CS 188: Artificial Intelligence

Introduction

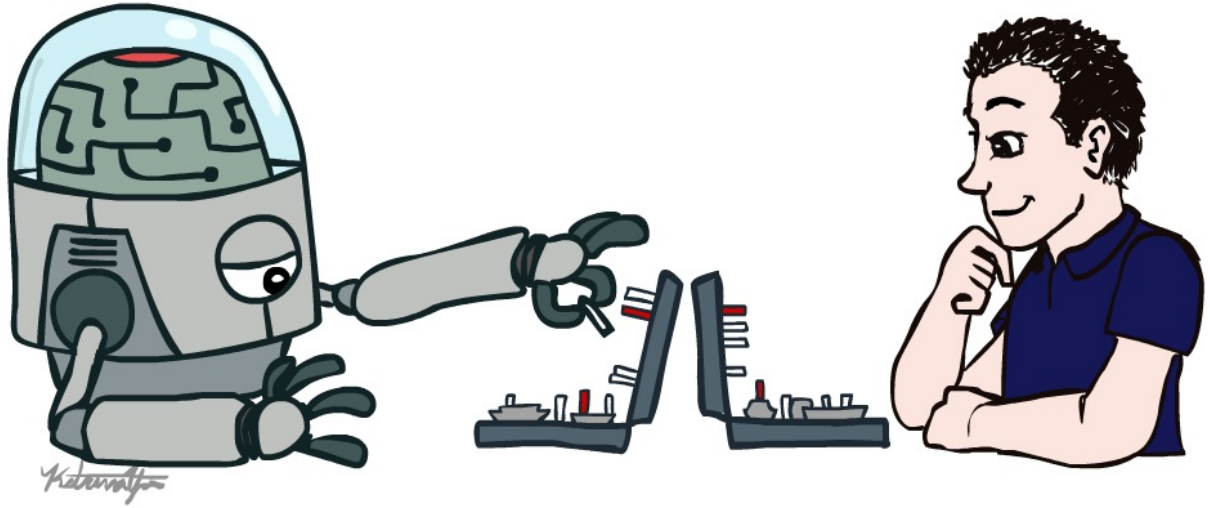

Spring 2023

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: Intros and Logistics

- Staff introductions: Peyrin, Igor 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: 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.

Staff Introductions: Igor (he/him)

- Currently: Lecturer and Researcher at Google DeepMind
- Previously: OpenAI, UC Berkeley, Pixar
- Research focus: robotics, large models, multi-agent AI

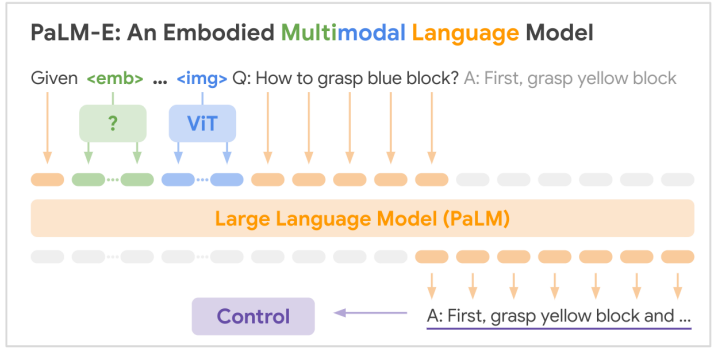

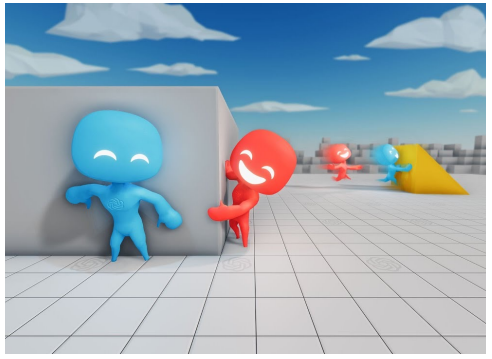

- Favorite animal: wombat
- Loves: feedback

Our talented course staff!

Evgeny Pobachienko
he/him

Sid Ijju
he/him

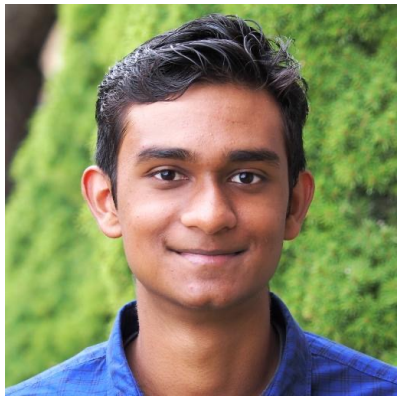

Ajay Sridhar
he/him

Alina Trinh
she/her

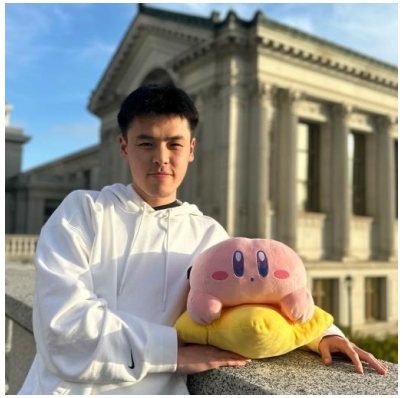

Austen Liao
he/him

Cham Yao
he/him

Jerry Sun
he/him

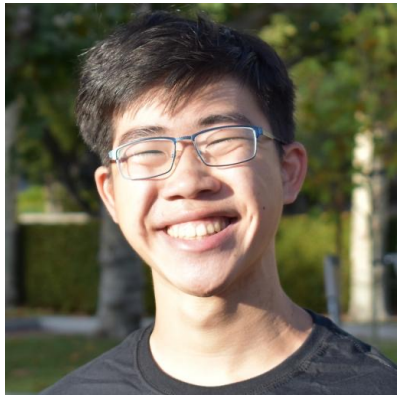

Joshua Liao
he/him

Our talented course staff!

Joy Liu
she/her

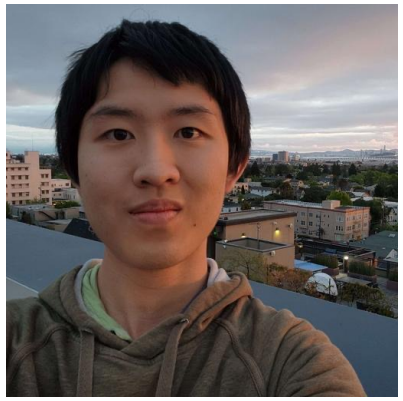

Kenny Wang
he/him

Michael Wu
he/him

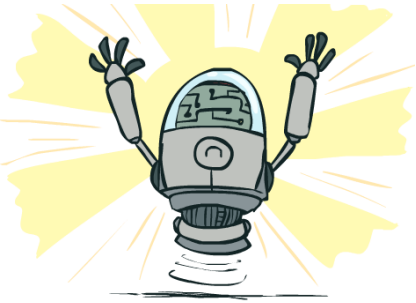

Nitish Dashora
he/him

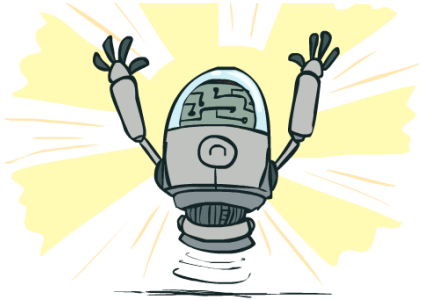

Pranav Muralikrishnan
he/him

Sashrika Pandey
she/her

Xavier Yin
he/him

Aidan Leung
he/him

Our talented course staff!

Allen Cao
he/him

Andres Lam
he/him

Dhruv Kumar
he/him

Emma Wu
she/her

Saathvik Selvan
he/him

Samantha Huang
she/her

Somu Patil
he/him

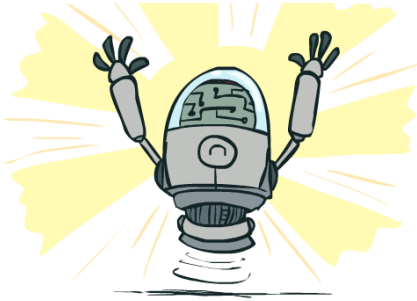

might have a few more people join so here's a placeholder

Enrollment

- Course staff does not control enrollment; we have to follow department policy
  - Only CS majors will be able to enroll this spring
  - More details on the course website

Course Structure: Lectures

- You are here!
- Tuesday/Thursday, 5:00–6:30 PM PT
- Attendance is not taken
  - But there may be a bit of extra credit for attending
- You can attend:
  - In-person in Pimentel 1
  - Remotely over Zoom
  - Asynchronously by watching recordings (posted next day 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 28)
- You can attend any discussion section you want (no need to enroll in a section)
  - A bit of extra credit available for attendance

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 (August 28)

Course Structure: Exams

- Save the dates!
  - Midterm: Monday, October 16, 2023, 7–9pm PT
  - Final exam: Thursday, December 14, 11:30am–2:30pm 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/fa23/
  - All resources (slides, notes, recordings, assignments, etc.) posted here
- Ed: Discussion forum
- 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 is this course?

