Nonlinear Adaptive Control
Professor Srikant Sukumar
Systems and Control
Indian Institute of Technology, Bombay
Week 12
Lecture No: 66
Discussion on Historical Developments in Adaptive Control and Learning

(Refer Slide Time: 00:17)

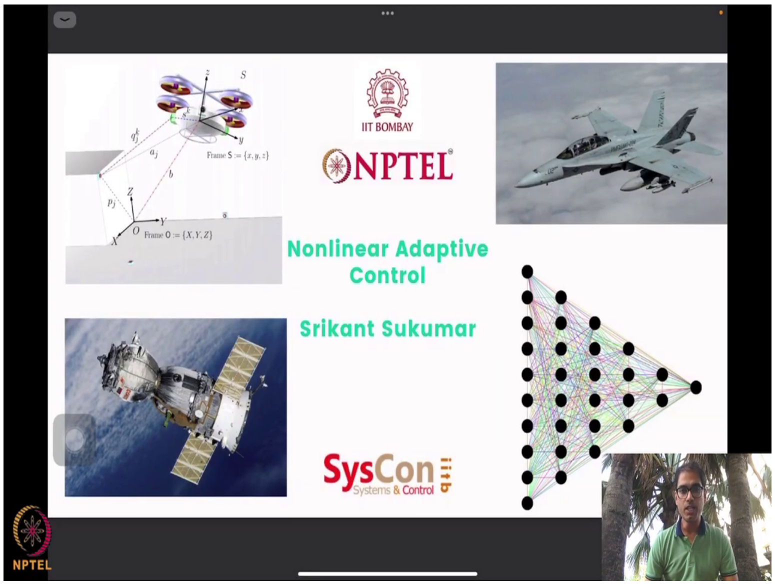

Hello, everyone. Welcome to yet another week of our NPTEL on Nonlinear and Adaptive Control. I am Srikant Sukumar, Systems and Control, IIT Bombay. So, we are embarking into the final week of our course on Nonlinear and Adaptive Control. I do really hope that you have learned quite a bit during this course on how to design and analyze robust adaptive control algorithms that can drive systems such as the drone, the aircraft, the spacecraft, the neural network that you see here in the background.

As I have mentioned before, I am always very, very open to hearing more about the sort of applications that you folks are working on or plan to working on, plan to work on using the tools and techniques that you have learned here and I am hoping to hear more about that from all of you.

(Refer Slide Time: 01:27)

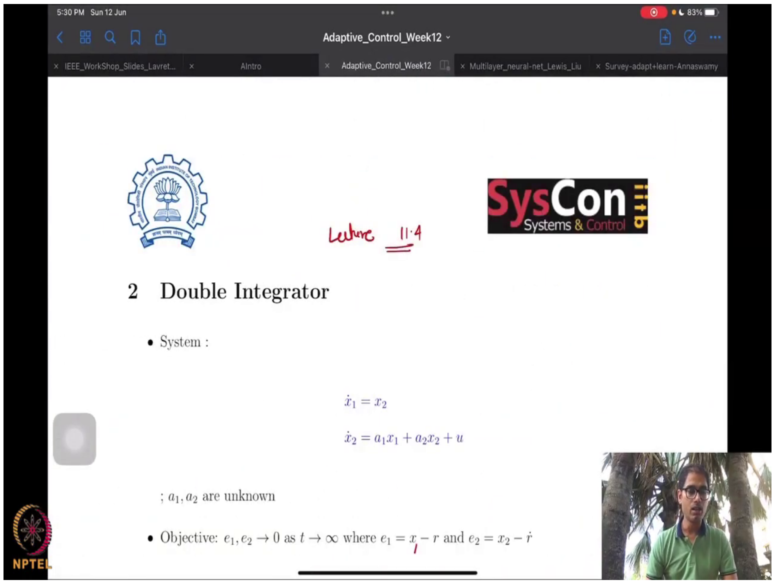

So, what we were doing until last time was to look at a particular kind of adaptive control design which does not rely on persistent excitation anymore. And only relies on having excitation for a finite amount of time.

(Refer Slide Time: 01:44)

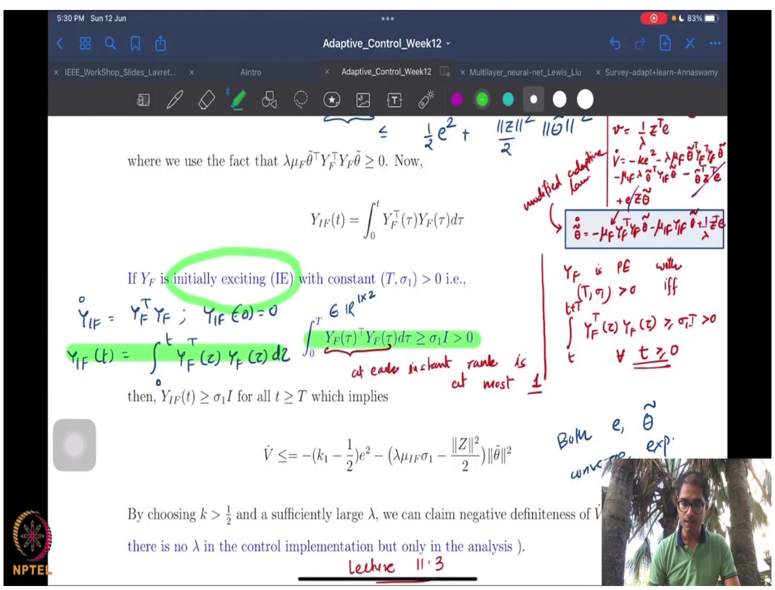

In this condition, it was of course called the initial excitation condition. And we of course saw how to do control design, we saw how it was a weaker requirement as opposed to your persistence excitation-based design. And we saw both the advantages and limitations of this method. We also saw that it is possible to alleviate some of the limitations by slightly modifying the initial excitation design to add a certainty equivalence adaptive control term.

(Refer Slide Time: 02:50)

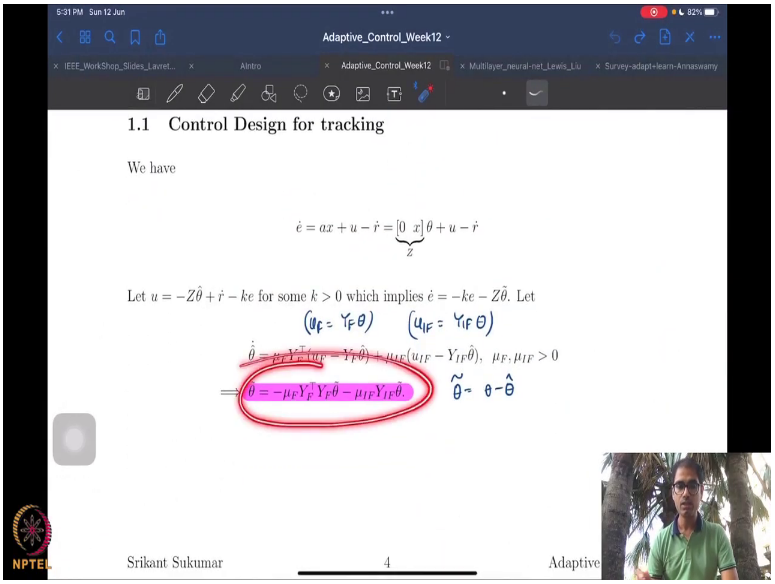

So, keeping that in mind if you do hope that this sort of design method helps us to give some improvements, not just in theory, of course, we have seen real theory that there are significant improvements especially in the sense that you have these negative terms in your adaptive update law which was of course not there earlier.

So, we do, of course, hope that this sort of improvement is also something that you see in the performance of the adaptive control, specifically, the numerical performance of this adaptive control. So, again, please try it out on actual systems and please do report back how it looks, how the performance looks compared with your traditional certainty equivalence controller and we would like to hear more from you.

(Refer Slide Time: 03:40)

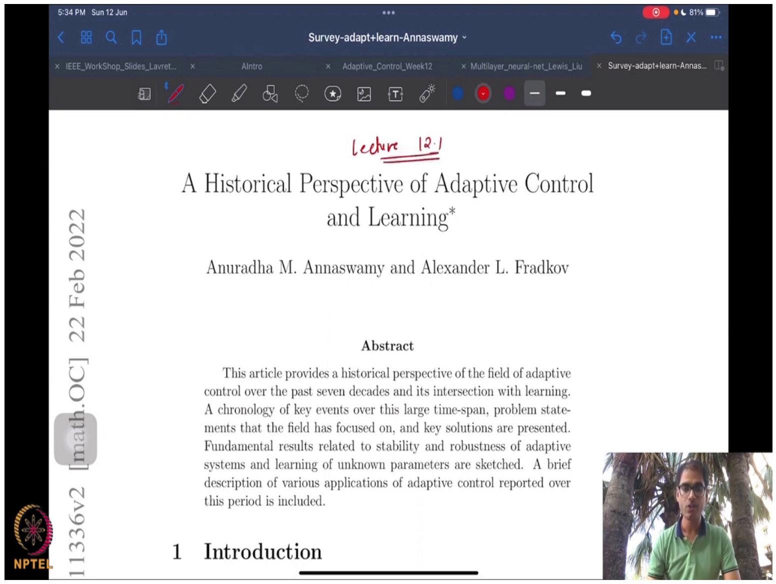

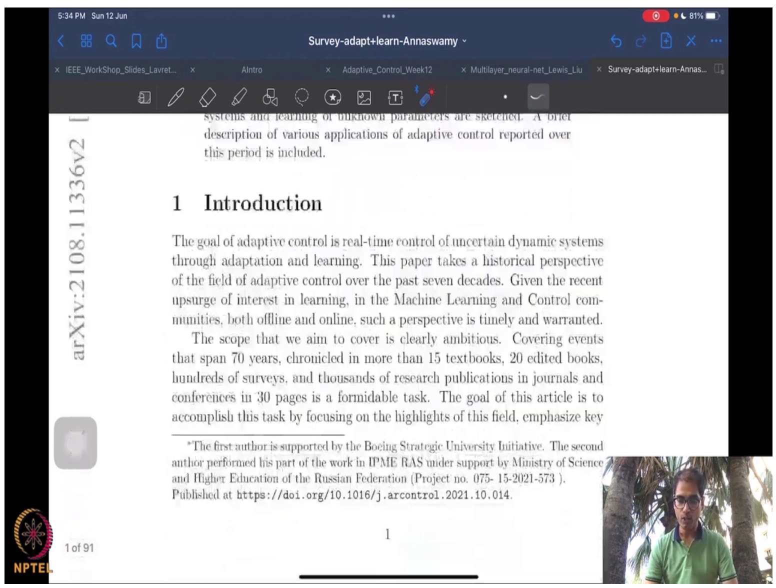

Now, where we want to go starting today is on discussing sort of modern ideas in adaptive control. And this, the aim is here to make some connections with learning theory, which has become very, very popular. And of course, adaptive control is also consisting of learning itself. So, this is another case in point that we will try to make that parameter learning is essentially an ingredient or a key ingredient of what you do in deep learning.

Albeit the algorithms that are used may differ and, but eventually the basic theory remains the same. So, the, so the key connection that we will try to make is with neural networks and deep learning. So, multi-layer neural networks which is essentially what deep learning involves, but before we do any of that, I want to point you and also talk a little bit about this particular paper.

As you can see this is a very, very recent article that let's come up in archive online and of course, I will put this up in the course material also. This is by Anuradha M. Annaswamy and Alexander Fradkov. These are rather senior researchers in the field of adaptive control and learning.

So, what I want to do is to sort of go over this paper, sort of give an overview of this paper in this session, and sort of get a feel for how adaptive control has evolved, what pieces of adaptive control we have covered in this course and of course what is left and what we have left out, due to again lack of time in any course such as this. And we also want to see what kind of connections with learning has been made in literature.

So, this is a what we call a survey paper. So, I am going to mark it here as lecture 12.1. So, our last week of lectures. So, like I said, this is a survey paper. So, we are not really going to do a lot of mathematics or discuss a lot of mathematics like we have been doing but we look at how things have evolved.

(Refer Slide Time: 06:27)

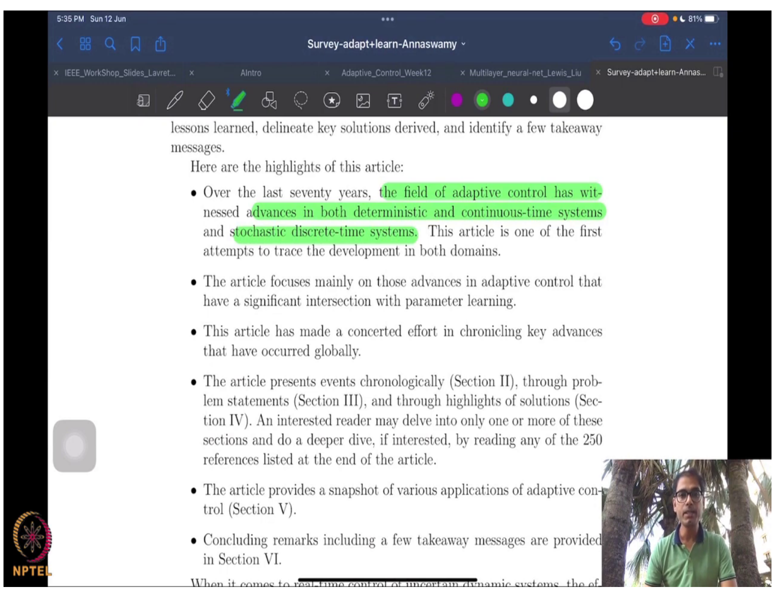

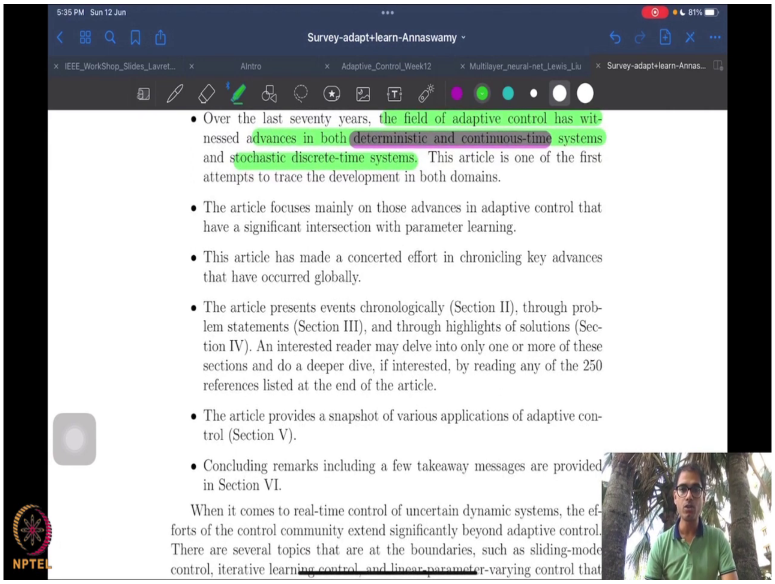

Survey-adapt+learn-Annaswamy

- Over the last seventy years, the field of adaptive control has witnessed advances in both deterministic and continuous-time systems and stochastic discrete-time systems. This article is one of the first attempts to trace the development in both domains.
- The article focuses mainly on those advances in adaptive control that have a significant intersection with parameter learning.
- This article has made a concerted effort in chronicling key advances that have occurred globally.
- The article presents events chronologically (Section II), through problem statements (Section III), and through highlights of solutions (Section IV). An interested reader may delve into only one or more of these sections and do a deeper dive, if interested, by reading any of the 250 references listed at the end of the article.
- The article provides a snapshot of various applications of adaptive control (Section V).
- Concluding remarks including a few takeaway messages are provided in Section VI.

