CSE 573:
Artificial Intelligence

Hanna Hajishirzi

https://courses.cs.washington.edu/courses/cse573/19wi

Several slides from Luke Zettlemoyer, Dan Klein, Dan Weld,
Stuart Russell, Andrew Moore

Logistics

Textbook:

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

- Useful: Mausam & Andrey Kolobov. Planning with Markov Decision Processes: An AI Perspective (free online version if accessed from UW)

- Useful: Richard Sutton & Andrew Barto, Reinforcement Learning: An Introduction, MIT Press. (limited chapters; freely available online)

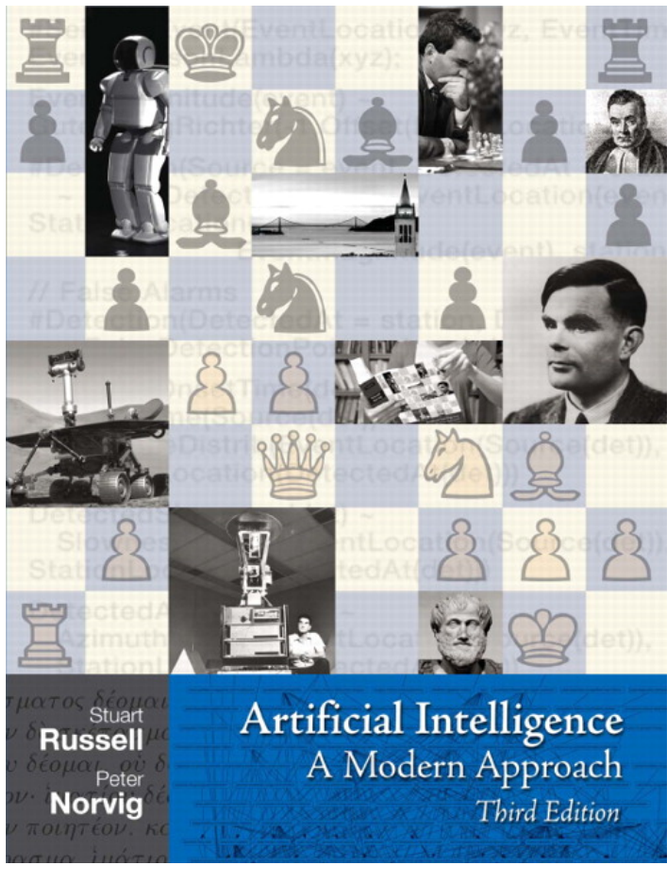

Logistics

Prerequisites:

- Data Structures or equivalent
- Basic exposure to probability, data structures, and logic
- Familiar with Python

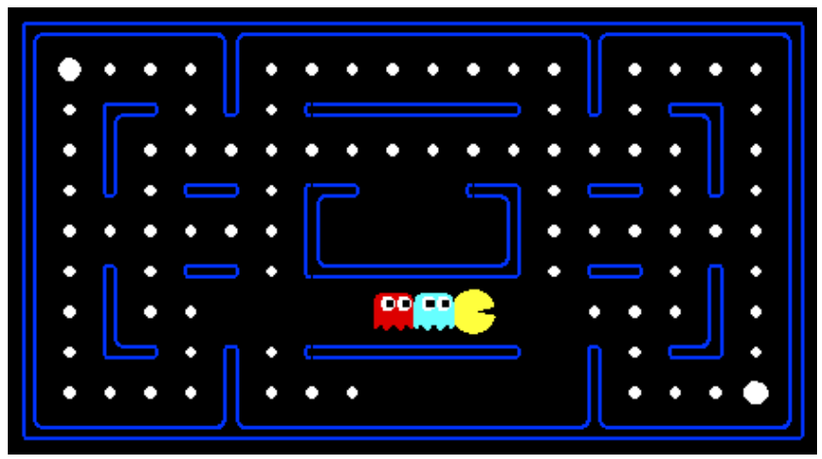

Pacman, autograder

Work:

- Class participation (10%)
- Programming assignment (30%),
- Paper reviews (10%)
- Final Project (25%),
- Two quizzes (25%),

Course Staff

- Instructor:
  - Hanna Hajishirzi
    Research: AI, NLP, ML, Reasoning

- TAs:
  - James Ferguson
  - Svetoslav Kolev
  - Dae Hyun Lee

- Introductions

What you will learn in this course

- 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

- Course Topics:
  - Inference algorithms
  - Decision making under uncertainty
  - Reinforcement learning
  - Representations: Sequential models, Graphical models
  - AI applications

Today

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

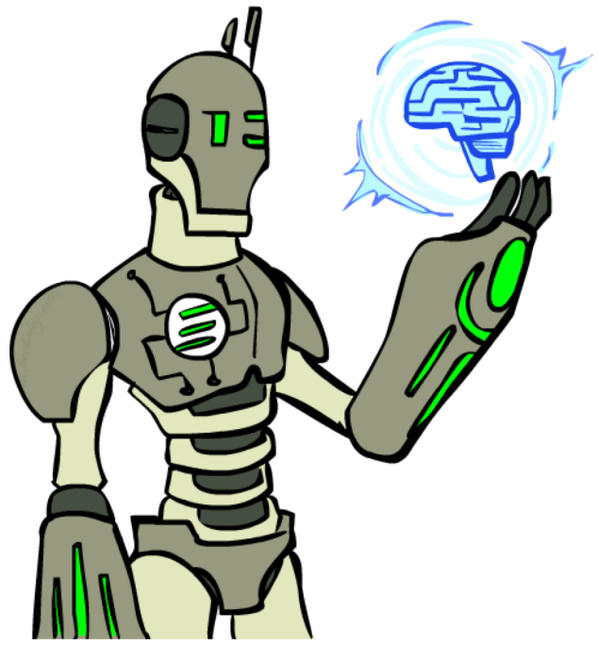

What is AI?

Science of making intelligent machines or computer programs

What Is AI?

The science of making machines that:

Think like humans
Act like humans

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
- 1988-2012: Statistical approaches
- 2012-now: Excitement again?

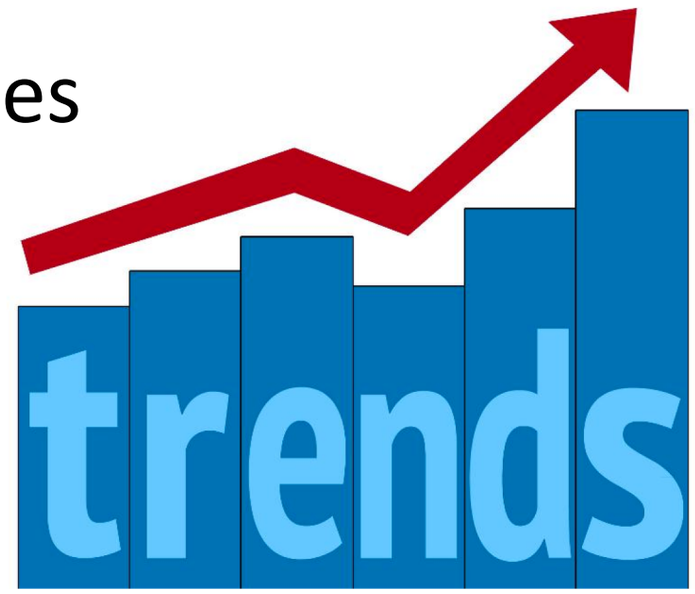

Prehistory

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

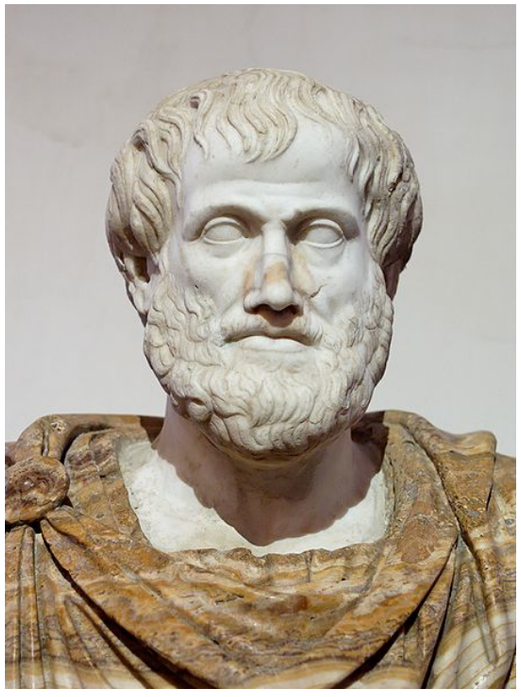

and

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:

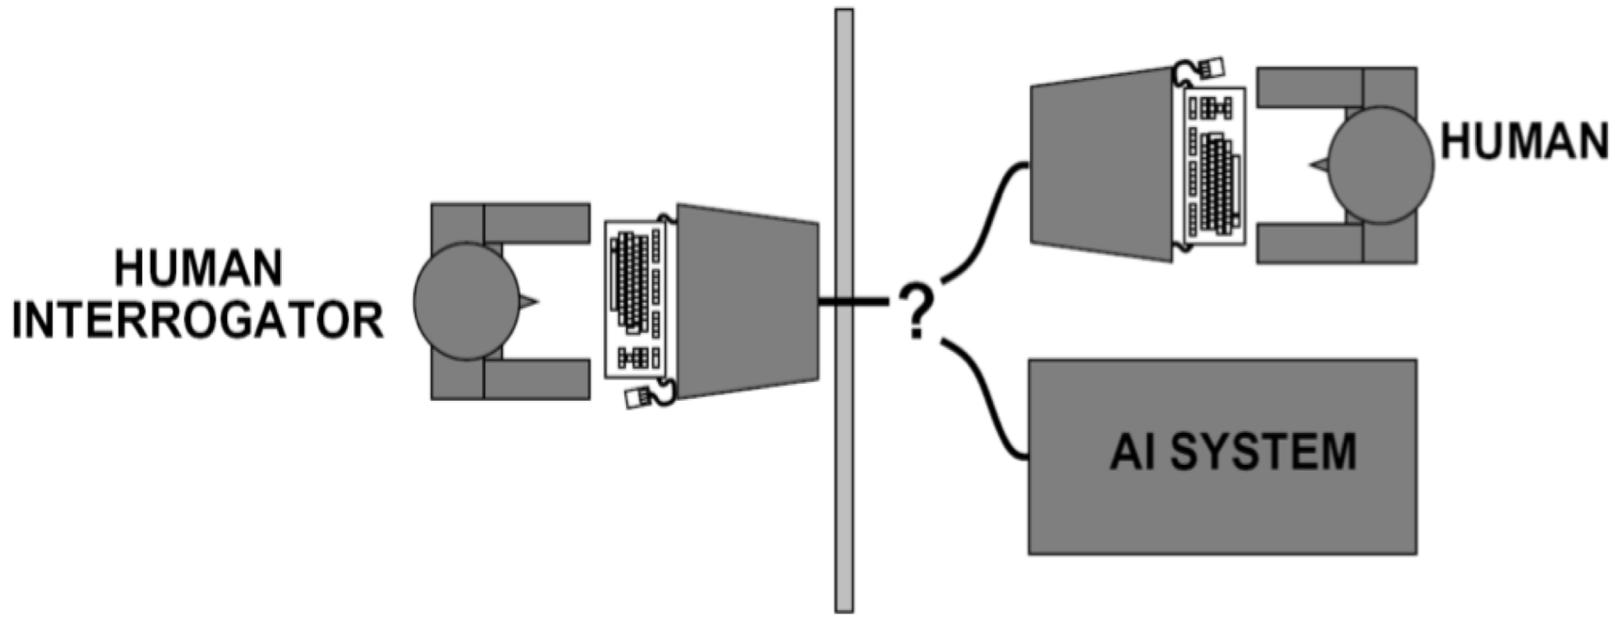

- Suggested major components of AI: knowledge, reasoning, language understanding, learning

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”

1988-2012: Statistical Approaches

- 1985-1990: Probability and Decision Theory win
  - Eg., Bayes Nets
    Judea Pearl - ACM Turing Award 2011

- 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

2012 -- Deep NN Tsunami

"Deep Learning waves have lapped at the shores of computational linguistics for several years now, but 2015 seems like the year when the full force of the tsunami hit the major Natural Language Processing (NLP) conferences." *– Chris Manning*

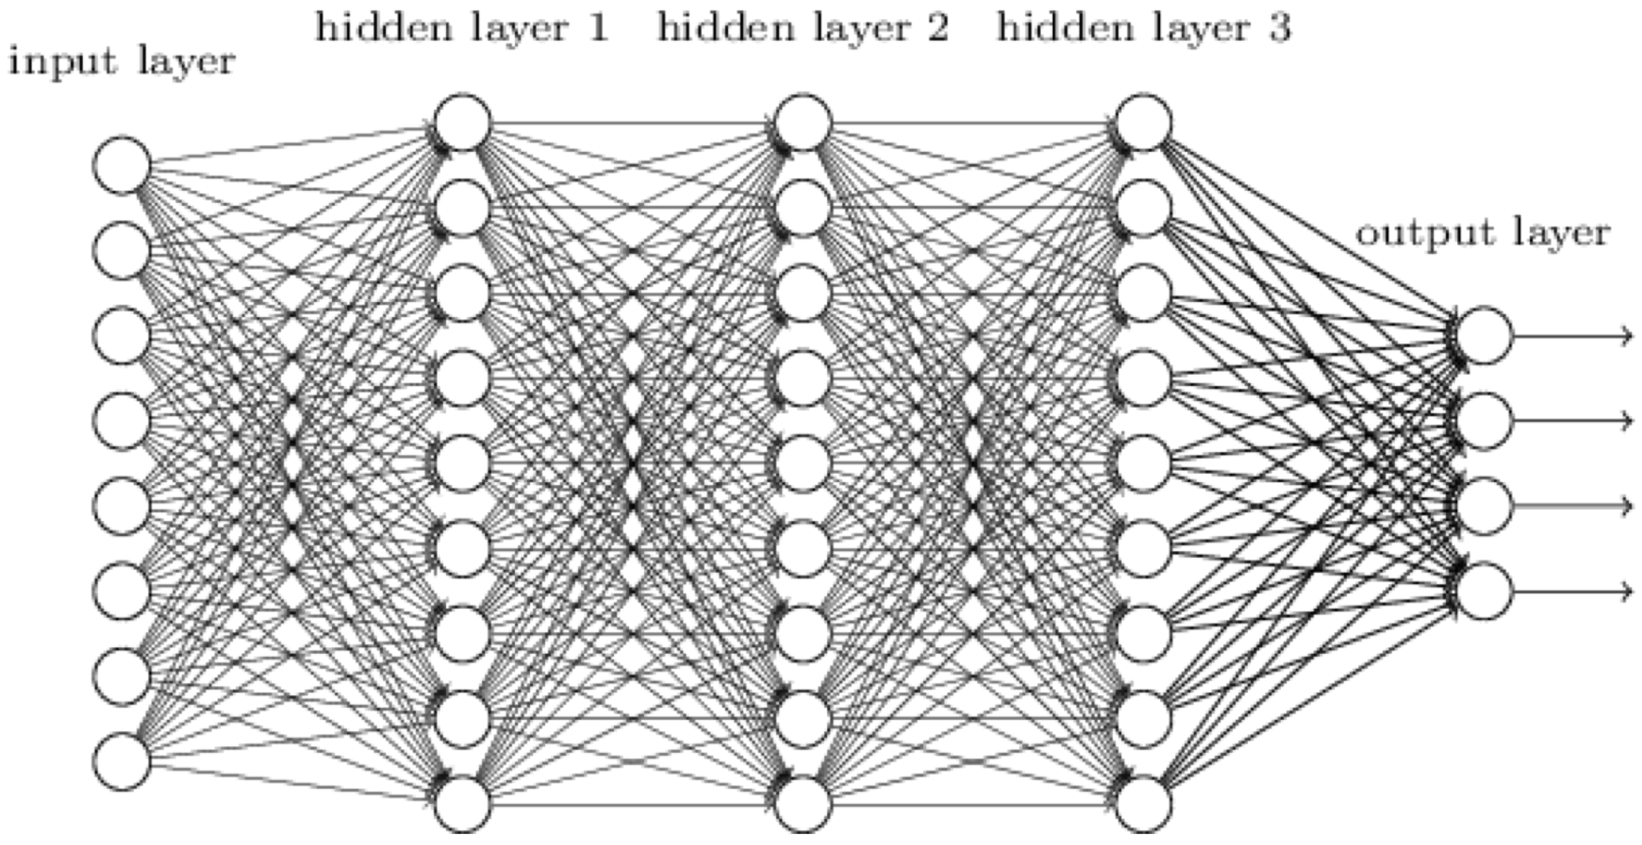

What Can AI Do?

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?
- Translate Chinese into English in real time?
- Write an intentionally funny story?

State of the art: 1997

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

-Gary Kasparov

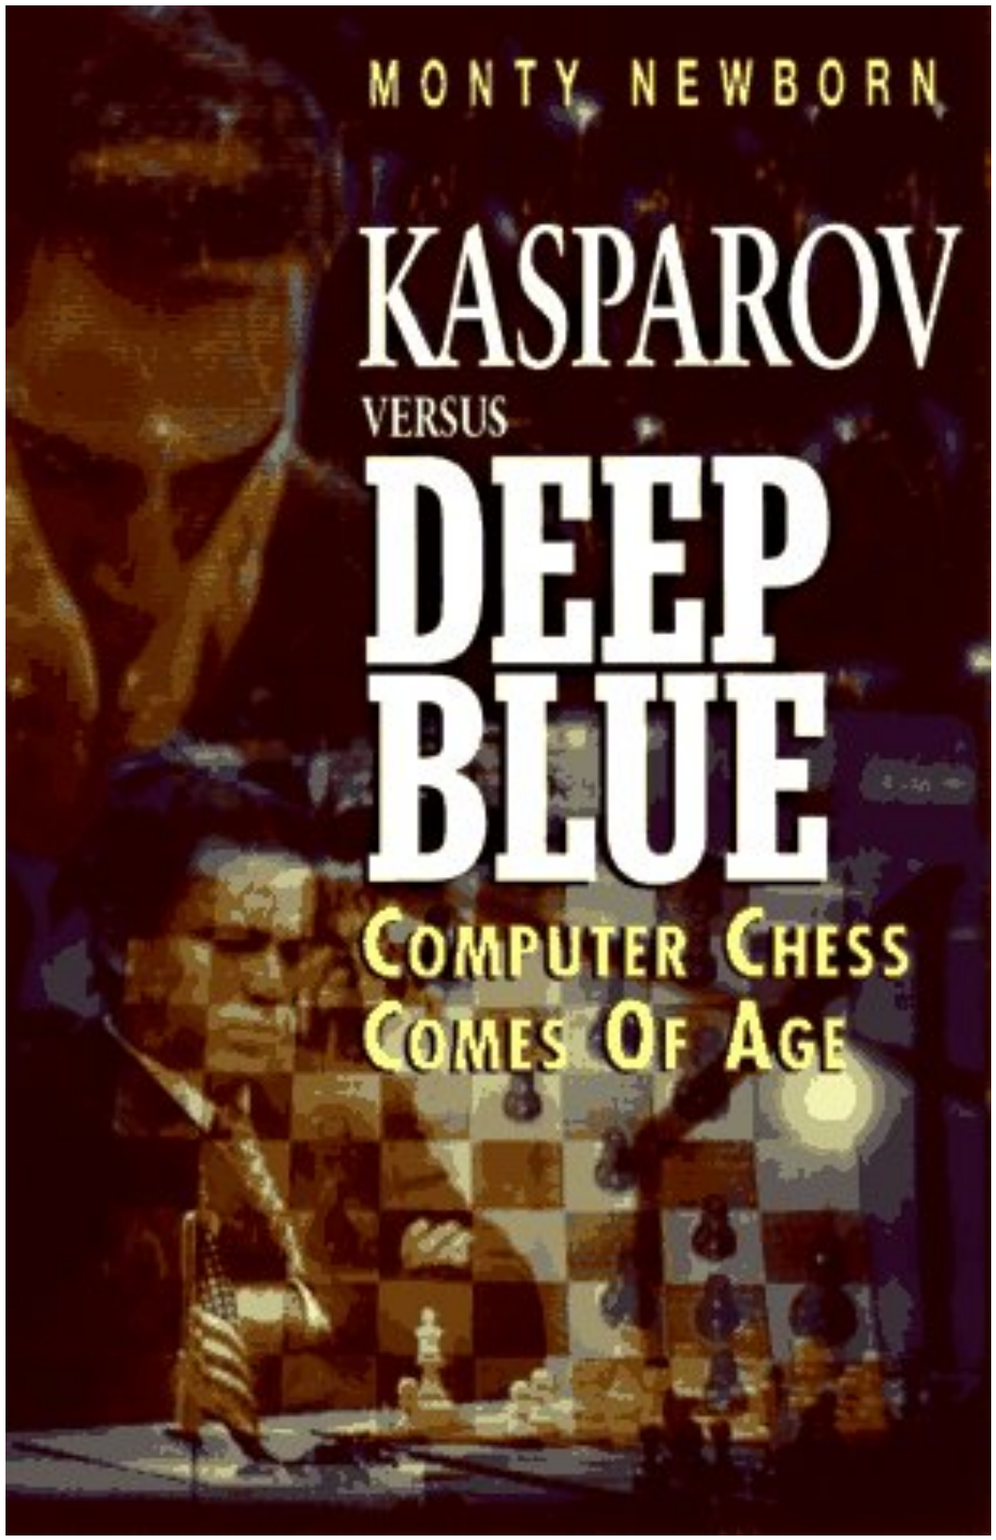

Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings.

