CSE 473: Artificial Intelligence

Instructor: Luke Zettlemoyer
Web: http://www.cs.washington.edu/cse473/11au/

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

What is AI?

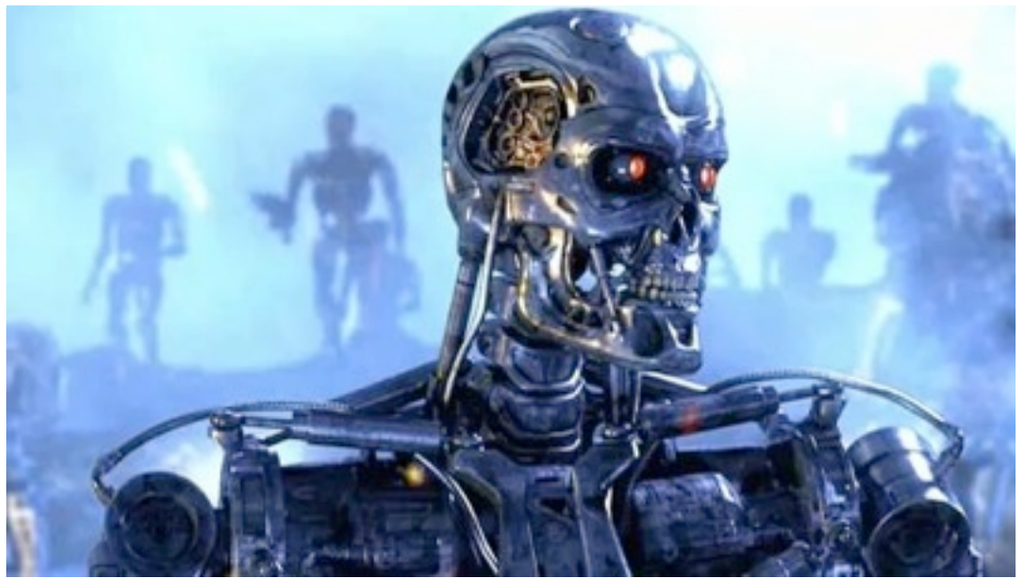

Could We Build It?

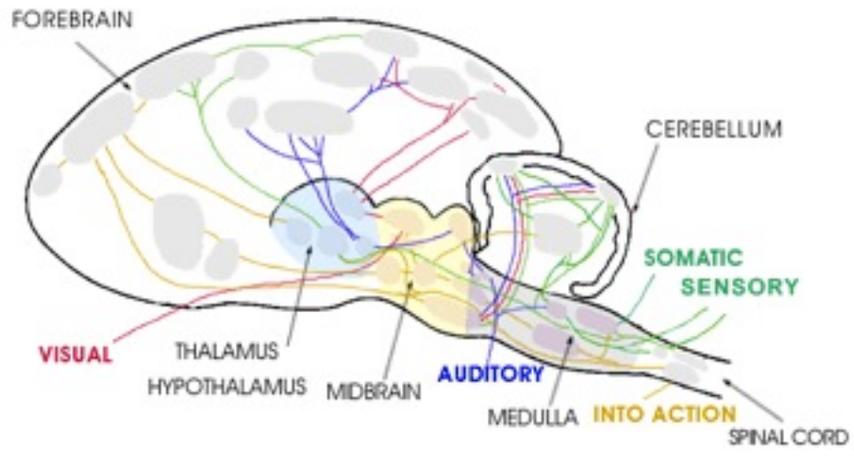

10^{11} neurons
10^{14} synapses
cycle time: 10^{-3} sec

vs.

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

What is CSE 473?

Textbook:

• Artificial Intelligence: A Modern Approach, Russell and Norvig (third edition)

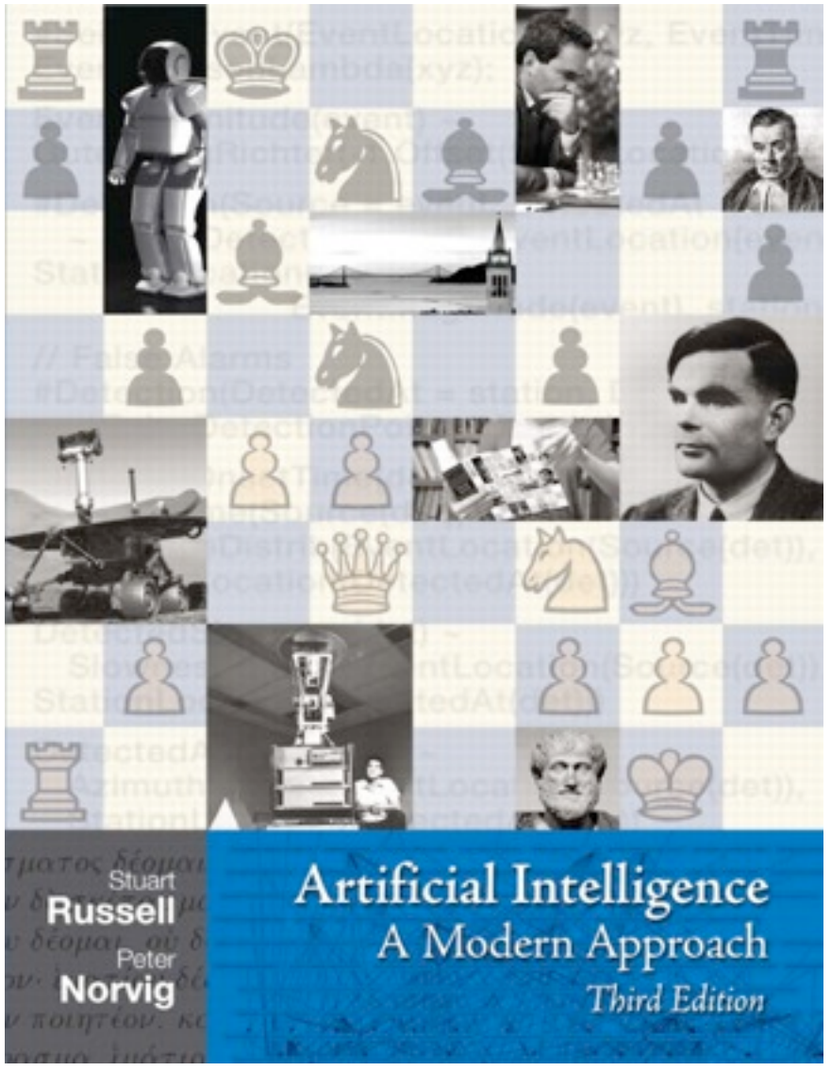

Prerequisites:

• Data Structures (CSE 326 or CSE 332), or equivalent

• basic exposure to probability, data structures, and logic

Work:

• Readings (mostly from text), Programming assignment (40%), written assignments (30%), final exam (30%)

Topics

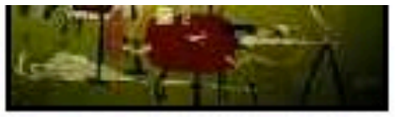

CSE 473 - Introduction to Artificial Intelligence - Autumn 2011
Mon, Wed, Fri 9:30-10:20 in MGH 231

Instructor: Luke Zettlemoyer (lsz *at* cs *dot* washington *dot* edu)
Office hours: Mondays 10:30-11:30, CSE 658

TA: Lydia Chilton (hmslydia *at* cs *dot* washington *dot* edu)
Office hours: Thursdays 3-4, CSE 220

