HERBERT GROSS: Hi. Well, I guess we're in the home stretch of the course. We're down to our last concept, and it will take us essentially two lectures to cover this new concept. Today we'll do the concept in general and next time we'll apply it specifically to the concept of series. But the concept we want to talk about is something called 'Uniform Convergence'.

Let me say at the outset that this is a very subtle topic. It is difficult. It seems to be beyond the scope of our textbook, because it's not mentioned there. Consequently, I will try to give the highlights as I speak with you, but the supplementary notes will contain a more detailed explanation of the things that we're going to talk about. Now, to set the stage properly for our discussion of uniform convergence, I think it's wise that we at least review the concept of convergence in general.

Well, let's take a look over here. Recall that when we write that the limit of ''f sub n' of x' as 'n' approaches infinity equals 'f of x' for all 'x' in the closed interval from 'a' to 'b'-- when this happens, another way of saying this is that the sequence 'f sub n', the sequence of functions 'f sub n', converges to the function 'f' on the closed interval from 'a' to 'b'. Now, again, these tend to be words unless you look at a specific example. Let's just pick one over here. Let ''f sub n' of x' be ''n' over '2n plus 1'' times 'x squared'.

Notice, of course, that the value of this particular number depends both on 'x' and 'n'. At any rate, let's pick a fixed 'x', hold it that way, and take the limit of ''f sub n' of x' as 'n' approaches infinity. In other words, as we let 'n' approach infinity in this case, notice that the limit of 'n' over '2n plus 1' becomes 1/2. 'x' has been chosen independently of the choice of 'n'. Consequently, the limit function in this case, 'f of x', is 1/2 'x squared'. And in this case, what we say is that the sequence of functions 'n 'x squared'' over '2n plus 1' converges to 1/2 'x squared'.

Now what does this mean more specifically? In other words, let's see if we can look at a few specific values of 'x'. For example, if I choose 'x' to be 2-- in other words, if I choose 'x' to be 2, if we look over here, this says that ''f sub n' of x' is '4n' over '2n plus 1'. If I now take the limit as 'n' approaches infinity, I'm going to wind up with what? 'x' is replaced by 2, 2 squared is 4, 1/2 of 4 is 2. In other words, the limit of ''f sub n' of 2', as 'n' goes to infinity, is 2.

In a similar way, if I replace 'x' by 4, 1/2 'x squared' becomes 8. And so what we have is, at the limit, as 'n' approaches infinity, ''f sub n' of 4' is 8. The key thing being that once you choose 'x', notice that for a fixed 'x', ''f sub n' of x' is a constant and you're now taking the limit of a sequence of numbers.

At any rate, here's what the key point is. What does this mean by our basic definition? By our basic definition, it means that we can find the number, capital 'N sub 1', such that when 'n' is greater than capital 'N sub 1', the absolute value of ''f sub n' of 2', minus 2, is less than epsilon. In a similar way, this means that we can find the number capital 'N sub 2' such that for any 'n' greater than capital 'N sub 2', the absolute value of ''f sub n' of 4' minus 8 is also less than epsilon. The key point is that 'N1' and 'N2' can be different. In other words, you may have to go out further to make this difference less than epsilon than you do to make this difference less than epsilon.

In other words, you see what's happening here-- and we're going to review this in writing in a few minutes, so that you'll see it in front you. What happens here is that you see that for different values of 'x', we get different values of 'n'. And since there are infinitely many values of 'x', it means that in general, we're going to be in a little bit of trouble trying to find one 'n' that works for everything. And let me show you what that means again, going more slowly.

I simply call this two basic definitions. In other words, if we have a sequence of functions 'f sub n', each of which is defined on the closed interval from 'a' to 'b', we say that that sequence of functions converges point-wise-- that means number by number-- to 'f' on [a, b] if this limit, ''f sub n' of x' as 'n' approaches infinity, equals 'f of x' for each 'x' in [a,b]. In other words, given epsilon greater than 0, we can find 'N1' such that 'n' greater than 'N1' implies that the absolute value of ''f sub n' of x1' minus 'f of x1' is less than epsilon for a given 'x1' in [a,b]. In general, the choice of 'N1' depends on the choice of 'x1', and there are infinitely many such choices to make in [a,b].

