ECSE 526
Artificial Intelligence (AI)

www.cim.mcgill.ca/~jer/courses/ai

Readings for this class

- Chapter 1-1.3
- Littman, 'Rise of the Machines' is Not a Likely Future

Learning Objectives

- gain a high-level awareness of what AI is
- overview some of the key developments in AI
- understand the expectations for the course
- gain some exposure to the learning tools we'll be using throughout the course

Course Website

- www.cim.mcgill.ca/~jer/courses/ai
  - course outline
  - lecture slides (treat as a summary)
  - videos for selected topics
  - readings (mostly from Russell & Norvig)
  - assignments
  - other course resources

Interaction beyond the classroom

- Office hours: M 11:45-12:45 in MC 424 or by appointment
- Moodle for on-line class communication

Evaluation

component	weight
assessment of learning	**30%**
• individual learning	25%
• group learning	5%
assignments x 3	**45%**
• submission mark	36%
• peer review activities	6%
• qualitative peer feedback	3%
project (2-3 students)	**25%**
• report and results	15%
• presentation	5%
• peer review activities	5%

In-Class Evaluation

- some classes will begin with an evaluation of your learning
- if you miss one of these classes, don’t panic – your worst grade is not counted
- we’ll be using Learning Catalytics both for discussion questions and for assessment

Class Format

- all classes will involve one or more of:
  - review and discussion of the assigned materials
  - exercises to reinforce the concepts
  - assignment competition (for fun and motivation)

What is AI?
The “Hollywood” view

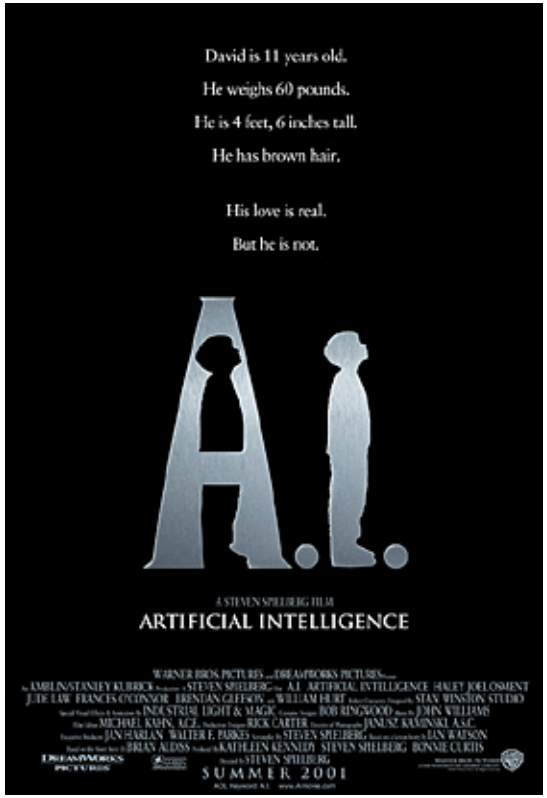

Definitions: What is AI?

Thinking humanly cognitive modeling, reasoning like a human, neural nets	**Thinking rationally** problem solving by logical inference and reasoning
Acting humanly passing the Turing Test	**Acting rationally** “doing the right thing”

What is AI?

Machines that can beat humans at intelligent tasks

What is AI?

Systems that perform complex human tasks

What is AI?

