[SQUEAKING]

[RUSTLING]

[CLICKING]

PAIGE BRIGHT: So far this class, we've talked about a lot of examples which were over and over again the same three concepts. And then we talked about the general theory, which is new because it's on metric spaces, but it's very similar to the stuff we've already seen for Euclidean space, right? Compact sets will be slightly different.

It's not that compact sets don't exist in Euclidean space. It's that they're pretty nice there. So I hope to highlight why it's so nice on Euclidean space, but also highlight the definitions that will become more difficult to manage on metric spaces.

All right. So what is so amazing about compact sets? Well, if you look up what's the intuition behind compact sets, the first thing you'll see on Stack Exchange is a conversation about finiteness.

Now, we have finiteness in two senses-- one in terms of integrals. One is an integral finite. And secondly, we have it in terms of finite sets.

We can consider what statements can I make about finite sets? These two notions drive the understanding of compact sets. And so today, my goal is to start off with motivation using norms to describe finiteness, and then I'm going to talk about finite sets and the analysis we could do on those.

OK, so let's start with norms. So to define a norm, which you might have heard about before, either in this class or in 1802, we're going to need the notion of a vector space. And so let's define what that is now.

A vector space is a set V with addition which takes in 2 vectors-- so V cross V-- and maps it to another vector. So when I add 2 vectors, I get another vector. And secondly, multiplication-- if I multiply by, let's say, a real number on a vector, I should get another vector back out.

So this is very fancy notation, but essentially, what this is saying, again, is that if I have 2 vectors, I can add them together and get another. And if I multiply it by a real number, I'll get another vector. And specifically, a vector space is a set V with these 2 operations on it with nice properties.

So those nice properties that we want it to define or want it to have are as follows. I'll just state a few of them just to point out why they're so nice. But essentially, it's the field axioms.

We want this vector space to be one, commutative. So the order in which I add two things doesn't matter. "-tative." Two, associative, meaning that A plus B plus C is A plus B plus C.

We want there to be a distributive law. "-butative." And I'll just state one more. We have the identity, by which I mean 0 and 1-- or sorry. I should say 0 is in your vector space, and minus V is in your vector space, so I can always get back to 0.

There's a few more axioms than this to define a vector space. This is mostly covered in 1806. But I'll just state a few of them right now.

Now, the reason that vector spaces are so helpful is because of this notion of addition, right? As we've talked about before, for metric spaces, there's not too much more we can do. We can't state facts about how you add two elements in a metric space because it doesn't always make sense. We don't have addition always defined.

So just to recap, today, we're talking about compact sets. And I just somewhat rigorously defined what a vector space is. So didn't miss too much.

All right, so now that we have this definition of a vector space, we're now going to define what a norm is. A norm. A norm is essentially going to be like a metric on your vector space. It's a map that's denoted like this-- which would make a bit more sense in a moment-- with map taking in vectors in a vector space V.

It's spitting out a positive number. And in particular, we want this norm to satisfy three properties, just like we did for metric spaces. One, of course, we want positive definiteness, where here we want V bigger than or equal to 0 for all V. And the norm of V is equal to 0 if and only if V equals 0. So that's the definition of positive definiteness.

Now, why didn't this work on metric spaces? Because there we didn't have a notion of what it meant for the number 0 to be in your vector space. But here we do, right?

In the vector space, we're assuming that we have an element 0 in our vector space. So we're all good there. Positive defniteness is essentially the same.

Secondly, we have homogeneity, which essentially just states that we can pull out constants. So the norm of lambda V where lambda is just a constant and V as a vector is equal to the absolute value of lambda times V, times the norm of V. And finally, we have the triangle inequalities. But here we don't need to have three vectors because this one's just defined in terms of 1.

So all we need is that V plus W is less than or equal to the norm of V plus the norm of W. This is it. These are the three properties that define what a norm space is.

And in fact, you can see right here-- the setup of this is very, very similar to metric spaces, the only difference being that we now know what it means to multiply by constants. We now know what it means for a vector to be 0. And we have addition defined on our space. OK?

Now, let's look at some example of norms. And the examples that we're going to look at are, in fact, actually very similar to metrics themselves. So example one-- suppose we're looking at the set of continuous functions on the integral a to b-- or actually, I'll just do 0 to 1 for now so that it's simpler.

Then consider the norm acting on this z0 0 to 1, two positive or non-negative numbers with the norm of a function f being equal to the supremum of all x in 0 to 1 of absolute value f of x. I claim that this is, in fact, norm. So let's show the three properties.

Proof-- well, the first one-- positive definiteness-- this one's mostly done. We know that it's positive or non-negative. More accurately, it's non-negative, but let's look at the case where it is equal to 0. So if the norm of f equals 0, which again, under the definition, is the supremum of x and 0 to 1 of absolute value f of x, this implies that the function must be 0 everywhere, right? Assume it wasn't 0 at one point.

