CSE 573: Artificial Intelligence

Instructor: Luke Zettlemoyer

Web: http://www.cs.washington.edu/cse573/10au/

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

What is AI?

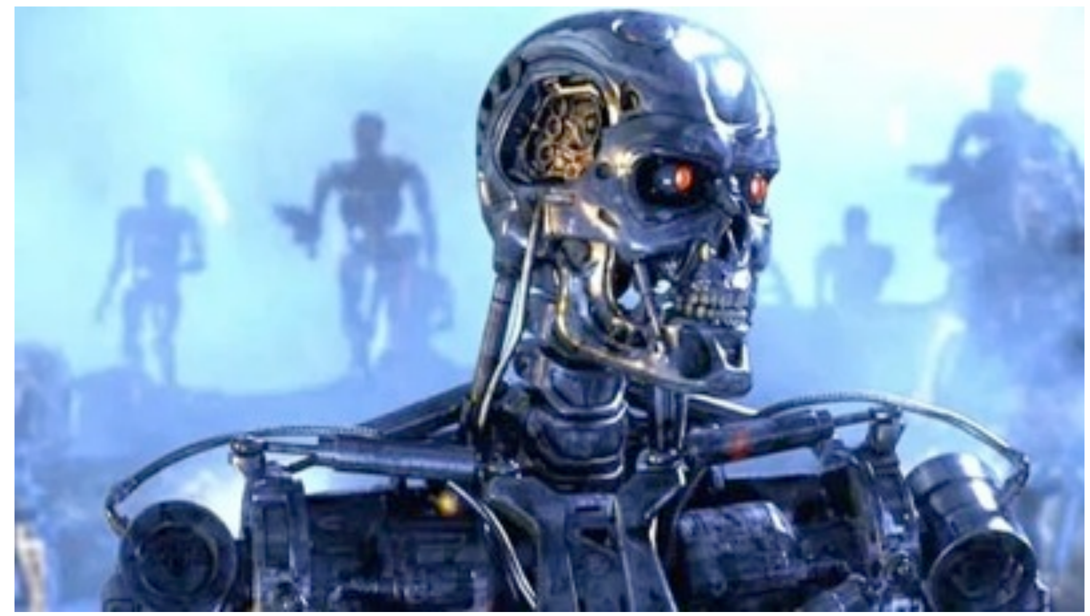

Could We Build It?

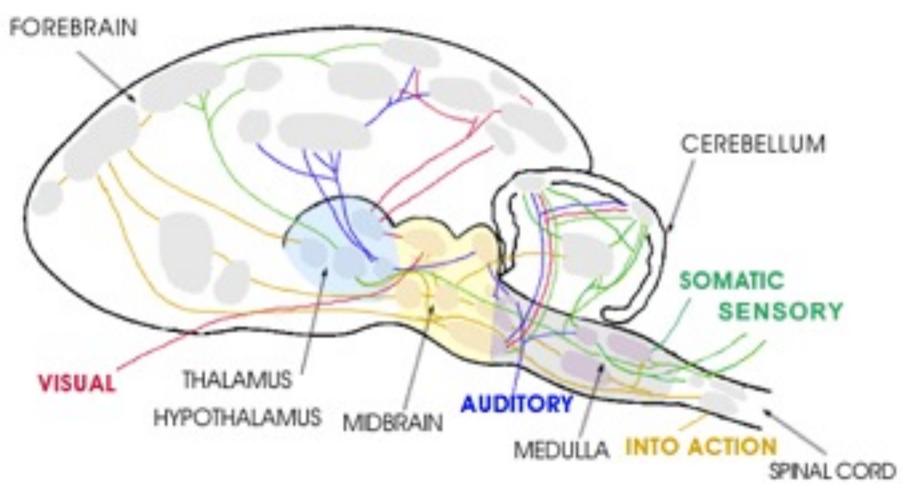

10^{11} neurons
10^{14} synapses
cycle time: 10^{-3} sec

vs.

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

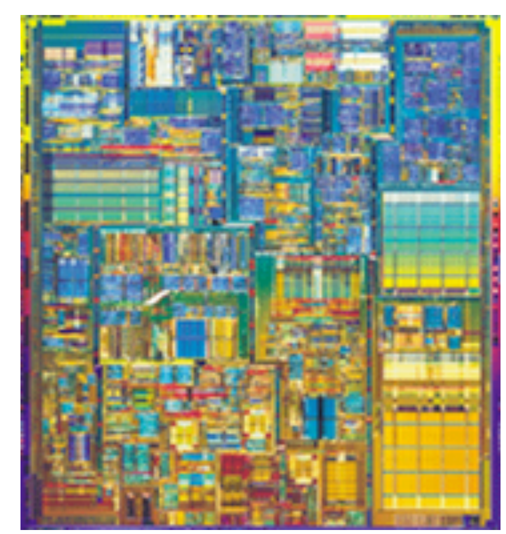

What is CSE 573?

Textbook:

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

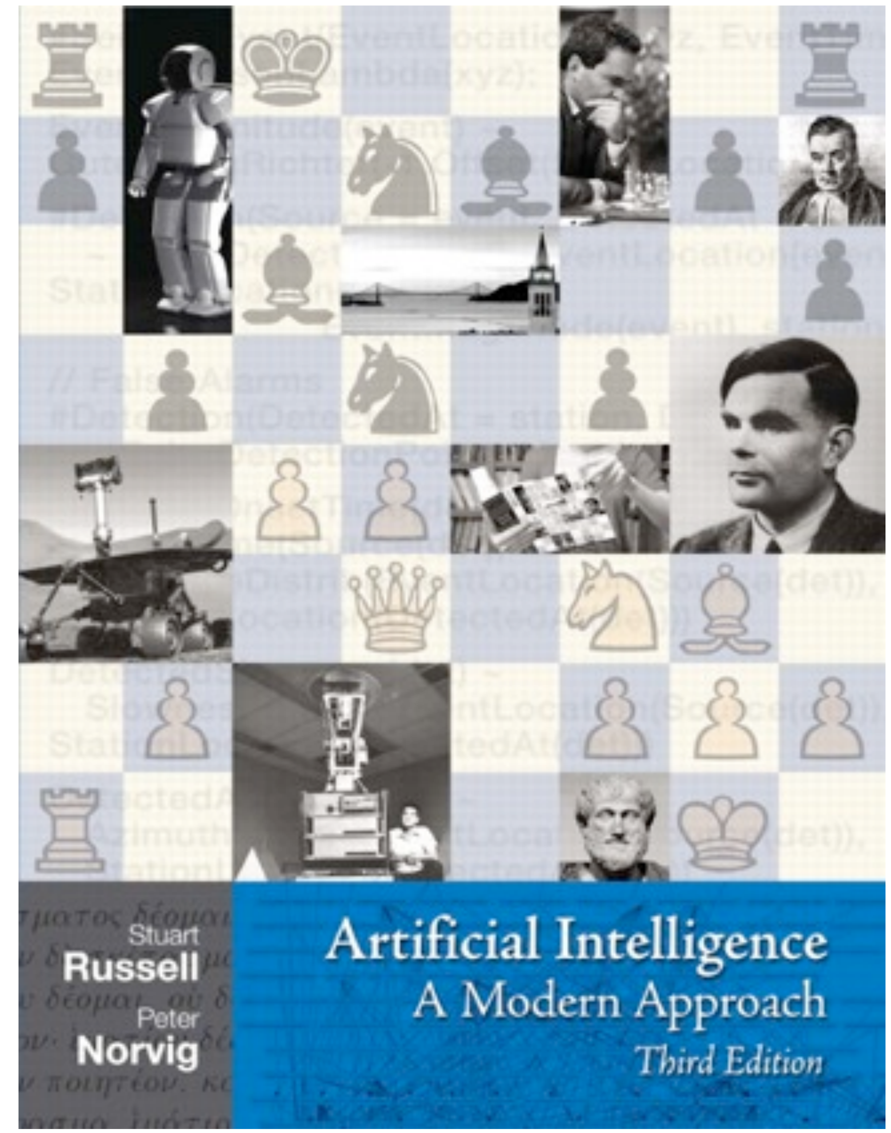

Prerequisites:

• CSE 421: Algorithms, or equivalent

• basic exposure to probability, data structures, and logic

Work:

• Readings (mostly from text), Programming / written homework (40%), In class exam (30%), Final mini-project (20%), Class participation (10%)

Topic Overview

CSE573 Lecture Slides

▷ CSE Home

Administrivia
- Home
- Using course email
- Email archive
- Policies

Content
- Topic Overview
- Slides & Reading

Assignments
- Reading Reports
- Problem sets

Lecture Slides [schedule subject to change]

Week	Dates	Topics & Slides	Readings	Notes
1	Sept 30	Introduction	R&N Ch 1, Ch 2	
2	Oct 5 Oct 7	Agents and Search Heuristic search	R&N Chap 3	
3	Oct 12 Oct 14	Game Trees Utility Theory		
4	Oct 19 Oct 21	Markov Decision Processes (MDPs) MDPs continued		
5	Oct 26 Oct 28	Reinforcement Learning (RL) RL continued		
6	Nov 2 Nov 4	Hidden Markov Models (HMMs) HMM Inference		
7	Nov 9 Nov 11	Bayes Net (BNs) No class (Veterans day)	N/A	
8	Nov 16 Nov 18	BN Inference In class mid-term		
9	Nov 23 Nov 25	First-order logic (FOL) No Class (Thanksgiving)		
10	Nov 30 Dec 2	FOL Inference Resolution		
11	Dec 7 Dec 9	TBD: Language / Vision / Robotics TBD: Language / Vision / Robotics		

Additional Info

ivia

urse email
hive

erview
Reading

nts
sets

Lectures:	Tues/Thu	12:00pm-1:20pm EEB 042

Who	**Office Hours**
Luke Zettlemoyer, Instructor lsz at cs-dot-washington-dot-edu	Wednesdays 10-11am in CSE 658, and by email
Adrienne Wang, TA axwang at cs-dot-washington-dot-edu	TBA and by email

Textbooks:

- *Stuart Russell & Peter Norvig, Artificial Intelligence: A Modern Approach, Prentice-Hall, Third Edition (2010).*

Grading: *Problem Sets 40%; midterm 30%; mini-project 20%; class participation 10%*

Assignments: Pac-man

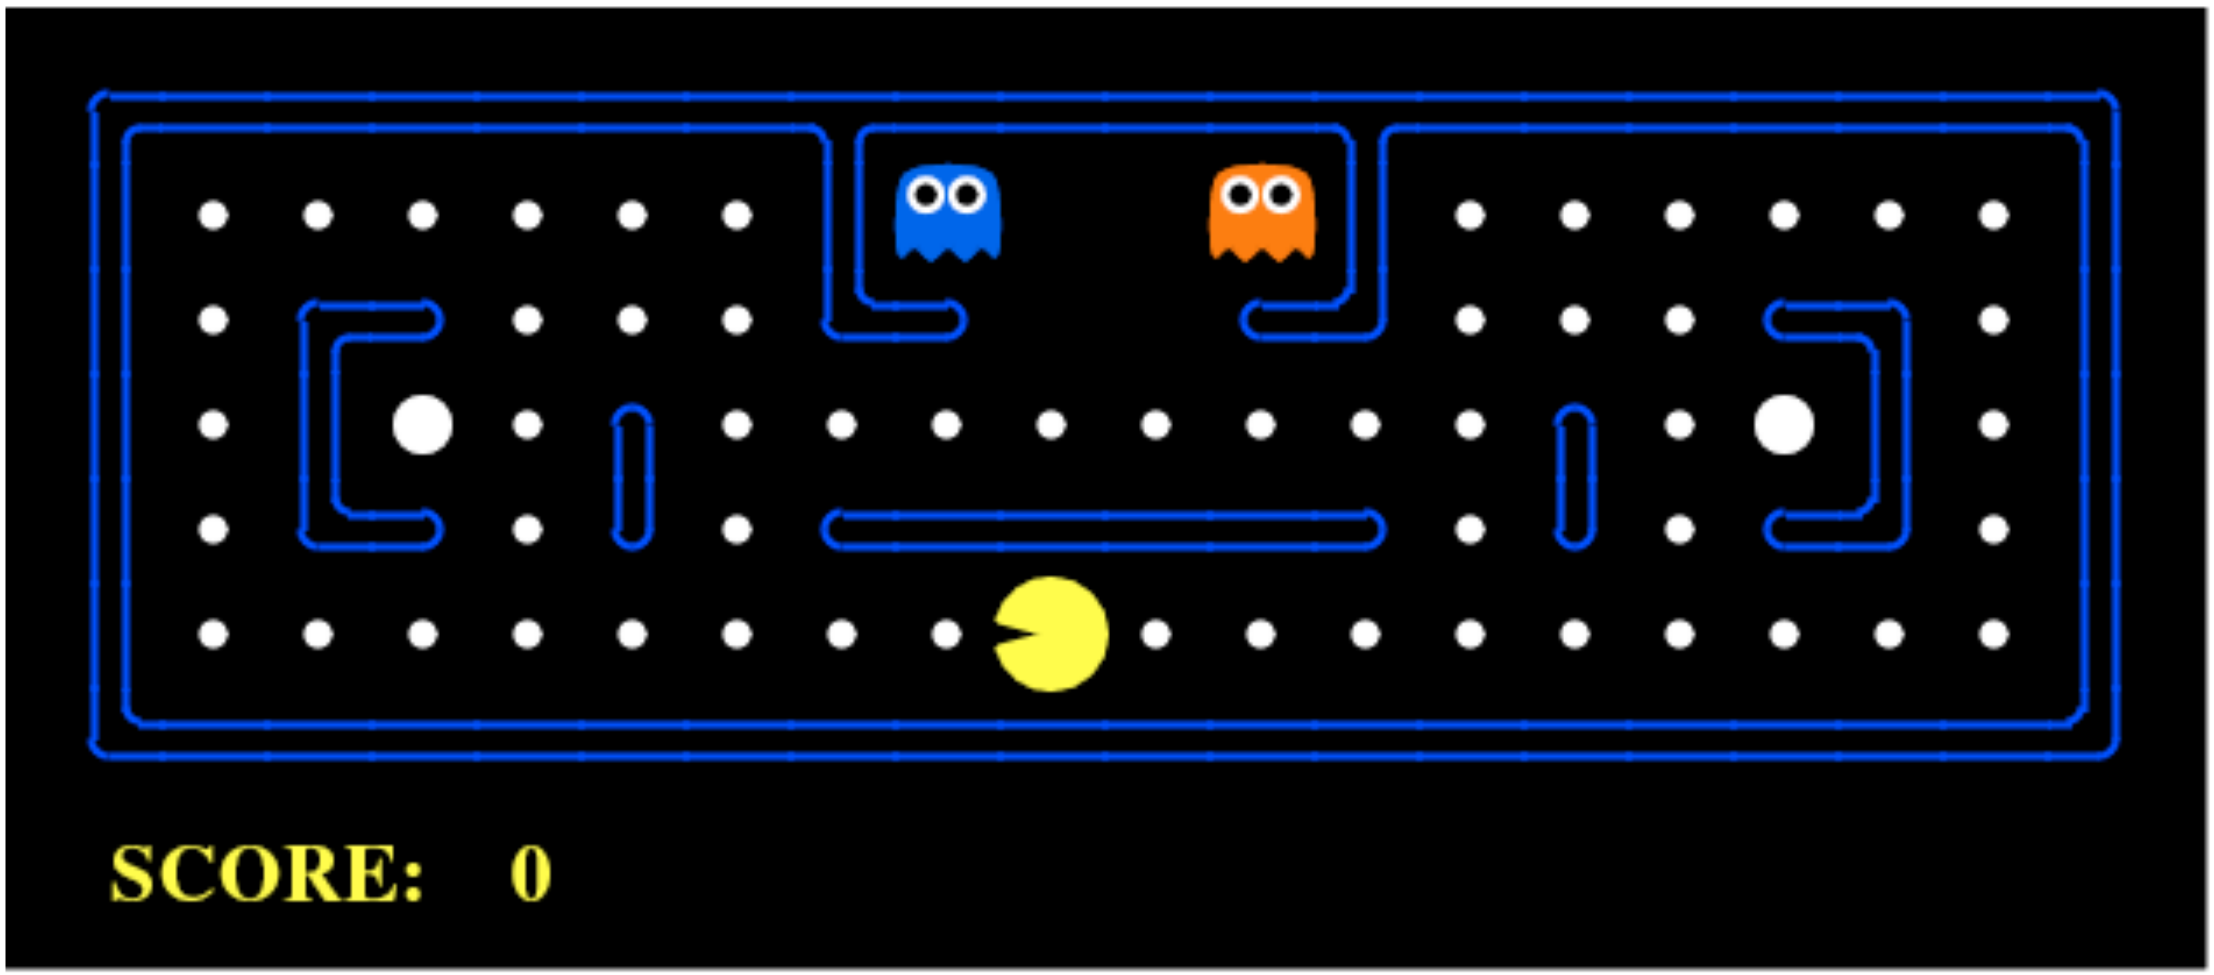

Originally developed at UC Berkeley:

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

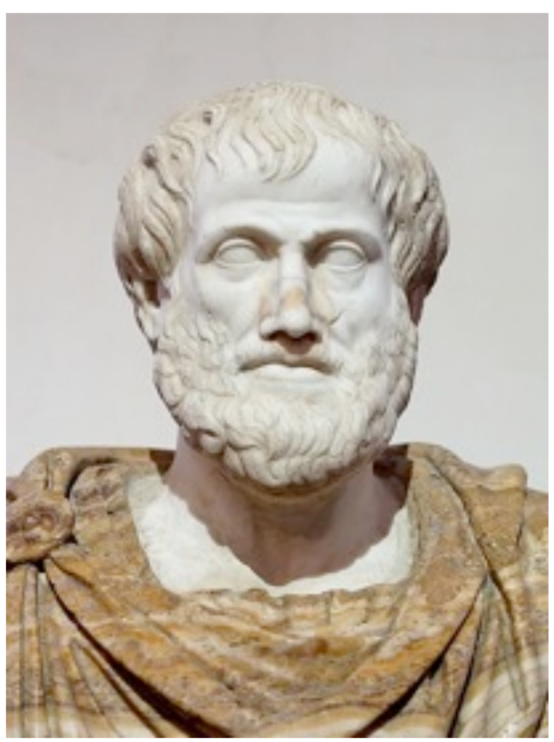

and

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:*

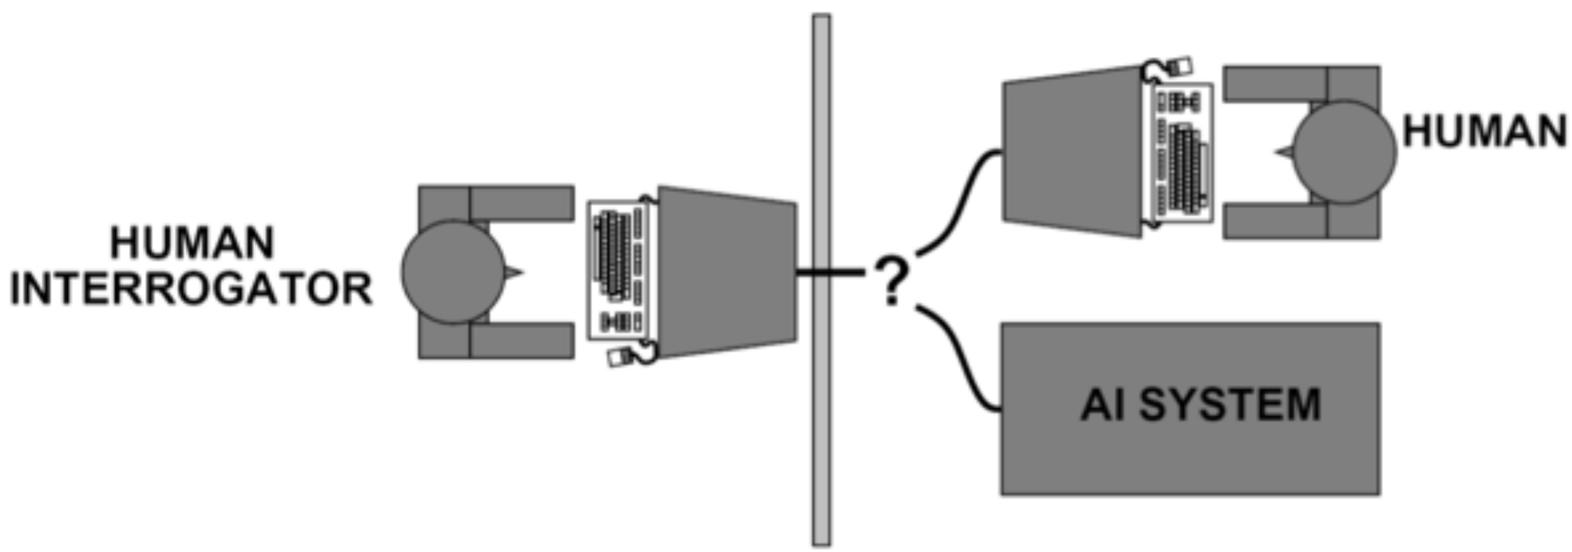

  - Suggested major components of AI: knowledge, reasoning, language understanding, learning

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

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.

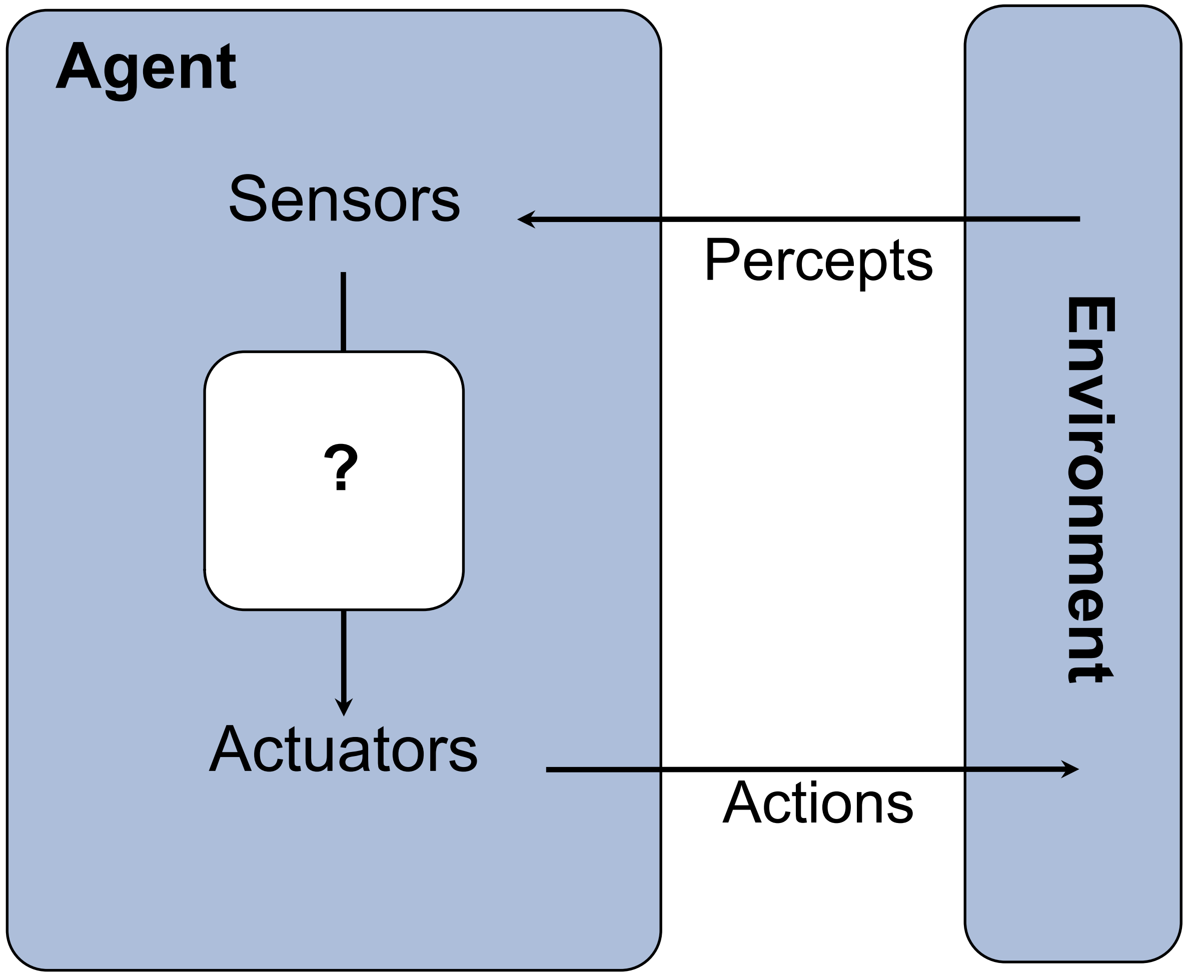

- This course is about:
  - General AI techniques for a variety of problem types
  - Learning to recognize when and how a new problem can be solved with an existing technique

