

Techniques of Artificial Intelligence

Geraint A. Wiggins
Professor of Computational Creativity
Department of Computer Science
Vrije Universiteit Brussel

- Topic 1: Introduction: Philosophy of Artificial Intelligence
 - ▶ What is AI?
 - Approaches, Foundations, History
 - Applications and the future
- Topic 2: Rational Agents
 - ▶ Computational entities that exhibit intelligent behaviour
- Topics 3-7: Knowledge Representation and Problem Solving
 - ▶ Formal Logic as a representation language
 - ▶ Representing problems and solving them
 - ▶ Building programs that can solve problems
 - ▶ Automated reasoning
 - ▶ Reasoning with Uncertainty

Module Overview (1st part)

- Weekly learning activity
 - ▶ Assisted theoretical work in labs

- Learning and Teaching

- ▶ Lectures:

- 2 hours of lectures each week

- ▶ Lab work (unassessed):

- 1.5-2 hours of supporting practical material in lab sessions each week

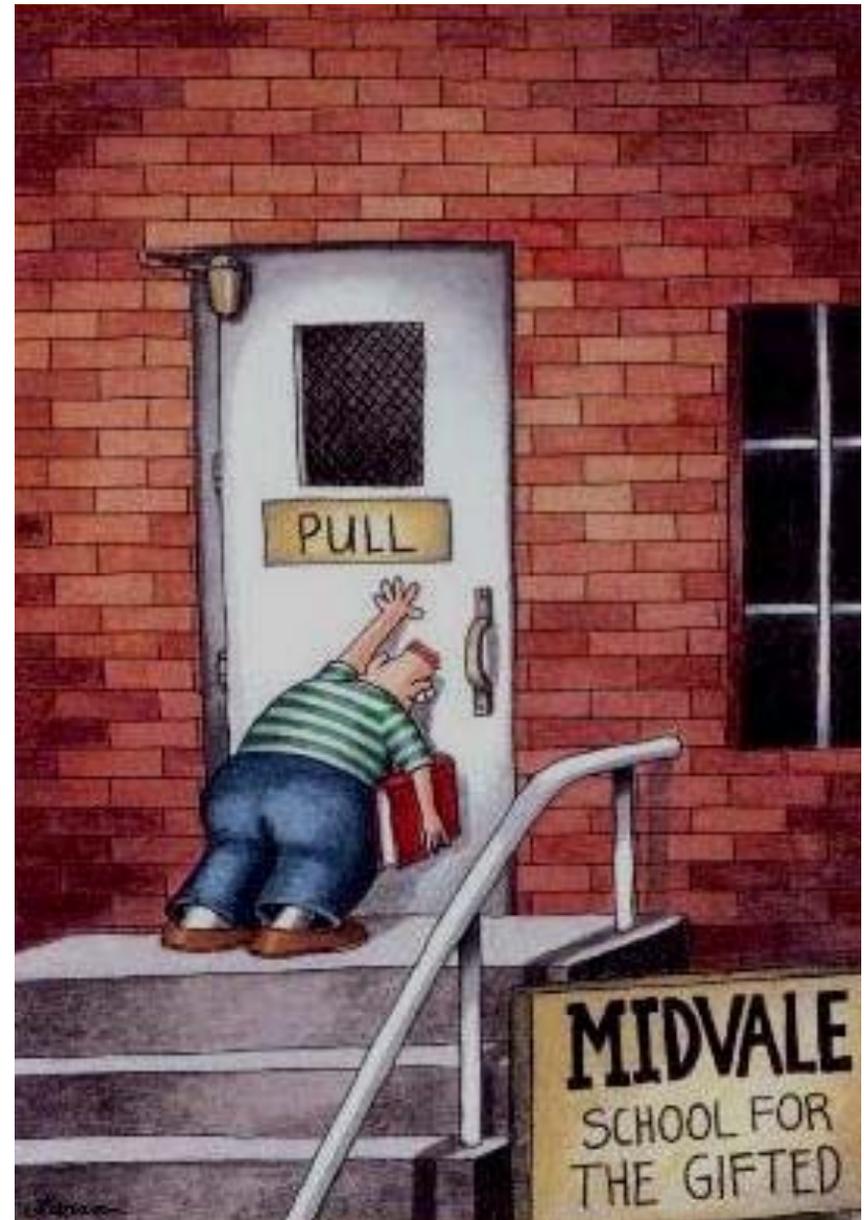
- ▶ Getting help

- TA contact in labs: Roxana Rădulescu
- Asking me
- Questions during and after lectures

- Stuart Russell and Peter Norvig
Artificial Intelligence: A Modern Approach
Prentice Hall, 2003+
 - ▶ Various editions; doesn't matter which for this module
- George F. Luger
Artificial Intelligence: Structures and Strategies for Complex Problem Solving
5th edition, Addison Wesley, 2005
 - ▶ Also available in earlier editions as Luger and Stubblefield
- Robert Kowalski
Logic for Problem Solving
Free on-line book, will be made available on line

What is Intelligence?

