Artificial Intelligence
Prof. Sudeshna Sarkar
Department of Computer Science and Engineering
Indian Institute of Technology, Kharagpur
Lecture - 1
Introduction

Today we start the lecture on the course Artificial Intelligence. This course will be delivered by me Sudeshna Sarkar and Professor Anupam Basu both from the Computer Science and engineering department IIT Kharagpur.

(Refer Slide Time: 01:10)

The goals for this course are to introduce you to the field of Artificial Intelligence. We want to explain to you the challenges that are inherent in building a system that can be considered to be intelligent. In this course we will be explaining the key paradigms of Artificial Intelligence, the core techniques and technologies used and algorithms for some of these techniques.

(Refer Slide Time: 01:43)

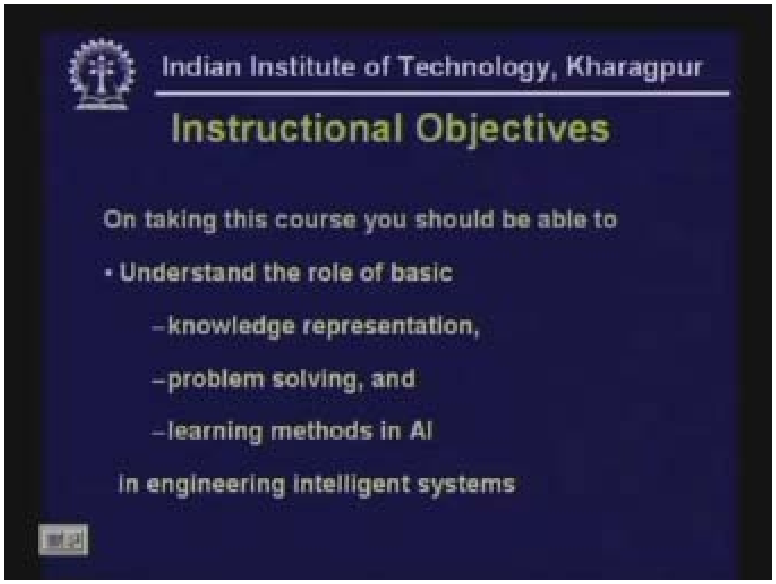

The instructional objectives of this course:
On taking this course you should b able to understand the role of basic knowledge representation, how to represent our knowledge about the world and knowledge of problem solving techniques and a knowledge of some of the learning methods in AI. And we will see how these are used to solve different problems and to build a complete intelligent system.

On taking this course you should be able to assess the applicability the strengths and weaknesses of these methods of the different techniques that we discuss. We will discuss the strengths of this method and situations where these methods can be applied to solve different types of problems that require intelligence. You will learn how to develop intelligent systems by assembling solutions to concrete computational problems.

(Refer Slide Time: 02:56)

The way we will do it is, we will look at the different components of intelligence and for each of these we will discuss ways of solving these problems. And then, depending on the functionality of the system that you wish to construct or engineer you can put together some of these solutions to get the full system. After you have taken the scores you should be able to appreciate the role of problem solving, the role of natural language processing, the role of computer vision etc in understanding human intelligence from a computational point of view.

Some more points on objectives of the course:
On taking this course you should be able to formulate certain types of problems as state space search problems and you should learn the efficient methods to solve them depending upon the characteristics of the problem space, you should be able to write programs that play games particularly two player games, you should be able to use learning to find patterns in data to find rules from data, you should be able to build expert systems for different diagnostic and other purposes.

(Refer Slide Time: 04:38)

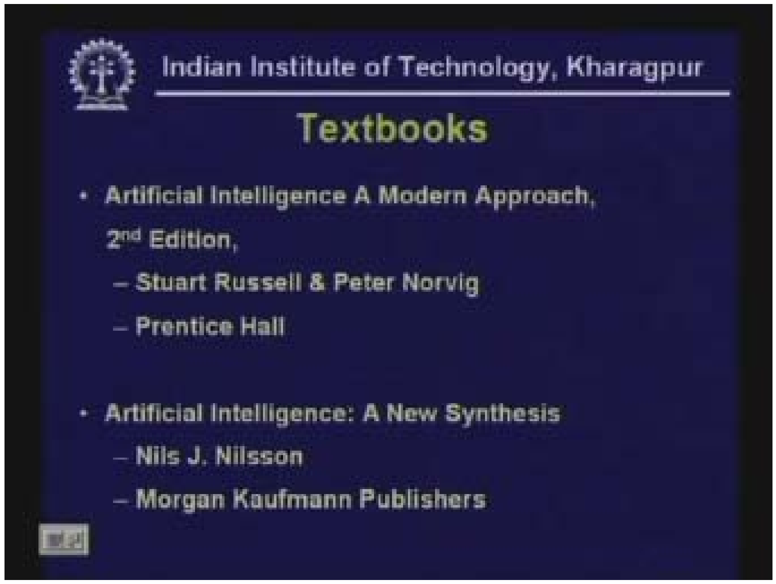

Some of the text books we will follow for this course are the two books which I will be referring to and professor Basu will also talk about the other books he will refer to. The books are Artificial Intelligence a modern approach second edition by Stuart Russell and Peter Norwich. This book is published by Prentice Hall and also by Pearson. The second book is Artificial Intelligence a new synthesis by Nilsson published by Morgan Kaufmann publishers. Today's lecture will be the first lecture in the series. The first module for this course is the introduction and today's lecture is the first lecture this module which is introduction to AI. Now let us come to the objectives of today's lecture.

(Refer Slide Time: 05:50)

The instructional objectives of today's lecture are:
To understand the definition of Artificial Intelligence, what Artificial Intelligence is? What is it about?
Secondly we will be discussing the different faculties involved with intelligent behavior, the different components that define intelligence. We will also be examining the different ways of approaching AI and finally we will also look at some example systems that have been constructed, that are popularly known which use AI and lastly we will also take a brief look at the history of AI. On taking this lesson you should become familiar with the different ways of defining Artificial Intelligence. As we will see different people may define AI differently and we will familiarize ourselves with these definitions.

Secondly, as I mentioned, we will try to understand the different components of intelligent behavior.

(Refer Slide Time: 07:06)

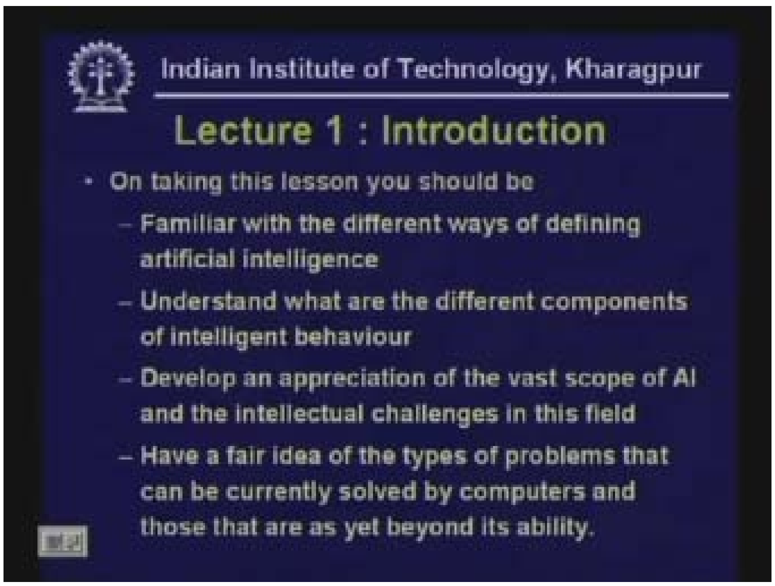

Another objective of today's lecture is also to let you develop an appreciation of the vast scope of Artificial Intelligence and the intellectual challenges that are there in the field. On talking today's course you should be able to have a fair idea of the types of problems that can be currently solved by the computers and today's techniques that we know. We will also have an idea of those problems that is still difficult or we cannot yet solve it by the techniques that we know today.

(Refer Slide Time: 07:54)

These are four main components of today's lecture. Definition of AI, example systems, approaches to AI and the brief history. First we will take up the definition of AI.

Now, what is Artificial Intelligence?
There are too many definitions of this term floating around.

