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

Lecture – 03
Introduction (2013 version)

(Refer Slide Time: 00:23)

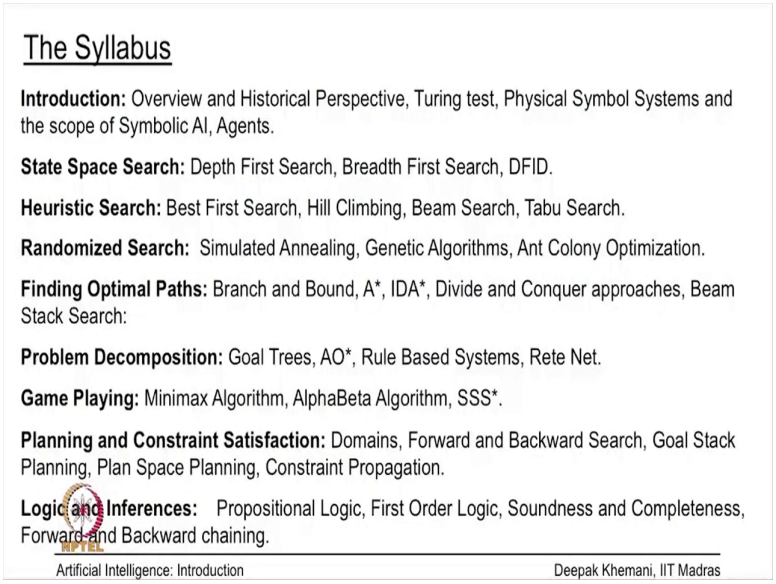

Welcome to this course on AI. As we mentioned let me first start today with the syllabus that we are going to cover and we did a little bit of this in the last class towards the end and I am just repeating this for continuity. So, we will spend the first few weeks not 2 or 3 lectures on the first part of the course which is the history and philosophy of AI.

And we as you will see goes back quite a bit in time and this is qualitatively going to be very different from the rest of the course which is going to be mostly algorithms. And we will start with the simplest algorithms like depth first search, breadth first search and so on.

Move on to heuristic search in which we look at how search can be guided towards the solution that we are trying to find and we look at algorithms like hill climbing, and tabu search and things like that. We will find that even that is not going to be good enough. So, we will try some randomize approaches like simulated annealing, genetic algorithms and ant colony optimization. These are basically optimization techniques, but we will try to see them from the search perspective. Then we will look at very well known algorithm called a star and it is variations which we will see.

Then as I mentioned earlier that we will look at something called goal trees or problem decomposition that if you want to solve a problem and you want to break it up into parts and solve each part separately that technique is called problem decomposition. It led to an area called rule based systems which we will look at briefly. We will also do game playing may perhaps not as late as this maybe somewhere here so that I can give you one assignment to start off with which is to implement a game playing program.

And finally, depending on how much time? We have left we should have something on planning and constraint satisfaction which is kind of a preview of the course that we offer next semester in which we will study this algorithms like alphabeta algorithm, minimax algorithm and a heuristic version called SSS star. And then depending upon how much time we have we will spend some time on these 2 topics planning and constraint satisfaction in which we look at general algorithms for planning.

And we will see by planning essentially we mean finding a sequence of actions which does something useful for you and now we will also look at logic and inferences. Because it is not that we are just solving problems of how to do things, but we also making inferences that if we know something then we know something else essentially. So, that is the process of

making inferences and the language that we use for representation is logic and we will spend some time in that.

So, these 2 topics are actually covered independently and completely in 2 different courses that we offer next semester. One is called planning and constraint satisfaction and the other one is called knowledge representation and reasoning which is not the title we are using here.

