STOCHASTIC APPROXIMATION: THEORY AND APPLICATIONS

Dr. Gugan Thope

Department of Computer Science and Automation

Indian Institute of Science

Week 3

Lecture 13

Doob's Forward Convergence Theorem for Super-martingales

Hello and Namaste, everyone. Welcome to Lecture 13 of this NPTEL course on Stochastic Approximation. So, let us do a quick recap of what we did in the previous two lectures. In the last but one lecture, we looked at martingales and some examples of martingales, and also looked at the connection between martingales and martingale difference sequences. In the previous lecture, we looked at two specific things.

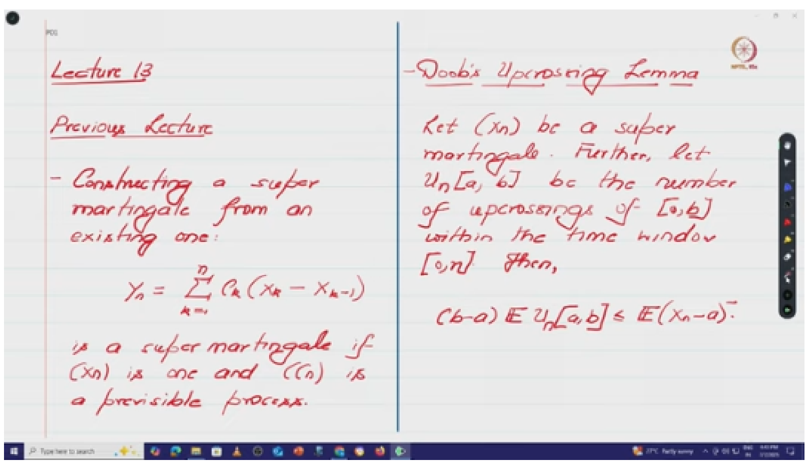

One of them was constructing a supermartingale from an existing one. In particular, we said that suppose Xn is a supermartingale and Cn is a pre-visible process, then the process Yn, which is defined in this way, is also a supermartingale.

$$Y_n = \sum_{k=1}^{n} C_k\left(X_k - X_{k-1}\right)$$

Apart from this, we studied something called Doob's upcrossing lemma. It is one of the fundamental results that is used toward proving the convergence of martingales. So, this result stated something like this.

It said that suppose Xn is a supermartingale, further suppose this variable un subscript ab. be the number of upcrossings of AB within the time window 0 to n, right? Then Doob's upcrossing lemma said that this inequality holds.

$$(b - a)EU_n[a, b] \leq E\left(X_n - a\right)^-$$

Okay, so what does this inequality say? b minus a times the expected value of un ab is less than or equal to the expected value of xn minus a minus. Okay, so as I said in the previous lecture, what is nice about this result is that, UN of AB, in some sense, tells you how many times the process XN did an upcrossing of the interval AB between the time window 0 to N.

that in some sense, UN of AB depends on what happened between the time, you know, what happened during the time interval 0 to N. Whereas, the right-hand side only concerns the value of the stochastic process Xn at time n. So, in other words, Doob's upcrossing lemma connects what happened in the time window 0 to n to the value of the process at time n, right? That is what is nice about this result. And these are the two things that we discussed in the previous class. In today's class, we will look at convergence of super martingales.

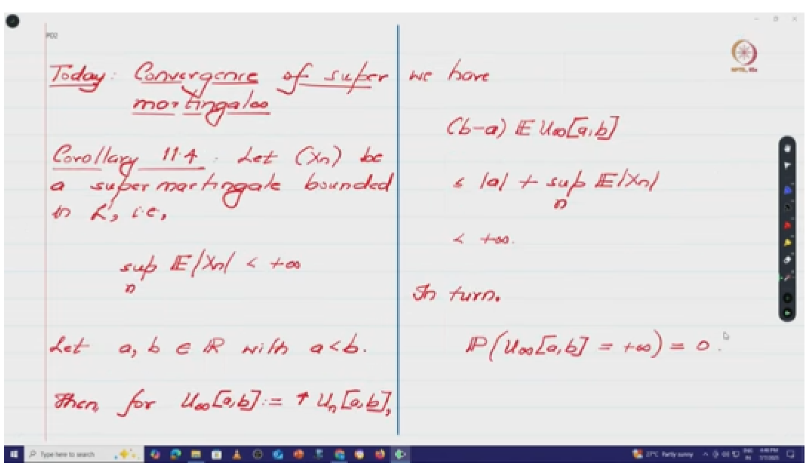

by building upon these results that we saw previously. So, towards this one of the first things that we will be studying is this corollary 11.4 again this number corresponds to a result in this textbook called as probability with martingales by David Williams. So, in case you want more details you can actually refer to this textbook and look up this result over there. So, what this result says in particular this result says something about the number of up crossings of a super martingale. So, this result says that suppose Xn is a super martingale that is bounded in L1.

So, there is a L superscript 1 over here if it is not clearly visible. So, we have a super martingale that is bounded in L1. So, when I say we have a super martingale bounded in L1, I would imply that a condition like this holds. So, let us look at this condition. If we just had expected value of absolute value of Xn is less than infinity then we will use that to conclude that your you know Xn is integrable.

Here we are saying something more: we are saying that the supremum of the expected value of Xn is less than infinity. That is, we are not just saying that every Xn is integrable, we are also saying—and importantly saying—that the supremum of these expectations is finite.

$$E|X_n| < +\infty$$

So, let us presume that such a condition holds for the given supermartingale. Furthermore, let us presume that there are two real numbers A and B such that A is strictly less than B. And let us define U infinity of AB to be the limit of UN of AB.