Schedule

Week	Dates	Topics & Lecture Notes	Readings
1	September 28, 30	Introduction, Search	R&N, Ch. 1 (Ch. 2 is optional)
2	October 3, 5, 7	Search (cont.), Heuristic Search	R&N, Ch. 3.1-3.7
3	October 10, 12, 14	Game Playing: Minimax and Expectimax	R&N, Ch. 5.1-5.9 (5.6 is optional)
4	October 17, 19, 21	Constraint Satisfaction	R&N, Ch. 6.1-6.5
5	October 24, 26, 28	Logic and Planning	R&N, Ch. 7, 8
6	October 31, November 2, 4	Markov Decision Processes	R&N, Ch. 16.1-16.3, Ch 17.1-17.3; S&B, Ch. 3-4
7	November 7, 9	Reinforcement Learning	R&N, Ch. 17.4 (also, finish previous reading)
8	November 14, 16, 18	Uncertainty, Bayesian Networks	R&N, 14.1-14.5
9	November 21, 23	Bayesian Network Inference, Hidden Markov Models	R&N, Ch. 15.1-15.6
10	November 28, 30, December 2	Machine Learning: Naive Bayes and Perceptron	R&N, Ch. 18.1, 18.4, 18.6
11	December 5, 7, 9	TBD: Advanced Topics	TBD

Textbooks

Assignments: Pac-man

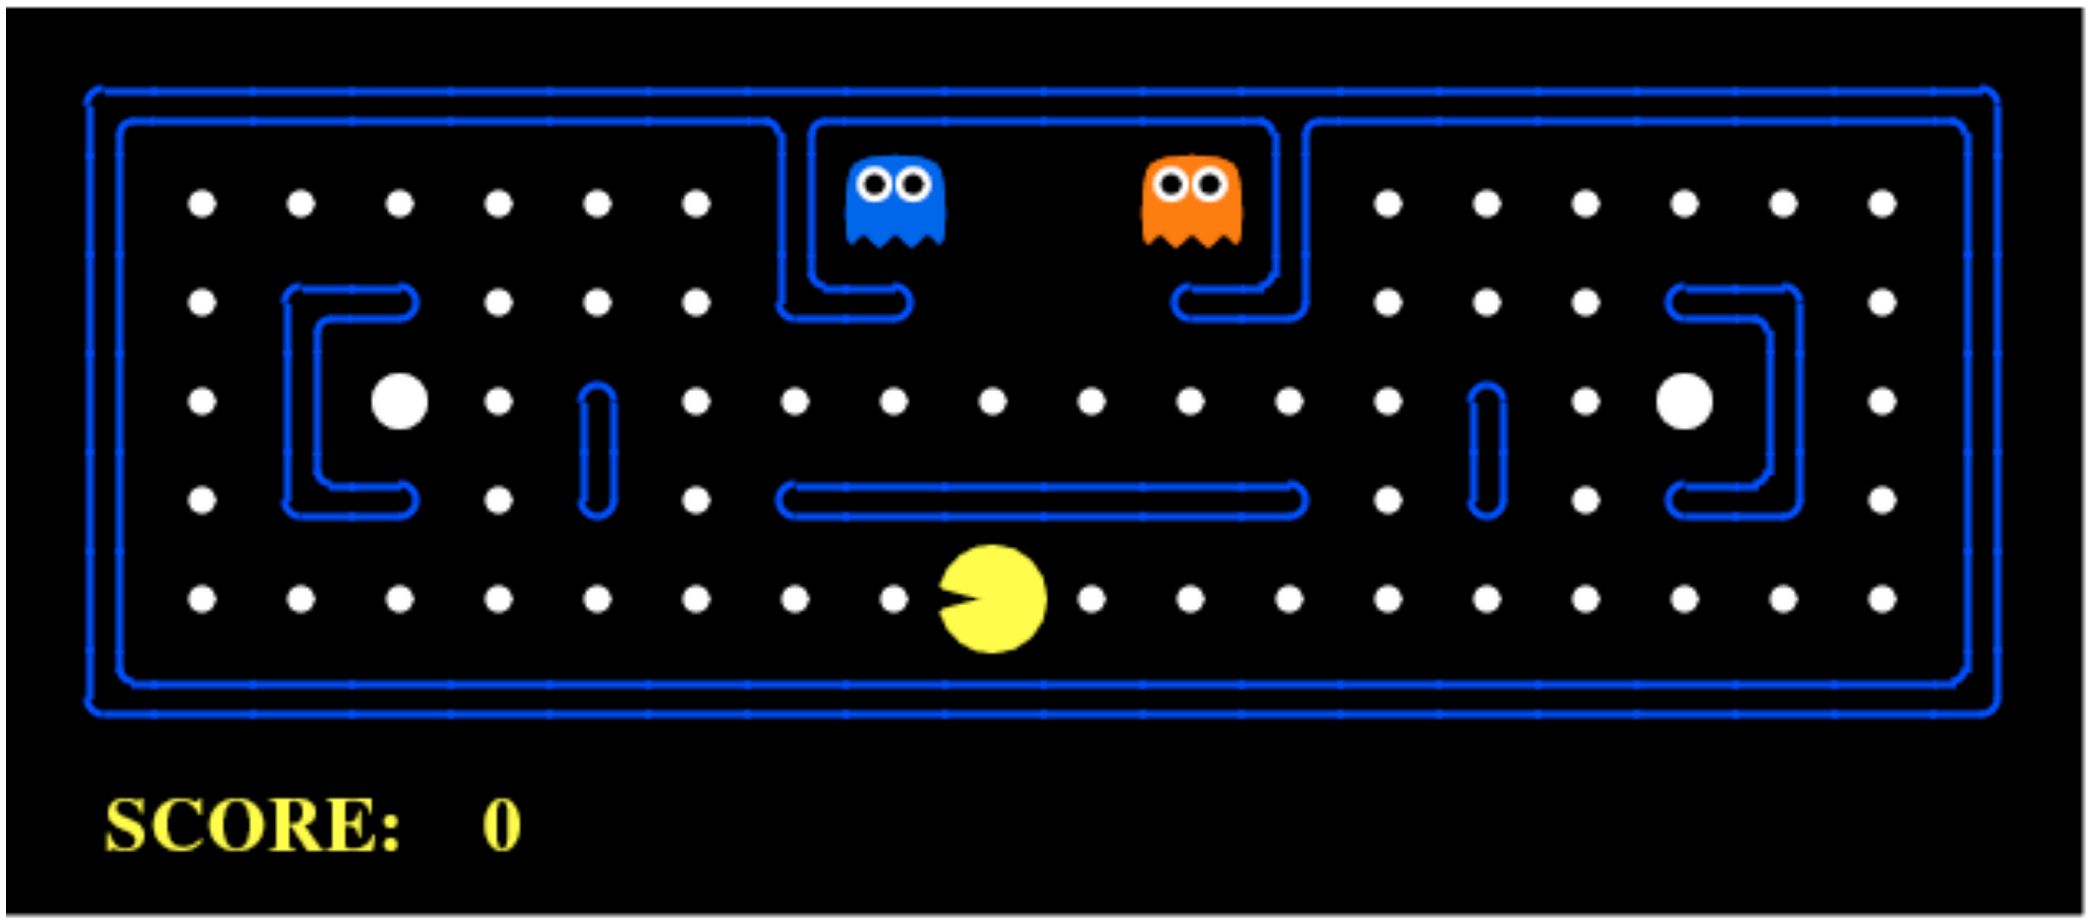

Originally developed at UC Berkeley:

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

Today

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

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

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

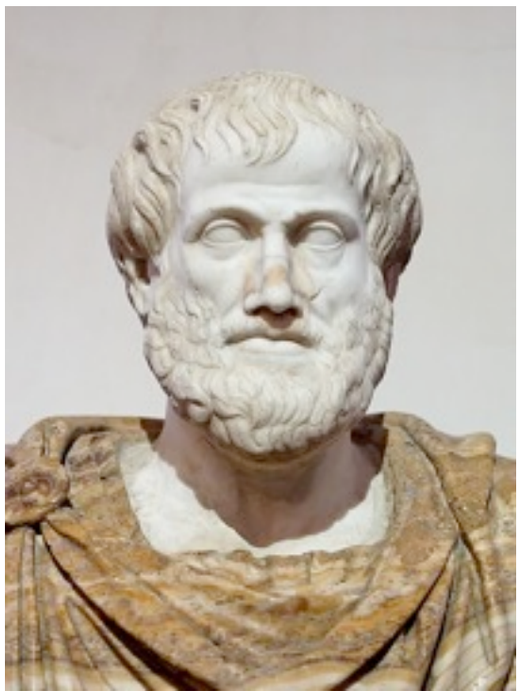

and

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?" → "Can machines behave intelligently?"
  - The *Imitation Game:*

