What is Artificial Intelligence?

Artificial Intelligence is a branch of *Science* which deals with helping machines to find solutions to complex problems in a more human-like fashion. This generally involves borrowing characteristics from human intelligence, and applying them as algorithms in a computer friendly way. A more or less flexible or efficient approach can be taken depending on the requirements established, which influences how artificial the intelligent behavior appears.

AI is generally associated with *Computer Science*, but it has many important links with other fields such as *Maths*, *Psychology*, *Cognition*, *Biology* and *Philosophy*, among many others. Our ability to combine knowledge from all these fields will ultimately benefit our progress in the quest of creating an intelligent artificial being.

Different definitions of AI are given by different books/writers. These definitions can be divided into two dimensions.

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

Top dimension is concerned with thought processes and reasoning, where as bottom dimension addresses the behavior.

The definition on the left measures the success in terms of fidelity of human performance, whereas definitions on the right measure an ideal concept of intelligence, which is called rationality.

Human-centered approaches must be an empirical science, involving hypothesis and experimental confirmation. A rationalist approach involves a combination of mathematics and engineering.

Acting Humanly: The Turing Test Approach

The Turing test, proposed by Alan Turing (1950) was designed to convince the people that whether a particular machine can think or not. He suggested a test based on indistinguishability from undeniably intelligent entities- human beings.

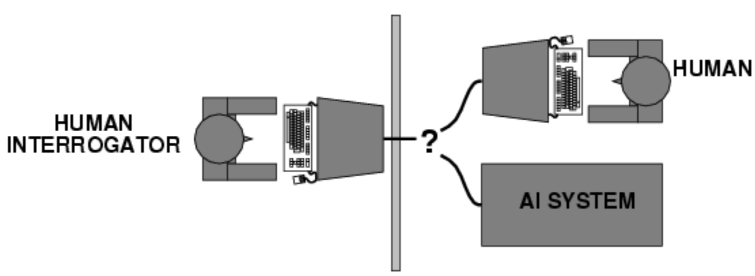

The computer passes the test if a human interrogator after posing some written questions, can not tell whether the written response come from human or not.

To pass a Turing test, a computer must have following capabilities:

- Natural Language Processing: Must be able to communicate in English successfully
- Knowledge representation: To store what it knows and hears.
- Automated reasoning: Answer the Questions based on the stored information.
- Machine learning: Must be able to adapt in new circumstances.

Turing test avoid the physical interaction with human interrogator. Physical simulation of human beings is not necessary for testing the intelligence.

The total Turing test includes video signals and manipulation capability so that the interrogator can test the subject's perceptual abilities and object manipulation ability. To pass the total Turing test computer must have following additional capabilities:

- Computer Vision: To perceive objects
- Robotics: To manipulate objects and move

Thinking Humanly: Cognitive modeling approach

Make the machines with mind. Cognition means the action or process of acquiring knowledge and understanding through thought, experience and senses. To make a machine that think like human brain, scientific theories of internal brain activities (cognitive model) are required. Two ways of doing this is:

- Predicting and testing human behavior (cognitive science)
- Identification from neurological data (Cognitive neuroscience)

Once we have precise theory of mind, it is possible to express the theory as a computer program. But unfortunately until up to now there is no precise theory about thinking process of human brain. Therefore it is not possible to make the machines that think like human brain

Thinking rationally: The laws of thought approach

Aristotle was one of the first who attempt to codify the right thinking that is irrefutable reasoning process. He gave Syllogisms that always yielded correct conclusion when correct premises are given.

For example:

Ram is a man

Man is mortal

⇨ Ram is mortal

Let
p(x) → x is man
q(x) → x is mortal

then above statement can be written as

p(x) => q(x)	Man is mortal
p(Ram)	Ram is man

Then from modus ponens q(x) ia also true. That is Ram is mortal

This study initiated the field of logic. The logicist tradition in AI hopes to create intelligent systems using logic programming.

Problems:
It is not easy to take informal knowledge and state in the formal terms required by logical notation, particularly when knowledge is not 100% certain.

Solving problem principally is different from doing it in practice. Even problems with certain dozens of fact may exhaust the computational resources of any computer unless it has some guidance as which reasoning step to try first.

Acting Rationally: The rational Agent approach:

Agent is something that acts. Computer agent is expected to have following attributes:

- ❖ Autonomous control
- ❖ Perceiving their environment
- ❖ Persisting over a prolonged period of time
- ❖ Adapting to change
- ❖ And capable of taking on another's goal

Rational behavior means doing the right thing. The right thing is that which is expected to maximize goal achievement, given the available information. Rational Agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome. In this approach the emphasis is given to correct inferences.

One way to act rationally is to reason logically to the conclusion and act on that conclusion. On the other hand there are also some ways of acting rationally that can not be said to involve inference. For Example, recoiling from a host stove is a reflex action that is usually more successful than a slower action taken after careful deliberation.

Advantages:

- ✓ It is more general than laws of thought approach, because correct inference is just one of several mechanisms for achieving rationality.

- ✓ It is more amenable to scientific development than are approaches based on human behavior or human thought because the standard of rationality is clearly defined and completely general.

Applications of AI

Game playing

You can buy machines that can play master level chess for a few hundred dollars. There is some AI in them, but they play well against people mainly through brute force computation--looking at hundreds of thousands of positions. To beat a world champion by brute force and known reliable heuristics requires being able to look at 200 million positions per second.

Speech recognition

In the 1990s, computer speech recognition reached a practical level for limited purposes. Thus United Airlines has replaced its keyboard tree for flight information by a system using speech recognition of flight numbers and city names. It is quite convenient. On the the other hand, while it is possible to instruct some computers using speech, most users have gone back to the keyboard and the mouse as still more convenient.

