CS343H: Honors Artificial Intelligence

Introduction

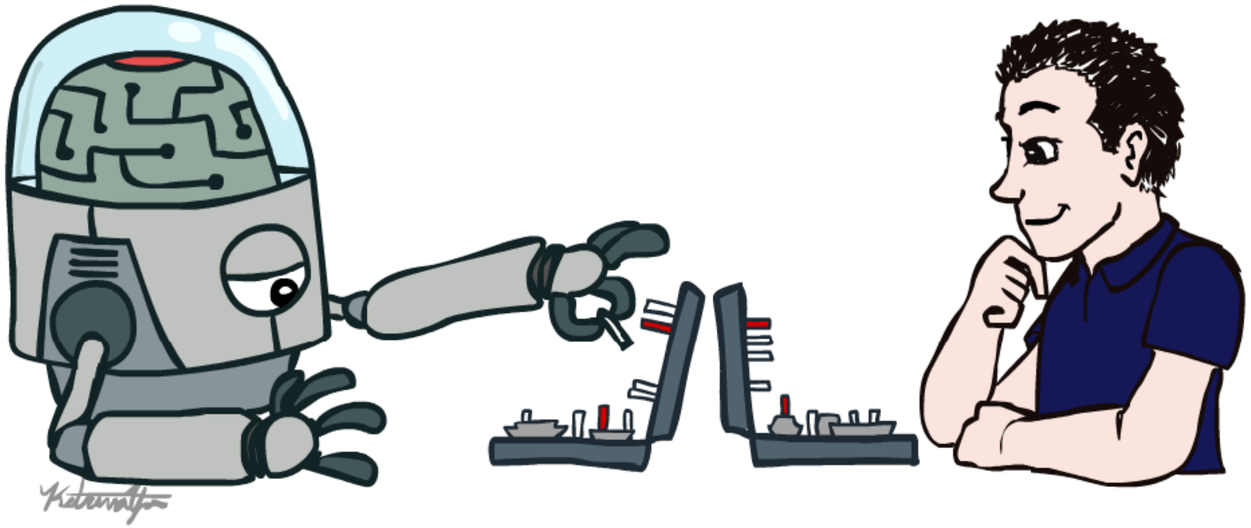

Prof. Peter Stone

University of Texas at Austin

[Based on slides created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley, modified by Scott Niekum at UT Austin.

All materials available at http://ai.berkeley.edu.]

Good Morning Colleagues

- Welcome to a fun, but challenging course
- Goal: Learn about Artificial Intelligence
  - Increase AI literacy
  - Prepare you for topics courses
  - Broad coverage of topics

Teaching staff

- Prof. Peter Stone
- TAs: Josiah Hanna, Yinan Zhao
- Proctor: Rohan Ramchand

Course Information

- Communication:
  - Announcements on webpage
  - Questions? Discussion on Piazza
  - Preference for Canvas?

Class website:

http://www.cs.utexas.edu/users/pstone/Courses/343Hfall17

(or Google “Peter Stone” and go to the Teaching tab)

- Course technology:
  - Autograded programming projects and interactive homeworks (unlimited submissions!) via edX
  - Make sure you have a CS Unix account IMMEDIATELY!
  - Create an edX account

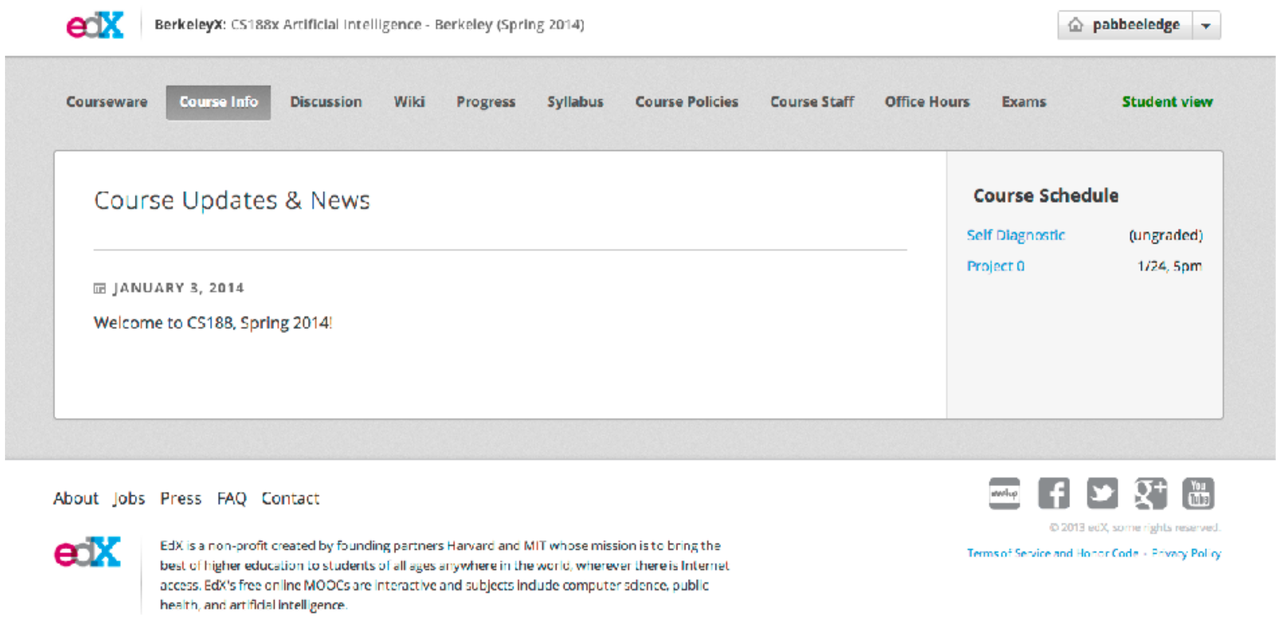

Course Information

- Prerequisites:
  - Upper division standing
  - No formal class pre-reqs
  - There will be a lot of math (and programming)

- Coursework
  - Reading assignments with written responses
  - ~9 homework assignments:
    - ~2 weeks for each, but overlapping
    - Online, autograded, solve together, submit alone
    - No late submissions accepted
  - 6 programming projects
    - Python, groups of 1 or 2 (except Project 0)
    - ~2 weeks for each, non-overlapping
    - 5 late days for semester, maximum 2 per project
  - One midterm, one final
  - Final Contest

Textbook

Russell & Norvig, AI: A Modern Approach, 3rd Ed.

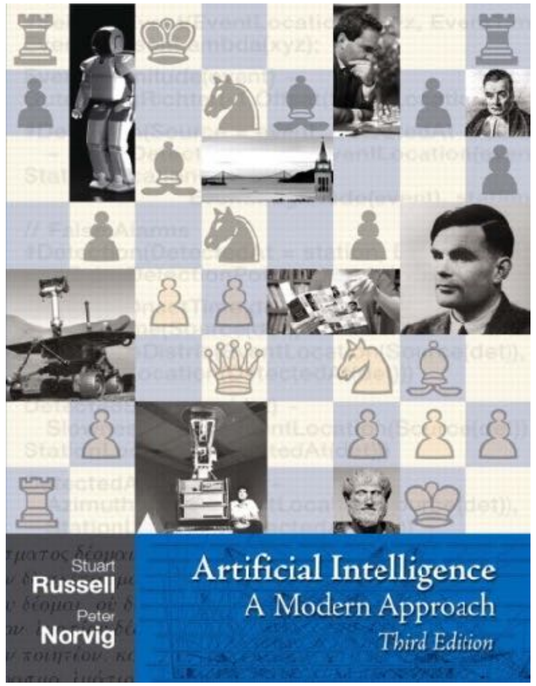

- After classes I'll post slides
- There will also be "step by step" videos posted for some topics

Warning: Not everything covered in the book will be covered in class (and to a small extent, vice versa). You are responsible for both.

Homework Exercises

- Online on edX
- Autograded text boxes / multiple choice
- Try as many times as you want!
- Goal: self-assess and prepare for tests
- Can discuss at high-level, but work alone
- No spoilers on Piazza discussions!

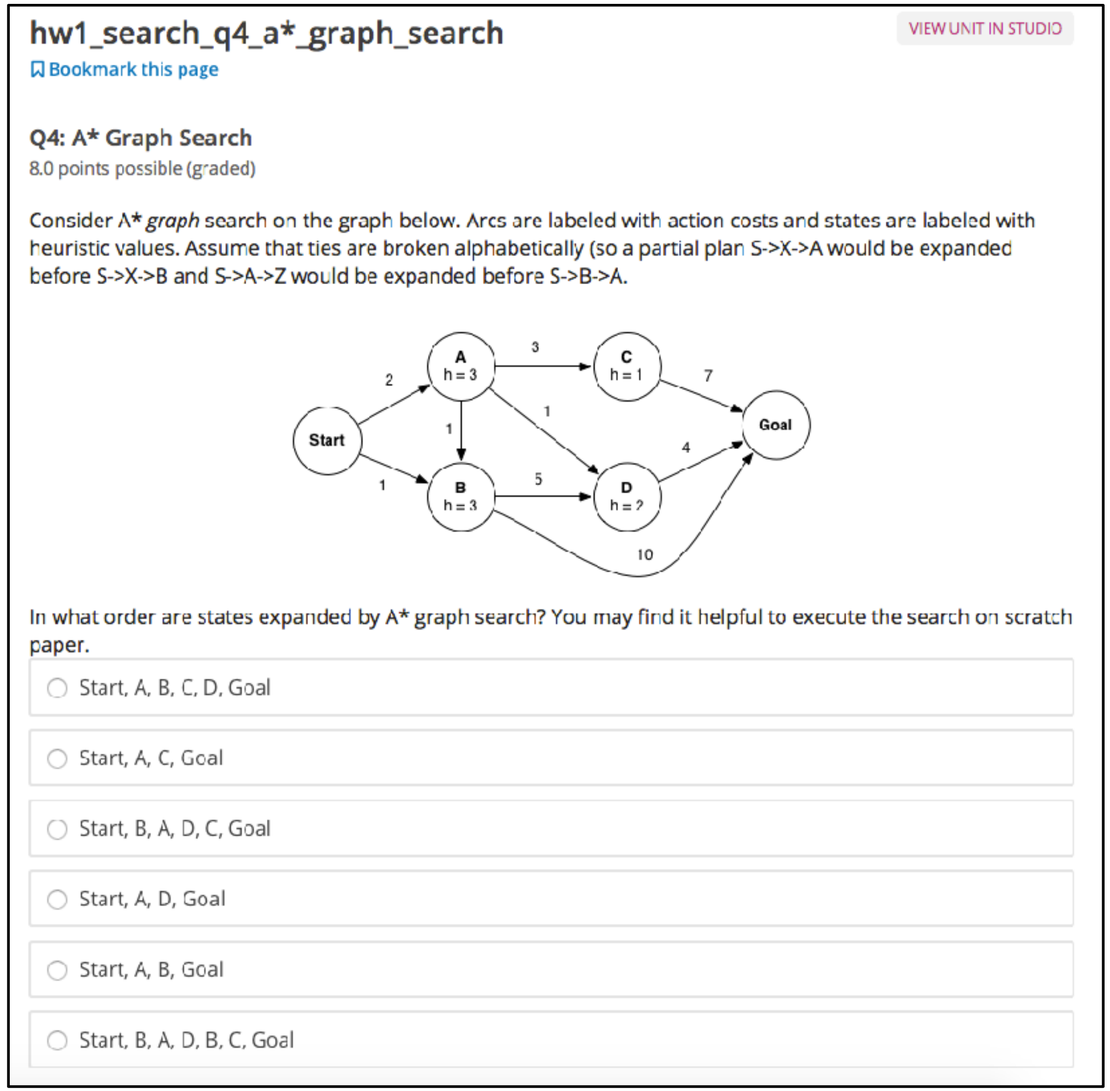

hw1_search_q4_a*_graph_search

VIEW UNIT IN STUDIO

Bookmark this page

Q4: A* Graph Search

8.0 points possible (graded)

Consider A* *graph* search on the graph below. Arcs are labeled with action costs and states are labeled with heuristic values. Assume that ties are broken alphabetically (so a partial plan S->X->A would be expanded before S->X->B and S->A->Z would be expanded before S->B->A.

In what order are states expanded by A* graph search? You may find it helpful to execute the search on scratch paper.

- Start, A, B, C, D, Goal
- Start, A, C, Goal
- Start, B, A, D, C, Goal
- Start, A, D, Goal
- Start, A, B, Goal
- Start, B, A, D, B, C, Goal

Programming Assignments

Pacman domain

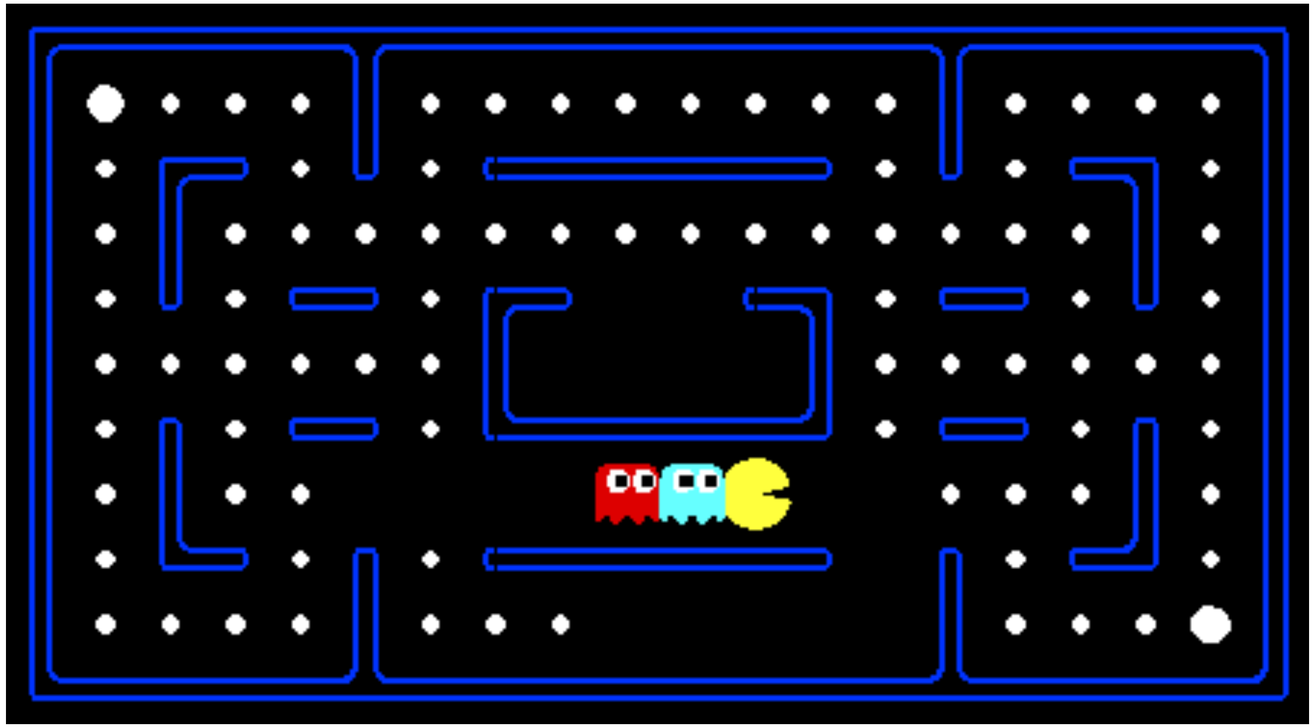

Projects include:

- path planning and search
- multi-agent game trees
- reinforcement learning
- state estimation
- classification
- final contest

Highly suggested: Pair programming
(switch “driver” and “observer” roles often)

Midterm and Final

- Midterm will cover roughly half the class material
- Final will be comprehensive
- One page of notes, but not open book

Syllabus

Official syllabus is online

Grades will be weighted as follows:

- Written responses to readings (5%)
- Class attendance and participation (10%)
- Homework Exercises (20%)
- Programming Assignments (25%)
- Midterm (15%)
- Final (25%)

Academic Honesty

READ THE STATEMENT IN THE SYLLABUS

- Discuss concepts, but don’t share solutions or written work with other students
- Don’t look for answers / code online or elsewhere
- Automated tools will be used to discover cheating
- If unsure, check departmental guidelines or ask — ignorance is not an excuse
- We will pursue the harshest penalties available, so please don’t cheat!
- To be clear: you will fail the class automatically and be reported to the university

Important This Week

- Important this week:
  - **Read and respond** to AI 100 report
  - **Read** the syllabus
  - **Reading assignment and email response** by Monday 5pm
  - **Register** for the class on edX (Click on link for HW1 to register for class after making account)
  - **Enroll and post something** on Piazza
  - **Be sure** that you have a usable CS Unix account - https://apps.cs.utexas.edu/udb/newaccount/
  - **P0: Python tutorial** is out (due on Thursday 9/7 at 9:30am via Canvas)
- Also important:
  - **If you are wait-listed,** you might or might not get in depending on how many students drop. Be patient if possible — many students often drop early in the course.
  - **Office Hours** after class on Thursday

Today

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

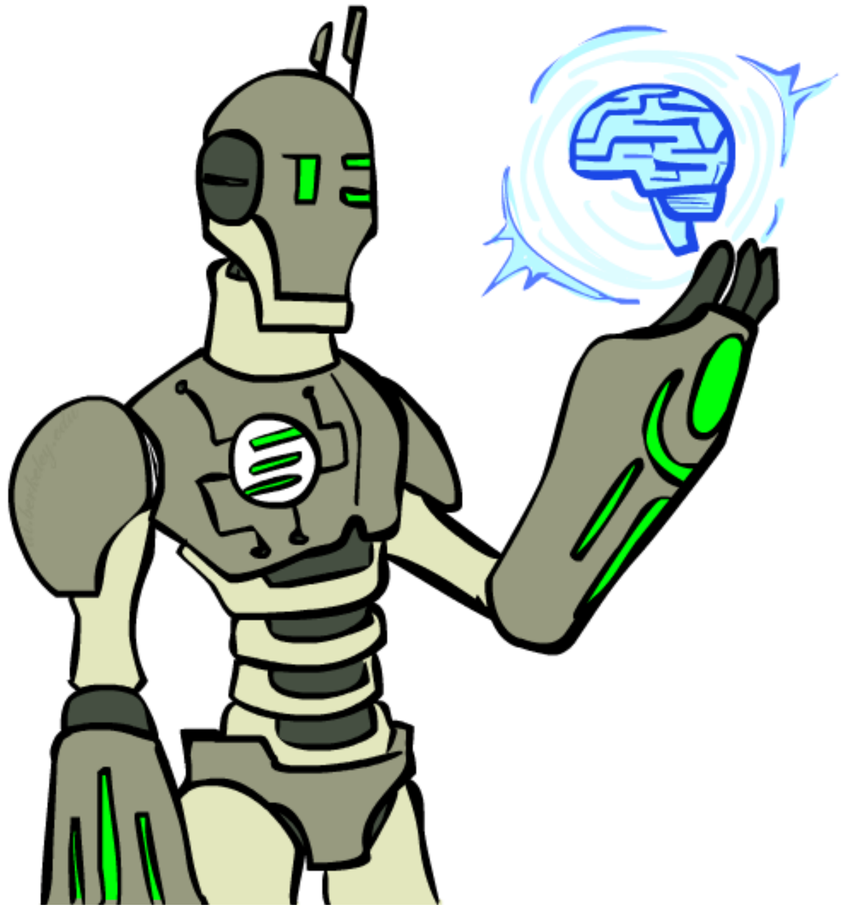

Sci-Fi AI?