Now the key point is this, and this is where uniform convergence comes in. If we can find one 'N' such that whenever little 'n' is greater than that capital 'N', ''f sub n' of x' minus 'f of x' in absolute value is less than epsilon for every 'x' in the closed interval, then we say that the convergence is uniform. In other words, if we can find one 'N' that makes that difference less than epsilon for the entire interval, then we call the convergence uniform.

Now, you see, convergence in general is a tough topic. In particular, uniform convergence may seem even more remote, and therefore what I'd like to do now is-- saving the formal proofs for the supplementary notes, let me show you pictorially just what the concept of uniform convergence really is.

So let me give you a pictorial representation. Let's suppose I have the curve 'y' equals 'f of x'. Now, to be within epsilon of 'f of x' means I retrace this curve displaced epsilon units above the original position, and epsilon units below. In other words, for a given an epsilon, I now draw the curve 'y' equals ''f of x' plus epsilon' and 'y' equals ''f of x' minus epsilon'. Now, what uniform convergence means is this, that for this given epsilon, I can find a capital 'N' such that whenever 'n' is greater than capital 'N', the curve 'y' equals ''f sub n' of x' lies in this shaded region. In other words, it can bounce around all over, but it can't get outside of this region. In other words, once I'm far enough out of my sequence, all of the curves lie in this particular region.

Now, of course, the question is what does this all mean? And the answer is-- well, look. Let's take epsilon to be, and I put this in quotation marks, "very, very small." Let's take epsilon, for example, to be so small that it's within the thickness of our chalk. If I now do this-- see, I draw the curve 'y' equals 'f of x'. Notice now that the thickness of my curve itself is the band width 2 epsilon. All I'm saying is that for this very, very small epsilon, when 'n' is sufficiently large, the curve 'y' equals ''f sub n' of x' appears to lie inside of this curve.

You see, in other words, what you're saying is that for a large enough 'n' and small enough epsilon-- loosely speaking what you saying is that 'y' equals ''f sub n' of x' looks like 'y' equals 'f of x', for a sufficiently large values of 'n'. In other words, it appears that we can't really tell the n-th curve in the sequence from the limit function.

And I want to make a few key observations about what that means. I've written the whole thing out on the blackboard so that you can see this after I say it, but what I want you to see is, can you begin to get the feeling that with this kind of a condition, for example-- if each 'f sub n' happens to be continuous, in other words, if each member of my sequence is unbroken, then the limit function itself must also be unbroken. Because you see I can squeeze this thing down to such a narrow width that there's no room for a break in here.

Also notice that if the curve 'y' equals ''f sub n' of x' is caught inside this curve-- if I, for example, were computing the area of the region 'R'-- for a large enough 'n', I couldn't tell the difference in area if I use 'y' equals 'f of x' for my top curve, or whether I use 'y' equals ''f sub n' of x'. Well, keep that in mind, and all I'm saying is this, that from our picture, it should seem clear that if the sequence 'f sub n' converges uniformly to 'f' on [a, b], and if each 'f sub n' is continuous on [a,b], then-- and this is a fundamental result. Fundamental result one, 'f' is also continuous on [a, b].

Now you say, look. That's what you'd expect, isn't it, if every member of the sequence is continuous? Why shouldn't the limit function also be continuous? The point is, well, maybe that's what you expect, but note this-- in one of our earlier lectures, we already saw that if we only had point-wise convergence, this did not need to be true. In particular, recall our example in which we defined ''f sub n' of x' to be 'x to the n', where the domain of 'f' was the closed interval from 0 to 1. Remember what happened in that case? Each of these 'f sub n's is continuous. But the limit function, you may recall, is what? It's 0 if 'x' is less than 1, and 1 if 'x' equals 1. In other words, the limit function was discontinuous at 'x' equals 1.

