

CPS 270 Introduction

Ron Parr

Contact Information

- Professor
 - Ron Parr
 - D209 LSRC, parr@cs, 660-6537
 - Office hours: T,W 10-11
- TA
 - Zheng Li
 - D125 LSRC, zheng@cs.duke.edu, 660-6576
 - Office hours: M 2-3, W 3-4, Th 1:30-2:30

About Me

- My seventh 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
 - matlab
- Some mathematical sophistication
 - Short proofs
 - Basic probability concepts
 - Basic algorithmic concepts
 - Complexity
 - Analysis of algorithms

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: 10%
 - Discussion OK, write-up must be your own
- Projects: 30%
 - Discussion OK, coding, write-up must be your own
- Midterm: 30%
 - Closed book, in class, no collaboration
- Final: 30%
 - 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



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 ☹️

About AI and the military

- Any technology can be used for good or harm
- Improved surveillance, recognition, tracking, control and guidance reduce:
 - Need for big explosions
 - Unintended death and destruction
- Don't worry about Terminator... Yet...
- Examples:
 - Soviet Missile Technology
 - DEFCON III in 1973 Arab Israeli Conflict



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 (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
- Intelligence is like...

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
- 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

Overpromising and the AI Winter

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

Modern AI

- 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

Eye on the prize

- Some senior researchers argue that modern AI has become too specific
- Making such claims is a privilege of tenure 😊
- Important point: Must not lose touch with the goal of general intelligence

Outline

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