When it comes to real-time control of uncertain dynamic systems, the efforts of the control community extend significantly beyond adaptive control. There are several topics that are at the boundaries, such as sliding-mode control, iterative learning control, and linear-parameter-varying control that

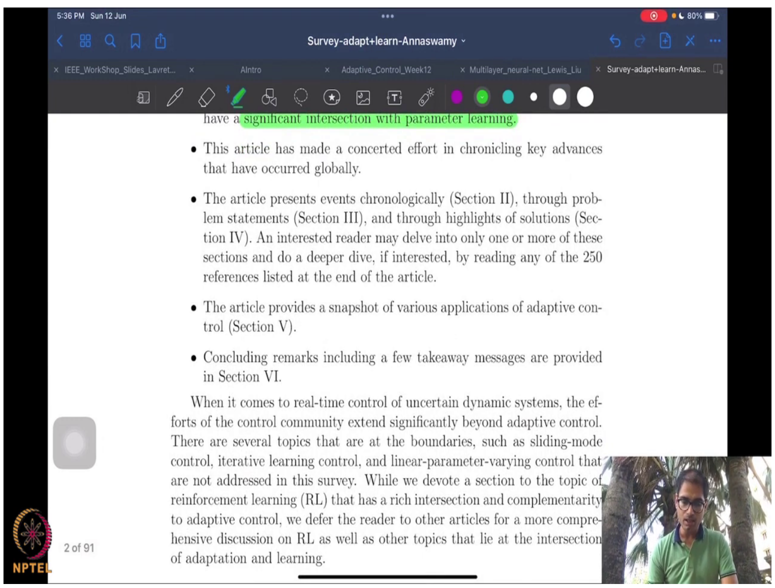

Survey-adapt+learn-Annaswamy

have a significant intersection with parameter learning.

- This article has made a concerted effort in chronicling key advances that have occurred globally.
- The article presents events chronologically (Section II), through problem statements (Section III), and through highlights of solutions (Section IV). An interested reader may delve into only one or more of these sections and do a deeper dive, if interested, by reading any of the 250 references listed at the end of the article.
- The article provides a snapshot of various applications of adaptive control (Section V).
- Concluding remarks including a few takeaway messages are provided in Section VI.

When it comes to real-time control of uncertain dynamic systems, the efforts of the control community extend significantly beyond adaptive control. There are several topics that are at the boundaries, such as sliding-mode control, iterative learning control, and linear-parameter-varying control that are not addressed in this survey. While we devote a section to the topic of reinforcement learning (RL) that has a rich intersection and complementarity to adaptive control, we defer the reader to other articles for a more comprehensive discussion on RL as well as other topics that lie at the intersection of adaptation and learning.

2 of 91

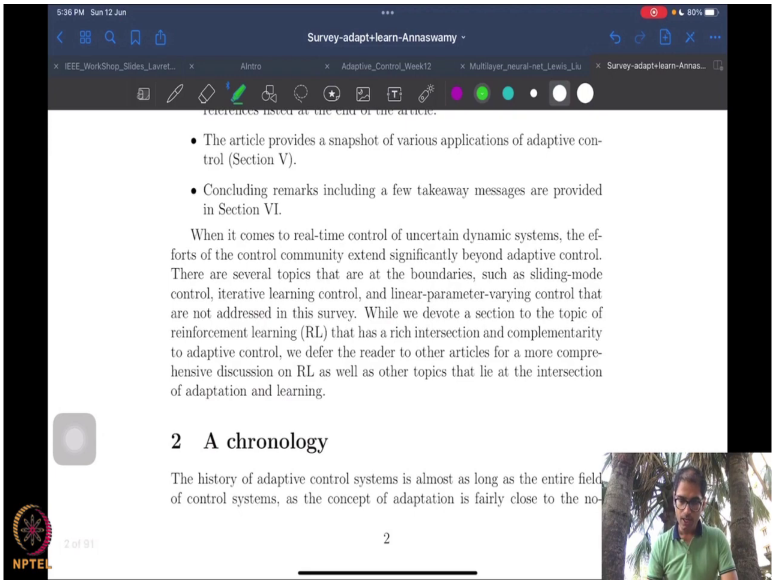

So, so of course these are some of the highlights of course of this article. So, there have been advances in the last 70 years in both, in adaptive control, in both deterministic and deterministic continuous time systems, and also stochastic discrete time systems. So, these are the two sort of dynamical systems that have been looked at extensively in the domain of adaptive control.

And as you can see, we have covered of course only deterministic and continuous time systems which is why I am going to mark it in a different color here, but of course there exists a lot of literature and work on stochastic discrete systems in adaptive control also. So, so this article of course is sort of trying to connect adaptive control to learning.

So, it sort of looks at the developments that have significant intersection with parameter learning. And of course, I mean as they usually do in survey articles, there is a lot of references that you can go through, and as part of different sections, and if you want to delve into a particular topic in much more detail and. So, you will of course get a very good, very good idea of how you want to proceed to learn a particular topic if you are interested in that from here.

(Refer Slide Time: 08:15)

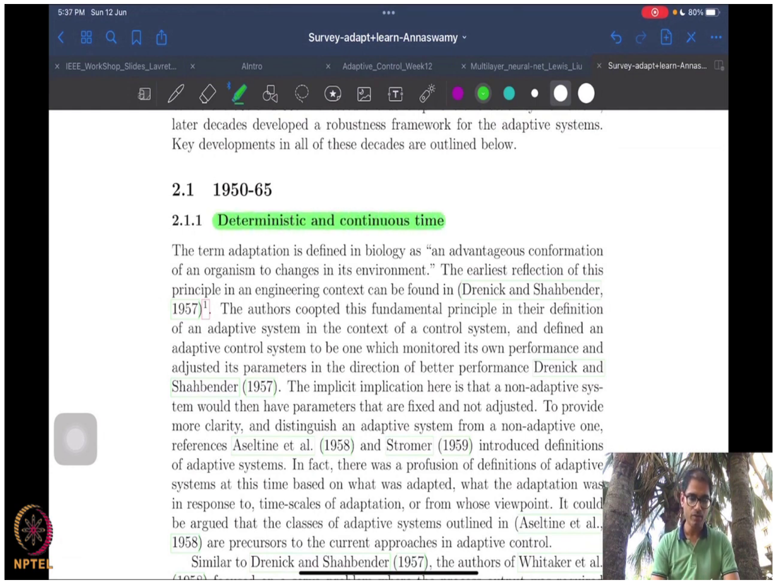
Survey-adapt+learn-Annaswamy

later decades developed a robustness framework for the adaptive systems. Key developments in all of these decades are outlined below.

2.1 1950-65

2.1.1 Deterministic and continuous time

The term adaptation is defined in biology as "an advantageous conformation of an organism to changes in its environment." The earliest reflection of this principle in an engineering context can be found in (Drenick and Shahbender, 1957)[1]. The authors coopted this fundamental principle in their definition of an adaptive system in the context of a control system, and defined an adaptive control system to be one which monitored its own performance and adjusted its parameters in the direction of better performance Drenick and Shahbender (1957). The implicit implication here is that a non-adaptive system would then have parameters that are fixed and not adjusted. To provide more clarity, and distinguish an adaptive system from a non-adaptive one, references Aseltine et al. (1958) and Stromer (1959) introduced definitions of adaptive systems. In fact, there was a profusion of definitions of adaptive systems at this time based on what was adapted, what the adaptation was in response to, time-scales of adaptation, or from whose viewpoint. It could be argued that the classes of adaptive systems outlined in (Aseltine et al., 1958) are precursors to the current approaches in adaptive control.

Similar to Drenick and Shahbender (1957), the authors of Whitaker et al.

NPTEL

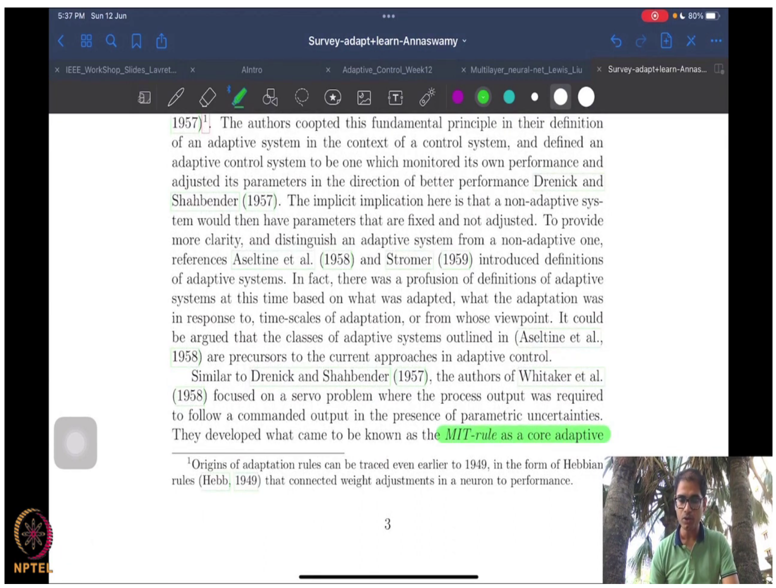
Survey-adapt+learn-Annaswamy

1957)[1]. The authors coopted this fundamental principle in their definition of an adaptive system in the context of a control system, and defined an adaptive control system to be one which monitored its own performance and adjusted its parameters in the direction of better performance Drenick and Shahbender (1957). The implicit implication here is that a non-adaptive system would then have parameters that are fixed and not adjusted. To provide more clarity, and distinguish an adaptive system from a non-adaptive one, references Aseltine et al. (1958) and Stromer (1959) introduced definitions of adaptive systems. In fact, there was a profusion of definitions of adaptive systems at this time based on what was adapted, what the adaptation was in response to, time-scales of adaptation, or from whose viewpoint. It could be argued that the classes of adaptive systems outlined in (Aseltine et al., 1958) are precursors to the current approaches in adaptive control.

Similar to Drenick and Shahbender (1957), the authors of Whitaker et al. (1958) focused on a servo problem where the process output was required to follow a commanded output in the presence of parametric uncertainties. They developed what came to be known as the *MIT-rule* as a core adaptive

[1] Origins of adaptation rules can be traced even earlier to 1949, in the form of Hebbian rules (Hebb, 1949) that connected weight adjustments in a neuron to performance.

3

NPTEL

So, so anyway, so, so we are, we start with, this article starts with some kind of a chronology on how adaptive control went. So, in from 50s to 65 you, you had a lot of work on deterministic continuous time systems, and this is the sort of work that we also discussed. So, there was a development of what is called the MIT rule which formed the core of adaptive control.

(Refer Slide Time: 08:42)

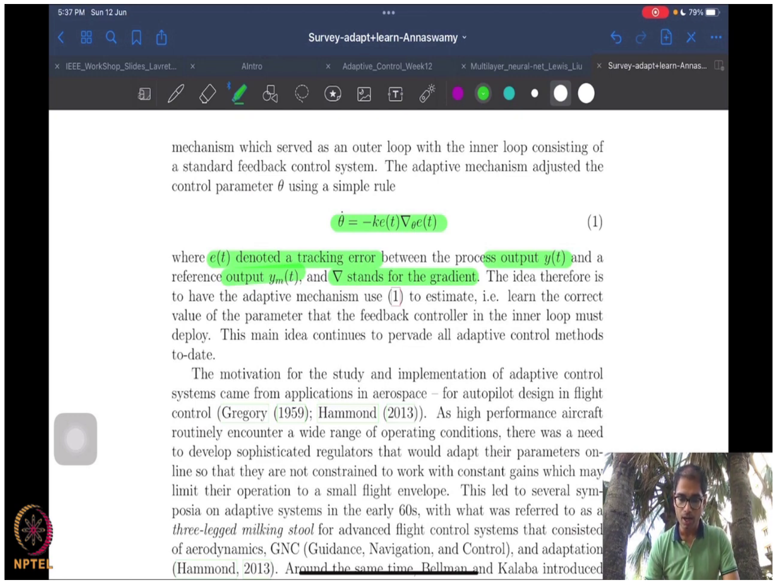

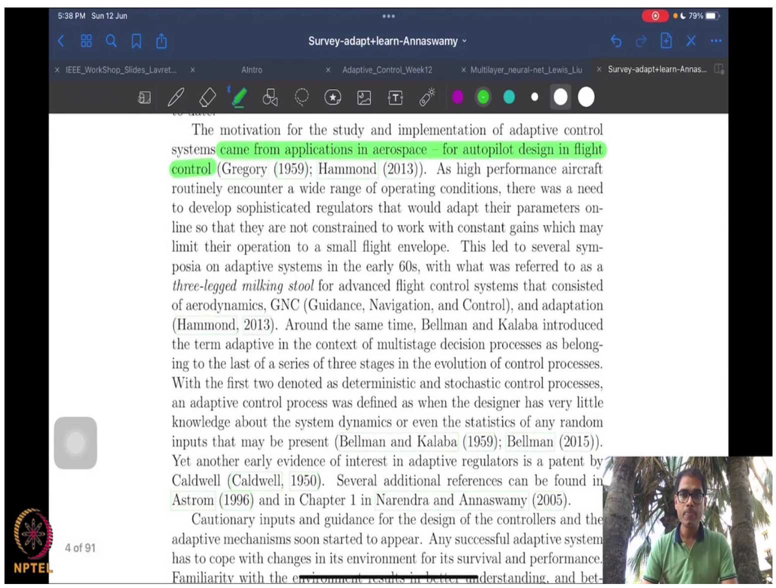

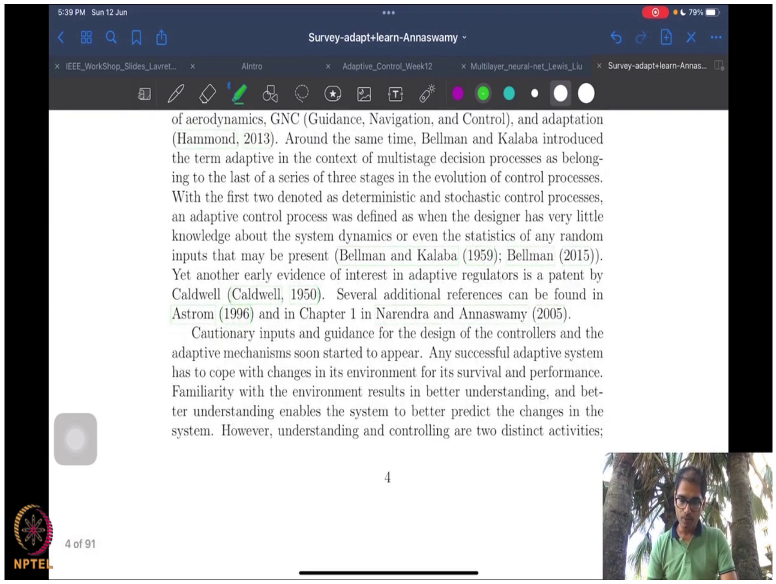

of aerodynamics, GNC (Guidance, Navigation, and Control), and adaptation (Hammond, 2013). Around the same time, Bellman and Kalaba introduced the term adaptive in the context of multistage decision processes as belonging to the last of a series of three stages in the evolution of control processes. With the first two denoted as deterministic and stochastic control processes, an adaptive control process was defined as when the designer has very little knowledge about the system dynamics or even the statistics of any random inputs that may be present (Bellman and Kalaba (1959); Bellman (2015)). Yet another early evidence of interest in adaptive regulators is a patent by Caldwell (Caldwell, 1950). Several additional references can be found in Astrom (1996) and in Chapter 1 in Narendra and Annaswamy (2005).

