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

Lecture-9
Introduction: Example Tasks, Phases of AI and Course Plan, Part-9

So now that I have given you this definition of, you know, ideal rationally, and you can say man, this is a good definition of AI, it makes sense. But it is too general.

(Refer Slide Time: 00:26)

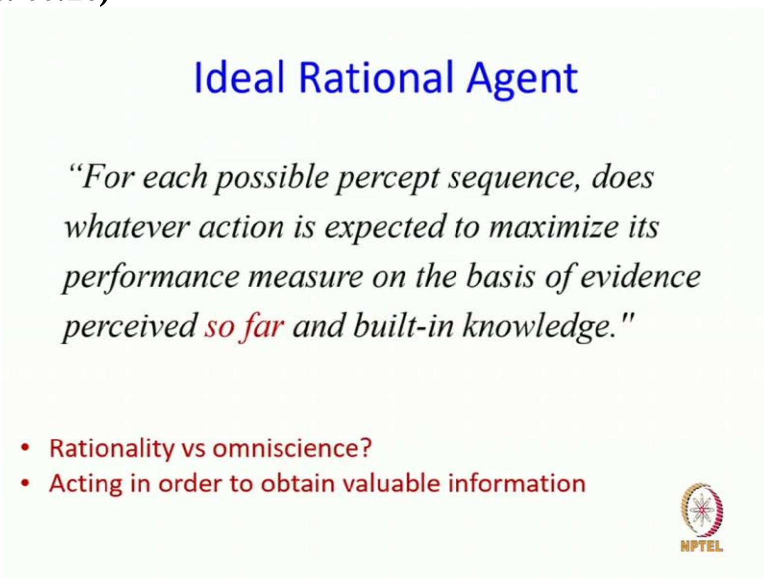

Technically, you can cast anything into some kind of optimization, you will have some knowledge that you know, and you will have something that you do not know and you will have to make some decisions. And then that is AI is that AI philosophically it is. But in practice, we are only interested in very hard problems.

(Refer Slide Time: 00:47)

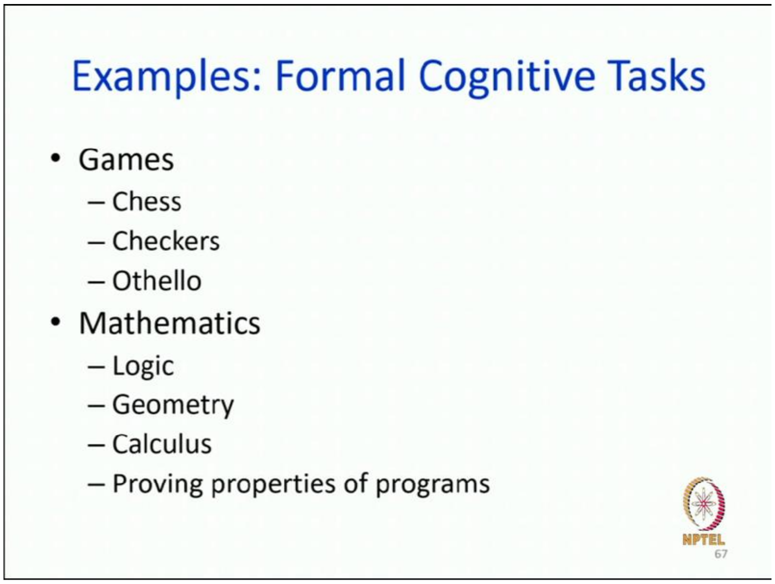

So we try to look at those problems where you know, most other people are not able to solve. Problems that are often NP hard and we will come to that in a few minutes. So for examples Games chess my objective function is to win my action is to make a move my observation sequences the history of the game. So far my built in knowledge is how the game works. For example, mathematics, you know, theorem proving, I am given some theorem proving actions like, you know, embolic deduction and applying rule number one and applying.