– Drew McDermott

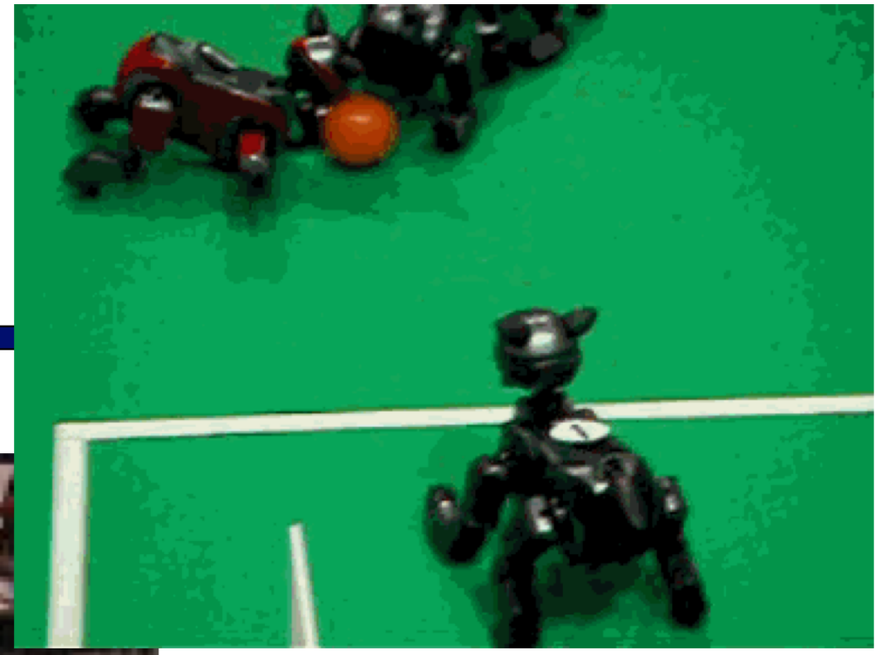


SONY
NOODLE
21

Robocup (Stockholm ’99)