(Refer Slide Time: 08:17)

As you see or what is clear from the two words Artificial Intelligence it is clear to see that AI is concerned with the design of intelligence. And in the first term AI is actually concerned with design of intelligence in artificial artifacts and artificial devices. Thus, artificial systems or man made systems are building intelligence into them. This term was coined by McCarthy in 1956 in a famous conference the Dartmouth conference.

(Refer Slide Time: 09:25)

Now, the term artificial is easy to understand but what is intelligence?
It is very difficult to define intelligence. Often we look at some people look at intelligence as something that characterizes humans. If you take human beings to be intelligent you can say Artificial Intelligence means having behavior which is like a human. In fact there are two schools of thought here. Here an idea is to have a machine or have a system that behaves like a human. Humans are not always completely intelligent even though humans are very good. Actually pretty intelligent but all the time humans do not behave intelligently.

So the other school of thought is that Artificial Intelligence concerns with intelligence which is the ideal or the best behavior or the most rational behavior. It is the machine that should behave in the best possible manner. There is another dichotomy in the definition.

When we talk about behavior what sort of behavior are we talking about?
There are two main types of behavior that people will like to talk about. Number one is thinking, thinking intelligently, reasoning properly and intelligently in order to come up with a solution. And the second approach is to talk about not thinking but acting how the system actually acts or behaves. We can talk about intelligence as something which characterizes humans or something that means behavior in the best possible manner or behaving rationally.

Again we can talk about intelligence in thought or intelligence in action. So, based on this criterion we can look at the different ways of defining AI. So we may look at thought processing or reasoning versus behavior, we may look at human like performance versus ideal rational performance. And this diagram shows the four different definitions that emerge from these two dichotomous.

(Refer Slide Time: 12:10)

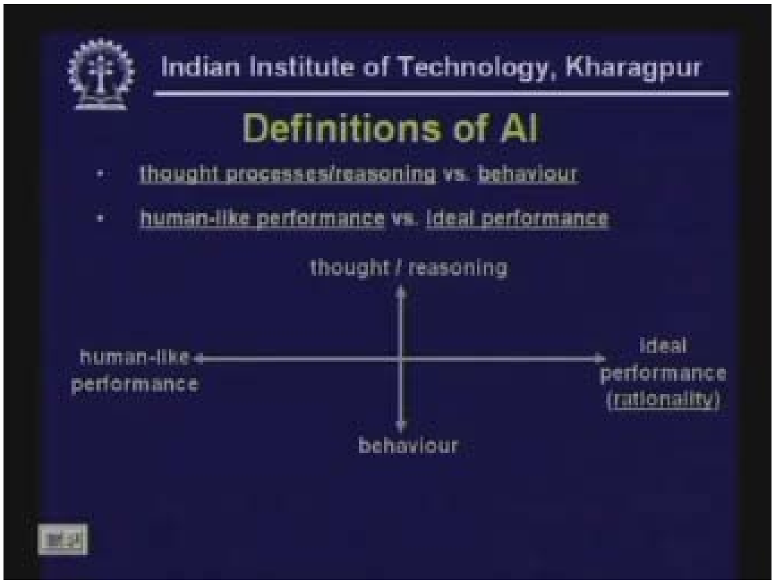

(Refer Slide Time: 12:12)

On the one hand we have thought or reasoning versus behavior and on the other hand we have human like performance versus ideal performance. So there are systems that think like humans. For example, we will discuss the famous Turing test which was devised by Alan Turing where the system which passed the Turing test would be a system that behaves like a human or thinks like a human. The second definition is systems that think rationally.

The school of thought were different philosophers, mathematicians and computer scientists who have worked on logic and laws of thought believe in this approach. Thirdly, there are systems that act like humans. Cognitive scientists look at the properties of systems that act like humans and finally we have the definition systems that act rationally or systems that act in the best possible manner. And for this we have the approach of constructing a rational agent an agent which acts rationally. Alan Turing considered by many to be the father of AI devised the Turing test.

In the Turing test this is the experimental set up that is devised. There will be a closed room and in this closed room there will be a being which may be a computer and it may be a human. There is an interrogator outside the room. The interrogator does not know whether the being inside the room is a computer or a human. So what the interrogator does is that the interrogator asks questions and the being inside the room processes these questions and returns some answer and the interrogator on the left room receives the answers on the screen.

Now the interrogator has to make out from the answers whether the being inside the room is a computer or human. Now, if there is a computer inside the room the computer tries to convince the interrogator that it is actually a human being in the way it answers to the questions and it is the task of the interrogator to decide who is human.

(Refer Slide Time: 15:30)

This is a schematic diagram of the Turing test, this is the interrogator sitting in front of the terminal, this is a wall room. The wall room may contain either a human or a computer and the interrogator has to decide whether what is inside is a computer or a human being.

(Refer Slide Time: 15:51)

(Refer Slide Time: 15:59)

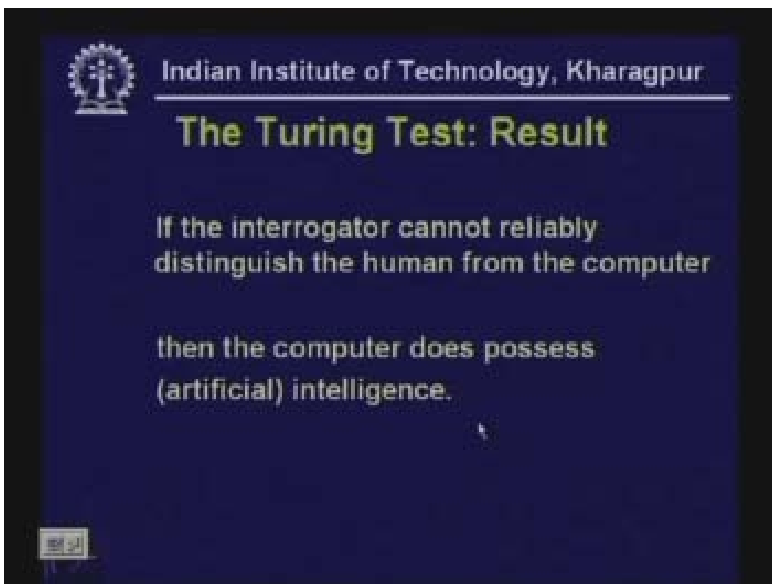

Now, if the interrogator cannot reliably distinguish between a human answerer and a computer answerer then we can say that the computer system possesses Artificial Intelligence. This is the test devised by Turing to find out whether the machine has been able to come up with a right amount of intelligence to match human intelligence in answering questions. Now let us look at typical AI problems.

Intelligent entities or agents need to be able to do different types of tasks. There are some tasks which are mundane tasks that we do as a matter of fact in out daily life and there are some tasks that we consider intelligent like solving difficult mathematical problems, playing games of chess in an expert fashion and other activities which intelligent people can do well. Now, examples of mundane tasks are planning route. Suppose you want to go to here from the market and you plan a path along which you will go. Or you want to go from here to let us a say a particular place in Delhi and you have to plan your journey and plan your path. Something that we do all the time is trying to recognize objects or recognize faces of people, this requires vision.

Thirdly, we communicate with each other through natural language. Fourthly, we navigate around obstacles on the street. So these are the tasks that we do routinely. In fact most animals do this task routinely. And then there are expert tasks like medical diagnosis which are only the doctor or the expert in the field does. And mathematical problem solving can be done effectively only by good mathematicians.

(Refer Slide Time: 18:31)

Now which of these problems are easy for the computer to do and which of these problems are hard. Surprisingly it has been much easier to mechanize many of the high level tasks which are so called expert tasks which has been easier in the history of AI and the history of computers. It has been easier to solve problems which are really the domain of experts but AI has not had the same amount of success in dealing with mundane tasks.

For example, AI systems can easily do symbolic integration. Some of the systems can prove some theorems. AI systems can play chess quite well. There are systems " " diagnosis in particular domains. However, there are certain things that humans and animals do quite effortlessly.

For example, walking around without running into things catching prey and avoiding predators, interpreting complex sensory information, modeling the internal states of other animals trying to understand what they are thinking about us and how to plan what to say and so on and also working as a team or collaborating.