(Refer Slide Time: 03:31)

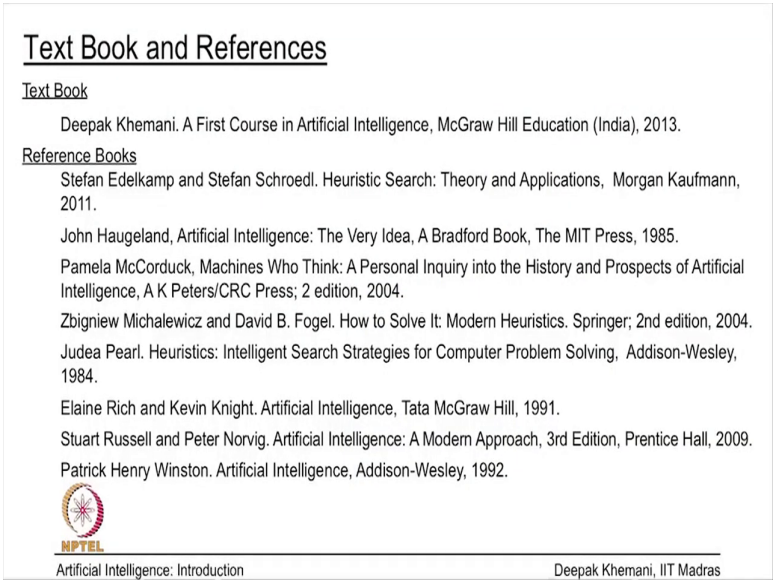

So, the textbook that we will follow is the book which I have just published it is just about come out and. There are some textbooks in AI which have been very popular and earlier I was using a lot of material from here. So, Rich and Knight book on AI Russell and Norvig which is probably the most well known textbook at this point of time and a by book by Winston which was written earlier.

Then there are certain specialized books. So, these 2 books were Fogel and Michalewicz is on certain aspects that we will cover and this book by Judea Pearl is something we will use while game playing essentially. And these 2 books which I will just mention again deal with the history and the philosophy part of AI essentially.

(Refer Slide Time: 04:24)

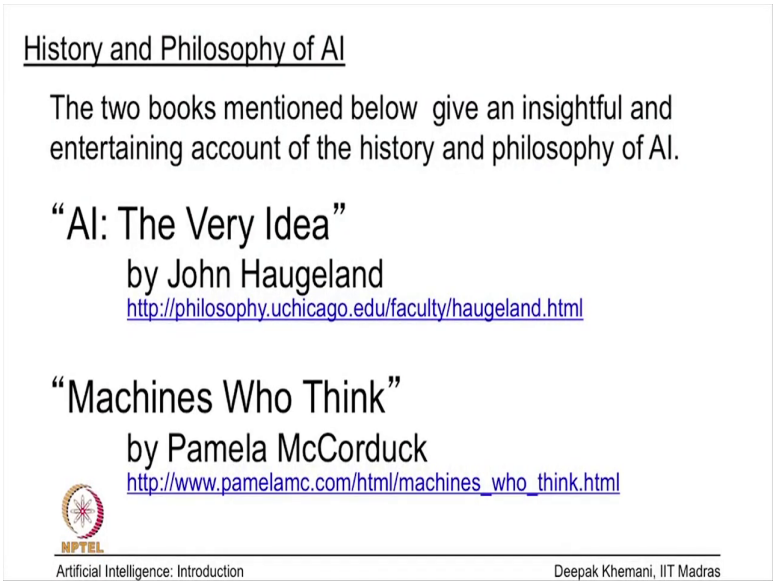

So, these 2 books that I mentioned and this is going to be the subject matter of the first few lectures the historical and the philosophical perspectives to AI. And you can see that it is a topic because we use this word intelligence here and that is something which has concerned people over a lot of time essentially and we want to see what has been the thoughts behind what is AI essentially. So, these 2 books are and I would recommend that you would read at least portions of this.

There is a book called AI The Very Idea and we will discuss shortly why this book is different from the rest. John Haugeland is a philosopher by profession not a computer scientist and he is looking at the philosophical side of things that one of the key questions we will ask and today we will start doing that is can machines think. I want you to start already thinking about this question and today we should discuss some of these basic concepts what is intelligence for example. And Haugeland looks into the philosophy behind this.

