Future of Neural Networks and Deep Learning

Danna Gurari
University of Colorado Boulder
Spring 2025

https://dannagurari.colorado.edu/course/neural-networks-and-deep-learning-spring-2025/

Review

- Last lecture:
  - Responsible deep learning activity
- Assignments (Canvas):
  - Final project presentation due in 1.5 weeks
- Next two lectures: Join Zoom (NOT IN-PERSON)
- Questions?

Today’s Topics

- Motivation
- How might future neural networks and deep learning look?
- How might future societies look?
- Future of neural networks and deep learning: what do “experts” say?

Today's Topics

- Motivation
- How might future neural networks and deep learning look?
- How might future societies look?
- Future of neural networks and deep learning: what do "experts" say?

Recall Origins of NN & DL: Began ~70 Years Ago

First programmable machine

Turing test

AI birth

Goal: mimic human abilities, such as speech recognition, natural language processing, and computer vision

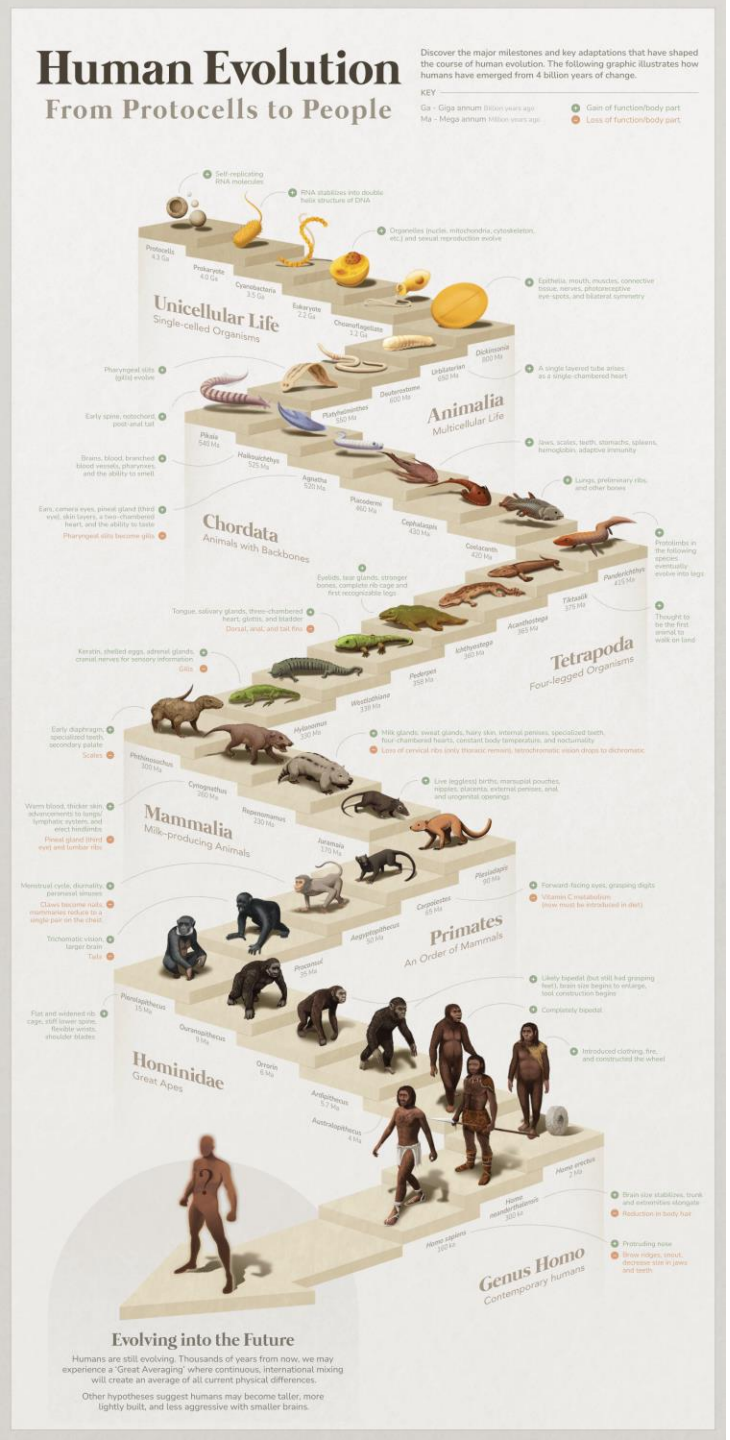

Context: Evolution to Modern Humans Began 4 Billion Years Ago!