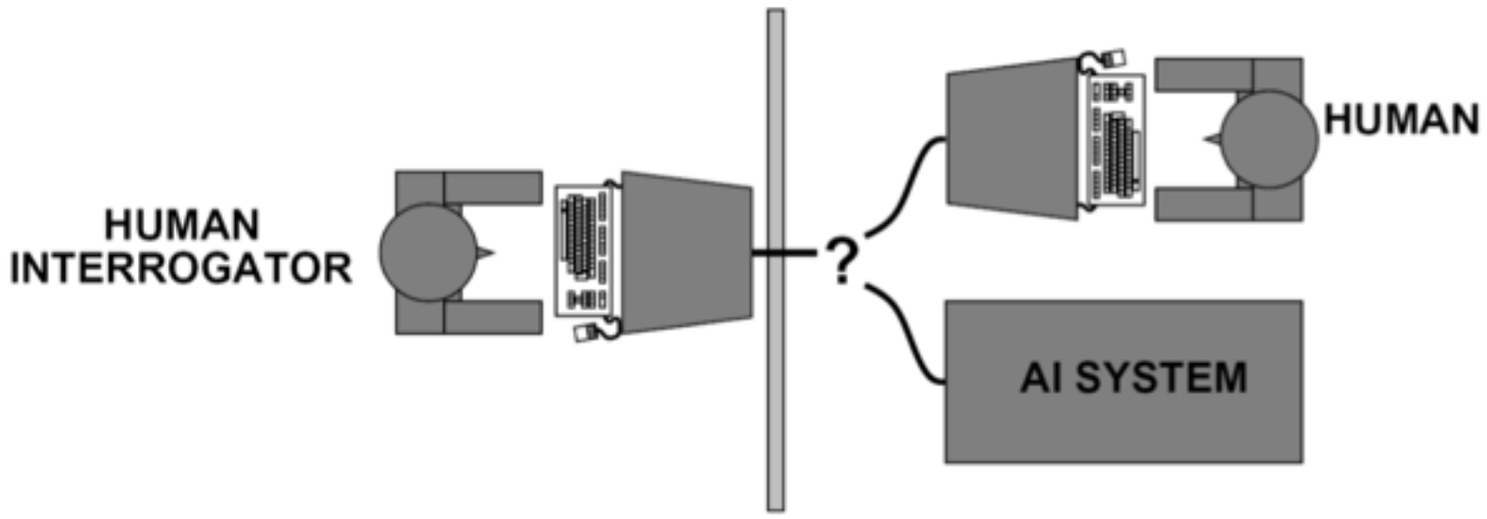

  - Suggested major components of AI: knowledge, reasoning, language understanding, learning

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?
- Drive safely along a curving mountain road?
- Drive safely along University Way?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at QFC?
- Make breakfast?
- Discover and prove a new mathematical theorem?
- Converse successfully with another person for an hour?
- 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?

- Drive safely along a curving mountain road?
- Drive safely along University Way?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at QFC?
- Make breakfast?
- Discover and prove a new mathematical theorem?
- Converse successfully with another person for an hour?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate Chinese into English in real time?

Google Car


AUTOMOBILE SHOW
LOOK, NO HANDS!
INTRODUCING THE DRIVERLESS CAR

What Can AI Do?

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

- ? Play a decent game of soccer?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along University Way?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at QFC?
- Make breakfast?
- Discover and prove a new mathematical theorem?
- Converse successfully with another person for an hour?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate Chinese into English in real time?

Pancakes Anyone?

Cookies?

BakeBot:
Motion Planning for Cooking

Mario Bollini and Daniela Rus
CSAIL, MIT

CSAIL

What Can AI Do?

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

- ? Play a decent game of soccer?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along University Way?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at QFC?
- ? Make breakfast?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person for an hour?
- ✗ 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.

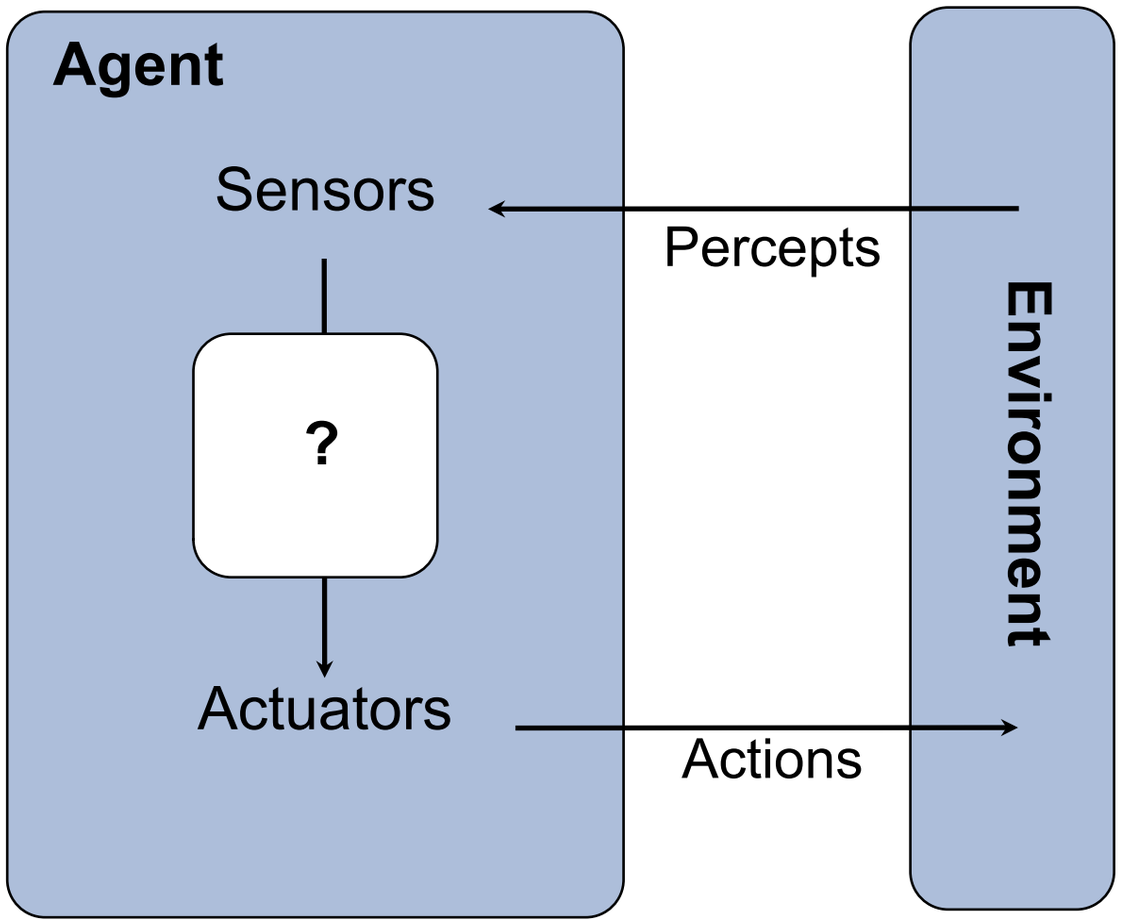

- This course is about:
  - General AI techniques for a variety of problem types
  - Learning to recognize when and how a new problem can be solved with an existing technique