Pacman as an Agent

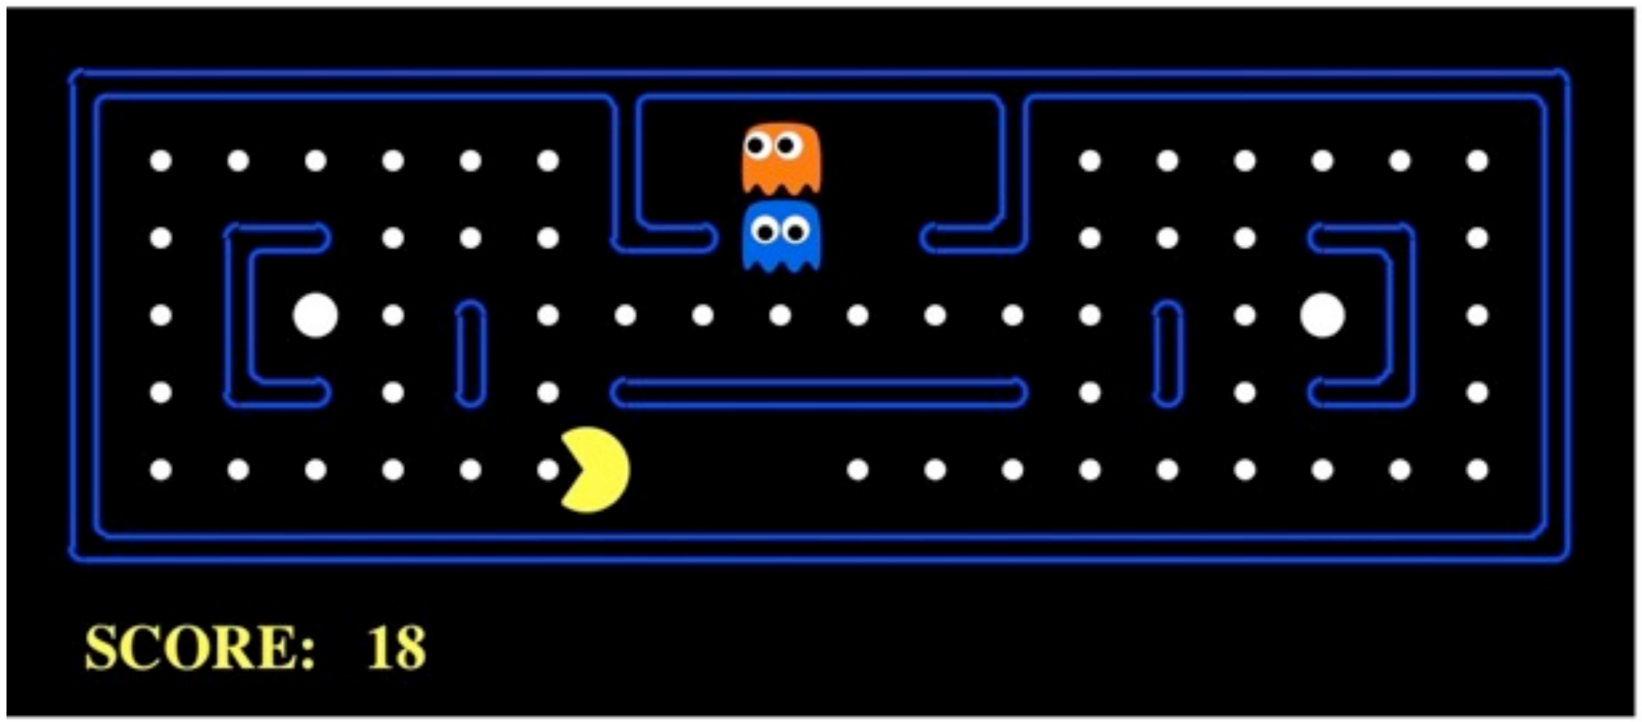

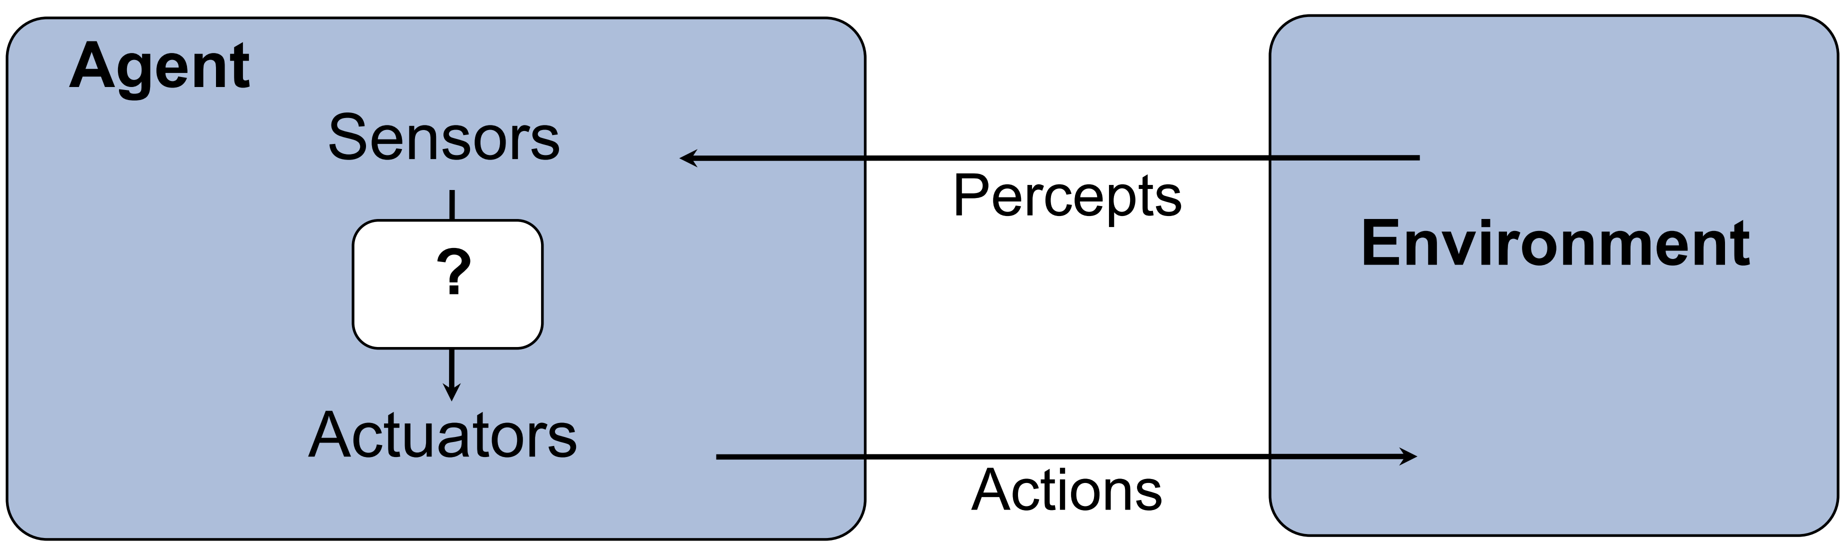

Types of Environments

- 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.