You know, to theorem number, rule number 2, whatever. And I am given a term that I want to blow my objective function is to be able to finally prove it. My action is to apply each rule one at a time. So, these are called formal cognitive tasks. These are difficult tasks. These are tasks that you and I will consider quite difficult. Not everybody can do theorems, not everybody can win in the game of chess.

These are tasks that we have learned over time and we also consider them into intelligence. And in the first generation of AI systems, these were the most important tasks let us, you know, to all theorem of Bertrand Russell Principia Mathematics, using AI or let us defeat guy the chess champion in chess.

(Refer Slide Time: 02:02)

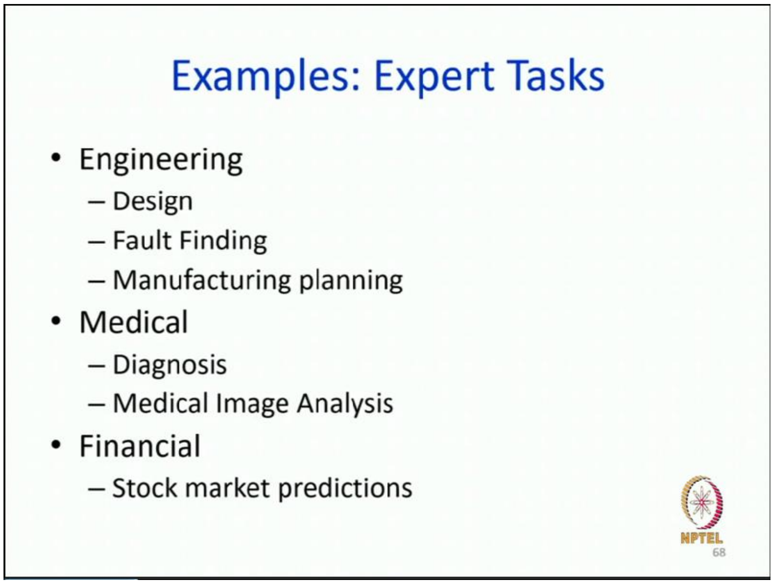

Then later in the 70s, and 80s, we started looking at expert tasks more than more, so expert tasks or tasks which require very expert knowledge in a specific sub area. No chess you can quickly teach. Mathematics is still somewhat expert. But then engineering design and medical diagnosis and financial stock market predictions are even more expert tasks. These require a lot of knowledge about your specific field, you need to know about diseases and treatments and symptoms and interactions in order to be able to do medical diagnosis.

And so these would be called expert systems, expert tasks. And so, financial stock market prediction would be a great a problem, the objective function will be to make a lot of money, the action will be to buy some amount of some shares at every point in time and sell some amount or some shares at the same amount of time, same point in time and my built in knowledge, maybe the historical data about how the stock market prices have gone and so these would be expert tasks.

(Refer Slide Time: 03:02)