Pacman as an Agent

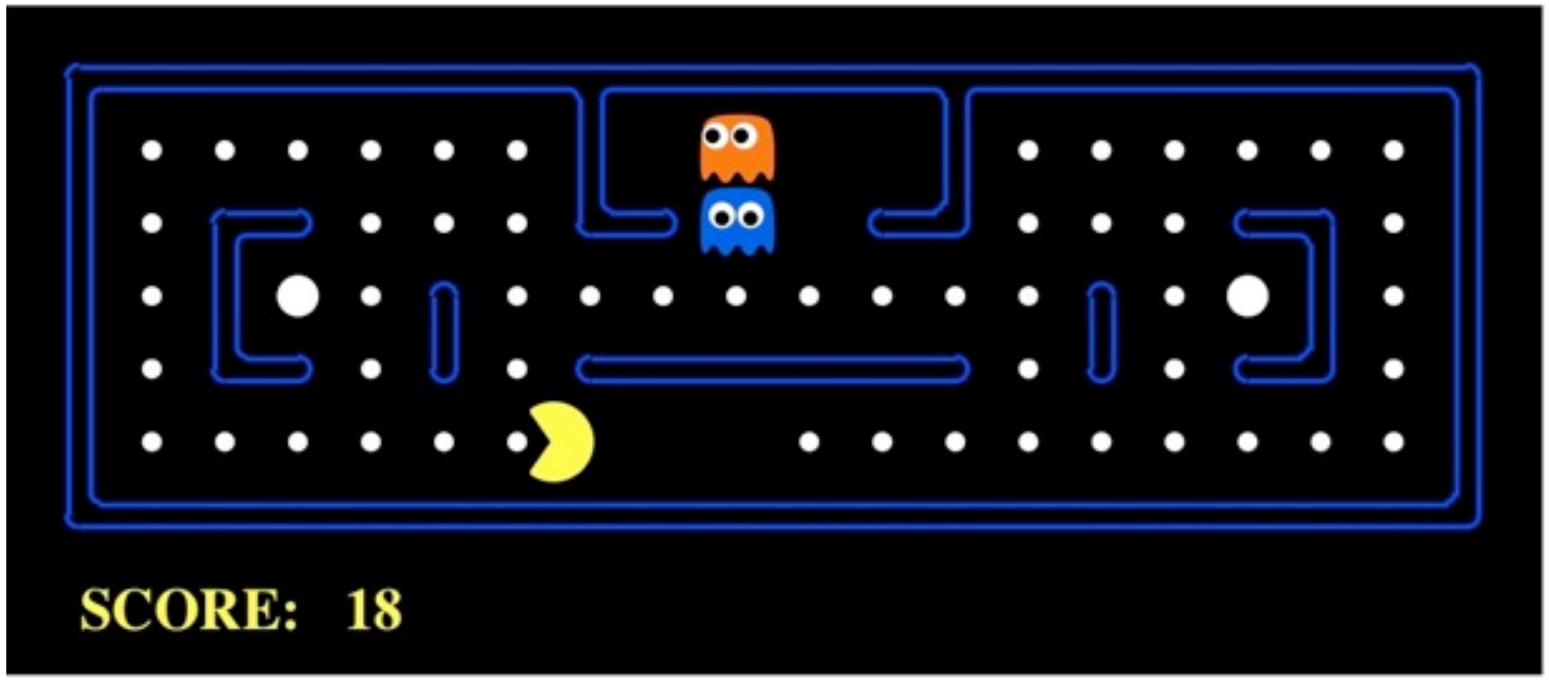

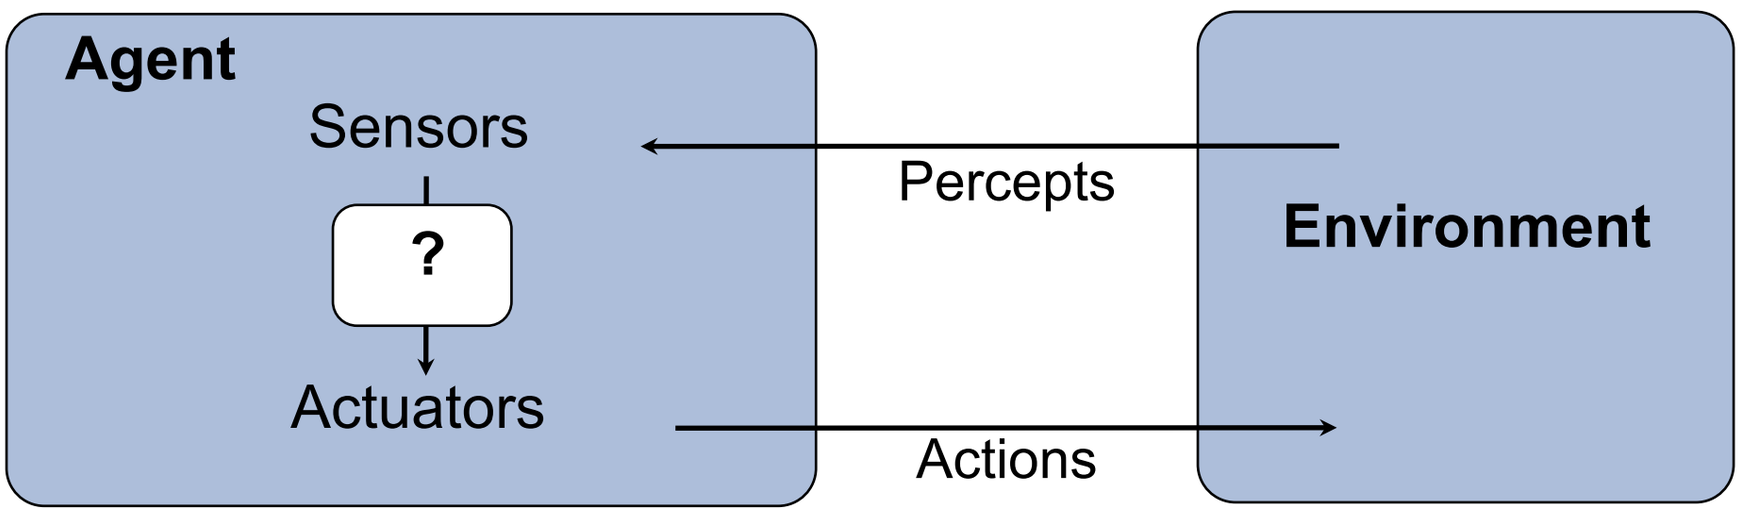

Types of Environments

- 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?

