

Foundations of Artificial Intelligence

CS472/CS473 — Fall 2004

Slide CS472 – Introduction 1

General Information for CS472/CS473

Where: Olin 155

When: Mon, Wed, Fri 11:15–12:05

Professor: Claire Cardie, Assoc Prof, Computer Science

Email: cardie@cs.cornell.edu

URL: www.cs.cornell.edu/home/cardie/

Office Hours: always posted on-line

5161 Upson, Mon 2–3, Thurs 3–4

Course web site: www.cs.cornell.edu/courses/CS472/2004fa/

Teaching Assistants: Kamal Aboul-Hosn, Alexa Sharp,
Andrew Kae, Eric Holmberg-Weidler.

Office hours posted on-line.

Slide CS472 – Introduction 2

Text:

Artificial Intelligence: A Modern Approach

Russell and Norvig, Prentice-Hall, Inc., second edition.

Class Notes and Handouts:

Available from course web site

Homework: approx. 6 homework assignments

Late assignments drop 10% per each late day.

Examinations: two prelims, one final exam.

Tuesday, December 14

Slide CS472 – Introduction 3

Tentative Grading Policy (CS472):

Assignments	50%
Prelim 1	10%
Prelim 2	10%
Final Exam	25%
Participation/Interest	5%

Slide CS472 – Introduction 4

Tentative Grading Policy (CS473):

Preliminary Project Proposal	0%
Project Proposal	10%
Status Report 1	10%
Status Report 2	10%
Final Code (20%), Write-up (40%), and Demo (10%)	70%

Slide CS472 – Introduction 5

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?

Slide CS472 – Introduction 6

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

Slide CS472 – Introduction 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 – Introduction 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 – Introduction 9

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, and so on.*

Slide CS472 – Introduction 10

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.*

Slide CS472 – Introduction 11

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], possibly by very different means.

Example: Deep Blue.

Our focus is on III (most recent progress).

Slide CS472 – Introduction 12

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.

Slide CS472 – Introduction 13

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.

Slide CS472 – Introduction 14

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).

Slide CS472 – Introduction 15

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?

Slide CS472 – Introduction 16

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 **exceptions** to the general rules.

Slide CS472 – Introduction 17

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 – Introduction 18

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.)

Slide CS472 – Introduction 19

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

Slide CS472 – Introduction 20

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.

Slide CS472 – Introduction 21

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)”

Slide CS472 – Introduction 22

Summary

Date: 10 NOV 88
Location: CHILE: SANTIAGO (CITY)
Type: MURDER
Weapon: STONE
Victim: “BIRDS”
Perpetrator: -
Organization: -

Slide CS472 – Introduction 23

Terrorism Text

SANTIAGO, 10 NOV 88 (QUE PASA) – [TEXT]
[CONTINUED]

...THE PLENUM OF THE SOCIALIST PARTY [PS]-ALMEYDA WAS, OF COURSE, THE MOST EAGERLY ANTICIPATED... THEY AMBITIOUSLY FELT THAT THIS WAS THE OPPORTUNITY TO REMOVE SOME STRATEGIC OBSTACLES, SORT OF LIKE KILLING TWO BIRDS WITH ONE STONE: REGISTRATION AND THE SOUGHT-AFTER SOCIALIST UNITY...

Slide CS472 – Introduction 24

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 (Doug Lenat et al.). Attempt to encode millions of facts.

Slide CS472 – Introduction 25

Goal of This Course

- To introduce you to the kinds of problems studied in AI.
- To introduce you to a set of key methods and techniques from AI, in the areas of reasoning, planning, learning, and natural language understanding.
- To teach you about the applicability and limitations of these methods.

Reading: Chapter 1, R&N.

Slide CS472 – Introduction 26