

Introduction

Artificial Intelligence

What is Intelligence?

- the ability to:
 - learn or understand or
 - deal with new situations or
 - apply knowledge to manipulate one's environment

AI is Exciting

- AI still has openings for several full-time Einsteins
- AI currently encompasses a huge variety of subfields, ranging from general-purpose areas, such as learning and perception to such specific tasks as:
 - playing chess
 - proving mathematical theorems
 - writing poetry
 - driving cars
 - creating puzzle
 - preparing and evaluating exams
 - and diagnosing diseases

What is AI?

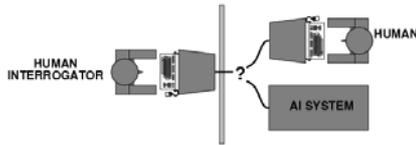
Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

The textbook advocates "acting rationally"

Acting humanly: Turing Test

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" → "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game



- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes

Acting humanly: 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 know 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

Acting humanly: Turing Test

Turing's test deliberately avoided direct physical interaction between the interrogator and the computer, because physical simulation of a person is unnecessary for intelligence. To pass the total Turing Test, the computer will need:

- **Computer Vision** to perceive objects, and
- **Robotics** to manipulate objects and move about

Thinking humanly: cognitive modeling

- requires some way of determining how humans think?
- We need to get inside the actual workings of human minds. There are two ways to do this:
 - through introspection (try to catch our own thoughts as they go by)
 - through psychological experiments
- Allen and Simon, who developed the "General Problem Solver" were more concerned with comparing the its reasoning steps to the trace of human subjects solving the same problem.

Thinking humanly: cognitive modeling

- The interdisciplinary field of cognitive science brings together computer models from AI and experimental techniques from psychology to try to construct precise and testable theories of the workings of the human mind
- Real cognitive science based on experimental investigation of actual humans or animals

Thinking rationally: "laws of thought"

- Aristotle was the first to attempt to codify "correct thinking". He provided pattern for argument structures that always yielded correct conclusions when given correct premises [Socrates is mortal]
- Logicians in the 19th century developed a precise notation for statements about all kinds of things in the world and about the relations among them. By 1965, programs existed that could, in principle, solve any solvable problem described in logical notation

Thinking rationally: "laws of thought"

There are two main obstacles to this approach:

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

Acting rationally: rational agent

- An agent is just something that acts
- Computer agents are expected to have other attributes such as operating under autonomous control, perceiving their environment, persisting over a prolonged time period, adapting to change, and being capable of taking on another's goals
- A rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty the best expected outcome
- In the laws of thought approach to AI, the emphasis was on correct inferences

Acting rationally: rational agent

Study of AI as rational-agent design has at least two advantages:

- **First**, it is more general than the “laws of thought” approach
- **Secondly**, it is more convincing to scientific development than are approaches based on human behaviour or human thought

Foundations of AI

We provide disciplines that contributes ideas, viewpoints, and techniques to AI:

- **Philosophy**
 - can formal rules be used to draw valid conclusions?
 - how does the mental mind arise from a physical brain?
 - where does knowledge come from?
 - how does knowledge lead to action?
- **Mathematics**
 - what are the formal rules to draw valid conclusions?
 - what can be computed?
 - how do we reason with uncertain information?

Foundations of AI

- **Economics**
 - how should we make decision so as to maximize payoff?
 - how should we do this when others may not go along?
 - how should we do this when the payoff may be far in the future?
- **Neuroscience**
 - how do brain process information?
- **Psychology**
 - how do humans and animal thinks and act?
- **Computer Engineering**
 - how can we build an efficient computer?

Foundations of AI

- **Control Theory and Cybernetics**
 - how do artifacts operate under their own control?
- **Linguistics**
 - how does language relate to thought?

State of the art

- **Game Playing** Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- **Autonomous Control** No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- **Logistics Planning** 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

State of the art

- **Autonomous Planning and Scheduling** NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- **Language Understanding and Problem Solving** solves crossword puzzles better than most humans