**Fundamentals of Artificial Intelligence**
**Prof. Shyamanta M Hazarika**
**Department of Mechanical Engineering**
**Indian Institute of Technology - Guwahati**

**Module - 1**
**Lecture - 1**
**Introduction to AI**

Welcome to the course of Fundamentals of Artificial Intelligence. This is the first lecture, An Introduction to AI. In this lecture, we will look at the quest for AI; the thoughts, ideas, work that have influenced the growth of artificial intelligence. Thereafter, we would look at the definition of AI and the different dimensions of artificial intelligence.

**(Refer Slide Time: 01:02)**

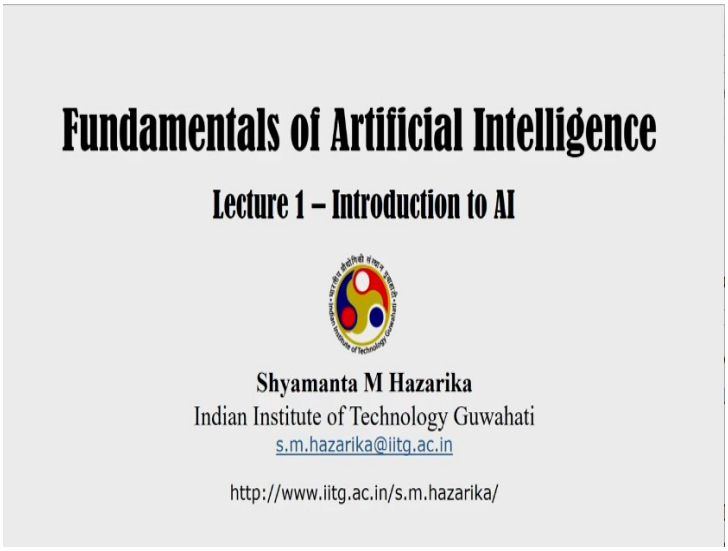


We would look at what is involved in an AI system and try to understand the difference between weak and strong AI. We would also very quickly go through different developments in the history of AI. The quest for artificial intelligence is a story of dreams and began with dreams as all quests do.

**(Refer Slide Time: 01:33)**

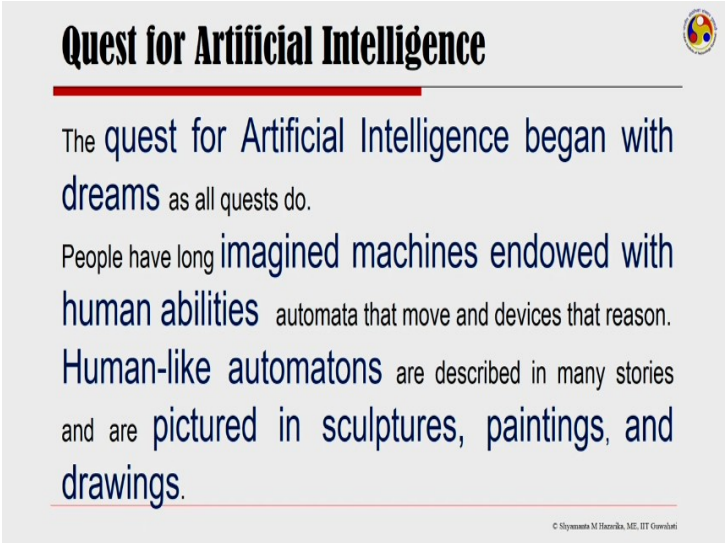


People have long imagined machines or devices that could reason and have abilities like the human; automatons that could move. In fact, human like automatons are described in many stories and are pictured in sculptures, paintings and drawings. Perhaps, the greatest of the thoughts was described in `The Politics' by Aristotle.

**(Refer Slide Time: 02:08)**

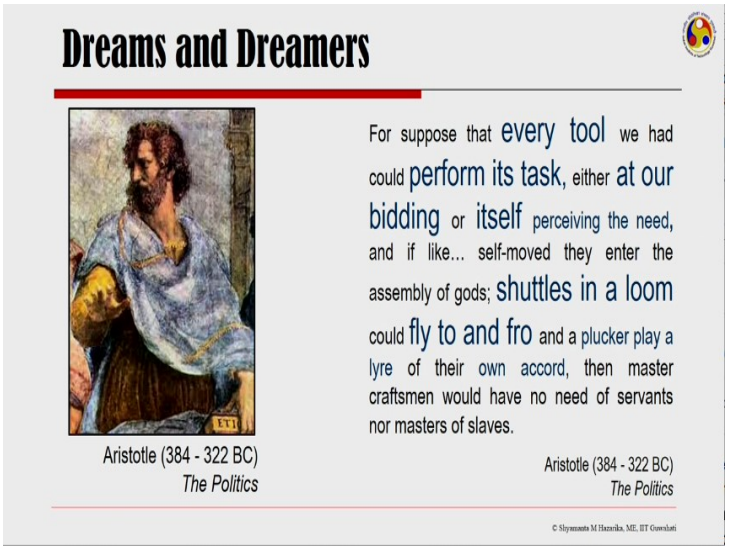


He thought of tools that could perform task at our bidding or itself. Shuttles in loom that could fly to and fro. And in fact, had generated the idea of having automated machines.

**(Refer Slide Time: 02:26)**


It was Leonardo Da Vinci who sketched designs for a humanoid robot in the form of a medieval knight around 1495. It is very difficult to be sure. And no one knows whether Leonardo Da Vinci or his contemporaries tried to build his design. Leonardo's knights, never the less, was supposed to be able to sit up, move its arms and head, opens its jaw. It could do everything that a human did in its mechanical form.

**(Refer Slide Time: 03:02)**

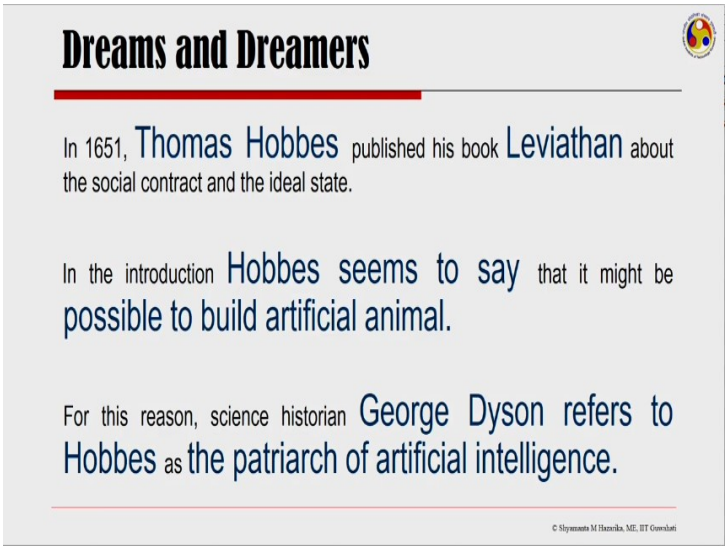


In fact, in 1651, Thomas Hobbes published his book Leviathan. And in that book Hobbes seems to say that it might be possible to build an artificial animal. And because of that, George Dyson refers to Hobbes as the patriarch of artificial intelligence.

**(Refer Slide Time: 03:21)**

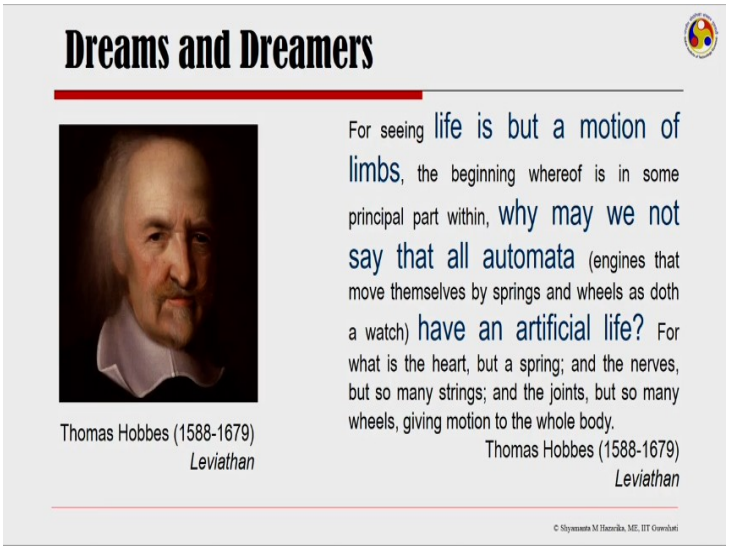


In Leviathan, Hobbes writes that life is but a motion of limbs. Why may we not say that all automatons have an artificial life? And he continues to say: what is the heart, but a spring; and the nerves, but so many strings; and the joints, but so many wheels, giving motion to the whole body. It seems, Thomas Hobbes in Leviathan, actually described the possibility of artificial life. **(Refer Slide Time: 03:55)**


Several people constructed actual automata that moved in startlingly lifelike ways. Perhaps, the most sophisticated of these was a mechanical duck which could quack, flap its wings, paddle, drink water, eat and digest grain.

**(Refer Slide Time: 04:13)**

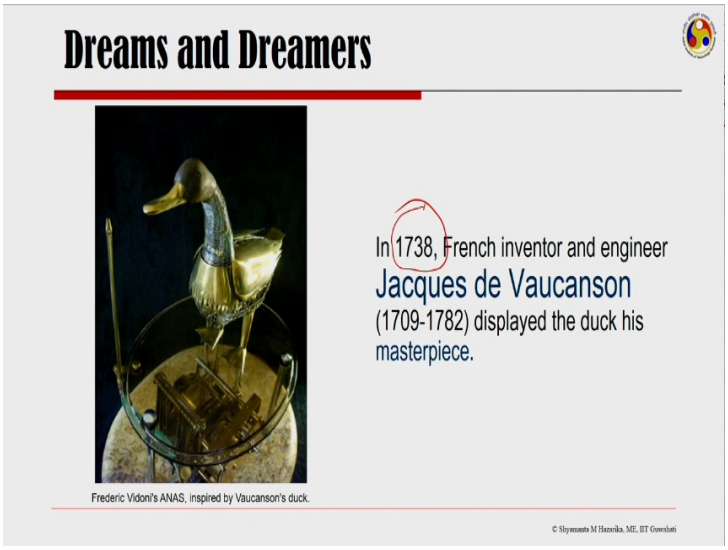


This is what the duck look like. So, in 1738, French inventor and engineer Jacques de Vaucanson displayed the duck, his masterpiece.

**(Refer Slide Time: 04:26)**

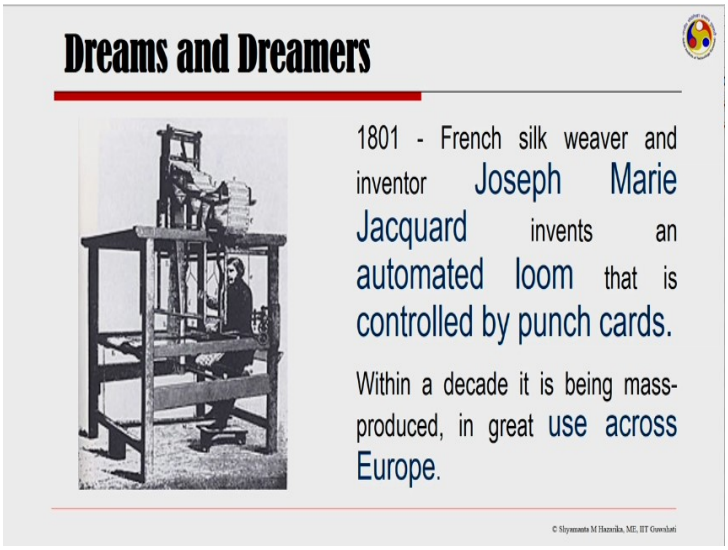


In fact, very shortly after that, 1801, French silk weaver and inventor Joseph Marie Jacquard invented the automated loom. Starting with Aristotle's idea in the politics, we had a real loom that is controlled by punch cards. It revolutionized the mass production and was in great use across Europe.

