CS 188: Artificial Intelligence

Introduction

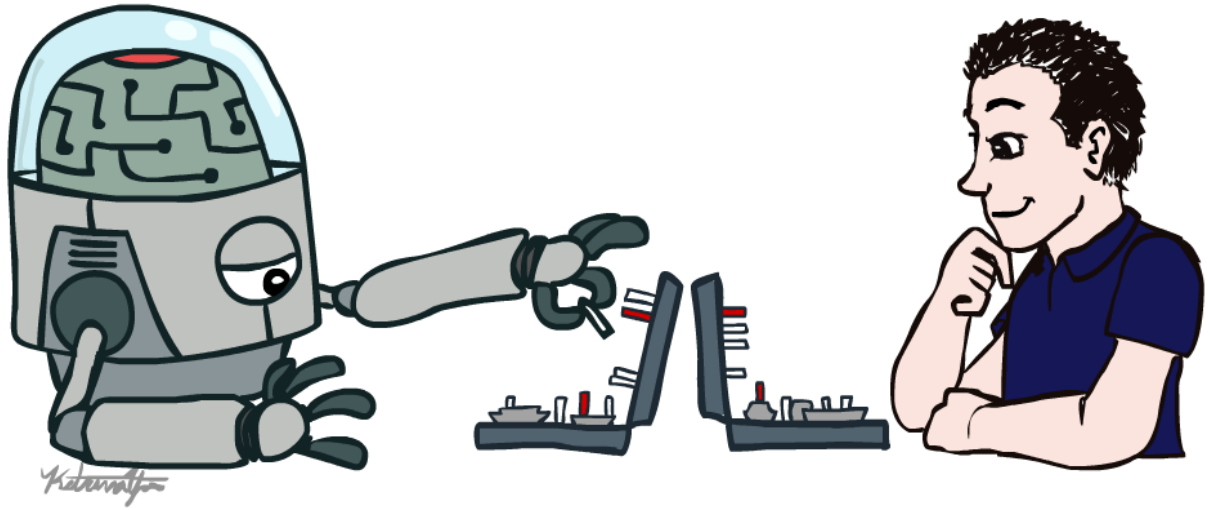

Fall 2022

University of California, Berkeley

[These slides were created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley (ai.berkeley.edu).]

First Half of Today: Intro and Logistics

- Staff introductions: Igor, Peyrin, and course staff
- Course logistics
  - Lectures, discussions, office hours, and exams
  - Resources and communication platforms
  - Collaboration and academic honesty
  - DSP and extenuating circumstances
  - Stress management and mental health

Staff Introductions: Igor Mordatch (he/him)

- Research scientist at Google Brain, with research interests in machine learning, robotics, and multi-agent systems
- Previously:
  - Research scientist at OpenAI
  - Postdoctoral fellow at UC Berkeley
  - PhD from University of Washington
  - Visiting researcher at Stanford and Pixar
- Co-organized OpenAI Scholars mentorship program
- Mentor and tutor for AI4All, Google CS Research Mentorship Program, and Girls Inc.

Staff Introductions: Peyrin (he/him)

- Did my undergrad at Berkeley (2017-2021)
  - TA for 10 semesters (8x CS 161, 3x CS 61C, 1x CS 188)
  - Also been on staff for CS 61A, EE 16A, EE 16B
- Did a 5th year MS at Berkeley (2021-2022)
  - Research focus: computer science education
  - Advisors: Nicholas Weaver and Dan Garcia
- First-year lecturer in EECS
  - I'm paid exclusively to care about students and staff
  - First time teaching a non-summer class as instructor, so your feedback/advice/complaints are appreciated!
- Please call me "Peyrin"!
  - No "professor", "Mr.", "sir", "doctor", etc. I'm not paid enough for that.

Actual real picture of me.

Our talented course staff!

Evgeny Pobachienko
he/him

Jason Wang
he/him

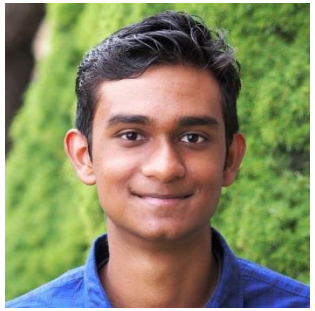

Ajay Sridhar
he/him

Alina Trinh
she/her

Andrew Wang
he/him

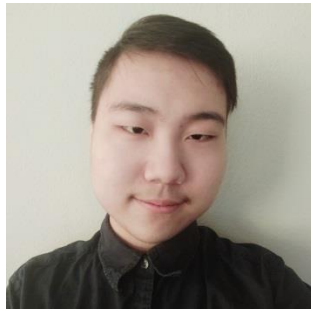

Anthony Han
he/him

Cham Yao
he/him

Joy Liu
she/her

Micah Carroll
he/him

Nitish Dashora
he/him

Our talented course staff!

Raiymbek Akshulakov
he/him

Rishi Parikh
he/him

Shizhan Zhu
he/him

Sid Ijju
he/him

Jerry Sun
he/him

Lucas Sosa
he/him

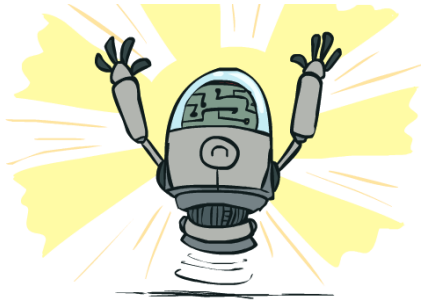

CS 188 Bots
any/all

Aviral Kumar (*he/him*)

Sherry Yang

Andrew Qin

Hrish Leen (*he/him*)

Enrollment

- Course staff does not control enrollment; we have to follow department policy
  - Only CS majors will be able to enroll this fall
- Waitlisted students should be enrolled once reserved seats disappear
  - Reserved seats should have disappeared yesterday
- Concurrent enrollment students are awaiting approval from the department to proceed
  - May need to wait until second week of classes
  - We'll add you to class platforms for now

Course Structure: Lectures

- You are here!
- Tuesday/Thursday, 5:00–6:30 PM PT
- Attendance is not taken
- You can attend:
  - In-person in Wheeler 150
  - Remotely over Zoom (we'll try our best to livestream)
  - Asynchronously by watching recordings (posted on website)

Course Structure: Discussions

- We offer three types of discussions
  - Regular discussions
  - Exam prep discussions
  - Extended-time discussions
  - We'll try to make recordings, but no promises
- Discussion schedule available on website
  - Discussions start next week (August 31)
- You can attend any discussion section you want (no need to enroll in a section)
  - Attendance is not taken, but a bit of extra credit for active participation

Course Structure: Office Hours

- Join in-person or remotely to talk to staff about content, ask questions on assignments, or raise any concerns you have
- Schedule and queue available on website
  - Office hours start next week

Course Structure: Exams

- Save the dates!
  - Midterm: Wednesday, October 12, 7:00–9:00 PM PT
  - Final exam: Thursday, December 15, 11:30 AM–2:30 PM PT
- If you can’t make it:
  - We’ll offer remote exams at the listed time
  - We’ll offer an in-person-only alternate exam right after the listed time
- More logistics closer to the exam

Resources

- Course website: https://inst.eecs.berkeley.edu/~cs188/fa22
  - All resources (slides, notes, recordings, assignments, etc.) posted here
- Ed: Discussion forum replacing Piazza
- Staff email for private concerns: cs188@berkeley.edu
  - Making a private post on Ed is easier/faster
- Gradescope: Submit assignments here

Grading Structure

- Projects (25%)
  - Python programming assignments, autograded
  - You can optionally work with a partner
  - Reduced credit for submitting late, unless you have an extension
- Homework (20%)
  - Electronic homework: Autograded on Gradescope
  - Written homework: One question per week, graded by TAs on correctness
  - Submit individually (but feel free to discuss with others)
  - No late submissions, unless you have an extension
- Midterm (20%), Final Exam (35%)

Extensions and Accommodations

- We'll drop your lowest homework score
- You have 5 slip days to use across the projects
  - See course policies page for details on how they work
- If you ever need an extension, please request one!
  - We're here to support you, and we understand that life happens.
  - Extension form will be posted on the website

DSP

- Disabled Students' Program (DSP)
  - There's a variety of accommodations UC Berkeley can help us set up for you in this class
  - https://dsp.berkeley.edu/
- Are you facing barriers in school due to a disability?
  - Apply to DSP!
  - We maintain proper access controls on this information: Only instructors, course managers, head TAs, and logistics TAs can access any DSP-related info
- Our goal is to teach you the material in our course. The more accessible we can make it, the better.