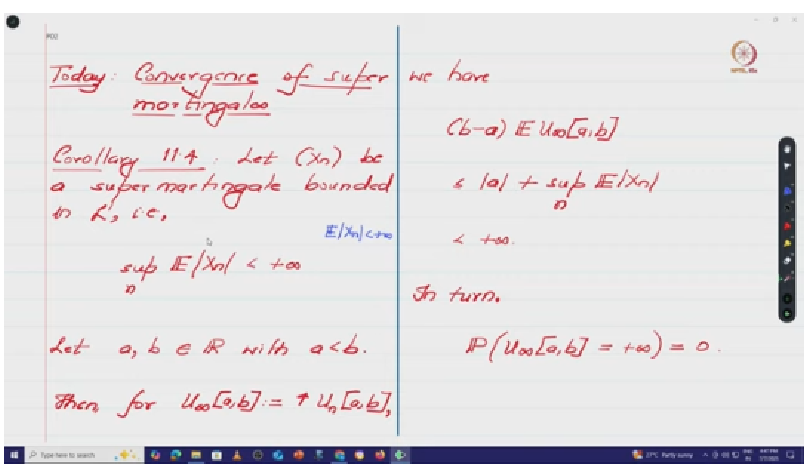

So, let us decipher this notation a bit. You can see that there is an up arrow here, right? So, what does this up arrow mean? Okay, so what does this up arrow mean? Well, UN of AB, in some sense, looks at the number of up crossings in the time window 0 to n, and it is easy to see that UN of AB will be less than or equal to UN plus 1 of AB, in particular for any omega.

In capital omega, UN of AB of omega is some integer, and UN plus 1 of AB of omega is another integer, right? And what we are saying is that the number of up crossings in the time window 0 to n will be less than the number of up crossings of the interval AB in the time window 0 to n plus 1. So this is obvious to see. This implies that your UN of AB, as a function of n, is actually a monotonically increasing sequence. And because it's a monotonically increasing sequence, the limit actually exists for every little omega, and hence there is a limit in terms of the random variable itself. Let's denote the limiting random variable by U infinity of AB, right? So U infinity of AB—I hope you are able to see—is again some function f omega to R. So, what does this result say?

This result says that suppose you have a supermartingale which is bounded in L1, and you have real numbers a and b such that a is less than b. Then your u infinity random variable satisfies the following relation. In particular, b minus a times the expected value of u infinity ab is less than the absolute value of a plus the supremum. of the expected value of the absolute value of xn, where the supremum is taken over the index n. And we have been given that this supremum is less than infinity. Furthermore, we have been given that a is a real number, which implies the absolute value of a is finite. And putting these two facts together, we can see that this plus this will also be finite, right? So what this result says is that if you have a supermartingale that is bounded in L1, then

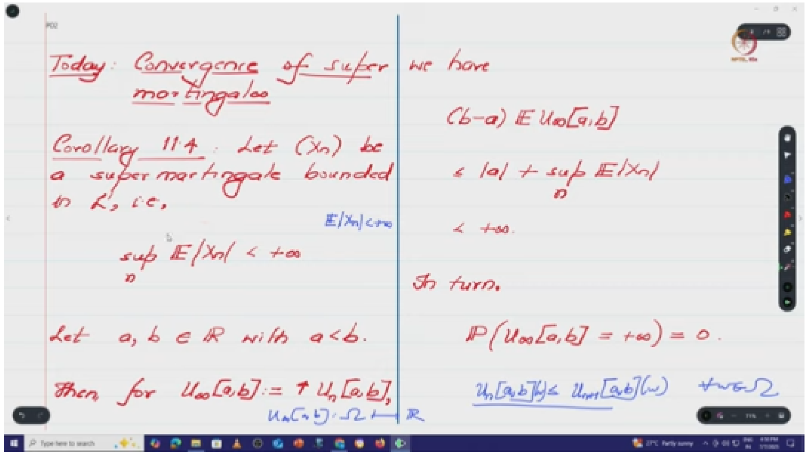

B minus A times the expected value of U infinity of AB is less than infinity.

$$(b - a)EU_{\infty}[a, b] \leq |a| + E\left|X_n\right| < +\infty$$

So, let us try to understand what U infinity of AB is. Well, as you may have already guessed, it is the number of upcrossings of the interval AB over the window starting from 0 and going all the way up to infinity. So, that is what U infinity of AB is, and this result says that the expectation of this random variable, which counts how many upcrossings you have over the entire 0 to infinity window, is actually less than infinity. And because this is a non-negative random variable and B minus A is finite, your expectation is finite. And because this is a non-negative random variable, from the definition of expectations, one can easily see that this fact implies that the probability that U infinity AB takes the value plus infinity is 0. In other words, your U infinity of AB

is less than infinity almost surely. That is what we can conclude from this fact. Okay, so we are now going to prove this result.

$$P\left(U_{\infty}[a, b] = +\infty\right) = 0$$

Okay, so let us go about proving this result. From the result that we saw in the previous class, in particular from Doob's upcrossing lemma, we can conclude that b minus a times the expected value of unab is less than or equal to the expected value of xn minus a minus. And separately, one can see that Xn minus A minus is trivially less than the absolute value of Xn plus the absolute value of A. And if you are a listener hearing this lecture, please figure out why this inequality holds. And the hint is that it is a consequence of the triangle inequality. And since you have the absolute value of xn here, we can take the supremum over all possible values of n. And to distinguish between this n, I have replaced the index as k. So, this absolute value is less than this absolute value.