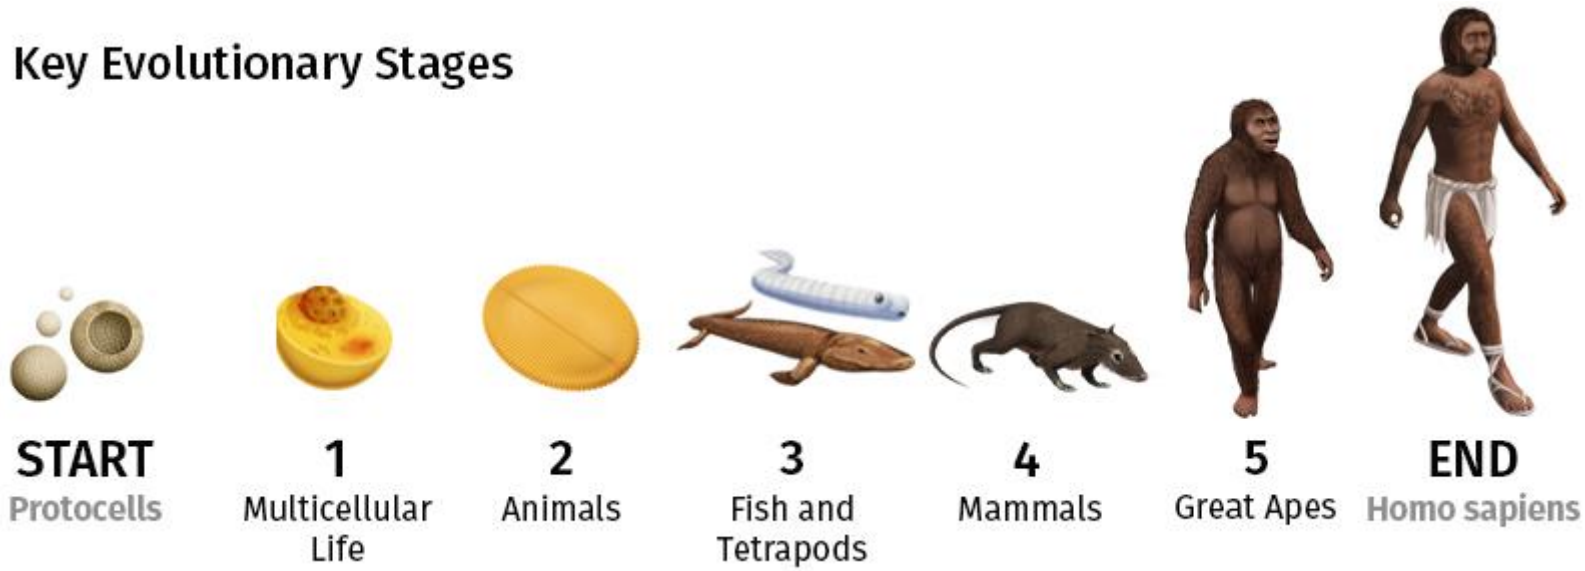

https://www.visualcapitalist.com/path-of-human-evolution/

Throughout the Process, Mammals Have Shaped Their Surroundings; e.g.,

Insects enable new plants to grow through pollination

https://www.britannica.com/topic/agriculture/Early-agricultural-societies

Humans control where and when water flows

https://www.usbr.gov/uc/rm/crsp/gc/

How Will Neural Networks and Deep Learning Evolve and How Will They Shape Our World?

Today's Topics

- Motivation
- How might future neural networks and deep learning look?
- How might future societies look?
- Future of neural networks and deep learning: what do "experts" say?

Biology-Inspired Designs? We Are Actively Learning About How Brains Work

Complete map of fruit fly brain circuitry unveiled

AI tools and hundreds of human proofreaders helped tease apart 140,000 neurons in a brain the size of a grain of sand

2 OCT 2024 · 11:05 AM ET · BY RODRIGO PÉREZ ORTEGA

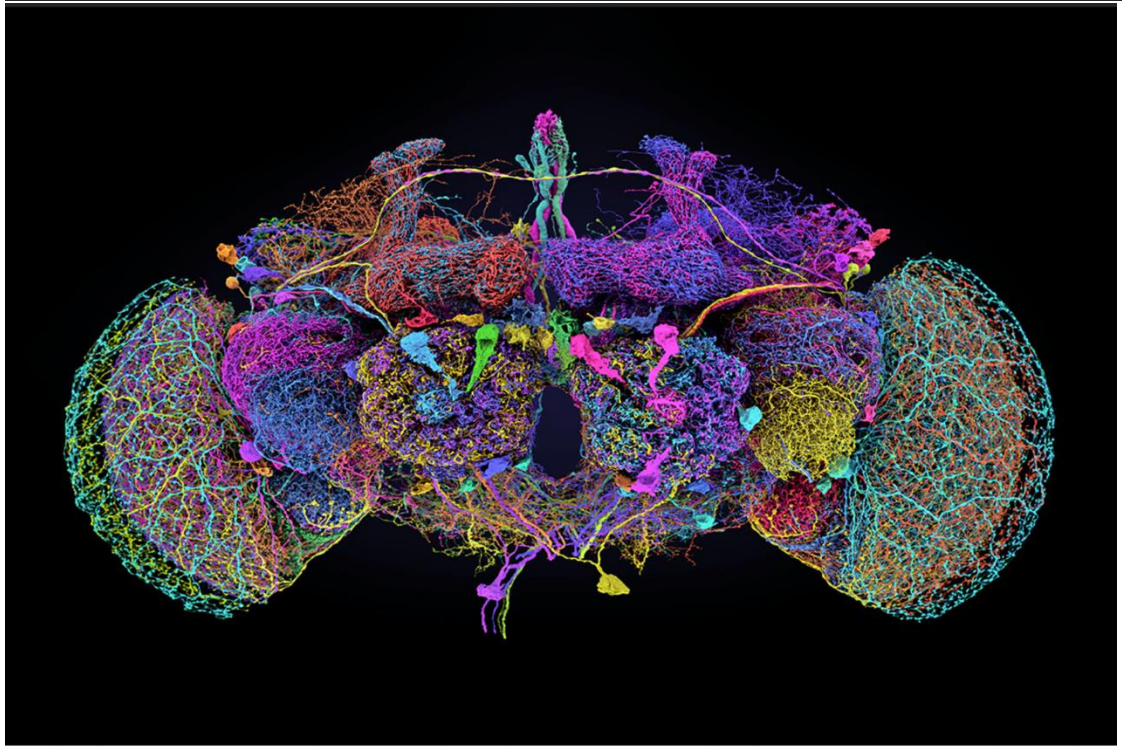

The 50 largest neurons in the fly's brain, shown here, are part of the largest, most complete neuronal map ever produced by scientists. TYLER SLOAN AND AMY STERLING FOR FLYWIRE/PRINCETON UNIVERSITY; DORKENWALD *ET AL.*, *NATURE*, 2024

Browse yourself: https://codex.flywire.ai/app/view_3d

Prior work mapped nervous systems of a fly larva (3,016 neurons) and worm (~300)

Fruit fly has ~140k neurons with 55M connections for an adult animal that can fly, eat, fight, flee, and more!