Pamela McCorduck is also from the social sciences and she wrote this book quite long time ago actually in 1974 or something like that. And I hope you will notice that the title is if nothing else at least a little provocative because she uses a pronoun who for machines.

So, she says Machines Who Think and who is something that we normally use for people essentially human beings and so on and so forth. So, she is talking about machines who think not machines which think for example. And therefore, already there is a suggestion that her own inclination is to believe that yes it is possible that machines can think.

And these 2 books we will follow and the slides that I have prepared are mostly from these 2 books and a little bit from Wikipedia. So, I will give you all these sources. From the rest of the course I will not use slides very much we will just discuss things on the board essentially. So, I want today's class to be a little bit interactive.

Well, not just today's class, but today's class will be more interactive and I wanted to start thinking about the question of what is intelligence and we will discuss that. But, before we do that let us just look at what are the classical definitions that people have given for this field of artificial intelligence.

(Refer Slide Time: 07:02)

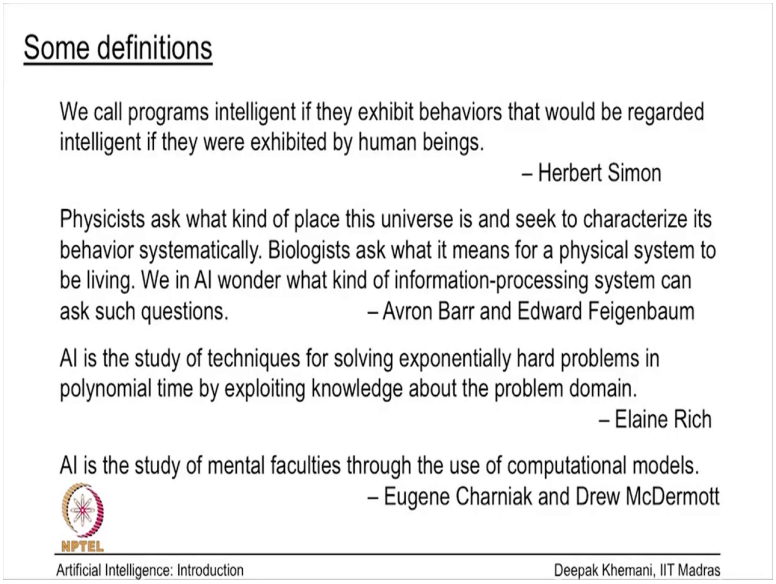

So, let us see first what Herbert Simon has to say. Herbert Simon was one of the founding persons in this area of AI starting in the 50s 1950s. He and his collaborator Allen Newell they founded the school at Carnegie Mellon University and we will see their contribution as we go along. Simon also one of the few people who works in AI who is got a Nobel Prize as you know we do not get Nobel Prize in computer science.

But, Simon got one for economics and he was a multifaceted person. He did many things as people used to be earlier. So, his definition is we call programs we call programs intelligent if they exhibit behaviors that would be regarded intelligent if they were done by human beings. So, this is the most common definition of AI that people use that it is concerned with writing programs or making machines do things which would be considered intelligent by if they were done by human beings especially.

So, one of the first things that AI people got into was things like chess playing essentially because chess playing was always considered to be a hallmark of intelligent behavior essentially. It is only the bright and the intelligent people who could play good chess. There is a long story of chess playing. The first programs were written in 1950s one of the first outline of the game was given by von Neumann. In the 60s, a grandmaster called David Levy I do not know whether I have it in my history, but maybe it will come later.

(Refer Slide Time: 08:59)