- What is meant by human intelligence?
 - ▶ exams, education...perform well in intelligence tests
 - ▶ skills – languages, chess, mathematical reasoning
 - ▶ cognitive, conative, emotional/affective intelligence
 - ▶ speed, efficiency and creativity of information processing



What is Intelligence?

artificial

- What might make computers smart...clever...intelligent... intellectual?
 - ▶ Beat human at chess ✓ *IBM's Deep Blue beat Gary Kasparov in 1997*
 - ▶ Discover mineral fields ✓ *PROSPECTOR closely matched human experts in 1978*
 - ▶ Diagnose disease ✓ *MYCIN diagnosed 68% of cases acceptably in 1975*
 - ▶ Hold a convincing conversation, social interaction ✗ *Turing "test"*

What is Artificial Intelligence?

- Marvin Minsky
 - ▶ "The science of making machines do things that would require intelligence if done by men."
- Eugene Charniak
 - ▶ "The study of mental faculties through the use of computational models."
- John L. Gordon
 - ▶ "The aim of Artificial Intelligence is to create intelligent machines and through this, to understand the principles of intelligence. At the moment, we can settle for creating less stupid machines."

Two kinds of Artificial Intelligence

- Weak AI

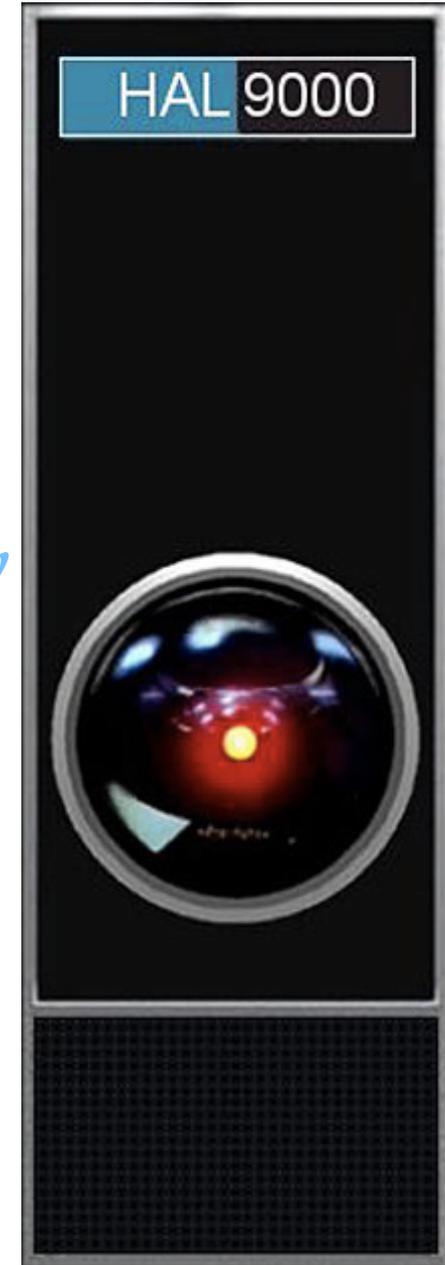
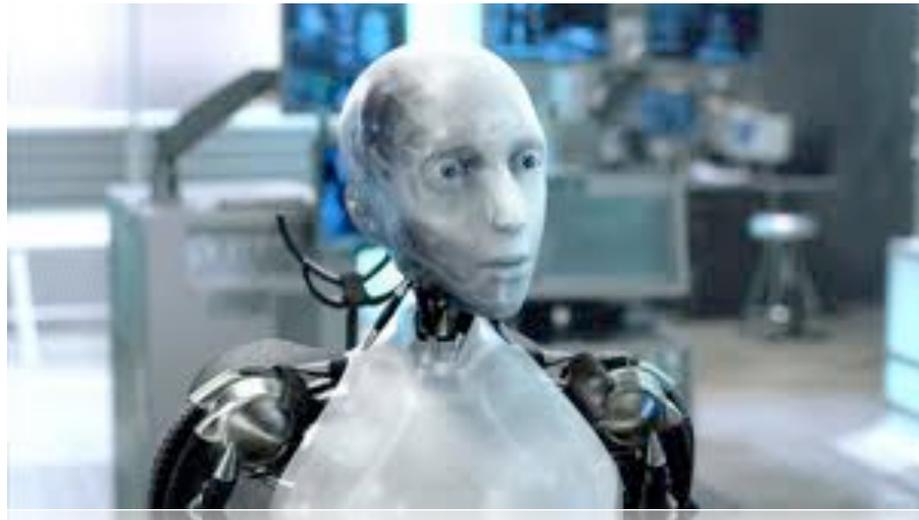
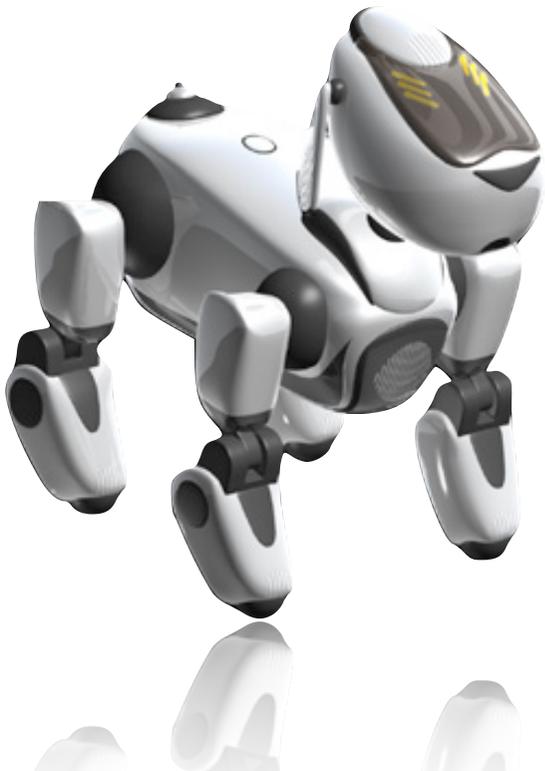
- ▶ Machines can be made to act *as though* they were intelligent

