STOCHASTIC APPROXIMATION: THEORY AND APPLICATIONS

Dr. Gugan Thope

Department of Computer Science and Automation

Indian Institute of Science

Week 6

Lecture 24

Proof of the Key Lemma - Part I

Hello and Namaste everyone, welcome to Lecture 24 of this NPTEL course on Stochastic Approximation. Let us recall what we have been doing so far. We have been trying to show that the linearly interpolated solution trajectory obtained from your stochastic approximation iterates closely tracks a suitable solution trajectory of your limiting ODE, and towards that, we said that there is a formal lemma, and that formal lemma says that the distance between these two trajectories over a window of length t goes to 0. So, we are proving this result, and we started a part of this result by showing that the cumulative martingale difference noise terms can be viewed as a martingale. Today, as a first part, we will be showing the convergence of this martingale sequence, and later we will be seeing how this part can be used to study the asymptotic behavior of this linearly interpolated trajectory.

So, let us formally begin. So, recall that this is the lemma that we are working with, that is, under assumptions A1 to A4, the following two statements hold almost surely, which is that for any t bigger than 0, if you look at the linearly interpolated trajectory and a suitable solution of your limiting ODE, then the distance between them over a window of length t starting from s goes to 0 as the starting point goes to infinity. And as a first step towards proving this item 1, we define zeta n to be the cumulative noise term. And in the previous class, we showed that this sequence zeta n is actually a martingale sequence with respect to your filtration Fn, where Fn is the sigma field generated by X0, M1 all the way up till Mn. Okay.

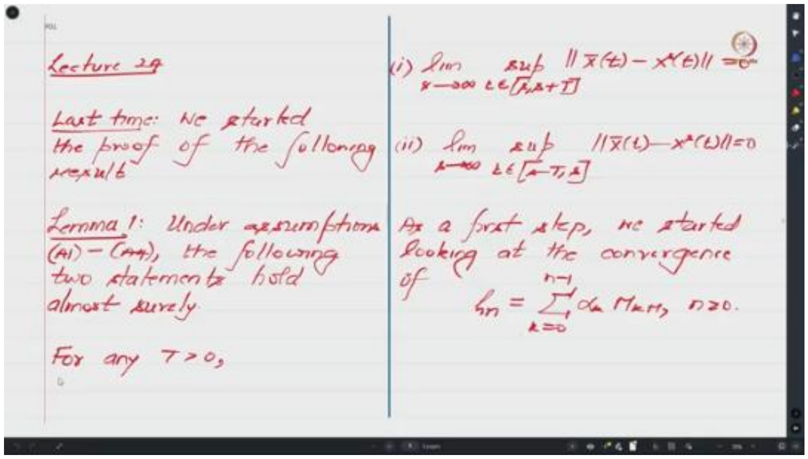

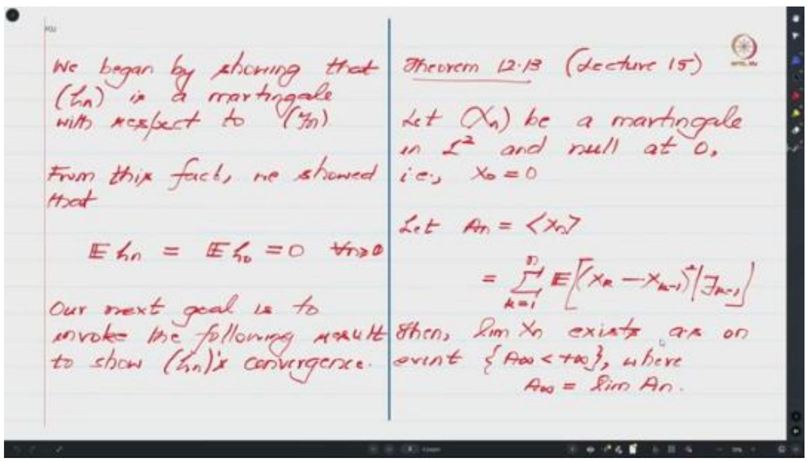

And using this fact, we managed to show that the expected value of zeta n equals the expected value of zeta 0, which is 0 for all n greater than or equal to 0.

$$\mathbb{E}\zeta_n = \mathbb{E}\zeta_0 = 0 \quad \forall \; n \geq 0$$

So, now our next goal is to show that zeta n converges, and for that, we will recall one of the results that we had discussed in Lecture 15 concerning the convergence of martingales. So, let us recall that result. This result is number 12.13 from this textbook called Probability with Martingales by David Williams, and the result goes something like this. Let Xn be a martingale in L2 and null at 0, that is, X0 is 0, and let An be the angle bracket process, that is, it is the sum of the expectation of the square of the successive differences conditioned on Fk minus 1.