Collaboration and Academic Dishonesty

- We're here to help! There are plenty of staff and resources available for you
  - You can always talk to a staff member if you're feeling stressed or tempted to cheat
  - Collaboration on homework is okay, but please cite collaborators
  - Do not post solutions online or share with others!
- Academic dishonesty policies
  - Reported to Center of Student Conduct
  - Negative points on assignments, and/or F in the class

Stress Management and Mental Health

- **Your health is more important than this course**
- If you feel overwhelmed, there are options
  - Academically: Ask on Ed, talk to staff in office hours, set up a meeting with staff to make a plan for your success this semester
  - Non-academic:
    - Counselling and Psychological Services (CAPS) has multiple free, confidential services
      - Casual consultations: https://uhs.berkeley.edu/counseling/lets-talk
      - Crisis management: https://uhs.berkeley.edu/counseling/urgent
    - Check out UHS's resources: https://uhs.berkeley.edu/health-topics/mental-health

Second Half of Today: What is AI?

- What is artificial intelligence?
- What can AI do?
  - What should we worry about?
  - What can we do about those things?
  - What should we not worry about?
- What is this course?

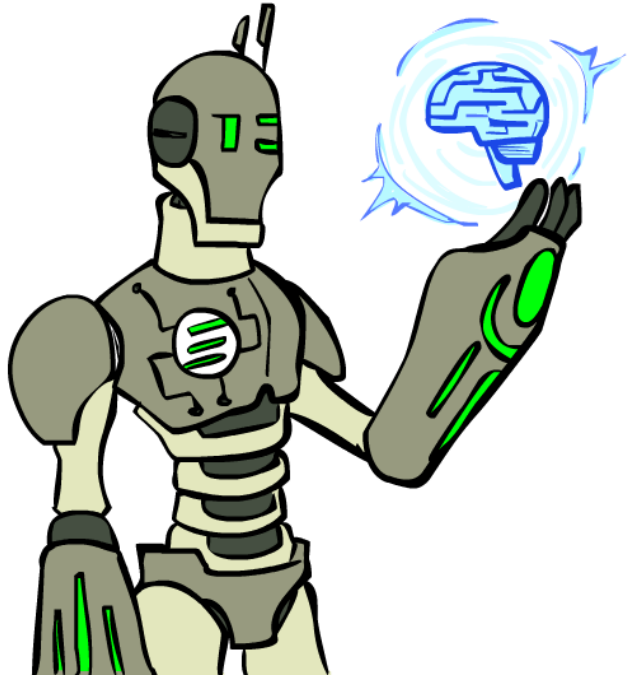

Sci-Fi AI?

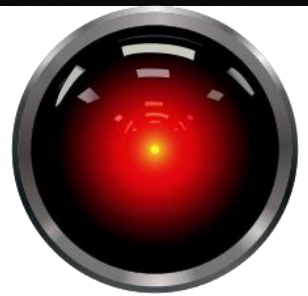

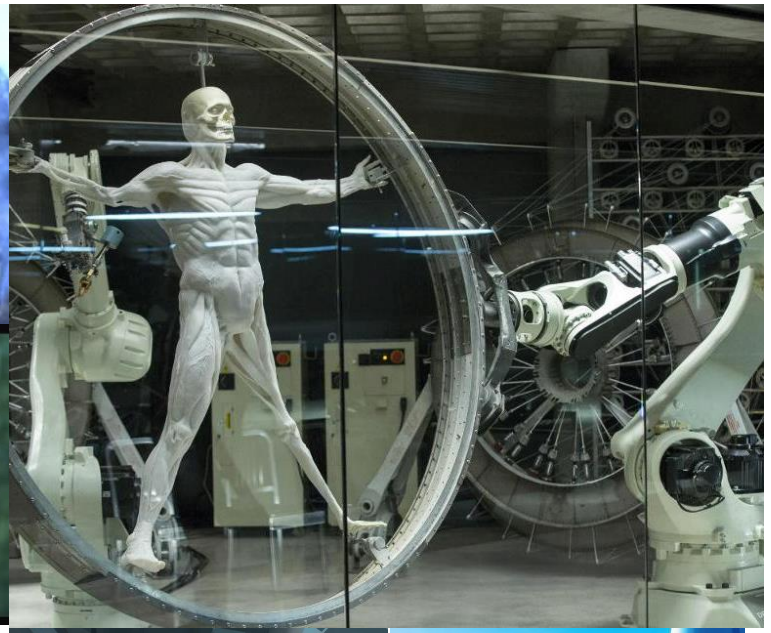

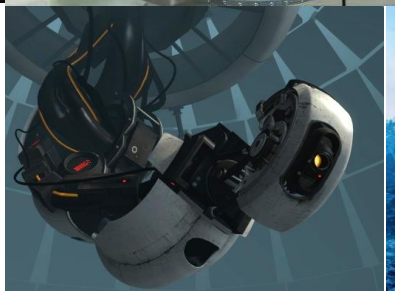

Real-World AI?

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**

A better title for this course would be:

Computational Rationality

Maximize Your Expected Utility

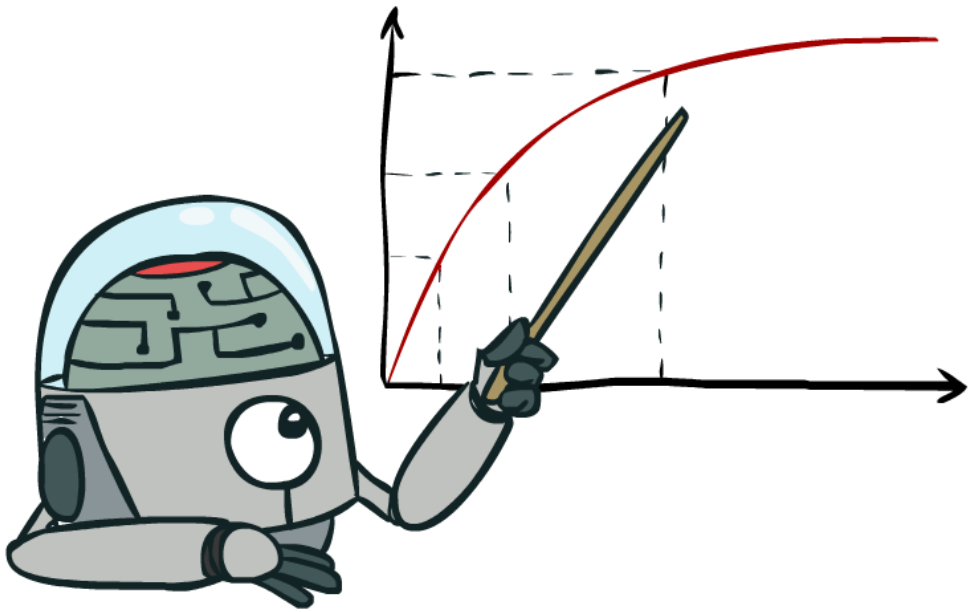

What About the Brain?

