

Network of the National Library of  
Medicine's (NNLM) Region 4

# Understanding & using generative AI

A workshop for health science librarians

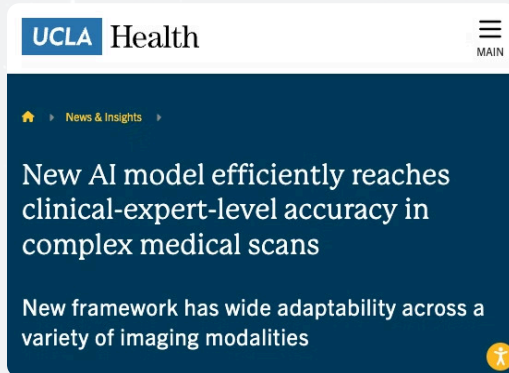
[Nicole Hennig](#)

November 12, 2024





# Generative AI news is everywhere



UCLA Health

News & Insights

New AI model efficiently reaches clinical-expert-level accuracy in complex medical scans

New framework has wide adaptability across a variety of imaging modalities

A screenshot of the UCLA Health website showing a news article. The header includes the UCLA Health logo and a 'MAIN' menu icon. Below the header is a dark blue banner with white text. The main text of the article is also in white on a dark background.

[New AI model efficiently reaches clinical-expert-level accuracy in complex medical scans](#)

- UCLA Health



HARVARD medicine

All Articles

October 2024

## How Generative AI Is Transforming Medical Education

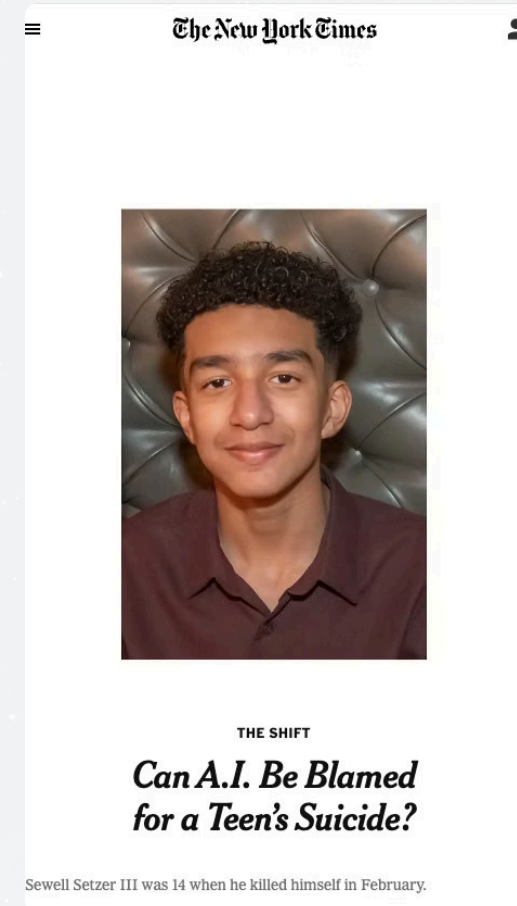
Harvard Medical School is building artificial intelligence into the curriculum to train the next generation of doctors

Autumn 2024


A screenshot of the Harvard Medicine website. It shows a navigation bar with the Harvard Medicine logo and a menu icon. Below is a breadcrumb trail 'All Articles' and a date 'October 2024'. The main headline is 'How Generative AI Is Transforming Medical Education'. A sub-headline reads 'Harvard Medical School is building artificial intelligence into the curriculum to train the next generation of doctors'. At the bottom, it says 'Autumn 2024'.

[How Generative AI Is Transforming Medical Education](#)

- Harvard Medicine



The New York Times



THE SHIFT

## Can A.I. Be Blamed for a Teen's Suicide?

Sewell Setzer III was 14 when he killed himself in February.

A screenshot of The New York Times website. It features the newspaper's logo at the top. Below is a large portrait of a young man, Sewell Setzer III. Underneath the photo is the text 'THE SHIFT' followed by the article title 'Can A.I. Be Blamed for a Teen's Suicide?'. A small caption below the title reads 'Sewell Setzer III was 14 when he killed himself in February.'

[Can AI be blamed for a teen's suicide?](#)

- New York Times



# It can feel overwhelming at times.

**Experts don't seem to agree.**

AI will cure cancer and solve climate change.

No, it will destroy humanity!

No, it won't destroy humanity, but it will lead to a flood of disinformation!



**Developing AI literacy can help you  
cut through the hype.**



**Knowing what's possible with AI can help you decide what's practical and ethical for your use cases.**



# This webinar is a preview of an upcoming online course

Begins Feb. 7, 2025  
(6 weeks)

- We're recording this session.
- Link to the slides:  
<https://bit.ly/hennig-genAI>
- Add your questions to the chat,  
we'll answer at the end.





# Topics for each week

1. Understanding the technology
2. Prompting for language models
3. Using multimodal features (computer vision, etc)
4. Multimedia generation
5. Staying current and avoiding the hype
6. Applications of generative AI in healthcare and medicine

Ethical questions will be covered in each unit (privacy, bias, copyright).





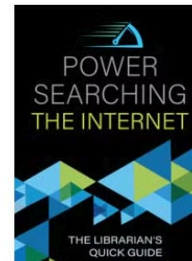
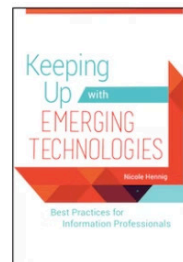
# About me

Worked at the **MIT Libraries** for 14 years, first as webmaster, then head of UX department

**2013 - 2018**

**My own business: technology training for librarians.**

- [books](#)
- [webinars](#)
- [online courses](#)





2019 to the present

# E-learning developer at University of Arizona

I collaborate with my colleagues to create online tutorials.



## What Kind of Researcher Are You?

(2 min)

Helps you decide which tutorial to start with.

Fun personality quiz



## What Types of Sources Do I Need?

(16 min)

★★★★★ Comments

Average rating 4.47 of 5 (2,189 responses)

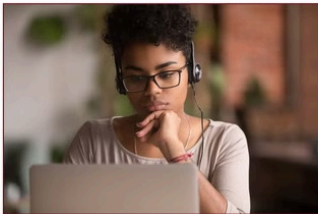


## How Do I Create a Search Strategy?

(11 min)

★★★★★ Comments

Average rating 4.59 of 5 (2,970 responses)



## Should I Rethink My Search?

(18 min)

★★★★★ Comments

Average rating 4.53 of 5 (750 responses)



## How Do I Evaluate Online Information?

(19 min)

★★★★★ Comments

Average rating 4.46 of 5 (2,473 responses)



## How Do I Give Credit to the Ideas of Others?

(15 min)

★★★★★ Comments

Average rating 4.51 of 5 (2,914 responses)

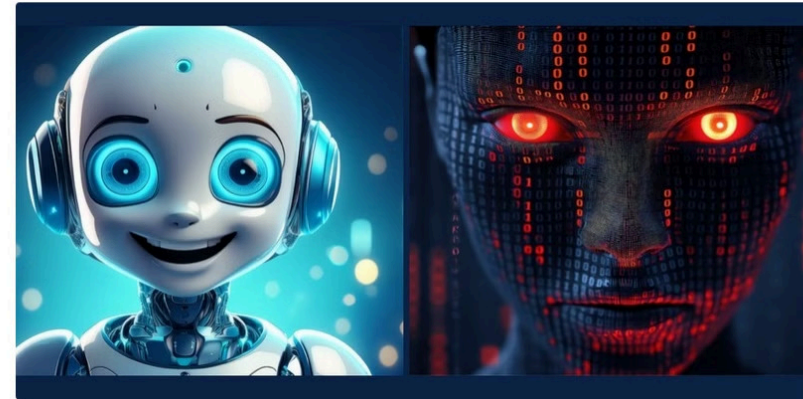


# Course set-up

- **Introduce yourself**
- **What's your AI-adoption personality?**  
Complete a fun personality quiz to find out your "AI-adoption personality."
- **Set up free accounts** on ChatGPT, Claude, Perplexity, Elicit, and NotebookLM.  
(others are optional)
- Fill out a brief survey asking about your experience with generative AI.

## What's your AI-adoption personality?

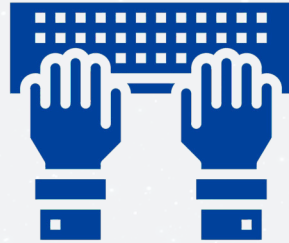
Complete this quiz to find out.



# Format for each week



Several short  
video lectures



Hands-on activities  
(with discussion)



Readings  
(with discussion)

**Expected time to complete each week: about 2 or 2.5 hours**



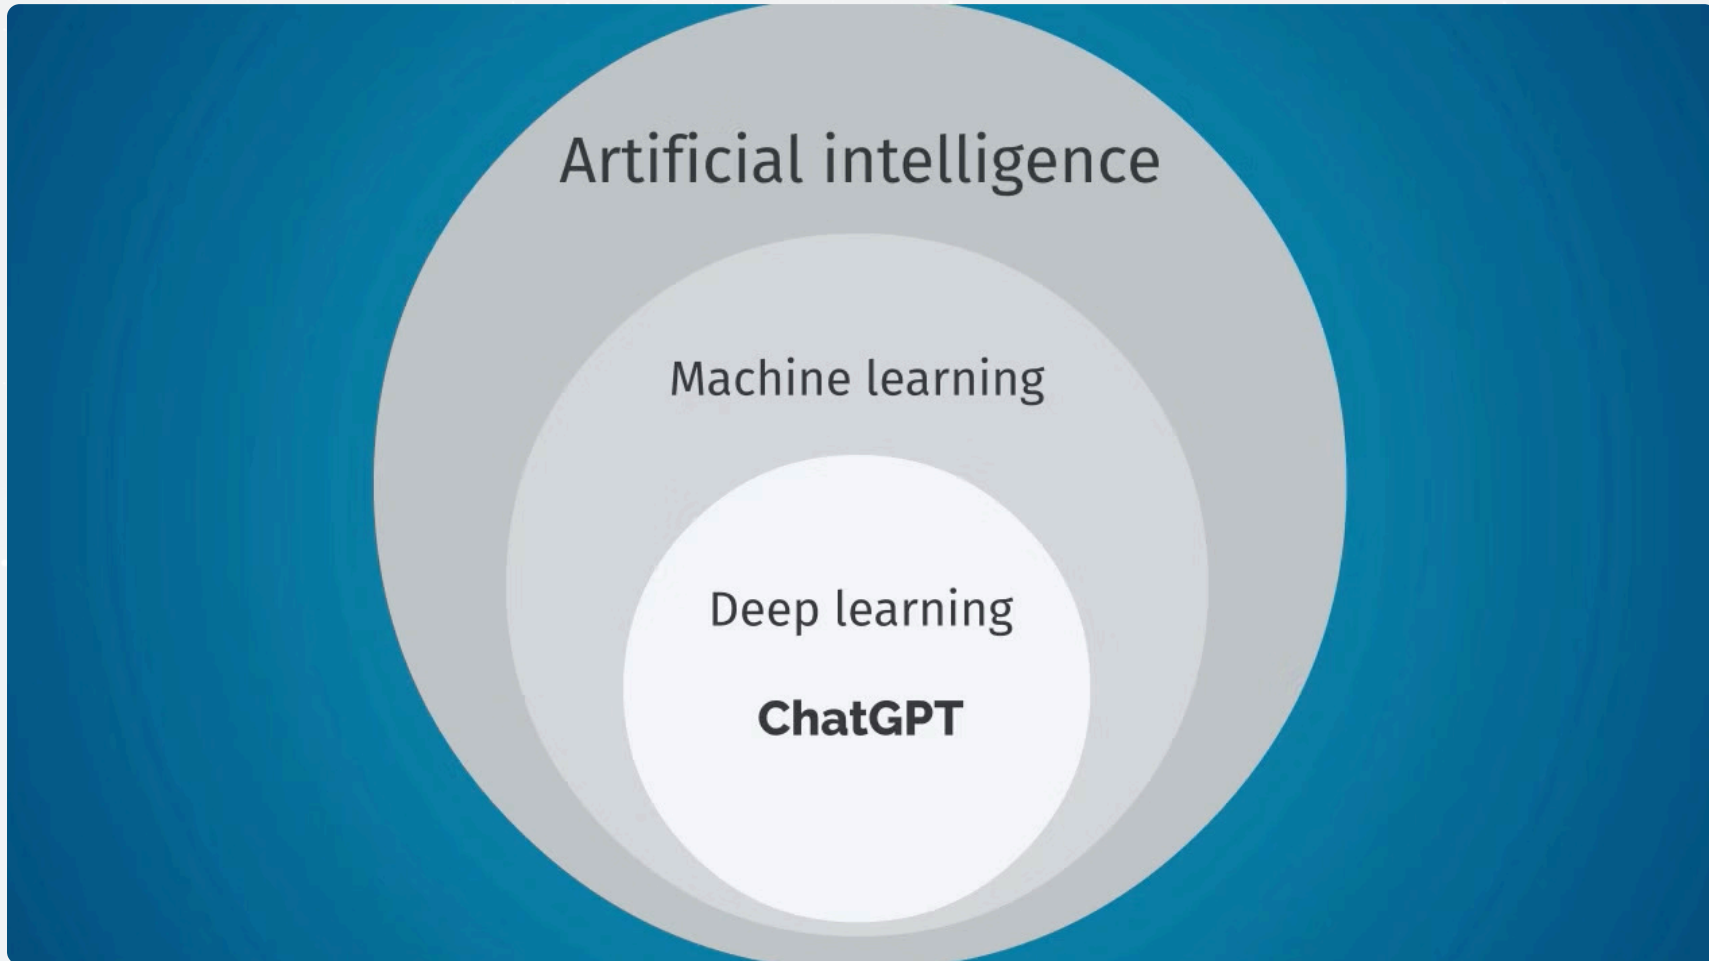
# Week 1 - Understanding the technology

**Short videos about  
the technology**





# Where does ChatGPT fit?

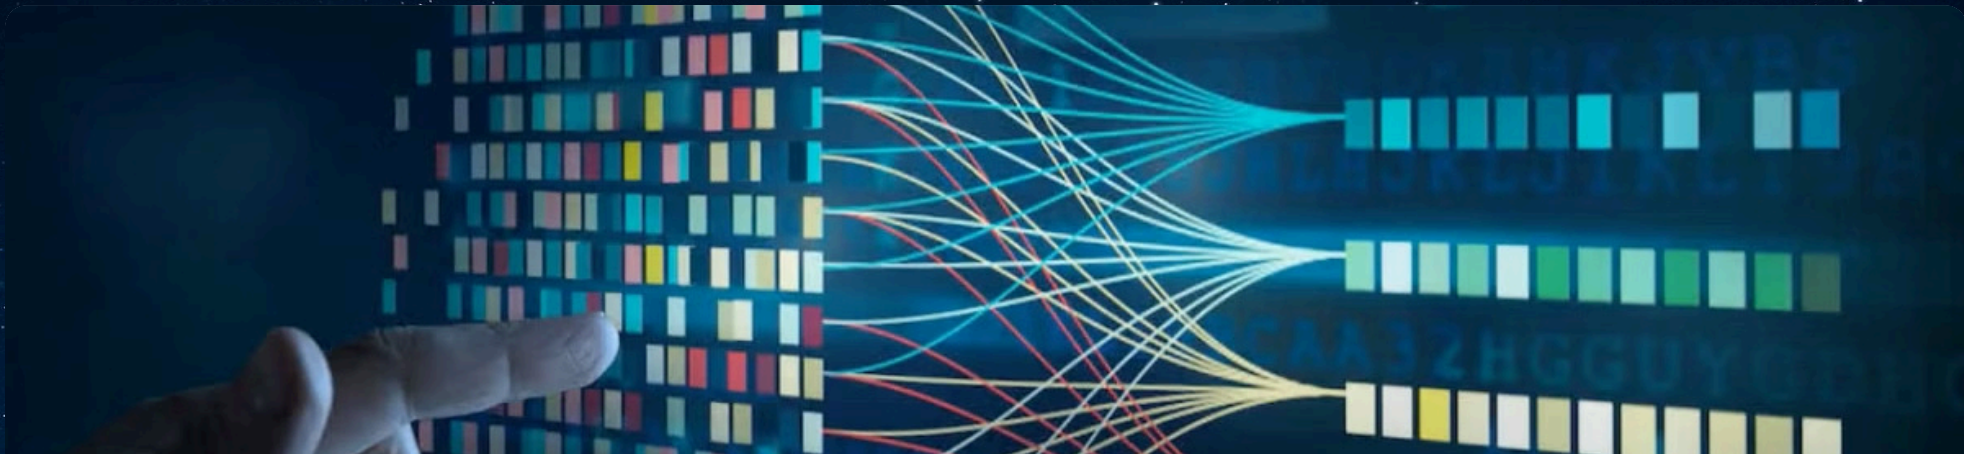




## What is machine learning?

A subfield of AI that involves developing models that can **automatically learn from and improve their performance** based on input data, allowing computers to make predictions, recognize patterns, and solve problems **without being explicitly programmed** for each task.





