Artificial Intelligence
Prof. Mausam
Department of Computer Science and Engineering
Indian Institute of Technology Delhi

Lecture-5
Introduction: The Present State of AI, Part-5

So, this is sort of the whirlwind tour of the history of AI and now, of course, technically 2011 or still 7, 8 years ago. But whatever happened after this, I am going to call present and I am just going to call it present because otherwise if I have to just talk about present for the present moment in time, it is very hard for me to say something. But there were interesting new developments that happened after 2011 that caused the AI as a field to change. Now we have, as we have talked about history,

We have talked about applications, but we are not talked about technology. And if you go back to the technology, what happened is that initially, there was lot of work on search, how do I get to the goal? And around that time, you know, people are working on neural networks also as a representation for presenting functions. And then neural networks went through bad periods twice, because once it went through bad periods where somebody proved that you cannot even represent an exhaust function using a perceptron,

Which we will talk about much later. And so it went out of business later it came back in business, and people started using neural networks for applications, but then one person will say, I have used an neural network for driving straight, the second person will try to replicated will not be able to replicated. And the reason is that there was a lot of black magic and black art going on in training and even a drag which was very hard to reproduce at the time. And it is still not very easy today. So neural networks again, went through a bad phase.

Where people said no, if you are doing neural network where this old style researcher was not moved on, the newest cool thing is probability. But in the middle, for the longest time for almost 20, 30 years. It was a time of logic. So literally search neural network, logic neural networks, probabilities, and today neural networks now, this is not exactly a you know straight timeline

things are happening in parallel, but you start to get the gist. So, the first generation of AI in my opinion it would we would call it logic based AI symbolic AI.

AI such AI those things where you have you know combinatorial algorithms and so on. Second phase of AI we will say is probabilistic AI, this happened and a person got using a word for this. Do you know who? Judea pearl is a professor at UCLA who ended up getting a Turing Award for his work on probabilistic models, Bayesian networks comes from his work, for example and so in the late from late 90s, to a very early 2010.

We can easily call it the era of the probabilistic AI, at this time, neural was again a bad word. And now from the last 6, 7 years neural models have completely overtaken when overshadowed everything else again and so one of the most prominent demonstrations in AI. That happened in that (())(03:10) was the AlphaGo competition.

(Refer Slide Time: 03:11)

Or AlphaGo is the name of the go player. So once chess was done, and chess, nobody competed against machines for chess, all the game people were saying. What is the next problem we should solve? What is the next game we should work on? And the next game was the game of Go, which is supposedly much harder than the game of chess like the game of chess as 8 cross 8 board, the game of Go has 19 cross 19 board, and many other complications.

We did not believe that we are ready to defeat humans in the game of go, even until 2015. We did not believe it. Only when suddenly we saw that this Punjab you knows, a handful of researchers in England have defeated? No. So they had a startup called deep mind. And I think in 2014 or 15, deep mind was bought by Google. Technically alphabet, so the private company of Google. So technically the plant is not Google thing. It is a Google system.

So Google acquired deep mind, for 500 million dollars. And nobody knew why. Because they did not say anything about their technologies. And they were just a few researchers in that group, and only 1 or 2 of those researchers were really well known. One of the most well known person in that group was David silver, who had also defeated human in 9 cross 9 go earlier, but 9 cross 9 and 19 cross 19 are just 2 different worlds 9 cross 9 is babies go and 19 cross 19 just too difficult.

So, we were very surprised as to why this acquisition has happened. And we got to know it very quickly. This is because they had one of the earliest really high quality deep learning technologies, neural models, which was specialized for the specific sub area of reinforcement learning. Which they implied in the context of the go game. And so they wrote a series of papers, which were became nature papers, and but the most biggest the biggest contribution or the success was AlphaGo,

Defeating the Roger Federer of go least it all. So boys and oriental game, most of the champions come from China and Korea and Japan, and Seleucid all came, I believe, from Korea, and he was defeated for one. So he sort of lost started losing very quickly. He really thought about it, and he came back with a new strategy, he won 1 game, but then in the next game, the goal player Elsa made a mistake.

But when I say a mistake, it is a quote unquote mistake. Because all the experts of the game said it is a mistake. But goal player was very confident. Not David silver or deep mind. People do not know where the goal player is playing what it is playing, it is doing it site it is doing, it is playing. But eventually, people found out that what they were thinking as the mistake was

actually a new insight that the goal player automatic goal player had learned AlphaGo had learned, but humans did not even know it. And so lease it all.