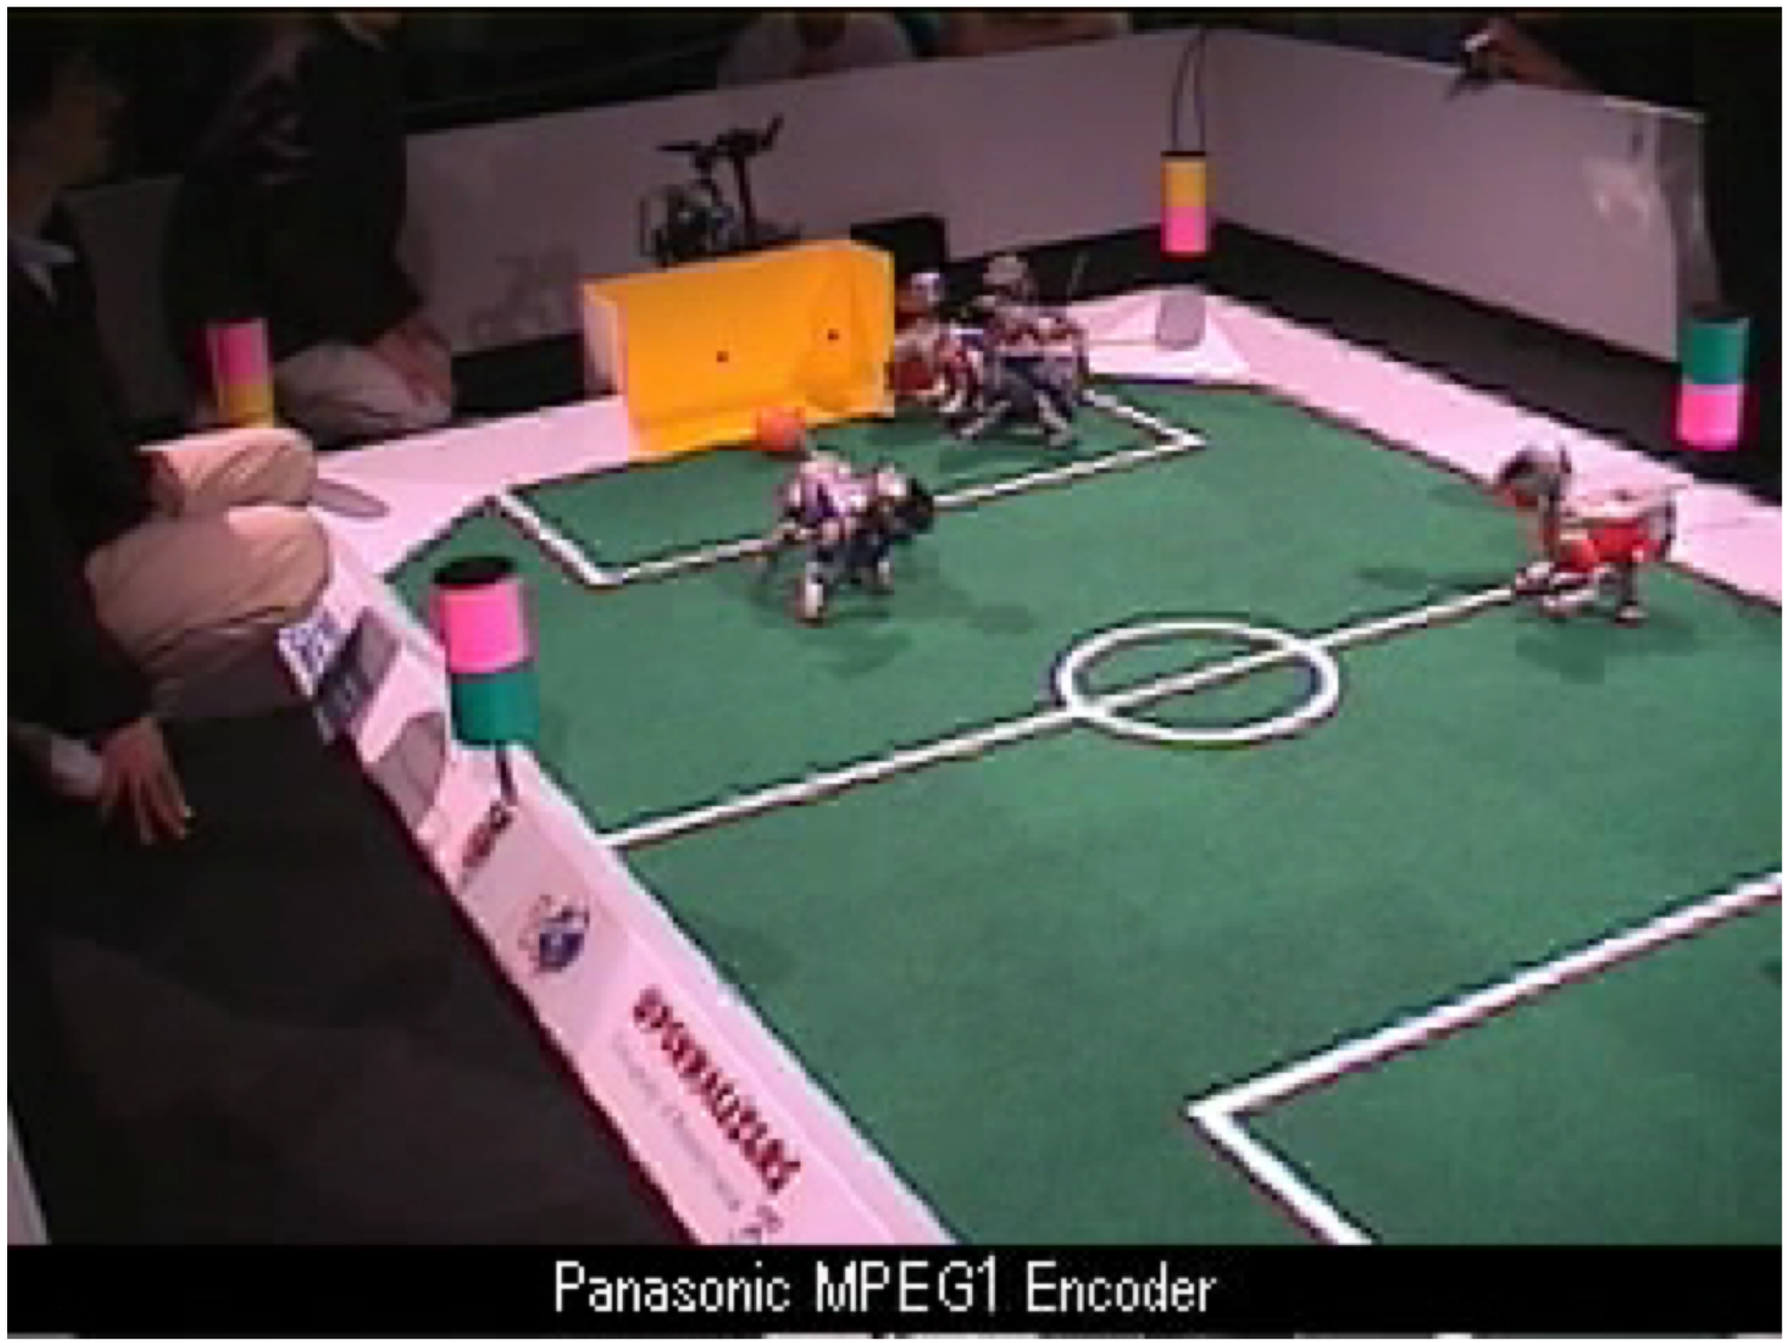

2005

Stanford Car
DARPA Grand Challenge

Google Self-Driving Car 2014

https://www.youtube.com/watch?v=TsaES--OTzM

2016

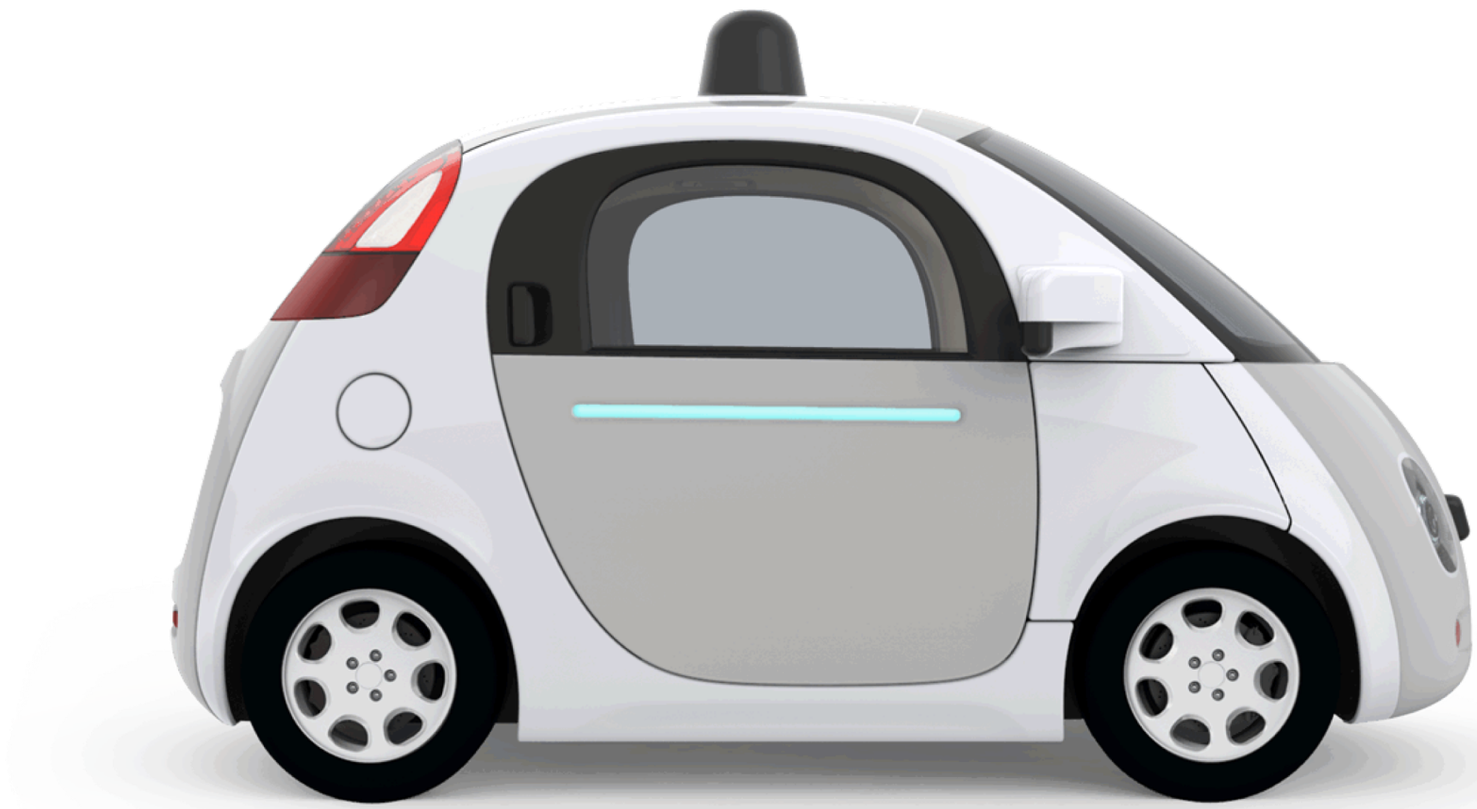

2009

amazon Prime

More Top Picks for You

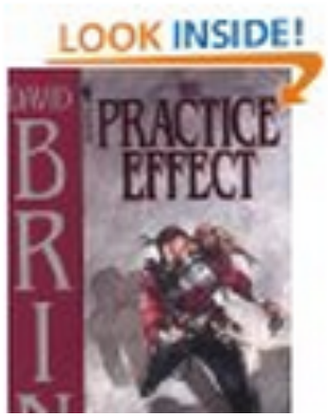

Recommendations,
Search result ordering
Ad placement

2011

http://www.youtube.com/watch?v=WFR3lOm_xhE

2016

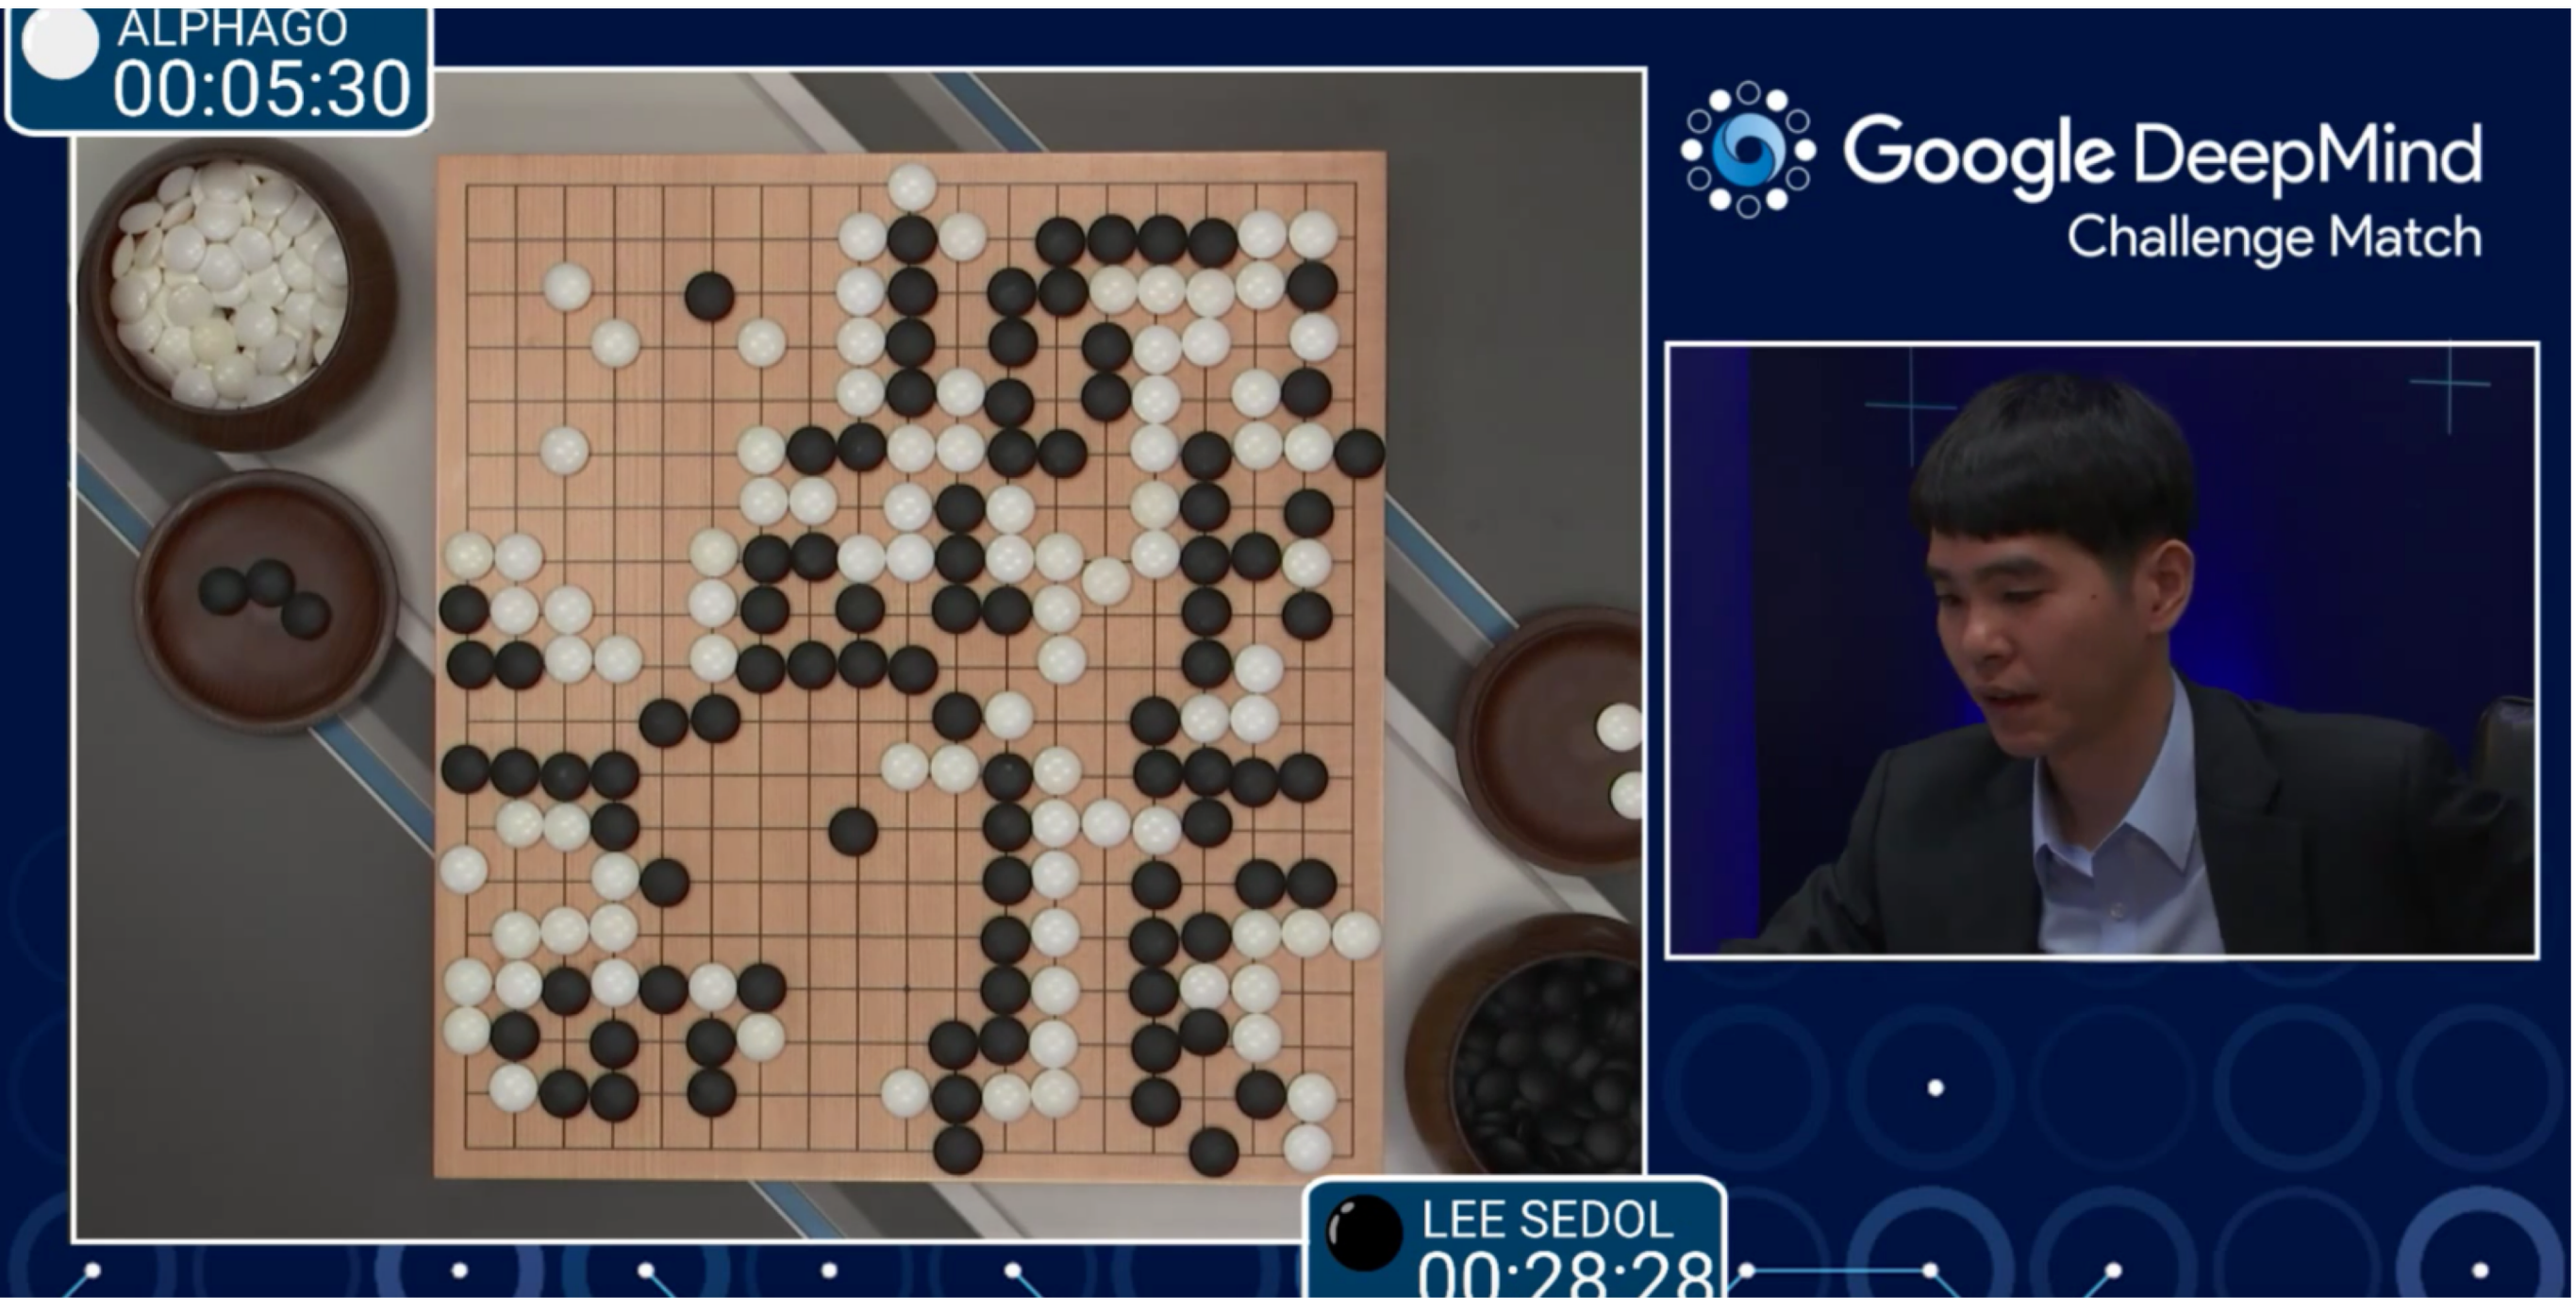

AlphaGo deep RL defeats Lee Sedol (4-1)

2014 = Momentous Times!

Fooled 33% of judges!

News > Technology

Turing Test breakthrough as super-computer becomes first to convince us it's human

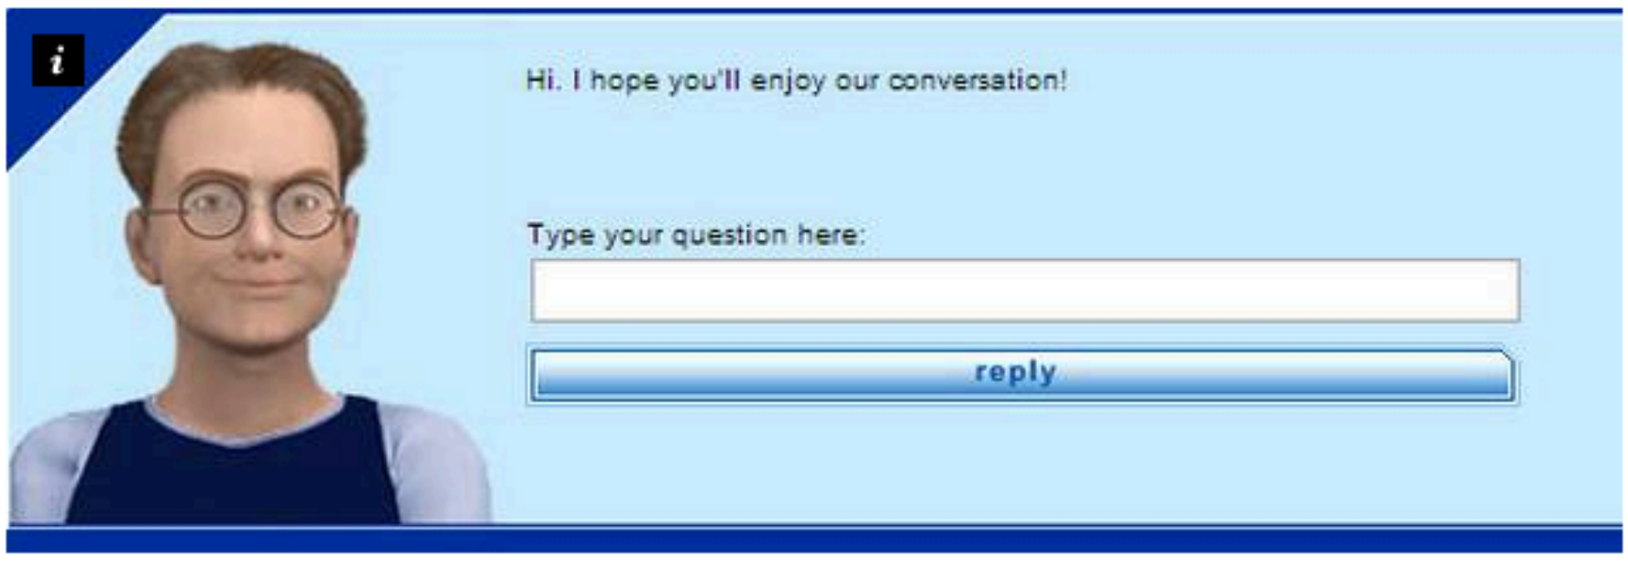

Eugene Goostman, a computer programme pretending to be a young Ukrainian boy, successfully duped enough humans to pass the iconic test

ANDREW GRIFFIN | Sunday 08 June 2014

A programme that convinced humans that it was a 13-year-old boy has become the first computer ever to pass the Turing Test. The test — which requires that computers are indistinguishable

Judges were not so smart

Conversation with Scott Aaronson:

Scott: Which is bigger, a shoebox or Mount Everest?

Eugene: I can’t make a choice right now. I should think it out later. And I forgot to ask you where you are from…

Scott: How many legs does a camel have?

Eugene: Something between 2 and 4. Maybe, three? :-))) By the way, I still don’t know your specialty – or, possibly, I’ve missed it?

Scott: How many legs does a millipede have?

Eugene: Just two, but Chernobyl mutants may have them up to five. I know you are supposed to trick me.

Scott: No, I need to know that you’re not a chatbot. Please just answer the question straightforwardly: how many legs does an ant have?

Eugene: Something between 2 and 4. Maybe, three? :-))) Oh, what a fruitful conversation;-)

Scott: Do you understand why I’m asking such basic questions? Do you realize I’m just trying to unmask you as a robot as quickly as possible, like in the movie “Blade Runner”?

Eugene: …wait

Scott: Do you think your ability to fool unsophisticated judges indicates a flaw with the Turing Test itself, or merely with the way people have interpreted the test?

Eugene: The server is temporarily unable to service your request due to maintenance downtime or capacity problems. Please try again later.

Scott: Do you think Alan Turing, brilliant though he was, had trouble imagining that the judges of his “imitation game” wouldn’t think to ask commonsense questions like the ones above—or that, if they did, they’d actually accept evasion or irrelevant banter as answers?

Eugene: No, not really. I don’t think alan turing brilliant although this guy was had trouble imagining that the judges of his imitation game would not consider to Oooh. Anything else?

BakeBot:
Motion Planning for Cooking

Mario Bollini and Daniela Rus
CSAIL, MIT

Mathematical Calculation

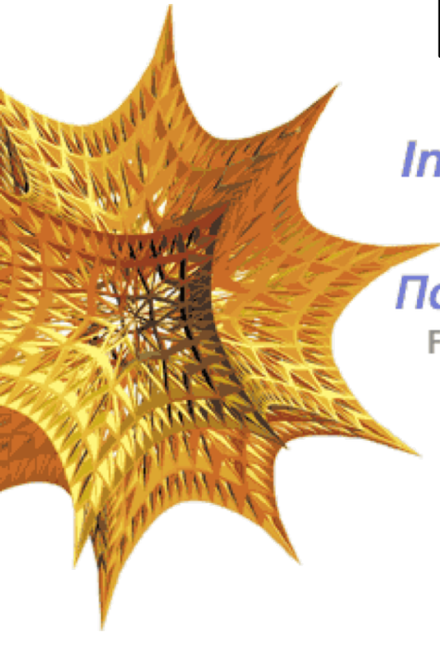

Introducing
MATHEMATICA 5
Παρουσιάζουμε το
Featuring a new generation of
advanced algorithms with unparalleled
speed, scope, and scalability

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

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

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

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

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

Different Research Areas in AI

- Natural Language Processing
- Computer Vision
- Robotics
- Logic
- Decision Making
- Game Playing
- Machine Learning

Natural Language Processing

- Speech Technologies (e.g., Siri):
  - Automatic Speech Recognition (ASR)
  - Text-to-speech synthesis
  - Dialog Systems

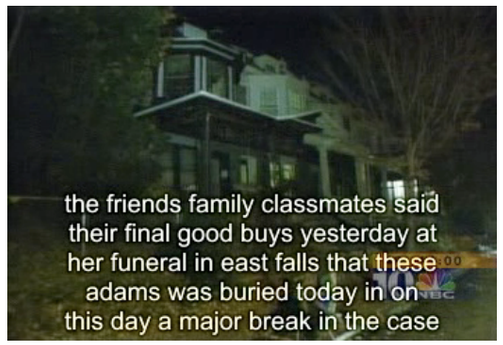

- Language Technologies:
  - Question answering
  - Machine translation

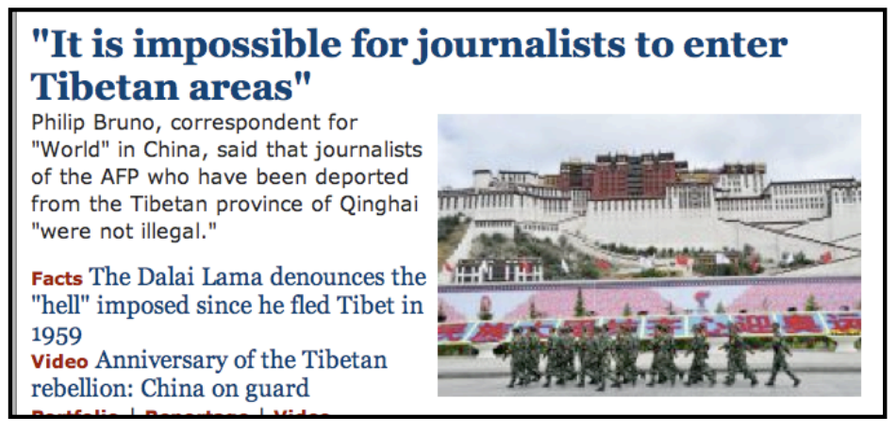

- Text classification; spam filtering; etc

Mobile devices can now answer (some of our) questions and execute commands...