- 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

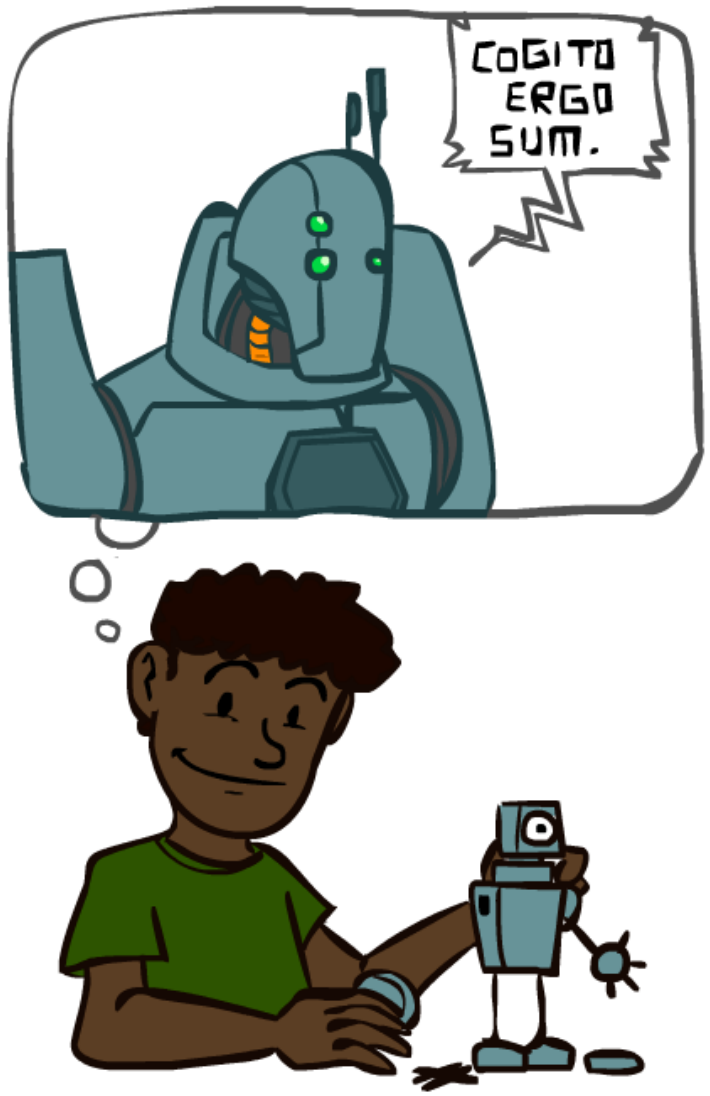

A (Short) History of AI

- 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

A (Short) History of AI

- 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”?
  - 1996: Kasparov defeats Deep Blue at chess
  - 1997: Deep Blue defeats Kasparov at chess

“I could feel --- I could smell --- a new kind of intelligence across the table.” ~Kasparov

A (Short) History of AI

- 2000—: Where are we now?
  - Big data, big compute, neural networks
  - Some re-unification of sub-fields
  - AI used in many industries
  - Chess engines running on ordinary laptops can defeat the world's best chess players
  - 2011: IBM's Watson defeats Ken Jennings and Brad Rutter at Jeopardy!
  - 2016: Google's AlphaGo beats Lee Sedol at Go

What Can AI Do?

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

- ✓ Play a decent game of Jeopardy?
- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of tennis?
- ✓ Grab a particular cup and put it on a shelf?
- ✗ Unload any dishwasher in any home?
- ? Drive safely along the highway?
- ✗ Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Berkeley Bowl?
- ? Discover and prove a new mathematical theorem?
- ✗ Perform a surgical operation?
- ✗ Unload a know dishwasher in collaboration with a person?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✗ Write an intentionally funny story?

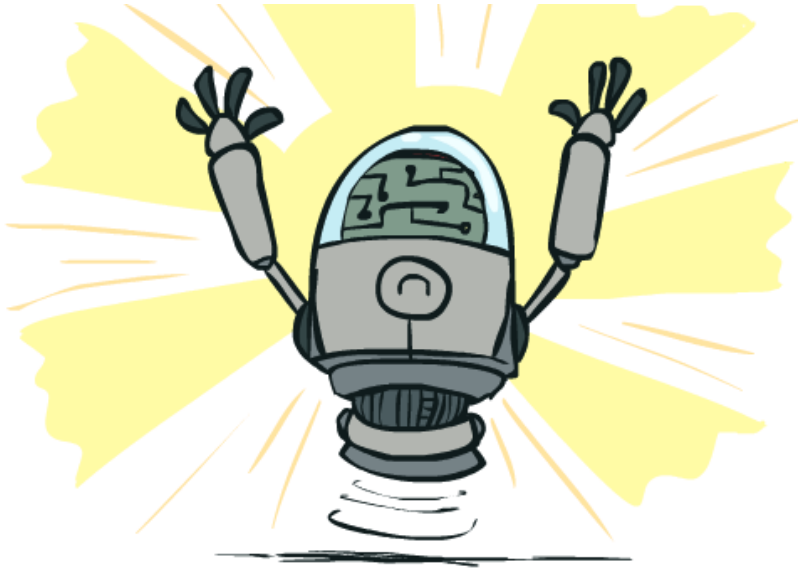

Unintentionally Funny Stories

- 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.

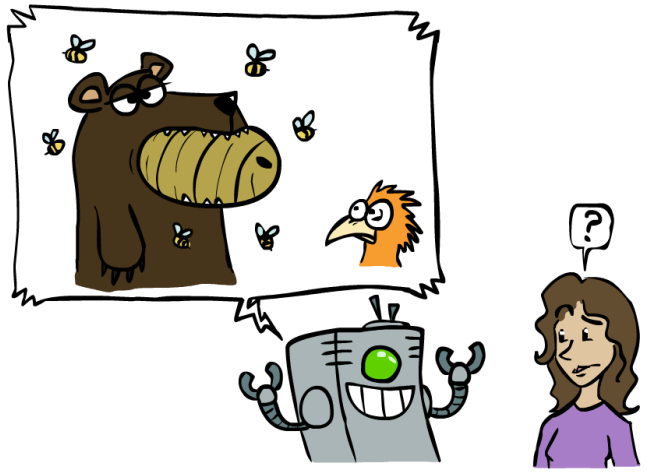

[Shank, Tale-Spin System, 1984]

Natural Language

- Speech technologies
  - Automatic speech recognition (ASR)
  - Text-to-speech synthesis (TTS)
  - Dialog systems
- Language processing technologies
  - Question answering
  - Machine translation
  - Web search
  - Text classification, spam filtering, etc…

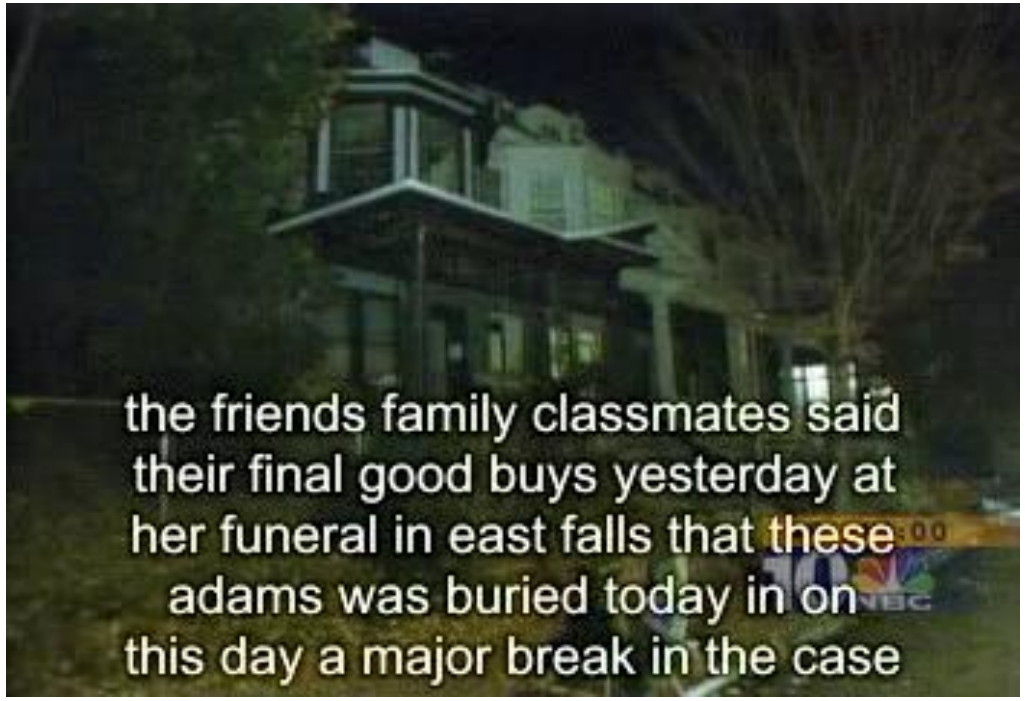

Computer Vision

"man in black shirt is playing guitar."

