[SQUEAKING]

[RUSTLING]

[CLICKING]

PAIGE BRIGHT: So welcome to 18.S190, Intro to Metric Spaces. My name is Paige BRIGHT, though sometimes you'll see it as Paige Bright online. That's just because I'm going to change my name, but you can just call me Paige. And, yeah, this is Intro to Metric Spaces, where today we're going to talk about the connection between what's covered in 18.100A and what's covered in 18.100B, or, respectively, p and q.

And for those of you who don't know, there's this nice little rubric that I like to draw every time I teach this class, because it was drawn for me in my first year, and it helped me really understand what the differences between the courses were, which were not CI-M, the communications class, communications class, real analysis on Euclidean space, and-- or I guess I should say more than Euclidean space. And then the classes go as follows 18.100A, B, P, and Q. But realistically what I hope to highlight today is the fact that there's not too much different between these two courses. It's just a conceptual leap.

But this conceptual leap is really important to have for next courses, so 18.101, 102, 103, 901. There's any number of classes that come after this that having intuition of metric spaces will be deeply helpful for. So, yeah, metric workspaces are going to be what lies on the interim between A, B, and P and Q.

A little bit about me before I jump right in. I'm a third year at MIT. I've been teaching this class for two years. This is my second year. And I started teaching it because one of my friends was in 18.102 and having a really hard time with things like norms, which is a version of a metric space which we'll talk about a little bit in this class. But it's rather unfortunate if you get to the next class, and it seems like you're unprepared for one of the introductory tools, but one that comes with a lot of baggage, a lot of conceptual baggage, to keep in mind. And hopefully today, through today's lecture, you'll see what some of those concepts are.

But let's start with a simpler example before we jump right into metric spaces. We're just going to talk about what makes real analysis work. What's the basic tool of real analysis that makes it work? And as I'm certain y'all have realized through the definitions that we use in real analysis, like convergence sequences, continuous functions, all of it relies on a notion of absolute values and Euclidean distance. And so let me write out what that Euclidean distance is.

So given two points x and y in Euclidean space, we call the distance between them x minus y and Rn to be the sum of the distances between each of the components squared to the 1/2. This is just the Pythagorean theorem. And this definition is known as Euclidean distance, as I'm certain y'all know from 18.02.

Mostly speaking, though, we focus on R, just because most of the theory breaks down to studying everything in R. So if you don't prefer thinking in n dimensions, you can always think in one dimension, and, for the most part, everything will be fine. And in fact, we'll see why it will be fine in a moment. But what are the most important properties of this distance?

Well, firstly, we have that it's symmetric, meaning that the distance from x to y is the same as the distance from y to x over Rn. And this is something we should expect. We should expect the distance for me to you is the same as the distance from you to me. We also have that it's positive or positive definite, specifically that the distances between points is always bigger than or equal to 0. And if the distance is exactly equal to 0, this is true only if-- if, and only if-- the two points are the same.

So the distance between you and someone else is only 0 if you are the same person. And if any of the notation I use doesn't make sense, I'm happy to elaborate. But this is the notation for if and only if. And then finally, the most important one, arguably, though the hardest one to show most of the time, is the triangle inequality, that the distance between x and z is less than or equal to the distance from x to y plus the distance from y to z. These are the three properties that make real analysis work.

Now if you think back to all those definitions that I was talking about earlier, convergence sequences, Cauchy sequences, and so on, all of them have something to do with these absolute values. And that's because these absolute values are, in fact, a metric. When I define what a metric space is, which will just be in a moment, it's simply going to be a set with these three properties. So let me go ahead and write that down.

A metric space is a set x with a function d, which will be called our metric, which takes in two points in x and spits out a real number. In fact, it doesn't just spit out a real number. It spits out one that is between 0 and infinity, and not including infinity. And it satisfies the three properties listed here.

And d satisfies-- I can rewrite these definitions if y'all would like, but just replace the-- what's really happening is you replace the absolute values between x and y with the distance from x to y. Does this make sense to everyone? And one point of confusion last year was, what does the notation with this times thing mean? It just means that we're taking in one point in x and another point in x, so distance from x to y.