- Strong AI

- ▶ Machines that act intelligently and have real, *conscious* minds

real

fictional



Four categories of AI

**Thinking
Humanly**

“The exciting new effort to make computers think ... machines with minds, in the full and literal sense”
(Haugeland, 1985)

“[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning ...”
(Bellman, 1978)

“The art of creating machines that perform functions that require intelligence when performed by people”
(Kurzweil, 1990)

**Acting
Humanly**

“The study of how to make computers do things at which, at the moment, people are better”
(Rich and Knight, 1991)

**Thinking
Rationally**

“The study of mental faculties through the use of computational models” (Charniak and McDermott, 1985)

“The study of the computations that make it possible to perceive, reason, and act”
(Winston, 1992)

“A field of study that seeks to explain and emulate intelligent behaviour in terms of computational processes”
(Schalkoff, 1990)

**Acting
Rationally**

“The branch of computer science that is concerned with the automation of intelligent behavior”
(Luger and Stubblefield, 1993)

**Thinking
Humanly**

**Thinking
Rationally**

**Acting
Humanly**

**Acting
Rationally**

Thinking
Humanly

Thinking
Rationally

Acting
Humanly

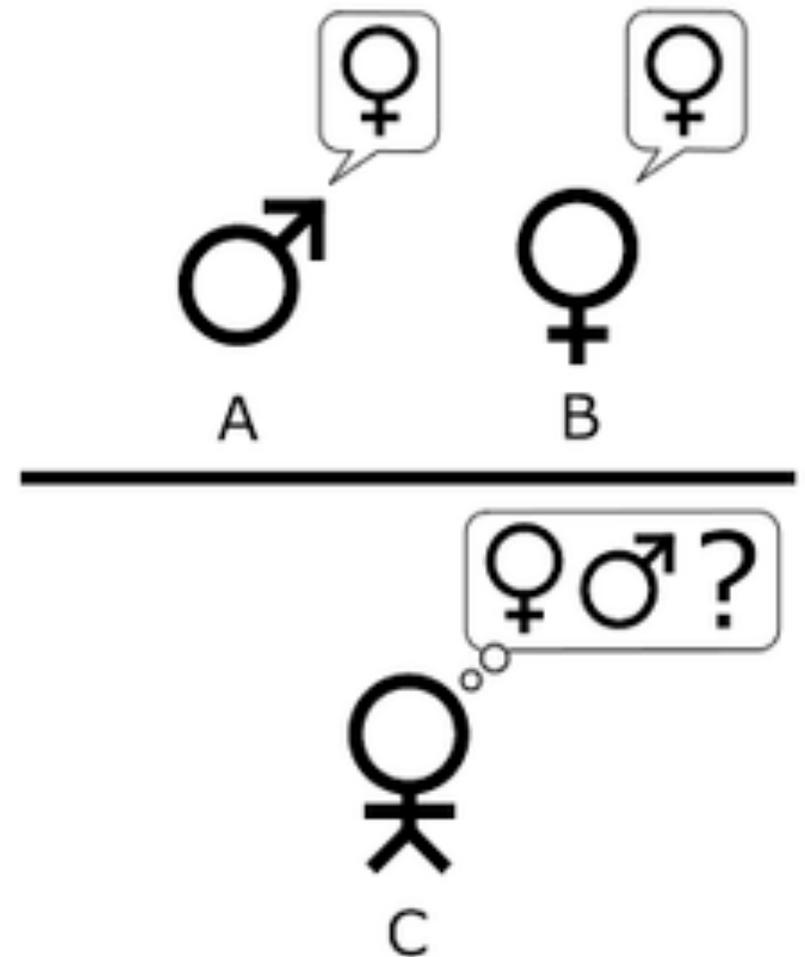
Acting
Rationally

- In 1950, Alan Turing's paper "Computing machinery and intelligence"
 - ▶ Instead of asking if machines can think, ask whether can pass a behavioural intelligence test
- Links:
 - ▶ www.turing.org.uk
 - ▶ www.alanturing.net



