Introduction to Artificial Intelligence

CMPT 310: SPRING 2011
HASSAN KHOSRAVI

topics

- Intelligent Agents
- uninformed and informed search
- Constraint Satisfaction Problems
- Game playing
- First-order Logic
- Reasoning under uncertainty
- Bayesian networks
- Learning

Grading

- Evaluation will be based on pair programming and individual written assignments, as well as midterm and Final exams.
- 40% Assignments
  - 4 Assignments
- 20% Midterm
- 40% Final Exam
- 5% class participation
  - Short talks
  - Summaries

Book

- Required
  - Artificial Intelligence: A Modern Approach (2nd Edition), Stuart Russell, Peter Norvig,Prentice Hall, 2002.
- REFERENCE:
  - Computational Intelligence - A Logical Approach, David Poole et al, Oxford University Press.
  - Artificial Intelligence (5th Edition). Structures and Strategies for Complex Problem Solving, George Luger, Addison Wesley.

Academic Honesty

- Academic Honesty plays a key role in our efforts to maintain a high standard of academic excellence and integrity. Students are advised that ALL acts of intellectual dishonesty are subject to disciplinary action by the School; serious infractions are dealt with in accordance with the Code of Academic Honesty (T10.02) (http://www.sfu.ca/policies/teaching/t10-02.htm). Students are encouraged to read the
- School's policy information (http://www.cs.sfu.ca/undergrad/Policies/)

- Midterm: Friday 4th of March 2011
- Course Webpage: http://www.cs.sfu.ca/~hkhosrav/personal/310.html
- My office hours:
  - Wed 3:30 -5:00

Course Aims

- Assumption:
  - You will be going off to industry/academia
  - Will come across computational problems
    - requiring intelligence (in humans and computers) to solve
- Two aims:
  - Give you an understanding of what AI is
    - Aims, abilities, methodologies, applications, …
  - Equip you with techniques for solving problems
    - By writing/building intelligent software/machines

- Why use computers for intelligent behaviour at all?
  - They can do things better than us
  - Big calculations quickly and reliably

What is AI?

Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

The textbook advocates "acting rationally"

Acting Humanly

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" → "Can machines behave intelligently?"

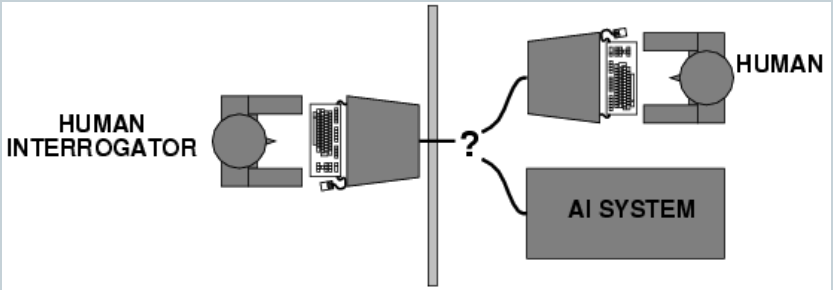

- Skills required:
  - Natural language processing
  - Knowledge representation
  - Automated reasoning
  - Machine learning

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
  - http://alice.pandorabots.com/

