Machine Learning for Engineering and Science Applications.
Professor Dr. Balaji Srinivasan.
Department of Mechanical Engineering.
Indian Institute of Technology, Madras.
Introduction to the Course
History of Artificial Intelligence.

Welcome to machine learning for engineering and science applications. This is the first video for the course, so in this course we will be looking at, in this video we will be looking at the introduction to the course and a brief history of artificial intelligence through the ages. So let us look at a few things that we are using that are essentially products of machine learning in real life today.

(Refer Slide Time: 0:39)

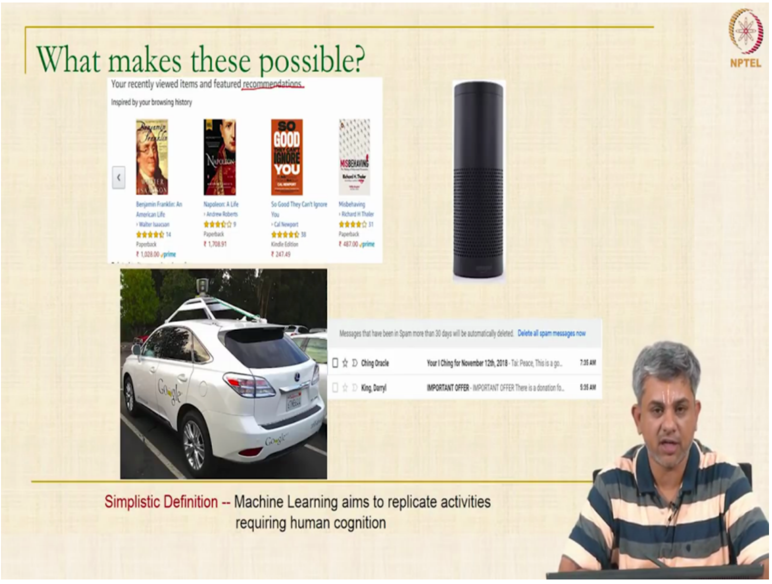

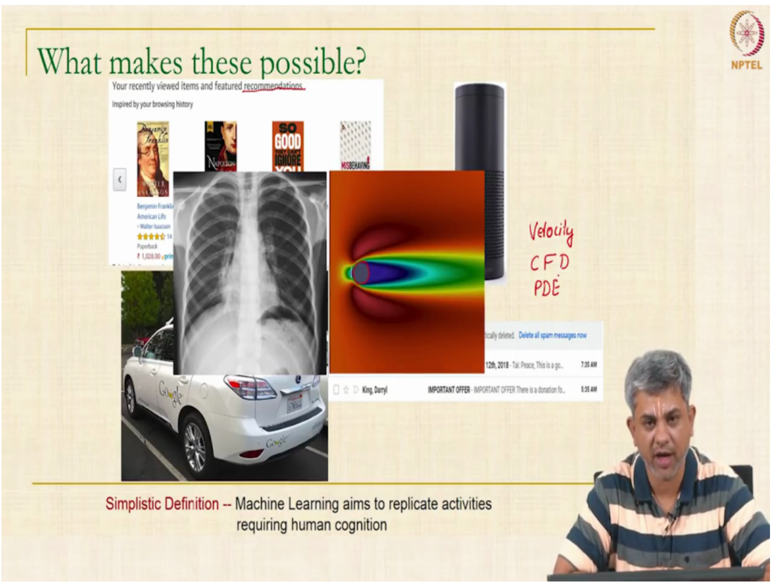

We have seen Amazon's recommendation system. You would have had several such things at using other software or other websites also. Basically you buy a bunch of stuff of this website and it recommends other things, you know maybe books, other books or other product that you might like. This is Amazon's Echo, it is effectively run by a speech recognition engine combined with website searches. We have been all using Google spam filter or any other company spam filter but Google works really really well as a part of our mail system it works very seamlessly nowadays.

This is Google's Lexus which is a self driving car. So now what is common between all these, all these 4, is essentially all of them use as an essential part of their algorithm, machine learning algorithms. Okay. So machine learning very simplistically speaking is a method or a set of algorithms that you can use to replicate activities typically that require human cognition. So for speech, humans recognise speech very very well, starting from a baby to another, we do it very quickly. We all drive, those of us that drive, drive fairly seamlessly, spam filter, most of us can look at email and almost instantaneously say whether we spam or not.

However in practice, in order to encode it into an algorithm is actually a difficult task. Because a number of rules expands very very rapidly, you cannot say that a mail if it comes from Nigeria is definitely going to be spam, if it involves money it is going to be spam, etc., etc. So there is a finite set of rules that you can make but nonetheless you would like a quick spam filter which works as well as a human being does. So in such circumstances we tend to use a large basket of techniques, these are called Machine learning models, these are very old models, several of them are at least half a century or even sometimes even a century old.

So we will be looking at many of these algorithms but as far as this course is concerned, this course, the course of machine learning can be part of electrical injury, it can be part of computer science, it is a branch of mathematics. We are going to treat it as if it is going to be primarily used for engineering and science applications. Just to give you a couple of applications in mind. When you have an x-ray or you have an MRI, that a radiologist looks like, looks at and stars giving you a diagnosis, okay, like the other tumour here, this is cancerous, not cancerous, etc., etc.

Can we replicate this sort of process using machine learning algorithm? Just like we are able to replicate driving, can we replicate this kind of judgement? Doctor Ganapathy for example is an expert in this field. Another application could be, we will see this in the middle of this course, that this is actually flow past a cylinder. So what we are trying to predict is not this thing is if you have a circular body, cylinder in this case kept within an external flow, you know it starts giving what you see here are velocity contours.

Traditionally even today we use several software and the whole process is called computational fluid dynamics, often abbreviated as CFD. You might use solid mechanics modelling or any continuum based, PDE-based modelling, the question is can we find out models that can do this more rapidly using machine learning. And we will see that indeed a part of it is possible and this is an exciting field. So our aim in this course is the following.

(Refer Slide Time: 4:31)

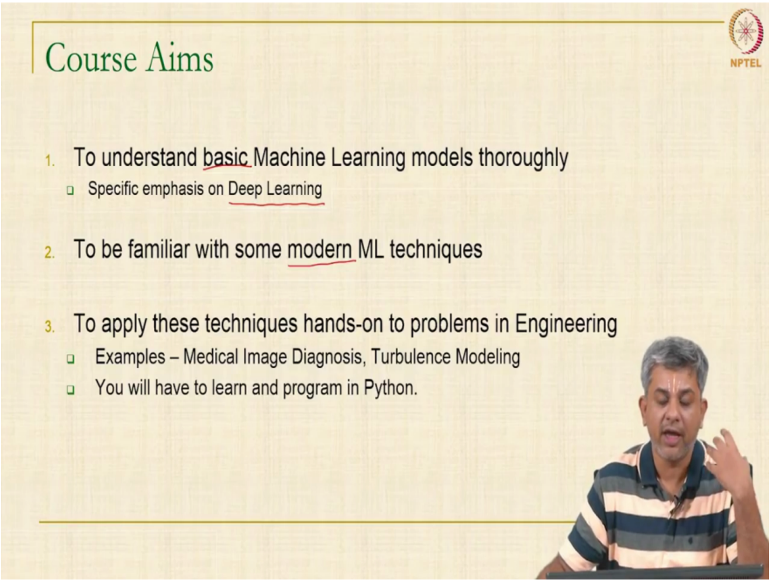

First thing is to want to understand basic machine learning models thoroughly. In specific we are going to look at what is now very popular Deep Learning, we will see what it is even later on today. Machine learning models thoroughly and in particular some very fundamental models that have been used for almost 50 years now, at various stages of development. We will also look at some modern machine learning techniques, which have caught on as the last of the last decade or sometimes even as recently as last few years.