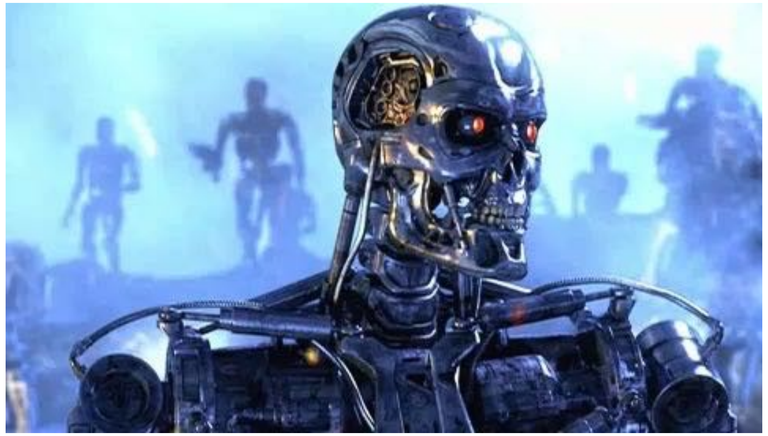

AI in the news

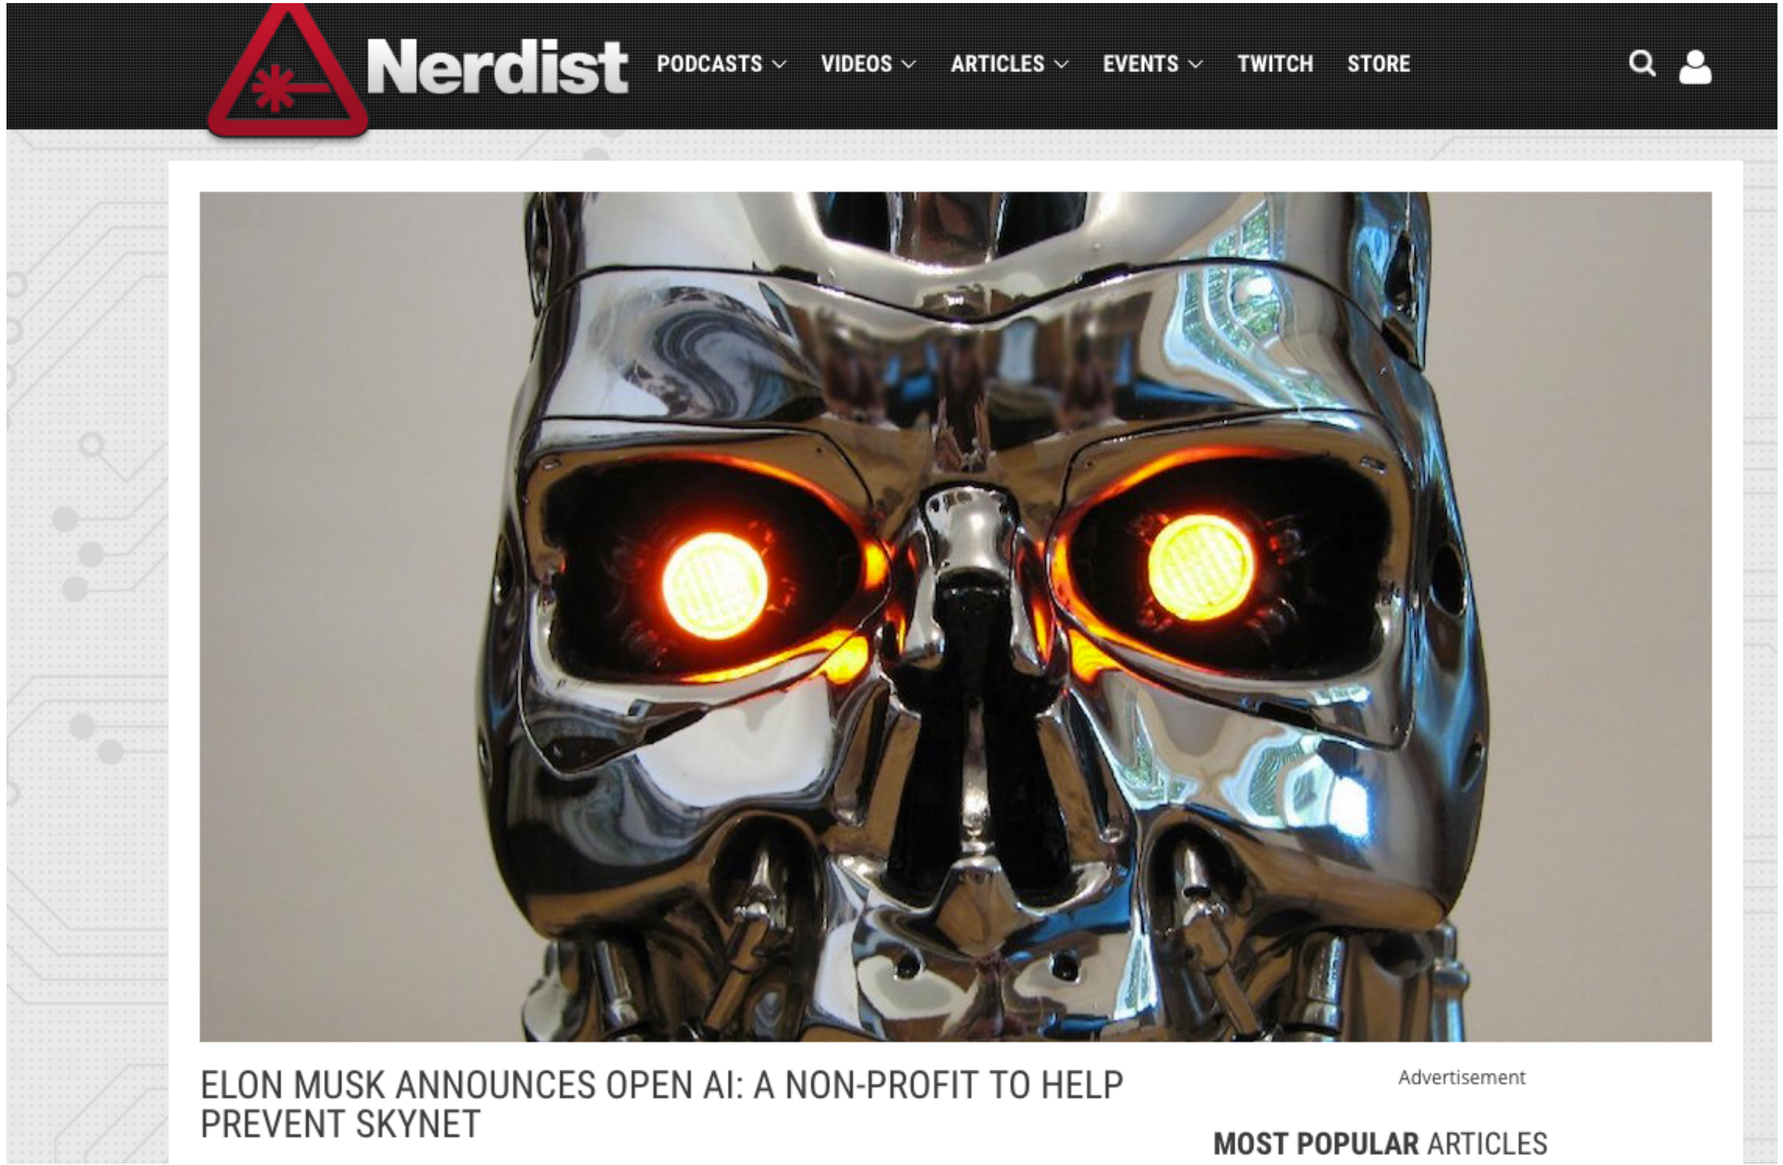

A definition for AI

“Artificial Intelligence (AI) is a science and a set of computational technologies that are inspired by — but typically operate quite differently from — the ways people use their nervous systems and bodies to sense, learn, reason, and take action.”

Philosophical questions

- AI is one of the great intellectual adventures of the 20th and 21st centuries.

  - What is a mind?
  - How can a physical object have a mind?
  - Can we build a mind?
  - Can trying to build one teach us what a mind is?

What is AI?

The science of making machines that:

What is AI?

The science of making machines that:

Think like people

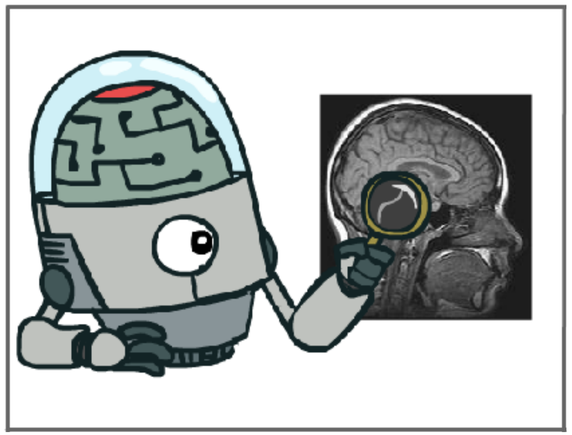

Thinking Like Humans?

- The cognitive science approach:
  - 1960s ``cognitive revolution'': information-processing psychology replaced prevailing orthodoxy of behaviorism (reflexive behaviors, classical conditioning, etc.)
- Scientific theories of internal activities of the brain
  - What level of abstraction? “Knowledge'' or “circuits”?
  - **Cognitive science:** Predicting and testing behavior of human subjects (top-down)
  - **Cognitive neuroscience:** Direct identification from neurological data (bottom-up)
  - Both approaches now distinct from AI
  - Both share with AI the following characteristic:

The available theories do not explain (or engender) anything resembling human-level general intelligence

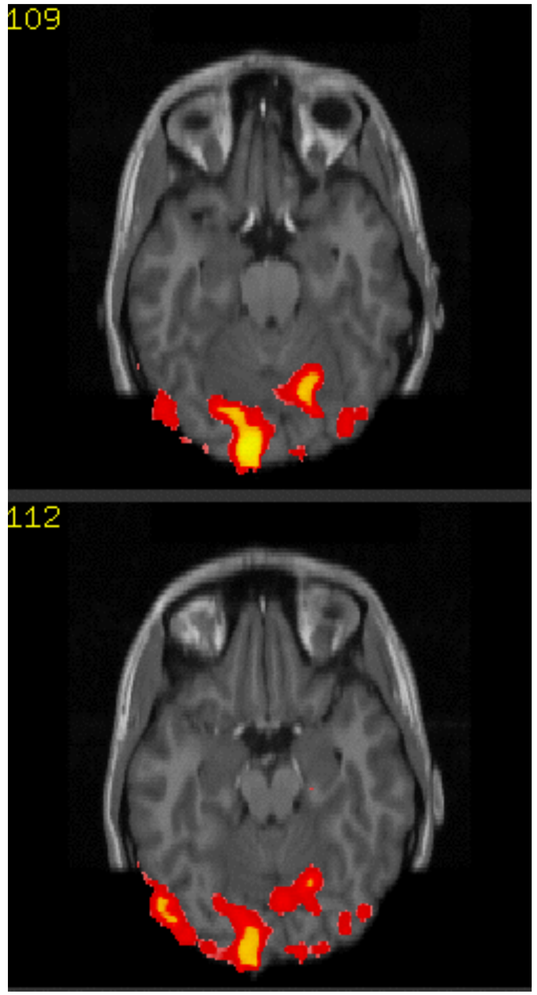

Images from Oxford fMRI center

What is AI?

The science of making machines that:

Think like people

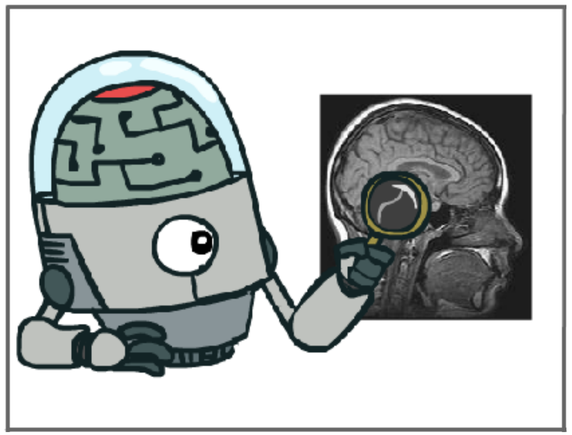

What is AI?

The science of making machines that:

Think like people

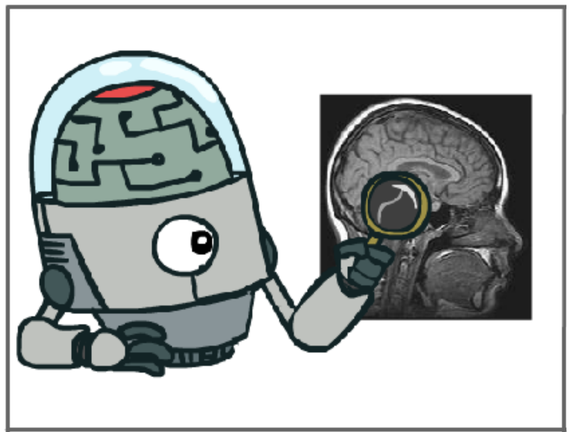

Act like people

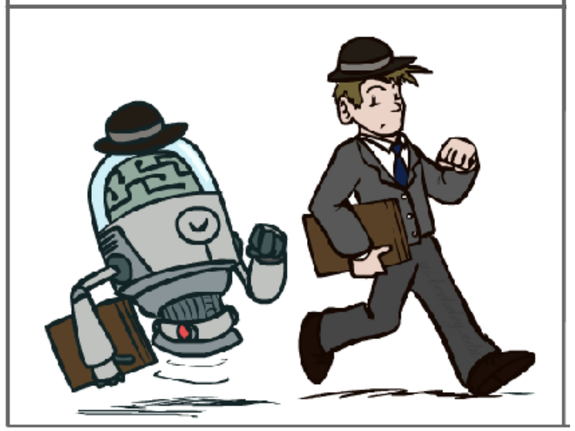

Acting Like Humans?

- Turing (1950) “Computing machinery and intelligence”
  - “Can machines think?” → “Can machines behave intelligently?”
  - Operational test for intelligent behavior: the *Imitation Game*

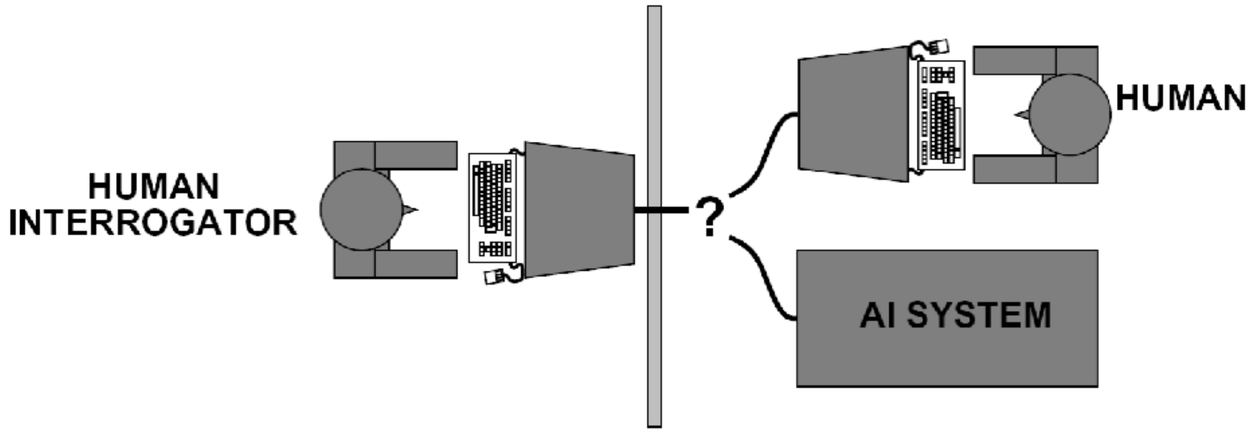

  - Predicted by 2000, a 30% chance of fooling a lay person for 5 minutes
  - Anticipated all major arguments against AI in following 50 years
  - Suggested major components of AI: knowledge, reasoning, language understanding, learning
- Problem: Turing test is not reproducible or amenable to mathematical analysis

What is AI?

The science of making machines that:

Think like people

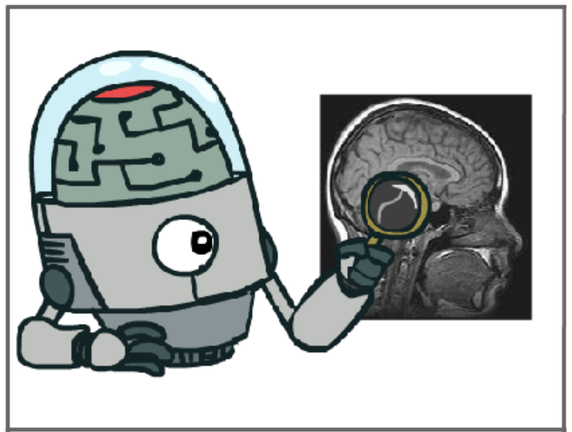

Act like people

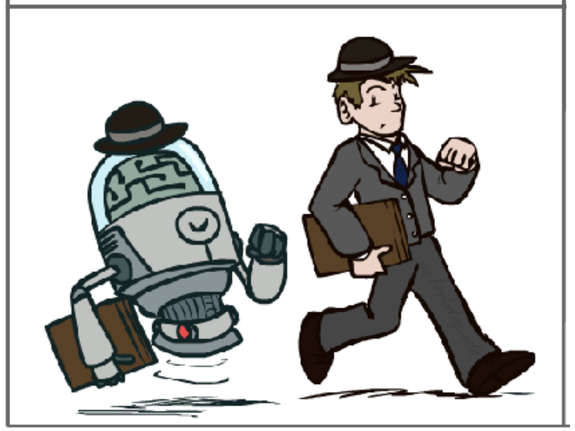

What is AI?

The science of making machines that:

Think like people

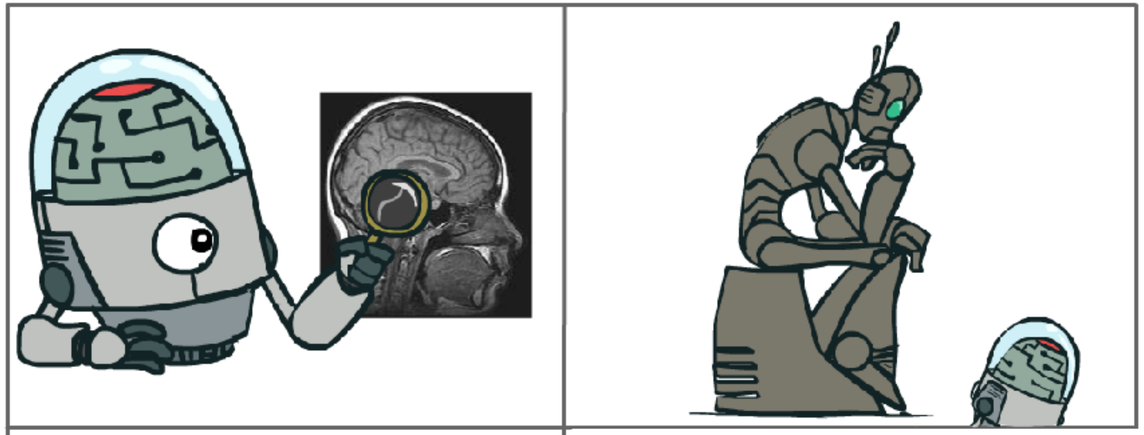

Think rationally

Act like people

Thinking Rationally?

- The “Laws of Thought” approach
  - What does it mean to “think rationally”?
  - Normative / prescriptive rather than descriptive

- Logicist tradition:
  - Logic: notation and rules of derivation for thoughts
  - Aristotle: what are correct arguments/thought processes?
  - Direct line through mathematics, philosophy, to modern AI

- Problems:
  - Not all intelligent behavior is mediated by logical deliberation
  - What is the purpose of thinking? What thoughts should I (bother to) have?
  - Logical systems tend to do the wrong thing in the presence of **uncertainty**

What is AI?