Systems that learn from and interact with humans

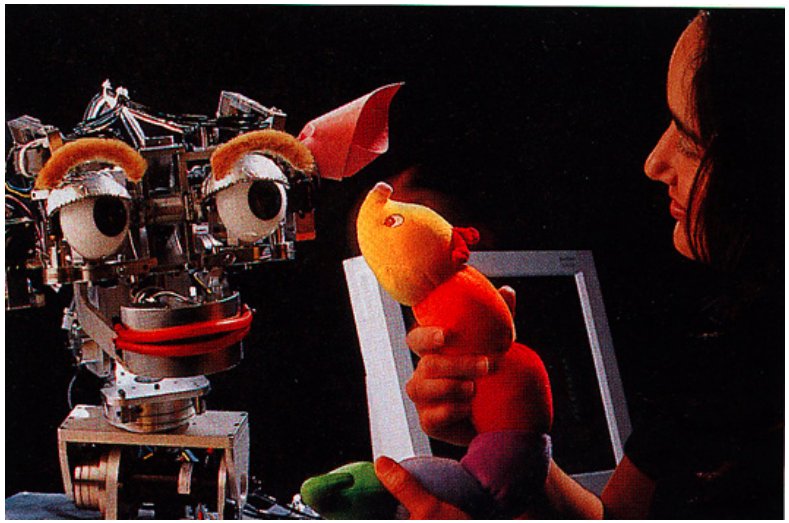

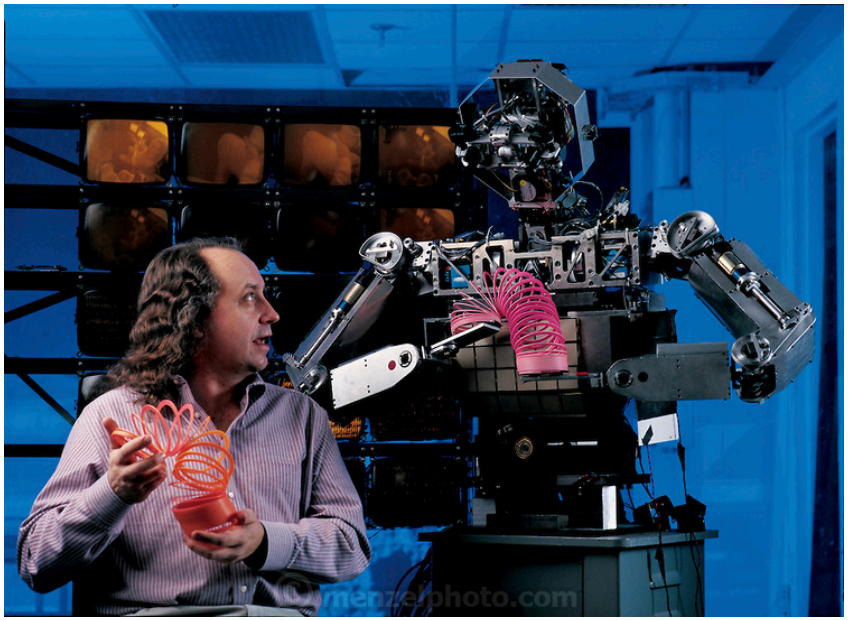

© 2000 Peter Menzel, www.menzelphoto.com, Robo Sapiens

What is AI?

Systems that learn from and interact with humans

Hiroshi Ishiguru (Japan) and android

What do you think? Can computers…

- write convincing poetry?
- compose pleasant music?
- learn to juggle?

Some AI topics

- Core areas
  - Knowledge representation
  - Reasoning/inference
  - Machine learning
- Perception
  - Vision
  - Natural language
  - Robotics
- Uncertainty
  - Probabilistic approaches
- General algorithms
  - Search
  - Planning
  - Constraint satisfaction
- Applications
  - Game playing
  - AI and education
  - Distributed agents
- Decision theory
  - Electronic commerce
  - Auctions
- Reasoning with symbolic data

History and Philosophical Underpinnings

René Descartes

cogito ergo sum (1637)

- it would never be possible to make a machine that thinks as humans do
- there could be no feeling without a conscious state of awareness, and no conscious state of awareness without a true mind to perceive it

McCulloch and Pitts

Neural Logical Calculus (1943)

- proposed highly simplified model of a biological neuron
- binary threshold neuron – influenced by logic

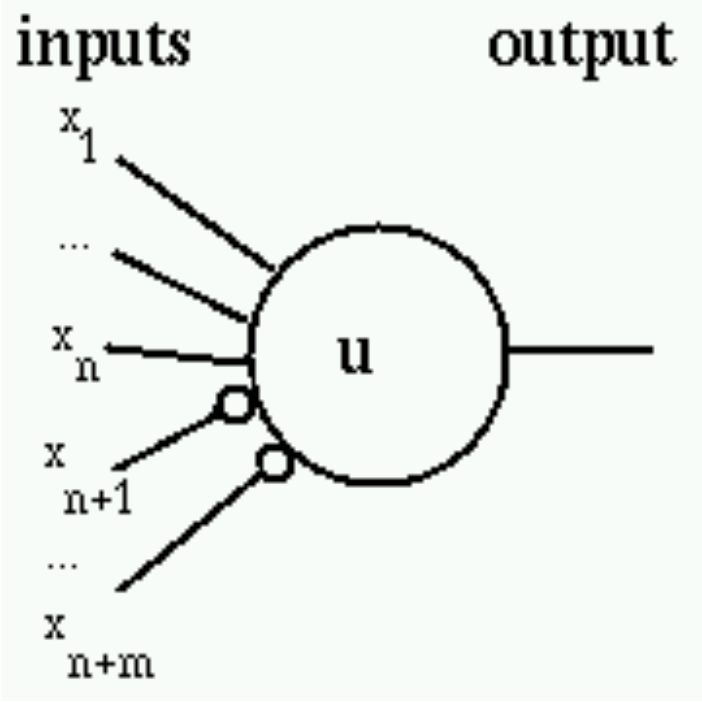

Alan Turing

Computing Machinery & Intelligence (1950)