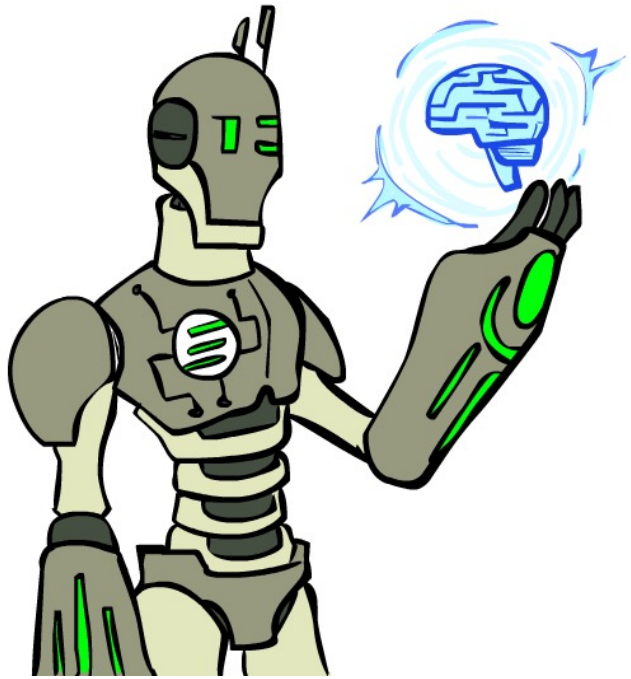

AI is having real-world impact

- Public imagination
  - Text assistants

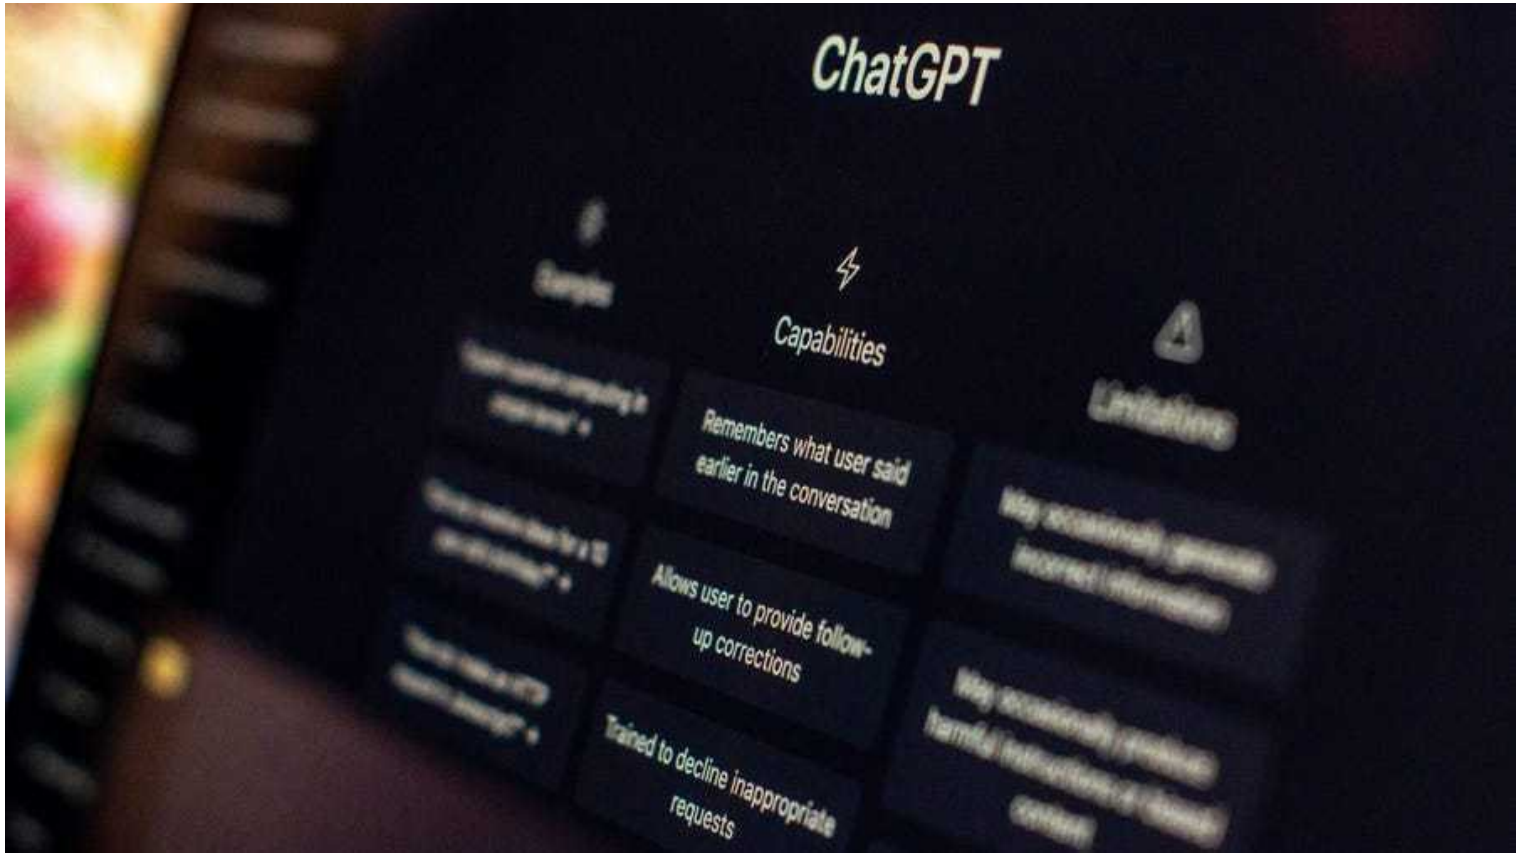

AI is having real-world impact

- Public imagination
  - Text assistants
  - Image generation

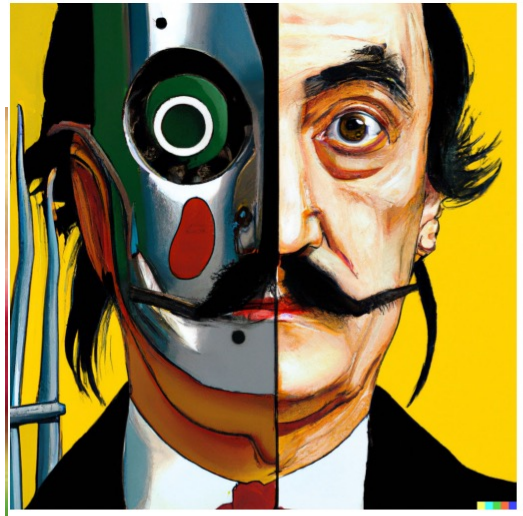

vibrant portrait painting of Salvador Dalí with a robotic half face

a shiba inu wearing a beret and black turtleneck

a close up of a handpalm with leaves growing from it

an espresso machine that makes coffee from human souls, artstation

panda mad scientist mixing sparkling chemicals, artstation

a corgi's head depicted as an explosion of a nebula

AI is having real-world impact

- Public imagination
- Economy
  - 454 billion USD globally

The global artificial intelligence (AI) market size was valued at USD 454.12 billion in 2022 and is expected to hit around USD 2,575.16 billion by 2032, progressing with a CAGR of 19% from 2023 to 2032. The North America artificial intelligence market was valued at USD 167.30 billion in 2022.

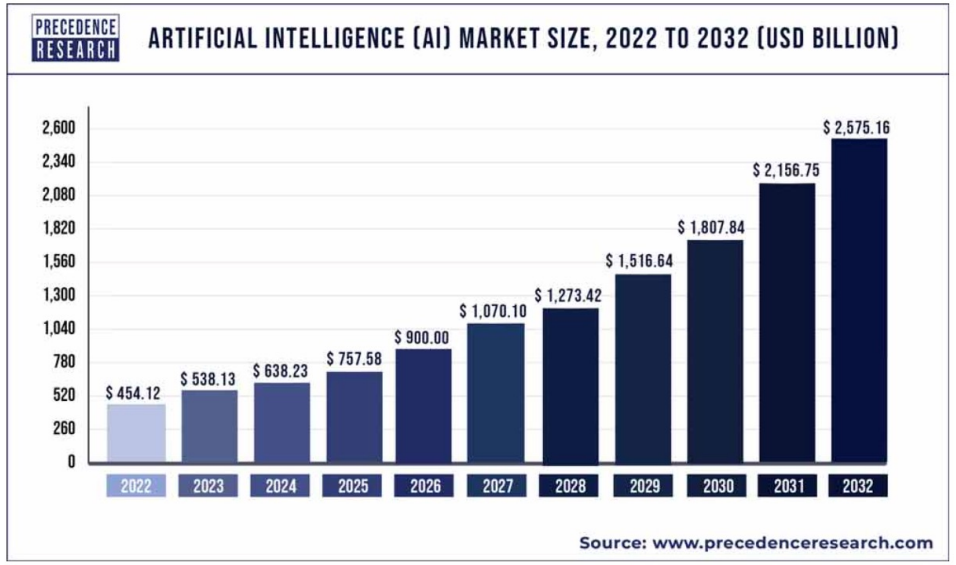

https://www.precedenceresearch.com/artificial-intelligence-market

AI is having real-world impact

- Public imagination
- Economy
- Politics

AI is having real-world impact

- Public imagination
- Economy
- Politics

AI is having real-world impact

- Public imagination
- Economy
- Politics

AI is having real-world impact

- Public imagination
- Economy
- Politics

AI is having real-world impact

- Public imagination
- Economy
- Politics
- Law

Aug. 18, 2023, 12:18 PM; Updated: Aug. 18, 2023, 12:48 PM

AI-Generated Art Lacks Copyright Protection, D.C. Court Says (1)

Riddhi Setty
Reporter

Isaiah Poritz
Legal Reporter

Bloomberg Law, 2023

AI is having real-world impact

- Public imagination
- Economy
- Politics
- Law
- Labor

Finance & economics | Free exchange

New research shows the robots are coming for jobs—but stealthily

Look beneath the aggregate economic numbers, and change is afoot

The Economist, 2021

The Optimist's Guide to Artificial Intelligence and Work

The focus of much discussion is on how it will replace jobs, but nothing is inevitable.

New York Times, 2023

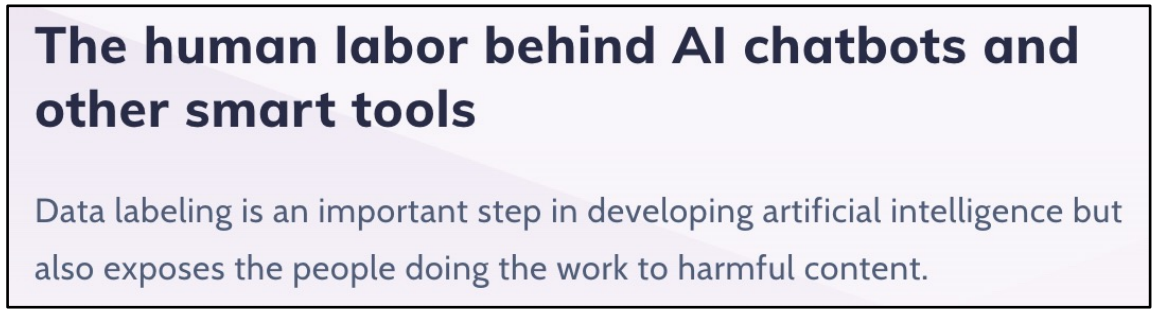
The human labor behind AI chatbots and other smart tools

Data labeling is an important step in developing artificial intelligence but also exposes the people doing the work to harmful content.

