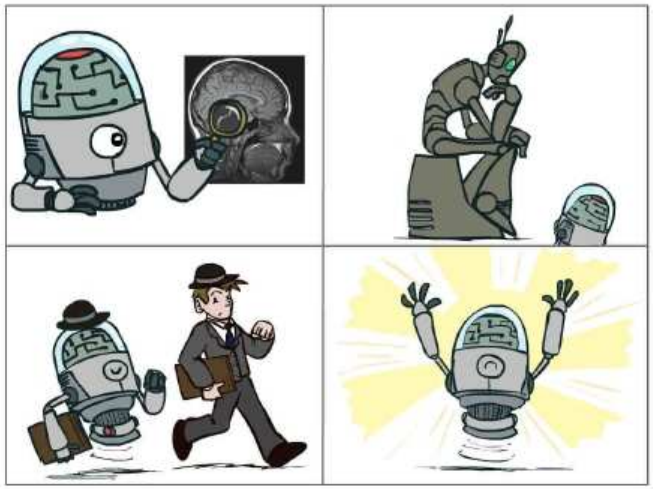
Artificial Intelligence

Chapter 1

Outline

- ♢ Class organisation
- ♢ What is AI?
- ♢ A brief history
- ♢ The state of the art

What is AI?



Thinking humanly: The automation of activities that we associate with human thinking, activities such as decision-making, problem solving, learning...(Richard Bellman, 1978)

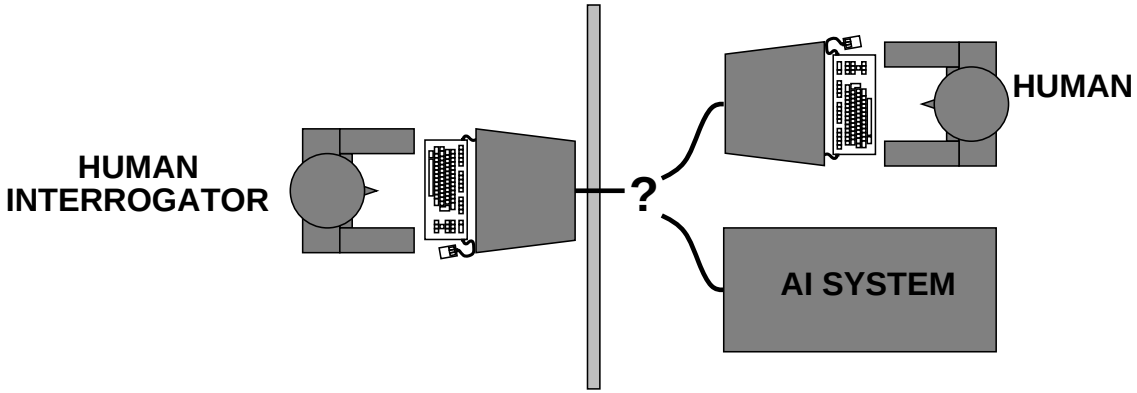
Acting Humanly: "The art of creating machines that perform functions that require intelligence when performed by people." (Ray Kurzweil, 1990)

Acting rationally: AI . . . is concerned with intelligent behavior in artifacts (Nilsson, 1998)

Acting humanly: The Turing test

Turing (1950) "Computing machinery and intelligence":

- "Can machines think?" $\longrightarrow$ "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
- Suggested major components of AI: knowledge, reasoning, language understanding, learning
- Total Turing test: computer vision, robotics

Thinking humanly: Cognitive Science

1960s "cognitive revolution": information-processing psychology replaced prevailing orthodoxy of behaviorism

Get inside: introspection, psychological experiments, brain imaging

Requires scientific theories of internal activities of the brain
- What level of abstraction? "Knowledge" or "circuits"?
- How to validate? Requires
  1) Predicting and testing behavior of human subjects (top-down)
  or 2) Direct identification from neurological data (bottom-up)

Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI

Both share with AI the following characteristic: **the available theories do not explain anything resembling human-level general intelligence**. Hence, all three fields share one principal direction!

Thinking rationally: Laws of Thought

Normative (or prescriptive) rather than descriptive

Aristotle: what are correct arguments/thought processes?

Several Greek schools developed various forms of logic:
notation and **rules of derivation** for thoughts;
may or may not have proceeded to the idea of mechanization

Logicist tradition withi AI hopes to build on such programs to create intelligent systems.

