33. How to Count

Exam Thursday. You have an at-home exam Thursday. It covers chapters 6-11, and 26-28, along with the currently released notes. You may refer to the book and to my notes, and any notes you took during class. Other references are off limits. I will email you the exam at 5pm. You will email me the answers by 6:30pm.

9/22/20

We start by re-learning how to count.

33.1 Counting

There are few mathematical concepts more fundamental than counting. Because most of us have known how to count since an early age, we rarely give it much thought. Nonetheless, it pays to examine counting closely as there are hidden complications when numbers become very large.

One way to count the elements of a set is to put them in a one-to-one correspondence with the elements of another set. Young children often do this using their fingers. They associate each finger with one of the objects they are counting, which works fine up to ten. We can do the same sort of thing by using a bijective (one-to-one and onto) mapping.

We formalize this by saying that A and B have the same *cardinality*, denoted $\#A = \#B$, if there is a bijective correspondence between them, a one-to-one function mapping A onto B. The concept of cardinality formalizes what it means for sets to have the same number of number of elements.[1]

The sets $A = \{1, 2, 3\}$ and $B = \{\text{red, white, blue}\}$ each contain 3 elements. Both have the same cardinality. One mapping that establishes this is $f(1) = \text{red}$, $f(2) = \text{white}$, $f(3) = \text{blue}$. This is not the only mapping that shows these sets have the same cardinality. Another that works is $g(1) = \text{white}$, $g(2) = \text{red}$, $g(3) = \text{blue}$. We could even use a mapping in the opposite direction, such as $h(\text{red}) = 3$, $h(\text{white}) = 1$, and $h(\text{blue}) = 2$. As you can, we can easily create a bijective correspondence between any set with three elements and $\{1, 2, 3\}$.

Instead of using their fingers, adults count using a mental reference set, the *counting numbers* $\mathbb{N} = \{1, 2, 3, \dots\}$. We count by mentally putting the set to be counted into bijective correspondence with an initial segment of $\mathbb{N}$. Thus if we count A and find it has n elements, we have put it into a bijective correspondence with $\{1, 2, 3, \dots, n\}$. We can define $n = \#(\{1, \dots, n\})$.

[1] If we want to consider the order of items, first, second, third, we would need ordinal numbers rather than cardinal numbers. Halmos, *Naive Set Theory* (1960) has a nice introduction to both cardinal and ordinal arithmetic.

33.2 Cardinality is an Equivalence Relation

We don't have to worry about counting producing inconsistent results because having equal cardinality is an equivalence relation.

What does that mean?

Equivalence Relation. A binary relation $\sim$ on a set X is an *equivalence relation* if:

1. It is *reflexive*. $x \sim x$.
2. It is *symmetric*. If $x \sim y$ then $y \sim x$.
3. It is *transitive*. If $x \sim y$ and $y \sim z$, then $x \sim z$.

Theorem 33.2.1. *For sets A and B, define $A \sim B$ if and only if $\#A = \#B$. Then $\sim$ is an equivalence relation.*

Proof. We have to show that all three properties are satisfied. Unpacking the definition, $A \sim B$ means there is a bijective mapping $f\colon A \to B$.

(1) $\sim$ is reflexive. Define $f\colon A \to A$ by $f(x) = x$. Then f is a bijective mapping from A to A, so $A \sim A$.

(2) $\sim$ is symmetric. If $A \sim B$, there is a bijective mapping $f\colon A \to B$. Since f is bijective, the unique inverse mapping, f^{-1}, is a bijective mapping $f^{-1}\colon B \to A$. This shows that $B \sim A$, proving symmetry.

(3) $\sim$ is transitive. Suppose we have three sets, with $\#A = \#B$ and $\#B = \#C$. Then there are bijective mappings $f\colon A \to B$ and $g\colon B \to C$. Define $h\colon A \to C$ by $h = g \circ f$. The composition of bijective maps is bijective, so $\#A = \#C$, establishing transitivity.

It follows that $\sim$ is an equivalence relation. $\square$

33.3 Larger and Smaller Sets

We can also define smaller and larger cardinality using one-to-one correspondences. A set A is *no bigger than* set B if there is a one-to-one mapping of A **into** B. In symbols, we write $\#A \le \#B$. In that case, there is a one-to-one mapping between A and $f(A)$, a subset of B. If A is no bigger than B and B is no bigger than A, they have the same cardinality, even if the sets are infinite. This is a consequence of the Schroeder-Bernstein Theorem, which holds for sets of any size.

Schroeder-Bernstein Theorem. *Let A and B be sets such that A is no bigger than B and B is no bigger than A, then there is a one-to-one mapping of A onto B. In other words, A and B have the same cardinality.*

Proof. See pg. 88 of Halmos (1960). □

When dealing with finite sets, proper subsets have fewer elements and so a smaller cardinality. That is no longer true when sets are infinite. Proper subsets of infinite sets can even have the same cardinality as the original!

Consider the counting numbers $\mathbb{N} = \{1, 2, 3, \dots\}$ and non-negative integers $\mathbb{Z}_0 = \{0, 1, 2, \dots\}$.[2] Even though there is clearly an extra element in $\mathbb{Z}_0$, both $\mathbb{Z}_0$ and $\mathbb{N}$ have identical cardinality. Here $f(n) = n + 1$ is the desired one-to-one mapping from $\mathbb{Z}_0$ onto $\mathbb{N}$. A similar mapping has some importance in monetary theory in Example .

[2] The names "counting numbers" and "natural numbers" are not standardized and definitions vary depending on whether or not zero is included. My own preference is to use "natural numbers" to refer to the counting numbers together with zero. Simon and Blume use $\mathbb{N}$ for the counting numbers, and I follow that so you will have a consistent notation.

33.4 Cardinality Examples

► Example 33.4.1: Fiat Money. Consider the following overlapping generations model. One agent (generation) is born at each time, and lives two periods. Each agent earns 1 unit of a consumption good when young, and none when old. (There is only one generation alive at time zero, two at all other times.) Each person consumes c_y when young and c_o when old, obtaining utility $c_y + c_o$.

Given prices $p_t = 1$, consuming the endowment maximizes utility. Since markets clear (no trade), we have an equilibrium. However, if the young generation at each time $t > 0$ gives $1 - 2^{-t}$ units of its endowment to the old generation, everyone is made better off since the generation young at time t gains $1 - 2^{-t-1}$ when old at a cost of $1 - 2^{-t}$ units when young. This yields a utility gain of 2^{-t-1} for generation t. The transfers are $(1/2, 3/4, 7/8, 15/16, \dots)$ starting at time $t = 1$.

This improvement cannot be an equilibrium. However, we can make it an equilibrium allocation by introducing a fiat currency. Give 1 unit of currency to the generation that is old at time zero. If goods have price one in each time period, and money has price $1/2$ at $t = 1$, and appreciates every period according to the schedule $(1/2, 3/4, 7/8, \dots)$, the old generation can pay for the transfer with cash. The young generation takes the cash, and spends it in the next period to buy a bit more when when old (due to the appreciation). The desired pattern of transfers is now an equilibrium. (This is not the only monetary equilibrium.) ◄

Another example of an infinite subset with the same cardinality as the set is the positive integers and the positive even integers.

► Example 33.4.2: Cardinality of Sets of Odd and Even Numbers. The sets of positive odd numbers and of positive even numbers each have the same cardinality as the set of counting numbers $\mathbb{N}$. Consider the following mapping defined for each counting number, $f(n) = 2n$. This is clearly one-to-one and maps the counting numbers onto the positive even numbers. The mapping $g(n) = 2n - 1$ accomplishes the same thing for the positive odd numbers. It follows that all three sets have the same cardinality, denoted $\aleph_0$. ◄

33.5 Degrees of Infinity

Defining cardinality via one-to-one correspondences allows us to distinguish different degrees of infinity, even though some infinities have the same cardinality. We say an infinite set is *countably infinite* if it has the same cardinality as the counting numbers $\mathbb{N} = \{1, 2, 3, \ldots\}$, which we denote as $\aleph_0$. We call a set *countable* if it is either finite or countably infinite. Sets that are not countable are called *uncountable*. Uncountable sets are necessarily infinite.

We will use the following lemma to show that certain sets are countable.

Lemma 33.5.1. *Let $A \subset \mathbb{N}$. If there is a mapping f of A onto a set B, then B is countable*

Proof. If B is empty we may take A to be empty and there is nothing to prove.

If B is non-empty, A must also be non-empty. For each $b \in B$ let $g(b) = \min\{a \in A : f(a) = b\}$. Since f is onto, the set will be non-empty and will have a minimum, so $g(b)$ exists. Then g maps B onto a subset of A so $\#B \geq \#A$.

The fact that we have a mapping from A onto B shows $\#B \leq \#A$. By the Schroeder-Bernstein Theorem, $\#A = \#B$. $\square$

33.6 The Rational Numbers are Countable

One big countable set is the set of rational numbers, $\mathbb{Q} = \{p/q : p, q \text{ are integers with } q \neq 0\}$.

Proposition 33.6.1. *The set of rational numbers is countable.*

Proof. Since the counting numbers are a subset of the rational numbers, the Schroeder-Bernstein Theorems tells us it is enough to find a mapping of the counting numbers onto the rational numbers.