MarketWatch, 2023

AI is having real-world impact

- Public imagination
- Economy
- Politics
- Law
- Labor
- Sciences

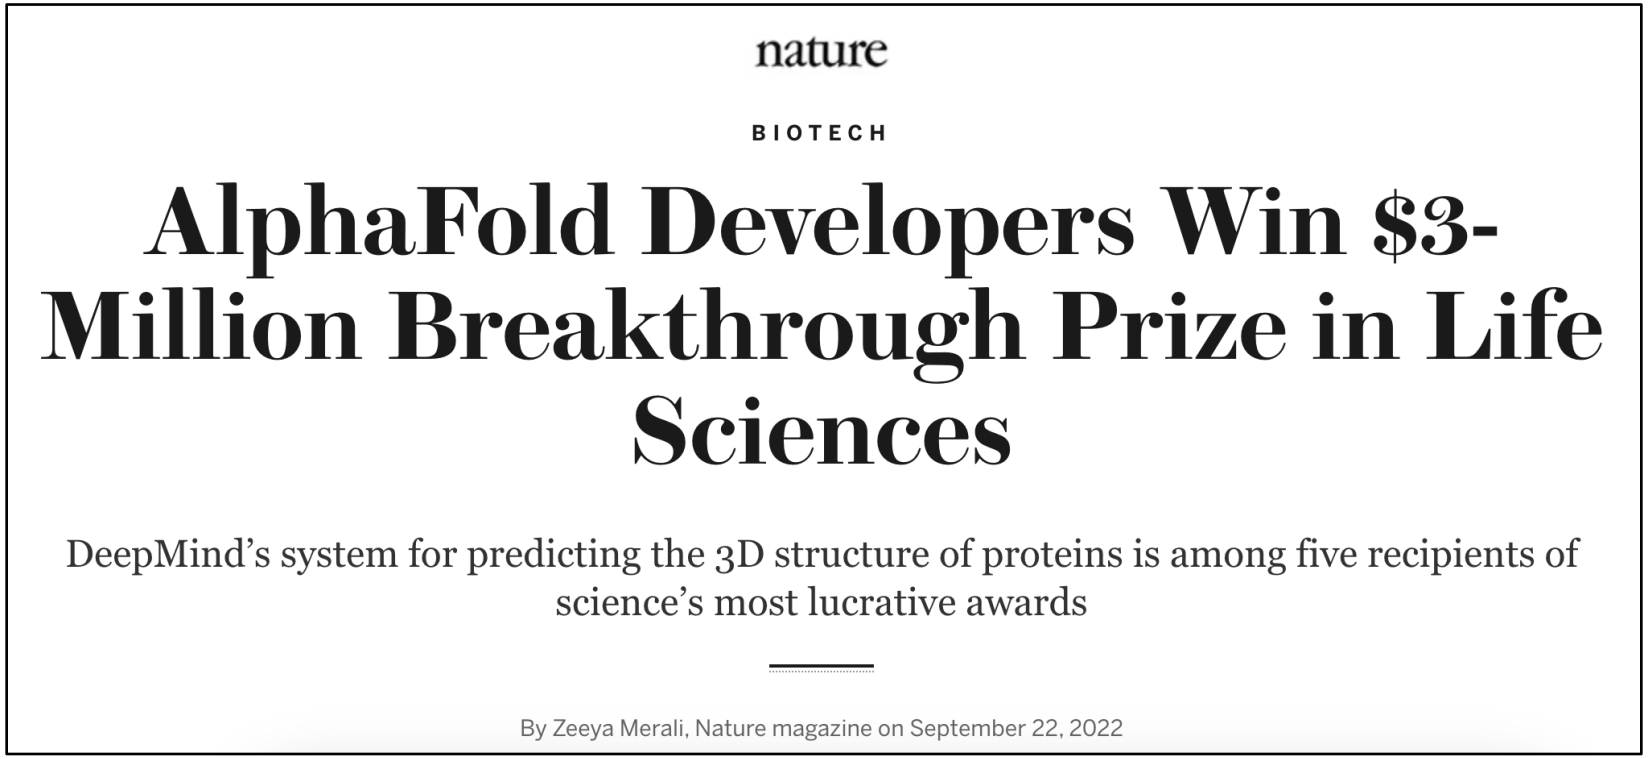

nature

BIOTECH

AlphaFold Developers Win $3-Million Breakthrough Prize in Life Sciences

DeepMind's system for predicting the 3D structure of proteins is among five recipients of science's most lucrative awards

By Zeeya Merali, Nature magazine on September 22, 2022

Nature, 2022

AI is having real-world impact

- Public imagination
- Economy
- Politics
- Law
- Labor
- Sciences

AMIT KATWALA SCIENCE FEB 16, 2022 11:00 AM

DeepMind Has Trained an AI to Control Nuclear Fusion

The Google-backed firm taught a reinforcement learning algorithm to control the fiery plasma inside a tokamak nuclear fusion reactor.

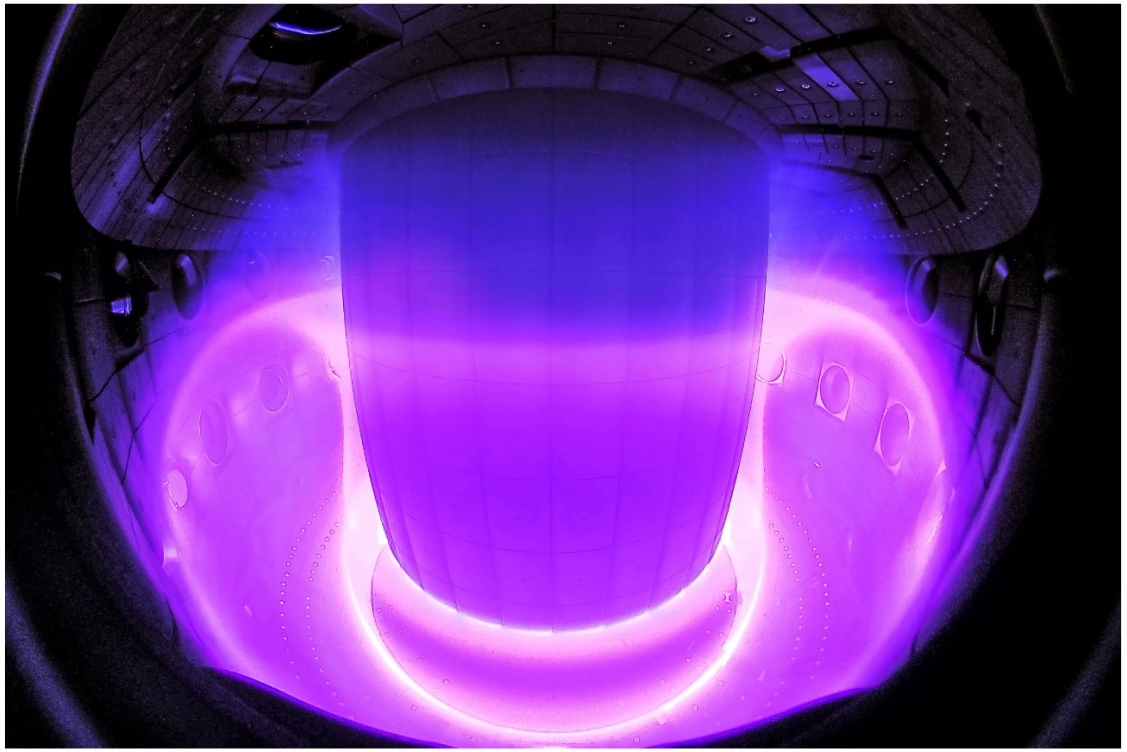

PHOTOGRAPH:CURDIN WÜTHRICH, SPC/EPFL

Wired, 2022

AI is having real-world impact

- Public imagination
- Economy
- Politics
- Law
- Labor
- Sciences
- Education

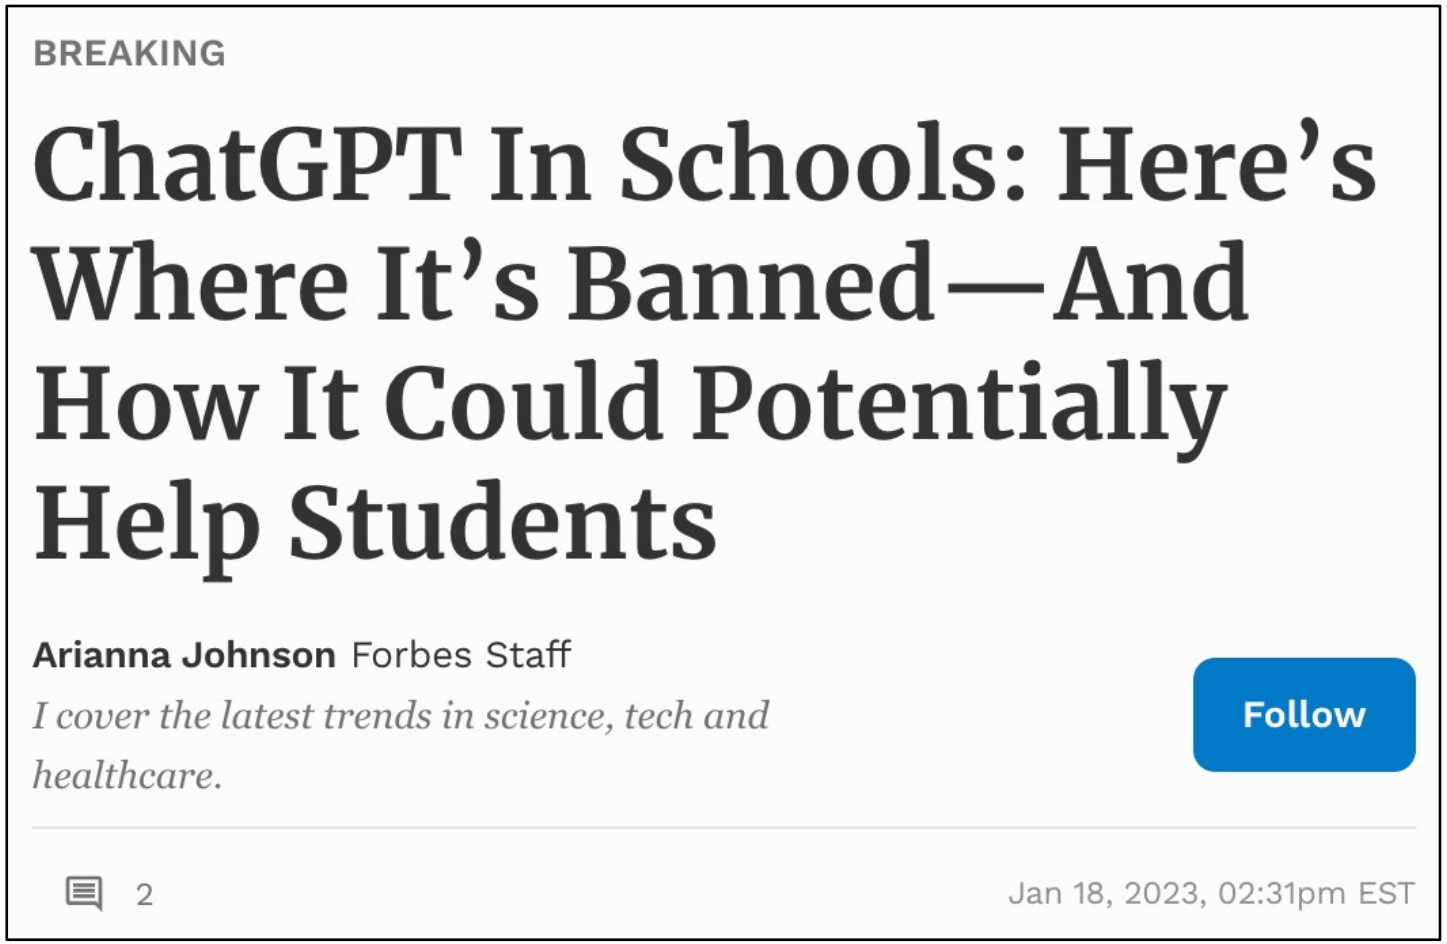

BREAKING

ChatGPT In Schools: Here's Where It's Banned—And How It Could Potentially Help Students

Arianna Johnson Forbes Staff

I cover the latest trends in science, tech and healthcare.

