

AI: Artificial Intelligence

Computer Literacy 1 Lecture 26
18/11/2008



Topics

- One definition of many
- Alan Turing and a bit of background history
- The beginning of AI and early problems
- Examples
- AI, a fragmented field
- Logic in AI, and program languages
- Problems of modern AI



AI: One definition

- *The capability of a device to perform functions that are normally associated with human intelligence, such as reasoning and optimization through experience.*



Alan Turing

- Alan Turing was the first to propose the possibility of Artificial Intelligence in 1950
- He invented the Turing test
- Is also considered to be the father of modern computer science



Turing Test



- Turing proposed an experiment that is now known as **Turing Test**
- The Turing Test is seen as an attempt to define a standard for a machine to be called “intelligent”
- If a computer can make a human think that she/he is interacting with another human, then the computer could be said to “think”

Chinese Room



- John Searle proposed in his 1980 paper *Minds, Brains and Programs* an argument against the Turing Test
- This argument is known as the Chinese room thought experiment:
 - Software could pass the Turing Test simply by manipulating symbols of which they had no understanding. Without any understanding they could not be described as “thinking” in the same sense as people do. Hence Turing’s test cannot prove that a machine thinks!

A(I) little bit of history



- AI started as early as the antiques
- In this class we will start with history from the 1950s
- You’ve already heard about Turing
- W. Pitts and W. McCulloch showed how networks of idealised artificial neurons might perform simple logical functions
- M. Minsky built the first neural net machine. He was one of the most versatile and important innovators in AI

Official Birth of AI



- Dartmouth conference 1956
 - Organised by Minsky, McCarthy and two senior scientists of IBM
 - This conference is seen as the birth of AI because here it gained its name, mission and the most important people to work on it for the next years
- “Every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it” (assertion of the conference proposal)

Problems in early years of AI



- Limited computer power
 - AI still needs a lot of memory
- Commonsense knowledge and reasoning
 - A computer must have at least a concept of vision or speech, like a child. A child already has a vast amount of information. In the early time there was no way to build a data base big enough to hold all this information
- Moravec's paradox
 - Solving difficult mathematical problems proves to be rather easy for a computer, while face recognition and walking is not

Deep Blue



- In 1997, Deep Blue became the first chess playing computer to beat a reigning chess champion
- Deep Blue's computer was 10 million times faster than Ferranti Marks I in 1951
 - This is measured by Moore's law, which states that the speed and memory capacity of computers doubles each year

Deep Blue insides



- Strictly spoken Deep Blue is not an AI he's more a supercomputer
 - He was built by IBM
 - Deep Blue, like other supercomputers, is using a form of UNIX as OS
 - Speed of processing and memory access is more important than compatibility
 - Supercomputers are forecasting the weather, model finance data and do other very complex calculation
- <http://www.research.ibm.com/deepblue/>

Kismet



- Kismet is a robot at the MIT
- He was made to socially interact with humans and demonstrate simulated human emotion as well as appearance
- He only has a head with auditory, visual and expressive systems

Kismet



- <http://www.ai.mit.edu/projects/sociable/videos.html>

Asimo



- Asimo was invented by Honda
- He's a so called humanoid robot
- He can recognise moving objects, postures and gestures, his environment and he can also recognise faces
- He is able to distinguish sounds and can go online
- <http://asimo.honda.com/>

RI-MAN



- http://www.bmc.riken.jp/~RI-MAN/index_us.html

AI: A fragmented field



- We can find different approaches to AI since it is still a young research area and is also a very wide area:
 - Cybernetics and brain simulation (connection between neurology and cybernetics)
 - Traditional symbolic AI (like vision and language recognition)
 - Sub-symbolic AI (higher thinking processes)
 - Intelligent agent paradigm (system that perceives its environment and takes action which maximises its chances of success)

Logic and other methods



- Logic programming - resolution and unification algorithm for logical deduction (1963)
- Logic is used for knowledge representation and problem solving
- Probabilistic methods for uncertain reasoning - Bayesian networks, Hidden Markov Models, etc...
- Search algorithms for search and optimization

Program languages



- IPL supports programs that could perform general problem solving
- Lisp, mathematical notation for programs based on lambda calculus
- Prolog, declarative language where programs are expressed in terms or relation

AI: Evaluation



- Categories
 - Optimal, not possible to perform better
 - Strong super-human, performs better than all humans
 - Super-human, performs better than most humans
 - Sub-human, performs worse than most humans

Problems of modern AI



- Problem solving
- Learning
- Knowledge representation
- Planning
- Motion
- Natural Language processing ...

The 3 Laws of Robotics



1. A robot may not injure a human being, or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given it by human beings except where such orders conflict with first law.
3. A robot must protect its own existence as long as such protection does not conflict with the first or second law.

Quoted from *I, Robot*, I. Asimov, 1950

Dark Star (John Carpenter, 1974)



- Philosophical bomb

http://www.youtube.com/watch?v=g_47mmt5SZY&feature=related

Key points



- The Turing test and the Chinese room that contradicts it
- Early problems in AI and current problems
- What can Kismet, Asimo and Ri-Man actually do
- The different fields of AI
- AI's program languages
- And it's evaluation scale