Cautionary inputs and guidance for the design of the controllers and the adaptive mechanisms soon started to appear. Any successful adaptive system has to cope with changes in its environment for its survival and performance. Familiarity with the environment results in better understanding, and better understanding enables the system to better predict the changes in the system. However, understanding and controlling are two distinct activities;

4

And, and this essentially was a way of updating unknown parameters by a very standard gradient type formula. And what is that gradiant type formula it is this equation 1 that you see here. So, now here, the e is of course the tracking error that we are so used to seeing between this process output y.

And the desired output y m, and the gradient is of course the gradient operator and the variant is taken with respect to the unknown parameter theta. So, so basically the idea is sort of a way of, sort of steepest descent kind of an idea that you have particularly in optimization type algorithms. So, this is essentially what was called the MIT rule and this formed the basis of a lot of adaptive control algorithms.

So, so basically a lot of motivation for adaptive control came from applications in aerospace, for automotive, for autopilot design and flight control, and this still continues to be the case. A lot of more recent work in adaptive control has been applied to autopilots in fighter planes and these have been tested, have been field tested successfully in fact. Of course, getting FAA clearances and security clearances et cetera is a different matter altogether, but they have undergone significant field advance.

(Refer Slide Time: 11:11)

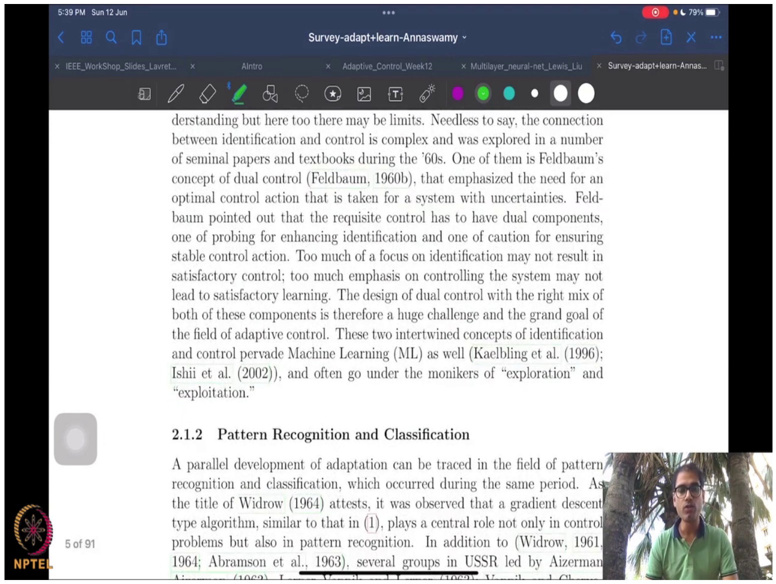

derstanding but here too there may be limits. Needless to say, the connection between identification and control is complex and was explored in a number of seminal papers and textbooks during the '60s. One of them is Feldbaum's concept of dual control (Feldbaum, 1960b), that emphasized the need for an optimal control action that is taken for a system with uncertainties. Feldbaum pointed out that the requisite control has to have dual components, one of probing for enhancing identification and one of caution for ensuring stable control action. Too much of a focus on identification may not result in satisfactory control; too much emphasis on controlling the system may not lead to satisfactory learning. The design of dual control with the right mix of both of these components is therefore a huge challenge and the grand goal of the field of adaptive control. These two intertwined concepts of identification and control pervade Machine Learning (ML) as well (Kaelbling et al. (1996); Ishii et al. (2002)), and often go under the monikers of "exploration" and "exploitation."

2.1.2 Pattern Recognition and Classification

A parallel development of adaptation can be traced in the field of pattern recognition and classification, which occurred during the same period. As the title of Widrow (1964) attests, it was observed that a gradient descent type algorithm, similar to that in (1), plays a central role not only in control problems but also in pattern recognition. In addition to (Widrow, 1961, 1964; Abramson et al., 1963), several groups in USSR led by Aizerman

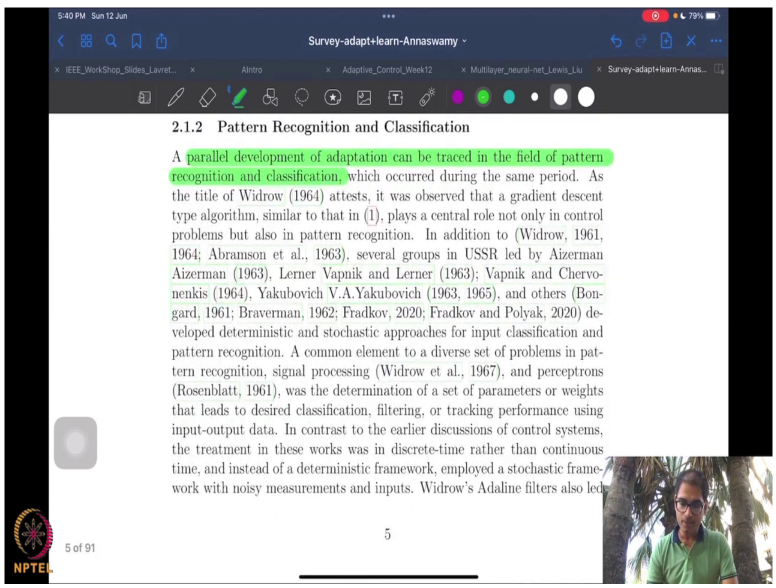

2.1.2 Pattern Recognition and Classification

A parallel development of adaptation can be traced in the field of pattern recognition and classification, which occurred during the same period. As the title of Widrow (1964) attests, it was observed that a gradient descent type algorithm, similar to that in (1), plays a central role not only in control problems but also in pattern recognition. In addition to (Widrow, 1961, 1964; Abramson et al., 1963), several groups in USSR led by Aizerman Aizerman (1963), Lerner Vapnik and Lerner (1963); Vapnik and Chervonenkis (1964), Yakubovich V.A.Yakubovich (1963, 1965), and others (Bongard, 1961; Braverman, 1962; Fradkov, 2020; Fradkov and Polyak, 2020) developed deterministic and stochastic approaches for input classification and pattern recognition. A common element to a diverse set of problems in pattern recognition, signal processing (Widrow et al., 1967), and perceptrons (Rosenblatt, 1961), was the determination of a set of parameters or weights that leads to desired classification, filtering, or tracking performance using input-output data. In contrast to the earlier discussions of control systems, the treatment in these works was in discrete-time rather than continuous time, and instead of a deterministic framework, employed a stochastic framework with noisy measurements and inputs. Widrow's Adaline filters also led

5

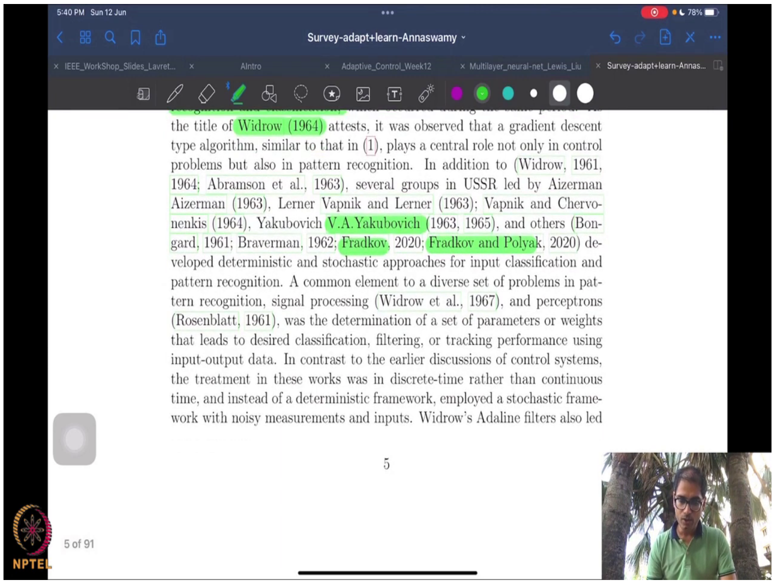

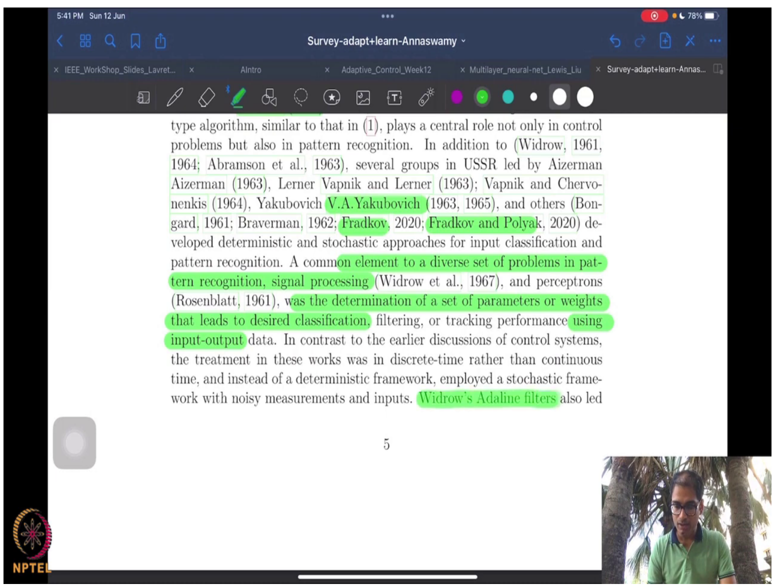

So, so anyway, so, so basically, so this is sort of what this section of this article talks about. So, so rather, rather, rather interesting, interesting set of topics on which there was a requirement to do adaptation of parameters and adaptation of unknown quantities. So, so anyway, so, so this is a sort of motivation for what why or how adaptive control for continuous time systems did come about.

There is also the, a discussion on this problem of pattern recognition and classification, and this was seen as a parallel development in adapt, to adaptation. And interestingly, they, and a lot of, of course, there were very, very different kinds of researchers who were working on these topics like Widrow, who was very, very famous. But then there were also folks from adaptive control and non-linear control like Yakubovich, who also were working on it. Fradkov, Polyak in more recent works.

So, so basically the idea is that there was a common element in, to a diverse set of problems in pattern recognition and signal processing which is the determination of, of a set of parameters or weights that leads to a desired classification. So, eventually there was also the notion of trying to identify parameters or weights and using basically just input, output data. There was no clear models per se, or the models were incomplete, to a large extent but the idea was to sort of identify parameters that would help describe the model.

(Refer Slide Time: 12:51)

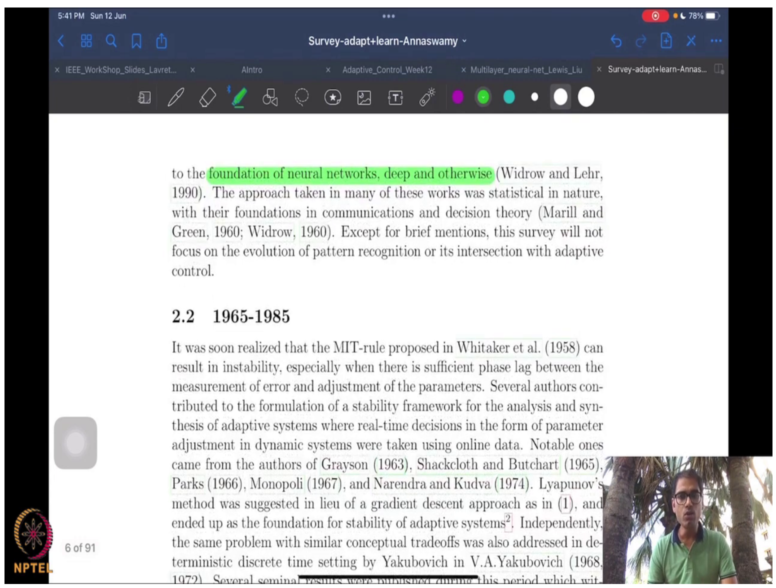

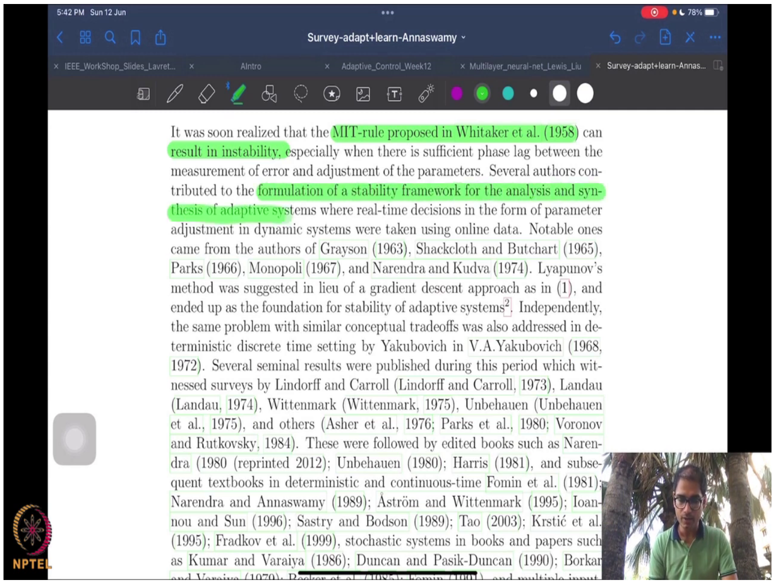

It was soon realized that the MIT-rule proposed in Whitaker et al. (1958) can result in instability, especially when there is sufficient phase lag between the measurement of error and adjustment of the parameters. Several authors contributed to the formulation of a stability framework for the analysis and synthesis of adaptive systems where real-time decisions in the form of parameter adjustment in dynamic systems were taken using online data. Notable ones came from the authors of Grayson (1963), Shackcloth and Butchart (1965), Parks (1966), Monopoli (1967), and Narendra and Kudva (1974). Lyapunov's method was suggested in lieu of a gradient descent approach as in (1), and ended up as the foundation for stability of adaptive systems[2]. Independently, the same problem with similar conceptual tradeoffs was also addressed in deterministic discrete time setting by Yakubovich in V.A.Yakubovich (1968, 1972). Several seminal results were published during this period which witnessed surveys by Lindorff and Carroll (Lindorff and Carroll, 1973), Landau (Landau, 1974), Wittenmark (Wittenmark, 1975), Unbehauen (Unbehauen et al., 1975), and others (Asher et al., 1976; Parks et al., 1980; Voronov and Rutkovsky, 1984). These were followed by edited books such as Narendra (1980 (reprinted 2012); Unbehauen (1980); Harris (1981), and subsequent textbooks in deterministic and continuous-time Fomin et al. (1981); Narendra and Annaswamy (1989); Åström and Wittenmark (1995); Ioannou and Sun (1996); Sastry and Bodson (1989); Tao (2003); Krstić et al. (1995); Fradkov et al. (1999), stochastic systems in books and papers such as Kumar and Varaiya (1986); Duncan and Pasik-Duncan (1990); Borkar