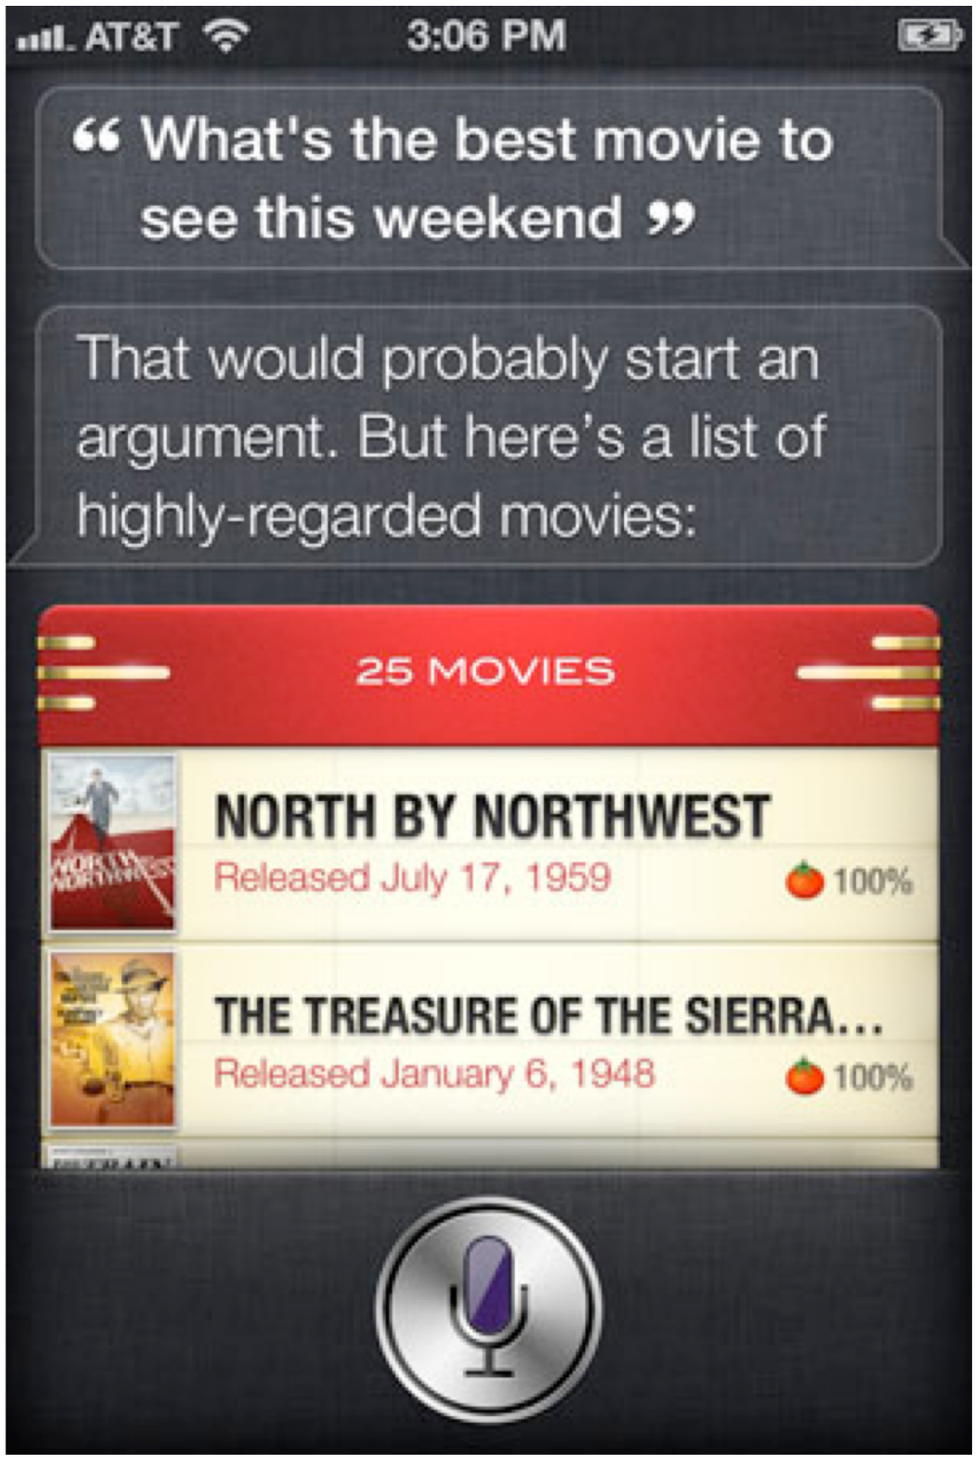

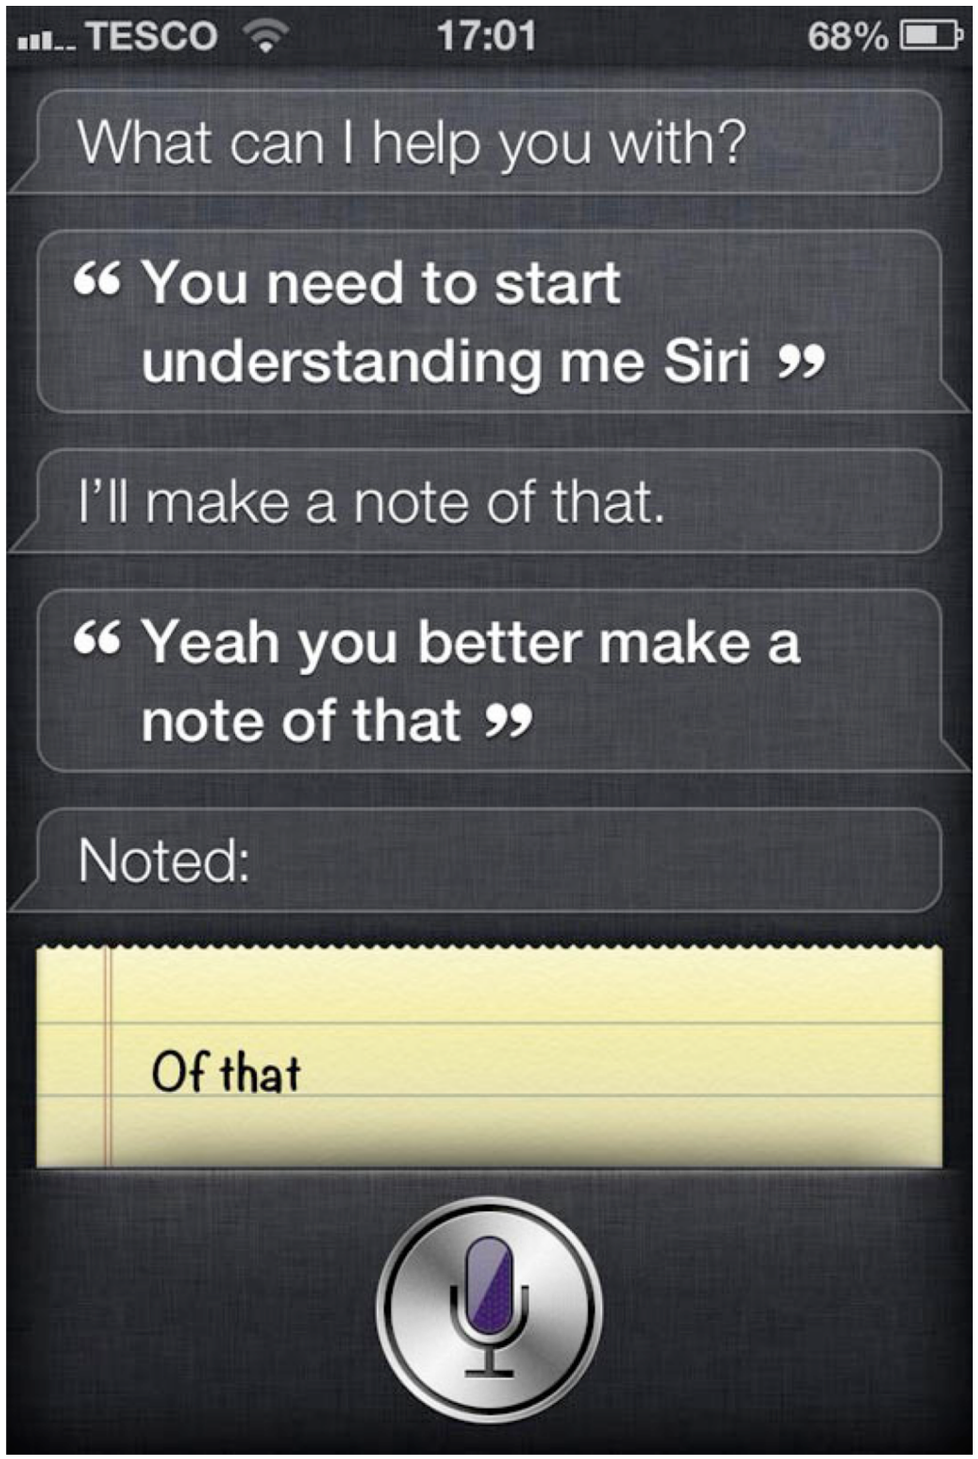

Vision

- Object Recognition
- Scene Classification
- Image Segmentation
- Human Activity Recognition

https://pjreddie.com/darknet/yolo/

Object Recognition

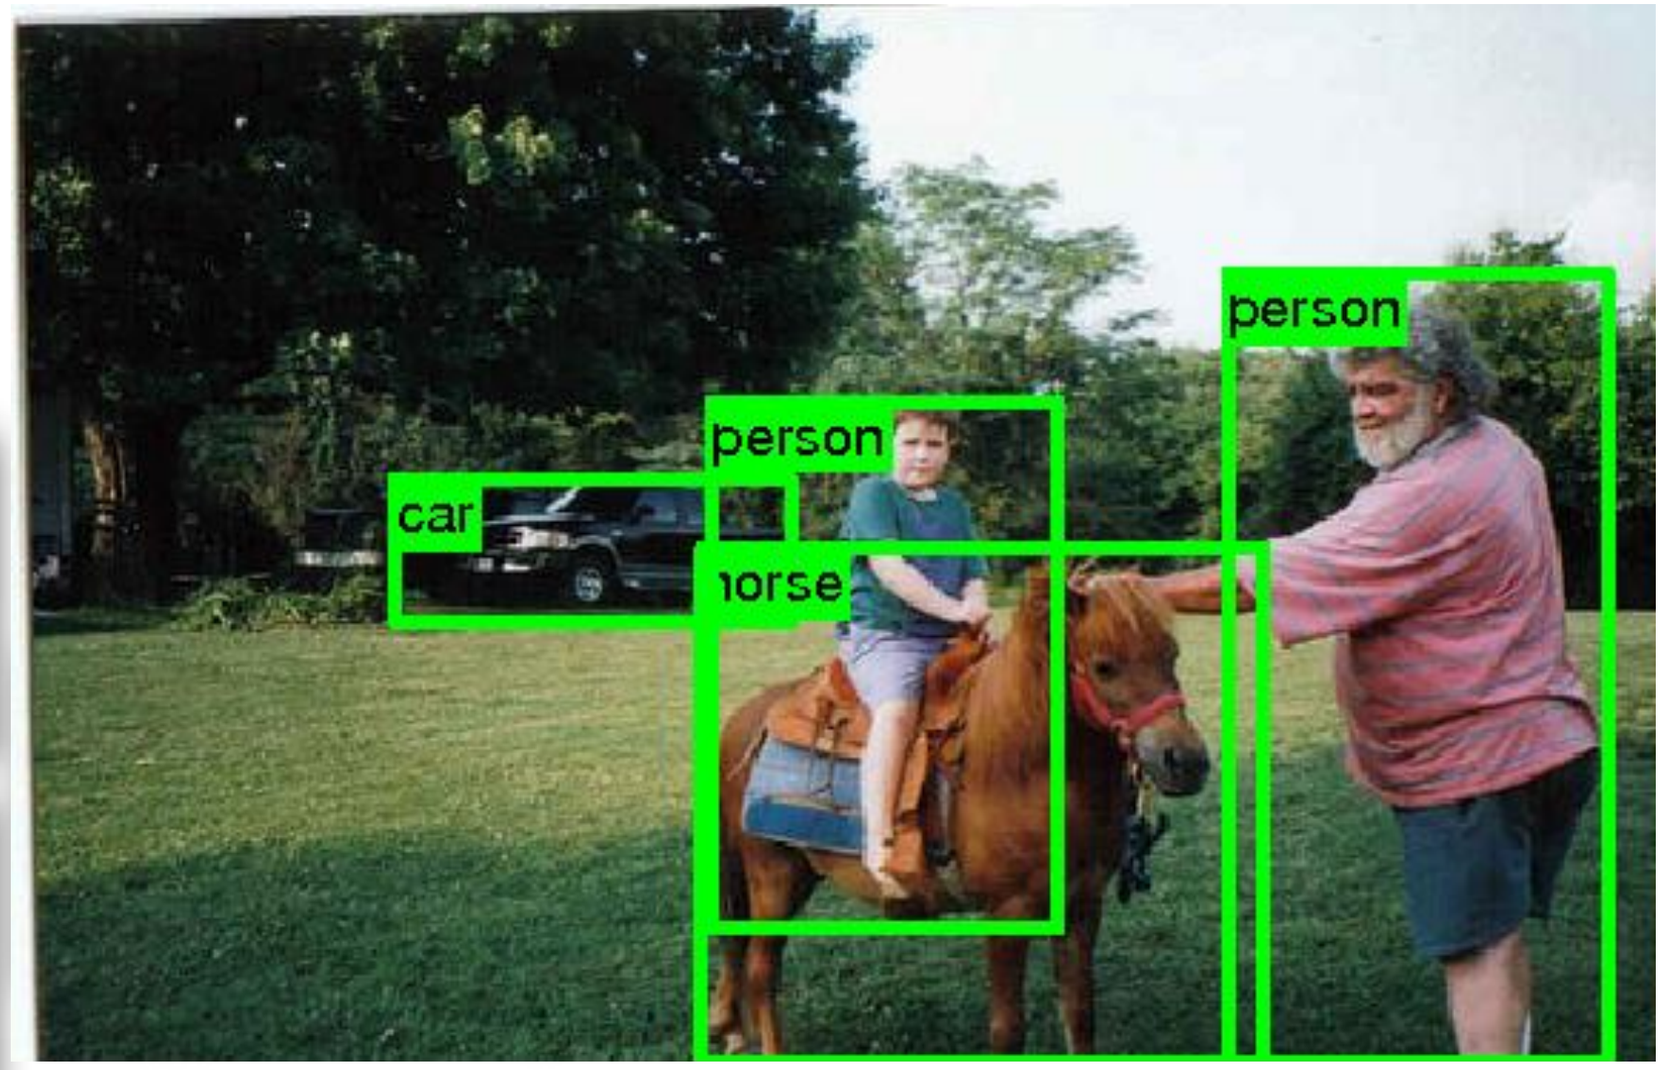

Scene Segmentation

(b) (c)

Google Goggles

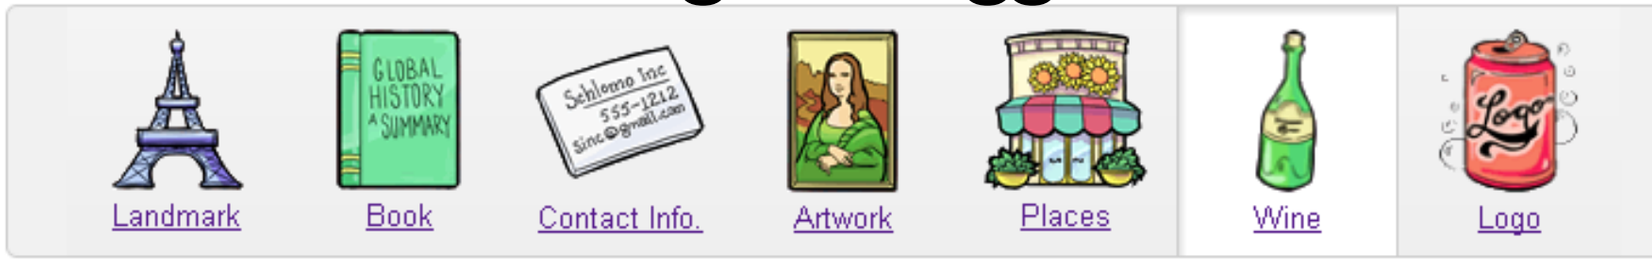

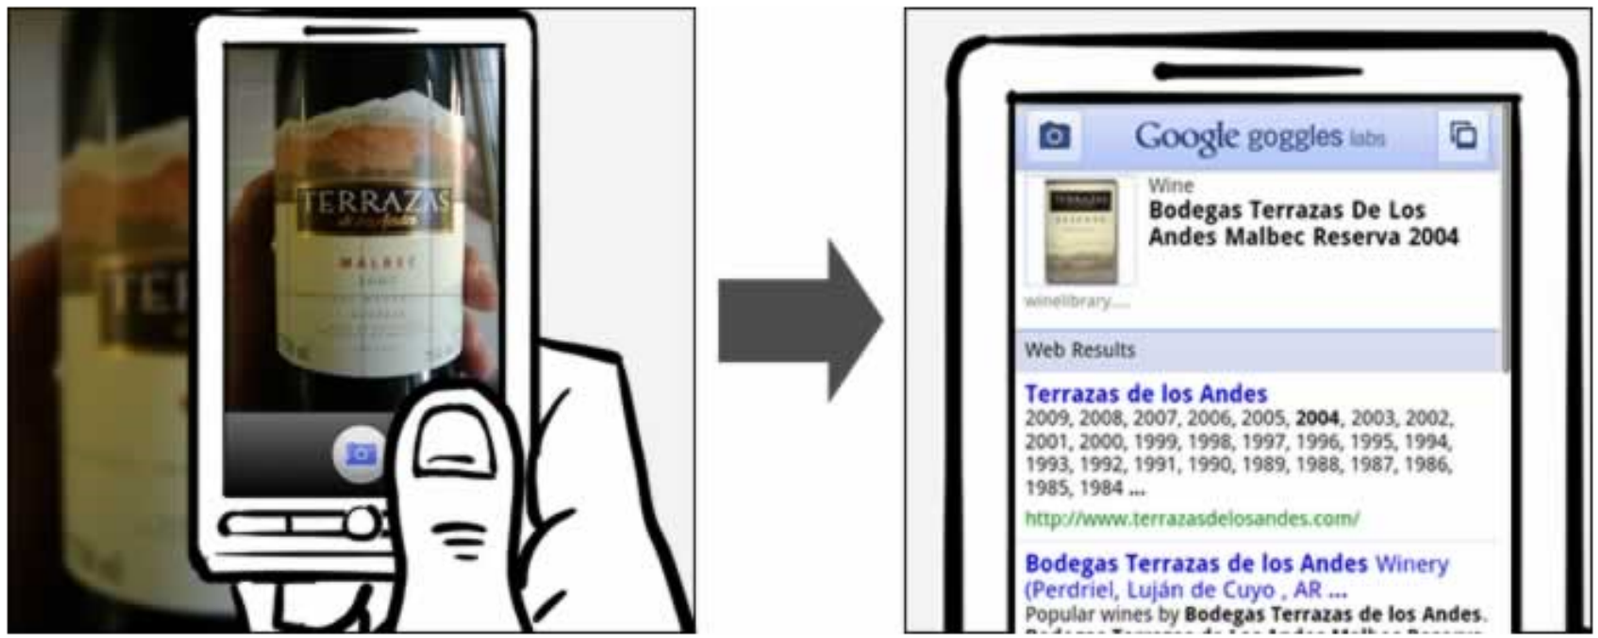

Smile Detection

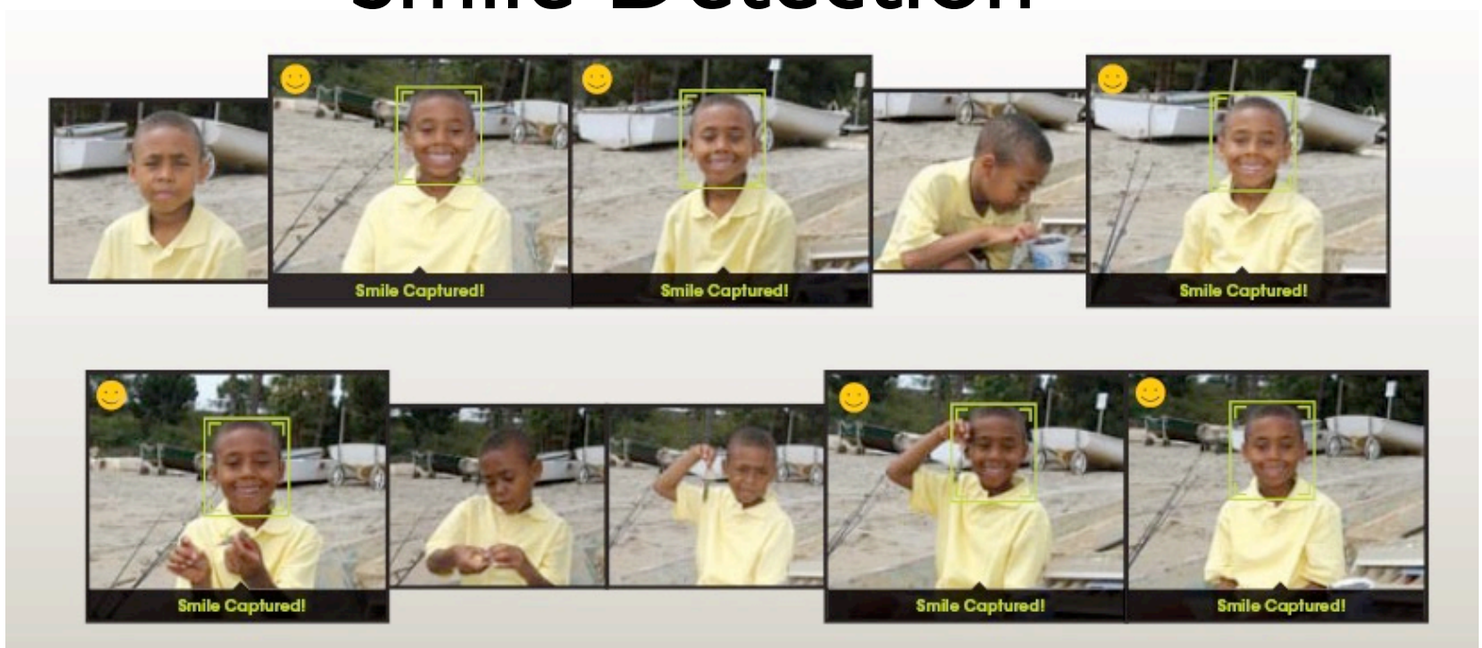

Leaf Snap

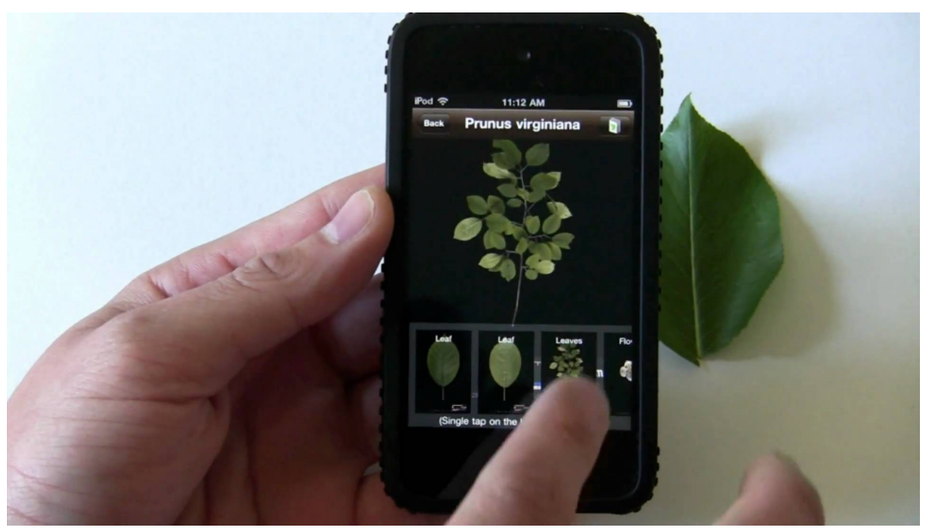

Image captioning: What begins to work

The flower was so vivid and attractive.

Blue flowers are running rampant in my garden.

We sometimes do well: 1 out of 4 times, machine captions were preferred over the original Flickr captions:

Spring in a white dress.

Blue flowers have no scent. Small white flowers have no idea what they are.

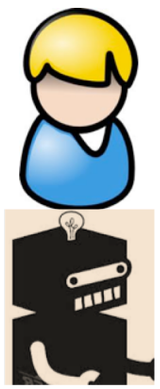

Scenes around the lake on my bike ride.

This horse walking along the road as we drove by.

But many challenges remain
(better examples of when things go awry)

The couch is definitely bigger than it looks in this photo.

Yellow ball suspended in water.

My cat laying in my duffel bag.

A high chair in the trees.

Robotics

- Robotics
  - Part mech. eng.
  - Part AI
  - Reality much harder than simulations!