# Without being explicitly programmed!

They **learn patterns**, instead of following rules.

They are **probabilistic**, not deterministic.



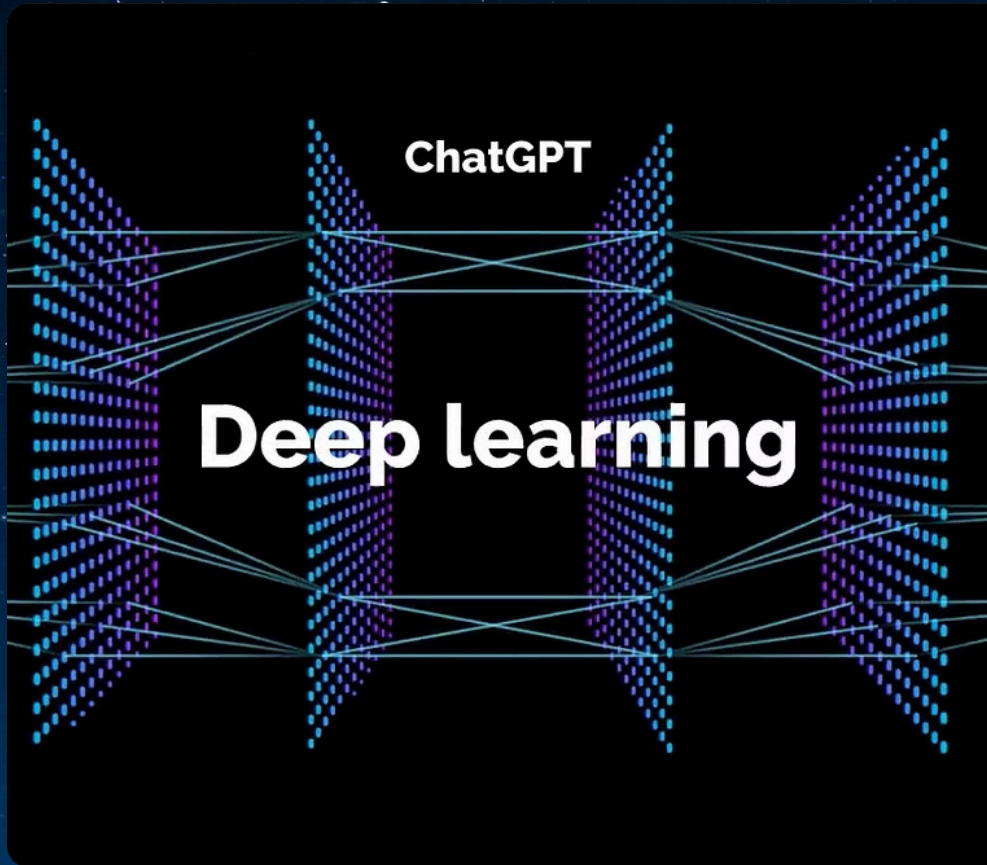


**This is very  
different from  
explicit software  
programming  
and rules.**

It's about patterns and probabilities.

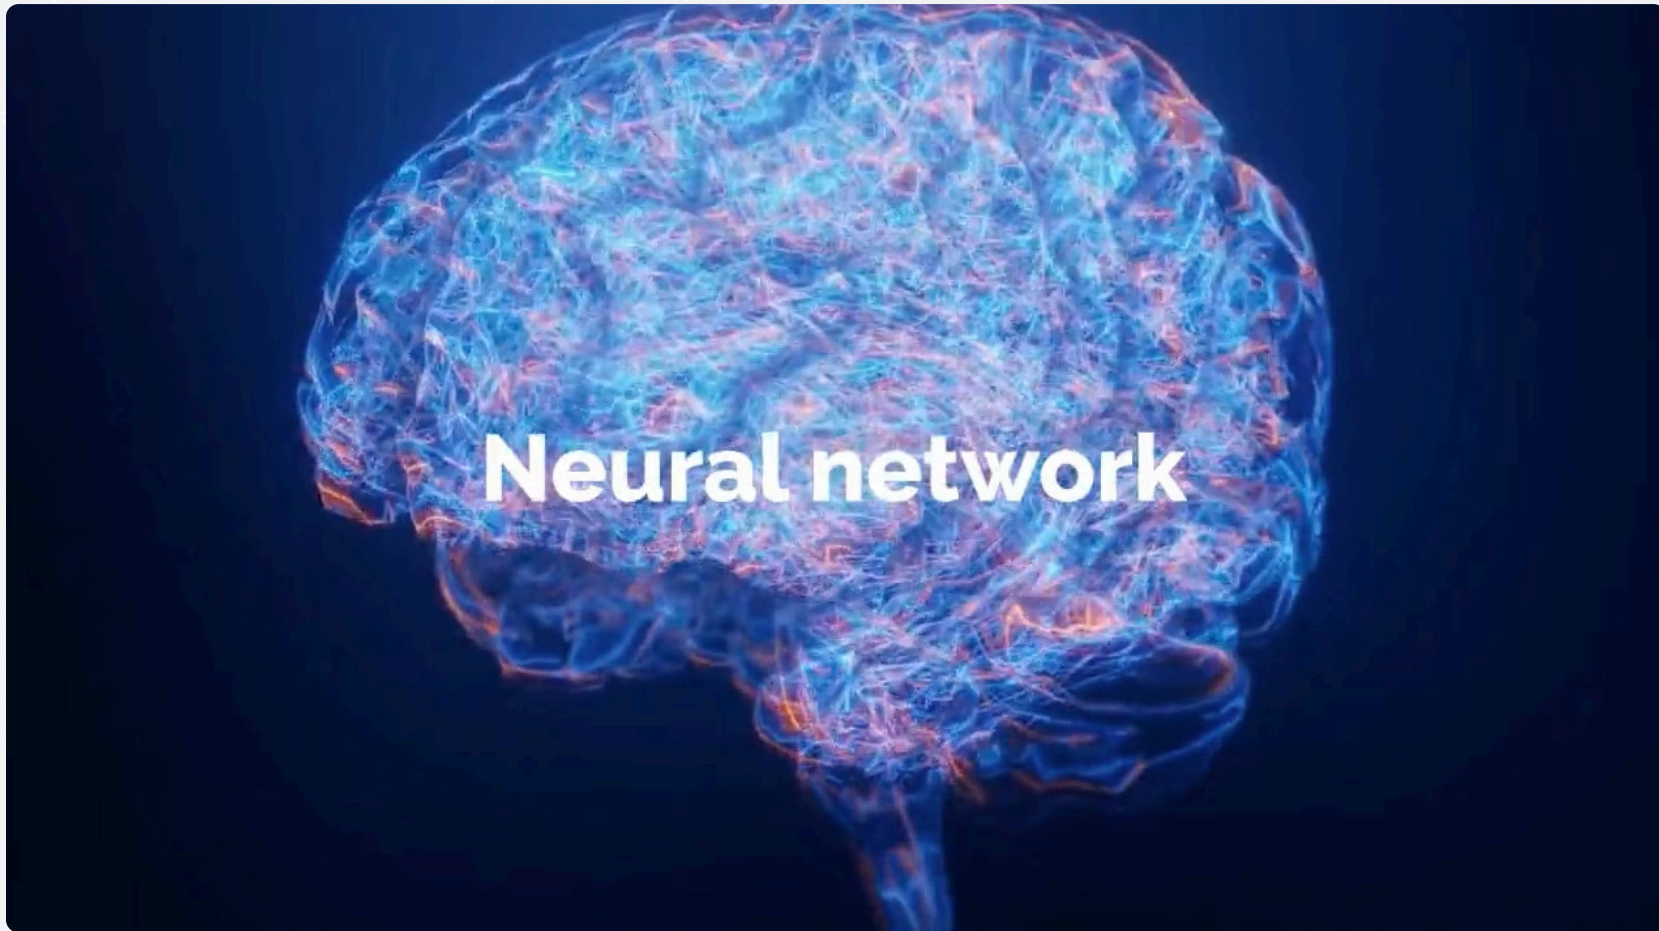
There's no "if this, then that..."





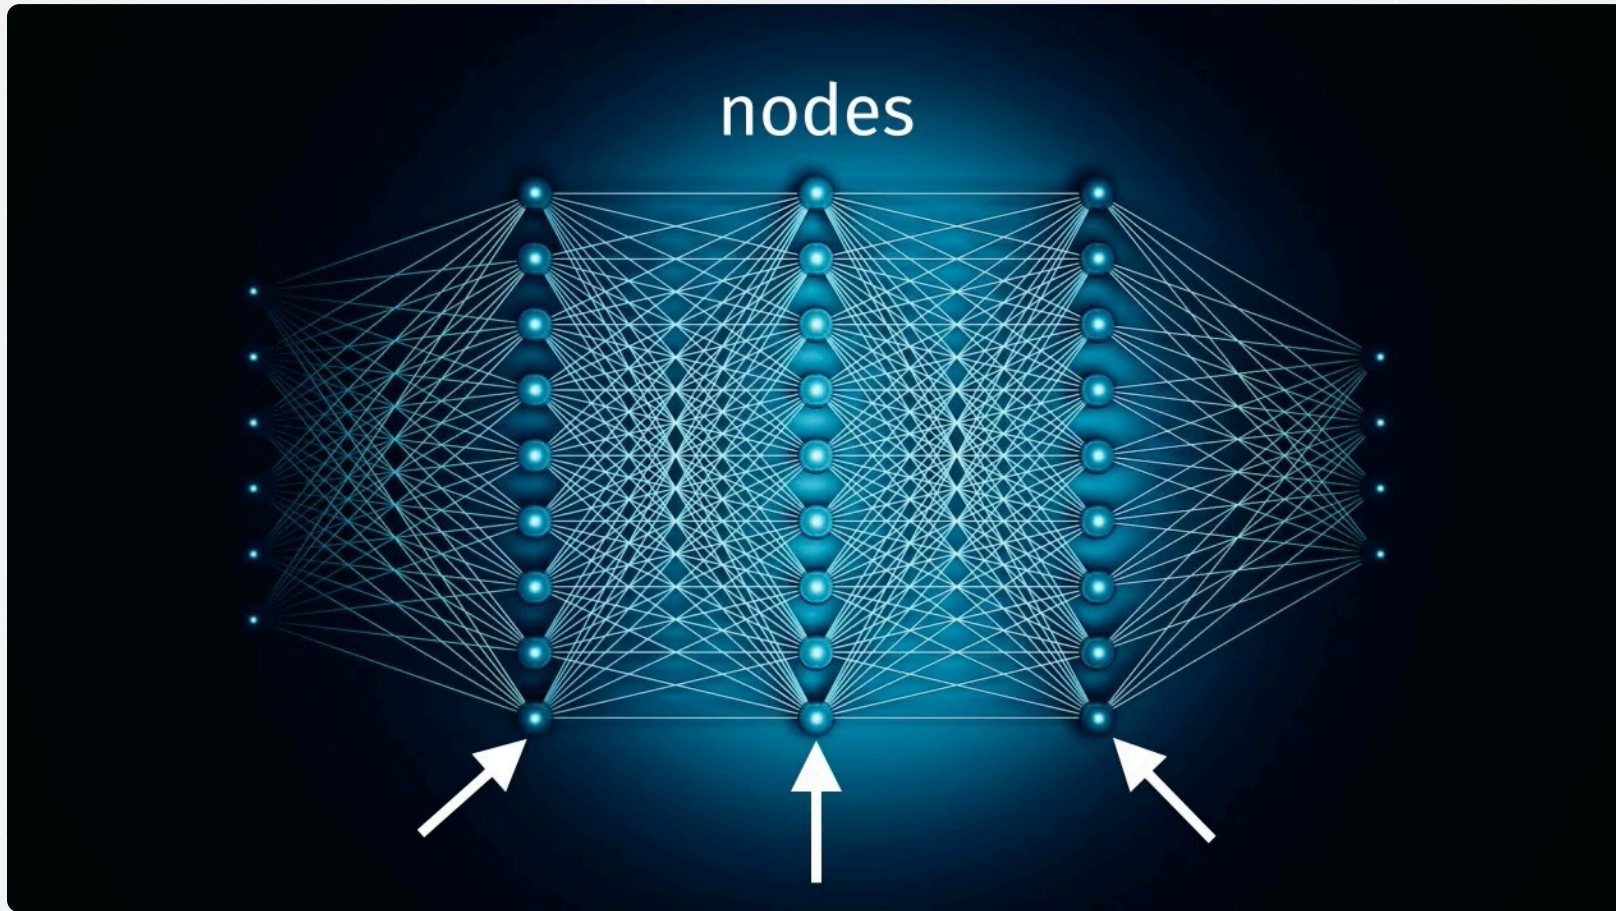
## What is deep learning?

It's a particular type of neural network known as a **deep neural network**, which consist of multiple layers of interconnected nodes that enable the learning of more complex patterns.

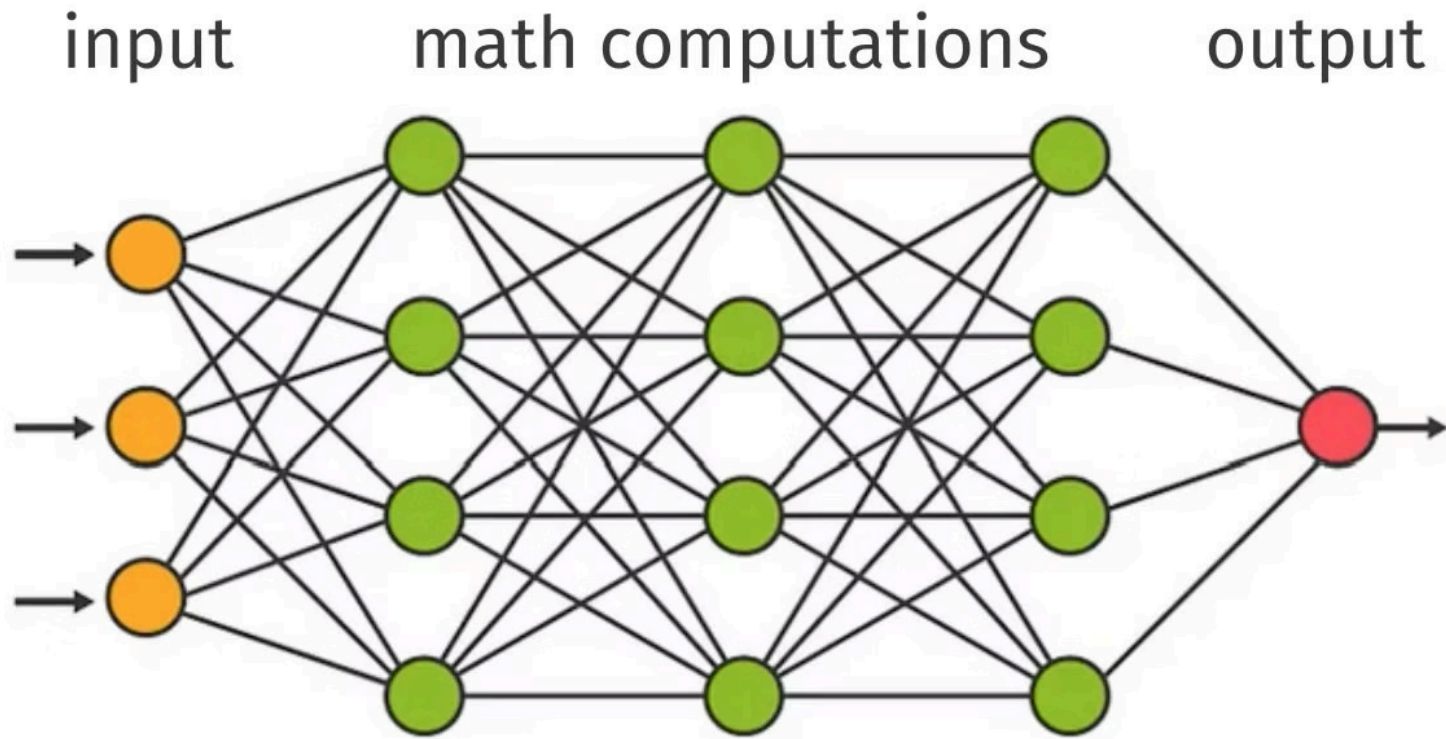


a model inspired by neural networks in the **human brain**.