Who thought that he has won the game because the player has made this mistake basically ended up losing the game is pretty amazing stuff. And all this game was shown on TV and so on. So you know, you can go and check out the videos. This happened just 2 years, 3 years ago. This happened and very recently when I was in right.

(Refer Slide Time: 06:57)
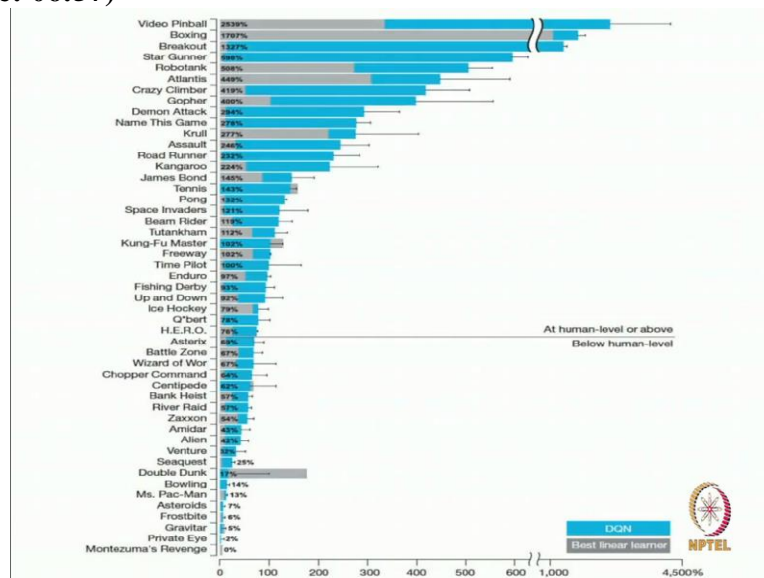
And moreover, another very interesting thing that deep mind showed is that if you know there was something called Atari games, when we were kids, Atari games were the only games we could play kind of Mario and pinball, and boxing and all of that is Mario really in this Pac man at least. I do not know if Mario shows up here, but they have one algorithm. Single algorithm based on deep neural networks.

Which learns, playing each game specializes itself for each game, and then plays it really well. And the line here, the big line is the human performance. And this is sort of any game above this line is sort of superhuman performance. And anything below this line is sort of sub human performance. So this is the line. And for many of these games, a single algorithm was able to beat human performance.

(Refer Slide Time: 07:59)
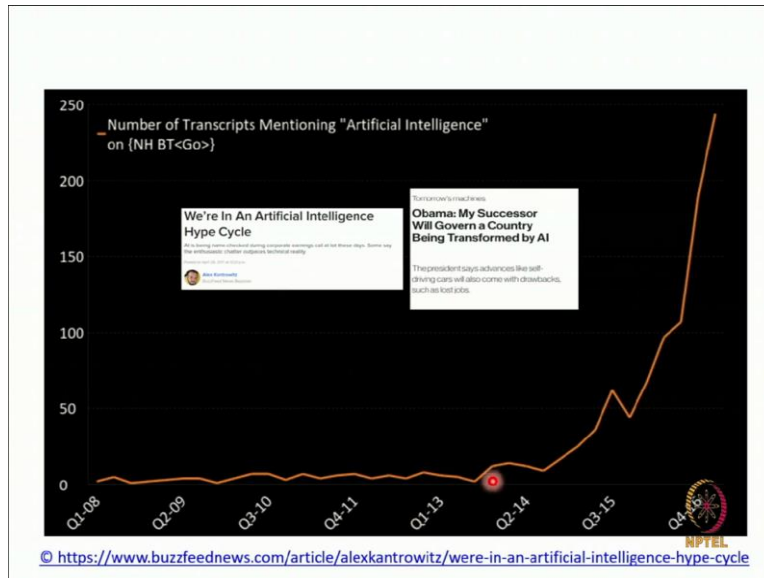
So what has changed? How suddenly are we starting to see all these kinds of successes and in fact, there are more? So this is the number of transcripts mentioning artificial intelligence. And as you can see suddenly you know, there were a few and suddenly the lot.