Then the supremum would be that maximum value. So this implies that f must be equal to-- this is true if and only if f equals 0 everywhere. So we're good on that front.

Secondly, homogeneity-- this one isn't that bad as well because then we have the norm of lambda f being equal to the supremum of lambda f of x. And of course, we can pull out this lambda now because it's just a constant times the function. So this is, in fact, equal to the supremum of lambda times f of x. And then we can pull out lambda. So then this would be lambda times the supremum of f of x, which is absolute value lambda times norm of f.

So all go down the parts with homogeneity. And I won't prove it again. Actually, I'll just state it because why not? Because we're going to look at one more example after this anyways.

So for the triangle inequality, here we have that f of x plus g of x-- or I should say minus, right? No, g, plus. That's fine.

This is less than or equal to absolute value f of x plus absolute value g of x. And then on the right-hand side, I can apply the supremum, right? So applying the supremum on the right-hand side of both f and g will get that it's less than or equal to the supremum of f of x plus the supremum g of x, which is, of course, the norm of f plus the norm of g.

And then I can take the supremum of the left-hand side to complete the proof, right? So this is what just having two extreme values here in the proof-- here all we used was bounds using the supremum which is how we could have done the other proof This is just a slightly more elegant way to do that.

All right. So that's it on this example. And this is, again, essentially what we looked at last time when we were looking at the metric except there we were looking at the distances between two functions itself. Here we can just look at the norm of a single function.

All right. Let's look at one more example, which would be akin to what we were looking at before. So I'm going to define an L1 norm. This notation will come up a bit later in lecture 5, but here we have the norm acting on c0 again-- so continuous functions on 0 to 1 and spitting out a non-negative number, where here specifically, the norm of a function f-- Let's say L1-- is equal to the integral from 0 to 1 of absolute value f of x dx.

And now, let's just talk through why this isn't back to norm. One-- positive definiteness, we've already talked about in regards to metrics, right? This is definitely going to be positive or non-negative because it's absolute values on the inside here, and we're integrating over 0 to 1.

But the fact that-- its positive definiteness comes from continuity, right? Because if at any point this is bigger than 0, there must be a point in which f of x is not 0. Homogeneity follows from homogeneity of the integral, right? Because we can pull out constants if we multiply them by f.

And finally, the triangle inequality is precisely the same as before because we can apply the triangle inequality to, let's say, f plus g, and then separate out using linearity of the integral. So again, three statements are very analogous. So you might be wondering, why do we care, right? Why are we looking at norm sets or norm spaces?

Oh, I should say if you have a vector space with the norm on it, then it's called a norm space. I'll write that down, in fact. Definition-- a vector space v with norm absolute values is called a norm space. Space-- just like a metric space.

So why do we even care? Well, firstly, 18102 talks about this concept quite a bit further, as you can note. We're doing all of this on vector spaces.

And what this lets us do, essentially, is linear algebra. Once we start studying norms, we can look at linear algebra on vector spaces, and in fact, look at infinite dimensional vector spaces. That's the major difference.

If that terminology does not make sense because y'all haven't done a linear algebra class, necessarily, totally fair. Don't worry about it. It's just what happens in 18102, which is a class I highly recommend. But secondly, the process of figuring out that a norm is, in fact, a norm is very similar to showing that a metric is, in fact, a metric.

It's just checking all three properties and going through the details. So that's the second reason why I bring it up here. But the third reason in particular with compact sets is it's going to help us understand this notion of finiteness.

So for instance, one question we could ask-- question-- when is the L1 or mono function finite? Well, here it's pretty straightforward, right? Because here we're only looking at functions on 0 to 1.

So if it's continuous on 0 to 1, we know that it has the maximum value. And we know that this term is always going to be finite. But I guess I should say more rigorously, this is clear if we're on 0 to 1. But what about when we're looking at functions on R?

Well, here we can similarly define L1 functions on R to be the set of functions-- this is loosely-- functions f such that the integral from minus infinity to infinity of f of x dx is less than infinity. So then this turns into a question of when are these integrals, in fact, finite? And there's a few answers to this, but the one that I want to point out is perhaps the most natural.

The L1 norm of a function f on R will be finite if f is 0 outside of some interval. So let's just say the interval is minus n to n. Right? If I have a function that is oscillating wildly in between two points, that doesn't matter as long as at a certain point-- let's say from n to minus n-- it goes to 0, right?

Then we're going to know that the integral is, in fact, finite. This is not an if and only if statement, of course. We could have functions like a Gaussian. So note-- this treatment is not if and only if. We can have functions like e to the minus x squared, where here it never quite goes to 0, but rapidly decays to 0.

So note that the L1 functions are not simply ones that eventually go to 0, but this is a pretty good answer to our question, right? This is a big set of functions. I can just take functions that are continuous and then eventually go to 0. That's a pretty big set.

To be more explicit here, this is how it's defined on R, but we can be even more general than functions on R. So if a function satisfies this, we say that it has compact support. But what is the support of a function?