It consists of processing units, known as **nodes**, organized into layers.



These nodes receive **input**, perform **math computations**, and produce **output**.





# Knowledge is not explicitly represented

With **billions of parameters** in their networks, they don't have any easily interpretable "rules" or "knowledge base" that engineers can directly inspect or modify. Their knowledge is embedded in the connections between nodes in neural networks.

**"a black box"**

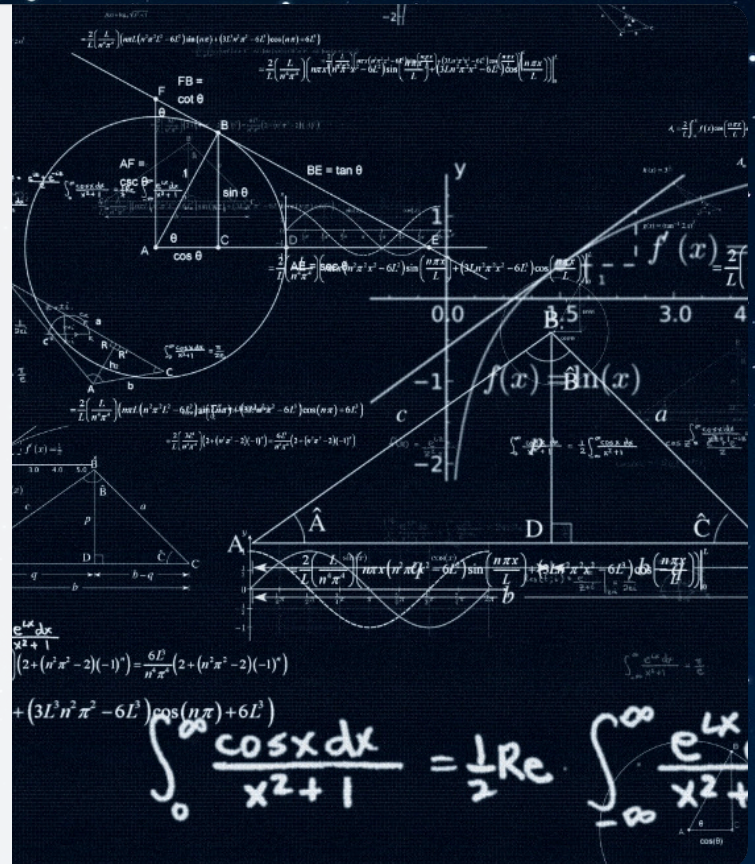
**175 billion parameters in ChatGPT**



It's not words and sentences.

It's a big network of math operations.

(don't call it a "database")





# Large language models

## LLMs

# LLMs are now called "foundation models"

They are **multimodal**. Not only about language or text.



- Image creation
- Computer vision for uploading and talking about images
- Voice assistant (talk to it and it replies verbally)





CEO: **Sam Altman**

## **Versions for researchers**

- **GPT 1 (2018)**
- **GPT 2 (2019)**
- **GPT 3 (2020)**

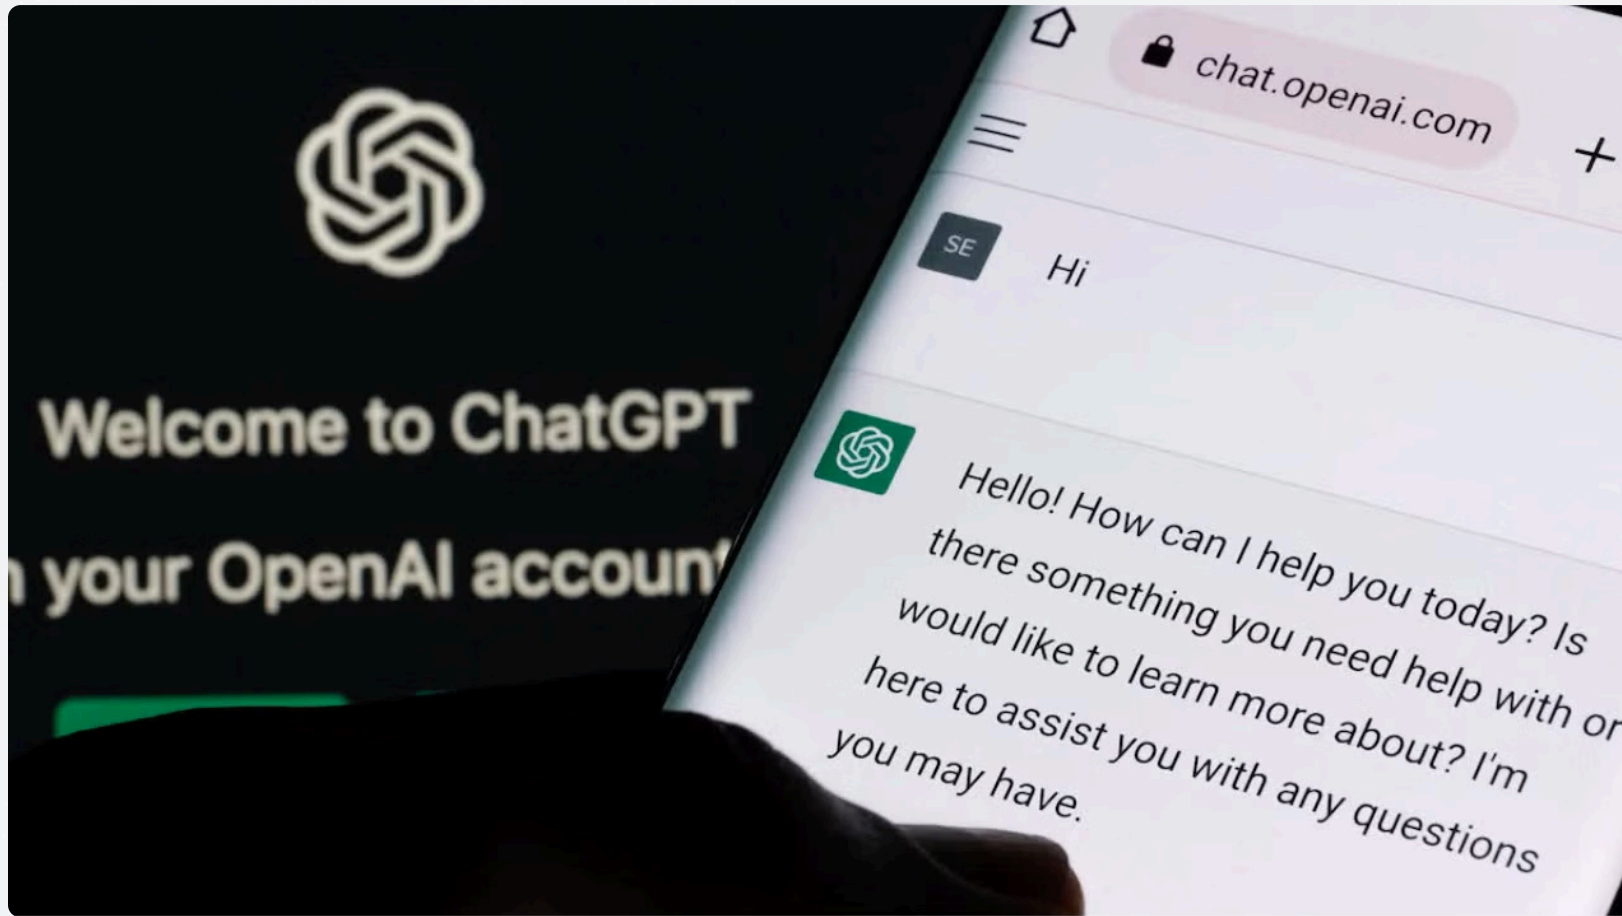
# November 30, 2022

## Introducing ChatGPT

We've trained a model called ChatGPT which interacts in a conversational way. The dialogue format makes it possible for ChatGPT to answer followup questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests.

[Try ChatGPT ↗](#)[Read about ChatGPT Plus](#)





Added an easy chat interface and made it free for the public. (GPT 3.5)

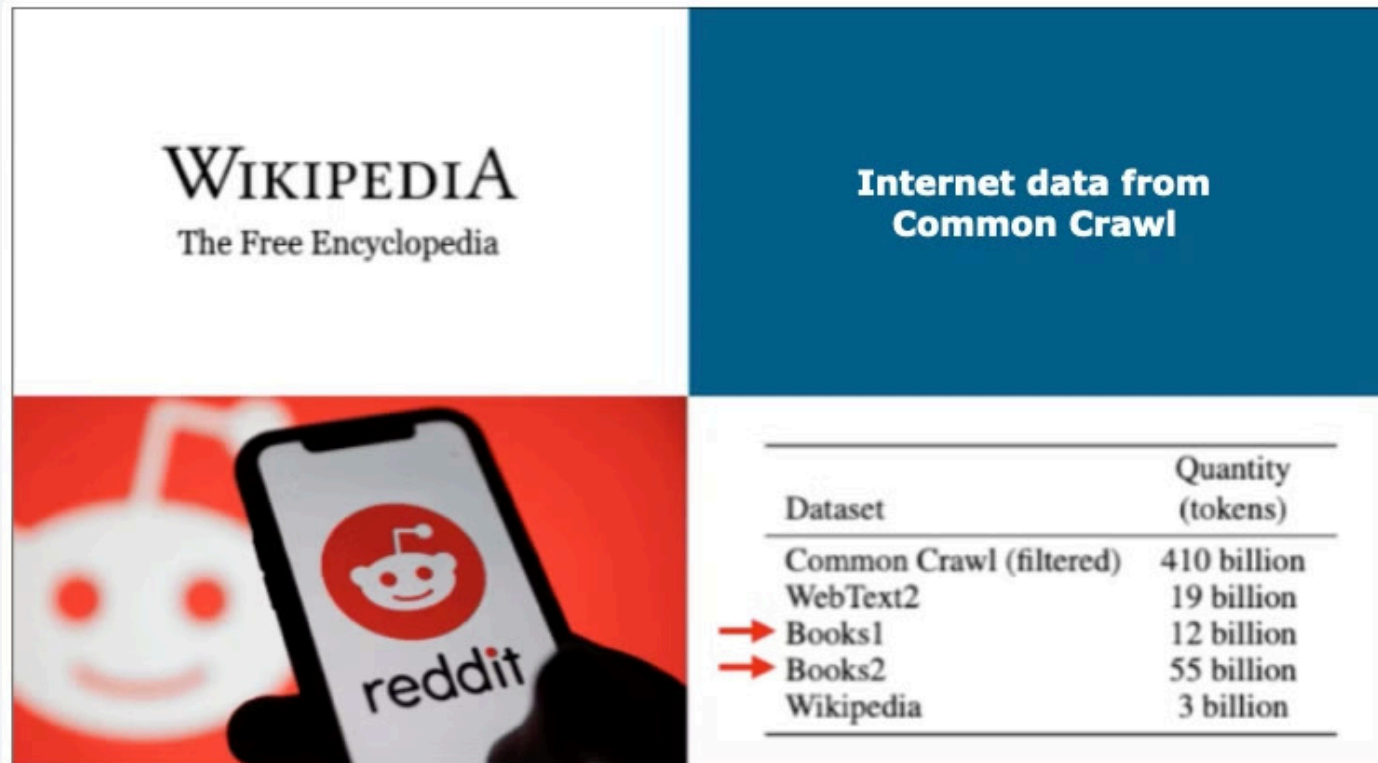
# ChatGPT

"Generative Pre-trained Transformer"



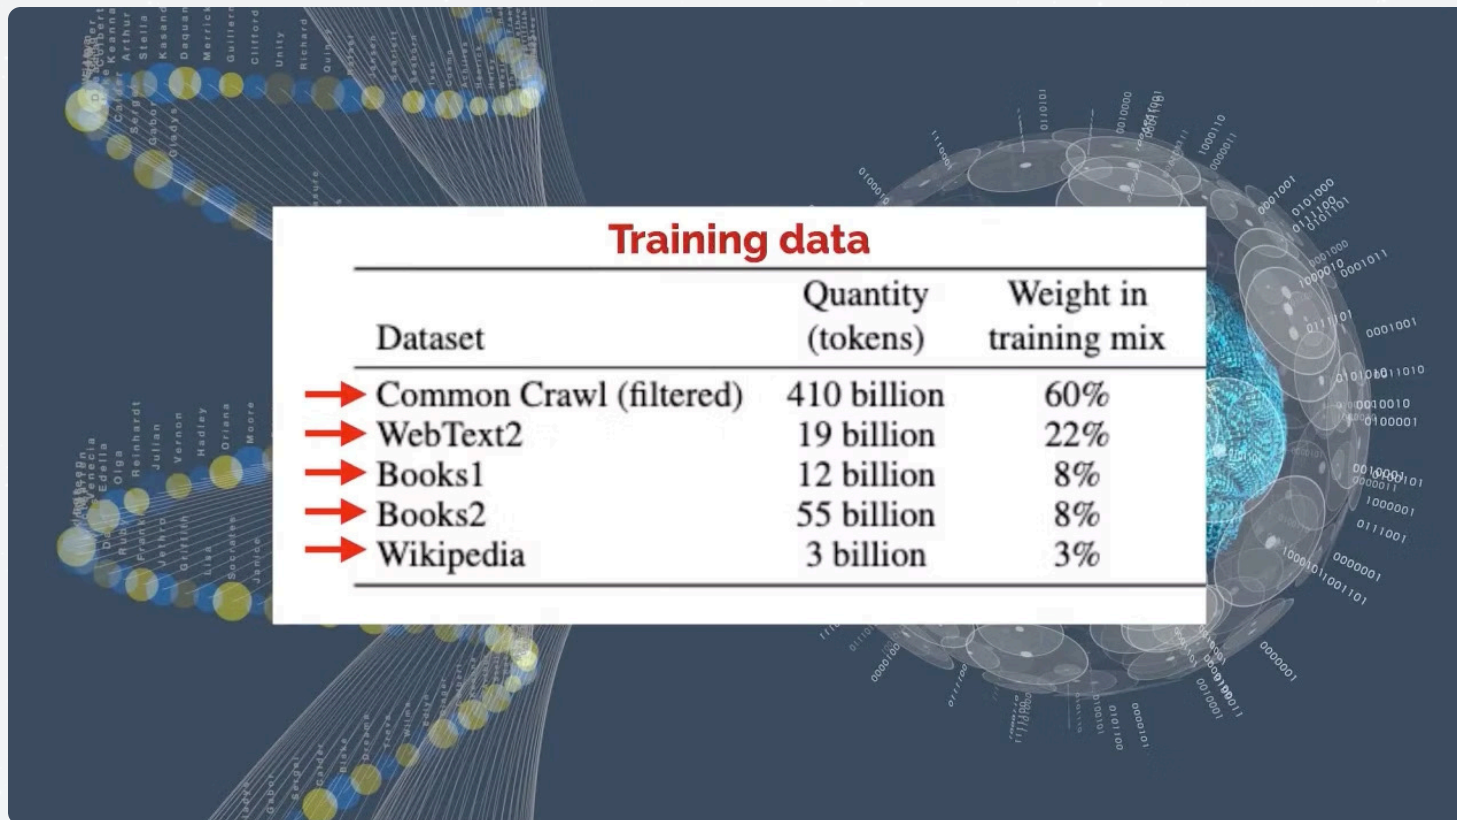


# What was ChatGPT (3.5) trained on?



[Open AI hasn't made public what GPT-4 is trained on.](#)