Let
$$f(x) = \begin{cases} 1/(4x-1) & \text{for } 0 < x < 1/4, \\ 4x-2 & \text{for } 1/4 \leq x \leq 3/4, \\ 1/(4x-3) & \text{for } 3/4 < x < 1. \end{cases}$$

Then $f(x)$ is a one-to-one mapping of the rational numbers between 0 and 1 onto all of the rational numbers (see Lemma 33.7.1 on the next page). It follows that we need only consider the rational numbers between 0 and 1. Consider the following array of rational numbers:
$$\begin{array}{ccccc} 1/2 & 1/3 & 1/4 & 1/5 & \ldots \\ & 2/3 & 2/4 & 2/5 & \ldots \\ & & 3/4 & 3/5 & \ldots \\ & & & 4/5 & \ldots \end{array}$$

All of the rational numbers in $(0, 1)$ are listed. The number p/q is found in row p, column q as well as other locations such as row $2p$, column $2q$. Now define a mapping $g(1) = 1/2$, $g(2) = 1/3$, $g(3) = 2/3$, $g(4) = 1/4$, $g(5) = 2/4$, $g(6) = 3/4$, $g(7) = 1/5$, etc. This maps $\mathbb{N}$ onto the rational numbers between 0 and 1. It follows that $f \circ g$ maps $\mathbb{N}$ onto $\mathbb{Q}$. This is our required mapping of $\mathbb{N}$ onto $\mathbb{Q}$. $\square$

33.7 The Function f

Lemma 33.7.1. *Define* $f\colon (0,1) \to \mathbb{R}$ *by*

$$f(x) = \begin{cases} 1/(4x-1) & \text{for } 0 < x < 1/4, \\ 4x-2 & \text{for } 1/4 \leq x \leq 3/4, \\ 1/(4x-3) & \text{for } 3/4 < x < 1. \end{cases}$$

Then

(a) *The function* f *maps* $(0,1)$ *onto* $\mathbb{R}$.
(b) *If* $r \in \mathbb{Q}$ *is rational, then* $f(r) \in \mathbb{Q}$.
(c) *If* $f(x) \in \mathbb{Q}$, *then* x *in* $\mathbb{Q}$.

Moreover, f *is a one-to-one mapping of the rationals in* $(0,1)$ *onto all of the rational numbers.*

Proof. We first prove each of the steps.

(a) When $0 < x < 1/4$, $-1 < 4x-1 < 0$, so $f(x) \in (-\infty, -1)$. When $1/4 \leq x \leq 3/4$, $-1 \leq 4x-2 \leq 1$, so $f(x) \in [-1, +1]$. When $3/4 < x < 1$, $0 < 4x-3 < 1$, so $f(x) \in (1, +\infty)$. The inverse function for f is

$$f^{-1}(x) = \begin{cases} 1/4 + 1/4x & \text{for } -\infty < x < -1, \\ (x+2)/4 & \text{for } -1 \leq x \leq +1, \\ 3/4 + 1/4x & \text{for } +1 < x < +\infty \end{cases}.$$

Its existence shows that f is a bijective mapping.

(b) Since f is defined by rational operations, $f(r)$ will be rational whenever r is rational.

(c) The inverse function f^{-1} is also defined by rational operations, so $f^{-1}(r)$ is rational whenever r is rational. Thus any $r \in \mathbb{Q}$ is in the image of f. In other words, f maps onto $\mathbb{Q}$.

Because (a) f maps onto $\mathbb{R}$, and both (b) maps rational numbers to rational numbers, and (c) every rational number is the image of a rational number, we can conclude that f is a one-to-one mapping of the rationals between $(0,1)$ onto all of the rational numbers. $\square$

33.8 Countability of Countable Unions

We can also show that countable unions of countable sets are countable.

Lemma 33.8.1. *A countable union of countable sets is countable.*

Proof. Let A_i be countable for each i and $A = \cup_{i \in I} A_i$ where I is a countable set. Since each A_i is countable, we can label the elements of A_i so that $A_i = \{a_1^i, a_2^i, \ldots\}$. We can also label the elements of I with $I = \{i_1, i_2, \ldots\}$. Note that A_i or I may have only finitely many elements.

Define a function $g \colon \mathbb{Q} \to \cup_i A_i$ by

$$g(r) = \begin{cases} a_{i_q}^p & \text{when } r = p + 1/q \text{ for } p, q \in \mathbb{N} \\ a_1^1 & \text{otherwise.} \end{cases}$$

This clearly maps $\mathbb{Q}$ onto $\cup_i A_i$, showing that $\cup_i A_i$ is countable. □

One consequence is that the set of open intervals with rational endpoints is itself countable. For each rational number q, consider the set of open intervals starting at q,

$$A_q = \{(q, r) : r \in \mathbb{Q}\}.$$

Each A_q is countable, and the set of rational intervals is the countable union $\cup_{q \in \mathbb{Q}} A_q$, which is countable by Lemma 33.8.1.

This set of intervals seems pretty big, but is still countable, no bigger than the set of counting numbers. This tells that uncountable sets have to be really big!

33.9 The Real Numbers are Uncountable!

We know such a set. The real numbers are **not countable**. We show that using *Cantor's Diagonal Argument* in the following example.

▶ Example 33.9.1: The Real Numbers are Uncountable. In fact, we will show that the unit interval is uncountable. Since it is a subset of the real numbers, the real numbers must also form an uncountable set.

Suppose $[0, 1]$ is countably infinite. Then we can put it into one-to-one correspondence with the counting numbers. This allows us to list the elements of $[0, 1]$ in the order of their correspondence with $1, 2, 3$, etc.

Write the real numbers on the list in decimal form. The list will look something like this:

$$\begin{array}{c} 0.182083820\ldots \\ 0.398020208\ldots \\ 0.972832711\ldots \\ 0.002820337\ldots \\ 0.137992365\ldots \\ \vdots \end{array}$$

We can now use an argument due to Georg Cantor (Cantor, 1891) to find a real number in the interval $[0, 1]$ that is not on the list. Take the n^{th} digit of the n^{th} number on the list. If this digit is less that 7, change it to 8. If it is more than 7, change it to 5. The diagonal has been shown in red above. In our case, it is $0.19289\ldots$ which we transform to $0.85855\ldots$. This number differs in digit n from the n^{th} number on the list, so it cannot itself be on the list.[3] This contradicts the fact that all real numbers in $[0, 1]$ are on the list and shows that the real numbers are not countable. ◀

The cardinality of $\mathbb{R}$ is denoted $\mathfrak{c}$ and referred to as the *cardinality of the continuum*. We now know that $\mathfrak{c} > \aleph_0$. You may wonder whether there are any sets that are larger than the rational numbers but smaller than the reals.

This is a deep question!

For many years, the answer was unknown. In 1963, Paul Cohen found a remarkable result. This question doesn't have an definitive answer! The question is undecidable. There are consistent mathematical systems, in which the axioms of set theory (and so all usual mathematics) hold, where there are such sets. Equally, there are other systems, consistent with normal mathematics, where such sets do not exist. Cohen showed how to construct such mathematical systems (Cohen, 1963).

[3] Notice that the way we changed the digits avoids problems with numbers that can be written two ways, such as 1 and $0.999\ldots$.

12. Limits and Open Sets

This chapter focuses on material from Simon and Blume's Chapter 12. It focuses on limits, open sets, closed sets, and related concepts. In other words, it's about topology.

A large chunk of topology is about points being either near or far. This holds true in normed and metric spaces where we have a precise notion of nearness, or more general spaces where nearness is not a numerical concept.

There are several ways to approach topology. One is to start with sequences and metric spaces. Then use the metric to define open balls and open sets, and convergent sequences to define closed sets.

Another is to start with the open sets themselves. It has the advantage of generality, but dodges the question of how we recognize a open set. To be sure, there are ways to answer the question: topological bases and subbases, neighborhood bases, nets, filters, etc. But the constructions are more involved.

We will build things up via the more concrete approach of metric spaces and sequences, and discuss more general topologies to see the overall structure.

12.1 Sequences

A sequence of vectors in $\mathbb{R}^m$ looks like this:[1]

$$\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots\}$$

Let $\mathbb{N} = \{1, 2, 3, \dots\}$ be the set of natural numbers, meaning the positive integers.

Sequence. A *sequence* in a set A is a mapping $n \mapsto x_n$ from $\mathbb{N}$ into A.

We usually indicate a sequence by $\{x_n\}_{n=1}^{\infty}$ or simply $\{x_n\}$. Sometimes we will write sequences using superscripts, $\{\mathbf{x}^n\}$, particularly when writing sequences of vectors. A sequence may be described by a formula, such as $x_n = n^2 + 1/n$, or by otherwise describing the sequence.

Examples include

► Example 12.1.1: Examples of Sequences.

$$\begin{aligned} x_n &= n & &\{1, 2, 3, 4, \dots\}, \\ x_n &= 1/n & &\{1, 1/2, 1/3, 1/4, \dots\}, \\ x_n &= (-1)^{n+1} & &\{+1, -1, +1, -1, \dots\}, \text{ and} \\ x_n &= (-1)^{n+1} n^2 & &\{1, -4, 9, -16, 25, -36, \dots\}. \end{aligned}$$

All of the above examples are sequences. They are mappings from $\mathbb{N}$ to $\mathbb{R}$. ◄

The examples above show that the long-run behavior of the sequence can vary. The first continually increases without bound. The second decreases toward zero, the third bounces back and forth between two values, and the fourth one makes bigger and bigger oscillations as n increases.

