Artificial Intelligence: Search Methods for Problem Solving
Prof. Deepak Khemani
Department of Computer Science & Engineering
Indian Institute of Technology, Madras

Lecture – 02
Prologue Winograd Schema Challenge

(Refer Slide Time: 00:16)

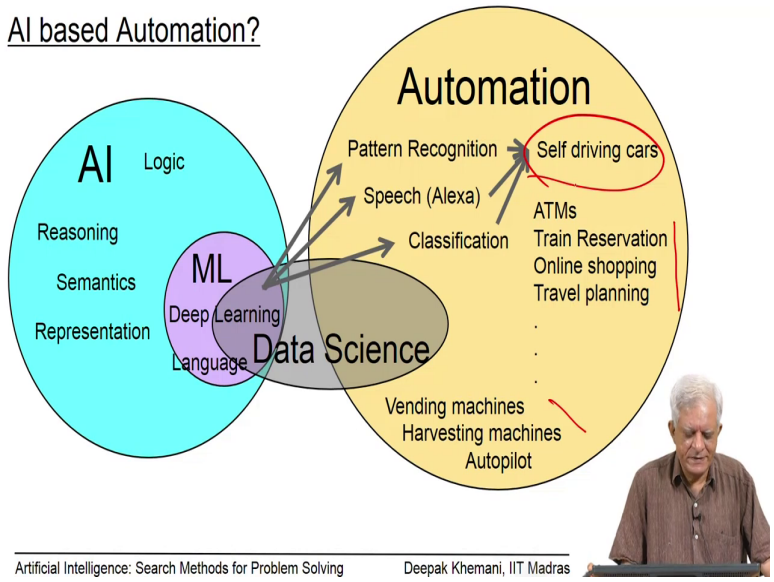

Now, we often confuse nowadays automation with AI and I would like to kind of emphasize here that they are quite different features you know. When we talk about AI, we want to talk about intelligent machines, we want to talk about logic and reasoning and semantics and representation and so on. We also want to talk about machine learning, deep learning, language processing and so on.

And we also want to talk about data processing because you know machine learning works on data. So, if you want to build an algorithm to learn a decision tree, you have to learn from

training examples so, data is the key there. But automation is a totally different cup of tea together essentially. I mean you have things which have nothing to do with AI. Things like train reservation, online shopping, travel planning. Travel planning may have a little bit of influence from AI. Vending machines, harvesting machines and so on and so forth.

But also that a lot of automation benefits from machine learning and data science specially things like self driving cars because they have to learn to recognize images and you know do pass planning and thing like that. So, I just want to in this slide differentiate between these different terms which are so popular nowadays machine learning, AI, data science and automation essentially. Our focus is again on the left part of the screen which is AI or symbolic AI essentially.

(Refer Slide Time: 01:48)

Robots is another field where the last 10 years, 20 years there has been tremendous amount of progress and for some reason, the progress has been kind of focused in a few pockets in the world.

A lot of progress in Japan and Korea, but also a lot of work in the US with but companies like Boston Dynamics now producing robots humanoid robots which can do Flipkarts and you can just look for a video from Boston Dynamics to be impressed by that. But what is happen in this current times which is when we have this covid epidemic coronavirus epidemic? There has been an increased demand for robots for various activities.

Now, this demand has come from many places for example, hospitals would like robots. They do not have to look like humans, but they have to be autonomous machines and they have to be intelligent enough to find their paths to the right place, serve food to the patients.

Even hotels are talking about serving food in restaurants by robots and robots are being considered for surveillance that you know looking at parks that if you the government has asked you to do social distancing, robots can go around, monitoring the parks doing surveillance to see people are actually doing social distancing and maybe send them a message or you know some something like that, even drones have been used for that kind of thing.

But robots have also been developed in especially in countries like Japan and perhaps in the US where the populations are ageing as companions. So, we have seen various examples and we will see some of them in the introduction that we will follow these particular lecture.

So, here is an example. In fact, this might be repeated in that talk as well about Nadine which is robot which looks like a human and for a person who is lonely and needs company may be such a robot especially if it can interact like a human would be a good idea essentially.

(Refer Slide Time: 04:02)

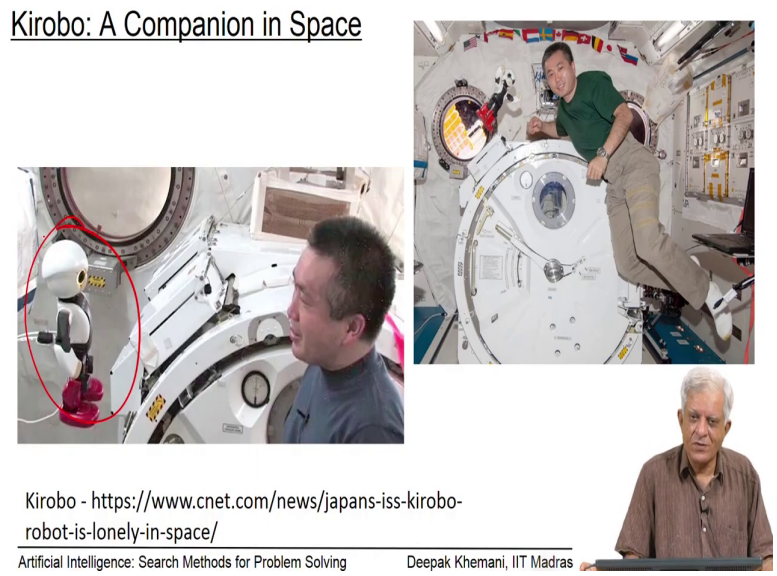

This idea was also there in this small tiny robot as you can see here called Kirobo which the Japanese Space Agency sent to into space to give company to robots and you can see two pictures here, when astronaut is enjoying the company of a robot essentially.