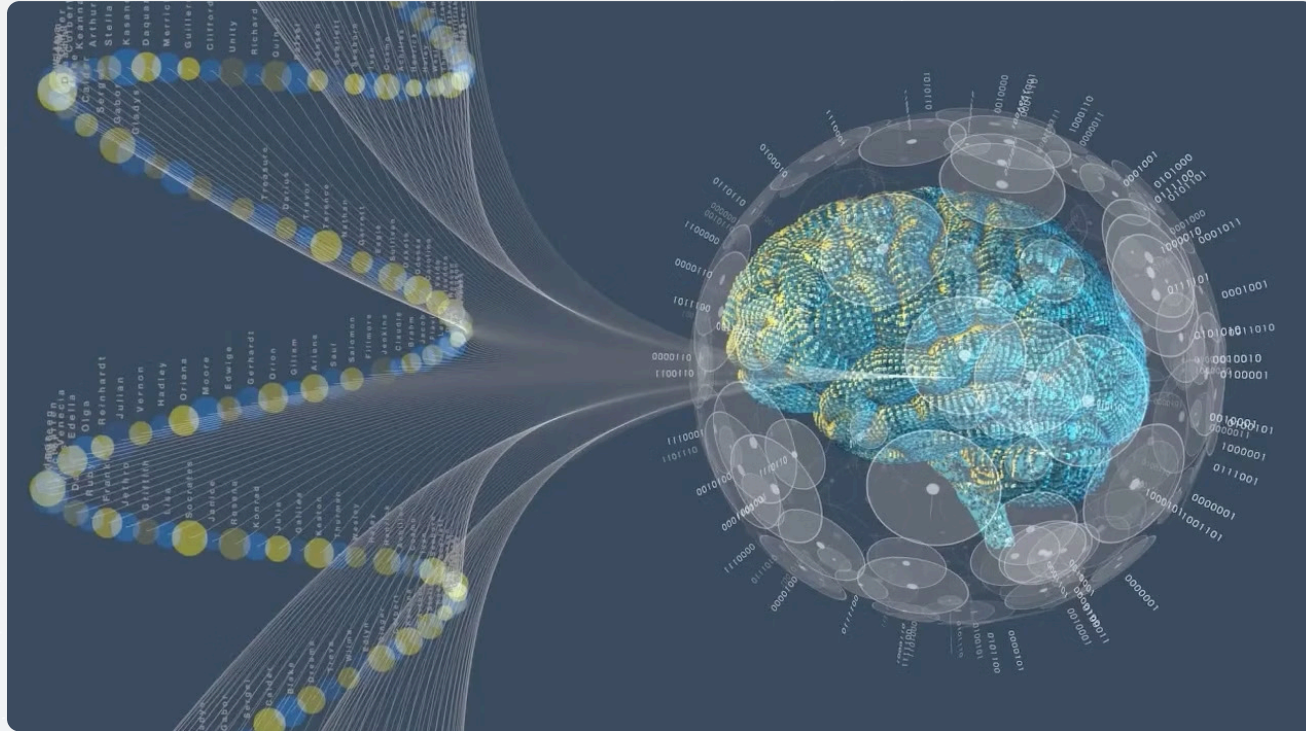
# When training is done: the data is set aside.



Training data		
Dataset	Quantity (tokens)	Weight in training mix
→ Common Crawl (filtered)	410 billion	60%
→ WebText2	19 billion	22%
→ Books1	12 billion	8%
→ Books2	55 billion	8%
→ Wikipedia	3 billion	3%

Once the training is done, it's not needed for the model to work.



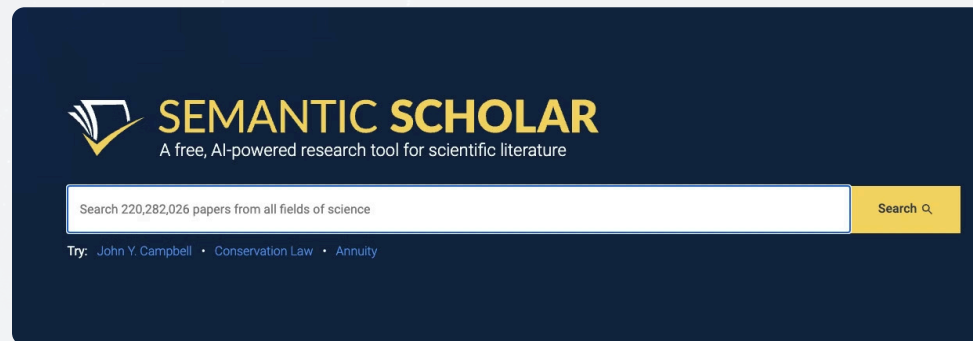


## Saved as patterns

It learned from all that data. It uses that to generate new text.

# Semantic searching

- Searching by meaning (using that math)
- Similar concepts don't need to include the same words, like in keyword searching.



[semanticscholar.org](https://www.semanticscholar.org)

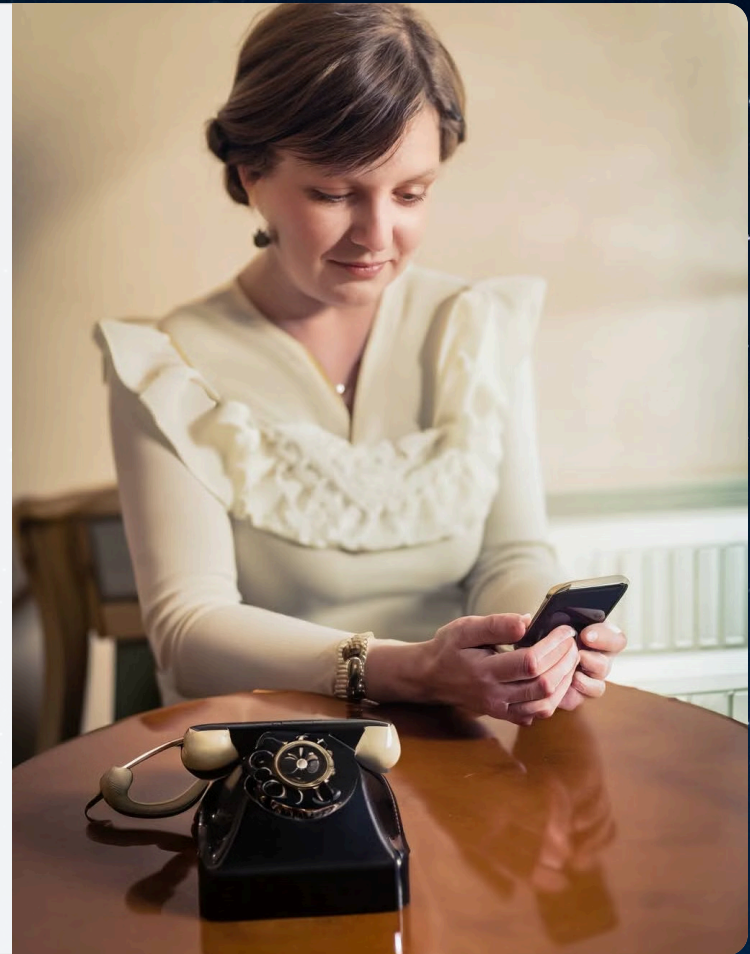


# What is "generative AI?"

AI that can generate new content:

- text
- images
- video
- music
- speech

 *This image was generated with Adobe Firefly.*





**It's important to know the difference between generative AI & other types of AI.**





# Discriminative AI

Classify or recognize patterns in existing data.

## Examples

- Spam filtering
- Netflix video recommendations
- Using AI to find medical images that might indicate health problems

# Why know the difference?

- These types are often lumped together in news stories, but they are very different.
- Their strengths and weaknesses are very different.



Facial recognition is **discriminative AI**.

[Time to Ban Facial Recognition from Public Spaces and Borders.](#)



# ChatGPT is not the only foundation model

All have both free and paid versions.



[ChatGPT](#)



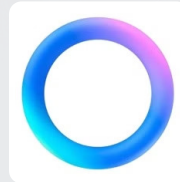
[Claude](#)



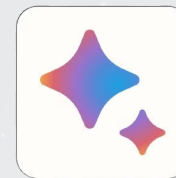
[ChatGPT Plus](#)



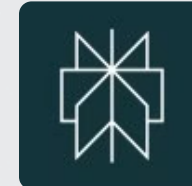
[Microsoft  
Copilot](#)



[Meta AI](#)



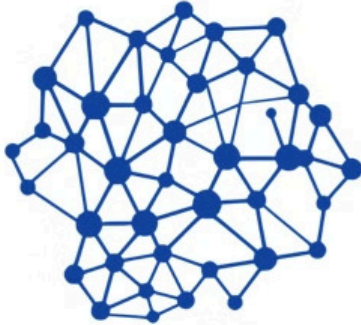
[Gemini](#)



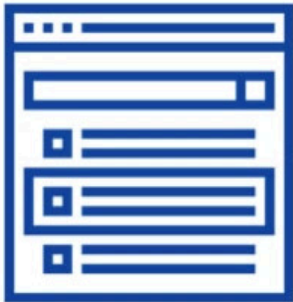
[Perplexity AI](#)  
(based on  
foundation models)

<https://libguides.library.arizona.edu/ai-literacy-instructors/beyond>

LLM



+ facts



## What is grounding?

**Grounding** is an important concept in AI research. It is about connecting the model to **external sources of knowledge** - such as:

- web search results  
(like Microsoft Copilot)



# Which models are grounded with search results?

## Not grounded



[ChatGPT](#)  
(exception:  
limited searching)

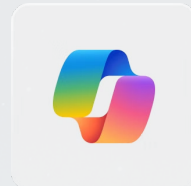


[Claude](#)

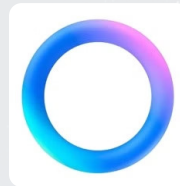
## Grounded



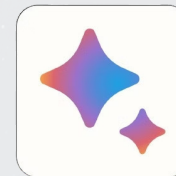
[ChatGPT Plus](#)



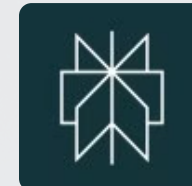
[Microsoft  
Copilot](#)



[Meta AI](#)



[Gemini](#)



[Perplexity AI](#)

# Un-grounded models



# Un-grounded models rely only their training data

- [ChatGPT](#) (free version)
- [Claude](#)

**These models will never be useful for very current information.**

- ChatGPT: up to January 2022
- Claude: up to [April 2024](#).



## Tasks for un-grounded models

- Generate ideas (for anything)
- Come up with examples
- Generate keywords for searching in research databases
- Summarize long documents
- Revise your writing in different styles
- Revise your writing for different audiences or levels
- Copy or upload parts of a research paper you don't understand and ask it to explain it in simpler terms



Think of these as **wordsmiths** and idea generators, not search engines.



# Grounded models

They have a search engine as a tool to work with.



