

Lecture 1: Foundations of Artificial Intelligence

CS472/3 — Fall 1999

Slide CS472-1

General Information for CS472/CS473

Where: Olin 155

When: Mon, Wed, Fri 11:15–12:05

Professor: Bart Selman

Phone: 5-5643 **Email:** selman@cs.cornell.edu

Office Hours: 4148 Upson: Mon 2–3 and Wed 2–3
(or by appointment)

Teaching Assistants: Alin Dobra and Dongmin Liang

Email: dobra@cs.cornell.edu / liang@cornell.edu

Office Hours: TBA

Slide CS472-2

Text:

Artificial Intelligence: A Modern Approach
Russell and Norvig, Prentice-Hall, Inc., 1995.

Class Notes and Handouts:

TBA (CS web site is being revised)

Homework: approx. 6 homework assignments

Late Assignments drop 20% per each late day.

Examinations: one prelim/midterm exam,
and one final exam.

Slide CS472-3

Tentative Grading Policy (CS472):

Assignments 50%

Midterm Exam 15%

Final Exam 35%

Tentative Grading Policy (CS473):

Project Proposal and Presentation 20%

Status Reports 10%

Final Code, Write-up and Presentation 70%

Slide CS472-4

Today's Lecture

What is Artificial Intelligence?

- the components of intelligence
- historical perspective

The current frontier.

- recent achievements

Challenges ahead.

- what makes AI problems hard?

Slide CS472-5

What is Intelligence?

Intelligence:

- “the capacity to learn and solve problems”
(Webster dictionary)
- the ability to act rationally

Artificial Intelligence

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

Slide CS472-6

philosophy

e.g., foundational issues (can a machine think?), issues of knowledge and believe, mutual knowledge (Joe Halpern)

psychology and cognitive science

e.g., problem solving skills

computer science and engineering

e.g., complexity theory, algorithms, logic and inference, programming languages, and system building.

mathematics and physics

e.g., statistical modeling, Hidden Markov Models, continuous mathematics, statistical physics, complex systems, chaos, and experimental CS.

Slide CS472-7

What's 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

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

Slide CS472-8

C) Learning and Adaptation

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

Slide CS472–9

Different Approaches

- I Building exact models of human cognition**
 - view from psychology and cognitive science
- II Developing methods to match or exceed human performance in certain domains, possibly by very different means.**

Example: Deep Blue.

Our focus is on II (most recent progress).

Slide CS472–10

Issue: The Hardware

The brain

- a neuron, or nerve cell, is the basic information 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?

- 10^7 or more transistors per CPU
- supercomputer: hundreds of CPUs, 10^{10} bits of RAM
- cycle times: order of 10^{-9} seconds

Slide CS472–11

Conclusion

In near future we can have computers with as many processing elements as our brain, but:

far fewer interconnections (wires or synapses)

much faster update

Fundamentally different hardware may require fundamentally different algorithms!

Very much an open questions.

Neural net research.

Slide CS472–12

Historical Perspective

Obtaining an understanding of the human mind is one of the final frontiers of modern science.

George Boole, Gottlob Frege, and Alfred Tarski

(formalizing the laws of human thought)

Alan Turing, John von Neumann, and Claude Shannon

(thinking as computation)

John McCarthy, Marvin Minsky,

Herbert Simon, and Allen Newell

(the start of the field of AI)

Slide CS472–13

Goal of This Course

- To introduce you to a set of key methods and techniques from AI, in the areas of reasoning, planning, learning, neural nets, and NLU.
- To teach you about the applicability and limitations of these methods.

Field has matured tremendously over the last 5 to 10 years.
The Russell & Norvig book is witness to that.

Slide CS472–14

The Current Frontier

Interesting time for AI

- (May, '97) **Deep Blue** vs. **Kasparov**
First match won against world-champion,
“intelligent & creative” play,
200 million board positions per move.
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?

Slide CS472–15

Deep Blue

An outgrowth of work started by early pioneers, such as,
Shannon and McCarthy.