We will be looking at several things that have been done in this field, even in the last year, the field is moving very very rapidly. We located in the context of engineering applications mostly. And finally what we want to do is to apply these techniques hands on to problems in engineering. Now within a video course of the sort that we are taking right now, there is only so much that we can do. We will look at some basic coding paradigms, we will also show you some examples but the hope is yourself motivated and you do learn and program in Python by yourself.

We will give you the basic rudiments and will give you several basic examples, you know for example from medical image diagnosis, from turbulence modelling, CFD competition, etc. But our expectation is that you will also do some work on your own, okay, whether you the course for credit or not, you can get the maximum out of this course, in case you tend to code yourself. So, we also hope that at the end of this course you should be able to read and understand research papers in machine learning, and especially applied research papers.

You might not be able to understand very hard-core machine learning theory but if somebody has applied machine learning to some practical problem, hopefully you be able to read a paper and understand it. This is also a primary aim of the course because most of the development that is happening today is not necessarily present in textbooks etc. It is mostly available as research papers, especially an archive. So our hope is you should be able to get this also, so this is a broad course aim, okay.

(Refer Slide Time: 6:42)

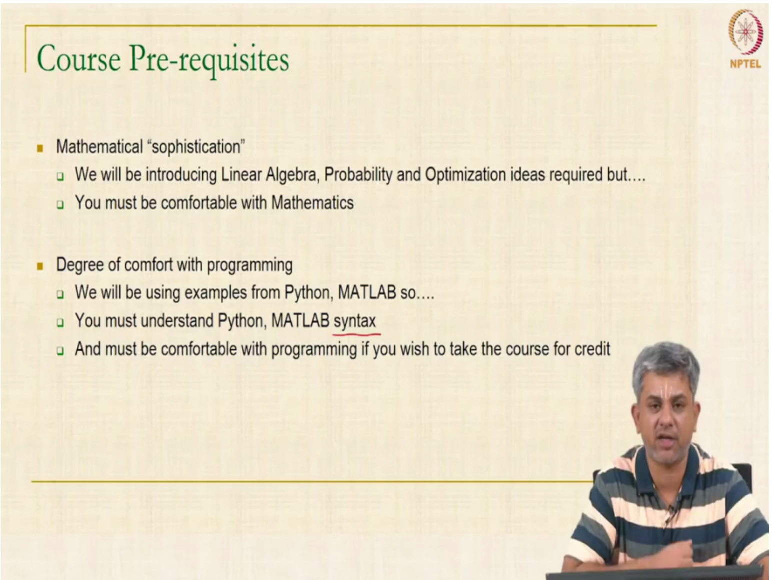

So it intersects with other courses from machine learning but the emphasis is a little bit more on the application side. And getting an overview and getting basic idea of the models that are in play. Okay. So, in terms of prerequisites for the course, what is it that you require in order make sure that you can complete the course successfully? One of the primary requirements that you would have is mathematical sophistication. This, sophistication is a vague term, what it means is that you are comfortable with mathematics, thinking about things in a mathematical framework.

Rather than just thinking of it in a vague framework which is just qualitative, we would like you to have a quantitative mindset. While we will be introducing ideas from you know whatever our essential ideas you need, not the whole of linear algebra or property or optimisation for that matter. But we will be introducing you to some basic ideas that are required for the course, okay.

Nonetheless that portion will be one week each, but more importantly you must be comfortable, okay, whenever we talk about something in mathematical terms, especially when it comes to probability, you should be comfortable in thinking through these things in a mathematical framework, that is one. Similarly, slightly similar in this respect is you should be comfortable with programming. Hopefully you have written programs at least in some language, we will be using examples from Python, examples from MATLAB. So we expect you to be able to at least understand syntax from Python and MATLAB.

But you will get the maximum out of the course if you are actually comfortable with programming itself and you can do some hands-on exercises. We will be giving suggested exercises throughout this course, so hopefully you should be able to do these, especially if you are trying to take a course for credit.

(Refer Slide Time: 8:24)

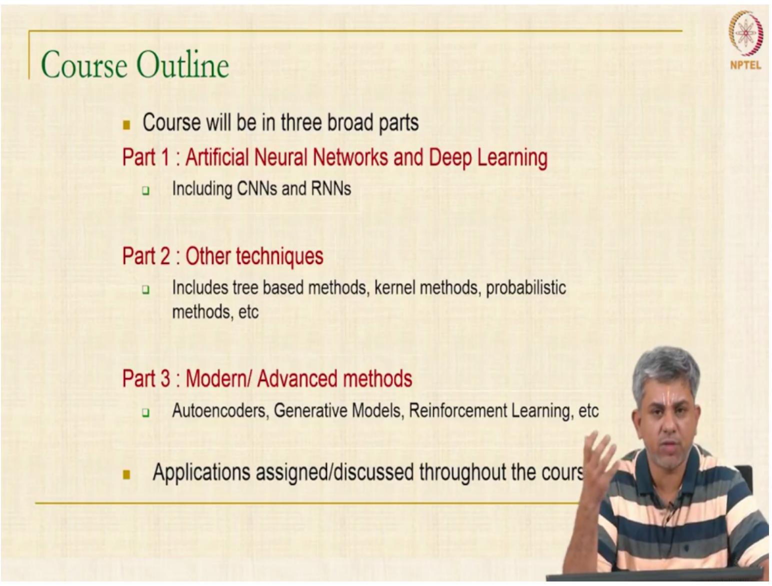

So here is an outline of the course. So we will have 3 broad parts of the course, the first is artificial neural networks and Deep learning, this includes what are called CNNs. CNNs are convolutional neural networks, these are used for vision. RNNs are used for typically sequential data. So we will be using RNNs, CNNs as well as what are typically called Simple ANNs and this is the first part of the course. The 2nd part of the course is other classical techniques that are being used for a long time.

They are still applied in various areas, depending on the complexity of the problem, Tree-based methods, support vector machines, probabilistic methods, etc, this is the 2nd part of the course. Finally we look at some modern techniques such as derivative, adversarial networks, etc. and reinforcement learning if time permits. As far as applications are concerned, within each module we will discuss various applications for each of those modules as we go for during the course.