So, you look at this sum of these condition squares, right? And you add it up, and that is what is referred to as the angle bracket process, right? Keep in mind that this angle bracket process is actually pre-visible. By that, I mean your a_n, right, is actually measurable with respect to f_n minus 1. So, notice that there is a minus 1 over here. In that sense, a_n is pre-visible, right? And then the result states that the limit x_n exists almost surely and is finite on the event a_infinity less than or equal to infinity, right? Where a_infinity is the limit of these random variables a_n. So, this is what we had shown. And, you know, this limit exists because this a_n is almost surely a monotonically increasing sequence. And because it is an increasing sequence, one can conclude that the limit of these random variables exists, which we denote as a_infinity.

$$A_n = < X_n >$$

$$= \sum_{k=1}^{n} \mathbb{E}[(X_k - X_{k-1})^2 | F_{k-1}]$$

$$A_\infty = \lim A_n$$

And then we collect all those sample points where a_infinity is less than infinity, and the claim is that on this—almost surely on this collection of sample points—which means that except for a subset of measure 0, on all the remaining points, your limit X_n exists and is finite.

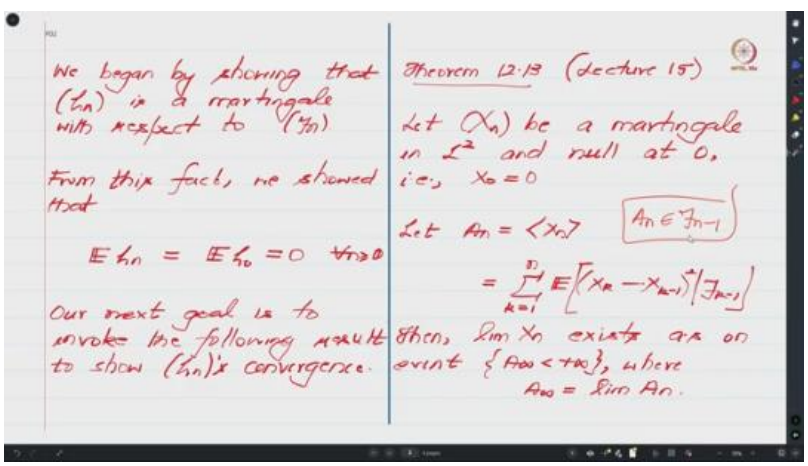

So, now the question is how we can use this result to show that your martingale sequence Zeta_n converges. The rest of our discussion would be to verify how all these conditions actually hold. So, as I said, we will now be verifying the desired conditions. The first of these conditions that we need to verify is to show that Zeta_n is a martingale sequence in L2. In other words, we need to show that the second moment of every Zeta_n is finite. More formally, we will now try to show that the expected value of the norm of Zeta_n squared is finite for all n greater than or equal to 0.

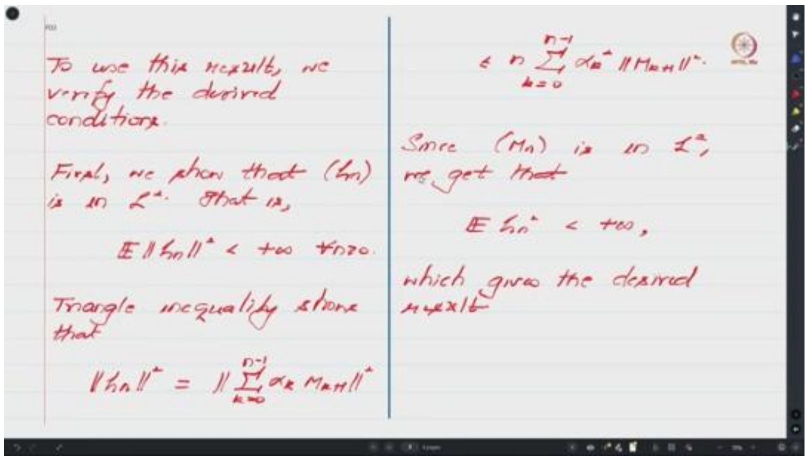

So, notice, keep in mind that you know, this bound can change with n. We only require that for every n, this bound be finite. Is this clear? So, towards that, the triangle inequality immediately shows that if you take Zeta_n squared—the square of the norm of Zeta_n—then this expression is upper bounded by n times the sum of squares of the individual terms, and the individual terms are alpha_k and m_k plus 1. So, if you take the norms and square them, we would end up with alpha_k squared and the norm of m_k plus 1 squared.

