CPS 270
Introduction

Ron Parr

Contact Information

- Professor
  - Ron Parr
  - D209 LSRC, parr@cs, 660-6537
  - Office hours: T,Th 2:30-3:30
- TA
  - Monika Schaeffer
  - N01 North Building, monika@cs, 660-4001
  - Office hours: Tuesday, 4:15 - 5:15
    Wednesday 3:15 - 5:15

About Me

- My fifth year at Duke
- Bachelor's degree in philosophy
  - Philosophy of mind
- Ph.D. in computer science
  - Hierarchical planning under uncertainty
- Current interests:
  - Planning under uncertainty
  - Probabilistic reasoning
  - Reinforcement learning
  - Mobile robotics

Requirements

- Prior AI experience useful but not required
- Reasonable programming skills:
  - C (mostly), other languages OK
- Some mathematical sophistication
  - Short proofs
  - Basic probability concepts
  - Basic algorithmic concepts
    - Complexity
    - Analysis of algorithms

Grading

- Homeworks (10%)
- Midterm (30%)
- Final (30%)
- 2 Projects (15% each)

- Discussion is encouraged.
- Final work must be your own.

Why Study AI?

- Important innovations that grew out of AI
  - Linked list manipulation (Lisp)
  - Timesharing
  - X
  - Formalization of search techniques
  - Heuristics for intractable (NP hard) problems
  - Pattern recognition methods
  - Speech recognition
  - Machine vision
- Cool tools
- Cool applications

Cool AI Applications

- AI is lurking in more places than you think:
  - My PDA
  - This Program
  - E-commerce
  - Dragon Naturally Speaking
  - Deep Blue
  - Mobile robotics
  - Space exploration
  - Logistics planning

Future Challenges for AI

- Systems
  - Networking
  - Self healing systems
  - Defending against attacks
- More engaging computer games
- Understanding the human genome
- Better medical treatment and diagnosis ☺
- Better killing machines ☹

So, what is this AI stuff?

- Make machines *think* like humans
- Is this enough?
- Is this too much?

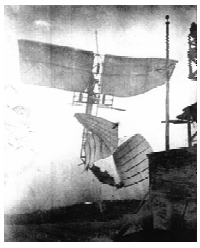

Turing Test

- Computer must be indistinguishable from a human based upon written exchanges
  (Actually more complicated than this)
  - Does this imply intelligence?
  - How could the computer cheat?
  - Does intelligence imply a certain type of computation?
  - Could an intelligent machine still fail the test?
- Does our notion of intelligence transcend our concept of humanity?

Ideal Intelligence

- Intelligence means making optimal choices
- Is anything truly intelligent?
- How do we define optimality?
- Is there a more modest goal?

Our Compromise

In Artificial Intelligence, we study algorithms for tasks typically associated with human intelligence to gain insight into the general question of intelligence.

The Moving Target

- What is human intelligence?
  - At one time, calculating ability was prized
    - Now it is deprecated
    - Calculators permitted earlier and earlier in school
  - Chess was once viewed as an intelligent task
    - Now, massively parallel computers use not very intelligent search procedures to beat grand masters
    - Some say Deep Blue wasn't AI
  - Learning once thought uniquely human
    - Now it's a well-developed theory
    - Best backgammon player is a learning program

Artificial Flight

- Even seemingly unambiguous terms such as "flight" were subject to biological chauvinism.

- Demonstrable, unambiguous success ended chauvinism

Intelligence: A web of abilities

- Intelligence is hard to define in isolation
- We are an odd mixture of special purpose and general purpose hardware
  - Special purpose
    - Recognizing visual patterns
    - Learning and reproducing language
  - General Purpose
    - Theorem proving
    - Learning and excelling at new tasks
- Seamless integration

Early Efforts: General

- Good news:
  - Many problems can be formalized as instances of
    - Search
    - Logical deduction
  - The space of all proofs is a (somewhat) searchable space
  - A theorem proving engine + a knowledge base provide a satisfying picture or reasoning, knowledge and learning
    - Tell PC:
      - All men are mortal
      - Socrates is a man
    - Ask:
      - Is Socrates mortal?

Bad news for general methods

- Searching in proof space is hard
- Representing knowledge is hard
  - What is a chair?
- Knowledge interconnected in strange ways
  - Chairs
  - People
  - Gravity
  - Customs…
- Early efforts were too general, ambitious

Early Efforts: Special Purpose Methods

- Neural networks attempted to reproduce the function of human neurons
- Proverbial wing-flapping flying machine?
- Success at reproducing low-level tasks
  - Pattern recognition, associative memory
- Nearly became a religion
- Huge gap between low level and high level
- Early efforts were too specific

Middle of the Road (This Class)

- Study broad classes of problems
- Restrict problem somewhat:
  - Develop a crisp input specification
  - Develop a well-defined success criterion
- Develop results with
  - provable properties
  - broad applicability
- Extract, study underlying principles
- Note: Some say modern AI is currently too specific

Modern AI: The worst case

- We fail to solve the general AI problem, but solve a bunch of less general problems
- We leave behind a many interesting and powerful tools for solving these problems
  (This is already happening.)
- Nobody gets upset with us for promising too much

Modern AI: The best case

- We achieve seamless integration
- Use pattern recognition and learning to shape and guide more general reasoning processes
- Work required for each new success goes down
- We can no longer explain away the intelligence we have created by calling it “artificial”.

Outline

- Search
- Intro probability, decision theory
- MDPs
- Reinforcement learning
- Games
- Logic
- CSPs
- Planning
- Advanced probability
- Learning
- Bayes nets
- Advanced decision theory

Major Topics *Not* Covered

- Natural Language
- Vision