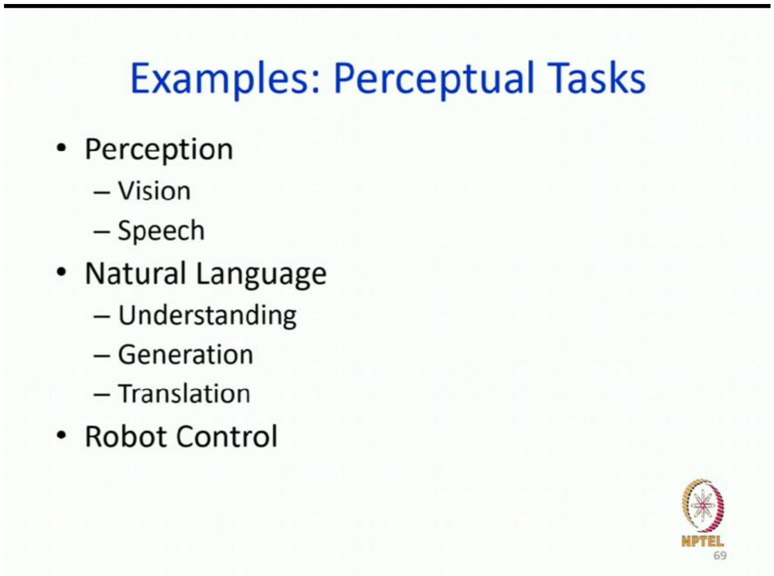

And last but not the least, we also consider perceptual tasks. Now, this is interesting because these are not the tasks that you and I will consider extremely intelligent. Like just being able to see a scene and make sense of it just you know, see, this is a face here, there is an eye here, there is a noses here, or just you listening to my words, and you are able to make sense of what are my words saying, all of these are sort of lower level perceptual tasks than you and I do very intuitively very innately. But a machine does not end up doing it very easily.

And so these tasks were traditionally considered extremely high. Because not that we think of them as very intelligent but we cannot write an algorithm for it. You cannot write an algorithm for why you see a room in this in through your eyes. How does your brain figure out that you know there is a face here their eyes here? Or how does your brain process language? These are questions that you do not know the answers to because you do not know the answers to it.

You cannot write down an algorithm therefore, giving these capabilities to a machine has become has remained a long standing challenge in the last many, many, many years. However, this is where we have made a lot of progress in the last 4 or 5 years. So that is the agent view of AI with several manifestations of the high level philosophical definition.

(Refer Slide Time: 04:13)

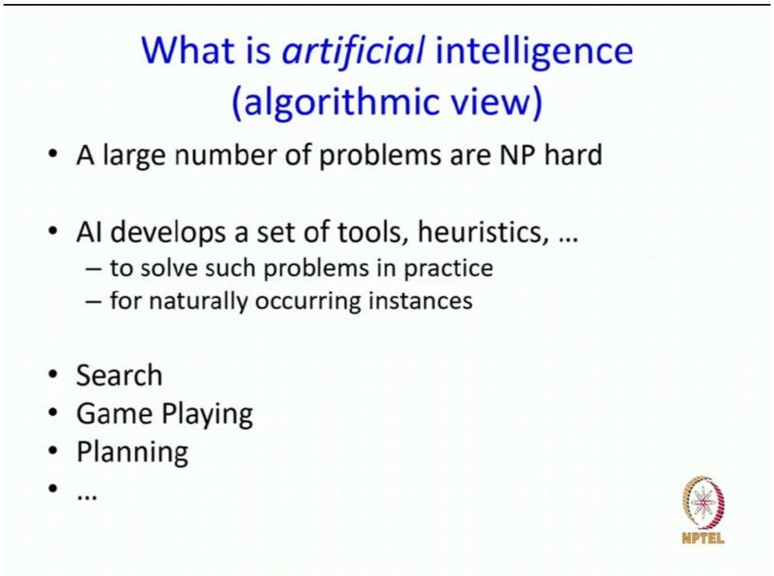

There are other views of AI one view that I quite like in practice is the view of which is the algorithmic view. So, suppose you have done a complexity course, on algorithms course, you must be introduced to the idea of NP hard NP complete problems. And the hard problems are those problems severe so far, we do not have a polynomial time solution to them. And, of course, these problems so hard that, you know, everybody tries to work on it theoreticians, they have to work on it, but their approach is slightly different than the AI people's approach.

And that is important in terms of the mindset. So they can come up with a new algorithm or add they could get the algorithm could be practical also. But their emphasis is to be able to prove something about it. Often you produce some proof, some approximation bound, that you know, it is a hard job. Therefore, I cannot expect to reach optimality. Therefore, I will create an algorithm which is within epsilon of optimal one minus epsilon of optimal or, you know some ratio of optimality, or whatever it is. That is what they like to focus on.

On the other hand, AI systems do not care that much about optimality. If we get to an optimal approximation bound, we do not mind it. Mean it is happy, we are happy to do the theory. But that is not our main focus. Our main focuses Well, these problems are important. And they happen again and again in practice, can we create an algorithm or a set of heuristics in with an algorithm such that in practice, many of the problems that we will encounter, we will be able to solve them effectively.

