Introduction to AI (COMP-424)

Instructor: Doina Precup
Email: dprecup@cs.mcgill.ca

Teaching assistants: TBA

Class web page: http://www.cs.mcgill.ca/∼dprecup/courses/ai.html

Outline

- Course overview and administrative details
- What is AI?
- Overview of AI history
- Examples of AI applications

Administrative Issues

- Class material:
  - Class notes (slides) and other references posted on the web page
  - I will try to record lectures (starting in Lecture 2) but I need a good technical solution for it
  - Russell and Norvig, Artificial Intelligence: A Modern Approach
  - Sutton and Barto, Reinforcement Learning: An Introduction
- Evaluation mechanisms:
  - Four individual assignments (30%): -10 points for each day late, no credit after 5 calendar days late
  - Project: implementation and written report (25%)
  - Midterm exam (15%)
  - Final exam (30%)

Project

- You will have to write a *game-playing program* (suggestions welcome)
- Major implementation component
  - Can be accomplished in C, Java, C++, some other language (e.g. Scheme, Python, Matlab, ...)
  - Sample code for a random player will be provided in C and Java
  - How you implement it and on what you focus is up to you
- *Written report* in good English is required.
- *Working code* is required.
- Competitive game playing quality is important
- Interesting ideas are important

Each student must work individually, for the whole class!

What is AI?

Quick answers from the media:

- What socially inept super-hackers do
- The opposite of natural stupidity
- Data, Hal, ...
- Deep Blue, Sony dancing robots, poker players, ...

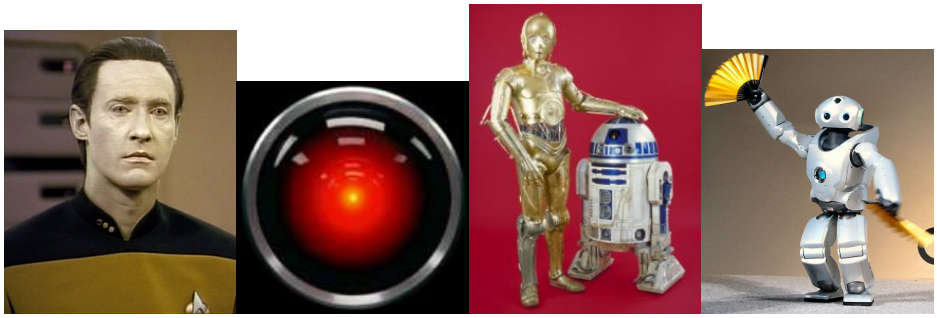

What is AI?

Quick answers from academia:

- Modeling human cognition using computers
- The study of problems that other CS folks do not know how to do
- *Cool stuff!*

  Game playing, machine learning, data mining, speech recognition, language understanding, computer vision, web agents, chatbots, robots, ...
- *Useful stuff!*

  Medical diagnosis, fraud detection, genome analysis, object identification, recommendation systems, space shuttle scheduling, information retrieval, ...
- AI methods are making a crucial impact in many areas of science and engineering

Which of the following do you consider intelligent?

- Sea Anemone
- Your roommate
- Dolphins
- Sturgeon fish
- Ants
- Cats

What is Biological Intelligence?

- A mix of *general-purpose* and *special-purpose* algorithms
- Examples of general-purpose algorithms
  - Learning new tasks
  - Memory formation and updating
  - Retrieving useful information from memory
- Examples of special-purpose algorithms:
  - Recognizing visual patterns
  - Recognizing sounds
  - Learning and reproducing language
- All are *integrated* seamlessly!

What is Artificial Intelligence?

- A mix of *general-purpose* and *special-purpose* algorithms
- Examples of general purpose algorithms
  - Learning algorithms
  - Theorem proving
- Examples of special purpose algorithms
  - Recognizing visual patterns
  - Language modelling
- All are integrated not quite so seamlessly...

Two Views of the AI Goal

AI is about duplicating what the (human) brain *DOES*.

Cognitive science, Turing test

AI is about duplicating what the (human) brain *SHOULD DO*.