Follow

2

Jan 18, 2023, 02:31pm EST

Forbes, 2023

AI is having real-world impact

- Public imagination
- Economy
- Politics
- Law
- Labor
- Sciences
- Education

Ok, but what does AI *do*???

Science fiction AI?

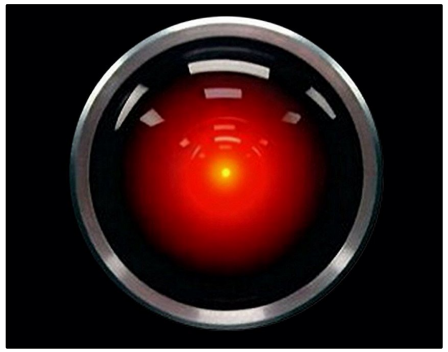

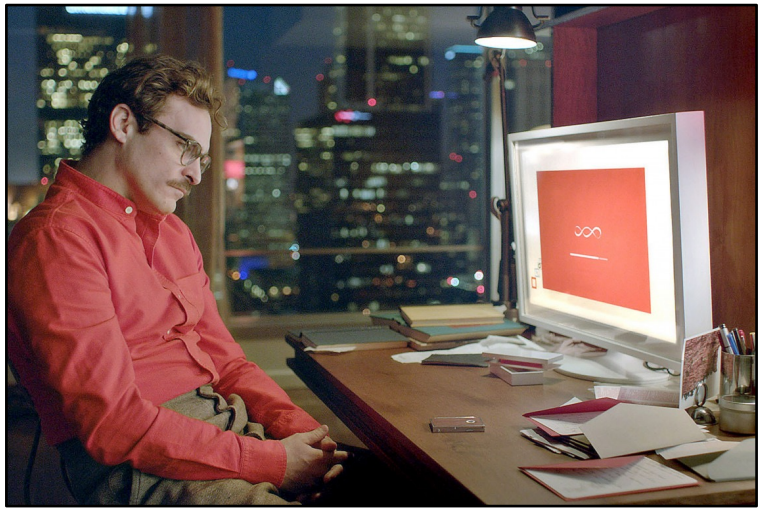

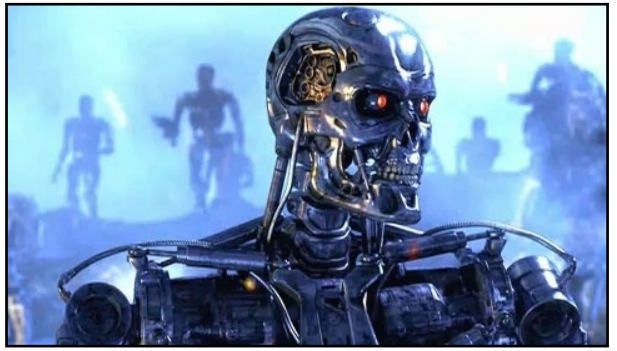

What is AI?

The science of making machines that:

Rational Decisions

- We'll use the term **rational** in a very specific, technical way:
  - Rational: *maximally achieving pre-defined goals*
  - Goals are expressed in terms of the **utility** of outcomes
  - World is uncertain, so we'll use **expected** utility
  - Being rational means acting to **maximize your expected utility**

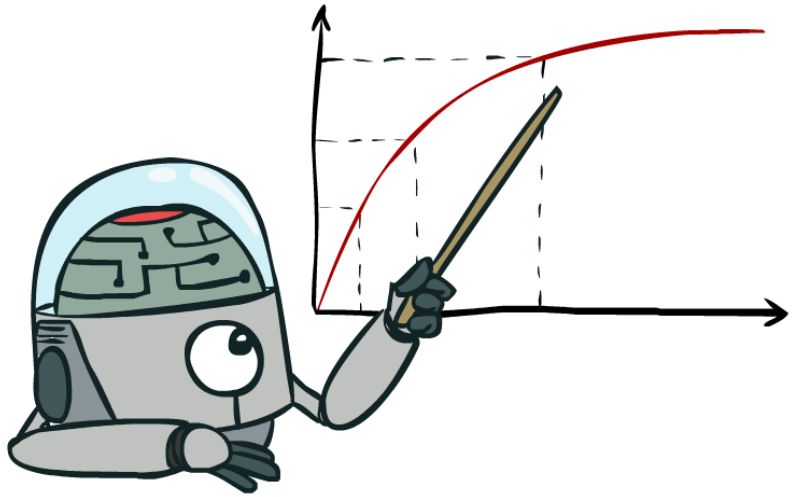

Rational Decisions

- We'll use the term **rational** in a very specific, technical way:
  - Rational: *maximally achieving pre-defined goals*
  - Goals are expressed in terms of the **utility** of outcomes
  - World is uncertain, so we'll use **expected** utility
  - Being rational means acting to **maximize your expected utility**

A better title for this course would be:

Computational Rationality

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!
- AI may be better than brains at some tasks
- *"Brains are to intelligence as wings are to flight"*
- We can't yet build AI on the scale of the brain
  - ~100T synapses in the human brain vs 500B weights in artificial neural networks
- Still, the brain can be a great inspiration for AI!

A (Short) History of AI

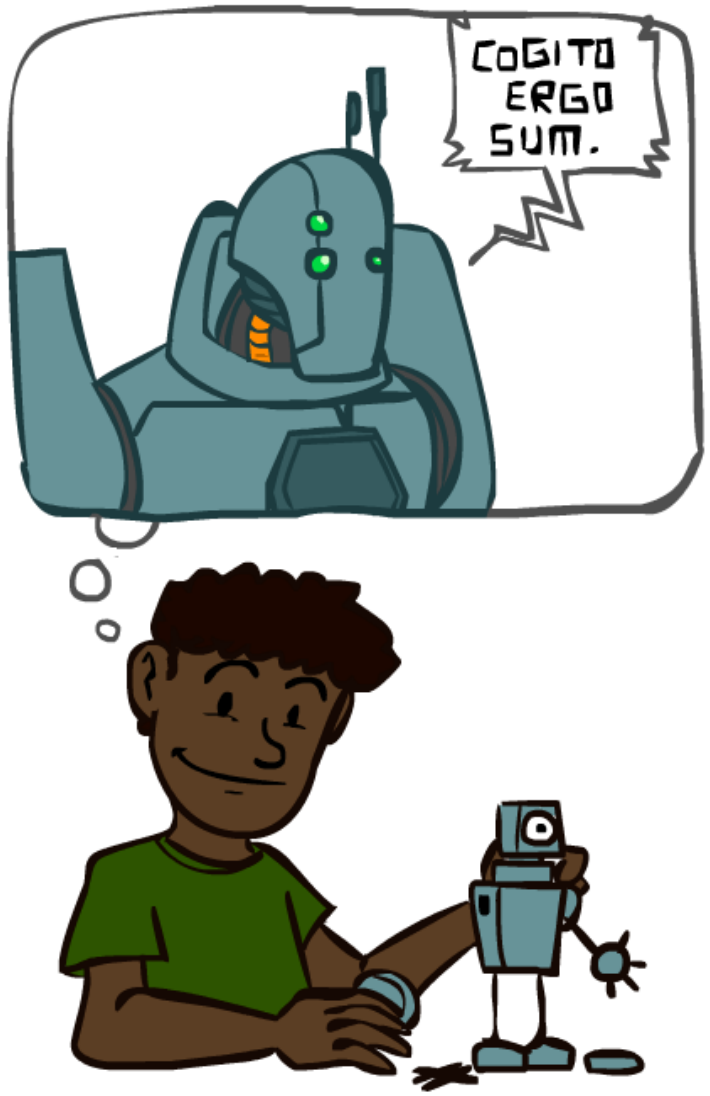

A (Short) History of AI

- 1940-1950: Early days: neural and computer science meet
  - 1943: McCulloch & Pitts: Boolean circuit model of brain
  - 1950: Turing's “Computing Machinery and Intelligence”
- 1950—70: Excitement! Logic-driven
  - 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