(Refer Slide Time: 20:03)

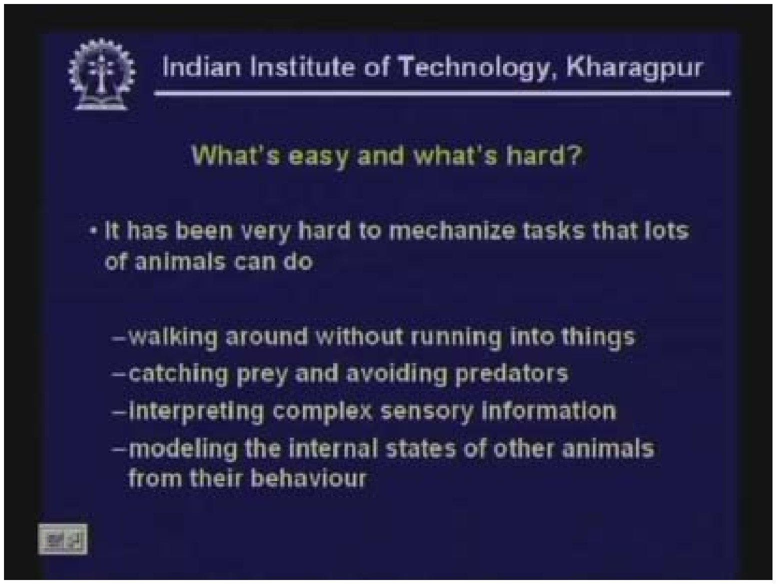

Then these tasks unfortunately have not all been easy to do by machines. Let us look at some of the basic intelligent behavior in human beings. Perception that is the ability to see, hear sensory information.

Reasoning: Reasoning with the information that we have.

Learning: Learning for new situations, understanding natural language, communicating in natural language, solving problems.

(Refer Slide Time: 20:50)

Hence these things namely perception, reasoning, learning, language, understanding and solving problems are examples of some of the things that we want our AI systems to solve. Having looked at the definition of AI let us have a look at some examples of AI systems that have been around.

(Refer Slide Time: 21:26)

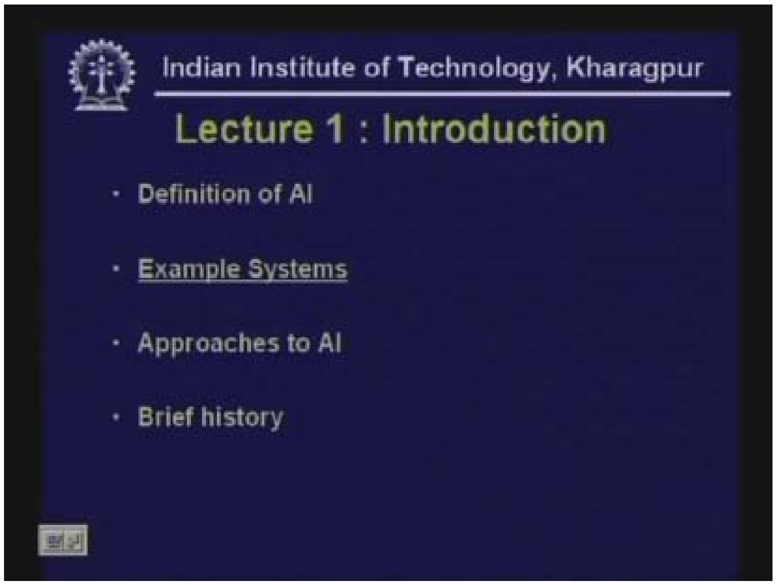

These are some of the applications of AI:
Computer vision, image recognition including face recognition, robotics, natural language processing and natural language understanding, speech processing, etc. Then if you look the practical impact of AI the AI components are embedded in numerous devices. Even in some copy machines there are AI components embedded. AI systems are in everyday use in detecting credit card fraud, in configuring products, in complex planning tasks, in advising physicians. Then intelligent tutoring systems provide students with personalized attention. These systems are there being used and they have a tremendous impact because they are so useful.

(Refer Slide Time: 22:22)

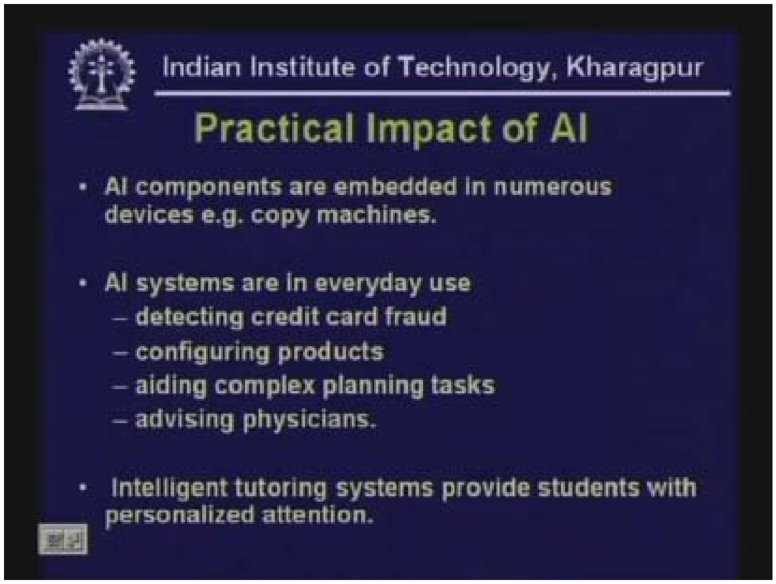

This is a system ALVINN which stands for Autonomous Land Vehicle in a Neural Network. It was designed in 1989 by Dean Pomerleau at Carnegie Mellon University. This system drove a car from the east coast to the west coast across United States of America using computer control. And it drove completely autonomously for most of the 2850 miles. Only for 50 miles especially at exits to freeways etc the human driver took charge. For 2800 miles the car drove it. And the idea behind the car is quite simple. In front of the car is a camera which takes a picture of the road in front.

(Refer Slide Time: 23:55)

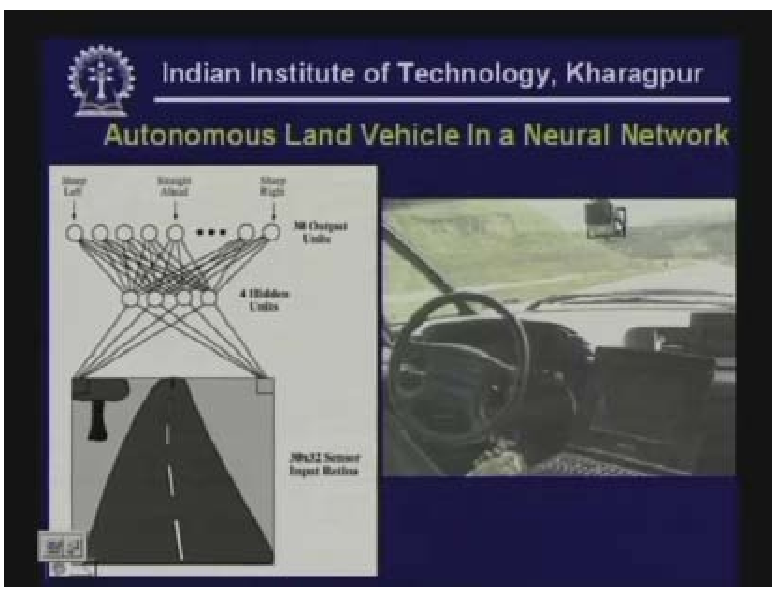

And this picture or this image is used in a neural network. This picture is captured into an image having 30/32 pixels. These pixels are fed into a neural network four hidden units and the output tells the processor which way to turn the wheel and decide the speed and so on. In 1997 the deep blue chess program developed at IBM be it the current world chess champion Gary Kasparov. This is the computer deep blue and this is Gary Kasparov after he lost the match accidentally.

(Refer Slide Time: 25:00)

In a machine translation if we could have immediate translations between people speaking different languages that would be a remarkable feat and it has very wide ranging economic and cultural implications. In the world today there are people speaking so many different languages and we do not understand the languages of many other people. Even in India as you know there are so many languages, there are more than 20 official languages and I cannot understand the language of each Indian. So would it not it be nice if we had a system which would do simultaneous machine translations so that we can effortlessly understand each other.

Full machine translation is not yet there but there has been quite some progress in the field of machine translation in a small way. For example, the US military is giving a simpler one way translation device they are using this in Iraq. US forces are using the Phraselator to communicate with injured Iraqi prisoners of war travelers at checkpoints and for other peace keeping duties. Carnegie Mellon University is working on a system called the speechlator for use in doctor patient interview.