The science of making machines that:

Think like people

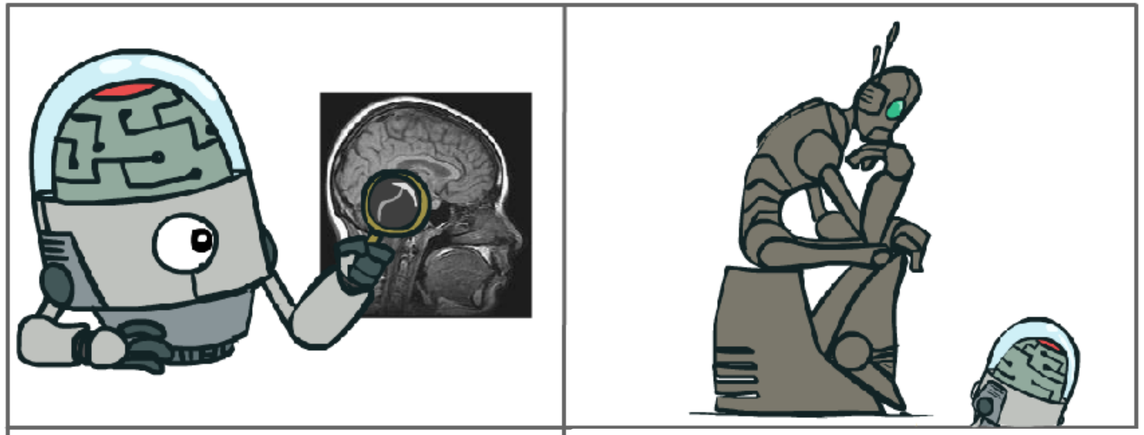

Think rationally

Act like people

What is AI?

The science of making machines that:

Think like people

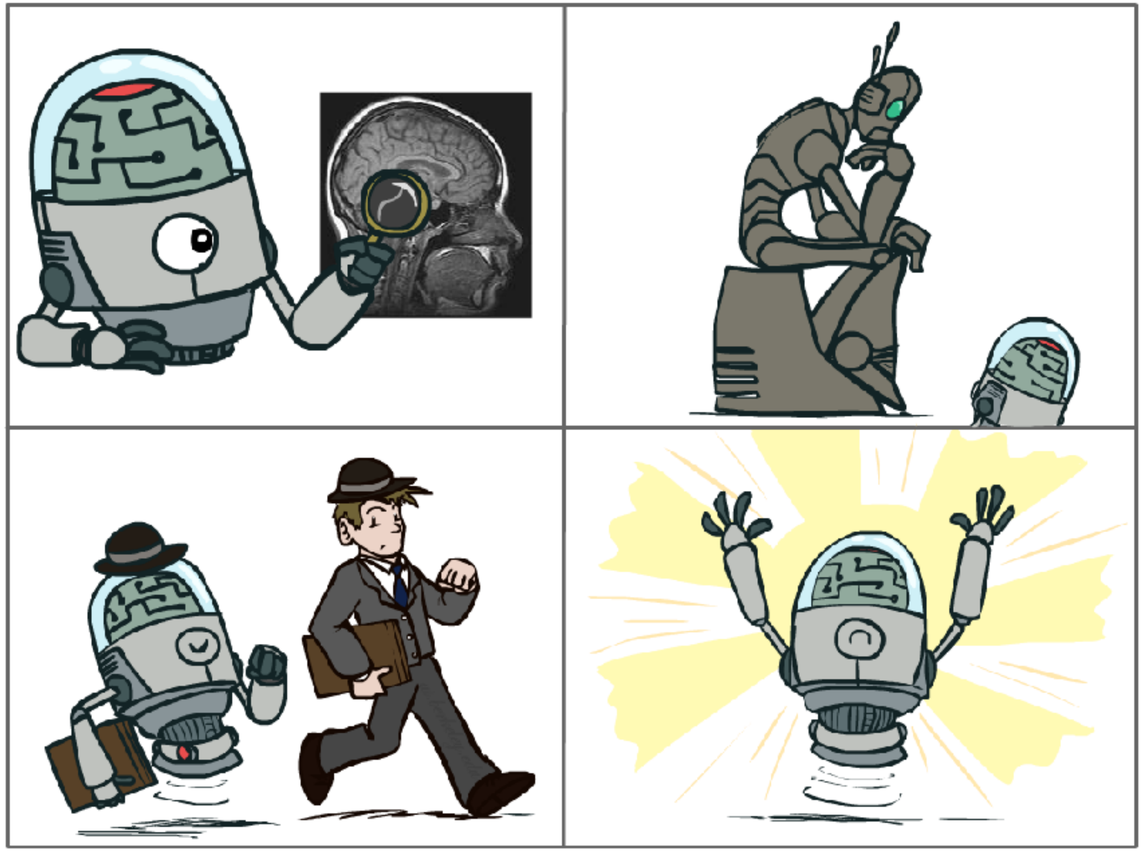

Think rationally

Act like people

Act rationally

Acting Rationally

- Rational behavior: doing the “right thing”
  - The right thing: that which is expected to maximize goal achievement, given the available information
  - Doesn't necessarily involve thinking, e.g., blinking
  - Thinking can be in the service of rational action
  - Entirely dependent on goals!
  - Irrational ≠ insane, irrationality is sub-optimal action
  - Rational ≠ successful

- Our focus here: rational agents
  - Systems which make the best possible decisions given goals, evidence, and constraints
  - In the real world, usually lots of uncertainty
    - … and lots of complexity
  - Usually, we’re just approximating rationality

Rational Decisions

We'll use the term **rational** in a very specific, technical way:

- Rational: maximally achieving pre-defined goals
- Rationality 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**

Another title for this course could be:

Computational Rationality

Maximize Your Expected Utility

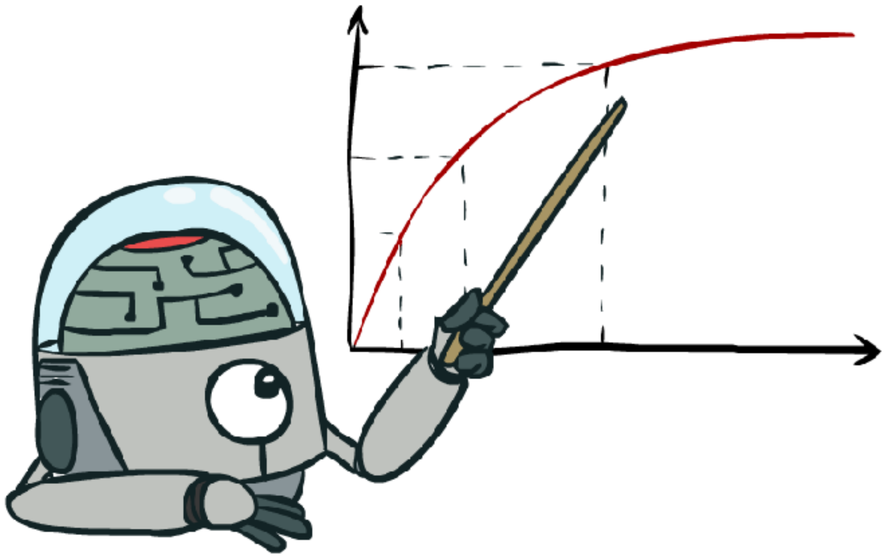

What About the Brain?

- Brains (human minds) are very good at making rational decisions, but not perfect
- Brains aren’t as modular as software, so hard to reverse engineer!
- “Brains are to intelligence as wings are to flight”
- Lessons learned from the brain: memory and simulation are key to decision making

A (Short) History of AI

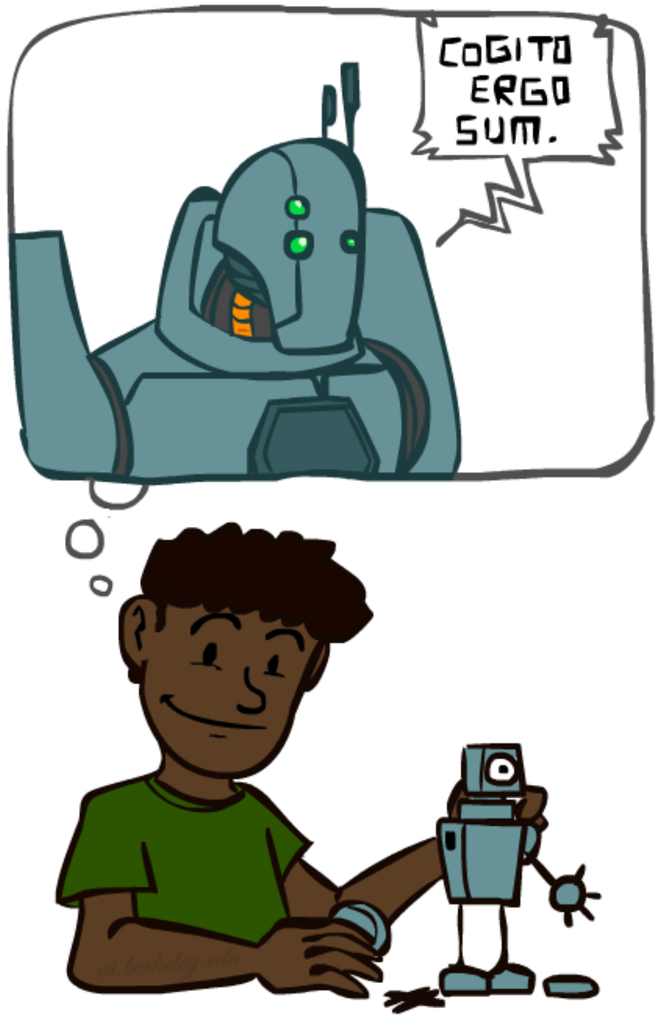

Demo: HISTORY – MT1950.wmv

A (Short) History of AI

A (Short) History of AI

- 1940-1950: Early days
  - 1943: McCulloch & Pitts: Boolean circuit model of brain
  - 1950: Turing's “Computing Machinery and Intelligence”
- 1950—70: Excitement: Look, Ma, no hands!
  - 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
- 1970—90: Knowledge-based approaches
  - 1969—79: Early development of knowledge-based systems
  - 1980—88: Expert systems industry booms
  - 1988—93: Expert systems industry busts: “AI Winter”
- 1990—: Statistical approaches
  - Resurgence of probability, focus on uncertainty
  - General increase in technical depth
  - Agents and learning systems… “AI Spring”?
- 2000—: Where are we now?

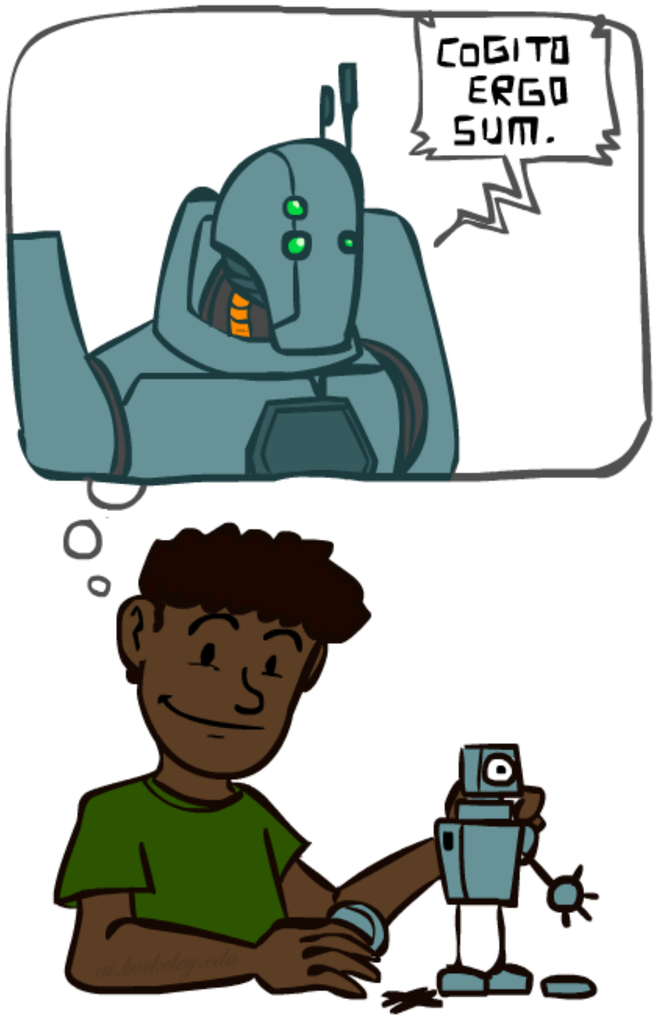

What Can AI Do?

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

- ✓ Play a decent game of table tennis?
- ✓ Play a decent game of Jeopardy?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along 6th Street on a Friday night?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at HEB?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person for an hour?
- ? Perform a surgical operation?
- ✓ Put away the dishes and fold the laundry?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✗ Write an intentionally funny story?

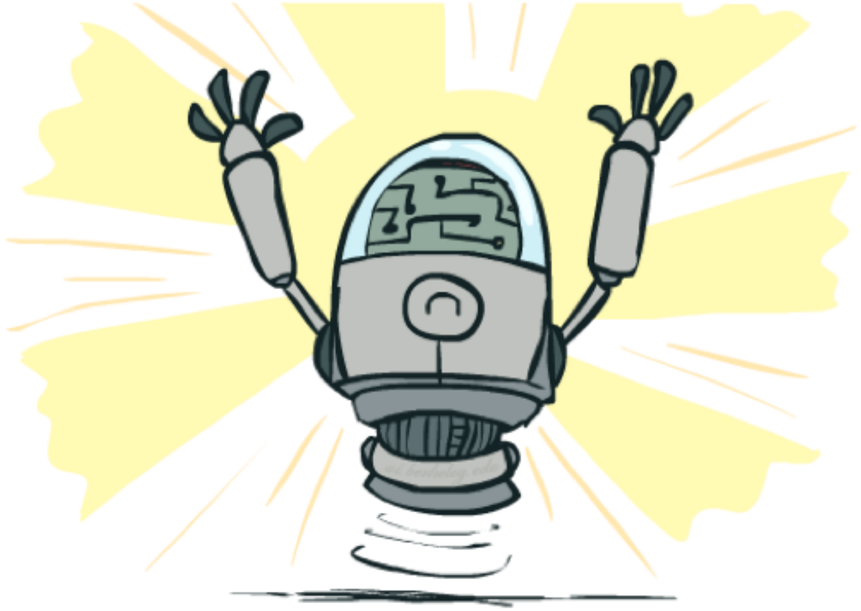

Unintentionally Funny Stories

- One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe walked to the oak tree. He ate the beehive. The End.

- Henry Squirrel was thirsty. He walked over to the river bank where his good friend Bill Bird was sitting. Henry slipped and fell in the river. Gravity drowned. The End.

- Once upon a time there was a dishonest fox and a vain crow. One day the crow was sitting in his tree, holding a piece of cheese in his mouth. He noticed that he was holding the piece of cheese. He became hungry, and swallowed the cheese. The fox walked over to the crow. The End.

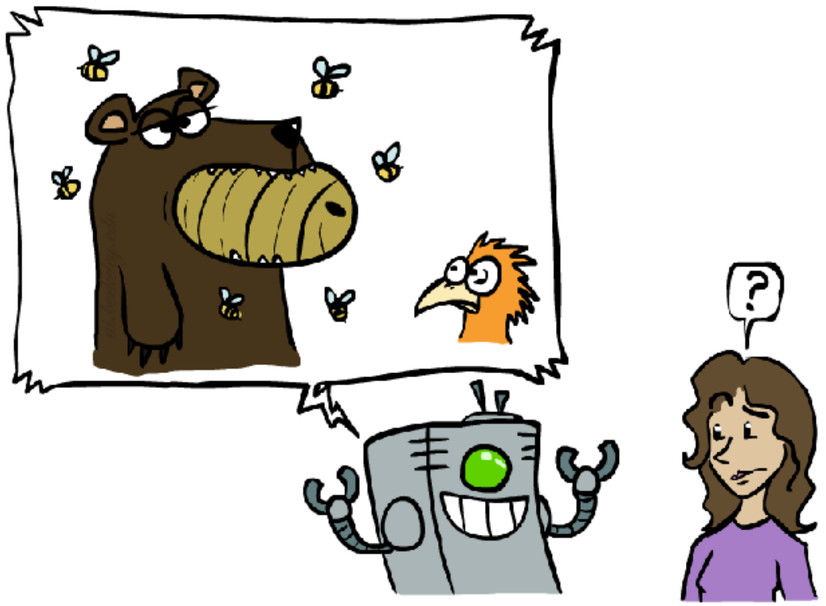

[Shank, Tale-Spin System, 1984]

Natural Language

- Speech technologies (e.g. Siri)
  - Automatic speech recognition (ASR)
  - Text-to-speech synthesis (TTS)
  - Dialog systems

Natural Language Video

Natural Language

- Speech technologies (e.g. Siri)
  - Automatic speech recognition (ASR)
  - Text-to-speech synthesis (TTS)
  - Dialog systems

- Language processing technologies
  - Question answering
  - Machine translation