[1] We use $\mathbb{R}^m$ here instead of $\mathbb{R}^n$ because we will be using n in sequences.

12.2 Convergent Sequences in $\mathbb{R}$

One type of sequence that plays an important role in mathematics is a convergent sequence. We start by considering convergent sequences of real numbers. After a little practice, we will graduate to $\mathbb{R}^m$ and beyond.

Convergent Sequences in $\mathbb{R}$. We say that x_n *converges to* x, written $x_n \to x$, or $\lim_{n\to\infty} x_n = x$, if for every $\varepsilon > 0$, there is a positive integer N with $|x_n - x| < \varepsilon$ for every $n \geq N$. A sequence is *convergent* if it converges to some x.

The definition means that no matter what we pick as a standard for closeness to x (call it ε), the rest of the sequence will eventually meet that standard. The *tails* of the sequence, $\{x_n\}_{n=N}^{\infty}$ must always be close, must eventually stay within ε of the limit x. Here "eventually" translates to when n is large enough, when $n \geq N$ for some N. Convergent sequences can't just occasionally visit a point, they have to settle down and stay near it.

▶ **Example 12.2.1: The Sequence $\{1/n\}$ Converges.** Let's see how this works with the sequence $x_n = 1/n$. I claim this sequence converges to $x = 0$. To prove it, we take any $\varepsilon > 0$. We must find an N so that $|1/n - 0| = 1/n < \varepsilon$ for $n \geq N$. Let's choose N with $N > 1/\varepsilon$. Then

$$n \geq N > 1/\varepsilon \Rightarrow x_n = 1/n < \varepsilon,$$

so

$$|x_n - 0| = |1/n - 0| = |1/n| < \varepsilon$$

for every $n \geq N$. ◀

NB: The N we pick depends on the ε we start with. Although many choices of N will work, we have to make sure it works for our particular ε.

▶ **Example 12.2.2: The Sequence $\left\{1 + 1/\sqrt{n}\right\}$ Converges.** This has limit 1. Set $x_n = 1 + 1/\sqrt{n}$ and consider $|x_n - 1| = 1/\sqrt{n}$. Choose any $\varepsilon > 0$. We want $1/\sqrt{n} < \varepsilon$, meaning $1/\varepsilon < \sqrt{n}$, or equivalently $1/\varepsilon^2 < n$. So we pick $N > 1/\varepsilon^2$. Then for $n \geq N$, $1/\sqrt{n} \leq 1/\sqrt{N} < 1/\sqrt{\varepsilon^{-2}} = \varepsilon$. It follows that for $n \geq N$,

$$|x_n - 1| = 1/\sqrt{n} < \varepsilon,$$

which tells us that $\lim_n x_n = 1$. ◀

12.3 Some Sequences Don't Converge

But do all sequences converge? What if we get one that doesn't? What if we try to show it converges to something?

▶ **Example 12.3.1: A Sequence without a Limit.** The sequence $x_n = (-1)^n$ does not converge to anything. Examination of the sequence reveals the problem. It is $\{-1, +1, -1, +1, -1, \ldots\}$. It never settles down but continues to bounce back-and-forth, back-and-forth, back-and-forth between -1 and $+1$.

There must be an ε where the definition fails. The key thing is that ε must be small enough that only only of $+1$ and -1 can be near the same point. We could try to guess an ε where the definition fails.

Or, we could try to show it converges and see where things go wrong. So we try to make the sequence converge. Suppose the sequence has limit x and take $\varepsilon > 0$.

In that case, there is an N with $|x_n - x| < \varepsilon$ for all $n \geq N$. Then also $|x_{n+1} - x| < \varepsilon$. Using the triangle inequality, we find

$$|x_n - x_{n+1}| \leq |x_n - x| + |x - x_{n+1}| < 2\varepsilon.$$

But $|x_n - x_{n+1}| = 2$, so $2 < 2\varepsilon$. It only works if $\varepsilon > 1$! This fails for any $\varepsilon \leq 1$.

We've shown that x_n cannot converge to any limit because we cannot squeeze the terms closer together than 1. Any $\varepsilon \leq 1$ will fail. The sequence $\{(-1)^n\}$ is a sequence without a limit. ◀

This is not the only way convergence can fail. Sequences such as $\{1, 2, 3, \ldots\}$ $(x_n = n)$ and $\{(-1)^n n\} = \{-1, 2, -3, 4, \ldots\}$ also don't have limits.

12.4 Limits are Unique

What about the sequences that bounce around? Can a sequence have two limits? No. A sequence can only converge to one limit. The reasoning is similar to that in Example , where $\{(-1)^n\}$ did not converge.

We will prove this for sequences in $\mathbb{R}$, but it is true in considerable generality, in metric spaces. The statement and proof of the general case are very similar to the following theorem.

Theorem 12.4.1. *Let $\{x_n\}$ be a sequence of real numbers. Suppose $x_n \to x$ and $x_n \to y$. Then $x = y$.*

Proof. By way of contradiction, **suppose** $x \neq y$.

Set $\varepsilon = |x - y|/2$. By the contradiction hypothesis, $\varepsilon > 0$. Since $x_n \to x$, there is an N_1 with $|x_n - x| < \varepsilon$ for $n \geq N_1$. Take $N_2 \geq N_1$ with $|x_n - y| < \varepsilon$ for $n \geq N_2$. Then for $n \geq N_2 \geq N_1$, both $|x_n - x| < \varepsilon$ and $|x_n - y| < \varepsilon$. It follows that

$$|x - y| \leq |x - x_n| + |x_n - y| < \varepsilon + \varepsilon = |x - y|.$$

But this is impossible, **so** $x \neq y$ **is impossible**.

The only possibility remaining is $x = y$. $\square$

12.5 Limits and Inequalities

One important property of limits in $\mathbb{R}$ is that they preserve the weak order relations, $\leq$ and $\geq$.

Theorem 12.5.1. *Let* $\{x_n\}$ *and* $\{y_n\}$ *be sequences of real numbers with* $x_n \leq y_n$. *Then* $\lim_n x_n \leq \lim_n y_n$.

Proof. Let $x = \lim_n x_n$ and $y = \lim_n y_n$. We will prove this by contradiction.

Suppose that $x > y$. Then we can set $\varepsilon = (x - y)/2 > 0$. Because $x_n \to x$, we can choose N_1 with $|x_n - x| < \varepsilon$ for $n \geq N_1$. Then we choose $N_2 \geq N_1$ so that $|y_n - y| < \varepsilon$ for $n \geq N_2$. Both inequalities hold for $n \geq N_2$.

Now for $n \geq N_2$, $|x - x_n| < \varepsilon$ and $|y - y_n| < \varepsilon$. These imply $x_n < x + \varepsilon$ and $y_n < y + \varepsilon$. Then for $n \geq N_2$,

$$\begin{aligned} x &< x_n + \varepsilon && \text{since } |x_n - x| < \varepsilon \\ &\leq y_n + \varepsilon && \text{by hypothesis} \\ &< y + 2\varepsilon && \text{since } y_n < y + \varepsilon \\ &= y + (x - y) && \text{definition of } \varepsilon \\ &= x. \end{aligned}$$

But $x < x$ is impossible. It follows **that** $x > y$ **cannot be true**.

Therefore $x \leq y$. □

The result also holds if we reverse the inequalities.

Theorem 12.5.2. *Let* $\{x_n\}$ *and* $\{y_n\}$ *be sequences of real numbers with* $x_n \geq y_n$. *Then* $\lim_n x_n \geq \lim_n y_n$.

Proof. Adapt the proof of Theorem 12.5.1. □

The corollary follows immediately.

Corollary 12.5.3. *Let* $\{x_n\}$ *and* $\{y_n\}$ *be convergent sequences of real numbers. with* $x_n < y_n$ *(or* $x_n > y_n$*). Then* $\lim_n x_n \leq \lim_n y_n$ *(or* $\lim_n x_n \geq \lim_n y_n$*).*

We only stated Corollary 12.5.3 to show you that it can't be strengthened to obtain a strong inequality such as $\lim_n x_n > \lim y_n$.

▶ **Example 12.5.4: Inequality Counterexample.** Let $x_n = 1 + 1/(n + 1)$ and $y_n = 1 + 1/(n + 1)^2$. Then $x_n > y_n$, but $\lim_n x_n = 1 = \lim_n y_n$. ◀

12.6 Convergence in Metric Spaces

Before we look at too many results for $\mathbb{R}$, let's upgrade the definition of convergence. We will skip the special cases of Euclidean $\mathbb{R}^m$ and even general normed spaces. We head directly to metric spaces, which are a natural home for convergent sequences.

Of course, every normed vector space is a metric space, so we are not really skipping them. Our new definition of convergence will work in normed spaces too. So how do we change the definition? In the old definition, we used $|x_n - x|$, which is the distance $d(x_n, x)$. Switching them is the only change we need!