(Refer Slide Time: 04:22)

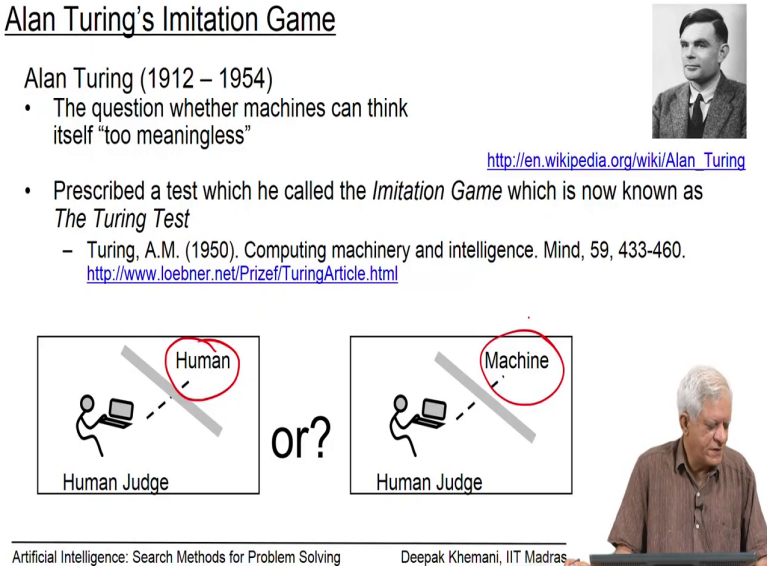

Now, coming back to this idea of thinking. So, I keep saying that, you know we are interested in machines which are thinking machines and so on and this is not something new, I mean this has been at the very genesis of this whole area of artificial intelligence and even earlier there was this debate about whether can machines think or not and we will talk about this again more in the introduction.

But Alan Turing introduced this thing imitation game and he started off by saying that the question whether machines can think itself is too meaningless that there is no definition clear cut definition of what is thinking to be able to say whether machines can think or not essentially.

Alan Turing was a British mathematician who played a seminal role in code breaking during the Second World War by designing a machine called enigma and there was a movie called the

Imitation Game which curiously enough was more about enigma and less about the imitation game itself.

But anyway, he created this game called imitation game and he said that if machines can pass a imitation game, then they are intelligent and of course, we call that game as Turing test essentially and he publish the paper in 1950 which is quite some time ago.

The idea of Turing test is simply that a human being is chatting with someone and it is not clear. So, there is a judge and the judge is chatting with someone. It is not clear whether the someone is a human or whether the someone is a machine and the idea behind the imitation game is that, if the machine can fool the human judge a sufficient number of times that it is talking that he is talking to a or she is talking to a human then, the machine has pass the Turing test essentially and we will discuss Turing test again little bit more. But it is essentially based on linguistic ability is that you know most of the time when somebody can produce a sequence of words, we often an ascribe lot of intelligence to them.

(Refer Slide Time: 06:31)

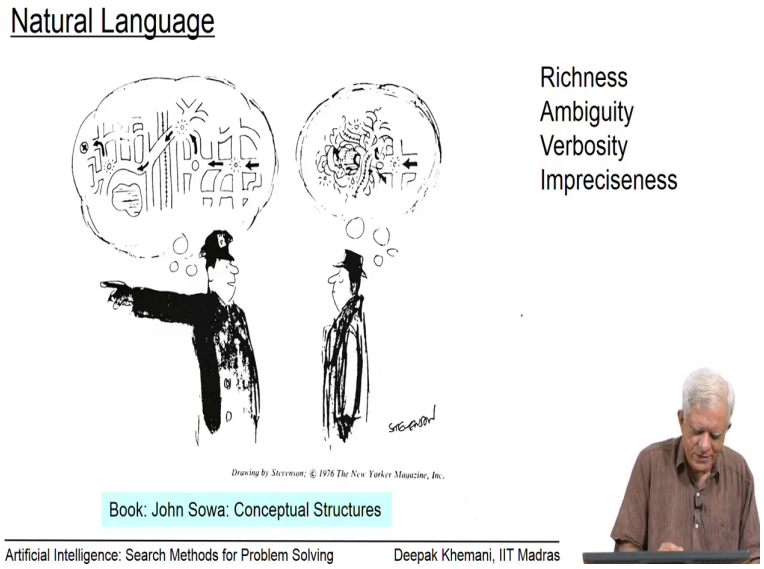

But processing language is we have made tremendous progress in processing language, but its still hard problem and that is largely because richness, ambiguity, verbosity, impreciseness of languages essentially. You know you can say the same thing in so many different ways and what you say may have different meanings.

So, languages of complex feature and we are still struggling how to deal with it even though we have programs like Alexa which can listen to certain commands by you and act on those commands still we are still far way off and this is a book which is not a which is not a new book this cartoon which was produced in this book called Conceptual Structure by John Sowa who worked in the area of knowledge representation and he introduce this notion of concept conceptual graphs. He reproduce this idea of natural language processing. Here, you can see a

policeman trying to give directions to stranger about how to get to a particular place and this kind of illustrates the kind of picture that the stranger is creating in his head essentially.