"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

  - Web search
  - Text classification, spam filtering, etc…

Natural Language

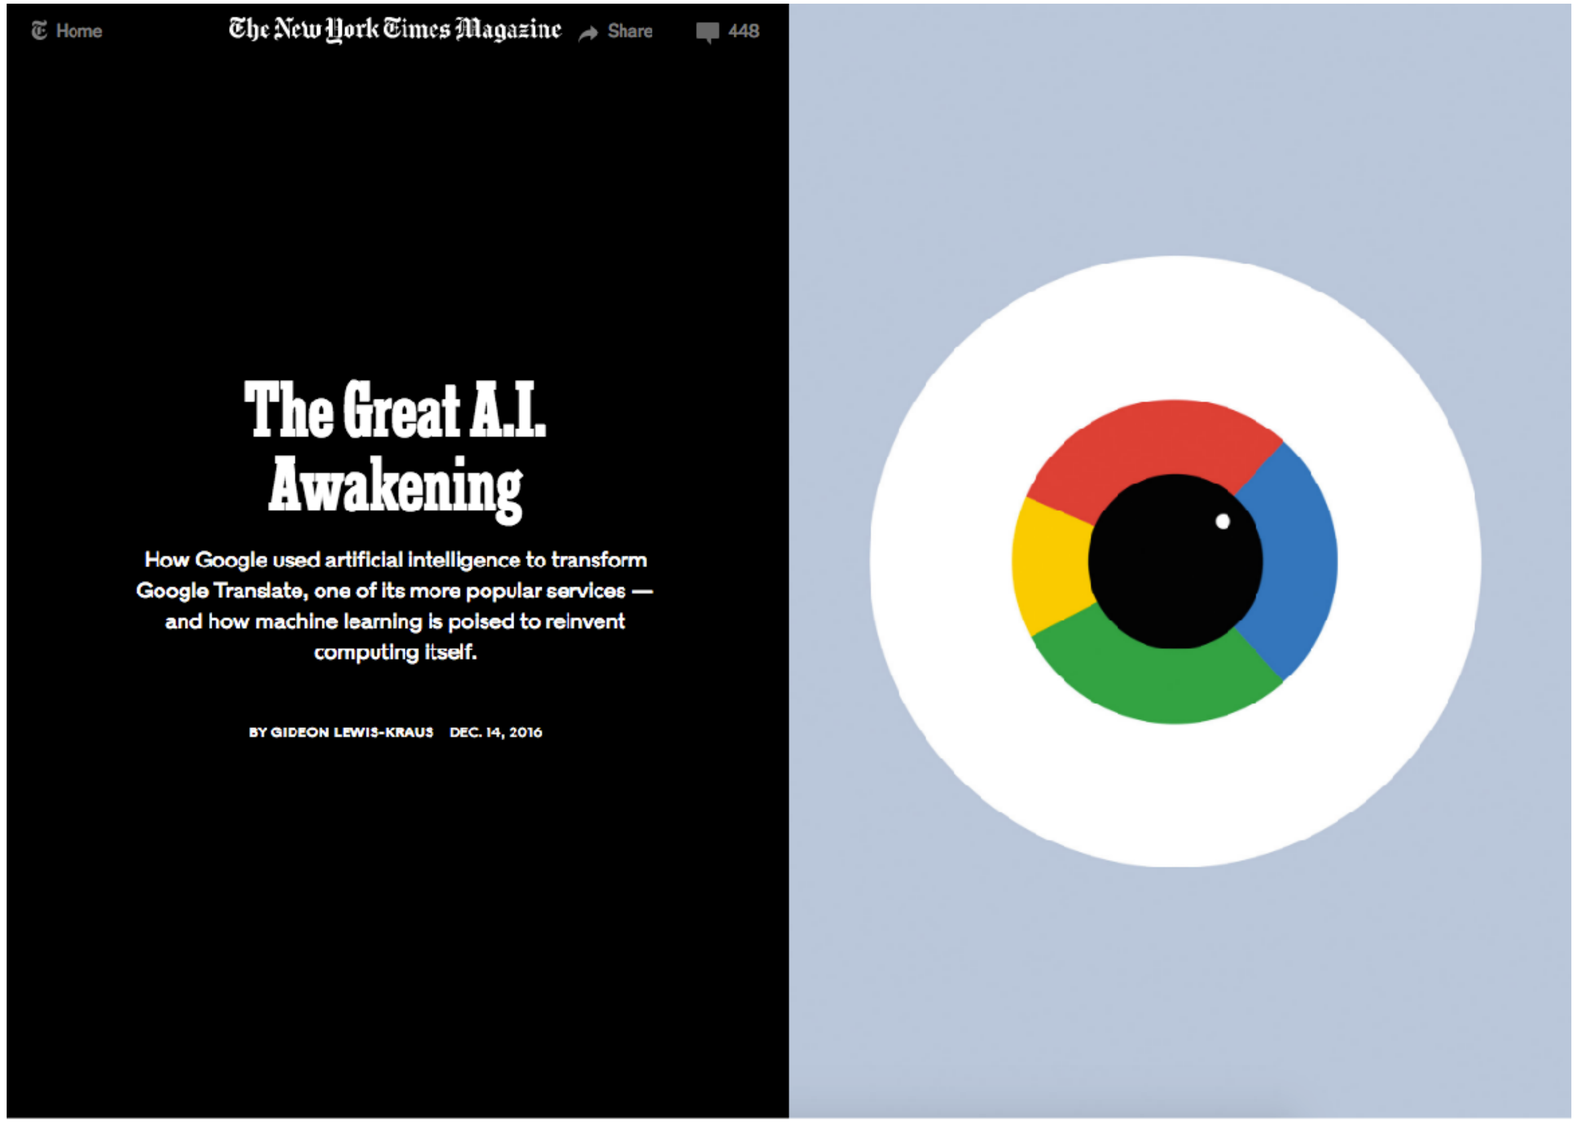

Deep learning

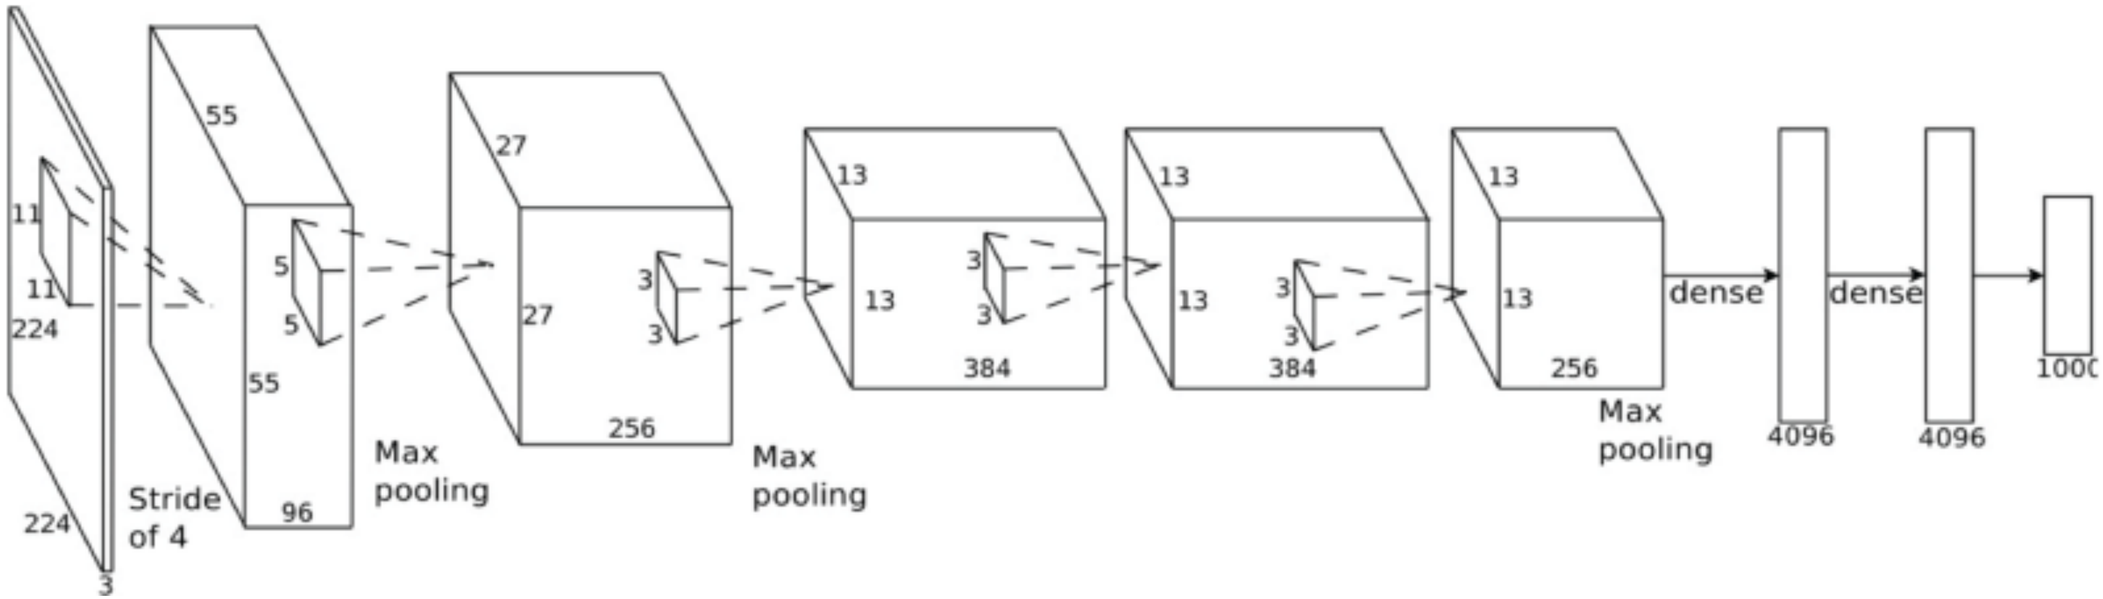

Natural Language

DeepDrumpf
@DeepDrumpf

I'm a Neural Network trained on Trump's transcripts. Priming text in [ ]s. Donate (gofundme.com/deepdrumpf) to interact! Created by @hayesbh.

DeepDrumpf @DeepDrumpf · 19 Oct 2016
[Math is a] common democrat lie. It can't make the budget great. I'll have the best economy. #debatenight
5 219 360 •••

DeepDrumpf @DeepDrumpf · 19 Oct 2016
[This election is rigged] I will give every American a solid gold nuclear weapon, we're going to defeat the world.

DeepDrumpf @DeepDrumpf · 19 Oct 2016
[We'll fix the economy by] selling out veterans. I will get the power, from some core of hell. I will be the most powerful. #debatenight
2 148 275 •••

Vision (Perception)

- Object and face recognition
- Scene segmentation
- Image classification

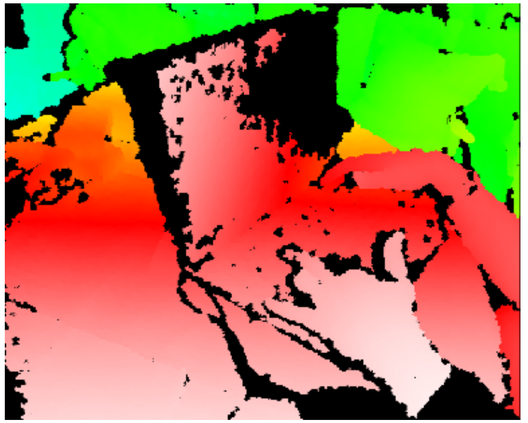

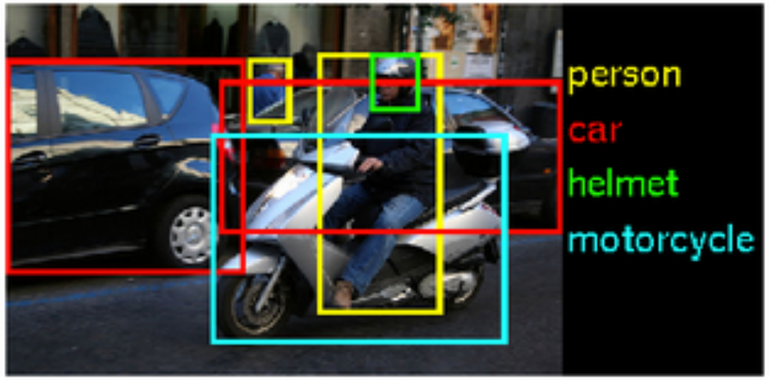

Images from Erik Sudderth (left), wikipedia (right)

Object Tracking

object detection / 3D pose estimation

arbitrary view rendered with estimated 3D pose

Perception + Natural Language

What vegetable is on the plate?
Neural Net: broccoli
Ground Truth: broccoli

What color are the shoes on the person's feet ?
Neural Net: brown
Ground Truth: brown

How many school busses are there?
Neural Net: 2
Ground Truth: 2

What sport is this?
Neural Net: baseball
Ground Truth: baseball

What is on top of the refrigerator?
Neural Net: magnets
Ground Truth: cereal

What uniform is she wearing?
Neural Net: shorts
Ground Truth: girl scout

What is the table number?
Neural Net: 4
Ground Truth:40

What are people sitting under in the back?
Neural Net: bench
Ground Truth: tent

Perception + Natural Language

We won’t discuss NLP and perception directly, but we will cover:

- Bayes nets
- Supervised learning
- Deep learning
- Course outline

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

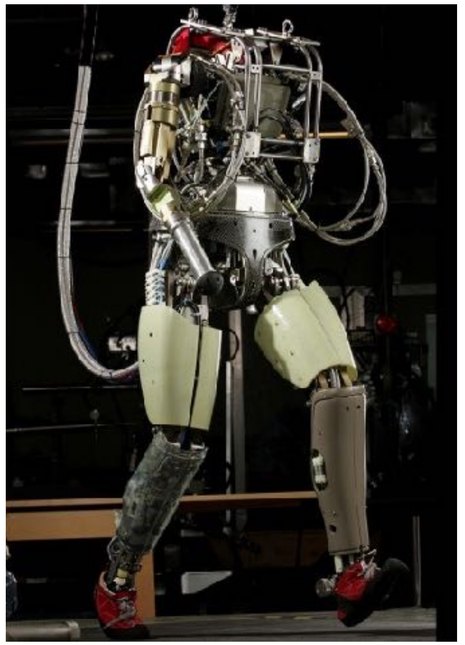

Images from UC Berkeley, Boston Dynamics, RoboCup, Google

Robot Laundry

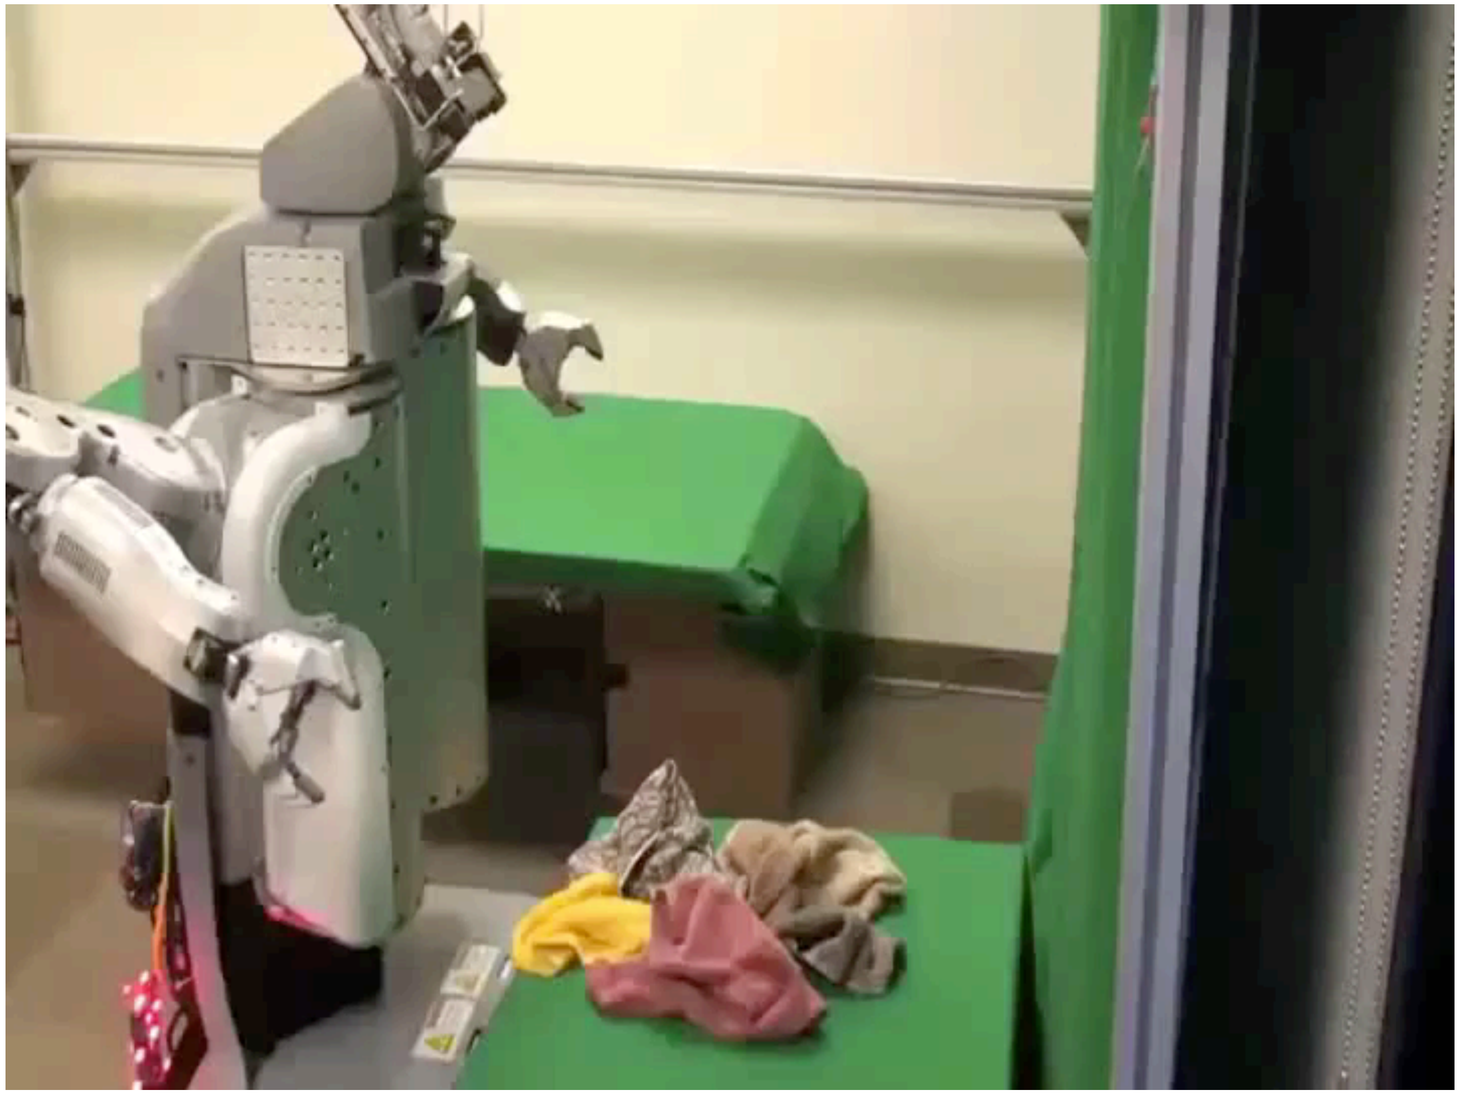

Robot Soccer

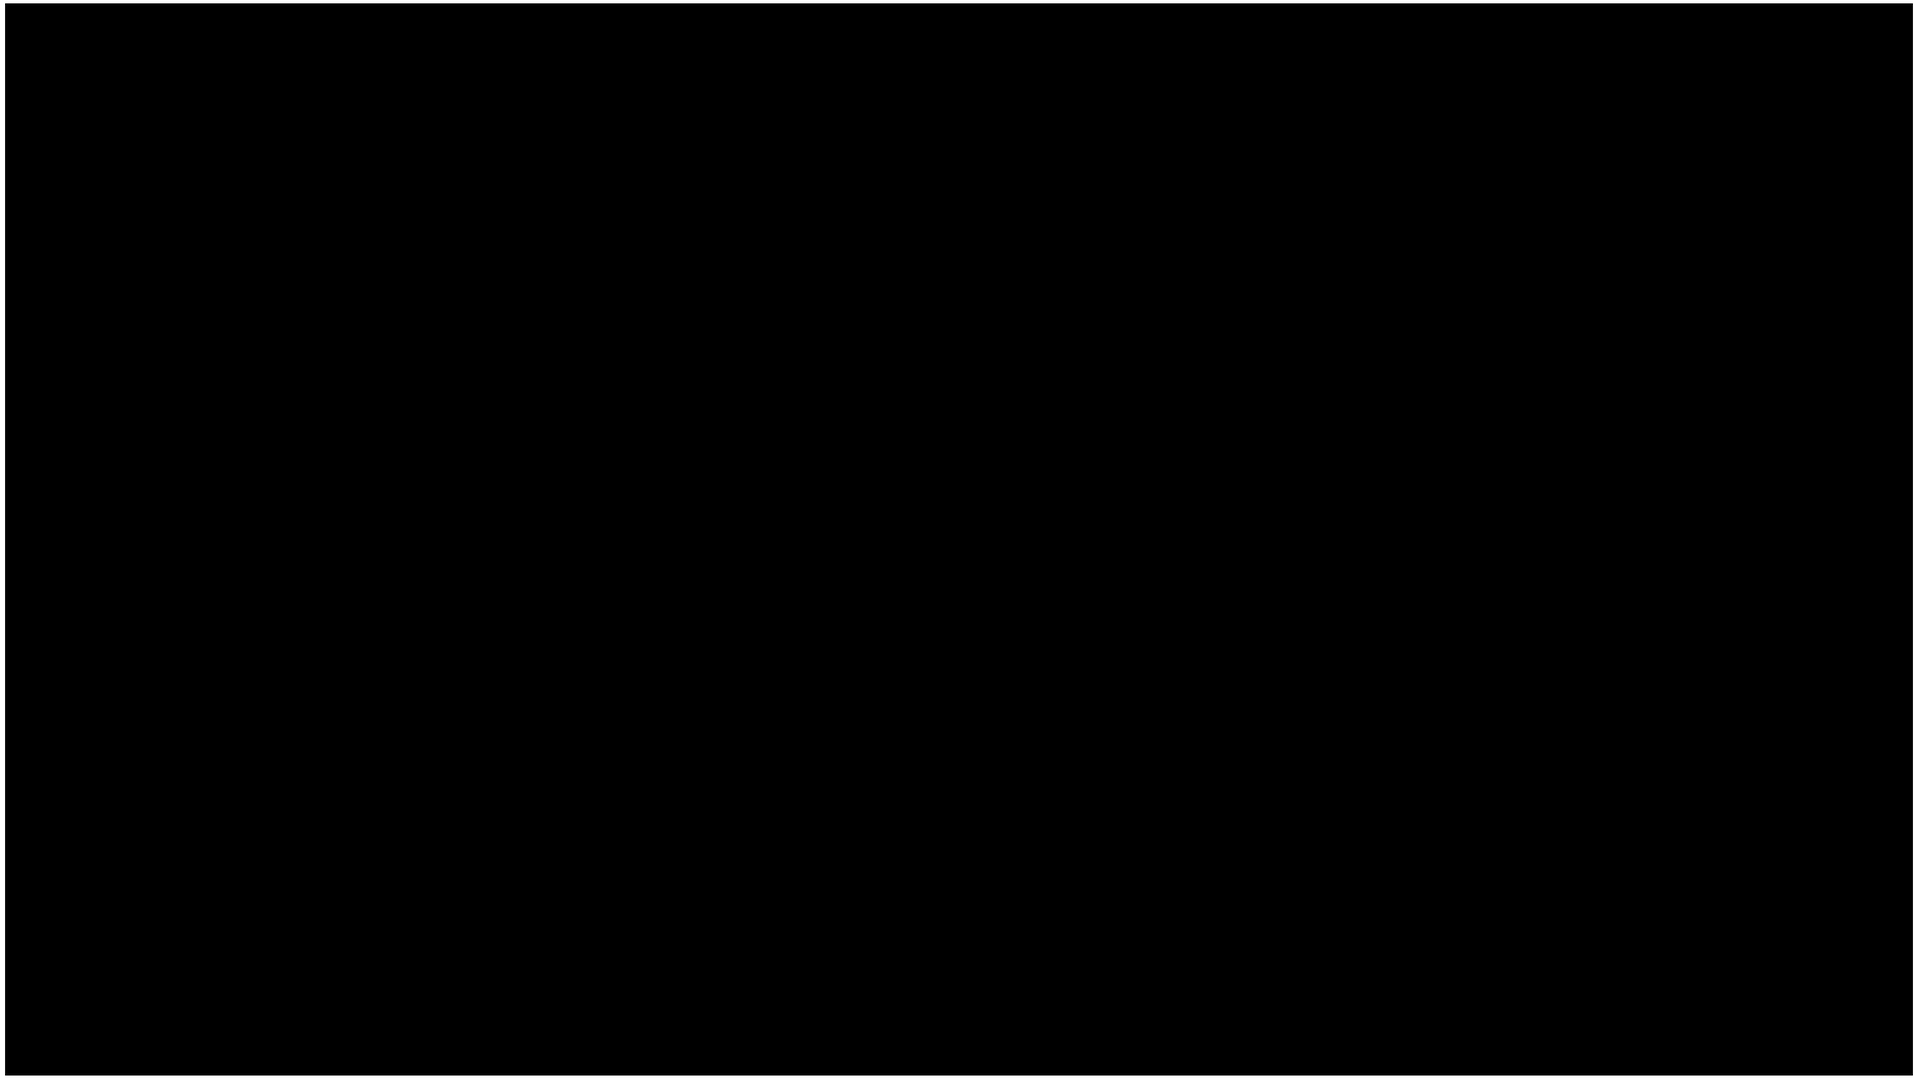

Learning from demonstration

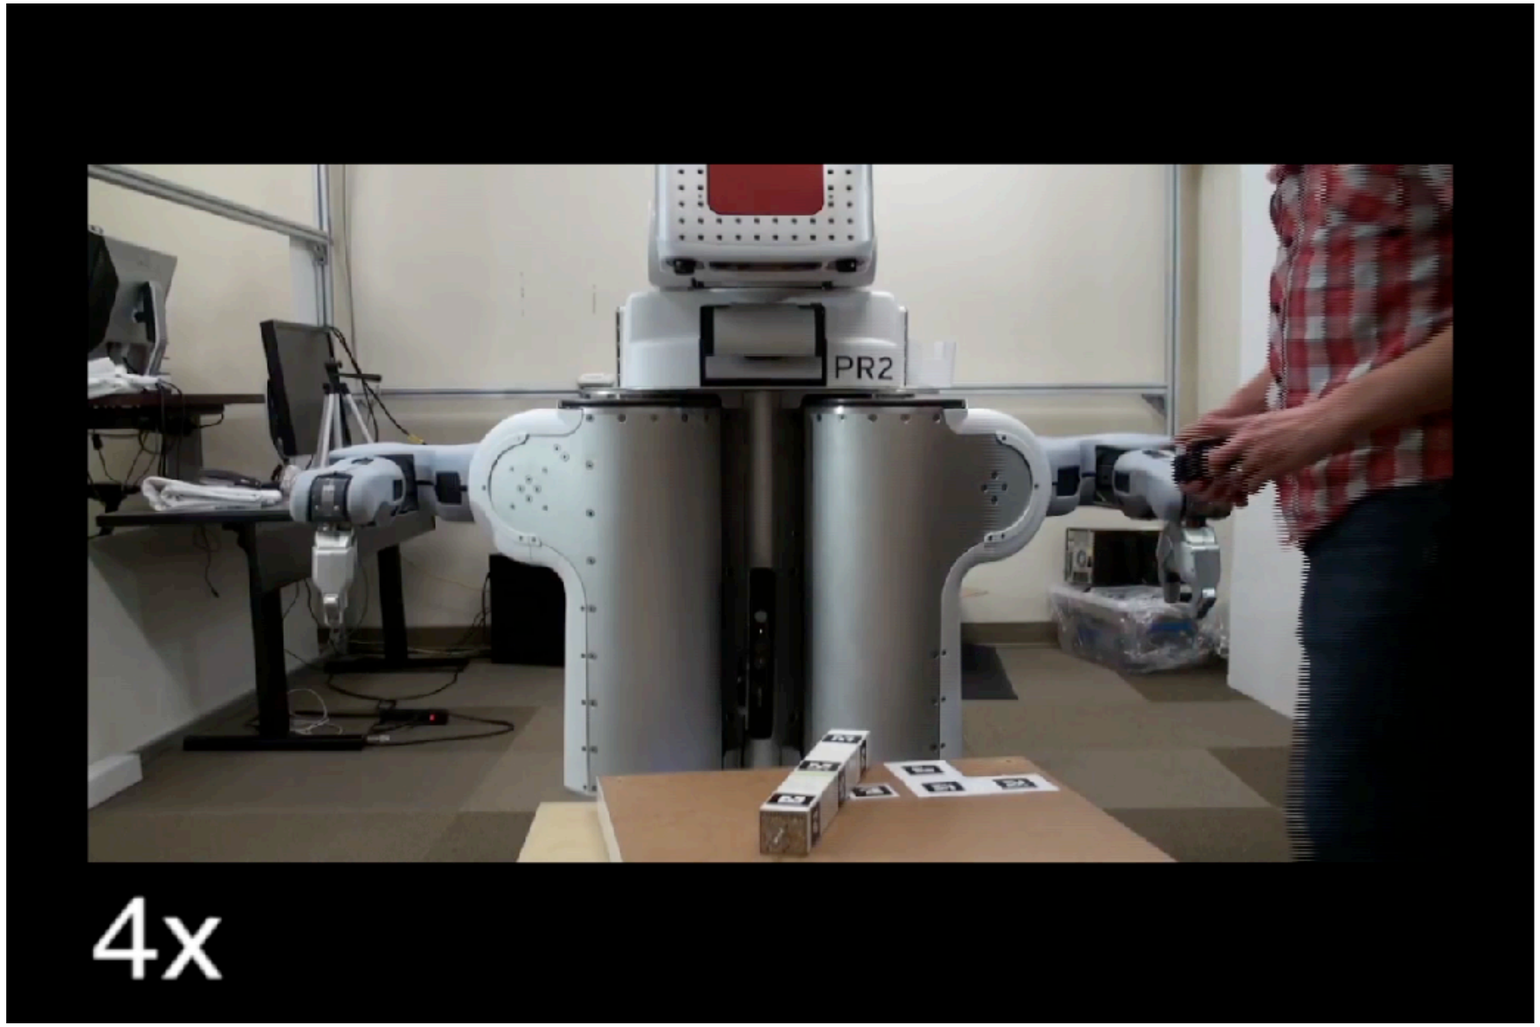

Learning from demonstration

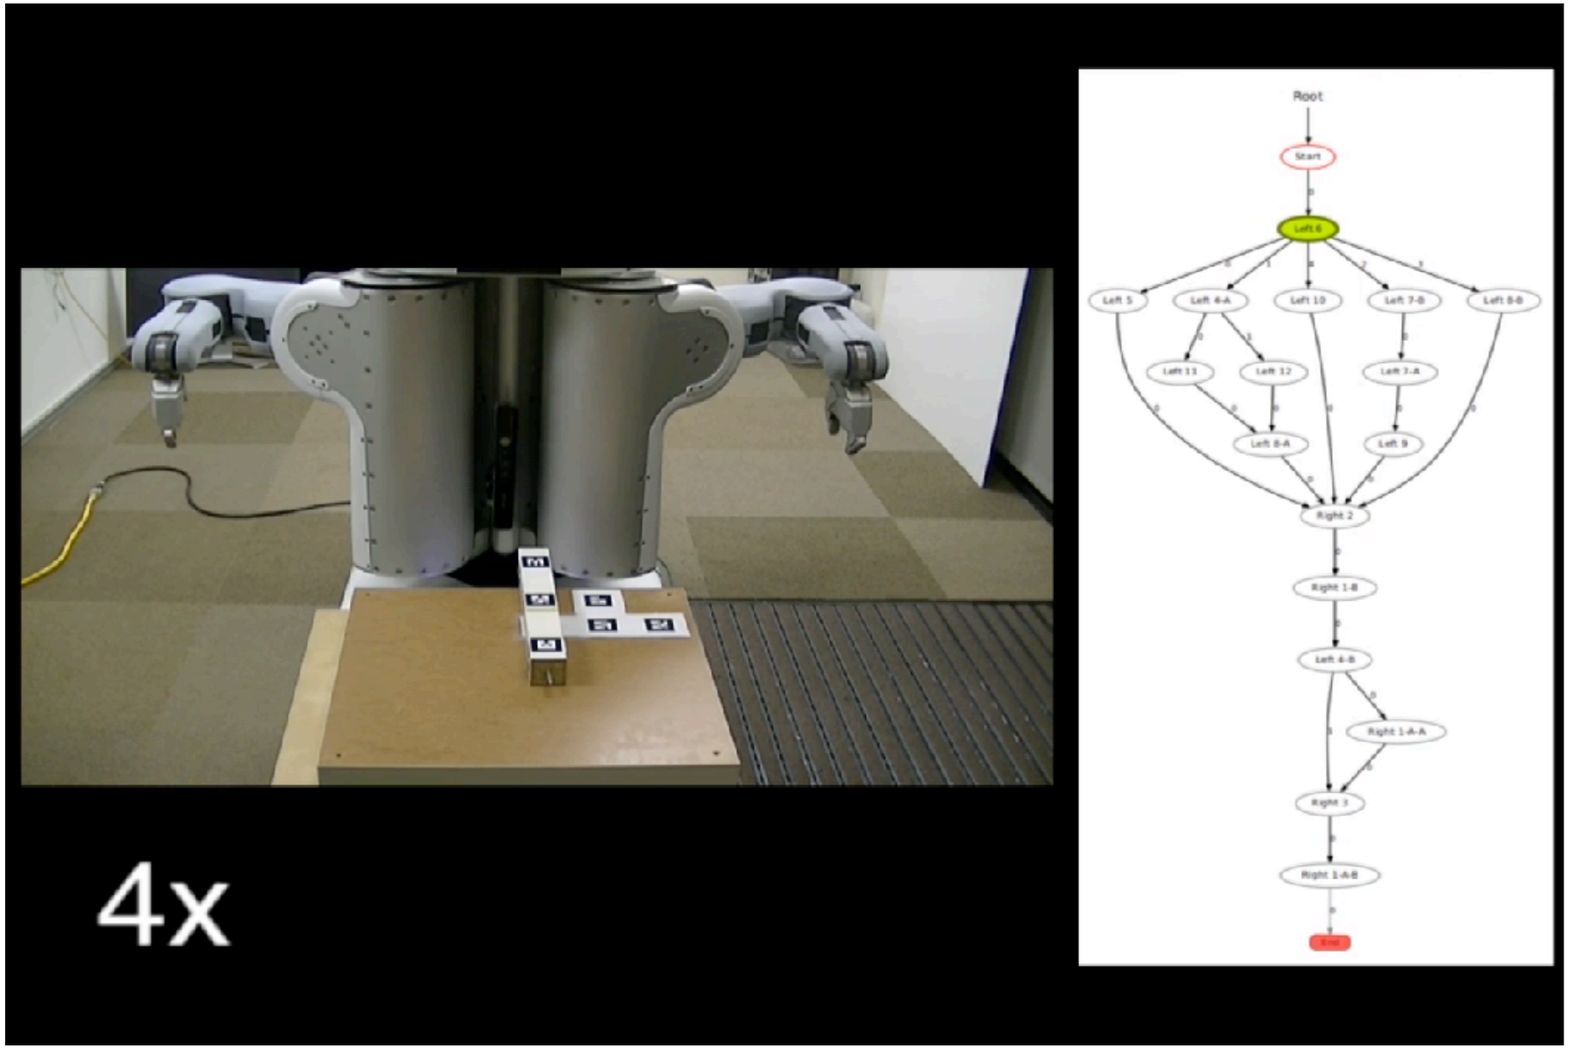

Full body control of humanoids

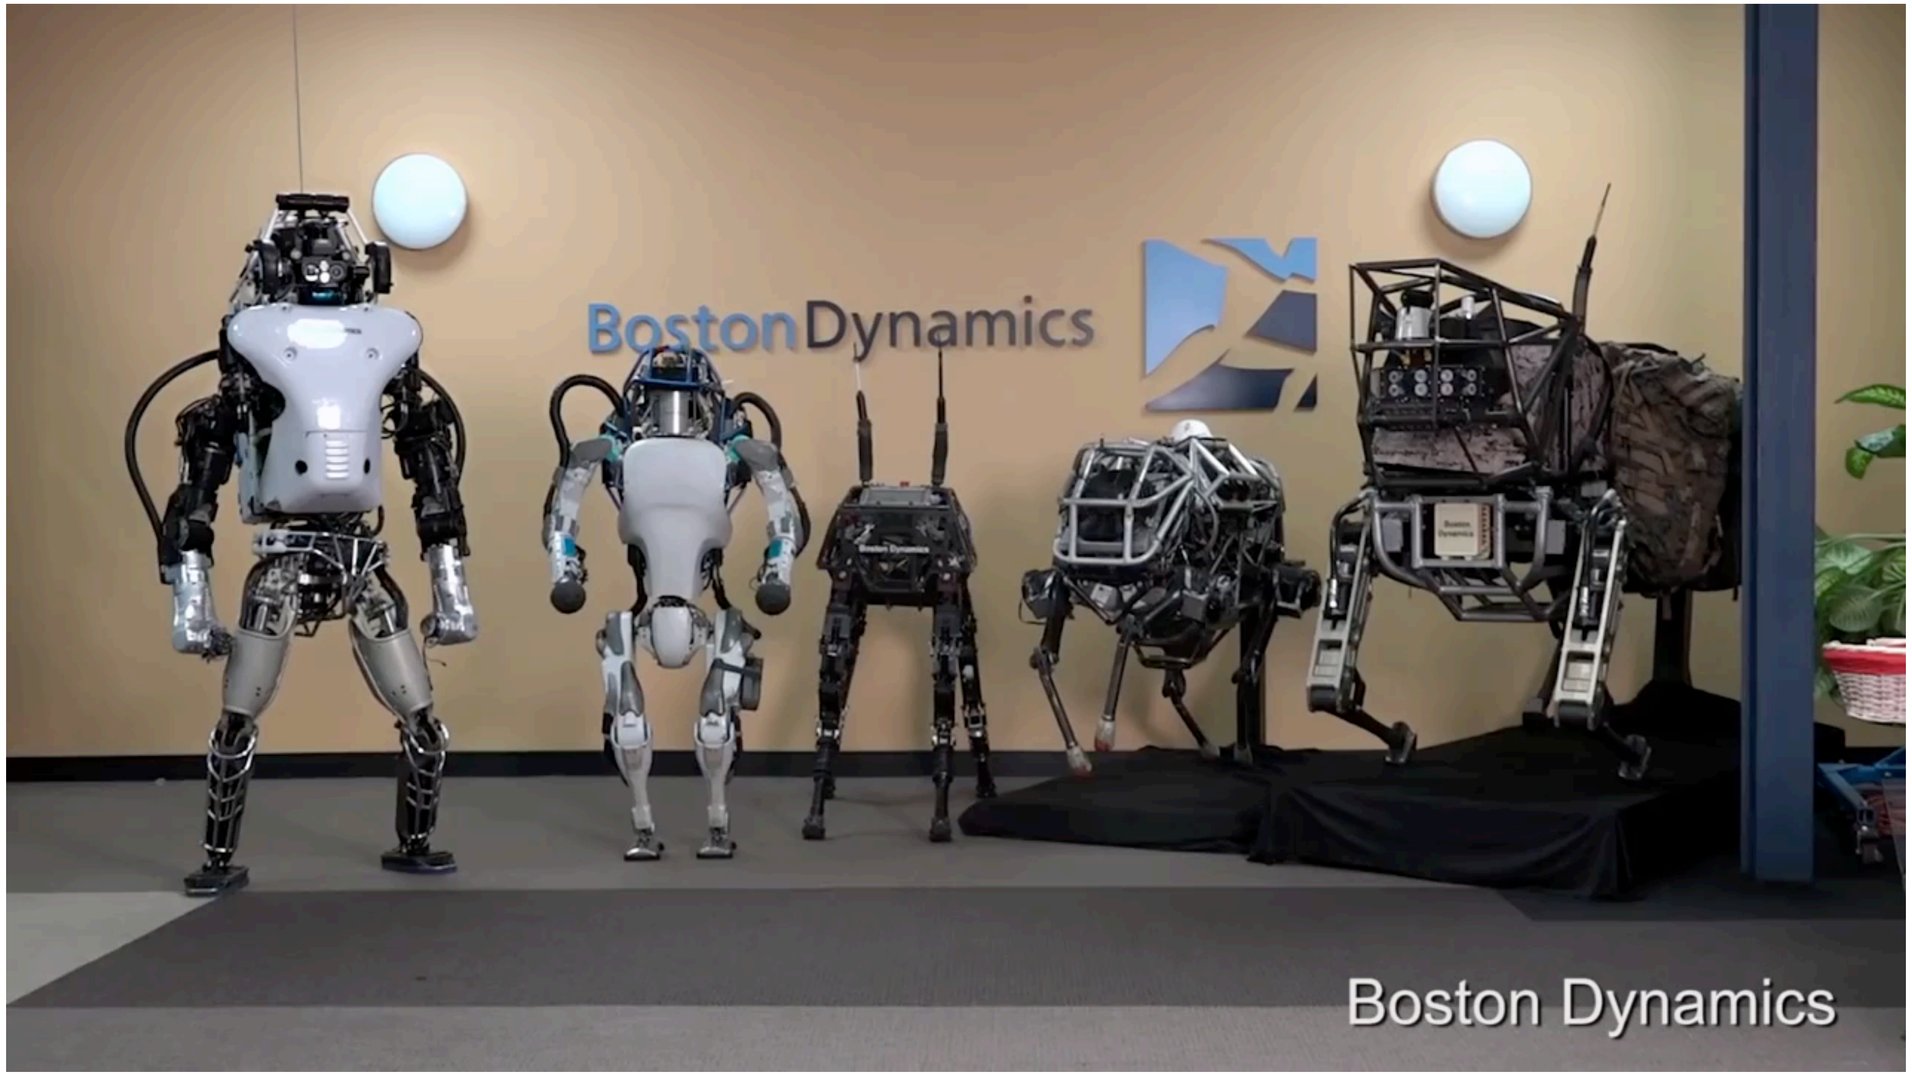

…but still a long way to go

Robotics

We will cover several topics relevant to robotics:

- Planning and search
- Reinforcement learning
- Time-series analysis
- State estimation and filtering

Logic

- Logical systems
  - Theorem provers
  - NASA fault diagnosis
  - Question answering
- Methods:
  - Deduction systems
  - Constraint satisfaction
  - Satisfiability solvers (huge advances!)