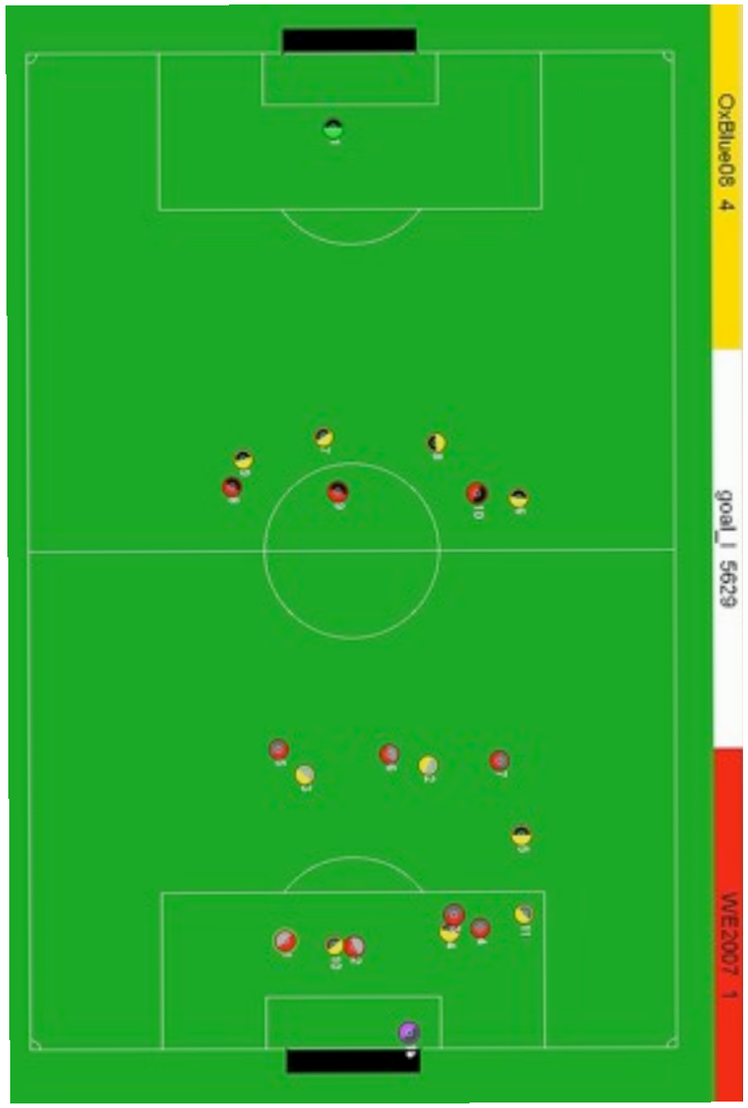

vs.

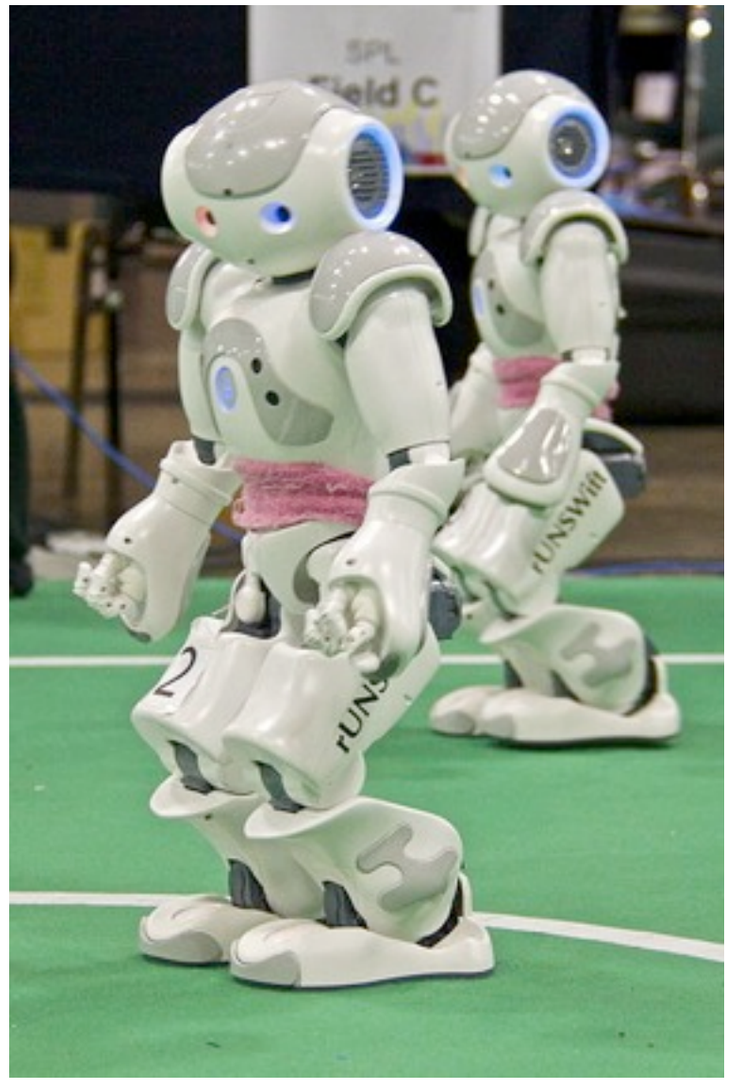

Single agent vs. Multiagent

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?

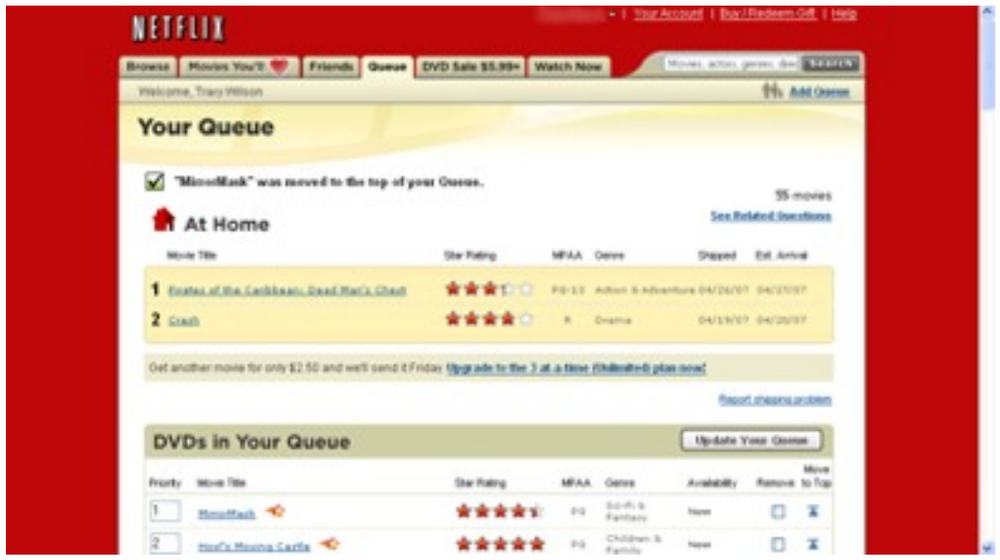

VS.

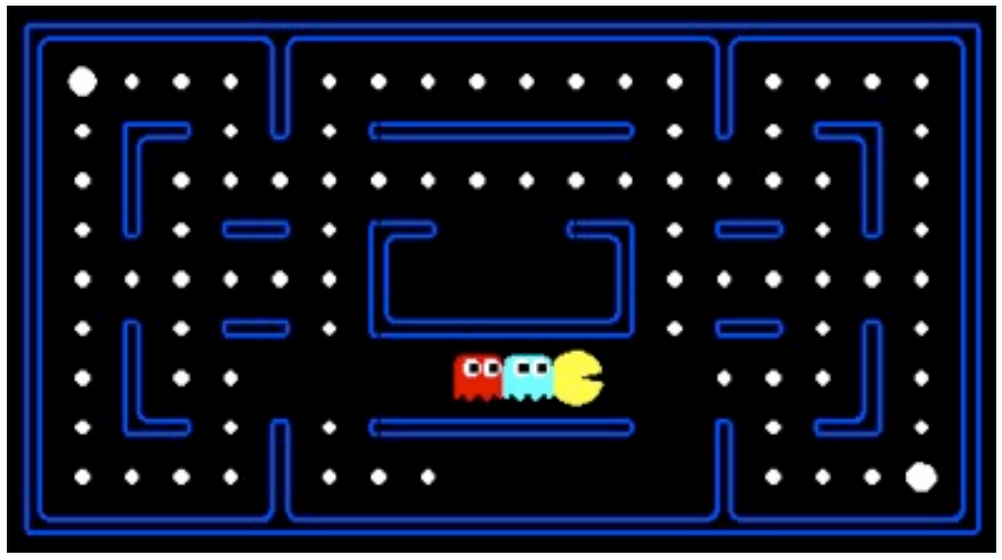

Discrete vs. Continuous

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

VS.