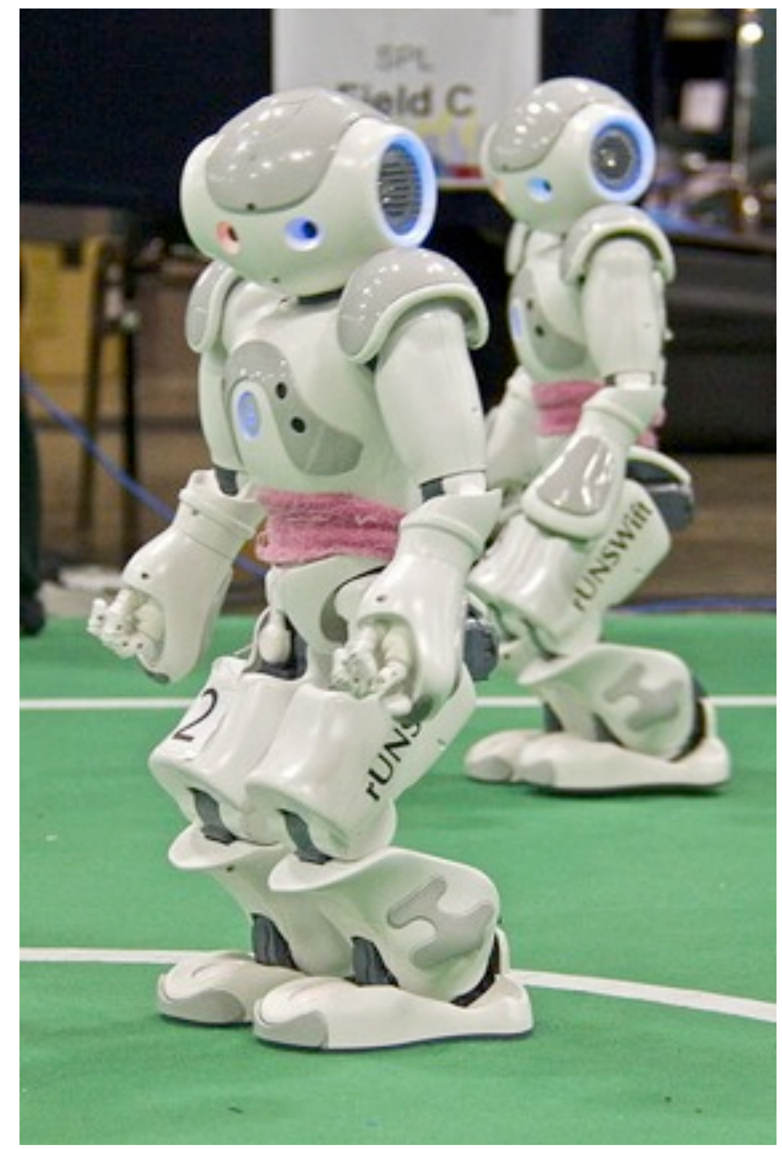

Single agent vs. Multiagent

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?

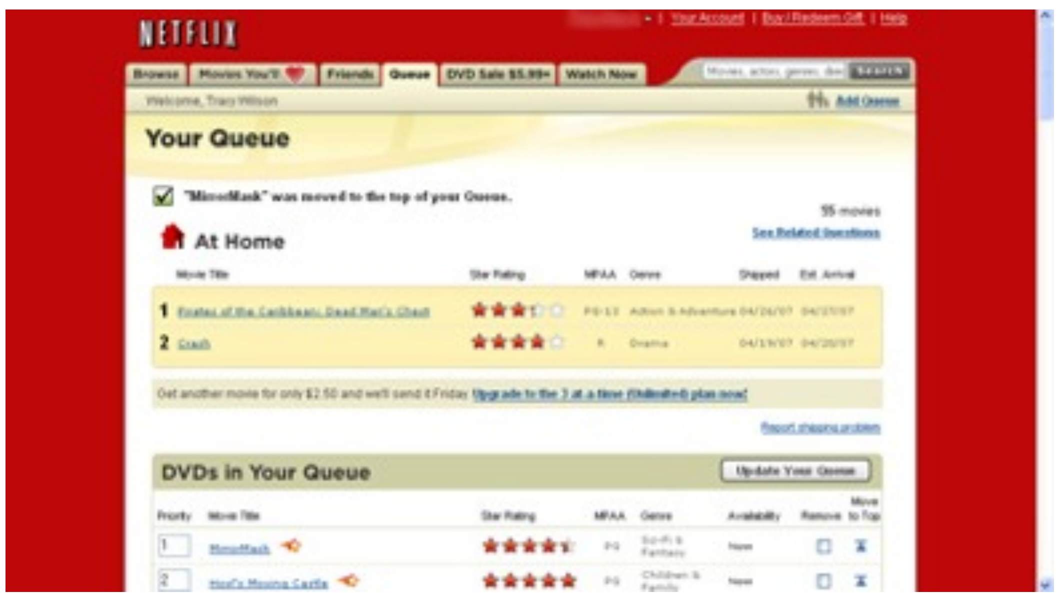

vs.

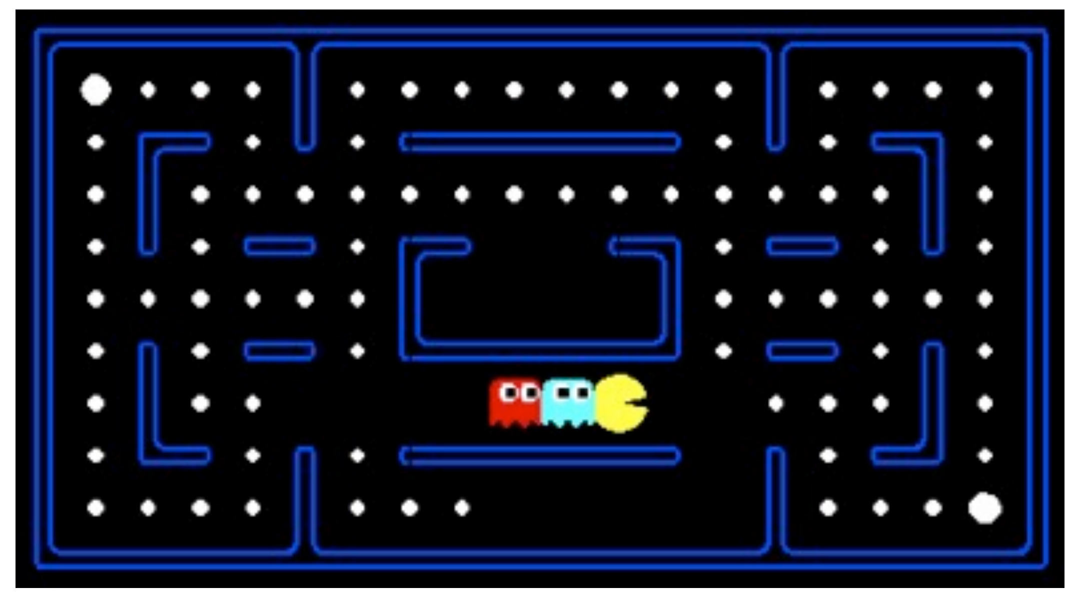

Discrete vs. Continuous

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

vs.