(Refer Slide Time: 9:24)

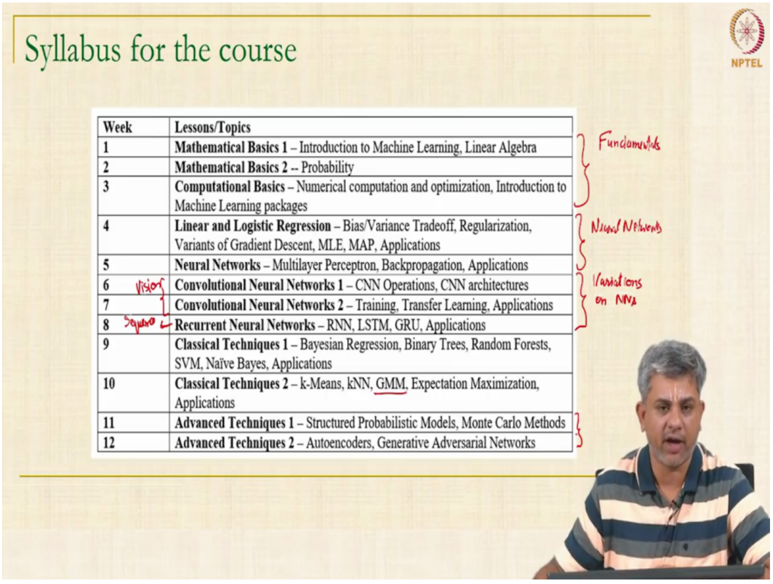

So here is the syllabus for the course, this was announced on the website also. So the first 3 weeks essentially are the basics, okay. So all the fundamentals that are required for the course. For the first week we are going to look at linear algebra primarily, then the 2nd week will be probability and statistics, whatever basics are required, visited the whole course in itself not just whatever basics are required for this course. The 3rd week would be whatever numerical computation and optimisation basics you require. And also popular machine learning packages that are available today, we look at an overview of those.

Week 4 and 5 is essentially neural networks. So it is possible to think of even linear regression are very very simplified neural networks. So we look at linear regression and logistical regression, these are 2 basic algorithms, the simplified algorithm, then complex neural networks and multilayer neural networks in the next couple of weeks. The next 3 weeks are essentially variations on neural networks. So this involves convolutional neural networks, this is for vision. Vision-based problems are usually solved using convolutional neural networks.

Recurrent neural networks are typically used for sequence-based problems, okay. So sequence that develop in time for example, So timeseries analysis in some sense can be used using recurrent neural networks. Then we look at classical techniques, techniques that have been around for a long time and they are still used in conjunction with Deep learning and neural networks. Some probabilistic techniques also we will be covering, for example Gaussian mixture models, etc. Okay.

Unsupervised learning will also be covered here, finally we will look at some advanced techniques, there might be some changes here as we go forth in the course, depending on how students are doing. We can also add a reinforcement learning if time permits towards the end of this.

(Refer Slide Time: 11:43)

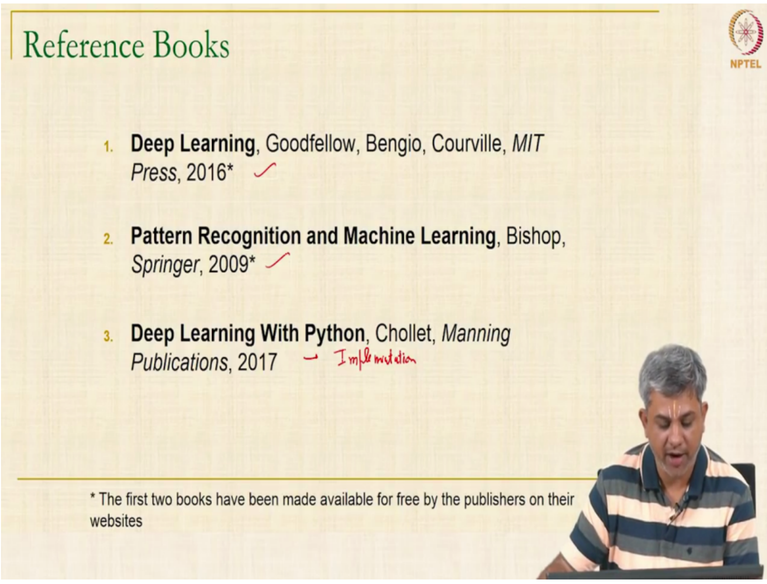

The reference books for this course. The first is by now, even though this is published only in 2016, it is already treated as with is a classic text. It is a very good text called Deep Learning by MIT press, Goodfellow and Bengio, etc are all researchers in the field. The 2nd is Pattern Recognition and Machine Learning, this is also a very very good text, though it is a dense text, it is a little bit harder to read but it is an exceptionally well-written text, very very thorough text by Christopher Bishop.

And the 3rd is towards the practical implementation side of deep learning, this is by Francois Cholle, it is called Deep Learning with Python. Now fortunately the first 2 texts are actually available for free. Okay. These have been made available by the publishers themselves, this is legal, you can search for these texts and you will find websites where these books are being checked by the publishers themselves. I would very highly recommend that you go forth and take a look at it and also read through these texts as the course progresses.

(Refer Slide Time: 13:11)

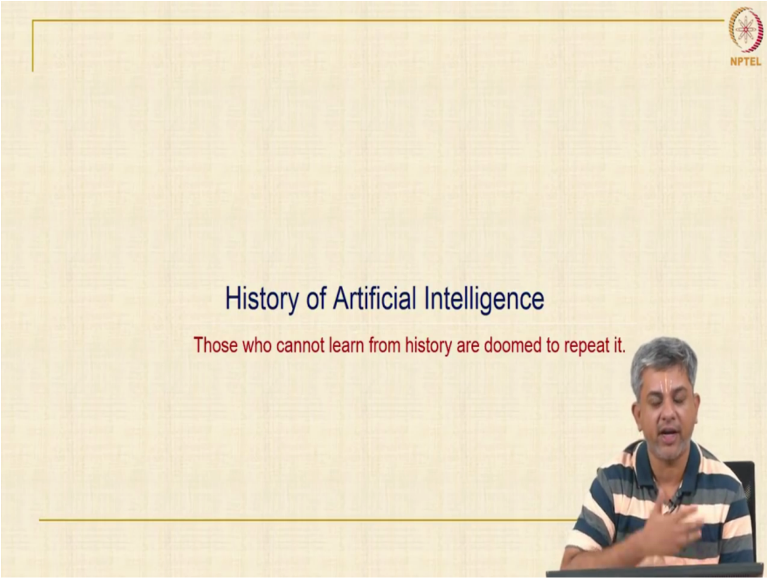

Please use the resources that have been generously made available by the publishers for those we are really grateful for them to having done that. So, now let us look at the history of artificial intelligence through the ages. Idea of artificial intelligence has been around for a really really long time. It is almost as old as whenever we started making tools. Why is it we are covering it, one, to see that many of the ideas that we are covering are actually quite old, also to see the ebb and flow of the ideas, when the ideas go up and when they come down and some of the ideas we might be covering in this course right now quite suddenly become unpopular for 5-10 years but suddenly might become popular again.