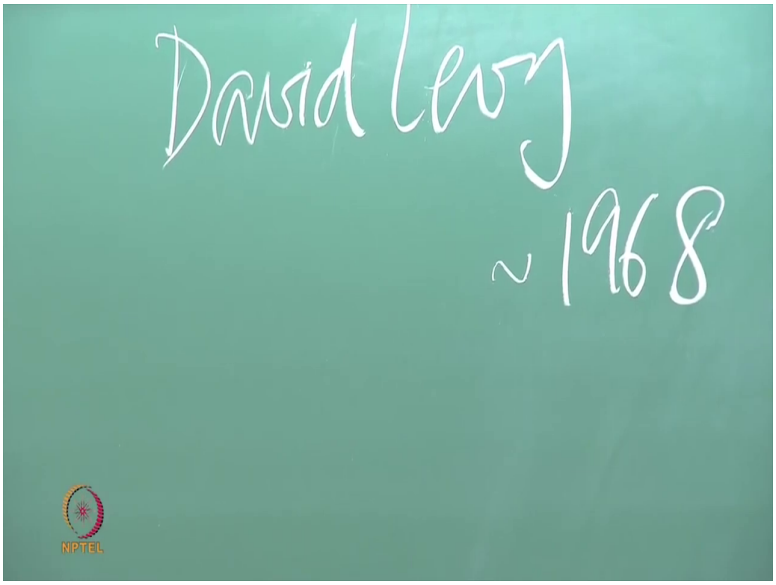

So, let me just write it here. I think around 1968 or so he wagered a bet that chess program cannot beat him for the next 10 years. Because chess was considered to be something which is very intellectual in nature. Well, luckily for him he won his bet which is because it ended in 1978, but many of you would know that in the mid 90s late 90s the then world champion Garry Kasparov was beaten by a chess playing program obviously.

Chess in fact is not so it is not so intellectual in the sense that we tend to talk about you know philosophical sense. Yes, it requires a lot of computing machinery and we will see that if you have a lot of computing machinery you can play good chess.

Let us look at another old definition. This is by Barr and Feigenbaum also 2 old timers in AI. So, his definition say that physicists ask what kind of place this universe is and seek to characterize a behavior systematically. Biologists ask what it means to be a physical system to be living and he says we in AI wonder what kind of information processing system can ask this such questions essentially.

So, in other words he is asking about talking about intelligence that physicists are asking questions about the physical world, biologists are asking questions about you know living creatures. What kind of information processing system could ask such questions. So, essentially he is saying what kind of system would be intelligent in that sense of the world essentially.

Then Elaine Rich as I mentioned one of the popular books in AI. He wrote one in 18s 3 or something or 86 and she gives a computer science flavor to the definition. She says that AI is a study of techniques for solving exponentially hard problems in polynomial time essentially by exploiting knowledge about the problem domain.

Now, of course, those of you who are diehard theory people would immediately object saying that you cannot solve a hard problem in polynomial time because by definition it is a hard problem.

But there are 2 counters to this one is that we may not necessarily be looking for solving them in polynomial time in the worst case. In certain situations like we will see traveling salesman problem is one of the hardest problems that people have encountered. But, given some constraints on the problem of how the edges are connected what are the weights on the edges you can have much faster solutions I think.

The second counter to this objection that you cannot sign that you cannot solve problems in polynomial time is that we are not seeking to find optimal solutions and this is something which many many people have observed that human beings are not optimizers.

We do not necessarily find what solutions the solutions that we consider to be optimal. We are about some people call as satisfiers us satisfies us essentially which says that you are happy with a good solution essentially, it does not have to be optimal essentially.

So, just to an example which sort of strikes me once in a while is living in Chennai that if you are walking along one of the roads in IIT maybe one thing that you want to optimize on the amount of shade that you walk through. But you do not have such dense trees at everywhere there is shade.

So, you have to choose a path essentially and even if one is conscious of the fact that one wants to walk through shade and with one does not mind walking a little bit longer. So, that our objective function is to maximize shade and not worry too much about the length of our path even then we do not go into zigzag path that we would if you were to really follow the shade essentially.

So, we do not optimize in that sense even when we want to be away from the sun we are happy that if the path that we are following has enough lot of shade not necessarily the maximum amount of shade essentially. So, in that sense we do not solve hard problems completely, we do not find the optimal solutions, but we tend to find good solutions essentially and that is what we do all the time.

If you go shopping you do not check in 10 places and then find the minimal cost price and then buy a product. Even though on the web nowadays you can do that sort of a thing, but in general if you think that the price is reasonable we go and buy the stuff essentially right.

