

Foundations of Artificial Intelligence

Introduction

CS472 – Fall 2007
Thorsten Joachims

Reading: R&N Chapter 1.

Lecture and Project

CS472: Foundations of Artificial Intelligence

- Instructor: Prof. Thorsten Joachims
- Lecture
- Introduction to AI techniques
- Agents, Search, CSP, Machine Learning, Planning, Reasoning, Knowledge Representation, Reinforcement Learning

CS473: Practicum in Artificial Intelligence

- Instructor: Prof. Hod Lipson
- Project
- Hands-on experience with AI methods
- Meets separate from CS472
- CS472 is co-requisite

Contact

Where: Philips 219

When: Mon, Wed, Fri 11:15-12:05

Professor: Thorsten Joachims, Computer Science

Email: tj@cs.cornell.edu

URL: www.cs.cornell.edu/People/tj

Office Hours: 4153 Upson, We 1:30-2:30

Course web site: www.cs.cornell.edu/courses/CS472/2007fa/

Teaching Assistants:

Thomas Finley, Alexander Chao, Ilya Sukhar,

Griffin Dorman, Rick Keilty

→Office hours posted on-line.

Syllabus

Problem solving

principles of search, uninformed search, informed (“heuristic”) search, constraint satisfaction, local search, genetic algorithms, game playing

Learning

inductive learning, decision tree learning, statistical approaches, support vector machines, kernels, neural networks

Knowledge representation and reasoning

knowledge bases and inference, propositional and first-order logic, theorem-proving, planning

Natural language understanding

syntactic processing, ambiguity resolution, text understanding

General Information

Text:

Artificial Intelligence: A Modern Approach
Russell and Norvig, Prentice-Hall, Inc., second edition.

Class Notes and Handouts:

Available from course web site / CMS

Homework: approx. 6 homework assignments

Examinations: two prelims, one final exam.

Prerequisites:

programming and data structures (CS211 or CS312),
basic discrete math (CS280), basic linear algebra

Grading Policy

Assignments	45%
Prelims	25%
Final Exam	25%
Participation/Interest	5%

Late assignments drop 5% per each late day.

Roughly: A=93-100; B=83-87; C=73-77; D=63-67; F= below 60

Today's Lecture

What is Artificial Intelligence (AI) anyway?

- the components of intelligence

The Current Frontier

- recent achievements

Current Challenges

- what makes AI problems hard?

What is Intelligence?

Intelligence:

- "the capacity to learn and solve problems" (Webster dictionary)
- the ability to think and act rationally

Goal in Artificial Intelligence:

- build and understand intelligent systems/agents
- synergy between
 - philosophy, psychology, and cognitive science
 - computer science and engineering
 - mathematics and physics

What is involved in Intelligence

A) Ability to interact with the real world

- to perceive, understand, and act
- speech recognition, understanding, and synthesis
- image understanding (computer vision)

B) Reasoning and Planning

- modeling the external world
- problem solving, planning, and decision making
- ability to deal with unexpected problems, uncertainty

C) Learning and Adaptation

- we are continuously learning and adapting
- Also: we want systems that adapt to us!
- Major thrust of industry research.

What is Artificial Intelligence

Rich and Knight: the study of how to make computers do things which, at the moment, people do better.

Handbook of AI: the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behavior - understanding language, learning, reasoning, solving problems, etc.

Dean, Allen and Aloimonos: the design and study of the computer programs that behave intelligently.

Russell and Norvig: the study of [rational] agents that exist in an environment and perceive and act.

Different Approaches

I Building exact models of human cognition

- view from psychology and cognitive science

II The logical thought approach

- emphasis on "correct" inference

III Building rational "agents"

- agent: something that perceives and acts
- emphasis on developing methods to match or exceed human performance [in certain domains]. Example: Deep Blue.

Our focus is on III (most recent progress).

Goals in AI

Engineering Goal

To solve real-world problems. Build systems that exhibit intelligent behavior.

Scientific Goal

To understand what kind of computational mechanisms are needed for modeling intelligent behavior.