"construction worker in orange safety vest is working on road."

"two young girls are playing with lego toy."

"boy is doing backflip on wakeboard."

"girl in pink dress is jumping in air."

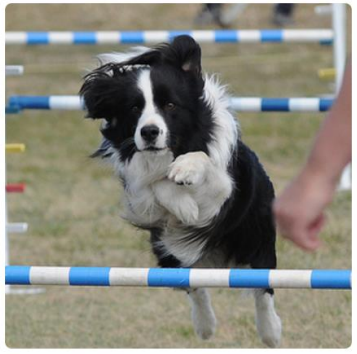

"black and white dog jumps over bar."

"young girl in pink shirt is swinging on swing."

"man in blue wetsuit is surfing on wave."

Karpathy & Fei-Fei, 2015; Donahue et al., 2015; Xu et al, 2015; many more

Course Topics

- Part 1: Intelligence from Computation
  - Fast search/planning
  - Constraint satisfaction (e.g. scheduling)
  - Adversarial and uncertain search (e.g. routing, navigation)
- Part 2: Intelligence from Data
  - Probabilistic inference with Bayes' nets (e.g. robot localization)
  - Decision theory
  - Supervised machine learning (e.g. spam detection)
- Throughout: Applications
  - Natural language, vision, robotics, games, etc.

Should I take CS 188?

- Yes, if you want to know how to design rational agents!
  - CS 188 gives you extra mathematical maturity
  - CS 188 gives you a survey of other non-CS fields that interact with AI (e.g. robotics, cognitive science, economics)
- Disclaimer: If you're interested in making yourself more competitive for AI jobs, CS 189 and CS 182 are better fits.
  - The last few CS 188 lectures (neural networks) are used by many modern state-of-the-art systems. CS 189 and CS 182 cover these in more depth

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 **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

