$$\overline{E} = \text{ Set of limit points of } E$$

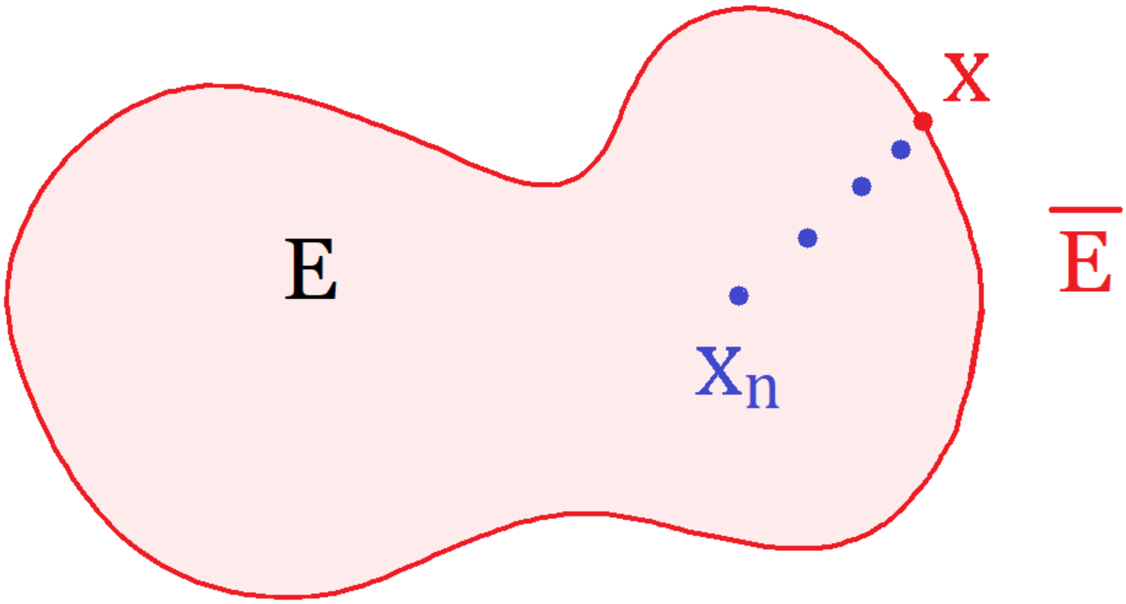

It is the smallest closed set containing E

Example:

If $E = (0,1]$, then $\overline{E} = [0,1]$

Think of it as the set of all possible destinations starting in E

Definition:

$$E^\circ = \text{ Set of all interior points of } E$$

It is the union of all open sets contained in E, as well as the largest open set contained in E

Example:

If $E = [0,1]$, then $E^\circ = (0,1)$

Because for any point *other than* 0 or 1, we can fit a ball inside $[0,1]$.

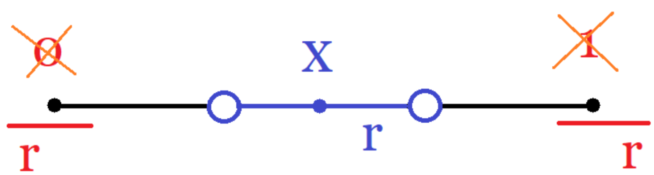

Fact:

E is open iff $E = E^\circ$

E is closed iff $\overline{E} = E$

Definition:

The boundary ∂E of set E is the set of points x such that every ball $B(x,\epsilon)$ contains at least one point of E and one point of $X \backslash E$

Equivalently, it is defined as $\partial E = \overline{E} \backslash E^\circ$

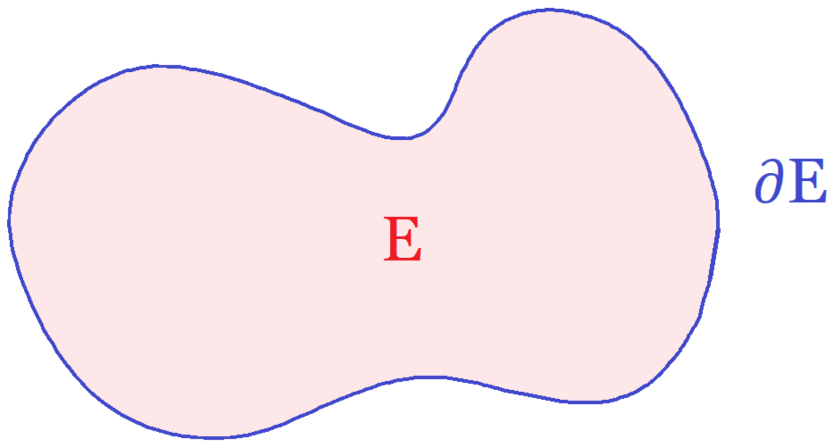

Think of the boundary as the edge of a cliff: You see both the cliff-part and the sea-part

8. Equivalent Metrics

Sometimes it does not matter much what metric we use, i.e. different metrics give us the same convergent sequences and the same open sets.

Definition:

Two metrics d_1 and d_2 on X are **equivalent** (or comparable) if there exist constants C_1 and C_2 such that for all $x, y \in X$,

$$C_1 d_1(x,y) \leq d_2(x,y) \leq C_2 d_1(x,y)$$

Example:

Because of the following identity

$$\max_{k=1,\ldots,n} |x_k - y_k| \leq \sum_{k=1}^{n} |x_k - y_k| \leq n \max_{k=1,\ldots,n} |x_k - y_k|$$

The Euclidean, taxicab, and maximum metrics in $\mathbb{R}^n$ are all strongly equivalent.

If two metrics on X are equivalent, the open sets and convergent sequences of X are the same, so for our purposes they are pretty much the same.