Well, definition-- consider some f a function on R. And then consider the set where it's non-0. So the x such that f of x is non-0. One question I could ask you before I get into what the actual definition is that we're looking at is, is this set open, or is it closed? Any ideas?

Yeah, it's open. We can write this equivalently as f inverse of the complement of 0, right? We're looking at the points such that the image of them is non-0. And we know that the singular set 0 is a closed set, which means that the complement is open and the inverse image of an open set is open.

So this is open. And then we define the support of a function to be the closure of the set. So I should say x such that f of x is non-0-- we define the support of the function to be the closure of the set. Now, I haven't defined what closure is. What closure is is the smallest closed set that contains it.

All this is doing is adding the boundary points, right? So if I have a function-- here's a little aside. If I have a function that's non-0 from minus a to a-- I can just draw out a little picture, a function that acts like this for minus a to a. Then the support will be adding the boundary points.

So this is a little confusing to anyone who's just watching the board right now, but verbally, that's what's happening. We're just adding the boundary points of our intervals. And this is a closed set.

So when I say that a function has compact support, I mean that the support of the function is a compact set. I haven't defined what a complex set is. I'm going to do that in just a moment, but this is how it's connected to the notion of finite nodes, right?

Again, the way that we got here was asking a very simple question. If I have a function on the real numbers, when is its integral going to be finite? And a nice answer is if it eventually goes to 0.

And that is somehow connected to the notion of a compact set. So any questions thus far? I know I'm being very hypothetical here. I know that I'm talking about a definition without saying what it is quite yet, but this is the logic that we're following. Any questions?

Cool. All right. So let's look at some analysis on finite sets. So this was the first part on norms. Now, we're looking at finite sets.

So let A be a subset of a metric space x with metric d. And suppose that A is finite. Then I claim that we have three properties of A.

Then one-- let me switch-- every convergent sequence-- actually, I'll just say every sequence of points of A-- so let's say xi-- contained in A has a convergent subsequence. And I'm going to prove these things, right? If you're looking at this and wondering how are these statements true, I'm going to quickly prove them, but they aren't too hard to show.

Two, A is closed and bounded. And three, if f is a function on A-- so f is a function that maps A to R-- then f has a maximum and a minimum. These are three properties of finite sets.

And let's actually prove these three properties, right? It's pretty nice to show them. Three, OK.

So let's prove these three properties. Proof-- consider some sequence xi in A. How can we find a convergent subsequence then? Well, because there's only finitely many points we can consider, and a sequence has infinitely many terms, we know that one of those terms must be repeated infinitely many times, right?

Well, let's just say there exists an xj in A with xj infinitely many times inner sequence. And because it appears infinitely many times, this is going to imply that the sequence of xnk's, which are simply the points in our sequence which are equal to xj-- this is a convergent subsequence. The fact that it's convergent is simply stating that if I have a sequence that's the same point over and over and over again, it's going to converge to that point.

Right? So again, to reiterate-- here we had a finite set. And we had a sequence on our finite set. And we wanted to say that it has a convergent subsequence.

We know that this has to be true because there is at least one term in our set A that's repeated infinitely many times in our sequence. And then we can just choose the subsequence to be that point over and over and over again. So that's our convergent subsequence.

Two-- how do we know that it's closed and bounded? Well, closure pulls from last time because we showed that if I'm looking at a single point, we know that that is a closed set. And we know that a union of closed sets, a finite union of closed sets, is closed.

So we know that it's closed. The fact that it's bounded comes from the fact that there's only finitely many points. So what we can consider is fix any point p in your metric space x-- so any point-- and consider our bound b to be the maximum of the distances from p to xi where xi is in set A.

And here we know that this maximum actually exists, right? Because it's a finite set. So if I'm looking at the maximum of a finite thing, this bound would be finite. So this implies that it's closed and it's bounded, right?

We were able to find an upper bound on our set. Cool. And lastly, three-- the fact that a function on A to R has a max and min follows from the fact that we can just consider the set f of A. This will be a finite set in R.

And because it's finite, I can order them, right? So I can just choose the point where the maximum is achieved and choose the point where the minimum is achieved. The thing that's important to note is that the maximum and the minimum is achieved at a single-- or at least a point of your set A. But yeah, these are the three properties of finite sets that we can do analysis on.

And notice here that we're not assuming that the sequence is convergent. We're not assuming that the function f is continuous. Everything is super nice when we're looking at finite sets.

Now, these three properties are ones that we're, in fact, going to show are true about compact sets. Either we're going to show that they're true or we're simply going to define them to be so. Let's go ahead and do that.

So this is for the motivation about compact sets. I'm going to define what a compact set is now and prove some nice facts about it. But before I do so, again, any questions about what I've done so far? A lot of motivation before we've actually defined what we're looking at.

All right. So I'll go ahead and erase the boards. And I'll move on to the definitions of compact sets. For the definition of compact sets, we're going to have two of them.