Convergent Sequences in Metric Spaces. Let $\{\mathbf{x}_n\}$ be a sequence in a metric space (X, d). We say that $\mathbf{x}_n$ *converges to* x, written $\mathbf{x}_n \to \mathbf{x}$, or $\lim_{n\to\infty} \mathbf{x}_n = \mathbf{x}$, if for every $\varepsilon > 0$, there is a positive integer N with $d(\mathbf{x}_n, \mathbf{x}) < \varepsilon$ for every $n \ge N$. A sequence is *convergent* if it converges to some $\mathbf{x} \in X$.

If we specialize to $\mathbb{R}^m$, the distance becomes $\|\mathbf{x}_n - \mathbf{x}\|$. Finally, in $\mathbb{R}$, the condition is exactly what we used earlier, that $|x_n - x| < \varepsilon$ for all $n \ge N$.

We will sometimes be interested in functions of two or more variables. We expand our notion of convergence to include such cases.

Product Space and Convergence. Let $X_1, X_2, \ldots, X_k$ be metric spaces. The *product space* $X_1 \times X_2 \times \cdots \times X_k$ is the set of all k-tuples $(\mathbf{x}_1, \ldots, \mathbf{x}_k)$ with $\mathbf{x}_i \in X_i$ for all $i = 1, \ldots, k$.

The sequence of k-tuples $\left\{\left(\mathbf{x}_1^n, \ldots, \mathbf{x}_k^n\right)\right\}_{n=1}^{\infty}$ *product converges to* $(\mathbf{x}_1, \ldots, \mathbf{x}_k$ if $\mathbf{x}_i^n \to \mathbf{x}_i$ for every $i = 1, \ldots, k$. We sometimes denote product convergence by $\mathbf{x}_n \xrightarrow{p} \mathbf{x}$.

► **Example 12.6.1**: $\mathbb{R}^m$. We previously treated $\mathbb{R}^m$ as a normed vector space using any of the ℓ_p norms. It can also be thought of as a product space,

$$\mathbb{R}^m = \overbrace{\mathbb{R} \times \mathbb{R} \times \cdots \times \mathbb{R}}^{m \text{ times}}$$

where the distance on each copy of $\mathbb{R}$ is $d(x, y) = |x - y|$. Then $\mathbf{x}_n \xrightarrow{p} \mathbf{x}$. ◄

12.7 Convergence in $\mathbb{R}^m$

We now have two definitions of convergence in $\ell_2^m = (\mathbb{R}^m, \|\cdot\|_2)$, one using the ℓ_2 norm and the other treating $\mathbb{R}^m$ as a product space. In fact, they are the same. Convergence in ℓ_2^m can be thought of as convergence in the ℓ_2 norm, or it can thought of as convergence in every coordinate.

Theorem 12.7.1. *Let $\{\mathbf{x}^n\}$ be a sequence in ℓ_2^m. Then $\lim_n \mathbf{x}^n = \mathbf{x}$ if and only if $x_i^n \to x_i$ where x_i^n is the i^{th} component of $\mathbf{x}^n$.*

Proof. Only if case: Suppose $\mathbf{x}^n \to \mathbf{x}$. Let $\varepsilon > 0$. We can choose N with $\|\mathbf{x}^n - \mathbf{x}\|_2 < \varepsilon$ for $n \ge N$. Then for $n \ge N$,

$$\begin{aligned} |x_i^n - x_i| &= \left(|x_i^n - x_i|^2\right)^{1/2} \\ &\le \left(\sum_{j=1}^m |x_j^n - x_j|^2\right)^{1/2} \\ &= \|\mathbf{x}^n - \mathbf{x}\|_2 \\ &< \varepsilon \end{aligned}$$

for all $n \ge N$. It follows that $\lim_n x_i^n = x_i$ for each $i = 1, \dots, m$.

If case: Suppose that for each $i = 1, \dots, m$, $\lim_n x_i^n = x_i$. Let $\varepsilon > 0$ be arbitrary. For each i, we can find an N_i with

$$\left|x_i^n - x_i\right| < \frac{\varepsilon}{m^{1/2}} \tag{12.7.1}$$

whenever $n \ge N_i$. It follows that equation (12.7.1) holds for every $i = 1, \dots, m$ when $n \ge N = \max_i N_i$. Then for all $n \ge N$,

$$\begin{aligned} \|\mathbf{x}^n - \mathbf{x}\|_2 &= \left(\sum_{j=1}^m |x_j^n - x_j|^2\right)^{1/2} \\ &< \left(\sum_{j=1}^m \left(\frac{\varepsilon}{m^{1/2}}\right)^2\right)^{1/2} \\ &= \left(m\left(\frac{\varepsilon^2}{m}\right)\right)^{1/2} \\ &= \varepsilon \end{aligned}$$

showing that $\mathbf{x}^n \to \mathbf{x}$. $\square$

12.8 Open Balls

9/29/20

As with sequences in $\mathbb{R}$, when $x_n \to x$ in a metric space (X, d), the sequence eventually stays near the point x. We'll now use a slightly different method to describe this that focuses a little more on where the sequence is.

Open Balls. Let (X, d) be a metric space. For $x \in X$ and $\varepsilon > 0$, we define the *open ball of radius ε about x*, $B_\varepsilon(x)$ by

$$B_\varepsilon(x) = \{y \in X : d(x, y) < \varepsilon\}.$$

The set of points with $d(x, y) = \varepsilon$ is **not** included in $B_\varepsilon(x)$. These points are included in the *closed ball* $\overline{B_\varepsilon(x)} = \{y \in X : d(x, y) \leq \varepsilon\}$. We make a distinction between the ball and its boundary, the *sphere*, $\{y \in X : d(x, y) = \varepsilon\}$. In $\mathbb{R}^m$ we can refer to the m-ball and $(m - 1)$-sphere. Here the dimension m refers to dimension of the set (which we have not covered) rather than the dimension of the space. In two dimensions, the ball is sometimes called a *disk* and its boundary a circle.

The $\mathbb{R}^2$ case, where $d(x, y) = \|x - y\|_2$, is illustrated in Figure 12.8.1.

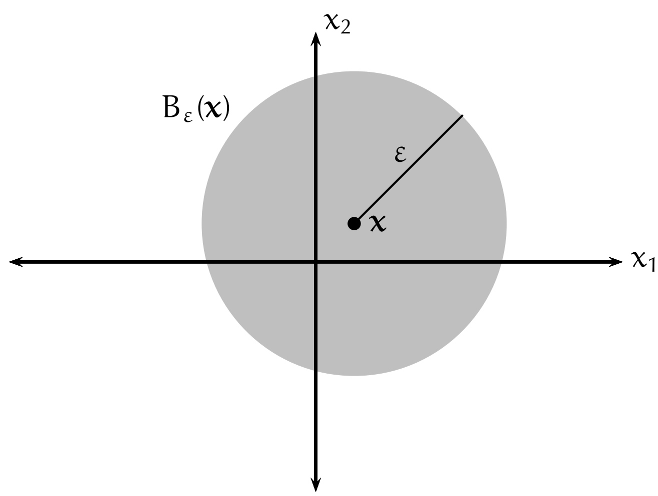

Figure 12.8.1: The open ε-ball about x in $\mathbb{R}^2$ is the disk of radius ε centered at x. The circle of points at distance ε from x are not included in the open disk, but are part of the closed disk.

12.9 Convergence in Metric Spaces

A set $N \subset X$ is a *neighborhood* of $x \in X$ if there is an $\varepsilon > 0$ with $B_\varepsilon(x) \subset N$. A neighborhood need not be open as it only has to contain a ball around x, not any other point. In particular, the closed ε-ball, $\overline{B_\varepsilon(x)}$ is a neighborhood of x.

We can now restate the definition of a convergent sequence that applies in any metric space (X, d).

Convergent Sequences in Metric Spaces II. Let $\{x_n\}_{n=1}^\infty$ be a sequence in a metric space (X, d). The sequence $\{x_n\}_{n=1}^\infty$ *converges to* x if for every $\varepsilon > 0$, there is a positive integer N with $x_n \in B_\varepsilon(x)$ for every $n \geq N$.

Alternatively, $\{x_n\}$ *converges to* x if for every neighborhood N of x, there is a positive integer N with $x_n \in N$ for every $n \geq N$.

Of course, saying that $x_n \in B_\varepsilon(x)$ for all $n \geq N$ is exactly the same as saying $d(x_n, x) < \varepsilon$ for all $n \geq N$. The point is that the new version makes you think about it a little differently, putting the focus on the ε-balls $B_\varepsilon(x)$ or even the neighborhoods of x.

12.10 Accumulation or Cluster Points

There are several ways that a sequence can fail to converge. It could head off to infinity, oscillate, or continually leave and return to the neighborhood of some point. The last two are our focus of interest.

Accumulation Point, Cluster Point. We say that $\mathbf{x}$ is xan *accumulation point* or *cluster point* of $\{\mathbf{x}_n\}_{n=1}^{\infty}$ if for every $\varepsilon > 0$ and every positive integer N, there is $n > N$ with $\mathbf{x}_n \in B_\varepsilon(\mathbf{x})$.

A point $\mathbf{x}$ is an accumulation point of a sequence if the sequence continually returns to any ε-ball about $\mathbf{x}$. It is not required to stay there, but does have to return there.