(Refer Slide Time: 07:48)

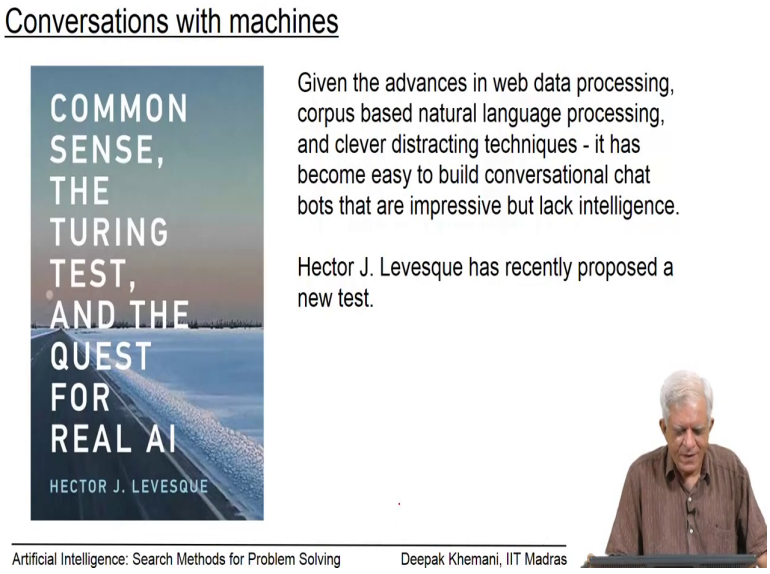

Now, quite recently ah, Hector Levesque who is known for his work and knowledge representation and in fact, he has written a very nice book along with Ronald Blackman which I use for the other course which I teach which is the knowledge representation and reasoning.

He has recently come up with this popular science book, it is meant for lay people and I would ask you to try and get hold of this book or look at some articles by him on the web you can find in which he again talks about the quest for real AI what do you mean by real AI? What are you looking for here? So, he talks about things like as you can see in the title of the book Common Sense, The Turing Test and the Quest for Real AI essentially.

Now, what he is saying essentially is that given the advances in web data processing, in corpus based natural language processing that you do not represent language along with its meaning, you do not represent word with their meanings. It is just that you have large collection of corpuses of text and language and you can somehow process them to do meaningful work and also clever distracting techniques.

So, when we discuss Turing test, we will see that a machine can easily fool a human being by suddenly you know changing the topic and say tell me more about your family or why did you do whatever you. So, if you say I feel like having a pizza it might simply say why do you feel like having a pizza.

So, it would you know appear that the machine is thinking and so on. So, you get there are lot of clever tricks that people have employed in implementing programs for this contest called the Loebner prize which we will also talk about in the introduction in the next few talks, but there are clever techniques they use essentially and specially now with the so much data available on the web and you know search engines going through the data it is become easy to build conversational chat bots that appear to be impressive, but they lack intelligence essentially.

So, in his book and elsewhere in the papers that he has published, Hector Levesque can essentially propose a new tests. So, I will just talk about that a little bit and it will kind of bring the focus back to the kind of thing that we are talking about when we talk about thinking machines.

(Refer Slide Time: 10:19)

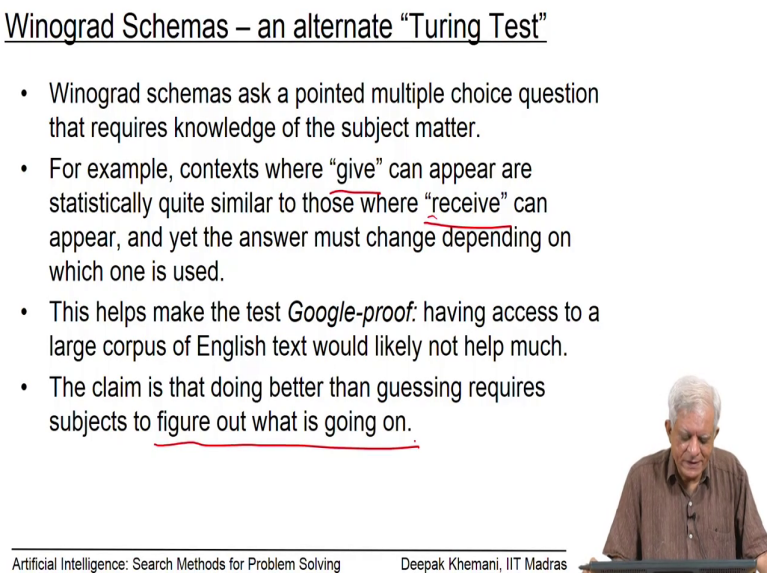

So, this new test is like an alternate Turing test it is called Winograd schemas and these are pointed multiple choice questions that requires knowledge of the subject matter.

So, Levesque's saying that let us not rely on the fact that machine can converse with you in language to be the criteria for deciding whether the machine is intelligent or not, let us rely on whether the machine can answer some pointed questions or not and these kind of very specialized, very structured questions the calls as Winograd's schemas essentially. So, we will talk about that in the next few minutes.

So, for example, if in some contexts, the word give can appear and this context may be statistically quite similar to those where the word receive appear. Now, as you can see, give and receive are in some sense opposites of each other.

You give and the other person receives essentially and the two words can achieve and the answer that the agent being tested must change depending on which one of the words that we have used. So, we will see examples of this essentially.