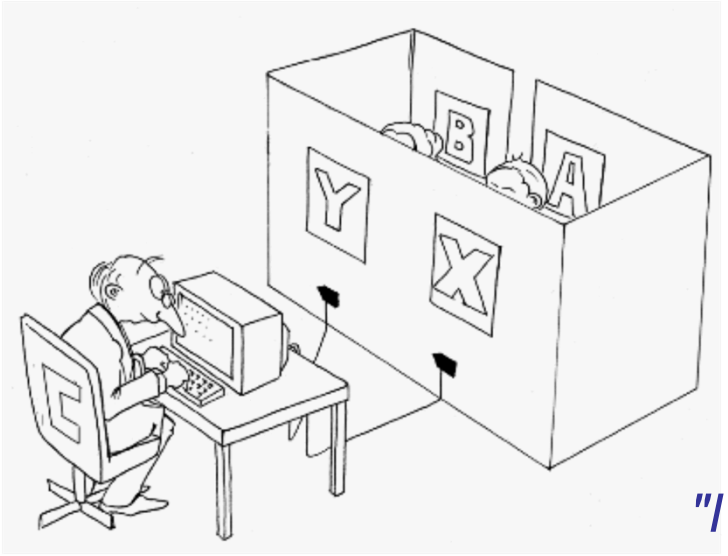

The Turing Test

"I believe that in about fifty years' time it will be possible to program computers ... to make them play the imitation game so well that an average interrogator will not have more than 70 percent chance of making the right identification after five minutes of questioning."

Has the Turing Test been passed?