I'll explain why we have two of them in a moment. But really, we're going to show that on metric spaces, they're the same. But before I define what a compact set is, I just have to say what a cover is. So a cover of a set A is simply a union of sets-- let's say those sets are called Ui-- such that A is contained in the union. So I should say a collection of sets rather than e again. It's a collection of sets such that A is contained in their union, all right?

This is the definition of a cover. And an open cover is the same exact thing, except then I'm going to assume that my Ui's are open. I'm messing with the mic too much. So same definition as a cover. I'm just assuming that the sets that are covering my set A are open.

All right. So let's define what a compact is. And this will become apparent when I find this first. Definition-- a set A contained in a metric space x is compact of-- because we're going to have two definitions, we're going to be careful here. I'm going to say it's topologically compact if every open cover-- so if every open cover of A has a finite subcover.

Now, what do I mean by finite subcover? I mean that I can just choose finitely many of these Ui's to still cover our set, all right? And notice that this only depends on open sets, which is why it's called the topological compact, right? This definition of compactness only requires that we know what open sets are, which is why we can do them in topology 18901.

All right. Our second definition is going to be one that has a much closer connection to finite sets-- at least apparently-- or it'll be more apparent. So suppose again that we have a set A in x. And this will be sequentially compact if every sequence of A has a convergent subsequence.

And now, this should become apparent-- how it directly relates to finite sets, right? Because there we had sequences. And we knew that every single sequence had a convergent subsequence, right?

So these are our two definitions of compactness. They are not equivalent always. In fact, in topological spaces, these are very different.

But on metric spaces, our ultimate goal for this class, it's to show that they're the same. But before we get to that on metric spaces, let's just talk about this on the real numbers because on the real numbers, there's quite a bit we can actually say. I just want to note one more notational thing.

As opposed to running compact sets every single time, I'm just going to use A is a subset of x if it's compact, where here I have two subsets inside of each other. And if you're using LaTeX to type up your p sets-- this is just slash capital S subset. So no worries there.

But that's the limitation for compact subsets that I'll be using. And it's, in fact, the common one to use. OK.

So what sets do we know that are sequentially compact from the real numbers? You might have an answer for this already in your head, especially if you've already read the lecture notes, but we'll go through them. An immediate example that we can consider, which is pretty nice, is simply the set of real numbers themselves.

So let's do that. So example-- consider the entire set of real numbers contained in R. Any guesses as to whether this will be compact or not?

Yeah, it will not be a compact set. It's not going to be compact in either sense of the terms. Proof-- the proof of sequential compactness is not that bad, right? I can just choose the sequence that I'm considering-- xn-- to simply be equal to n for every natural number.

Firstly, this is not convergent, which is how we should expect, right? If the sequence that I'm considering is convergent, then every subsequence will be convergent. So start off with a divergent set or a divergent sequence.

But furthermore, any subsequence of this I choose is going to go off to infinity, right? So no subsequence of this will converge. So this shows not sequentially compact. But furthermore, how do we show that the entire set of real numbers is not compact topologically?

Just choose an open cover that's nice. So consider the open cover given by the union of open intervals minus n to n or n in the natural numbers. So here I'm just considering nested open intervals, right? 1 instead of the other instead of the other.

We know that this is going to contain the real numbers. Any real number that you choose is going to be in one of these open intervals, right? Just choose one large enough.

But does there exists a finite subcover? The answer is no. Suppose there existed a subsequence nk such that the union of nk in the natural numbers-- or I should say the union of k equal to 1 to some m of minus nk to nk-- so this is our finite subcover. Suppose that this contains the real numbers.

This cannot be the case, right? Choose the largest possible nk, right? If I choose the largest possible nk, then I know that there exists a real number larger than that.

And so then I know that the real numbers are not contained in this finite subcover. And this is true for every possible finite subcover. Realistically, yeah, so in statements like this, you need to show it's false for every subcover of the open cover you chose to choose. OK?

But this shows both not compact and not sequentially compact. Any questions here? These examples are very important to understand, so I'm happy to reiterate any points about them. Cool.

Second example-- what about instead of considering all the real numbers, we consider just nice intervals? Let's consider one that's half open and half closed. So 0 to 1, but not exceeding 0-- this is also not compact.

And we can get this, right? Because I'm claiming that topological compactness and sequential compactness are going to be false for metric spaces. And we know that the set is not sequentially compact.

How do we know that? Because 0 is not contained. So just consider this xn to be 1 over n for n in the natural numbers. This is a sequence that converges in R, right?

What's it going to converge to? 0. So xn-- so the sequence does converge to 0. It doesn't converge in our subset, but it does converge.

What this tells us is that every possible subsequence xnk also contributes to 0. But this is not in our open interval 0 to 1, which implies not sequentially impact. Right?

