Artificial Intelligence
Prof. Mausam
Department of Computer Science and Engineering
Indian Institute of Science Education and Research-Delhi

Lecture-66
Decision Theory:
Steps in Decision Theory
PART-1

(Refer Slide Time: 00:17)

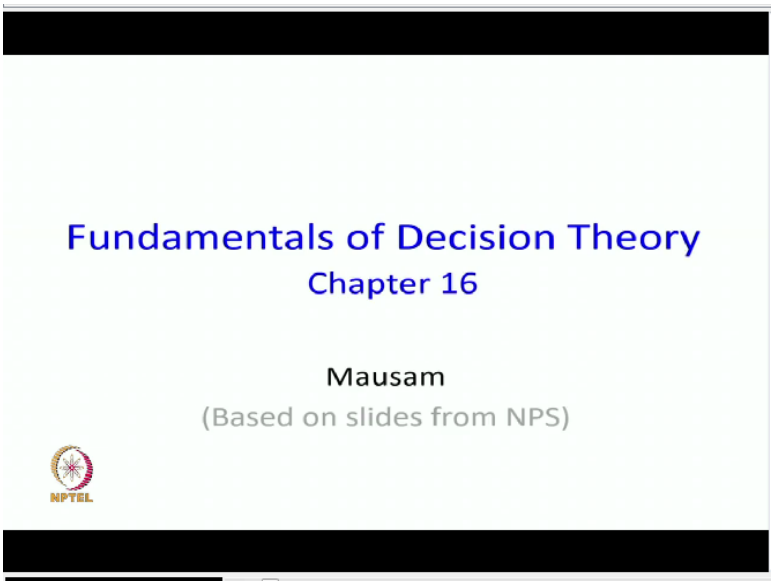

Okay, so here is where we are. We have done a lot of work in this course so far, talking about how to solve a traditionally AI problems, what is a traditionally AI problem, you cast it has a search problem and the reason you cast it as a search problem is because you want to reach some goal state or you want to reach some state where some constraints are satisfied. You can think of it as a goal state.

In some cases, the path to the goal is important, in some case the state itself is important. The world is mostly deterministic you know the consequences of each action exactly. And then the question is how do we solve such a problem. And why we saw such a problem is because it is a general problem, it can, you can use that framework to formulate pretty much any deterministic problem in the world, single agent.

And we also did one extension of it in the case of adversarial setting where there was uncertainty in what the opponent is going to play. But this uncertainty was of a very specific nature that we knew that whatever we were trying to maximize the opponent is trying to minimize. So that was a very specific situation there we did the search again. Then we said, that is good. But of course, this would not lead us to a very strong AI system.

We need to give knowledge to the AI system. And we started giving knowledge to the AI system and we did 2 main knowledge representation languages. The 2 I would say with relatively most popular knowledge representation languages, one was the language of logic, and one was the language of probabilistic graphical models or Bayesian networks, right. And in logic, we did not do first order logic.

But and there are many, many different kinds of logics, there is a full, you know, 70 year history in logic. And so there is a lot that goes on in the field of AI in logic, but we touched on that. And then in the last 20, 25 years, there has been a lot of work in probabilistic models, and we talked about that. And in this knowledge representation language, we will express the world or how the world behaves, how the state variables in the world random variables in the world interact with each other.

And what should be true and what should be, you know, how should we change our belief, and then we could do inference on top of it. And we learned several inference algorithms that given a new query can tell you whether it is true or false or what is the probability of that depending upon which knowledge representation language we are talking about. I would say these are 2 different pieces that we have studied.

One is the search algorithm and decision making in deterministic environments and one is knowledge representation inference algorithms in both deterministic and probabilistic environments. So, therefore, what remains the most critical part that remains in terms of the big picture is decision making in probabilistic environments or uncertain environments right.