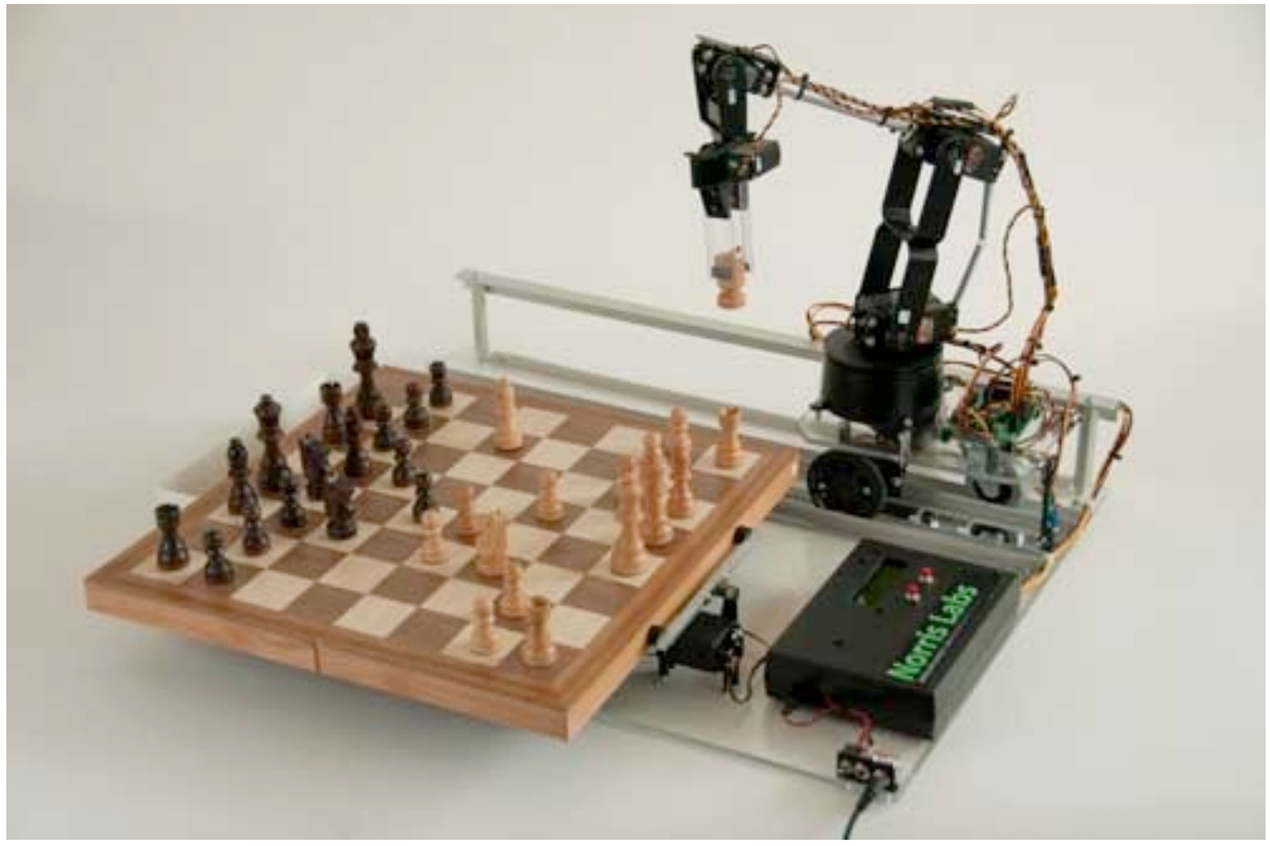

Assignments: Pac-man

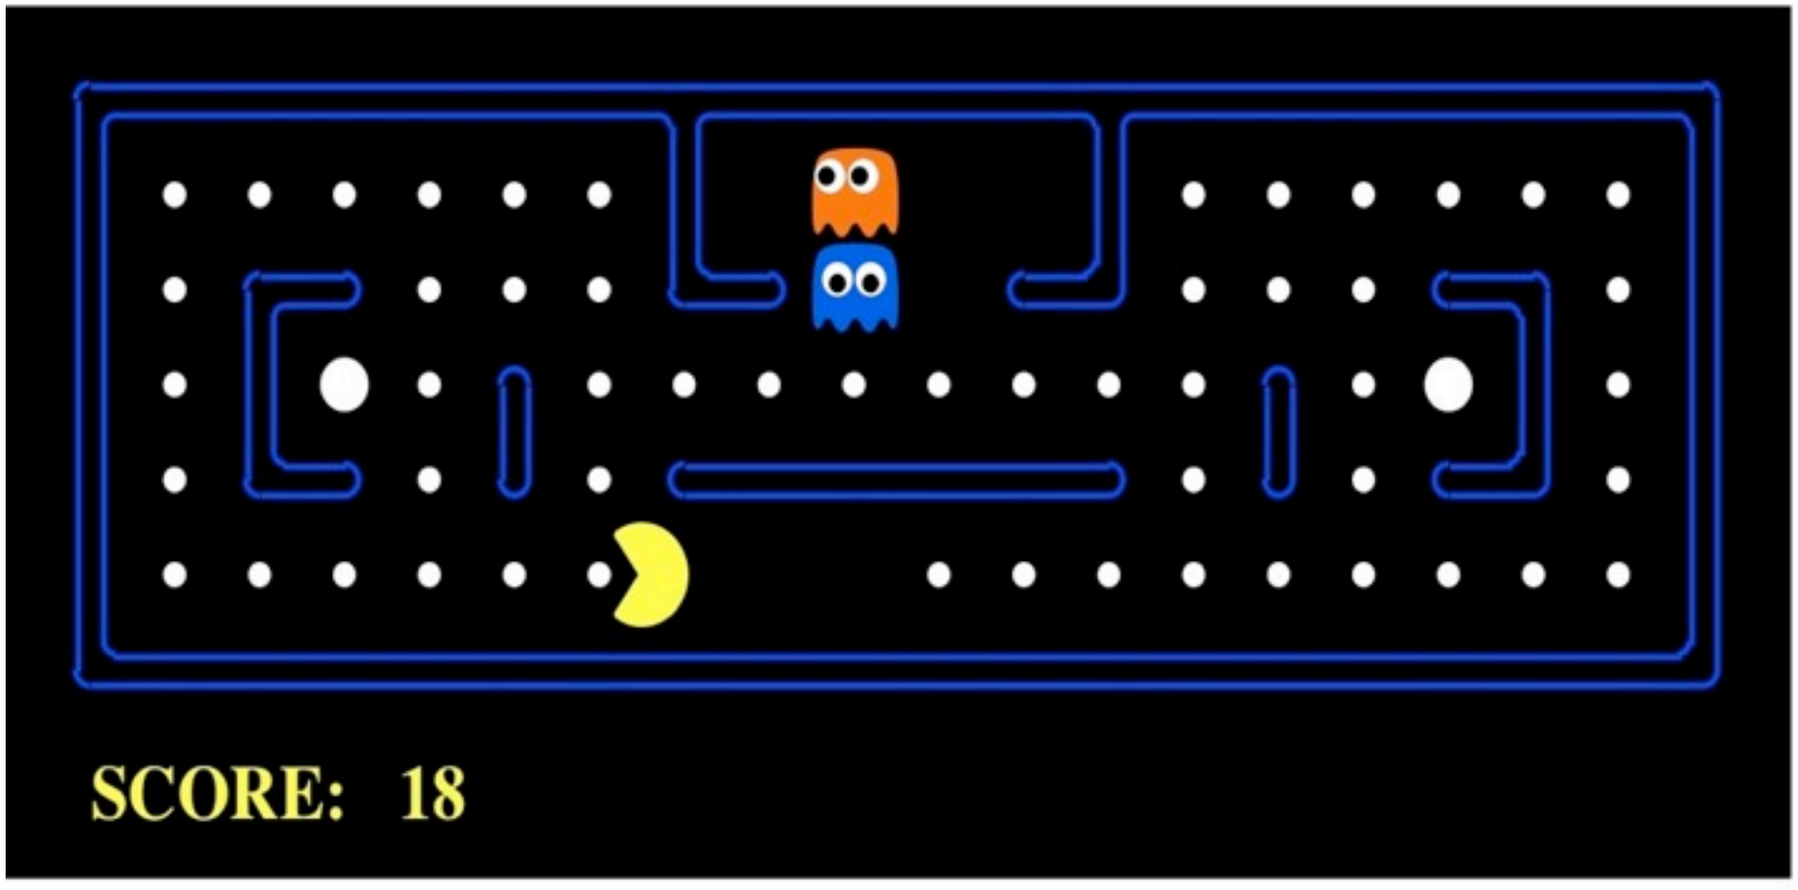

Originally developed at UC Berkeley:

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

PS1: Search

Goal:

- Help Pac-man find his way through the maze

Techniques:

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

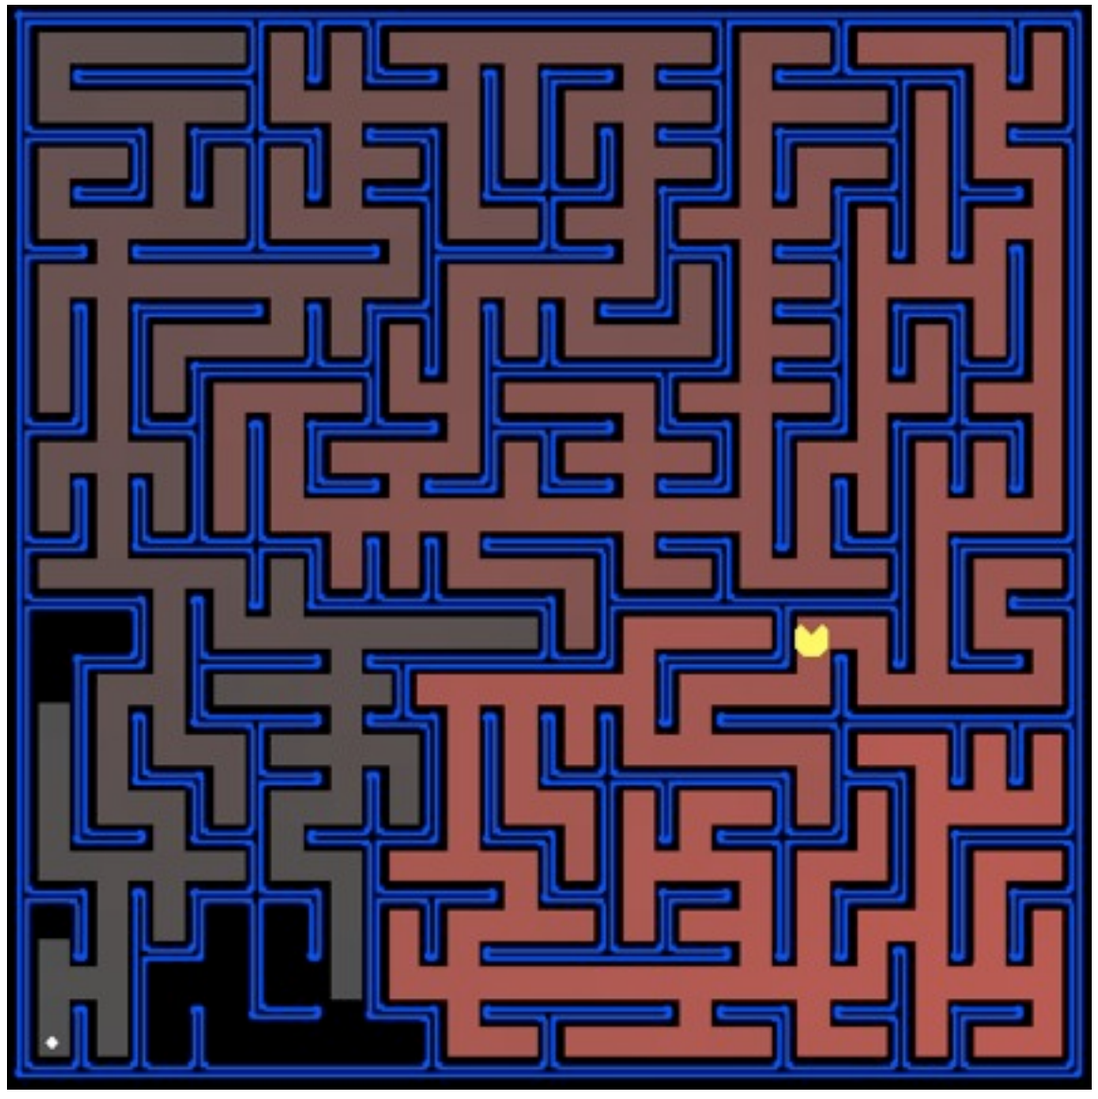

PS2: Game Playing

Goal:

- Play Pac-man!

Techniques:

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

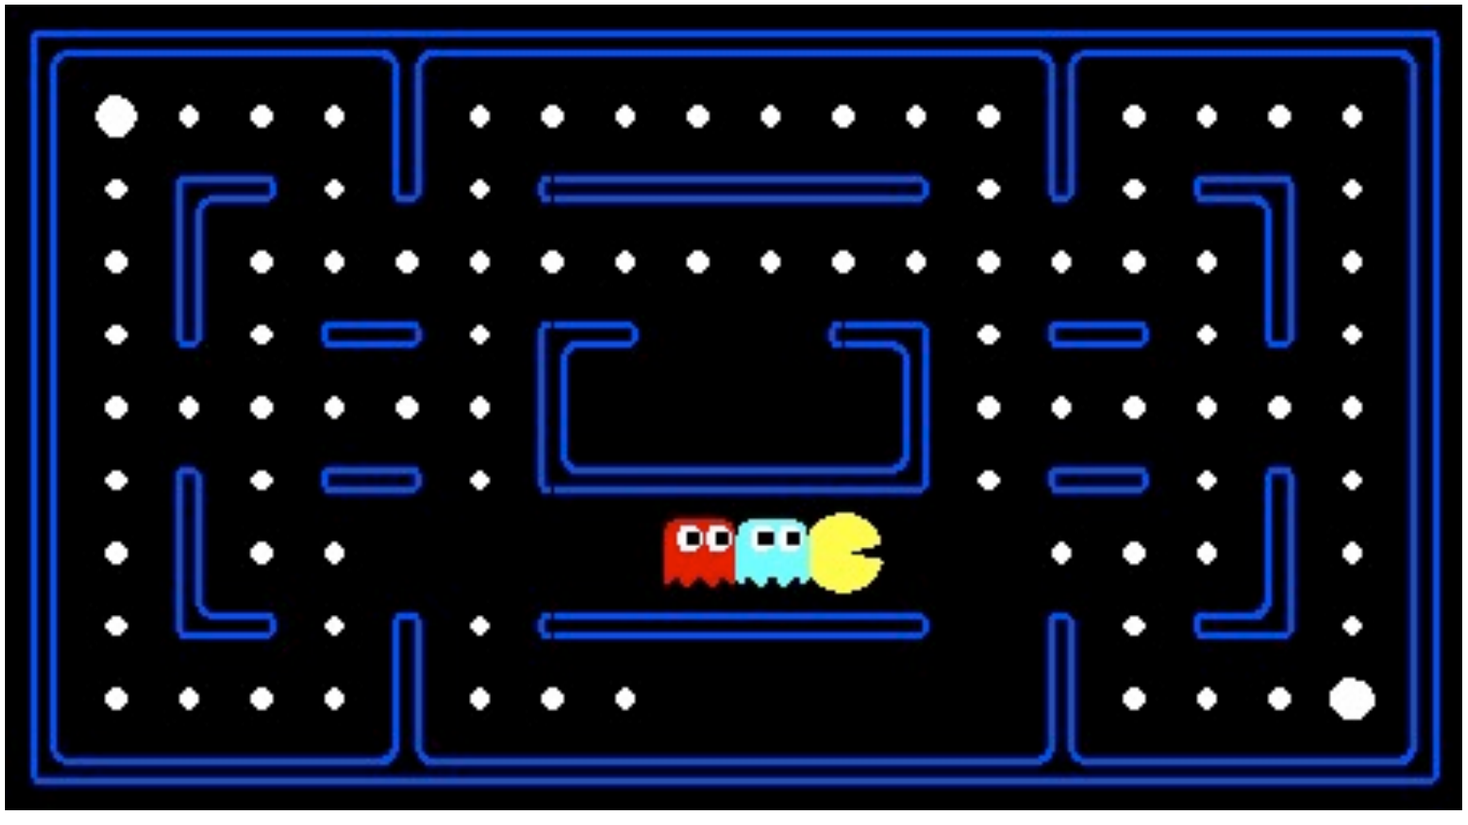

PS3: Planning and Learning

Goal:

- Help Pac-man learn about the world

Techniques:

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

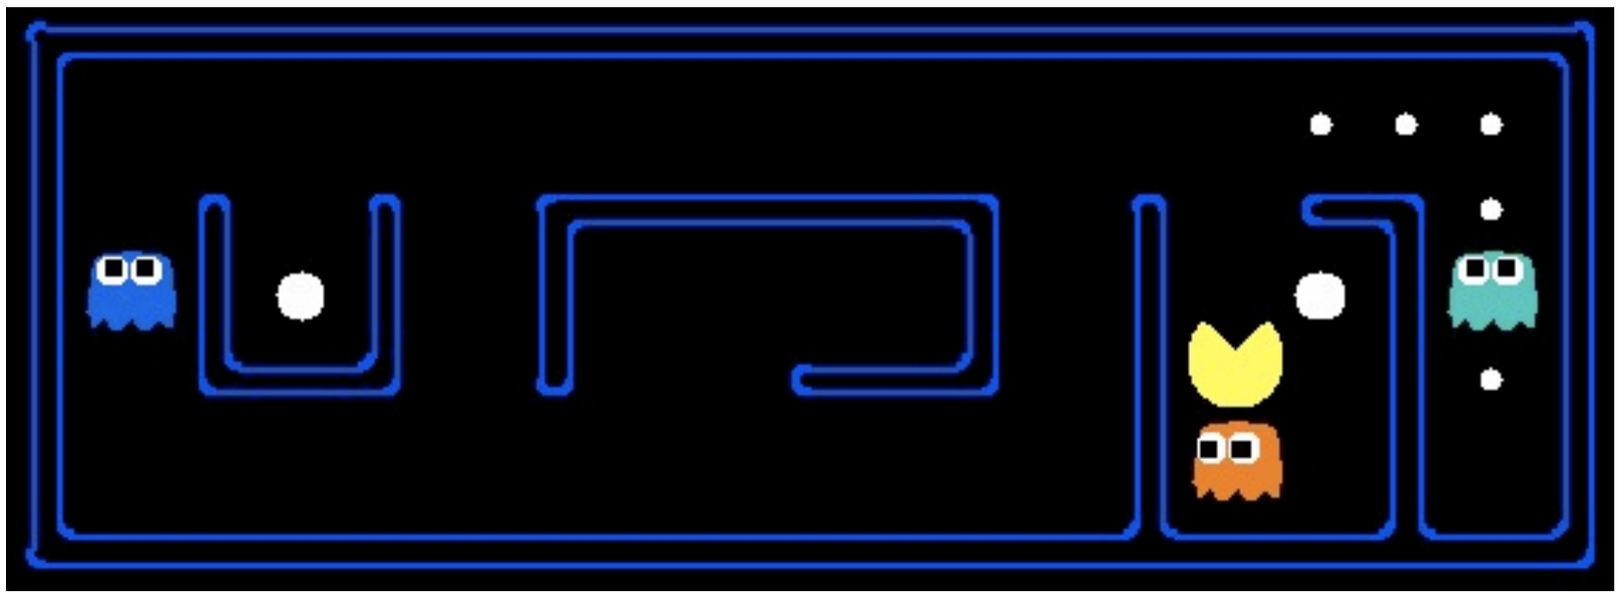

PS4: Ghostbusters

Goal:

• Help Pac-man hunt down the ghosts

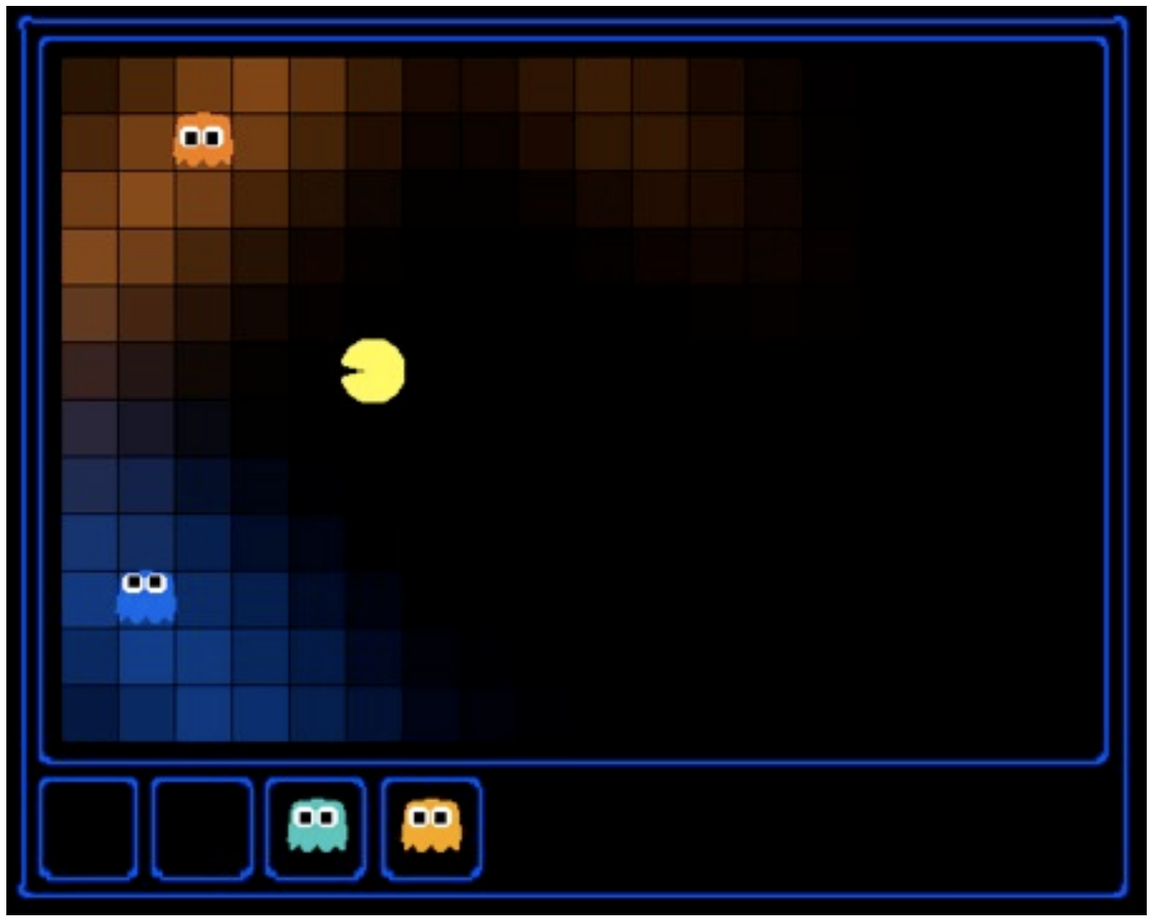

Techniques:

• Probabilistic models: HMMS, Bayes Nets

•Inference: State estimation and particle filtering

To Do:

- Look at the course website:
  - http://www.cs.washington.edu/cse473/10au/
- Do the readings
- Do the python tutorial