(Refer Slide Time: 13:53)

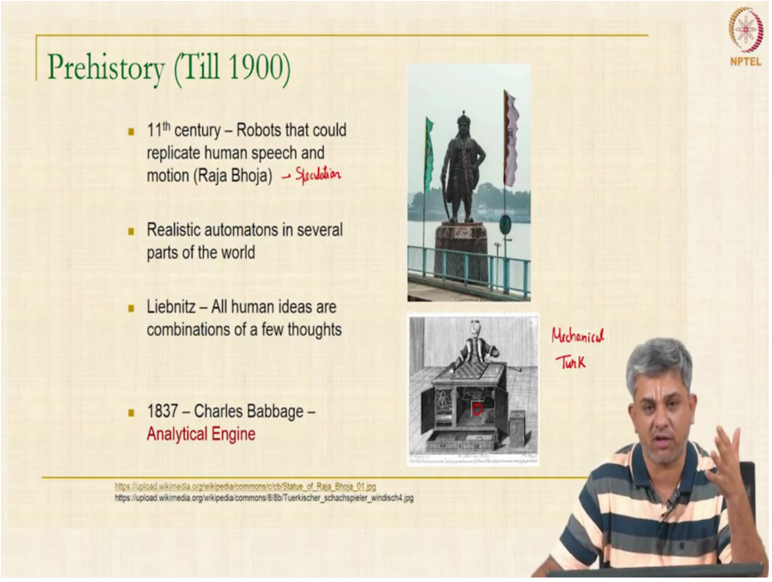

So it is a good idea to know where a topic comes from. So as I said history of artificial intelligence is really really old. Humankind has been fascinated with tools, they have been fascinated with what we can do, by looking at our hands we have started looking at you know what kind of machine tools can I make so that I can replicate the motion of my hand. So similarly just like we have been thinking about mechanical tools, we have also been thinking about thinking tools.

So can I not only replicate whatever I am, motions I am making, the wheel for transportation, the hand for working, the lever for lifting, etc, etc, not just that but can we actually make tools that can ease our thinking. This is Raju Bhoj of Parmar, this is from Bhopal. So this is speculation, I do not want to say that he actually made a machine, he did not. So within his works, rather bought was very very accomplished person, a poet, an engineer, he has a fantastic civil engineering work, all sorts of things.

He also speculated that you could have proposed that imitated or replicated human speech and motion. And even before him and after him there have been several people throughout the world increase, in Rome, etc. who have been doing this. There were realistic automatons in several parts of the world, right from prehistory till date. One such example, though it is a fake example, is what is called the Mechanical Turk. This was claimed by essentially a con man that he could make a machine that could play chess.

What he actually had was a person inside, a person who could play chess inside. But nonetheless we know that there were several automatons, that is things that could move automatically and replicate at least human emotion. Now Leibnitz, he is also the father of calculus along with Newton had an idea of a calculus of human ideas. We will come to this entry go on later. But Leibnitz idea was that every thought that we have is a combination of a few axiomatic simple basic units of thoughts.

We do not go that way but you can see an analogy with how modern machines are working. How do modern versions work? Modern machines or modern computers work on the basis that everything can be broken down zeros and ones. So the basic question that machine learning people have been asking in a very broader sense is to see whether all that we do in terms of thinking, I terms of creativity, can it be broken down into a few elemental products or a few elemental processes.

So Leibnitz speculation was that this is indeed possible, okay. Charles Babbage made the first or at least conceptualised the first analytical engine, modern-day computers are very very similar to whatever Babbage actually conceptualised. This was right back in 1837, the fructification of this was in 1940s when the first computers were made. My point is whatever we think of computation today, it is actually amazing, if even great work from the 1940s, they could have conceptualised what computers are doing today.

In some ways you can even think of computers today as artificial intelligence. In that you can book a ticket, you can record speech, you can play music, you can see movies, this is a wide variety of tasks, mind you. All of those are done on one single humble computer. In some sense this is already artificial intelligence, what we are going to do at least in this course is a little bit further. We know that while this is going on, it is not really thinking, it is not really learning. Our idea is to see if we can make algorithms which can actually learn.

(Refer Slide Time: 17:50)

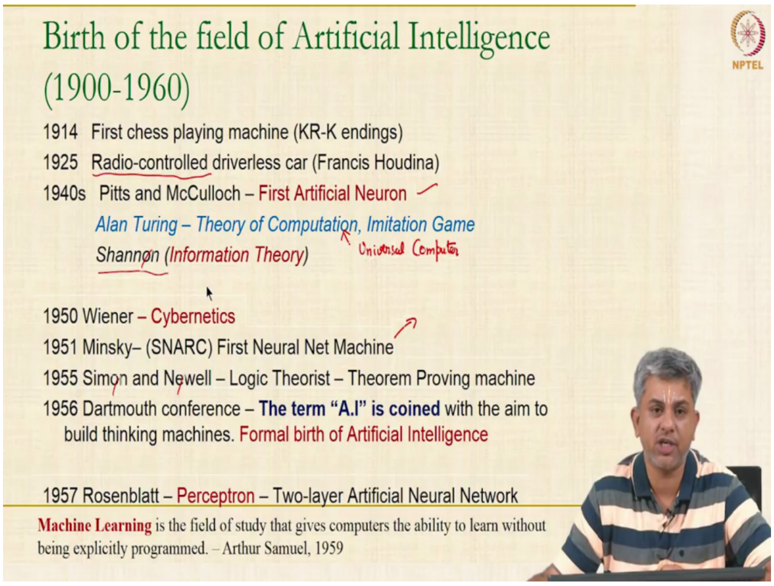

So here is the birth of artificial intelligence. 1914 was the first chess playing machine, all it was doing was king rook and king ending, those of you know how to play chess can, we will know that if one party has a king and rook and another party just has a king, you can checkmate And, always. So here was a machine which would actually do that, okay. Surprisingly enough, the first driverless car was right back in 1925, this was made with the help of the U.S. Army and I think Francis Houdina, this is not Houdini of the magician fame, this was a different person altogether.

So he made the first driverless car, however this was radio controlled. See, even this was quite amazing to people right back then. The first theoretical progress happened in the 1940s, we had the first artificial neurons, we will see this when we will come to neural networks. Turing also first proposed the idea of theory of computation and the idea of a universal computer. The idea of a universal computer is the idea that one single computer can do all computable tasks.

This seems obvious to us, we have sort of grown used to it but it was not obvious right in the beginning that every single computation can actually be done. You can think of a ticket booking as a competition, you can think of playing a video as also a computation, all the of that can be done on a single universal computer was not obvious at all. So Alan Turing was the person who actually pioneered this idea. Shannon also came up with information theory, this is now also being used extensively within ideas in machine learning and of course in a lot of places like signal processing, etc.

1950s was in some sense a place when artificial intelligence took off whether you read science-fiction from that period or whether you read just normal research people writing. So Norbert Wiener came with this idea of cybernetics, it was very very popular. Lewinski is a very famous researcher in the field, he made the first neural net machine, stochastic neural, automatic reinforcement calculator a SNARC. And Simon and Newell, Simon was a Nobel Prize winner, he worked for all his life for decision theory and affectively artificial intelligence, what we call artificial intelligence today.

So they made an automatic theorem proving machines. So the first machine was, this was not suffer but it was actually hardware, the first neural network machine was made by Lewinski in 1950s. 1956 was the first coining of the term artificial intelligence, a famous conference call the Dartmouth conference, Simon, Newell and Shannon all 3 of them participated is this. And the sentiment was really really positive. Rosenblatt was the first person who came up with a 2 layer artificial neural network called the Perceptron, he unfortunately died really really young, I think he died in the 1960s at the age of 40 or something.