How do we show it's not topologically compact? Well, we just choose an open cover that works out in our favor again. Consider here is the one I'm going to consider is the open cover 5 1 over n to 2.

Why do I go to 2? So that I know that it contains 1, right? If I'm going to cover an open interval, it has to contain 1. So I'm going to go slightly past it.

In fact, you could choose this 2 to be 1.5 if you really wanted do from n equal to 1 to infinity. I know that this contains my open interval 0 to 1. How do I know this is true? Choose any point. We know that there exists an n large enough such that that point is in one of these open intervals.

But again, if you choose any finite subcover of this, there's going to exist a number between 0 and that largest nk that doesn't exist in our finite subcover. So I can write this out again, but it's the same as this one, right? Here the issue was any finite subcover I have, I can choose the real number larger than that maximum number.

Here any finite subcover I have will have a finite lower bound. And I can find a real number smaller than that. So this implies not topologically compact, which is all great, right?

These are all things that we want to be true about our definitions because if they weren't true, then they wouldn't be equivalent, which is what we're hoping to show. Do I want to use this one? I'll use this one.

So now, we're going to look at an example that is compact. But notice-- that's going to be a little bit harder to show, at least on the face of it, because here for these examples, we only had to come up with examples of sequences and of open covers that fail to be sequentially and topologically compact. So when we're looking at an actual compact set, we're going to have to consider every open cover and every sequence, which would be a little bit harder.

So example-- here I'm just going to add in the point 0. It'd be 0 to 1. This is a compact subset of the real numbers.

How do we actually show this? Well, I'm going to show topologically compact, but how do we show sequentially compact? Sequentially compact, I'm not actually going to show right now, but I'll make a note of it.

Consider the sequence xn, which is contained in 0 to 1. How do I know that this has convergent subsequences? This is a fact from 18100A, but does anyone remember the name of it?

No worries. It's one of the ones that's not used too often. But it's Bolzano-Weierstrass theorem. The Bolzano-Weierstrass theorem tells us-- so Bolzano-Weierstrass-- tells us there exists a convergence subsequence, where here the only assumptions Bolzano-Weierstrass has is that your set be closed and bounded.

So again, to reiterate, Bolzano-Weierstrass says if I have a closed and bounded subset of a metric space-- or not a metric space-- I should say it's the real numbers. If I have a closed and bounded subset of the real numbers, Bolzano-Weierstrass says that every sequence has a convergent subsequence. Unfortunately, unless y'all would like me to next lecture, I will not be presenting this proof.

This is one that I'm going to black-box for the moment being because it's one that we did in 100A. But if y'all would like me to, feel free to send you an email. And I'm happy to do it next time. All right?

But the one that I think is slightly more interesting is the finite subcover, right? We want to show that for every open cover of the closed and bounded set 0 to 1, we want to show that there exists a finite subcover. And I think that's slightly more interesting.

So let me go through that proof. Let 0 to 1 be contained in the union of Ui's from i equal to 1 to infinity. And I'm going to assume that this is an open cover.

So each of these Ui's are open. In fact, I can assume slightly more because last time what we talked about is that every open set in a metric space can be covered in open balls. This is a fact that is showing up on your second problem set, but is one that I feel is generally safe to assume.

So if you prefer, you can think about this as an open cover by open intervals, OK? Now, how do I show that there is this finite subcover? Well, to do so, what I'll do is I'll consider the following set.

Consider the set of elements c such that the interval 0 to c-- or I should say c between 0 and 1 such that equal 0 to 1-- such that the interval 0 to c has a finite step subcover. Right? So I want to consider the subset of element c between 0 to 1 such that the closed interval 0 to c has a finite subcover.

What is my goal here? My goal is to show the supremum of such c is 1, right? Because if I show it that that c is 1, then I'm done, right? But how do we know that a supremum exists?

How do we know the supremum exists-- in particular, the supremum of this set? This is a property of real numbers.

And I want to recall this because it's one that-- it's been months at this point, if you've done real analysis. It's one that you might have forgotten. This is a bounded set, right?

It's bounded because it's just a subset of 0 to 1. But it's also-- yeah, it's a bounded subset of the real numbers, which means that we have the least upper bound property, which guarantees that the supremum exists. That's how we define the real numbers.

So that's the answer to the question. We know that a supremum exists by the least upper bound property. And because I know that one exists, I might as well give it a name. So I'm going to call this bound c prime.

And I want to show that c prime is, in fact, equal to 1. Suppose for the sake of contradiction that c prime was less than 1. What would this tell us?

Well, let's consider this pictorially because this will be the easiest way to see what's happening. Here we have the interval from 0 to 1. And I have some c prime less than that, right?

And I'm covering the interval from 0 to c prime in finitely many open intervals. So maybe it looks something like that. Union with that. Union with that. Union with something like that.

The issue is that when I'm covering this in open intervals, c prime cannot be the maximum of such elements. How do we rigorously see this? This is a point that I want to be careful about.