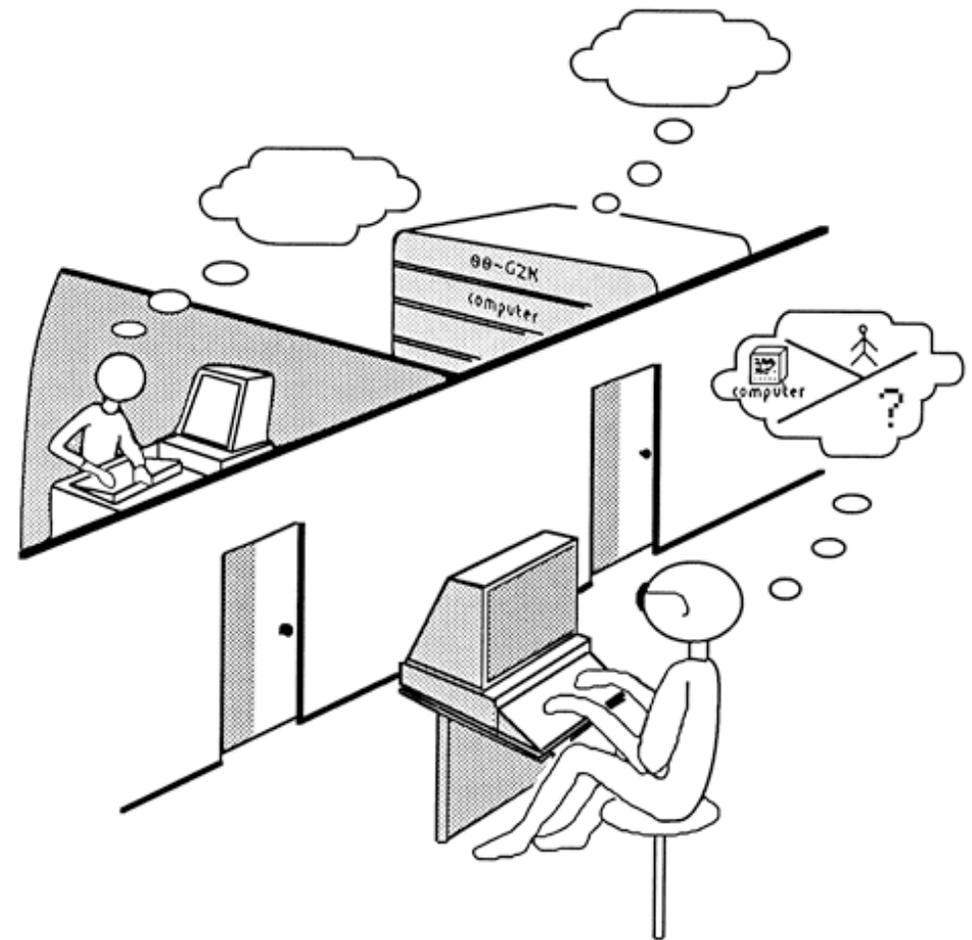
The Imitation Game

- Turing asked the question “Can machines think?”
 - ▶ but because “think” is an ill-defined term, he proposed a thought experiment instead:
The Imitation Game
 - ▶ based on Victorian parlour game where interrogator guesses gender of unseen person