**(Refer Slide Time: 04:52)**

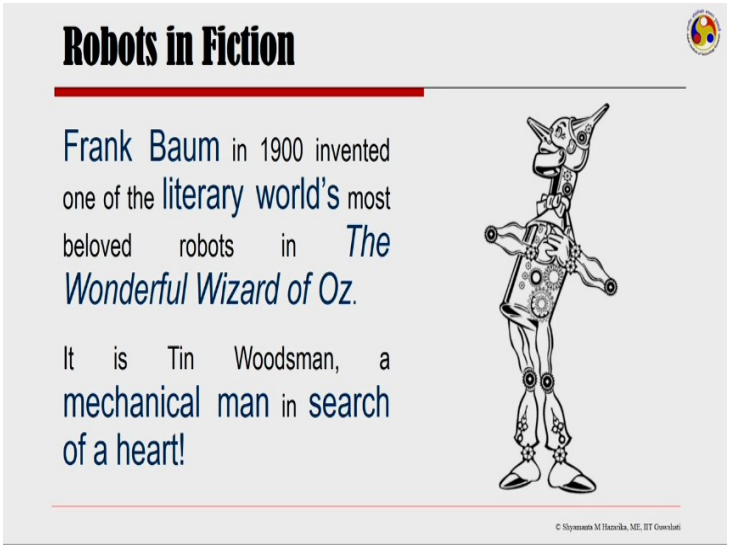


Frank Baum in 1900 invented one of literary world's most beloved robots, The Wonderful Wizard of Oz. This was a mechanical man and in search of a heart.

**(Refer Slide Time: 05:07)**

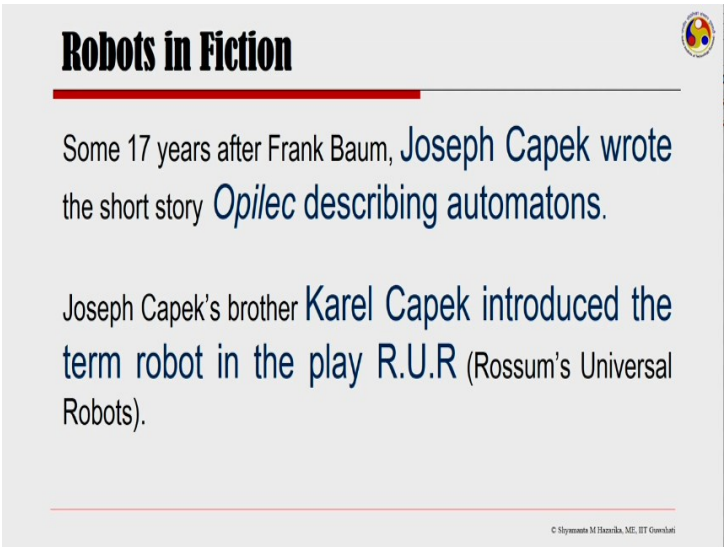


Not only in fiction, but in real life, these things were supposed to become reality soon. Some 17 years after Frank, Joseph Capek wrote the short story Opilec, describing automatons. And Joseph's brother Karel introduced the term robot in the play Rossum's Universal Robots. This is when we started having humanoids with brains, albeit in fiction.

**(Refer Slide Time: 05:41)**

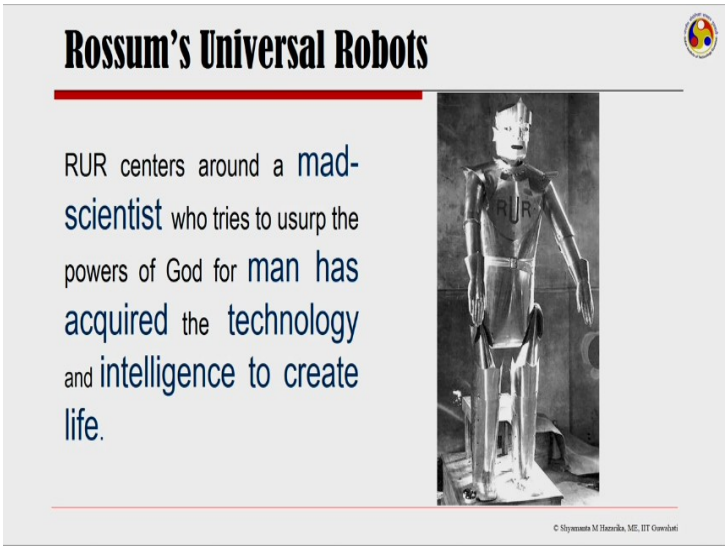


Rossum's Universal Robot centers around a mad scientist who tries to get the powers of God for man has acquired the technology and intelligence to create life.

**(Refer Slide Time: 05:54)**

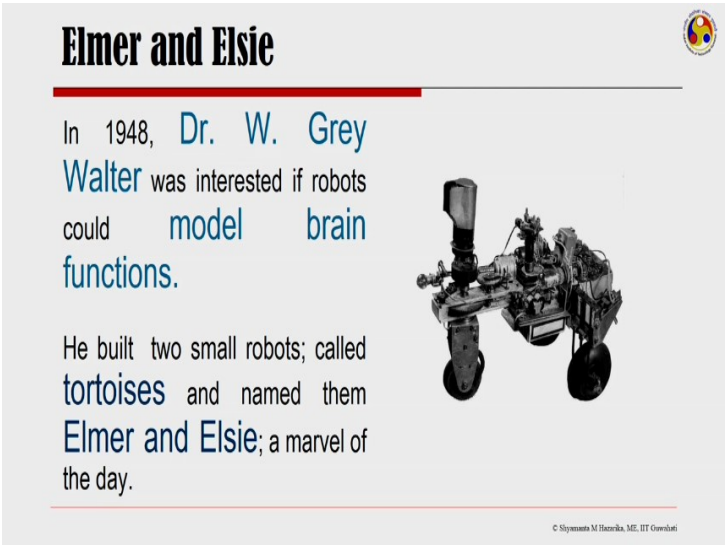


Not in fiction but in reality; two important tortoises like robots, Elmer and Elsie revolutionized the growth of AI. In 1948, Dr. W. Grey Walter was interested if these robots could model brain functions. And he built 2 small robots which we called tortoises and he named them Elmer and Elsie. Elmer and Elsie were a marvel of the day.

**(Refer Slide Time: 06:31)**

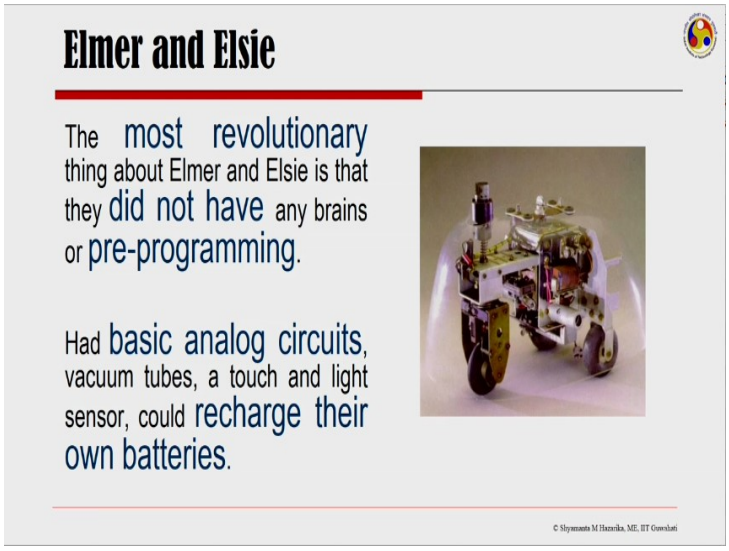


They were the most revolutionary things because they did not have any pre-programming. And the most interesting thing was that, in spite of having basic analog circuits, they could go and recharge their own batteries when they felt that the batteries were running down. Perhaps, this was a turning point in the history of the quest for artificial intelligence. What then is artificial intelligence?

**(Refer Slide Time: 06:58)**

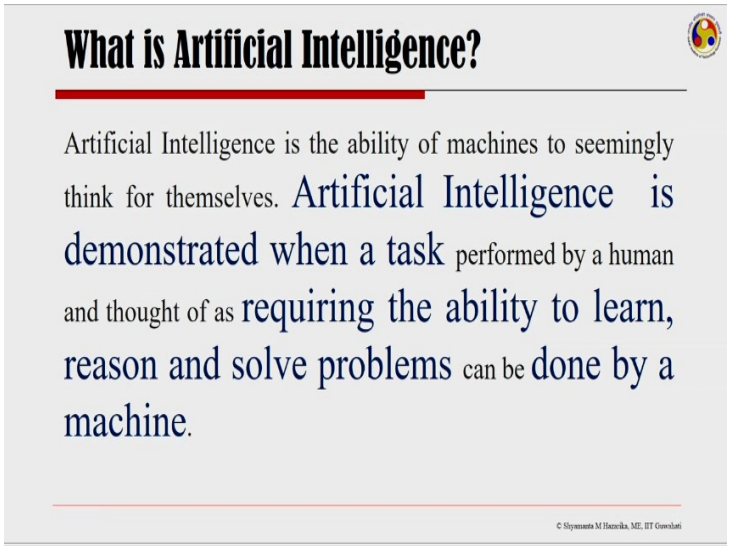


We say artificial intelligence is demonstrated when a task that is performed by a human and thought of as requiring the ability to learn, to reason and solve problems can be done by a machine. This is itself a deep philosophical question.

**(Refer Slide Time: 07:16)**

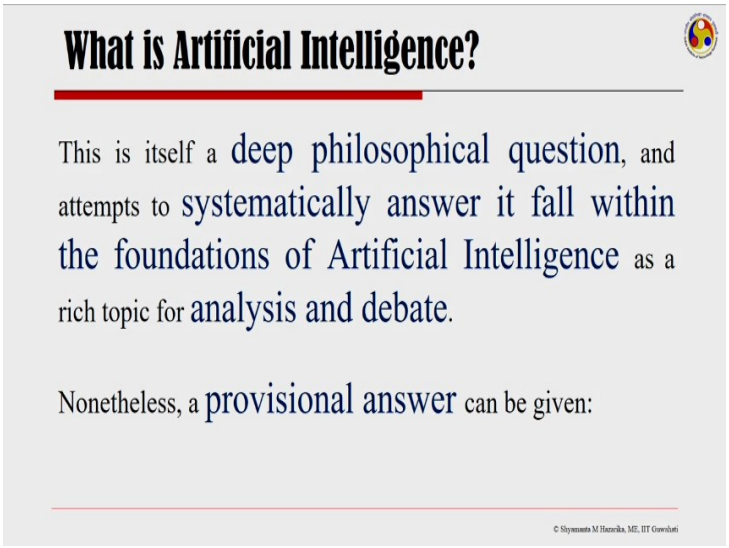


Attempts to systematically answer it will fall within the foundations of artificial intelligence as a rich topic for analysis and debate. This is not the aim of this lecture. So, we will for the time being have a provisional answer. A provisional answer could be the following.