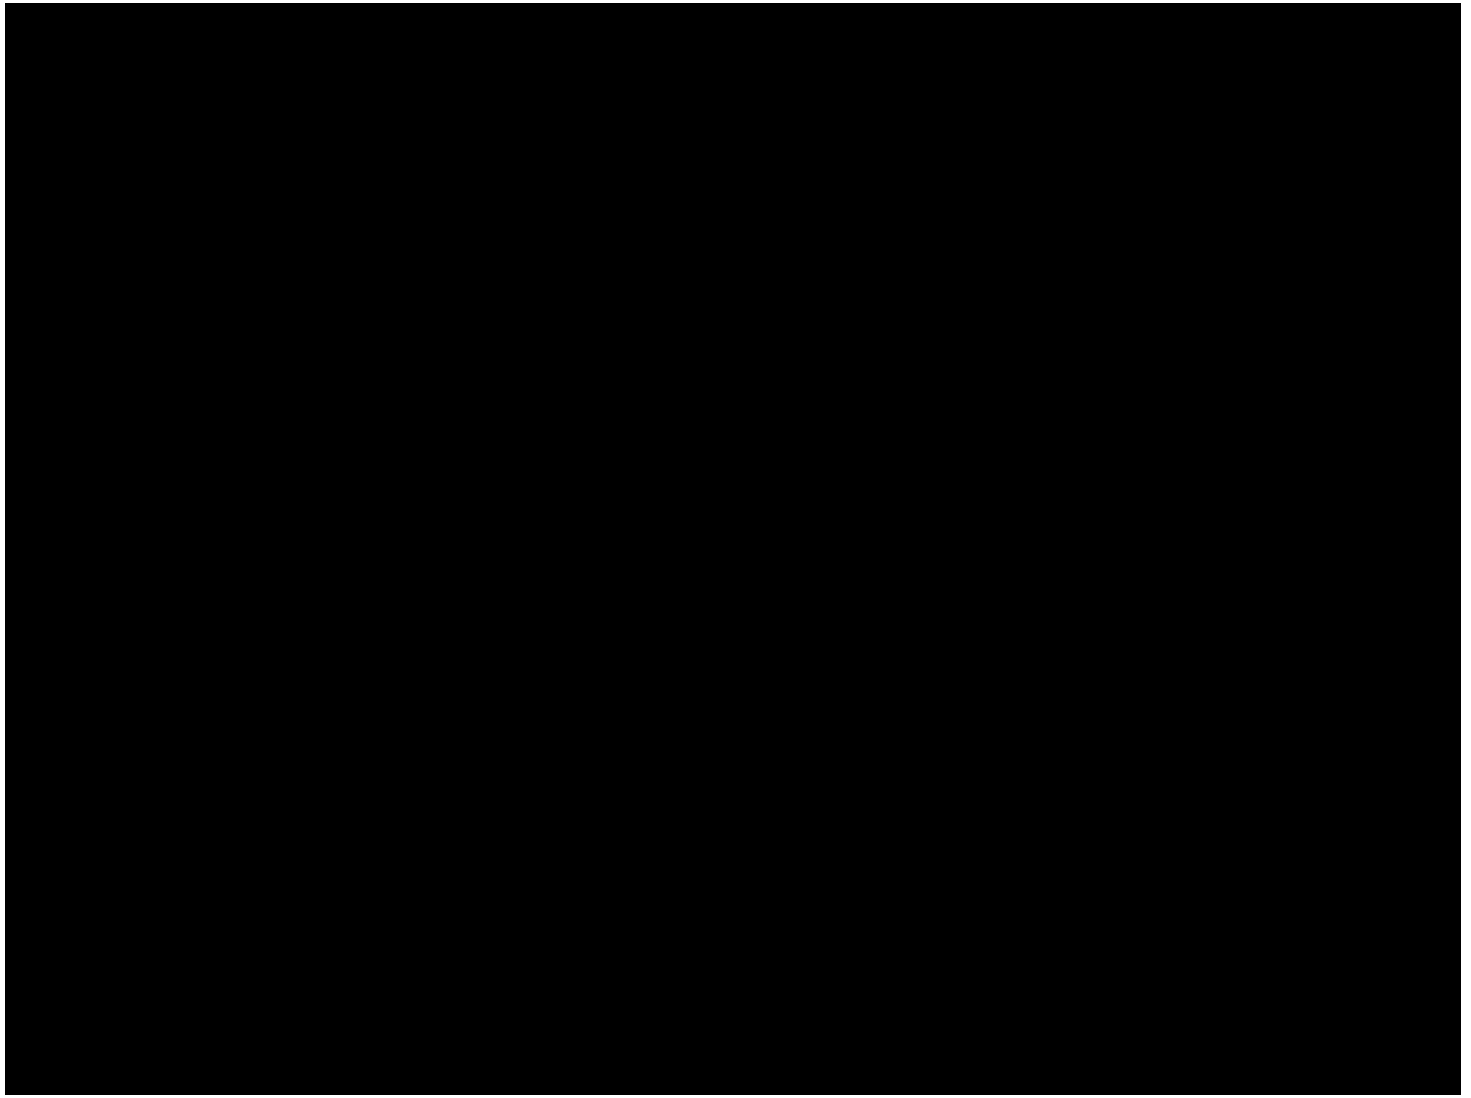

A (Short) History of AI

- 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
  - What can AI do now?

What Can AI Do?

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

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- Play a decent game of tennis?
- Unload any dishwasher in any home?
- Drive safely along the highway?
- Drive safely along streets of San Francisco?
- 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?
- Translate spoken Chinese into spoken English in real time?
- Win an art competition?
- Write an intentionally funny story?
- Construct a building?

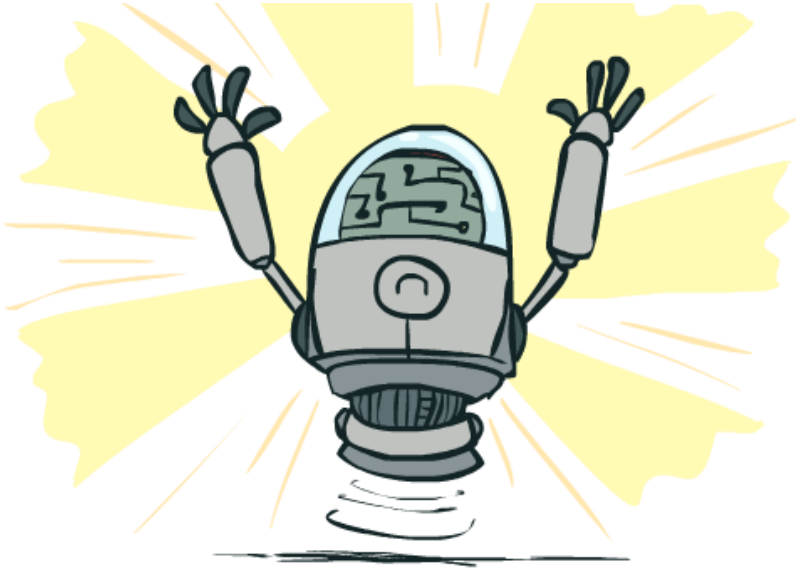

What Can AI Do?

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

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- Play a decent game of tennis?
- Unload any dishwasher in any home?
- Drive safely along the highway?
- Drive safely along streets of San Francisco?
- 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?
- Translate spoken Chinese into spoken English in real time
- Win an art competition?
- Write an intentionally funny story?
- Construct a building?

Google's AlphaGo beats Lee Sedol at Go in 2016

But … a plot twist in 2023!

DEEP BLUE WAS JUST THE START —

Man beats machine at *Go* in human victory over AI

Amateur exploited weakness in systems that have otherwise dominated grandmasters.

RICHARD WATERS, FINANCIAL TIMES - 2/19/2023, 4:51 AM

What Can AI Do?

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

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of tennis?
- Unload any dishwasher in any home?
- Drive safely along the highway?
- Drive safely along streets of San Francisco?
- 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?
- Translate spoken Chinese into spoken English in real time?
- Win an art competition?
- Write an intentionally funny story?
- Construct a building?

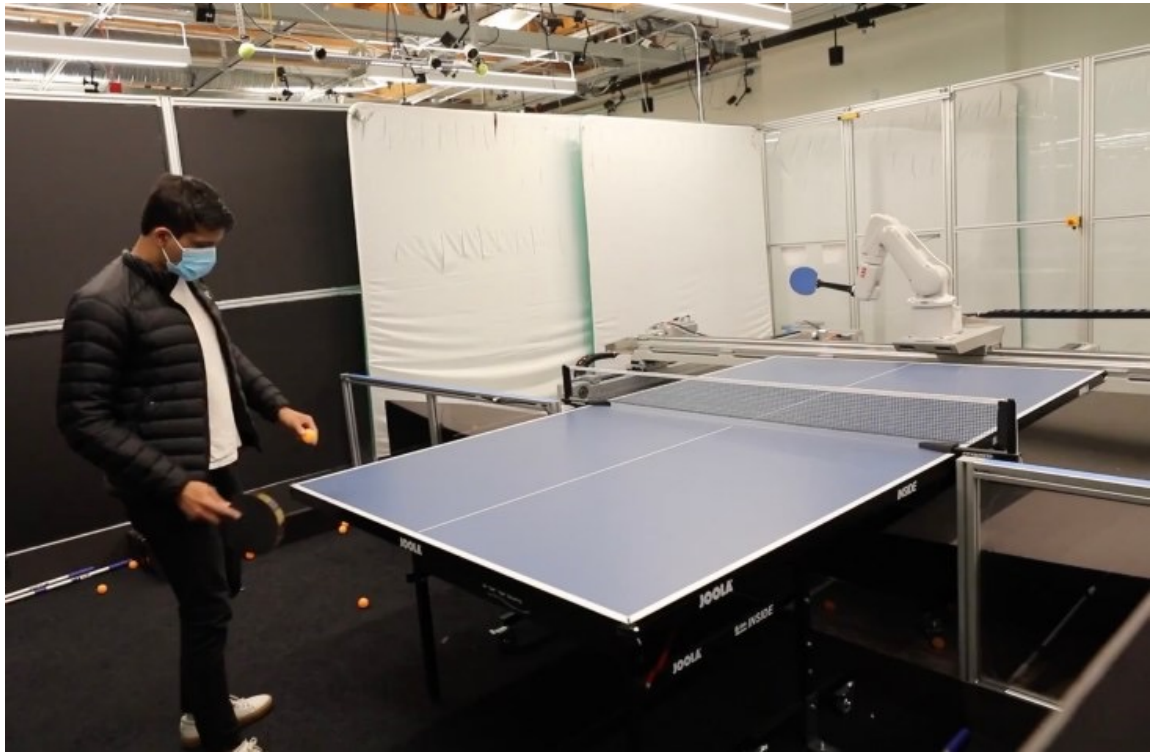

What Can AI Do?

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

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of tennis?
- ✗ Unload any dishwasher in any home?
- ✓ Drive safely along the highway?
- ? Drive safely along streets of San Francisco?
- 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?
- Translate spoken Chinese into spoken English in real tim
- Win an art competition?
- Write an intentionally funny story?
- Construct a building?

California just opened the floodgates for self-driving cars

After a state regulatory board ruling, San Francisco will have 24/7 robotaxis on its streets. A Washington Post analysis shows how it's a pivotal moment for the industry.

By Trisha Thadani and Jeremy B. Merrill

Updated August 10, 2023 at 10:06 p.m. EDT | Published August 10, 2023 at 9:15 p.m. EDT

The Register®

California DMV hits brakes on Cruise's SF driverless fleet after series of fender benders

50% chop effective 'immediately' as department investigates traffic and safety issues

Jude Karabus

Mon 21 Aug 2023 // 14:28 UTC

What Can AI Do?

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

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of tennis?
- ✗ Unload any dishwasher in any home?
- ✓ Drive safely along the highway?
- ? Drive safely along streets of San Francisco?
- ✓ 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?
- ✓ Translate spoken Chinese into spoken English in rea
- ✓ Win an art competition?
- Write an intentionally funny story?
- Construct a building?

What Can AI Do?

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

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of tennis?
- ✗ Unload any dishwasher in any home?
- ✓ Drive safely along the highway?
- ? Drive safely along streets of San Franc
- ✓ Buy a week's worth of groceries on th
- ✗ Buy a week's worth of groceries at Be
- ? Discover and prove a new mathematic
- ✗ Perform a surgical operation?
- ✓ Translate spoken Chinese into spoken
- ✓ Win an art competition?
- ✓ Write an intentionally funny story?
- Construct a building?