This is the definition of a metric space. It's not terribly insane. But what it allows us to do is study all the tools that we did in real analysis all of a sudden in a completely new setting. So now the first few examples I'm going to talk about are ones on Euclidean space, just because we have a lot of intuition about that already. But soon, as we'll see, we can study this on plain old sets. You can study this on topological spaces. You can study this on functions. It's a very powerful tool that generalized a lot of math when it was first invented.

All right, so let's start off with some more of these examples. So another distance on our end that we can consider-- hi-- is one that takes in two points, x and y. In fact, I'll call this one d infinity for the metric-- that takes in x and y and spits out the maximum over every component of the distances.

So in other words, if I have a vector here in 3-dimensional space x, and one that goes here, the largest distance would probably be the vertical one, I will just say. And that would be the Euclidean distance. It's a maximum distance between the components. So if I write x as x1 through xn, symmetric just gives out the maximum value of the differences of these components.

Since you just walked in, let me just briefly explain what's happening. So we define what a metric space is as simply as set with a function that acts like a distance. And the three properties we want it to have is we want it to be symmetric. So the distance from me to you is the same as you to me. We want it to be positive definite. We don't want our distances to be negative, and the triangle inequality, which we should be a little bit more familiar with. So, yeah, that's the definition of a metric space, and I'm just defining a new example.

But if I ask you on a problem set, which I will do, to prove that something is a metric, you have to prove the three properties, which isn't, in this case, terrible to do, because-- I guess prove, technically-- one, it's definitely positive-- or it's definitely positive because we're using absolute values. It's not going to give us something negative. But the thing that's a little bit harder to check sometimes is you have to check it's positive definite, so definite. So if the distance d infinity from x to y is 0, what does that tell us?

Well, that tells us that for all i, the distance between xy and yi must be 0. Why is this true? Well, if it wasn't 0, if it was something slightly bigger than 0, then the supremum metric on this d infinity would make this have to be bigger than zero, right? So this implies the x is, in fact equal, to y. That's how we define equality on Euclidean space.

And the opposite direction, usually a little bit easier, if x is equal to y, then the same conclusion holds. So usually that direction is a little bit easier if you assume equality showing that it's 0. But the other direction can be a little bit harder. Assume it's 0. What can you say? Or assume it's not 0. What can you say to do a proof by contradiction?

And then lastly, we have to-- or, sorry, no, two more. It's definitely Symmetric the fact that it's symmetric just follows from the fact that you can just swap any of the two in the maximum. It's not too bad. And three, you have to check the triangle inequality. And this is where it can be a little bit more tricky because you might be inclined to just say, oh, it's the maximum of things that satisfy the triangle inequality. But we have to be a little bit more careful than that, and I'll explain why in a moment.

So let's consider two-- three points, x, y, and z in Euclidean space, Rn. And let's suppose we're considering the distance d infinity from x to z, which I'll write down again is the maximum of i between 1 and n of the different distances between the components. Now we want to go from these to information about y.

And what we want to do is exploit the fact that we know that the absolute value satisfies the triangle inequality. But this maximum function makes it a little bit more difficult. Any ideas what we could do before we use the triangle inequality? Well, in this case, there's only finitely many terms that we're considering. So we know that a maximum has to exist.

So I'll just call that term j. This is as opposed to looking at the maximum, I'll just consider the distances of xj to zj. And then I can apply the triangle inequality. And this is where it's important to double-check that everything is working out. We had to know that a maximum existed before we could apply the triangle inequality.

So then this is less than or equal to the distance from xj to yj plus the distance from yj to zj. And now we can apply the maximum operator again to both of these because we know that this one is going to be less than or equal to the maximum over i of xi to yi. And we know that this term will be less than or equal to the maximum of that side, so plus the maximum yi to zi. And then we're done. Any questions? Happy to reiterate any of these points.

So that shows that it is, in fact, a metric. And it's a pretty useful one at that. What this is telling you is if I want to study things on Rn, I can essentially just study them on R, and things will work out the same, because we can just study the maximum of the differences between their components. Let's look at one more example on Euclidean space.

Now as opposed to looking at the maximum, I could instead sum over all of the terms, which will, in fact, be a little bit easier. So I'll call this the d1 distance between x and y in Euclidean space. This is the sum of the distances between the components, so yi, xi to yi. And I'll leave you to prove that this is, in fact, a metric. But let's run through the checklist one time.