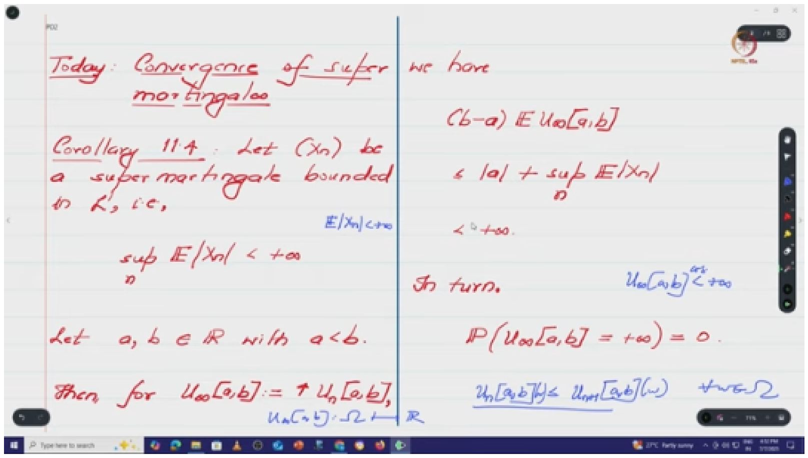

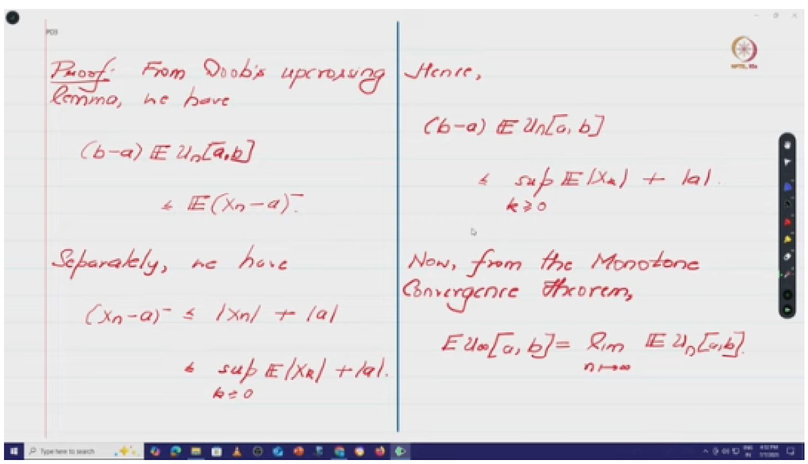

So, there is a typo over here; this expectation should not be there. Right, so xn minus a minus is less than the absolute value of xn, and that is less than the supremum over k of the absolute value of xk plus the absolute value of k, right? And because you have this inequality over here, right? Um, you have this inequality over here, right? We can take expectations on both sides to conclude that b minus a times the expected value of un ab is less than the supremum of k greater than or equal to 0, the expected value of xk plus a. So, there is a small typo that I made. Actually, what I should say is this inequality over here actually implies the expected value of xn minus a minus. is less than or equal to the expected value of Xn plus the absolute value of A, and then one can actually conclude that this is less than or equal to this, right? And from this, one can conclude that because this relationship holds, B minus A times the expected value of Un AB is less than or equal to the supremum of the expected value of the absolute value of Xk plus A.

Now, because this relationship holds and U infinity is a monotone limit of Un's, we can invoke the monotone convergence theorem on the one hand to conclude the expected value of U infinity is the limit of these expectations.

$$(b - a)EU_n[a, b] \leq E\left(X_n - a\right)^-$$

$$\left(X_n - a\right)^- \leq \left|X_n\right| + |a| \leq E\left|X_k\right| + |a|$$

$$(b - a)EU_n[a, b] \leq E\left|X_k\right| + |a|$$

$$EU_\infty[a, b] = E\ U_n[a, b]$$

And because each of these expectations is upper bounded by this, we can then conclude that B minus A times the expected value of U infinity of AB. This is less than this, and because of the given condition that your Xn supermartingale is bounded in L1, one can conclude that this is less than infinity to finally infer that B minus A times the expected value of U infinity AB is less than infinity, as desired. So, recall that we wanted to show this part, and we have shown that part. So, the key idea is to use the monotone convergence theorem to relate the expected value of U infinity to the expected value of Un and connect the expected value of Un to the supremum of these expected values of absolute values of xk.

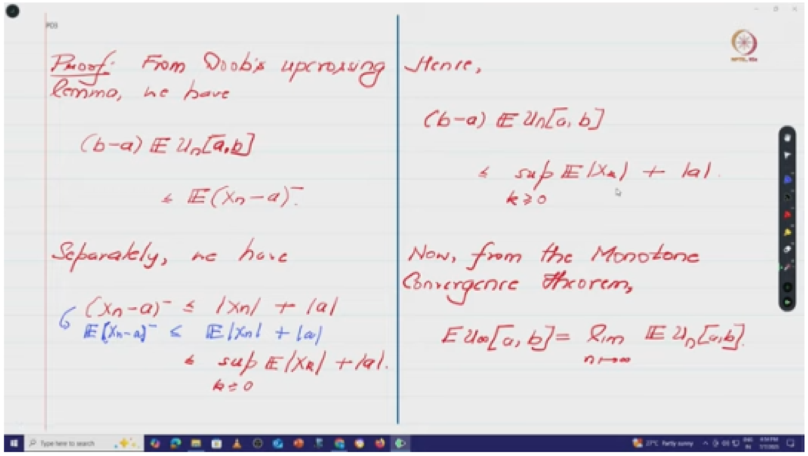

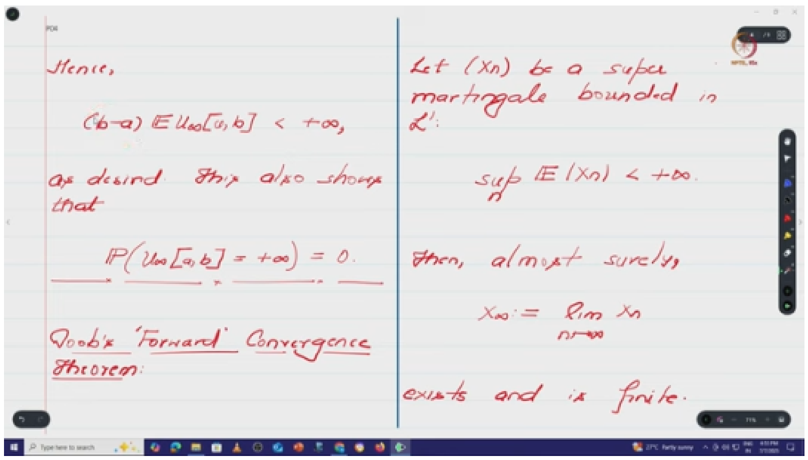

So, we can conclude this, and then, as I said, since this is a non-negative random variable and its expectation is finite, we can conclude that the probability that U infinity of AB equals plus infinity is actually 0. So, what does this mean?