(Refer Slide Time: 26:29)

Imagine how difficult it is when a doctor does not understand the language of the patient. And when the patient does not understand the language of the doctor the patient will not be able to communicate his symptoms to the doctor. So speechlator is used in order to help doctors do so. In space exploration robotic space probes autonomously monitor their surroundings, make decisions and act to achieve their goals. This is the homepage of Mars Exploration Rover Mission.

(Refer Slide Time: 27:35)

If you have a look at this page hosted by jet propulsion laboratory this page brings us live the explorations that are being carried on by the Mars Rover. There are two Mars Rovers spirit and opportunity that have been sent to Mars. They have already finished their primary assignments and are continuing with exploratory duties. These two pages contain updates of spirit and this page contains the update of opportunity.

(Refer Slide Time: 28:32)

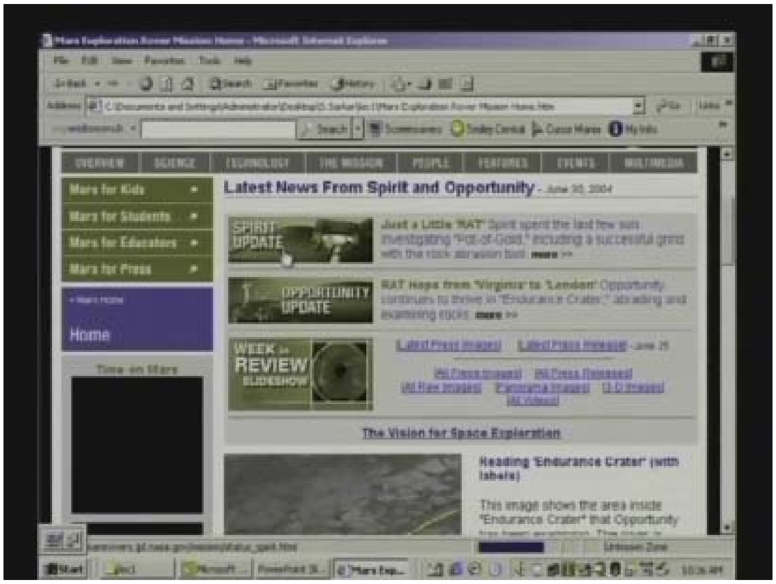

For example, let me read for you an excerpt from the spirit update. Just a little RAT, Spirit spent the last few salts investigating pot of gold including a successful grind with the Rock Abrasion Tool that is what RAT is, a RAT is a Rock Abrasion Tool. So, what spirit is doing is using a rock abrasion tool and getting samples of rocks from the surface of Mars and it is trying to find out what chemicals are present in the rock. So, one of the objectives of this mission is to find out whether there is water in Mars.

In fact these Mars Rovers have been able to trace the presence of water from the rock samples in Mars. Then opportunity is going from Virginia to London. These are different locations defined on the mars' surface and opportunity is currently in a crater called the endurance crater and it is abrading and examining rocks. This image shows the area inside endurance crater that opportunity has been examining. The Rover is investigating the distinct layers of rock that make up this region. And this image taken by Rover highlights the nodular nuggets that cover the rock that has been named the Pot of Gold. These nuggets appear to stand on the end of stalk like features.

The surface of the rock is dotted with fine scale pits. And there are so many other news about these two Space rovers. The Spirit rover is currently exploring a range of marsh and hills that took two months to reach. It is finding curiously eroding rocks that may be new pieces to the puzzle of the region's past. Spirit's twin opportunity is also negotiating sloped ground examining exposed rock layers inside a crater informally named endurance.

(Refer Slide Time: 31:06)

We have intelligent agents that are going to unknown territory where no human has been before and they are carrying on explorations and making new inferences.

(Refer Slide Time: 31:29)

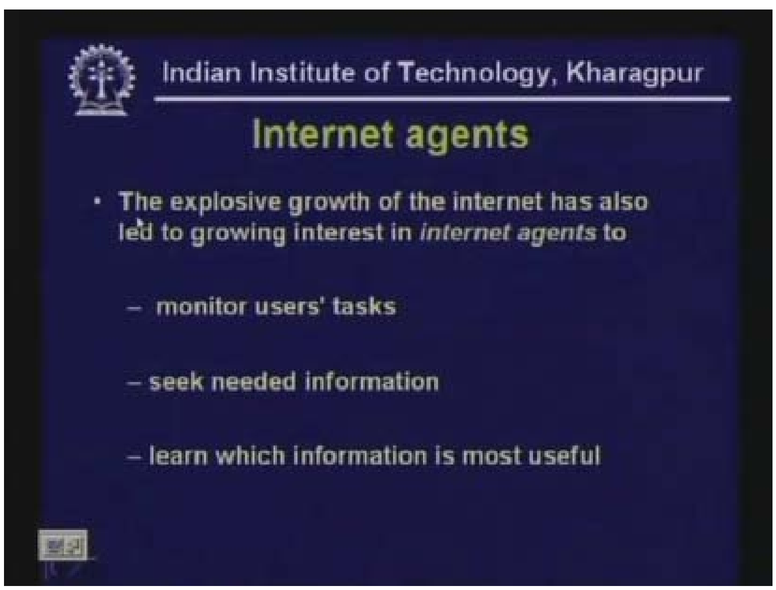

Then there are internet agents. All of you are familiar with the explosive growth of the internet in recent years and there is a growing interest in internet agents that can monitor users' tasks, seek information that is needed from the web and learn which information is most useful for a particular user. Now that we have looked at different examples of systems that use Artificial Intelligence we will briefly look at some approaches to AI and some approaches to solving AI tasks. One way of looking at AI is strong AI or weak AI.

Strong AI aims to build machines that can truly reason and solve problems. Strong AI are machines that are self aware and whose overall intellectual ability is indistinguishable from that of a human being.

(Refer Slide Time: 32:54)

So strong AI proponents want to develop systems that are completely intelligent and that can do things fully using their own intelligence. Such systems can be human like can be non-human-like but rational. When AI was conceived in the 1950s and 1960s there was a huge optimism about AI and there was a prediction that very soon AI systems will be able to overtake humans and able to everything that a human can do and can do them much better and do tasks that humans cannot do within a short time.

However, such optimism has been ill founded and this was partly the reason why some people lost faith in the techniques of AI. But now after research into AI has taken place for over 50 years now we are in a position to understand and appreciate the true difficulty of the different problems that AI face. And we known what we can aim to solve now and what is more difficult and we will need different techniques, different hardware and different paradigms to be able to solve.

Weak AI unlike strong AI deals with the creation of some Artificial Intelligence that cannot truly reason and solve problems but act as if it were intelligent. So the proponents of weak AI claim that machines which have been suitably programmed can simulate human cognition, appear to behave intelligently, appear to do tasks well and intelligently without really having the same intelligence or understanding as humans possess. Therefore, strong AI really deals with machines that really have mental states that think, reason, understand their behavior whereas weak AI is involved in simulating human behavior or simulating intelligent behavior without really claiming that the reasoning process behind it is intelligent.

(Refer Slide Time: 35:45)

The goal of applied AI is to produce viable smart systems. For example, it will be nice to have a security system that is able to recognize the faces of people who are permitted to enter a particular building. There are certain applications which are useful to us and applied AI aims to solve these applications intelligently, not necessarily to construct a complete intelligent agent but an agent which is intelligent in doing a specific task.

For example, recognize people, detect credit card fraud, drive a vehicle autonomously. So they take up specific tasks and develop systems that solve those tasks. Fourthly, cognitive AI deals with the studies where computers are used to test theories about how the human mind works. Cognitive scientists want to understand how humans act, how humans behave, how humans think and these theories can be tested by building these theories into machines and watching and testing how well the machines function using those theories.

(Refer Slide Time: 37:49)

For example, one may have a theory about how humans recognize faces. We do not know how we recognize faces, how our brain recognizes faces, how we store all the different faces or some of the many different faces that we have seen in our lifetime and how we look at a person and recognize them. So, cognitive scientists have come up with different theories about how people recognize faces or how people solve different types of problems. And some of these theories can be tested by building similar mechanisms which are machine and testing how well the machines perform. Here I have outlined some of the topics of AI.

(Refer Slide Time: 38:37)

In the core areas we talked about knowledge representation, reasoning, machine learning. General algorithms: search, planning, constraint, satisfaction.