One, we know that it's definitely positive. Positive definite is a little bit harder, right? I mean, if they're equal, then we definitely have that the distances between these ones are all 0. If the distance is 0, then note that we're taking the sum over non-negative things. So if the thing on the left-hand side is 0, every single term in the sum must be 0. And that allows us to use the same conclusion. And then you can apply the triangle inequality right now because we are just summing over terms. So the triangle inequality automatically applies. So if you want to write through that, feel free to.

The thing that's more important that I want to note here is that this is known as the l1 metric. In fact, if we wanted to, I could just replace this d1 with a dp, raise this to the pth power, and raise this to the 1 over p. And that would give me the lp metric, which is deeply important in functional analysis. It's a tool that-- it's a surprise tool that will help us later.

So if you take a class on functional analysis, this is likely one of the first ones that you'll consider, though to check that the lp metric is, in fact, a metric is a little bit more difficult. There's a problem on the first problem set that's optional if you want to work through those details. And there's a lot of hints, so that's helpful.

So now that we've gone through all of these examples on Euclidean space, before I go much further, I want to explain how this actually relates to the definitions that we already have, because you might be sitting here, thinking to yourself, why does this matter? We've been studying real analysis, and now we're going back to studying a function that acts like a distance. So what?

Well, these four definitions we're going to be able to rewrite in terms of sequences of a metric space. So let's write out what a conversion sequence is. Let xi be a sequence the metric space xd. And let x be just a point in the metric space. Here our sequence is just as we defined it in real analysis, it's a bijection between this set of points and the natural numbers.

And then we defined convergent sequence. For all epsilon bigger than 0, there exists an n in the natural numbers such that for all n bigger than or equal to n, the distance between xn and x is less than epsilon. So this is what we mean for converge-- a sequence that converges to the point x. And this is essentially just the same definition that we've been dealing with in real analysis. Just replace the distance with absolute values.

And I'm just going to go through these other definitions as well. So, same setup let xi be a sequence in xd. Then a Cauchy sequence is such that for all epsilon bigger than 0, there exists an n in the natural numbers such that for all n and m bigger than or equal to n, the distance between individual points of a sequence is less than epsilon. So here's the definition of a Cauchy sequence.

Who here hasn't heard of a definition of an open set before? No worries if not. Some classes don't cover it. Yeah, let me just briefly explain up here what an open-- actually, I'll do it down here-- what an open set is. Actually, no, I'll do it up here so that y'all can keep taking notes if you need to.

So an open set is just a generalization of what an open interval is. So an example of an open set in R is, for instance, 0 to 1. It's such that if I'm considering the interval from 0 to 1, the thing that makes it an open interval is the fact that for any point I consider, let's say, that one, that I can choose a ball of radius epsilon around it, so a distance epsilon on both sides, such that the ball of radius epsilon around the point x is contained between 0 and 1. And I can do this for every single point in the interval 0 to 1. This is what makes it an open interval.

And the definition of an open set is going to be essentially the same, where instead, the definition of the ball of radius epsilon, this is defined as the set of points in your metric space such that the distance from x to y is less than epsilon. So that's the only difference between what's happening on a metric space rather than Euclidean space. Before, again, this would be absolute values.

So to define an open set, a set A contained in x where x is a metric space is open if for all points A and A, there exists an epsilon bigger than 0 such that the ball of radius epsilon around A is contained in A. Pictorially, if you prefer to think about it this way, this essentially means that there is no boundary on your set.

So if I were to draw a conceptual diagram of what's happening here, here's my set x. Here's my subset A that lies in it. And I'm just cutting out the boundary. And this allows us to say, OK, I can get for any point, even as close as we want to the boundary of A, I can still squeeze in a little ball there.

And in fact, thinking about it in this conceptual drawing is deeply important. And I'll give an example of this in a moment. But this is the definition of an open set. Does everyone feel a little bit comfortable with what this is? Cool. I know it's just a little bit weird, because some classes don't cover it, because on 100A, you don't particularly need it as much.

Now continuous functions are going to be a little bit weirder, because before when we studied continuous functions, we studied ones from R to R. It takes in values in the real numbers, and it spits out a real number. But now that we have a notion of a metric space, we can now study continuous functions between them, because all we really need again, is absolute values, right?