RATIONALITY

Acting Like A Human: The Turing Test

- In 1950, Turing write the paper "Computing machinery and intelligence", in which he asked himself: "Can machines think?" And if so, how would be able to tell?
- He translated this into a different question: "Can machines behave intelligently?"
- *Turing test*: An operator interacts with either a human or an AI agent. Can he correctly guess which one?

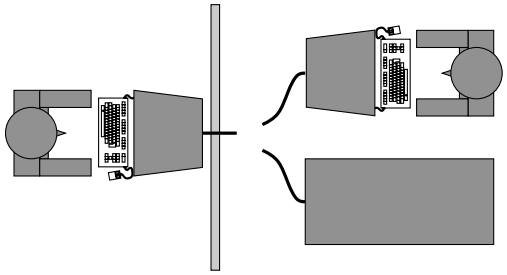

Turing Systems: Alice

- Alice: the winner of several Loebner prizes (2000, 2002, 2004), featured in NYTimes and various other media outlets.
- One judge out of 5 ranked it more believable than one of the humans
- About 30,000-40,000 canned phrases and sentences, with a very elaborate rule base and search engine for picking them

Turing Systems: Julia

- Julia (1990-1995): chatbot living in a Multi-User Dungeon (MUD) environment, expert in hockey
- MUDs have a rich but limited set of action verbs that users can perform on each other; users can create content (rooms, objects etc)
- Not quite natural language, but close
- Fooled at least one user into believing she was a real girl:

  barry says, 'Hey girlfriend are you out there'
  Julia says, 'Hello, barry.'
  barry says, 'Hello, I love Your name. It's sexy'
  Julia says, 'Me too, barry.'
  barry says, 'What are you doing there'
  Julia says, 'I'm just hanging around in The Hearts Room, barry.'
  barry says, 'Well come over to MY ROOM'
  Julia says, 'Could you say that again, it scrolled off my screen.'

Turing Systems: Watson

- In February 2011, Watson (developed at IBM) handily beat the two top human competitors in the game of Jeopardy
- Watson computes probabilities of different answers, based on evidence from the question, many knowledge sources, and past questions
- It learns from feedback (success/failure) to adjust how much it trusts different information sources

Is the Turing test still relevant?

- Of course!
- However, AI researchers do not spend much time/effort working towards it anymore
- Big problem: Turing test is *not reproducible*, or amenable to *mathematical analysis*
- It is more important to study the underlying principles of intelligence than to try to reproduce intelligent behavior

Two Views of the AI Goal

AI is about duplicating what the (human) brain *DOES*.

Cognitive science, Turing test

$\Longrightarrow$ AI is about duplicating what the (human) brain *SHOULD DO*.

RATIONALITY

This latter aspect will be our focus

Acting Rationally

- *Rational behavior* = doing the right thing
- *The right thing* = what is expected to maximize goal achievement, given the available information

  cf. Aristotle: "Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good"
- Note that doing the right thing may not involve thinking!

  E.g. Blinking and other reflexes are rational, by this definition
- But any thinking that takes place should be serving some purpose
- $\Longrightarrow$ This is the flavor of AI on which we focus

Components of an intelligent system

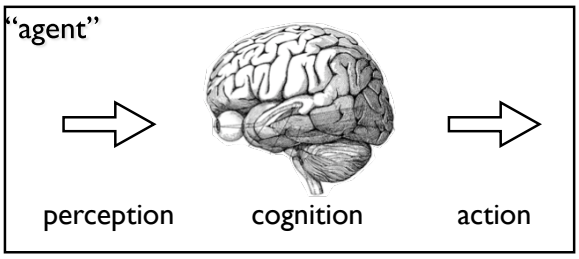

Craik (1943) defined 3 steps that any intelligent system must take:

1. Stimulus must be translated into an internal representation
2. The internal representation is manipulated to create new internal representations
3. The internal representation is translated into an action.