$$(b - a)EU_{\infty}[a, b] < +\infty,$$

$$P\left(U_{\infty}[a, b] = +\infty\right) = 0$$

So, let us try to, you know, give a physical interpretation of this statement. So, it means that if you have a supermartingale that is bounded in L1, so the supremum of these expected values of the absolute value of Xn is less than infinity. So, if that is the case, then

You know the number of up-crossings of this interval [a, b] in the infinite time window from 0 to infinity. Okay, the number of up-crossings being infinity has 0 probability, which means that with probability 1, the number of up-crossings will be a finite number. Okay, so of course, on different sample paths or different sample points Omega, the number of up-crossings could either be 100, 500, 5 trillion, or any big number. However, what this result says is that, you know, with probability 1, this number of up-crossings will be a finite number, right? So, this finishes one of the key results that we would be needing to discuss the convergence of supermartingales. We will now use this result to prove the convergence of supermartingales.

In particular, this result is known as Doob's forward convergence theorem, and the 'forward' over here refers to the fact that, you know, we are looking at what happens as the time index n is taken to infinity in the forward sense. So, that is what 'forward' over here means. So, what does Doob's forward convergence theorem state? It states that suppose you have a supermartingale which is bounded in L1, so that you know the previous corollary that we discussed applies. So, bounded in L1 means this condition.

Then Doob's forward convergence theorem says that almost surely, the limit of the supermartingale exists and is finite. Is this okay? So, almost surely, the limit exists and is finite. So, this is what the convergence theorem is. So, what it means is that if you have a supermartingale that is bounded in L1, then with probability 1, if you look at the limit of your supermartingale, the limit will be a finite number with probability 1.

So, again, you know this notation over here means that your x infinity of omega is defined to be the limit as n tends to infinity of xn of omega. So, that is what this definition means. Alright, so now let us see how the previous corollary can be used to prove this result, and this result is actually a very clever application of real analysis. Hopefully, you will enjoy it the way I also enjoyed it when I read or saw it for the first time. So, we want to show that almost surely the limit exists and is finite. We will break the proof into two parts: first, we will show that the limit exists, and then we will show that the limit is actually a finite number. To show that the limit exists, what we will do is first define a capital lambda set. This is an event which consists of all those small omegas for which your sequence Xn does not have a limit. When I say it does not have a limit, it does not have a limit in the extended interval from minus infinity to plus infinity, which means we are also not claiming that your Xn could actually go to plus infinity. So, that convergence is also presumed not to hold.

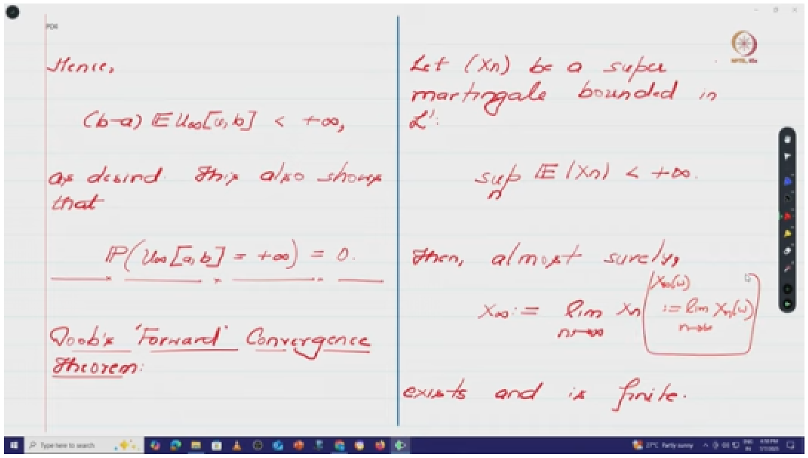

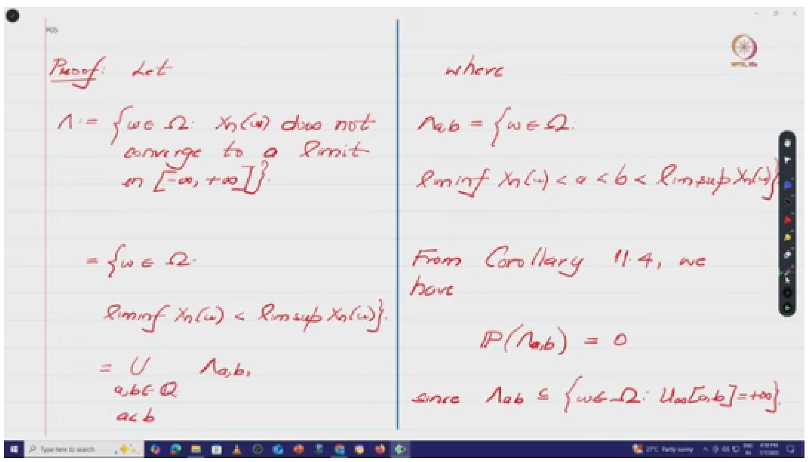

So, this event is very easy to see from some basics of real analysis that this limit is equivalent to this limit. So, it basically says that lambda is the collection of all little omegas such that the lim inf is strictly less than the lim sup. So, those who have seen limits of real numbers or real sequences know that the limit is said to exist only when the lim inf is equal to the lim sup. We also know that the lim inf is always less than or equal to the lim sup, and hence when we say the limit does not exist, it is equivalent to saying that the lim inf is strictly less than the lim sup. And what we are going to do now is break down this event into a union of some very interesting events.