So let f be a function between x and y where both of these are metric spaces. In fact, I'll write that out now. I'm going to say that x has metric dx, and y has metric dy. And this is the appropriate notation for noting which one goes with which one. Then I say that f is continuous, continuous, if for all epsilon bigger than 0 there exists a delta bigger than 0, such that if the distance between two points x and y is less than delta-- or, sorry, dx-- then the distances in the space y of f of x and f of y is less than epsilon.

Now let's double-check that this definition actually makes sense. This is a very helpful way to remember what metric goes where. x and y, again, are two points in x, right? So it only makes sense to consider the distance on the metric space x acting on them. And then f takes points in x to points in y So here on the right-hand side, we should have the metric on y.

And this is the definition of continuous. I'm going to prove that a very specific special operator is continuous today. And that will be helpful conceptually. But this is far more general. If you prefer, you can just let y be the set of real numbers. And this is still a really powerful tool already because now we can study them between metric spaces x and the real numbers.

So I just wanted to highlight what these four definitions convert over to in this class. Tomorrow-- not tomorrow, sorry. On Thursday, we'll be proving quite a few theorems about all of these spaces. Yeah, on Thursday we'll be proving quite a bit of properties about these definitions. In fact, I'm calling that the general theory of metric spaces. It will be a very intense class in terms of proof writing. But I wanted to bring them up today so that y'all all saw that it wasn't just random stuff.

So though this is all good and dandy, we've only talked about Euclidean metrics. And this is fine. But things get a little bit weirder once you go from finite dimensions to infinite dimensions. And in fact, the definitions of, for instance, a continuous function will look a little bit different, right?

So let's give an example of one that's just a little bit weirder than a set of points in Euclidean space. Actually, I'm going to rearrange this a little bit. First, I want to give a really weird example that's pretty simple, but interesting, nonetheless. Let x be any set. Then I define the metric on x, so the metric between two points x and y, could be as follows-- 1 if x is not equal to y, and 0 otherwise.

So if x is equal to y, then it's going to be 0. Why is this important? Because this is telling you that every single set you can give a metric to. Now that-- this metric isn't all too interesting, but it is a useful example to keep in mind, right?

And let's prove that this is a metric, because even though it's a simple function, it's going to be a little bit weird. One, definitely positive or specifically non-negative. I guess I should say that, non-negative. Now we just have to check positive definiteness. So if x is equal to y, then the distance from x to y is 0. And the distance from x to y is 0 only if x is equal to y. That's the definition of the metric.

Two, it's definitely symmetric because equality is symmetric. And finally, three, the hardest one, triangle inequality. Why is this one the hardest? Because now we have three points, and we have a binary acting on them, right? So let x, y, and z be in x.

Then we have three possibilities, or quite a few more, but I'll state it in generality. One, x is not equal to y, y is not equal to z, and z is not equal to x. Two, x is equal to y, but y is not equal to z. Or three, all of them are, in fact, equal.

Why are these the only three cases? Well, because x, y, and z that we're choosing are just arbitrary, right? So if I wanted to decide the case where y is equal to z, but x is not equal to y, I would just relabel them. And so these are the three cases we have to check to make sure that our metric is actually a metric, which happens in cases like this, where our metric is defined via a binary or similarly simple cases.

So let's check each of these. Actually, I'll just do it here. I can squeeze it in here. If none of them are equal, then the distance from x to z is just going to be 1, by definition. And this is certainly less than or equal to 2, which is the sum of the other two distances. So we're good there. Two, the distance from x to z-- what's this going to be? Anyone? If x is equal to y, and y is equal to z, what's the distance from x to z? Yeah?

STUDENT: 1?

PAIGE BRIGHT: 1, exactly, because x being equal to y means x is also not equal to z. And this is equal to 1, which is the distance from x to y plus the distance from y to z, because only one of them is 1. And then, of course, the last case, when you plug in the x equal to y is equal to 0, we're just going to get 0 is equal to 0 plus 0. So we're all good there.

Now you're going to have an example on your homework that is similarly based off of a binary, where instead, the binary is going to be based off of if three points x, y, and z or specifically-- sorry, I'll just say x and y-- where if x and y are in R2, then the distance between x and y is just going to be the regular distance in R2 if x, y, and 0 are colinear and the sum of the two otherwise. So to draw a picture of what I mean here, if x and y lie on the same line, then the distance is just as normal. If they're not, then I have to add up the two distances between them.