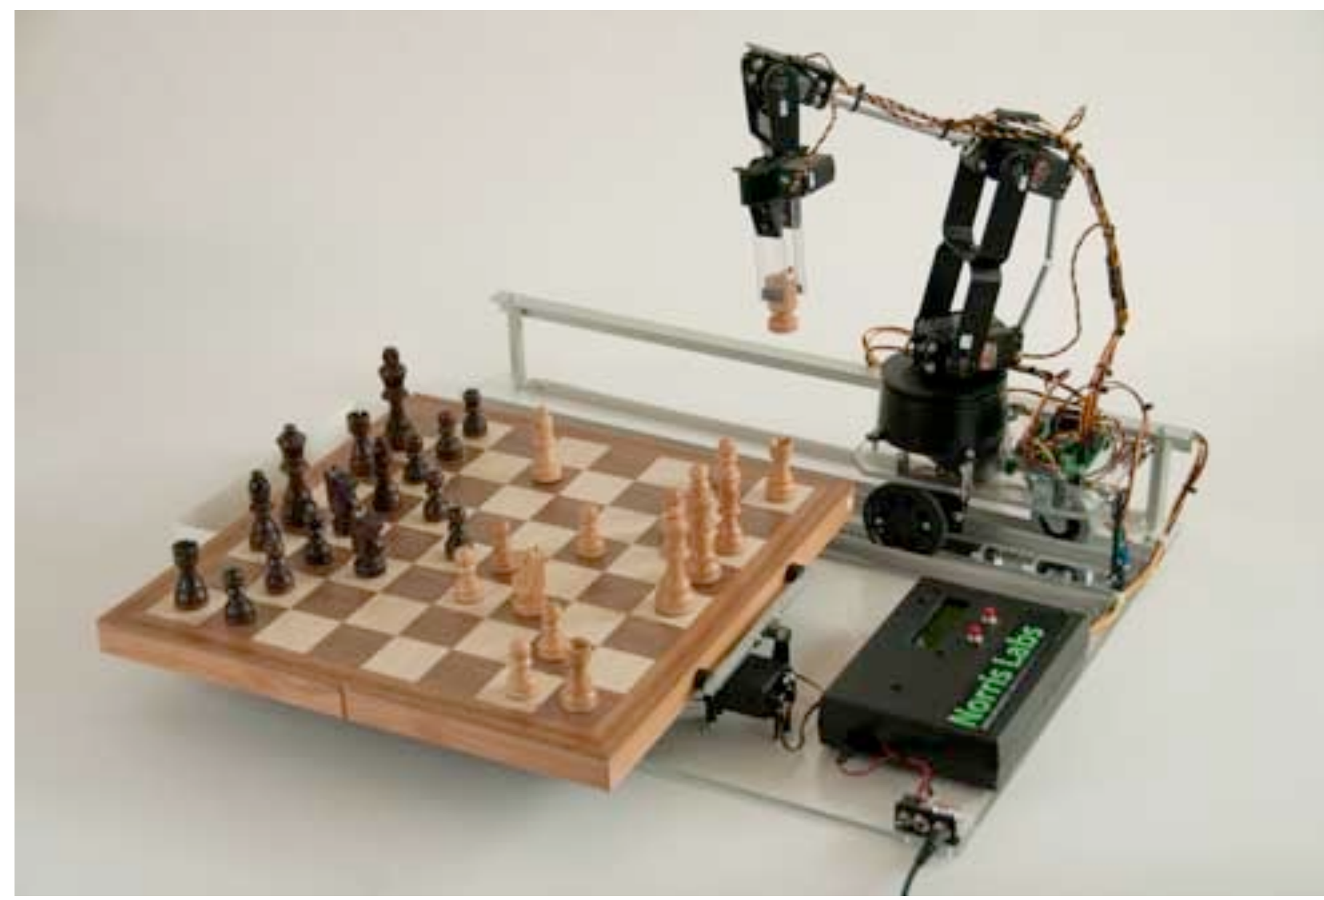

Assignments: Pac-man

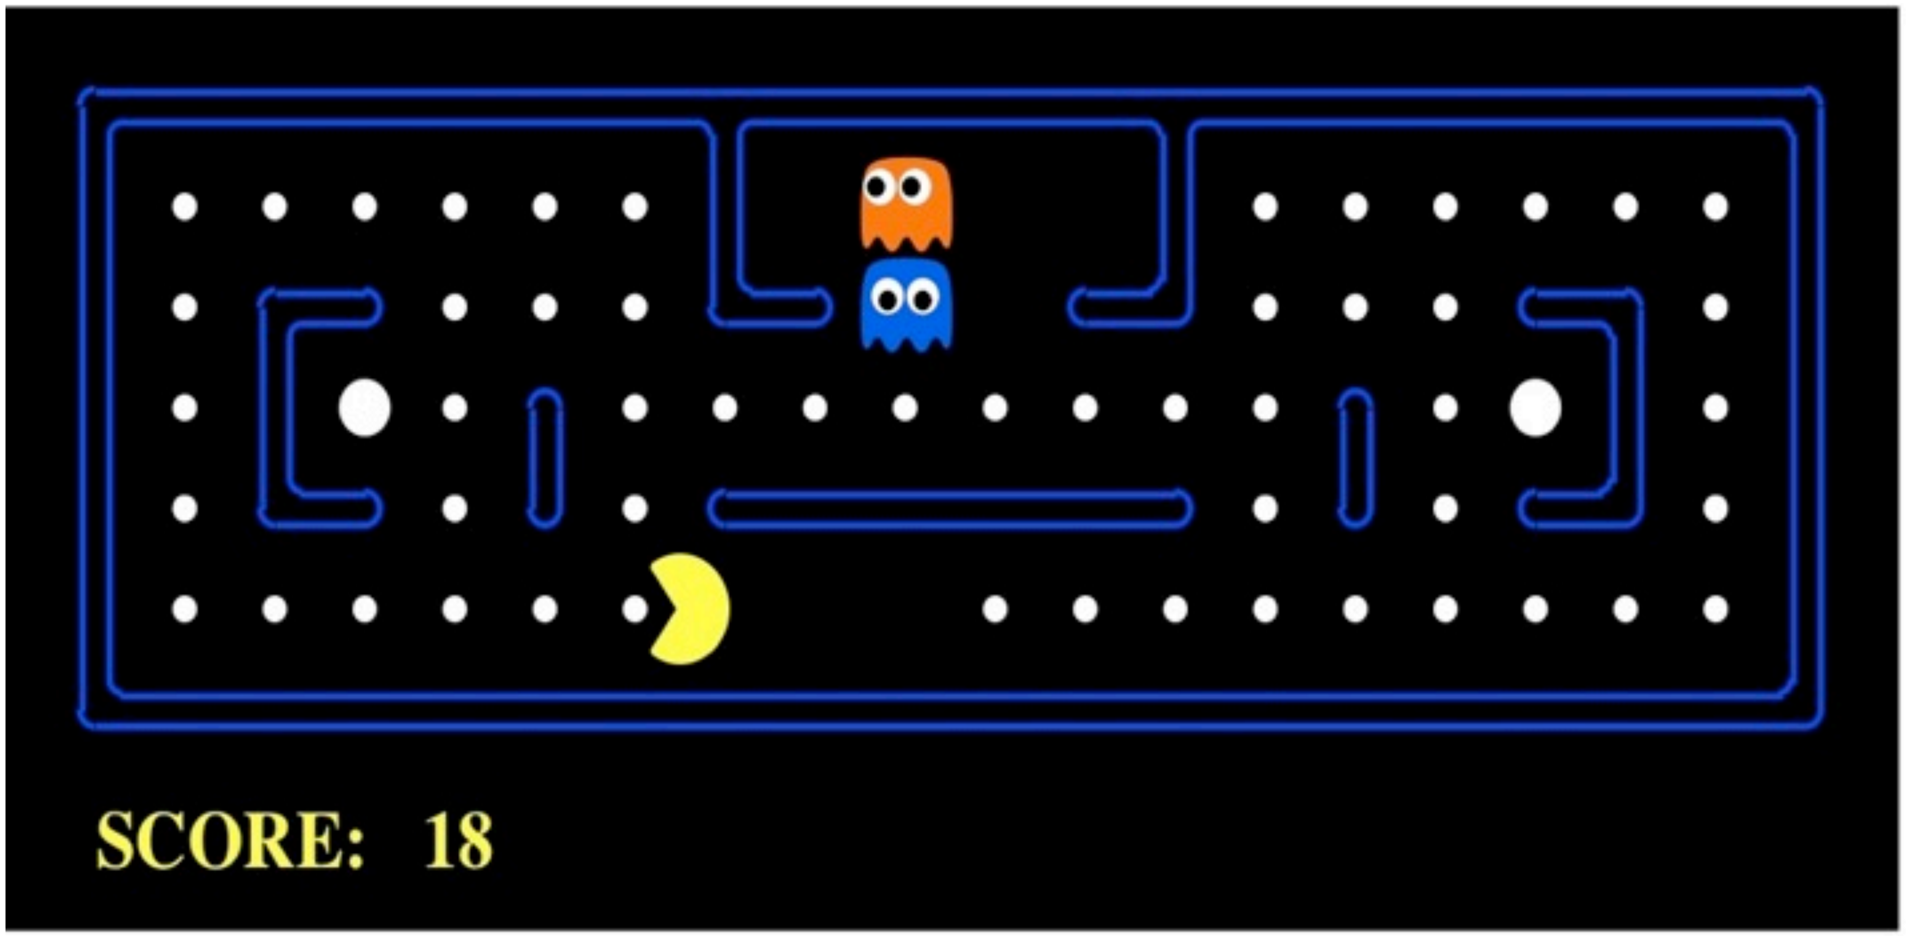

Originally developed at UC Berkeley:

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.