So, we will break down this event as a union of events which are indexed by a and b, where a and b are some rational numbers with the property that a is less than b. So, what is lambda AB now? Lambda AB is the collection of all those omegas where the lim inf is strictly less than A, and A is strictly less than B, and B is strictly less than the lim sup. So, let us try to understand what we are doing over here. Because your lim inf and lim sup are

Have this strict inequality relation. I hope you agree that we can squeeze in two real numbers, a and b, such that this relation holds, right? So, we are saying take all these omegas and break them up into this union so that whenever this inequality holds, right, you can find two numbers, a and b, such that this inequality holds, right? And you know you can put all of them together. And notice that I am not saying this union is a disjoint union. So, some of the omegas, you know, could overlap multiple lambda ABs. It does not matter. What matters is that this set of rational numbers is a countable set, and hence this is a countable union.

So, that is the key part. So now, from our previous corollary, we know that this random variable with probability 1, right, takes a finite value, or this random variable U infinity takes the value infinity with probability 0, right? And lambda AB, which is defined over here—my claim is that it is a subset of— So, let us try to understand why this subset relation holds. So, lambda AB says that lim inf is less than A and lim sup is bigger than B. And from the basics of real analysis, we know that a lim inf of a sequence is a point which is infinitely often visited by Xn. So, or more generally, or more specifically, any neighborhood of the lim inf of a sequence is visited infinitely often by that sequence.

And from that observation, and also a similar conclusion for lim sup—that is, any neighborhood of lim sup will also be visited infinitely often by the sequence—one can see that if you visit— You know, this neighborhood infinitely often and some neighborhood of this infinitely often, and if you choose the neighborhood so that you know that neighborhood is below A and you know the neighborhood of this lim sup is above B, one can conclude that, you know, because these neighborhoods are visited infinitely often, the up-crossings of AB will be plus infinity. So, for every omega that belongs to lambda AB, one can conclude that on those omegas, the number of up-crossings of the interval AB will be plus infinity—that is, you know, you will have infinitely many up-crossings of the interval AB, and you will have this infinitely many up-crossings because you will infinitely often visit a point below A and you will infinitely often visit a point above B. So, that is the reason, and hence one can conclude that lambda AB is a subset of this. And in the previous corollary, we have seen that the probability of, you know, this event is 0 for any AB, right? And hence, one can conclude that the probability of AB must be 0, right? And from this, and the fact that you know your set of rational numbers is countable, one can conclude that the probability of lambda itself is zero, right? So, what is lambda? Lambda is the set of all omegas where—

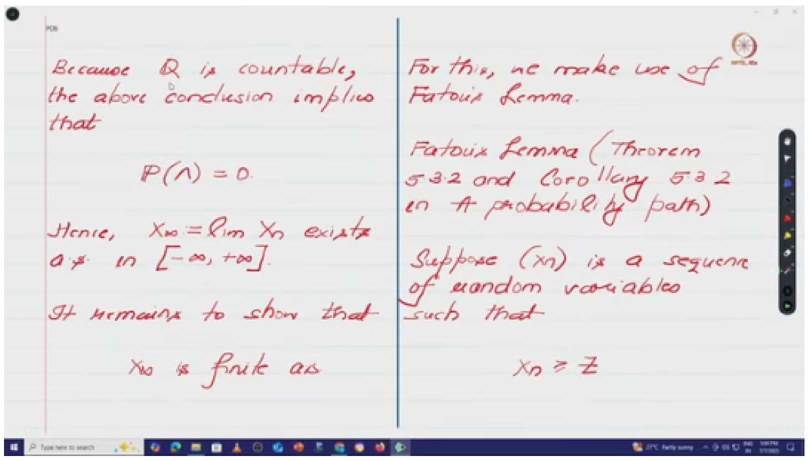

The sequence Xn does not have a limit, and what we have managed to show is that the probability of lambda is 0, which means that this event has 0 probability, and hence the complement of this event, which says that the limit exists, has probability 1. So, that is what we have written over here: you know, x infinity, which is defined as the limit of xn, exists almost surely in this extended interval from minus infinity to plus infinity. So, I am, you know, working with this extended interval. So, the extended real line basically means that you also allow, you know, x infinity to take minus infinity and plus infinity. So, we are allowing that at this point because

you could also have a scenario where, you know, your xn's take larger and larger values, and hence their limit would actually be plus infinity, right? And so on and so forth, right. So, now, in the result, though, it states that the limit exists and it is finite. So, we have now shown that the limit exists, and now it remains to show that the limit actually is finite almost surely. And in order to prove this second part, we are going to make use of another very, very important result from probability theory, which is referred to as Fatou's lemma.

So, I am going to state this Fatou's lemma, but I won't be proving it. However, I will tell you how one can go about proving this result. Those who are interested in knowing about the proof of Fatou's lemma can refer to Theorem 5.3.2 and Corollary 5.3.2 in this textbook called 'Probability Path' by Resnick. So, what does Fatou's lemma say? At a very high level, Fatou's lemma relates to the expected value of the lim inf of Xn and the expected value of the lim sup of Xn. So, let us try to understand what Fatou's lemma says. So, Fatou's lemma says that suppose Xn is some sequence.

So, for Fatou's lemma, you do not need a supermartingale or any special property of the sequence. All we require is that Xn be some sequence of random variables. With the property that all Xn's are lower bounded by Z. So, again, when I write something like this, I actually mean that Xn of omega is greater than or equal to Z of omega for all omega in capital omega. So, we require this condition. So, suppose Xn is a sequence of random variables such that Xn is greater than or equal to Z, and Z is some random variable in L1, which basically means that the expected value of Z is less than infinity.