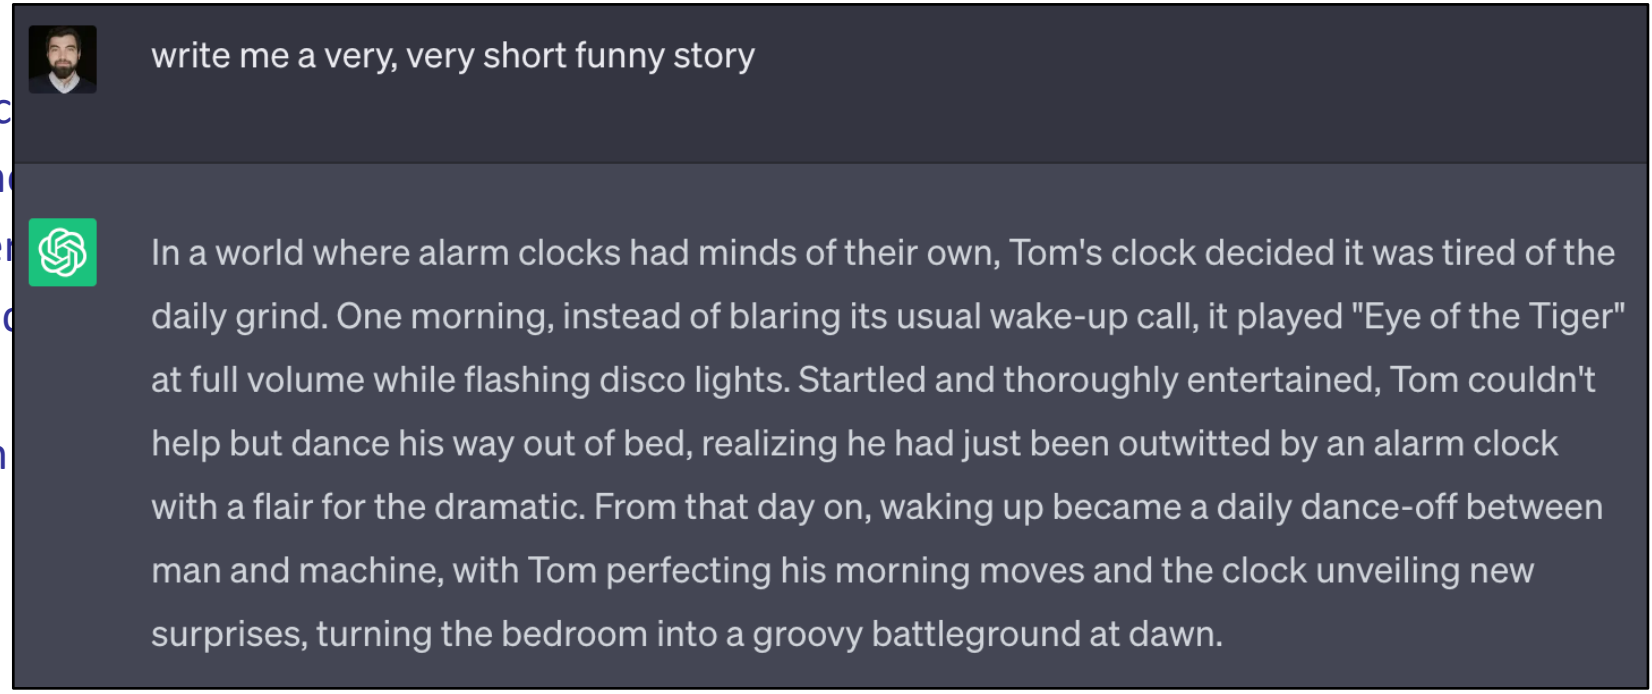

What Can AI Do?

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

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of tennis?
- ✗ Unload any dishwasher in any home?
- ✓ Drive safely along the highway?
- ? Drive safely along streets of San Francisco?
- ✓ 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?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✓ Win an art competition?
- ✓ Write an intentionally funny story?
- ✗ Construct a building?

This Course: 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

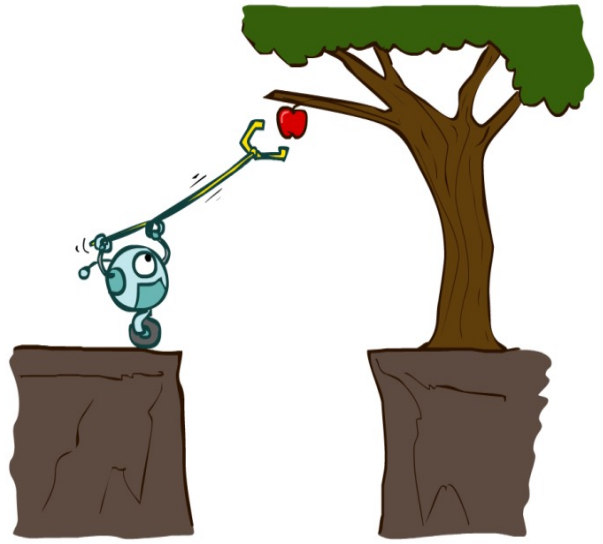

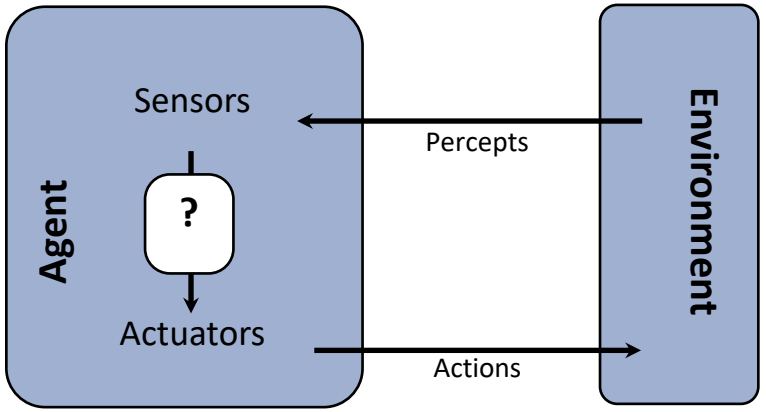

Pac-Man as an Agent

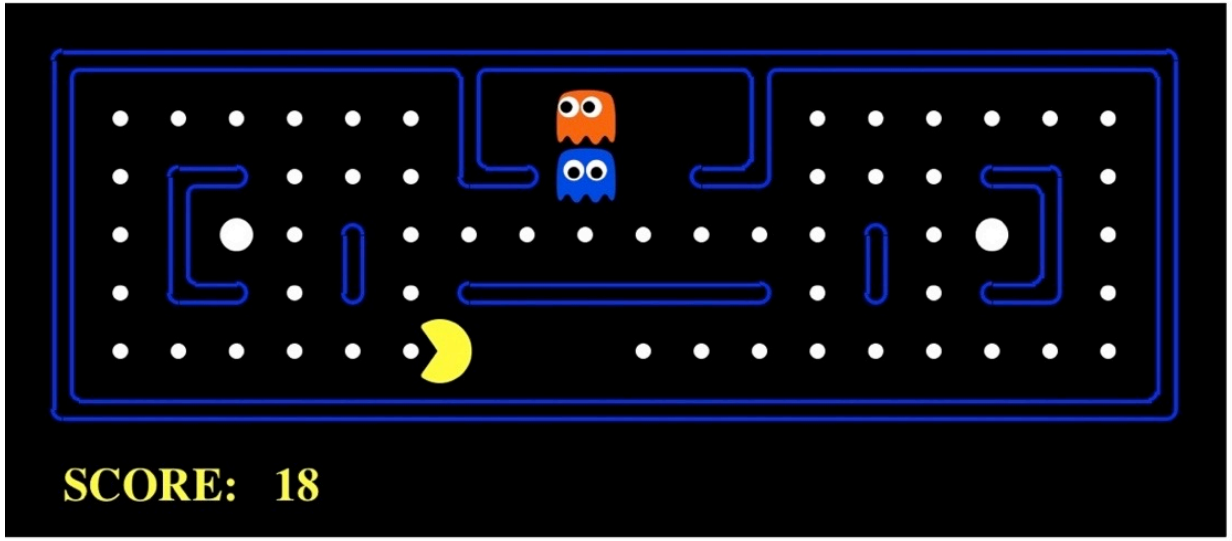

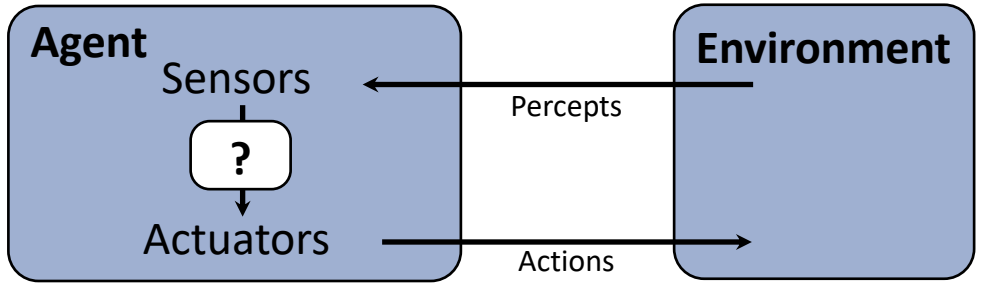

Pac-Man is a registered trademark of Namco-Bandai Games, used here for educational purposes

Course Topics

Core Components of Rational Agents:

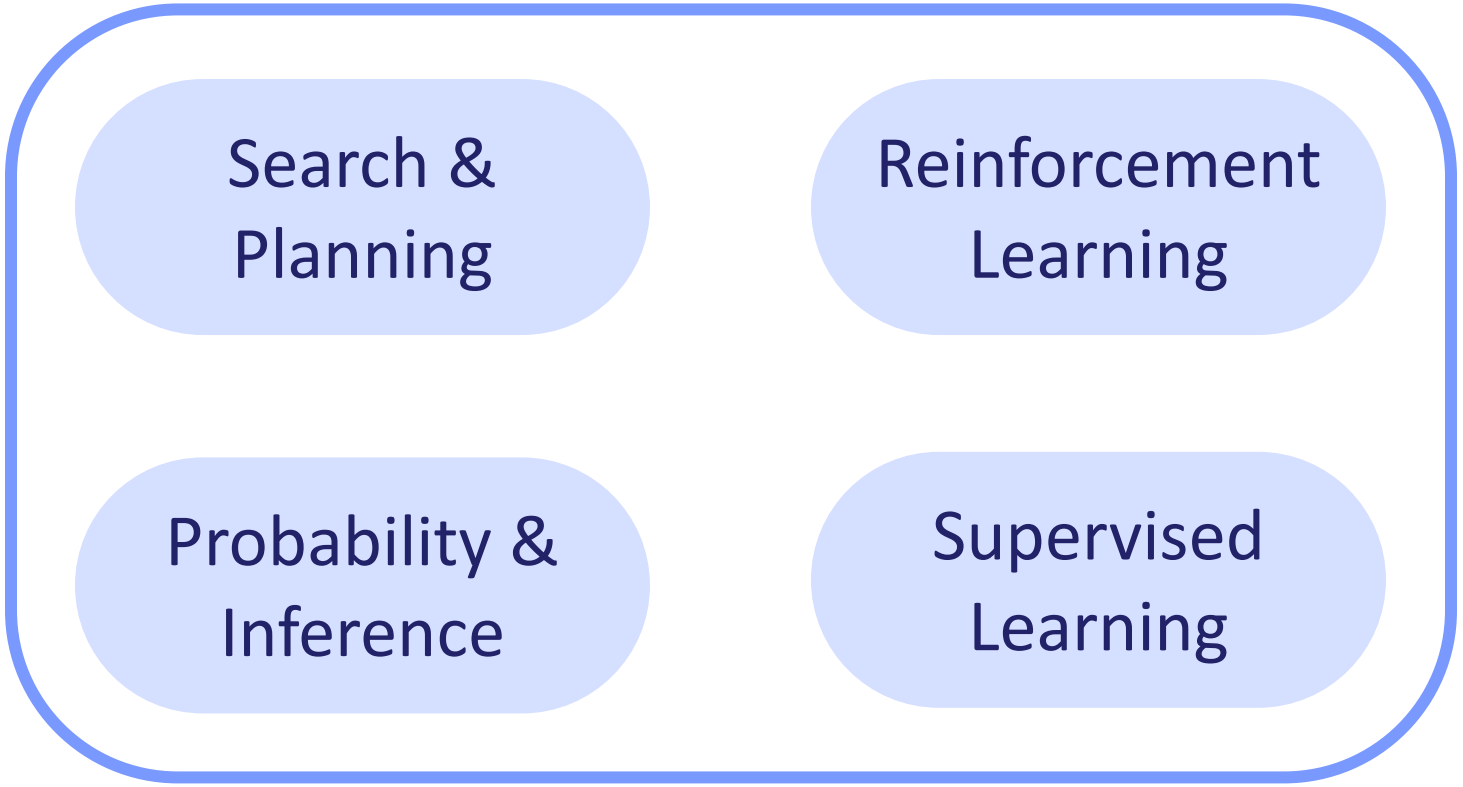

Course Topics

Search & Planning

Reinforcement Learning

Probability & Inference

Supervised Learning

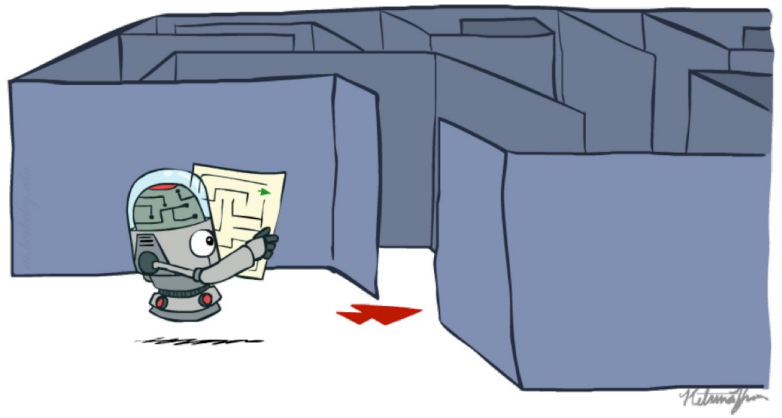

How can I find a ***sequence of best decisions*** for a ***particular*** situation?

Course Topics

Search & Planning

Reinforcement Learning

Probability & Inference

Supervised Learning

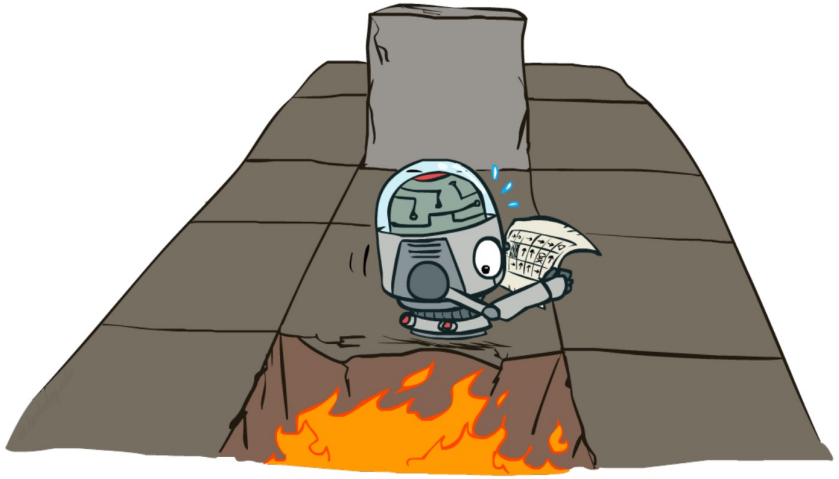

How can I find ***rules*** **(*policy)*** to make best decisions for ***any*** situation?

Course Topics

Search & Planning

Reinforcement Learning

Probability & Inference

Supervised Learning

How can I make sense of ***uncertainty*** in the world?

Course Topics

Search & Planning

Reinforcement Learning

Probability & Inference

Supervised Learning

How can I learn a ***model*** of the world from ***data***?

