Introduction to Artificial Intelligence

CMPT 310
OLIVER SCHULTE

Topics

- Intelligent Agents.
- Multi-agent decision making, game theory.
- Search
- Probability
  Reasoning under uncertainty
- Bayesian networks
- Learning
- Reinforcement Learning: Acting and Learning

Course Aims

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

Computers and Intelligence

- Why use computers for intelligent behaviour at all?
  - They can do some things better than us.
    - Big calculations quickly and reliably
    - Search through many options.
    - Avoid common mistakes.
  - Cognitive Science: building intelligent machines helps us understand the nature of intelligence.
- Informal Definition of AI: “Things that humans are good at, but computers are not (yet).”

Intelligent Behavior: Examples (?)

- Siri, Google Voice Search
- Google Translate
- AlphaZero, AlphaGo
- Soccer Goalie Robot
- roboclean action
- Learn to flip pancakes
- Watson Game Show
- Watson U.S. cities
- Self-Driving Car. No Hands Across America

AI Research

AI Research at SFU

- Various opportunities for funding:
  - NSERC Undergraduate Research Award. Full-time research in the summer.
  - Work-study SFU.
  - RAships from professors.
- AI researchers
  - Richard Vaughan. Robotics. Demo
  - Anoop Sarkar. Fred Popowich.Linguistics, Machine Translation.
  - James Delgrande. Logic and AI.
  - David Mitchell. Eugenia Ternovska. Logic, Theorem Proving, Constraint Satisfaction.
  - Greg Mori. Vision, Tracking.
  - Oliver Schulte. Machine Learning, Network Analysis.
  - Hang Ma. Multi-agent Planning
  - Ke Li . Machine Learning

What is AI?

Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

- Modern view (ie. Since 1990s): Acting rationally.
- In economics and statistics, since the 1920s or earlier.

Thinking

HUMANLY AND RATIONALLY

Thinking humanly: cognitive modeling

- Validate by comparing with thinking in humans

- Cognitive science brings together
  - computer models from AI
  - experimental techniques from psychology

to construct the working of the human mind.

- Example of MIT Research

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.

Acting

HUMANLY AND RATIONALLY

THE TURING TEST

THE CHINESE ROOM

Acting Humanly

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" → "Can machines behave intelligently?"
- The Imitation Game
- Turing's article

- 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
  - Mitsuku
  - Loebner Prize