By the way, you might like to see what this means from another point of view. And let me show you what it does mean from another point of view. Since 'f' continuous at 'x' equals 'x sub 1' means that the limit of 'f of x' as 'x' approaches 'x1' is 'f of x1', and since 'f of x' itself, by definition, is the limit of ''f sub n' of x' as 'n' approaches infinity, we may rewrite our first condition in the form-- what? We may rewrite our first equation, this equation here, in what form? Limit as 'x' approaches 'x1' of limit 'n' approaches infinity, ''f sub n' of x'. And that equals the limit as 'n' approaches infinity, ''f sub n' of x1'.

Now keep in mind that since each 'f sub n' is given to be continuous, by definition of continuity ''f sub n' of x1' is the same as saying the limit as 'x' approaches 'x1', ''f sub n' of x'. The point I'm making is, if you now put this together with this, it says the rather remarkable thing that unless you have uniform convergence when you interchange the order in which you take the limits over here, you may very well get a different answer.

In other words, when you're dealing with non-uniform convergence, you must be very, very careful to perform every operation in the given order. What we're saying is what? This will give you an answer. This will give you an answer. But if the convergence is not uniform, the answers may be different, and consequently by changing the order you destroy the whole physical meaning of the problem.

Well, again, that's reemphasized in the supplementary notes. Let me continue on here. Let me tell you another interesting property of uniform convergence. Suppose the sequence ''f sub n' of x' converges uniformly to 'f of x' on [a, b]. The point that's rather interesting is that you can reverse the order of integration and taking the limit in this particular case. In other words, suppose you want to compute the integral of the limit function from 'a' to 'b'. What you can do instead is compute the integral of the n-th number of your sequence, and then take the limit as 'n' goes to infinity. In other words, rewriting this, it says that if you have uniform convergence, you can take the limit inside the integral sign.

And again, these results are proven in our supplementary notes. We also say a few words about corresponding results for differentiation in our supplementary notes. And I should point out that differentiation is a far more subtle thing than integration. See, remember that for integration, all you need is continuity. For differentiation, you need smoothness.

The point is that as you put a thin band around the function 'y' equals 'f of x' when you have uniform convergence, that's enough to make sure that the limit function must be continuous if each of the members in the sequence is continuous. But without going into the details of this thing, it does turn out that for the degree of smoothness that you need, these things can jump around enough so that for differentiation, we do have to be a little bit more careful. Rather than to becloud the issue, I will stick with integration topics for our lecture today.

Now, the other thing that I want to mention is, again, in terms of doing what comes naturally, I think we are tempted to look at something like this and say, look, all I did was bring the limit inside. Why can't I do that? And instead of saying, look, you can't, I think the best thing to show is that when you don't have uniform convergence, you get two different answers. Again, the idea being what? We are not saying that you can't do this. We are not saying that you can't compute this. All we're saying is that if the convergence is not uniform, these two expressions may very well name different numbers.

Now, to show you what I have in mind here, let me give you an example. See, to show you that 2 need not be true if the convergence is not uniform, consider the following example. Now, in the supplementary notes, I repeat this example both the way I have it on the board and also from an algebraic point of view, without using the pictures. But in terms of the picture, here's what we do.

We define a function on the closed interval from 0 to 2, which I'll call 'f sub n', as follows. For a given 'n', I will locate the points '1/n' and '2/n'. For example, if 'n' happened to be 50, this would be 1/50 and this would be 2/50. Now what I do, is at the 'x' value '1/n', I take as the corresponding y-value, 'n squared'. And I draw the straight line that goes from the origin to this point, '1/n' comma 'n squared'. Then I draw the straight line that comes right back to the x-intercept, '2/n', and I finish off the curve by just letting it hug the x-axis till we get over to 'x' equals 2.

I'll come back to this on the next board, to show you why I chose this, but let's make a few observations just to make sure that you understand what this function looks like. I'll just make a few arbitrary remarks about it. First of all, for each 'n', ''f sub n' of '1/n'' is 'n squared'. That's just another way of indicating a label for this point.

Secondly, my claim is that for any number 'x sub 0', if 'x sub 0' is greater than '2/n', ''f sub n' of x0' must be 0. And the reason for that is quite simple. I'm just trying to show you how to read this picture. Namely, notice that as soon as 'x' gets to be as great as '2/n', the 'f' value is 0, because the function is hugging the x-axis. And just as a final observation, notice that when 'x sub 0' is 0, ''f sub n' of 0' is 0 for every 'n', meaning that every member of my family of functions goes through this particular point. In other words, let me just label this. This is 'y' equals ''f sub n' of x'.