Both $+1$ and -1 are accumulation points of the sequence $x_n = (-1)^n$, even though the sequence has no limit. The accumulation point doesn't have to be part of the sequence. Consider the sequence defined by

$$x_n = \begin{cases} 1/n & \text{if } n \text{ is odd} \\ 1 + 1/n & \text{if } n \text{ is even} \end{cases}$$

The points 0 and 1 are accumulation points even though they are not part of the sequence.

Can we get a sequence that converges to an accumulation point? Yes! If $\mathbf{x}$ is an accumulation point of $\{\mathbf{x}_n\}$, we can find an $\mathbf{x}_{n_1}$ with $d(\mathbf{x}_{n_1}, \mathbf{x}) < 1$. Then we can find $n_2 > n_1$ with $d(\mathbf{x}_{n_2}, \mathbf{x}) < 1/2$. We continue this with $n_k > n_{k-1}$ and $d(\mathbf{x}_{n_k}, \mathbf{x}) < 1/k$ for every k. If we then set $\mathbf{y}_k = \mathbf{x}_{n_k}$, $\mathbf{y}_k$ is a sequence that converges to $\mathbf{x}$.

12.11 Subsequences

So what have we done? We have constructed a new sequence by plucking out certain elements of the original sequence, and we have done this is a way that guarantees that as we go out the new sequence, we are also moving out in the original sequence. This is an example of a subsequence.

Subsequence. A sequence $\{\mathbf{y}_k\}_{k=1}^{\infty}$ is a *subsequence* of $\{\mathbf{x}_n\}_{n=1}^{N}$ if there is an increasing set of positive integers,

$$n_1 < n_2 < n_3 < \cdots$$

with $\mathbf{y}_k = \mathbf{x}_{n_k}$ for every $k = 1, 2, 3, \dots$.

We may sometimes not bother writing down $\mathbf{y}_k$, but instead just use $\mathbf{x}_{n_k}$ directly as the subsequence. For example, if $x_n = (-1)^n$ defines a sequence, it is not convergent but has two cluster points, -1 and $+1$. The subsequence x_{2k+1} (i.e., $n_k = 2k + 1$) converges to -1 and x_{2k} converges to $+1$. These are not the only convergent subsequences. The subsequence x_{2k^2+1} also converges to -1.

Finally, an alternative way to characterize cluster points is that a point $\mathbf{x}$ is a cluster (accumulation) point of $\{\mathbf{x}_n\}$ if and only if there is a subsequence of $\{\mathbf{x}_n\}$ with limit $\mathbf{x}$.

12.12 Open Sets

A set U is *open* if it contains a ball about each point in the set. Equivalently, U is open if it contains a neighborhood of every point.

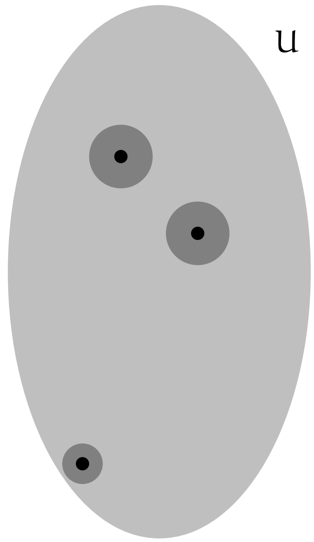

Figure 12.12.1: The ellipse U is an open set. Three points are illustrated, along with darker ε-balls that both contain the points and are themselves contained in the ellipse.

A singleton, $S = \{\mathbf{x}\}$ in $\mathbb{R}^m$ is not open in the usual topology on $\mathbb{R}^m$ (using the Euclidean norm) as any ε-ball about $\mathbf{x}$ will also contain points not in S.

The closed ball of radius 1, $\overline{B_1(\mathbf{x})} = \{\mathbf{x} \in \mathbb{R}^m : \|\mathbf{x}\| \le 1\}$ is not open. The problem is that there no room to fit in a ball about any point on the boundary. For example, if $\mathbf{x} = (1, 0)$, any $B_\varepsilon(1, 0)$ will contain the point $(1 + \varepsilon/2, 0) \notin \overline{B_1(\mathbf{x})}$.

The open interval $(a, b) \subset \mathbb{R}$ is open. If $a < x < b$, let $\varepsilon = \min\{b - x, x - a\}$. Then $B_\varepsilon(x) = (x - \varepsilon, x + \varepsilon) \subset (a, b)$. The choice of ε insures that $B_\varepsilon(x) \subset (a, b)$.

Like open balls, open sets absorb the tails of convergent sequences.

Theorem 12.12.2. *Let U be an open set and $\mathbf{x} \in U$. If $\mathbf{x}_n \to \mathbf{x}$, then there is an N with $\mathbf{x}_n \in U$ for $n \ge N$.*

Proof. Since U is open and $\mathbf{x} \in U$, there is an $\varepsilon > 0$ with $\mathbf{x} \in B_\varepsilon(\mathbf{x}) \subset U$. Because $\mathbf{x}_n \to \mathbf{x}$, there is an N with $\mathbf{x}_n \in B_\varepsilon(\mathbf{x})$ for $n \ge N$. But $B_\varepsilon(\mathbf{x}) \subset U$, so $\mathbf{x}_n \in U$ for $n \ge N$. $\square$

12.13 Open Balls are Open!

Open balls are open sets.

Theorem 12.13.1. *Any open ball* $B_r(x)$*,* $r > 0$*, is an open set.*

Before proving this, it is helpful to draw a diagram.

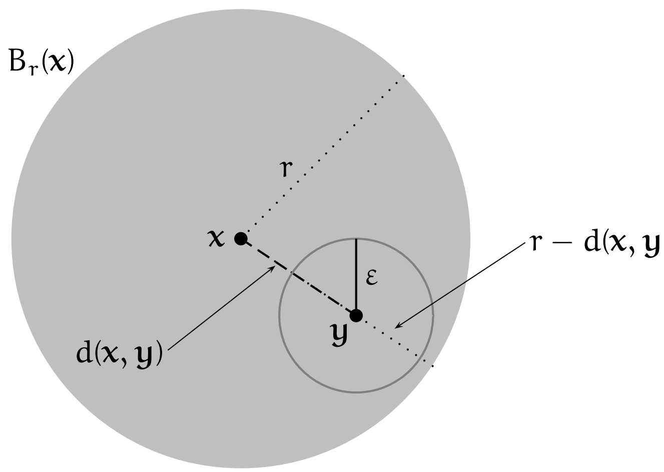

Figure 12.13.2: To show $B_r(x)$ is open, we need to show that for any point $y \in B_r(x)$ we can find an ε-ball about y that fits inside $B_r(x)$, as illustrated here. Since r is the radius of $B_r(x)$ and $d(y, x)$ is the distance between x and y (dashed), the distance from y to the boundary (dotted) is $r - d(y, x)$.

Proof. Let $y \in B_r(x)$. Based on Figure 12.13.2, we should be to show that $B_\varepsilon(y) \subset B_r(x)$ for any $\varepsilon < r - d(y, x)$. So take ε with $r - d(y, x) > \varepsilon > 0$ and let $z \in B_\varepsilon(y)$. We must show $z \in B_r(x)$.

Now

$$\begin{aligned} d(z, x) &\le d(z, y) + d(y, x) && \text{triangle inequality} \\ &< \varepsilon + d(y, x) && z \in B_\varepsilon(x) \\ &< \big(r - d(y, x)\big) + d(y, x) && \text{choice of } \varepsilon \\ &= r. \end{aligned}$$

It follows that any $z \in B_\varepsilon(y)$ is also in $B_r(x)$, establishing that $B_\varepsilon(y) \subset B_r(x)$, completing the proof. □

12.14 Basic Properties of Open Sets

There are three fundamental properties that open sets have. These properties are so important, that in more general works they are part of the definition of open sets. Let $\mathcal{T}$ be the collection of all open sets.

Topology. Let $\mathcal{T}$ be a collection of subsets of X. We say that $\mathcal{T}$ is a *topology on* X if:

1. Any union of sets in $\mathcal{T}$ is in $\mathcal{T}$.
2. Any finite intersection of sets in $\mathcal{T}$ is in $\mathcal{T}$.
3. The empty set and X are in $\mathcal{T}$.

The sets in a topology $\mathcal{T}$ are referred to as *open sets*.

The simplest possible topology on X is $\mathcal{T} = \{\emptyset, X\}$. Any other topology must contain this one. It is called the *trivial topology*, the *indiscrete topology*, or the *coarsest* or *weakest possible topology*.

The largest possible topology on X is to let $\mathcal{T} = \{\text{subsets of } X\}$. Any other topology is a subset of this one. It is called the *discrete topology*, or the *finest* or *strongest possible topology*.

The discrete topology is a metric topology using the discrete metric

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

The open balls in the discrete topology are a bit surprising. They consist of a single point whenever the radius is less than one. Then $B_{1/2}(x) = \{x\}$. It follows that every set S is open because if $x \in S$, $B_{1/2}(x) \subset S$.

12.15 Basic Properties of Open Sets

One important result is that the open sets as defined by ε-balls form a topology.