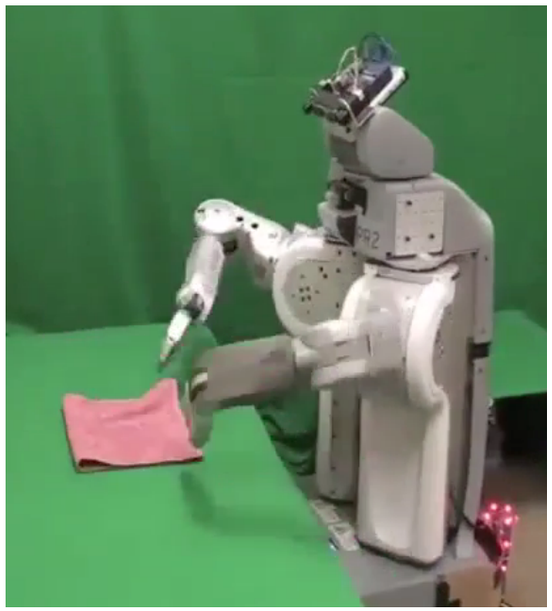

- Technologies
  - Vehicles
  - Rescue
  - Soccer!
  - Lots of automation…

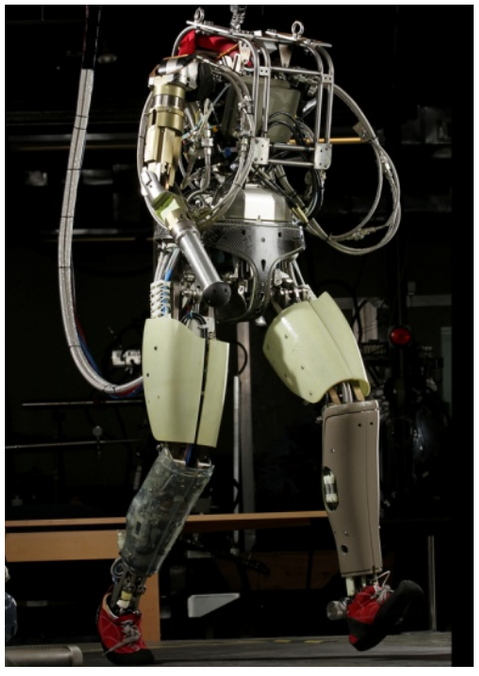

- In this class:
  - We ignore mechanical aspects
  - Methods for planning
  - Methods for control

Images from UC Berkeley, Boston Dynamics, RoboCup, Google

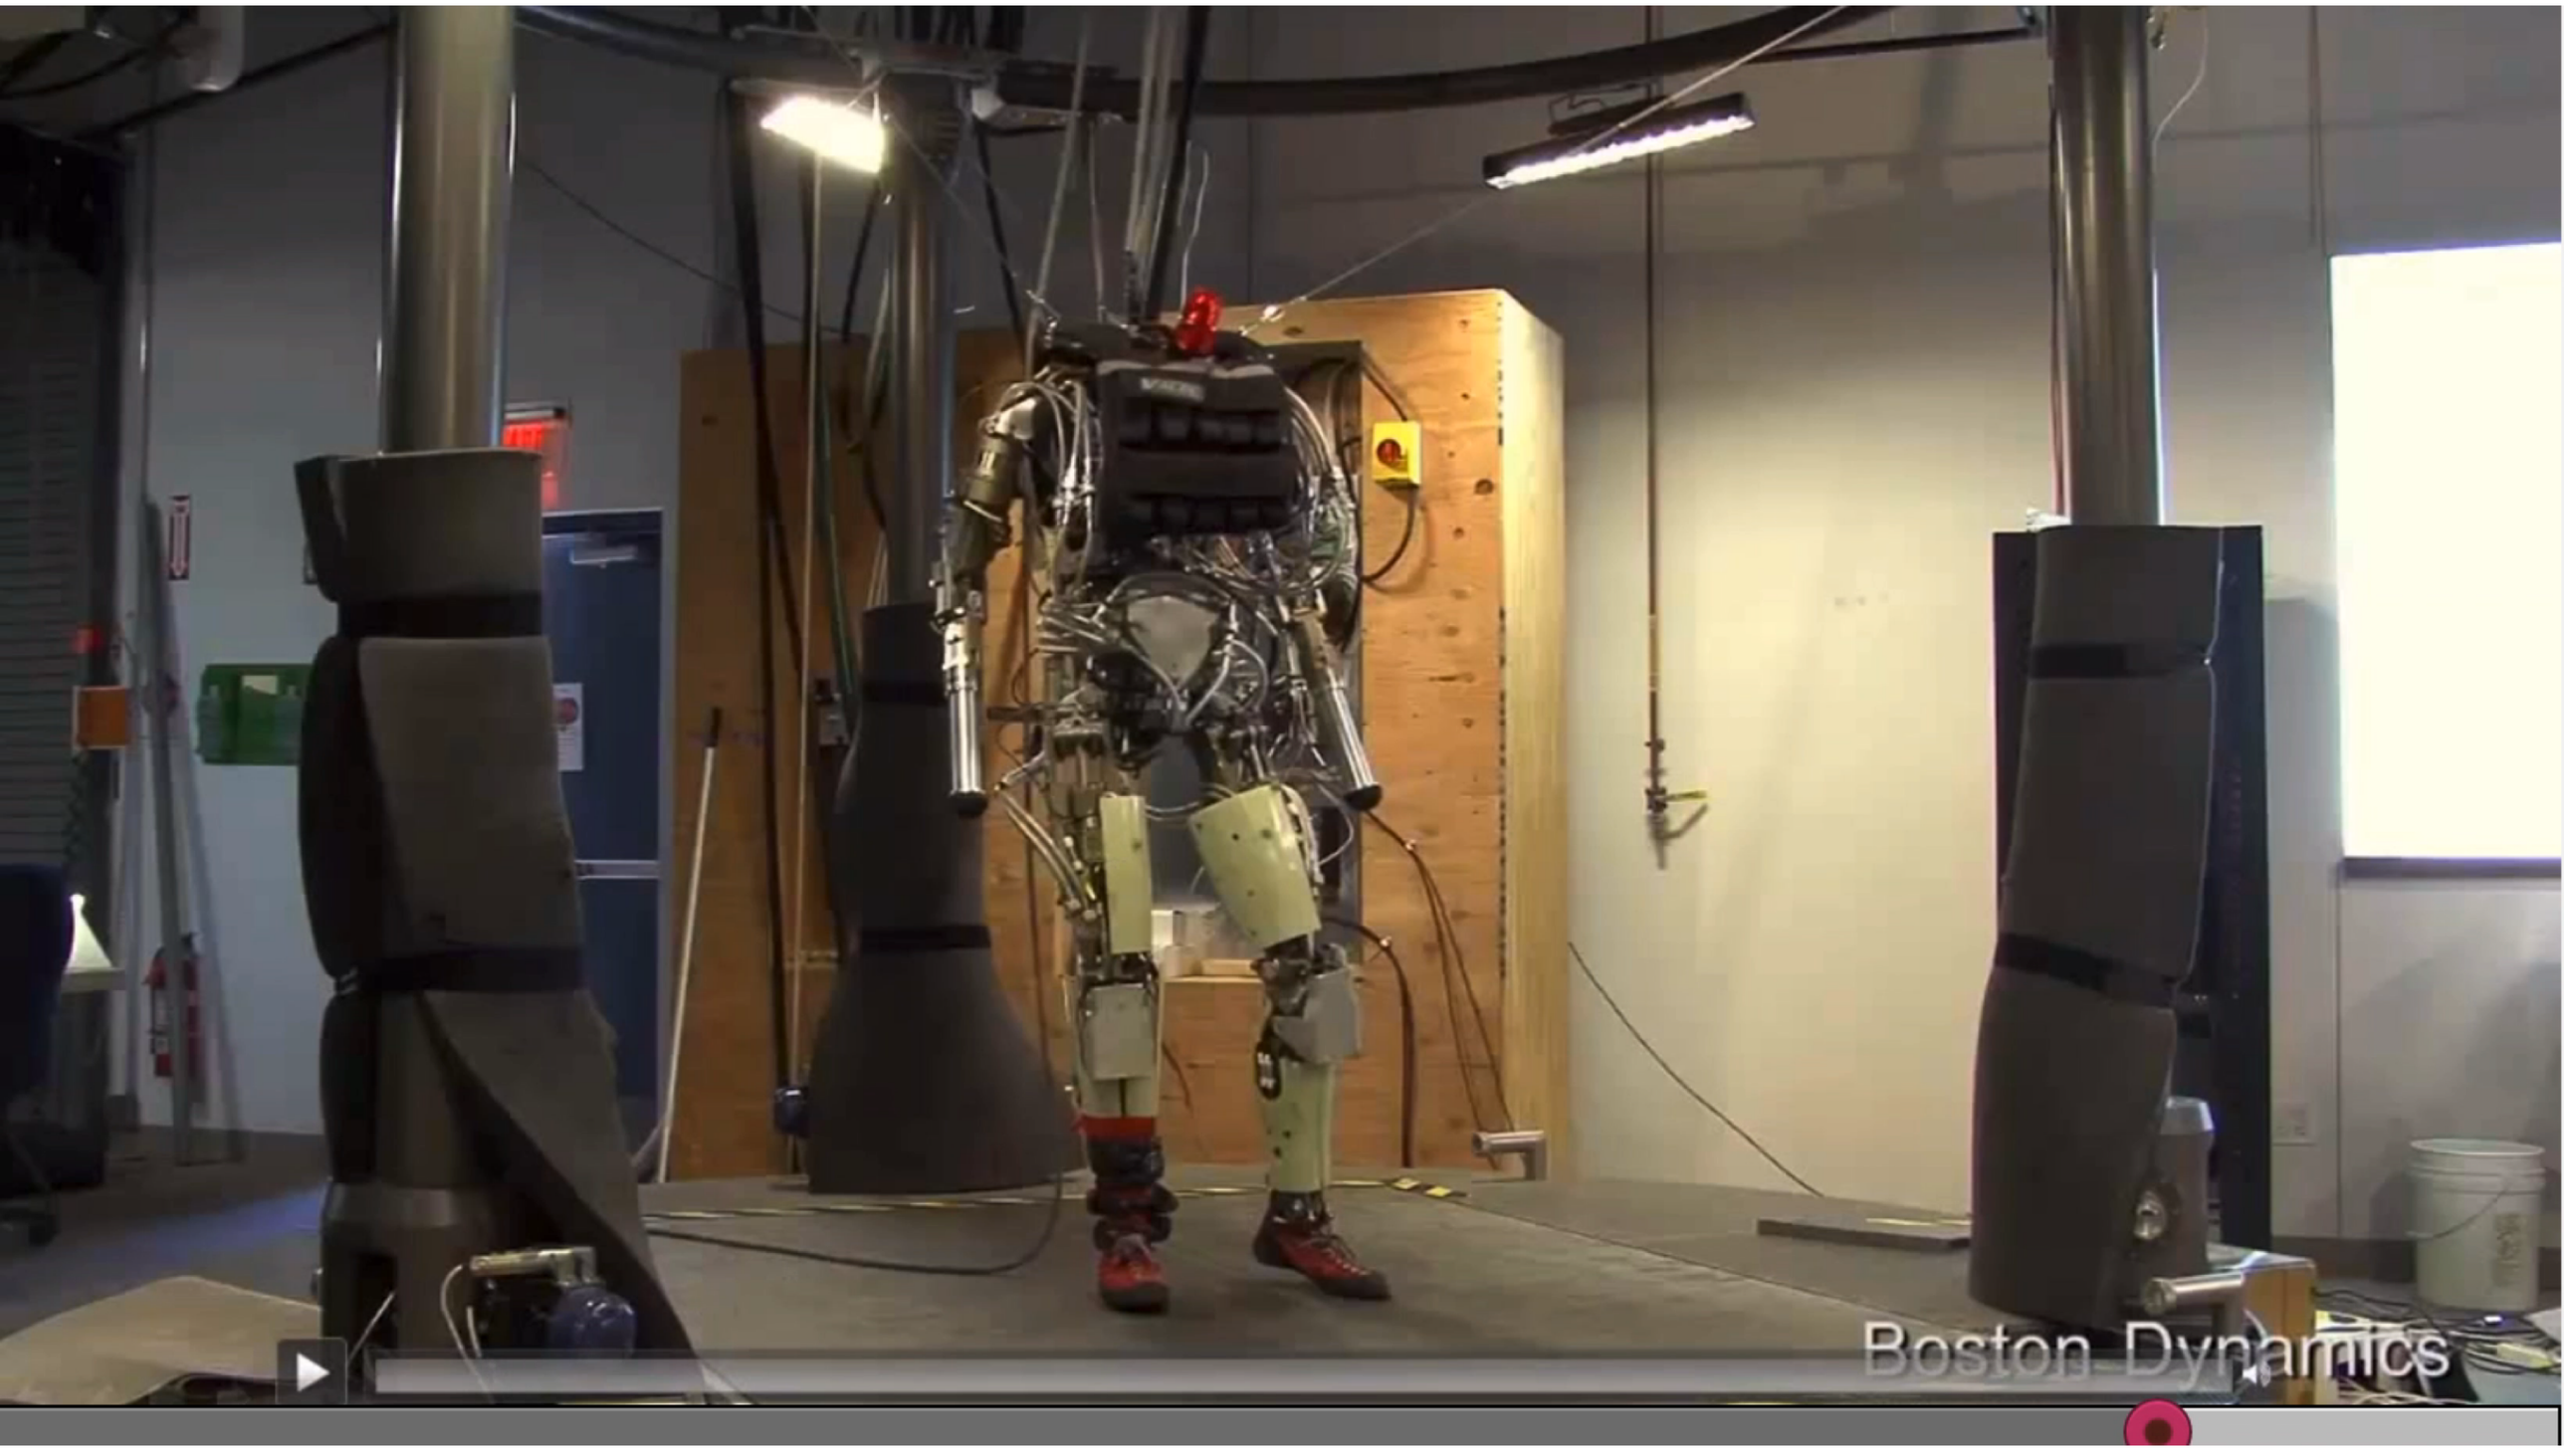
Boston Dynamics

Logic

- Logical systems
  - Theorem provers
  - NASA fault diagnosis
  - Question answering

- Methods:
  - Deduction systems
  - Constraint satisfaction
  - Satisfiability solvers (huge advances!)

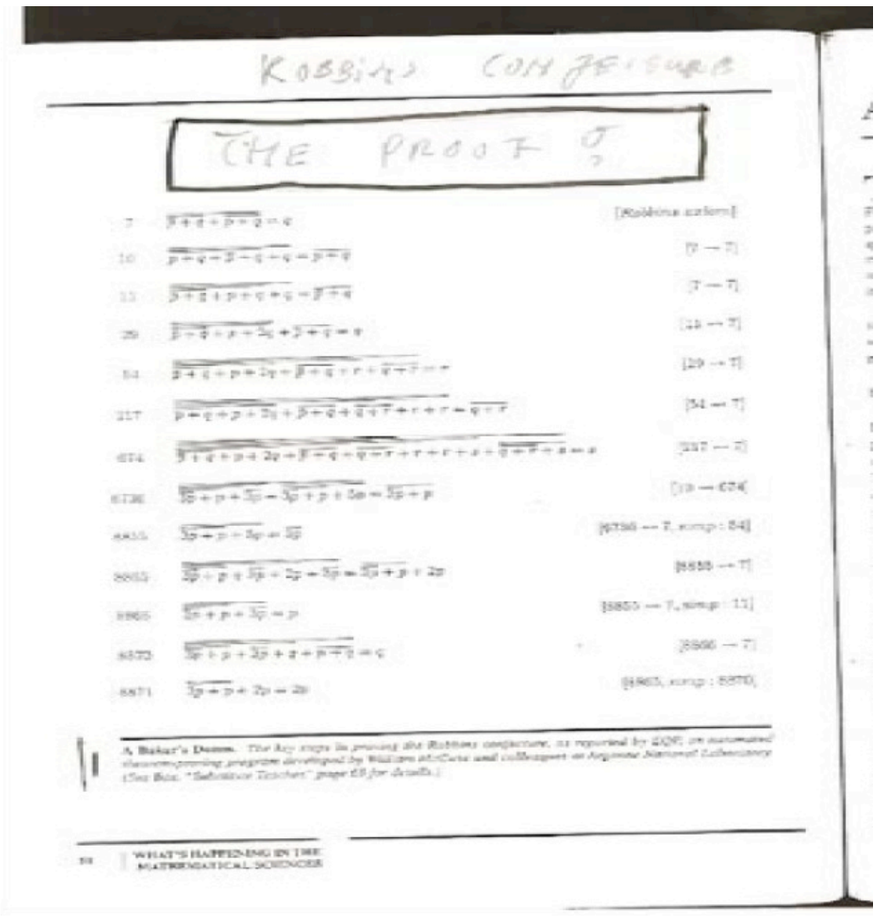

Image from Bart Selman

AI is everywhere now!

- Scheduling, airline routing
- Route planning
- Medical diagnosis
- Web search
- Spam classification
- Automated help desks
- Smarter devices, like cameras
- Fraud detection
- Product recommendation
- … Lots more!

Machine Learning

- Most current AI systems

"Il est impossible aux journalistes de rentrer dans les régions tibétaines"

Bruno Philip, correspondant du "Monde" en Chine, estime que les journalistes de l'AFP qui ont été expulsés de la province tibétaine du Qinghai "n'étaient pas dans l'illégalité".

Les faits Le dalaï-lama dénonce l'"enfer" imposé au Tibet depuis sa fuite, en 1959

Vidéo Anniversaire de la rébellion

"It is impossible for journalists to enter Tibetan areas"

Philip Bruno, correspondent for "World" in China, said that journalists of the AFP who have been deported from the Tibetan province of Qinghai "were not illegal."

Facts The Dalai Lama denounces the "hell" imposed since he fled Tibet in 1959

Video Anniversary of the Tibetan rebellion: China on guard

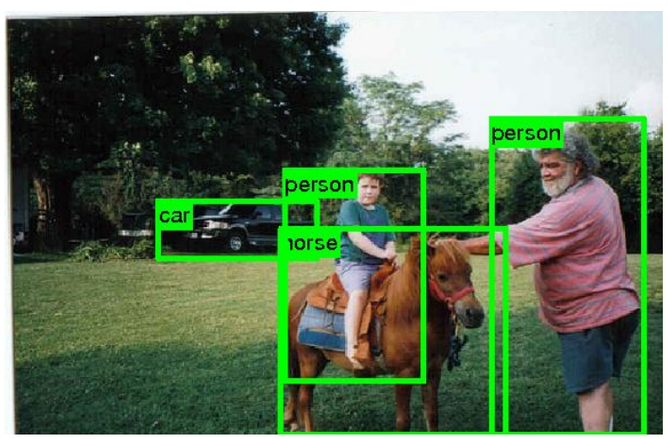

Game Playing

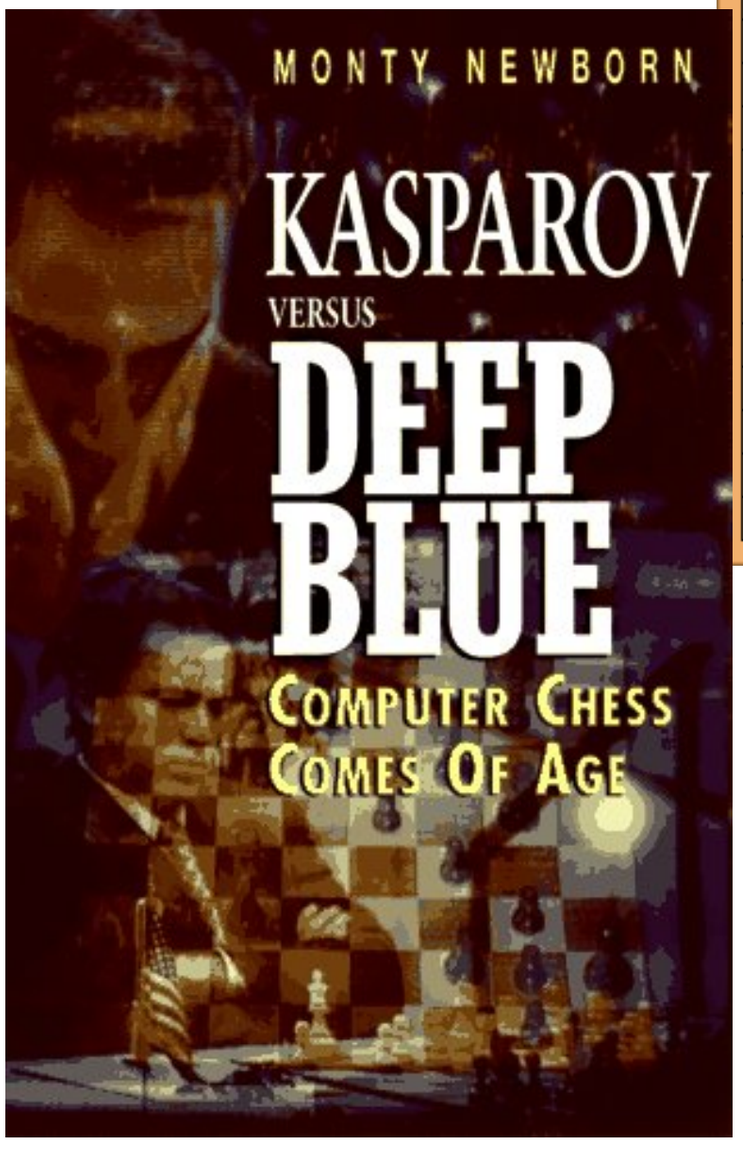

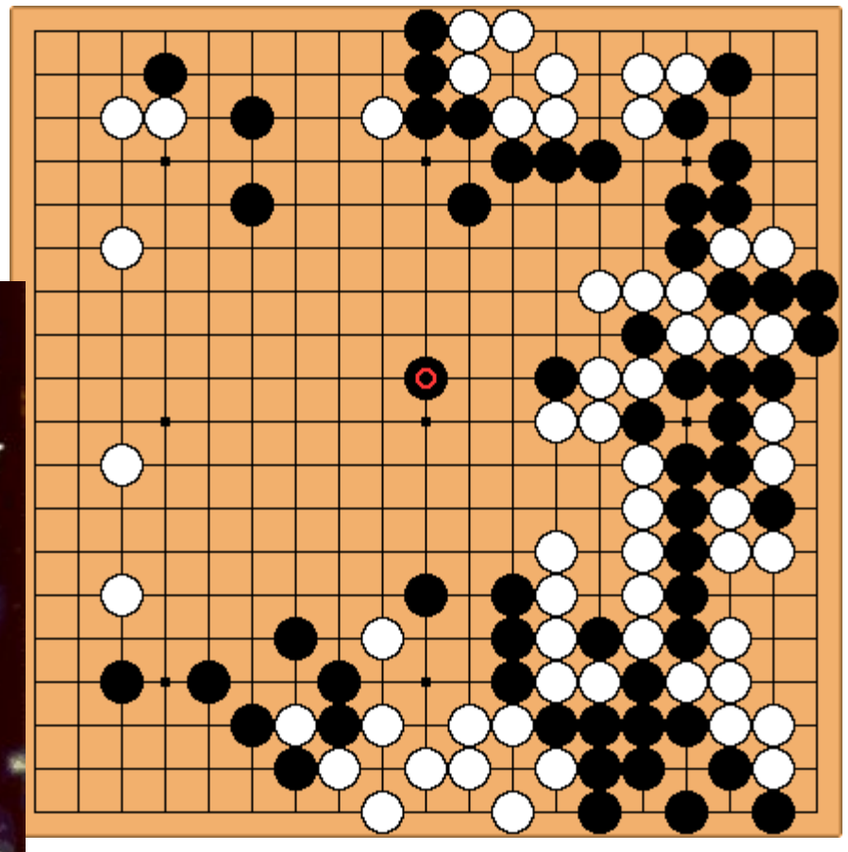

What is this course about?

