Artificial Intelligence

Topic 1

Introduction

- ◇ What is AI?
- ◇ Contributions to AI
- ◇ History of AI
- ◇ Modern AI

Reading: Russel and Norvig, Chapter 1

1.1 AI in the Media - the glitz and glamour

- ◇ sci-fi
  — Kubric, Spielberg,...
- ◇ "science" programs
  — "Towards 2000"
- ◇ news/current affairs
  — Kasparov
- ◇ advertisements
  — intelligent washing machines, TVs, cars, molecules,...?

⋮

Don't believe a word you hear!

(...without proof)

1.2 The AI Literature

"[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . ." (Bellman, 1978)	"The study of mental faculties through the use of computational models" (Charniak+McDermott, 1985)
"The study of how to make computers do things at which, at the moment, people are better" (Rich+Knight, 1991)	"The branch of computer science that is concerned with the automation of intelligent behavior" (Luger+Stubblefield, 1993)

Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

1.3 Thinking humanly: cognitive modelling

- determine how humans think
- develop theory of human mind — psychological experiments
- model theory using computer programs

  eg. General Problem Solver (GPS) [Newel & Simon, 1961]

Requires scientific theories of internal activities of the brain

- What level of abstraction? "Knowledge" or "circuits"?
- How to validate?

  1. Predicting and testing behavior of human subjects (top-down)
     $\Rightarrow$ *Cognitive Science*
  2. Direct identification from neurological data (bottom-up)
     $\Rightarrow$ *Cognitive Neuroscience*

1.4 Acting humanly: The Turing test

Alan Turing (1950) "Computing machinery and intelligence":

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

intelligence = ability to act indistinguishably from a human in cognitive tasks

♢ Operational test for intelligent behavior ⇒ Turing Test

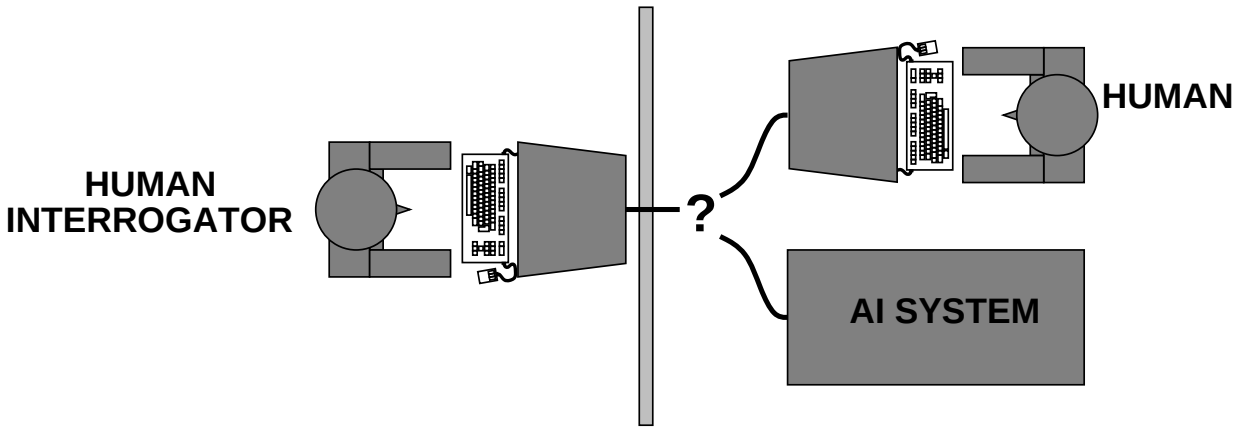

– human interrogates computer via teletype
– passes test if human cannot tell if there's a human or computer at the other end

What else might today's turing test include...?