Now another researcher, Arthur Samuel, just to define machine learning a little bit more precisely than we did in the beginning, is to say that machine learning is a field that gives the computers the ability to learn. This is what is key, the ability to learn, we will define learning itself a little bit later, without being explicitly programmed. So we will come to this distinct and shortly, okay.

(Refer Slide Time: 21:32)

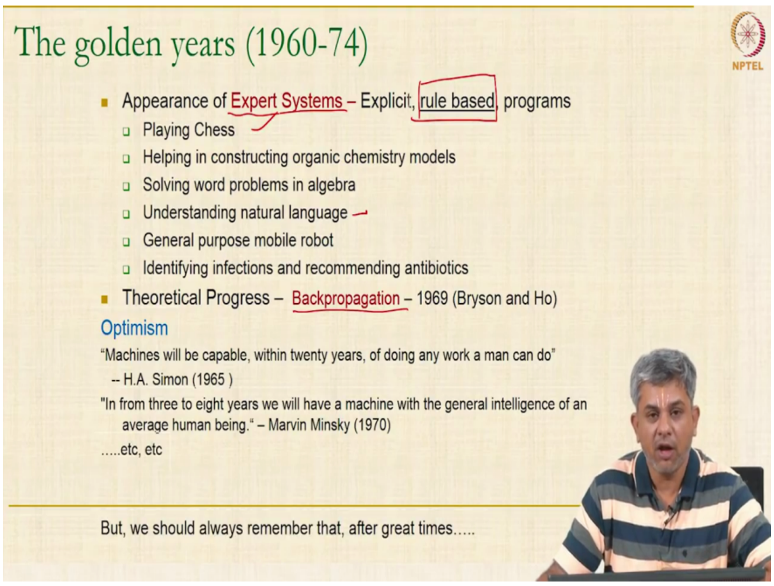

So here is the idea of being explicitly programmed. This is the idea of an expert system, so let us say you want to find out or you want to make an algorithm that detects grammar errors, okay. So some way to say it is you start putting all the rules that you know of English grammar, let us say, into the machine saying that if this follows that, you know if it is a singular person, then you put S, if it is multiple people, put R. But what happens shortly is in many cases it is a really really hard to program all the rules. And we will see some examples as we go on forth.

You will see that it is quite hard to do it even in grammar. It is not called clear how it is that human beings are able to recognise different grammars for different languages. Most of us in Indians speak at least 2 to 3 languages and most of us can seamlessly switch from the grammar of one language to the grammar of another language and it is not clear what sort of roles before low. So, expert systems work well when the rules are clear and when the rules are not clear is typically when we would like to use machine learning, okay.

So when in 1960s, when people thought of artificial intelligence, mostly they were thinking of rule-based systems. Even in chess, for simple games like tic-tac-toe you can use usually, quickly give the rule that will let you win or at least not lose. But in more complex such as chess or Go, it is actually hard and that is what people have been finding out. So, but early days of artificial intelligence, whether it was playing chess, making organic chemistry models, solving world problems in algebra, etc., etc., and even in understanding natural language, some progress was made but it was not good enough precisely because of this.

Because of the fact that these were completely rule-based, you had a rule for every single case. And if you did not give a rule, the company does not know what to do. There was also theoretical progress, we will see this algorithm is what makes machine learning algorithms or at least neural networks learn. Back propagation, this was available way back in 1969, that is nearly 50 years ago today in 2018. But the progress was good and people were very optimistic, here is Herbert Simon again, he is a Nobel Prize winner as I said. He thought that the machines should be capital by 1980s of doing any work a man can do.

Unfortunately, till today machines cannot even do what let us say mosquitoes or rats can do. So we are nowhere close. Nonetheless we have made a lot of progress, which is why the course is here. We also had Lewinski saying something extremely positive it was very very similar. But whenever you have such kind of a hype cycle, you should always know that you are going to get into problems. So what was known as the first AI winter happened between 1974 and 1980.

(Refer Slide Time: 24:28)

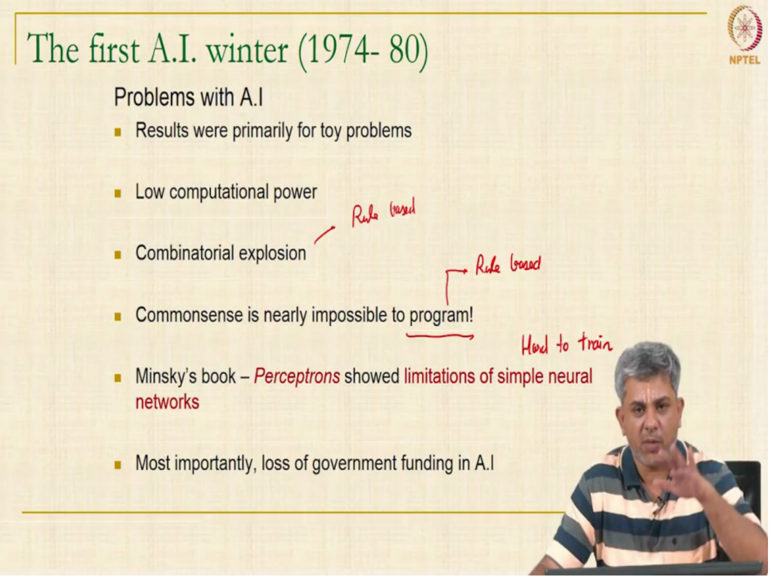

The problems were that all the results that were there even within chess were primarily for simple toy problems. The kind of exercise problems we do in any course, simple problems. The computational power was also exceedingly low, today Cellphones probably have greater power than most of the big mainframe machines had back then. So the computational problems we are looking at, in those days were really really low. Combinatorial explosion, especially for rule-based systems, okay. So like I said there is no simple finite set of rules within which you can explain every single case of grammar, okay.

It is really really difficult. Even for chess you cannot give a rule for every single situation, it is just too expensive. Combinatorial explosion means when you go from a small problem to a larger problem, the number of choices which are expanding is really large. If I only have one knight rook or knight ending, I have only 3 pieces, that is slightly less complex. If I have 5 pieces, you have things start growing in terms of factorial is, in terms of power laws and that is very very hard for a computer to handle. Okay.

So the key to even modern-day machine learning is this idea, that commonsense is nearly impossible to program. A baby today can look at a face and recognise that it is father or it is its mother. But you actually program in and to say why it is that it is its father or why it is that is the same person, regardless of this person changing their clothes, changing the expression, changing the way they speak, then growing older, having a beard or not having a beard. It is really impossible for you to program every single case in. But, somehow magically human beings tend to do this really really rapidly, okay.

So how is it that this happens is, of course it is a long-standing problem in cognition, it is still an open problem. Nonetheless we do know that it is nearly impossible to program in this explicitly. That is at least when we have rule-based programs, you cannot do this that easily. In fact probably, I will be bold enough to say that you cannot do it at all, which is where machine learning steps in. One other thing that happened during 1974 to 1980, was Lewinski's Perceptron showed up.