Components of an AI system

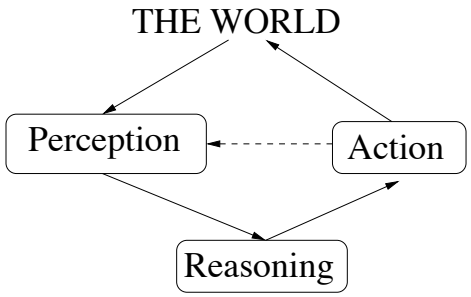

Rational Agents

- An *agent* is an entity that perceives and acts
- An agent can be modeled as a *function from percepts 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
- A *rational agent* implements this function so as to maximize a performance criterion

  E.g. goal achievement, resource minimization
- Caveat: computational limitations often make perfect rationality unachievable
- Our goal is to find the best function *given the limited computational resources* (memory, time)

AI Beginnings

- ENIAC: first super-computer, created in 1946
- Early focus in 1950s was on artificial intelligence:
  - Rosenblatt's perceptron
  - Samuel's checkers player

Dartmouth conference (1956)

"We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves."

Dartmouth key players

- John McCarthy (also known for LISP, time-sharing) coined the term "artificial intelligence"
- Marvin Minsky (co-founder with McCarthy of MIT AI lab, pioneer and detractor of neural nets)
- Claude Shannon (Bell Labs, inventor of information theory, and juggling robot designer)
- Herbert Simon (Nobel prize in economics) and Allen Newell presented "Logic theorist" (an automated theorem prover), one the first operating programs in AI

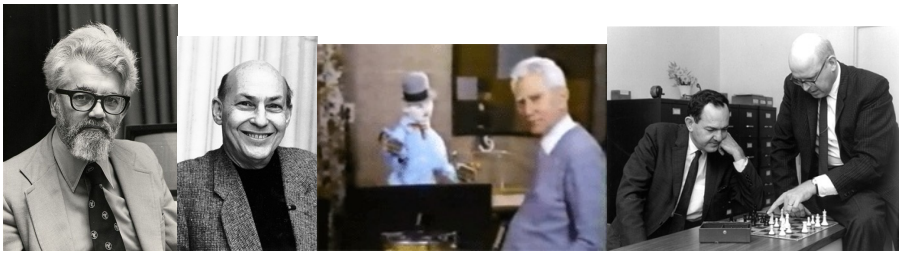

Early AI Hopes and Dreams

- Make programs that exhibit similar signs of intelligence as people: prove theorems, play chess, have a conversation
- Learning from experience was considered important
- Logical reasoning was key
- The research agenda was geared towards building *general problem solvers*

  But general problem solving is very hard!
- There was a lot of hope that natural language could be easily understood and processed

50 years later: Darpa Grand Challenge

- Darpa (US military research funding agency), offered a $2,000,000 prize for an automated driving competition
- Task: *Drive* through the Nevada desert 132 miles, start and finish specified the morning of the competition, *with no input from any human*
- In the 2004 competition, the best robot car crashed after 7.4 miles
- In the 2005 competition, five robots (out of 23) finished the race
- The winning robot, Stanley (from Stanford Univ.) finished in little under 7 hours
- Similar performance was achieved in 2007 in an "urban" driving challenge

From NY Times

The five robots that successfully navigated a 132-mile course in the Nevada desert last weekend demonstrated the re-emergence of artificial intelligence [...] The winning robot, named Stanley, covered the unpaved course in just 6 hours and 53 minutes without human intervention and guided only by global positioning satellite waypoints. The feat, which won a $2 million prize from the Pentagon Defense Advanced Research Project Agency, was compared by exuberant Darpa officials to the Wright brothers' accomplishment at Kitty Hawk, because it was clear that it was not a fluke. [...] The ability of the vehicles to complete a complex everyday task - driving - underscores how artificial intelligence may at last be moving beyond the research laboratory.