Course Topics

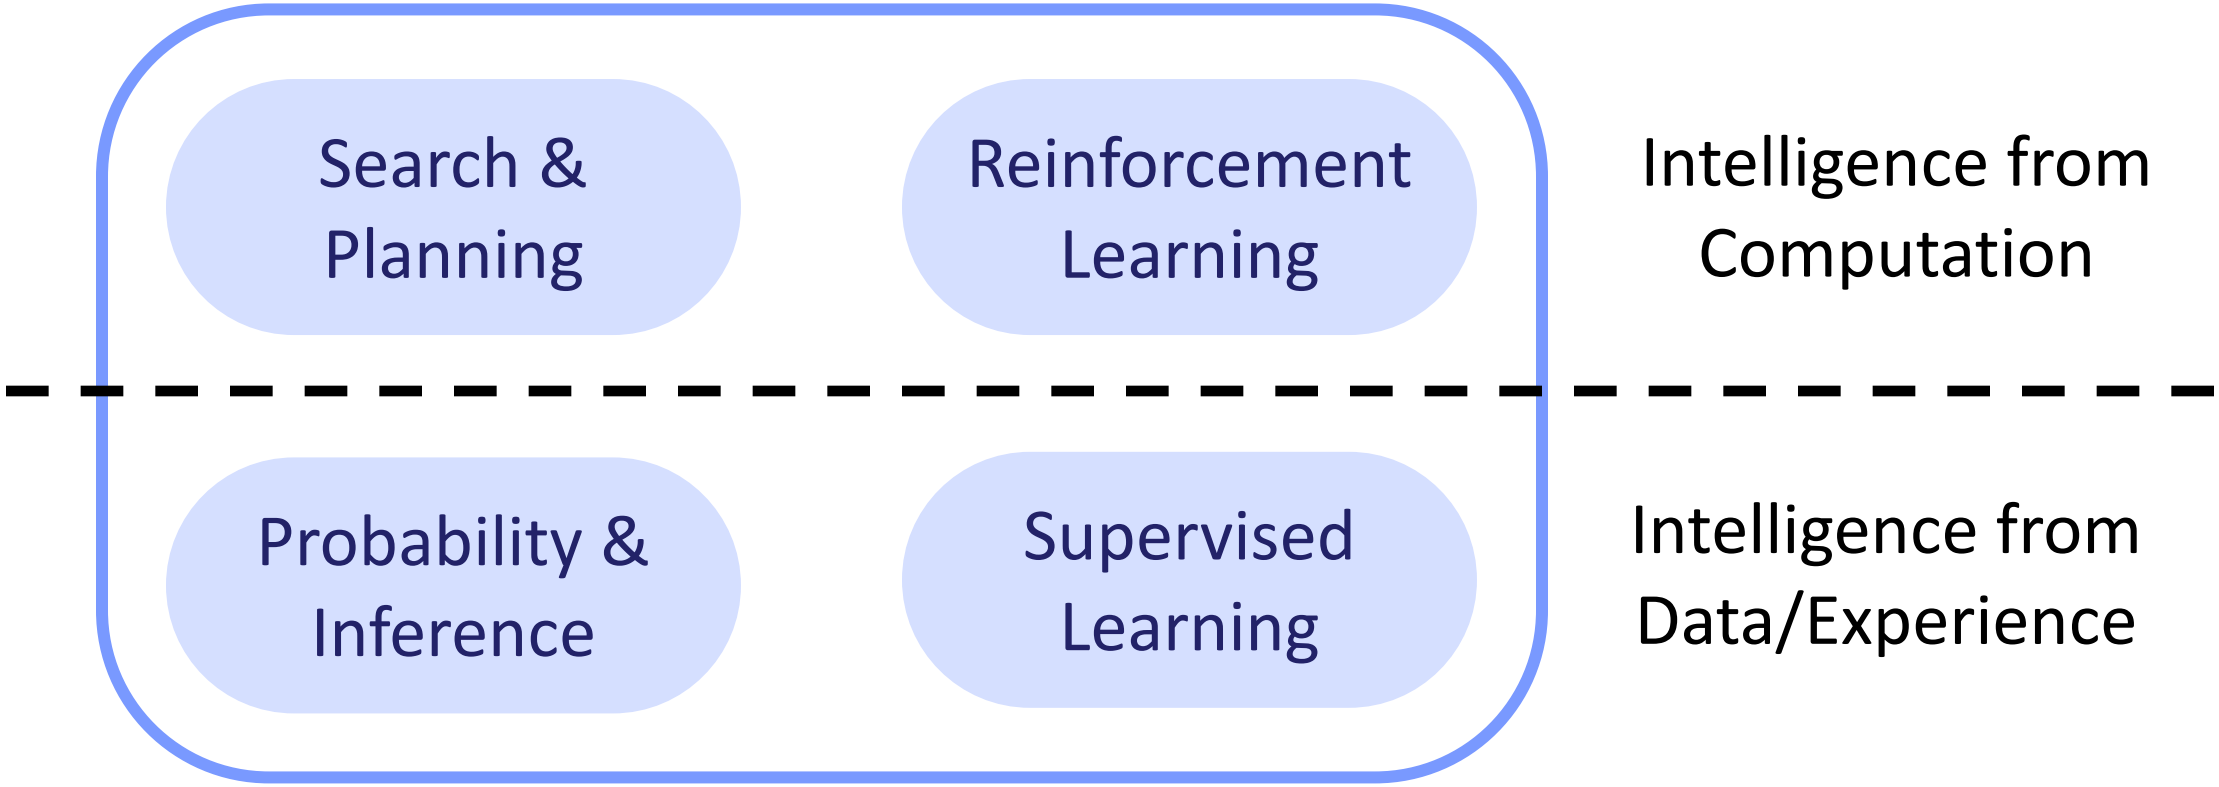

Course Topics

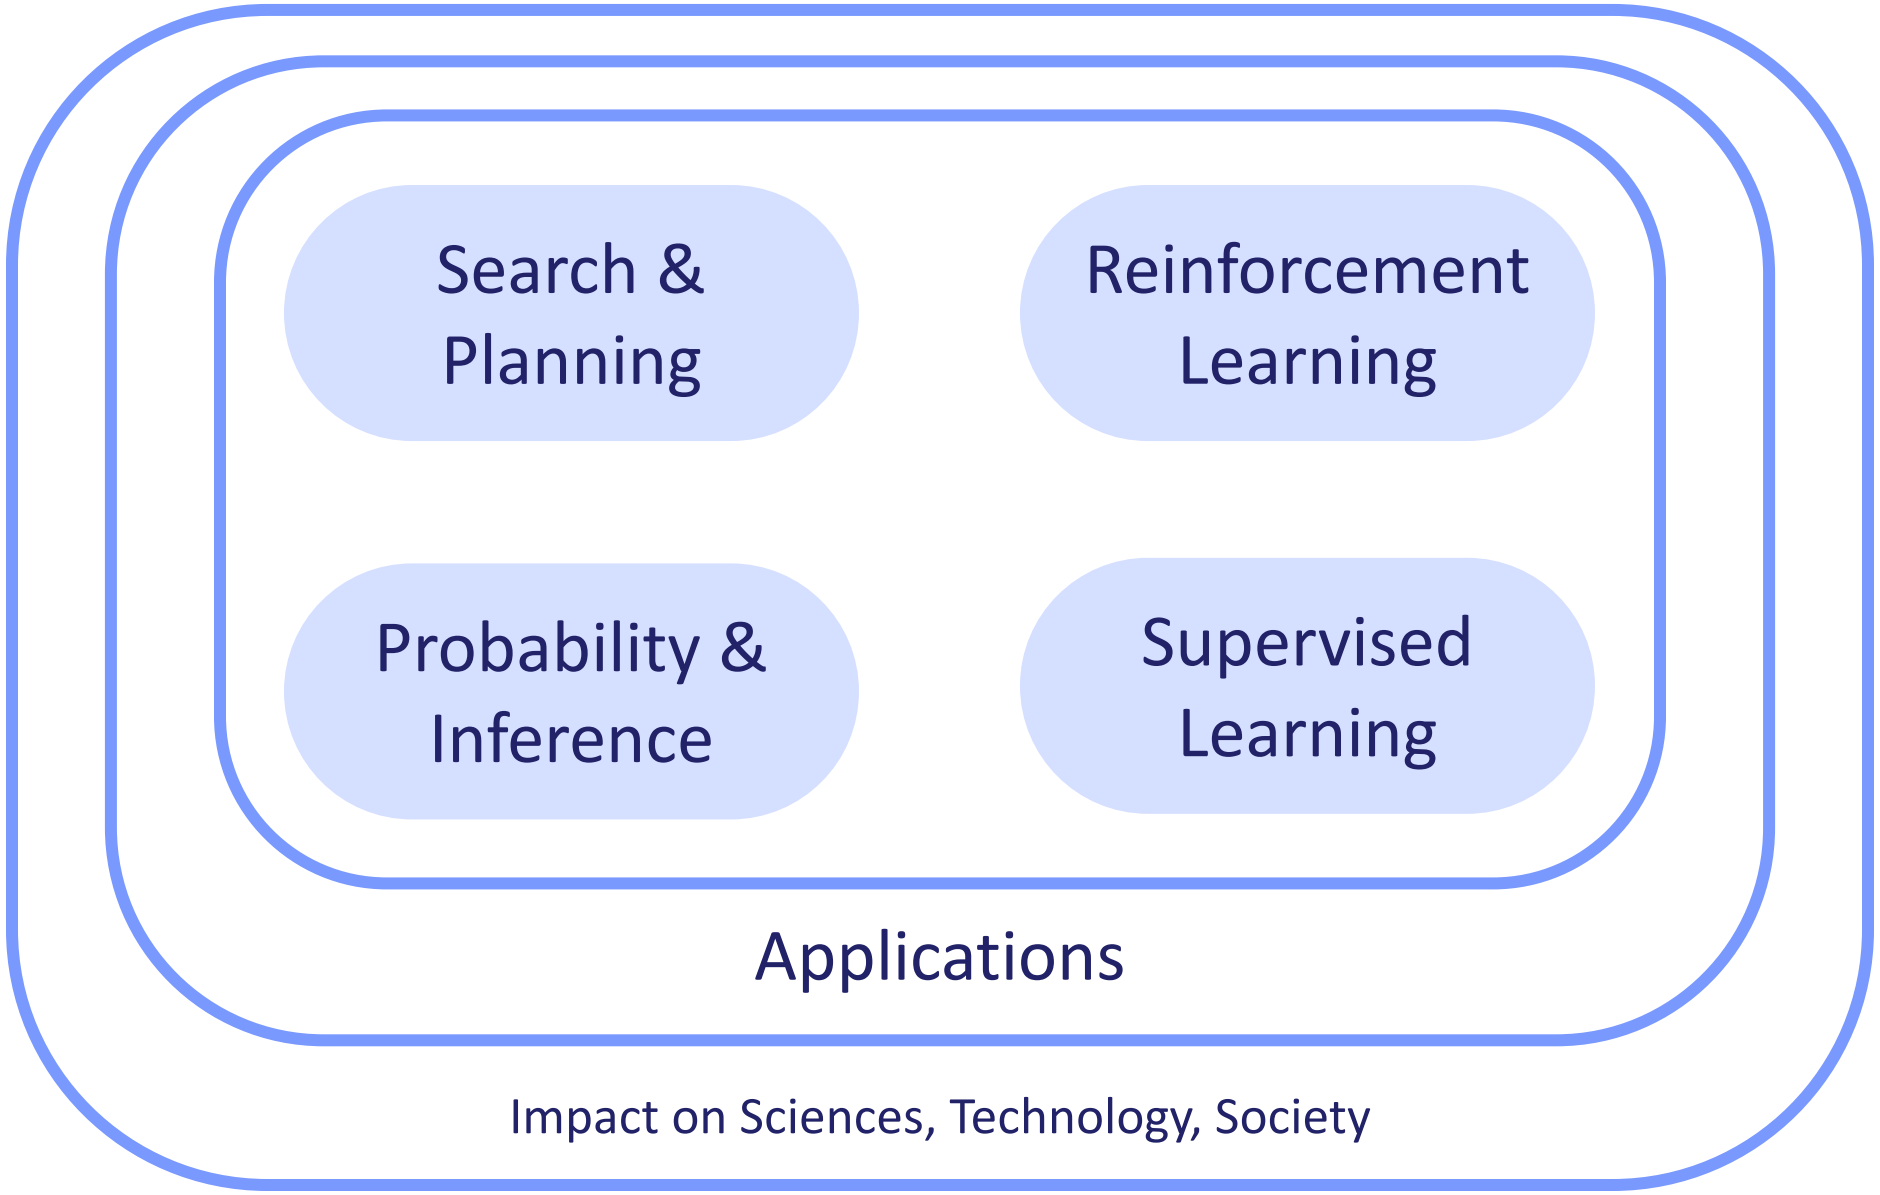

Should I take CS 188?

- 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

By the end of this course you'll:

- Build and understand math of rational, learning agents
- Select and apply the right AI methods for wide range of problems
- Recognize how these methods are used in modern AI systems
- Be prepared to make decisions on how AI is used in society

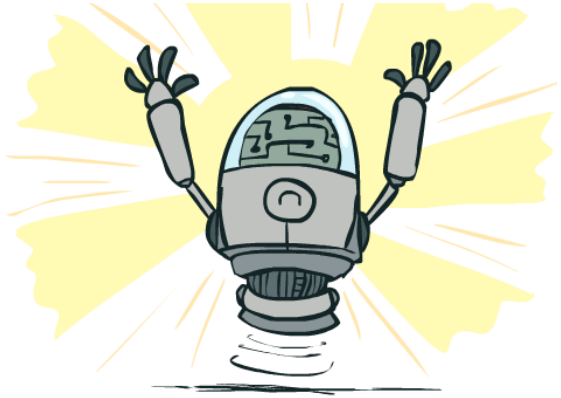
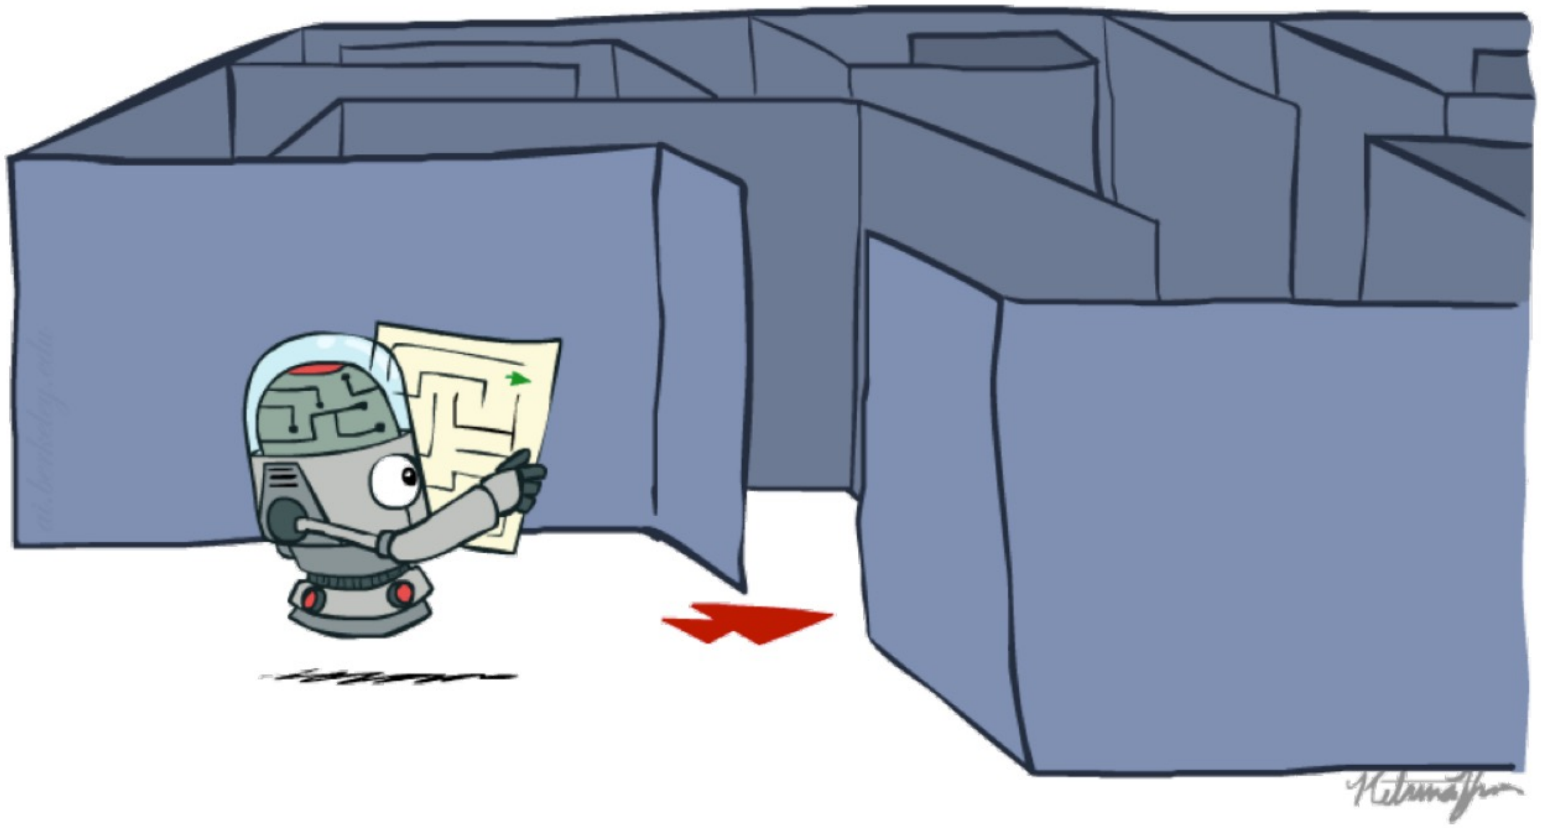
Next Week: Search