The fact that there exists this wiggle room between them follows from openness. So notice then that c prime is an element of the union of this subcover-- so k equal to 1 to m. Right? So I take this finite cover from 0 to c prime.

And I know that then c prime is an element of that. What this tells us-- this set is open. So there exists an epsilon bigger than 0 such that the ball of radius epsilon around c prime is contained in our set. So it's contained in the union of these Ui.

Right? This is the definition of open. We have a union of open set, so I know that it's open. c prime is an element of this open set, so there exists a ball of radius epsilon around that c prime. But then notice that c prime plus epsilon over 2 is bigger than c prime, but it also has a finite subcover, right?

Because the finite subcover that works for c prime will also work for c prime plus epsilon over 2. So this is a contradiction, right? Because this implies that we have an element bigger than the supremum essentially that it has a finite subcover. And that's our contradiction.

OK, so what does this imply? This implies that c prime must be equal to 1, which is what we wanted to show. All right?

So I'll move this up so y'all can see. Any questions about this proof? Feel free to shout them out as I'm erasing the board, if any.

OK. So this shows that 0 to 1 is going to be-- or the closed interval 0 to 1 is compact. And in fact, I'll make a small remark, which is-- let me just get some new chalk. And I'll say this how out loud.

The new remark thing to note is that this would have worked for any closed interval a to b, right? The proof would work by considering the set of c between a to b such that the interval a to b has a finite subcover. And the proof would work exactly the same.

So remark-- a to b is compact. In fact, if I'm considering this in Rn-- so let's say an R2, I know that a, b cross c, d is compact. The proof of the second statement would also work out the same way.

One, we can prove it directly. I can prove-- for sequential compactness, I can apply Bolzano-Weierstrass again. For topological compactness, I can just do one element at a time, just choose one active set at a time and go through the proof.

So these two problems-- to show these two things rigorously are on your problem set, which is an optional problem. But it is a helpful thing to work through. So if you're interested in working through it, I would highly recommend this problem. But it, in fact, gives us the next direction to go, right?

Because what can we know about all of these sets? They are both closed and bounded, which on the Bolzano-Weierstrass theorem is a good thing, right? Because the Bolzano-Weierstrass theorem, which is over there, says that if your set is closed and bounded, it has a convergent subsequence.

So it would be great if topologically compact sets in Rn are both closed and bounded, right? Because that will be halfway to proving that the two are the same, at least on Euclidean space. And so let's actually start proving that because in fact, that will be true.

So proposition-- now, we're going back to proofs or the general theory. Compact sets in R are closed and bounded. Let's prove this.

Proof-- the fact that it's bounded is not too bad. Sorry, let me just-- yeah, the fact that it's bounded is not too bad to show. The proof of bounded-- just fix any point p in x-- oh, sorry. I should be more careful.

I should say let a be a compact subset of x. Now, we're going to prove closed and boundedness of the set A. Boundedness is nice. Just consider the union-- or I should say fix p in x and consider the union from i equal to 1 to infinity of the balls of radius i around p, right?

Then what do we know about this union? We know if we choose our p correctly, that this is going to contain our set a, right? Because again, this is just because we're taking the unions of infinitely large balls, right?

But this is compact, right? So I know there exists a finite subcover. And the finite subcover will look like this. A is the subset of the union from i equal 1 to, let's say, m of balls of radius i around p. But each of these balls are contained in one another, right?

As I make the radius larger, it's going to be contained to the next one. So this is a subset of the ball of radius m in our p, or in fact, is equal to p, I should say. So what this tells us is our set A is bounded, right? This is the definition of boundedness.

There exists point p in a finite radius such that A is contained in the ball around p of that radius. So that implies boundedness. And notice here nowhere did I use the real numbers. In fact, this statement is going to be true about compact sets in a metric space.

So good things there. Let's prove closure now. Or I should say closed. Proving closedness is going to be a little bit harder, as one would expect, right?

There should be some point in this proof set that it becomes quite a bit harder. To prove closure, what we're going to want to do is show that R not including A is open. Realistically, if you feel comfortable enough with it now, you can replace this R with any metric space. The proofs that I present today will be true for any metric space.

And now, let me draw a picture of what we're going to do essentially in our proof. So let's say that this blob is our metric space x. And I'm going to consider my set A should be this yellow blob.

I want to show that x minus A, i.e. the region outside of A, is going to be open. So let's consider some p out here. How do I show that there exists a ball around p of radius epsilon such that it doesn't intersect with A? That's what we want to show for openness, right?

Well, what I can do is compare p to every single q in A. So let's say q1, q2, q3. And what I can do is just cover this and the ball of 1/2 of that radius.

So let's say that's the ball of around p of 1/2 the radius and another ball of 1/2 this radius, and so on and so forth. My goal is to take the intersection of all of these such that they don't intersect with A, right? I know that this ball won't intersect with q1, and so on and so forth.