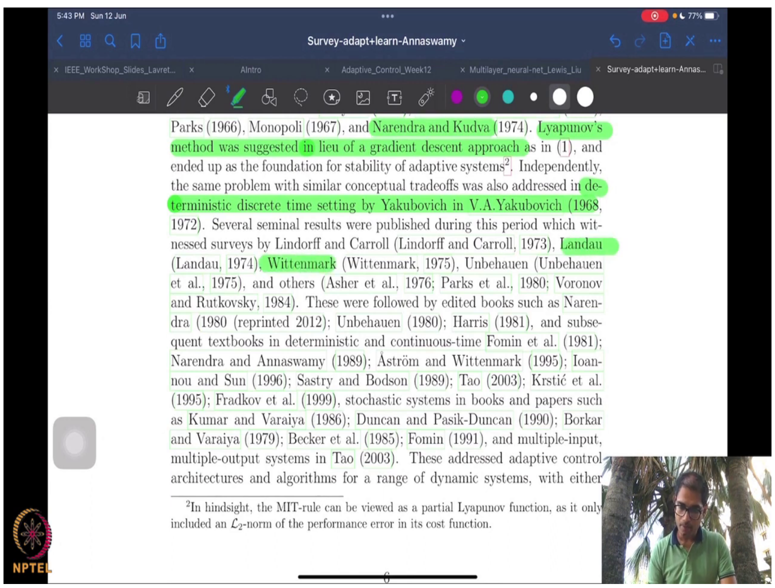

Parks (1966), Monopoli (1967), and Narendra and Kudva (1974). Lyapunov's method was suggested in lieu of a gradient descent approach as in (1), and ended up as the foundation for stability of adaptive systems[2]. Independently, the same problem with similar conceptual tradeoffs was also addressed in deterministic discrete time setting by Yakubovich in V.A.Yakubovich (1968, 1972). Several seminal results were published during this period which witnessed surveys by Lindorff and Carroll (Lindorff and Carroll, 1973), Landau (Landau, 1974), Wittenmark (Wittenmark, 1975), Unbehauen (Unbehauen et al., 1975), and others (Asher et al., 1976; Parks et al., 1980; Voronov and Rutkovsky, 1984). These were followed by edited books such as Narendra (1980 (reprinted 2012); Unbehauen (1980); Harris (1981), and subsequent textbooks in deterministic and continuous-time Fomin et al. (1981); Narendra and Annaswamy (1989); Åström and Wittenmark (1995); Ioannou and Sun (1996); Sastry and Bodson (1989); Tao (2003); Krstić et al. (1995); Fradkov et al. (1999), stochastic systems in books and papers such as Kumar and Varaiya (1986); Duncan and Pasik-Duncan (1990); Borkar and Varaiya (1979); Becker et al. (1985); Fomin (1991), and multiple-input, multiple-output systems in Tao (2003). These addressed adaptive control architectures and algorithms for a range of dynamic systems, with either

[2] In hindsight, the MIT-rule can be viewed as a partial Lyapunov function, as it only included an $\mathcal{L}_2$-norm of the performance error in its cost function.

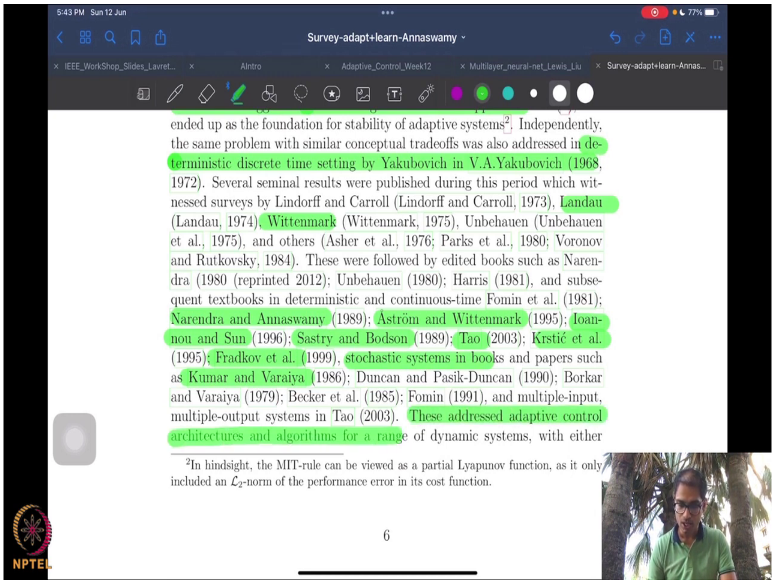

ended up as the foundation for stability of adaptive systems[2]. Independently, the same problem with similar conceptual tradeoffs was also addressed in deterministic discrete time setting by Yakubovich in V.A.Yakubovich (1968, 1972). Several seminal results were published during this period which witnessed surveys by Lindorff and Carroll (Lindorff and Carroll, 1973), Landau (Landau, 1974), Wittenmark (Wittenmark, 1975), Unbehauen (Unbehauen et al., 1975), and others (Asher et al., 1976; Parks et al., 1980; Voronov and Rutkovsky, 1984). These were followed by edited books such as Narendra (1980 (reprinted 2012); Unbehauen (1980); Harris (1981), and subsequent textbooks in deterministic and continuous-time Fomin et al. (1981); Narendra and Annaswamy (1989); Åström and Wittenmark (1995); Ioannou and Sun (1996); Sastry and Bodson (1989); Tao (2003); Krstić et al. (1995); Fradkov et al. (1999), stochastic systems in books and papers such as Kumar and Varaiya (1986); Duncan and Pasik-Duncan (1990); Borkar and Varaiya (1979); Becker et al. (1985); Fomin (1991), and multiple-input, multiple-output systems in Tao (2003). These addressed adaptive control architectures and algorithms for a range of dynamic systems, with either

[2]In hindsight, the MIT-rule can be viewed as a partial Lyapunov function, as it only included an $\mathcal{L}_2$-norm of the performance error in its cost function.

6

And basically, Widrow’s Adaline filters is one of the big, big foundations for neural networks deep and otherwise. So, these Adaline filters became the foundations for what became neural networks. So, so very, very old, very, very classical ones. So, anyway, so very, very interesting. Now, it, it was of course realized soon that this MIT rule proposed by Whitaker can result in instability.

And of course, then there was a requirement to study the stability of adaptive controllers. I mean, it was not just enough to figure out a way to adapt for the parameters, but of course to find a way to formulate a stability framework for analysis and synthesis of adaptive systems. And this is where there was a discussion or there was a lot of delving into dynamics, dynamical system notions.

Many, many authors, I would, I would again point you more towards Narendra and Kudva, who, are one of the seminal researchers in this area. And of course. Lyapunov's method was suggested in lieu of the gradient descent approach. And, and this is what we learned in this course. Most of what we did in adaptive control, in this course, was we got the use of Lyapunov’s method.

And, and you can see the origin for this is going as far as the 70s. So, not just Narendra, of course, this, this in, in a deterministic discrete time setting, this was being parallelly addressed by Yakubovic in Russia, so 68 and 72. So, so a lot of a lot of different works. I mean there was work by Landau, you have Wittenmark amongst others.

I mean you will find these names in very, very good books in adaptive control, Narendra and Annaswamy, Astrom and Wittenmark, Ioannou and Sun, so Sastry and Bodson, Tao, Krstic, so, Fradkov, Kumar, Varaiya, so, who did the stochastic systems. So, so these basically addressed adaptive control architectures for a range of dynamical systems, with of course, full or partial measurements and so on and so forth.

(Refer Slide Time: 15:18)

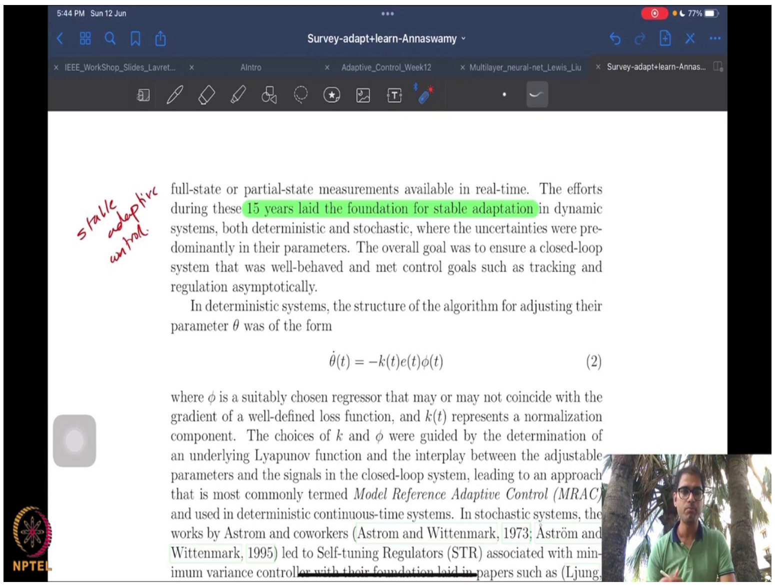

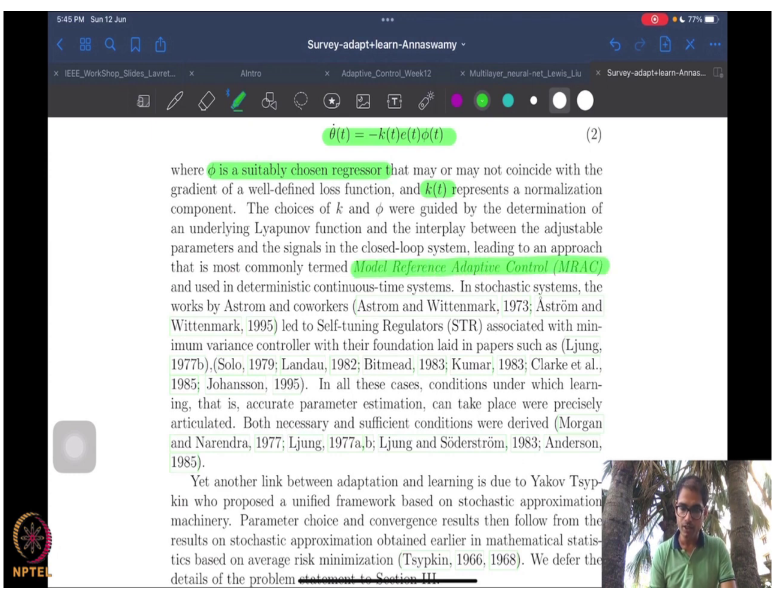

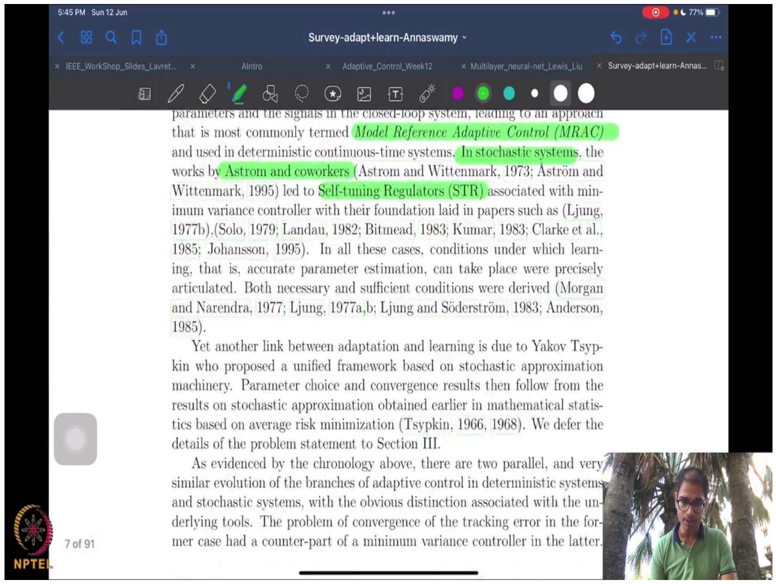

Survey-adapt+learn-Annaswamy

parameters and the signals in the closed-loop system, leading to an approach that is most commonly termed *Model Reference Adaptive Control (MRAC)* and used in deterministic continuous-time systems. In stochastic systems, the works by Astrom and coworkers (Astrom and Wittenmark, 1973; Åström and Wittenmark, 1995) led to Self-tuning Regulators (STR) associated with minimum variance controller with their foundation laid in papers such as (Ljung, 1977b),(Solo, 1979; Landau, 1982; Bitmead, 1983; Kumar, 1983; Clarke et al., 1985; Johansson, 1995). In all these cases, conditions under which learning, that is, accurate parameter estimation, can take place were precisely articulated. Both necessary and sufficient conditions were derived (Morgan and Narendra, 1977; Ljung, 1977a,b; Ljung and Söderström, 1983; Anderson, 1985).

Yet another link between adaptation and learning is due to Yakov Tsypkin who proposed a unified framework based on stochastic approximation machinery. Parameter choice and convergence results then follow from the results on stochastic approximation obtained earlier in mathematical statistics based on average risk minimization (Tsypkin, 1966, 1968). We defer the details of the problem statement to Section III.

As evidenced by the chronology above, there are two parallel, and very similar evolution of the branches of adaptive control in deterministic systems and stochastic systems, with the obvious distinction associated with the underlying tools. The problem of convergence of the tracking error in the former case had a counter-part of a minimum variance controller in the latter.

7 of 91

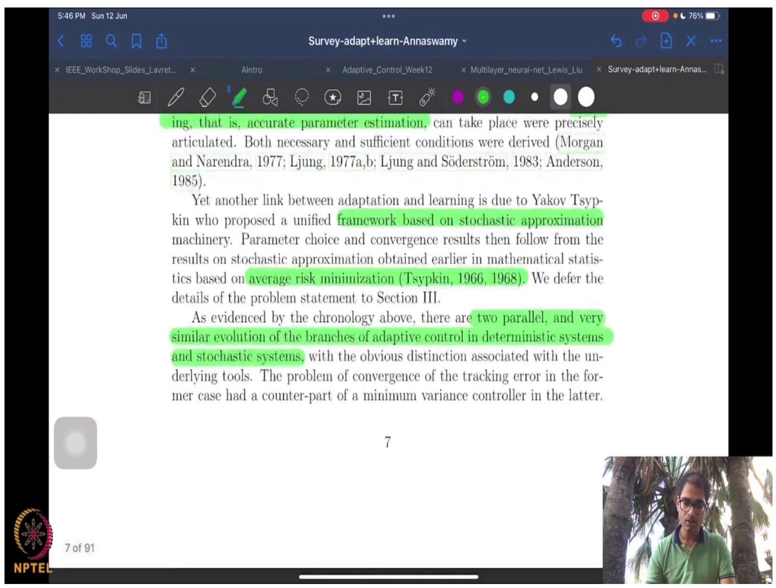

Survey-adapt+learn-Annaswamy

ing, that is, accurate parameter estimation, can take place were precisely articulated. Both necessary and sufficient conditions were derived (Morgan and Narendra, 1977; Ljung, 1977a,b; Ljung and Söderström, 1983; Anderson, 1985).

Yet another link between adaptation and learning is due to Yakov Tsypkin who proposed a unified framework based on stochastic approximation machinery. Parameter choice and convergence results then follow from the results on stochastic approximation obtained earlier in mathematical statistics based on average risk minimization (Tsypkin, 1966, 1968). We defer the details of the problem statement to Section III.