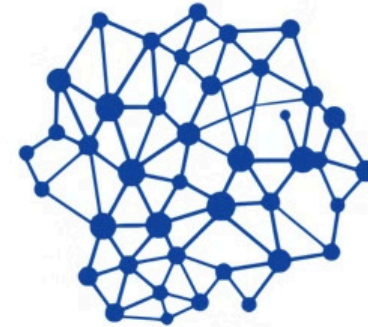
# Models *with* grounding

Models that answer based on their training PLUS .... a source of facts

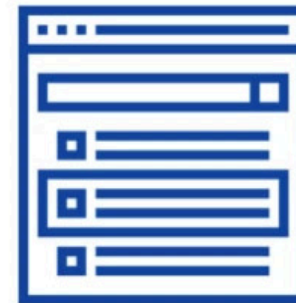
Examples: web search results

- Microsoft Copilot
- Perplexity AI
- Gemini

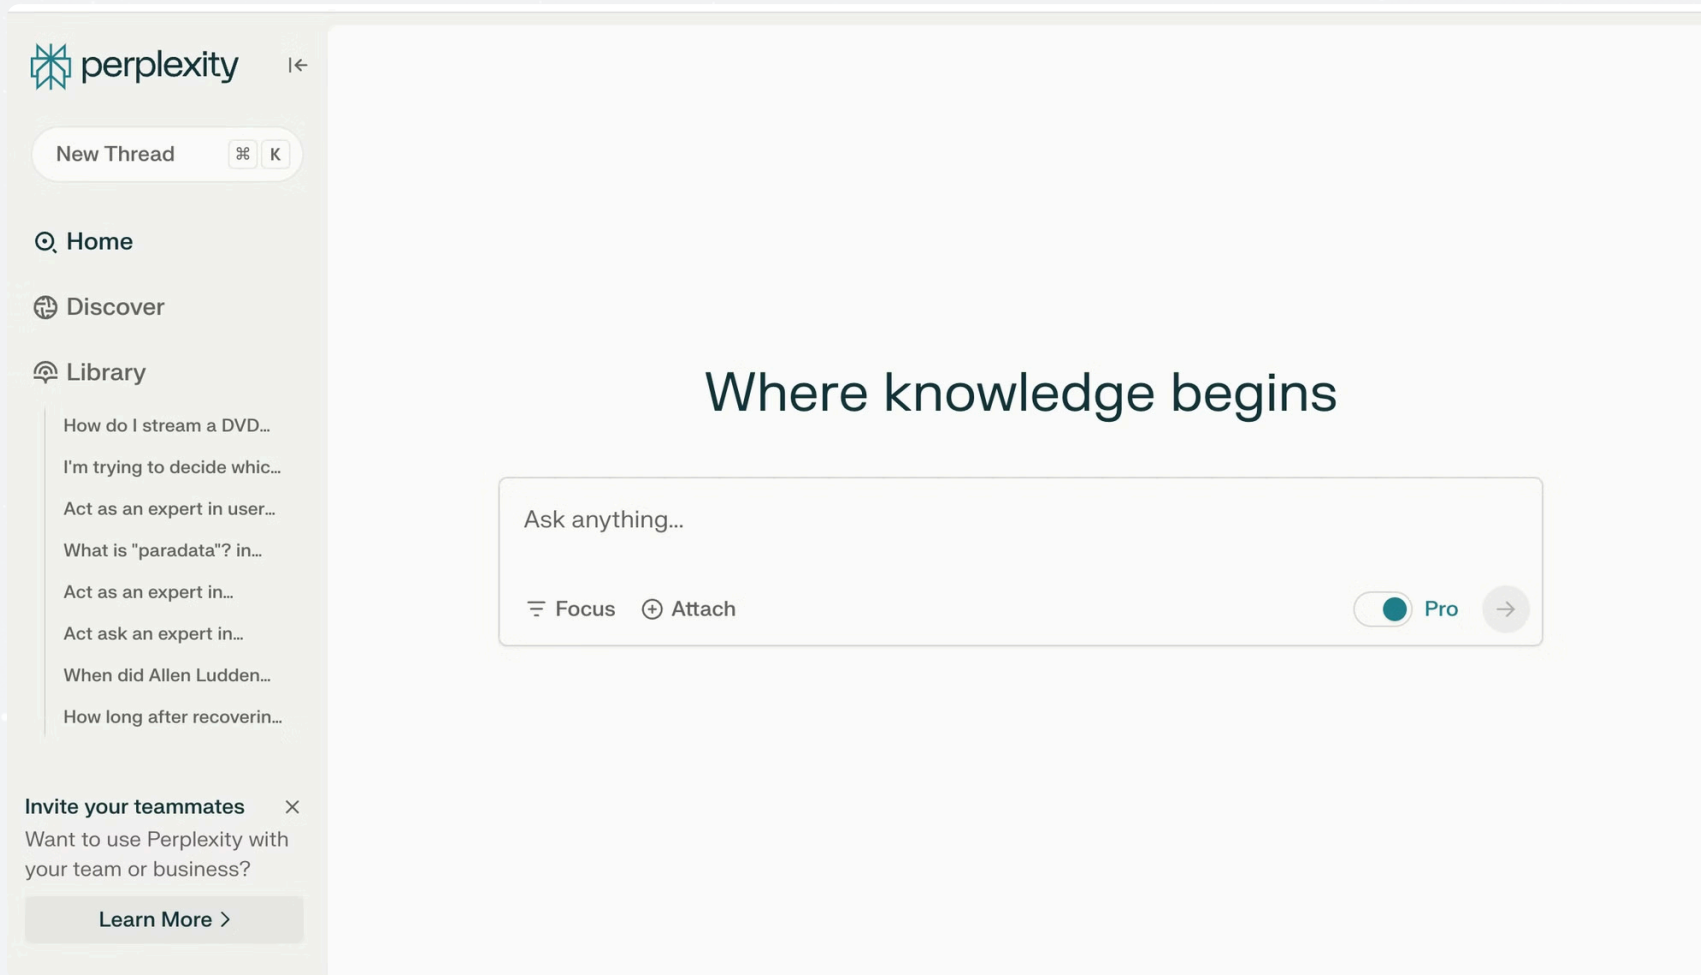
LLM



+ facts



# Grounded with web search

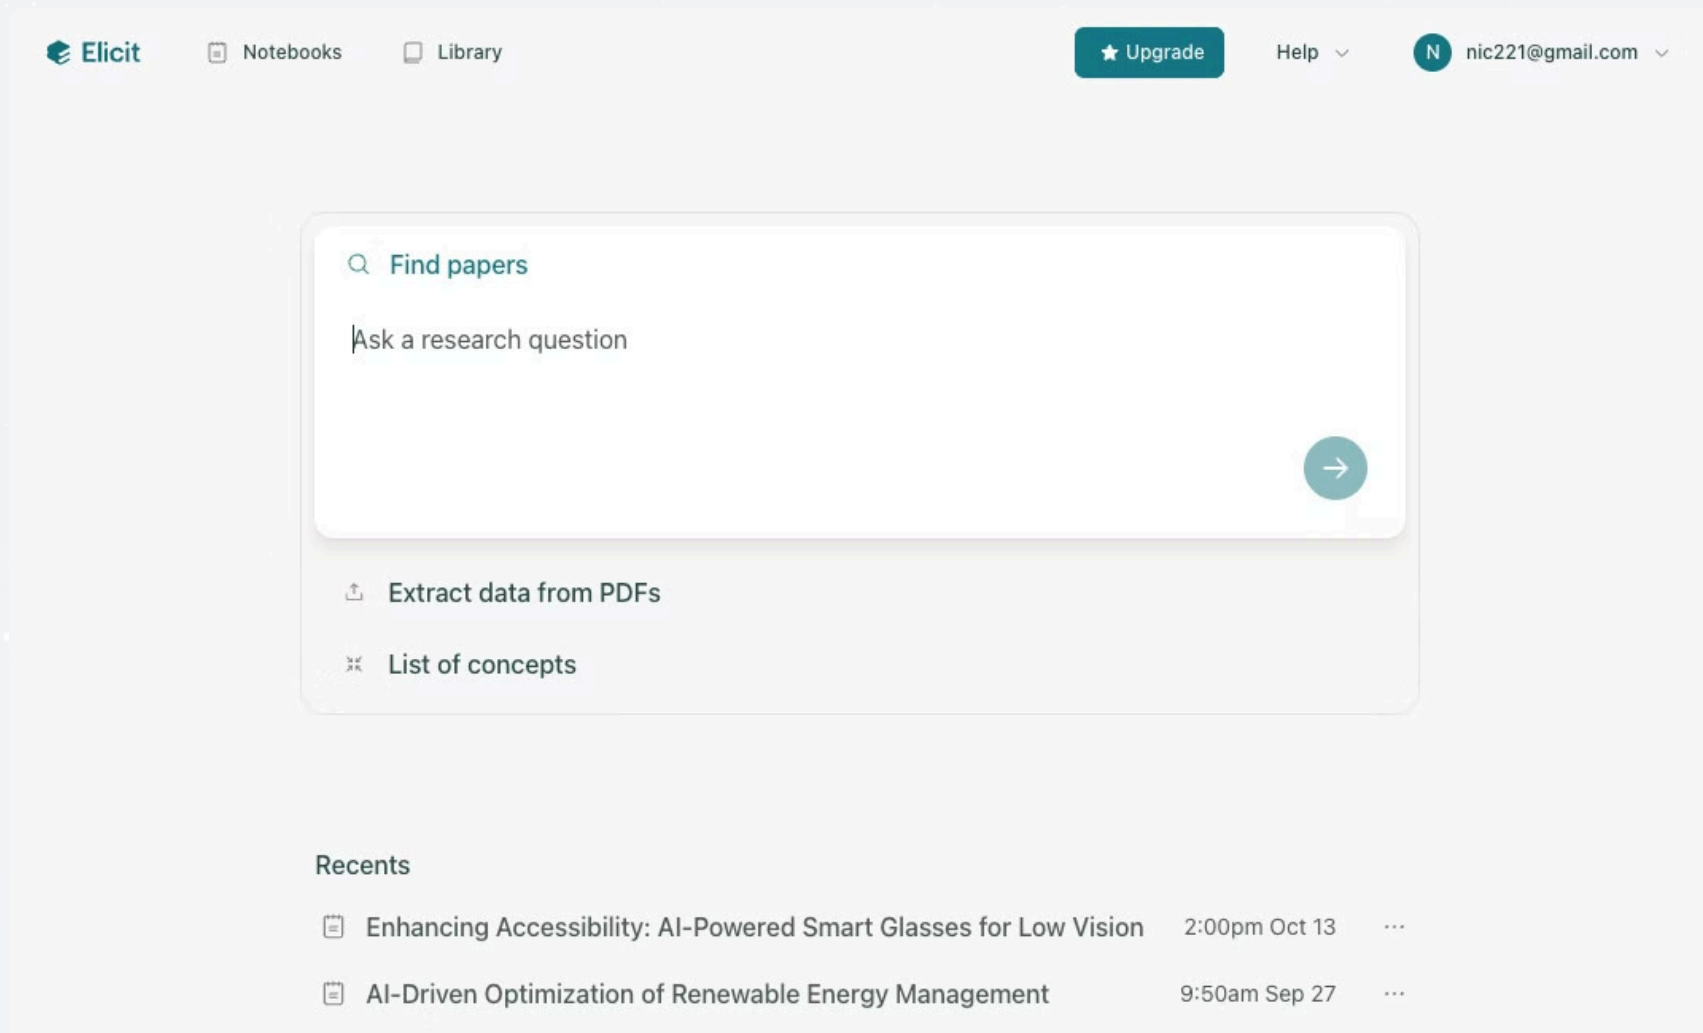


perplexity.ai

Keep track of these models: <https://libguides.library.arizona.edu/ai-literacy-instructors/beyond>



# You can also ground a model with scholarly papers



The screenshot displays the Elicit website interface. At the top, there is a navigation bar with the Elicit logo, links for 'Notebooks' and 'Library', an 'Upgrade' button, a 'Help' dropdown, and a user profile for 'nic221@gmail.com'. Below the navigation bar is a large search input field with the placeholder text 'Ask a research question' and a search icon. Below the search field are two options: 'Extract data from PDFs' and 'List of concepts'. At the bottom of the interface, there is a 'Recents' section with two entries: 'Enhancing Accessibility: AI-Powered Smart Glasses for Low Vision' (dated 2:00pm Oct 13) and 'AI-Driven Optimization of Renewable Energy Management' (dated 9:50am Sep 27).

Elicit

Notebooks Library

Upgrade Help

nic221@gmail.com

Find papers

Ask a research question

Extract data from PDFs

List of concepts

Recents

Enhancing Accessibility: AI-Powered Smart Glasses for Low Vision 2:00pm Oct 13

AI-Driven Optimization of Renewable Energy Management 9:50am Sep 27

<https://elicit.com/>

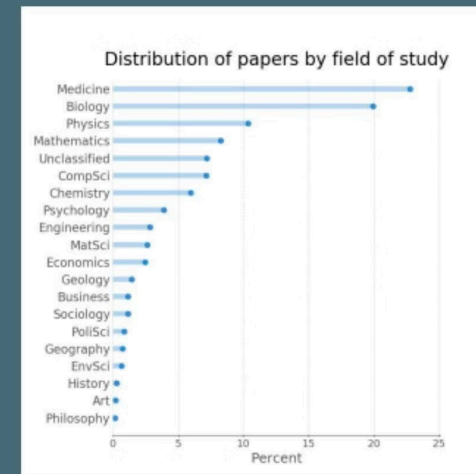
# Elicit is grounded with Semantic Scholar data

Academic discipline	All papers	Full text available
Medicine	12.8M	1.8M
Biology	9.6M	1.6M
Chemistry	8.7M	484k
n/a	7.7M	583k
Engineering	6.3M	228k
Comp Sci	6.0M	580k
Physics	4.9M	838k
Mat Sci	4.6M	213k
Math	3.9M	669k
Psychology	3.4M	316k
Economics	2.3M	198k
Poli Sci	1.8M	69k
Business	1.8M	94k
Geology	1.8M	115k
Sociology	1.6M	93k
Geography	1.4M	58k
Env Sci	766k	52k
Art	700k	16k
History	690k	22k
Philosophy	384k	15k

## Semantic Scholar's Corpus

Used by:  
  
Elicit  
Consensus  
LitMaps

Source:  
<https://arxiv.org/pdf/1911.02782.pdf>



But still not nearly as big as the coverage of Google Scholar or library databases.



# Hallucination can still happen, even with grounded models.

But it's easy to recognize when it links to a source that's not relevant.

And of course, web search results can contain misinformation.

Keep a "[human in the loop](#)"


# What else we'll cover in week 1

- Choosing an AI tool for your task
- Understanding tokens and context windows, and how words get turned into math
- Openly licensed vs proprietary models

## Which generative AI tool for your task?

People often ask, which AI tool can I use to do X? This aims to help.

Version 1, November 2024  
[nicolehennig.com](https://nicolehennig.com)

 by Nicole Hennig  
Last edited about 2 hours ago

---

*Which generative AI tool for your task?* © 2024 by [Nicole Hennig](https://nicolehennig.com) is licensed under [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/)



# Hands-on activities, week 1

- Configuring your privacy settings
- Writing an explanation of the technology for a 13 year old
- Sorting virtual flash cards representing various types of AI into generative vs discriminative AI
- Using Perplexity effectively to search the web
- Using Elicit effectively to search scholarly articles

**Discriminative AI or generative AI?**  
There are 15 cards in the stack. Drag each one to the category it belongs in.

**Remember**  
Discriminative AI **classifies** things while generative AI **creates content**.

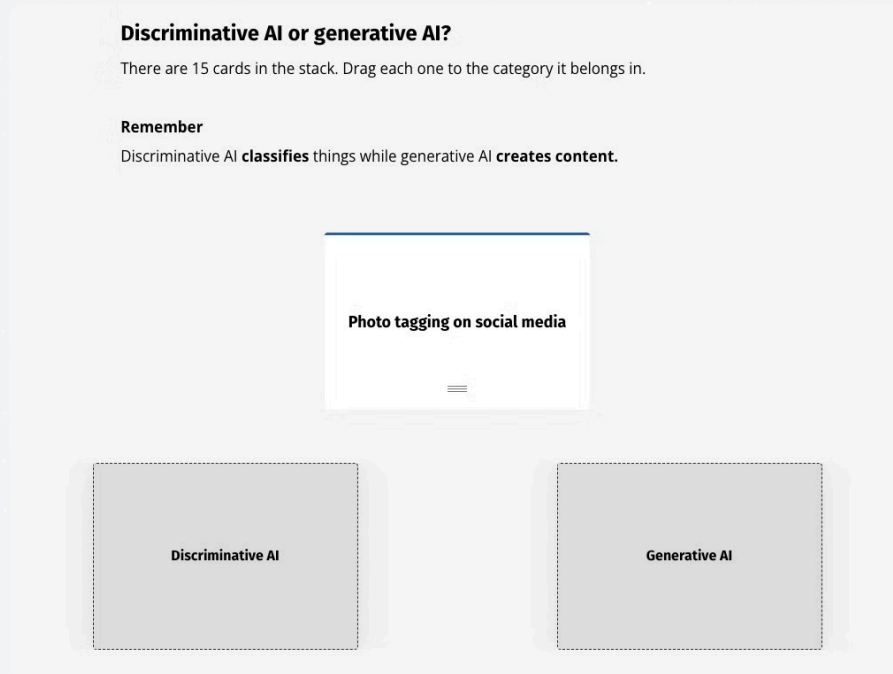


Photo tagging on social media

Discriminative AI

Generative AI

# Readings, week 1

- Privacy (especially related to HIPPA and other regulations)
- Transparency and citing
- Openly licensed (open source) models

---

## Discussion forum

AI Policy & Governance, CDT AI Governance Lab

### **CDT Joins Mozilla, Civil Society Orgs, and Leading Academics in Urging U.S. Secretary of Commerce to Protect AI Openness**

March 25, 2024 / [Kevin Bankston](#)

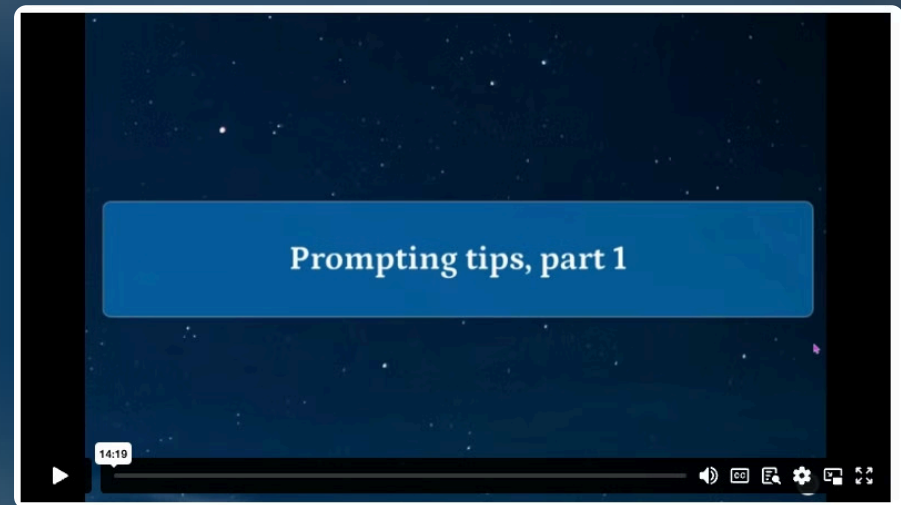
The Center for Democracy & Technology (CDT) joined Mozilla and nearly fifty civil society organizations and scholars in a letter to U.S. Secretary of Commerce Gina Raimondo, urging her to protect openness and transparency in AI. The joint letter is in response to a [public consultation process](#) run by the National Telecommunications and Information Administration (NTIA) at the Department of Commerce, to examine the risks, benefits, and potential policy approaches related to open models for AI.

