CPS 170
Introduction

Ron Parr

Contact Information

- Professor
  - Ron Parr
  - D209 LSRC, parr@cs.duke.edu, 660-6537
  - Office hours: Tuesday 1-2, Wednesday 2-3

- TA
  - Zheng Li
  - D125 LSRC, zheng@cs.duke.edu, 660-6576
  - Office hours: Tuesday 2-3, Wednesday 3-4

About Me

- My sixth year at Duke
- Bachelor's degree in philosophy (Princeton)
  - Philosophy of mind
- Ph.D. in computer science (Berkeley)
  - Hierarchical planning under uncertainty
- Current interests:
  - Planning under uncertainty
  - Probabilistic reasoning
  - Game theory
  - Reinforcement learning
  - Robotics
  - Sensing & Vision

Requirements

- Good programming skills:
  - C
  - Other languages OK, but will require extra work from you b/c you won't be able to use our code

- Prequesites
  - Short proofs
  - Basic probability concepts
  - Basic algorithmic concepts
    - Complexity - O()
    - Analysis of algorithms
  - Math
    - Partial derivatives

Major Topics Covered

- Search
  - A*, Games, SAT, CSPs
- Logic and Knowledge Representation
  - Propositional Logic
  - First Order Logic
- Planning
  - Classical, stochastic
- Reasoning under uncertainty
  - Bayes nets, decision theory, HMMs, tracking
- Introduction to robotics
- Learning
  - Decision trees, Neural nets, Reinforcement learning

Major Topics *Not* Covered

- Natural Language

- Vision

Class Mechanics

- Textbook: *Artificial Intelligence, A Modern Approach*, Russell & Norvig (second edition)
- Homeworks: 25%
  - Discussion OK, write-up must be your own
- Projects: 25%
  - Discussion OK, coding, write-up must be your own
- Midterm: 25%
  - Closed book, in class, no collaboration
- Final: 25%
  - Closed book, finals week, no collaboration

Why Study AI?

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

Cool AI Applications

- AI is lurking in more places than you think:
  - PDAs
  - This Program
  - E-commerce
  - Voice/language recognition
    - Voice jail
    - My car
    - Dragon naturally speaking
  - Deep Blue
  - Mobile robotics
  - Space exploration
  - Logistics planning

So, what is this AI stuff?

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

Turing Test

- Computer must be indistinguishable from a human based upon written exchanges
  - 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?

Ron's Compromise Definition

Artificial Intelligence is the task of developing general purpose algorithms with which machines can accomplish tasks which, if performed by a human, would be considered indicative 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
  - Knowledge base + theorem proving 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
- Even if searching were easy, representing knowledge is hard
  - What is a chair?
- Knowledge is interconnected in subtle ways
  - Chairs
  - People
  - Gravity
  - Customs…

Early Efforts: Special Purpose Methods

- Neural networks attempted to reproduce the function of human neurons
- Wing-flapping mechanical flying machine?
- Success at reproducing low-level tasks
  - Pattern recognition, associative memory
- Huge gap between low and high level
- Nearly became a religion

Overpromising and the AI Winter

- Years of
  - Naïve optimistic
  - Unrealistic assessments of challenges
  - Poor scientific/academic discipline
- Lead to (early 90's)
  - Backlash
  - Reduced government funding
  - Reduced investment from industry
  - The "AI Winter"

AI Moving Forward

- 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 and study underlying principles behind successful methods

Conclusion

- We want to solve hard problems that would traditionally require human-level intelligence.
     (Most we consider are at least NP-hard.)

- We want to be good computer scientists, so we force ourselves to use well-defined input/output specifications.

- We aim high, but we let ourselves simplify things if it allows us to produce a general-purpose tool with well-understood properties.