Future work will likely move to mice and eventually to humans

Biology-Inspired Designs? We Are Actively Learning About How Brains Work

- The computer as a brain that comprises **specialized accelerators**
- **Low power** – the brain consumes only about 20W
- **Fault tolerant** – the brain loses neurons all the time

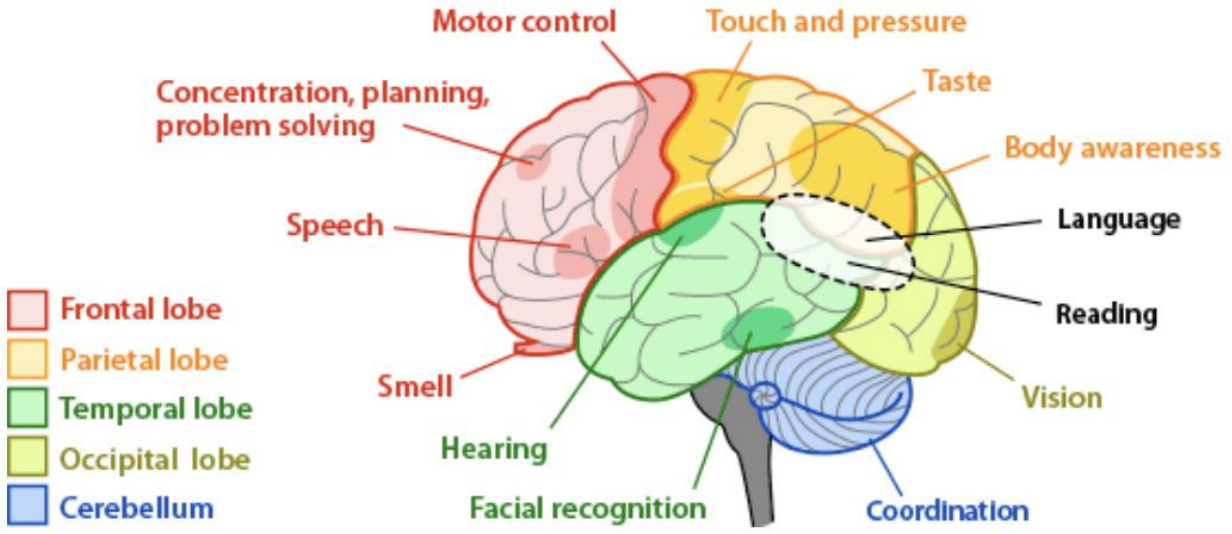

2400 kcal/24 hr = 100 kcal/hr = 27.8 cal/sec = 116.38 J/s = 116 W
20% x 116 W = 23.3 W

Yang, Eric. Think Dinner. *Mac Evolution, 1998*

https://askabiologist.asu.edu/sites/default/files/resources/articles/nervous_journey/brain-regions-areas.gif

Animals vs Neural Networks

- What does an animal need in order to grow and survive?
  - Habitat with sufficient air, water, and food (e.g., rotting food for flies)
- What does a modern NN system need in order to grow and survive?
  - More power than needed to support small cities
  - Massive amounts of data
- What is needed to bridge the gap between animals and NNs? e.g.,
  - Run at exaflop speed (a billion billion calculations per second) with 20W (current models need ~8 MW, or 8 million watts)
  - Simultaneously support many senses/modalities
  - Learn without backpropagation
  - Continuously learn, adapting from and overcoming injuries

Holistic Perspective Development: HW to SW

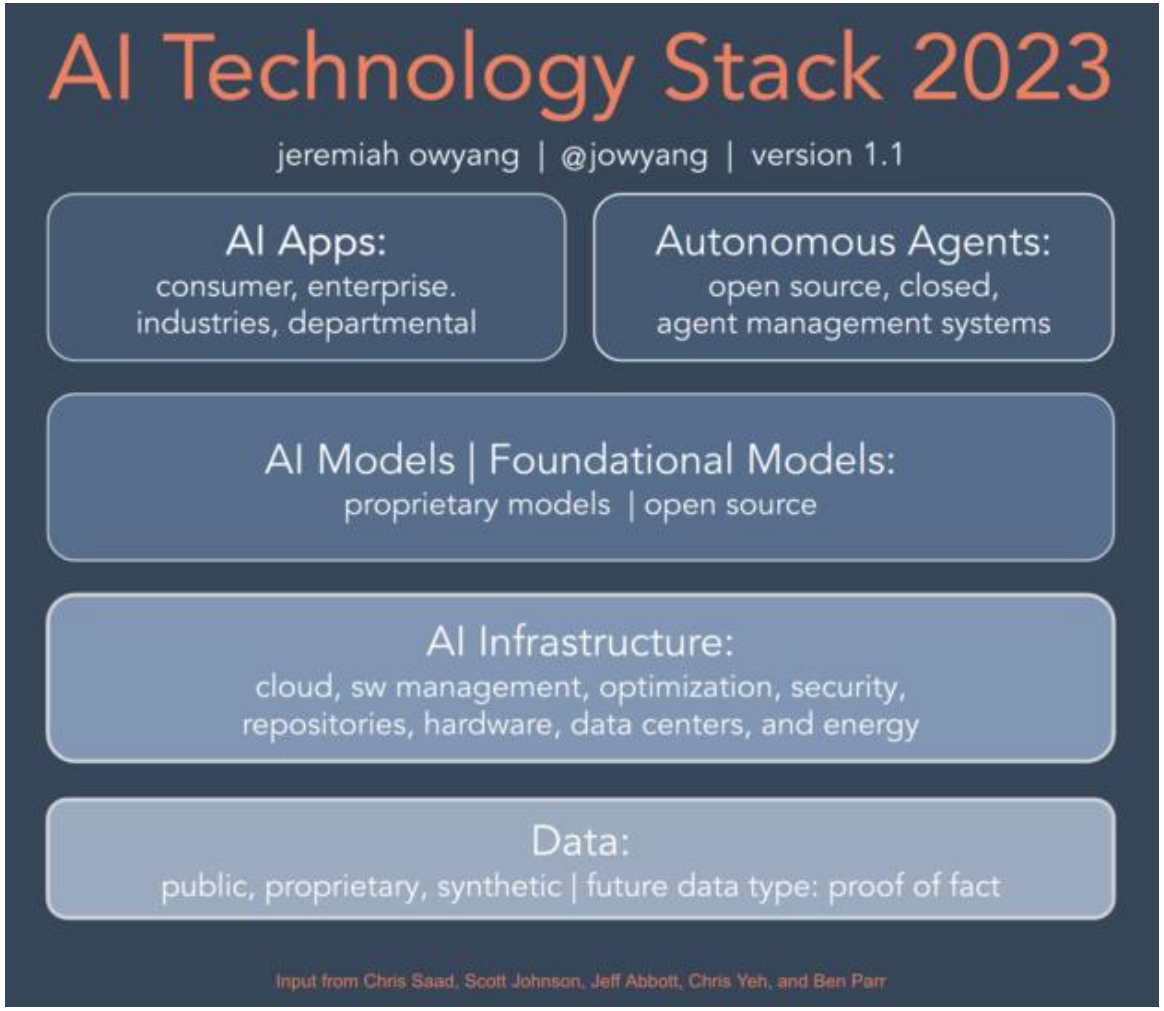

We discussed in the past few lecture ways models will improve via SW, improving their learning and efficiency