Matches expert level performance, while doing (most likely)
something very different from the human expert.

Dominant direction in current AI: we’re interested
in **overall performance** .

So far, attempts at incorporating more expert specific chess
knowledge to **prune** the search have failed: the game
evolves around the **expectations** to the general rules.

Slide CS472–16

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`
How intelligent is Deep Blue?

Slide CS472-17

Examples, cont.

- (Nov., '96) first **“creative”** proof by computer
60 year open problem.
Robbins' problem in finite algebra.
Qualitative difference from previous brute-force results.
Compare with computer proof of four color theorem.
`http://www.mcs.anl.gov/home/mccune/ar/robbins`
Does technique generalize?
(Our own expert: Robert Constable.)

Slide CS472-18

- **NASA: Autonomous Intelligent Systems.**

Engine control next generation spacecrafts.

General deductive approach from **first principles** .

Automatic planning and execution model.

Fast **real-time, on-line** performance.

Compiled into 2,000 variable, 10,000 clause SAT problem.

<http://ic-www.arc.nasa.gov/ic/projects/mba/index.html>

Contrast: current approach customized software with ground control team. (E.g., Mars mission \$50 million.)

Slide CS472–19

- **Decision theory and statistical user-models.**

Microsoft Office '97.

Probabilistic reasoning; diagnosis; Bayesian models.

<http://www.research.microsoft.com/research/dtg/>

Also, restricted natural language parsing.

Key issue: attempt to **adapt** to individual user.

- **Configuration systems** and expert-system style fault diagnosis and monitoring of telephone switching networks (**AT&T**) .

Slide CS472–20

Learning

- **TD-Gammon** (Tesauro 1993; 1995)
World-champion level but **learns from scratch**
by playing millions of games against itself!

Has changed human play.

Slide CS472–21

Challenges ahead

Note that the examples we discussed so far all
involve a quite **specific tasks** .

The systems lack a level of **generality** and
adaptability . They can't easily (if at all)
switch **context** .

Current work on **“intelligent agents”**
— integrates various functions (planning, reasoning,
learning etc.) in one module
— goal: to build more flexible / general systems.

Slide CS472–22

A Key Issue

- The **knowledge-acquisition** bottleneck
 - Lack of general **commonsense** knowledge.
 - CYC project (Doug Lenat et al.).
 - Attempt to encode millions of facts.

Reasoning, planning, learning can compensate to some extent for lack of background knowledge by deriving information from first-principles.

But, presumably, there is a hard limit on to how far one can take this. (open question)

Slide CS472–23

Example: Natural Language Understanding

Task rapidly becomes more difficult the more context & background knowledge is required.

Slide CS472–24

Role of Knowledge in Language Understanding

Slide CS472–25

Language understanding depends heavily on knowledge of the domain of discourse. Consider the knowledge required to understand the different meanings of “give” — one of the most common English verbs:

- John gave Pete a book
- John gave Pete a hard time
- John gave Pete a black eye
- I give him a week before he has a breakdown
- It is 300 miles, give or take 10
- John gave up
- John gave in
- John’s legs gave beneath him
- ...

Slide CS472–26

Or what about:

- **Time flies like an arrow.**
- **Fruit flies like bananas.**

Difficulty in understanding “unconstrained” natural language led to consideration of limited domains that require little domain knowledge.

Slide CS472–27

Natural Progression

Speech:

- “word spotting” — feasible today (e.g., AT&T / IBM)
- continuous speech — rapid progress / almost feasible
- translation / understanding — very limited progress

Infamous example:

The spirit is willing but the flesh is weak. (English)

The vodka is good but the meat is rotten. (Russian)

Slide CS472–28

Summary

- Discussed characteristics of intelligent systems.
- Gave series of example systems.
Involving e.g. **game playing**, **automated reasoning**
and **diagnosis** , **decision theory** , and **learning**.

Suggested readings: Chapter 1, R&N.