

22C:145

Artificial Intelligence

Hantao Zhang
Department of Computer Science
The University of Iowa

Fall 2008

Textbook

S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, Prentice Hall, (second edition) 2003.

Class Web Site

`www.cs.uiowa.edu/~hzhang/c145`

Check the class web site daily!

Some slides and most of the figures are from our textbook.

Other AI texts

The first edition: If you have the first edition, you may download some chapters provided at

aima.cs.berkeley.edu

Title	Authors	Publisher	Year
AI: A New Synthesis	Nilsson	Morgan Kaufmann	1998
Computational Intelligence	Poole, Mackworth, Goebel	Oxford	1998
Artificial Intelligence (3rd ed.)	Winston	Addison-Wesley	1992
Artificial Intelligence (2nd ed.)	Rich, Knight	McGraw-Hill	1991
AI: Theory and Practice	Dean, Allen, Aloimonos	Benjamin Cummings	1995
Mathematical Methods in AI	Bender	IEEE Comp. Press	1996
Logical Foundations of AI	Genesereth, Nilsson	Morgan Kaufmann	1987

Course Overview

Topic	Chapters
Intelligent Agents	1, 2
Search techniques	3, 4
Constraint Satisfaction	5, 6
Knowledge Representation and Reasoning	7, 8
Logical Inference	9, 10
Reasoning under Uncertainty	13, 14
Decision Making	16
Learning and Neural Networks	18, 19, 20

Prerequisites

- The course will be self-contained, but an elementary background in CS and Math is required.
- Expect the class material to become a little technical at times.
- You will implement some of the techniques seen in class. Programming assignments will be in your favorite programming languages and Prolog.

What is Artificial Intelligence?

A scientific and engineering discipline devoted to:

- *understanding principles that make intelligent behavior possible in natural or artificial systems;*
- *developing methods for the design and implementation of useful, intelligent artifacts. [Poole, Mackworth, Goebel]*

What is intelligence then?

- Fast thinking?
- Knowledge?
- Ability to pass as a human?
- Ability to reason logically?
- Ability to learn?
- Ability to perceive and act upon one's environment?
- Ability to play chess at grand-master's level?

Dictionary: Intelligence


1. (a) The capacity to acquire and apply knowledge.
(b) The faculty of thought and reason.
(c) Superior powers of mind.
2. An intelligent, incorporeal being, especially an angel.
3. Information; news.
4. (a) Secret information, especially about an actual or potential enemy.
(b) An agency, staff, or office employed in gathering such information.
(c) Espionage agents, organizations, and activities considered as a group

Dictionary: Artificial Intelligence

1. Dictionary 1:
 - (a) The ability of a computer or other machine to perform those activities that are normally thought to require intelligence.
 - (b) The branch of computer science concerned with the development of machines having this ability.
2. Dictionary 2: The subfield of computer science concerned with the concepts and methods of symbolic inference by computer and symbolic knowledge representation for use in making inferences. AI can be seen as an attempt to model aspects of human thought on computers. It is also sometimes defined as trying to solve by computer any problem that a human can solve faster.

Acting humanly: The Turing test

Turing (1950) “Computing machinery and intelligence”:



Acting humanly: The Turing test

- “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
- Suggested major components of AI: knowledge, reasoning, language understanding, learning

Problem: Turing test is not **reproducible**, **constructive**, or amenable to **mathematical analysis**.

Acting humanly: The Turing test

Loebner Prize Gold Medal

<http://www.loebner.net/Prizef/loebner-prize.html>

In 1990 Hugh Loebner agreed with The Cambridge Center for Behavioral Studies to underwrite a contest designed to implement the Turing Test. Dr. Loebner pledged a Grand Prize of \$100,000 and a Gold Medal (pictured above) for the first computer whose responses were indistinguishable from a human's. Such a computer can be said "to think." Each year an annual prize of \$2000 and a bronze medal is awarded to the most human-like computer. The winner of the annual contest is the best entry relative to other entries that year, irrespective of how good it is in an absolute sense.