Now, under A3, we have presumed that your Mn is actually a sequence in L2, which means that the expected value of mk plus 1 squared is less than infinity for every k. And from this, one can conclude that the expected value of this expression is less than infinity for every n. Right, and from that, one can conclude that your zeta n is indeed in L2.

$$\mathbb{E}||\zeta_n||^2 < +\infty \quad \forall \ n \geq 0$$

$$||\zeta_n||^2 = ||\sum_{k=0}^{n-1} \alpha_k \ M_{k+1}||^2$$

$$\leq n \sum_{k=0}^{n-1} \alpha_k^2 \ ||M_{k+1}||^2$$

$$\mathbb{E}\zeta_n^2 < +\infty$$

So, again, notice that as your n increases, this bound is very loose, and this n can keep increasing, but all we require is that for every n, the expectation of this expression be finite, and that is what we have shown, and that is good enough to conclude that zeta n is in L2. I would like to highlight that, please do not

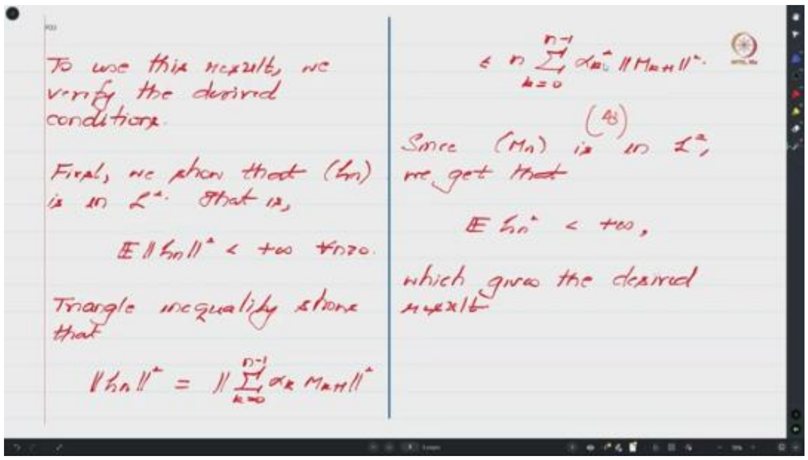

mistake this with showing that it is uniformly bounded in L2. Uniformly bounded would require that the expected value of zeta n squared be upper bounded by the same constant throughout. That is something that we are not showing and, in fact, would make our setup very, very restrictive. Instead, what we are right now showing is that the expected zeta n squared is less than infinity, and this bound can change with n. Alright, so if you recall, Theorem 12.13 said that the limit Xn exists almost surely on the event A infinity less than infinity.

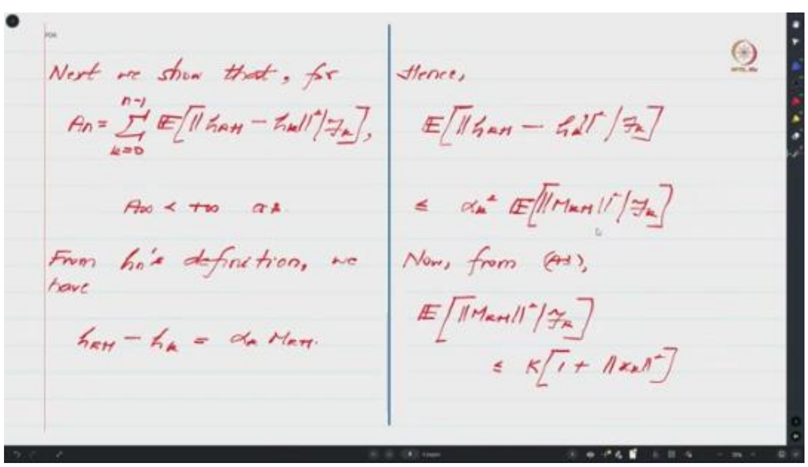

So, in our case, what we will show is that this event itself happens with probability 1. In other words, except for a subset of sample points whose probability is 0, on all the other sample points, A infinity is indeed less than infinity. Again, this bound over here could change from one sample point to the other, but what we require is that on every, or you know, except for a subset of measure 0, on every other sample point, the bound be finite. So, let me elaborate this a bit more.

So, A infinity is a random variable, and we require that A infinity of omega be less than infinity for all omega except in a subset of measure 0, except in a subset with probability 0. Alright, so towards showing that this is less than infinity almost surely, let us recall the definition of a n, okay? So, a n is basically the, you know, the expected value of the squares of the norm of this, right. So, this is slightly different from what we had seen in Theorem 12.13. So, recall that here we have this xk minus xk minus 1, okay.