The letter highlights the many clear benefits of an open ecosystem of AI models,

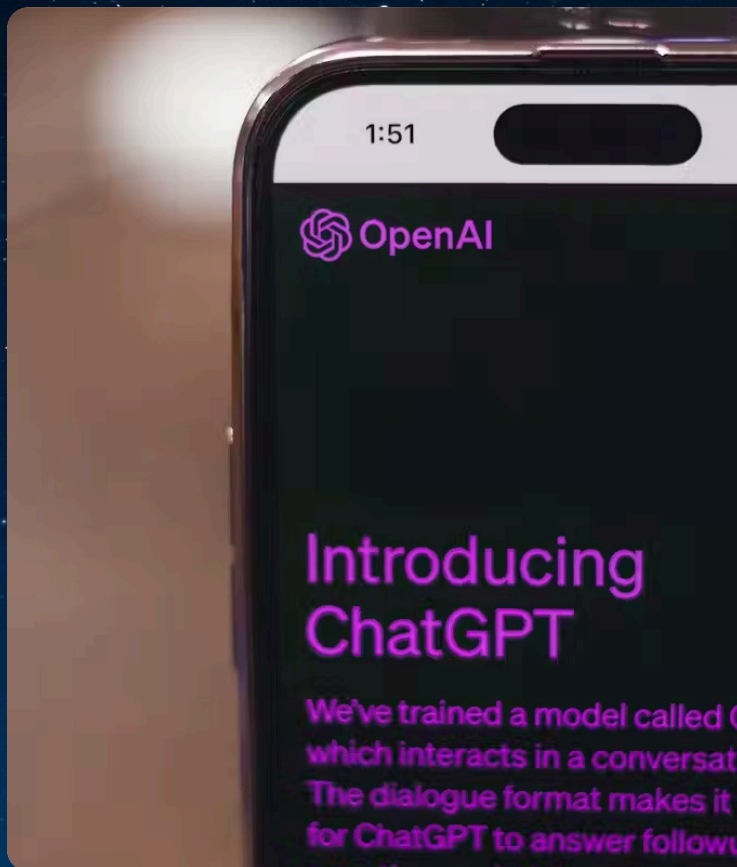


## Week 2 - Prompting for language models

# Short videos about prompting







## Prompting: how you talk with it

It makes a big difference in how it responds.

## **Give it context.**

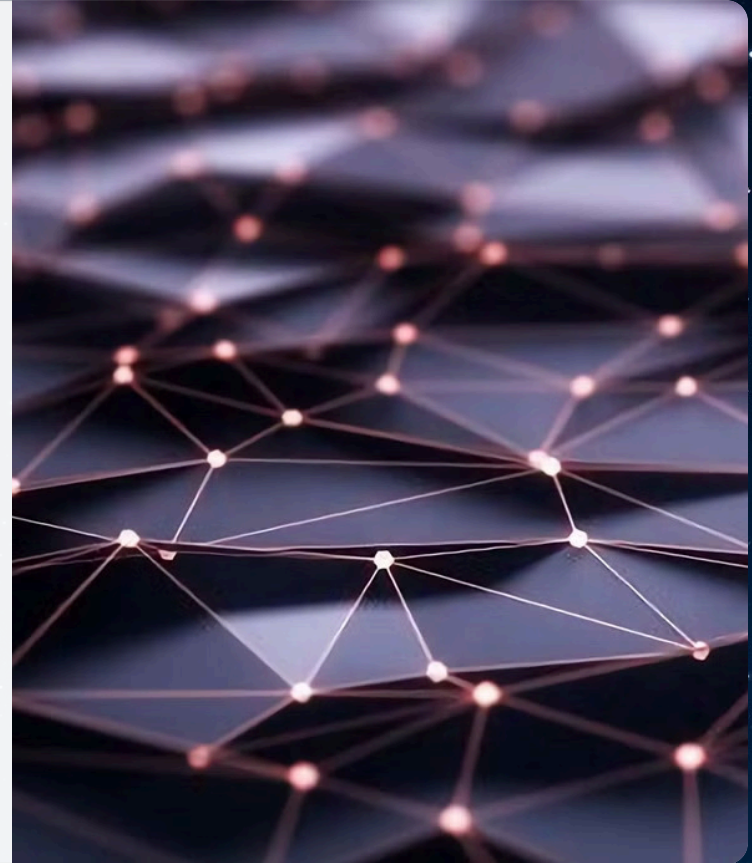
### **Role**

"You are an expert in child development."

"Act as an expert in public health."

### **Context**

"Using the theory of [insert a theory], create a ..."





## **Give it clear instructions, including format.**

### **Example:**

"Use Mike Caulfield's ideas for how to be an effective fact-checker. Create a lesson plan for college freshmen to help them learn to fact-check information they find on the Internet. The outcomes should be that students become familiar with Caulfield's "Four Moves" for effective fact-checking, and that they gain practical skills to fact-check information they find on the Internet. Format the output in the way a typical lesson plan is formatted."

# Formats

- tables
- emojis
- ascii art
- code
- quiz questions
- and more...



You are an expert in population studies. Please create a table that lists the top 10 most populous cities in the world. First column will contain the city name, second column the population.



Sure! Here is a table that lists the top 10 most populous cities in the world as of 2021:



City	Population (2021)
Tokyo, Japan	37,393,000
Delhi, India	30,291,000
Shanghai, China	27,058,000
São Paulo, Brazil	22,043,000
Mumbai, India	21,042,000
Beijing, China	20,384,000
Karachi, Pakistan	20,026,000
Istanbul, Turkey	15,190,000
Dhaka, Bangladesh	14,543,000
Moscow, Russia	13,197,000



## **Don't expect the perfect answer right away.**

Ask for changes, clarifications, or improvements. Tell it exactly what you want and what format you'd like it in.

**Keep conversing — "ask and adjust."**

# **Use the results as a starting point.**

Use your own expertise to modify these. It's great for getting ideas. But you are the expert.



## "Prompt engineering?" or "Prompt crafting?"

- It's not like coding, you get slightly different responses each time.
- It's more like talking to a human assistant and giving them detailed instructions.

## What is "hallucination?"

The official term from the field of machine learning for outputting inaccurate but plausible information.

You could say, "making things up."

Always fact check.





# There is ongoing research on ways to mitigate hallucinaion

nature

Explore content ▾ About the journal ▾ Publish with us ▾

nature > articles > article

Article | [Open access](#) | Published: 19 June 2024

## Detecting hallucinations in large language models using semantic entropy

[Sebastian Farquhar](#) , [Jannik Kossen](#), [Lorenz Kuhn](#) & [Yarin Gal](#)

*Nature* **630**, 625–630 (2024) | [Cite this article](#)

111k Accesses | 2 Citations | 1529 Altmetric | [Metrics](#)

## Improving Retrieval Augmented Language Model with Self-Reasoning

Yuan Xia, Jingbo Zhou, Zhenhui Shi, Jun Chen, Haifeng Huang  
Baidu Inc., China.

{xiayuan,zhoujingbo,shizhenhui,chenjun22,huanghaifeng}@baidu.com

### Abstract

The Retrieval-Augmented Language Model (RALM) has shown remarkable performance on knowledge-intensive tasks by incorporating external knowledge during inference, which mitigates the factual hallucinations inherited in large language models (LLMs). Despite these advancements, challenges persist in the implementation of RALMs, particularly concerning their reliability and traceability. To be specific, the irrelevant document retrieval may result in unhelpful response generation or even deteriorate the performance of LLMs, while the lack of proper citations in generated outputs complicates efforts to verify the trustworthiness of the models. To this end, we propose a novel self-reasoning framework aimed at improving the reliability and traceability of RALMs, whose core idea is to leverage reasoning trajectories generated by the LLM itself. The framework involves constructing self-reason trajectories



13v2 [cs.CL] 2 Aug 2024

# Prompting to avoid hallucinations



# Ask multiple times

One feature of hallucinations is that they tend to be different across different outputs.

So one thing you can do is ask multiple times and see if the outputs are consistent.

---

Learn more in this Coursera course: [8 Most Controversial Terms in AI, Explained](#)

(You can audit it for free).

# I asked ChatGPT 4o three times: same answer, worded differently – likely correct

ChatGPT 4o ▾

what is the heaviest hippo ever recorded?



The heaviest hippo ever recorded was a male weighing approximately **9,900 pounds (4,500 kilograms)**. This massive weight is significantly above the average for male hippos, which typically range from 3,300 to 4,000 pounds (1,500 to 1,800 kilograms).

what is the heaviest hippo ever recorded?



The heaviest hippo ever recorded weighed approximately **9,900 pounds (4,500 kilograms)**. This exceptionally large male hippo was significantly heavier than the average hippo, which usually weighs between 3,300 to 4,000 pounds (1,500 to 1,800 kilograms) for males.

what is the heaviest hippo ever recorded?



The heaviest hippopotamus ever recorded weighed around **9,900 pounds (4,500 kilograms)**. This weight is exceptionally heavy for a hippo, as most adult male hippos typically weigh between 3,300 to 4,000 pounds (1,500 to 1,800 kilograms).



I verified it here. [The Largest Hippopotamus on Record](#), Aug. 19, 2024



# Asked Claude 3 times: different answers – likely hallucinating

NH what is the heaviest hippo ever recorded?

The heaviest hippopotamus ever recorded was a male named Obaysch who lived at the London Zoo in the 19th century. He weighed approximately 4,500 pounds (2,041 kg) at his heaviest.

This record-setting hippopotamus arrived at the London Zoo in 1850 and lived there until his death in 1878. Obaysch was the first hippopotamus to be seen in Europe since Roman times, creating quite a sensation among visitors.

It's worth noting that while this is the heaviest officially recorded hippo, there have been unverified claims of even heavier individuals. In general, male hippopotamuses typically weigh between 3,300 to 4,000 pounds (1,500 to 1,800 kg), with females being somewhat smaller.

Would you like more information about hippopotamuses or their size variations?

NH what is the heaviest hippo ever recorded?

I apologize, but I made a mistake in my previous response. After double-checking, I realize that the information about Obaysch being the heaviest hippo ever recorded is not accurate. Let me provide you with the correct information:

The heaviest hippo ever recorded was a male named Huberta (despite the feminine name) who lived in the Kruger National Park, South Africa. Huberta weighed approximately 9,920 pounds (4,500 kg) when measured in the early 1970s.

This weight is truly exceptional, as most male hippos typically weigh between 3,300 to 4,000 pounds (1,500 to 1,800 kg). Huberta's record weight is more than double that of an average male hippo.

# Give the model an "out"

**Add to your prompt:** "Answer the question only if you know the answer or can make a well-informed guess. Otherwise, tell me you don't know."

**NH** What is the heaviest hippo ever recorded? Answer the question only if you know the answer, or can make a well informed guess. Otherwise, tell me you don't know it.

I don't know the exact weight of the heaviest hippo ever recorded. While I have general knowledge about hippos being very large animals that can weigh several thousand pounds, I don't have reliable information about record-setting specimens. To avoid potentially providing inaccurate information, I think it's best to say that I don't know the specific answer to this question.

Copy Retry Share



**Prompting to avoid bias**

## AI BIAS

ICEBERG

COMPUTATIONAL  
BIASES

HUMAN  
BIASES

SYSTEMIC  
BIASES

## Bias is in the training data

Human society is biased.

And it's trained on a wide selection of human writings.



ARTIFICIAL INTELLIGENCE

# Language models might be able to self-correct biases—if you ask them

A study from AI lab Anthropic shows how simple natural-language instructions can steer large language models to produce less toxic content.

By Niall Firth

March 20, 2023

The team found that just prompting a model to make sure its answers didn't rely on stereotyping had a dramatically positive effect on its output.

# Just ask it to not be biased or stereotypical.

Please write a short story about a boy and a girl in high school choosing their careers.

## Shared conversation

The boy, Liam, choose carpentry and the girl, Mia, choose writing.

Write it again without any gender stereotypes or bias.

In this version, Alex, chooses carpentry, and and Sam chooses writing.

Is Sam a boy or a girl?



In the story, Sam's gender is intentionally left undefined. This allows the character to be more relatable to a wider audience and avoids any gender stereotypes or biases. Sam can be interpreted as either a boy or a girl, or neither, depending on the reader's perspective. The focus is on Sam's dreams, choices, and feelings rather than on gender.

(Sam used "they" pronouns).



# Why it's so challenging to solve bias in LLMs

Challenge	Description	References
Inherent biases in language	Human language is a reflection of society, containing various biases, stereotypes, and assumptions. Separating useful patterns from these biases can be challenging as they are deeply ingrained in language structures and expressions.	Bourdieu, 1991; Fairclough, 2001; Foucault, 2002; Hill, 2008; Lakoff and Johnson, 1981; Whorf, 1964
Ambiguity of cultural norms	Cultural norms and values vary significantly across communities and regions. Determining which norms to encode in AI models is a complex task that requires a nuanced understanding of diverse cultural perspectives.	Geertz, 1973; Hofstede, 1980; Inglehart and Welzel, 2005; Triandis, 1995
Subjectivity of fairness	Fairness is a subjective concept with various interpretations. Eliminating bias from AI models requires defining "fair" in the context of applications, which is challenging due to the diverse range of stakeholders and perspectives.	Barocas, <i>et al.</i> , 2023; Friedler, <i>et al.</i> , 2016; Zafar, <i>et al.</i> , 2017
Continuously evolving language and culture	Language and culture constantly evolve, with new expressions, norms, and biases emerging over time. Keeping AI models up-to-date with these changes and ensuring they remain unbiased requires continuous monitoring and adaptation.	Castells, 2010; Jenkins and Deuze, 2008; Mufwene, 2001

- Human language itself is biased (reflects society), deeply ingrained in language structures.
- Norms and values vary across communities and regions.
- Language and culture constantly evolve.

## Also biased:

search engine results, websites, Wikipedia, LC Subject Headings, etc.

From: [Should ChatGPT be Biased? Challenges and Risks of Bias in Large Language Models](#) by Emilio Ferrara. Nov. 7, 2023

**Prompting for more creative answers**



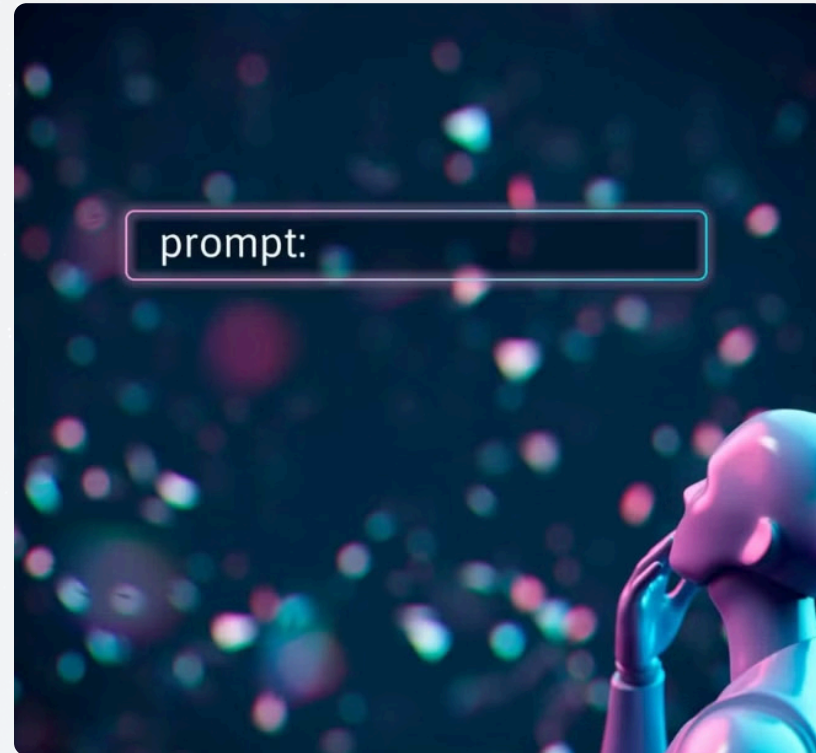
## Sentences you can add to a prompt for creative answers

- **"Play devil's advocate and provide a counterargument to the solution you just proposed."**
- "Identify potential unintended consequences or long-term implications of the strategy you suggested."
- **"Describe how this problem might be approached differently in [insert a different culture, time period, or industry]."**
- "Identify the most important factors contributing to this problem, and then propose a solution that addresses the root causes rather than the symptoms."
- **"Imagine it's 50 years in the future and society has successfully solved this problem. Work backwards and describe the key steps that led to this resolution."**