He made a very simple argument which we will make later to, this was not an unknown argument. The fact that very simple neural networks which are called single layer neural networks and not solve some knowledge problems. It is kind of obvious as you will see later on in this course, it is a very obvious argument. He also made an argument that multilayer neural networks are hard to train, so that was also made. What we mean by training is that somehow it is difficult to kind of Automatically program, which is what we are going to go into this course.

That this is very difficult to do, if you have more than one layer. If you do not understand that a player, that is okay, we will see this later on during the course. Anyway, this actually set up a panic but it will be set up a panic in conjunction with the fact that already results were hyped. And what happened was there was a lot of loss of government funding in AI and obviously most of the research work stopped.

Nonetheless some brave pioneers continued and as usual, you know, you have this kind of boom bust, boom cycle that keeps on going on in many fields, machine learning is just one example. One boom was between 80 and 87, we actually had the first driverless car then, this was, if I remember right, this is by Mercedes. And there was huge funding, if I understand it currently, about 750 million pounds were invested in driverless cars right back then. And it did not come to anything as we do now, now of course Tesla is taking over, Google Lexus, etc.

(Refer Slide Time: 28:30)

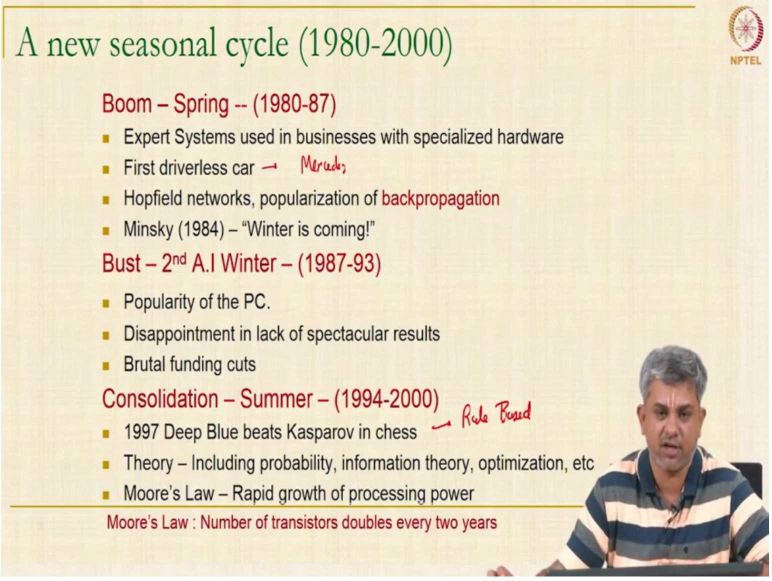

There was a book because primarily some builders are using expert systems, popularisation of back propagation and as usual Lewinski said that you know, we are going to probably going to go into first and it did happen. Between 87 and 93, the PC became very popular, you are not looking at large computation but individual people are able to do small word processing and, etc. for their needs. They are not looking at grand aims like General Machine becoming smart, etc.

Again there were a total funding cuts, usually precedes long winter in AI. Now between 94 and 2000, people got a little bit smarter and there was a long period of consolidation. Some of you might recall or you might know that in 1997 IBM's Deep Blue beat Kasparov in chess, this was still a rule-based system, it was not a machine learning system. A machine learning system playing chess has come only last year, again by Google's people, it is an extension of Alpha Go, I think it is called Alpha Zero.

So in 1997, Deep Blue beat Kasparov in chess, almost all stock fish etc., all chess engines that exist today are still rule-based. There was also simultaneously development in theory, including probability theory, information theory, optimisation theory, good optimisation algorithms which will be using in this course. And of course there was the stupendous power of Moore's Law. Moore's law is the law that number of transistors doubles every 2 years.

(Refer Slide Time: 30:12)

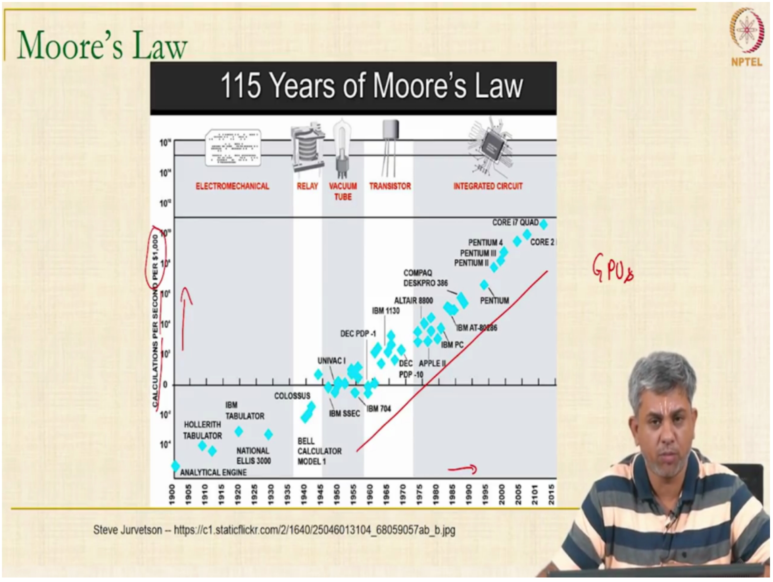

Now an adapted version of that is, you know we have sort of made it a tool. On y-axis here is the number of calculations that you can do per second per unit money that you are spending. That, this is a log log scale on the y-axis you have to log scale, this is a linear scale. And you can see that you have exponential growth of computational power. At least computational power in terms of the cost it takes for you to do this. So integrated circuits are there, we are now predicting quantum computation which is at least supposed to help or not necessarily interpretation but at least in some types of algorithm it is supposed to help even 40 planning.

(Refer Slide Time: 31:08)

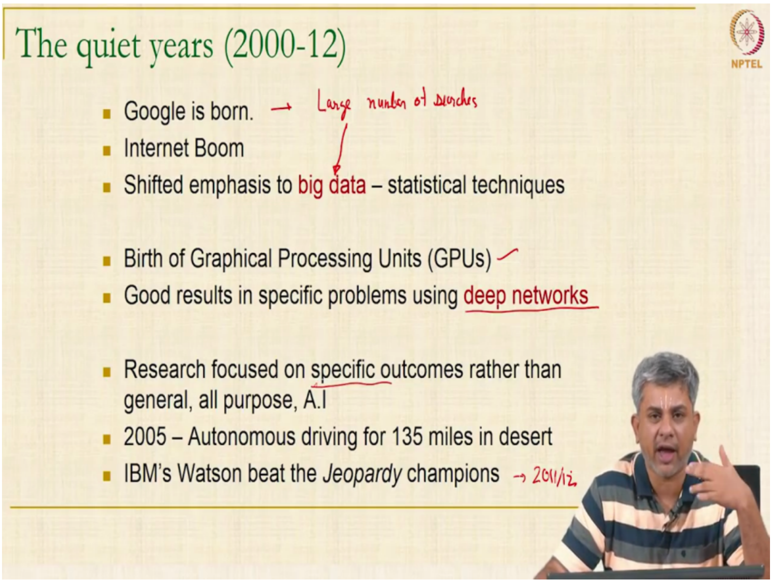

Google and other companies have invested very deeply into it. We also have GPUs, that allows much cheaper computation. So Moore's Law in terms of growth of computational, exponential growth of computational power has really helped. So 2000 to 2012, I would call the quiet years but they were quiet in terms of artificial content but there were very significant developments. The company Google was born, not only Google, there were several search engines, of course Google used it really well and came up with good algorithms for that.