Understanding natural language

Just getting a sequence of words into a computer is not enough. Parsing sentences is not enough either. The computer has to be provided with an understanding of the domain the text is about, and this is presently possible only for very limited domains.

Computer vision

The world is composed of three-dimensional objects, but the inputs to the human eye and computers' TV cameras are two dimensional. Some useful programs can work solely in two dimensions, but full computer vision requires partial three-dimensional information that is not just a set of two-dimensional views. At present there are only limited ways of representing three-dimensional information directly, and they are not as good as what humans evidently use.

Expert systems

A ``knowledge engineer'' interviews experts in a certain domain and tries to embody their knowledge in a computer program for carrying out some task. How well this works depends on whether the intellectual mechanisms required for the task are within the present state of AI. When this turned out not to be so, there were many disappointing results. One of the first expert systems was MYCIN in 1974, which diagnosed bacterial infections of the blood and suggested treatments. It did better than medical students or practicing doctors, provided its limitations were observed. Namely, its ontology included bacteria, symptoms, and treatments and did not include patients, doctors, hospitals, death, recovery, and events occurring in time. Its interactions depended on a single patient being considered. Since the experts consulted by the knowledge engineers knew about patients, doctors, death, recovery, etc., it is clear that the knowledge engineers forced what the experts told them into a predetermined framework. In the present state of AI, this has to be true. The usefulness of current expert systems depends on their users having common sense.

Heuristic classification

One of the most feasible kinds of expert system given the present knowledge of AI is to put some information in one of a fixed set of categories using several sources of information. An example is advising whether to accept a proposed credit card purchase. Information is available about the owner of the credit card, his record of payment and also about the item he is buying and about the establishment from which he is buying it (e.g., about whether there have been previous credit card frauds at this establishment).

Foundations of AI

Different fields have contributed to AI in the form of ideas, viewpoints and techniques.

***Philosophy*:**
Logic, reasoning, mind as a physical system, foundations of learning, language and rationality.

***Mathematics*:**
Formal representation and proof algorithms, computation, undecidability, intractability, probability.

***Psychology*:**
adaptation, phenomena of perception and motor control.

***Economics*:**
formal theory of rational decisions, game theory.

***Linguistics*:**
Knowledge representation, grammar

***Neuroscience*:**
Physical substrate for mental activities

***Control theory*:**
Homeostatic systems, stability, optimal agent design

Brief history of AI

What happened after WWII?

- 1943: Warren Mc Culloch and Walter Pitts: a model of artificial boolean neurons to perform computations.
  - First steps toward connectionist computation and learning (Hebbian learning).
  - Marvin Minsky and Dann Edmonds (1951) constructed the first neural network computer

- 1950: Alan Turing's "Computing Machinery and Intelligence"
  - First complete vision of AI.

The birth of AI (1956)

- Dartmouth Workshop bringing together top minds on automata theory, neural nets and the study of intelligence.
  - Allen Newell and Herbert Simon: The logic theorist (first nonnumeric thinking program used for theorem proving)
  - For the next 20 years the field was dominated by these participants.

Great expectations (1952-1969)

- Newell and Simon introduced the General Problem Solver.
  - Imitation of human problem-solving
- Arthur Samuel (1952-)investigated game playing (checkers) with great success.
- John McCarthy(1958-) :
  - Inventor of Lisp (second-oldest high-level language)
  - Logic oriented, Advice Taker (separation between knowledge and reasoning)

- Marvin Minsky (1958 -)
  - Introduction of microworlds that appear to require intelligence to solve: e.g. blocks-world.
  - Anti-logic orientation, society of the mind.

Collapse in AI research (1966 - 1973)

- Progress was slower than expected.
  - Unrealistic predictions.
- Some systems lacked scalability.
  - Combinatorial explosion in search.
- Fundamental limitations on techniques and representations.
  - Minsky and Papert (1969) Perceptrons.

AI revival through knowledge-based systems (1969-1970)

- General-purpose vs. domain specific
  - E.g. the DENDRAL project (Buchanan et al. 1969)
    First successful knowledge intensive system.

- Expert systems
    - MYCIN to diagnose blood infections (Feigenbaum et al.)
        - Introduction of uncertainty in reasoning.

- Increase in knowledge representation research.
    - Logic, frames, semantic nets, …

AI becomes an industry (1980 - present)
- R1 at DEC (McDermott, 1982)
- Fifth generation project in Japan (1981)
- American response …

Puts an end to the AI winter.

Connectionist revival (1986 - present): (Return of Neural Network)
- Parallel distributed processing (RumelHart and McClelland, 1986); backprop.

AI becomes a science (1987 - present)

- In speech recognition: hidden markov models
- In neural networks
- In uncertain reasoning and expert systems: Bayesian network formalism

The emergence of intelligent agents (1995 - present)
- The whole agent problem: "How does an agent act/behave embedded in real environments with continuous sensory inputs"

Knowledge and its importance

Knowledge is information that helps us to solve problems in particular domain, to predict what will happen next and to explain why and how something has happened? Knowledge contains facts, procedures and judgmental rules. Knowledge is important in AI for making intelligent machines. Key issues confronting the designer of AI system are:

Knowledge acquisition: Gathering the knowledge from the problem domain to solve the AI problem.

Knowledge representation: Expressing the identified knowledge into some knowledge representation language such as propositional logic, predicate logic etc.

Knowledge manipulation: Large volume of knowledge has no meaning until up to it is processed to deduce the hidden aspects of it. Knowledge is manipulated to draw conclusions from knowledgebase.