

CSE 473: Artificial Intelligence

Luke Zettlemoyer

<http://www.cs.washington.edu/cse473/13sp/>

Slides from Dan Klein, Dan Weld, Stuart Russell, Andrew Moore

What is CSE 473?

Textbook:


Artificial Intelligence: A Modern Approach,
Russell and Norvig (3rd ed)

Prerequisites:

- Data Structures (CSE 326 or CSE 322)
or equivalent
- Basic exposure to probability, data
structures, and logic

Work:

Readings (mostly from text),
Programming assignment (40%),
Written assignments (20%),
Take Home Final Exam (35%),
Class participation (5%)




Topics

- Introduction
- Search
- Game Playing (minimax, alpha beta, expectimax)
- Constraint satisfaction
- Logic & Planning
- Markov Decision Processes
- Reinforcement Learning
- Uncertainty, Bayesian networks, HMMs
- Supervised Machine Learning
- Natural Language Processing

Today

- What is artificial intelligence (AI)?
- What can AI do?
- What is this course?

What is AI?



What is AI?

The science of making machines that:

Think like humans	Think rationally
Act like humans	Act rationally

Rational Decisions


We'll use the term rational in a particular way:

- Rational: maximally achieving pre-defined goals
- Rational only concerns what decisions are made
(not the thought process behind them)
- Goals are expressed in terms of the utility of outcomes
- Being rational means maximizing your expected utility

A better title for this course would be:

Computational Rationality


Can We Build It?



10^{11} neurons
 10^{14} synapses
cycle time: 10^{-3} sec

VS.

10^9 transistors
 10^{12} bits of RAM
cycle time: 10^{-9} sec




A (Short) History of AI

- Prehistory
- 1940-1950: Early days
- 1950—70: Excitement: Look, Ma, no hands!
- 1970—88: Knowledge-based approaches
- 1988—: Statistical approaches
- 2000—: Where are we now?

Prehistory

- **Logical Reasoning:** (4th C BC+) Aristotle, George Boole, Gottlob Frege, Alfred Tarski
- **Probabilistic Reasoning:** (16th C+) Gerolamo Cardano, Pierre Fermat, James Bernoulli, Thomas Bayes



and



1940-1950: Early Days


- 1943: McCulloch & Pitts: Boolean circuit model of brain
- 1950: Turing's "Computing Machinery and Intelligence"

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed...

-Alan Turing

The Turing Test

- Turing (1950) “Computing machinery and intelligence”
 - “Can machines think?” \diamond “Can machines behave intelligently?”
 - The Imitation Game:



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

1950-1970: Excitement

- 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

“Over Christmas, Allen Newell and I created a thinking machine.”

-Herbert Simon

1970-1980: Knowledge Based Systems

- 1969-79: Early development of knowledge-based systems
- 1980-88: Expert systems industry booms
- 1988-93: Expert systems industry busts
“AI Winter”

The knowledge engineer practices the art of bringing the principles and tools of AI research to bear on difficult applications problems requiring experts' knowledge for their solution.

- *Edward Felgenbaum* in “The Art of Artificial Intelligence”

1988--: Statistical Approaches

- 1985-1990: Probability and Decision Theory win
Pearl, Bayes Nets
- 1990-2000: Machine learning takes over subfields:
Vision, Natural Language, etc.
- Agents, uncertainty, and learning systems...
“AI Spring”?

"Every time I fire a linguist, the performance of the speech recognizer goes up"

-Fred Jelinek, IBM Speech Team

What Can AI Do?

Quiz: Which of the following can be done at present?

- Play a decent game of Soccer?
- Play a winning game of Chess? Go? Jeopardy?
- Drive safely along a curving mountain road? University Way?
- Buy a week's worth of groceries on the Web? At QFC?
- Make a car? Make a cake?
- Discover and prove a new mathematical theorem?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate Chinese into English in real time?

Robocup



What Can AI Do?

Quiz: Which of the following can be done at present?


- Play a decent game of Soccer?
- Play a winning game of Chess? Go? Jeopardy?
- Drive safely along a curving mountain road? University Way?
- Buy a week's worth of groceries on the Web? At QFC?
- Make a car? Make a cake?
- Discover and prove a new mathematical theorem?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate Chinese into English in real time?

State of the Art

May 1997

“I could feel –
I could smell –
a new kind of
intelligence
across the
table”


-Gary Kasparov



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.

– Drew McDermott

Other Games?



What Can AI Do?

Quiz: Which of the following can be done at present?