We will also see improvements to HW, potentially with:

- Quantum computing for faster computation
- Memristors that preserve a non-binary state (e.g., analog) after current/voltage passes through it, and so simultaneously support processing and storage (current devices use different modules)

Notably, it is faster to make changes in SW than HW

https://web-strategist.com/blog/2023/06/12/ai-technology-stack-2023-v1-1/

Open-Source vs Closed-Source Models

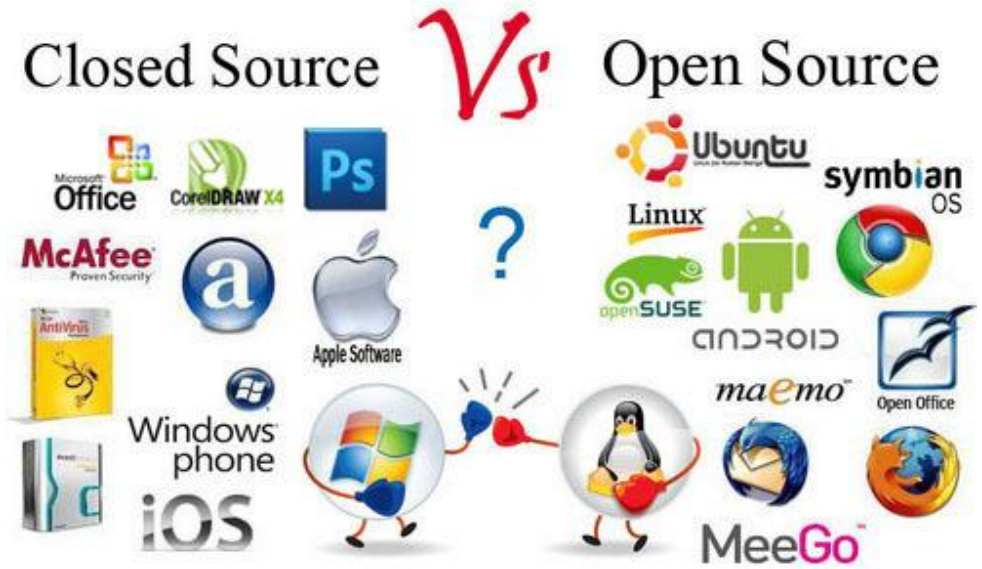

https://medium.com/@mccccc/is-open-source-software-more-reliable-or-secure-than-closed-source-software-5470732d7b90

May 4, 2023

Google "We Have No Moat, And Neither Does OpenAI" // Leaked Internal Google Document Claims Open Source AI Will Outcompete Google and OpenAI

- Recommended article: https://semianalysis.com/2023/05/04/google-we-have-no-moat-and-neither/
- Do you think AI should be open-source or closed-source?

What advice would you provide to companies (who may even be your future employers ☺)?