So, essentially the goal is to move away from the felicity of language, the ability to you know generate colorful sentences, but more to being able to answer certain questions correctly. So, Levesque says that this makes a test Google prove. That you cannot have a program which has access to Google and can search for data and for similar sentences or something and you know somehow generate an answer which you may think is a good answer essentially.

And the claim he makes is that doing better than guessing requires a subjects to figure out what is going on; to figure out what is going on and that is what we will take to be the goal of an intelligent program.

(Refer Slide Time: 12:30)

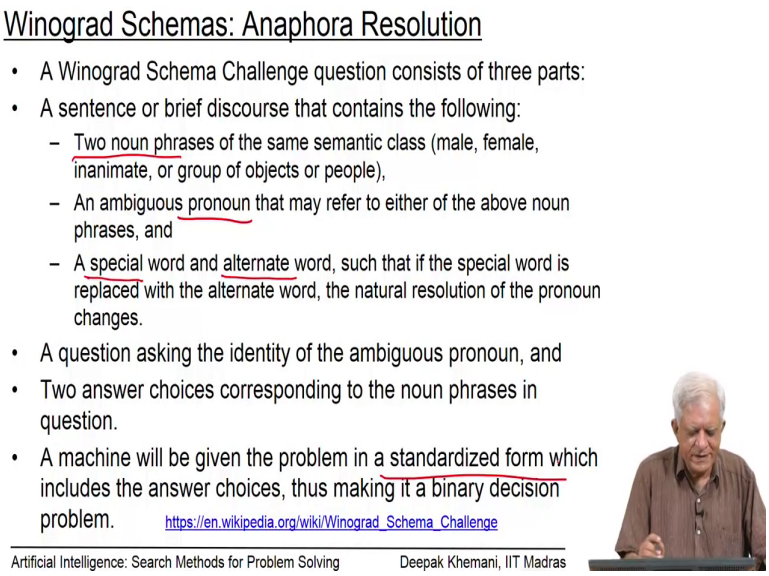

So, Winograd schemas are what a linguist would called as Anaphora resolution. An Anaphora basically is a pronoun which needs to be you need to identify what it is referring to.

So, a Winograd schema challenge consists of three parts. The first part is a sentence or a brief discourse that contains the following; two noun phrases of the same semantic class. For example, male, female or you know other kinds of semantically similar things.

An ambiguous pronoun so, the keys on the pronoun that may refer to either of the above noun phrases. So, remember that there are two noun phrases and the pronoun can refer to any one of them and the question we ask the machine is what does it referred to?

A special word and then alternate word. So, for example, give and receive such that the when you replace a special word with the alternate word, the answer will change essentially.

So, you need to be able to figure out what is happening to be able to answer such questions fluently and then you ask a question asking the identity of the ambiguous pronoun. Remember that we are talking about a pronoun which is ambiguous and we give two choices in this particular schema corresponding to the noun phrases in the question.

And a machine will be given the problem in a standardized form which includes the answer choices. So, all the machine has to do is to be able to say whether the first answer is a correct answer, or the second answer is the correct answer. So, it is like a multiple-choice questions essentially. The thing is that to answer correctly, you need to know what is happening.

(Refer Slide Time: 14:21)

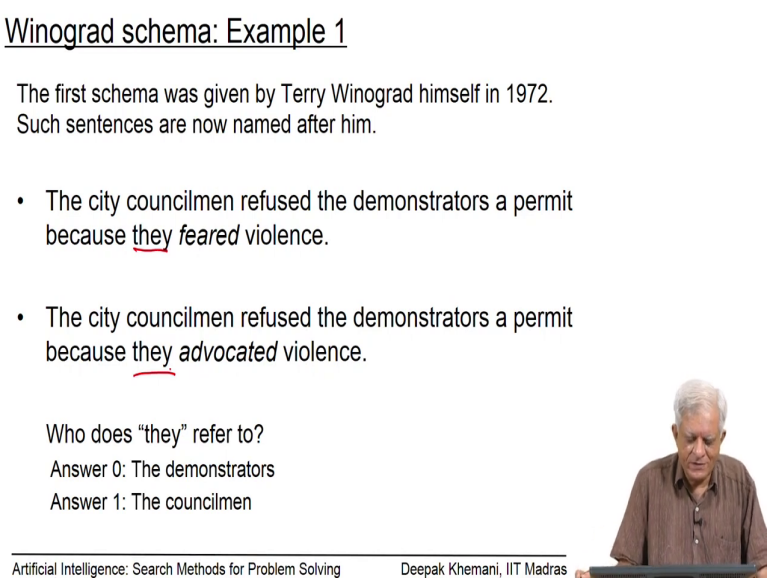

So, let us look at some examples. The first example was in fact, given by Terry Winograd, who was a researcher in AI in the early 70's he was working on natural language processing programs, he wrote a program called SHRDLU and so on and the first schema was in fact, given by him 1972. So, here is a sentence the city council refuse the demonstrators a permit because they feared violence.

Now, remember what did we say? We said that there is a special word and then alternate word essentially. So, here the special word is feared essentially the city council refused the demonstrators a permit because they feared violence. The alternate word gives us this new sentence the city council refuse the demonstrators a permit because they advocated violence.

So, you can see that feared and all advocated they change the meaning of the sentence essentially. Then the simple question that you ask the machine is that what does that they referred to. So, when you say they feared the violence or they advocated the violence what does the pronoun they refer to essentially. So, this is called Anaphora resolution does it refer to the demonstrators or does it refer to the council. So, this is an example of a Winograd schema which in which the you cannot do guesswork and you cannot find sentences like this on the web so that you can answer them easily essentially.