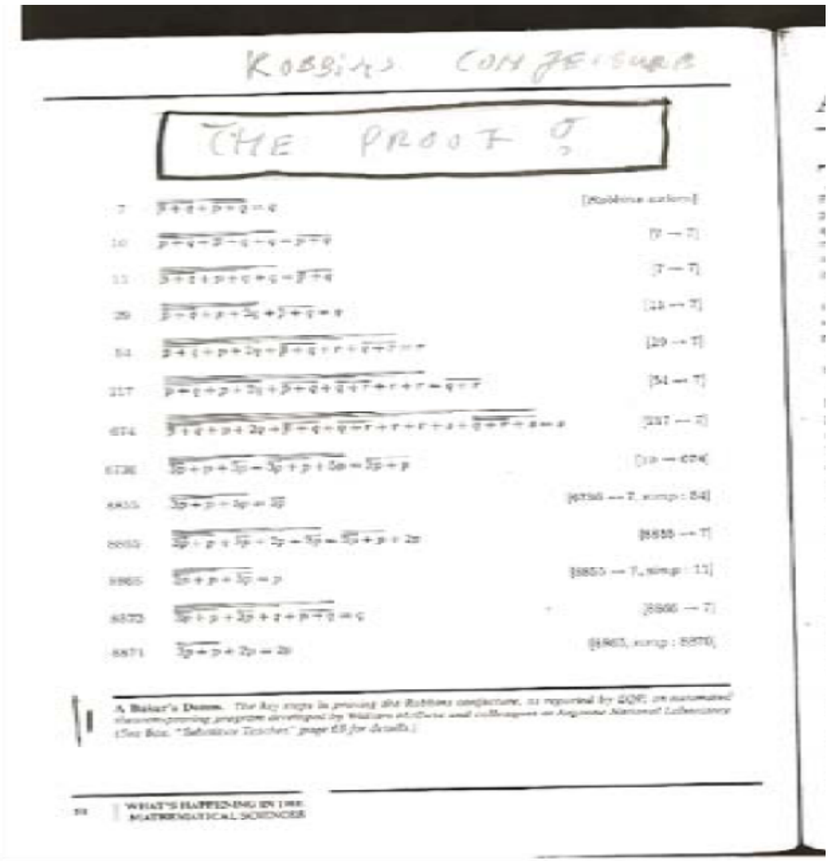

Image from Bart Selman

Game Playing

- Classic Moment: May, '97: Deep Blue vs. Kasparov
  - First match won against world champion
  - “Intelligent creative” play
  - 200 million board positions per second
  - Humans understood 99.9 of Deep Blue's moves
  - Can do about the same now with a PC cluster
- Open question:
  - How does human cognition deal with the search space explosion of chess?
  - Or: how can humans compete with computers at all??
- 1996: Kasparov Beats Deep Blue

  “I could feel --- I could smell --- a new kind of intelligence across the table.”
- 1997: Deep Blue Beats Kasparov

  “Deep Blue hasn't proven anything.”
- Huge game-playing advances recently, e.g. in Go!

Text from Bart Selman, image from IBM’s Deep Blue pages

Decision Making

- Applied AI involves many kinds of automation
  - Scheduling, e.g. airline routing, military
  - Route planning, e.g. Google maps
  - Medical diagnosis
  - Web search engines
  - Spam classifiers
  - Automated help desks
  - Fraud detection
  - Product recommendations
  - … Lots more!

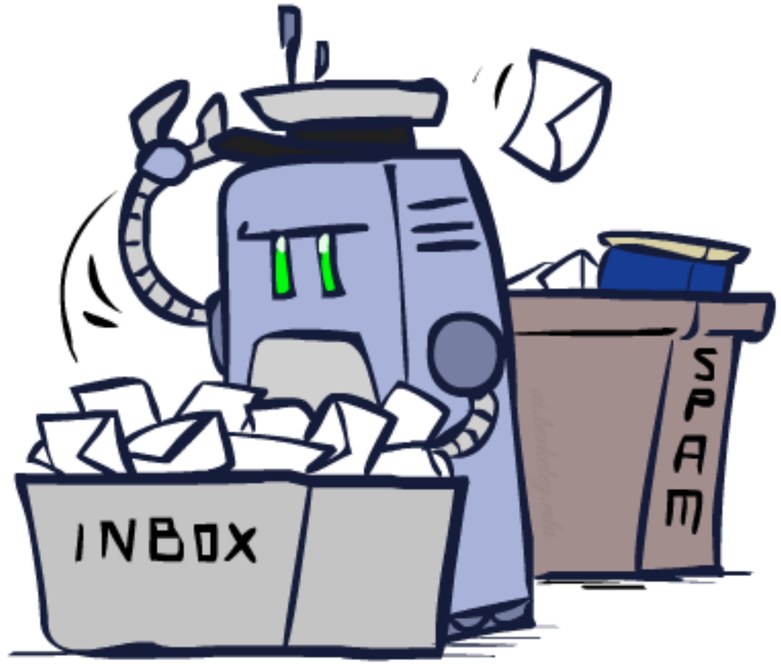

Course Topics

- Part I: Making Decisions
  - Fast search / planning
  - Constraint satisfaction
  - Adversarial and uncertain search
  - MDPs and Reinforcement learning
- Part II: Reasoning under Uncertainty
  - Bayes nets
  - Decision theory and value of information
  - Statistical Machine learning
- Throughout: Applications
  - Natural language, vision, robotics, games, …

Bringing it all together

- A goal of AI: Robust, fully autonomous agents in the real world

- Bottom-up metaphor:

  Russell, ‘95: *“Theoreticians can produce the AI equivalent of bricks, beams, and mortar with which AI architects can build the equivalent of cathedrals.”*

Bottom-up approach

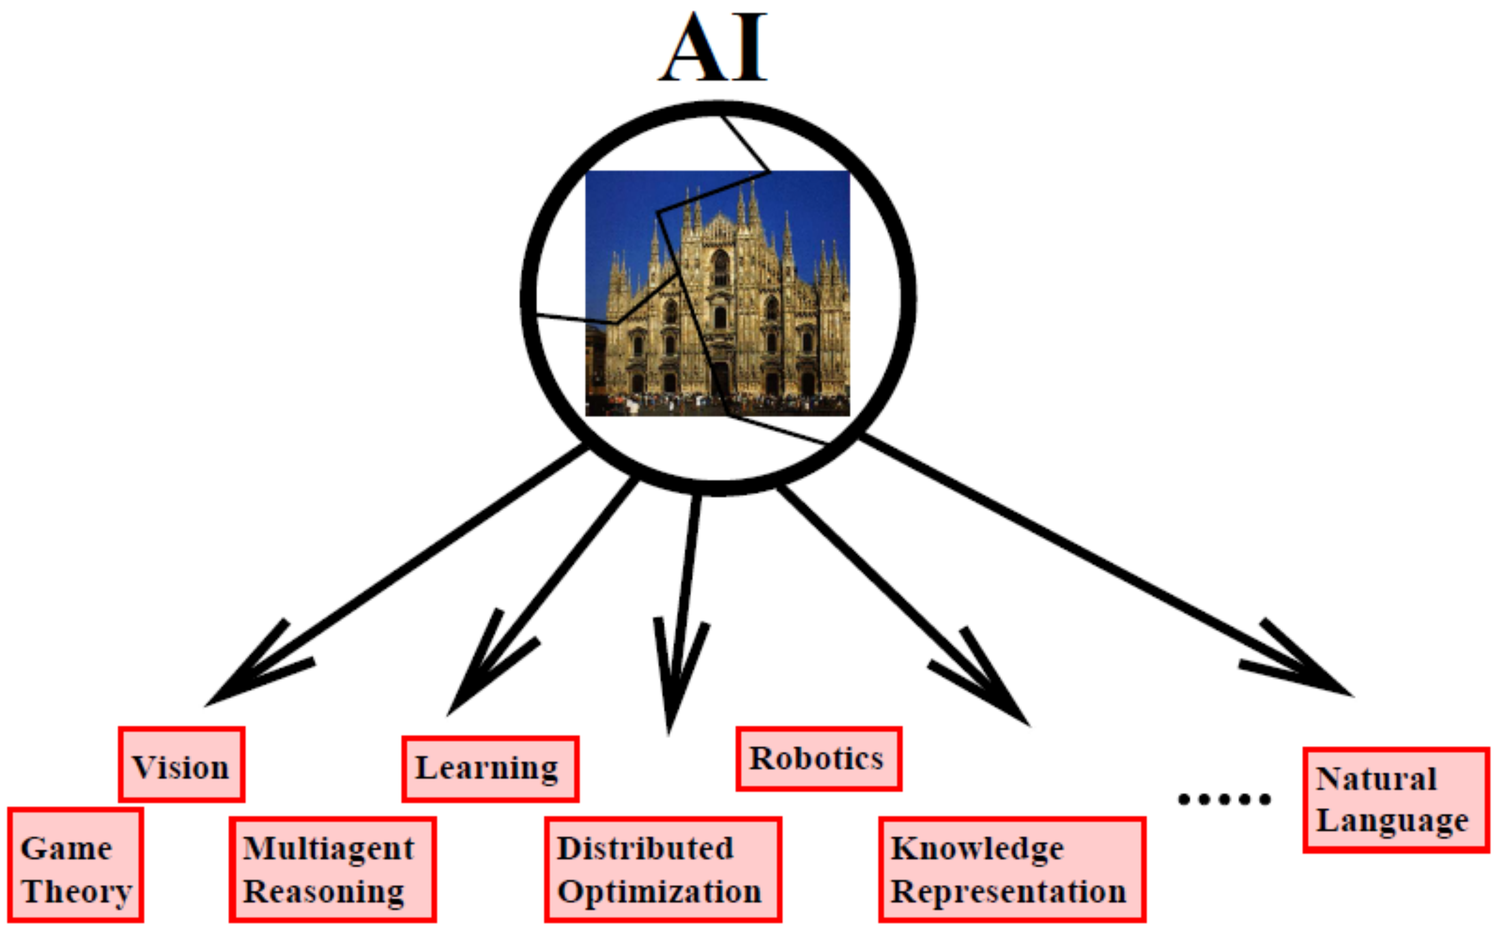

The bricks

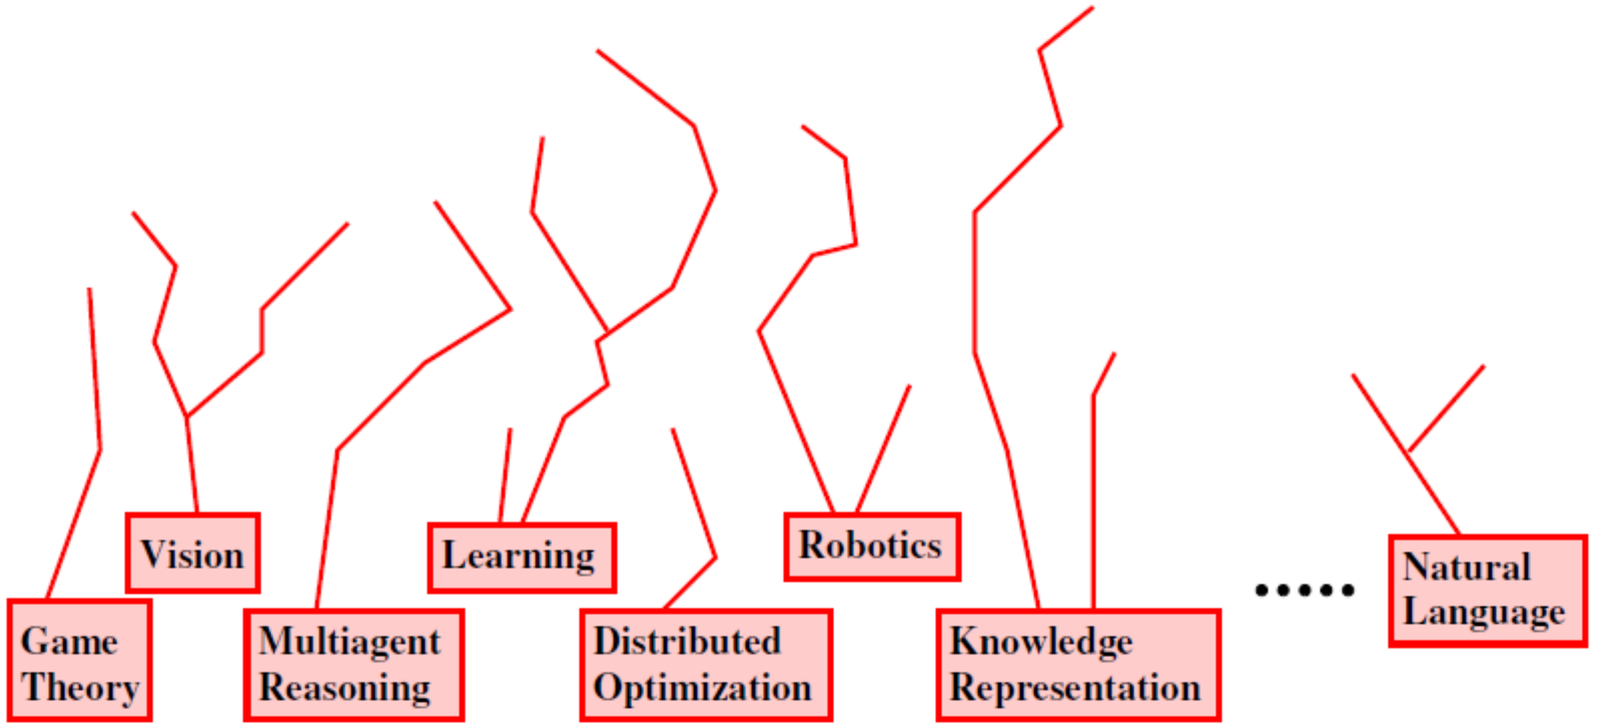
Game Theory
Vision
Multiagent Reasoning
Learning
Distributed Optimization
Robotics
Knowledge Representation
.....
Natural Language

The beams and mortar

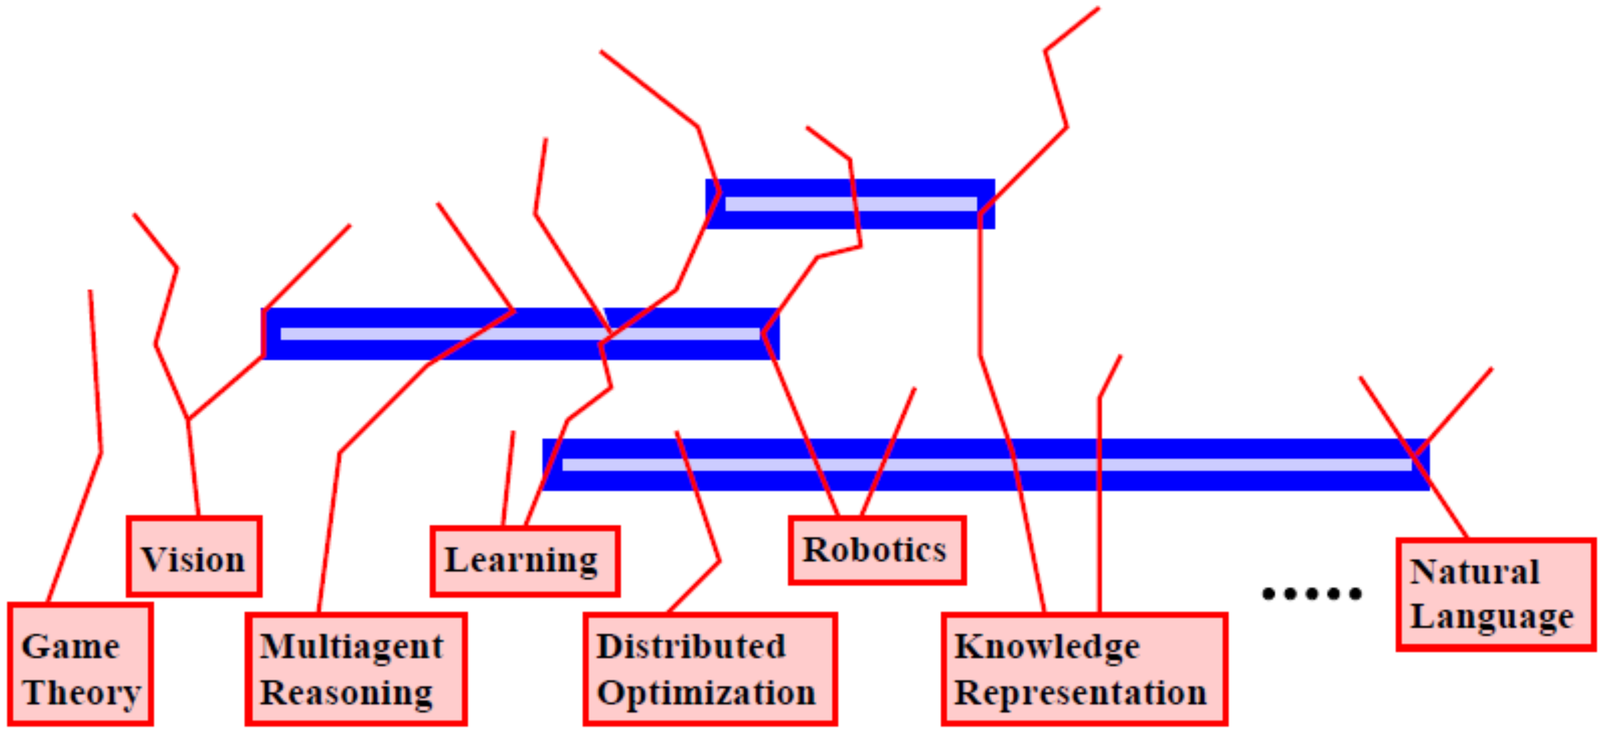

Towards a cathedral?