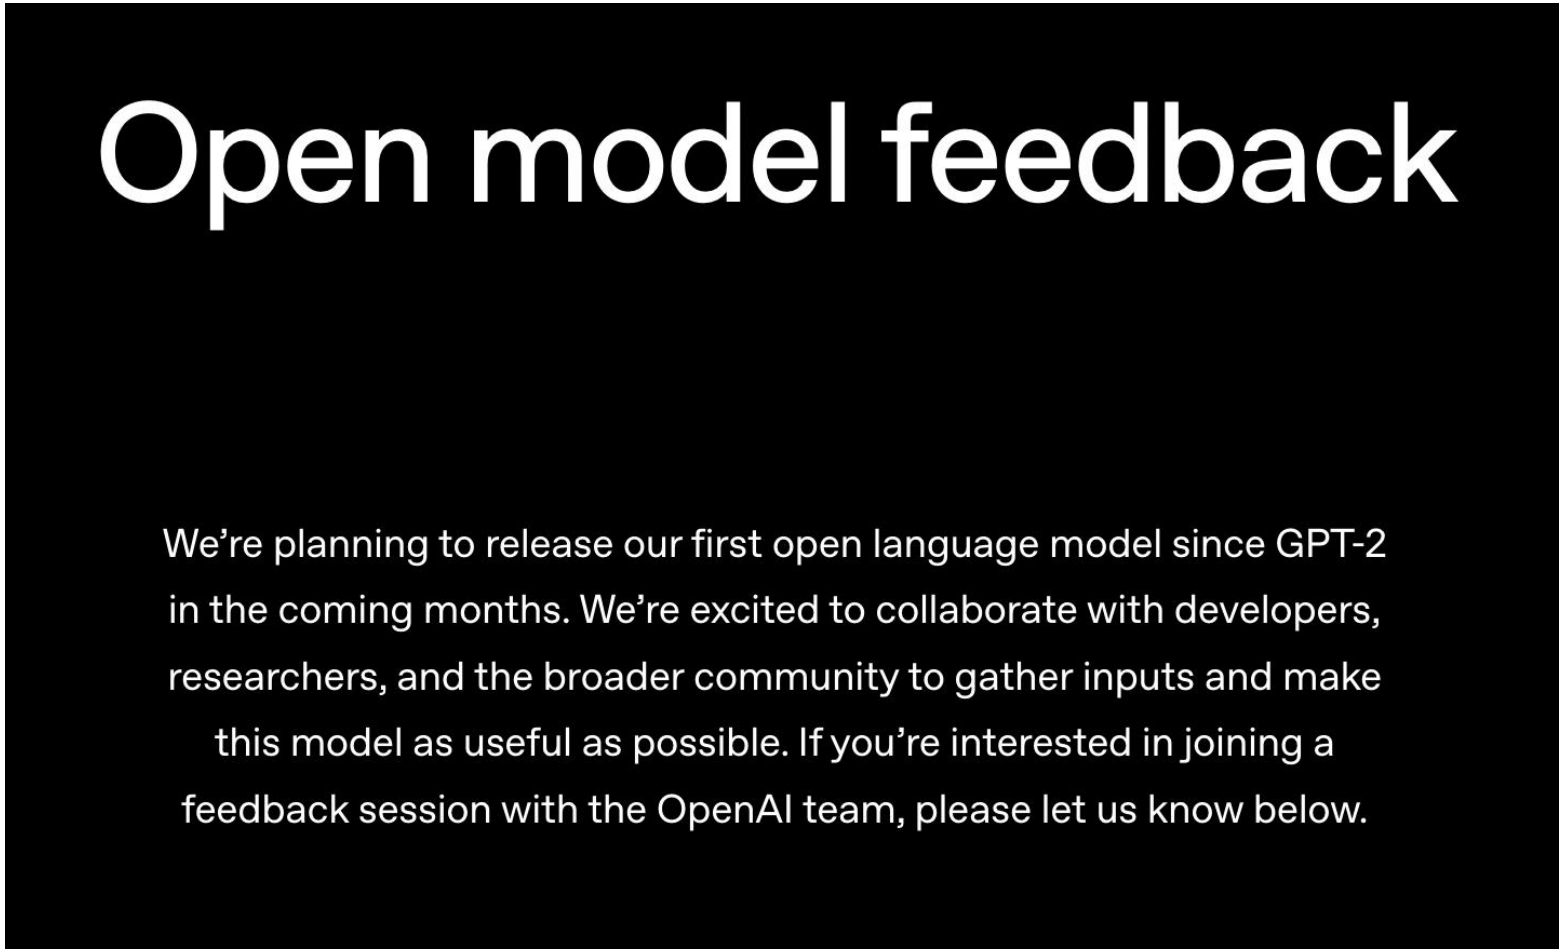

https://openai.com/open-model-feedback/

Today's Topics

- Motivation
- How might future neural networks and deep learning look?
- How might future societies look?
- Future of neural networks and deep learning: what do "experts" say?

How Might We Interact with Future Models?

- Prompts aren't always natural interactions

- What are more natural/desired ways to interface with machines?
  - Gestures
  - Physically (e.g., tactiley pointing to a screen or object)
  - Brain-computer interfaces
  - Combinations of multiple interfaces
  - Constraint-based graphical user interfaces (e.g., menus)

https://cacm.acm.org/opinion/prompting-considered-harmful/

When Will We Interact With Future Models?

- *Eventually* nearly every domain of life may be touched, but companies are still struggling how to leverage/monetize them

Co-existence with Future Models Will Change How We Think/Behave, But How?