As evidenced by the chronology above, there are two parallel, and very similar evolution of the branches of adaptive control in deterministic systems and stochastic systems, with the obvious distinction associated with the underlying tools. The problem of convergence of the tracking error in the former case had a counter-part of a minimum variance controller in the latter.

7

7 of 91

So, this was essentially, I think you can trace it to about, I would say, I mean there is of course much more recent 2003, but you can trace it from 1975 to ’85 ’90 and so on. So, almost 15 years of work actually laid the foundation for stable adaptive control. So, I would use the word stable adaptive control, came about. Until then, it was more like, there was a update law for the parameters which was coming out of some kind of an optimization rule but it did not really guarantee stability of the system.

So, anyway. So, this is, so, you are I mean we are already aware of this kind of a structure for adjusting the parameter. This is what came out of most of the adaptive control laws. And in fact, we have also looked at update laws which looked pretty much like this equation over here. And then of course, there was this, where phi is the suitably chosen regressor, and, and you have k with some kind of a normalization component or a weight, if you may, which we call the adaptive control gain.

And then of course, there was the development of the notion of model reference adaptive control, again something we studied. In the stochastic systems domain, based on the work by Astrom and coworkers, you, we, the outcome of that was the development of what is called self-tuning regulators.

And again, this is also been worked on by many other authors like Landan, Kumar, Clarke, et cetera. So, of course, then there were connections, in all these cases, there were a requirement for learning, that is accurate parameter estimation, just like in signal processing, Adaline filters in neural networks you had the requirement for learning parameters and weights, here also you are, we know we have parameters if we wanted to estimate. So, we in fact saw that in our, in several of our lectures.

So, this sort of connection was also made in stochastic approximation framework by Tsypkin in his work. So, so the idea is that there was there was this parallel and similar evolution into branches of adaptive control in deterministic and stochastic systems. So, therefore they were distinct tools, but this development was happening sort of parallel.

(Refer Slide Time: 18:08)

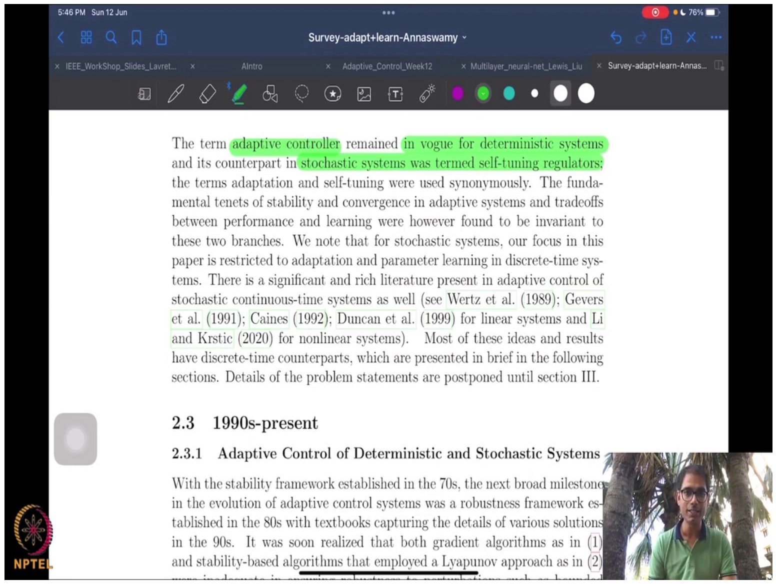

The term adaptive controller remained in vogue for deterministic systems and its counterpart in stochastic systems was termed self-tuning regulators; the terms adaptation and self-tuning were used synonymously. The fundamental tenets of stability and convergence in adaptive systems and tradeoffs between performance and learning were however found to be invariant to these two branches. We note that for stochastic systems, our focus in this paper is restricted to adaptation and parameter learning in discrete-time systems. There is a significant and rich literature present in adaptive control of stochastic continuous-time systems as well (see Wertz et al. (1989); Gevers et al. (1991); Caines (1992); Duncan et al. (1999) for linear systems and Li and Krstic (2020) for nonlinear systems). Most of these ideas and results have discrete-time counterparts, which are presented in brief in the following sections. Details of the problem statements are postponed until section III.

2.3 1990s-present

2.3.1 Adaptive Control of Deterministic and Stochastic Systems

With the stability framework established in the 70s, the next broad milestone in the evolution of adaptive control systems was a robustness framework established in the 80s with textbooks capturing the details of various solutions in the 90s. It was soon realized that both gradient algorithms as in (1) and stability-based algorithms that employed a Lyapunov approach as in (2)

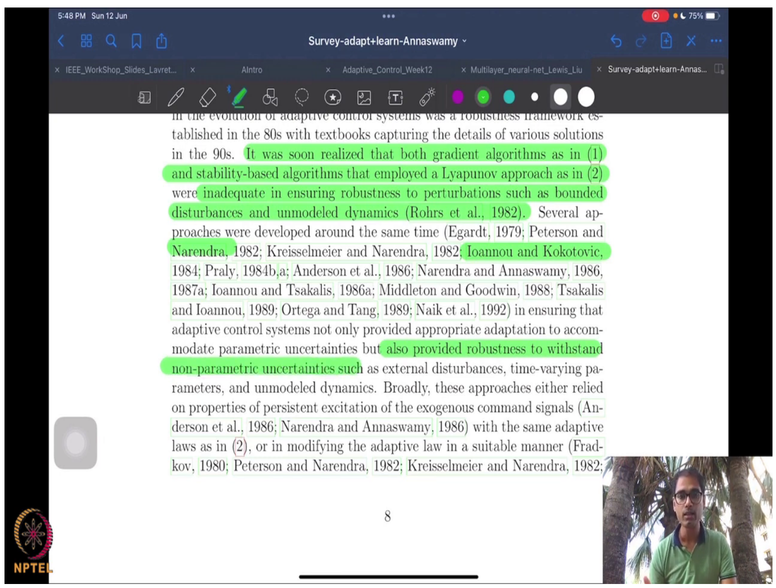

tablished in the 80s with textbooks capturing the details of various solutions in the 90s. It was soon realized that both gradient algorithms as in (1) and stability-based algorithms that employed a Lyapunov approach as in (2) were inadequate in ensuring robustness to perturbations such as bounded disturbances and unmodeled dynamics (Rohrs et al., 1982). Several approaches were developed around the same time (Egardt, 1979; Peterson and Narendra, 1982; Kreisselmeier and Narendra, 1982; Ioannou and Kokotovic, 1984; Praly, 1984b,a; Anderson et al., 1986; Narendra and Annaswamy, 1986, 1987a; Ioannou and Tsakalis, 1986a; Middleton and Goodwin, 1988; Tsakalis and Ioannou, 1989; Ortega and Tang, 1989; Naik et al., 1992) in ensuring that adaptive control systems not only provided appropriate adaptation to accommodate parametric uncertainties but also provided robustness to withstand non-parametric uncertainties such as external disturbances, time-varying parameters, and unmodeled dynamics. Broadly, these approaches either relied on properties of persistent excitation of the exogenous command signals (Anderson et al., 1986; Narendra and Annaswamy, 1986) with the same adaptive laws as in (2), or in modifying the adaptive law in a suitable manner (Fradkov, 1980; Peterson and Narendra, 1982; Kreisselmeier and Narendra, 1982;

8

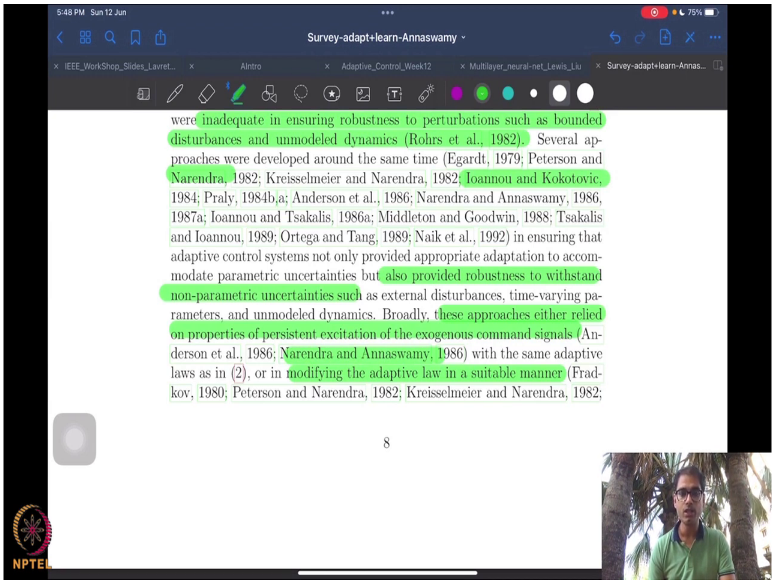

So, of course this, this adaptive control terminology remained in vogue for deterministic systems but then in the stochastic systems framework, it was termed as self-tuning regulators. We of course, did not cover any stochastic systems self tuning regulators in this course, but I would sort of strongly encourage you to look at that.

Then, of course 90s to present, there were many, a lot of new novel flavors in adaptive control of deterministic and stochastic systems. So, one of the things that was evident, was that both the gradient algorithm and the stability-based algorithm in the Lyapunov approach had robustness issues, perturbations such as bounded disturbance and unknown dynamics. This was figured out in 80s itself. And, so, therefore, there were several approaches that were developed in order to nullify this issue.

This was due to Narendra, we saw, we saw the sigma modification which was due to Ioannou and Kokotovic, which was, sorry, which was due to Narendra, and the epsilon modification which was due to Ioannou and Kokotovic. So, so basically, we wanted to ensure that these adaptive algorithms also provided robustness to withstand non-parametric uncertainties, such as bounded external disturbances.

So, so either of course, they really, I mean there was this connection to persistent excitation of exogenous signals or of course, we modify the adapter controller in a suitable manner, which is what is the sigma epsilon modification. So, if you have persistent excitation, none of this was ever a problem, there was always nice convergence properties and all that. We saw this Narendra and Annaswami's results and so on. And if you did not have persistent excitation, we had to modify the adaptive law by sigma and epsilon multiplication. So, this is what we learned, and this is what was a sequence of development.

(Refer Slide Time: 20:16)

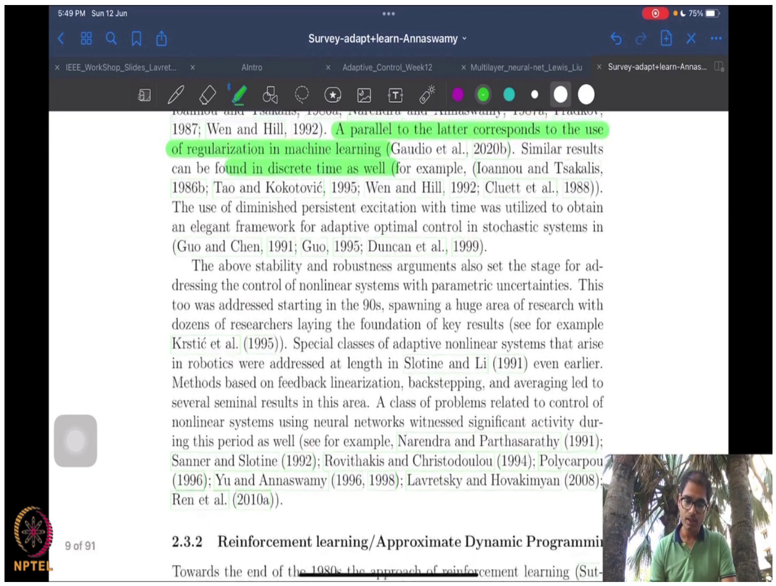

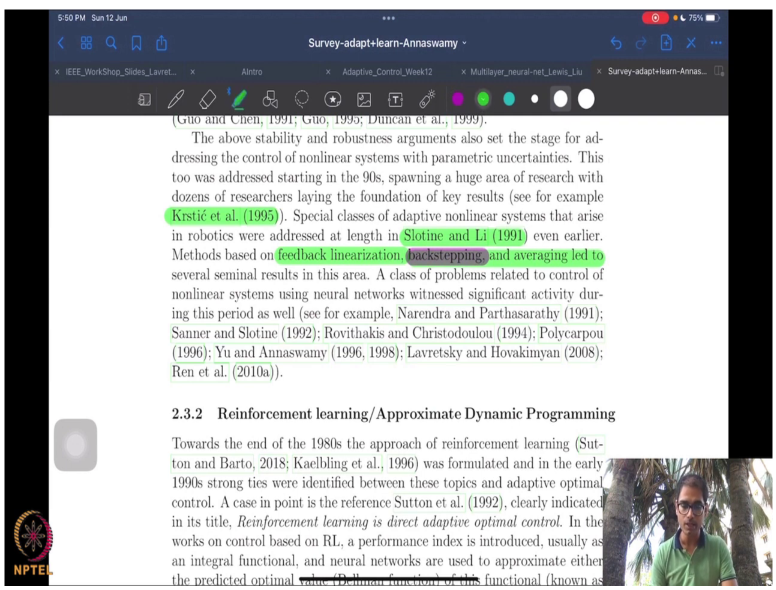

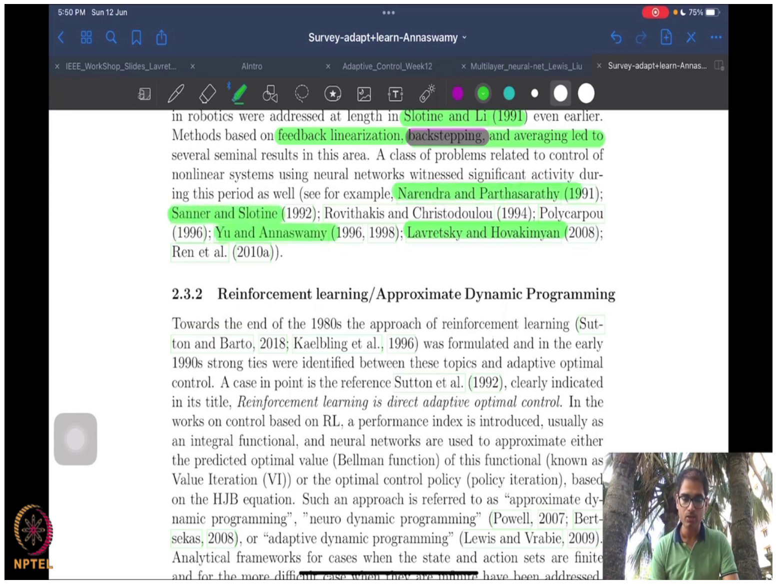

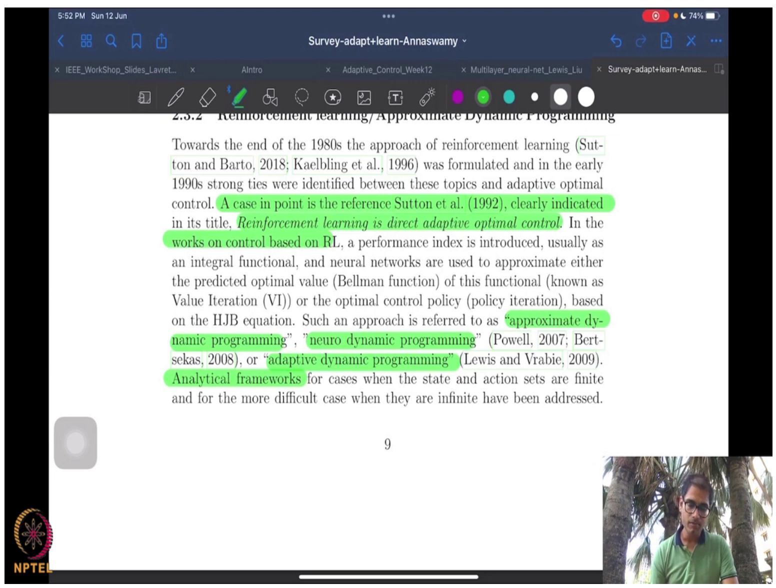

So, again this, there was again a parallel in, in, in the learning framework in the learning domain on regularization and machine learning, and then again similar results in discrete time as well. So, anyway, so, so this sort of highlights the parallels that you have in the discrete, continuous, stochastic framework that is happening. So, of course then, there was a large, large requirement for addressing parametric uncertainties that appeared in different, sort of, cases, different sort of situations. And a lot of special case, non-linear systems were started to be considered in the 90s.