And one more definition which is due to Charniak and McDermott who also wrote a very famous book on AI very popular book which I use for part of my section. I do not think I

mentioned it maybe I should add it to the list there. They talk about AI being the study of mental faculties through the use of computational models.

So, we had said earlier that there are 2 approaches to AI one is the cognitive approach which says which I we are trying to understand intelligence and the other is the engineering approach which say that we want to build smart systems or smart apps if you want to say nowadays essentially. So, what this definition says is that we want to study mental faculties and to do that we will be computational models and use them for the studies essentially.

The definition which I like most before I come to that look at these definitions they are saying if a human being does this then it is intelligent and we want to sort of do something similar. So, you want to mimic human intelligence.

(Refer Slide Time: 15:08)

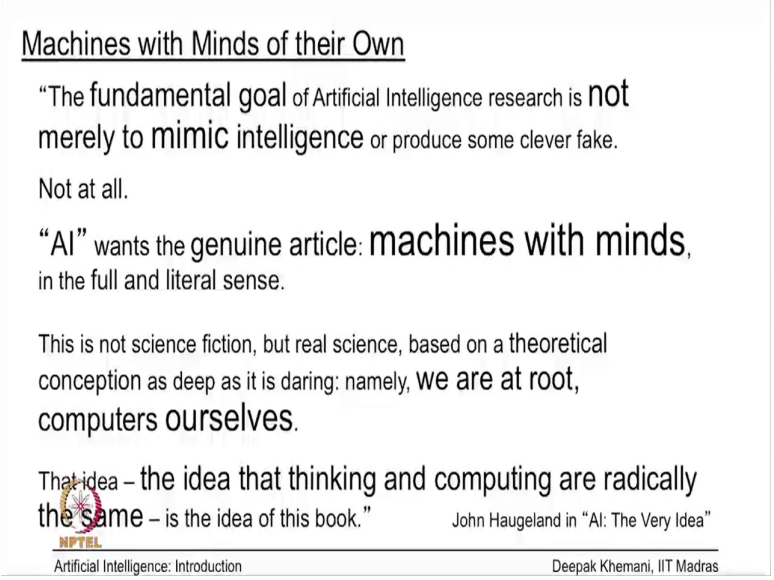

So, the definition which I like most it comes from not a computer scientist, but from a philosopher that we mentioned John Haugeland in the book AI The Very Idea. He says that the fundamental goal of AI is not merely to mimic intelligence or produce some clever fake of intelligence. He says that is not the goal at all AI wants the genuine article, machines with minds of their own in the full and the literal sense.

Now, it is a very interesting question and we will debate it today a little bit in the class as to what we mean by intelligence and can machine have machines have it. And then he goes on to say and all this is in this book here that this is not science fiction, but real science based on a theoretical conception as deep and daring namely that we are at the root computers ourselves essentially.

So, if we are at the root computers ourself which means if you are at the root machine ourselves then to answer the questions can machine think has been trivially solved essentially. Because yes human beings can think and therefore, machines can think essentially. But, the idea that we want to pursue is that the idea that thinking and computing are radically the same.

His idea in his book which is the AI The Very Idea that is very interesting book and for those of you who are philosophically inclined should go and have a look at it. And this idea that thinking and computing are kind of tied up together goes back must before haugeland and we will see either in today's class or in the next class that the British philosopher Thomas Hobbes was one of the first person to put forward this idea.

Hobbes of course, was not a computer scientist in those days there was no computer science. He was a political scientist and this kind of stuff.