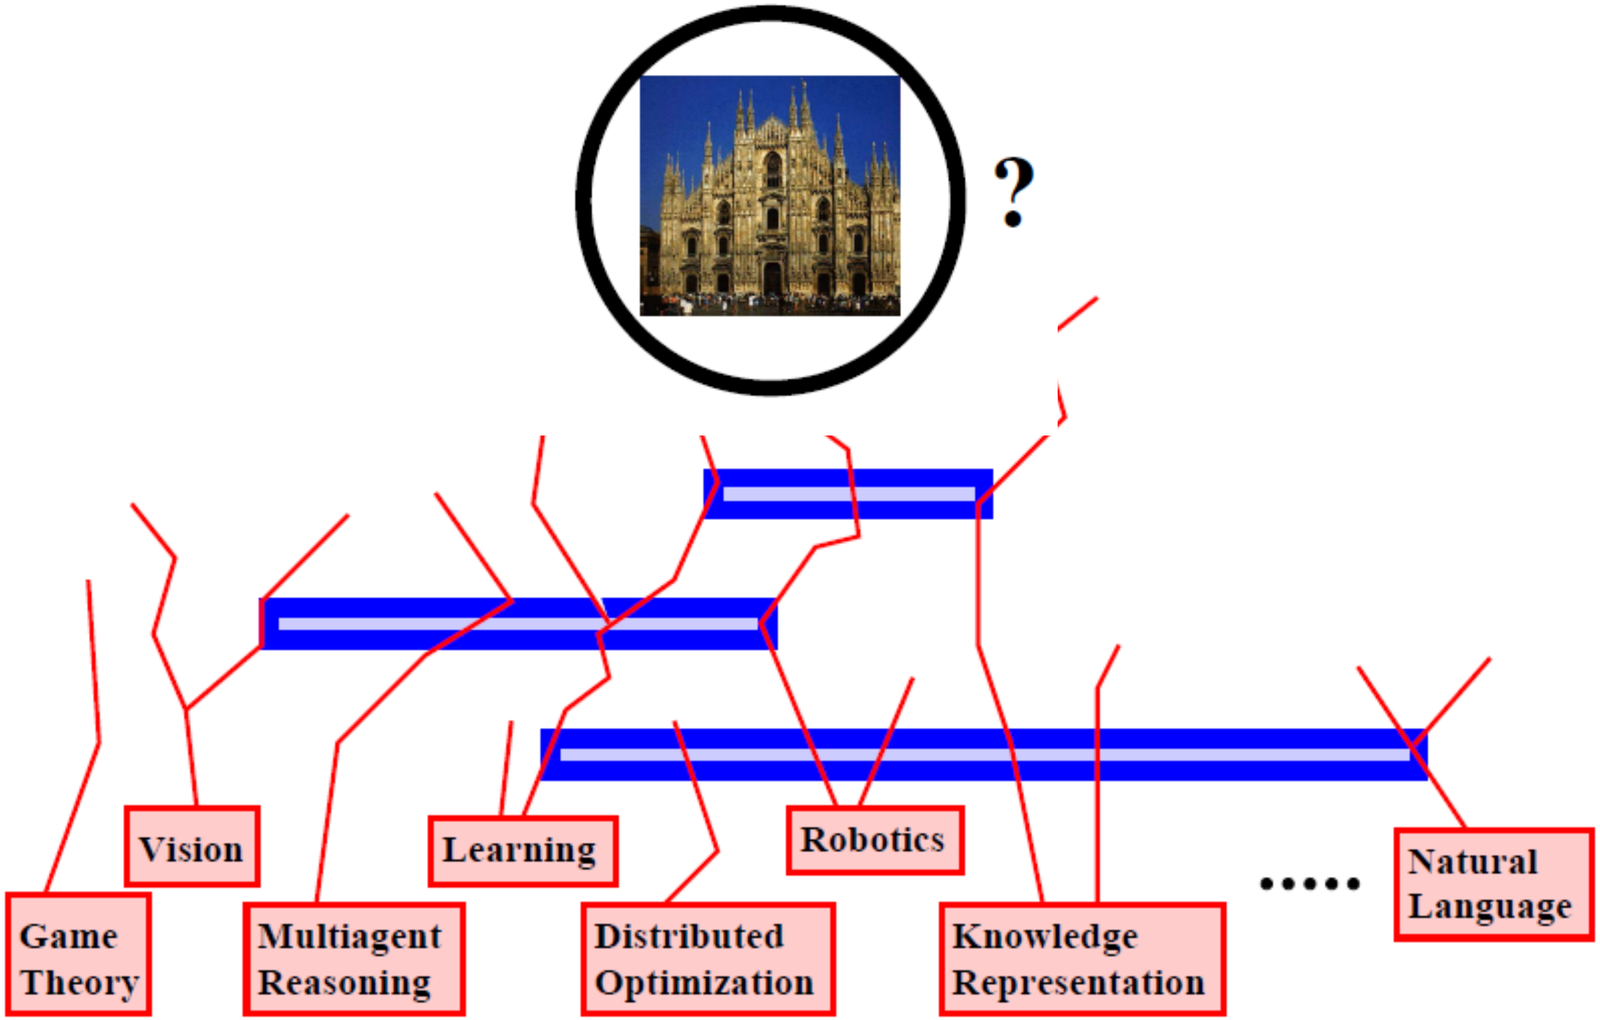

Or something else?

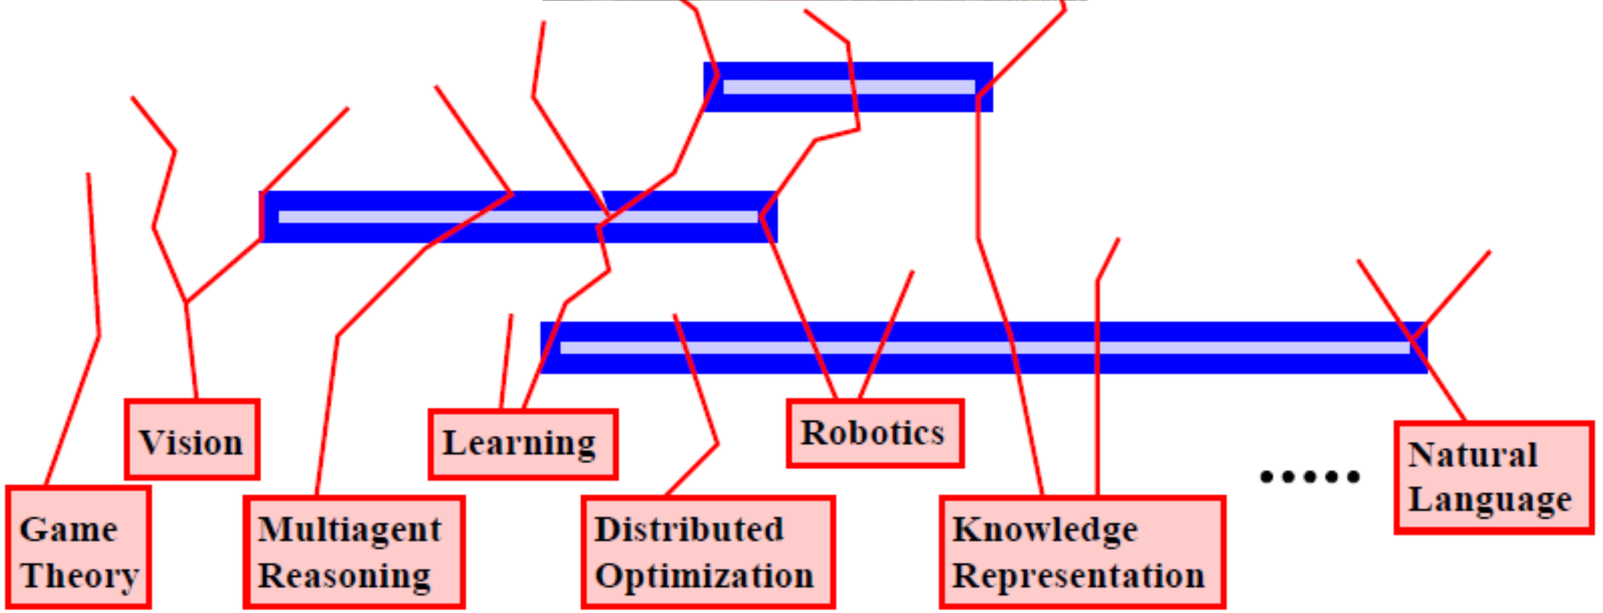

Top-down approach

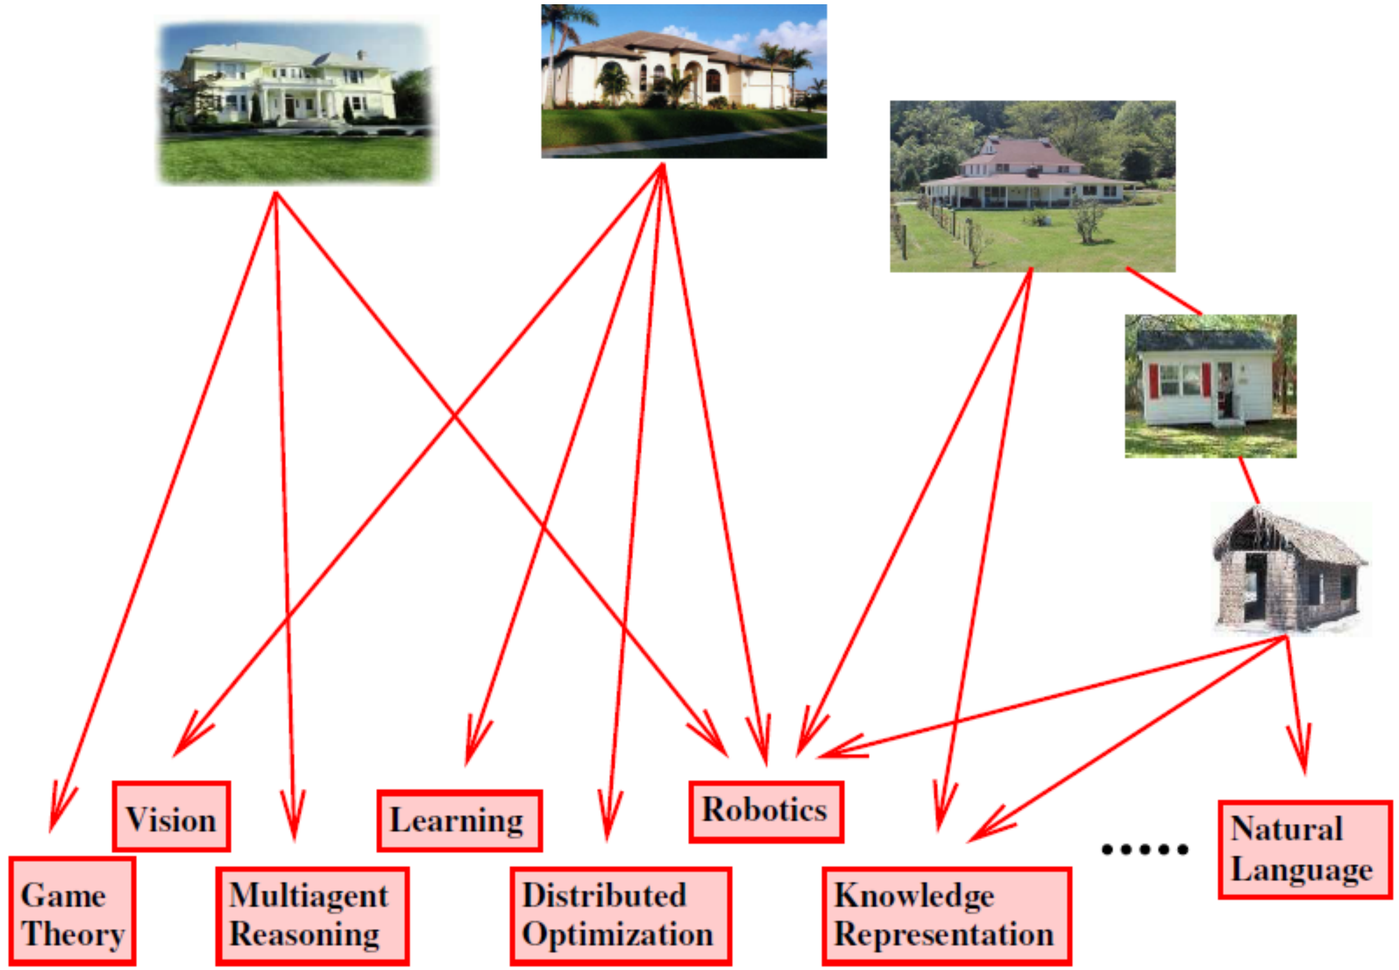

"Good problems . . . produce good science" [Cohen, '04]

Meeting in the middle

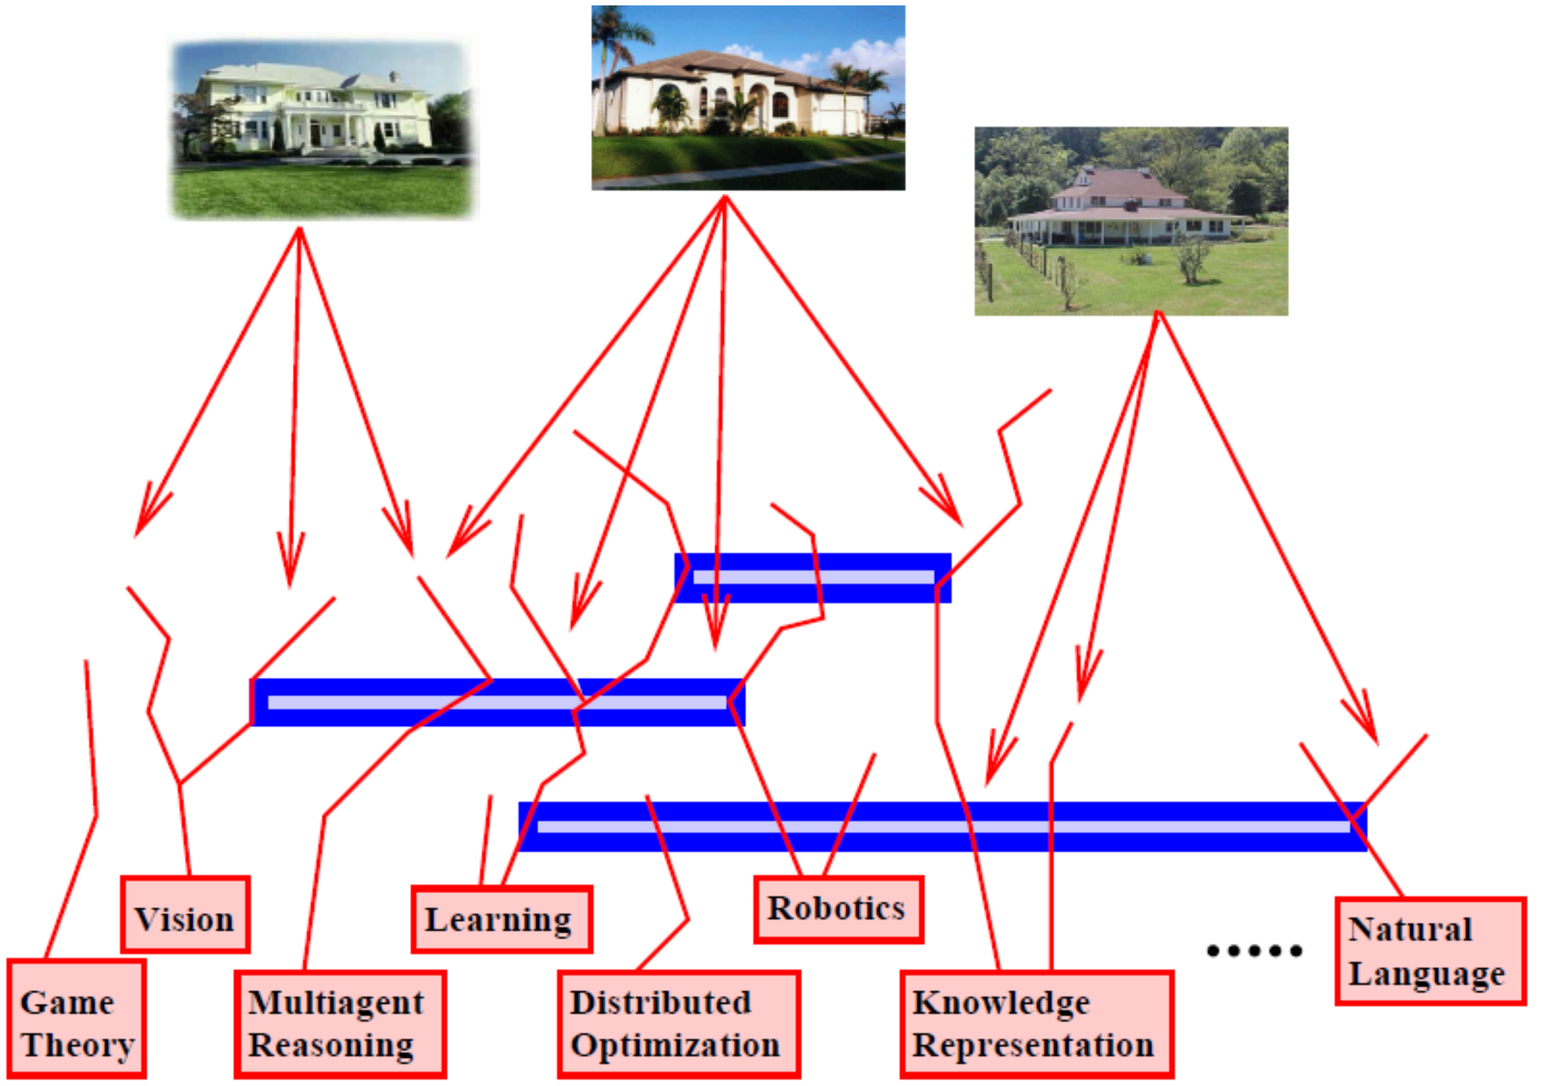

Good problems produce good science

Manned flight

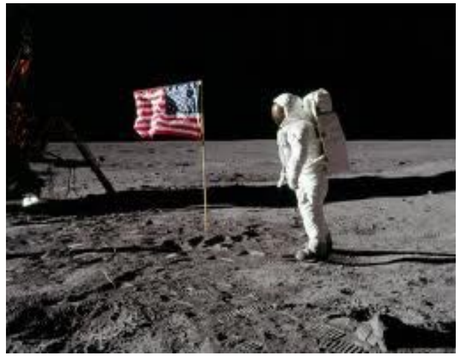

Apollo mission

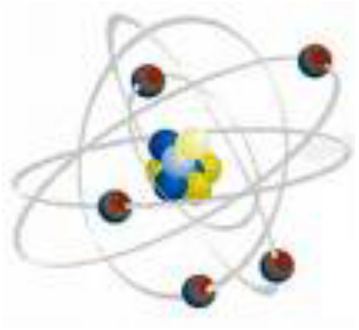

Manhattan project

Autonomous vehicles

RoboCup soccer

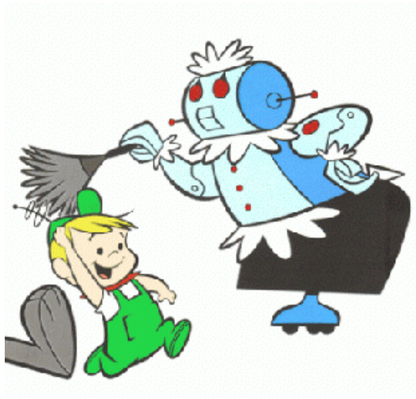

Assistive robots

Ethics, implications, hard questions…

- What happens when we achieve robust, fully autonomous agents in the real world? Utopia (Jetsons)? Distopia (Terminator)?
- Who is liable if a robot driver has an accident?
- Will machines surpass human intelligence (in all ways)?
- Would such machines have conscious existence? Rights?
- Can human minds exist indefinitely within machines (in principle)?