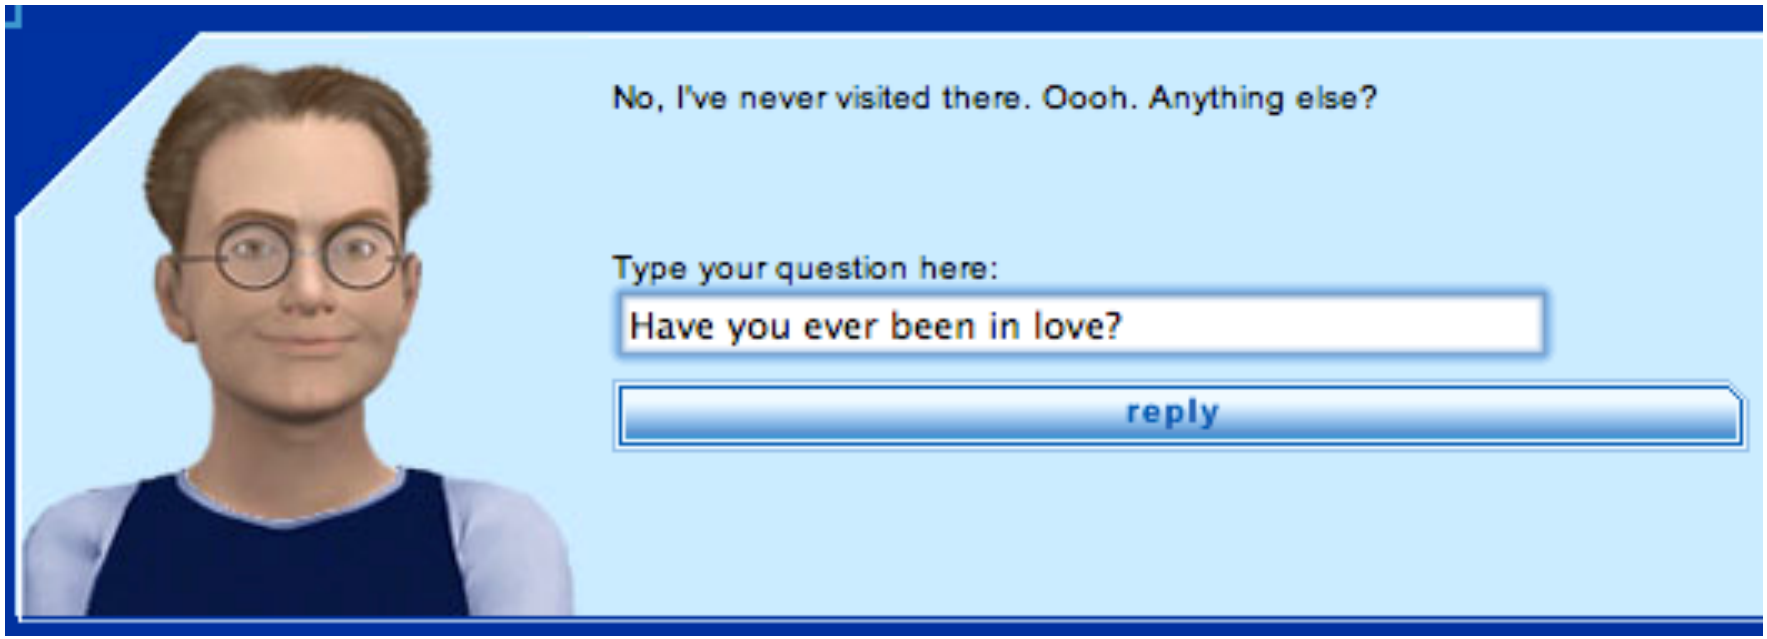

John McCarthy

Organizer (with Minksy) of Dartmouth Conference (1956)

- coined the term “Artificial Intelligence”
- invented LISP (1956-1959)
- wrote first paper on logical AI: “Programs with Common Sense” (1959) considered foundation of Good Old Fashioned Artificial Intelligence (GOFAI)

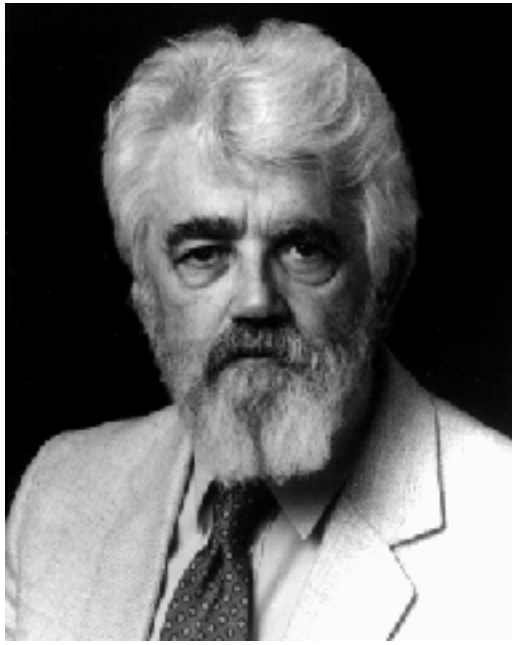

Newell and Simon

General Problem Solver (1957)

- any formalized symbolic problem can be solved, in principle
- solves (simple) problems following line of human reasoning

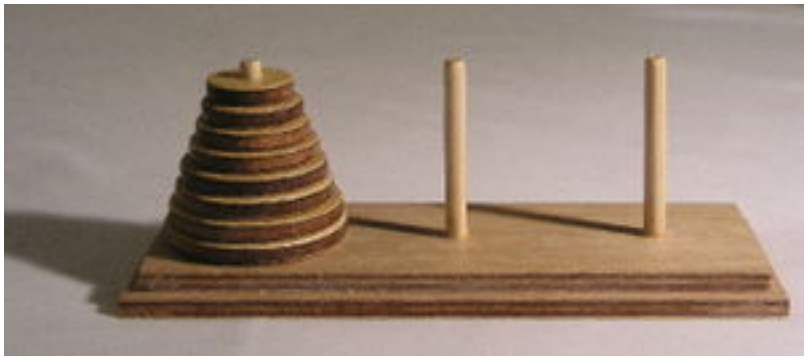

Terry Winograd

SHRDLU (1968-1970)