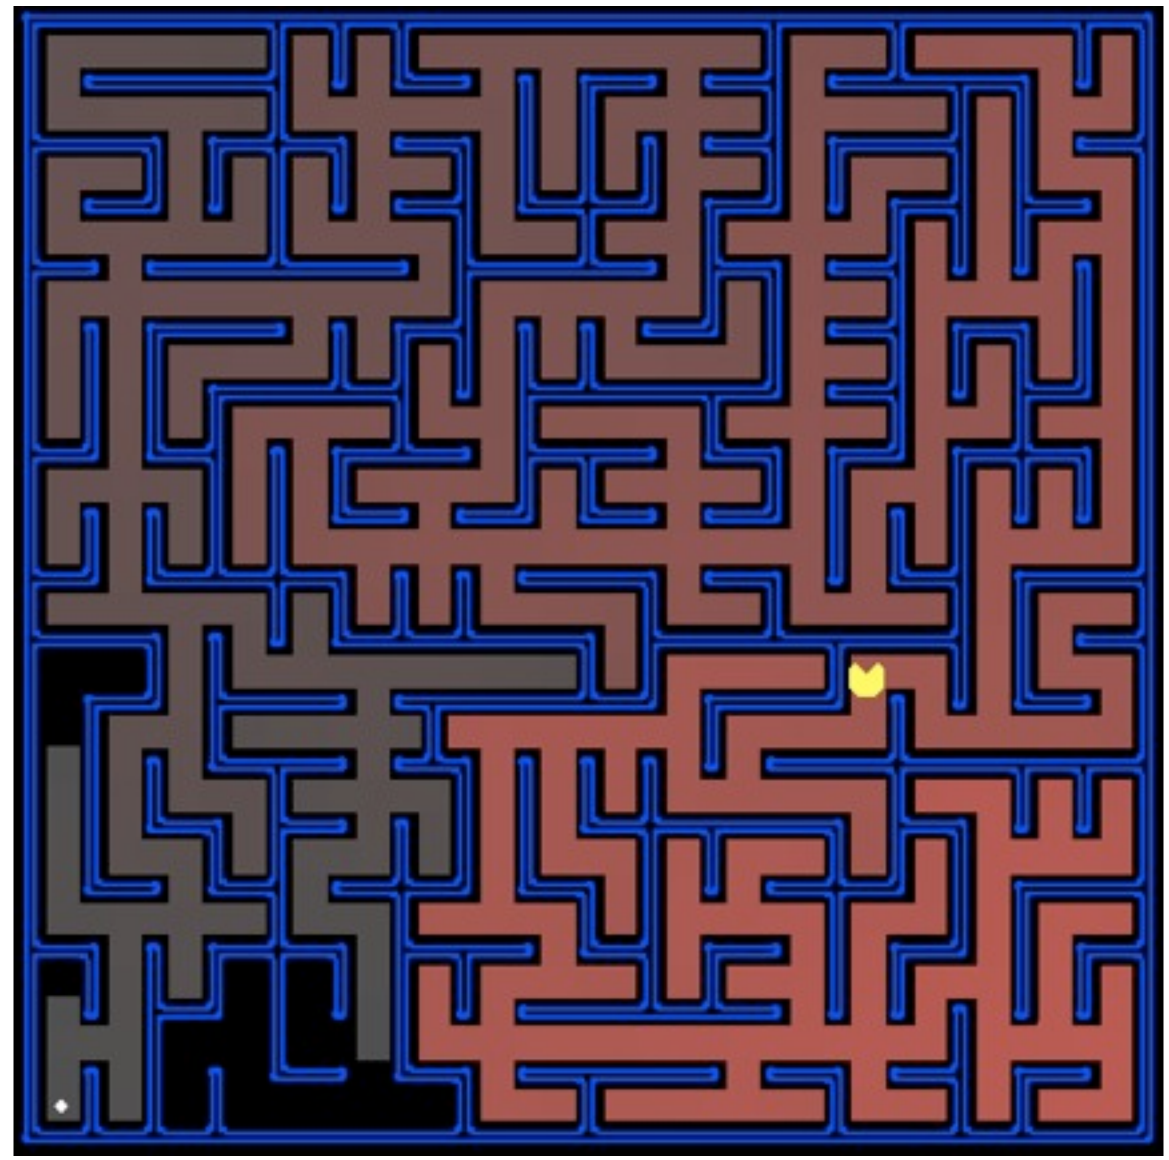

PS2: Game Playing

Goal:

- Play Pac-man!

Techniques:

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

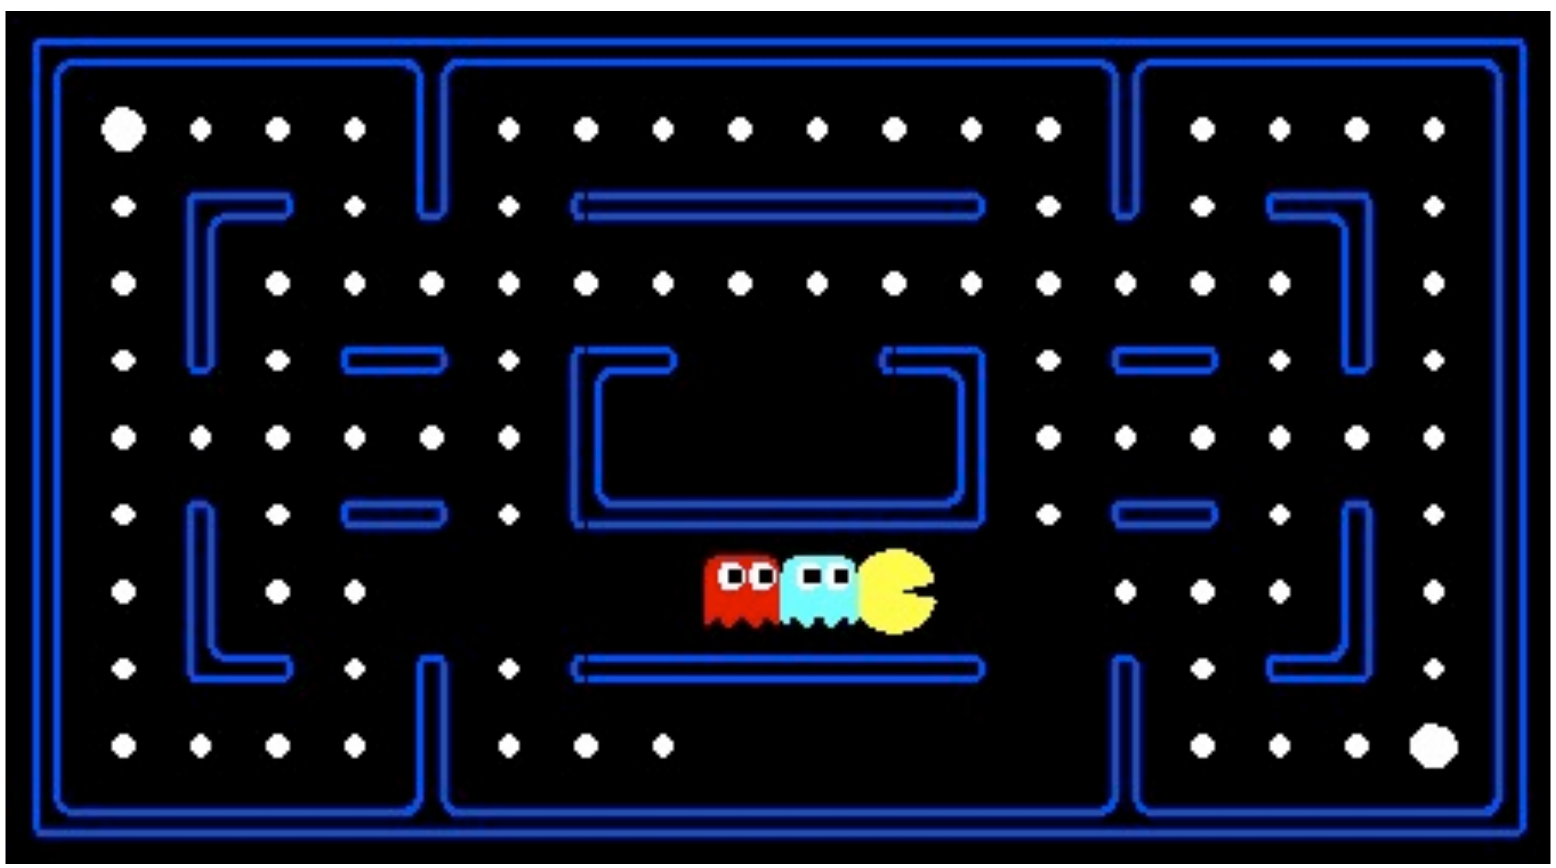

PS3: Planning and Learning

Goal:

- Help Pac-man learn about the world

Techniques:

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

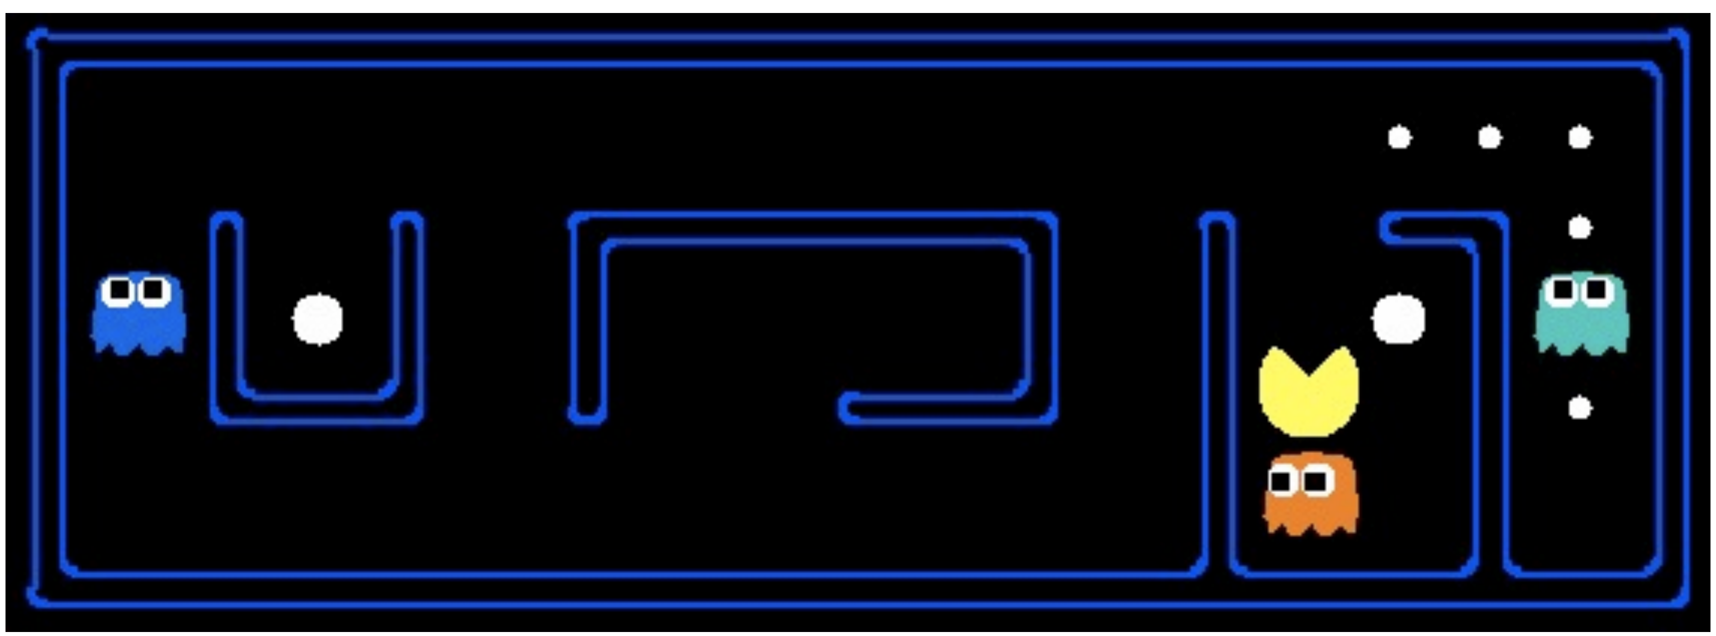

PS4: Ghostbusters

Goal:

- Help Pac-man hunt down the ghosts

Techniques:

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

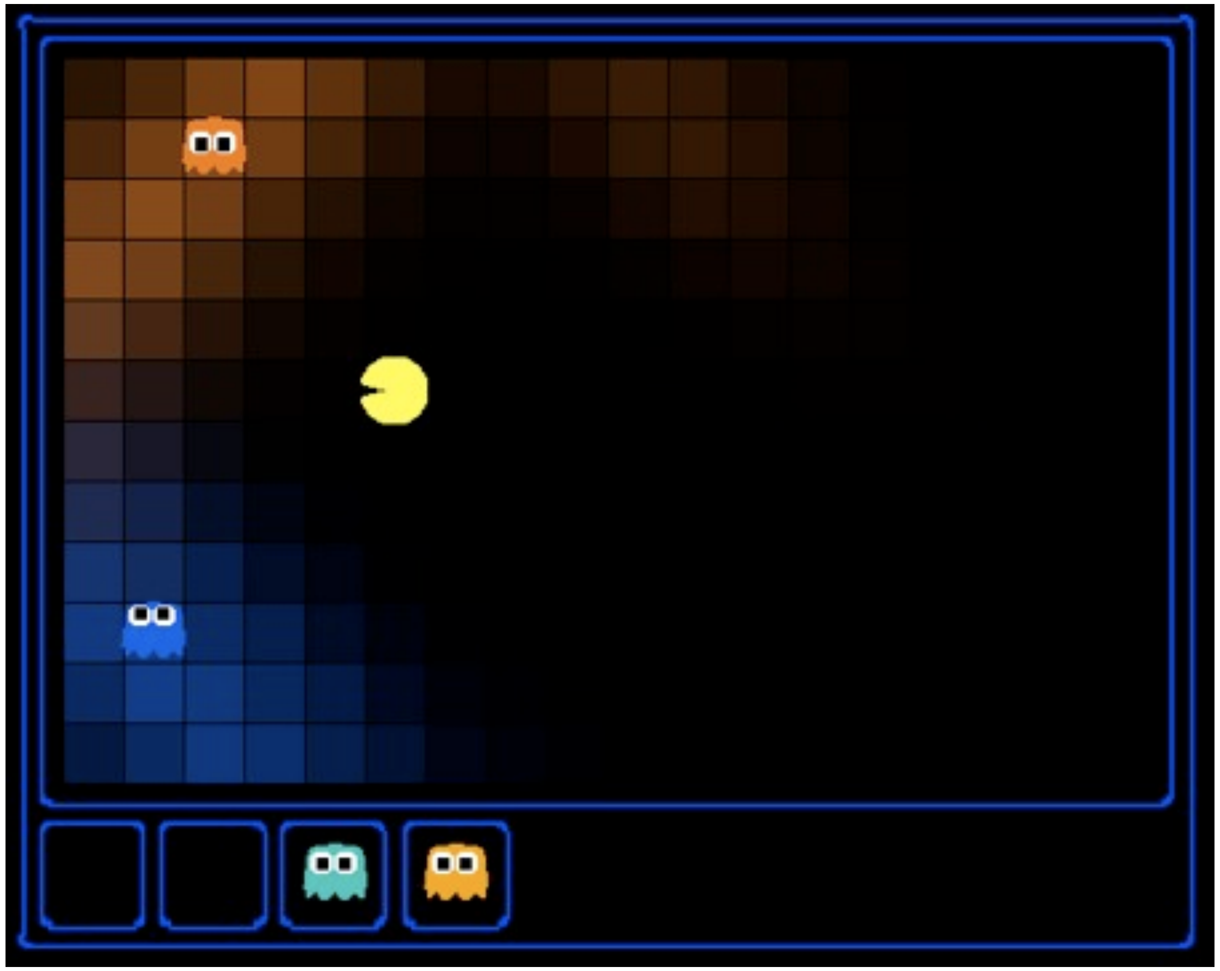

To Do:

- Look at the course website:
  - http://www.cs.washington.edu/cse573/10au/
- Add yourself to the email list
- Do the readings
- Do the python tutorial