So, this was started with Krstic's work and then Slotine and Li. So, a lot of focus on non-linear system because, of course, most real systems are non-linear. So, just focusing on linear systems or linearizing systems was not always the solution. And these were based on methods such as feedback linearization, backstepping and averaging. So, we did of course look at the backstepping based results, we did look at the backstepping based results, mostly due to Krstic and of course some of them due to Slotine and Li.

But then, there were many, many authors there is Lavrestsky and Hovakimyan, there is Slotine, there is Annaswamy, Narendra, Parthasarathy. So, basically, a lot of authors who did contribute to this area. So, by the end of 80s, there was this, 80s and of course and 90s, there was a focus on what is called reinforcement learning or approximate dynamic programming.

So, and this sort of trend was highlighted by this paper in 92 by Satton et al, which is basically reinforcement learning in direct adaptive control. So, in usually, the works on control based on reinforcement learning typically a performance index is introduced. And then the neural network is used to approximate, either the predicted optimal value of this functional or the optimal control policy.

So, basically a neural network is a function approximator, and we will see something later on. And the idea is that in this performance index which is like a cost, which is usually an integral, integral cost, the idea is that we, in reinforcement learning, is to use a function approximator to figure out the optimal value or the optimal policy that will give the optimal value.

So, so these, there are many a lot of different terminology for these, and these are called either an approximate DP, dynamic programming or neurodynamic programming or adaptive dynamic programming. So, all of these are sort of terminologies that have been parallelly used. The thing to remember is that it is very difficult to give good analytical results in this setup because of obviously the complications arising due to dynamic programming.

(Refer Slide Time: 23:53)

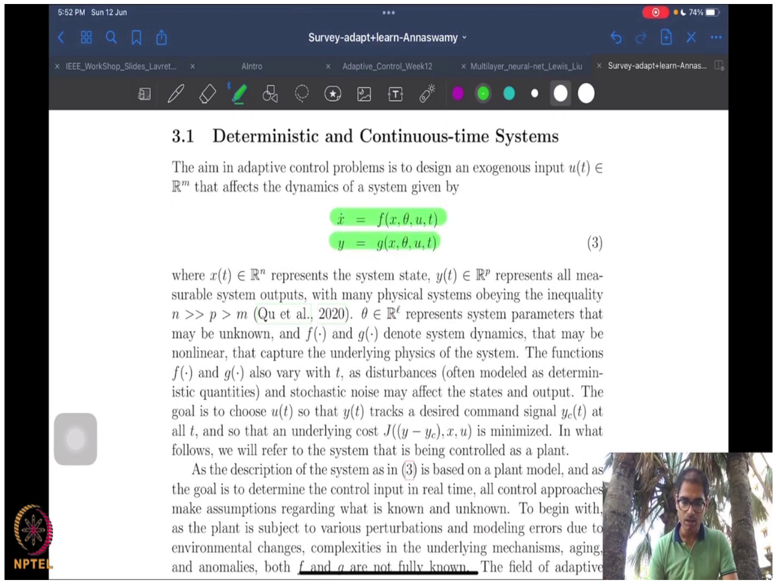

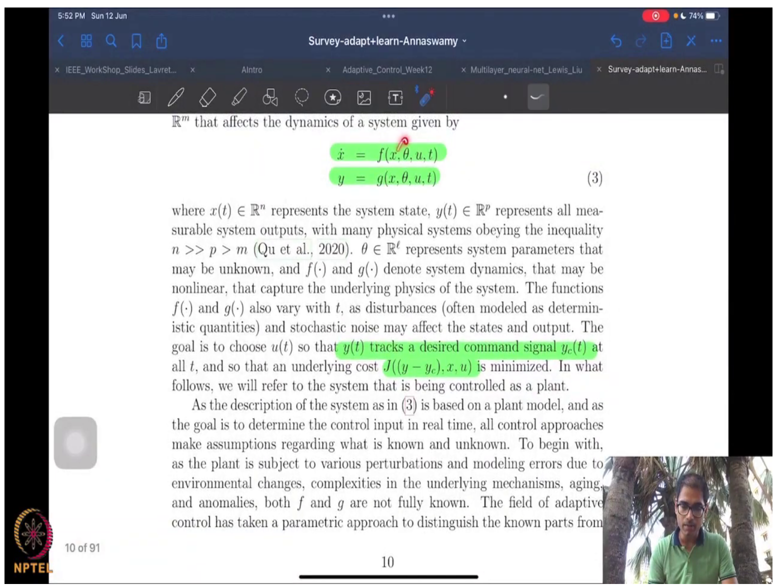

So, different sort of, as we are aware, I mean we have seen, there are many different kinds of problem statements, and I mean this is the typical deterministic continuous time system where you want y to track some kind of a desired y c. I mean, or you may have some kind of a cost to be minimized, we usually do the first. I mean, we have looked at y tracking some y c. And then of course we have these unknown parameters theta right here. So, this is the sort of framework we have.

(Refer Slide Time: 24:27)

the unknown. In particular, it is assumed that f is a known function, while the parameter θ is unknown. A real time control input is then designed so as to ensure that the tracking goals are achieved by including an adaptive component that attempts to estimate the parameters online. A linearized version of the problem in (3) is of the form

$$y = W(s, \theta)[u] \tag{4}$$

where s denotes the differential operator d/dt, $W(s, \cdot)$ is a rational operator of s, and θ is an unknown parameter, and the goals of tracking and regulation are the same as above.

In the following subsections, four broad categories of subproblems that have been addressed in the context of adaptive control in deterministic continuous-time systems are described.

3.1.1 Boundedness and real-time decision making

As mentioned above, the control goal is to ensure that

$$\lim_{t\to\infty} e(t) = 0 \tag{5}$$

where $e(t) = y(t) - y_c(t)$. As these decisions are required to be made in real time, the focus of the solutions is to have them lead to a closed-loop dynamic

$$\lim_{t\to\infty} e(t) = 0 \tag{5}$$

where $e(t) = y(t) - y_c(t)$. As these decisions are required to be made in real time, the focus of the solutions is to have them lead to a closed-loop dynamic system that has bounded solutions at all time t and a desired asymptotic behavior. The central question, therefore, is if this can be ensured even when there are parametric uncertainties in θ and several other non-parametric uncertainties that may due to unmodeled dynamics, disturbances, and the like. Once this is guaranteed, the question of learning, in the form of parameter convergence, is addressed. As a result, *control for learning* is a central question that is pursued in the class of problems addressed in adaptive control rather than *learning for control* (Krstic, 2021).

3.1.2 Certainty Equivalence Principle and Adaptive Control Solutions

The well known certainty equivalence principle (CEP) proceeds with the following mantra: first, optimize under perfect foresight, then substitute optimal estimates for unknown values. This philosophy underlies all adaptive control solutions by first determining a controller structure that leads to an

11

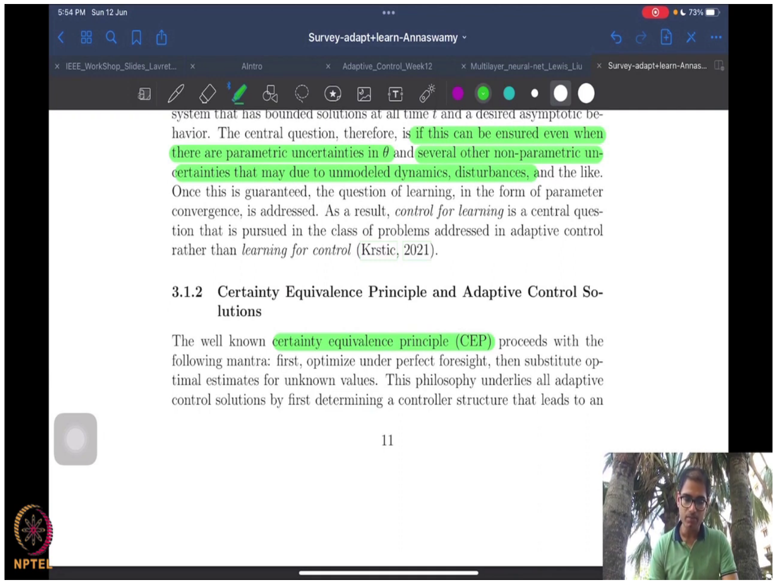

And, and we, we have also looked at some kind of linearized versions of these with, with some regressor parameter type structure. Of course, the goal is usually to send the error to zero over time, and then, this is, what is the idea. Now, and the question is that if this, this can be ensured, even when there are parametric uncertainties in theta, and also non-parametric uncertainties and several non-parametric uncertainties like modeling disturb, modeling, unmodeled dynamics disturbances et cetera.

So, learning is always an important aspect of an adaptive control because we are always trying to learn some parameter theta. We of course, look at this, the certainty equivalence principle which forms the basis of most of what we discussed in adaptive control. So, so that is the idea. I mean, we basically use a certainty equivalence principle is what came about of this.

(Refer Slide Time: 25:33)

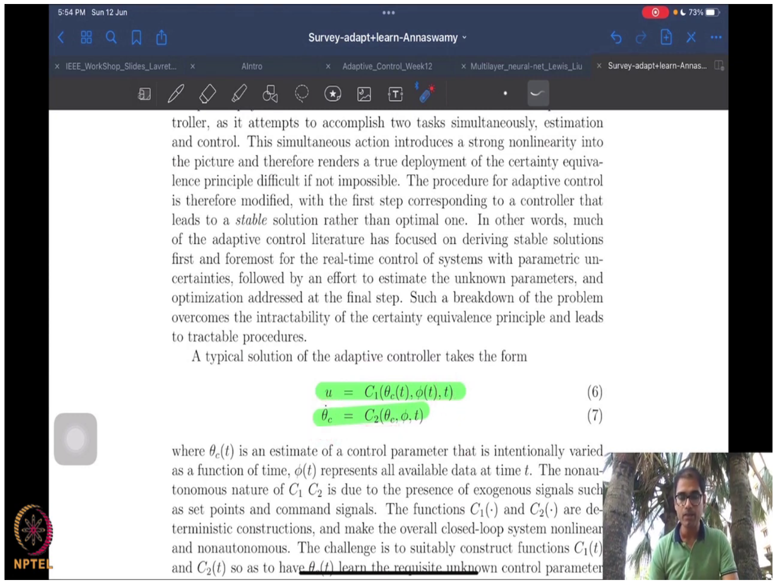

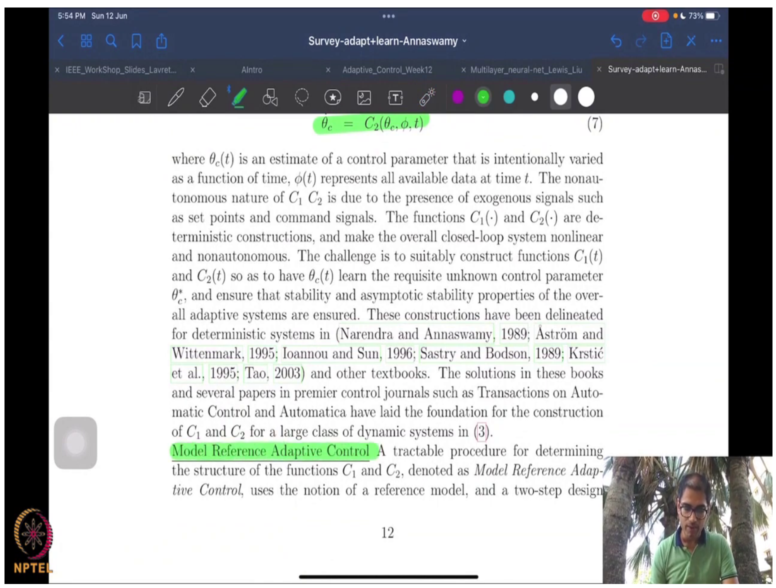

And a standard solution of an adaptive controller looks something like this. I mean you have a controller which depends on some estimates of the parameters, some regressor and time possibly, and then there is an update law, the parameter law, which of course again depends on the, possibly depends on the estimate itself and the regressor. So, so this is sort of the structure of what we have done.

(Refer Slide Time: 26:06)

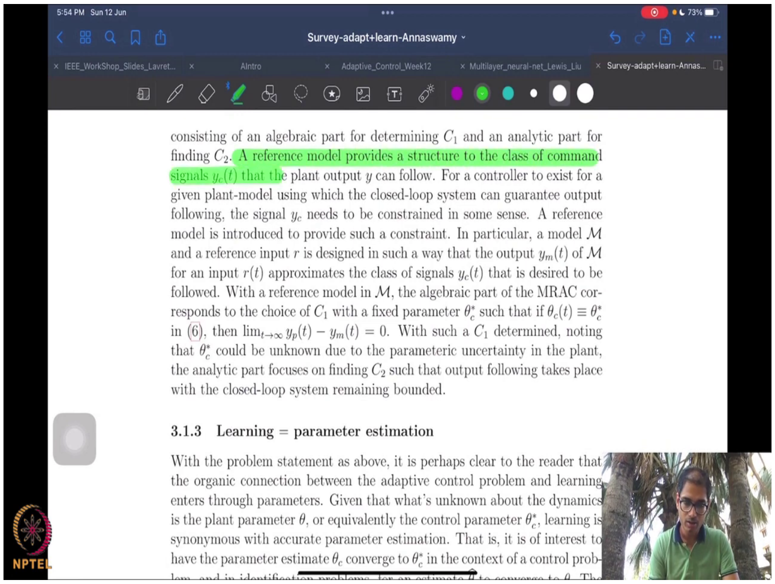
consisting of an algebraic part for determining C_1 and an analytic part for finding C_2. A reference model provides a structure to the class of command signals $y_c(t)$ that the plant output y can follow. For a controller to exist for a given plant-model using which the closed-loop system can guarantee output following, the signal y_c needs to be constrained in some sense. A reference model is introduced to provide such a constraint. In particular, a model $\mathcal{M}$ and a reference input r is designed in such a way that the output $y_m(t)$ of $\mathcal{M}$ for an input $r(t)$ approximates the class of signals $y_c(t)$ that is desired to be followed. With a reference model in $\mathcal{M}$, the algebraic part of the MRAC corresponds to the choice of C_1 with a fixed parameter θ_c^* such that if $\theta_c(t) \equiv \theta_c^*$ in (6), then $\lim_{t\to\infty} y_p(t) - y_m(t) = 0$. With such a C_1 determined, noting that θ_c^* could be unknown due to the parameteric uncertainty in the plant, the analytic part focuses on finding C_2 such that output following takes place with the closed-loop system remaining bounded.