Turing Test

- Interrogator asks questions of two “people” who are out of sight and hearing. One is a person; the other is a machine.
- 30 minutes to ask whatever he or she wants.
- Task: to determine, only through the questions and answers typed into a computer terminal, which is which
- If can't reliably distinguish the human from the computer, then the computer is deemed intelligent.

Artificial intelligence is the enterprise of constructing an artifact that can pass the Turing test.

Objections to Turing Test?

Newell and Simon [1976]

- Turing test is as much a test of the judge as it is of the machine.
- Promotes the development of artificial con-artists, not artificial intelligence (Loebner competition).

The Current Frontier

Interesting time for AI

Deep Blue vs. Kasparov (May, '97)

- first match won against world-champion
- “intelligent & creative” play
- 200 million board positions per second

Kasparov: “**I could feel - I could smell - a new kind of intelligence across the table.**”

... still understood 99.9% of Deep Blue's moves.

Intriguing issue: How does human cognition deal with the **combinatorics** of chess?

Different Algorithm, Similar Behavior

Drew McDermott (New York Times, May, 1997):

Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings.
<ftp://ftp.cs.yale.edu/pub/mcdermott/papers/deepblue.txt>

The brain

- a neuron is the basic processing unit (10^{11})
- many more synapses (10^{14}) connect the neurons
- cycle time: 10^{-3} seconds (1 millisecond)

How complex can we make computers?

- 10^8 (Itanium) or more transistors per CPU
- supercomputer: thousands of CPUs, 10^{11} bits of RAM
- cycle times: order of 10^{-9} seconds

Examples, cont.

- First “**creative**” proof by computer (Nov, '96)
 - 60 year open problem.
 - Robbins' problem in finite algebra.

Qualitative difference from previous brute-force results.

Does technique generalize?

(Our own expert: Robert Constable.)

Machine Learning

- **TD Gammon** (Tesauro 1993; 1995)
 - World champion level but learns from scratch by playing millions of games against itself!
 - Has changed human play
- **ALVINN** (Pomerleau 1993)
 - Neural net used to steer vehicle in coast-to-coast highway driving
 - Speeds of up to 90 mph
 - DARPA Grand Challenge / Urban Challenge

Natural Language Processing

BOGOTA, 9 JAN 90 (EFE) - RICARDO ALFONSO CASTELLAR, MAYOR OF ACHI, IN THE NORTHERN DEPARTMENT OF BOLIVAR, WHO WAS KIDNAPPED ON 5 JANUARY, APPARENTLY BY ARMY OF NATIONAL LIBERATION (ELN) GUERRILLAS, WAS FOUND DEAD TODAY, ACCORDING TO AUTHORITIES. CASTELLAR WAS KIDNAPPED ON 5 JANUARY ON THE OUTSKIRTS OF ACHI, ABOUT 850 KM NORTH OF BOGOTA, BY A GROUP OF ARMED MEN, WHO FORCED HIM TO ACCOMPANY THEM TO AN UNDISCLOSED LOCATION.

Summary:

- **Date:** 05 JAN 90
- **Location:** COLOMBIA: BOLIVAR (DEPARTMENT): ACHI (TOWN)
- **Type:** KIDNAPPING
- **Weapon:** *
- **Victim:** "RICARDO ALFONSO CASTELLAR" (MAYOR OF ACHI)
- **Perpetrator:** "GROUP OF ARMED MEN"
- **Organization:** "ARMY OF NATIONAL LIBERATION (ELN)"

Challenges Ahead

- Note that the examples we discussed so far all involve quite **specific tasks**.
- The systems lack a level of **generality** and **adaptability**. They can't easily (if at all) switch **context**.
- Key issue: **knowledge acquisition** bottleneck
 - Lack of general **commonsense** knowledge.
 - CYC project small (Doug Lenat et al.). Attempt to encode millions of facts.

Goal of This Course

- Introduce you to the kinds of problems studied in AI.
- Introduce you to a set of key methods for
 - problem solving,
 - knowledge representation and reasoning,
 - learning,
 - natural language understanding.
- Teach you about the applicability and limitations of these methods.

Reading: Chapter 1, R&N.