Now, there are many fuzzy words here. We are not saying that we have an optimality bound we should just come up with some effective solution what is effective, well, better, better than what we could uglier or better than humans? Whatever, it is does not have to be proven it can also be heuristics in there. And moreover, we are not always interested in the random problems. In fact, we are rarely interested in random problems. We are mostly interested in problems which are naturally occurring, which occur in practice due to some application or the other.

Because they have structured and this is something we will talk about a little later, but random problems are hard. And traditions typically prove about any problem therefore, they are also applicable to random problems. AI systems most of more often than not work, try to work on problems which occur naturally, and therefore, they try to exploit the structure that natural problems have. And that natural versus random will become even more important and we will talk about local search.

So, all the problems that we will study in the course will be at least NP hard. At least, and if you do not know what is NP hard, the these are words that you hear for the first time or if you feel that NP stands for non polynomial stop, go home and you know read up on empty first. This is very important. So NP is an important concept in the field of computer science. In fact, the biggest question in the field of computer science is P versus NP, and you must know what NP is in order to even understand that question, and it is not non polynomial.

(Refer Slide Time: 07:25)

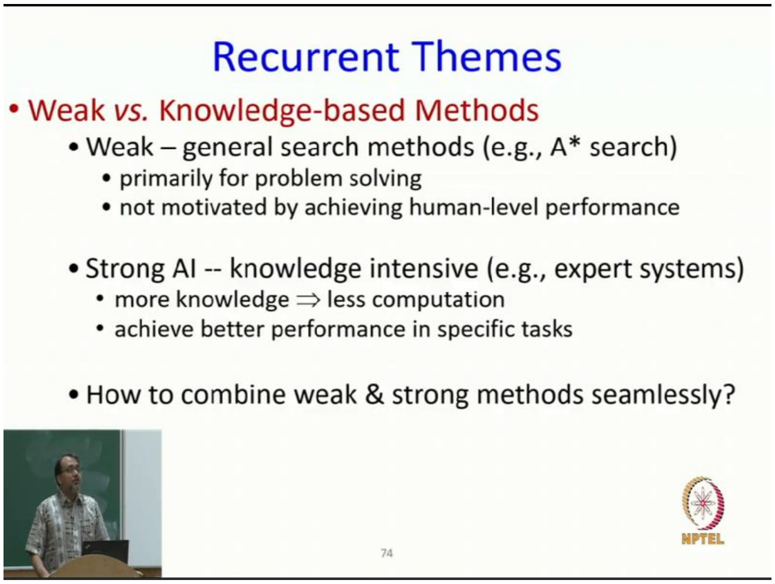

So, there is one team that is going to come up again and again, which is this a weak AI solution? Or is it a knowledge based solution or a strong AI solution? The idea is that we AI consider itself focuses itself on general search methods, general methods which can solve any problem and every problem. Think about the goal is to create this one algorithm that can solve everything and that would be considered week AI.

More often than not this this emphasis the emphasis is not to achieve human level performance. The emphasis is to create these general purpose algorithms. On the other hand strong AI algorithms where the goal is to create you achieve human level of superhuman level performance, they can be guided by week AI solutions, the general purpose AI solutions, but because those general purpose AI solutions in the context of a specific application may not be that effective.

We put in more knowledge in it, the knowledge can come in many forms and that leads to knowledge intensive AI systems and they often have stronger performance and expert systems as an example of it and all your assignments would be example of that. The other the recurrent theme that is very important.