- Play a decent game of Soccer?
- Play a winning game of Chess? Go? Jeopardy?
- Drive safely along a curving mountain road? University Way?
- Buy a week's worth of groceries on the Web? At QFC?
- Make a car? Make a cake?
- Discover and prove a new mathematical theorem?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate Chinese into English in real time?

Google Car



What Can AI Do?

Quiz: Which of the following can be done at present?

- Play a decent game of Soccer?
- Play a winning game of Chess? Go? Jeopardy?
- Drive safely along a curving mountain road? University Way?
- Buy a week's worth of groceries on the Web? At QFC?
- Make a car? Make a cake?
- Discover and prove a new mathematical theorem?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate Chinese into English in real time?

Brownies Anyone?

BakeBot: Motion Planning for Cooking

Mario Bollini and Daniela Rus
CSAIL, MIT



What Can AI Do?

Quiz: Which of the following can be done at present?

- Play a decent game of Soccer?
- Play a winning game of Chess? Go? Jeopardy?
- Drive safely along a curving mountain road? University Way?
- Buy a week's worth of groceries on the Web? At QFC?
- Make a car? Make a cake?
- Discover and prove a new mathematical theorem?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate Chinese into English in real time?

Mathematical Calculation

Introducing
MATHEMATICA⁵

Παρουσιάζουμε το

Featuring a new generation of
advanced algorithms with unparalleled
speed, scope, and scalability •

$$\partial_r^2 u = - \left[E' - \frac{l(l+1)}{r^2} - r^2 \right] u(r)$$

$$e^{-2s} (\partial_s^2 - \partial_s) u(s) = - [E' - l(l+1)e^{-2s} - e^{2s}] u(s)$$

$$e^{-2s} \left[e^{\frac{1}{2}s} \left(e^{-\frac{1}{2}s} u(s) \right)'' - \frac{1}{4} u \right] = - [E' - l(l+1)e^{-2s} - e^{2s}] u(s)$$

$$e^{-2s} \left[e^{\frac{1}{2}s} \left(e^{-\frac{1}{2}s} u(s) \right)'' \right] = - \left[E' - \left(l + \frac{1}{2} \right)^2 e^{-2s} - e^{2s} \right] u(s)$$

$$v'' = -e^{2s} \left[E' - \left(l + \frac{1}{2} \right)^2 e^{-2s} - e^{2s} \right] v$$

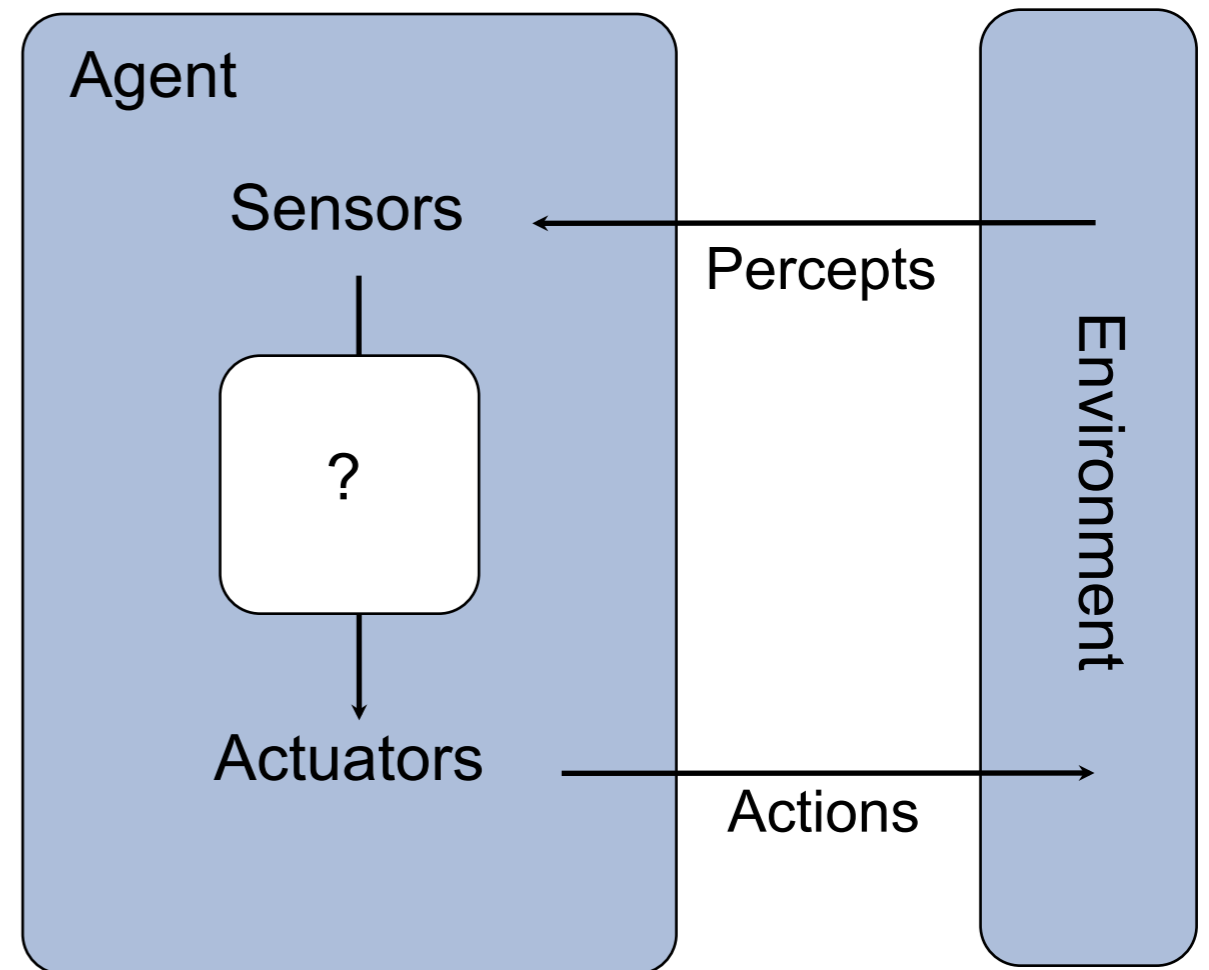
What Can AI Do?