Designing Rational Agents

An agent:

- Perceives and acts
- Selects actions that maximize its utility function
- Has a goal

Environment:

- Input and output to the agent

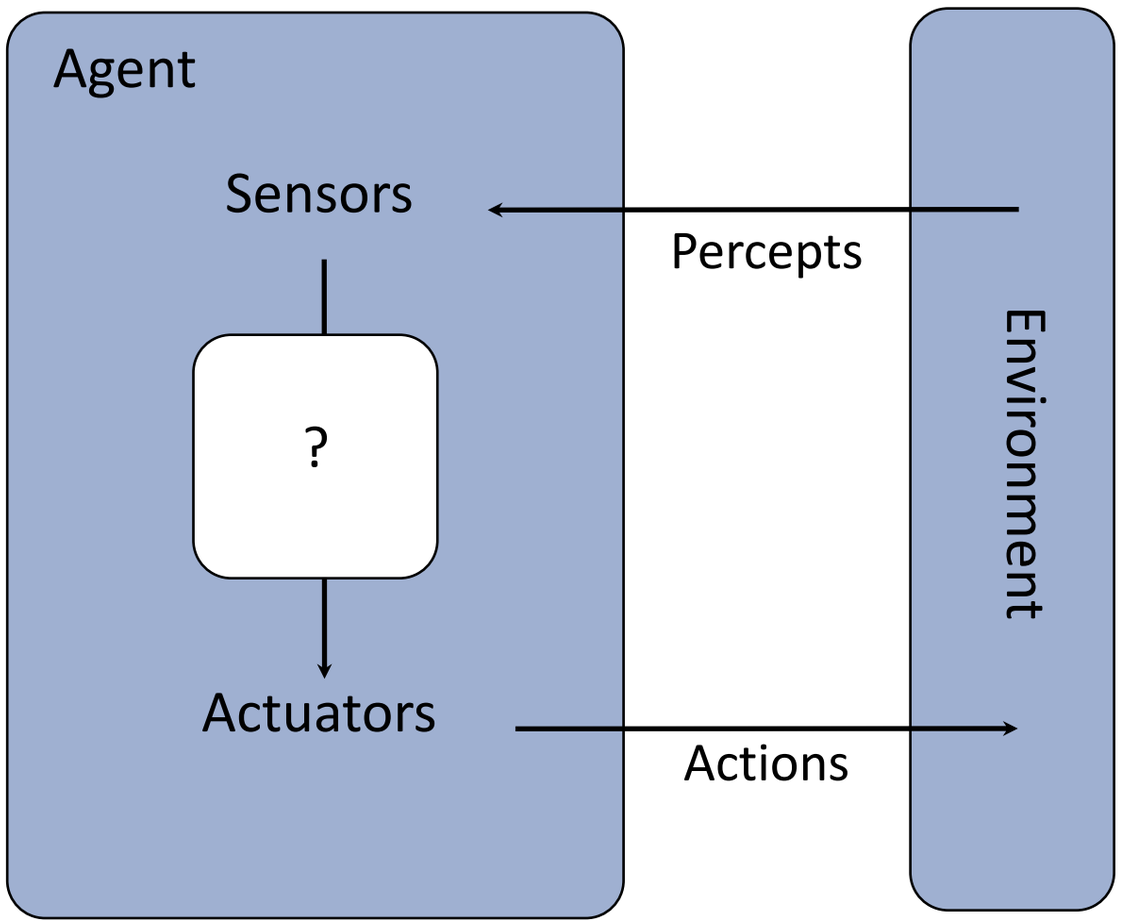

Designing Rational Agents

Characteristics of the percepts, environment, and action space dictate techniques for selecting rational actions.

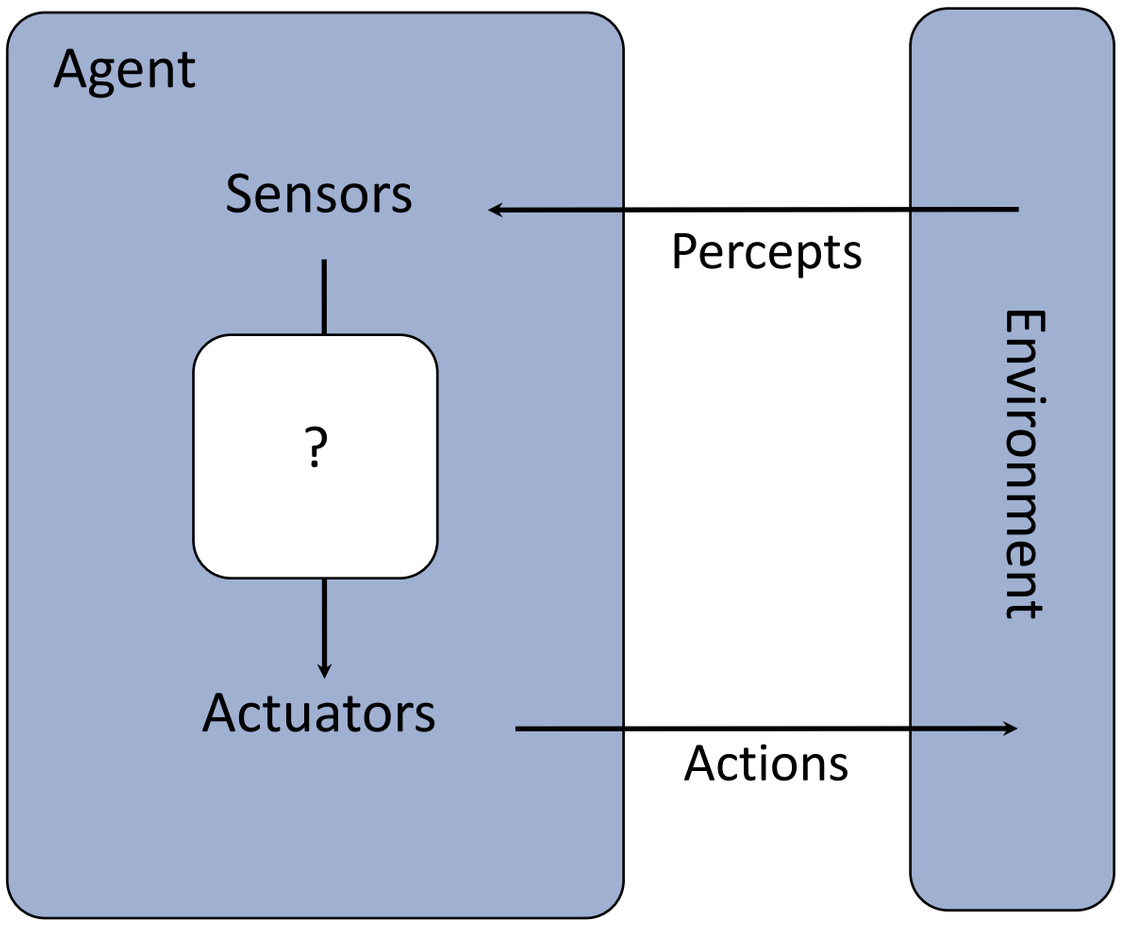

This course is about:

- 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

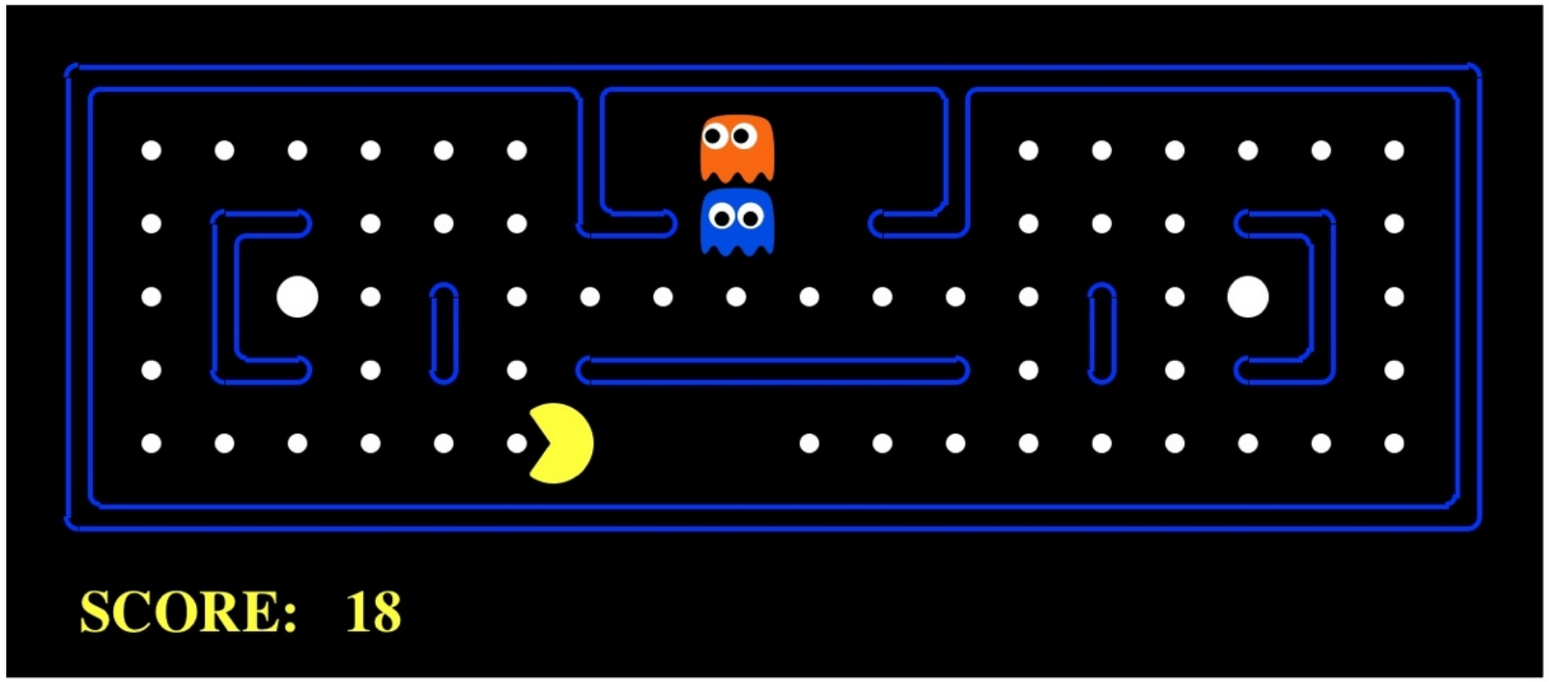

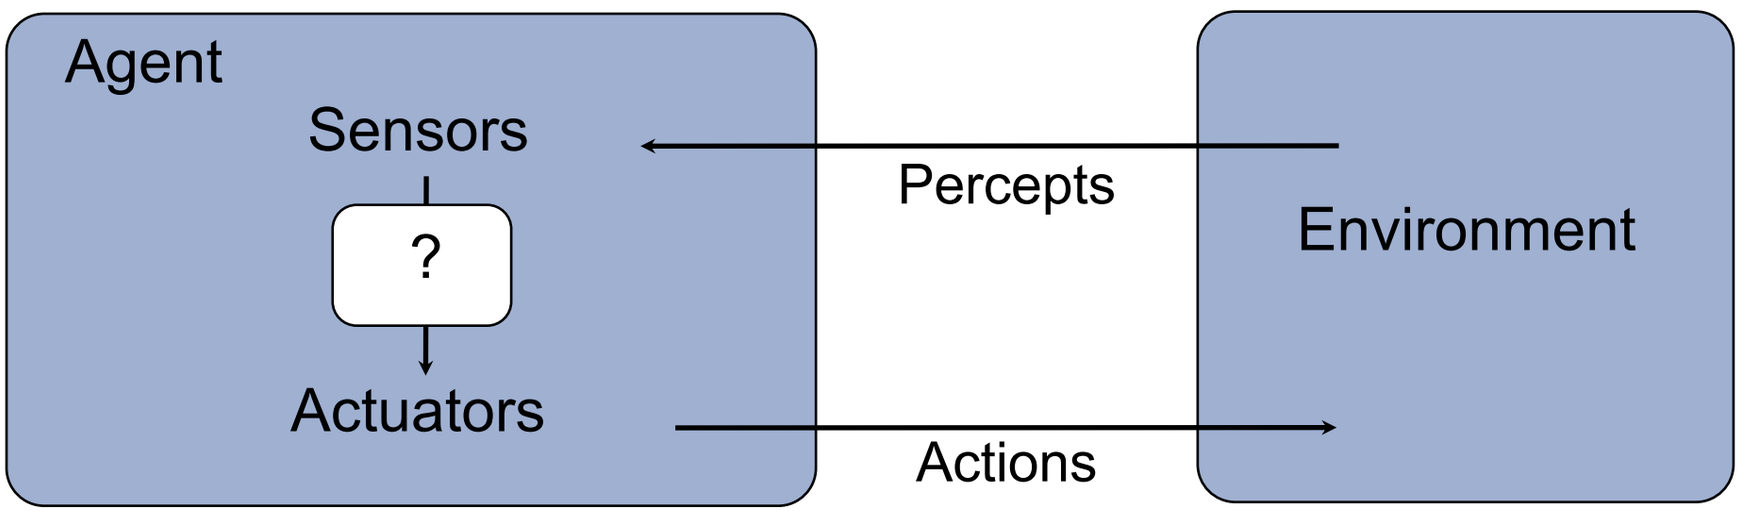

Actions? Percepts?

Actions? Percepts?

Recommender System

Types of Environments

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

Fully observable vs. Partially observable

Can the agent observe the complete state of the environment?

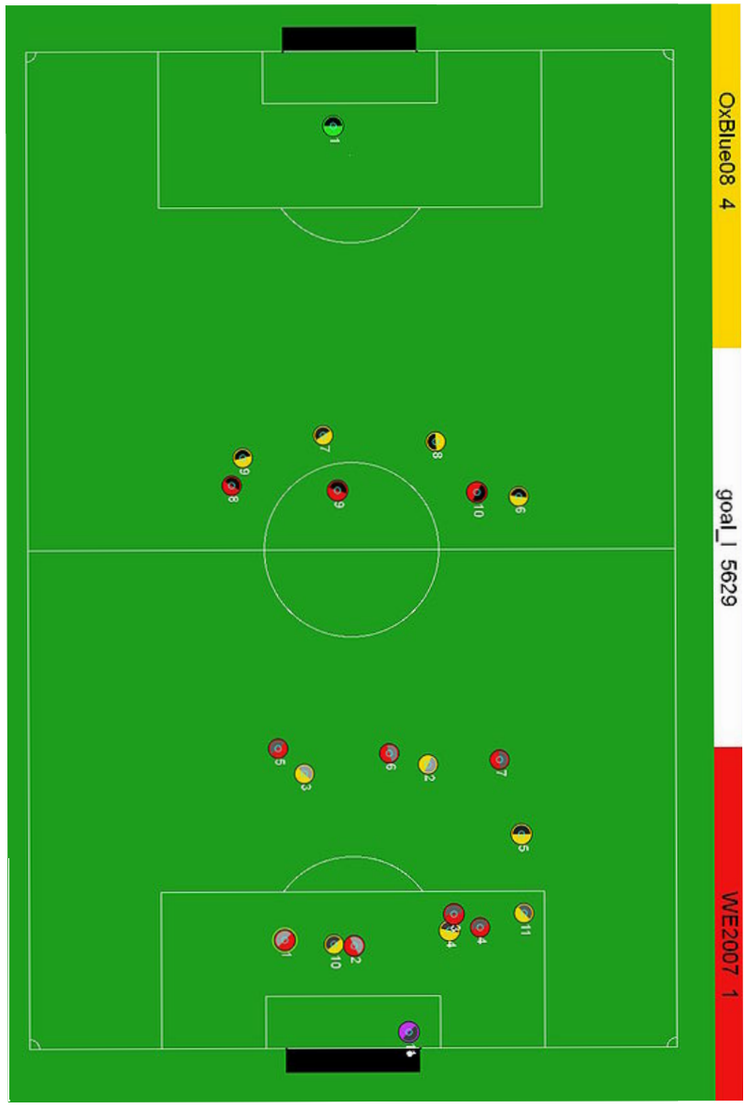

VS.