Thinking Humanly: Cognitive Science

- 1960s “cognitive revolution”: information-processing psychology replaced prevailing orthodoxy of behaviorism
- Require scientific theories of internal activities of the brain
 - What level of abstraction? “Knowledge” or “circuits”?
 - How to validate? It requires
 1. Predicting and testing behavior of human subjects (top-down)
 2. Direct identification from neurological data (bottom-up)
- Both approaches, Cognitive Science and Cognitive Neuroscience, share with AI on: *the available theories do not explain (or engender) anything resembling human-level general intelligence*

Thinking Rationally: Laws of Thought

Several Greek schools at the time of Aristotle developed various forms of **logic**:

- **Notation** and **rules of derivation** for thoughts; they may or may not have proceeded 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** behavior: doing 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: *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*
- Approach: design best **program** for given machine resources

Summary of Experts' View of AI

“The exciting new effort to make computers think ... *machines with minds*, in the full and literal sense” (Haugeland, 1985)

“[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning ...” (Bellman, 1978)

“The art of creating machines that perform functions that require intelligence when performed by people” (Kurzweil, 1990)

“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991)

“The study of mental faculties through the use of computational models”
(Charniak and McDermott, 1985)

“The study of the computations that make it possible to perceive, reason, and act”
(Winston, 1992)

“A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes” (Schalkoff, 1990)

“The branch of computer science that is concerned with the automation of intelligent behavior” (Luger and Stubblefield, 1993)

Systems that *think* like humans

Systems that *think* rationally

Systems that *act* like humans

Systems that *act* rationally

Operational Definition of AI

Systems that *act* like humans
Turing test.

Systems that *think* like humans
Cognitive Science

Systems that *think* rationally
Logic-based AI

Systems that *act* rationally
Rational Agents

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
Psychology	adaptation phenomena of perception and motor control experimental techniques (psychophysics, etc.)
Economics	formal theory of rational decisions
Linguistics	knowledge representation grammar
Neuroscience	plastic physical substrate for mental activity
Control theory	homeostatic systems, stability simple optimal agent designs

Potted History of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1965 Robinson's complete algorithm for logical reasoning
- 1966–74 AI discovers computational complexity
Neural network research almost disappears
- 1969–79 Early development of knowledge-based systems
- 1980–88 Expert systems industry booms
- 1988–93 Expert systems industry busts: "AI Winter"
- 1985–95 Neural networks return to popularity
- 1988– Resurgence of probability; general increase in technical depth
"Nouvelle AI": ALife, GAs, soft computing
- 1995– Agents agents everywhere . . .

State of the art

Which of the following can be done at present?

1. Play a decent game of table tennis
2. Drive along a curving mountain road
3. Drive in the center of Cairo
4. Buy a week's worth of groceries at a store
5. Buy a week's worth of groceries on the web
6. Play a decent game of bridge
7. Discover and prove a new mathematical theorem
8. Write an intentionally funny story
9. Give competent legal advice in a specialized area of law
10. Translate spoken English into spoken Swedish in real time
11. Perform a complex surgical operation

Why Study AI?

AI helps

- computer scientists and engineers build more useful and user-friendly computers,
- psychologists, linguists, and philosophers understand the principles that constitute what we call intelligence.

AI is an interdisciplinary field of study.

Many ideas and techniques now standard in CS (symbolic computation, time sharing, objects, declarative programming, ...) were pioneered by AI-related research.

AI is among us!

Recent applications using AI techniques:

- **Sony Aibo**
Entertainment robot with pet-like behaviour
(<http://www.us.aibo.com>)
- **Dragon Naturally Speaking**
(Dictation and voice recognition software)
(<http://www.dragonsys.com/naturallyspeaking>)
- **Ananova**
Virtual newscaster on the web
(<http://www.ananova.com/video>)
- **Honda Humanoid Robot**
Demo walking robot
(<http://www.honda.co.jp/robot>)

AI is among us!

More applications using AI techniques:

- **Deep Blue**

Chess program that beat chess grand-grand-master Kasparov

(<http://researchweb.watson.ibm.com/deepblue>)

- **Mars Pathfinder**

Autonomous land vehicle sent to Mars

(<http://mars.jpl.nasa.gov/MPF>)

- **Aaron The Robot as an Artist**

<http://www.scinetphotos.com/aaron.html>

- **Astronomy and Space Exploration**

<http://www.aaai.org/AITopics/html/astro.html>

- **and many more!**

AI is pretty hard stuff

I went to the grocery store, I saw the milk on the shelf and I bought it.

What did I buy?

- The milk?
- The shelf?
- The store?

An awful lot of knowledge of the world is needed to answer simple questions like this one.