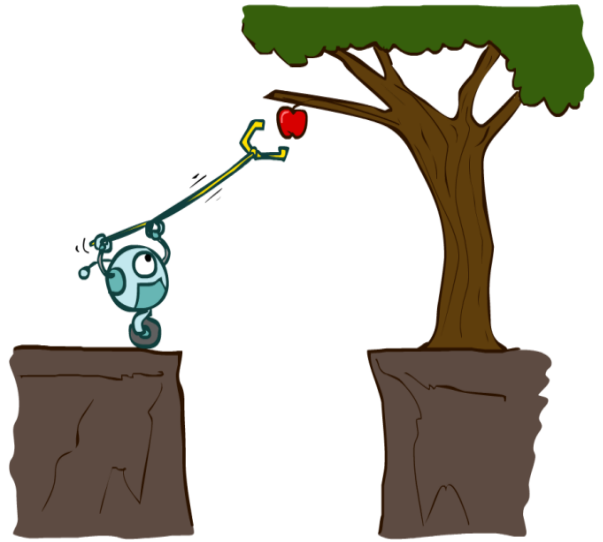

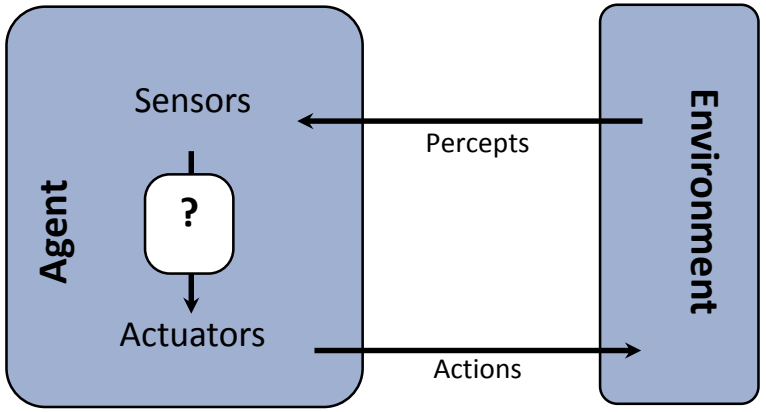

Pac-Man as an Agent

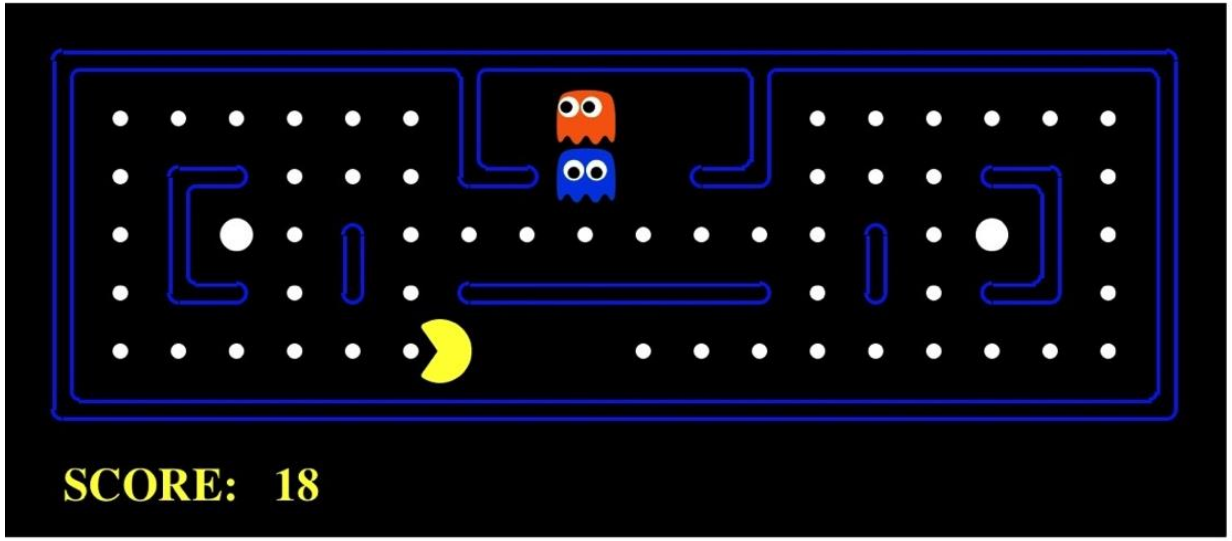

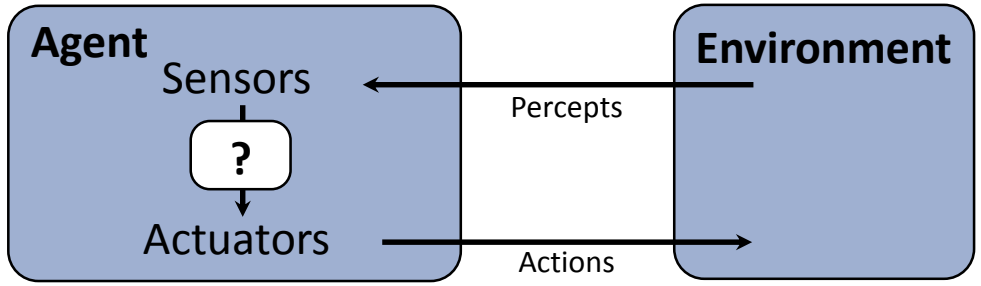

CS 188: Artificial Intelligence

Search

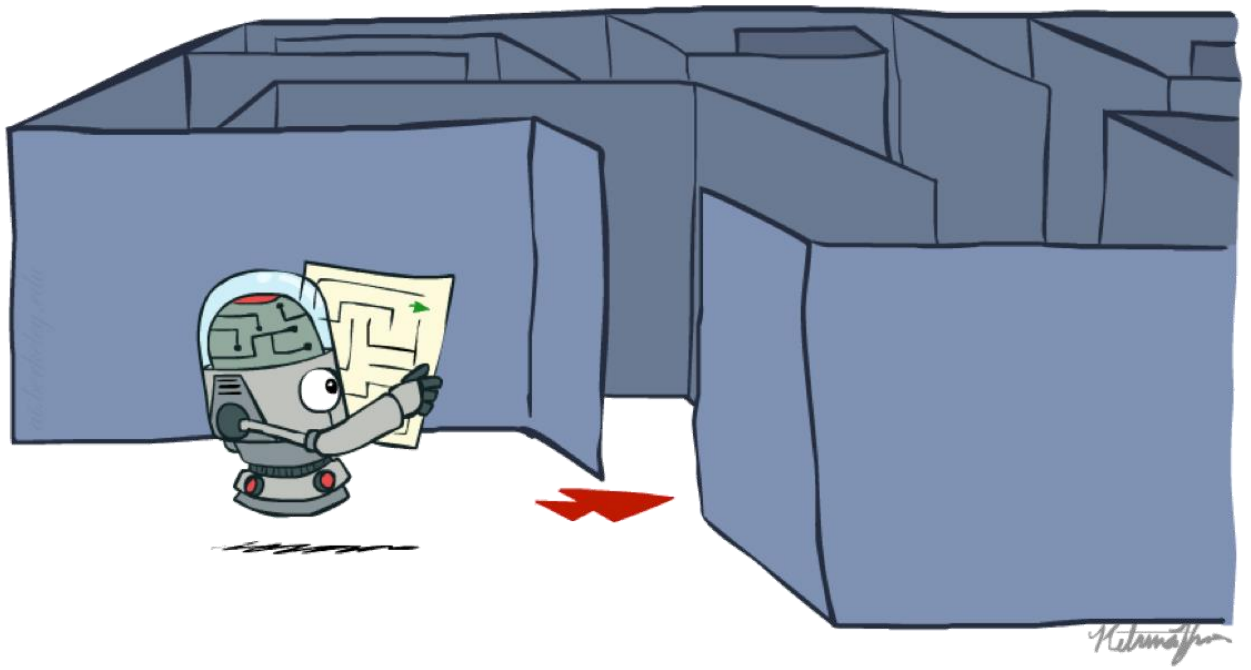

Fall 2022

University of California, Berkeley

[These slides were created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley (ai.berkeley.edu).]

Today

- Agents that Plan Ahead
- Search Problems
- Uninformed Search Methods
  - Depth-First Search
  - Breadth-First Search
  - Uniform-Cost Search

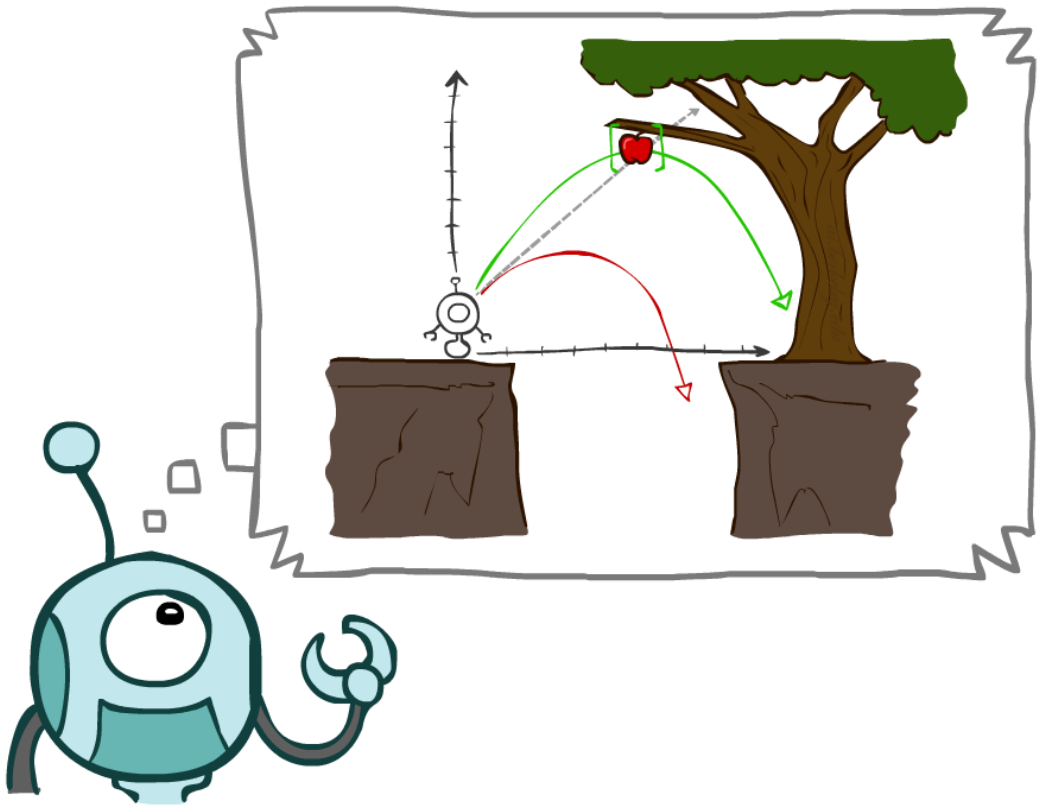

Agents that Plan

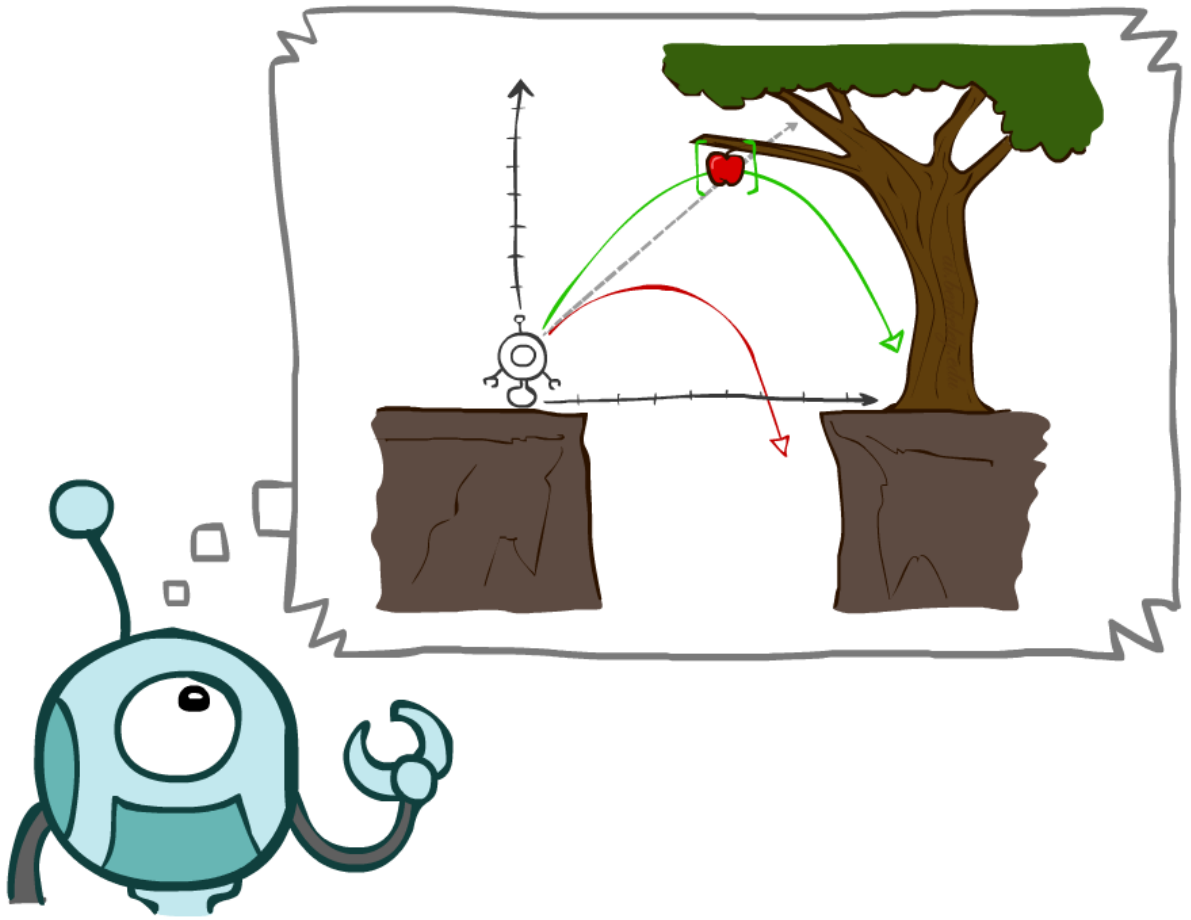

Reflex Agents

- Reflex agents:
  - Choose action based on current percept (and maybe memory)
  - May have memory or a model of the world's current state
  - Do not consider the future consequences of their actions
  - Consider how the world IS
- Can a reflex agent be rational?