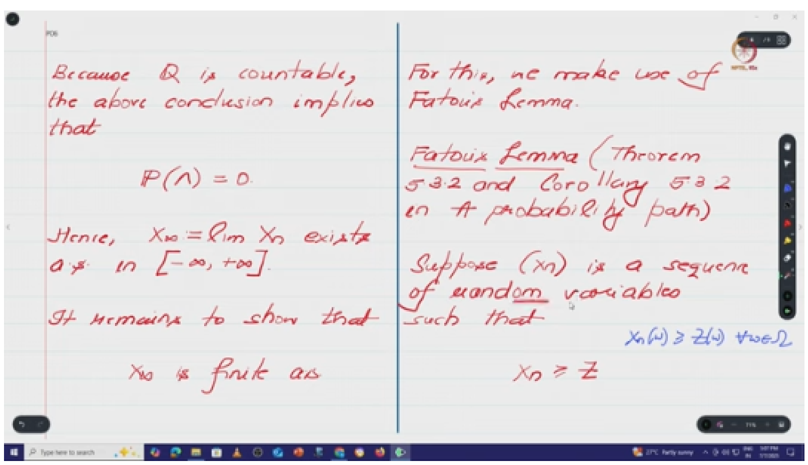

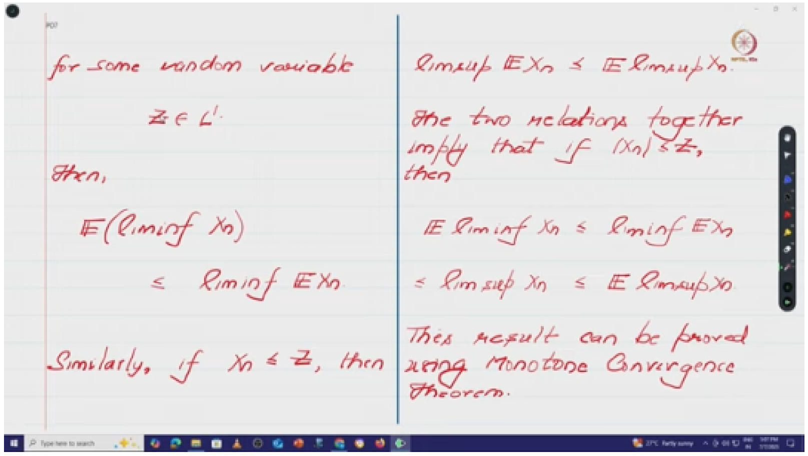

So, the setup is that we have a sequence of random variables which is lower bounded by an integrable random variable Z, and Z can be arbitrary. So, if such a setup—I mean, such a random variable exists—then one can conclude that the expected value of the liminf of Xn is less than or equal to the liminf of the expected value of Xn. So, notice that liminf over here is applied to Xn, right? So, which basically means that this is the random variable which, on a sample point omega, is expected. The liminf of Xn of omega.

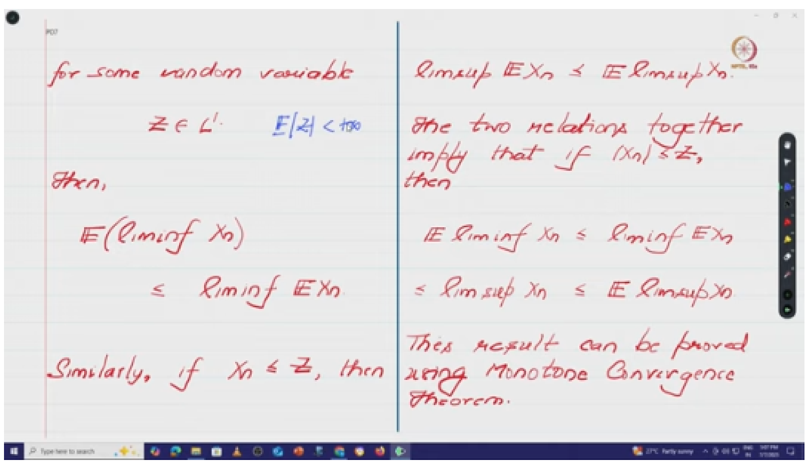

So, this random variable at omega is basically taking this value. Now, Xn of omega will be a sequence of real numbers, and liminf for a sequence of real numbers always exists. So, this liminf is always well-defined. Right, and by invoking the monotone convergence theorem, one can conclude that the expected value of the liminf also exists. And what Fatou's lemma says is that you can interchange the liminf and the expectation. So, the result says that the expected value of the liminf is upper bounded by the liminf of the expected value of Xn. So, notice that in this Fatou's lemma, Z does not appear anywhere in the result, right? So, that is very nice. In some sense, in the language of chemistry, it acts like a catalyst. Its presence is needed for this inequality to hold true, but as such, it does not play any role in the result, right.

So, this is what Fatou's lemma says on the Lim-Inf side and there is a similar outcome on the Lim-Soup side. So, the result states that suppose you have Xn less than equal to Z. So, in the Lim-Inf case notice that we wanted Xn greater than equal to Z whereas in the Lim-Soup case we require that Xn be upper bounded by Z and Z again is required to be L1 integrable. Then, one can conclude that limb soup of expected value of xn is upper bounded by expected value of limb soup of xn. And if you combine both these relations together, we can conclude that if the absolute value of xn is upper bounded by z and z is integrable, which implies that minus z is less than or equal to the absolute value of xn—sorry, minus z is less than or equal to xn. So, this thing holds, allowing us to apply both parts of Fatou's lemma.

We can conclude that the expected value of lim inf is less than the lim inf of the expected value of Xn, and for any sequence of real numbers, we know that the lim inf is less than or equal to the lim sup. From the second part of Fatou's lemma, we know that the lim sup of X—so I should say— The lim sup of the expected value of Xn—this is from the second part of Fatou's lemma. Sorry, from the second part of Fatou's lemma, we know that this inequality holds: that is, the lim sup of the expected value of Xn is less than the expected value of the lim sup of Xn, right? So, this is, in some sense, the broad conclusion of Fatou's lemma: that is, if the absolute value of Xn is less than or equal to Z for every n and Z is integrable, then— The expected value of lim inf is less than the lim inf of the expected value of Xn.