1.4 Acting humanly: The Turing test

- Loebner Prize in Artificial Intelligence
  `http://www.loebner.net/Prizef/loebner-prize.html`
- BotPrize
  `http://botprize.org/`

1.4 Acting humanly: The Turing test

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

♢ Anticipated all major arguments against AI in following 50 years

♢ Suggested major components of AI: knowledge, reasoning, language understanding, learning

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

1.5 Thinking rationally: Laws of Thought

Began with Greeks in 4th century BC (1st century BE)

— e.g. Aristotle's logical *syllogisms*

All men are mortal,
Socrates is a man,
therefore Socrates is mortal

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

1.5 Thinking rationally: Laws of Thought

Direct line through mathematics and philosophy to modern AI, e.g.

- Boole (1815–1864, *Laws of Thought* 1854, *Boolean logic*)
- Frege (1848–1925, *Begriffsschrift* 1879, *first-order logic*)
- Hilbert (1862–1943), (*Hilbert systems*)
- Russell & Whitehead (*Principia Mathematica* 1918)
- Tarski (1902–1983, *Tarski semantics* 1933)

Problems:

1. Normative (or prescriptive) rather than descriptive
2. Not all intelligent behavior is mediated by logical deliberation

1.6 Acting rationally

$\Rightarrow$ act in such a way to achieve goals, given beliefs

(Doesn't necessarily involve thinking—e.g., blinking reflex.
Is a thermostatically controlled heater "intelligent"?)

Define *agent* — entity that percieves surroundings and acts accordingly

AI = study and construction of *rational agents*.

AI as intelligent agent design

- incorporates aspects from the other three approaches
- currently the dominant view of AI

1.7 An Engineering Viewpoint

How can we develop programs/systems that

- do more useful stuff?
- do it better?

We have a proof of concept!

> "*a bunch of methodologies, inspired by socio-biological analogy, for getting machines to do cool things*"

2. Pre-history — Contributions to AI

Philosophy (428 BC $\longrightarrow$ present)

- logic, methods of reasoning
- mind as a physical system
- foundations of learning, language, rationality

Mathematics (c. 800 $\longrightarrow$ present)

- formal representation and proof
- computation, (in)tractability, (un)decidability
- probability

Psychology (1879 $\longrightarrow$ present)

- perception and motor control
- cognitive neuroscience
- learning (reinforcement)

2. Pre-history — Contributions to AI

Computing (1940 → present)

- provision of programmable machines
- algorithms
- declarative languages (PROLOG and LISP)
- neural computing

Linguistics (1957 → present)

- language theory: grammar, semantics

3. The Chequered History of AI

In the beginning...

1943 McCulloch & Pitts: Boolean circuit model of brain
on/off neurons, corresponding to propositions
even suggested networks could learn!
$\Rightarrow$ forerunner of symbolic and connectionist traditions

1950 Turing's "Computing Machinery and Intelligence"
Turing and Shannon writing chess programs
— with no computers!

1951 Minsky & Edmonds, first neural net computer, SNARK

1950s Newell & Simon's Logic Theorist (machine coded by hand!)

> "*We have created a computer program capable of thinking non-numerically, and thereby solved the venerable mind-body problem.*" — Simon

Able to prove theorems in Russell and Whitehead's *Principia Mathematica*
— even came up with a shorter proof!

1956 McCarthy's Dartmouth meeting: "Artificial Intelligence"

3. The Chequered History of AI

"Look, Ma, no hands!" era

1952 Samuel's checker playing programs
— eventually tournament level
$\Rightarrow$ disproved computers can only do what they are told

1958 McCarthy
— defined LISP $\Rightarrow$ dominant AI programming language
— he and others invented time-sharing $\Rightarrow$ birth of DEC
— published "Programs with Common Sense"
$\Rightarrow$ defined hypothetical program Advice Taker
AI program that includes general knowledge, axioms
forerunner to knowledge representation and reasoning today

1959 McCarthy & Hayes "Philisophical Investigations from the Standpoint of AI" $\Rightarrow$ KR, reasoning, planning,...

1961 Newell and Simon's General Problem Solver (GPS)
imitate human problem solving — "thinking humanly"

3. The Chequered History of AI

1965 Robinson's *resolution principle*
complete theorem proving algorithm for 1st-order logic
made PROLOG possible

1971 STRIPS — practical logic-based planning system
SHAKEY — integration of logical reasoning and
physical activity

3. The Chequered History of AI

Falling off the bike

Bold predictions prove elusive...