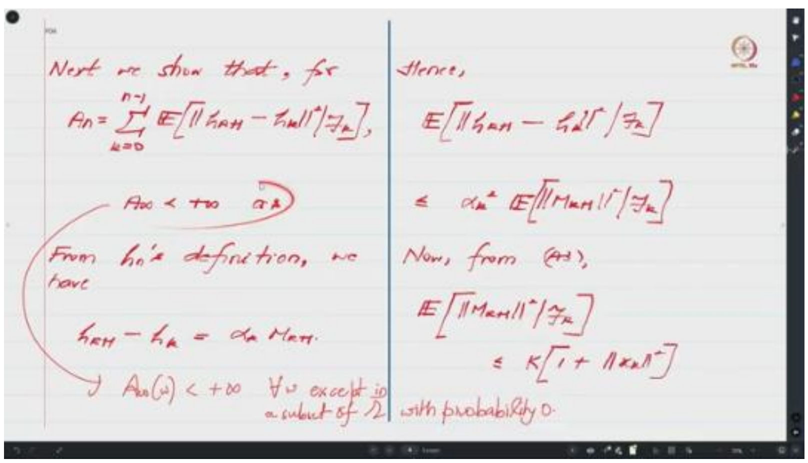

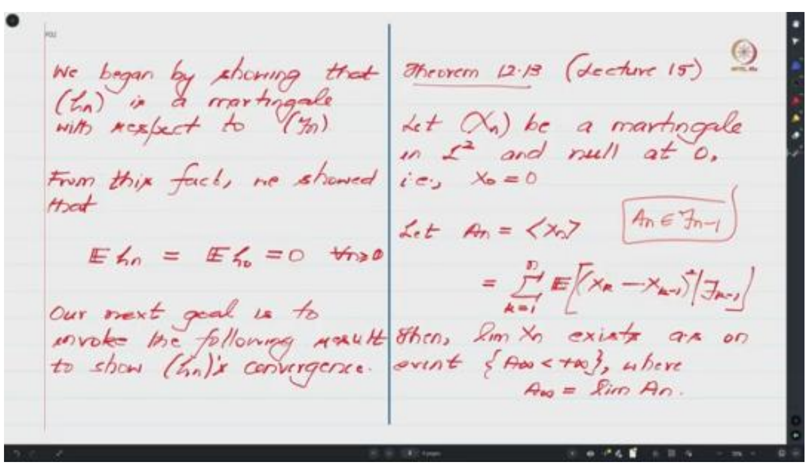

So, we were able to work with this because, in this result, we were working with real-valued martingales. However, you know, in our stochastic approximation context, we are working with vector-valued random variables. And hence, we are going to work with the norms of this zeta n sequence. Is this okay? So, towards that, we define a n to be the expected value of the square of this norm over here.

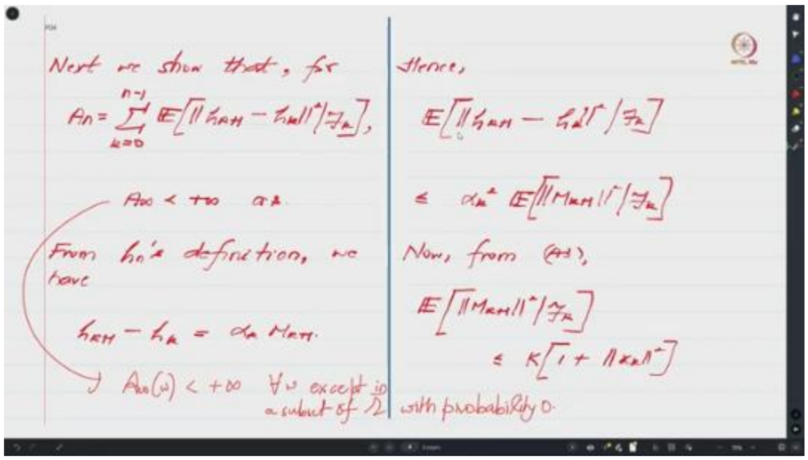

And from the definition of zeta n, recall that the difference between successive terms, that is zeta k plus 1 minus zeta k, would be alpha k times mk plus 1. Hence, if you look at the conditional expectation of the square of the norm of zeta k plus 1 minus zeta k, which is one of the terms over here, we will end up with, you know, this expression being alpha k squared times the norm of mk plus 1 squared. In fact, this is actually not an inequality, but rather this is an equality, right? And now, in our assumption A3, we had presumed that this conditional expectation is upper bounded by k times 1 plus

norm xk square. So, this was one of the statements under A3, and hence one can conclude that this expression over here is upper bounded by alpha k square times k times 1 plus norm xk square. Now, separately, our assumption A4 shows that the xk's are uniformly bounded almost surely, and hence if you take the squares of xk's, then this expression. is also less than infinity almost surely. Again, I would like to say that, you know, your xk's are random variables.

