Introduction to Artificial Intelligence

Amitava Chattopadhyay

- This class is a broad introduction to artificial intelligence (AI)
  - AI is a very broad field with many subareas
    - We will cover many of the primary concepts/ideas
    - But in 10 weeks we can't cover everything
    - Other classes in AI you may want to consider:
      - Belief Networks, 276
      - Winter: Probabilistic Learning, 274A
      - Spring: Machine Learning, 273A
  - If you have taken another class (e.g., undergrad) in AI, you may want to consider waiving this class and taking a more specialized AI class (feel free to ask me about this).

- What is intelligence? What is artificial intelligence?
- A very brief history of AI
  - Modern successes: Stanley the driving robot
- An AI scorecard
  - How much progress has been made in different aspects of AI
- AI in practice
  - Successful applications
- The rational agent view of AI

- **What is artificial intelligence?**

  It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.

- **Yes, but what is intelligence?**

  Intelligence is the computational part of the ability to achieve goals in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines.

- **Isn't there a solid definition of intelligence that doesn't depend on relating it to human intelligence?**

  Not yet. The problem is that we cannot yet characterize in general what kinds of computational procedures we want to call intelligent. We understand some of the mechanisms of intelligence and not others.

What's involved in Intelligence?

- Ability to interact with the real world
  - to perceive, understand, and act
  - e.g., speech recognition and understanding and synthesis
  - e.g., image understanding
  - e.g., ability to take actions, have an effect
- Reasoning and Planning
  - modeling the external world, given input
  - solving new problems, planning, and making decisions
  - ability to deal with unexpected problems, uncertainties
- Learning and Adaptation
  - we are continuously learning and adapting
  - our internal models are always being "updated"
    - e.g., a baby learning to categorize and recognize animals

History of AI

- 1943: early beginnings
  - McCulloch & Pitts: Boolean circuit model of brain
- 1950: Turing
  - Turing's "Computing Machinery and Intelligence“
- 1956: birth of AI
  - Dartmouth meeting: "Artificial Intelligence“ name adopted
- 1950s: initial promise
  - Early AI programs, including
  - Samuel's checkers program
  - Newell & Simon's Logic Theorist
- 1955-65: “great enthusiasm”
  - Newell and Simon: GPS, general problem solver
  - Gelertner: Geometry Theorem Prover
  - McCarthy: invention of LISP

Success Stories

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- AI program proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- 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
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- `Proverb` solves crossword puzzles better than most humans
- Robot driving: DARPA grand challenge 2003-2007
- 2006: face recognition software available in consumer cameras

Example: DARPA Grand Challenge

- Grand Challenge
  - Cash prizes ($1 to $2 million) offered to first robots to complete a long course completely unassisted
  - Stimulates research in vision, robotics, planning, machine learning, reasoning, etc
- 2004 Grand Challenge:
  - 150 mile route in Nevada desert
  - Furthest any robot went was about 7 miles
  - … but hardest terrain was at the beginning of the course
- 2005 Grand Challenge:
  - 132 mile race
  - Narrow tunnels, winding mountain passes, etc
  - Stanford 1st, CMU 2nd, both finished in about 6 hours
- 2007 Urban Grand Challenge
  - This November in Victorville, California

Consider what might be involved in building a computer like Hal….

- What are the components that might be useful?
  - Fast hardware?
  - Chess-playing at grandmaster level?
  - Speech interaction?
    - speech synthesis
    - speech recognition
    - speech understanding
  - Image recognition and understanding ?
  - Learning?
  - Planning and decision-making?

Can we build hardware as complex as the brain?

- How complicated is our brain?
  - a neuron, or nerve cell, is the basic information processing unit
  - estimated to be on the order of 10^{12} neurons in a human brain
  - many more synapses (10^{14}) connecting these neurons
  - cycle time: 10^{-3} seconds (1 millisecond)

- How complex can we make computers?
  - 10^8 or more transistors per CPU
  - supercomputer: hundreds of CPUs, 10^{12} bits of RAM
  - cycle times: order of 10^{-9} seconds

- Conclusion
  - YES: in the near future we can have computers with as many basic processing elements as our brain, but with
    - far fewer interconnections (wires or synapses) than the brain
    - much faster updates than the brain
  - but building hardware is very different from making a computer behave like a brain!

Can Computers Talk?

- This is known as “speech synthesis”
  - translate text to phonetic form
    - e.g., “fictitious” -> fik-tish-es
  - use pronunciation rules to map phonemes to actual sound
    - e.g., “tish” -> sequence of basic audio sounds
- Difficulties
  - sounds made by this “lookup” approach sound unnatural
  - sounds are not independent
    - e.g., “act” and “action”
    - modern systems (e.g., at AT&T) can handle this pretty well
  - a harder problem is emphasis, emotion, etc
    - humans understand what they are saying
    - machines don’t: so they sound unnatural
- Conclusion:
  - NO, for complete sentences
  - YES, for individual words

Can Computers Recognize Speech?

- Speech Recognition:
  - mapping sounds from a microphone into a list of words
  - classic problem in AI, very difficult
    - “Lets talk about how to wreck a nice beach”
    - (I really said “_____”)
- Recognizing single words from a small vocabulary
    - systems can do this with high accuracy (order of 99%)
    - e.g., directory inquiries
      - limited vocabulary (area codes, city names)
      - computer tries to recognize you first, if unsuccessful hands you over to a human operator
      - saves millions of dollars a year for the phone companies

Recognizing human speech (ctd.)

- Recognizing normal speech is much more difficult
  - speech is continuous: where are the boundaries between words?
    - e.g., “John’s car has a flat tire”
  - large vocabularies
    - can be many thousands of possible words
    - we can use **context** to help figure out what someone said
      - e.g., hypothesize and test
      - try telling a waiter in a restaurant:
        “I would like some dream and sugar in my coffee”
  - background noise, other speakers, accents, colds, etc
  - on normal speech, modern systems are only about 60-70% accurate

- Conclusion:
  - NO, normal speech is too complex to accurately recognize
  - YES, for restricted problems (small vocabulary, single speaker)

Can Computers Understand speech?

- Understanding is different to recognition:
  - "Time flies like an arrow"
    - assume the computer can recognize all the words
    - how many different interpretations are there?
      - 1. time passes quickly like an arrow?
      - 2. command: time the flies the way an arrow times the flies
      - 3. command: only time those flies which are like an arrow
      - 4. "time-flies" are fond of arrows

Can Computers Learn and Adapt ?

- Learning and Adaptation
  - consider a computer learning to drive on the freeway
  - we could teach it lots of rules about what to do
  - or we could let it drive and steer it back on course when it heads for the embankment
    - systems like this are under development (e.g., Daimler Benz)
    - e.g., RALPH at CMU
      - in mid 90's it drove 98% of the way from Pittsburgh to San Diego without any human assistance
  - **machine learning** allows computers to learn to do things without explicit programming
  - many successful applications:
    - requires some "set-up": does not mean your PC can learn to forecast the stock market or become a brain surgeon
- Conclusion: YES, computers can learn and adapt, when presented with information in the appropriate way