The issue here, though, the issue that we're going to have to work around, is the fact that the intersection of infinitely many open sets is not inherently an open set, right? And we know that the intersection of finitely many sets is open, but we don't know that the intersection of infinitely many sets is open. And that's the issue.

So in our proof of showing openness, we're going to essentially use topological compactness to go from this definition to-- to complete the proof, to get a finite subcover. And that will be enough for our proof. So that's the outline of how this proof is going to go, but let's do it 10 times as rigorously as the proof.

OK. So fix your point p in x minus A. And what we're going to do is start constructing an open cover of the set A and go from there.

And this group will highlight where we use topological compactness very directly. So let p be in x minus A, or again, you can just let this be the real numbers if you prefer. And consider for all q in A-- I'm going to consider the open cover of A by-- I'm going to consider vq to be the ball around p of radius distance from p to q over 2 and wq to be the ball around q of the same radius, right?

So what I'm doing now is I'm creating both an open cover of A and an open cover of p-- or sorry, yeah, an open cover of p. You want to find a neighborhood of p. So is this, in fact, an open cover of A?

Is A subset of the union of wq's for q in A? The answer is yes. It's pretty much on the face of it, right? Because every element q is in wq.

So the answer is yes. And so what this tells us is I can choose a finite subcover because A, we're assuming, is compact. So A is a subset of i equal to 1 to, let's say, m of wq i, which I'll again write out what this is. This is the union from i equal 1 to m of balls around qi of radius distance p to qi over 2.

Everyone follow so far? Now, the goal is to claim that intersection of-- now, the corresponding v qi is not going to intersect A. So I claim that that is true.

Claim-- the intersection of v qi from i equal 1 to m does not intersect A. How do I know that this is true? Well, suppose-- I'll just state this verbally. Suppose that there existed an element of this intersection that was also in A.

Then that would have to be in one of these balls, right? The ball from qi to p of 1/2 the radius. We know that this cannot be true by the triangle inequality, right?

Because this ball is open. And we know that it must be-- if it's contained in one of these balls, it's not going to intersect the neighborhood of p. Right? I would highly suggest working through this detail.

It's a very specific point, but it's one to double-check and make sure you fully understand. But what this is telling us is that one, the set doesn't intersect A; two, our point p is definitely in this intersection because it's in all of the balls; and three, this tells us that we're done because this set is a finite intersection of open set. So this is open.

So we're done because we've found an open neighborhood around p that doesn't intersect with A. And that tells us that A is a closed set. So we're done. This implies closedness.

Any questions? This proof was certainly difficult, especially if you don't have much experience with topology specifically. But I hope that this picture helps quite a bit.

OK. So we've shown that compact sets are both closed and bounded. Now, is the converse true? Are closed and bounded sets compact?

The answer will be not all of the time. In fact, it will be true for the real numbers, but it will not be true all of the time. To prove that it's true with real numbers-- because again, we're looking at compact subsets of R today-- it's not too bad. I'll just note one small lemma.

If f is a subset of a compact set k, say in x, is closed-- so here I'm assuming that f is closed-- then f is itself compact. I wrote this out in notationally, but what this tells you, again in words, is that a closed subset of a compact set is compact itself. Let's go through the proof.

Well, because f this closed, we know that the complement is open. That's the definition of closed. So in fact, this is if and only if. f complement is open.

And we want to show topological compactness of f. So let Ui be an open cover of f. We want to go from this to a finite subcover, right? That will show compactness of f.

Well, how do we go from this to using the fact that k is topologically compact? Well, let's draw out the picture. Here we have our set x, which doesn't really matter too much in this position.

Here we have k, which we know is compact. And let's consider our set f. What we've done now is we've covered our f in an open cover.

And in fact, how do we go from this to an open cover of k? Well, we know that f complement-- everything outside of f-- is going to be open. So what this tells us is that k is contained in the union of these Ui's from i equal 1 to, let's say, infinity union with f complement. And this is an open cover of k.

This is the open cover that we're interested in, right? So we want to go from this to a finite subcover of f or to a finite subcover. And we know that we can do so because k is compact. So because k is compact, we know that k is contained in the union from i equal 1 to m of Ui.

Potentially, union this f complement-- it's not going to hurt the union one more set, so let's keep it for now. And what we know is that f is a subset of k. So we've gone from an open cover of f to an open cover of k.

We've gone from that open cover of k to a finite subcover. And we note that that finite subcover of k is also going to cover f. So we're done, right? Because we showed that every open cover of f-- I'll just take one more line to be extra clear-- f is therefore contained in the union from i to 1 to m of Ui because no element of f is in f complement.

So we've gone from an open cover of f to a finite subcover, which implies topological compactness. right? That's the definition of topological compactness. And we're going to use this very directly to show that closed and bounded sets in Euclidean space are compact.

So how do we do that? Well, we're just going to shove our compact set in a closed and bounded interval. So proposition-- compact subsets of R are precisely closed and bounded sets of R, right? This is the statement.

