

CSCI-4150

ARTIFICIAL INTELLIGENCE

Dave Hollinger
hollingd@cs.rpi.edu

Course Web Pages

www.cs.rpi.edu/~hollingd/ai

- syllabus
- announcements
- office hours
- assignments
- notes
- links

webct

- submissions
- grades
- discussions

Survey Course

- A broad range of ideas, algorithms, approaches, data structures, problem domains and issues from the field of AI.
- We will look at some topics in detail, but in general we will not.
- Term project is your chance to become an expert on something.

Grades

- Two tests (30%)
 - October 30th, December 7th
- Homework (40%)
 - at least 4 serious programming projects
 - done individually
- Term Project (30%)
 - research paper or programming project
 - can be done in groups.
 - due December 8th (last day of classes).

Grading

- No changes are made to grades one week after you have received your grade.
 - don't expect me to regrade stuff old projects at the end of the semester just because you think it might change your letter grade...
- TAs will do most of the grading.
 - If you don't feel the grading is right, see the TA and if there is still a problem – see Dave.

Textbook

- Widely used, widely respected authors.
- Most lectures will be based on the textbook.
 - there will be exceptions (Scheme)
- Some assignments/problems may be from the book (but we will spell them out, not assign as "problem 3.21").
- The 1st edition is similar, although the 2nd edition has some significant changes (additional material).

Scheme

- Part of the first 3 weeks will involve learning the Scheme programming language.
 - *A simple* version of LISP
 - Traditional AI programming language (LISP).
 - Strengths in symbolic computation.
- No textbook for Scheme, but many websites (check the course home page).
- You need MIT/gnu scheme.

Some Topics

- Scheme Programming
- Search
- Game playing
- Learning
- Reasoning and Logic
- Knowledge Representation

What is AI?


Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

The textbook advocates "acting rationally"

Acting humanly: Turing Test

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" → "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game



- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning.

Thinking humanly: cognitive modeling

- 1960s "cognitive revolution": information-processing psychology
- Requires scientific theories of internal activities of the brain
 - What level of abstraction?
 - How to validate? Requires
 - 1) Predicting and testing behavior of human subjects (top-down)
 - or 2) Direct identification from neurological data (bottom-up)
- Now know as Cognitive Science and Cognitive Neuroscience
 - are now distinct from AI.

Thinking rationally: "laws of thought"

- Aristotle: what are correct arguments/thought processes?
- Several Greek schools developed various forms of *logic: notation* and *rules of derivation* for thoughts; may or may not have evolved to the idea of mechanization
- Direct line through mathematics and philosophy to modern AI

Problems:

1. Not all intelligent behavior is mediated by logical deliberation
2. What is the purpose of thinking? What thoughts should I have?

Acting rationally: rational agent

- **Rational** behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information
- Doesn't necessarily involve thinking – e.g., blinking reflex – but thinking should be in the service of rational action

- Aristotle (Nicomachean Ethics):
- Every art and every inquiry, and similarly every
every
- action and pursuit, is thought to aim at
some good

Aristotle

(Nicomachean Ethics)

"Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good"

Rational agents

- An **agent** is an entity that perceives and acts
- This course is about designing rational agents
- Abstractly, an agent is a function from percept histories to actions:

$$[f: \mathcal{P}^* \rightarrow \mathcal{A}]$$

- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: computational limitations make perfect rationality unachievable
 - design best **program** for given machine resources

AI prehistory

- Philosophy** Logic, methods of reasoning, mind as physical system foundations of learning, language, rationality
- Mathematics** Formal representation and proof algorithms, computation, (un)decidability, (in)tractability, probability
- Economics** utility, decision theory
- Neuroscience** physical substrate for mental activity
- Psychology** phenomena of perception and motor control, experimental techniques
- Control theory** design systems that maximize an objective function over time
- Linguistics** knowledge representation, grammar

Abridged history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1952-69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1965 Robinson's complete algorithm for logical reasoning
- 1966-73 AI discovers computational complexity
Neural network research almost disappears
- 1969-79 Early development of knowledge-based systems
- 1980- AI becomes an industry (Expert Systems)
- 1986- Neural networks return to popularity
- 1988- Resurgence of probability;: ALife, GAs, soft computing
- 1995- The emergence of intelligent agents
- 2003- Human-level AI back on the agenda

State of the art

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- Proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- `Proverb` solves crossword puzzles better than most humans