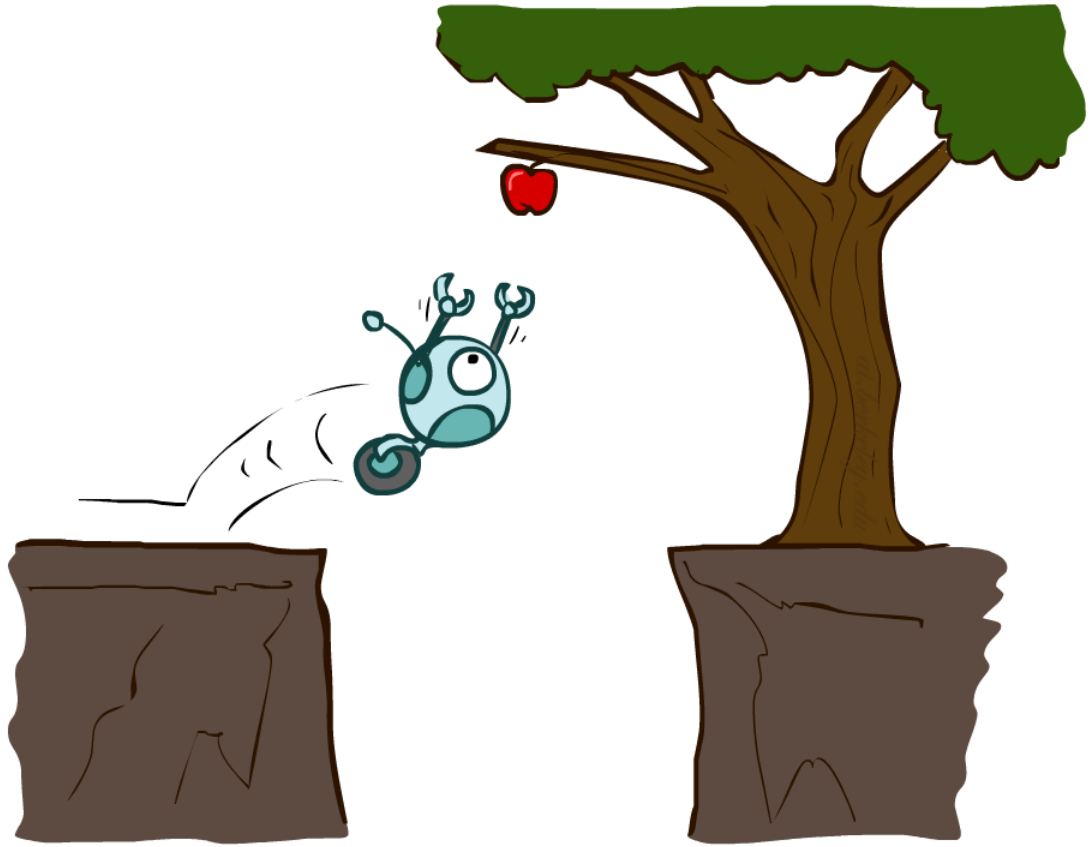

[Demo: reflex optimal (L2D1)]

[Demo: reflex optimal (L2D2)]

Planning Agents

- Planning agents:
  - Ask “what if”
  - Decisions based on (hypothesized) consequences of actions
  - Must have a model of how the world evolves in response to actions
  - Must formulate a goal (test)
  - Consider how the world WOULD BE
- Optimal vs. complete planning
- Planning vs. replanning

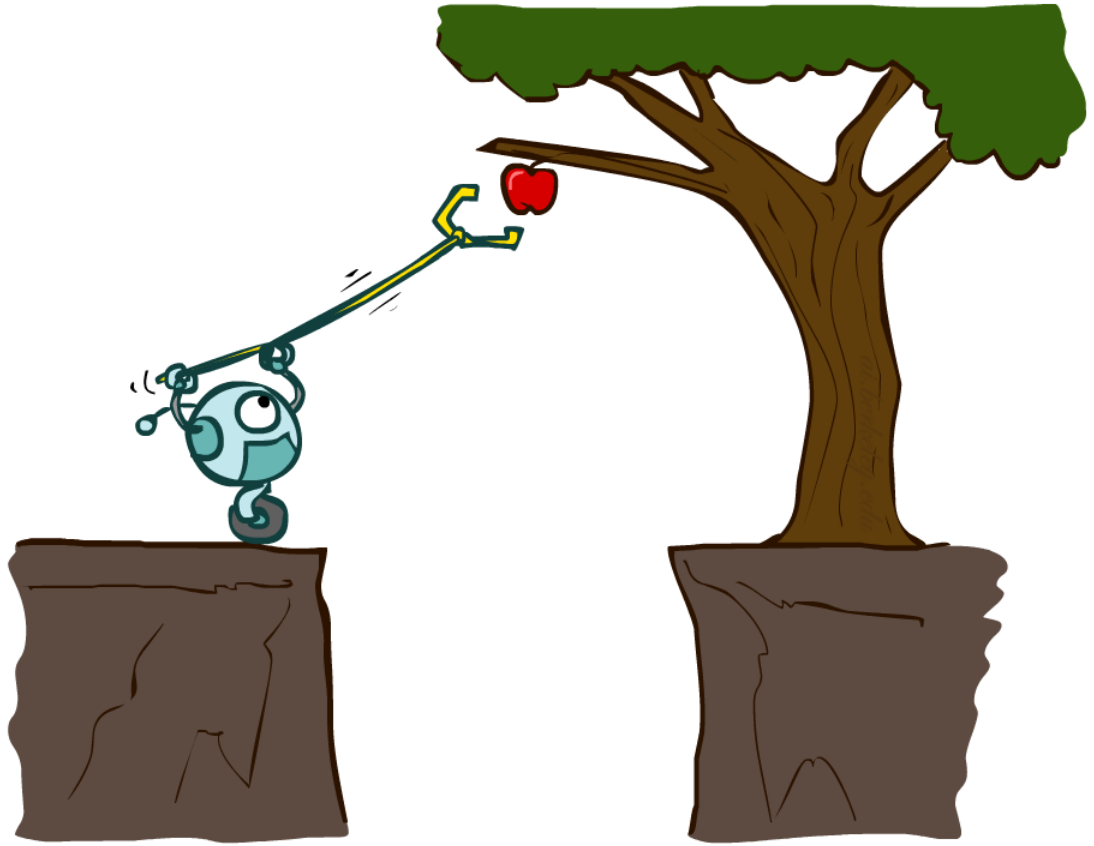

[Demo: re-planning (L2D3)]

[Demo: mastermind (L2D4)]

Search Problems

Search Problems

- A search problem consists of:
  - A state space

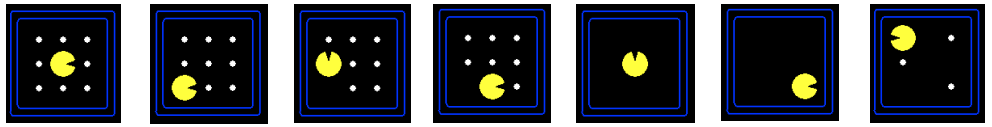

  - A successor function (with actions, costs)

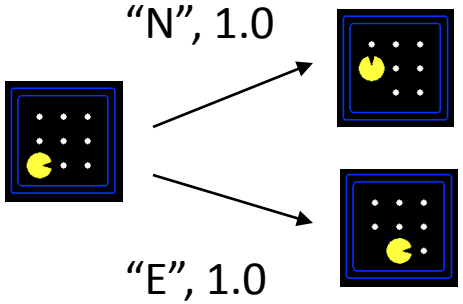

  - A start state and a goal test

- A solution is a sequence of actions (a plan) which transforms the start state to a goal state