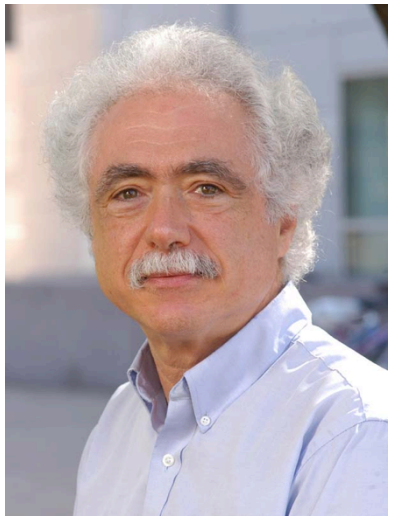

- computer programs could solve spatial and logic problems

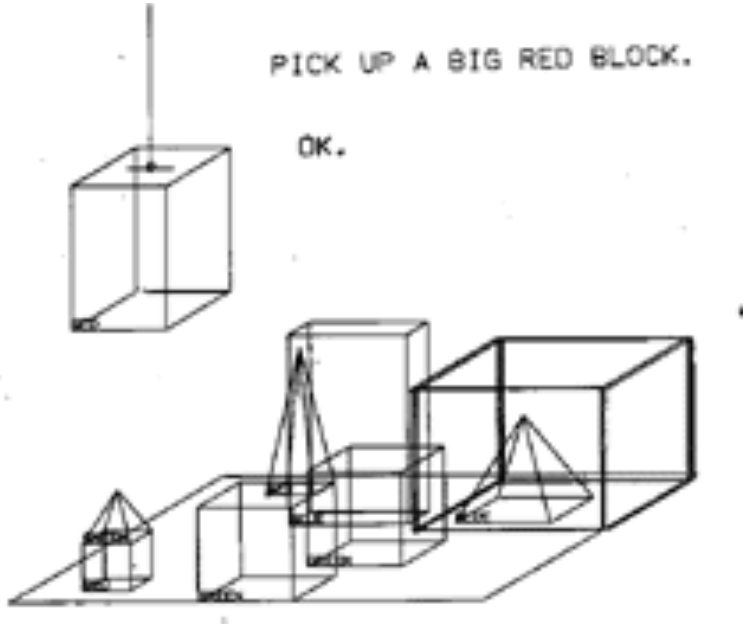

Marvin Minsky

- (SNARC, 1951)
  built first neural network simulator: Stochastic Neural-Analog Reinforcement Computer
- (Society of Mind, 1986)
  what we call intelligence could be a product of the interaction of non-intelligent parts

Lenat

Cyc (1984)

- "Intelligence is 10 million rules."
  - "every tree is a plant": (#$genls #$Tree-ThePlant #$Plant)
  - "plants die eventually"
  - Question: Do trees die?
- massive knowledge engineering effort

John Searle

The Chinese Room (1980)

A man is in a room with a book of rules. Chinese sentences are passed under the door to him. The man looks up in his book of rules how to process the sentences. Eventually the rules tell him to copy some Chinese characters onto paper and pass the resulting Chinese sentences as a reply to the message he has received.

Strong vs. Weak AI
(Searle 1980)

- Strong AI
  - the supposition that some forms of artificial intelligence can truly reason and solve problems, achieve self-awareness and demonstrate a wide range of human-level cognitive abilities
- Weak AI
  - machines can demonstrate intelligence but do not necessarily have a mind

Rodney Brooks

Subsumption Architecture & Embodiment (mid-1980s)

decomposing complicated intelligent behaviour into many "simple" behaviour modules

Deep Blue

First Computer Chess Victory over a reigning world champion (1996)

"A machine would be world chess champion within ten years"
(Herbert Simon, 1957)

Hiroaki Kitano

RoboCup (1995)

By the year 2050, develop a team of fully autonomous humanoid robots that can win against the human world soccer champion team.

DARPA Grand Challenge (2004)

- race for a $2 million prize where cars drive themselves across several hundred miles of challenging desert terrain without any communication with humans, using GPS, computers and a sophisticated array of sensors

Watson wins “Jeopardy!” (2011)

The computer's techniques for unraveling Jeopardy! clues sounded just like mine. That machine zeroes in on key words in a clue, then combs its memory (in Watson's case, a 15-terabyte data bank of human knowledge) for clusters of associations with those words. It rigorously checks the top hits against all the contextual information it can muster: the category name; the kind of answer being sought; the time, place, and gender hinted at in the clue; and so on. And when it feels "sure" enough, it decides to buzz. This is all an instant, intuitive process for a human Jeopardy! player, but I felt convinced that under the hood my brain was doing more or less the same thing.

— Ken Jennings (human opponent)

Homework

- Read **before next class**:
  - Ch. 1-1.3, Ch. 3 of Russell and Norvig
  - Littman, 'Rise of the Machines' is Not a Likely Future
- simple AI:
  - read and run the “Hello World” example in `/opt/linux64/simpleai/samples/search/hello_world.py`
  - can you identify the main characteristics of a problem that can be solved using search techniques?