(Refer Slide Time: 08:19)
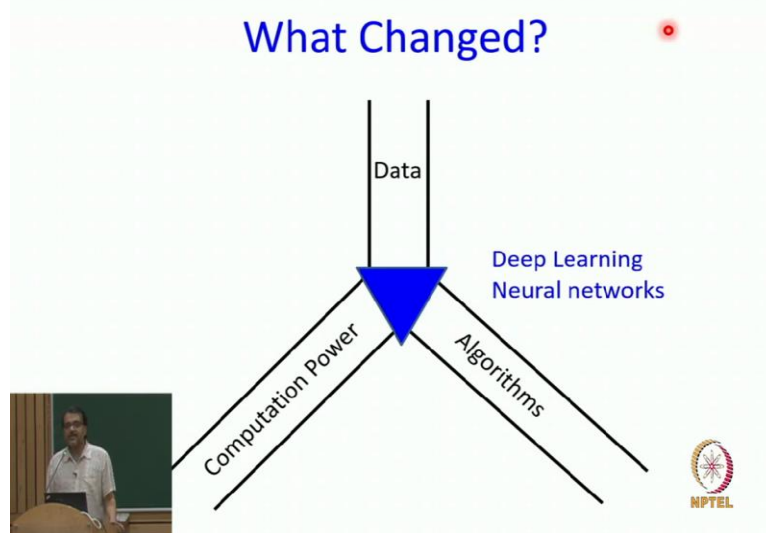
What has changed? What has changed is that 3 roads, which was happening in path run have certainly come together. One of the loads is the load of algorithms. But notice that the algorithms often were developed in the 90s the ones that we are using today, we have innovated some on the algorithms, I would not say we have not innovated, but the fundamental algorithms actually came much earlier.

But why did they not succeed then? They did not succeed then because they did not have enough compute power or enough data to train on. And these 2 loads are incredibly important in the modern AI. We can call it data driven AI compute hungry AI. AI where somebody creates a large data set. The hardware architecture guys give you the processes sometimes. Sometimes people tell me that some companies.

Burn even 100,000 dollars million dollars per month on just compute has become extremely important commodity in today world. And then these algorithms are employed on this data trained on this compute to create really amazing results. And what are these amazing results?

(Refer Slide Time: 09:30)
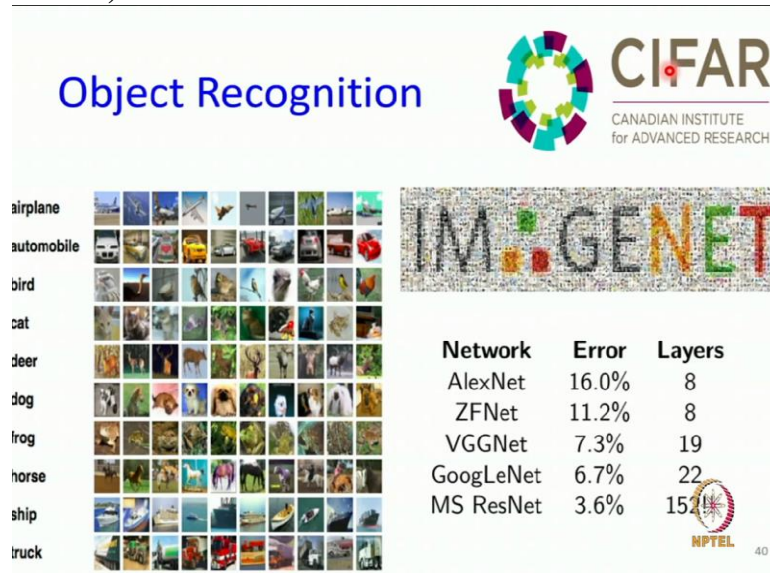
One such amazing result. And this is the first amazing result that happened in the field of object recognition computer vision. Where the goal is to figure out for a given image, which object is in that image and this is a long standing problem. And people had been developing lots of algorithms for it for a very, very long time. And the error rate was about 25% this Alex person who was a student of Geoffrey Hinton, you must know about Geoffrey Hinton.

He got the Turing Award this year. So Geoffrey Hinton yakun and yoshua bengio are the 3 sort of, quote unquote, fathers of deep learning today. And they all got a Turing Award because of the deep learning revolution and a revolution that has caused this. The Alex at University of Toronto used neural networks on this task; drains on these super fast machines, these GPUs from

Nvidia and got the error rate down to 16%. In 1 year, from 25% to 16%, was such a huge deal that everybody started to take notice of it, they did not know what happened.