Search Problems Are Models

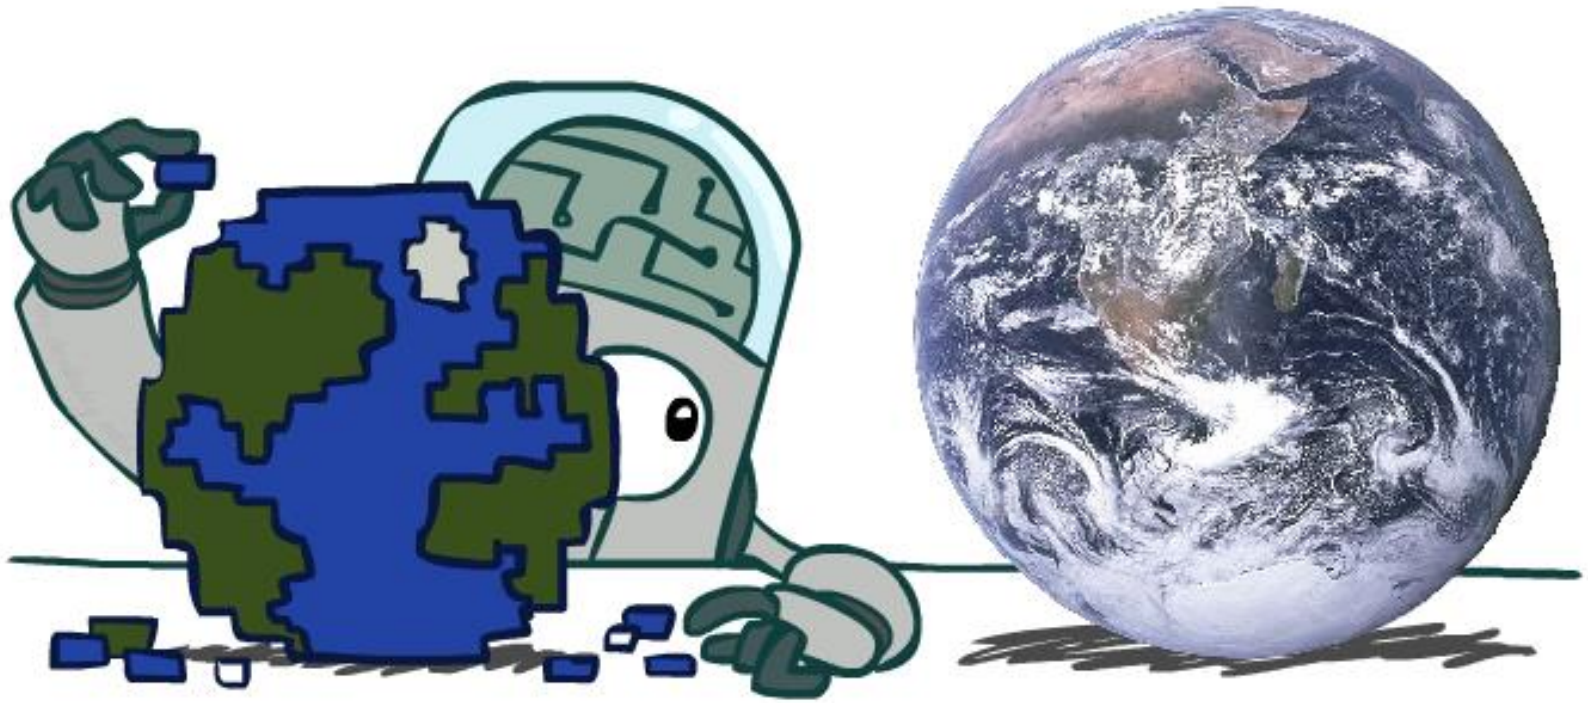

Example: Traveling in Romania

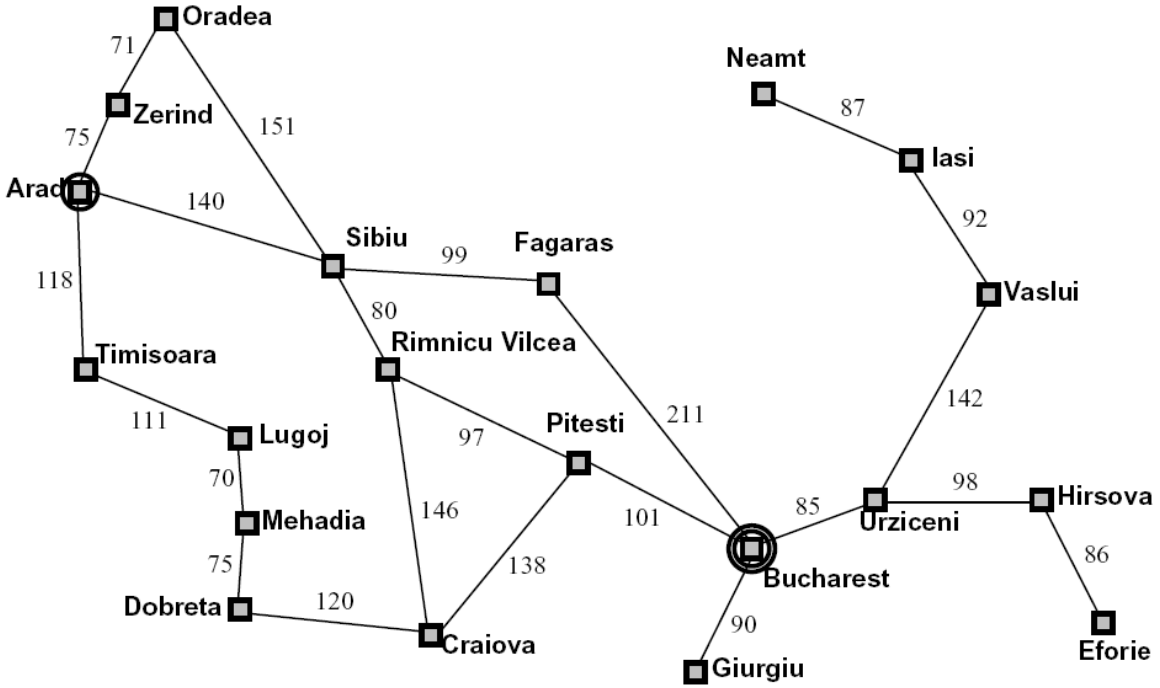

- State space:
  - Cities
- Successor function:
  - Roads: Go to adjacent city with cost = distance
- Start state:
  - Arad
- Goal test:
  - Is state == Bucharest?
- Solution?

What's in a State Space?

The world state includes every last detail of the environment

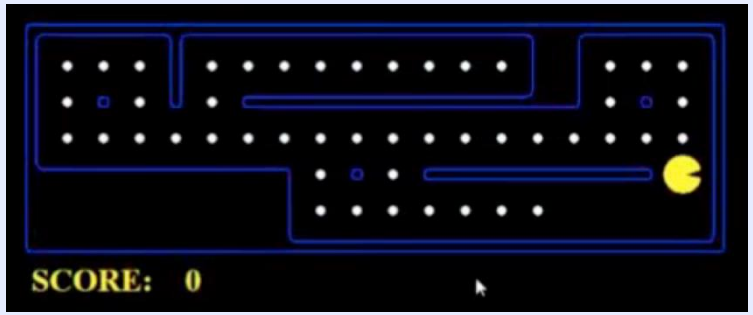

A search state keeps only the details needed for planning (abstraction)

- Problem: Pathing
  - States: (x,y) location
  - Actions: NSEW
  - Successor: update location only
  - Goal test: is (x,y)=END

- Problem: Eat-All-Dots
  - States: {(x,y), dot booleans}
  - Actions: NSEW
  - Successor: update location and possibly a dot boolean
  - Goal test: dots all false

State Space Sizes?

- World state:
  - Agent positions: 120
  - Food count: 30
  - Ghost positions: 12
  - Agent facing: NSEW

- How many
  - World states?

    $120x(2^{30})x(12^2)x4$
  - States for pathing?

    120
  - States for eat-all-dots?