(Refer Slide Time: 03:26)

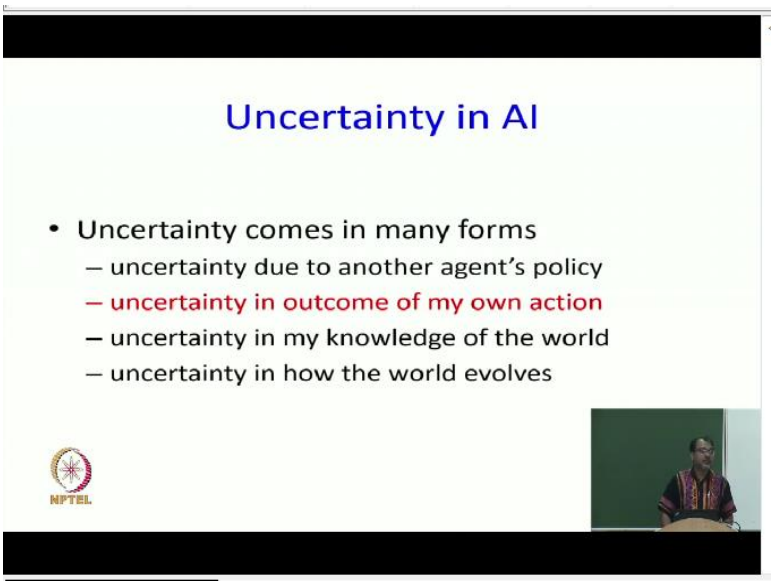

And uncertainty in AI divides traditionally AI and modern AI. So, if you go back, you know starting 55 and so on so forth. A lot of the work that kept happening until 80s was using logic representations was using deterministic environments in deterministic scenarios. It is not exactly true, but there is an approximate sort of high level bit, but starting the 90s starting with Judea Pearl started giving his points of view starting with them.

And then leading up together full all of 90s, a lot of 90s, all of 2000s, the world was mostly probabilistic. Life was probabilistic, AI was probabilistic. And that I would consider modern AI. If you really want to fine grandly splitted further, I would say, phase 1, phase 2, phase 3, phase 1 was logic, phase 2 was probabilistic. And now you can say that phase 3 in the last 5 years is neural, right, neural networks, deep learning.

When we will touch on that towards the very end of the course. This is also goes approximation because neural networks has had 2 other lives earlier. It came up, it went down, it came up, it went down, and then it came up again in the last 5 years. Maybe this time it is here to stay, who knows right. We do not know this has happened with neural networks many, many times, then we will talk about that in weekend.

So today's world is not only probabilistic it is probabilistic, plus, plus and this plus, plus part comes from the neural models. But we will get to that when we get to that. Let us talk about

uncertainty in AI, which has been a significant theme in the last 10 20 years. And you have to realize that uncertainty in an environment comes from many, many forms and what was happening. And we discussed this some time at the time of Bayesian networks also, that logic is brittle.

You cannot describe everything you know, in your language. And if you cannot describe everything in your language, then your logical systems cannot handle many scenarios that actually occur in the world. And because of that, they are just not able to scale up, right or scale up or as in give good performance, right. And uncertainty comes in many, many, many forms. So if there is a multi agent scenario, you do not know what the other agent is going to do.

In fact, in minimax, we assume that what we are maximizing the agent is minimizing but in practice, the agent can do anything, right. If you make a game bot of yourself, you do not know exactly what the opponent is going to do and so you have to be able to respond to them. In fact, some people try to learn what kind of mistakes the opponent's bot is doing right. And so that also is part of you know, there is uncertainty.

And let us try to figure this out. Uncertain the outcome of my own action, I try to do something and of course, I fail or I end up with unusual consequences. And we have discussed this example, several times that I have a physical robot, a physical robot, which is acting in the world, it wants to get to location 10, 20, it may not reach exactly 10, 20. It may reach 9, 19, right or 21, 21. Oh, no, that is too bad 11, 21 right.

So you do not know exactly where you are. If you are playing robot soccer, you do not know exactly where the ball is right. Did I give you this example. We have once upon a time. This dog, Sony dog, which was playing soccer, kept going in circles like this. And nobody understood what is going on. Why is it not playing Football, why is it going in circles like this and suddenly it was just basically rotating on its Z axis.

And then later people realized at that look, the ball was orange and it was programmed to track the ball that was orange. So far, so good. But in the audience, there was a guy wearing orange

shirt, which was just moving around, you know, doing something and the dog thought that, that orange shirt person is the ball and it was rotating on its axis trying to track the ball, there is uncertainty in the real environment.

Not necessarily because the environment is uncertain, although sometimes it is, you know, you may change seats, you may talk to the other person, you know there is uncertainty, you have to constantly keep track of it, but there is also a uncertainty in the sensors. Even if the world is deterministic, you may not get to know exactly right, you may only have some probability distribution about where you are or where the ball is and so on so forth.