What it helped was there is a large number of searches and when there is a large amount of searches, there is a large amount of data. So the key thing for machine learning, machine learning is very very data hungry as you will see was the amount of data that subjected itself through statistical analysis and statistical techniques. This is what happened between 2000 and 2012, there was an Internet boom, once again a lot of people offering a lot of products, a lot of people offering a lot of data, images, videos, all these came together and you had a large database on which you could train.

By train you will see what we mean later on in the course. Also we had Nvidia, with GPUs, these are very packed computational power horses. And we had specifically good results, if you brave researchers were continuing their work using deep networks. Another thing that people did very pragmatically was instead of looking at some bold aims like Machine becoming intelligent on its own, they started looking at very specific outcomes.

Instead of even saying stuff like I want a machine that will recognise everything when it looks at it, which would be the computer vision problem. It would say can I have a machine that can read a postcard and read the pin code. Just in that case it needs recognise only 10 digits 0 through 9. So, can, such specific outcomes really helped in getting good results. That is what led to the boom, even today you will see specific algorithms for vision, specific algorithms for natural language processing in a specific area.

Instant of saying something that can read your mail and understand it, can I have a spam filter. That is a specific outcome and in such cases you can actually find out which kind of machine learning model works better than the others. So this led to good, very good results and led to a positive growth of the field. 2005, once again we had autonomous driving for about 135 miles, of course without any interruption. What people are trying for now is more sophisticated. Can you actually drive on a street while people are moving around and we have had good results with Tesla.

Also one important result, I think this was in 2011 or 2012, I am not sure, IBM's Watson beat the Jeopardy champion, this is a quiz show and it is a nontrivial quiz show, it is not a simple language show. It has puns, it has plays on words, so the machine needs to understand more than how human sense in some sense, okay. And after this period of consolidation, we are now within what would be called AI spinning. We do not know whether it is ended, some people are already saying it is not it but anyway, all of us know that we are known growth cycle right now.

(Refer Slide Time: 34:20)

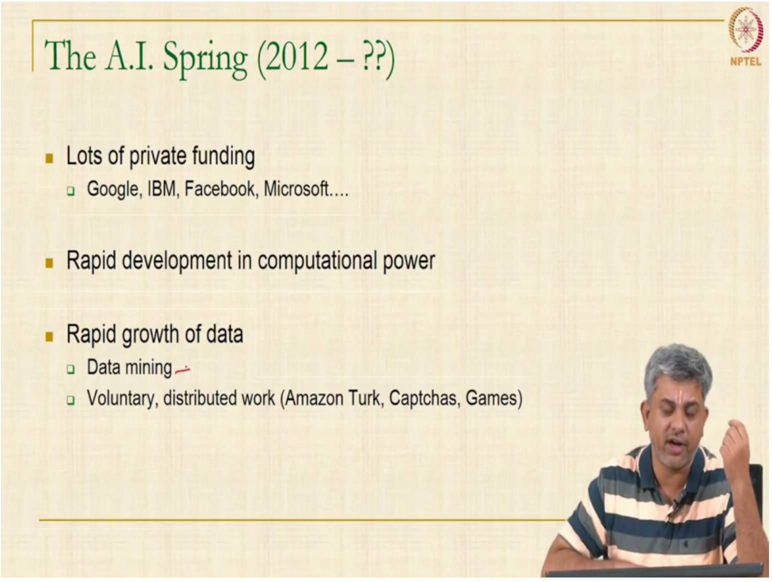

There has been a lot of what has sort of pushed it, there is a lot and a lot of private funding, not just government funding, in fact governments are catching up. Google has its deep mind, IBM has its own thing has its Watson, Facebook has its things, Microsoft has its things, and all of them have been fairly democratic about sharing their resources also. There has been rapid growth in computational power as I said earlier GPUs, etc. And a very important portion of it has been rapid growth in data.

Widow Facebook has data, Google has data, Microsoft has our data and they have been doing a lot of data mining, legally Hopefully and this has led to a lot of growth in machine learning itself. A lot of people have done voluntary distributed network in let us say tagging images, some of us have done it even though semi-voluntarily choosing captchas. Captchas are these things, you know you will have digits like N123 etc., that is just popped up in order to identify whether your human being or a Robot. But what it has done is it has also helped machines being trained.

Each time you say this, a machine recognises that this kind of image probably means N, this kind of image probably means 1. So that has been used also in training. So, voluntarily and semi-voluntarily we have been doing a lot of training for these machines and that has led to a lot of data. Games also have led to a lot of data.

(Refer Slide Time: 36:02)

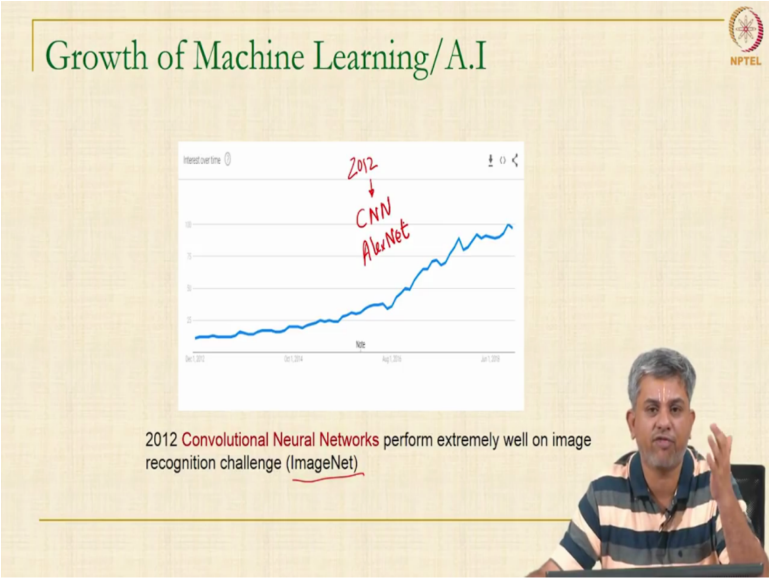

The inflection point, you know the point where a lot of people identified as the real growth of machine learning in AI, at least the modern boom cycle is sometimes in 2012. There is a challenge called the image net challenge which we will see them become to CNNs. So this is just a vision recognition challenge, out of a thousand categories of images you have to say which one is which, is this a cat, is it a dog, is it a building, etc. So, most of the algorithms at that point, all of the algorithms at that point which had the winning work in some sense traditional vision based, rule-based algorithm.

So 2012 was the first time that a machine learning algorithm called convolutional neural networks which one, this is called the algorithm which one is called Alex net, which we will be covering in detail later on in the course. This was 2012 and since then every single year, the algorithm which has won is only a machine learning algorithm. Alex net showed a huge jump in performance based from, about 12 percent jump from previous algorithms. And this is when people sat up and took notice and since then the field is just taken off very rapidly.

