

CS 486/686

Artificial Intelligence

Sept 9th, 2008

University of Waterloo

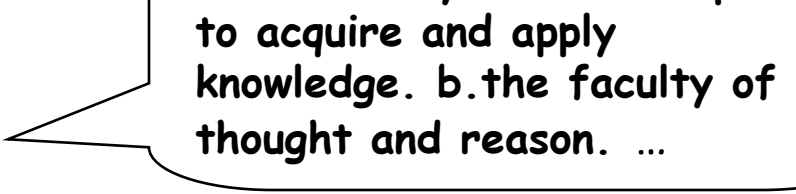
Course Info

- Instructor: Pascal Poupart
 - Email: ppoupart@cs.uwaterloo.ca
 - Office Hours: Wed 10am-12pm (DC2514) or by appt.
- Lectures: Tue & Thu, 14:30-15:50 (MC1056)
- Textbook: *Artificial Intelligence: A Modern Approach* (2nd Edition), by Russell & Norvig
- Website
 - <http://www.student.cs.uwaterloo.ca/~cs486>

Outline

- What is AI? (Chapter 1)
- Rational agents (Chapter 2)
- Some applications
- Course administration

Artificial Intelligence (AI)

- What is **AI**?
- What is **intelligence**? 

Webster says: a. the capacity to acquire and apply knowledge. b. the faculty of thought and reason. ...
- What features/abilities do humans (animals? animate objects?) have that you think are indicative or characteristic of intelligence?
- *abstract concepts, mathematics, language, problem solving, memory, logical reasoning, emotions, morality, ability to learn/adapt, etc...*

Some Definitions (Russell & Norvig)

The exciting new effort to make computers that think... machines with minds in the full and literal sense
[Haugeland 85]

[The automation of] activities that we associate with human thinking, such as decision making, problem solving, learning [Bellman 78]

The art of creating machines that perform functions that require intelligence when performed by a human [Kurzweil 90]

The study of how to make computers do things at which, at the moment, people are better [Rich&Knight 91]

The study of mental faculties through the use of computational models
[Charniak & McDermott 85]

The study of computations that make it possible to perceive, reason and act
[Winston 92]

A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes [Schalkoff 90]

The branch of computer science that is concerned with the automation of intelligent behavior
[Luger&Stubblefield93]

Some Definitions (Russell & Norvig)


Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

What is AI?

- **Systems that think like humans**
 - Cognitive science
 - Fascinating area, but we will not be covering it in this course
- **Systems that think rationally**
 - Aristotle: What are the correct thought processes
 - Systems that reason in a logical manner
 - Systems doing inference correctly

What is AI?

- Systems that behave like humans
 - Turing (1950) “Computing machinery and intelligence”



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

What is AI?

- **Systems that act rationally**
 - Rational behavior: “doing the right thing”
 - Rational agent approach
 - Agent: entity that perceives and acts
 - Rational agent: acts so to achieve best outcome
 - This is the approach we will take in this course
 - General principles of rational agents
 - Components for constructing rational agents

Intelligent Assistive Technology

- Let's facilitate aging in place
- Intelligent assistive technology
 - Non-obtrusive, yet pervasive
 - Adaptable
- Benefits:
 - Greater autonomy
 - Feeling of independence




COACH project

- Automated prompting system to help elderly persons wash their hands
- Collaborators: Geoff Fernie, Alex Mihailidis, Jennifer Boger, Jesse Hoey and Craig Boutilier



System Overview



Video Clip #1



Video Clip #2



Topics covered

- **Search**
 - Uninformed and heuristic search
 - Constraint satisfaction problems
 - Propositional and first order logic
- **Reasoning under uncertainty**
 - Probability theory, utility theory and decision theory
 - Bayesian networks and decision networks
 - Markov networks and Markov logic networks
- **Learning**
 - Decision trees, statistical learning, ensemble learning
- **Specialized areas**
 - Natural language processing and robotics

A brief history of AI

- 1943-1955: Initial work in AI
 - McCulloch and Pitts produce boolean model of the brain
 - Turing's "Computing machinery and intelligence"
- Early 1950's: Early AI programs
 - Samuel's checker program, Newell and Simon's Logic Theorist, Gerlenter's Geometry Engine
- 1956: Happy birthday AI!
 - Dartmouth workshop attended by McCarthy, Minsky, Shannon, Rochester, Samuel, Solomonoff, Selfridge, Simon and Newell


A brief history of AI

- 1950's-1969: Enthusiasm and expectations
 - Many successes (in a limited way)
 - LISP, time sharing, Resolution method, neural networks, vision, planning, learning theory, Shakey, machine translation,...
- 1966-1973: Reality hits
 - Early programs had little knowledge of their subject matter
 - Machine translation
 - Computational complexity
 - Negative result about perceptrons - a simple form of neural network

A brief history of AI

- 1969-1979: Knowledge-based systems
- 1980-1988: Expert system industry booms
- 1988-1993: Expert system busts, AI Winter
- 1986-2000: The return of neural networks
- 1988-present:
 - Resurgence of probabilistic and decision-theoretic methods
 - Increase in technical depth of mainstream AI
 - Intelligent agents

Agents and Environments



Agents include humans, robots, softbots, thermostats...

The **agent function** maps percepts to actions $f:P^* \rightarrow A$

The **agent program** runs on the physical architecture to produce f

Rational Agents

- Recall: A rational agent “does the right thing”
- Performance measure – success criteria
 - Evaluates a sequence of environment states
- A **rational agent** chooses whichever action that maximizes the **expected** value of its performance measure **given the percept sequence to date**
 - Need to know performance measure, environment, possible actions, percept sequence
- Rationality \neq Omniscience, Perfection, Success
- Rationality \rightarrow exploration, learning, autonomy

PEAS

- Specify the **task environment**:
 - **P**erformance measure, **E**nvironment, **A**ctuators, **S**ensors

Example: COACH system

Perf M: task completion, time taken, amount of intervention

Envir: Bathroom status, user status

Actu: Verbal prompts, CallCaregiver, DoNothing

Sens: Video cameras, microphones, tap sensor

Example: Autonomous Taxi

Perf M: Safety, destination, legality...

Envir: Streets, traffic, pedestrians, weather...

Actu: Steering, brakes, accelerator, horn...

Sens: GPS, engine sensors, video...

Properties of task environments

- Fully observable vs. partially observable
- Deterministic vs. stochastic
- Episodic vs. sequential
- Static vs. dynamic
- Discrete vs. continuous
- Single agent vs. multiagent

Hardest case: Partially observable, stochastic, sequential, dynamic, continuous and multiagent.
(Real world)

Examples

Solitaire	Backgammon	Internet Shopping	Taxi
Fully Observable	Fully Observable	Partially Observable	Partially Observable
Deterministic	Stochastic	Stochastic	Stochastic
Sequential	Sequential	Sequential	Episodic
Static	Static	Dynamic	Dynamic
Discrete	Discrete	Discrete	Continuous
Single agent	Multiagent	Multiagent	Multiagent

Many Applications

- credit card fraud detection
- printer diagnostics, help in Windows, spam filters
- medical assistive technologies
- information retrieval, Google
- Intelligent Systems Challenge
- scheduling, logistics, etc.
- aircraft, pipeline inspection
- language understanding, generation, translation
- Mars rovers
- and, of course, cool robots

Mobile Robotics



Next Class

- Uninformed search
- Sect. 3.1-3.5 (Russell & Norvig)