Perception: vision, natural language processing, robotics.
Applications: game playing, AI and education, distributed agents.
Uncertainty: probabilistic approaches, decision theory, reasoning with symbolic data. These are some of the topic that people study in AI and in this course also we are going to study most of these topics. Today successful AI systems operate in well defined specific domains employing narrow or specialized knowledge.

(Refer Slide Time: 39:55)

However, if you want to artifact a system that has general intelligence that can work intelligently in any domain we need to have a lot of things. For example, such a system must have common sense knowledge which is needed to function in open ended worlds. We use such a huge amount of common sense knowledge or background knowledge to do our tasks well. If we really start thinking and try to note them down then it is a huge effort. There is an effort at Stanford University by Tog Lenat Guha and others called the psyche project whose objective is to document all common sense knowledge so that one can have a system that can use all these common sense knowledge for their reasoning.

Secondly, a general unconstrained AI system must be able to understand natural language, in fact unconstrained natural language. Though there has been a lot of stride in natural language understanding, then understanding unconstrained natural language in general is a very difficult problem which will require a lot of expertise to solve completely.

What can today's AI systems do?

We have systems that can recognize faces, we have almost autonomous vehicles, our natural language processing systems can do simple machine translation. Our expert systems can do medical diagnosis in a narrow domain.

(Refer Slide Time: 42:01)

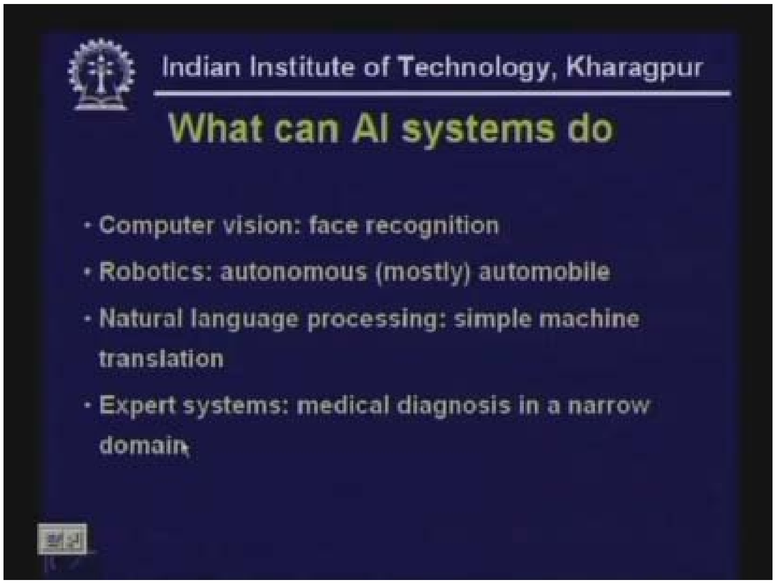

Our spoken language systems are capable of thousand word continuous speech. Planning and scheduling systems are used in Hubble telescope experiments. Hubble telescope is one of the most well known telescopes which have been around for several years. Now there is a talk of dismantling the Hubble telescope because it has become quite old and the cost of maintaining it has become huge. But it is for a long time the Hubble telescope has been the most important telescope for gathering a lot of data and there are so many people who want to use the Hubble telescope. There is a complex planning and scheduling problem to schedule these tasks on the telescope which has been done by AI systems.

In learning our text categorization systems can work and categorize the text at about thousand topics. In games AI system has achieved grand master level in chess where the noise…world champions we have good programs playing checkers.

(Refer Slide Time: 43:22)

But there are many limitations to what AI cannot do yet. AI systems currently cannot understand natural language robustly. AI systems cannot surf the web yet or interpret an arbitrary visual scene. We have seen that they can recognize facial images or work in a narrow domain of recognition. AI systems cannot fully learn a natural language. They cannot construct plans in all sorts of dynamic real time domains in general. And AI systems do not yet exhibit true autonomy and intelligence.

(Refer Slide Time: 44:05)

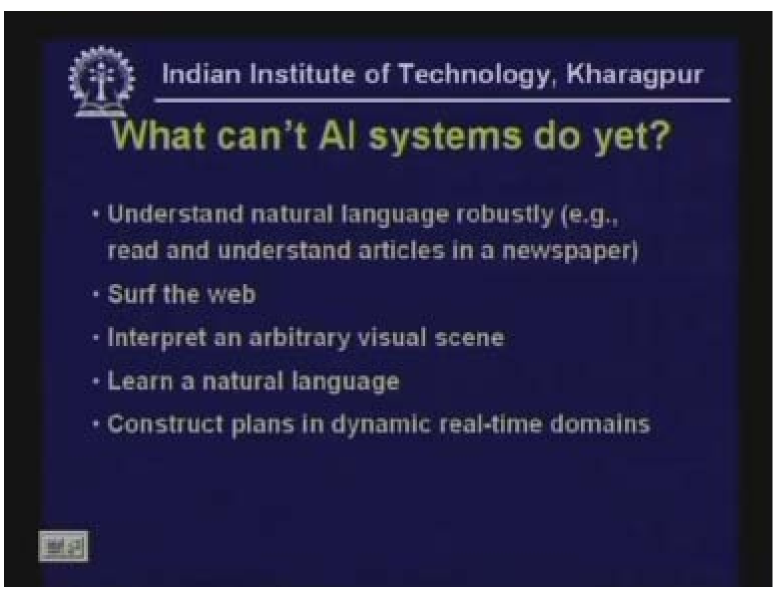

Now that we have looked at some of the approaches of AI and what AI can do and not do at present let us have a look at the brief history of Artificial Intelligence. The dream of making a computer imitate us began many centuries ago. Intellectual roots of AI stretch back thousands of years into the earliest studies of nature of knowledge and the nature of reasoning.

(Refer Slide Time: 44:58)

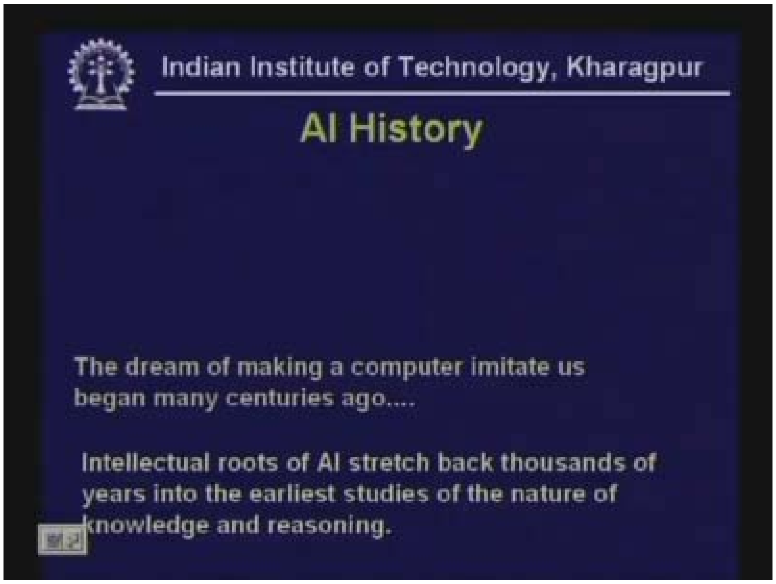

The concept of intelligent machines is found in Greek mythology. In 8th century Pygmalion is credited to have asked the goddess and obtained an ivory statue of a woman built after the fashion that he liked.

(Refer Slide Time: 45:22)

Hephaestus created a huge robot Talos to guard Crete. So this robot used to go around the island of Crete hurling stones at invaders and to detract invaders and if found an opponent it would squeeze him to death. Artificial Intelligence draws from many areas from philosophy, from mathematics, from economics, biology, and psychology and from computer engineering and also from linguistics.