Quiz: Which of the following can be done at present?

- Play a decent game of Soccer?
- Play a winning game of Chess? Go? Jeopardy?
- Drive safely along a curving mountain road? University Way?
- Buy a week's worth of groceries on the Web? At QFC?
- Make a car? Make a cake?
- Discover and prove a new mathematical theorem?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate Chinese into English in real time?

Designing Rational Agents



- An agent is an entity that perceives and acts.
- A rational agent selects actions that maximize its utility function.
- Characteristics of the percepts, environment, and action space dictate techniques for selecting rational actions.



This course is about:

- General AI techniques for a variety of problem types
- Learning to recognize when and how a new problem can be solved with an existing technique

Pacman as an Agent



Types of Environments


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

Fully observable vs. Partially observable

Can the agent observe the complete state of the environment?



vs.



Single agent vs. Multiagent

Is the agent the only thing acting in the world?



vs.

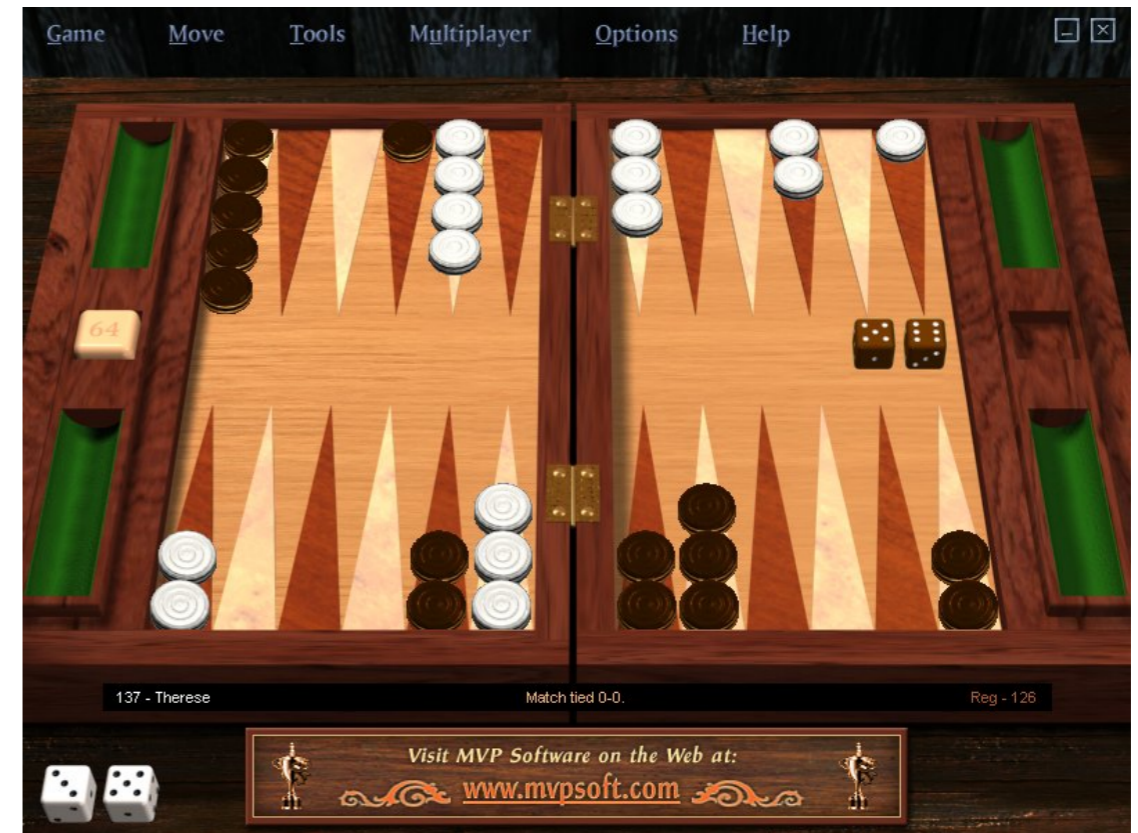