Problems:
1) Not all intelligent behavior is mediated by logical deliberation
2) What is the purpose of thinking? What thoughts **should** I have
out of all the thoughts (logical or otherwise) that I **could** have?

Acting rationally

Rational behavior: doing the right thing

The right thing: that which is expected to maximize goal achievement, given the available information

Doesn't necessarily involve thinking—e.g., blinking reflex—but thinking should be in the service of rational action

Aristotle (Nicomachean Ethics):

> **Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good**

Rational agents

An agent is an entity that perceives and acts

This course is about designing rational agents

Abstractly, an agent is a function from percept histories to actions:

$$f : \mathcal{P}^* \to \mathcal{A}$$

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**

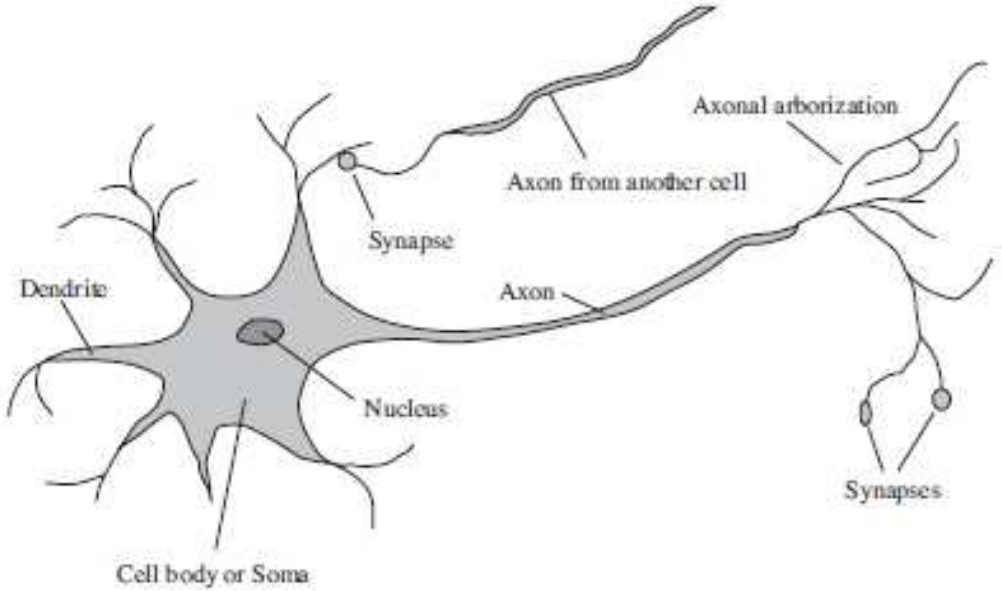
$\to$ design best program for given machine resources

Interplay of reasoning and perception:

AI prehistory

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

Neuron



	Supercomputer	Personal Computer	Human Brain
Computational units	10^4 CPUs, 10^{12} transistors	4 CPUs, 10^9 transistors	10^{11} neurons
Storage units	10^{14} bits RAM	10^{11} bits RAM	10^{11} neurons
	10^{15} bits disk	10^{13} bits disk	10^{14} synapses
Cycle time	10^{-9} sec	10^{-9} sec	10^{-3} sec
Operations/sec	10^{15}	10^{10}	10^{17}
Memory updates/sec	10^{14}	10^{10}	10^{14}

The history of AI

The gestation (1943-1955):

♢ 1943: McCulloch & Pitts: model of neurons → Boolean circuit of the brain

♢ 1949: Donald Hebb - updading rule for modifying the connection strengths (Hebbian learning)

♢ 1950: Turing's Computing Machinery and Intelligence: introduces Turing Test, machine learning, genetic algorithms, and reinforcement learning.

The birth (1956):

♢ McCarthy (1927 -2011): 2 month, 10 man study of AI, to make machine use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.

♢ Main actors for the next 20 years from MIT, CMU, Standford and IBM

The history of AI

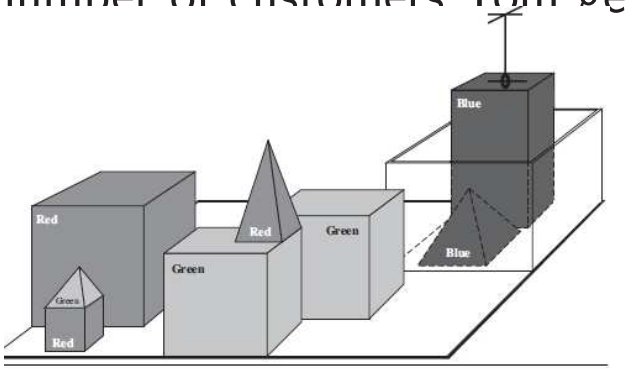
Early enthusiasm, great expectations (1952-1969):