    $120x(2^{30})$

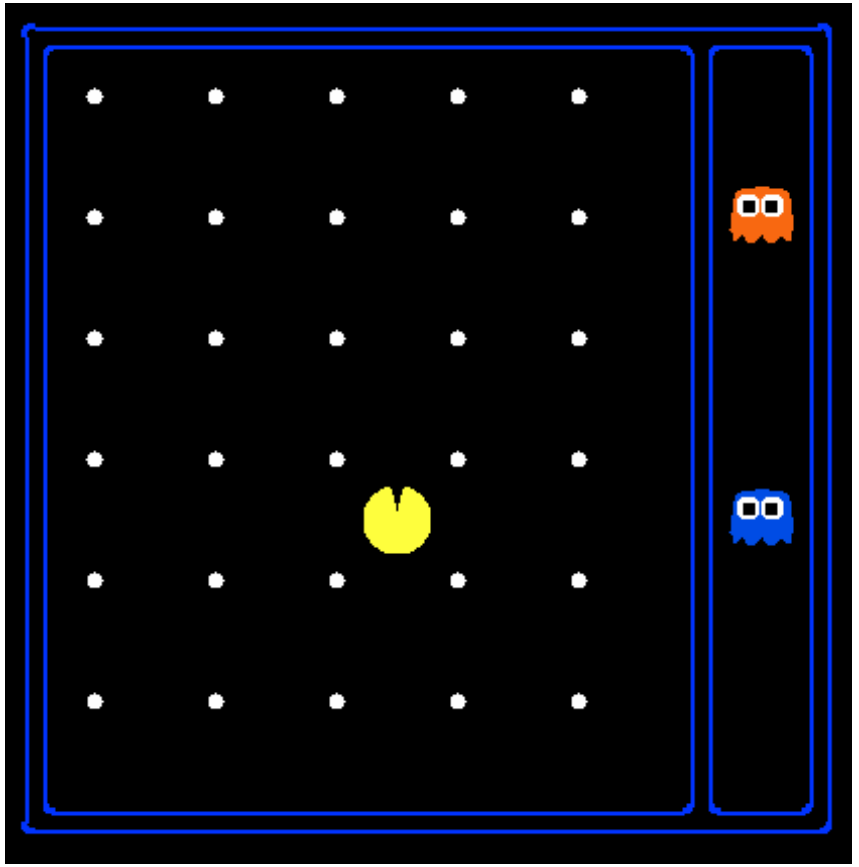

Quiz: Safe Passage

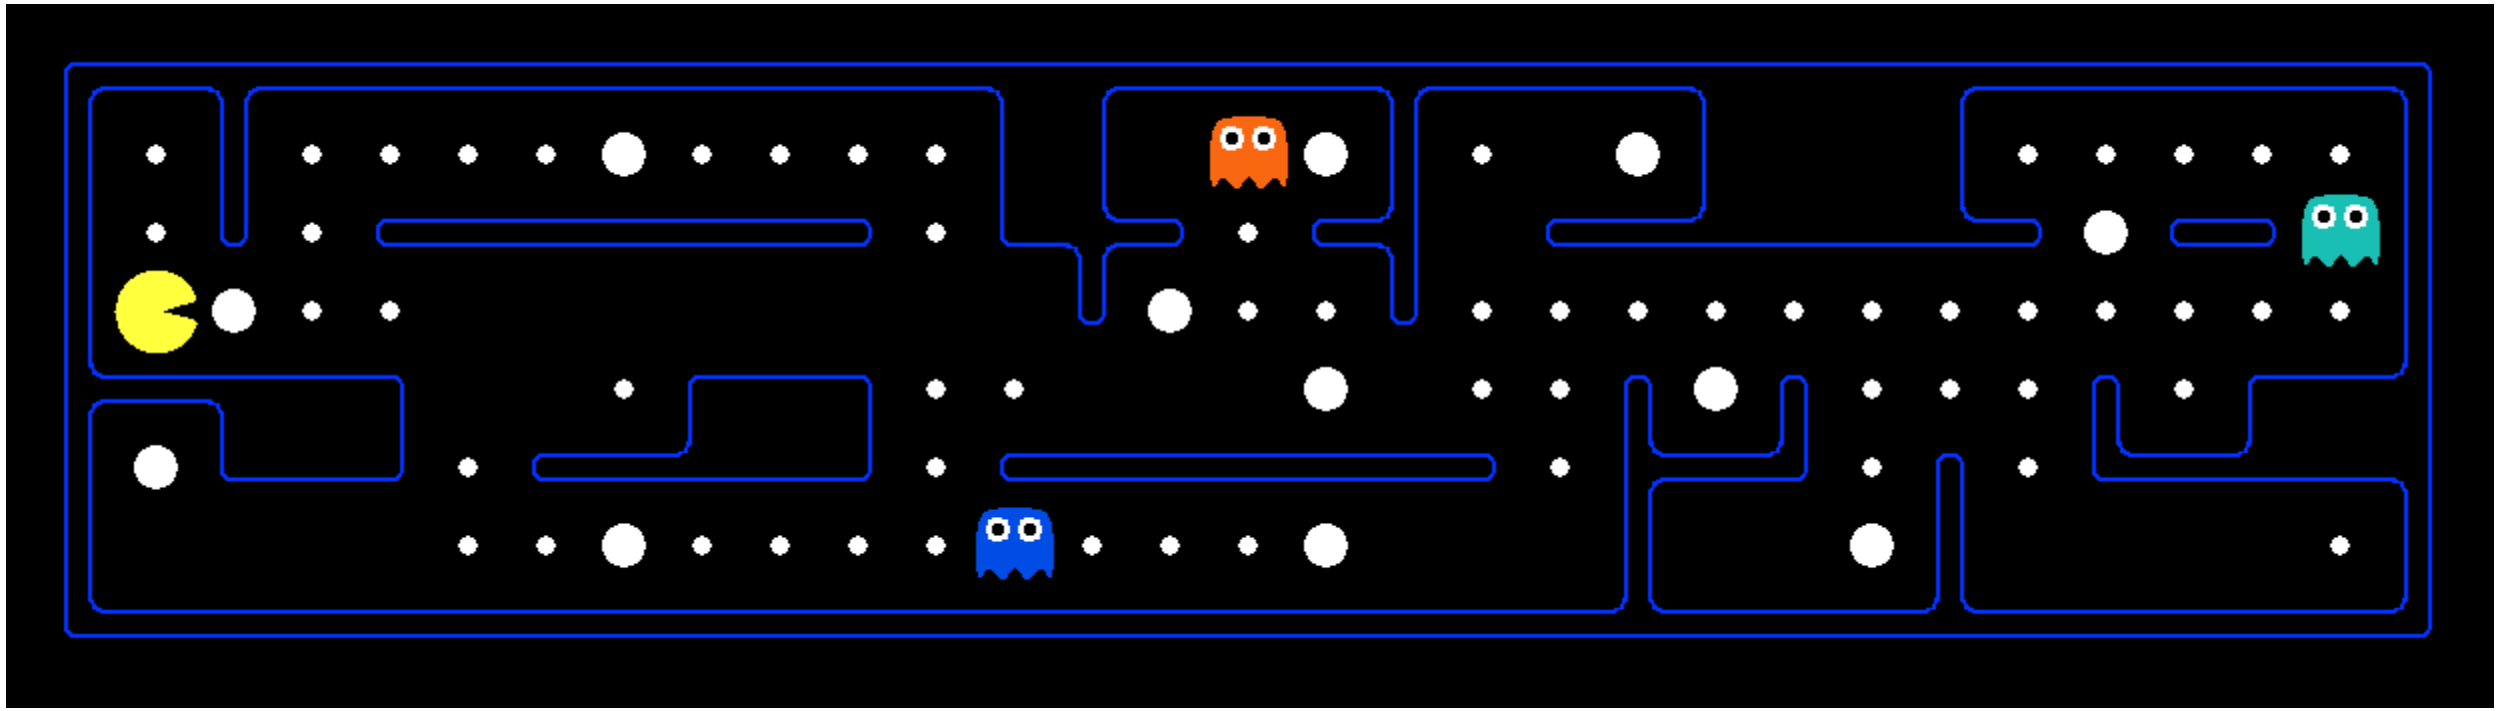

- Problem: eat all dots while keeping the ghosts perma-scared
- What does the state space have to specify?
  - (agent position, dot booleans, power pellet booleans, remaining scared time)