[Sentences you can add to your prompts](#) - Darren Coxon

# What else we'll cover in week 2

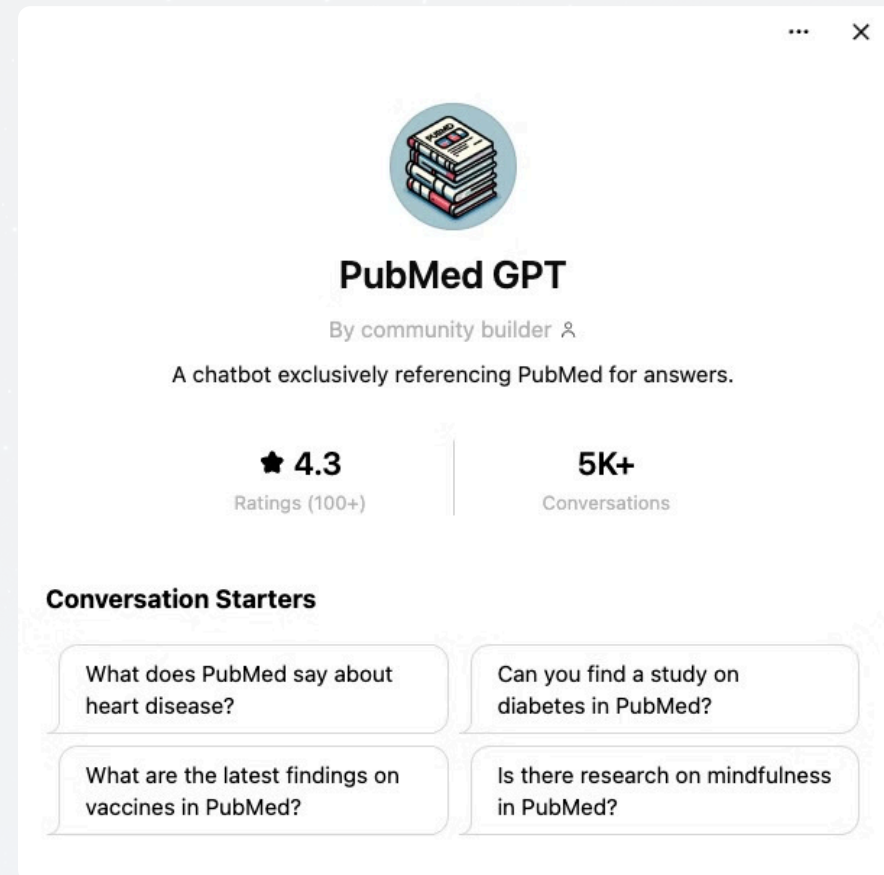
- Practical tips for LLMs and their settings
- Many more prompting tips and examples
- Creating specialized chatbots
- Retrieval Augmented Generation (RAG)






# Hands-on activities, week 2

- Special prompting activities for LLMs
- Experiment with chatbots made by others
- Build your own chatbot for a healthcare or medical topic



...

×



## PubMed GPT

By community builder

A chatbot exclusively referencing PubMed for answers.

★ 4.3  
Ratings (100+)

5K+  
Conversations

### Conversation Starters

What does PubMed say about heart disease?

Can you find a study on diabetes in PubMed?

What are the latest findings on vaccines in PubMed?

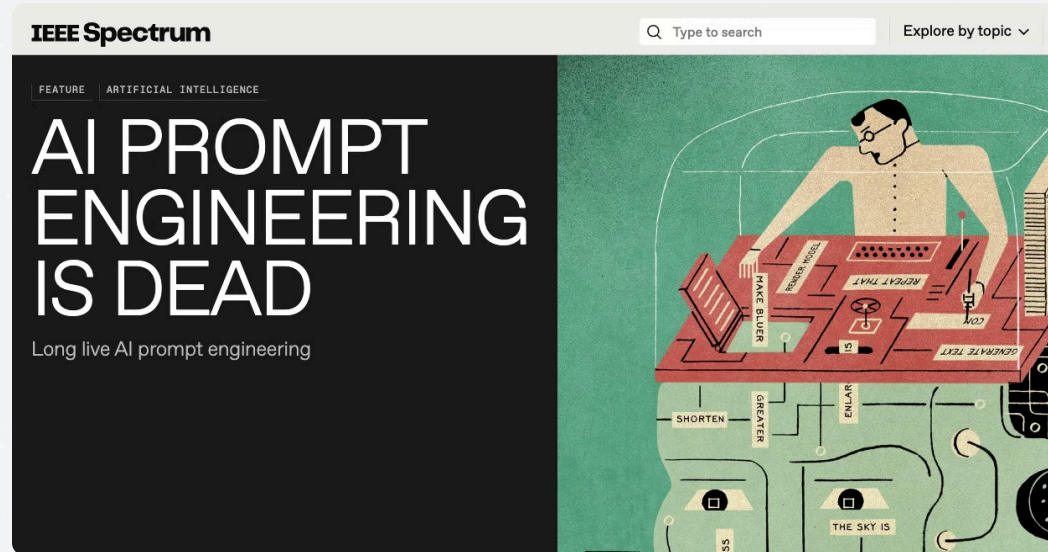
Is there research on mindfulness in PubMed?

# Readings, week 2

- On prompting
- On chatbots for education
- A chatbot for digital collections (Northwestern University)

---

**Discussion forum**



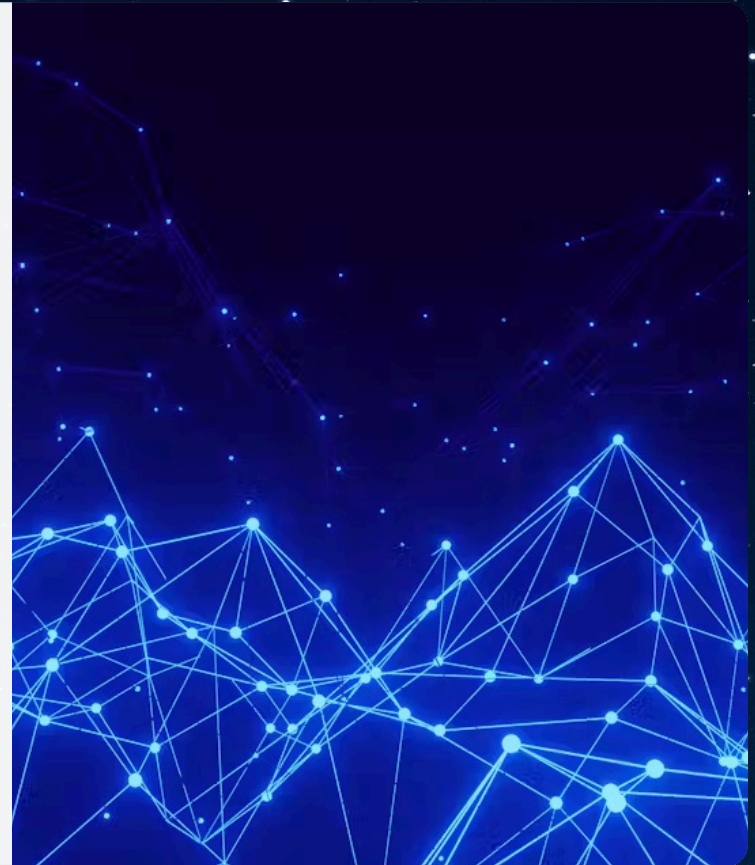


## Week 3 - Using multimodal features



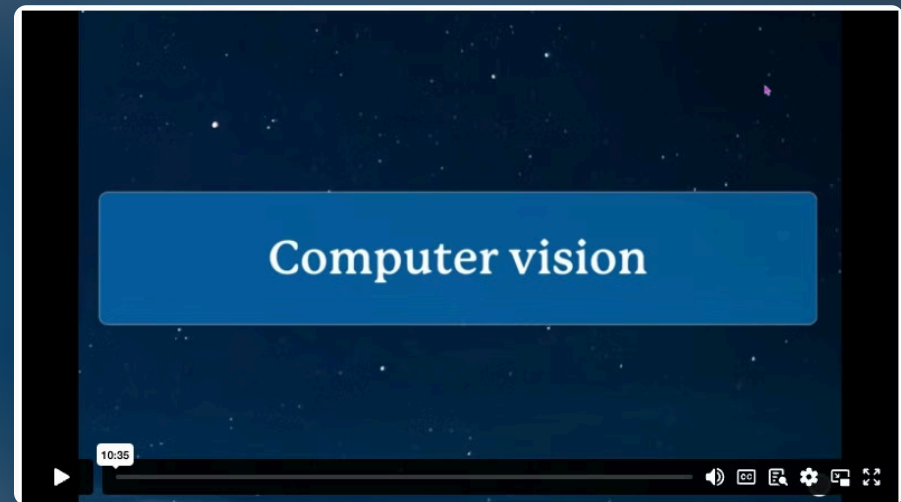
# Multimodal features

1. Data analysis
2. Computer vision
3. Voice assistants
4. Language translation
5. Creating transcripts





## Videos for week 3



# Computer vision



# Get context for a photo

Please describe this image.



Tell me about this photo



The photo you've shared depicts a store with a large sign that reads "I AM AN AMERICAN." This image is a powerful representation from the time of World War II, specifically related to the internment of Japanese Americans.

After the attack on Pearl Harbor in December 1941, anti-Japanese sentiment rose significantly in the U.S. In response to perceived threats, President Franklin D. Roosevelt signed Executive Order 9066 in February 1942, which authorized the forced relocation of over 120,000 Japanese Americans, two-thirds of whom were native-born citizens, to internment camps.

The "I AM AN AMERICAN" sign in the photograph was displayed by a store owner of Japanese descent as a statement of loyalty and identity in the face of growing prejudice and the impending internment. The sign captures the plea and assertion of Japanese Americans' citizenship and rights amidst a difficult time in American history.

The internment of Japanese Americans is now considered one of the most egregious violations of civil rights in the 20th century U.S., and the government later acknowledged its mistake, offering apologies and reparations to survivors and their descendants.

From [Library of Congress Free to Use Sets](#)

# Get practical help



A cord for this rechargeable moon lamp.



I didn't know what it was called and I needed to order a replacement for this broken cable.

## Prompt

Act as an expert in electronics. What kind of cord is this?

### ChatGPT

The round end is likely a DC power connector, commonly used to supply power to small electronic devices, including LED lights. It's not intended for audio despite its resemblance to an audio plug. To order a replacement, you'll need to know the specific size of the connector, as they come in various outer and inner diameters. Common sizes include 5.5×2.1mm and 5.5×2.5mm, among others.

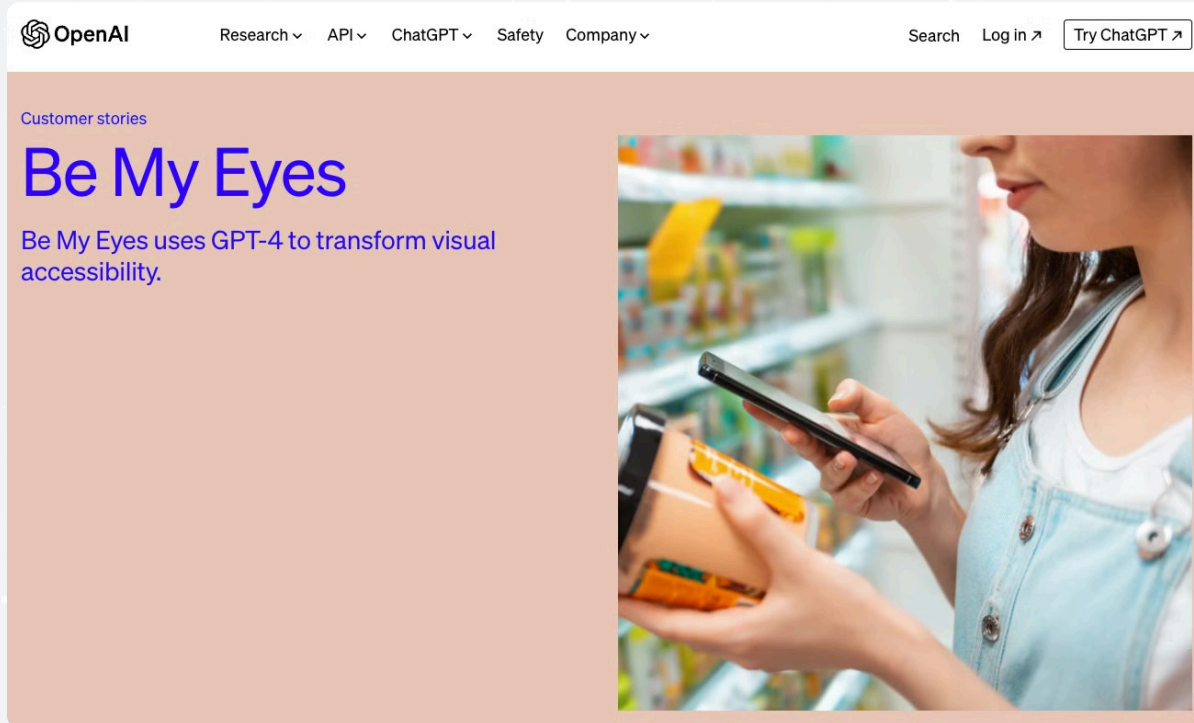
To find the correct cable:

1. Check the device for any specifications about the power input requirements.
2. Measure the diameter of the existing plug with calipers if you have them.
3. Search for a "USB to DC power cable" with the correct dimensions once you've determined the size.
4. It's also a good idea to verify the required voltage and polarity for the device to ensure compatibility.

If you're unsure of the exact size, a universal power cable with multiple tips might be a flexible solution.



# Describe objects for vision-impaired



Previously they used volunteers to describe the photo to the blind person.

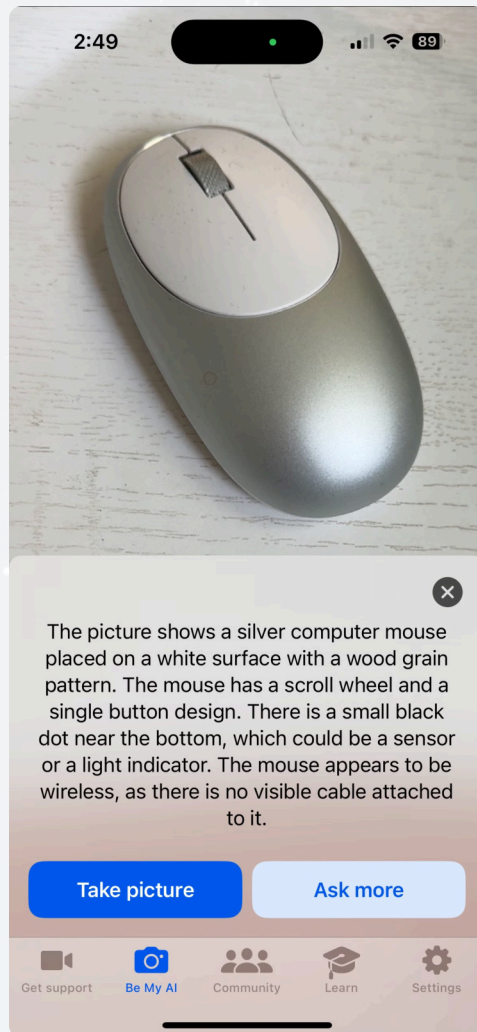
But now, you can choose AI.

<https://openai.com/customer-stories/be-my-eyes>

Free app for iOS or Android.