This is similarly a binary. And when you're doing this on your homework, one thing I would suggest is breaking it up into similar casework. Either they all are colinear, or none of them are, or only one of them is. That's the sort of thing that you should do on a problem like this on the homework. This casework it can be a little bit annoying, but is useful to do.

So even though this metric is vaguely interesting in that it makes any set a metric space, let's start studying a set that we care about a little bit more, continuous functions. And we can, in fact, define a metric on them. So I'm going to define c0 of ab to be the set of continuous functions from ab to R. And I'm going to define a metric on them.

And the metric-- so, example-- or I guess I should say this is the definition. Example, if I'm considering two functions f and g in the set of continuous functions on the interval a to b, then the distance between these two functions is going to be the supremum over all x in a to b of the distances between the two points or between the two functions. This is what I'm claiming is a metric.

And it's going to be a little bit difficult to prove in one of the steps, if you had to guess, probably the triangle inequality, because that's where everything messes up. But let's go ahead and check that this is a metric. Well, specifically, first, it's definitely going to be symmetric because the distance from f to g is equal to the supremum of the distances between the functions f and g, which is certainly equal to the distances of the supremum from g of x minus f of x.

Now one thing to be careful about here, though, is-- you don't have to be as careful about it with symmetry, because symmetry is a little bit more clear that you can mess around with these operations. With the triangle inequality, it becomes a little bit harder. There is a technical thing I'll come back to you later about this step. But essentially, symmetry is done.

Two, the positive definite. Well, it's definitely going to be positive because it's absolute value of zn or non-negative. And we have to check definiteness. Well, if f is equal to g, then by definition, f of x is equal to f of g of x everywhere for all x, which implies that the distance between f and g is 0.

Now what if the distance between f and g is 0? How do we conclude that f and g are, in fact, equal? You know that because it's a supremum, everywhere, at every single point, they have to be equal. You can also make an argument via continuity, which is-- which makes use of the extreme value theorem. But we don't have to get into that much detail. The harder part, though-- so the positive definiteness.

The hardest part, though, is going to be the triangle inequality, because here, we do have to use the extreme value theorem. A triangle inequality-- so let f, g, and h be a continuous functions on a to b. Then let's consider the distance from f to h minus h of x. We want to go from this to information about f and g and g and h so we can apply the triangle inequality.

Well, to do so, this is the only part where we use the fact that the functions are continuous, because knowing that they're continuous, we can apply the extreme value theorem to them, right? So we know that this supremum has to exist somewhere. And I'm going to just let that point be y. So this is going to be equal to the distance from f of y to h of y.

So what now? Well, now we can just directly apply the triangle inequality because these are absolute values. And we know that this is going to be less than or equal to the distance from f of y to g of y plus absolute value of g of y minus h of y. And then we can take the supremum of both of these individually, and we'll conclude the proof.

But the thing we had to make sure of first was we couldn't just apply the triangle inequality right away. The supremum operator acting on it made it so that we had to go through these steps individually. And in fact, if we wanted to be really careful, it might have made sense to do it here as well, where we know the supremum exists, and it's going to be equal to if I swap the two orders. But really, it's important in the triangle inequality.

And the thing to note here is that this is precisely the same sort of argument we had to do right here for the infinity metric. We knew that maximum existed, so we just went with it and ran. And that's, generally speaking, good advice. If we're dealing with something that's continuous, try narrowing your focus down as much as possible to a single point or to a single function, whatever you can do to finish off your proof.

Now, depending on your background with analysis, you might be wondering to yourself, why do we only care about c0? Why don't we care about-- what is c, c0? Why is that 0 there? Well, in fact, the reason it's there is because we can study differentiable functions, specifically continuously differentiable functions. So, definition-- actually, I'll just define it for ck.

This is the set of continuous functions on ab such that the first k derivatives of f1 exist, and two are continuous. This is the set of what are known as continuously differentiable functions, so ones in which when I take the derivative, this-- the derivative is going to be continuous. And once we have this new set defined, we can find even weirder metrics. Well, I guess not all that much weirder, but ones that are slightly more complicated, I guess I should say.

So, example-- let's just consider the ones on c1 of ab. We want to show that-- the following, dc1 of f and g. I'm going to define what the metric is first. And what we want to show is that this is, in fact, a metric. This is going to be the supremum over x in a to b of f of x minus g of x plus the supremum of x and ab of f prime of x plus g prime-- or, sorry, minus g prime of x. This is going to be our new metric.