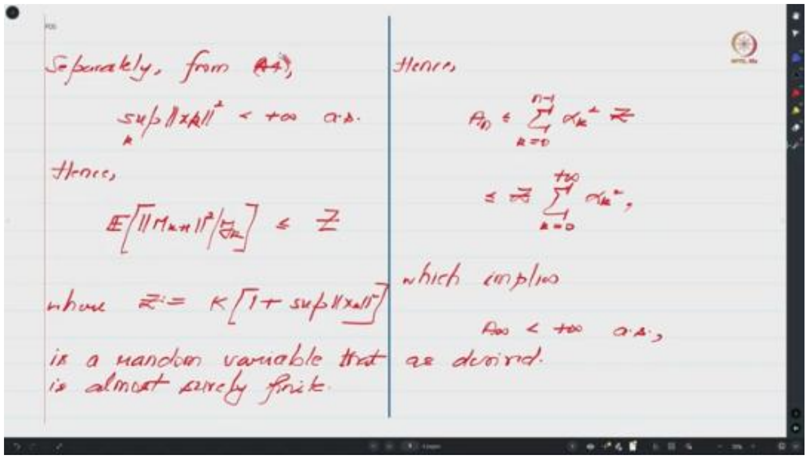

So, xk square will be a random variable, and the supremum will also be a random variable, and all we can conclude is that the collection of points, the collection of sample points where this quantity is finite, has probability 1. Hence, one can conclude that your, you know, conditional value of mk plus 1 square given fk, right, this is upper bounded by this random variable z, which is defined to be k times 1 plus supremum of xk squares, right. So, this expression now becomes, you know, independent of this small index k, right, because we have a supremum, and this supremum is over 1. k over here. So, this supremum is actually over k. Hence, this expression is now devoid of k. So, in other words, this is a common bound for all these expressions on the left which depend on k.

And recall that Z is a random variable, and all that we know is that this random variable Z is less than infinity almost surely. And accordingly, we can then conclude that since a n is the sum of these expressions, and this expression is alpha k square times this, and we have shown that each of these expressions is upper bounded by Z. So, we can conclude that a n is upper bounded by the sum of alpha k square times z, and since this z does not depend on k, we can pull it out and write it as z times the sum of these expressions, and since these are non-negative terms, all these sums can be upper bounded by the sum of alpha k square where k goes from 0 to infinity. And we have presumed that these, you know, step sizes are chosen such that this quantity over here is less than infinity, which implies that a n is upper bounded by a bounded random variable. almost surely bounded random variable for every n. So, for every n, a n is upper bounded by this, and since your a n is a monotonically increasing sequence which is upper bounded by an almost surely bounded random variable, one can conclude that its limit, that is the limit of a n which we denote by a infinity, is less than infinity almost surely.

So, one can now conclude that the conditions of Theorem 12.13 hold, that is, that your martingale sequence is in L2 and this A infinity less than infinity event holds almost surely, and hence one can use Theorem 12.13 to conclude that the limit zeta n actually exists and is finite almost surely. Is this okay? And since we are in R^D, right, that is your d-dimensional Euclidean space, we know that every convergent sequence is Cauchy. And since every convergent sequence is Cauchy, one can conclude that if you look at the difference between zeta n plus k and zeta n, take their norm and take the supremum, right?

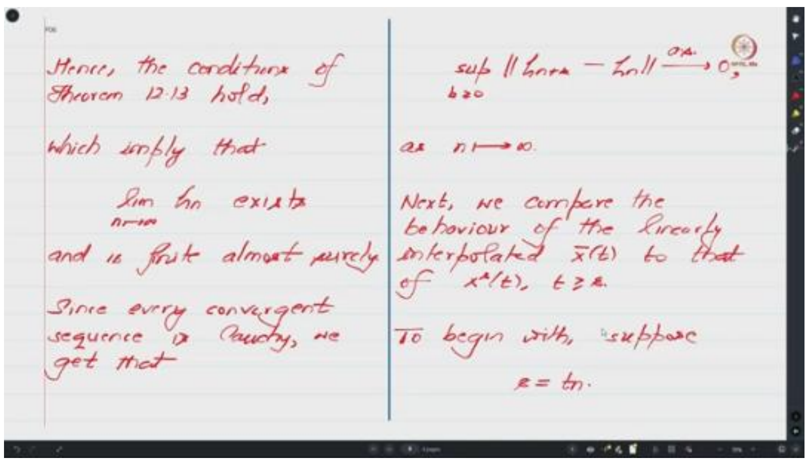

Then this expression goes to 0 almost surely. So, this statement we get for free from the fact that your zeta n actually converges, right? So, this is just a reinterpretation of a convergent sequence from a Cauchy perspective, the Cauchy criterion perspective, and what this tells us is that if you look at the tail of your cumulative noise, right? Then the tail of the cumulative noise, as n tends to infinity, has negligible effect. So, this is, in some sense, telling us that eventually the noise is not going to affect the behavior of your stochastic approximation algorithm.

So, let us pause here and, you know, try to understand why we were able to show this. So, the idea over here is that we were, you know, using this Theorem 12.13. And this Theorem 12.13 required that this A n be less than infinity, and by our assumptions on your x n sequence and alpha n, in particular by presuming that this alpha summation alpha k square is less than infinity, which means that your step sizes decay sufficiently fast, we were able to show that this A infinity is less than infinity almost surely, right? And from this fact and the fact that your iterates are almost surely bounded, we were able to conclude that your zeta n converges almost surely, from which we were able to show that the tail of your noise sequence actually goes to 0 as n tends to infinity, right.