(Refer Slide Time: 46:14)

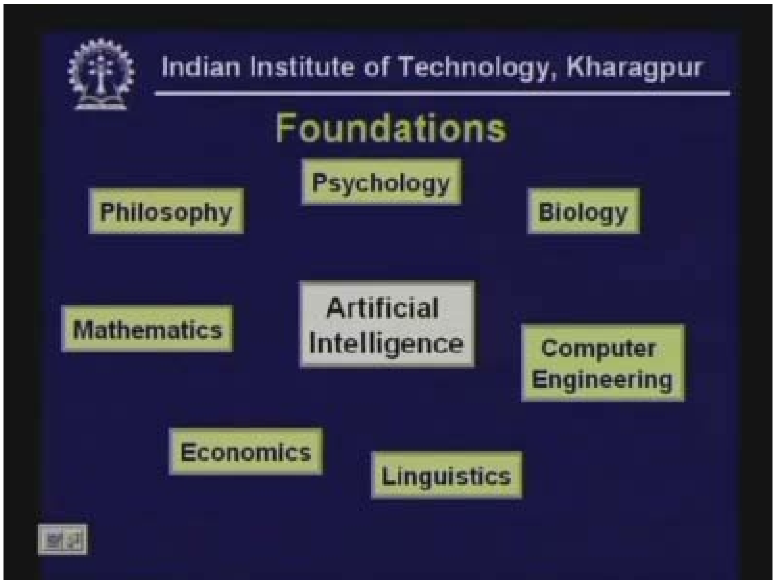

Philosophers have analyzed the nature of knowledge and have explored formal frameworks for developing conclusions. There have been mathematical formalizations in logic, in computation and probability. Economists have developed decision theory and biologists have reasoned about how the brain processes information. Psychologists have long studied human cognition and they require knowledge about the nature of human intelligence.

(Refer Slide Time: 47:02)

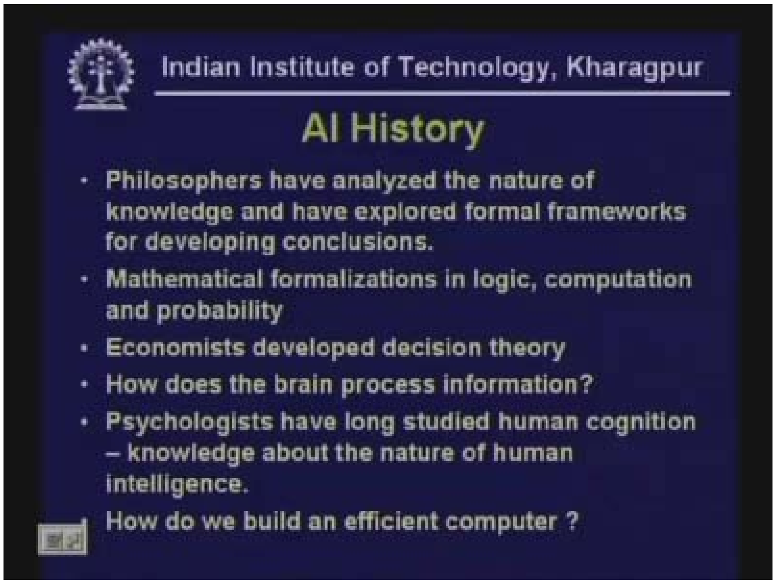

And finally we want to know how to build an efficient computer. So, in the ancient days Aristotle in the 4^{th} century B.C. developed an informal system of logic which was the first formal deductive reasoning system.

(Refer Slide Time: 47:26)

In the 13^{th} century we have Ramon Lull a Spanish theologian who invented the idea of a machine that would produce all knowledge by putting together words at random. He even tried to build such a machine as concept wheel.

(Refer Slide Time: 47:43)

Then early in the 17th century Descartes proposed that bodies of animals are nothing more than complex machines.

(Refer Slide Time: 47:54)

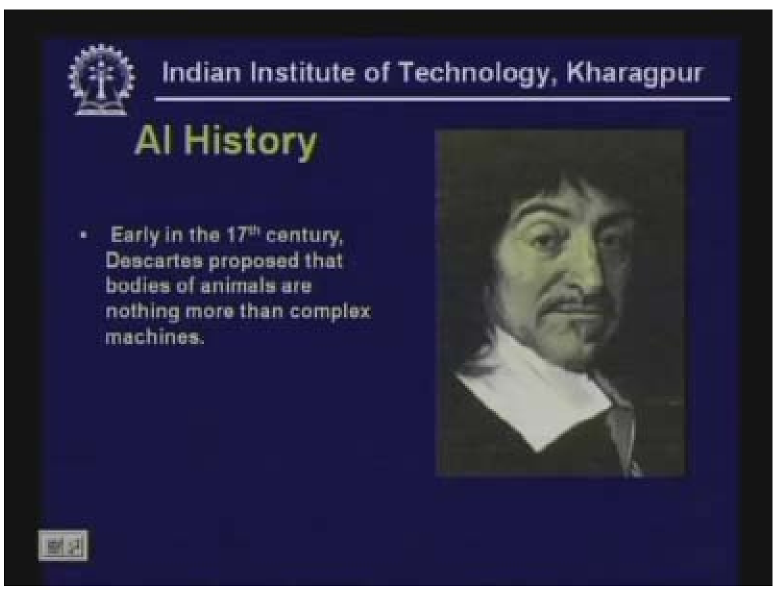

Blaise Pascal in 1642 built the first mechanical digital calculating machine. Leibniz in 1673 improved Pascal's machine.

(Refer Slide Time: 48:05)

So that was the first step in building a mechanical computing device. In 19th century George Boole developed a binary algebra representation which laid the foundation of Boolean algebra. Charles Babbage and Lady Ada Byron worked on programmable mechanical calculating machines.

(Refer Slide Time: 48:37)

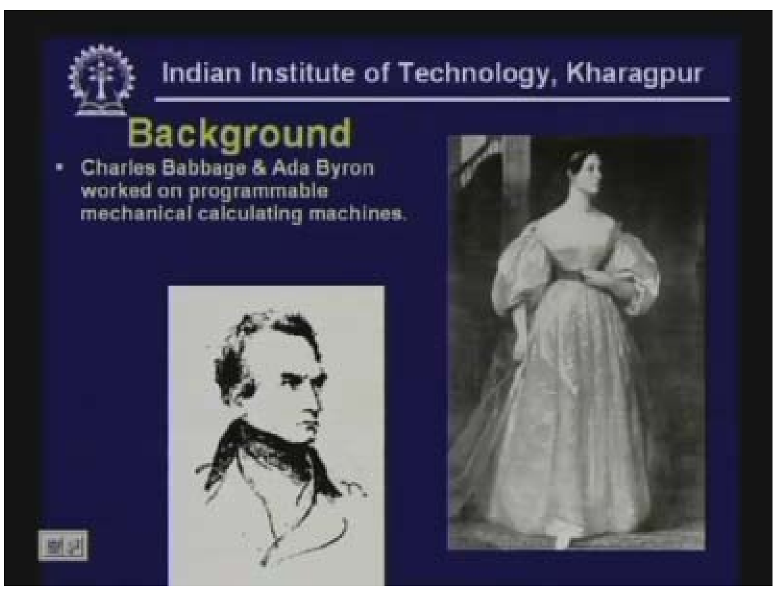

In the late 19th century and the early 20th century mathematical philosophers like Gottlob Frege, Bertram Russell, Alfred Whitehead and Kurt Gödel built on Boole's initial logic concepts to develop mathematical representations of logic problems.

(Refer Slide Time: 49:00)

The advent of electronic computers really provided a revolutionary advance in our ability to study intelligence.

(Refer Slide Time: 49:12)

In 1943 McCulloch and Pitts built a Boolean circuit model of the brain. A Logical Calculus of Ideas Immanent in Nervous Activity was published and it explained for the first time how it is possible for neural networks to compute.

(Refer Slide Time: 49:32)

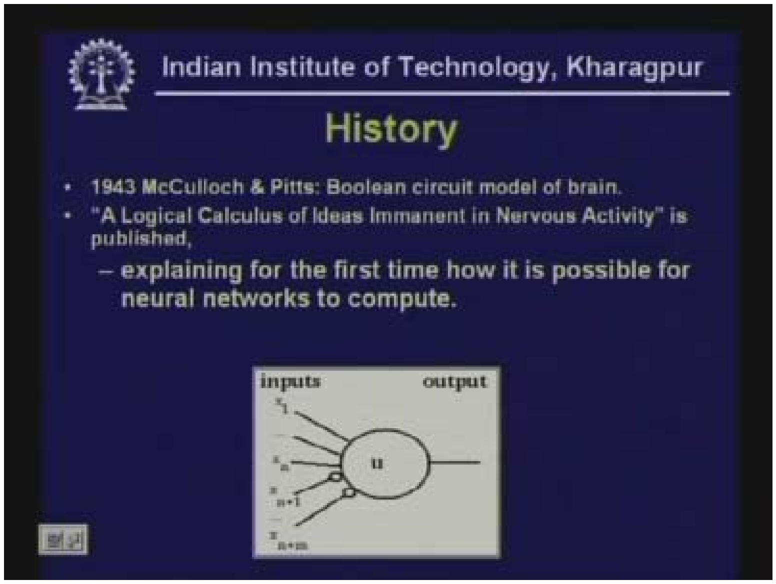

Marvin Minsky and Dean Edmonds built the SNARC in 1951 which is a neural network computer.