Theorem 12.15.1. *Let (X, d) be a metric space. The open sets defined using ε-balls form a topology for (X, d). That is:*

1. *Any union of open sets is an open set.*
2. *Any finite intersection of open sets is an open set.*
3. *The empty set and X are open sets.*

Proof. (1) Let $\mathfrak{U}$ be a collection of open sets and

$$V = \bigcup_{U \in \mathfrak{U}} U$$

be the union of sets in $\mathfrak{U}$. If $x \in V$, there is $U \in \mathfrak{U}$ with $x \in U$. Since U is open, there is $\varepsilon > 0$ with $B_\varepsilon(x) \subset U \subset V$, showing that V is open.

(2) Let $U_1, \ldots, U_k$ be open sets and

$$V = \bigcap_{i=1}^{k} U_i.$$

If $x \in V$, then for every $i = 1, \ldots, k$, $x \in U_k$. As each U_i is open, there is $\varepsilon_i > 0$ with $B_{\varepsilon_i}(x) \subset U_i$. Let $\varepsilon = \min_i \varepsilon_i > 0$. Then $B_\varepsilon(x) \subset U_i$ for all $i = 1, \ldots, k$, so $B_\varepsilon(x) \subset B_{\varepsilon_i}(x) \subset V$, showing that V is open.

(3) It is vacuously true that $\emptyset$ is open. What does that mean? It means you can't show $\emptyset$ is not open because that would require it have something in it, which it doesn't. Suppose $\emptyset$ was not open. There there would be $x \in \emptyset$ with $B_\varepsilon(x) \not\subset \emptyset$ for all $\varepsilon > 0$. However, because $\emptyset$ is empty, there can't be any $x \in \emptyset$.

As for X, let $x \in X$. Since X contains all balls around x for every $x \in X$, X is open. $\square$

12.16 The Interior of a Set

Every set has an open set associated with it, the interior of the set. However, the interior may be empty.

Interior. Let S be a set. The *interior of* S, written S^0 or $\operatorname{int} S$, is the union of all open sets contained in S.

As the union of open sets, S^0 is open. In fact, it is the largest open set contained in S.

Theorem 12.16.1. *Let S be a set. Then the interior S^0 is the largest open set contained in S.*

Proof. Let S_1 be the largest open set contained in S. Since S^0 is the union of all open sets contained in S, including S_1, $S^0 \supset S_1$.

But S_1 is the largest open set contained in S, so if $U \subset S$ and U is open, $U \subset S_1$. Otherwise, $U \cup S_1$ would be a larger open set contained in S. Because S_1 contains all open sets in S, it contains their union S^0.

Now both $S^0 \subset S_1$ and $S_1 \subset S^0$, so $S_1 = S^0$. □

12.17 More on Interiors

When a set S is open, $S^0 = S$.

For sets in $\mathbb{R}$, $(a, b]^0 = [a, b]^0 = [a, b)^0 = (a, b)$.

The interior in $\mathbb{R}$ of the rational numbers $\mathbb{Q}$ is empty. To see this note that if $x \in \mathbb{Q}$, $B_\varepsilon(x)$ will always contain irrational numbers of the form $x + 1/n\sqrt{2}$ for n large enough whenever $\varepsilon > 0$. This means that $\mathbb{Q}$ has no open subsets, and so the interior of $\mathbb{Q}$, the union of open sets contained in $\mathbb{Q}$ is empty.

Informally, we can think of the interior of S as S with the boundary removed. Be careful, your intuition can fool you in arbitrary $\mathbb{R}^m$.

It's not a good practice to think about interiors, or even whether a set is open, as being an intrinsic property of the set. The topological properties of a set can change depending on what topological space it is part of.

For example, consider an open interval $(a, b) \subset \mathbb{R}$. In $\mathbb{R}$, $(a, b) = B_r\big((a + b)/2\big)$ where $r = (b - a)/2$. As an open ball it is an open set. Intuitively, an open set is a set that does not include its boundary points. Here the boundary points are a and b, and are not part of the set.

Let's take that same set and embed it in $\mathbb{R}^2$. One way to do it is to define $S = \{(x, y) : a < x < b, y = 0\}$. But this set is not open at all. In fact, $S^0 = \emptyset$. See the Figure 12.17.1.

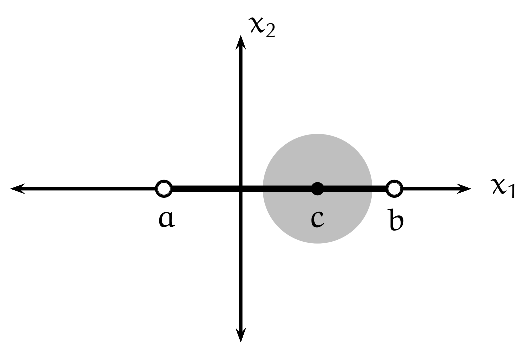

Figure 12.17.1: Here (a, b) has been embedded in $\mathbb{R}^2$ as $\{(x, y) : a < x < b, y = 0\}$. Any point such as $(c, 0)$ is not in the interior as any ball around it will include points above and below the embedded interval.

12.18 Closed Sets

Open sets are not the only way to characterize a topology. We can also use closed sets.

Closed Set. A set S is closed if whenever $\{x_n\}$ is a convergent sequence in S, $\lim_n x_n \in S$.

Any singleton is a closed set.

Theorem 12.18.1. *A singleton $S = \{x\}$ is a closed set.*

Proof. If $\{x_n\}$ is a sequence in S, then $x_n = x$ and $\lim_n x_n = x \in S$. □

Suppose a set S is contained in the space X. The *complement of* S is $S^c = \{x \in X : x \notin S\}$.

Theorem 12.18.2. *Let (X, d) be a metric space. A set $S \subset V$ is closed if and only if $T = S^c$ is open*

Proof. **Only if case:** Suppose S is closed. We must show that $T = S^c$ is open. That is, for all $x \in T$, there is $\varepsilon > 0$ with $B_\varepsilon(x) \subset T$. Take $x \in T$.

If not, then for $\varepsilon > 0$, $B_\varepsilon(x) \not\subset T$. That means that for every n, we can find $x_n \in B_{1/n}(x)$ with $x_n \notin T$. In other words, there are $x_n \to x$ with $x_n \in S$. Since S is closed, $x \in S$, implying $x \notin T$. This contradiction shows that T must be open.

If case: Suppose T is open and $S = T^c$. Let $x_n \in S$ with $x_n \to x$. We must show that $x \in S$. **If not**, $x \in T$. Since T is open, there is $\varepsilon > 0$ with $B_\varepsilon(x) \subset T$. Since $x_n \to x$, there is $N > 0$ with $x_n \in B_\varepsilon(x) \subset T$ for $n \geq N$. **This contradicts** the fact that $x_n \in S$ for all n. Therefore $x \in S$, showing that S is closed. □

Theorem 12.18.2 does **not** say that any set is either open or closed. The whole space is both. Many sets are neither.

Let $a, b \in \mathbb{R}$, $a < b$. In $\mathbb{R}$, the interval (a, b) is open, $[a, b]$ is closed, and both of the half-open, half-closed intervals $(a, b]$ and $[a, b)$ are **neither open nor closed**.

We can find such sets in other spaces. In $\mathbb{R}^2$, consider the punctured disk $S = \{x \in \mathbb{R}^2 : 0 < \|x\| \leq 1\}$. This set is neither open nor closed. It is not closed because $x_n = (1/n, 1/n)$ is a sequence of points in S whose limit $((0, 0))$ is not in S. It is not open because points such as $(1, 0) \in S$, and any ε-ball around $(1, 0)$ will contain points not in S (e.g., $(0, 1 + \varepsilon/2)$).

12.19 Properties of Closed Sets

Just as open sets obey certain properties, so do closed sets.

Theorem 12.19.1. *Let* (X, d) *be a metric space. The closed sets obey:*

1. *Any finite union of closed sets is a closed set.*
2. *Any intersection of closed sets is a closed set.*
3. *The empty set and* X *are closed sets.*

Proof. Take complements and use Theorems 12.18.2 and 12.15.1. The result follows immediately. □

Since singletons are closed sets, any finite union of singletons is closed too, meaning that all sets of the form $\{x_1, \ldots, x_n\}$ are closed. An infinite set of isolated points in $\mathbb{R}^2$ that is closed is the set of grid points,

$$\left\{(x, y) \in \mathbb{R}^2 : \text{both } x \text{ and } y \text{ are integers}\right\}.$$

This is closed because if (x_n, y_n) are grid points converging to (x, y), then if $\|(x_n, y_n) - (x, y)\| < 1$, $x_n = x$ and $y_n = y$. It follows that any convergent sequence of grid points must obey $(x_n, y_n) = (x, y)$ for n large enough, so that (x, y) is also a grid point.

12.20 Example: The Cantor Set

▶ **Example 12.20.1: Cantor Set.** One famous closed set is the Cantor set. Start with the closed interval $C_0 = [0, 1]$ and remove its open middle third, $(1/3, 2/3)$. That leaves the closed set $C_1 = [0, 1/3] \cup [2/3, 1]$. Then remove the open middle thirds of each part of C_1 to get another closed set, C_2. At each set, C_n is closed as a finite union of closed intervals. Each time, we form the next C_{n+1} by removing the open middle thirds of each interval. The *Cantor set* $\mathfrak{C}$ is the intersection of the C_n.

$$\mathfrak{C} = \bigcap_{n=1}^{\infty} C_n.$$

As the intersection of non-empty nested closed sets, the Cantor set is closed.

