Artificial Intelligence

Chapter 1

Outline

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

What is AI?

The four approaches of AI:

Systems that think like humans	**Systems that think rationally**
Systems that act like humans	**Systems that act rationally**

Can you guess the trend today?

Acting humanly: The Turing test

Turing (1950) "Computing machinery and intelligence":

♢ "Can machines think?" ⟶ "Can machines behave intelligently?"

♢ Operational test for intelligent behavior: the Imitation Game

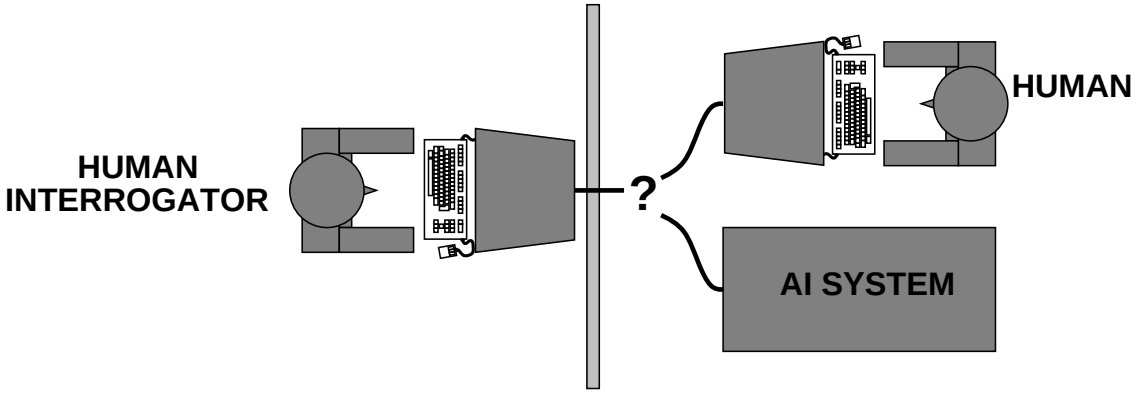

♢ Suggested major components of AI:
Natural language processing, knowledge representation, automated reasoning, and machine learning.

Problem: Turing test is not **reproducible**, **constructive**, or amenable to **mathematical analysis**

Thinking humanly: Cognitive Science

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

Requires scientific theories of internal activities of the brain

- What level of abstraction? "Knowledge" or "circuits"?
- How to validate? Requires
  1) Introspection - trying to catch our thoughts as they go by
  2) Psychological experiments - observing a person in action
  3) Direct identification from neurological data (brain image like fMRI)

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 (or engender) anything resembling human-level general intelligence

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

Direct line through mathematics and philosophy to modern AI

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:

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

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**

$\rightarrow$ design best program for given machine resources

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

Potted history of AI

1943–55	The gestation of AI: first of neural network
1956	Dartmouth meeting: "Artificial Intelligence" adopted
1952-1969	early enthusiasm, great expectations: "Look, Ma, no Hands" era
1965	Robinson's complete algorithm for logical reasoning
1966–74	The reality: AI discovers computational complexity
	Neural network research almost disappears
1969–79	Early development of knowledge-based systems: weak methods
1980–	Expert systems in industry
1985–95	Neural networks return to popularity
1987–	AI adopts scientific methods: HMM, data mining, machine learning
1995–	Agents, agents, everywhere . . .
2003–	Human-level AI back on the agenda
2001–	Availability of very large data

State of the art

- ♢ Robotic vehicles
- ♢ Speech recognition
- ♢ Autonomous planning and scheduling
- ♢ Game playing
- ♢ Spam filtering
- ♢ Machine translations