(Refer Slide Time: 49:44)

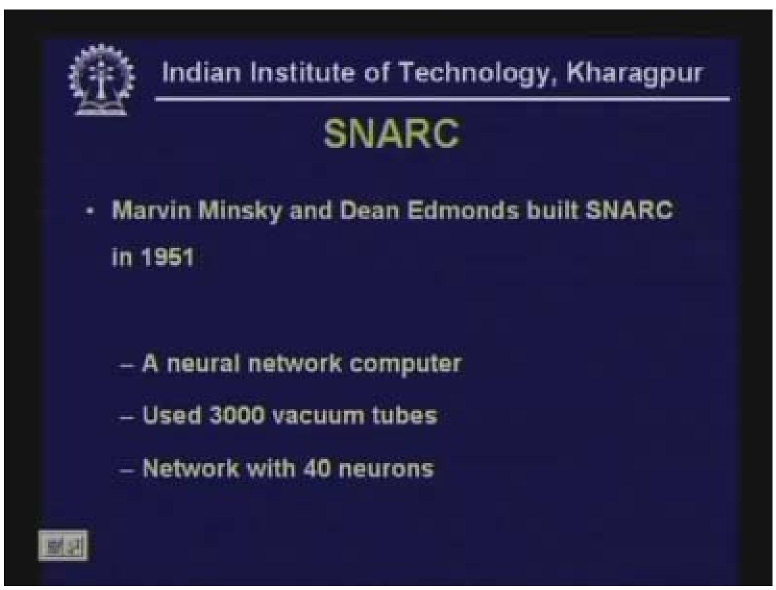

We have already seen Alan Turing. In 1950 Turing published his computing machinery and intelligence and this article articulated a complete vision of AI of solving problems, how AI systems can solve problems by searching through a space of possible solutions guided by heuristics.

(Refer Slide Time: 50:09)

He illustrated his ideas on machine intelligence by reference to chess. He propounded the possibility of letting the machine alter its own instructions so that machines can learn from experience. In 1952 to 56 Samuel designed a checkers playing program. In 1956 Allen Newell and Albert Simon designed the logic theorist. Then the general problem solver was built by the same people. In 1959 Gelernter developed the geometry engine for solving plane geometry problems.

(Refer Slide Time: 50:48)

In 1956 a meeting was held in Dartmouth where the first researchers in AI met. And in this month long meeting the term Artificial Intelligence was adopted. This conference brought together the founding fathers of AI for the first time. In 1961 James Slagle wrote the first symbolic integration program. This program saint could solve calculus problems at the college freshman level. In 1963 Thomas Evan's program analogy was designed, it could solve IQ test problems. In 1963 Feigenbaum and Feldman wrote a collection of important articles about AI.

(Refer Slide Time: 51:45)

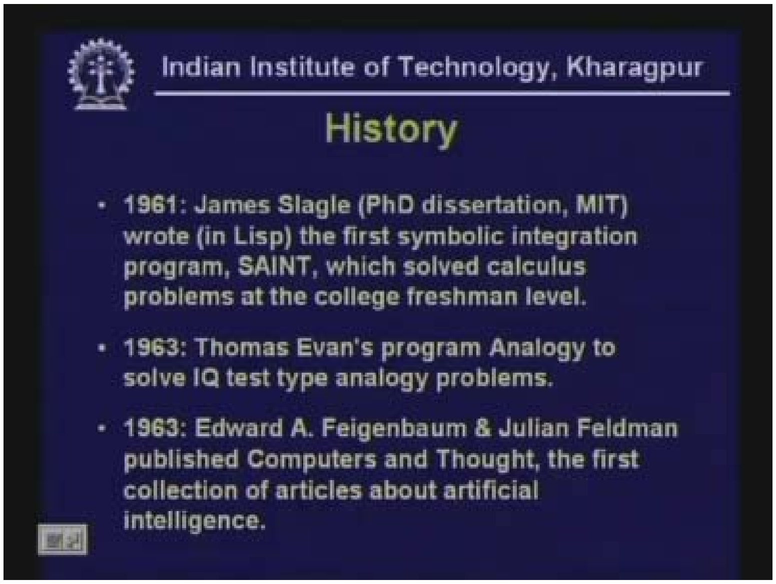

Then we have Danny Bobrow in 64 who worked with algebra word problems and in 1965 Allen Robinson developed a resolution method. In 1966 to 74 there was a lot of work on computational complexity by not really AI researchers but by computer theorists which had a tremendous impact on the field of AI.

(Refer Slide Time: 52:05)

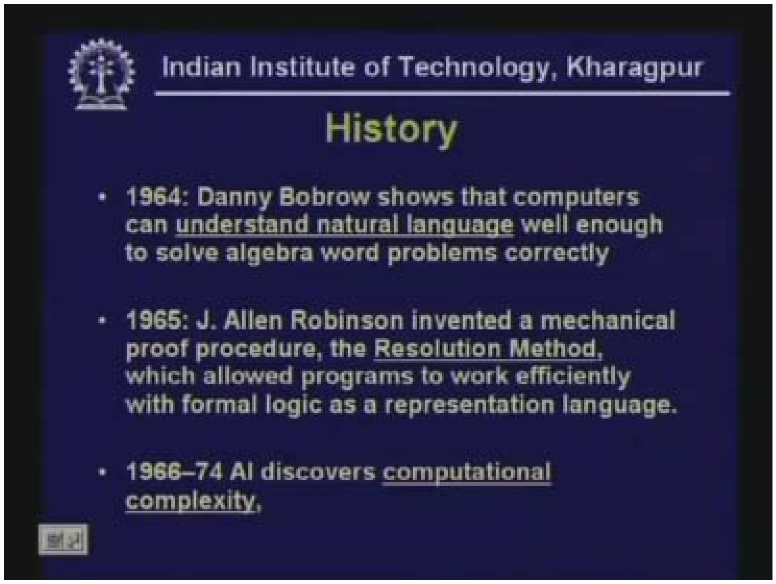

Before that people felt that a lot of things were possible by AI and we will soon have an extremely intelligent computer. But the limitations to the computational power was discovered when computational complexity was understood. In 1967 Feigenbaum and others developed a general program which was demonstrated used to demonstrate and interpret mass spectrum on organic chemical compounds.

In 1968 there was a very significant paper by Minsky and Papert which demonstrated the limits of simple neural net. This paper had a tremendous negative effect in discouraging the field of neural network for the time being. And later of course people realized that there are ways of coming out of this problem. In 1969 SRI robot, Shakey in Stanford demonstrated locomotion perception and problem solving.

(Refer Slide Time: 53:40)

In 1969 to 79 knowledge based systems were developed. In 1976 Doug Lenat handled the program called AM and Heurisko demonstrated the discovery model. In 1978 Herbert Simon from CMU won the Nobel Prize in Economics for his theory of bounded rationality.

(Refer Slide Time: 54:12)

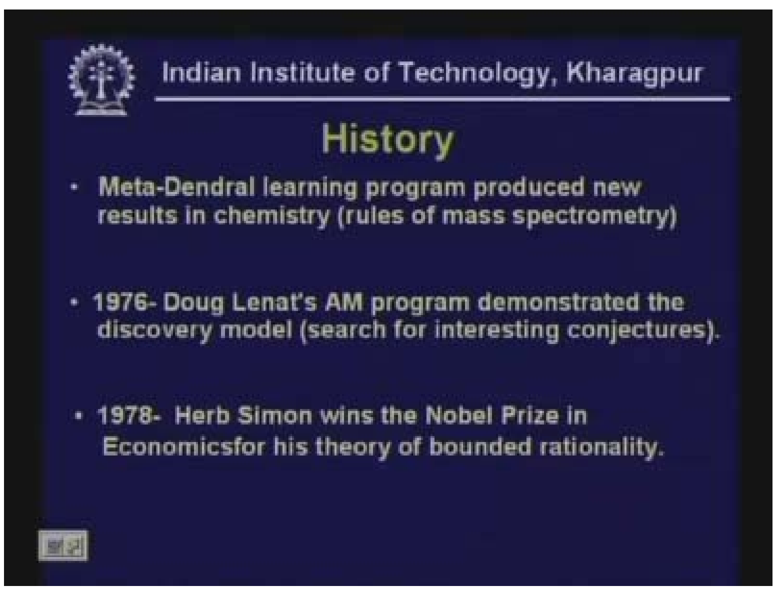

In 1980 lisp machines were developed and marketed. In 1985 to 95 neural networks returned to popularity. In 1988 there was a resurgence of probabilistic and decision theoretic methods.