Proof that this is, in fact, a metric-- proof-- one, definitely non-negative. But we have to check positive definiteness. Well, if f is equal to g, then f of x is equal to g of x everywhere. And in fact, f prime will be equal to a g prime, which implies the same result, because you can just take the derivative of this. So that includes one direction.

If the distance on c1 of f to g is 0, what can we say? Well, again, we're summing over non-negative things. So if the sum of two non-negative things is 0, that implies that both terms must be 0. This fact, I cannot iterate it enough, is deeply important. When you're working on your problem set, you will have to use this fact repeatedly, that the sum of non-negative things being 0 implies that the individual terms must be 0 as well, to prove positive definiteness. That's where it mostly comes at.

So what this implies is that the supremum of x in ab of f of x minus g of x equals 0. And we've already explained above how this implies that this metric or that f equals g. So this is one of those examples where you want to boil it down to the examples you've already done before. So that proves positive definiteness. Symmetry is pretty immediate. And the triangle inequality follows by the same argument up here, right? We know that the-- we know that it-- the supremum exists for each of the terms, because it is-- because all of the terms are continuous.

But the thing I want to note here is, what stops us from doing this at c1? Can we do this at ck? The answer is, yes, it's not too much more difficult. You instead sum over the 0th derivative, the first derivative, so on and so forth, up until the kth derivative. And that's fine.

What if I want to study it on smooth functions, which I'll define right now? We define c infinity ab to be the set of smooth functions, i.e., infinitely differentiable. Notice that if it's infinitely differentiable, each of the derivatives must be continuous, because if it wasn't, then the next derivative wouldn't exist, or the next derivative wouldn't be continuous.

So this is the set of smooth functions. What stops us from just taking the sum over all of these terms, over all infinitely many terms? The issue is that there are infinitely many, right? Again, a metric must be-- it must be the case that for a metric that the value that you get out is not infinity. I guess in theory you could mess around with this a little bit, and you'd get weirder types of metrics. But for our purposes, you don't want it to be infinity.

But remark, there is a not-too-bad addendum to this. On the P set, there's an example that you can work through that's an optional one, where you define a metric on this, where what you do is you sum over this interesting fraction, dck(fg), 1 plus dck(fg). So it's summing over a bunch of metrics. You have to, in fact, show that this is a metric, which is one of the other problem set problems. But it's not that you can't define a metric on smooth functions. It's that you just have to be careful of the fact that there are infinitely many terms.

And one small thing to note-- is it possible that we could have gotten away with not including this term in the sum? The answer is no. And I'll let you think about this some more. But the idea ends up coming from if we take this away, and the distances is 0, then what's the distances between f and g prime imply about f and g? So I'll let y'all think about that some more because that is a problem set problem. But the answer is no. And it's a bit interesting to think about why. And that reasoning is exactly why we have to be careful about the infinitely differentiable case.

So now that we've done this, now that we've defined a metric on c1 and c0? Yeah?

STUDENT: Can we use the very first metric as a metric on all of these spaces? Does it still read valid?

PAIGE BRIGHT: Yeah, yeah, there's nothing wrong with that. The issue is that it's not encapsulating as much information as we want it to. So as we'll see in a moment, we want to understand differentiability and integration as functions that are continuous. So that's-- great question. Yes, we could have just considered the first one as a metric on all of them in the same way that we can consider the trivial metric, this one-- this is known as the trivial metric-- to be a metric on all of them, right? But we want more information, when possible. Great question. So I will come back over here.

So now that we have c0 and c1 defined, we can, in fact, define differentiation and integration-- or not define. We can state that differentiation and integration are, in fact, continuous. And I'll do that right now. So I guess this is a proposition.

If I consider differentiation as a map between c0 of ab, so continuous functions to-- or, sorry, c1-- to continuous functions, my claim is that differentiation is continuous as a function-- or I guess map is a better word-- but a map between metric spaces. Does this notation make sense to everyone, the differentiation as a map? Cool.

So let's check that this is, in fact, a continuous function, which is pretty nice to do. By this setup, it's made to be nice, which addresses your question about why we sum over the two terms. So to prove continuity, what's often best to do is just to consider what the distances, in fact, are. Let's write out the left-hand side and the right-hand side of this implication.