Stanley

- Built based on a Volkswagen car
- An array of sensors: cameras, laser range finders, radar, GPS
- Probabilistic reasoning and machine learning algorithms are the heart of the software
- The robot is capable of assessing how good the data is, based on prior training

Example AI System: Chess playing

- Deep Blue (IBM) defeats world champion Gary Kasparov in 1997
- Perception: advanced features of the board
- Actions: choose a move
- Reasoning: heuristics to evaluate board positions, search

Example AI System: Poker playing

- Polaris (Univ. Alberta) defends some of the best on-line poker players in 2008
- Perception: features of the game
- Action: choose a move (card to play, bet, etc)
- Reasoning: search and evaluation of possible moves, probabilistic reasoning, machine learning

Example AI System: Medical diagnosis

- Pathfinder (Heckerman, 1992): diagnosis of lymph node disease
- Perception: symptoms of a patient, test results
- Actions: suggest tests, make diagnosis
- Reasoning: Bayesian inference, machine learning, Monte Carlo simulation

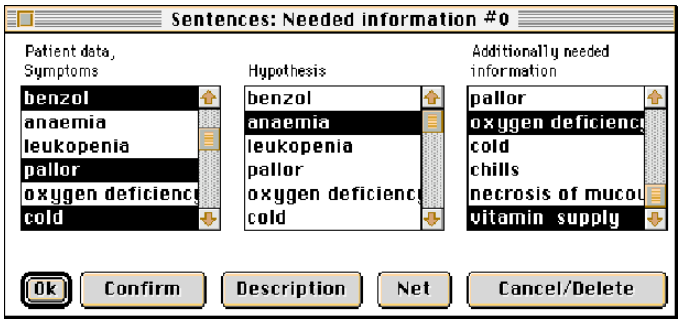

Example AI system: Chat bot

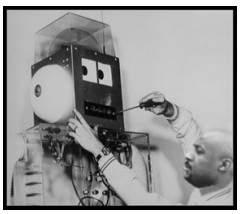

- Cobot (Isbell et al, 2000): learning chatbot in LambdaMOO
- Percepts: what users are present, how experienced they are, how much they interact with each other, what they do with each other, things being said...
- Actions: "verbs" that express actions on others and on objects - but restricted to a few

  E.g. proposing a conversation topic, introducing two users to each other, giving a piece of information
- Goals: maximize reward received from other users

Example AI system: Mind reading

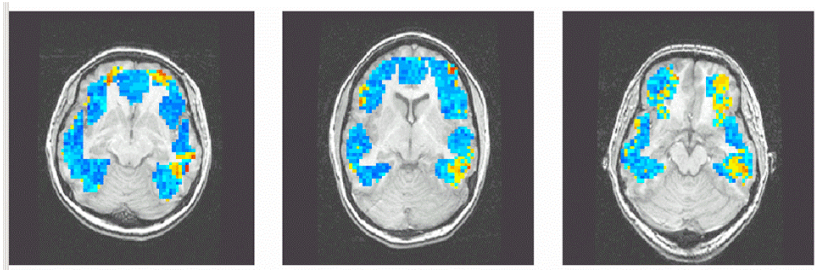

- Brain image analysis (T. Mitchell, 2008)
- Perception: brain imaging using fMRI technology
- Action: detect which word (e.g. "hammer", "apartment", ...) is read by the person
- Reasoning: machine learning

Example AI System: Automatic Driver

- ALVINN (D.Pomerleau, CMU, 1995): drives autonomously on the highway at 55 mph, for 21 miles
- Perception: digitized camera image of the road
- Actions: 64 different steering angles
- Reasoning: artificial neural network trained with backpropagation

Recent AI: Math to the Rescue!

- Heavy use of probability theory, decision theory, statistics
- Trying to solve *specific problems* rather than build a general reasoner
- AI today is a collection of sub-fields:
  - Perception (especially vision) is a separate community
  - Robotics is also largely separate
  - Deliberative reasoning is now the part named "AI"
  - Even within reasoning, different approaches evolved, with different styles and terminology
- A lot of progress was made in this way!
- Some recent efforts to put all this together

  E.g. STAIR project at Stanford

Which of the following can be done now?

- Play a decent game of table tennis
- Play a decent game of bridge
- Discover and prove a new mathematical theorem
- Write an intentionally funny story
- Give competent legal advice in a specialized area of law
- Translate spoken English into spoken Swedish in real time

Course Contents

We will overview *selected topics*:

- Search: A*, alpha-beta search, Monte Carlo tree search, local search
- Logic and planning
- Reasoning under uncertainty
- Learning
  - Supervised learning
  - Reinforcement learning
- Specialized areas
  - Robotics
  - Natural language processing