(Refer Slide Time: 15:56)

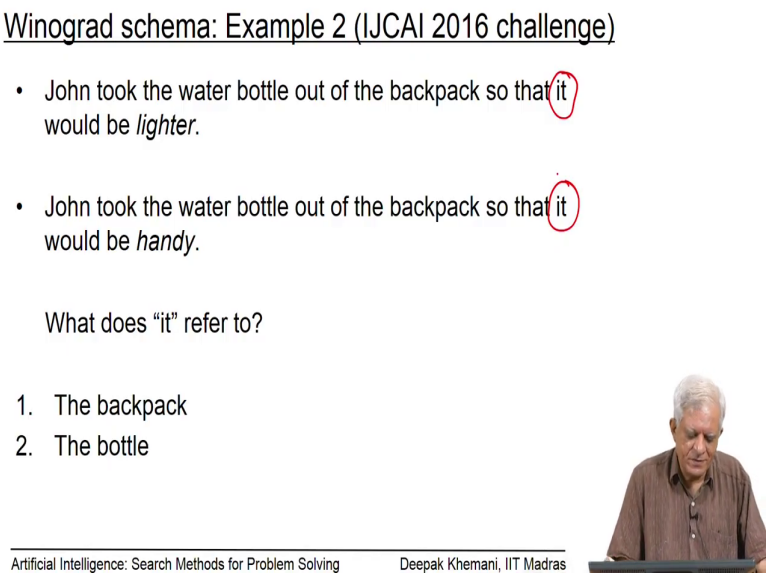

Here is another example which was used in the International Joint Conference on AI in 2016 challenge. The first sentences is John took the water bottle out of the backpack so that it would be lighter. The second is John took the water bottle out of the backpack so that it would be handy.

So, notice that the two sentence are identical except for the special word and the alternate word. The special word is lighter or handy in this case and depending on which word you put in a question so obviously, you are not going to give both questions to the machine which is being tested. You will give one of these two sentences and expect them to produce the answer and then you would ask what does it refer to/ So, that we said it would be lighter right. What is this it? Is it the backpack or is it the bottle?

Again of course, you need to understand that what is happening that if we take out something from ah backpack, the backpack will become lighter or if you take it out then it will be more easily accessible to you. So, you have to be able to in some sense create this whole imaginary of what is happening to be able to answer this questioning, you cannot just do it by you know doing some language processing by you know searching on the web and things like that. People do not probably talk about such things on the web.

(Refer Slide Time: 17:12)

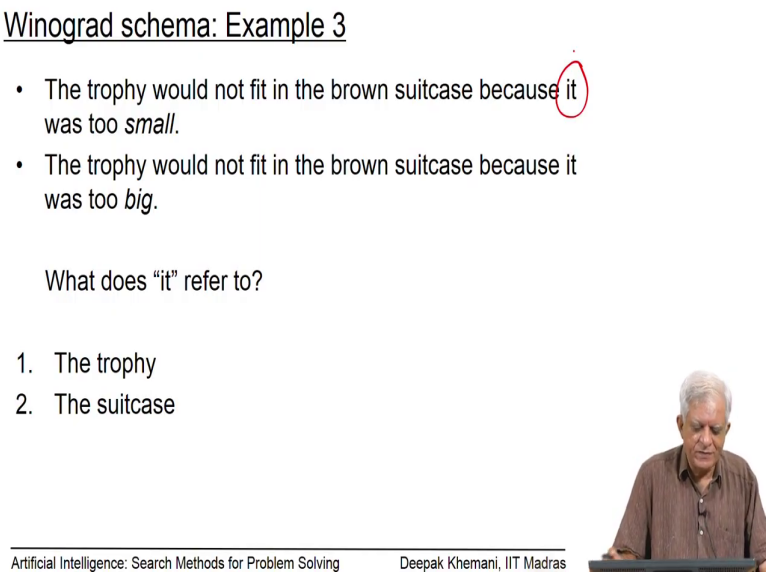

Here is the another example. The trophy would not fit in the brown suitcase because it was too small. Again, remember it is a pronoun and you need to figure out what does it referred to. The trophy would not fit into the brown suitcase because it was too big. So, what does it referred to? Is it the trophy or is the suitcase? As you can see, for one sentence it will be the

trophy and for the other sentence it would be the suitcase. So, if a machine can answer this questions, then we can say that yes it is doing some intelligent processing in the background.

(Refer Slide Time: 17:44)

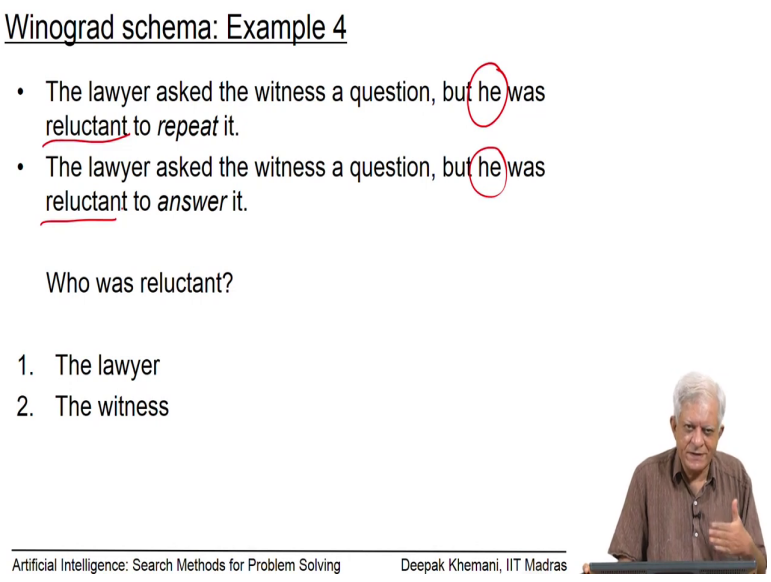

Here is one last example. The lawyer asked the witness a question, but he was reluctant to repeat it. The lawyer asked the witness a question, but he was reluctant to answer it. So, who are we talking about? Who is his he or who was reluctant? That is a question.