**(Refer Slide Time: 07:38)**

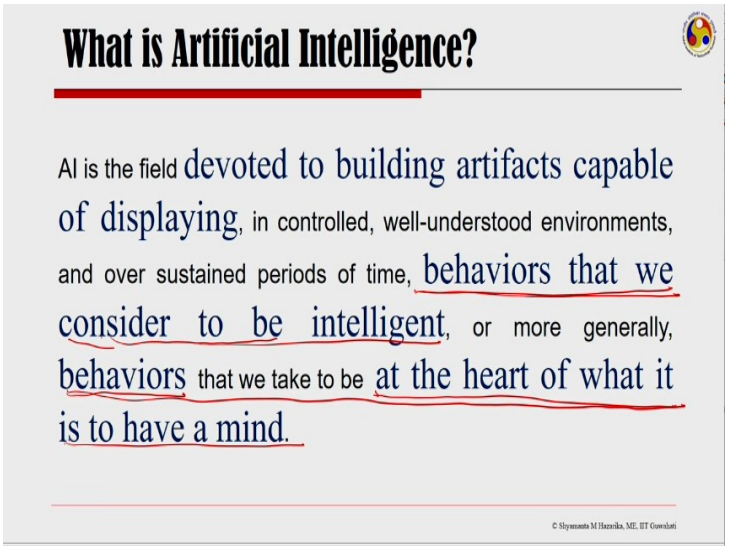


Artificial intelligence is the field devoted to building artifacts capable of displaying in controlled, well-understood environments and over sustained periods of time, behaviors that could be considered to be intelligent; or more generally, behaviors that are at the heart of what it is to have a mind. 2 things are very important in this working definition of AI that we go forward with.

We need to look at behaviors that we consider to be intelligent. There may not be anything of a human brain there. But then, we also look at behaviors that we believe that at the heart of it, we need to have a mind or a brain. This is itself a deep philosophical question.

**(Refer Slide Time: 08:41)**

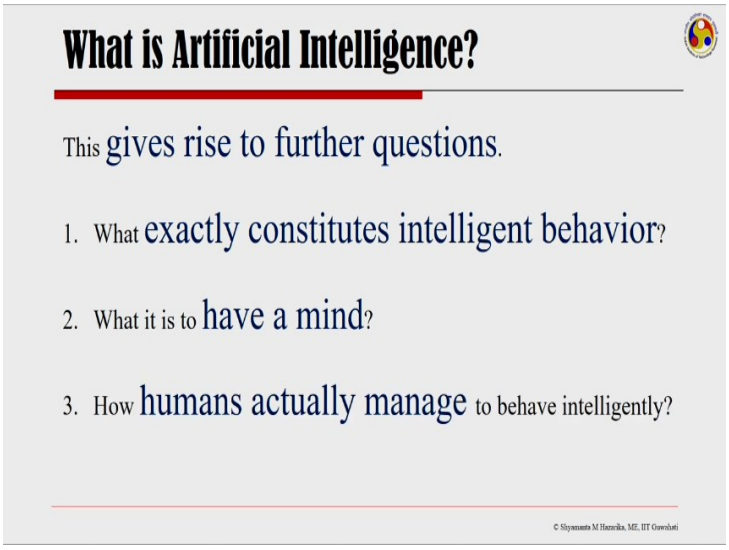


This gives rise to further questions. 1. What exactly constitutes intelligent behavior? 2. What is to have a mind? 3. How humans actually manage to behave intelligently? Let us look at each of this question one by one.

**(Refer Slide Time: 09:03)**

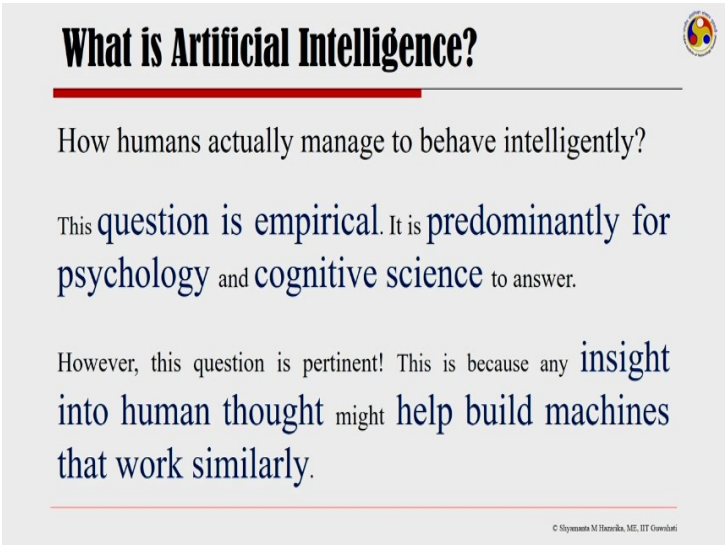


How humans actually manage to behave intelligently? This question is empirical. It is predominantly for psychology and cognitive science to answer. However, this is a very pertinent question for AI. This is because, any insight into how the human thought process work may help build machines that work similarly. In fact, it is very difficult to separate out

the growth of cognitive science and the growth of AI. Each have lent help to the other and their history in intertwined.

**(Refer Slide Time: 09:42)**

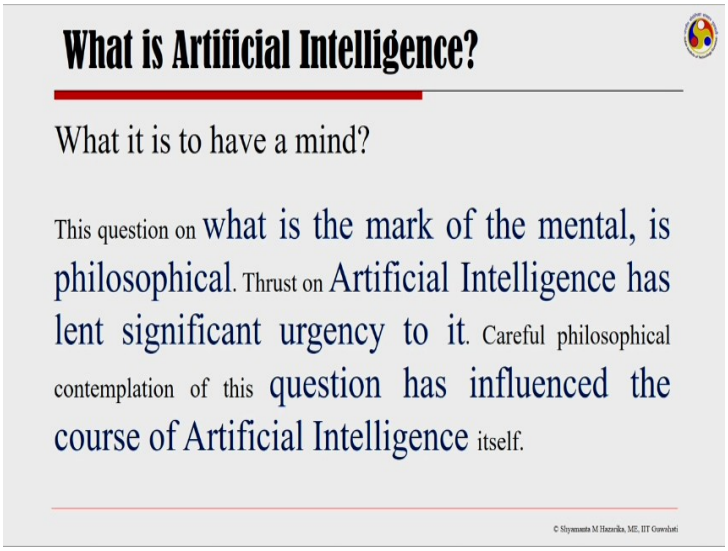


The next question on what it is to have a mind is a question that what is the mark of the mental; and is a philosophical question. The trust on artificial intelligence has lent significant urgency to it. In fact, very careful philosophical contemplation of that question has influenced the course of artificial intelligence itself. The third question, what exactly constitutes intelligent behavior is:

**(Refer Slide Time: 10:12)**

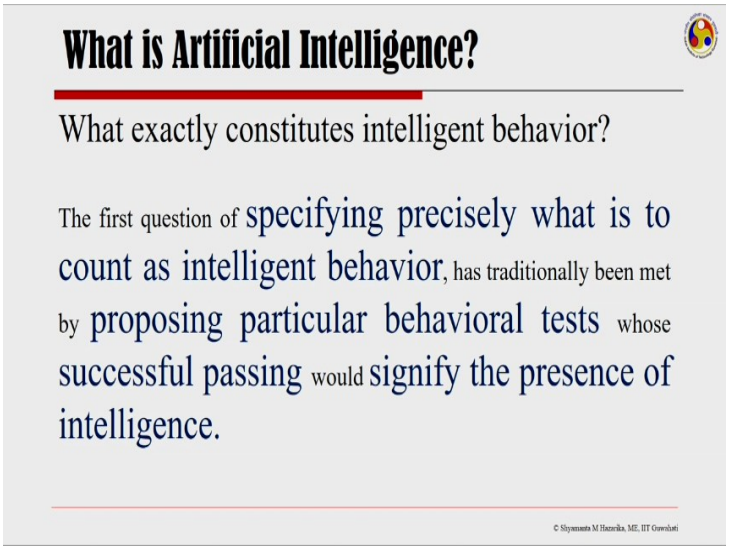


The question specifying precisely what it is to count as intelligent. This has been traditionally met by proposing particular behavioral tests whose successful passing would signify the presence of intelligence.

**(Refer Slide Time: 10:31)**

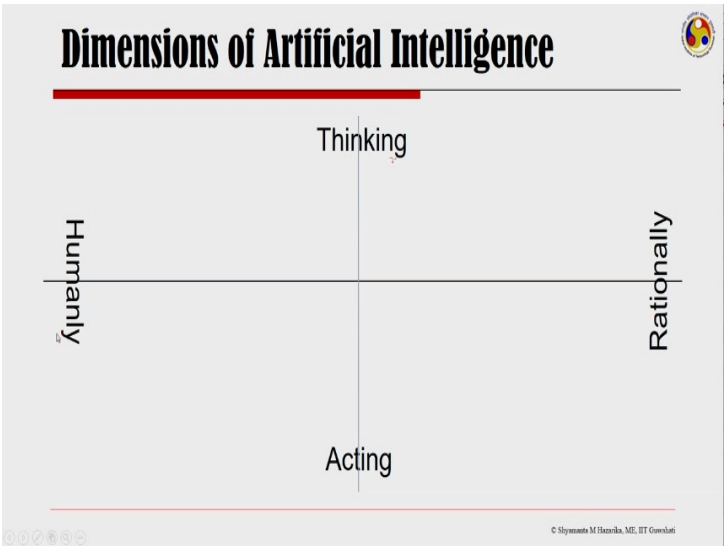


Let us now look at the dimensions of artificial intelligence vis-a-vis these three questions. So, we could have four different things involved. It would be like, whether something is thinking, acting; or it is giving rational behavior; or it is giving human behavior. So, an artificial system will either fall into 1 of these 4 ways of looking at things.

**(Refer Slide Time: 11:02)**

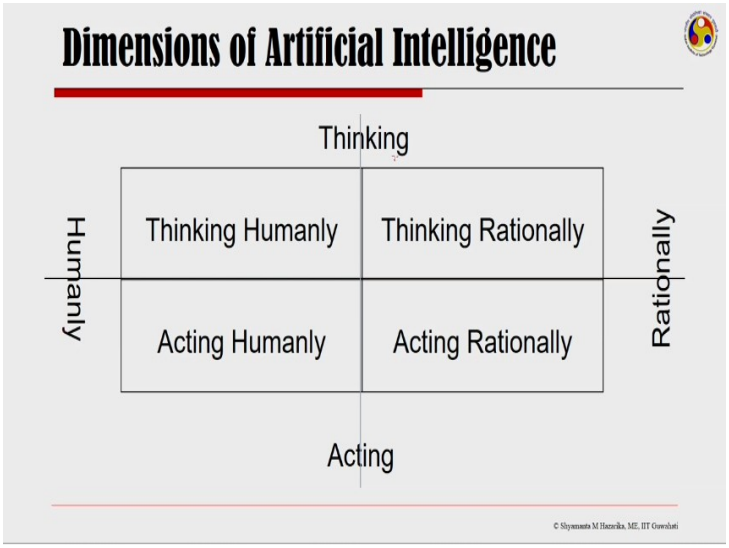


And therefore, we could have systems which would be thinking humanly or thinking rationally or acting rationally or acting humanly.