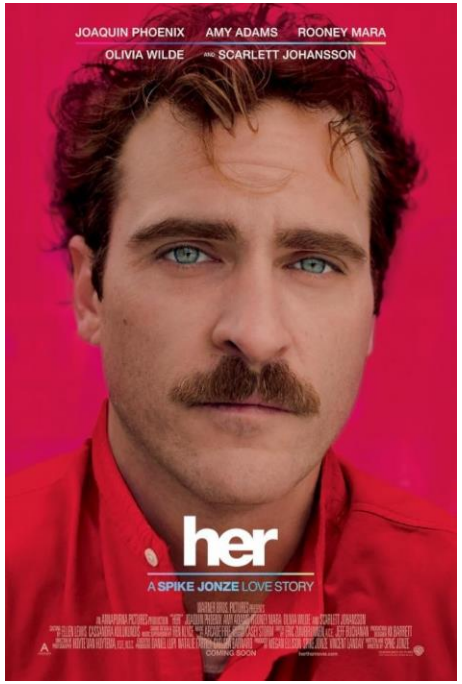

A love story

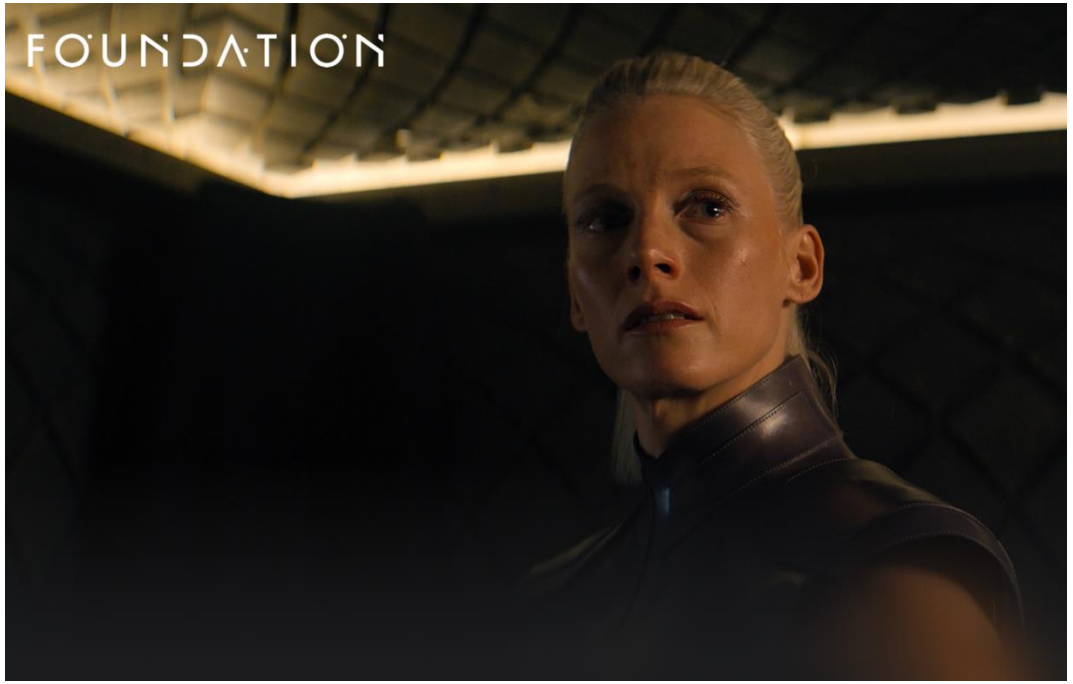

A protective, violent bodyguard

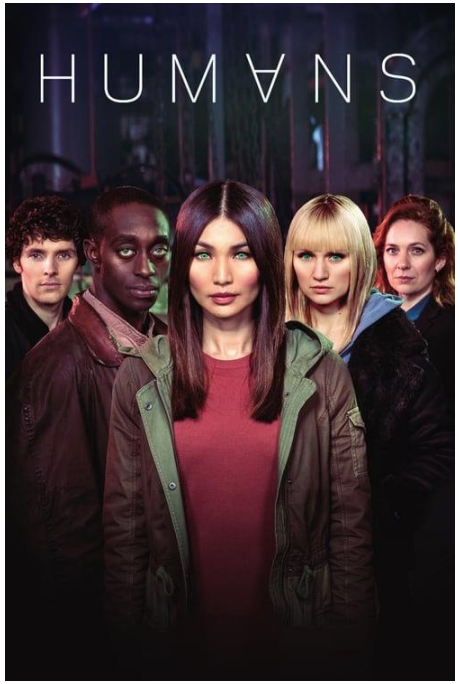

Always-present “slaves”

Many great science fiction movies, tv shows, and books explore possibilities!

Co-existence with Future Models Will Change How We Think/Behave, But How?

Singularity: point at which AI has abilities beyond human control, potentially even surpassing human intelligence:

Do you think the singularity will ever occur?

If so, will the singularity be an existential threat to humans?

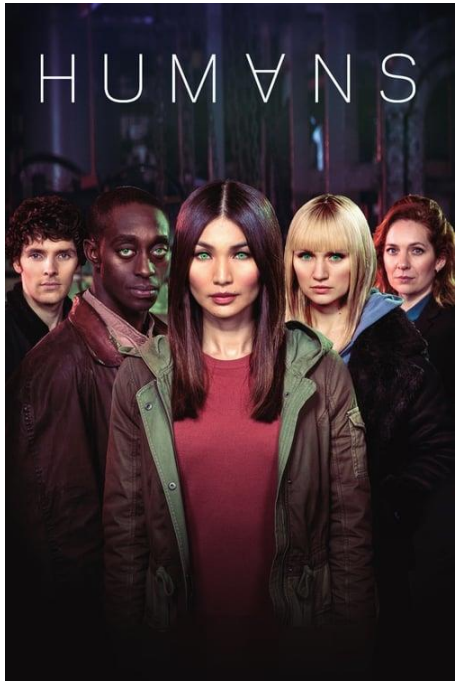

Always-present “slaves”

Today's Topics

- Motivation
- How might future neural networks and deep learning look?
- How might future societies look?
- Future of neural networks and deep learning: what do "experts" say?

Individuals

Dario Amodei

Machines of Loving Grace[1]

How AI Could Transform the World for the Better

October 2024

I think and talk a lot about the risks of powerful AI. The company I'm the CEO of, Anthropic, does a lot of research on how to reduce these risks. Because of this, people sometimes draw the conclusion that I'm a pessimist or "doomer" who thinks AI will be mostly bad or dangerous. I don't think that at all. In fact, one of my main reasons for focusing on risks is that they're the only thing standing between us and what I see as a fundamentally positive future. **I think that most people are underestimating just how radical the upside of AI could be**, just as I think most people are underestimating how bad the risks could be.

SITUATIONAL AWARENESS

The Decade Ahead

Introduction I. From GPT-4 to AGI: Counting the OOMs II. From AGI to Superintelligence: the Intelligence Explosion

IIIa. Racing to the Trillion-Dollar Cluster IIIb. Lock Down the Labs: Security for AGI IIIc. Superalignment IIId. The Free World Must Prevail

IV. The Project V. Parting Thoughts Full series as PDF About Dwarkesh podcast

(Leopold Aschenbrenner, June 2024)

"By 2025/26, these machines will outpace many college graduates. By the end of the decade, they will be smarter than you or I; we will have superintelligence, in the true sense of the word."

From 5 Members
(Led by Concerned Former OpenAI Employee)

AI Futures Project

About Blog

The **AI Futures Project** is a small research group forecasting the future of AI, funded by charitable donations and grants. (More about us.)

Our work includes:

- **AI 2027**. Daniel Kokotajlo, Scott Alexander, Thomas Larsen, Eli Lifland, Romeo Dean. April 3, 2025.
- Compute Forecast. Romeo Dean. April 3, 2025.
- Timelines Forecast. Nikola Jurkovic, Eli Lifland. April 3, 2025.
- Takeoff Forecast. Daniel Kokotajlo, Eli Lifland. April 3, 2025.
- AI Goals Forecast. Daniel Kokotajlo. April 3, 2025.
- Security Forecast. Romeo Dean. April 3, 2025.
- 4 Ways to Advance Transparency in Frontier AI Development. Dean Ball & Daniel Kokotajlo. *Time*, October 15, 2024.
- Our Tabletop Exercise, which we've facilitated for 200+ people, including experts and policymakers.

To stay up to date with our work, **subscribe to our blog** and **follow us on X**.

Summary of interview with Kokotajlo:
"… sometime in early 2027, if current trends hold, A.I. will be a superhuman coder. Then, by mid-2027, it will be a superhuman A.I. researcher — an autonomous agent that can oversee teams of A.I. coders and make new discoveries. Then, in late 2027 or early 2028, it will become a super*intelligent* A.I. researcher — a machine intelligence that knows more than we do about building advanced A.I., and can automate its own research and development, essentially building smarter versions of itself. From there, he said, it's a short hop to artificial superintelligence, or A.S.I., at which point all bets are off."