Captcha

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

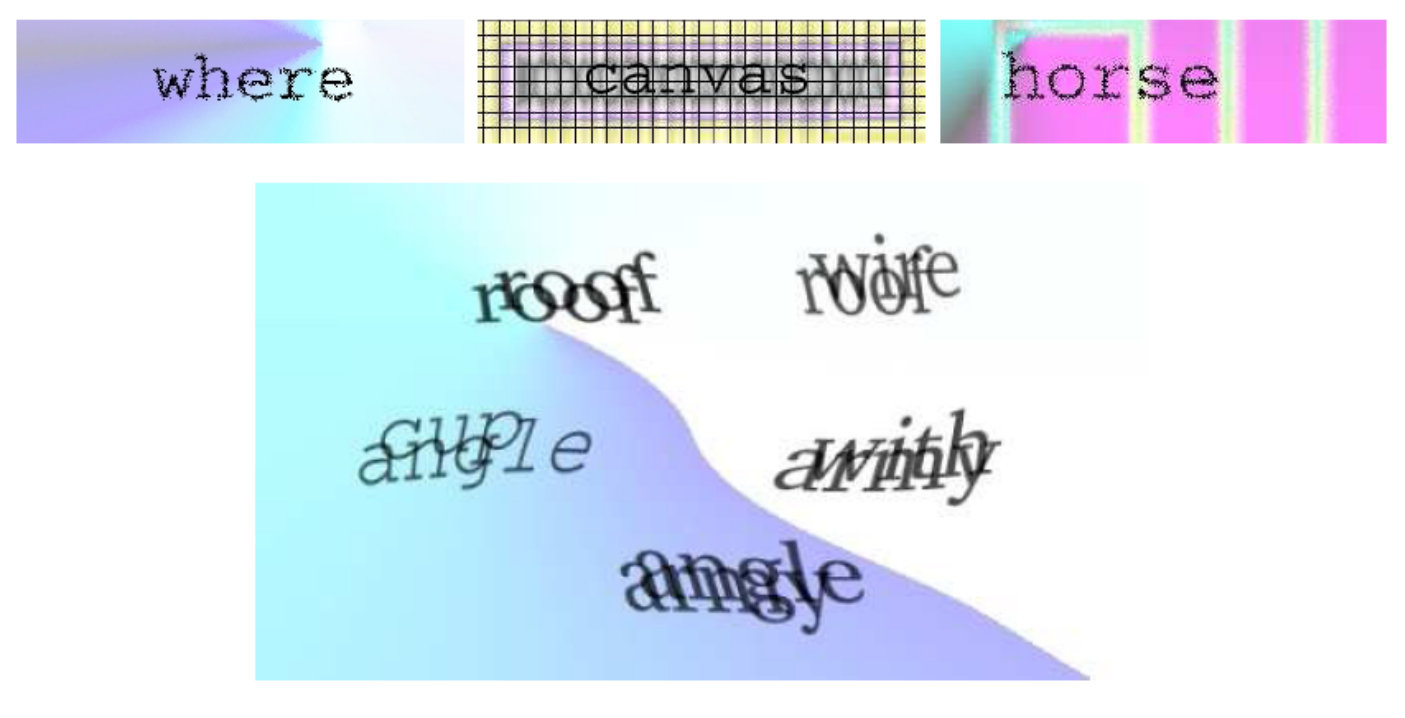

Searle's Chinese Room

- Person sits in Chinese room
- The room has a book with rules for mapping Chinese input sentences to output sentences.
- This allows the person in the room to carry on a conversation with a Chinese speaker.

The Translation Room

German	Spanish
Schach	ajedrez
königliche	real
das	el
Spiel	juego
ist	es

- Jane is given a translation table like the one shown
- We ask her to translate “Schach ist das königliche Spiel” into Spanish
- Her answer “ajedrez es el juego real”
- Correct!
- Does this mean that she speaks German and Spanish?

Chinese Room Conclusion

- Modest conclusion: it is possible for a program to engage in speech recognition, conversation, translation without understanding language.
- Stronger conclusion (controversial): it is possible for a program to pass the Turing test without understanding language.
- Strongest conclusion (very controversial): computer programs can only apply rules, not understand the meaning of language.
- Infinite AI Loop

Rational Action

- Rational behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information
- What is the right thing?
  - Whatever the designer specifies (in the current "standard" model)
  - This can be problematic for complex applications (e.g. self-driving car)

Acting vs. Thinking

Does acting require thinking?

- Not always.
  - Iroboclean? Dyson cleaner?
  - blinking reflex.
  - Insects. Do dung beetles think?
  - Siri? Watson?
- What are the advantages of thinking? Why would a thinking animal have evolved?
- Thinking seems to lead to
  - **flexibility** and
  - **robustness**.

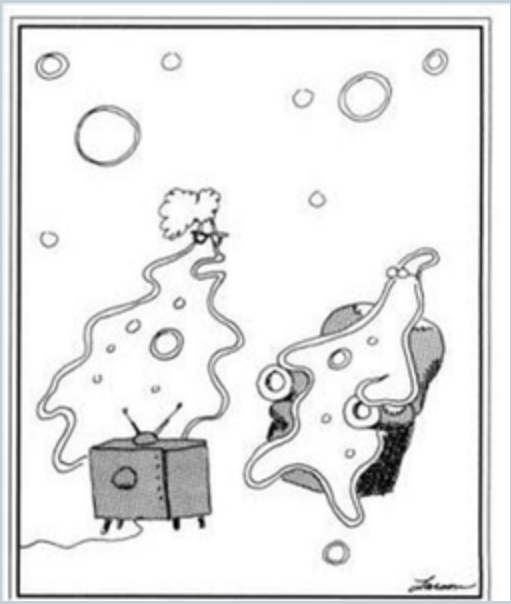

Stimulus, response! Stimulus, response! Don't you ever THINK?"

State-of-the-art

- Autonomous planning and scheduling
  - NASA's Mars Rover 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
  - Alphago defeated top player in 2016
- Autonomous control
  - Self-driving cars
- Language understanding and problem solving
  - solves crossword puzzles better than most humans
  - automated speech assistant (Siri)

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:
   - Not designed for uncertainty.

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
      - Neurons Firing

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
   - 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
      - Used in movies too.

Decision Theory and Rational Agents

- For any given class of environments and task, we seek the agent (or class of agents) with the **best performance**.
- The primary goal is performance, *not*
  - thinking
  - consciousness
  - intelligence
  - autonomy
  - These may be means to achieve performance.
- Performance measure is usually given by the user or engineer. Economics: rationality = maximize utility (performance).
- computational limitations make perfect performance unachievable
  - design best program for given machine resources