**(Refer Slide Time: 11:12)**

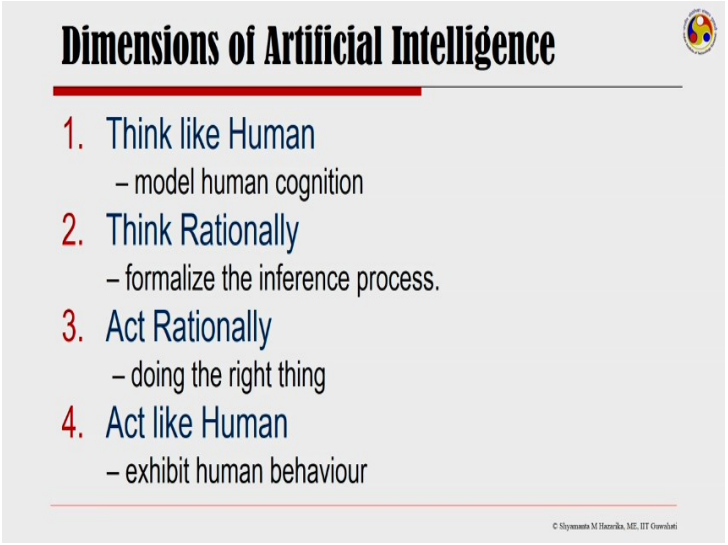


Think like human. When we say, we mean we need to model human cognition. Think rationally would involve formalizing the inference process. To act rationally would mean doing the right thing always. And the fourth, act like human, would involve exhibiting human behavior.

**(Refer Slide Time: 11:35)**

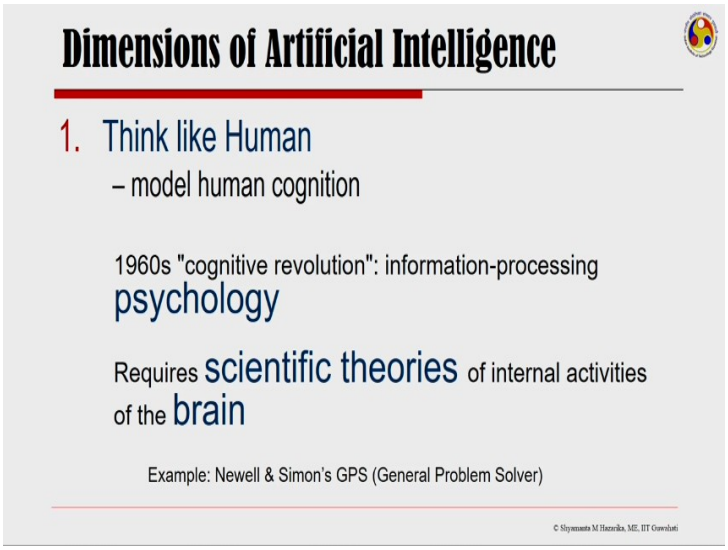


Let us look at each of this dimensions one after another. The first one, thinking like the human involves modeling human cognition; and is an area that has been looked at very closely by the information processing community within psychology and has got huge push from the cognitive revolution. It requires scientific theories of internal activities of the brain. In fact, Newell and Simon's General Problem Solver is a kind of this intelligences.

**(Refer Slide Time: 12:18)**

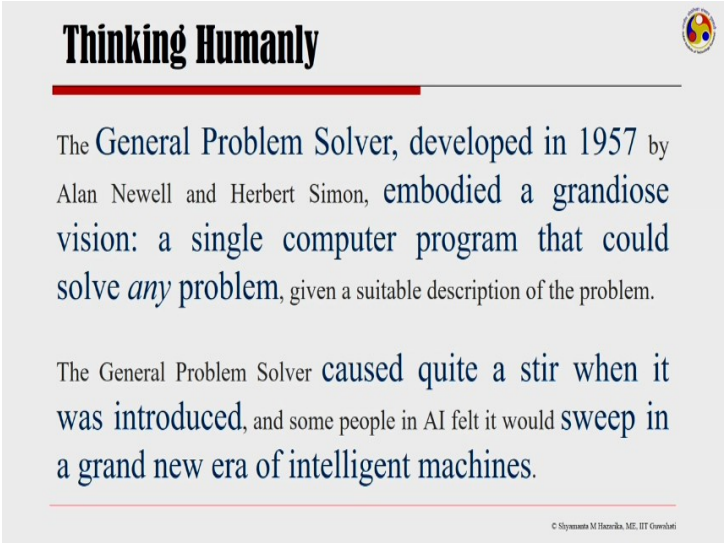


The General Problem Solver developed in 1957 by Alan Newell and Herbert Simon embodied a grand vision. It thought of a single computer program that would solve any problem; given a sustainable description of the problem. It did cause a quite a stir when it was introduced. And many people felt it would sweep in a grand new era of intelligent machines. It is a different story that it did not happen the way it was hoped to make its contribution.

**(Refer Slide Time: 12:55)**

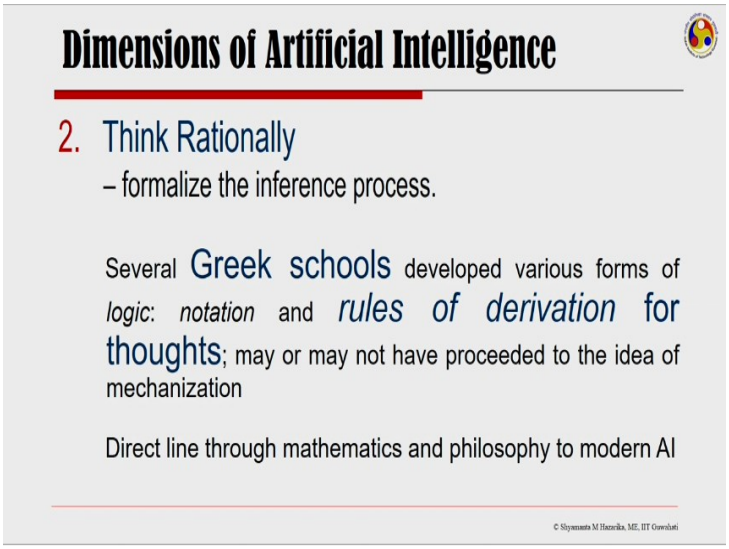


The next dimension of artificial intelligence is about thinking rationally, which involves formalizing the inference process. In fact, there are several Greek schools which developed various forms of logic; notation and rules of derivation of thoughts. They may or may not have proceeded to the idea of mechanization of these processes. But the very idea of

formalizing the inference process as a direct line through mathematics and philosophy to modern AI. In fact, Aristotle has a huge contribution.

**(Refer Slide Time: 13:34)**

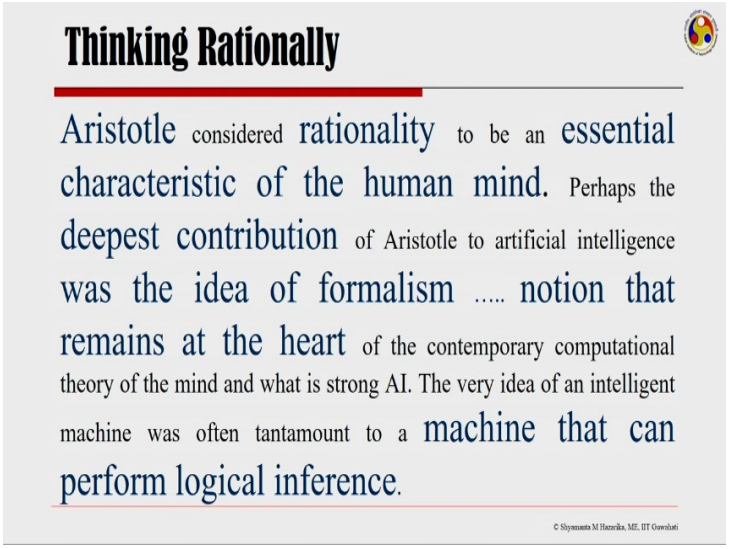


Aristotle considered rationality to be an essential characteristic of the human mind. Perhaps the deepest contribution of Aristotle to AI was the idea of formalism. Notion that remains at the heart of contemporary computational theory of mind and what we will see later to be strong AI. The idea of an intelligent machine was often thought to be a machine that can perform logical inference.

**(Refer Slide Time: 14:07)**

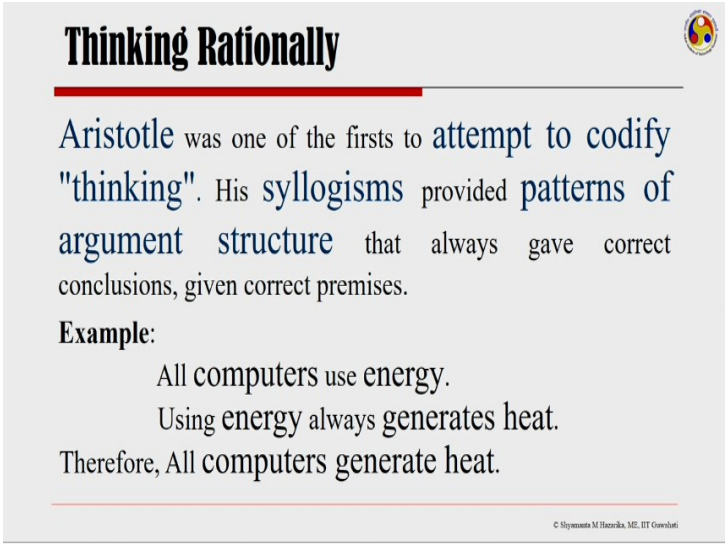


Aristotle was one of the first to attempt and codify thinking. His syllogisms provided patterns of argument structure that always give the correct conclusion, given correct premises. Let us take an example and see what it means. Here is an example. This says that all computers use

energy. We could then say using energy always generates heat. And from these 2 sentences, we could try conclusion that all computers generate heat. Aristotle told us how to do it in a pattern of argument structure.

**(Refer Slide Time: 14:51)**

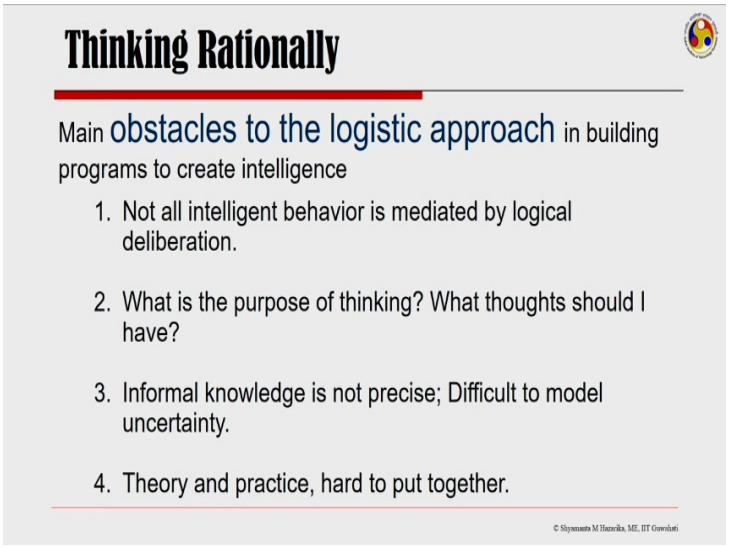


But then, there are many obstacles to the logistic approach in building programs to create intelligence. They are: Number 1, not all intelligent behavior is meditated by logical deliberation. Number 2, if everything is logical deduction, question is, what is the purpose of thinking? What thoughts should I have? Informal knowledge is not precise. And therefore, under a logical premise, it is difficult to model uncertainty.