(Refer Slide Time: 08:37)

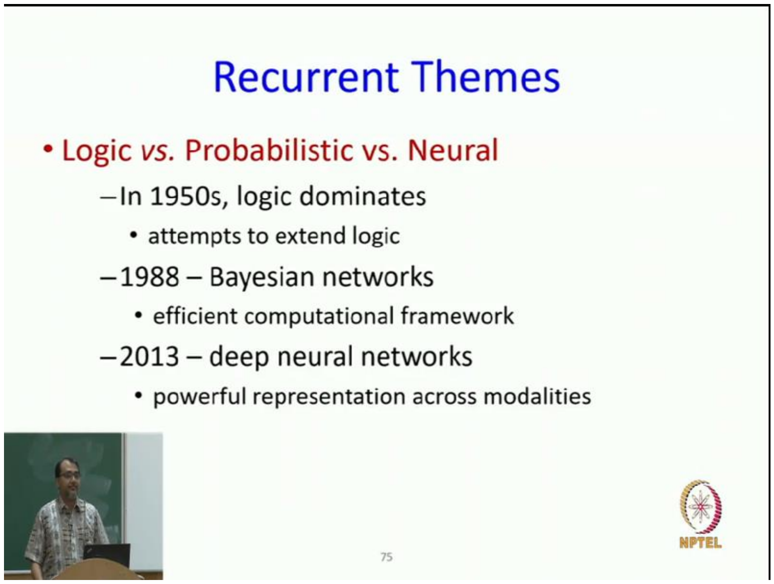

This dichotomy between logical probabilistic and neural. And this is how sort of act of 36,000 feet level this is how AI systems have evolved in their first generation, second generation in the most modern generation of AI. So in the very early days in the 50s, and 60s and 70s. There was a lot of emphasis on search a lot of emphasis on logic, first order logic different kinds. Logic like non monotonic logic, temporal logic, linear tempo, etc. fuzzy logic, many different kinds of logics were developed.

It was thought that a logic based system is the right way to formulate and formulate an AI system. Over time that view went on the side, it sort of got out of circulation because there were limitations of the logical system. And the next versions of systems were probabilistic systems.
These were systems often based on probabilistic graphical models like Bayesian networks, vision because of Bayesian networks.

You know, one person got a Turing award you know who got that you Judea pearl, at UCLA, another one person you must absolutely know about amazing person, spent a lot of time created amazing students. These students were not just good researchers, but they became thought leaders in various sub communities of AI. So Judea pearl has a lot of respect in the field of AI, Judea pearl. So there was a lot of work going on in the 90s and 2000s in probabilistic graphical models.

But then neural networks, which had already always had a parallel existence in the field of AI, they went up, they went down, they went up, they went down, we have talked about it came into prominence around 2013. With the work of, you know, Geoff Hinton and his student's yakun, and his students and yoshua, bengio, and his students, and all 3 of them very recently, in the last month; month and a half got the latest Turing award.

So we are really talking about 3 eras here. And of course, people in the 50s and the AI folks have gotten their Turing awards long, long time back. So there were a lot of people working in logic, a lot of people working in probability, and now a lot of people working in neural AI. And of course, these things are not exactly in conflict with each other. So for example, a lot of research works us and my colleague, Professor paroxysmal doing today is that we are trying to combine neural systems with logical systems.

So we call it neural symbolic AI we are not the only people in the world to do this many people are doing it, but the point is that these systems each era has some strengths and some weaknesses. And so, can we bridge the gap is something that people have been thinking about, even in the earlier era, they were probabilistic logical systems. They were called statistical relational learning, for example, so one of the leading members of statistical relational learning is Professor (())(11:33).

Who worked on Markov logic networks one particular way to combine first order logic and probability. So in this course, now, starting now, we are going to study all these topics. From the phase one, the old AI, we are traditional AI, we will study search constraint, satisfaction, logic and games for topics. From phase 2, we will study uncertainty, Bayesian networks, reinforcement learning and Markov decision processes.

And in phase 3, we will study deep neural networks. And if time if there is time we will talk try to connect deep neural networks with reinforcement learning to get to deep reinforcement learning. So this is our core structure course plan. So, this is a good point to stop and we will meet tomorrow thanks