3.1.3 Learning = parameter estimation

With the problem statement as above, it is perhaps clear to the reader that the organic connection between the adaptive control problem and learning enters through parameters. Given that what's unknown about the dynamics is the plant parameter θ, or equivalently the control parameter θ_c^*, learning is synonymous with accurate parameter estimation. That is, it is of interest to have the parameter estimate θ_c converge to θ_c^* in the context of a control prob-

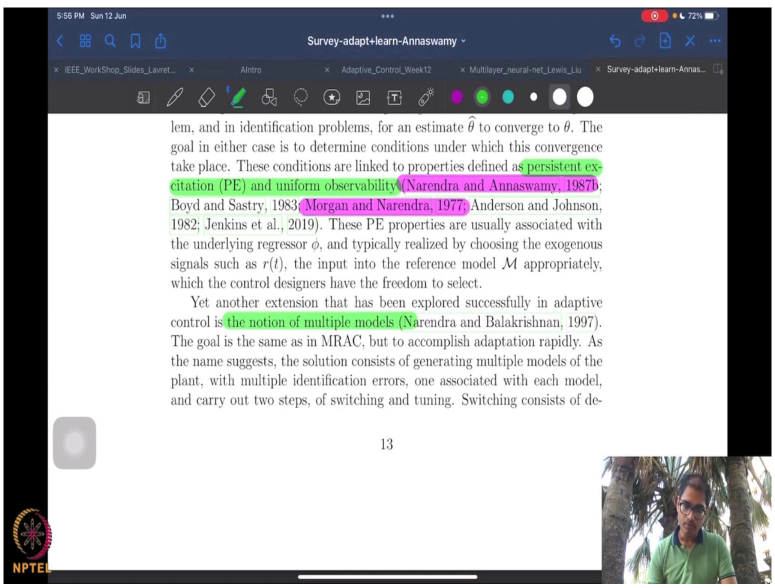
lem, and in identification problems, for an estimate $\widehat{\theta}$ to converge to θ. The goal in either case is to determine conditions under which this convergence take place. These conditions are linked to properties defined as persistent excitation (PE) and uniform observability (Narendra and Annaswamy, 1987b; Boyd and Sastry, 1983; Morgan and Narendra, 1977; Anderson and Johnson, 1982; Jenkins et al., 2019). These PE properties are usually associated with the underlying regressor ϕ, and typically realized by choosing the exogenous signals such as $r(t)$, the input into the reference model $\mathcal{M}$ appropriately, which the control designers have the freedom to select.

Yet another extension that has been explored successfully in adaptive control is the notion of multiple models (Narendra and Balakrishnan, 1997). The goal is the same as in MRAC, but to accomplish adaptation rapidly. As the name suggests, the solution consists of generating multiple models of the plant, with multiple identification errors, one associated with each model, and carry out two steps, of switching and tuning. Switching consists of de-

13

Then you of course have the linear version, it is the model reference adaptive control where instead of following a reference signal, you, you follow a reference model. So, that is, we also, this is also something that we have seen. One of the points that this article tries to make is that learning is essentially equal to parameter estimation. In adaptive control, very much so.

In typical setting of learning for control also, this is the case. More often than not, whether it be reinforcement learning or deep learning, you are trying to learn some parameters. So, and, and we did, I mean, many researchers established when such learning can happen. And these are connected to the persistence of excitation condition and uniform observability condition. We did look at this extensively, we looked at the work here.

And in fact, in fact mostly, the Morgan Narendra paper in 77, so this is what we looked at. So, this article is what we focused on mostly, the results that you saw, but we also saw results, further results. Then there is, there have been some results on the notion of multiple models in adaptive control also. But anyway, that is an again another parallel sort of a framework, is what I would say.

(Refer Slide Time: 27:31)

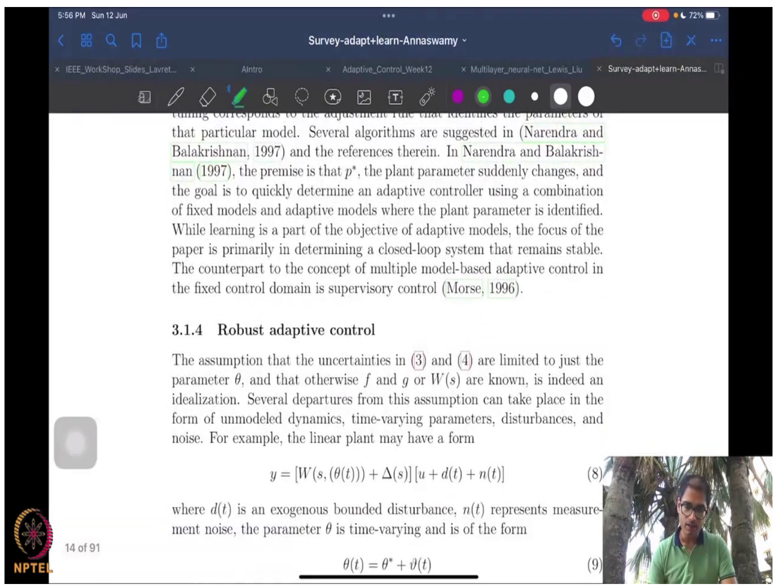

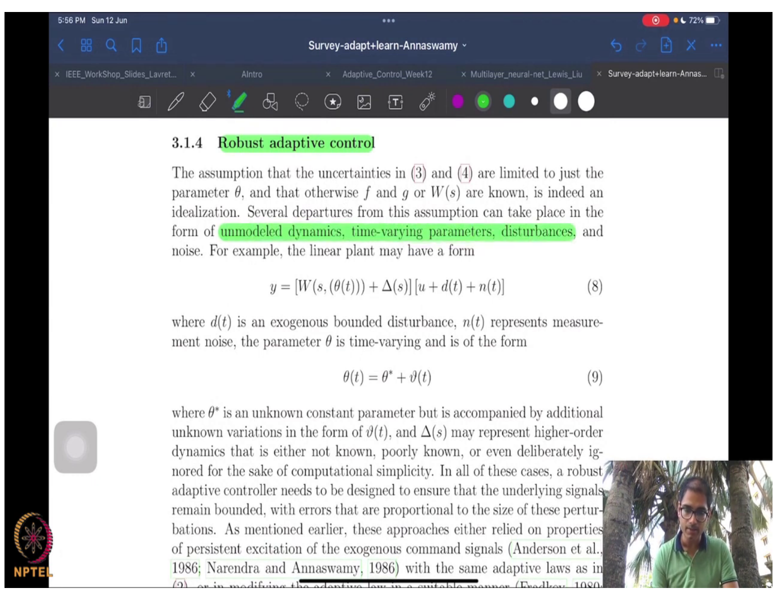

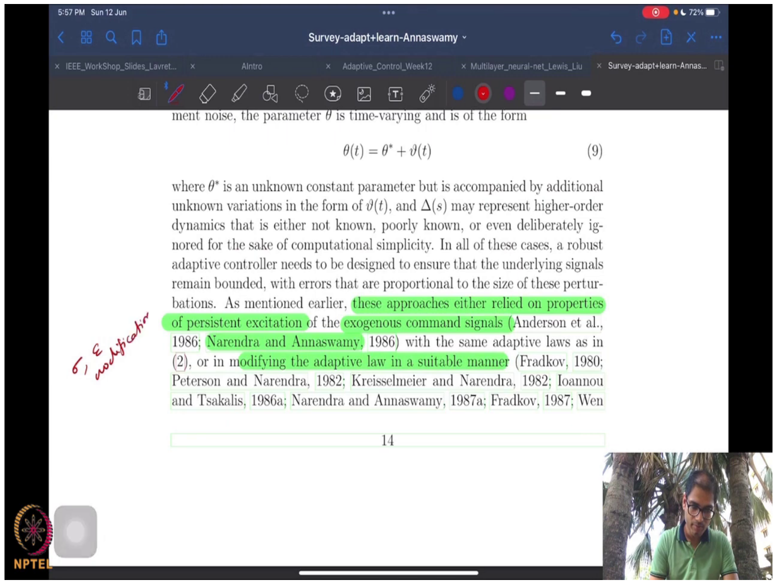

So, so, yeah, so, several, several interesting results on when parameter learning can happen, and we did look at this. I mean again, so, these are results from 70s and 80s and we did look at some of these. Beyond that, there was a notion of robustness in adaptive control also. So, the idea being that if there are unmodeled dynamics, time varying parameters and disturbances how does the plant look.

So, in, in, in the frequency domain, you would have something like this where of course d is some kind of a disturbance and n is some kind of a noise, and theta is some kind of a time varying quantity. So, no, again, most of these results, these approaches on, relied on properties of persistent excitation of exogenous commands. This is again Anderson, Narendra et cetera, or you modify the adaptive law. So, these are again the results that connect to, I would, this is basically sigma epsilon modification. This is what helped impart robustness.

(Refer Slide Time: 28:54)

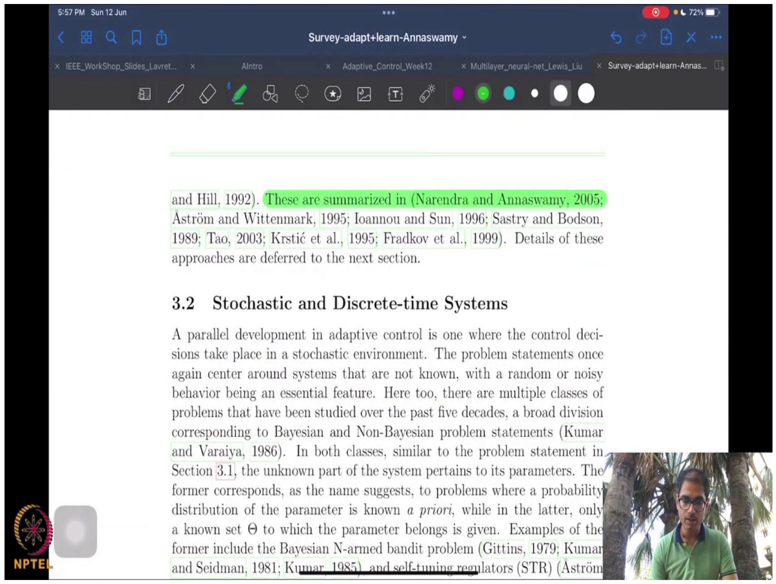

and Hill, 1992). These are summarized in (Narendra and Annaswamy, 2005; Åström and Wittenmark, 1995; Ioannou and Sun, 1996; Sastry and Bodson, 1989; Tao, 2003; Krstić et al., 1995; Fradkov et al., 1999). Details of these approaches are deferred to the next section.

3.2 Stochastic and Discrete-time Systems

A parallel development in adaptive control is one where the control decisions take place in a stochastic environment. The problem statements once again center around systems that are not known, with a random or noisy behavior being an essential feature. Here too, there are multiple classes of problems that have been studied over the past five decades, a broad division corresponding to Bayesian and Non-Bayesian problem statements (Kumar and Varaiya, 1986). In both classes, similar to the problem statement in Section 3.1, the unknown part of the system pertains to its parameters. The former corresponds, as the name suggests, to problems where a probability distribution of the parameter is known *a priori*, while in the latter, only a known set Θ to which the parameter belongs is given. Examples of the former include the Bayesian N-armed bandit problem (Gittins, 1979; Kumar and Seidman, 1981; Kumar, 1985), and self-tuning regulators (STR) (Åström

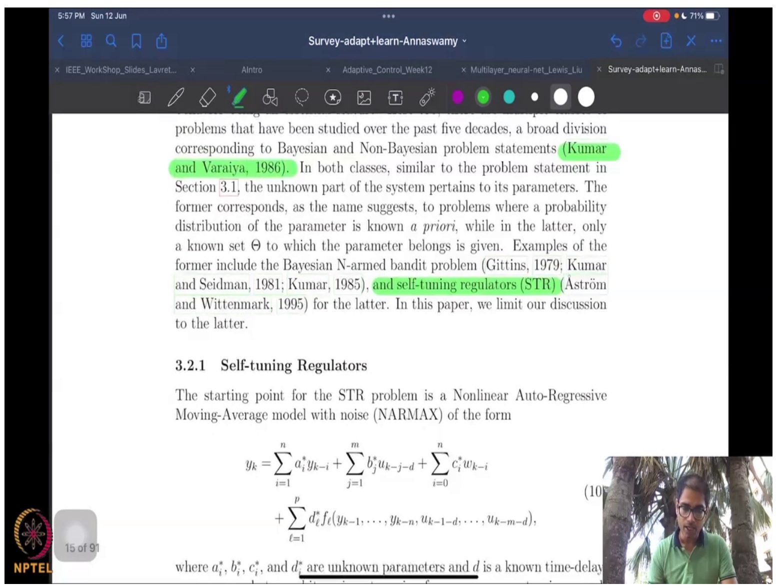

problems that have been studied over the past five decades, a broad division corresponding to Bayesian and Non-Bayesian problem statements (Kumar and Varaiya, 1986). In both classes, similar to the problem statement in Section 3.1, the unknown part of the system pertains to its parameters. The former corresponds, as the name suggests, to problems where a probability distribution of the parameter is known *a priori*, while in the latter, only a known set Θ to which the parameter belongs is given. Examples of the former include the Bayesian N-armed bandit problem (Gittins, 1979; Kumar and Seidman, 1981; Kumar, 1985), and self-tuning regulators (STR) (Åström and Wittenmark, 1995) for the latter. In this paper, we limit our discussion to the latter.

3.2.1 Self-tuning Regulators

The starting point for the STR problem is a Nonlinear Auto-Regressive Moving-Average model with noise (NARMAX) of the form

$$y_k = \sum_{i=1}^{n} a_i^* y_{k-i} + \sum_{j=1}^{m} b_j^* u_{k-j-d} + \sum_{i=0}^{n} c_i^* w_{k-i} + \sum_{\ell=1}^{p} d_\ell^* f_\ell(y_{k-1}, \dots, y_{k-n}, u_{k-1-d}, \dots, u_{k-m-d}), \quad (10)$$

where a_i^*, b_i^*, c_i^*, and d_ℓ^* are unknown parameters and d is a known time-delay

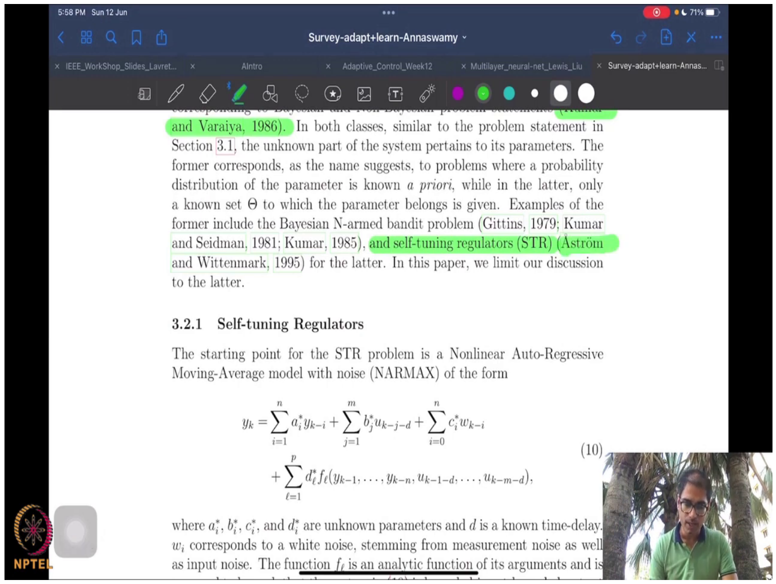

And of course, these are summarized in this result. There is also this, I mean I am not going to discuss this in too much detail again here, but there is also been a lot of development on stochastic and discrete time systems, adaptive control in parallel and as I mentioned these are called, these have been termed as self-tuning regulators and a lot of this work has been due to Kumar and Varaiya. A lot of different were connected to Varaiya, Astrom et cetera. So, if you search these authors, you will usually find some results.

