

Introduction to AI

Instructor: Dr. Wei Ding
Fall 2009

What is Artificial Intelligence?

- Views of AI fall into four categories:

Thinking Humanly

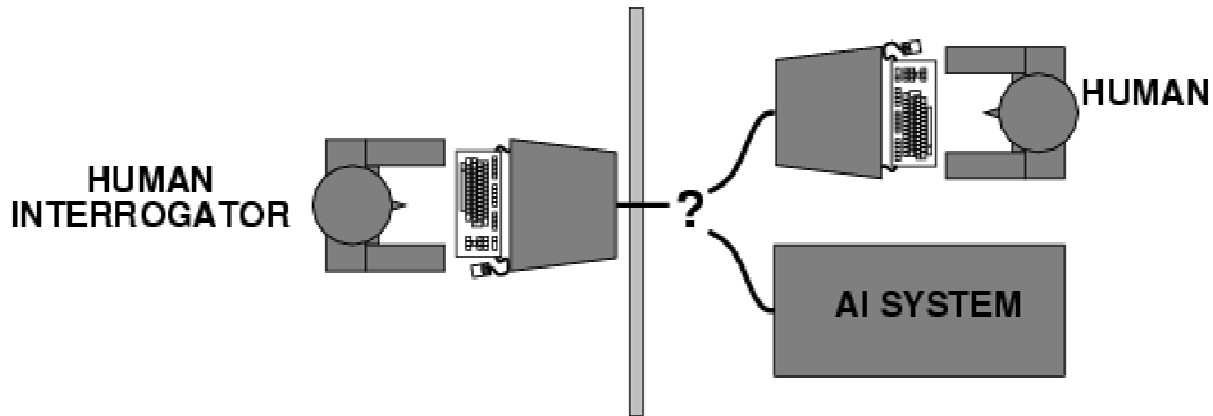
Thinking Rationally

Acting Humanly

Acting Rationally

- The textbook advocates "acting rationally"

Acting humanly: Turing Test



The computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses from a person or not.

However,

Acting humanly: Turing Test

- Yet AI researchers have devoted little effort to passing the Turing test, believing that it is more important to study the underlying principles of intelligence than to duplicate an exemplar.

The quest for "artificial flight" succeeded when the Wright brothers and others stopped imitating birds and learned about aerodynamics. Aeronautical engineering tests do not define the goal of their field as making "machines that fly so exactly like pigeons that they can fool even other pigeons."

Acting humanly: Turing Test

- To pass the Turing Test, the computer would need to possess the following capabilities:
 - **Natural language processing** to enable it to communicate successfully in English
 - **Knowledge representation** to store what it knows or hears
 - **Automated reasoning** to use the stored information to answer questions and to draw new conclusions
 - **Machine learning** to adapt to new circumstances and to detect and extrapolate patterns
 - **Computer vision** to perceive objects
 - **Robotics** to manipulate objects and move about

Thinking humanly: cognitive modeling

- Comparison of the trace of computer program reasoning steps to traces of human subjects solving the same problem.
- **Cognitive Science** brings together computer models from AI and experimental techniques from psychology to try to construct precise and testable theories of the working of the human mind.
- Now distinct from AI
 - AI and Cognitive Science fertilize each other in the areas of vision and natural language.

Thinking rationally: "laws of thought"

- The Greek philosopher Aristotle was one of the first to attempt to codify “right thinking.”
- His syllogisms provided patterns for argument structures that always yielded correct conclusions when given correct premises.
 - For example, “Socrates is a man; all men are mortal; therefore, Socrates is mortal.” → initiated the field called **logic**.

Two main obstacles:

1. It is not easy to take informal knowledge and state it in the formal terms required by logical notation, particularly when the knowledge is less than 100% certain.
2. There is a big difference between being able to solve a problem “in principle” and doing so in practice.

Acting rationally: rational agent

- **Rational** behavior: doing the right thing
- The **right thing**: that which is expected to maximize goal achievement, given the available information
- We will concentrate on general principles of rational agents and on components for constructing them.

- Achieving perfect rationality – always doing the right thing – is not feasible in complicated environments.
- Limited rationality – acting appropriately when there is not enough time to do all the computations one might like.

Rational agents

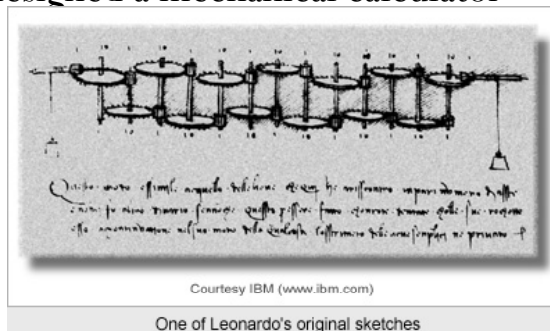
- An **agent** is an entity that perceives and acts
Abstractly, an agent is a function from percept histories to actions:

$$[f: \mathcal{P}^* \rightarrow \mathcal{A}]$$

- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: computational limitations make perfect rationality unachievable
→ design best **program** for given machine resources

The foundations of AI

- Philosophy (428 B.C. – present): Logic, methods of reasoning, mind as physical system foundations of learning, language, rationality
 - Aristotle (384-322 B.C.) developed an informal system of syllogism for proper reasoning, which in principle allowed one to generate conclusions mechanically, given initial premises.
 - The automation of computation. Leonardo da Vinci (1452-1519) designed a mechanical calculator



Courtesy IBM (www.ibm.com)

One of Leonardo's original sketches

The foundations of AI

- Mathematics (c. 800 – present): mathematical formalization in logic, computation, and probability.
 - Intractability: a problem is called intractable if the time required to solve instances of the program grow exponentially with the size of the instance.
 - NP-completeness: Nondeterministic Polynomial time complete.
 - The most notable characteristic of NP-complete problems is that no fast solution to them is known
 - An expert programmer should be able to recognize an NP-complete problem so that he or she does not unknowingly waste time trying to solve a problem

The foundations of AI

- Economics (1776-present): utility, decision theory
 - Utility: how people make choices that lead to preferred outcomes.
 - Decision theory & game theory: how should we do this when others may not go along.
- Neuroscience (1861-present): How do brains process information?
- Psychology (1879-present): How do humans and animals think and act?
- Computer engineering (1940-present): How can we build an efficient computer?
- Control theory and cybernetics (1948-present): How can artifacts operate under their own control?

The foundations of AI

- Linguistics (1957-present): How does language relate to thought?
 - Computational linguistics or natural language processing.
 - knowledge representation: the study of how to put knowledge into a form that a computer can reason with.

Abridged history of AI

- 1943-1955: The gestation of artificial intelligence
- 1956: The birth of artificial intelligence: The Dartmouth workshop
- 1952-1969: Early enthusiasm, great expectations
- 1966-1973: A dose of reality, AI discovers computational complexity
- 1969-1979: Early development of knowledge-based systems
- 1980: AI becomes an industry
- 1987: AI becomes a science: to be accepted, hypotheses must be subjected to rigorous empirical experiments, and the results must be analyzed statistically for their importance.
- 1995: The emergence of intelligent agents: understand the workings of agents embedded in real environments with continuous sensory inputs.

State of the art

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997,
http://en.wikipedia.org/wiki/Deep_Blue_versus_Garry_Kasparov
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego),
http://www.cs.cmu.edu/afs/cs/usr/tjochem/www/nhaa/nhaa_home_page.html
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- Proverb solves crossword puzzles better than most humans, <http://www.oneacross.com/proverb/>