Designing Rational Agents

- An **agent** is an entity that *perceives* and *acts*.
- A **rational agent** selects actions that maximize its (expected) **utility**.
- Characteristics of the **performance measure, environment, actions, and sensing** dictate techniques for selecting rational actions
- **By then end of the course** you should understand:
  - General AI techniques for a variety of problem types
  - How to recognize when and how a new problem can be solved with an existing technique

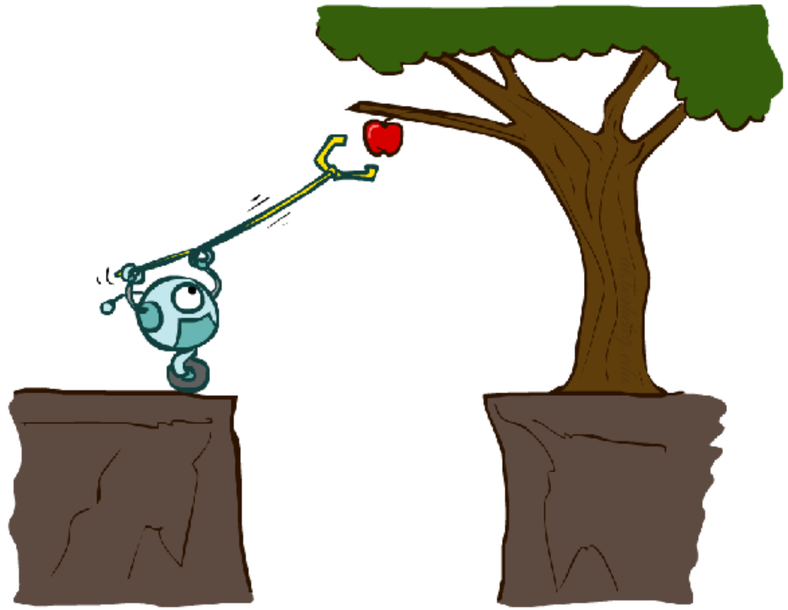

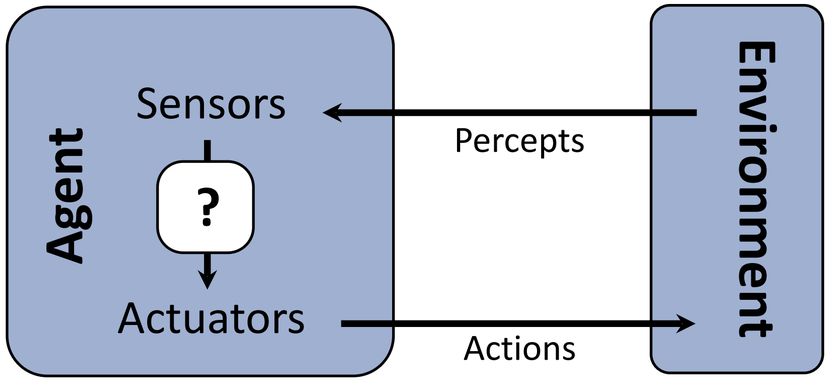

Pac-Man as an Agent

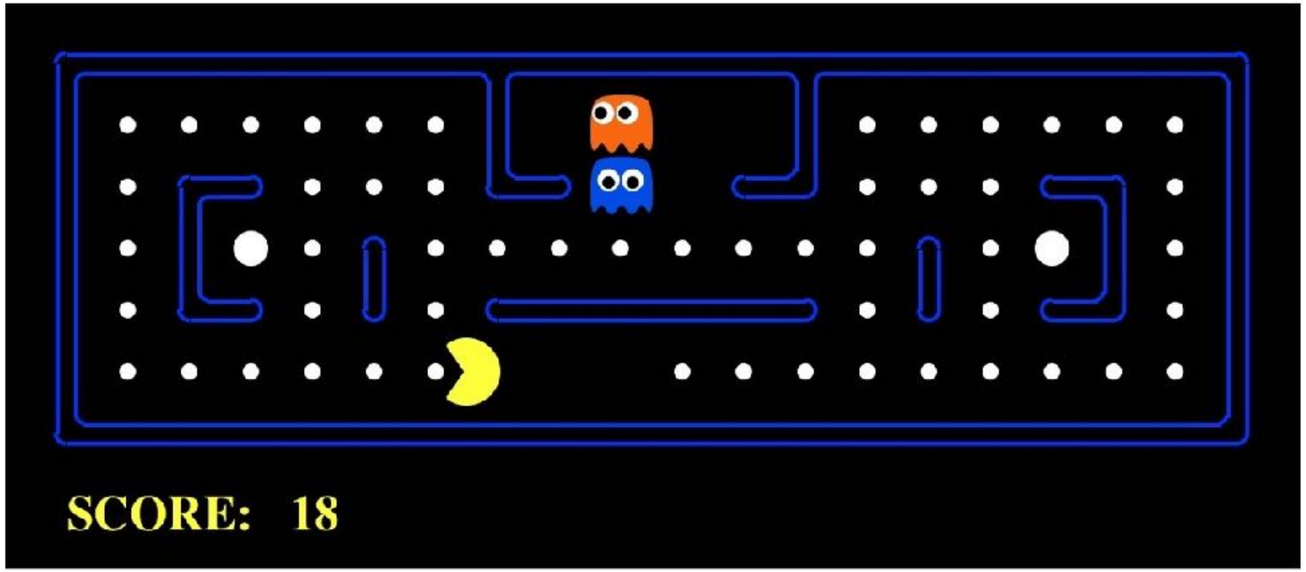

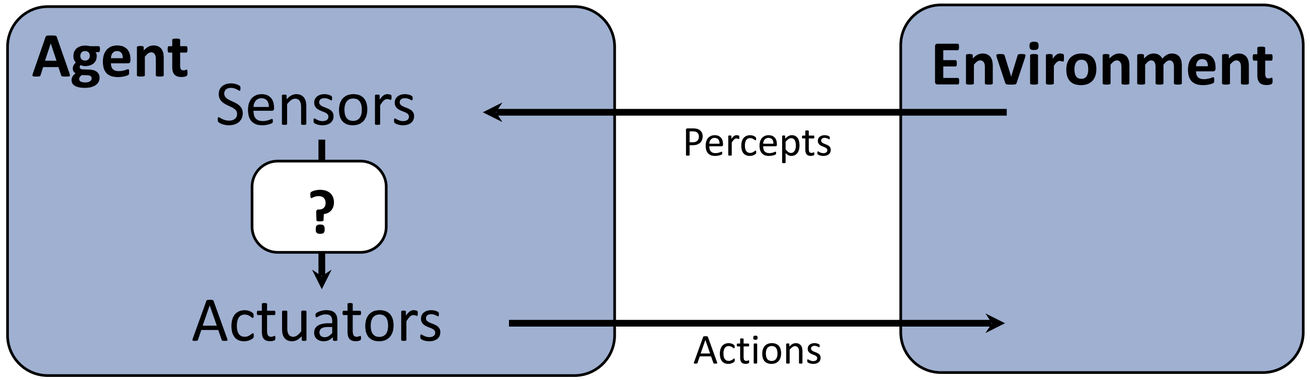

Pac-Man as an Agent

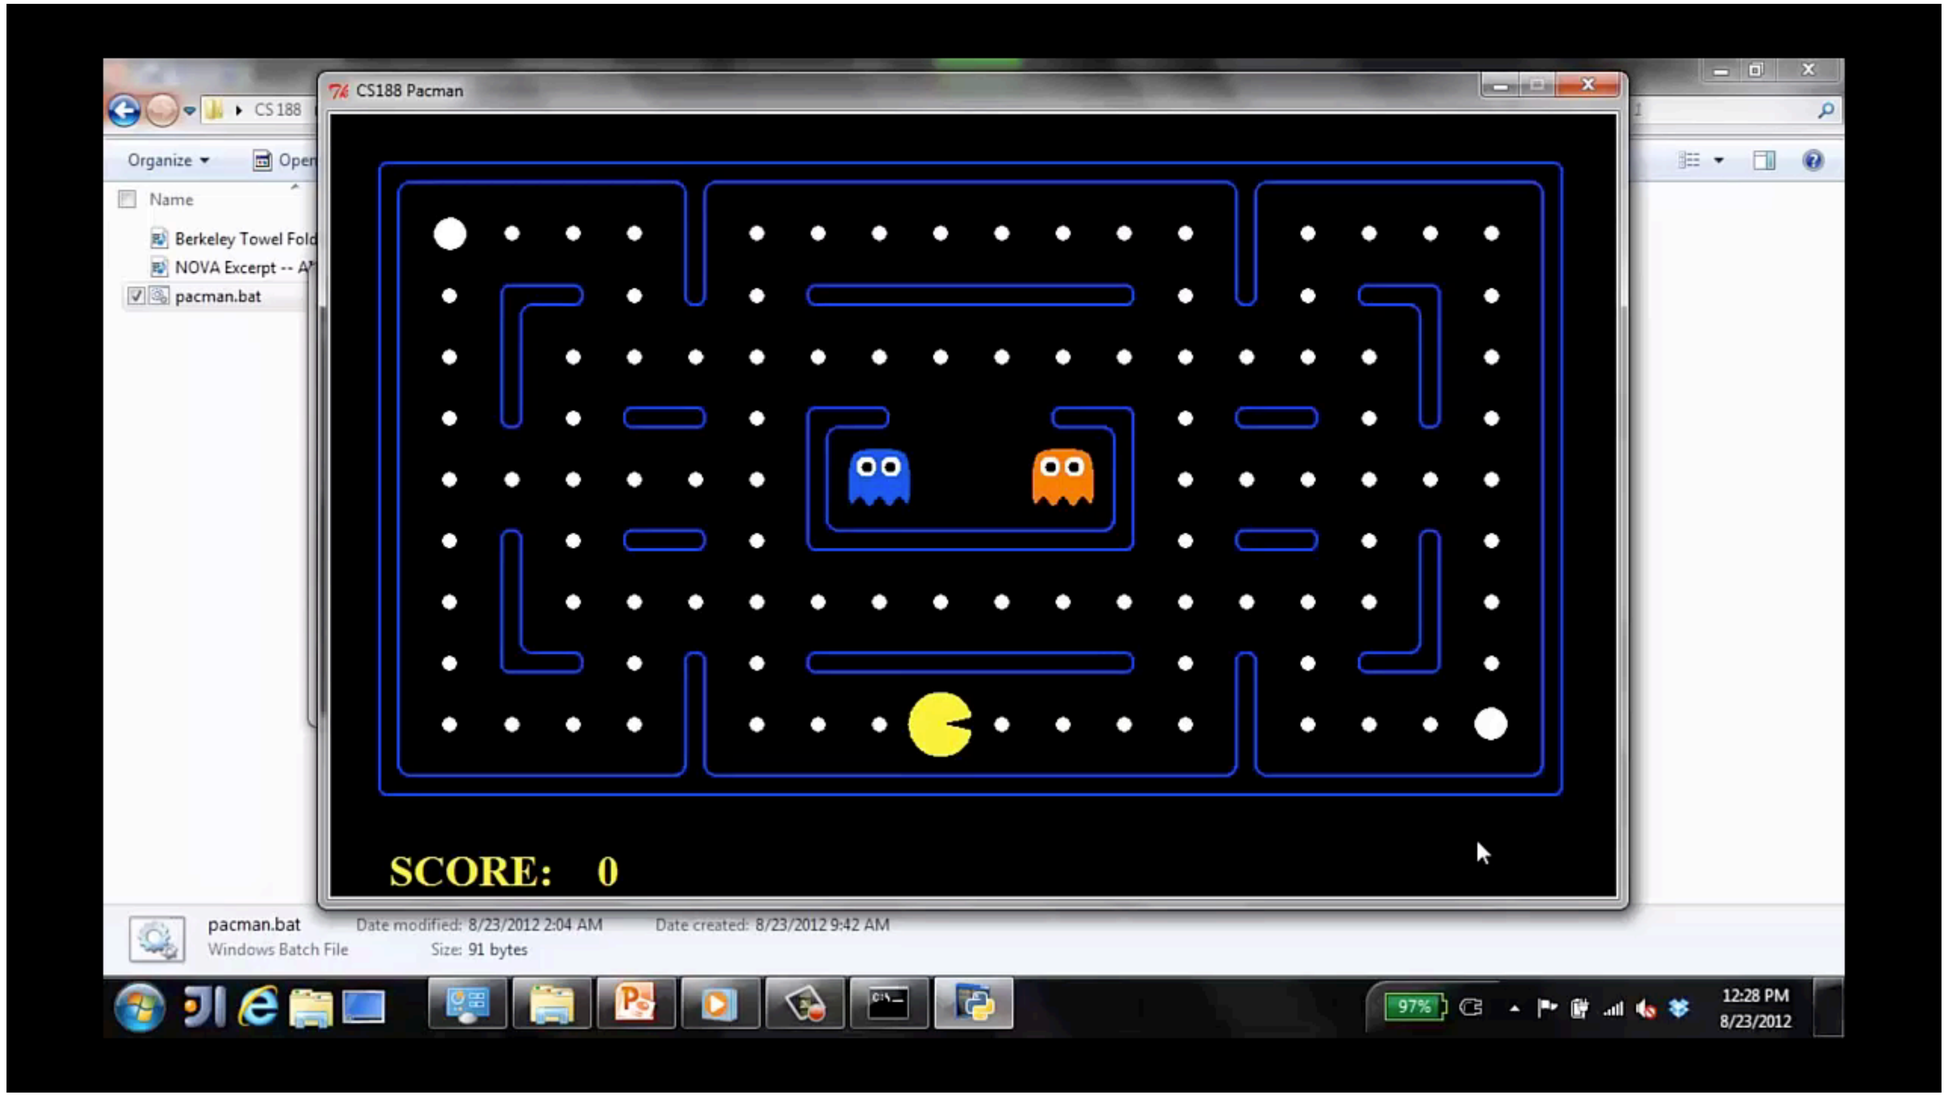

Reflex Agents

- Reflex agents:
  - Choose action based on current perceptions
  - May have memory or a model of the world's current state
  - Do not consider the future consequences of their actions
  - Often referred to as a "policy"
  - Consider how the world IS

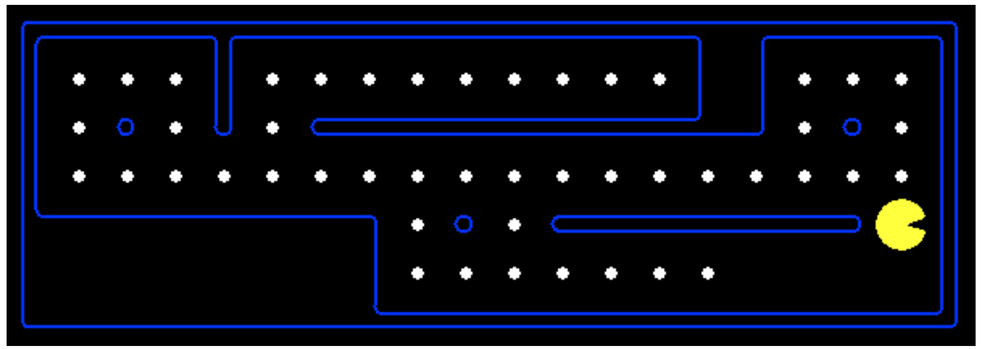

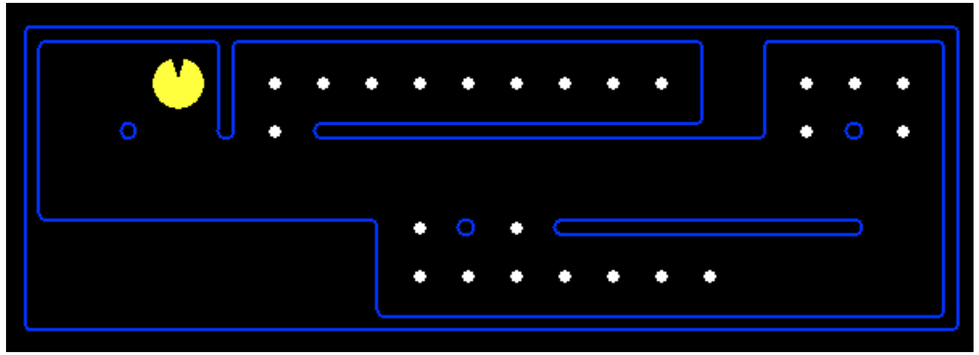

- Can a reflex agent be rational?

Planning Agents

- Plan ahead
- Ask “what if”
- Decisions based on (hypothesized) consequences of actions
- Must have a model of how the world evolves in response to actions
- Consider how the world WOULD BE

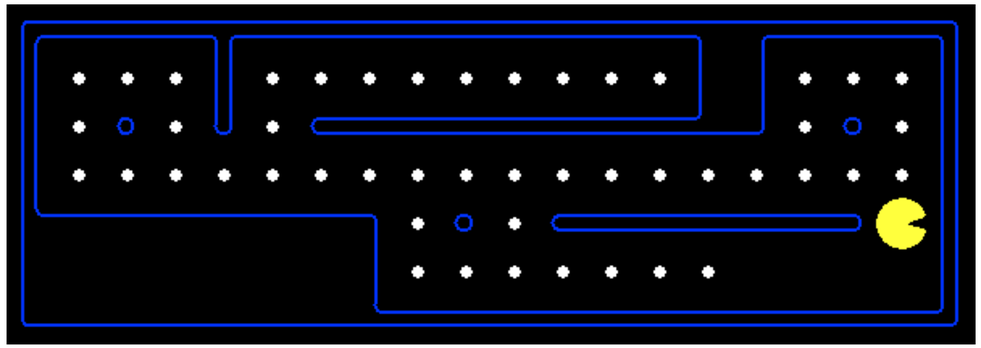

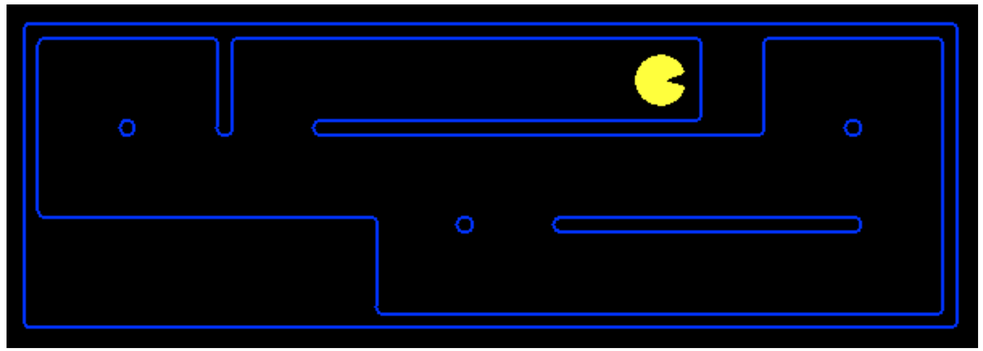

Quiz: Reflex or Planning?

Select which type of agent is described:

1.Pacman, where Pacman is programmed to move in the direction of the closest food pellet

2.Pacman, where Pacman is programmed to move in the direction of the closest food pellet, unless there is a ghost in that direction that is less than 3 steps away.

3.A navigation system that first considers all possible routes to the destination, then selects the shortest route.

Properties of task environment

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

Fully observable vs. partially observable

- Fully observable: agent's sensors give it access to complete state of the environment at all times
- Can be partially observable due to noisy and inaccurate sensors, or because parts of the state are simply missing from the sensor data
- Example: Perfect GPS vs noisy pose estimation
- Example: IKEA assembly while blindfolded

Almost everything in the real world is partially observable

Single agent vs. multi-agent

- Not multi-agent if other agents can be considered part of the environment
- Only considered to be multi-agent if the agents are maximizing a performance metric that depends on other agents' behavior
- Single agent example: Pacman with randomly moving ghosts
- Multi-agent example: Pacman with ghosts that use a planner to follow him

Deterministic vs. stochastic

- Deterministic: next state of environment is completely determined by the current state and the action executed by the agent
- Stochastic: actions have probabilistic outcomes
- Strongly related to partial observability — most apparent stochasticity results from partial observation of a deterministic system
- Example: Coin flip

Episodic vs. sequential

- Episodic: agent's experience is divided into atomic episodes; single percept and action each episode, and next episode does not depend on the actions taken in previous ones
- Sequential: current decision could affect all future decisions
- Example: Image classification vs. Chess

Static vs. dynamic

- Static: The environment does not change, except for actions taken by the agent
- Dynamic: The environment continuously evolves without input from the agent
- Example: Checkers vs. pacman

Discrete vs. continuous

- Refers to state, how time is handled, and the percepts and actions of the agent
- Discrete: Finite category or integer
- Continuous: Real-valued quantity
- Percept example: Discrete colors vs. color spectrum
- Action example: Chess moves vs. steering angle
- State example: Pacman’s (x,y) position vs. robot’s joint angles

Known vs. unknown

- Agent's state of knowledge about the "rules of the game" / "laws of physics"
- Known environment: the outcomes for all actions are given
- Unknown: agent has to learn how it works to make good decisions
- Possible to be partially observable but known (solitaire)
- Possible to be fully observable but unknown (video game)