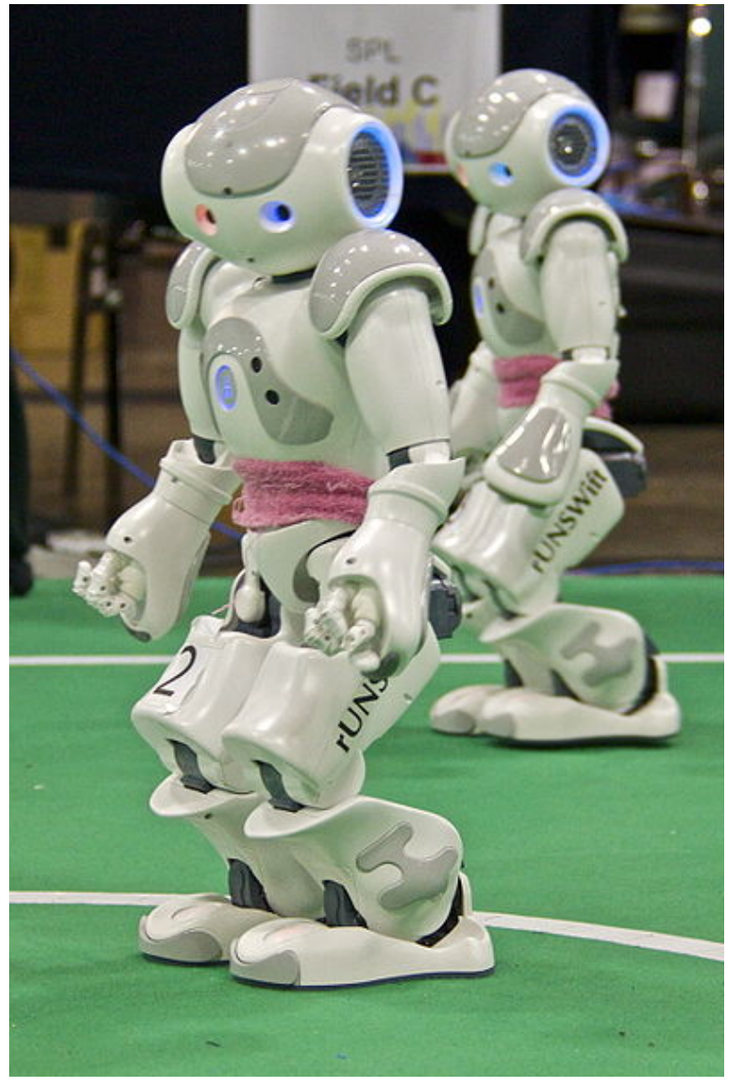

Single agent vs. Multiagent

Is the agent the only thing acting in the world?

vs.

Deterministic vs. Stochastic

Is there uncertainty in how the world works?

VS.

Static vs. Sequential

Does the agent take more than one action?

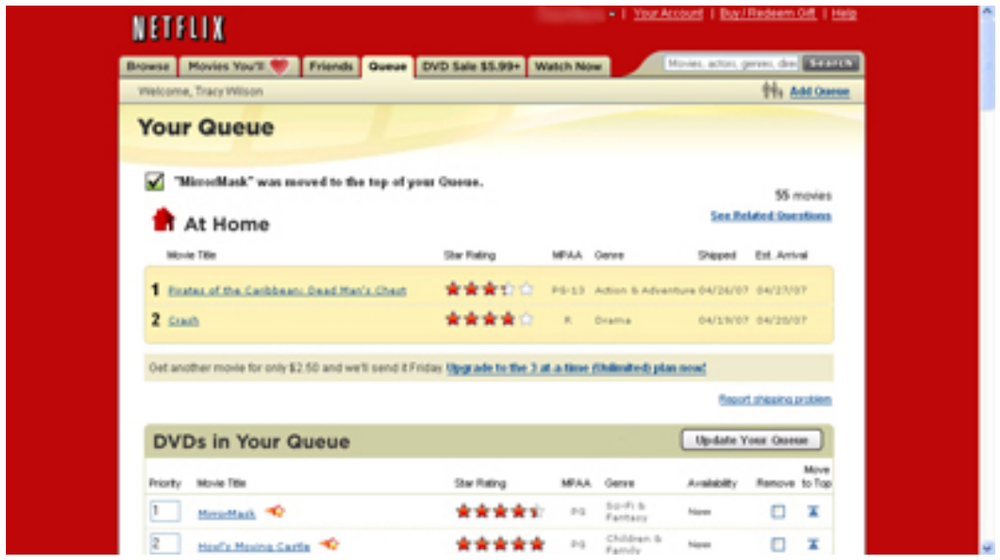

vs.

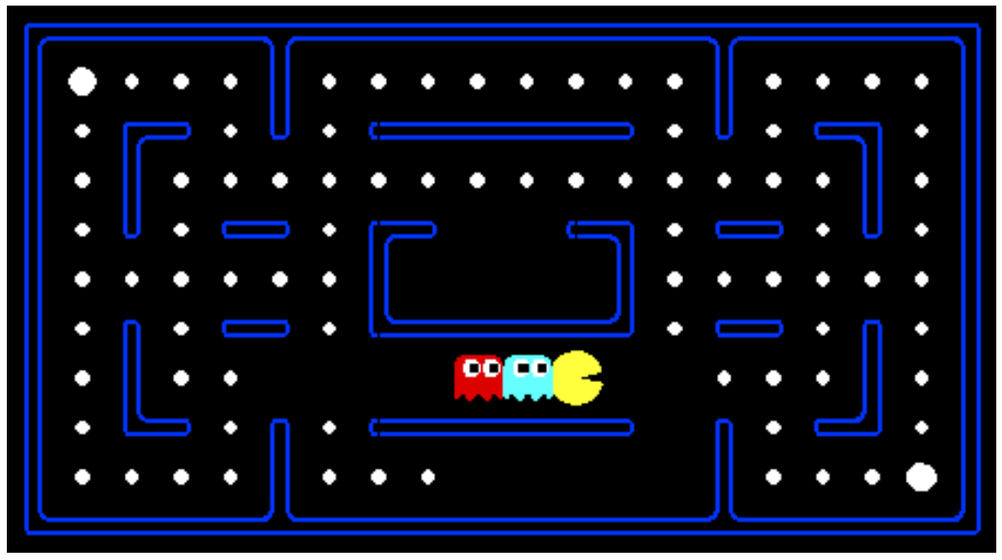

Discrete vs. Continuous

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

VS.

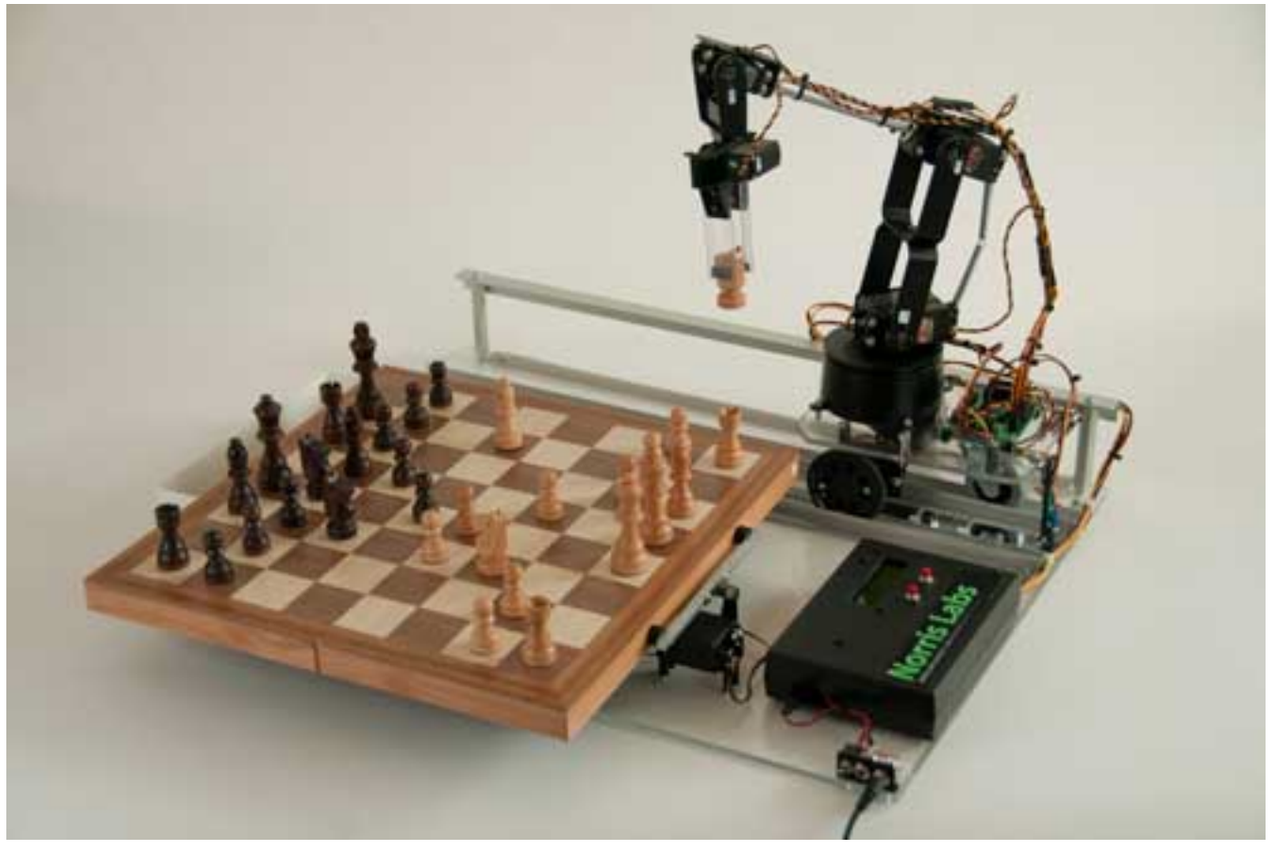

Types of Agent

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.

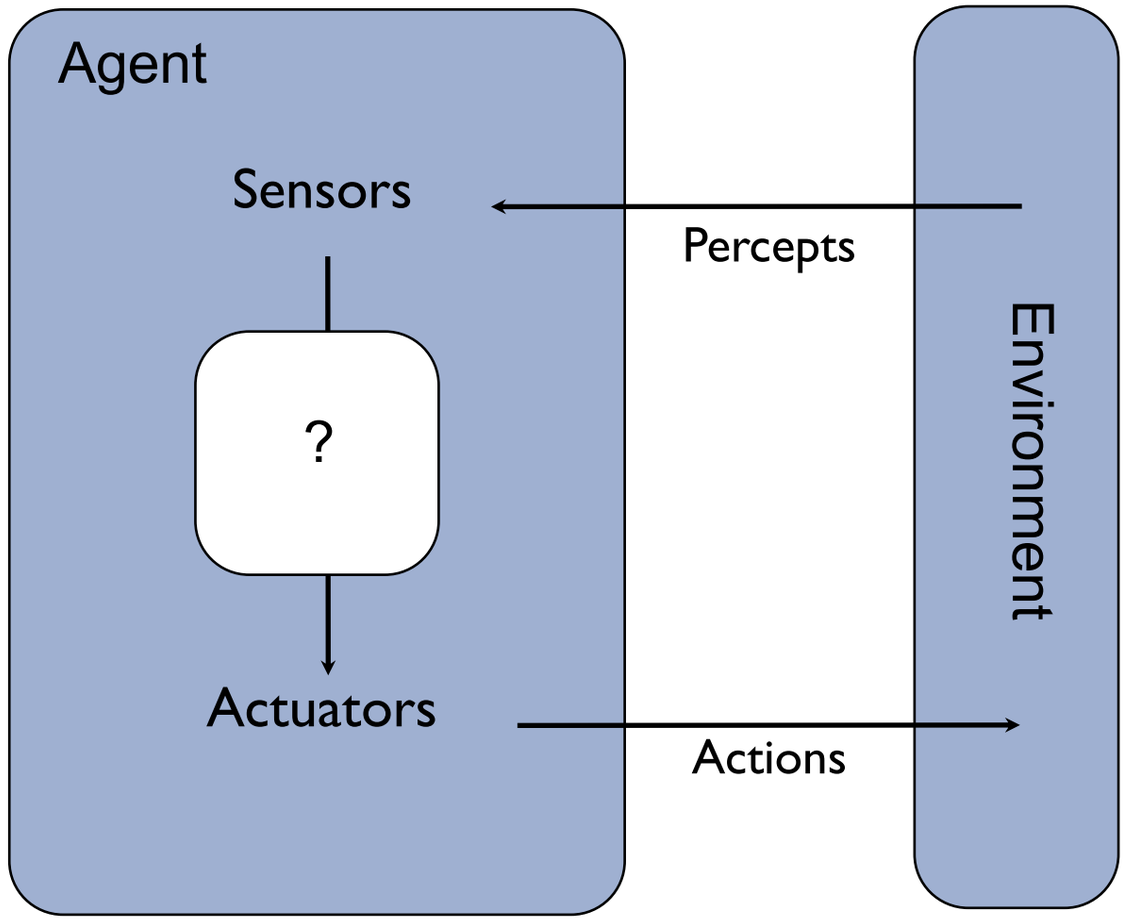

Goal Based Agents

- Plan ahead
- Ask “what if”

- Decisions based on (hypothesized) consequences of actions
- Uses a model of how the world evolves in response to actions

- Act on how the world WOULD BE

Utility Based Agents

Like goal-based, but

Trade off multiple goals

Reason about probabilities of outcomes

Act on how the world will LIKELY be

Perceive

Act

AI

techniques

applications

Knowledge Representation

Inference

Machine Learning

Vision

NLP

Robotics

This course vs. others

CSE 515 – Stat methods
CSE 517 – NLP
CSE 546,7 – ML
CSE 571 – Robotics
CSE 574
CSE 576,7 – Vision

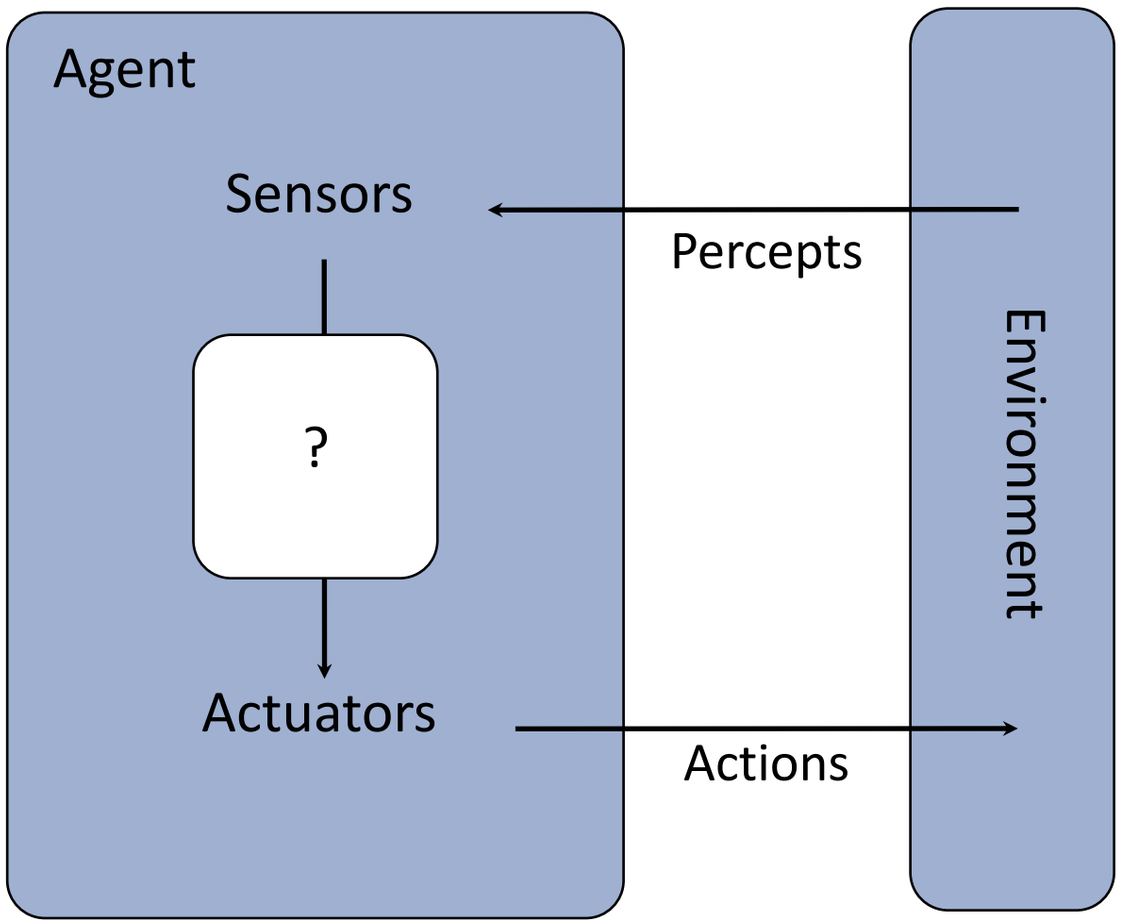

Topics in This Course

- Introduction
- Search
- Game Playing (minimax, alpha beta, expectimax)
- Markov Decision Processes
- Reinforcement Learning, deep reinforcement learning
- Uncertainty, Bayesian networks, HMMs
- Applications: Natural Language Processing, Computer Vision