Deterministic vs. Stochastic

Is there uncertainty in how the world works?

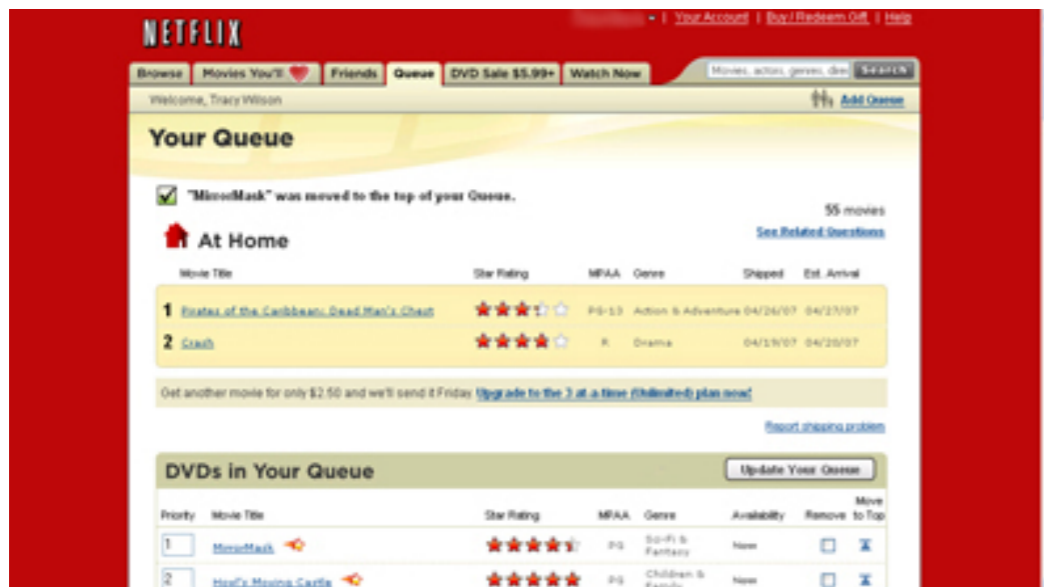


vs.




Episodic vs. Sequential

Does the agent take more than one action?



vs.




Discrete vs. Continuous


- Is there a finite (or countable) number of possible environment states?



vs.



Assignments: Pac-man



Originally developed at UC Berkeley:

<http://www-inst.eecs.berkeley.edu/~cs188/pacman/pacman.html>


PS I: Search

Goal:

- Help Pac-man find his way through the maze

Techniques:

- Search: breadth-first, depth-first, etc.
- Heuristic Search: Best-first, A^* , etc.




PS2: Game Playing

Goal:

- Play Pac-man!

Techniques:

- Adversarial Search: minimax, alpha-beta, expectimax, etc.



PS3: Planning and Learning

Goal:

- Help Pac-man learn about the world

Techniques:

- Planning: MDPs, Value Iterations
- Learning: Reinforcement Learning




PS4: Ghostbusters

Goal:

- Help Pac-man hunt down the ghosts

Techniques:

- Probabilistic models: HMMS, Bayes Nets
- Inference: State estimation and particle filtering



To Do:

- Look at the course website:
 - <http://www.cs.washington.edu/cse473/13sp/>
- Do the readings
- Do the python tutorial