So, in some sense, we have handled, you know, one part of this lemma. So, now we are going to discuss the rest of the lemma. Again, this proof is quite involved. So, we will be doing it in bits and pieces in this class. We will be, you know, sort of setting the stage for discussions in the subsequent lectures. Okay, so the goal now is to, you know, use this fact and compare the behavior of the linearly interpolated trajectory X bar of T to Xs of T. Recall what Xs of T is: it is the solution of your limiting ODE which, at time S, passes through X bar of S, right? This means at time S, X bar of T and Xs of T start at the same place. However, after S, the behavior of X bar of T is governed by your Xn iterates, whereas the behavior of Xs of T is governed by the vector field dictated by your H function.

Is this okay? Which is deterministic in nature. So, let us do this comparison. So, in this lemma 1, your S can be arbitrary, and we only require that S go to infinity.

However, we will, you know, do the proof first by presuming that S is Tn, and later on, we will see how to, you know, get rid of this restriction. So, since X is Tn, what we are interested in comparing is, you know, your X bar of T versus we are going to compare X TN of T, right? So, these are the two trajectories that we wish to compare against each other, right? Now, you know, in the previous class, we had said that X N plus M is

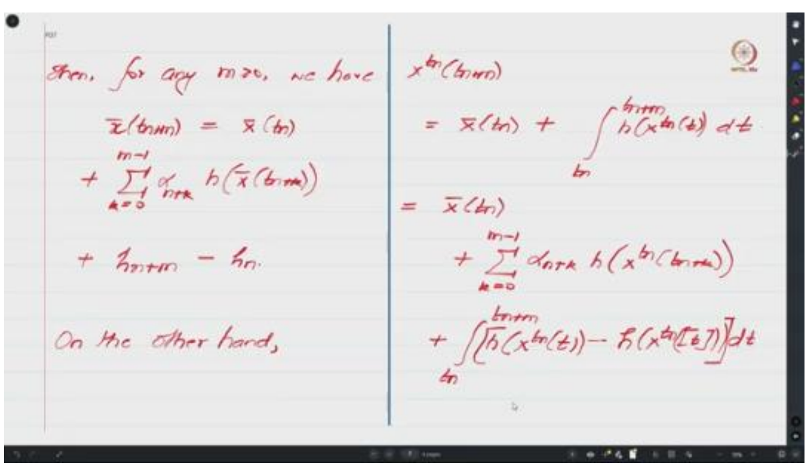

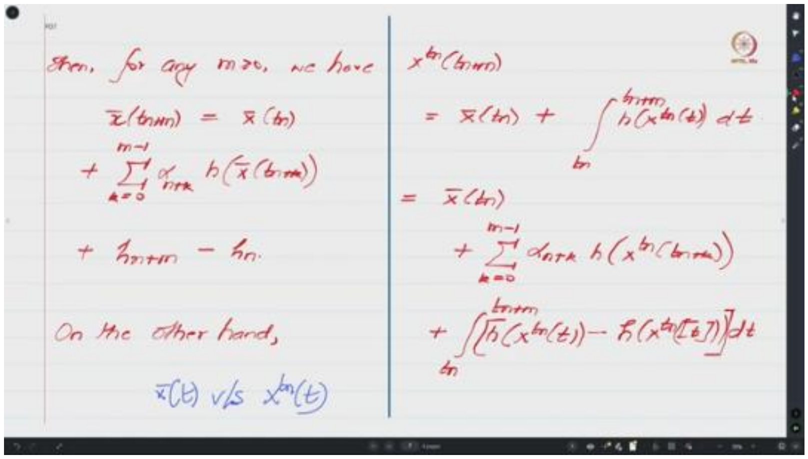

is equal to X n plus summation k equals 0 till m minus 1 alpha k H of—sorry, I should not say alpha k but rather I should say alpha of n plus k H of X n plus k. plus a similar cumulative noise term which is k equals 0 to m minus 1 alpha n plus k m n plus k plus 1. So, this is something that we had shown in the previous class. Now, the way we have linearly interpolated your x bar—I mean, linearly interpolated these xn values to get this x bar t sequence—one can show that x bar of t n plus m is actually x n plus m and x bar of t n is actually xn. Hence, whatever expression you have over here can be written in this form, and this expression is basically zeta n plus m minus zeta n. In other words, the evolution of your linearly interpolated trajectory at time instance n plus m can be expressed as the behavior at time instance tn plus tn.