So let f and g be in c1 ab. And then we consider the distances on c1 of f and g. And we want to say that if this is less than delta, then the distances on c0 of the derivatives-- so I'll just say f prime and g prime-- is less than epsilon. We want to find for all epsilon bigger than 0, there exists a delta such that this is true.

Well, let's state out what these two metrics actually are. And that'll make it clear that the reduction is not that bad. So again, this is equal to the supremum over x and ab of f of x minus g of x plus the supremum of x in a to b of f prime of x plus g prime-- or, sorry, minus g prime of x. And here, this is simply the supremum over the distances of f prime of x minus g prime of x.

The reason why this step is nice to do, first, is because we notice that this term is precisely the same as this one on the left-hand side. And all the terms are non-negative. So let delta equal epsilon. If you let delta equal epsilon, then this term being dc1 of f to g being less than delta implies that this term must be less than delta, or, in fact, less than epsilon by its construction. So that implies that this term is less than epsilon, which is what we wanted to show.

So this shows-- or that's the end of the proof. This shows that differentiation is a continuous map. And here, we use the fact that we're summing over the two terms. Now on your problem set what you'll do is you'll show that integration from a to a point t from a to b is, in fact, also a continuous operator. So that, I think, is pretty interesting. It's going to be a little bit harder than this one because this direction doesn't include both terms. But I think it will be a worthwhile exercise to work through. Let's see what's next.

So while you're going to show that integration is a continuous operator, this doesn't stop us from studying integration as a worthwhile metric right now, because we can define integration not too bad. So we define i1 to be a metric on c0 ab times c0-- actually I'll change this a little bit-- from ab 0, 1, 0, 1, to 0 to infinity given by i1 of f and g. It's simply the integral from 0 to 1 of f of x minus g of x dx. All right, we're going to show that this is, in fact, a metric.

Now what are the three components? One, symmetry. Here we have to use the fact from Riemann integration that because f of x minus g of x is equal to g of x minus f of x under absolute values, then we get that the integrals are, in fact, the same from 0 to 1. You can just apply-- has everyone seen this fact about integration? Cool.

One way you can prove it is that if they're less than or equal to each other, then the integration still follows, and you can apply the inequality both ways. So I just wanted to say that. But symmetry works out pretty nicely.

Two, we know it's non-negative. But positive definiteness is going to be a little bit harder this time. If f is equal to g, then clearly i1 of f to g is 0. This comes from the fact that the distances between the two points is just literally 0. But how to do the other direction? Because this is where it can be a little bit more complicated.

What if, for instance, the two functions differ by a point? Then their integrals are still the same, but what does this actually tell us? Let me write this out. What if i1 fg equals to 0? Well here's-- yeah, go for it.

STUDENT: Well, it would continuous, though. So it would have a little ball where it all is, like where the [INAUDIBLE] is.

PAIGE BRIGHT: Exactly, right. So because it's continuous-- suppose that this was not 0, proof by contradiction. And I'll draw a nice picture of what's happening here. So assuming that the integral of them is non-zero, then there must at least be one point. Let's say this is f, and this is g. It's pretty close by.

Then there must exist at least one point such that they're different. Otherwise, if there was no singular point such that they were different, then they must be equal everywhere. And we reduce to the first case, right? Or, then, in fact, that's what we want to show. But it not being equal to 0 means that there must exist a point where they're not the same.

And what this tells you is there must exist a little ball around x, ball of radius epsilon, let's say, such that they're, in fact, not equal on that entire ball. But we can use this to show that the integrals over that, the integral over the entire interval 0 to 1, must therefore not be 0, and reach a contradiction. So specifically we would reach the conclusion that f is not equal to g. And I'll let you all work through the statements of that.

But, yeah, it is precisely the fact that if there's a point where the integral is non-zero, then there must be a ball around that point such that the distances is not 0. Cool. And this just highlights why the fact-- why continuity is important here yet again, not just to show that nice things that are continuous.

Lastly, the triangle inequality-- I just want to state this, because it's not too bad. We just note before even integrating that the distances between f of x and h of x is less than or equal to f of x minus g of x plus g of x minus h of x. And then integrate on both sides of the inequality. And that allows us to reach our conclusion.