Number 4, the theory and practice of putting everything within a logical system is hard to put together for real life problems. Let us now focus attention on the third dimension of artificial intelligence, which is act rationally.

**(Refer Slide Time: 15:48)**

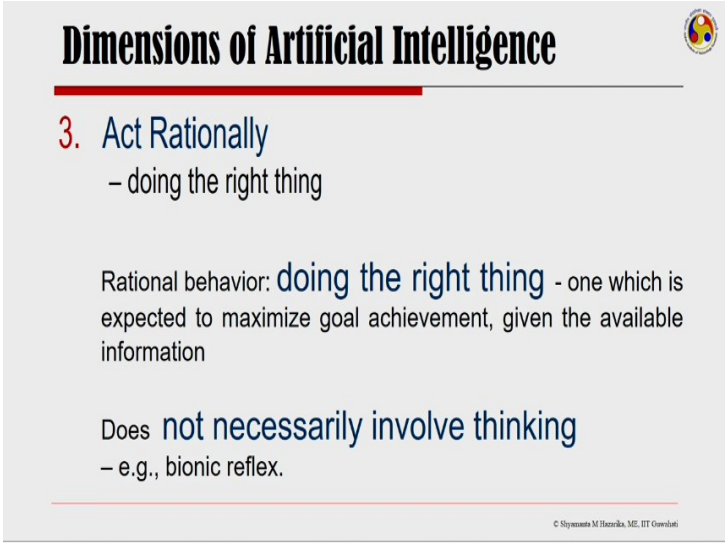


That is, doing the right thing. Rational behavior or doing the right thing is where one is expected to act to maximize goal achievement, given the available information. Now, many of our rational acts do not necessarily involve thinking. Most of the bionic reflexes that we do; like if your glass is slipping out of your hand, you just re-grasp it with more pressure. We do it without thinking. So, they do not necessarily involve thinking.

**(Refer Slide Time: 16:24)**

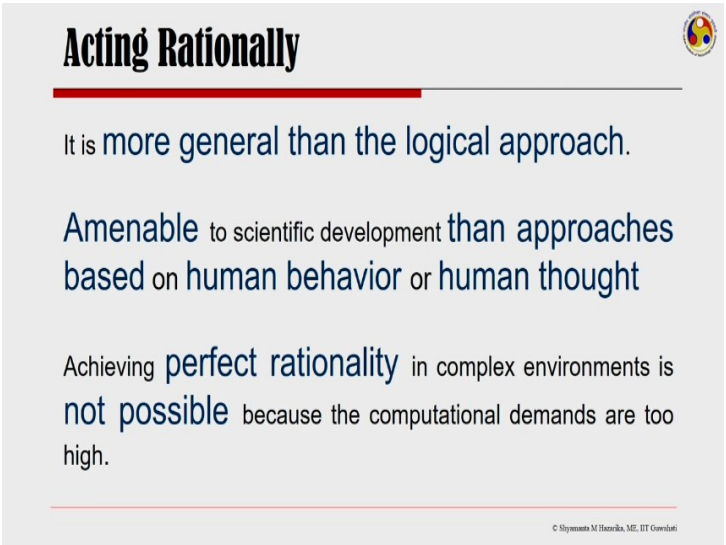


It is definitely more general than the logical approach and amenable to scientific development than approaches based on human behavior or human thought. However, achieving perfect rationality in complex environments is not possible. This is because of the high computational demands that such systems put.

**(Refer Slide Time: 16:50)**

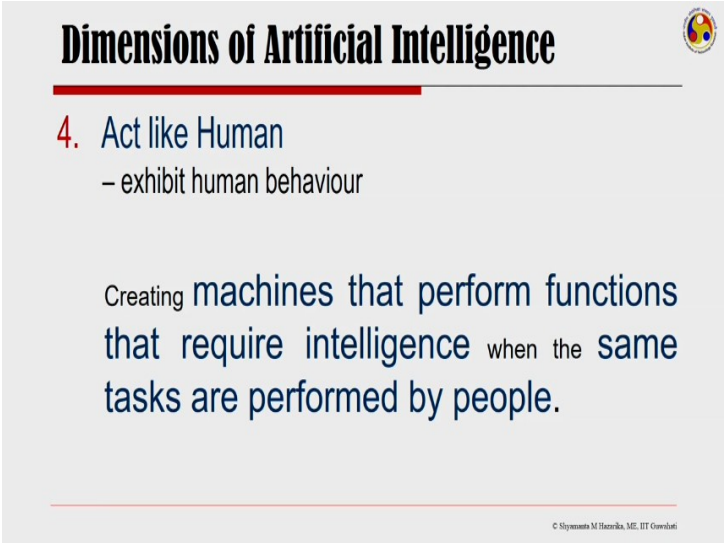


We now focus attention on our fourth dimension which is, act like human. That is, exhibiting human behavior. Creating machines that perform functions that require intelligence when the same functions are performed by people.

**(Refer Slide Time: 17:08)**

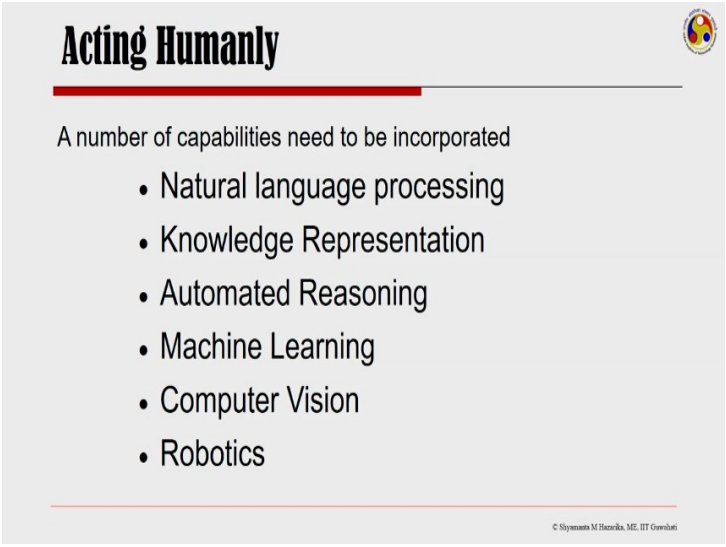


There are a number of capabilities to be incorporated in such systems, that need to act like human. They include things like natural language processing, knowledge representation, automated reasoning, machine learning, computer vision and of course robotics.

**(Refer Slide Time: 17:27)**

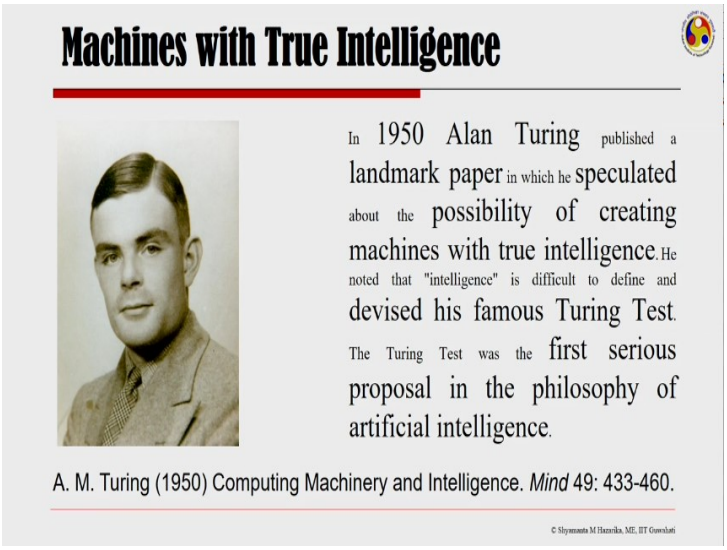


It was Alan Turing who published his landmark paper in which he speculated about the possibility of creating machines with true intelligence. He noted that intelligence is difficult to define. And therefore, he devised the famous Turing Test. In fact, the Turing Test was the first serious proposal in the philosophy of artificial intelligence.

**(Refer Slide Time: 18:00)**

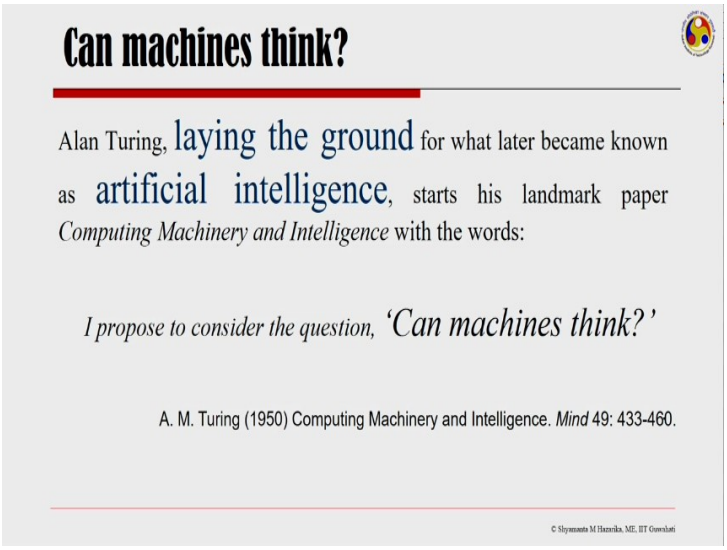


In his paper Computing Machinery and Intelligence, Alan Turing starts with the words, I propose to consider the question, can machines think? And he laid the ground for what later became to be known as artificial intelligence.

**(Refer Slide Time: 18:14)**

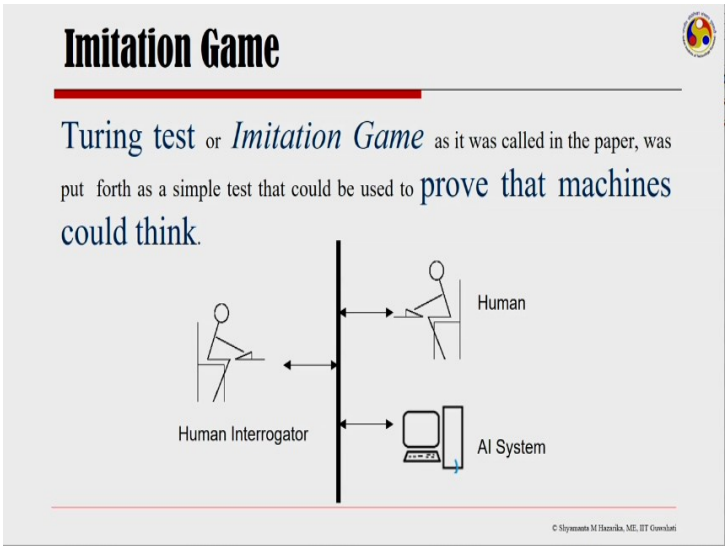


The Turing Test or the Imitation Game, as it was called in the paper, was put forth as a simple test that could be used to prove that machines could think. The Turing Test involves the following. There is a human interrogator talking to a human or to a machine, unaware of whom he gets responses from. If the interrogator is not able to distinguish the responses received from the man or the system, the system is said to be artificially intelligent. Turing also has lot of contribution apart from this, in the history of AI.