♢ Geometry Theorem Prover, Logic Theorist, General Problem Solver, Playing checkers (Given the primitive computers and programming tools)

♢ McCarthy: Lisp (1958), Time sharing (1959), AdviceTaker

♢ Minsky: Microwords- algebra story problems, blocks world

"If the number of customers Tom gets is twice the square of 20 percent of the number of advertisements he runs, and the number of advertisements he runs is 45, what is the number of customers Tom gets?"



The history of AI

A dose of reality (1966-1973):

♢ Translation of Russian scientific paper in context of Sputnik (Alpack report 1966)

"the spirit is willing but the flesh is weak"

"the vodka is good but the meat is rotten"

♢ Lighthill report (1973) most successful algorithms would halt on real world problems and were only suitable for solving "toy" versions.

false optimism on: cobinatorial explosion to be solved by faster hardware and larger memories, no progress on genetic algorithms

2 input perceptron cannot be trained to recognize that the inputs are different

The history of AI

Knowledge-based systems: The key to power? (1969-1979)

♢ expert systems: Dendral, Mycin (certainty factor)

♢ Prolog, 1972 (EU), Planner (US)

♢ Minsky: frames - facts about a particular object, taxonomy of types
roots for OOP

AI becomes an industry (1980-present)

♢ The first successful commercial expert system, R1 (savings of 40milion a year)

♢ 1981, the Japanese announced the "Fifth Generation project, a 10-year plan to build intelligent computers running Prolog.

♢ hundreds of companies building expert systems, vision systems, robots,

♢ the return of neural networks: complements the symbolic approaches

The history of AI

AI adopts the scientific method (1987-present)

- ♢ build on existing theories than to propose brand-new ones
- ♢ to base claims on rigorous theorems or hard experimental evidence rather than on intuition
- ♢ and to show relevance to real-world applications rather than toy examples
- ♢ speech recongition (HMM), datamining, bayesian networks

The emergence of intelligent agents (1995-present)

- ♢ Internet, the most important environment

The availability of very large data sets (2001-present)

Potted history of AI

1943	McCulloch & Pitts: Boolean circuit model of brain
1950	Turing's "Computing Machinery and Intelligence"
1952–69	Look, Ma, no hands!
1950s	Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
1956	Dartmouth meeting: "Artificial Intelligence" adopted
1965	Robinson's complete algorithm for logical reasoning
1966–74	AI discovers computational complexity Neural network research almost disappears
1969–79	Early development of knowledge-based systems
1980–88	Expert systems industry booms
1988–93	Expert systems industry busts: "AI Winter"
1985–95	Neural networks return to popularity
1988–	Resurgence of probability; general increase in technical depth "Nouvelle AI": ALife, GAs, soft computing
1995–	Agents, agents, everywhere . . .
2003–	Human-level AI back on the agenda

State of the art

Which of the following can be done at present?

- ♢ Play a decent game of table tennis
- ♢ Drive safely along a curving mountain road
- ♢ Drive safely along Telegraph Avenue
- ♢ Buy a week's worth of groceries on the web
- ♢ Buy a week's worth of groceries at Berkeley Bowl
- ♢ Play a decent game of bridge
- ♢ Discover and prove a new mathematical theorem
- ♢ Design and execute a research program in molecular biology
- ♢ Write an intentionally funny story
- ♢ Give competent legal advice in a specialized area of law
- ♢ Translate spoken English into spoken Swedish in real time
- ♢ Converse successfully with another person for an hour
- ♢ Perform a complex surgical operation
- ♢ Unload any dishwasher and put everything away

Summary

Approach AI with different goals: Are you concerned with thinking or behavior? Do you want to model humans or work from an ideal standard?

Rational action: intelligent agent takes the best possible action in a situation

Philosophers: "mind is in some ways like a machine"

Mathematicians: logical and probabilistic statements

Economists: decisions that maximize the expected outcome.

Neuroscientists: how the brain works and the ways in which it is similar to and different from computers.

Psychologists: humans can be considered information-processing machines.

Control theory: devices that act optimally on the basis of feedback.

The history of AI has had cycles of success, misplaced optimism, and resulting cutbacks in enthusiasm and funding.