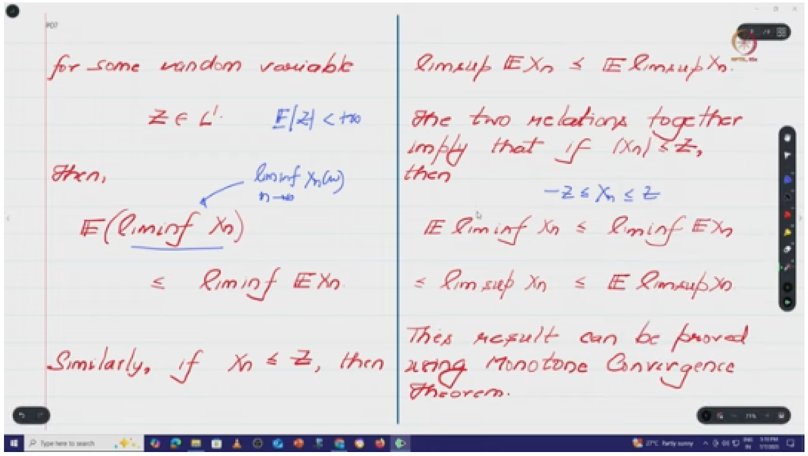

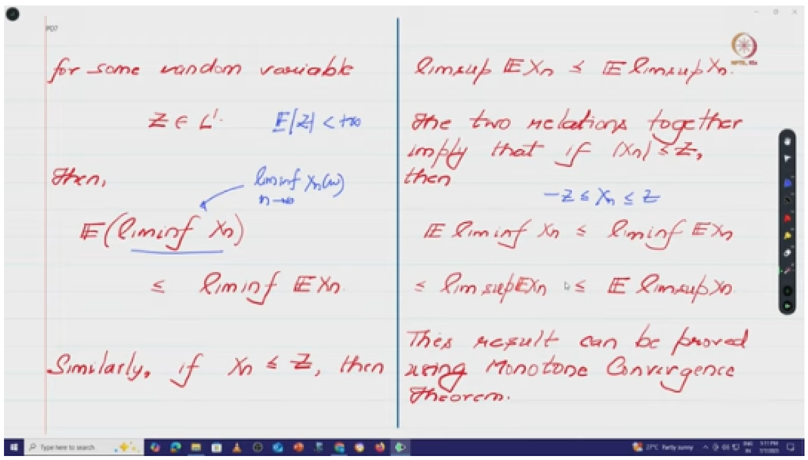

Now, this is less than or equal to this from a standard fact about sequences of real numbers, which says that the lim inf of a sequence is always less than the lim sup, and the third inequality again follows from the second part of Fatou's lemma.

$$E\left(X_n\right) \leq E\, X_n$$

$$E\ X_n \leq EX_n$$

$$EX_n \ \leq \ E\ X_n$$

$$\leq\ E\ X_n \leq EX_n$$

So, I am not going to prove this result. However, for those who are interested, I encourage you to try proving it yourself, and I am only going to say that it is a very easy consequence of the monotone convergence theorem. For those who need some additional help, I suggest that you look up this textbook by Resnick called Probability Path. So, now what remains to show is that this X infinity is a finite random variable and it takes a finite value almost surely.

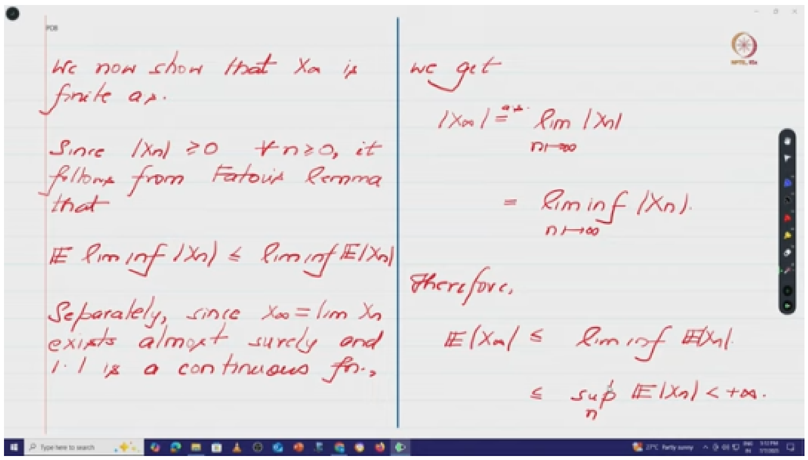

So, how can we invoke Fatou's lemma? The way to invoke Fatou's lemma is to make use of the fact that you know the absolute value of Xn is greater than or equal to 0 for all n greater than or equal to 0. So, this 0 over here is playing the role of Z, and hence from Fatou's lemma, one can conclude that the expected value of lim inf of the absolute value of Xn is less than the lim inf of the expected value of the absolute value of Xn. Is this okay? Now, in the first part of this proof, we have already shown that the limit exists right almost surely, and furthermore,

you know, we know that this absolute value function is actually a continuous function, so we can conclude using these two facts that this holds right. So, in other words, the absolute value of X infinity is equal to the absolute value of the limit as n tends to infinity of Xn, and now this just follows from the definition of X infinity. Okay, and from this, one can conclude that this is equal to this, and this follows from the continuity of the absolute value function. Is this okay? And because the limit exists, the limit of this is equal to the limit of this. Is this okay? So, one can easily show right, and one can then use this fact that the absolute value of X infinity is this right, and we know that the limit expectation is less than the lim inf of expectations to conclude that the expected value of the absolute value of X infinity is less than the lim inf of the expected value of the absolute value of Xn. So, we are just putting this over here and using these facts over here to write this left-hand side.