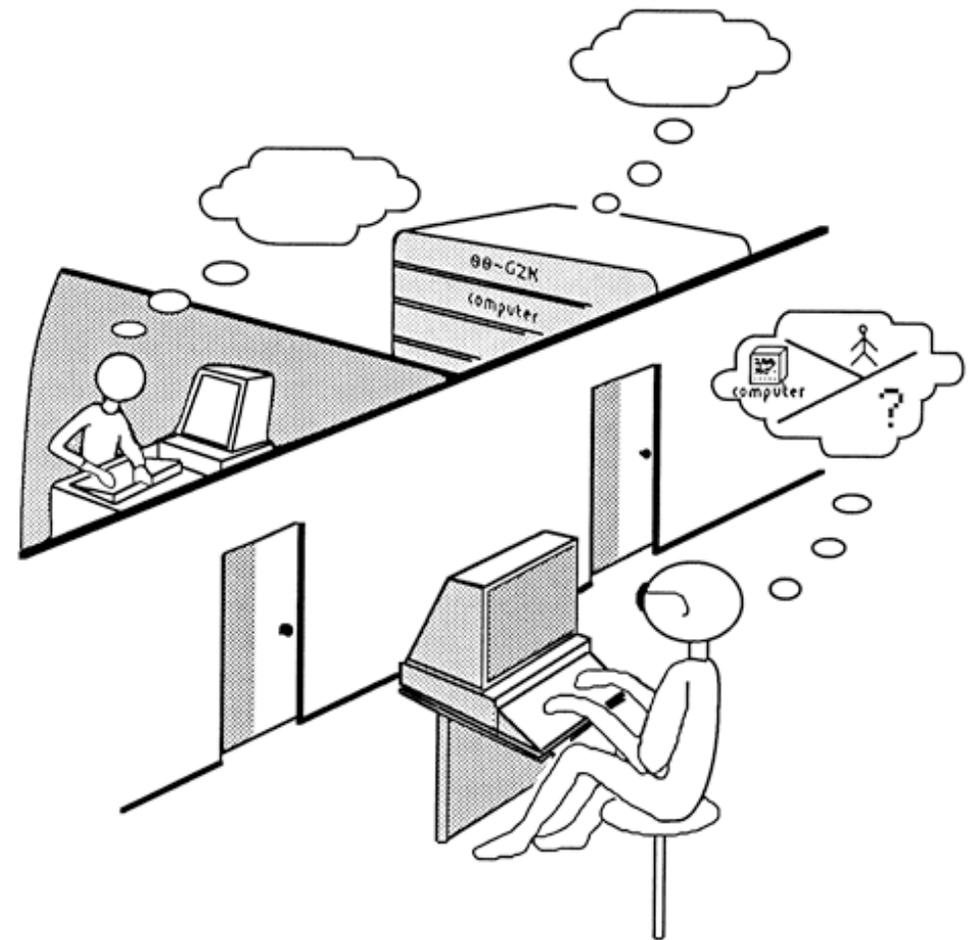
The Imitation Game

- Turing asked the question “Can machines think?”
 - ▶ but because “think” is an ill-defined term, he proposed a thought experiment instead:
The Imitation Game
 - ▶ based on Victorian parlour game where interrogator guesses gender of unseen person
- In Turing’s Game, human interrogator must decide if unseen chat partner is a human or a computer



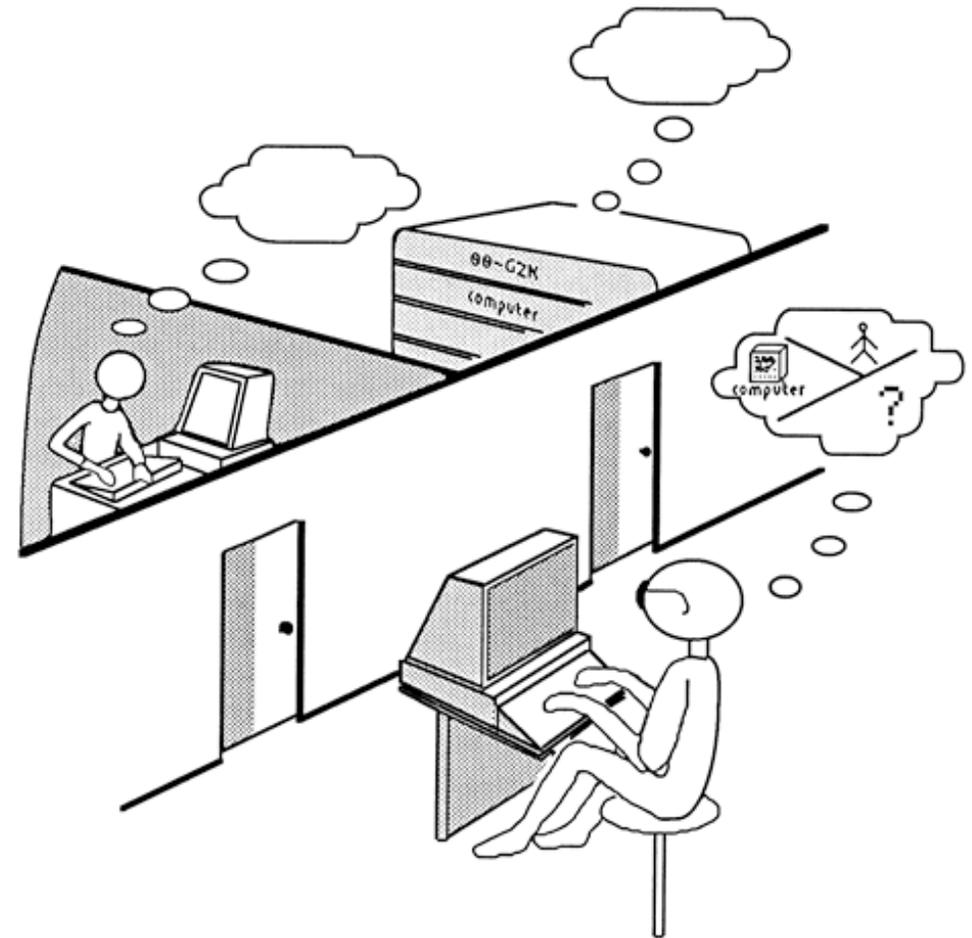
The Imitation Game

- To convince the interrogator of its intelligence, the computer must
 - ▶ understand and generate language
 - ▶ know about the world
 - ▶ reason about the world
 - ▶ learn about the dialogue and the interrogator
 - ▶ combine all this knowledge and reasoning instantaneously



The Imitation Game

- Turing's paper is usually interpreted as though he is proposing a test for intelligence
 - ▶ The Turing Test
- However, he doesn't actually say that in his paper
 - ▶ Another interpretation is that this is a thought experiment intended to demonstrate that you can't answer the question "is a machine thinking?"
- Loebner prize for first success
 - ◉ [See Guardian article for more info](#)



**Thinking
Humanly**

**Thinking
Rationally**

**Acting
Humanly**

**Acting
Rationally**

Thinking
Humanly

Thinking
Rationally

Acting
Humanly

Acting
Rationally

- To build a machine that thinks like a human we must first know how we think, e.g.,
 - ▶ through introspection (thinking about how we do it – unscientific)
 - ▶ through psychological experiments (proper science)
- A precise theory of mind can be expressed as a computer program:
 - ▶ similar I/O and timing behaviour to humans is evidence that similar mechanisms to the program are operating in humans
 - ▶ compare trace of reasoning steps of program to that of humans, e.g., the “General Problem Solver” (Newell and Simon, 1961)

- Cognitive science unites

- ▶ computer models from AI
- ▶ experimental techniques from psychology

to try to construct precise and testable theories of how the mind works.