One way to understand the Cantor set a bit better is to write the $x \in [0, 1]$ in ternary form (base-3). I.e., we interpret the expression $0.d_1 d_2 d_3 \ldots$ as meaning $\sum_{n=1}^{\infty} d_n/3^n$ rather than $\sum_{n=1}^{\infty} d_n/10^n$. Notice that there may be two ways to write such numbers as $1/3 = 0.1000\ldots = 0.02222\ldots$.

Whenever we *must* use a 1 to write the ternary form of a number in $[0, 1]$, we are in one of the middle thirds that is removed to form the Cantor set. That means that the Cantor set consists of all $x \in [0, 1]$ that can be written in ternary form using only 0 and 2. Since 1 is not involved, this expression is unique. Numbers starting with 0.0 are in the first third (including $1/3 = 0.0222\ldots$), while numbers starting 0.2 are in the last third (including $2/3 = 0.2$), and so forth.

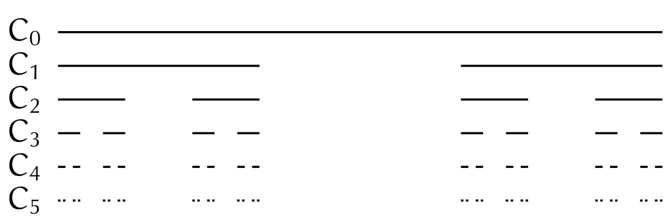

Figure 12.20.2: Here are the $C_0, \ldots, C_5$, the first six sets defined by eliminating middle thirds, and whose intersection comprises the Cantor set $\mathfrak{C}$.

◀

12.21 Closures

Every set has a closed set associated with it, the closure of the set.

Closure. Let S be a set. The *closure of* S, written $\overline{S}$ or $\operatorname{cl} S$, is the smallest closed set containing S. Equivalently, it is the intersection of all closed sets containing S.

As with open sets, we could prove the equivalence of the two definitions of closure, but the arguments are pretty similar, so we won't. I do suggest you give it a try.

The closure always includes S itself.

Lemma 12.21.1. *If $x \in S$, then $x \in \overline{S}$.*

Proof. Consider the sequence $x_n = x$. This obviously converges to x, so $x \in \overline{S}$. □

Besides S itself, $\overline{S}$ may include some points that arise as limits.

Limit Point. Let S be a set. A point x is a *limit point of* S if for all $\varepsilon > 0$, $B_\varepsilon(x)$ contains some point of S other than x.

Equivalently, a point x is a limit point of S if there is a sequence of points in S other than x that converge to x.

The closure of S consists of S together with its limit points.

Theorem 12.21.2. *The closure of S is the union of S and the set of limit points of S.*

Proof. Suppose $T \supset S$ is closed. By definition of closed, every limit point of S is in T. It follows that the closure of S contains all limit points of S as well as S itself.

Suppose that $x \notin S$ is not a limit point of S. Then there is an $\varepsilon > 0$ so that $B_\varepsilon(x)$ contains no points of S. Then $B_\varepsilon(x)^c$ is a closed set containing S but not containing x. It follows that $x \notin \overline{S}$. Thus $\overline{S}$ contains nothing other than S and its limit points. □

► **Example 12.21.3: Closure of the Open Ball.** Let's find the closure of the open unit ball in a normed vector space $\big(V, \|\cdot\|\big)$. Set $S = B_1(y) = \{x \in V : \|x - y\| < 1\}$. First, if $x_n \in S$ converges to x, $\|x_n - y\| < 1$. By Theorem , $\|x - y\| \le 1$. The only question is whether all of the points with $\|x - y\| = 1$ are in the closure of S.

For x with $\|x - y\| = 1$, consider $x_n = y + (1 - 1/n)(x - y)$. Then $\|x_n - y\| = (1 - 1/n)\|x - y\| < 1$, so $x_n \in S$. Also $x_n \to x$, so x is a limit point of S.

In general,

$$\overline{B_\varepsilon(x)} = \{y \in X : \underset{\cdot}{(}x, y) \le \varepsilon\}.$$

◄

12.22 Boundary of a Set

Boundary. The *boundary of a set* S is the intersection of the closure of S and the closure of its complement. The boundary is written ∂S or bdy S.

The boundary of the open ball $S = B_\varepsilon(\mathbf{y})$ is $\partial S = \{\mathbf{x} : \|\mathbf{x} - \mathbf{y}\| = \varepsilon\}$. We saw earlier that the closure of the ball includes all points with $\|\mathbf{x} - \mathbf{y}\| \le \varepsilon$. The complement of the ball is closed, and consists of all points with $\|\mathbf{x} - \mathbf{y}\| \ge \varepsilon$. Taking the intersection yields ∂S.

If $S = (a, b)$, an interval in $\mathbb{R}$, $\partial S = \{a, b\}$. To see this, note that the closure is $\bar{S} = [a, b]$ and the complement, which is closed, is $S^c = (-\infty, a] \cup [b, +\infty)$. Their intersection is $\partial S = \{a, b\}$.

The boundary is not exactly the edge of a set, although clear edges are generally part of the boundary. Boundaries can sometimes be a little surprising. Consider the boundary of the set of rational numbers, $\mathbb{Q}$. Every irrational number can be written as a limit of rational numbers, so $\overline{\mathbb{Q}} = \mathbb{R}$. Every rational number can be written as the limit of irrational numbers (e.g., $x_n = x + (\pi/n)$), so $\overline{\mathbb{Q}^c} = \mathbb{R}$. Then $\partial\mathbb{Q} = \mathbb{R}$.

Let's take another case where it might not be obvious what the boundary is. In $\mathbb{R}$, consider $S = \{1, 1/2, 1/4, \ldots\}$. The boundary is $S \cup \{0\}$, which is also the closure. Here the closure of the complement is $\mathbb{R}$ itself.

► **Example 12.22.1: Closures and Interiors.** Informally, we think of the closure as adding the boundary and the interior removing it. However, that is sometimes misleading. You may not always be adding and removing the same boundary. Consider the set $S = (0, 1) \cup (1, 2)$. It has boundary $\{0, 1, 2\}$. The closure adds $[0, 2]$, but if we take the interior, we don't get the original set. Instead, we get the interval $(0, 2)$. We also see this in the complement $S^c = (-\infty, 0] \cup \{1\} \cup [2, \infty)$, which has interior $(-\infty, 0) \cup (2, \infty)$. Closing that gives us $(-\infty, 0] \cup [2, \infty)$ which is not S^c. ◄

12.23 Closures and Complements

In 1922, Kuratowski showed that you can potentially form up to 14 different sets by repeatedly taking the closure and complement of a set S. This can even be done in the real line with the usual topology.

It's easy to get four. Let $S = (0, 1)$. Then $\overline{S} = [0, 1]$, $\overline{S}^c = (-\infty, 0) \cup (1, +\infty)$, and $\overline{(\overline{S}^c)} = (-\infty, 0] \cup [1, +\infty)$. Taking the complement again brings us back to S. Starting with a half-open finite interval yields six sets.

As for 14, one example is

$$S = \{0\} \cup (1, 2) \cup (2, 3) \cup \big(\mathbb{Q} \cap (4, 5)\big).$$

The seven sets starting with the complement are

$$\begin{aligned}
S^c &= (-\infty, 0) \cup (0, 1] \cup \{2\} \cup [3, 4] \cup \big(\mathbb{Q}^c \cap (4, 5)\big) \cup [5, +\infty) \\
\overline{S^c} &= (-\infty, 1] \cup \{2\} \cup [3, +\infty) \\
(\overline{S^c})^c &= (1, 2) \cup (2, 3) \\
\overline{(\overline{S^c})^c} &= [1, 3] \\
\left(\overline{(\overline{S^c})^c}\right)^c &= (-\infty, 1) \cup (3, +\infty) \\
\overline{\left(\overline{(\overline{S^c})^c}\right)^c} &= (-\infty, 1] \cup [3, +\infty) \\
\left(\overline{\left(\overline{(\overline{S^c})^c}\right)^c}\right)^c &= (1, 3).
\end{aligned}$$

There are six more starting from the closure:

$$\begin{aligned}
\overline{S} &= \{0\} \cup [1, 3] \cup [4, 5] \\
\overline{S}^c &= (-\infty, 0) \cup (0, 1) \cup (3, 4) \cup (5, \infty) \\
\overline{\overline{S}^c} &= (-\infty, 1] \cup [3, 4] \cup [5, \infty) \\
(\overline{\overline{S}^c})^c &= (1, 3) \cup (4, 5) \\
\overline{(\overline{\overline{S}^c})^c} &= [1, 3] \cup [4, 5] \\
(\overline{(\overline{\overline{S}^c})^c})^c &= (-\infty, 1) \cup (3, 4) \cup (5, +\infty)
\end{aligned}$$

That makes 14 in all, including S itself.