Was it the lawyer or was it the witness? As you can see for one sentence it is a lawyer and for the other sentence its the witness and the machine can give the correct choice ah. So, let us say you ask a 100 such questions and if the machine produces a score of 50, then you can kind of you know say that it is just doing guess work, it is not it does not know what is happening.

If he does much more than you can you know say that it can answers some of these questions. This is just to kind of highlight the fact that that intelligent activity requires much deeper stuff than simply being able to process data which is available at the surface level on the web.

(Refer Slide Time: 18:36)

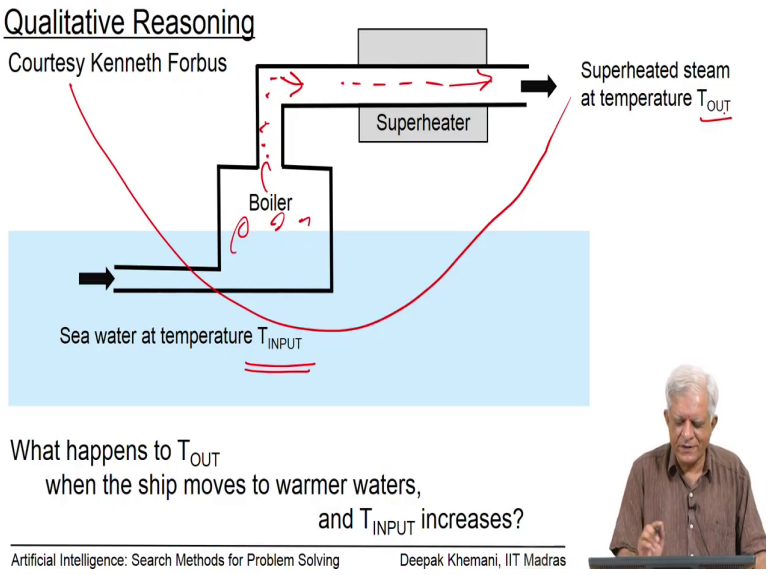

Here is another of my favorite examples ah. It is also not a new example qualitative reasoning was a flourishing area in the early 80's, mid 80's and this example is due to one of its practitioners Kenneth Forbus and the example is like this that the we are talking about the ship.

So, there is a ship which I have not quite drawn. So, imagine that there is a ship and you know this is part of a ship or something like that and the ship has a boiler which takes in seawater. The seawater is at an input temperature of T input we call it and in the boiler the water boils and then steam is generated and steam goes up in the pipe essentially. Then, this steam which

is flowing through this pipe is super heated that it is heated beyond its boiling point which is 100 degrees. So, what we get outside is superheated steam which that some temperature called T out essentially.

So, we are not talking about numbers here. You know most of the time in physics, people work they you know they say this is a specific heat of water and this is the latent heat of boiling and things like that.

We are not talking about numbers here, we are simply asking questions at a qualitative level and the question that Forbus had ask was supposing the ship moves from arctic to the tropical area which means that it moves from a colder area to a warmer area which means that temperature of the input water goes up essentially. In arctic it may be close to 0 in somewhere near Chennai it may be close to I do not know 25 or something.

And so, what happens when T input changes essentially and you can answer this question by doing reasoning without numbers which is called as qualitative reasoning and I will leave this as a small exercise for you to do. The only hint I will give you here is that when the input water is hotter, it will boil faster essentially hm. So, take that and try to figure out what will happen to T out.

(Refer Slide Time: 20:51)

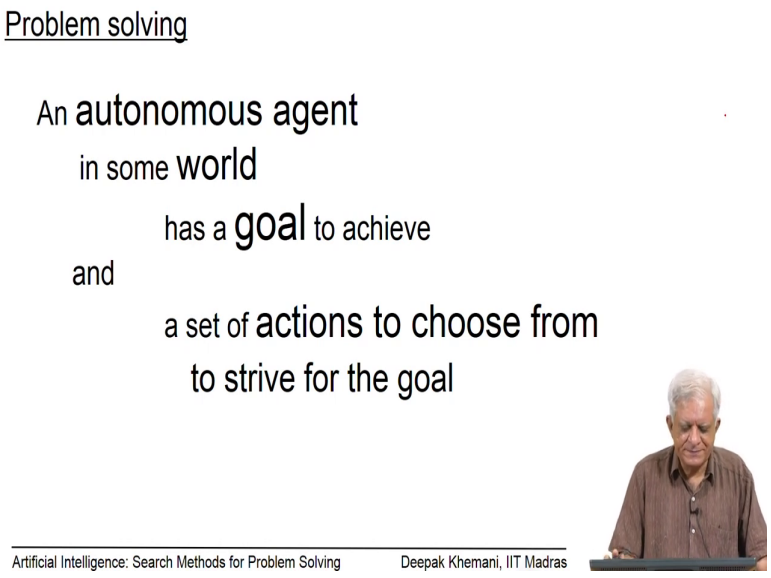

The focus of this course that we are doing is going to be problem solving and we say that an autonomous agent in some world has a goal to achieve and the agent has access to a certain set of actions which it has to choose from to strive for the goal. So, this is what we understand by or we mean by the word problem solving and we want to build systems which will be able to do a problem solving.