- Now a distinct subject from AI

- ▶ Indeed, what once was a unified subject is now a collection of specialised sub-fields, e.g.
 - machine vision
 - computational creativity
 - automated reasoning
 - etc.

**Thinking
Humanly**

**Thinking
Rationally**

**Acting
Humanly**

**Acting
Rationally**

Thinking
Humanly

Thinking
Rationally

Acting
Humanly

Acting
Rationally

- Aristotle's syllogisms provided patterns for argument structure that always gave correct conclusions from correct premises, e.g.,

Socrates is a man

All men are mortal

∴ Socrates is mortal

- Precursor of formal logic that provides a precise notation for statements about things in the world and relations between them.

- By 1965, programs existed that could solve soluble logical problems, given enough time
- Logician tradition in AI hopes to build on such programs to create intelligent systems
 - ▶ logician emphasis in this module
- Emphasis is on correct inferences. However,
 - ▶ formalising knowledge into logic is not easy, especially if that knowledge is uncertain
 - ▶ soluble in principle but, in practice, time and memory are limited
 - ▶ some problems are *undecidable*, or only *semi-decidable*, and can only be solved by *search*

Thinking and Acting Rationally

**Thinking
Humanly**

**Thinking
Rationally**

**Acting
Humanly**

**Acting
Rationally**

Thinking
Humanly

Thinking
Rationally

Acting
Humanly

Acting
Rationally

- To act rationally means
 - ▶ to act so as to achieve one's goals, given one's beliefs
- An agent is something that *acts* (i.e. does things) in the world
 - ▶ to do so usefully it also needs to *perceive* the world
- AI can be viewed as the study and construction of rational agents
 - ▶ many connections to economics and control theory
- What is the impact of self-consciousness, emotions, desires, love of music, fear of dying, etc ... on human intelligence?

Primates expect rationality in primates

- When trying to understand someone's intentions, non-human primates expect others to act rationally by performing the most appropriate action allowed by the environment

Rhesus macaques were one of three types of primate that responded to an experimenter's actions



The rationality experiment

- Primates were presented with two potential food containers
- Two scenarios: an experimenter
 - ▶ touched a container with his elbow when his hands were full
 - ▶ touched a container with his elbow when his hands were empty
- The primates looked for the food in the container indicated with the elbow more often when the experimenter's hands were full
- The monkeys considered, just as a human being would, that if someone's hands are full then it is rational for them to use their elbow to indicate the container with food, whereas if their hands are empty it is not rational for them to use their elbow, because they could have used their unoccupied hand.

- Extends laws of thought approach:
 - ▶ one way to act rationally is to reason logically that a particular action will achieve one's goals and then to perform it
- But acting rationally goes further:
 - ▶ when no action is provably correct but some action must be taken
 - ▶ logical reasoning is unnecessary (reflex actions, e.g., pulling hand away from hot stove)
 - ▶ not all cognitive skills are purely logical (e.g., visual perception, language comprehension)

- Cognitive skills needed in Turing's Imitation game allow rational actions
 - ▶ communicating effectively helps an agent get by in a complex society
 - ▶ representing knowledge and reasoning with it enables good decision-making in a wide variety of situations
 - ▶ learning about how the world works enables agents to generate more efficient strategies for dealing with it
- This approach sums up AI, because
 - ▶ it generalises thinking rationally, which is useful but not always necessary
 - ▶ it is scientifically easier than reproducing human thought or human behaviour
 - ▶ it equates AI with rational behaviour (clearly defined), rather than human intelligence (much harder to define)
- This module focuses mostly on rational agent capabilities

- Philosophy
 - ▶ logic, methods of reasoning, mind as physical system, foundations of learning, language, rationality
- Mathematics
 - ▶ formal representation and proof, algorithms, computation, decidability, tractability, probability
- Psychology
 - ▶ adaptation, phenomena of perception and motor control, experimental techniques (psychophysics, etc.)
- Linguistics
 - ▶ knowledge representation, grammar
- Neuroscience
 - ▶ physical substrate for mental activity
- Control theory
 - ▶ stability, simple optimal agent designs, homeostatic* systems

* **Homeostasis:** *The ability or tendency of an organism or cell to maintain internal equilibrium by adjusting its physiological processes.*

- McCulloch and Pitts (1943)