**(Refer Slide Time: 18:58)**

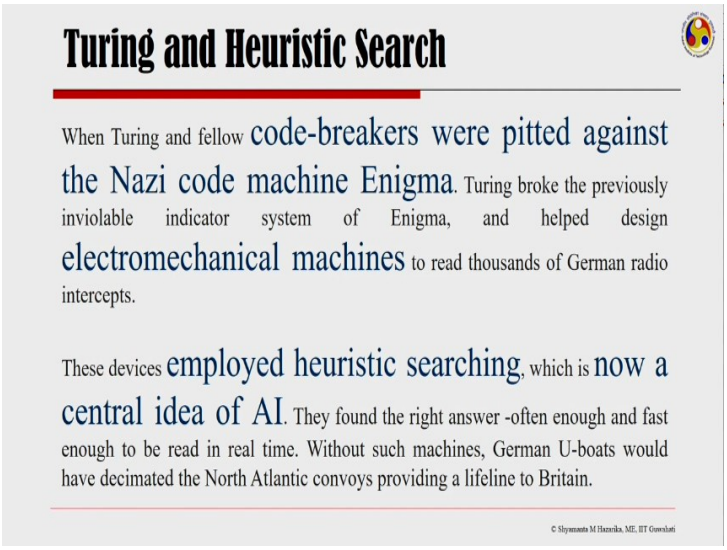


Turing and his fellow code-breakers were actually pitted against the Nazi code machine Enigma. And Turing put up electromechanical machines to read thousands of German radio intercepts. In fact, these machines employed something which we now have as a central idea of AI called heuristic searching. They found a right answer, often enough and fast to be read

in real time. In fact, without such machines, German U-boats would have decimated the North Atlantic convoys.

**(Refer Slide Time: 19:34)**

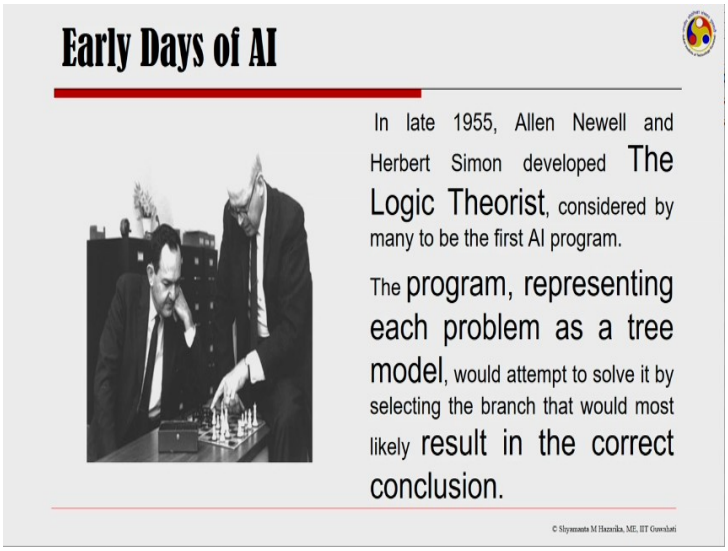


We now are able to talk of the early days of AI. Having looked at the dimensions; having looked at the general growth of AI; in late 1955, Alan Newell and Herbert Simon developed the logic theorist. This can be considered as the first AI program. The program representing each problem as a tree model, attempt to solve it by selecting the branch that would most likely result in the correct conclusion.

**(Refer Slide Time: 20:10)**

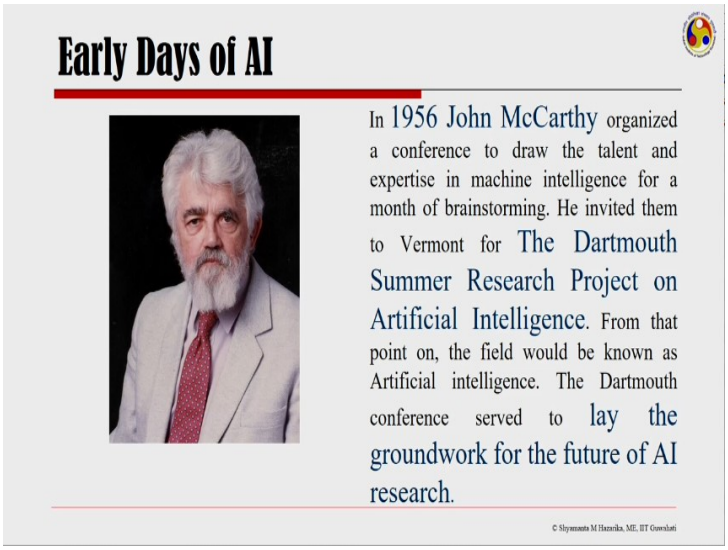


However, it was in 1956 that John McCarthy organized the Dartmouth Summer Research Project on artificial intelligence. And in fact, from that point on, the field came to be known

as artificial intelligence. The Dartmouth conference served to lay the groundwork for future of AI research.

**(Refer Slide Time: 20:36)**

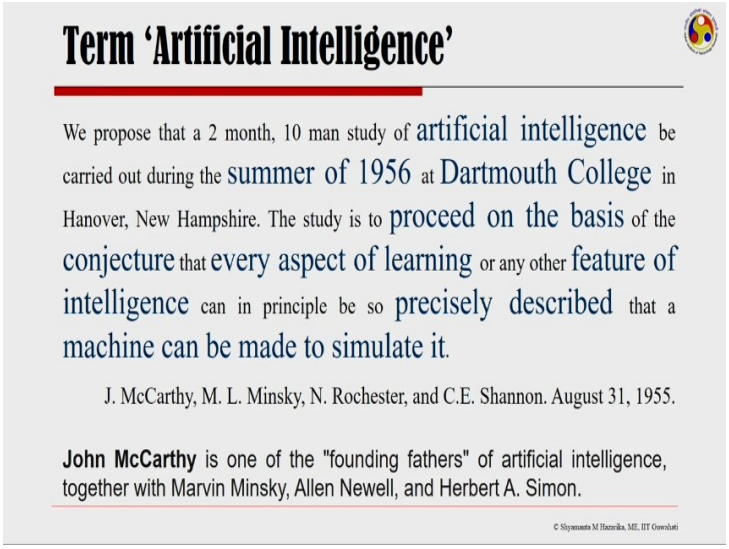


The term artificial intelligence was first used in their document. They say it and I read here for you. *We propose that a two month, 10-man study of artificial intelligence be carried out during the summer of 1956, at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture, that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. McCarthy, Minsky, Rochester and Shannon; August 31, 1955.*

In fact, John McCarthy is one of the founding fathers of artificial intelligence together with Marvin Minsky, Alan Newell and Herbert Simon.

**(Refer Slide Time: 21:34)**

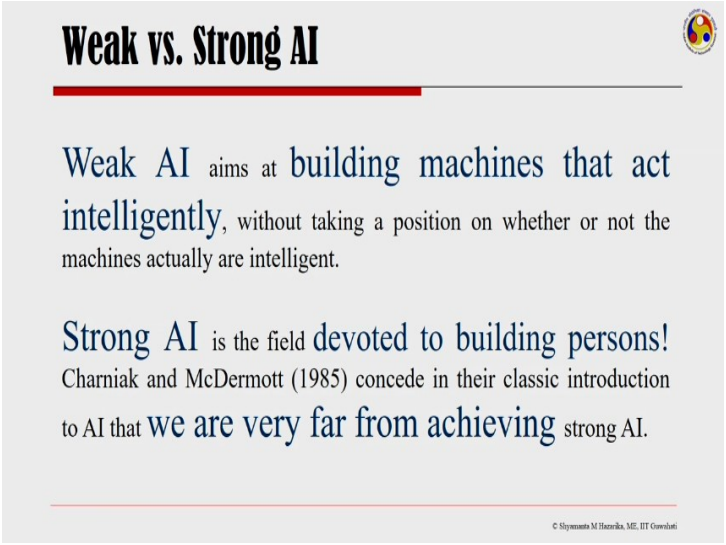


We now have looked at the dimensions of AI and have realized that we could be either thinking humanly, thinking rationally, acting rationally or acting humanly. This leads us to 2 distinct concepts of AI. The weak versus the strong AI. The weak AI is aimed at building machines that act intelligently. Now, they do not take a position on whether or not the machines actually are intelligent.

Whereas strong AI is a field which is devoted to building persons. As Charniak and McDermott concede in their classic introduction to AI, that we are very far from achieving what is strong AI. However, the ultimate goal of AI which we are very far from achieving is to build a person or more humbly an animal.

**(Refer Slide Time: 22:33)**

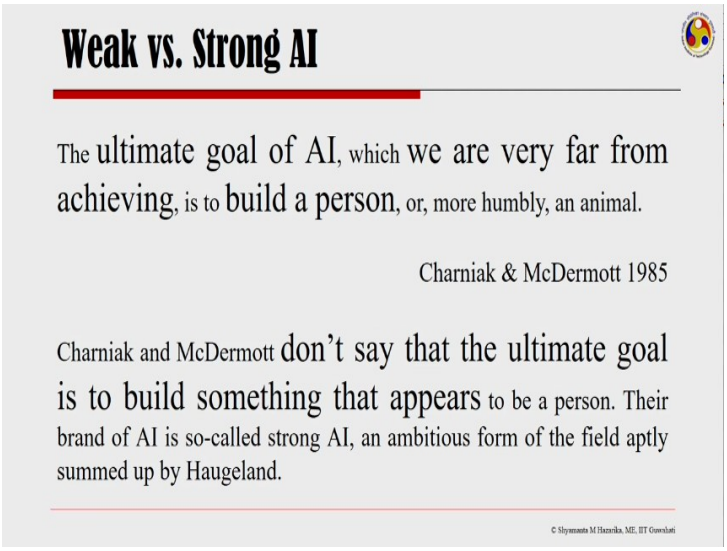

Charniak and McDermott do not say that the ultimate goal is to build something that appears to be a person or acts like a person. Their brand of AI is so called strong AI. An ambitious form of the field aptly summed up by Haugeland.