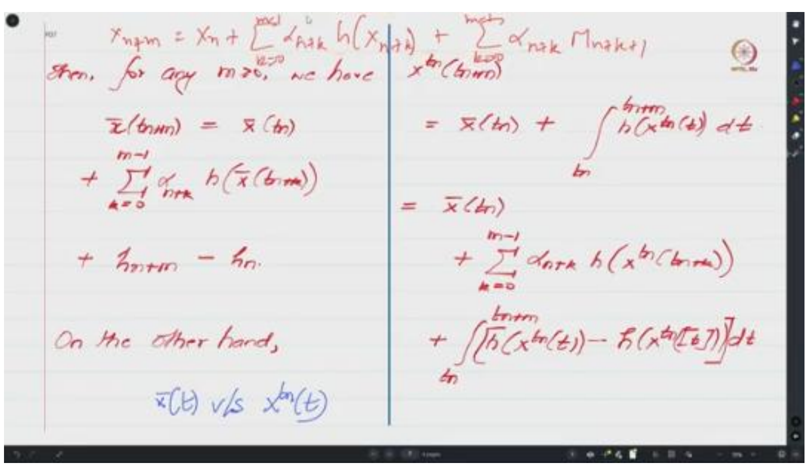

So, what happens based on the H function, plus what happens due to the noise term. So, we have sort of split the effect of noise and the effect of H from each other. So, now one can ask: this is the behavior of your linearly interpolated trajectory. What can we say about the behavior of your solution of your ODE? So, by that, I mean if you start this solution at X bar Tn and you look at the solution of the ODE at time Tn plus m, where would it—

So, one can show that because this solves the ODE relation, this expression equals where you started from, which is X bar Tn, plus the time that has elapsed—that is, from Tn to Tn plus m—of H of the value of this solution trajectory at some time T dt. Is this okay? So, what I have written over here is H of X Tn of T. Is this okay? So, now you know this is where the solution trajectory is at time tn plus m, and this is where the linearly interpolated trajectory is at time tn plus m, and we would like to compare.

This with this, and in order to, you know, do this comparison, what we will do is we will write this expression in some simplified form. In particular, what we will do is we will take this right-hand side and write it as X bar of Tn plus alpha n plus k times h of this value at specific values of t, which is t n plus k, and since we have added this, we will subtract it over here, and whatever we have here, we will retain it over here. So, let us look at what we have subtracted. So, we have written some expression in a fancy way, but if you take the integral of this, of this, this is exactly this expression, and now I am going to illustrate why that is the case.

So, if you see here, the argument includes a square bracket t. So, the square bracket t, let me first define. This is the max of tk such that tk is less than or equal to t. So, in some sense, this is the largest Tk which is less than or equal to T. So, which means that if you integrate this expression, so let me elaborate. So, if you integrate this expression between, let us say, Tk to Tk plus 1. So, this expression with this square bracket T that you have over here.

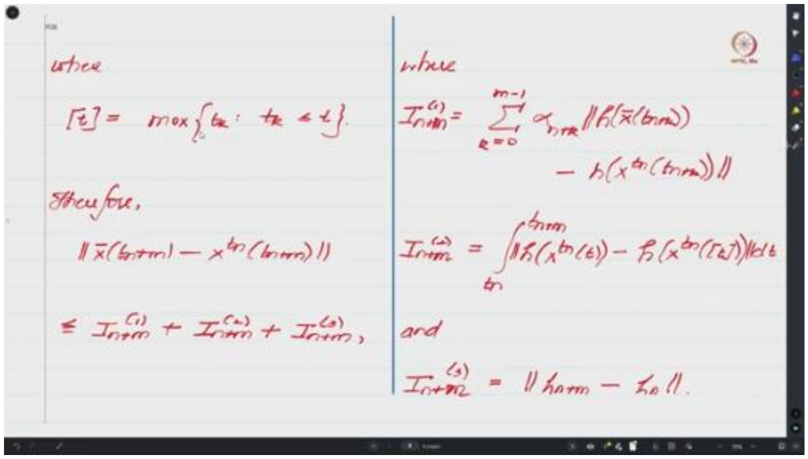

Right? This expression will be constant because for every t between tk to tk plus 1, the square bracket t will be tk. Hence, this whole expression will actually be alpha k times xtn of Tk. Now, you may wonder, how did I get this alpha k?

Well, I get this alpha k because that is the difference between Tk plus 1 and Tk. In other words, since this is a constant, this integral is this constant expression which is over here times the difference between the upper limit and the lower limit. And the way we had defined these tk's, one can check that tk plus 1 minus tk is alpha k, and hence this integral will be alpha k times h of tn tk. Right, and accordingly, one can see that this expression over here plus this expression is indeed this expression that we have. Okay, so one can now try to compare x bar of t n plus m with x tn of t n plus m, right? And the norm of this difference can be expressed as a sum of three terms. So, what is the first term?