(Refer Slide Time: 16:58)

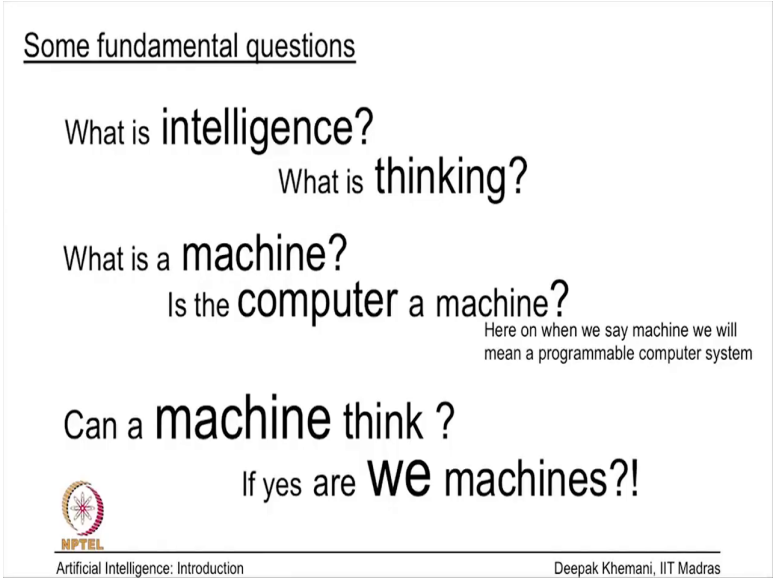

So, let us get to the fundamental questions and this is the part that I want you to give answers to or what do you think about this question. So, I have not written any answers for this I have just written the questions and I will write the answers on the board as and when they come out from the class essentially.

So, the question we want to ask is what is intelligence. I mean if there is going to be ever a debate about whether machines can be intelligent or not machines can think or not first we should be clear as to what do we mean by intelligence. I mean if I write a program is let us say the singular value decomposition of a matrix would that is that program intelligent. Well, I do not know.

So, can I have some responses from the class? What is intelligence? What is let us forget about what is thinking let us say. So, thinking is this thing, but when is when would something be

called intelligent or what is intelligence. What would you require in a system or in agent for you to call it intelligent?

Student: That.

What are the fundamental characteristics of intelligent behavior? Please.

Student: Maybe ability to take decisions.

That is very generic.

(Refer Slide Time: 18:26)

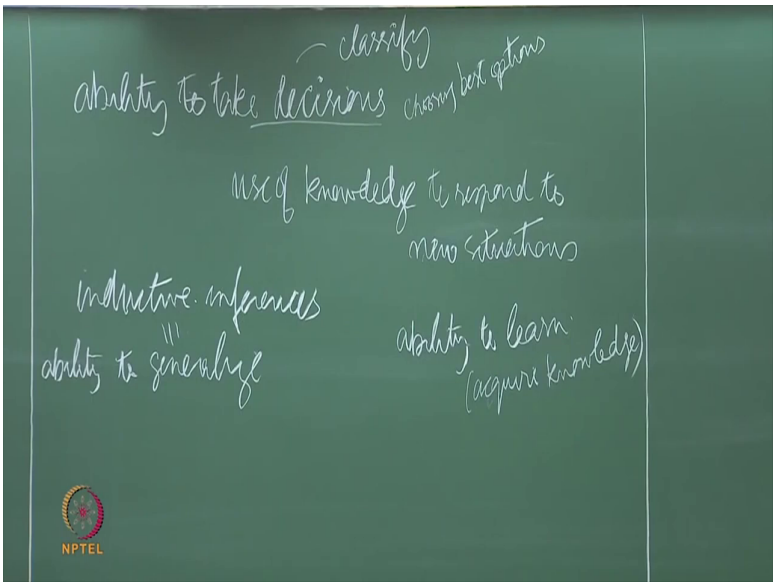

Yes, definitely a part of the intelligence, but maybe if you could expand on that one a little bit. For example, you know you have a small program which says if something then something

else it is also doing taking some decision by looking at some data. Obviously, you are looking at something and taking a decision.

Student: Use of knowledge to respond the new situations.

Use of knowledge. Of course, you will have to tell me what do you mean by knowledge and this definitely has a little bit of inconsistency built into it in the sense that most of the time when we use knowledge or experience exploit experience we use them in situations which are similar which are not entirely new in that sense.

Well, if by new situation you mean a new problem then one has to ask the question what do you mean by that essentially. You know there is this old saying which says that you can never step into the same river twice essentially I think and it is never the same thing.