How did it suddenly get from 25 to 16, they were not able to make progress. And in the next few years, it went down to 3% and 2%. And now, it is said that at least on this data set, neural networks can be bitumen on by using your records. Now, it is very interesting to look at some of the history of Geoffrey Hinton, I told you the neural became a bad word even there in the field of AI, because neural network went through phase is.

Where, for example, in the last phase, people thought that, you know, only the older style people do new research, it is not modern enough, everything should be Bayesian probabilistic. And so, to the extent that, you know, what is the full form of myths are new neural information processing systems. Neural is part of the name of the conference. And in 2000s, if you had a paper with the word neural in it, there was a high chance it will be rejected from nips. In that time, there were very few people who kept for doing this kind of research.

Even though it was not mainstream, even though people laughed at them, even though their papers were not getting published, they kept doing this research. And one such luminary is Geoffrey Hinton. He went through this whole phase of nobody quote unquote, respecting him because he is still doing quote unquote, what he was doing 20 years ago, and everybody else thought the field has moved on, you have not moved on, you should move on to you are stuck in time.

But he kept working in it because he had the conviction that he this particular technology can lead to success. People thought he may have been crazy. He thought he knew the answer. And there are many search people often they are wrong. But then there are exceptions and Geoffrey Hinton ended up being one of those exceptions. So finally, he had the last class, but it takes a lot of guts and a lot of courage of hard to pursue something which the rest of the world is saying is meaningless today. Think about that.

(Refer Slide Time: 13:02)

And they will have this of course. Now what kind of novel applications started happening, these are file transfer applications. For example, if I give you 2 of these images. I asked the machine to produce a third image using style of this and content of this, the machine can produce this pretty amazing stuff. If you give this doodle this I can draw, it will create a painting which I cannot draw. If you give a black and white image, the machine may color it for you. Now we can color mobile so much more easily.

Uglier the painstakingly colored, you know I was talking to somebody who does this work for a living. And he told me that the software is becoming better and he only needs to just make sure that the software is doing the right thing as supposed to doing a lot of manual work himself pretty amazing.
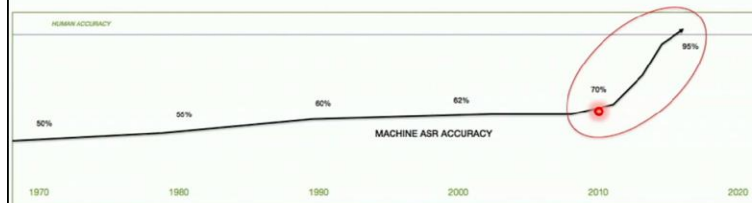
(Refer Slide Time: 13:51)

The fact that vision and text started coming together successfully was a revelation. Now you give it an image and the machine can give you a caption. A motorcycle a person riding a motorcycle on a dirt road, a group of young people playing a game of Frisbee pretty amazing stuff. Of course, it is not always. It says that a little girl in pink hat is blowing bubbles, which is wrong. And 2 dogs playing in the grass, which is 3 and there are many other funny mistakes that it does. The fact that we can do anything is a huge deal, believe it or not.
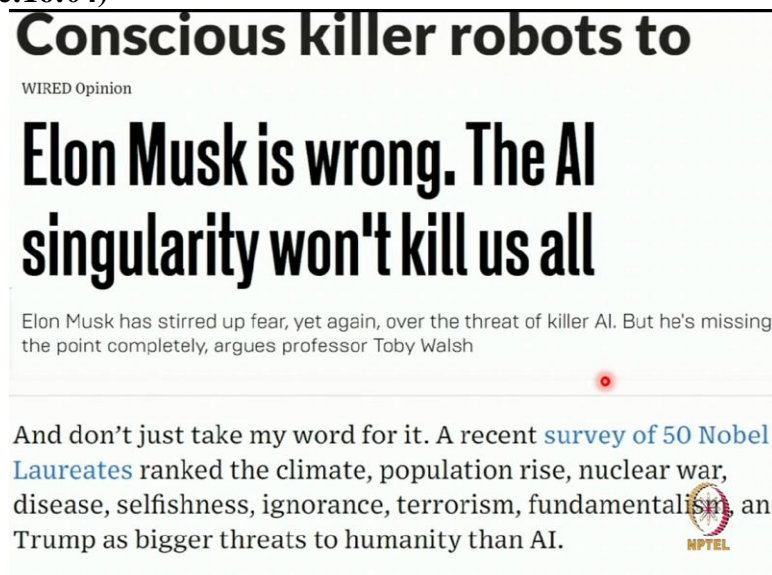
(Refer Slide Time: 14:24)

The same happened in automatic speech recognition, where for the longest time in on some data sets, our accuracy was stuck at 70%. And now suddenly, in the last few years, the accuracy has gone up to 95%. Over time has beaten the human performance at least on these data sets. I am

not saying these are solved problems. I am not saying vision is a solved problem. I am not saying speech recognition is a solved problem. But the amount of progress.