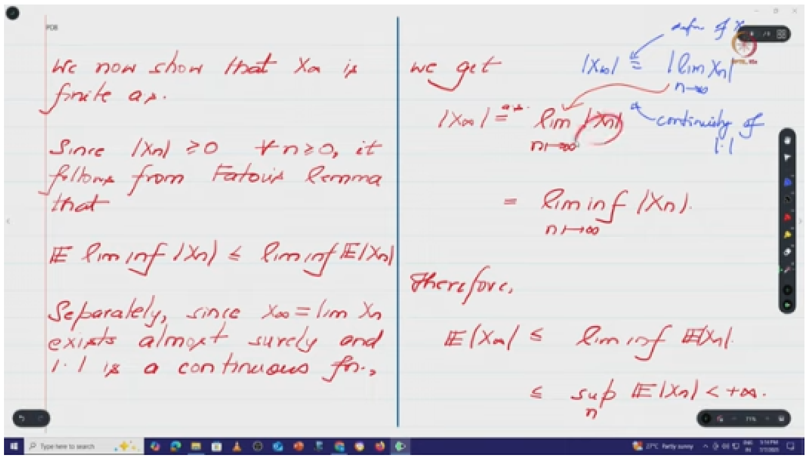

Right. And we have, I mean, one can trivially see that the lim inf of the expected value of Xn, just from the definition of lim inf, is upper bounded by the supremum of the expected value of Xn's right. And you know, the supremum of the expected value of the absolute value of Xn is presumed to be less than infinity because we have presumed that your sequence Xn is bounded in L1. And we can use all these facts to conclude that the expected value of the absolute value of X infinity is less than infinity. And we can use the definition of the absolute value of X infinity.

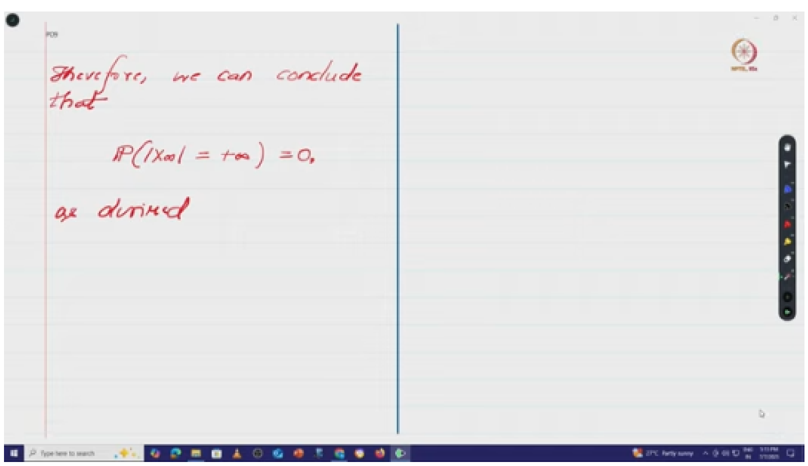

To conclude that—or, sorry—by using the definition of expectation of non-negative random variables, one can now conclude that a relation like this holds. Because if this probability had a positive value, then that would immediately imply that the expected value of the absolute value of X infinity is infinity. Now, since we have shown it is strictly less than infinity, it can only happen when this probability is 0. Now, because this probability is 0, we can conclude that the absolute value of X infinity is a finite number almost surely. So, on different omegas, X infinity can take different values.

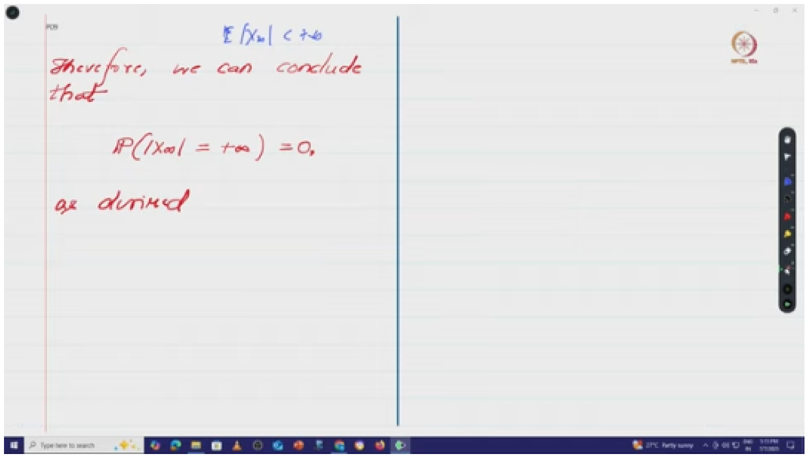

So, sometimes it can take 5, sometimes it can take minus 100, and as large a number as you can imagine. But what we are saying is that the set of omegas where the absolute value of X infinity is plus infinity has probability 0. So, with this, I come to the end of this class. Let me now summarize what we have discussed in this class. So, as I told you, in the analysis of stochastic approximation algorithms, we have to understand the behavior of noise. And in the first week—during the discussion of the first week—I told you that the noise in stochastic approximation algorithms is typically viewed as a martingale difference sequence.

And a martingale difference sequence is, in some sense, a difference of a martingale sequence. And we will later see that we will use the convergence of martingale processes to show that the cumulative effect of noise in a stochastic approximation algorithm is negligible in an asymptotic sense. So, that is the relation between what we are studying—or what we studied today—and how we will be using them in the study or analysis of stochastic approximation algorithms. So, let me quickly summarize what we did today. In today's class, we looked at the convergence of supermartingales.

Since a martingale is a special case of a supermartingale, whatever we have studied today also applies to martingales. And we showed that one of the sufficient conditions under which a supermartingale converges is that the supermartingale is bounded in L1. That is, the supremum of the expected value of the absolute value of Xn is less than infinity. And the key ingredient that was used to show this convergence result was Doob's upcrossing lemma. With this, let me stop.

Thank you. Namaste. I hope to see you in the next class.