So uncertainty in my knowledge of the world, therefore, where is the ball. That is also the kind of uncertainty and then later we will study uncertainty in the way the world evolves. We do not even know how world evolves. If I give you an assignment, whether I am going to do hard grading or easy grading, you do not know this. Your seniors may have told you something. But that changes over time also, right. I also evolved.

So the world evolve around you, you know, the dean in the institute changes and suddenly they bring new rules. You cannot predict exactly how the world evolves, you have to somehow make sure that you are safe, right, etc., etc. So there is lots of different uncertainties in the world that we somehow have to grapple with deal with and handle. And today, we are going to start talking about the uncertainty in the red primarily that is uncertainty in the outcome of my own action.

And coupled with uncertainty in the knowledge of the world okay and so, we will take a very simple example. But then after a couple of classes we will build it into a full blown model of Markov decision process. But today we will not.

(Refer Slide Time: 09:09)

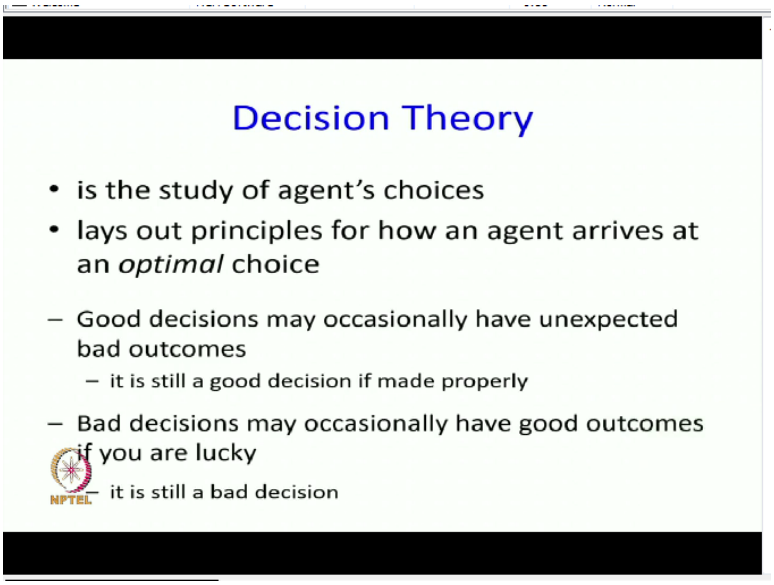

We will just start talking about decision theory, now decision theory is this full field which studies how agents should make choices. It also studies sometimes how people make choices, but that is a different story in general, because people do not use decision theory in making their choices. We have discussed a time and again that people humans sometimes emotionally make a decision. And then later if you ask them, why did you make a decision they will cook up a story.

You have done it, I have done it, we all have done it. This is the way we work. I am not saying this is good. This is bad. This is just who we are. But we cannot give this to the machine. We cannot have the machine take random decisions. Kill itself you know commit suicide, break the glass, throw water on the floor. You know, we do not want machines around us like this we have people for that.

So, machines should be taking reasonable rational decisions and this is what we had decided long time back, what is the definition of AI it is rationality, long term rationality, in turn in its actions is sort of what we are decided once upon a time, right. So, in light of that, we want some theory on how should the agent take decisions right. And decision theory lays out the principles for how an agent arrives at some optimal choice or an optimal choice or the choice near optimal or whatever best it can do in terms of bounded rationality.

And again, I want to re-emphasize I have done this twice already in the course. So, I will not belabor it but good decisions may lead to bad outcomes, that does not make the decision bad and bad decisions may lead to good outcomes and that does not make the decision good. So, you took a completely crappy decision suddenly stars became lucky I mean for you and gave you the right best cards and you won or you did something or you got lots of mark, you only studied one chapter for the exam out of 5.

And for whatever reason 60% of questions came from one chapter. And initially you did not know that this is going to happen. But then you are very happy I did the right thing. I only started one chapter, that is foolish thing. You still got lucky, you still got probably good marks. But that does not make a decision correct. And so it is important to distinguish how the world worked out for you, or whether it did or whether it did not.