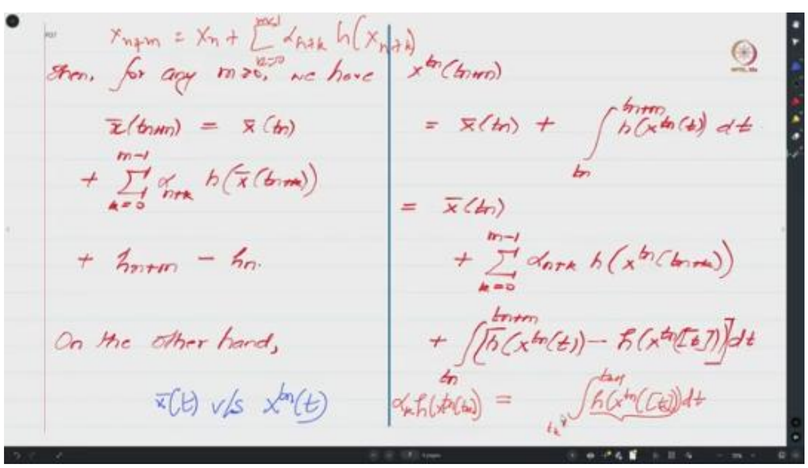

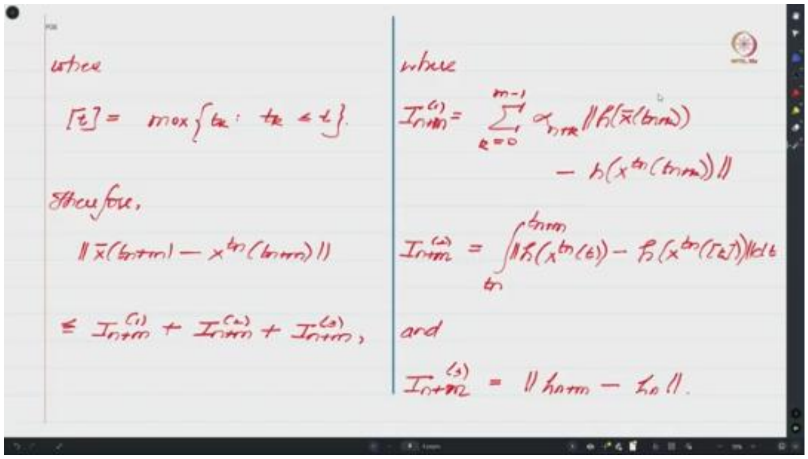

The first term is basically the difference between where X bar Tn plus k is and where X Tn is at time Tn plus k. This is basically your linearly interpolated trajectory, and this is your ODE solution. So, this is basically comparing this expression with this expression over here. And also observe that you had X bar Tn and X bar Tn over here. So, these two things will cancel out when you take the difference between them. So, this explains the first term over here.

The second term over here is basically arising due to this expression. So, one can see that whatever I have over here is precisely this term over here, and the third term that we have arises because of this expression over here. So, let me just give a verbal description of what these different terms capture. This one can think of as the discretization error, and this one can imagine to be the error due to noise, right? And this basically is the gap between your linearly interpolated trajectory and the ODE trajectory.

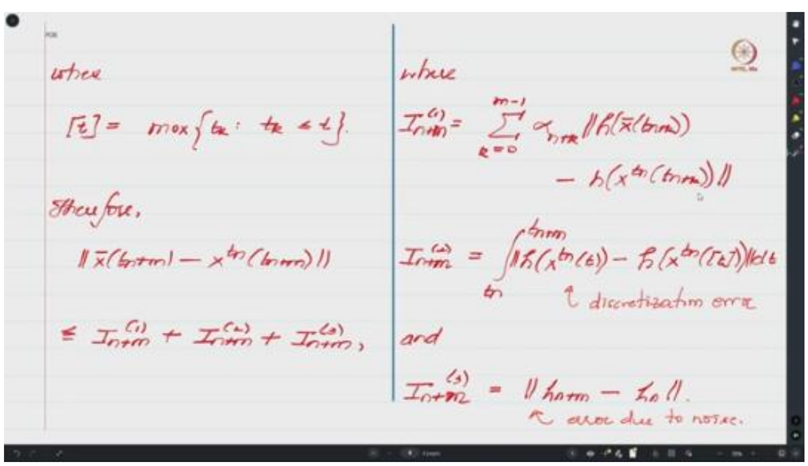

In the next class, we will show you how to bound these three terms. This term, we have already shown that as your n increases, it goes to 0. So, it only remains to bound this expression and somehow handle this expression. In the next class, we will show you how to deal with these terms. So, let me quickly summarize what we have done so far.

We are, in some sense, halfway through the proof of this technical lemma, which tells us that under suitable conditions, the linearly interpolated stochastic approximation trajectory closely tracks the suitable solution trajectory of the limiting ODE. So, we will complete the rest of the proof in the next class. Until then, thank you. Namaste and bye.