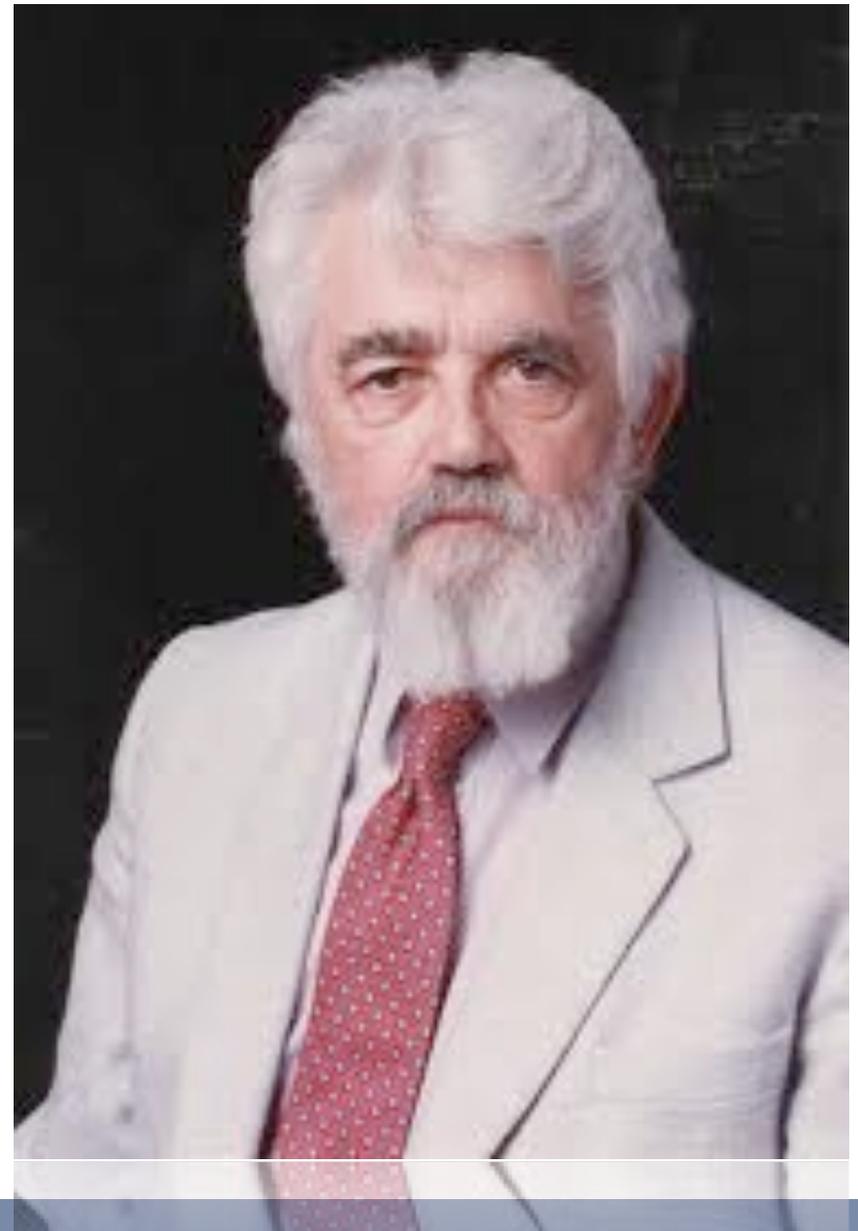
- ▶ proposed model of artificial neurons characterised as either on or off; switched on in response to stimulation of sufficient number of neighbouring neurons
- ▶ showed that any computable function is computable by some network of connected neurons; logical connectives implemented by simple net structures

- Early 1950s

- ▶ Christopher Strachey wrote first working AI program for Ferranti Mark I computer: draughts
- ▶ Claude Shannon & Alan Turing wrote chess programs (1950ish)
- ▶ Marvin Minsky & Dean Edmonds built first neural network (SNARC) (1951)

- 1952 Arthur Samuel's checkers (draughts) program
 - ▶ wrote a program to play checkers that reached tournament level
 - ▶ disproved idea that computers could not do more than they are told since the program soon beat its creator
- 1956 conference at Dartmouth College, New Hampshire
 - ▶ John McCarthy coined phrase "Artificial Intelligence"
 - ▶ Newell and Simon presented a reasoning program (Logic Theorist) that could prove theorems (from Newton's Principia Mathematica)

- 1958 John McCarthy
 - ▶ defined high-level symbolic language LISP
 - ◎ still dominates AI programming (used more in US)
 - ▶ invented time-sharing operating systems to enable better access to scarce and expensive computing resources
 - ▶ hypothesised the *Advice Taker*, a complete AI system that used knowledge to search for solutions to problems
 - ▶ Stressed representation and reasoning in formal logic



- 1963–1968 Marvin Minsky’s “microworlds”
 - ▶ solved closed-form integration problems from first-year calculus courses
 - ▶ solved geometric analogy problems from IQ tests
 - ▶ solved algebra story problems (e.g., if my father were twice as old as I am now...)
 - ▶ More interested in getting programs to work than in theory (cf McCarthy)



- 1966–1974, the wildly optimistic predictions for AI failed to materialise
 - ▶ programs only worked for simple examples, e.g., Weizenbaum’s ELIZA program apparently engages in serious conversations but actually just borrows and manipulates sentences typed in by humans. (<http://www-ai.ijs.si/eliza/eliza.html>)
 - ▶ problems with scalability weren’t properly understood
 - ▶ Minsky and Papert’s book “Perceptrons” (1969) showed that single layer networks were incapable of representing XOR
 - ▶ general purpose problem solving methods were too weak to solve problems in complex domains
 - ▶ UK National Research Council cut AI funding after Lighthill Report (1973)