(Refer Slide Time: 21:17)

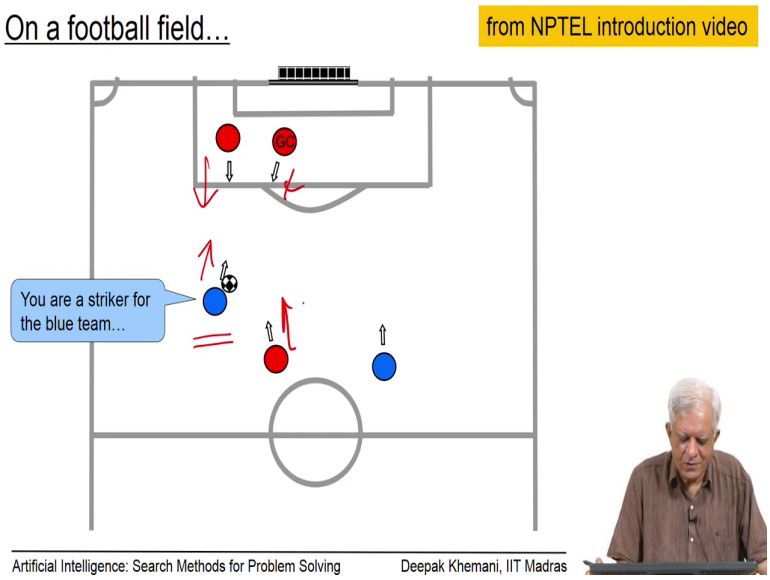

So, in my introduction video which you must have seen on the Swayam website or the NPTEL website, I have used this example of you being a football player here with the ball and you are running towards the opponents goal and the opponents are closing in upon you. So, what is the right action for you to do? So, essentially you are in some situation and you have to make some decisions and how do you make those decisions is this whole area that we will call as problem solving.

(Refer Slide Time: 21:50)

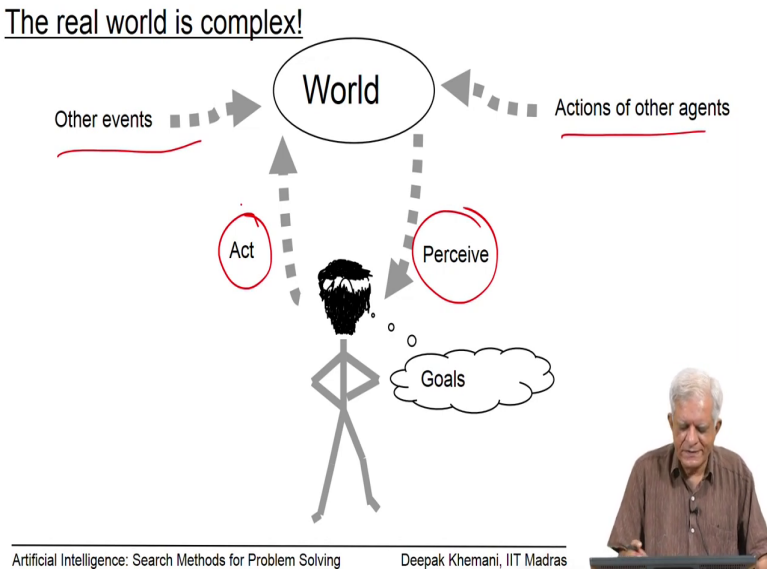

Now in the real world it is very complex ah. There are many agents in active, there are many things happening for example, rain and storms and viruses and so on. You know on top of it your perception of the world may be imperfect, you only know certain things, you do not know for example, sitting here in this room in Chennai whether at this moment it is it is raining in Delhi or not.

So, our perception is incomplete. Our actions also may be imperfect essentially that we may say that do this action, but in the real world that action may not happen. You know if you are a basketball player, you would know this very well essentially.

(Refer Slide Time: 22:36)

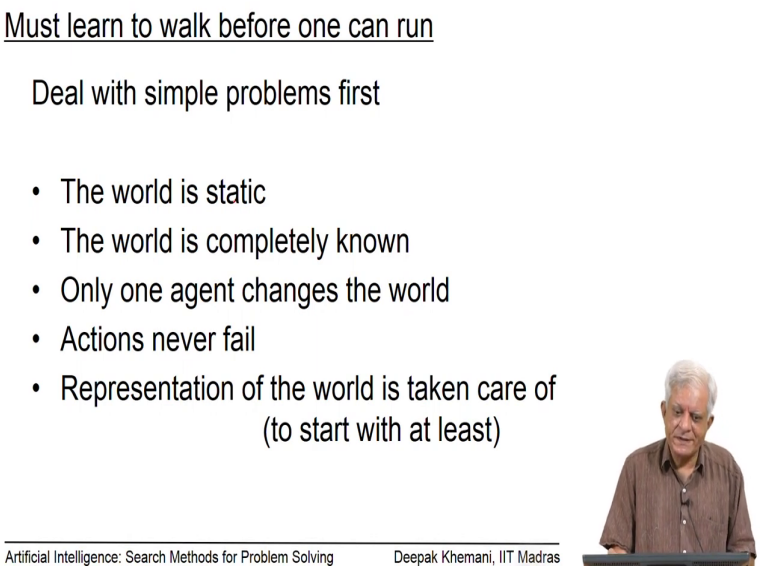

So, the real world is quite complex. So, what we would like to do is to learn to walk before we can run and we will deal with simple problems first. We will assume that the world is static that nobody else is changing the world except your agent ah. It is completely known, the agent knows exactly what is there in the world and the only the agent changes the world.