(Refer Slide Time: 29:35)

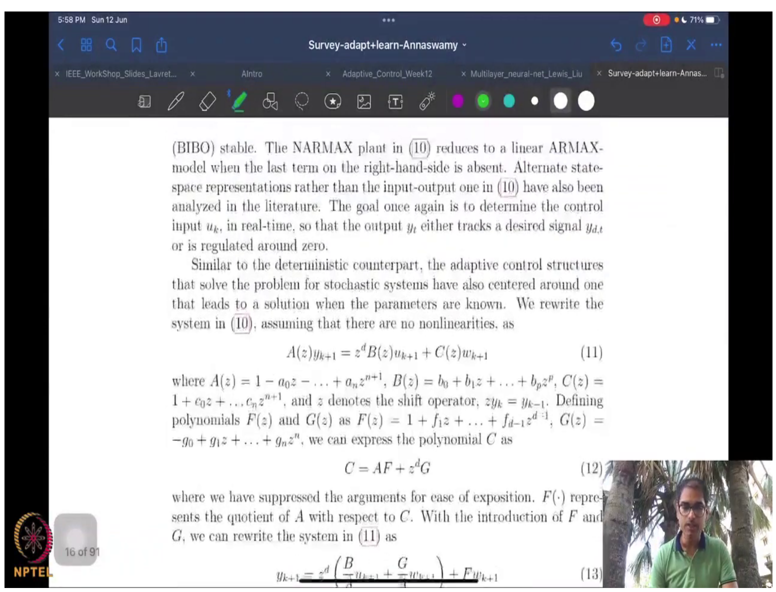

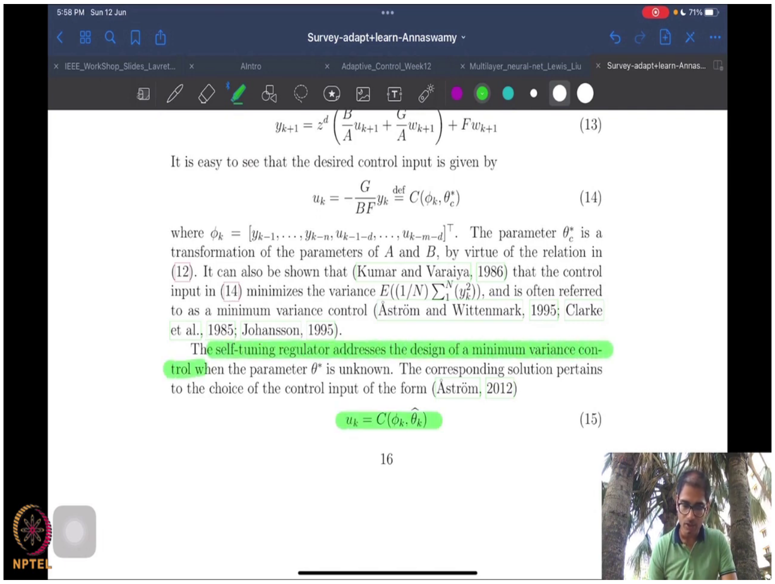

And so, there is also some little bit of a summary on what the self-tuning regulators that is stochastic discrete time adaptive control results is. So, this is sort of what you have. And the idea is that they, the stochastic the self-tuning regulator addresses the design of a minimum variance control which looks something like this.

(Refer Slide Time: 29:53)

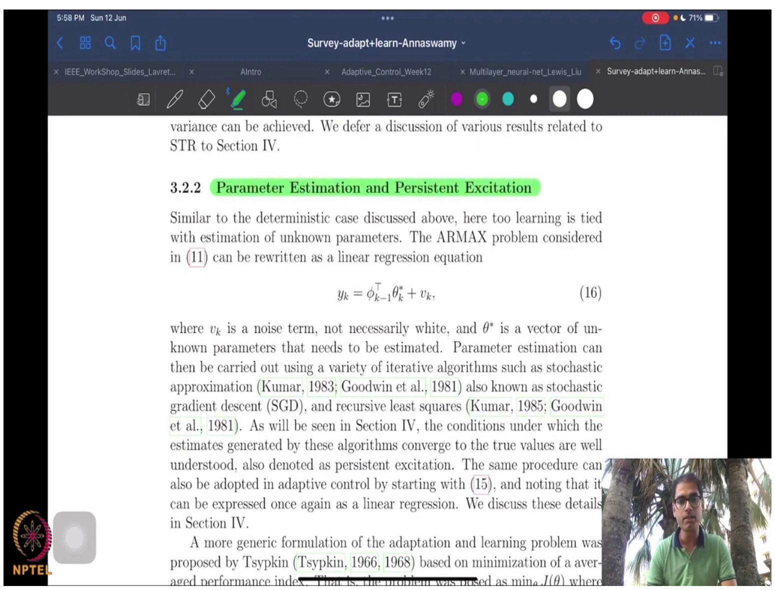

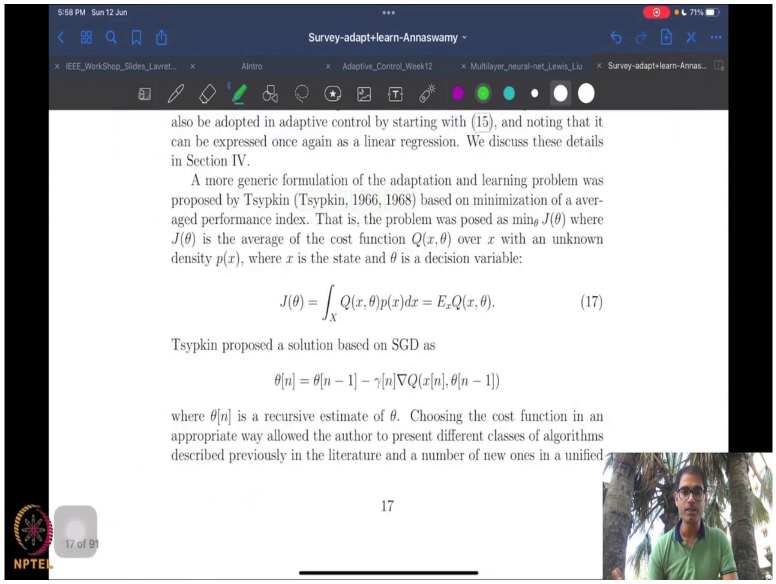

And there is still, the connection of parameter estimation and persistent excitation. I mean that cannot of course be sort of given up, even in this case. I mean, the conditions for persistence may look different but the requirement for persistent excitation still works.

(Refer Slide Time: 30:13)

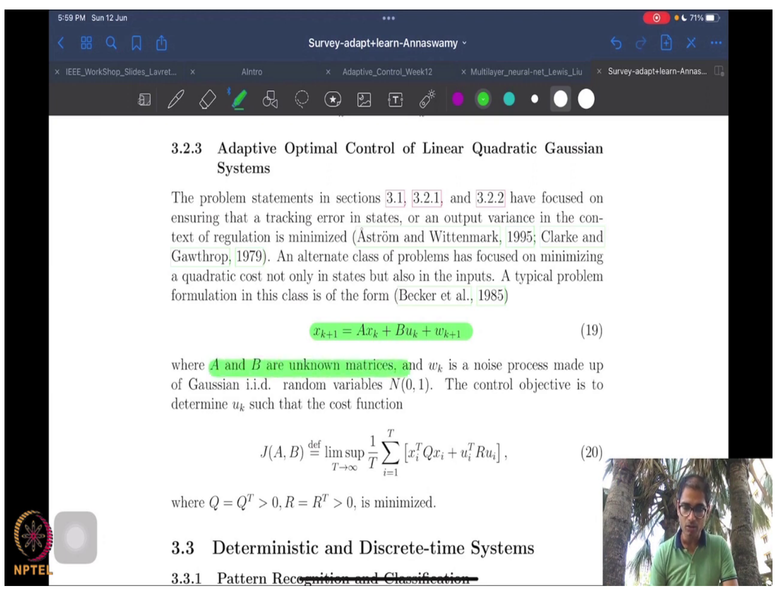

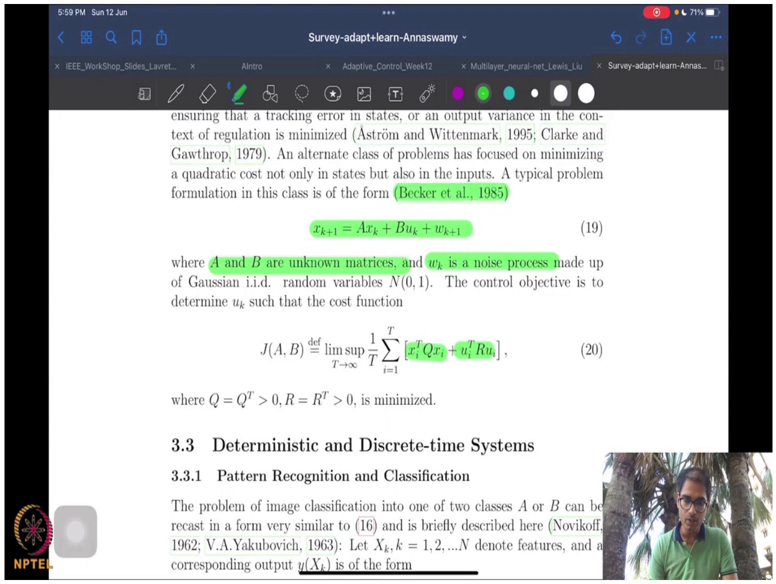

ensuring that a tracking error in states, or an output variance in the context of regulation is minimized (Åström and Wittenmark, 1995; Clarke and Gawthrop, 1979). An alternate class of problems has focused on minimizing a quadratic cost not only in states but also in the inputs. A typical problem formulation in this class is of the form (Becker et al., 1985)

$$x_{k+1} = Ax_k + Bu_k + w_{k+1} \quad (19)$$

where A and B are unknown matrices, and w_k is a noise process made up of Gaussian i.i.d. random variables $N(0,1)$. The control objective is to determine u_k such that the cost function

$$J(A,B) \overset{\text{def}}{=} \limsup_{T\to\infty} \frac{1}{T}\sum_{i=1}^{T}\left[x_i^TQx_i + u_i^TRu_i\right], \quad (20)$$

where $Q = Q^T > 0, R = R^T > 0$, is minimized.

3.3 Deterministic and Discrete-time Systems

3.3.1 Pattern Recognition and Classification

The problem of image classification into one of two classes A or B can be recast in a form very similar to (16) and is briefly described here (Novikoff, 1962; V.A.Yakubovich, 1963): Let $X_k, k = 1, 2, \ldots N$ denote features, and a corresponding output $y(X_k)$ is of the form

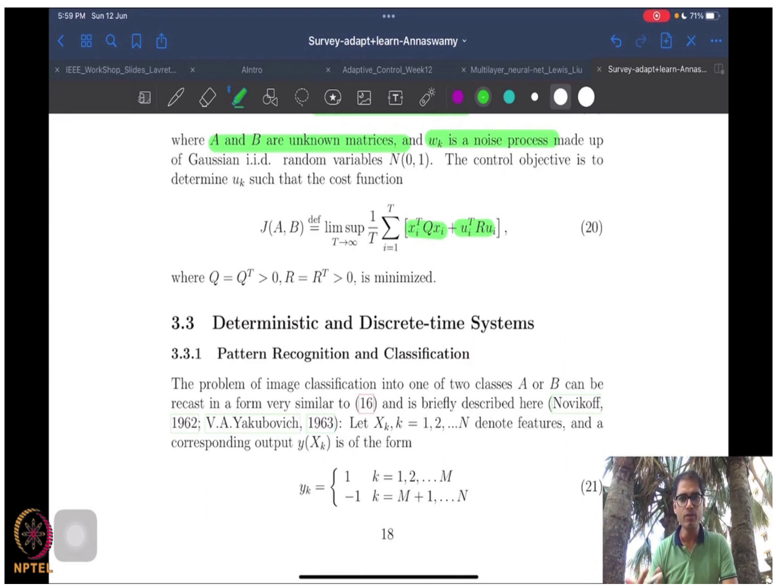

where A and B are unknown matrices, and w_k is a noise process made up of Gaussian i.i.d. random variables $N(0,1)$. The control objective is to determine u_k such that the cost function

$$J(A,B) \overset{\text{def}}{=} \limsup_{T\to\infty} \frac{1}{T}\sum_{i=1}^{T}\left[x_i^TQx_i + u_i^TRu_i\right], \quad (20)$$

where $Q = Q^T > 0, R = R^T > 0$, is minimized.

3.3 Deterministic and Discrete-time Systems

3.3.1 Pattern Recognition and Classification

The problem of image classification into one of two classes A or B can be recast in a form very similar to (16) and is briefly described here (Novikoff, 1962; V.A.Yakubovich, 1963): Let $X_k, k = 1, 2, \ldots N$ denote features, and a corresponding output $y(X_k)$ is of the form

$$y_k = \begin{cases} 1 & k = 1, 2, \ldots M \\ -1 & k = M+1, \ldots N \end{cases} \quad (21)$$

18

So, and then finally, there is also adaptive optimal control for LQG type of situations. And this was also studied, where you have a discrete time system with, with unknown A B matrices, and you want to do some kind of an LQG control, that is, you want to minimize a cost which relies on the state x and also the control signal u, and you want to figure out how to minimize this cost whilst knowing that A and B are unknown, and in the presence of noise.

So, this is also something that has been looked at, so has been looked at and formulated here. So, so these are the sort of rather interesting sets of, or the interesting progress of how adaptive control has moved forward now what we want to do in a subsequent session is to also look a little bit more at how this progress has happened.

So, I mean, forgive me if you may, for introducing you to some history in adaptive control, but what I did want you all to see is how the thought process of this innovation went. I mean, there was a need based on autopilot aircraft flight systems, then came about some optimization-based rules, like the MIT rule for parameter updation, and then it was figured out that these were inadequate so then stability mechanisms were developed.

And then stable adaptive controllers came about, then there was a realization that there is stable adaptive controllers also do not guarantee robustness against disturbance and unmodelled dynamics. So, then there was the notion of sigma epsilon modifications persistent excitation for parameter learning et cetera. And again, in parallel to all of this there was the learning theory that was getting developed Adaline filters, and where also parameter identification was essential persistent excitation was formulated.

And then there was of course the self-tuning regulators in stochastic, no, discrete time systems which did not, we did not look at in our course. But that is also very, very large set of literature that is out there in adaptive control. So, I hope you got a good scene for how this adaptive control literature has evolved. I am going to look at a little bit more at this article before we start into more discussion on the learning aspects in adaptive control. Thank you, and see you again soon.