**(Refer Slide Time: 22:54)**

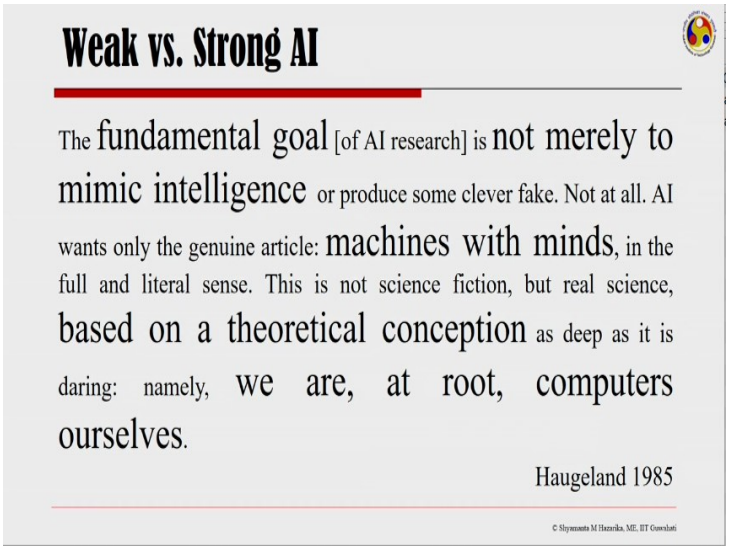


The fundamental goal according to Haugeland is not merely to mimic intelligence or produce some clever faith; not at all. AI wants only the genuine article; machines with minds, in the full and literal sense. This is not science fiction but real science based on theoretical conception as deep as it is daring; namely, we are at a root, computers ourselves.

**(Refer Slide Time: 23:29)**

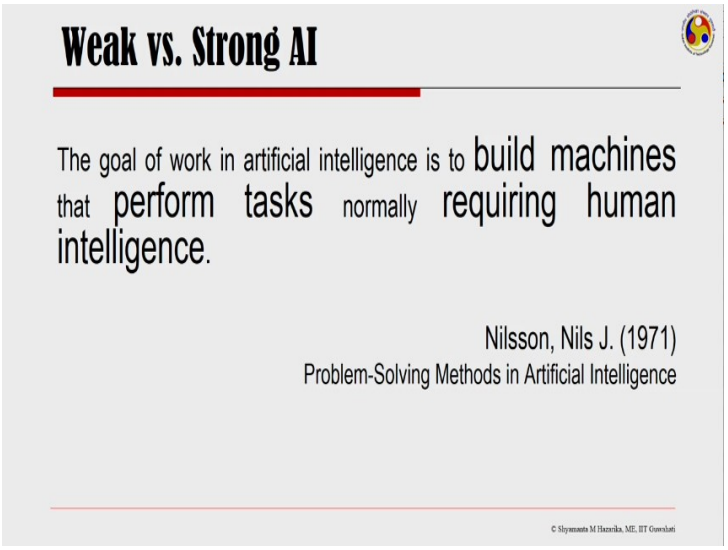


And this is where Nilsson's definition of AI is interesting to note. The goal of work in AI according to Nilsson is to build machines that perform tasks normally requiring human intelligence. Whether it copies the way human intelligence works is not clearly specified.

But, we will love to have it perform tasks that would require intelligence. So, what is involved if some AI system has to work, that requires intelligence for a human being to do the same work.

**(Refer Slide Time: 24:13)**

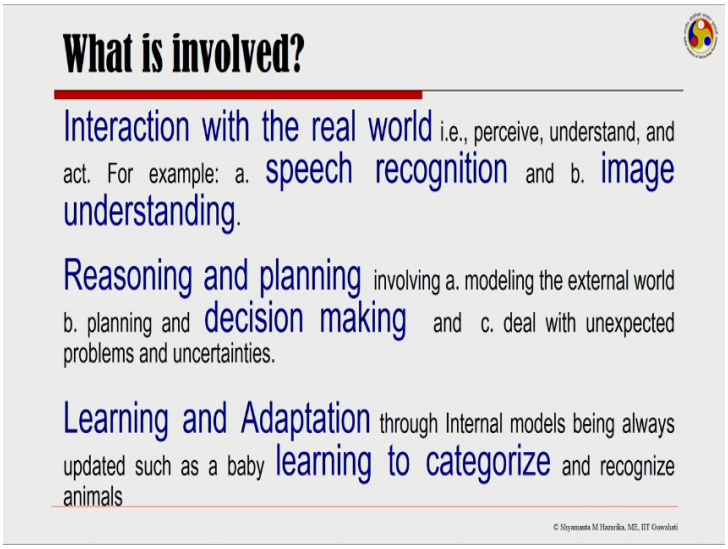


We would need interaction with the real world. That is, the system should be able to perceive, understand and act. It would involve things such as speech recognition, image understanding. The next thing would be reasoning and planning involving modeling the external world, planning and decision making and deal with unexpected problems and uncertainties. Thereafter, any intelligence system should be able to learn and adapt. That is, should have internal models which should be always updated. For example, a baby learning to categorize and recognize animals. Such a work would therefore involve a huge number of disciplines contributing and interacting with each other. One of the disciplines that has made large contribution to the growth of AI is philosophy;

**(Refer Slide Time: 25:15)**

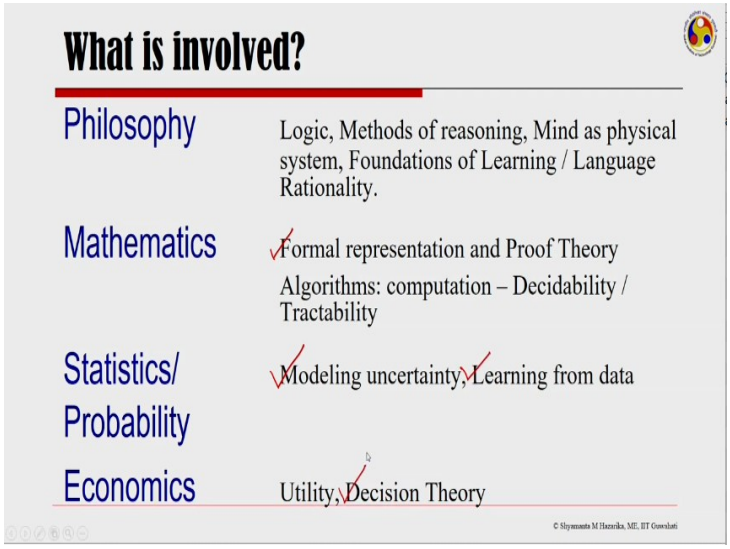


where we had looked at one of these ideas of logic, then they have contributed to method of reasoning, mind as a physical system to the foundations of learning and language rationality. The next area that has contributed hugely to the growth of AI is mathematics; involving formal representation and proof theory; algorithms, computational decidability, tractability. The other area that has contributed to the growth of AI is statistics and probability which involve modeling uncertainty and learning from data. Of course, economics has had a huge impact on AI in theories such as theories of utility and decision theory. In this course, we would definitely look at some formal representation. And we will look at modeling uncertainty and learning. We would of course look at decision theory.

**(Refer Slide Time: 26:21)**

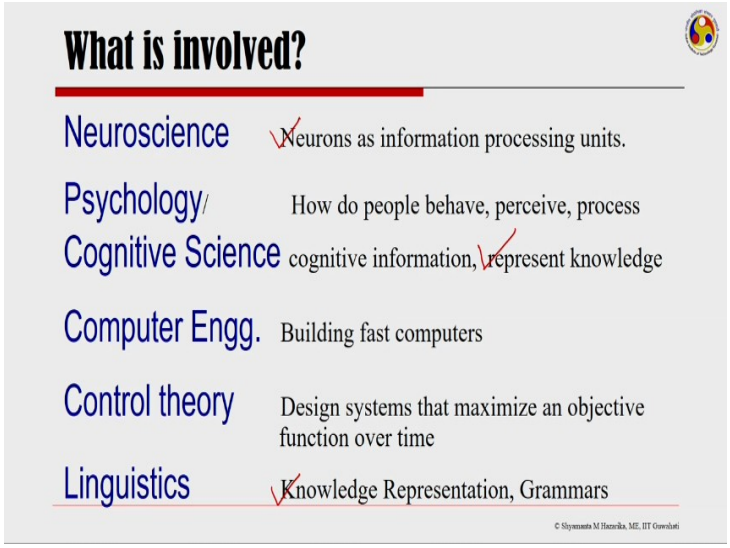

The other area that has influenced AI in a huge way is the growth of the area of neuroscience or neurons as information processing units. As we will see later, these developments in this area has in fact moved the course of AI in a huge way. Psychology and cognitive science, as I told you before, has made a huge impact in the area of AI. We look at how do people behave, perceive, process cognitive information, represent knowledge. Our understanding of these processes allow us to develop artificial intelligence systems that can do the same thing. In fact, the growth of cognitive science has also been pushed by the growth of AI, particularly machine learning.

Next area of study that has made AI possible and pushed the frontiers of AI is the area of computer engineering. Having fast computers has made a large impact.

And then, control theory that talks of designing systems that maximize an objective function over time has pushed the growth of AI. One area which have made a huge impact in the way artificial intelligence systems are developed over the years is the area of linguistics; particularly, knowledge representation and grammars. In this course, we would look at knowledge representation; we would look at neurons as information processing units. We would definitely do representation of knowledge in cognitive science and psychology.

**(Refer Slide Time: 28:18)**

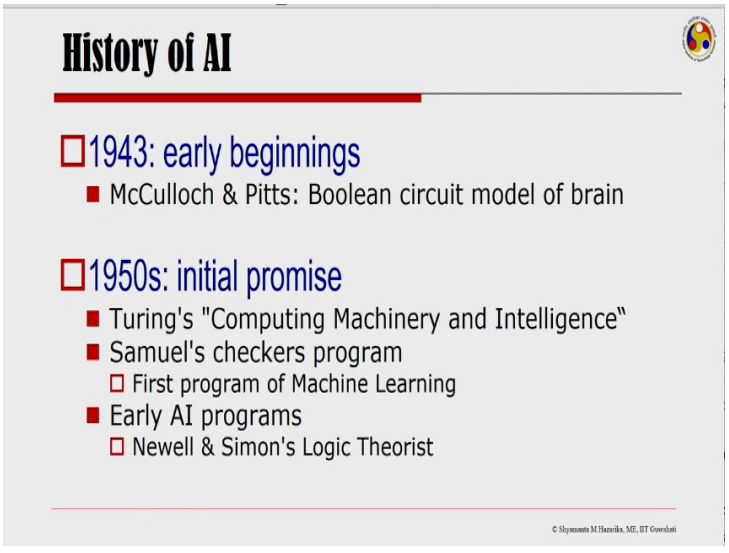


Let us now look at a brief history of AI and try to understand how the area has evolved over the years. As I told you in the previous slide, the earliest beginning of AI was in 1943 with the McCulloch and Pitts, Boolean circuit model of the brain, which is about the neural model.

In 1950s, we had the initial promise with of course, Turing's interesting and seminal paper on computing machineries and intelligence.

We have Samuel's checker program which we will talk of in our module on machine learning, which was the first program of machine learning. In fact, the Samuel's checker program running on an IBM machine, playing checkers with a human, could defeat a human and was very popular at that point of time.

We had very early AI programs like Newell and Simon's logic theorists. These were very initial promising enterprises in the area of AI. The enthusiasm continued for the decade 1955 to 1965.