Actions never fail. So, if the agent says I will take two steps forward and three steps to the right, that is what the robot will do without fail and the representation of the world is taken care. In this course, which is on search methods we will not worry so much about representation, but I do a separate course as I mentioned earlier on knowledge representation in which we focus more on representation.

(Refer Slide Time: 23:22)

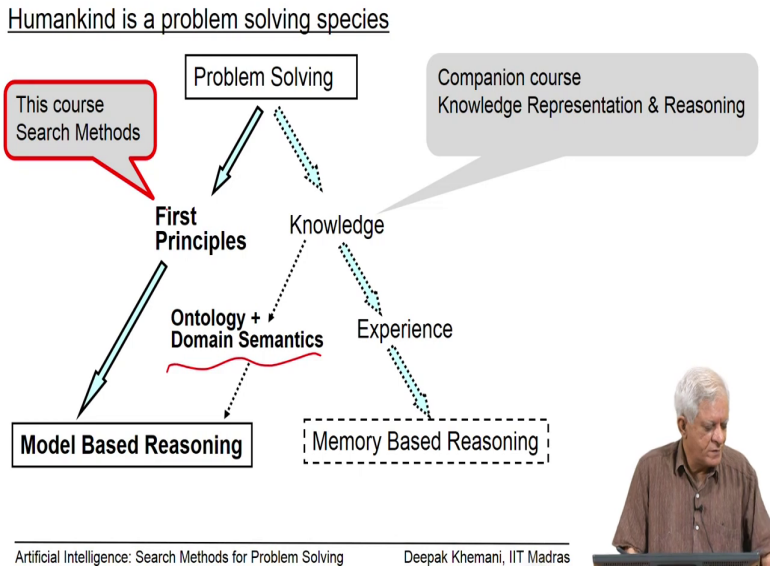

Humankind is a problem-solving species. We approach we take different approaches to problem solving ah. We can highlight two approaches; one is a first principles approach that every time you are faced with the problem by using some kind of first principles, you try to explore what could be the solution and come up with a solution.

Of course, you still have to have certain kind of representation and so certain kind of knowledge modeling of the domain is decide. But the problem-solving approach is by first principles and by first principles we mean search based methods and often this whole field is called model based reasoning.

The other approach of course, human beings learn from experience all the time and you do not want to reinvent the wheel all the time essentially. So, once you have for example, figure out

how to make a cup of tea, then you simply use your knowledge about how you; how you learnt how to make cup of tea.

Whether you learnt it by yourself or whether somebody taught you or whether you read it from some book or from some website it does not matter. Once you know how to solve problem you can use that solving a problem. So, that is an approach of knowledge based reasoning and in particular, memory based reasoning where you kind of exploit your own experiences.

In this course, we will focus on search methods which is first principles methods that we will assume that we are just given the problem and we have to figure out what the solution is. In the companion course which I mentioned which is knowledge representation and reasoning, we will more focus on how to represent knowledge for solving problems essentially.

(Refer Slide Time: 25:06)

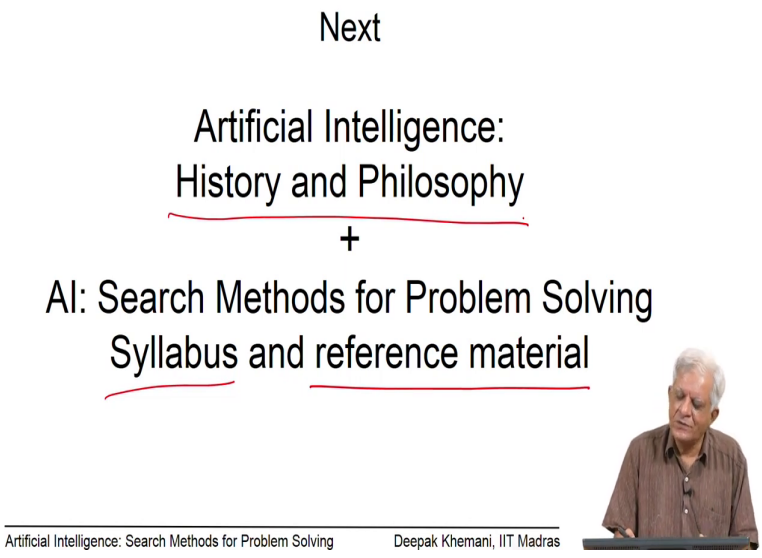

So, that is all for the first lecture in this course. What we will do next is we will look at the history and philosophy of AI. AI has got roots to times much before machines were invented or even computers were invented and we will look at some of this history, it is quite fascinating.

These are lectures which I have recorded about 5-6 years ago and nothing much has changed. So, I am going to reuse those lectures again ah. In those lectures, we will also look at the syllabus and the reference material that we are using and after using those five odd lectures which deal with history and philosophy, we will revert back to the new lectures that I have recorded in the last few months. So, I hope to see you for the rest of this course here.

So, welcome again to this course on AI recorded partly in the times of this corona virus. So, see you next time when we will look at the history and philosophy of AI.