But of course, nevertheless I will I am not disputing what you are saying I am just trying to get people to respond more. We do as human beings you make extensive use of knowledge and we spend close to what should I say 22 years, 25 years acquiring knowledge which we will later use essentially in our lives essentially.

Human being humans are a very different kind of a species I think. I mean we are the only species which has schools up to 12th standard and then college 4 years after that and then masters and maybe you even a Ph.D in some cases essentially. No other species spend so much time acquiring knowledge essentially.

Student: To be able to make inductive deeps I mean something which does not just follow from your input, but to be able to make some new.

Ok.

Student: Assumptions.

So, I will just use the term inductive inferences or in other words to generalize ability to generalize right.

Student: Ok.

So, you go to the some hotel and you eat the masala dosa and you are happy you come back next day you go there and you have something else let us say uttapam and you come back and then you generalize that this hotel gives you good food or you might say that you know South Indian food is very good.

These kind of inferences that we come to is making inductive inferences. We look at a few instances of something and then from where we generalize that you know it holds for a certain class of things essentially.

I see a few leaves and all of them are green then I conclude that all leaves are green essentially which of course, is not true at least not all the time. Maybe in Chennai yes when we have leaves, but not in the rest of the world.

Student: Basically extending that definition ability to generalize and classify.

Classify would come in this making decisions right.

What else I mean is that all that we do as human beings is that all we lay our claim to for being intelligent.

Student: Choosing the best available option.

Well.

Student: Yes.

Would not that come here choosing best options?

Student: Sir, ability yes sir or the ability to learn.

Ability to learn yes which is a little bit difference from here and we can say by learn we mean acquire knowledge. One can learn from ones own experience. You will you do something which gives you a little bit of a pain maybe you touch a hot stove or something like that 2 3 times and then you learn that is again inductive influence essentially ok, but to learn all kinds of things to learn facts to learn relations between things is something that we do quite effectively.

(Refer Slide Time: 23:23)

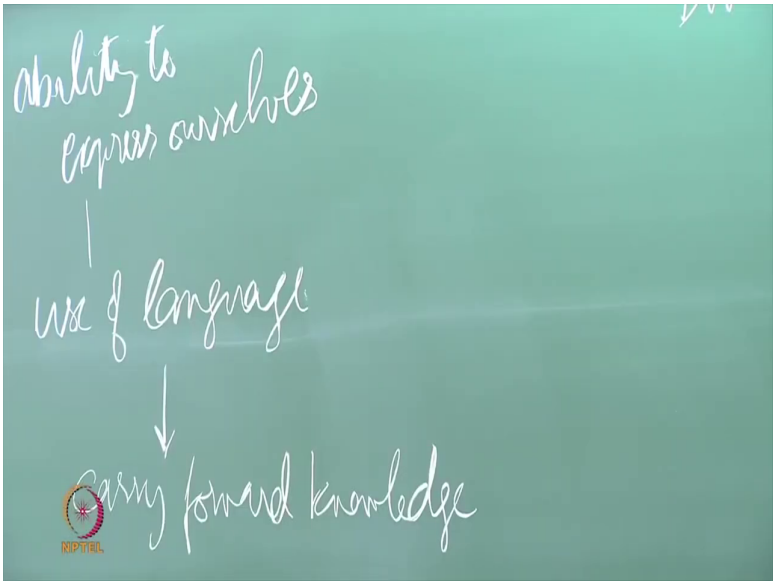

So, what do you mean by this; what do you mean by this?

Student: Communication.

Communication.

Student: Not.

So, but there is a more fundamental thing to I mean expressing well something. Incidentally is something which is a feedback we get from all the companies which come to hire people here. They say that our students are not good at communications essentially, but that is not the idea that is not what you are talking about I think. The very fact that we can communicate something. So, let me go to the fundamental thing what does this rely on something which is specific to the human species.

Student: Speech.

Speech the speech before speech.

Student: (Refer Time: 24:18).