**(Refer Slide Time: 29:40)**

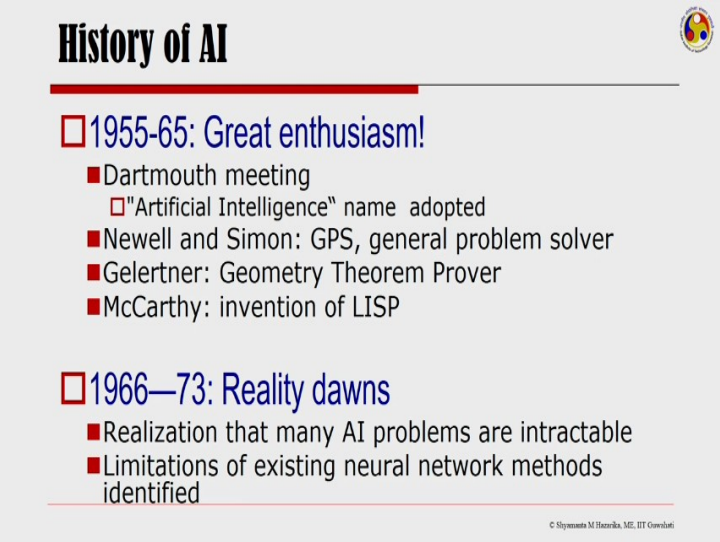


With the Dartmouth meeting came the term artificial intelligence. Very shortly thereafter, Newell and Simon came up with the General Problem Solver which we have discussed in one of the previous slides. Thereafter, there was the Geometry Theorem Prover and John McCarthy came up with the programing language called LISP. However, very soon the reality dawned. There was a realization that many artificially intelligent problems are intractable.

People soon realized the limitations of existing neural network methods. And in fact, people left neural networks as something that would not make any great impact.

**(Refer Slide Time: 30:34)**

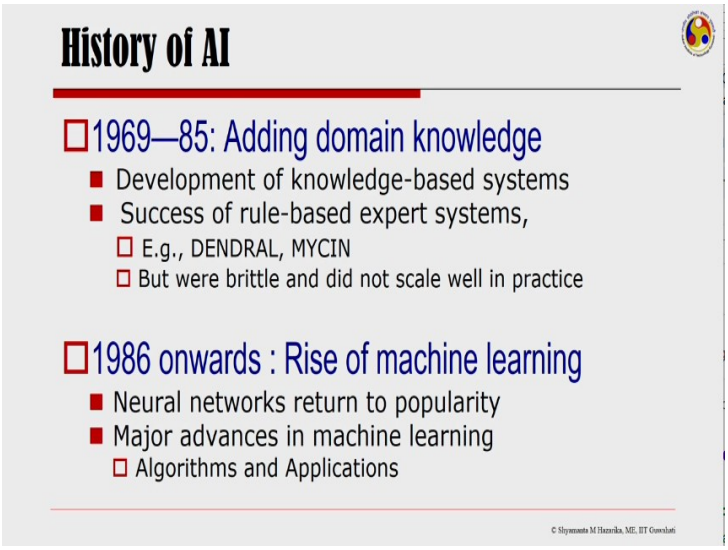


However, between 1969 and 1985, people were starting to realize that domain knowledge is something that could change the tilt. Development of knowledge based system was growing. And there are lot of success stories; success of rule based expert systems such as DENDRAL, MYCIN. But these were too brittle and they did not definitely scale well in practice. However, starting with 1986, we had the rise and rise of machine learning.

Neural networks again returned to popularity, albeit with lot of other modifications, which we will look at in our module on machine learning. There were mature advances in machine learning, both in terms of algorithms and applications.

**(Refer Slide Time: 31:36)**

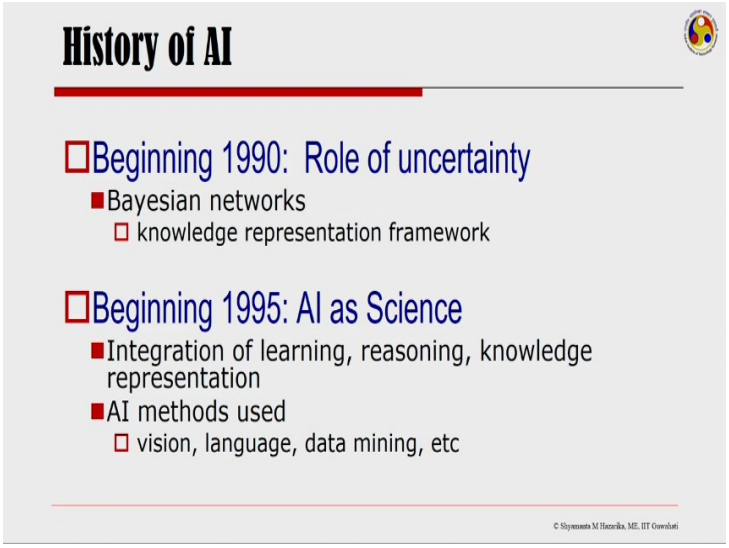


Beginning of 1990, people started to talk of uncertainty; the role of uncertainty in machine learning and in AI in general. They looked at Bayesian networks for knowledge

representation. 1995 is the turning point when artificial intelligence started to be accepted as science. During this period, there was integration of learning, reasoning, knowledge representation. AI methods started to be used in vision, language, data mining and lot of other real-life problems. If we now look at the history of AI; let us focus on this slide.

**(Refer Slide Time: 32:15)**

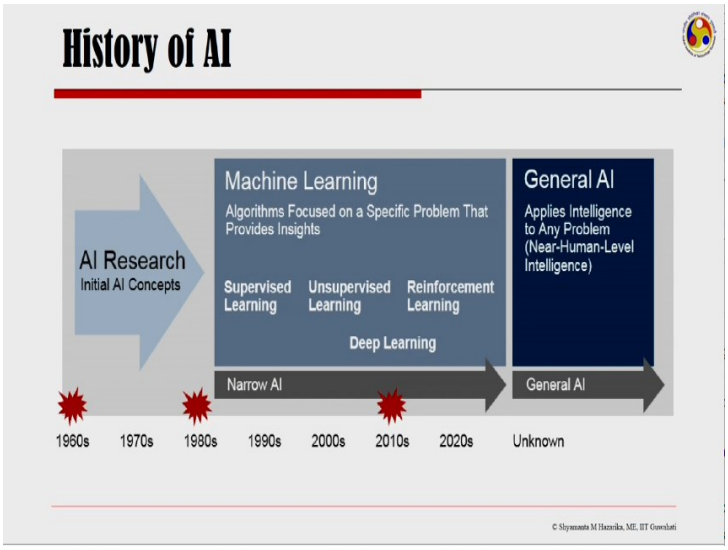


Up till 1980-85, we would see that AI research was involved in its initial euphoric concepts. Beyond 85, machine learning took over the area of AI roughly. This period of AI, we were no longer concerned with the general AI problem and can be categorized as narrow AI; AI with very focused attention - problems which are really really focused and do not talk of a general AI concept.

It is expected that as we make more progress in machine learning, we would be able to apply intelligence to any problem; near human level intelligence and would go back to our general AI problem. But it is very uncertain when that would happen.

**(Refer Slide Time: 33:11)**

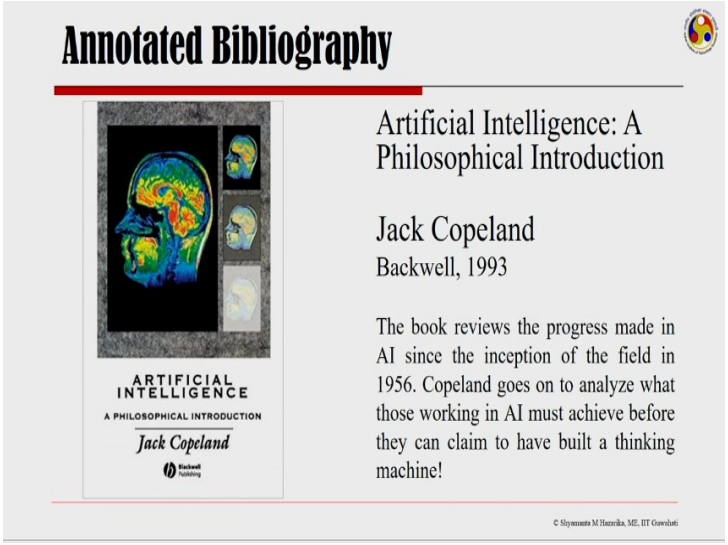


Here is a very important book that you need to pursue in this first module of the course. This is Artificial Intelligence: A Philosophical Introduction. This is by Jack Copeland. The book reviews the progress made in AI since the inception of the field in 1956. Copeland goes on to analyze what those working in AI must achieve before they can claim to have built a thinking machine.

**(Refer Slide Time: 33:47)**

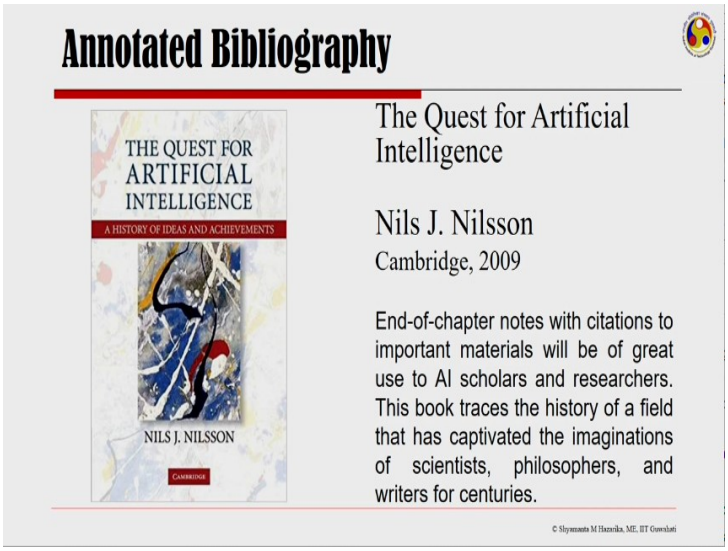


The next book is The Quest for Artificial Intelligence by Nils J. Nilsson. The most important thing on this book is the end of chapter notes, which has citations to important materials; which is of great use to AI scholars and researchers. In fact, the book traces the history of the field that has captivated the imagination of scientists, philosophers and writers for centuries. A third book that I would love to refer to for this module of the course is:

**(Refer Slide Time: 34:21)**

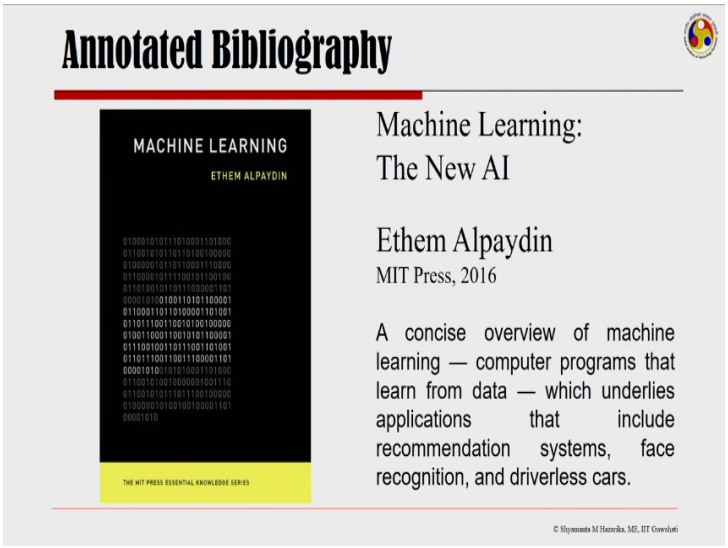


Machine Learning, the new AI. A concise overview of machine learning- computer programs that learn from data, in order to understand what lies in the new challenges for artificial intelligence. So, this was all about this lecture of the first module of the course, Introduction to AI. Thank you very much.