A lot of people you will find a very unhappy life because the world has not worked out for them the way they wanted it to work. You know, we will keep that on the side when the world works sometimes and it does not work sometimes for different people. But the question we are asking is not whether the world works for you. The question we ask is whether you took the right decision at the time or not.

And if there is some information that you get after taking the decision, which worked or did not work, in your favour, that does not mean your original decision, right or wrong. That has no bearing on your decision. Your decision needs to be taken based on what you know at that point in time. And still, if you make a terrible decision, it is still possible to take a bad decision. So let us look at what is the formal mechanism for how we take a decide on a decision in decision theory.

(Refer Slide Time: 12:19)

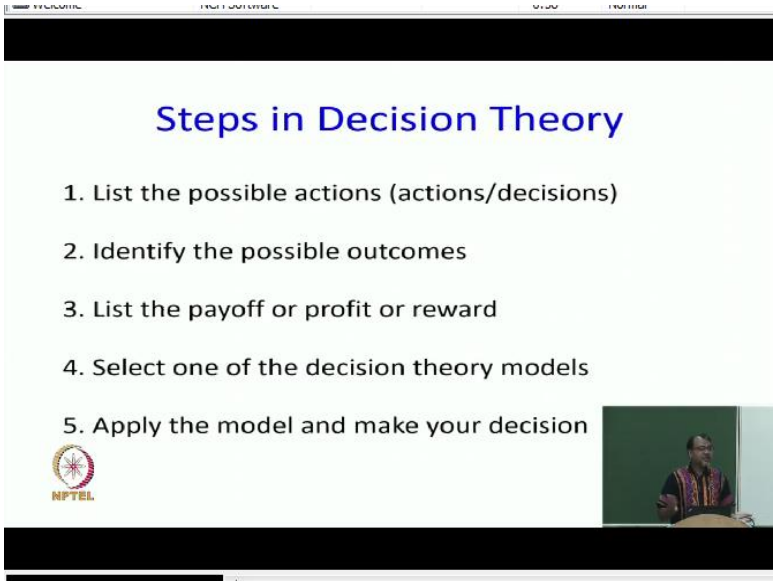

And the steps are very simple. We will first list all the possible actions. This is something I have found as a difference between a rational is people and emotional is people, right. There are people who would say I cannot do anything about it, therefore I am unhappy. It is a very good question for you to think about. Is this a fair enough statement. Should you alternatively not say that because I cannot do anything about it, let us forget the weather I should be happy or unhappy about it.

This is not in my hands, will let the life play out as it plays out. If on the other hand there is something in your hand that you can do about then you do the right thing. See in our life the only thing that is in our control is our actions. So if there is something you cannot do anything about that should not take you know too much of your thought process. It is not important anymore but it is going to play out the way it is going to play out.

The cards have been dealt, I have a friend whose son cannot process fats. It just cannot process fats. Think about a little child who cannot process fat at all. Like now this child is 12 or 13 I believe plays games and everything, has to be given food completely zero oil. Now you cannot know if you can do something about it. Yes. Do you see the doctor see this, see that, you know figure this out.

Once you figure out that there is nothing you can do anything about the only thing you can do is to figure out how to make the best give the best life to the child. You know, I have a friend he is now 80 plus, he has only one hand, he lost his hand when he was 18. Maybe for a few months he was unhappy. But then today, if you meet him within a few minutes, you will forget that he does not have a left hand.

Just forget, he has figured out absolutely everything that he drives the car by talking US so he does not have to change gears. It is automatic, but that is all right. The only thing I have seen him not doing and I doing well, okay, I can play tabla, he cannot, 1 and 2, I can talk on my phone while driving the car which he cannot. Everything else I have never felt that there is something he cannot do.

So if as they say life gives you lemons, you make lemonade. So if there is something you cannot do anything about, who cares. Let us figure out what you can and work and focus your energies on that. So the first thing you need to do is list your possible actions because that is your space. That is the space in which you are playing, that is the space in which you are making a choice. If it is not there, it should not even be part of your thought process ideally.

Of course, now with each action you identify the possible outcomes and list the payoff profit or reward, same thing we can call it in different ways payoff or profit reward, cost, cost is negative of a reward. You can list that out. And then based on the knowledge that you have select one of the decision theory models and apply the model and arrive at a decision. This is how at least you will do it in the machine.