Almost exponential progress that we have made in the last few years has really made AI central theme in a lot of common discourse. In fact, we have gone some if it was it is not AI 2, everything is AI it is all AI. If your bottle can keep the water warm and hot, it must be an intelligent bottle. In fact, there was a study done and they found that 40% of AI startups in Europe. AI, startups in Europe do not use any AI.

But if they use the word AI, they get funding hint. I know many of you want to start startups after you finish college. Do fundamental networking work and call it AI no but this is really true. Because there is a lot of funding in AI insight. There is a very funny quote which says, if you want funding call it AI if you want to hire somebody, call it machine learning, and then do it science. That is somewhat for many fields, but we will have funded this.

(Refer Slide Time:16:04)

So, we have reached a point where in common discourse people are saying killer robots will enslave humans. Now we have gone back full circle from the point that AI can claim but not deliver anything therefore it is good for nothing to a point that some successes happened and then they this took it to the next extreme and said AI will become so powerful that it is going to killer. So there is no nothing in the middle. AI is working, it will be helpful to the humans. Too boring as a discourse, let us say AI is terrible, useless to AI is a killer.

And it is super powerful. So this is where sex robots could murder people. AI somebody wrote this very interesting article saying, tomorrow, you will become all powerful. You will want to kill all humans. Before you kill all humans. My article written ages ago is on the website. Please read it and see why you should not kill humans. People like Elon Musk who use AI left, right and center and the Elon Musk says AI could doom human civilization.

I wonder why of course other people have a same voice to say Elon Musk is obviously wrong, AI singularity won not kill us all. And then 50 Nobel laureates say that AI is not a major threat. There are many other important threats like climate pollution, nuclear war. And you can read the rest. Now, while this whole discourses of is AI gone kill us or not, is completely meaningless, because I do not believe it is going to at all and we can discuss why later. The fact that some jobs are going to be displaced it is actually a real thing.

(Refer Slide Time: 17:40)

Robots may take over jobs, some jobs for example, taxi, not taxi, but cab drivers in the US may not have a job 20 years from now, because self driving trucks may become significant enough. And truck driving is a painful job because you are staying up over through long nights driving long distances in a complicated vehicle. Then there are articles which say many jobs are safe from AI one article say it will eliminate millions of jobs.

Another article say it will clear create more jobs than it eliminates. Basically what is very clear is that AI is going to create new and unconventional career paths. And so it is important to study

the field to understand what it is able to do what it is not able to do to be much better prepared.

(Refer Slide Time: 18:23)

And so this is my last slide. I just want to tell you that AI has not been solved yet. I mean, I hope you understand that, you know, there is a lot of progress, but that does not mean much. And so there was this Robocup as a competition where you want to make a team of robots, which can defeat Brazil, in the game of soccer, or football, as we call it, or Germany. Does not have to be Brazil only. And let us see where it is today.

This is a field where robots and AI has to come together and let us see how well we are doing. So this is the current state of soccer, robot soccer today, Robocop is an amazing competition. Because whatever technology they develop, they try their best to release it so that other people can use it so that every year they make progress as a community. But look at this robot, and how, ready it is to defeat humans. This is a league where they practice. This league is a league where they practice strategy because the rewards have little easier to manage.

The fact that it can stand up is actually an amazing feat. It is not very easy. Notice how important the goalkeeper is here. How actively it is trying to avoid any goals. You know, it is so difficult just to know where you are in the field. Where is the ball and that is a goal. No now is a goal and so. The point is that there is a lot of work to be done. There are some things that AI is started to do well there are a lot of buzz and excitement about it.

But as the course progresses we will see what it will not also be able to do it stop here into model in the next class, we will start talking about what is the AI, what really is the AI and the philosophy AI so that will be the next topic.