Well, by the way if ''f sub n' of 0' is 0 for each 'n', in particular the limit of ''f sub n' of 0' as 'n' approaches infinity is 0. What happens if we pick a non-zero value? For example, suppose I pick 'x0' to be greater than 0 but less than or equal to 2? The key point is this, that since the limit of '2/n' as 'n' approaches infinity is 0, given a value of 'x0' which is not 0, I can find the capital 'N' such that when 'n' is greater than capital 'N', '2/n' is less than 'x0'. In other words, if 'x0' is greater than 0, and '2/n' approaches 0, for large enough values of 'n', '2/n' be less than 'x0'.

In particular, when that happens, if we couple this with our earlier observation-- what earlier observation? Well, this one. If we couple that with our earlier observation, we see that when 'n' is greater than capital 'N', ''f sub n' of x0' is 0. Correspondingly, then, the limit of ''f sub n' of x0' as 'n' approaches infinity, by definition, is 0. In other words-- I'm going to reinforce this later, but notice that the limit function here is the function which is identically 0.

Now, since this may look very abstract to you, let me take a few minutes-- and I hope this doesn't insult your intelligence, but let me just take a few minutes and redraw this for a couple of different values of 'n', just so that you can see what's starting to happen here. Keep that picture in mind, and now look what this means. For example, when 'n' is 1, '1/n' is 1, '2/n' is 2, 'n squared' is 1. In other words, the graph 'y' equals 'f1 of x' is just this triangular-- just this. Why give it a name?

Well, let's try a tougher one. Let's see what the member 'f sub 20' looks like. Recall how you draw this, now. With the subscript 20, what do you do? You come in to the point 1/20, and at that point, you do what? You locate the point 1/20 comma 'n squared'. In this case, it's 400. And I have obviously haven't drawn this to scale, but you now do what? Draw the straight line that goes from the origin to this point. Then from this point, you draw the straight line that comes back to the x-axis, hitting it at 'x' equals 1/10. And then you come across the x-axis all the way to 'x' equals 2. This would be the graph of 'y' equals ''f sub 20' of x'.

And by the way, do you sense what's happening over here? See, notice that as 'n' gets very, very large, the curve hugs the x-axis, starting in closer and closer to the y-axis. But what happens is someplace in here, no matter how close 'x sub 0' is to 0, there comes a very high peak. In fact, what is that high peak? It's 'n squared'. In other words, when this number is very close to the y-axis, the peak is very, very high. In other words, no matter how you put the squeeze on over here, this particular peak jumps out. This is why this particular sequence of functions is not uniformly convergent.

Again, this is done more slowly in the notes. But at any rate, let me show you an interesting thing that happens over here. Let me redraw this now for a general 'n'. In other words, let me draw 'y' equals ''f sub n' of x' for any old 'n'. Recall what our definition was, now, especially with this as review. We locate the point '1/n' comma 'n squared'. We then draw the line that goes from the origin to that point. Then we draw the line that goes from that point back to the x-axis at the point '2/n'. And then we come across to 'x' equals 2.

Let's try to visualize what the integral from 0 to 2, ''f sub n' of x', 'dx', means. After all, in a case of a continuous curve, which this is, isn't the definite integral interpreted just as the area under the curve? Well, you see, the curve coincides with the x-axis from '2/n' on to 2. Consequently, this triangular region which I call 'R' is the area under the curve. In other words, the integral from 0 to 2, ''f sub n' of x', 'dx', is the area of the region 'R'.

But here's the point. We can compute the area of the region 'R' very easily. It's a triangle, right? What is the area of a triangle? Well, it's 1/2 times the base-- but the base is just '2/n'-- times the height. The height is 'n squared'. In other words, the area of the region 'R' simply is 'n'. And that's rather interesting. In other words, for each 'n', this particular integral just turns out to be 'n' itself. That's what's interesting about this particular diagram. In other words, this thing rises so high that even though the base gets very, very small as 'n' gets large, the height increases so rapidly that the area under this curve, numerically, is always equal to 'n'.