(Refer Slide Time: 54:32)

Earlier AI systems used very general systems of little knowledge but recent AI systems use specialized knowledge to perform specific tasks.

(Refer Slide Time: 54:40)

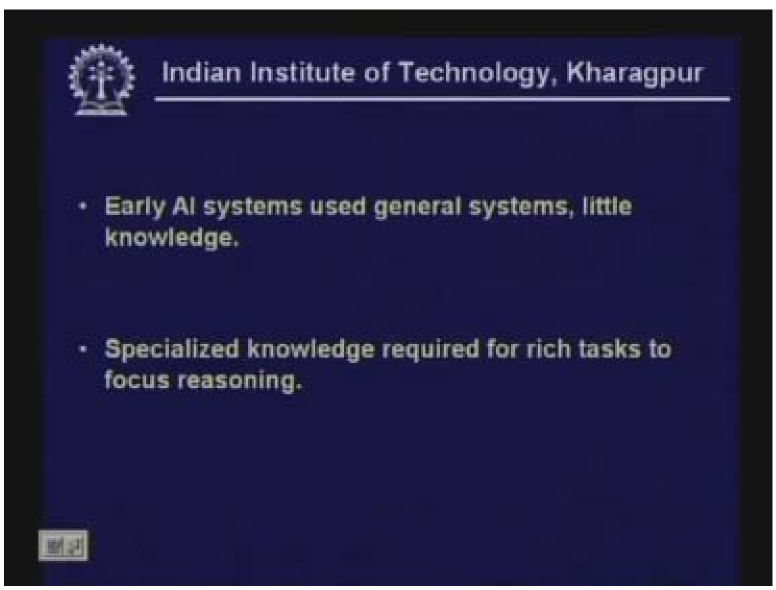

In 1990s there have been major advances in all areas of AI including machine learning, intelligent tutoring, multi agent planning, uncertain reasoning, natural language understanding, translation, vision and other topics. Rodney Brooks worked on the cog project at MIT which made significant progress in building a humanoid robot.

(Refer Slide Time: 55:10)

We have already looked at the deep blue chess playing program and we have interactive robot pets which have become commercially available realizing the vision of the 18th century toy makers. In 2000 the nomad robot explored remote regions of Antarctica and AI is a popular topic which is constantly in the news.

(Refer Slide Time: 55:38)

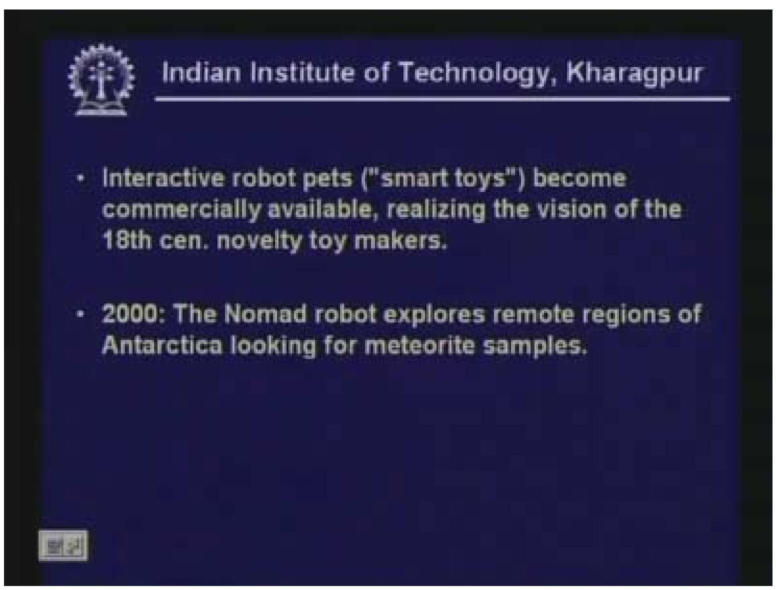

So this is the triple AI site which publishes news about AI and if you visit the site you will find that at any time there is a lot of interesting news on AI.

(Refer Slide Time: 56:00)

With this we will end today's lecture and before we end we have a few questions.

(Refer Slide Time: 56:16)

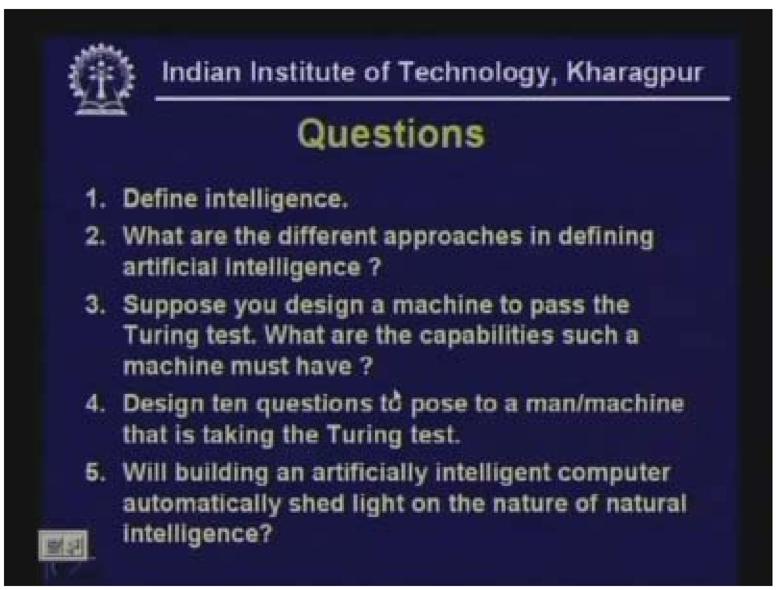

Question 1 is, define intelligence.
Question 2: What are the different approaches in defining Artificial Intelligence?
Question 3: Suppose you design a machine to pass the Turing test what are the capabilities such a machine must have?
Question 4 is, design ten questions to pose to a man or a machine that is taking a Turing test.

Question 5 is, will building an AI computer automatically shed light on the nature of natural intelligence, do you think so?

(Refer Slide Time: 57:01)

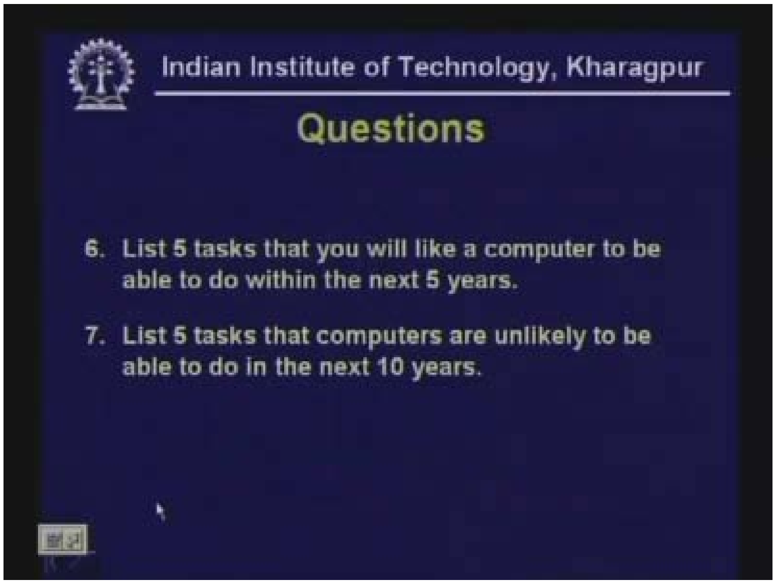

Question 6 is, list five tasks that you will like a computer to be able to do within the next five years. The last question, question 7, list five tasks that computers are unlikely to be able to do in the next ten years. With this we come to the end of today's lecture, thank you.