-https://www.nytimes.com/2025/04/03/technology/ai-futures-project-ai-2027.html

From 22 Scientists

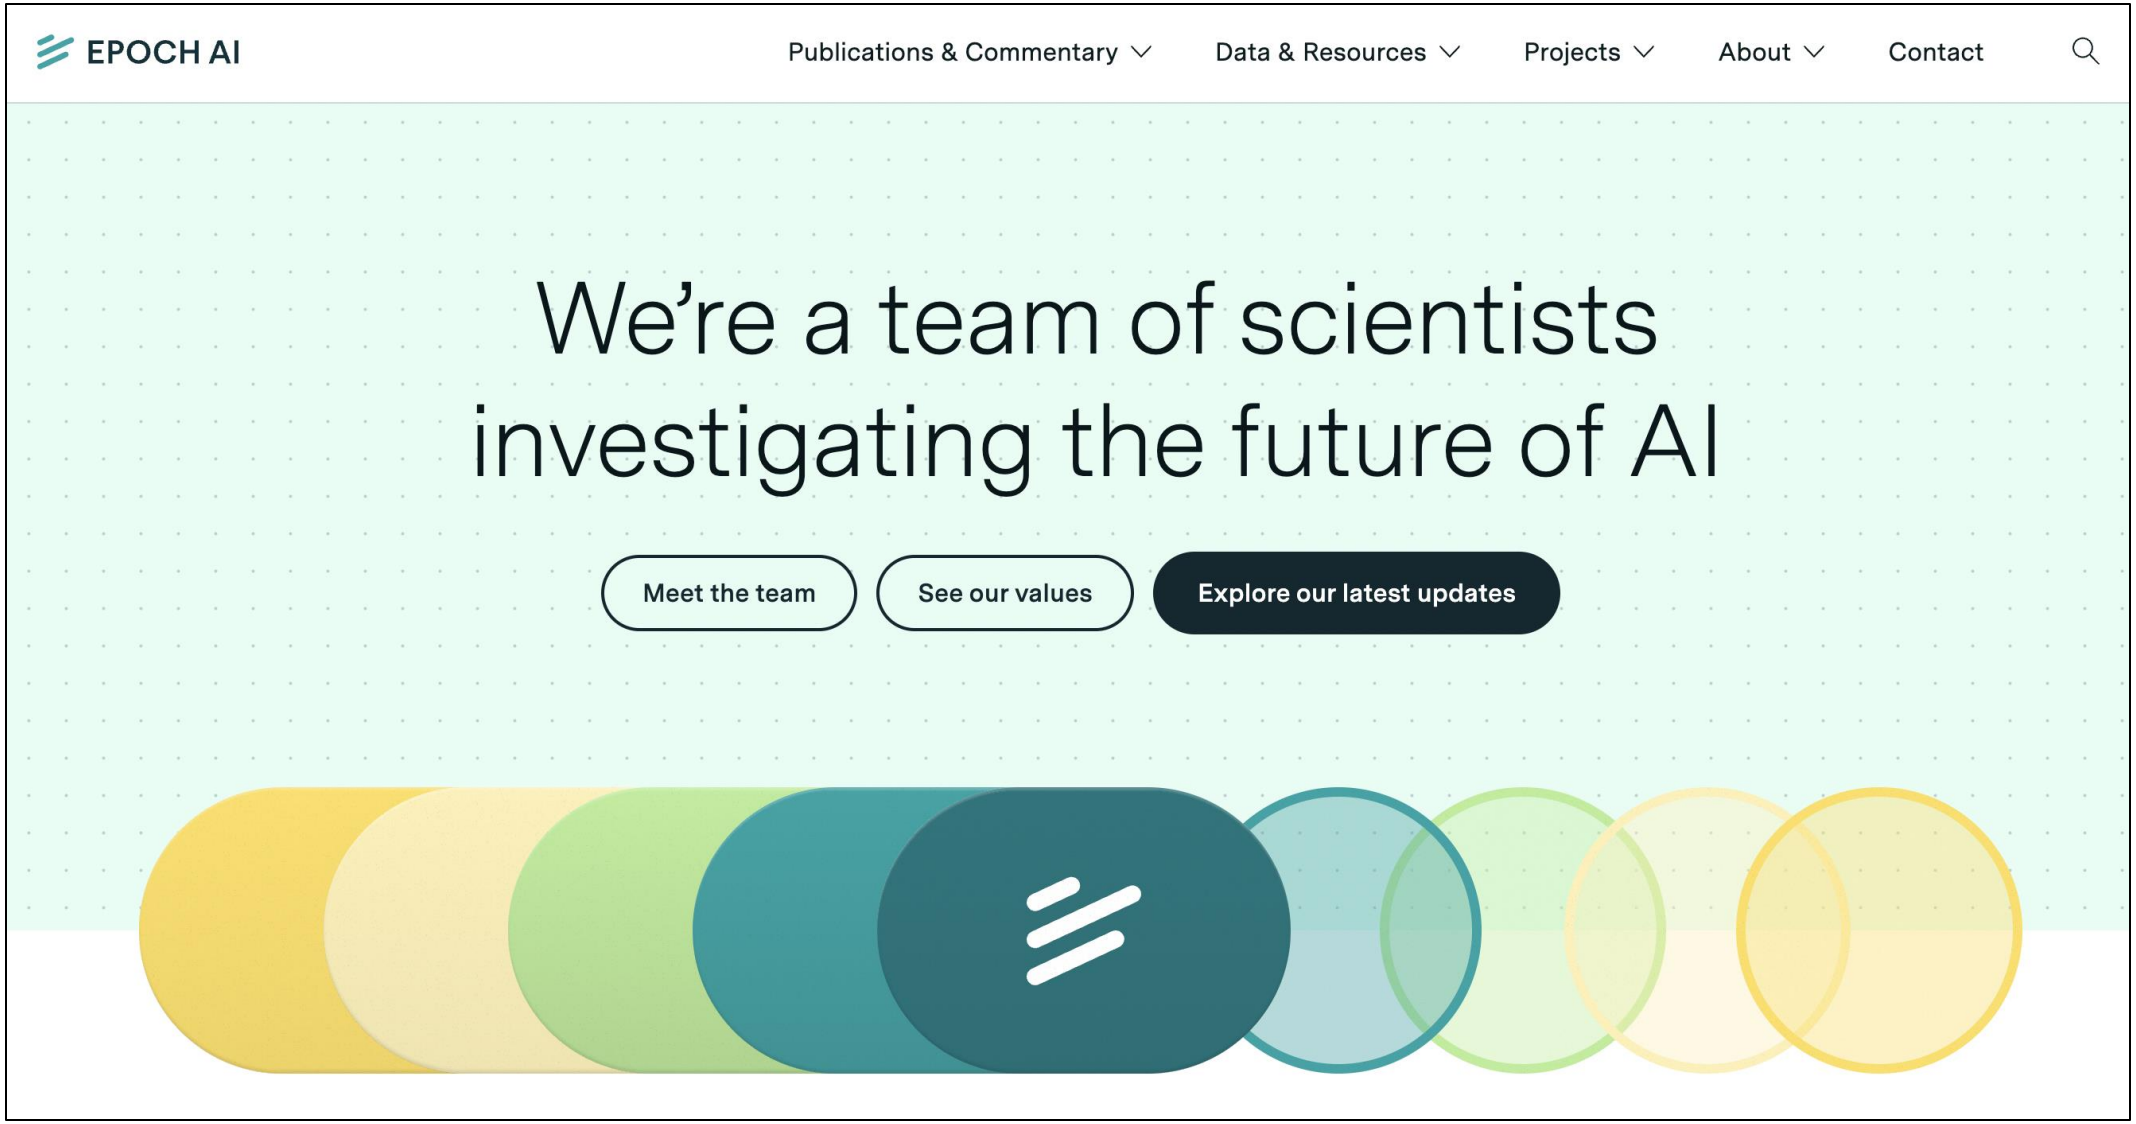

From ~30 Researchers

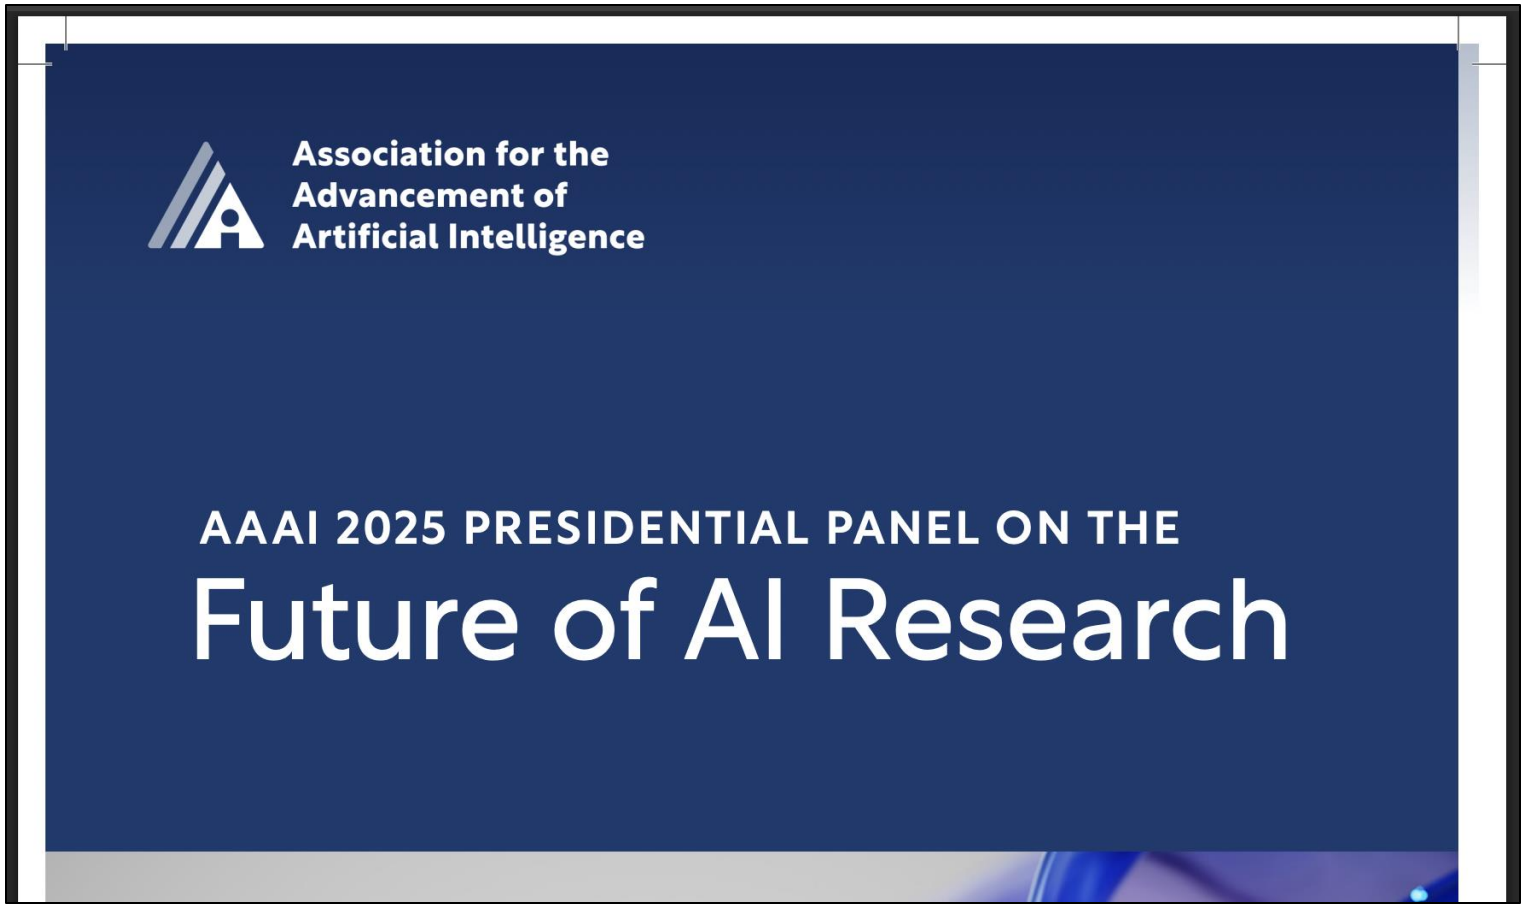

Table of Contents

Class Activity

- When did you start leveraging foundation/frontier models?
  - Never
  - Between 0 and 12 months ago
  - Between 12 months and 2 years ago
  - Between 2 years ago and 5 years ago

- Create a small group, discuss, and add your responses to a shared document:
  - How has your thinking or behavior changed since partnering with foundation/frontier models?
  - How do/can you differentiate yourself from modern computers powered by NNs in the workforce?
  - What will your professional role look like in 1, 5, and 10 years, given continual advancements in NNs?

Today’s Topics

- Motivation
- How might future neural networks and deep learning look?
- How might future societies look?
- Future of neural networks and deep learning: what do “experts” say?


The End