The number of people who have got in within the last 5 or 6 years is just huge. People who started out doing their Ph.D.'s in 2012 without a machine learning algorithm have done things that are surprisingly chaste and the course of their Ph.D. This is not very long but a lot of the material that will be covering in the courts will actually be from the last 3 or 4 or 5 years. We will be covering classical techniques but will be also covering what has been specifically done in the last few years.

Which is another reason that we are asking that you also understand how to read research papers because the field is still developing, it is still in some sense early days and you need to know how to keep up with the literature. So, part of the language and the techniques that we will be introducing you through the course is for us to make sure that you can actually read the papers, understand and maybe implement them yourself in some application of your interest.

(Refer Slide Time: 38:16)

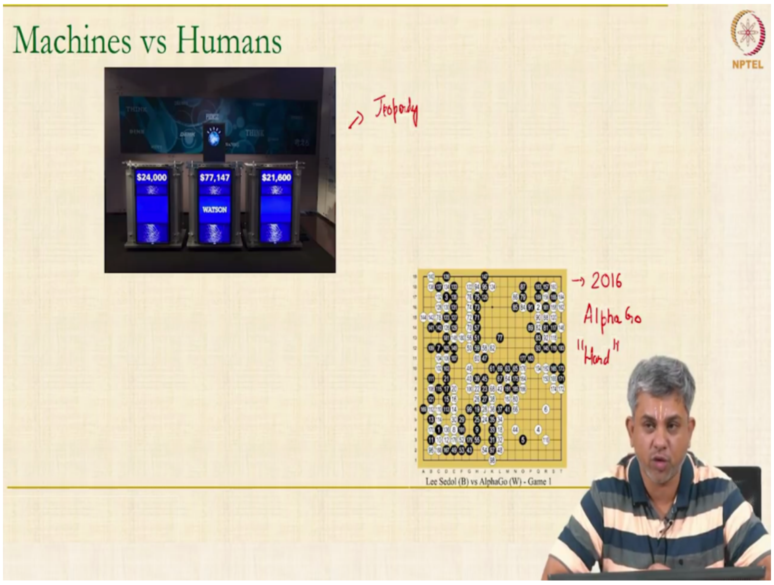

So here are 2 results that have been the reasons for people getting kind of worries. This is a machine versus humans, this is the same thing, Jeopardy, this is a quiz show in United States, I do not believe this is there in India as yet. But it is not quite sure and it is an involved language based quiz show, of course it is knowledge-based also. But there are puns, etc. and you have to understand allusions very clearly. So the thing that wants was IBM's Watson and it is a machine learning-based algorithms or at least semi-rule-based semi-machine learning based.

Another thing that was frankly a shock for many people in the field was Alpha Go. Go is a game which has simpler rules than just but it is known to be combinatorially much more harder to solve. It is a 19 by 19 board, where people simply place a white or a black piece. But nonetheless, it was often known to be a hard problem problem in AI, that is chess was thought to be essentially solved by a rule-based system. People thought that there would be no machine which will beat a Go champion in maybe another 10 years.

Even Google when they came up with Alpha, were not sure that they would actually win, their aim while playing Lee Sedol, the Go champion at that point was simply to maybe win a game or 2. And then to learn and make the system better but it actually beat Lee Sedol handsdown. And after that I think they have retired Alpha Go records it is just has been so good in betting every single human being. Practically it is unbeatable at this point. So one thing that is also true about machine learning algorithms is sometimes it is hard to know how good they will be.

(Refer Slide Time: 40:31)

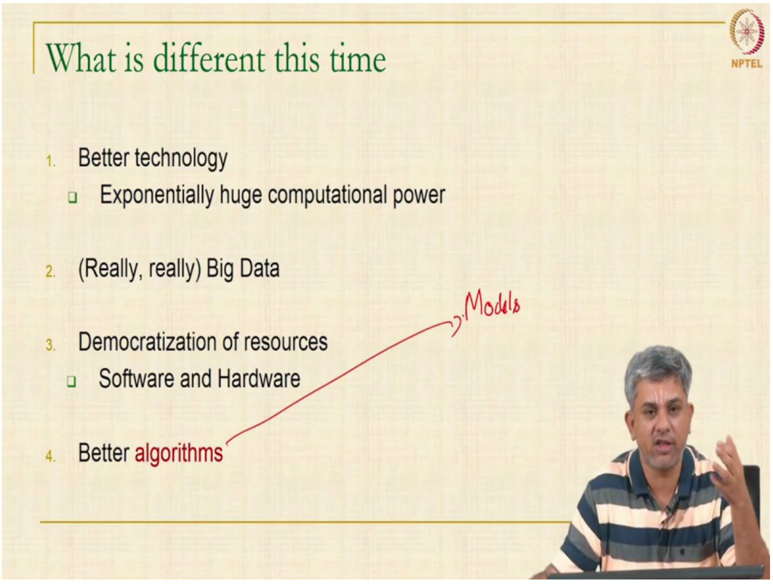

So these are the 2 of the recent results and they have been cause for again a recent hype cycle. So the question is, is there anything different this time and the answer is yes, at least there are few tangible things that are different this time. We generally have better technology, okay, and computational problem that is available today is exponentially such as compared to the difference between let us say 1950s and 1970s, between 2 cycles. So we have GPUs, then we have all sorts of futuristic computational technology that people are proposing.

Whether they come through or not, Moore's Law is still holding it is breath and so it is kind of running its course already. It is going for but we have different architecture. We have really really big data, we are practically drowning in data and probably we need better algorithms to handle this kind of data that we have. An important portion of the current boom cycle is we have had democratisation of resources. A lot of people have, a lot of algorithms can be run on a simple laptop that is accessible to most people today.

Between 75,000 rupees to 1 lakh Laptop can actually have a simple GPU, a simple card that can actually do a very good job and even some of the simpler algorithms can work. Also many of the commercial companies have been very very generous with their software. All the packages, whether this Google's Tensor flow, Facebook's pie torch, etc., etc. IBM, Microsoft, all of them have made a lot of resources available to the common public. So the open source movement has also taken off and this has led to a lot of software which is being available to the common public.

We will be using a few of these through this course which actually makes a person come up to speed very quickly. Even if they do not know how to code something from scratch, they can use existing packages are at least use a view of those algorithms. We will be seeing that later on in this course. And there have been generally better algorithms, even though there are variations of prior algorithms, we do genuinely have better algorithms today. So, our focus for the rest of the course is the algorithms portion. What we are going to really look at is what organisms work and under what circumstances do they work.

We are going to look at algorithms as if they are models, okay. So, if you have done any engineering problem at all, you will know that sophisticated processes we have various models. If you are in fluid mechanics you will have various models of how a fluid behaves, you will have various models of how turbulence we have. Similarly if you are an solid mechanics, you would have seen various models for how stress strain should be model, etc., etc. In every field, an ideal gas law is a model, Ohm's Law is a model for how current and voltage and resistance play together.

So all these are models, with think of algorithms as if they are models. Modelling what, the specific input and output relationship in any problem that you are looking at. So, suppose I have a mail and I am going to classify this as spam or not spam. So, there is something that is going on in my range which is modelling this. What sort of model will work best is what we are going to look at through this course. Under what circumstances, if you have a vision problem, what kind of range of models do we have.

If you have a time sequence problem, what kind of models we have, this is what we are going to look at for the rest of this course. Thank you.