Captcha

- Completely Automated Public Turing test to tell Computers and Humans Apart

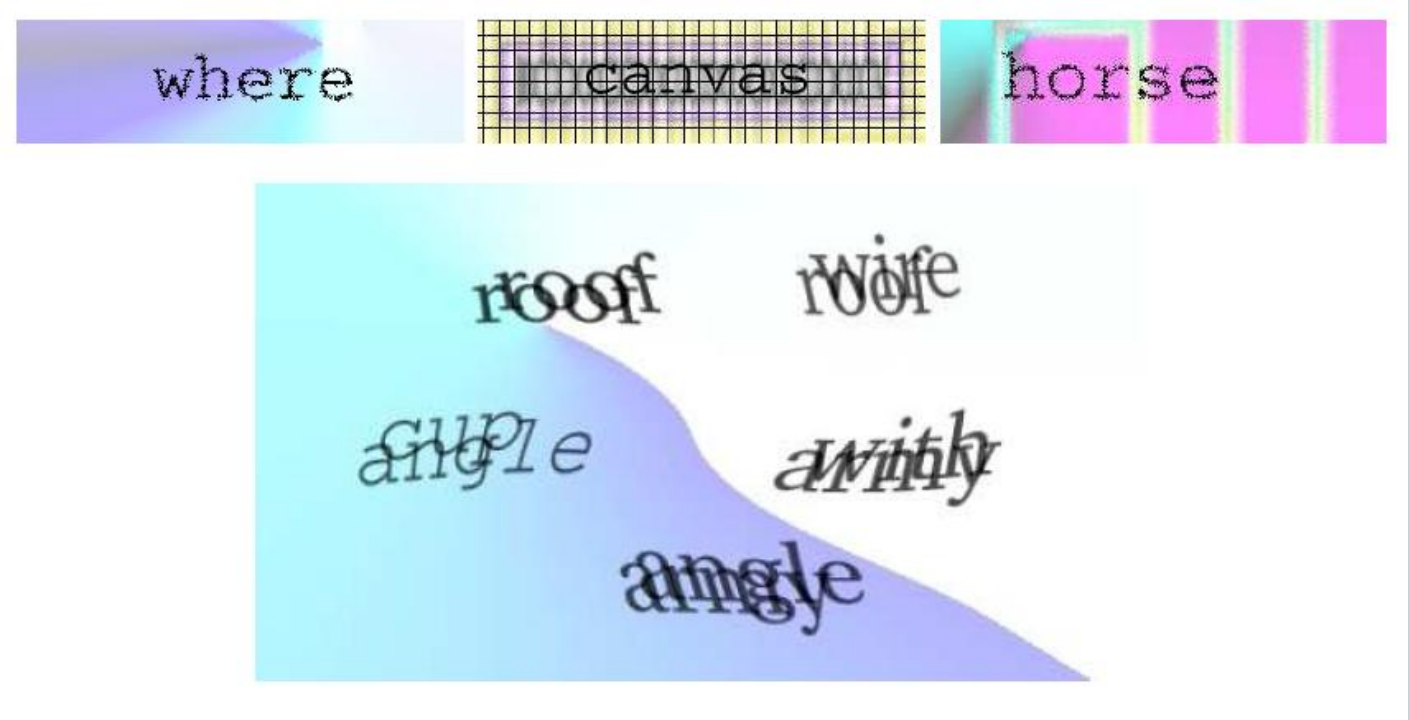

Thinking humanly: cognitive modeling

- Validate thinking in humans
- Cognitive science brings together computer models from AI and experimental techniques from psychology to construct the working of the human mind.

Thinking rationally

- Aristotle: what are correct arguments/thought processes?

- Several Greek schools developed various forms of logic:
  - notation and rules of derivation for thoughts;
- 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?

Action rationally

- Rational behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information
- Does it require thinking?
  - No – e.g., blinking reflex – but thinking should be in the service of rational action

Inspirations for AI

- Major question:
  - "How are we going to get a machine to act intelligently to perform complex tasks?"

Inspirations for AI

1. Logic
   - Studied intensively within mathematics
   - Gives a handle on how to reason intelligently
- Example: automated reasoning
   - Proving theorems using deduction
   - http://www.youtube.com/watch?v=3NOS63-4hTQ
- Advantage of logic:
   - We can be very precise (formal) about our programs
- Disadvantage of logic:
   - Theoretically possible doesn't mean practically achievable

Inspirations for AI

2. Introspection
   - Humans are intelligent, aren't they?
- Expert systems
   - Implement the ways (rules) of the experts
- Example: MYCIN (blood disease diagnosis)
   - Performed better than junior doctors

Inspirations for AI

3. Brains
   - Our brains and senses are what give us intelligence
- Neurologist tell us about:
   - Networks of billions of neurons
- Build artificial neural networks
   - In hardware and software (mostly software now)
- Build neural structures
   - Interactions of layers of neural networks
      - http://www.youtube.com/watch?v=r7180npAU9Y&NR=1

Inspirations for AI

4. Evolution
   - Our brains evolved through natural selection
- So, simulate the evolutionary process
   - Simulate genes, mutation, inheritance, fitness, etc.
- Genetic algorithms and genetic programming
   - Used in machine learning (induction)
   - Used in Artificial Life simulation

1.2 Inspirations for AI

5. Society
   - Humans interact to achieve tasks requiring intelligence
   - Can draw on group/crowd psychology
- Software should therefore
   - Cooperate and compete to achieve tasks
- Multi-agent systems
   - Split tasks into sub-tasks
   - Autonomous agents interact to achieve their subtask
      - http://www.youtube.com/watch?v=1Fn3Mz6f5xA&feature=related
      - http://www.youtube.com/watch?v=Vbt-vHaIbYw&feature=related

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:

$$[f: \mathcal{P}^{\star} \rightarrow \mathcal{A}]$$

- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
-
- computational limitations make perfect rationality unachievable
  → design best program for given machine resources

AI prehistory

- Philosophy
  - Can formal rules be used to draw valid conclusions?
  - Where does knowledge come from?
  - How does knowledge lead into action?
- Mathematics
  - What are the formal rules to draw valid conclusion?
  - How do we reason with uncertain information?
- Economics
  - How should we make decisions to maximize payoff?
  - How should we do this when others don't get along?
- Psychology
  - How humans and animals think?
- Computer
  - How can we build efficient computers
- Linguistics
  - How does language relate to thoughts
  - knowledge representation, grammar

Abridged history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence“
- 1950s Early AI programs, including Samuel's checkers
- 1965 Robinson's complete algorithm for logical reasoning
- 1966—73 AI discovers computational complexity
  Neural network research almost disappears
- 1969—79 Early development of knowledge-based systems
- 1980-- AI becomes an industry
- 1986-- Neural networks return to popularity
- 1987-- AI becomes a science
- 1995-- The emergence of intelligent agents

State-of-the-art

- Autonomous planning and scheduling
  - NASA's on-board program controlled the operations for a spacecraft a hundred million miles from Earth
- Game playing:
  - Deep Blue defeated the world chess champion Garry Kasparov in 1997
- Autonomous control
  - No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- Logistic planning
  - During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- Language understanding and problem solving
  - solves crossword puzzles better than most humans