So in statements like this, either choosing a point such that the supremum exists. If it's based off the supremum, that's a good way to do it. The other way to do it is try to utilize facts about the absolute values before you even apply the metric. So before you even apply integration, see what you can say. These are the few techniques that are deeply helpful.

And there's one thing I want to note here. Do I still have it up? I do not because it was a while ago. We can view integration as a sum. It's an infinite sum. Sure, it's Riemann integration, but a sum nonetheless. What this is called-- I'll note it here-- this is known as the capital L1 metric on c0 ab. Or specifically, the notation would be it's L1 ab. This is symmetric.

The reason it's called L1 might remind you of the fact that we had a little l1 that we'd find earlier, where instead there, we're summing over finitely many terms. It's the same exact notion. It's just the only difference here is that we're summing over infinitely many terms. And that's why the notation is the same.

Now this notion is deeply important for 18.102 because this is the definition-- this leads to definitions of Lebesque integration, which we'll talk a little bit about in this class. But really, it's a question of integration as we already understand it via Riemann integration. It's really, really bad, right? What if we don't want to study things on continuous functions but instead ones with finitely many discontinuities? The Riemann integral just becomes so much more annoying to deal with. So the Lebesque integral is the way around that. And it's called capital L1 because of Lebesque, the person who invented it.

But one thing I want to note is that all of the spaces that we've considered so far are vector spaces. So for those of you all who have studied linear algebra, this is-- this will be slight review. But I just want to note it here because it is somewhat important and will come up a little bit later in the class.

A vector space is simply a space in which you can add two terms together, and it'll stay in the same space. This isn't a class on linear algebra, so there's not going to be too many times when the definition of a vector space comes up. But it's a useful thing to keep in mind and to know about.

So, vector space. This is the TLDR. I'm not going to write out the entire definition because it's, in fact, quite lengthy. But it's a space V such that we have addition, which maps from V and V to V, and multiplication, which maps from real numbers across V to V. And ways to think about this is, for instance, the set of continuous functions, c0 ab. We have addition defined on this such that we define f plus g to just be f of x plus g of x everywhere. And we define a constant times a function, which is what R here is doing, as just being the constant times f of x everywhere.

The basic idea for vector space is that addition acts how you would want it to. It maps from two points in the vector space to a point in the vector space. And scalar multiplication acts the same way. It takes a scalar and a point in your vector space and maps it to another point in your vector space.

Why do I note this? Because metric spaces do not only have to be on vector spaces. We can have it be much weirder. A good example of this-- so if you did not catch the basic idea of what a vector space is, I'll-- this example might highlight it. Example-- consider the sphere S1 to be the set of x in, let's say, R3 such that-- actually, you know what? I'll just do a circle, the circle of radius 1 in R2.

This is not a vector space under usual pointwise addition, because if I take a point on my vector space-- let's say that one-- let's call it x. And I do x plus x, I'm going to end up with a point that's not on the circle. This is not a vector space in the usual sense because we want addition to be contained in the vector space.

Why is this important? Because we can still define distances on the circle, right? We can define distances on the circle via just the regular distances, so the distance from x to y being the usual distance on R2. You can also define it via the shortest distance between them which is known as a geodesic. I don't like this chalk, but anyways.

So the point that I'm trying to highlight here is you don't always have addition defined, which is going to limit the number of theorems we can actually say about metrics, right? Nowhere in our definition of a metric does it tell us how to add two functions together because sometimes we just simply can't. The set of-- or when we can, it starts becoming more like functional analysis, which is 18.102. We'll talk briefly about that in two lectures from now. But this notion of it not being a vector space is pretty important. Was there anything else I wanted to note? We might end today a little bit early, which will be nice.

The main other thing I wanted to note here was something that I noted earlier for the little l1 spaces, is you can also define lp spaces f of x minus h of x to the p-- to the 1 over p. This is known as lp spaces. So functions in which this term is finite is an lp space.

All right, so I actually went through this 20 minutes earlier-- faster than last time. So if y'all have any questions, I'm happy to talk about more of the material. But we might just end today a little bit early. I don't want to get into the general theory quite yet.

What I would highly suggest is trying to sit with this notion a little bit, because though it seems like a simple notion, what we've really done here today is gone from understanding, for instance, functions as things that take in points and spit out points as things that can be manipulated, as things that have a distance between them. And once we study sequences of functions, which is deeply important in analysis, it's going to be-- become more and more important.