12.24 Nesting Balls

Now that we know something about topologies, let's go back and explore the relation between a metric and its topology a little further. We know how to get the topology from a metric by using open balls. But can different metrics generate the same topology?

In fact, they can. We prove in Theorem 12.27.1 that all of the ℓ_p^n norms generate the same topology on $\mathbb{R}^m$.[2] The point is that you can nest balls for the various norms within one another, both up and down. So any ℓ_p ball contains an ℓ_q ball about each of its points, and vice-versa. That means that ℓ_p^n open sets are ℓ_q^n open, and vice-versa for $1 \le p, q \le \infty$.

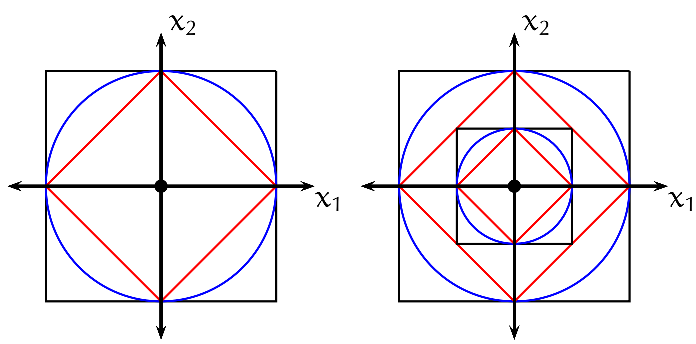

Figure 12.24.1: The left panel shows balls of radius one in ℓ_∞ (black), ℓ_2 (blue), and ℓ_1 (red) norms. The right panel shows how smaller balls of the same type nest inside them. Here the smaller balls have half the radius of the larger balls.

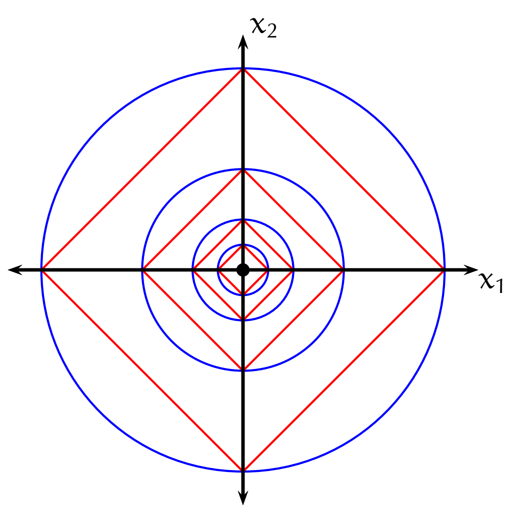

Figure 12.24.2: Only the ℓ_2 (blue) and ℓ_1 (red) balls are illustrated, but at four different sizes, differing by factors of 2, See how they successively nest within one another. Each ℓ_2 ball has a smaller ℓ_1 ball inside it, and vice-versa

[2] See also section 29.4 of Simon and Blume.

12.25 Equivalent Norms

So let's look at the details about how different metrics can generate the same topology.

Equivalent Norms. On a vector space V, norms $\|\cdot\|_1$ and $\|\cdot\|_2$ are *norm equivalent* or *equivalent norms* if there are positive numbers a and b with

$$a\|x\|_2 \leq \|x\|_1 \leq b\|x\|_2.$$

Of course, $a \leq b$. You may be bothered by the asymmetric treatment of the two norms. If so, good. We need to show that norm equivalence is an equivalence relation, putting each norm on an equal footing.

Theorem 12.25.1. *For any vector space V, norm equivalence is an equivalence relation on the set of norms on V.*

Proof. We need to show that norm equivalence is reflexive, symmetric, and transitive.

(1) The norm $\|\cdot\|_1$ is equivalent to itself. Just set $a = b = 1$ in the definition. So norm equivalence is reflexive.

(2) Suppose $\|\cdot\|_1$ is equivalent to $\|\cdot\|_2$. Then there are $a, b > 0$ with $a\|x\|_2 \leq \|x\|_1 \leq b\|x\|_2$. We take this apart and rewrite each piece separately.

$$a\|x\|_2 \leq \|x\|_1 \quad \text{so} \quad \|x\|_2 \leq \frac{1}{a}\|x\|_1$$

and

$$\|x\|_1 \leq b\|x\|_2 \quad \text{so} \quad \frac{1}{b}\|x\|_1 \leq \|x\|_2.$$

Then put them back together to obtain

$$\frac{1}{b}\|x\|_1 \leq \|x\|_2 \leq \frac{1}{a}\|x\|_1,$$

showing that $\|\cdot\|_2$ is equivalent to $\|\cdot\|_1$. This means norm equivalence is symmetric.

(3) Now suppose $\|\cdot\|_1$ is equivalent to $\|\cdot\|_2$ and $\|\cdot\|_2$ is equivalent to $\|\cdot\|_3$. Then there are $a, b, c, d > 0$ with

$$a\|x\|_2 \leq \|x\|_1 \leq b\|x\|_2$$

and

$$c\|x\|_3 \leq \|x\|_2 \leq d\|x\|_3$$

Again consider both halves of the each equation separately, then multiply, and reassemble, to obtain

$$ac\|x\|_3 \leq a\|x\|_2 \leq \|x\|_1 \leq b\|x\|_2 \leq bd\|x\|_3$$

so

$$ac\|x\|_3 \leq \|x\|_1 \leq bd\|x\|_3$$

This shows $\|\cdot\|_1$ is equivalent to $\|\cdot\|_3$, that norm equivalence is transitive, and completes the proof. □

12.26 The Nesting Ball Property

Equivalent norms have a nesting ball property as was illustrated in Figures 12.24.1 and 12.24.2.

Lemma 12.26.1. *If, for all $x \in V$,*

$$a\|x\|_2 \leq \|x\|_1 \leq b\|x\|_2.$$

Then

$$B^2_{\varepsilon/b}(x) \subset B^1_\varepsilon(x) \subset B^2_{\varepsilon/a}(x)$$

where $B^i_\varepsilon(x)$ is the ε-ball about x computed using the i-norm.

Proof. Suppose $x \in B^1_\varepsilon(x)$. Then $\|x\|_1 \leq \varepsilon$, so $a\|x\|_2 \leq \varepsilon$, meaning that $\|x\|_2 \leq \varepsilon/a$. It follows $B^1_\varepsilon(x) \subset B^2_{\varepsilon/a}(x)$.

Now suppose $x \in B^2_{\varepsilon/b}(x)$. Then $\|x\|_2 \leq \varepsilon/b$, so $\|x\|_1 \leq b(\varepsilon/b) = \varepsilon$. It follows that $B^2_{\varepsilon/b}(x) \subset B^1_\varepsilon(x)$. $\square$

12.27 Equivalent Norms, Same Topology

We can use the nesting ball property to show that equivalent norms generate the same topology, the same collection of open sets.

Theorem 12.27.1. *Suppose two norms are equivalent on a vector space V. Then a set is open under norm $\|\cdot\|_1$ if and only if it is open under $\|\cdot\|_2$.*

Proof. Suppose U is 1-open. Then for every $\mathbf{x} \in U$, there is an ε_x with $B^1_{\varepsilon_x}(\mathbf{x}) \subset U$. But then

$$B^2_{\varepsilon_x/b}(\mathbf{x}) \subset B^1_{\varepsilon_x}(\mathbf{x}) \subset U$$

by Lemma 12.26.1. That means U is 2-open because there is a 2-ball about each $x \in U$ that is contained in U.

Now suppose U is 2-open. Then for every $\mathbf{x} \in U$, there is an ε_x with $B^2_{\varepsilon_x}(\mathbf{x}) \subset U$. But then

$$B^1_{a\varepsilon_x}(\mathbf{x}) \subset B^2_{\varepsilon_x}(\mathbf{x}) \subset U$$

by Lemma 12.26.1. That means U is 1-open because there is a 1-ball about each $x \in U$ that is contained in U. □

We can also show that any of the ℓ_p norms are equivalent on $\mathbb{R}^m$.

Theorem 12.27.2. *For any p, $1 \le p < \infty$, the ℓ_p and ℓ_∞ norms are equivalent.*

Proof. First,

$$\|\mathbf{x}\|_p = \left(\sum_{i=1}^m |x_i|^p\right)^{1/p} \le \left(\sum_i \|\mathbf{x}\|_\infty^p\right)^{1/p} = m^{1/p}\|\mathbf{x}\|_\infty.$$

Second, $\|\mathbf{x}\|_\infty^p \le \sum_{i=1}^m |x_i|^p$, so $\|\mathbf{x}\|_\infty^p \le \|\mathbf{x}\|_p^p$, implying $\|\mathbf{x}\|_\infty \le \|\mathbf{x}\|_p$.

Combining the results shows they are equivalent. □

Combining the last two theorems shows that all of the ℓ_p norms define the same open sets.

Corollary 12.27.3. *For a given m, the open sets in $\mathbb{R}^m$ are the same for every ℓ_p norm, $1 \le p \le \infty$. The open sets are also the same in the product topology.*

Proof. For the first part, combine Theorems 12.27.2 and 12.27.1. Theorem 12.7.1 shows that a sequence product converges if and only if it converges in ℓ_2. This implies that both topologies have the same closed sets, and by taking complements, the same open sets. □

October 11, 2020