♢ Lack of domain knowledge

eg. machine translation

US goverment funding to translate Russian to English after Sputnik in 1957

Initially thought syntactic transformations using grammar and electronic dictionary

"the spirit is willing but the flesh is weak"

"*the vodka is good but the meat is rotten*"

1966 report — no immediate prospect of success

All government funding cancelled

3. The Chequered History of AI

♢ Intractibility — early solutions did not scale up!

development of computational complexity theory and NP-completeness

> *program finds solution in principle* $\not\rightarrow$ *has any of the mechanisms needed to find it in practice*

difficulties of combinatorial explosion one of main criticisms in Lighthill Report in 1973

$\Rightarrow$ British government decision to end support for AI research in all but 2 universities

♢ limitations of basic AI structures

1969 Minsky and Papert's book *Perceptrons*

— two-input perceptron could not be trained to recognise when its two inputs were different (*exclusive-or* problem)

$\Rightarrow$ research funding for neural nets all but disappeared

3. The Chequered History of AI

AI goes specialist

1969 Buchanan et al, "Heuristic DENDRAL: a program for generating explanatory hypotheses in organic chemistry."
Arguably first *knowledge-intensive* system.
Later incorporated McCarthy's ADVICE TAKER approach — clean separation of knowledge from reasoning.
$\Rightarrow$ birth of *expert systems*.

1976 MYCIN — diagnosis of blood infections.
$\sim$ 450 rules. Performed as well as some experts, better than junior doctors.
No theoretical model — acquired rules from interviewing experts.
Early attempt to deal with *uncertainty*.

3. The Chequered History of AI

1979 Duda et al., PROSPECTOR
Probabilistic reasoning system — recommended drilling at geological site that proved to contain large molybdenum deposit!!

1970s Recognition that language understanding also required knowledge and a means of using it. (Charniak)
"*There is no such thing as syntax.*" — Schank

1973 Woods' LUNAR system — allowed geologists to ask questions in English about Apollo's rock samples.
— first NLP program used by others for real work.

3. The Chequered History of AI

Expert systems industry booms...

1982 McDermott's R1 began operation at DEC
— first successful commercial expert system.
Helped configure orders for new systems — saved estimated $40 million a year.

1988 DEC: 40 deployed expert systems
Du Pont: 100 in use, 500 in development, est $10m year

Increased demand for AI languages — eg Prolog

1980s Japanese ambitious "Fifth Generation" project.
"Prolog machines" — millions of inferences per sec.

US funding increased accordingly.
British Alvey report reinstated funding cut by Lighthill report (under new name "Intelligent Knowledge-Based Systems")

Boom in Expert System development tools,
dedicated Lisp workstations (eg Symbolics), etc...

few million sales in 1980 $\longrightarrow$ $2 billion in 1988

3. The Chequered History of AI

and busts. . . ?

∼1986 Recognition of limitations and some disillusionment
— buying expert system shell and filling it with rules
not enough.

1985–95 Neural networks return to popularity
— rediscovery of back-propagation learning algorithm.
Brooks' insects.
Symbolic vs sub-symbolic argument intensifies!
(battles for funding)

⇒ predictions of "AI winter".

4. Modern AI

AI matures (at least a bit!)

- Better understanding of difficulty!
- Increase in technical and theoretical depth.
- Recognition (by most) of the need for both symbolic and sub-symbolic approaches, working together.
- Resurgence and incorporation of probabilistic and decision-theoretic methods.

$\Rightarrow$ emphasis on solid foundations and a more wholistic approach

The new environment!

- Fast distributed hardware
- New languages (eg Java), distributed programs
- Increased communications (WWW).
- New possibilities, eg ALife, GAs,...
- Advances in understanding of biological and neural systems (biologically-inspired computing)
- New applications, eg Web agents, Mars rovers, ...

$\Rightarrow$ emphasis on intelligent "agents", incorporating a range of AI technologies

The End