We've proven one of the directions so far. We've shown that compact implies closed and bounded. Let's go the other direction, that closed and bounded implies compact.

Proof-- let A be subset of R be closed and bounded. Because A is closed and bounded, we know that A must be contained in some minus n to n because of boundedness. And here n is finite.

And now, what can we do to show that A is compact from here? We can use the lemma we just dated. A is a closed subset of this integral.

So by closedness, this implies that A is compact because it's a closed subset of a compact-- it is a closed subset of a compact set. And the fact that minus n to n is compact is precisely the proof we did earlier, right? So I won't go for that proof again.

But this implies the result that we wanted, so we're done. I'll move it back down in case you want to see it again. This statement is known as the Heine-Borel theorem. It's not, on the face of it, the most easy thing to show, but we're able to cover it in a lecture.

But the issue is that this statement is not going to be true about metric spaces, right? We are not going to have that closed and bounded sets are the same as compact sets in a metric space. But it's so important on Euclidean space that I wanted to bring this proof up today because it does give a good example of how we think about things in Euclidean space and why metric spaces are so important.

All right? So now what we're going to do is show that closed and bounded is the same as sequentially compact in Euclidean space. And that will prove that the same-- all three are equivalent. All right?

So note-- theorem-- by Bolzano-Weierstrass-- I'm just going to state it again-- closed and bounded implies sequentially continuous. Sorry, not sequentially continuous-- sequentially compact. So to show that sequentially compact is the same definition as compact subset R, we want to prove the other direction, right?

Yeah. Actually, yeah. Yeah. So let's prove the other direction.

Proposition-- sequentially compact implies closed and bounded. Well firstly, to show closure, what I'll note is that this is going to show up on your second problem set. On problem set 2, you'll show that closed sets are the same as ones which contain all of their limit points, meaning if I take a sequence, and I know that it converges, the thing it converges to must be in your closed set.

This is a fact that you'll prove on the second problem set, so I won't write it out right now. But what this tells you is that by sequential compactness, if I take a sequence that converges anywhere, then it's going to have a-- sorry, let me say this again. If I take a sequence in my closed set, I know that it's going to-- if I take a sequence, I know that it's going to have a convergent subsequence, which tells us this contains all of its limit points, which tells us that it's closed, right?

This is a fact that I will leave to you to show because I believe that you all can do it. Two, to show boundedness, it's going to be a little bit easier to do because assume A is unbounded, where A here is the sequentially compact set I'm considering. Then what this tells us is I can find a sequence xi which goes to infinity.

And this is a subset of A. Right? So if I'm assuming my set A is unbounded, I can find a sequence which goes off to infinity. What this tells us is that there's not going to be any convergent subsequence of the sequence, right? This is the same thing that we did before, right?

We wanted to show that the real numbers were not sequentially compact. Just take a sequence that goes off to infinity, and we're done. And the same is true here.

Every subsequence of xi or xn diverges, which is a contradiction-- or sorry, not a contradiction. We showed the unbounded implies not sequentially compact. So we know that the converse is true. Sequentially compact implies bounded.

So we've proven the contrapositive. All right? And this shows closed and bounded.

So I know that I'm leaving part of the proof left to you, but I would highly recommend working through it, working through that portion yourself. OK? And in fact, we'll talk more about it on Thursday.

So what things do I want to note now? We've shown compact subsets are the same as closed and bounded sets in R. And we've in fact showed that closed and bounded is the same as sequential compactness.

So in the case of Euclidean space, we're all done, right? We've done what I've set off to do. I want to show that topological compactness is the same as sequentially compact in R. Great.

The issue, again, is closed and bounded sets are not always the same as compact subsets of R because here we inherently used the fact that I can shove A into a compact subset, right? I can set A to be a subset of minus n to n if it's bounded. This fact is not going to be true about metric spaces, right?

Because what does minus n to n mean in a metric? Space it doesn't make sense. But I'll make one small note.

Note-- metric spaces with the following, following being the fact that closed and bounded implies compact, if I have closed and bounded implies compact, then I say that the metric space has the Heine-Borel property. Property.

So it's not true all the time, but if it's true that closed and bounded implies topologically compact, then it's a specific name. It's a special name known Heine-Borel, which makes sense because the theorem is called the Heine-Borel theorem. But that's what's going to fill for metric spaces.

Notice that we should get for that point to be the case because sequential compactness, once you've fully worked through the proof-- the proof of sequential compactness being the same as closed and bounded, that's going to work out in metric space as well. So we should guess that if anything is going to break, it's going to be closed and bounded by topologically compact, OK? So that's it for today.

I do want to iterate that this is not meant to be the easiest portion of this class. This is, in fact, one of the hardest portions of this class. But it's, in fact, probably the most important. If you ask any analysis professor what's the major difference between 100B and 100A, it's experience with compact sets. It's hard, and it's brutal sometimes, but it is worthwhile to do.