In fact, we can check that if you'd like. Come back to this particular case. Look at this particular triangle. The base is 2, the height is 1. The area is 1 unit. Look at this particular triangle. The base is 1/10, the height is 400. 400 times 1/10 is 40, and half of that is 20. The area of this triangle is 20, which exactly matches this subscript. That's what's going to happen here all the time. In particular, then, if we compute the integral from 0 to 2, 'f of n', 'x dx', and then let the limit as 'n' goes to infinity be computed, what do we get for an answer? We get that this limit is the limit of 'n' as 'n' approaches infinity, and that of course is infinity.

On the other hand, suppose we bring the limit inside? In other words, suppose we compute this. Well, the point is that we have already shown that this is identically 0 for all 'x'. Consequently, this integral is the integral from 0 to 2, 0 'dx', which is 0. In other words, if you first take the limit and then integrate, you get 0. On the other hand, if you first integrate and then take the limit, you get infinity. And this should be a glaring example to show you that the answer that you get indeed does depend on the order in which you do the operation.

Again, let me emphasize-- which of these two is wrong? The answer is, neither is wrong. All we're saying is that if you were supposed to solve this problem and by mistake you solve this one, you are going to get a drastically different answer.

OK. Let's not beat this to death. So far, so good. Let me make one more remark, namely, what does all of this have to do with the study of series? See, now we're just talking about sequences of functions. And you see, the answer to this question is essentially going to be our last lecture of the course. But for now, what I'd like to do is to give you a preview of that. Namely, the application of uniform convergence to series is the following.

Recall that when we write summation 'n' goes from 0 to infinity, 'a sub n', 'x to the n', that's an abbreviation for what? A polynomial, 'k' goes from 0 to 'n', 'a sub k', 'x sub k', as 'n' goes to infinity. In other words, recall that the sum of the series is a limit of a sequence of partial sums, and this is the n-th member of that sequence of partial sums. Again, if the sigma notation is throwing you off, all I'm saying is to observe that 'a0' plus 'a1 x' plus 'a2 'x squared'' plus-- et cetera, et cetera, et cetera, forever, just represents the limit of the following sequence. 'a0', next member is 'a0 plus a1 x'. Next member is 'a0' plus 'a1 x' plus 'a2 'x squared''. The next member is 'a0' plus 'a1 x' plus 'a2 'x squared'' plus 'a3 'x cubed''.

By the way, what is each member of the sequence? It's a polynomial. And polynomials have very nice properties, among which are what? Well, a polynomial is a continuous function. A polynomial is an integral function, et cetera. The idea, therefore, is that if this sequence of partial sums converges uniformly to the limit function, then, for example, the limit function, namely the power series, must be continuous since each partial sum that makes up the sequence of partial sums is also continuous. Namely, every polynomial is continuous.

Also, if, for some reason or other, you want to integrate that power series from 'a' to 'b', if the convergence is uniform, I can then do what? I can then take the summation sign outside, integrate the n-th partial sum, and add these all up. You see, the idea being what? That the n-th partial sum is a polynomial, and a polynomial is a particularly simple thing to integrate. That's one of the easiest functions to integrate, in fact.

OK. Now, here's the wrap up, then. What we shall show next time is that within the interval of absolute convergence, the sequence of partial sums, which we already know converges absolutely to the limit function, also converges uniformly. In other words, within the radius of convergence, the power series-- and I don't know how to say this other than to say, it enjoys the usual polynomial properties associated with a polynomial such as summation 'k' goes from 0 to 'n', 'a sub k', 'x to the k'.

In other words, then, this about wraps up what our introductory lecture for today wanted to be, namely the concept of uniform convergence. What I would like you to do now is to study this material very carefully in the supplementary notes, go over the learning exercises so that you become familiar with this. Then we will wrap up our course in our next lecture, when we show what a very, very powerful tool this particular concept of absolute convergence is in the study of the mathematical concept of convergence.

At any rate, until next time, then, goodbye.

Funding for the publication of this video was provided by the Gabriella and Paul Rosenbaum Foundation. Help OCW continue to provide free and open access to MIT courses by making a donation at ocw.mit.edu/donate.