Use of language. Language is something which is unique to at least we think it is unique to our species. Now, there are doubts that you know maybe whales communicate over long distances and dolphins can communicate and that kind of stuff, but we are not quite sure and we do see that there are other creatures which make sounds which are obviously aimed or directed at least towards their own species, but it is not clear to us what they are doing essentially. So, it is a use of language which has enabled us to carry forward knowledge.

So, if you have a brilliant scientist like Newton who is thinking about the universe and the world around him and coming to conclusions and arriving at some understanding of how the world operates. The fruit of his effort is available to us and it is available to us only through

the medium of language essentially because, we can talk to other people because we can write books.

So, printing of course, was another invention which helped this process. But, just simply be to be able to communicate to tell stories this whole idea of folklore you know that stories are passed on from one person to the next like all the stories that we hear in our subcontinent the Ramayan, the Mahabharat and so on.

Overall sort of orally conveyed from generation to generation the and all that is possible entirely through the use of language. It is language which has allowed us to hold on to whatever knowledge we get from our interactions with the world and pass it on to other people essentially. Anything else can one think of ok. So, we will take this as part of thing and then we will see whether machines can be intelligent. So, let me move on a little bit and ask the next question.

This is not a very complicated question. I just want to be sure that we are all on the same page because when we talk of machines thinking and so on. So, what do we mean by a machine?

(Refer Slide Time: 27:03)

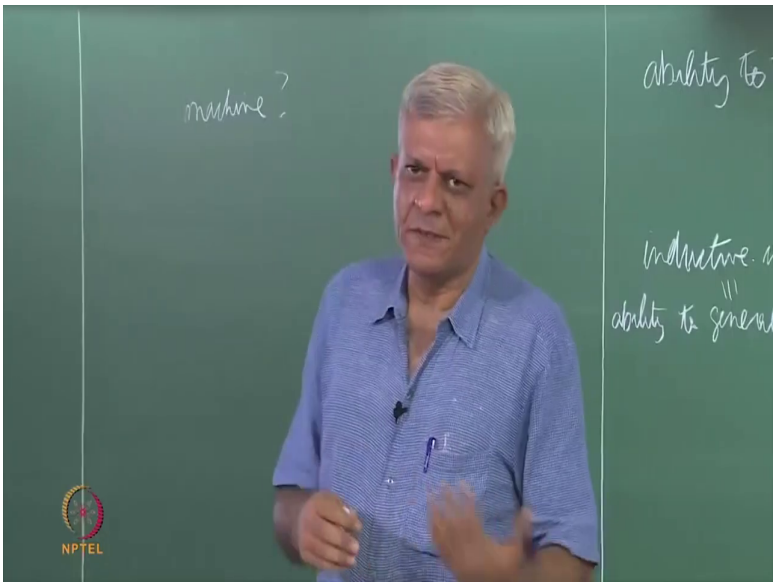

Otherwise we will be stuck with trying to answer the question that can machines think without knowing what we mean by thinking and without knowing what exactly we mean by machines essentially. So, both these terms we should know. So, what do we mean by that essentially?

(Refer Slide Time: 27:19)

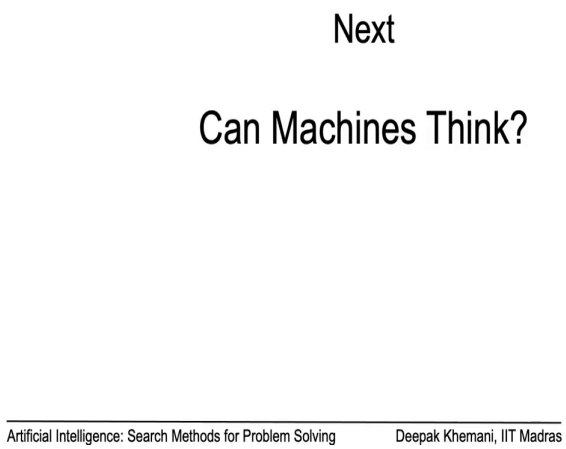

Student: Device may.