A Dose of Reality (Maybe)

- 1966–1974, the wildly optimistic predictions for AI failed to materialise
 - ▶ programs only worked for simple examples, e.g., Weizenbaum’s ELIZA program apparently engages in serious conversations but actually just borrows and manipulates sentences typed in by humans. (<http://www-ai.ijs.si/eliza/eliza.html>)
 - ▶ problems with scalability weren’t properly understood
 - ▶ Minsky and Papert’s book “Perceptrons” (1969) showed that single layer networks were incapable of representing XOR
 - ▶ general purpose problem solving methods were too weak to solve problems in complex domains
 - ▶ UK National Research Council cut AI funding after Lighthill Report (1973)



- The maturing of AI came through the development of knowledge-based systems that use:
 - ▶ knowledge more suited to making larger reasoning steps
 - ▶ more problem-specific domain knowledge
 - ▶ narrow area of expertise
 - ▶ rules obtained from expert practitioners
- It continues in the various sub-fields mentioned earlier
- Some older AI technology is now considered mainstream computer science
 - ▶ e.g., machine learning

- The ups and downs of AI have resulted in changes to both content and methodology of AI research
- There is now more emphasis on
 - ▶ building on existing theories rather than proposing brand new ones
 - ▶ basing claims on rigorous theorems or hard experimental evidence rather than intuitions
 - ▶ showing relevance to real world applications rather than toy examples

- AI in Entertainment

- ▶ Movies

- Metropolis (1927)
- 2001: A Space Odyssey (1968)
- Bladerunner (1982)
- The Terminator (1984)
- AI (2001)
- I, Robot (2004)

- ▶ TV

- Star Trek: The Next Generation (Lieutenant Commander Data; 1987-94)
- Battlestar Galactica (The Cylons; 1978, 2004-5)



- Books

- ▶ Do Androids Dream of Electric Sheep (1968) by Philip K. Dick
- ▶ Neuromancer (1984) by William Gibson (coined term “cyberspace”)
- ▶ Hyperion (1989) by Dan Simmons
- ▶ Cryptonomicon (1997) by Neal Stephenson (nanotechnology)
- ▶ The Diamond Age (1998) by Neal Stephenson (early wifi, cryptography)
- ▶ Thinks... (2001) by David Lodge

- Robot pets

- ▶ Tamagotchi
- ▶ Aibo

- Video Games

- ▶ Halo

- ◎ Non Player Characters with real time perceptions of their environment
- ◎ Knowledge of the state of the world, as last perceived
- ◎ Emotions based on events
- ◎ Decision making capability



- Video Games

- ▶ Halo

- ◎ Non Player Characters with real time perceptions of their environment
- ◎ Knowledge of the state of the world, as last perceived
- ◎ Emotions based on events
- ◎ Decision making capability

- ▶ Black and White

- ◎ Characters learn from feedback from player
- ◎ Emergent unscripted behaviour



- Video Games

- ▶ Halo

- Non Player Characters with real time perceptions of their environment
- Knowledge of the state of the world, as last perceived
- Emotions based on events
- Decision making capability

- ▶ Black and White

- Characters learn from feedback from player
- Emergent unscripted behaviour

- ▶ The Sims

- Toggle 'free will' on and off
- Agents climb peaks of a 'happiness landscape'



- Strategy games
 - ▶ Successes
 - (3D) Noughts and Crosses
 - Othello/Reversi
 - Connect-4
 - Chess (can still improve)
 - Go
 - ▶ Still a challenge
 - Backgammon
 - Poker
 - Bridge

- Everyday applications
 - ▶ cars – cruise control, fuel injection
 - ▶ planes – autopilots and lower-level control systems
 - ▶ lawnmowers & vacuum cleaners
 - ▶ washing machines
 - ▶ environmental control – light, thermostats, etc.



- **Everyday applications**

- ▶ cars – cruise control, fuel injection
- ▶ planes – autopilots and lower-level control systems
- ▶ lawnmowers & vacuum cleaners
- ▶ washing machines
- ▶ environmental control – light, thermostats, etc.



- **Large scale applications**

- ▶ military strategy planning – Desert Storm
- ▶ prevention of mid-air collisions between planes
- ▶ disaster recovery services – 9/11
- ▶ Deep Space I – remote agent experiment



The Robocup Challenge

- Design a team of robots that can play soccer
 - ▶ and beat other (robot) teams



DARPA Grand Challenge

- Autonomous vehicles – not remote controlled!
 - ▶ 131.6 miles through the Mojave desert
 - ▶ \$2,000,000 (US) Cash Prize



- What is AI, really?
 - ▶ Can machines really be said to think?
 - ▶ What does an intelligent system look like?
 - ▶ Does an AI need—and can it have—emotions and consciousness?
- Can we ever achieve AI, even in principle?
 - ▶ How will we know if we've done it?
 - ▶ If we can do it, should we?

A warning?

- In many sci-fi stories, the robots take over
- But we can always unplug...