Assignments: Pac-man

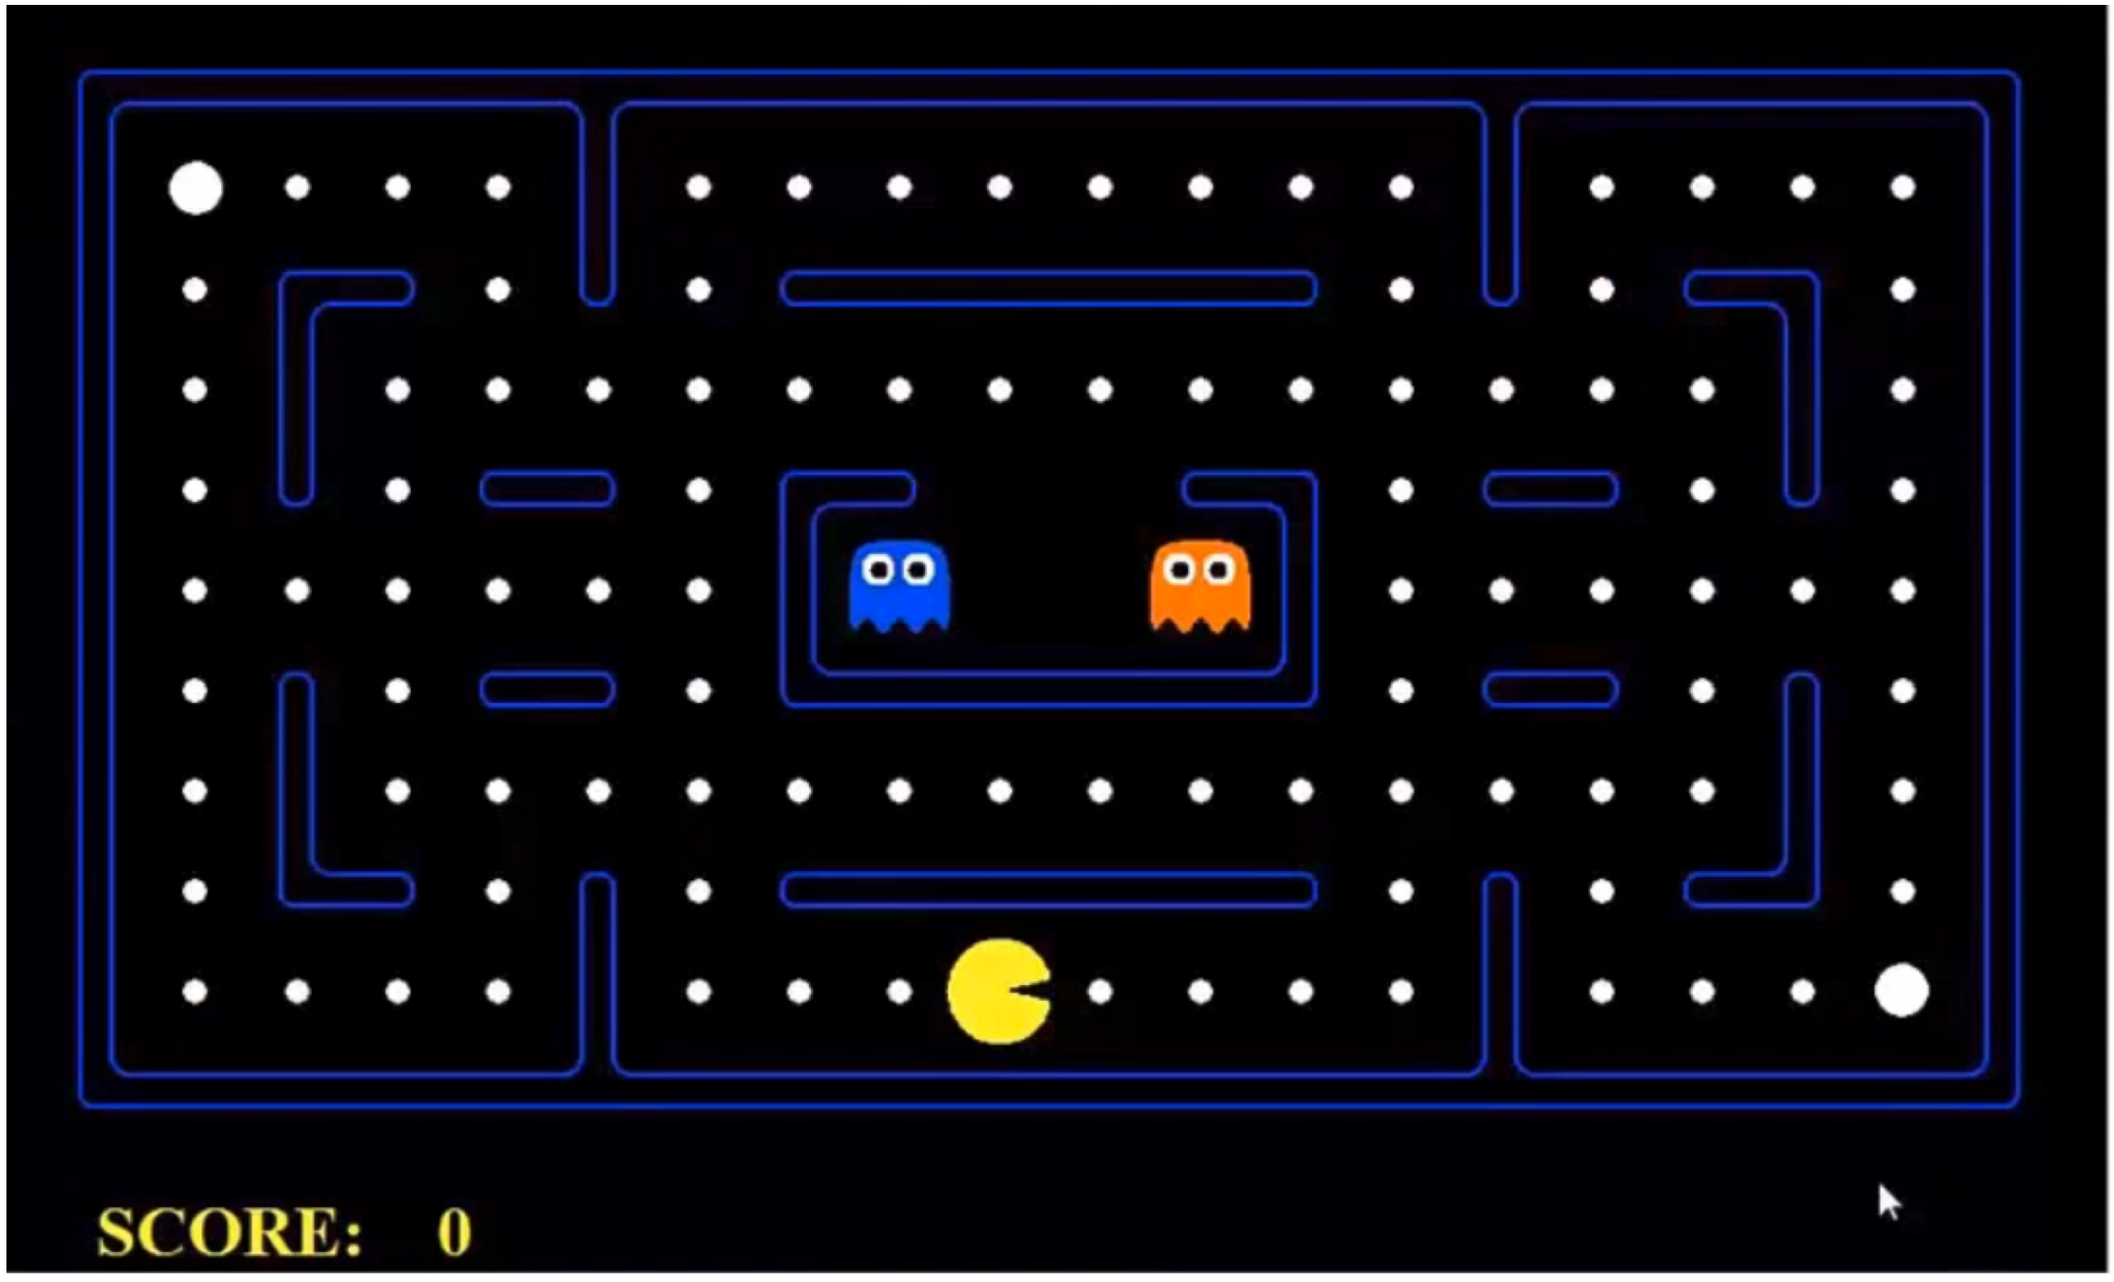

Originally developed at UC Berkeley:

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.

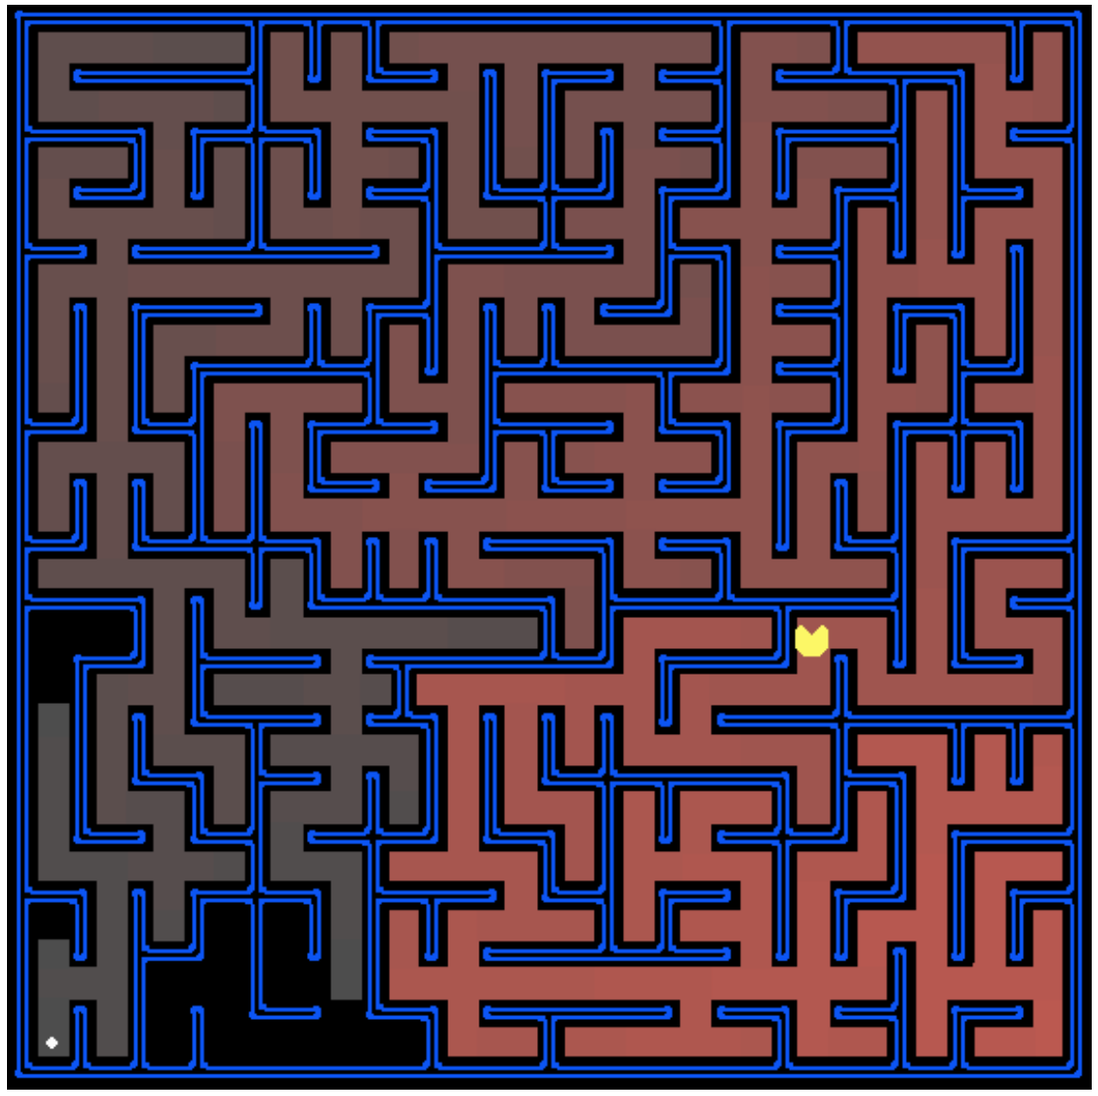

PS2: Game Playing

Goal:

- Play Pac-man!

Techniques:

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

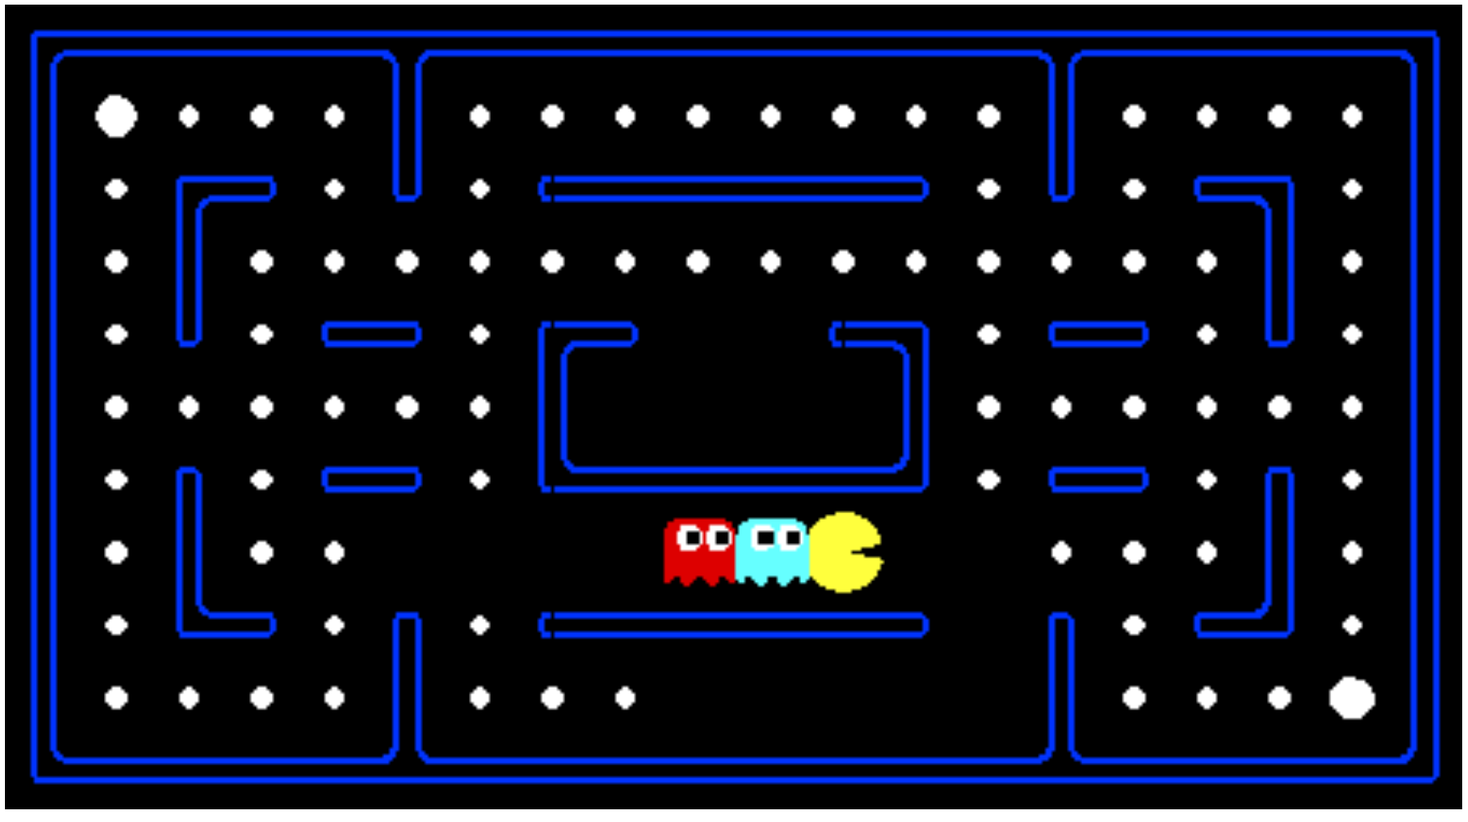

PS3: Ghostbusters

Goal:

- Help Pac-man hunt down the ghosts

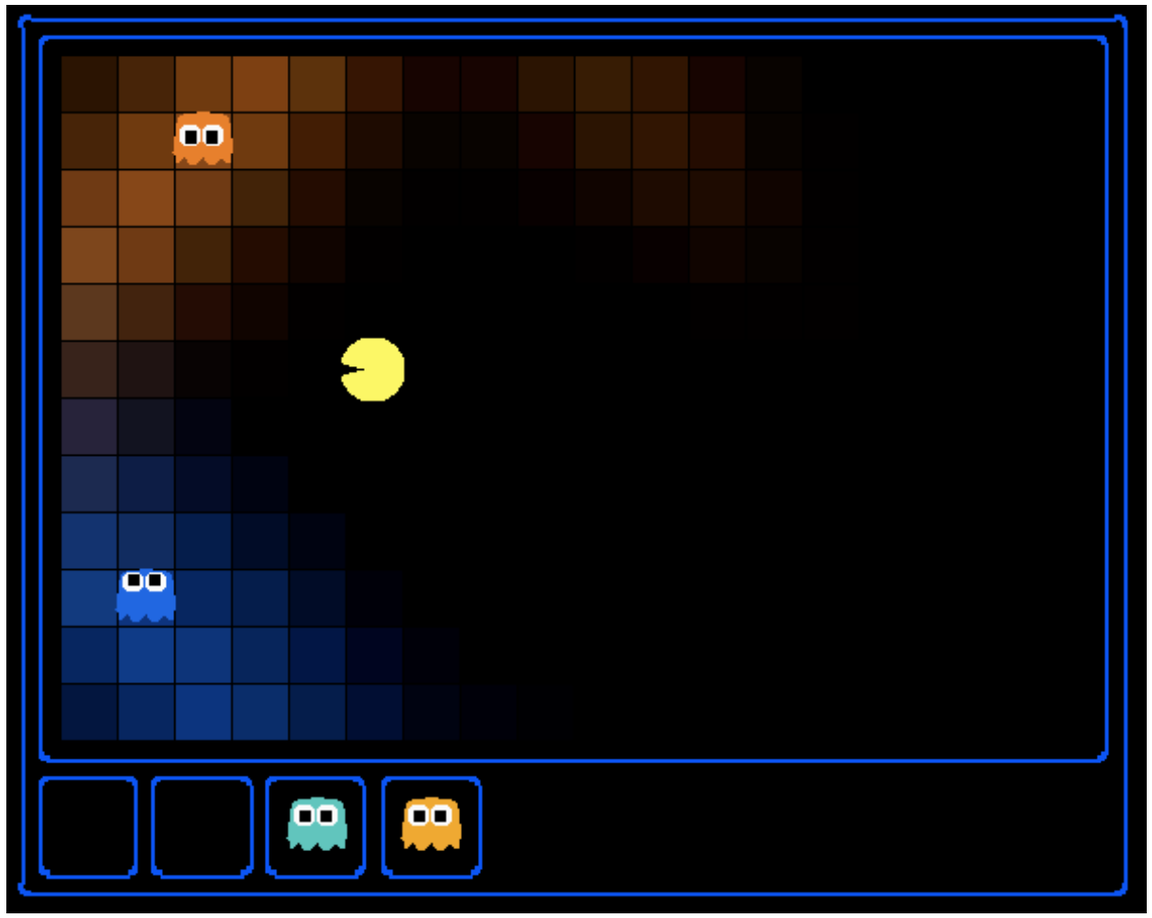

Techniques:

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

PS4: Reinforcement Learning

Goal:

- Help Pac-man learn about the world

Techniques:

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

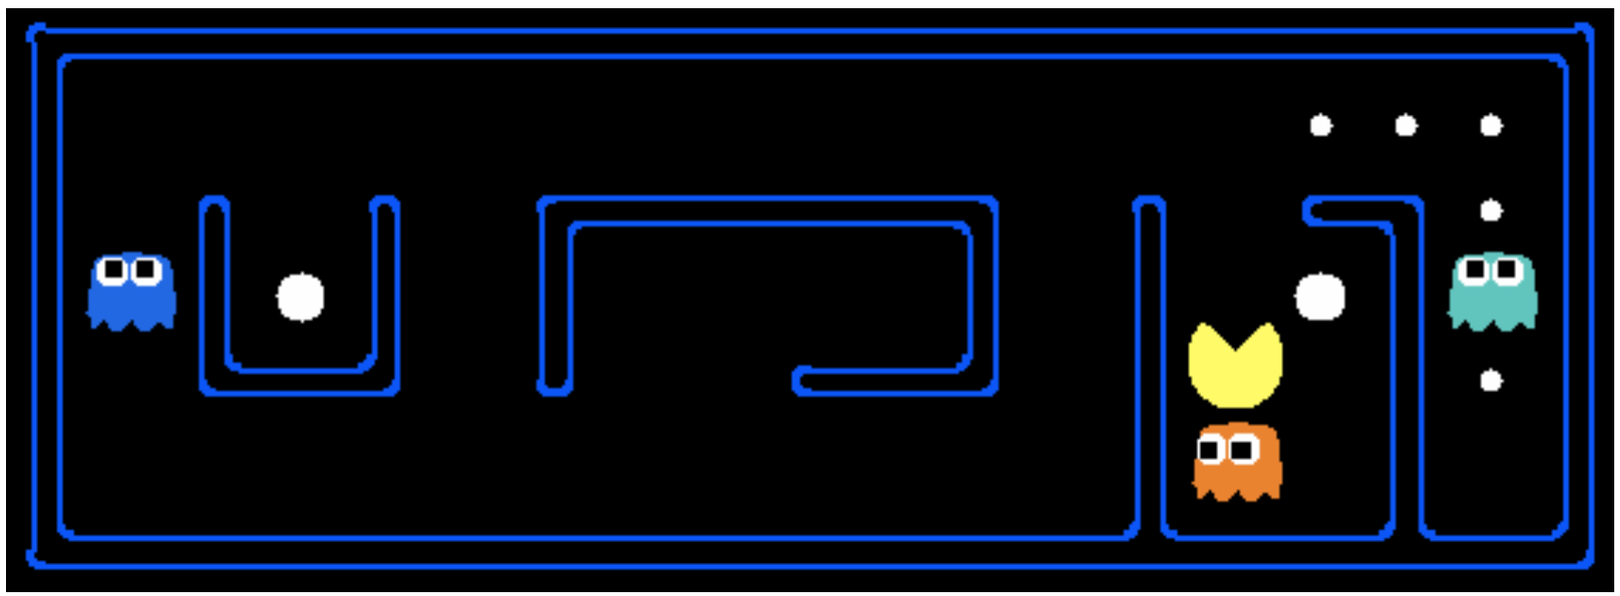

Project

Navigation and finding objects in scenes in the AI2Thor simulator using deep reinforcement learning algorithms

https://ai2thor.allenai.org/

What is CSE 573?

Prerequisites:

- Data Structures or equivalent
- Basic exposure to probability, data structures, and logic
- Familiar with Python

Work:

- Class participation (10%)
- Paper reviews (10%)
- Programming assignment (30%),
- Final Project (25%),
- Two quizzes (25%),

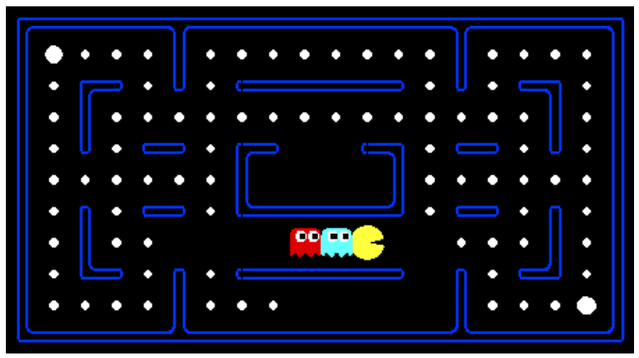

Pacman, autograder

To Do:

- Look at the course website: https://courses.cs.washington.edu/courses/cse573/19wi/