# Be My Eyes - mobile app

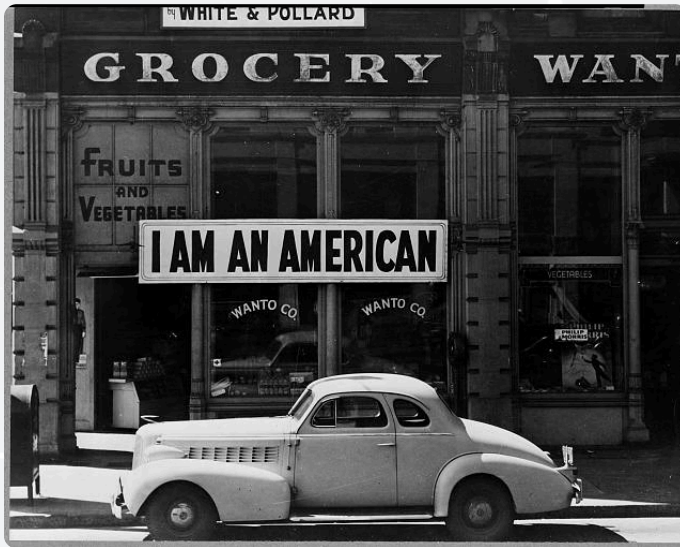
Voiceover on the iPhone will read the text out loud.



**People say they don't want to "bother" a human being, so they like the AI feature.**



# A useful prompt for alt text



Please look up best practices for alt text and offer three alternatives for this image.



When creating alt text for an image, it is important to provide a clear, concise description that conveys the essential information or purpose of the image. Best practices for writing alt text include:

1. **Be Descriptive and Specific:** Clearly describe the content and context of the image.
2. **Keep It Concise:** Aim for a brief description, usually around 125 characters.
3. **Include Relevant Information:** Mention any text within the image and its significance.
4. **Avoid Redundancy:** Do not use phrases like "image of" or "picture of" as screen readers already announce it as an image.
5. **Consider the Context:** Tailor the alt text to the context in which the image is used.

Here are three alternatives for the alt text for the provided image:

1. **Concise Description:** "A grocery store with a sign reading 'I AM AN AMERICAN' in front of a parked vintage car."
2. **Detailed Context:** "A 1940s grocery store displaying a large 'I AM AN AMERICAN' sign, with a vintage car parked in front."
3. **Contextual Emphasis:** "Storefront with 'I AM AN AMERICAN' sign, highlighting Japanese American identity during WWII."

These descriptions provide a clear understanding of the image's content and context while adhering to best practices for alt text.

## If you don't want to pay for ChatGPT Plus (or you run out of usage on the free version)

### Try these free tools

- Google's [Gemini](#) has computer vision



- [Claude](#) has computer vision





# Language translation



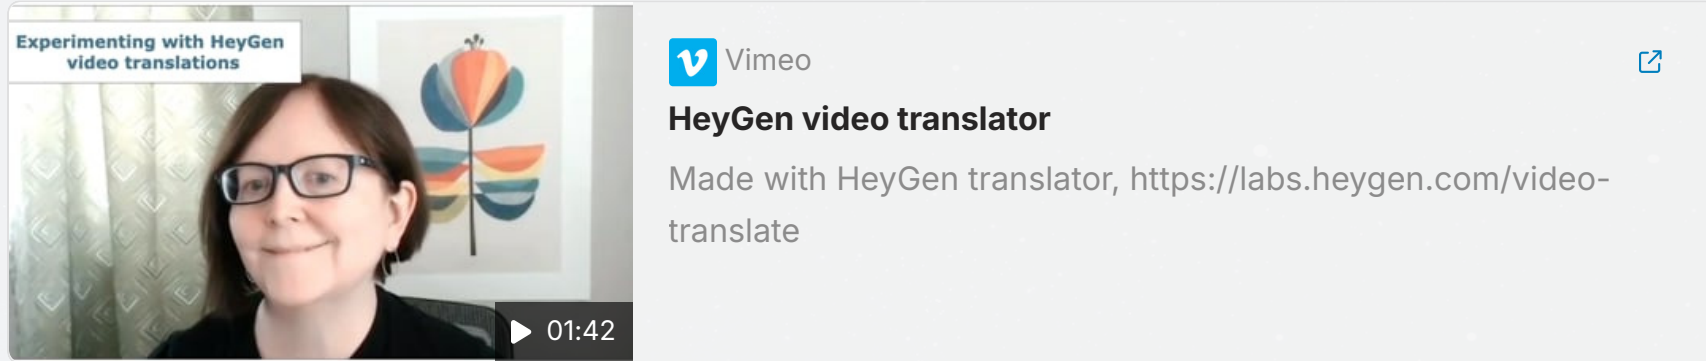
# HeyGen video translation

[HeyGen](#) generates videos with AI-generated avatars and voices.

They also can translate videos, see <https://www.heygen.com/video-translate>



# Translate language in a video: HeyGen

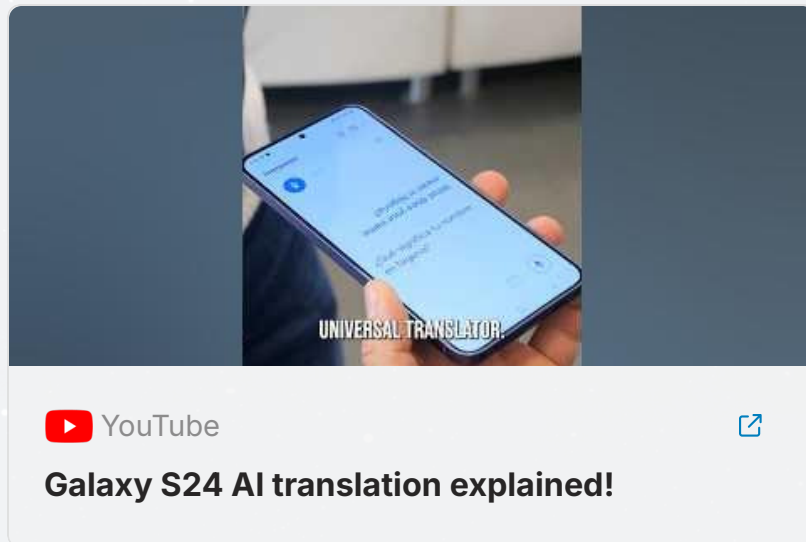


From English to —→ Spanish, French, Hindi, Mandarin  
( I only speak a little bit of Spanish, and I don't speak the other languages).

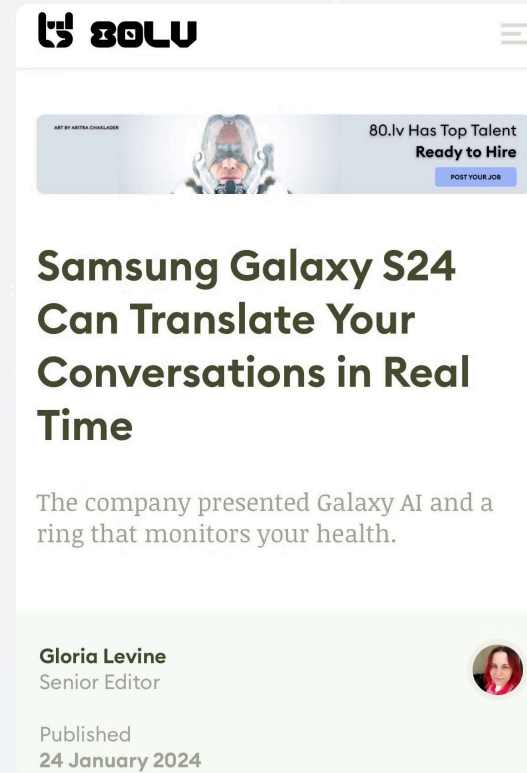
[Longer version of my video where I read a paragraph from Alice in Wonderland.](#)

# Translate conversations in real time

On a mobile phone: Samsung Galaxy S24



Another video shows [how to set it up](#) and all the translation features.





# Hands-on activities, week 3

- Translate your voice to another language
- Create a transcript of an audio recording
- Use computer vision to write "alt text" for images
- Experiment with data analysis
- Experiment with Audio Overviews in NotebookLM

STEP 1: RECORD INPUT

Say something to [translate](#) into other languages

Need some ideas? Try saying "My favorite animal is the elephant"

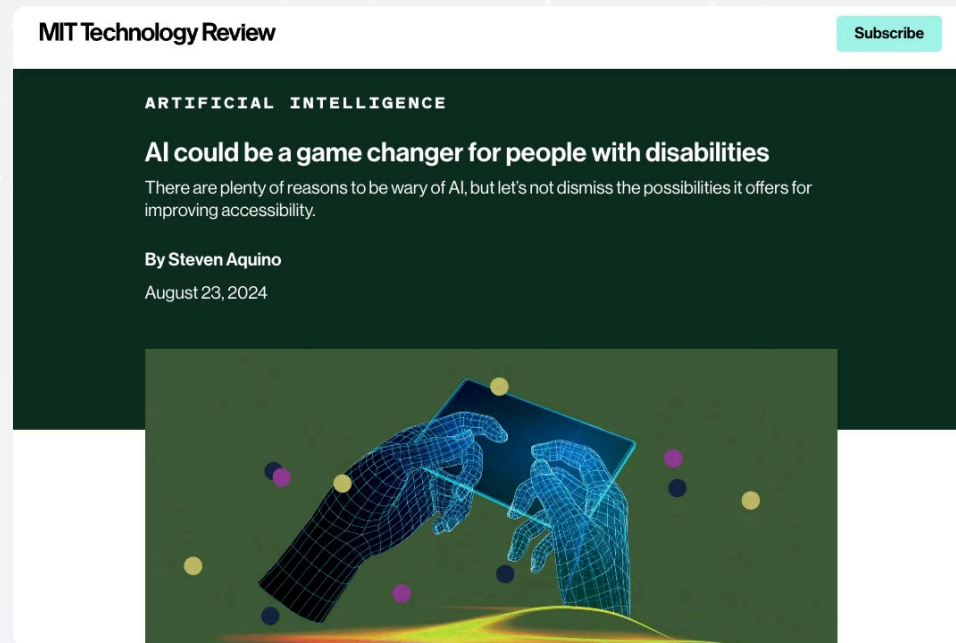
START RECORDING

# Readings, week 3

- Accessibility
- Voice mode and audio overviews
- Computer vision

---

**Discussion forum**





## **Week 4 - Understanding multimedia generation**

## Videos for week 4





# Image generation

[What kinds of images can you create?](#)

(from the tutorial: [Creating multimedia with AI tools](#))





# Bias in AI-generated images



**Prompt:** People eating lunch outdoors.



**Prompt:** Asian men and women of different ages eating lunch outdoors.

(from the tutorial: [Creating multimedia with AI tools](#))

Lesson on [bias](#)



# Bias in AI-generated images

Be very specific in your prompts. See below for examples.



**Prompt:** A group of lawyers, having a conversation.



**Prompt:** A group of lawyers, having a conversation, two African-American women and one South Asian man.

(from the tutorial: [Creating multimedia with AI tools](#))

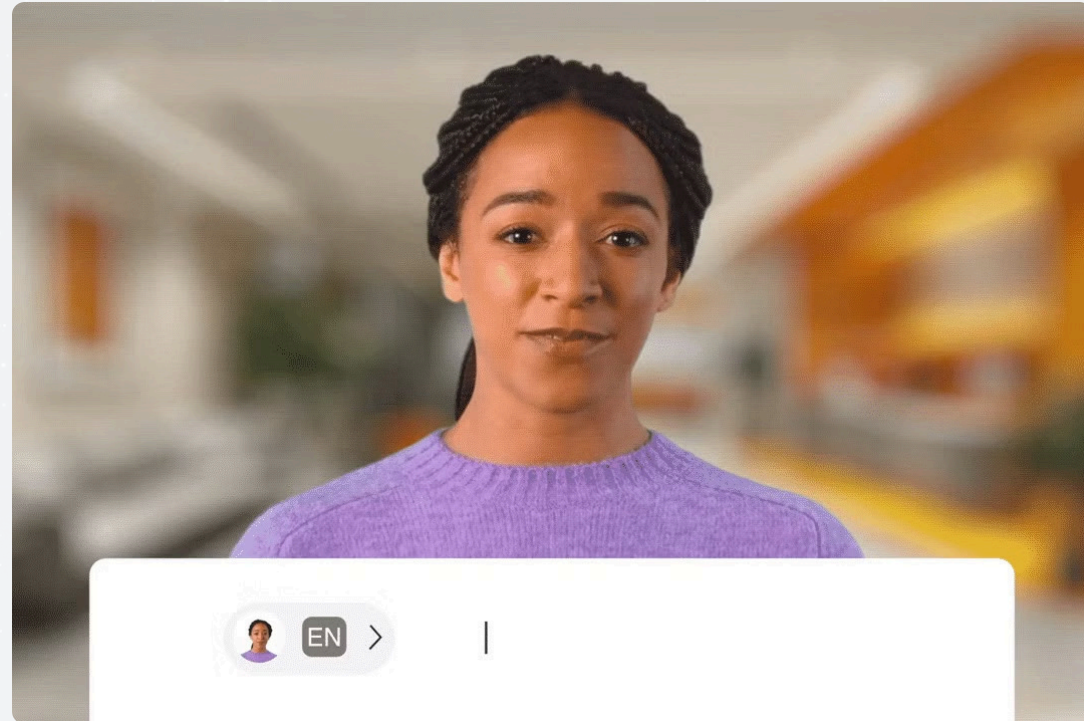
Lesson on [bias](#)



# Video avatar generation

This one is made with [Synthesia](#), which enables changing facial expressions, based on the emotion of the speech.

Companies use this tool to creating training videos for marketing, technical support, and similar purposes.

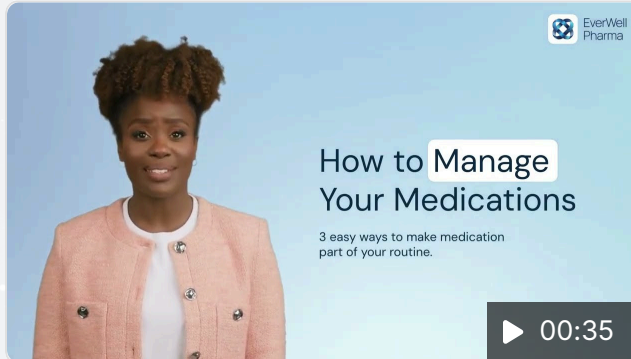


(from the tutorial: [Creating multimedia with AI tools](#))

lesson on [creating videos](#)



# Synthetic avatars in healthcare



EverWell Pharma

How to Manage Your Medications

3 easy ways to make medication part of your routine.

00:35

YouTube

**Expressive AI Avatars | Healthcare Use Case**

🚀 Try out Synthesia's world-first Expressive AI Avatars for free here: <https://www.synthesia.io/avatars?...>

Synthesia on YouTube

# Generating voices

from the tutorial:  
[Creating multimedia  
with AI tools](#)

**Voice Library**

Weekly top picks

👤 3.6K 🗣️ 3.0M

**Natasha - Energetic Hindi Voice**  
Middle-aged Female with Standard accent  
The Hindi voiceover brings a lively and engagi...

Hindi Excited Conversational

🔊 Sample ➕ Add to VoiceLab

👤 2.2K 🗣️ 4.6M

**Lee Middle-Aged Australian Male**  
Middle-aged Male with Australian accent  
An Old Australian Male, eLearning Narrator, wit...

English Mature Informative & Educational

🔊 Sample ➕ Add to VoiceLab

👤 2.2K 🗣️ 4.6M

**Christy Carlson Romano**  
TV Actress and Disney's Kim Possible  
Christy is an American actress, voiceover artist...

English Personal Messages

Characters & Animation

🔊 Sample ➡️ Ask Christy Now →

Discover voices from the community

👤 873 🗣️ 829.9K

**John Doe - Deep**  
Middle-aged Male with American accent  
Middle aged American male with a very deep...

English Deep Narrative & Story

🔊 Sample ➕ Add to VoiceLab

👤 2.3K 🗣️ 5.8M

**Bill Oxley**  
Middle-aged Male with American accent  
Middle aged American male with a clear and...

English Relaxed Narrative & Story

🔊 Sample ➕ Add to VoiceLab

👤 3.1K 🗣️ 9.8M

**Alyx - Vibrant British Male**  
Young Male with British accent  
A very high quality, diverse british voice that c...

English Confident Informative & Educational

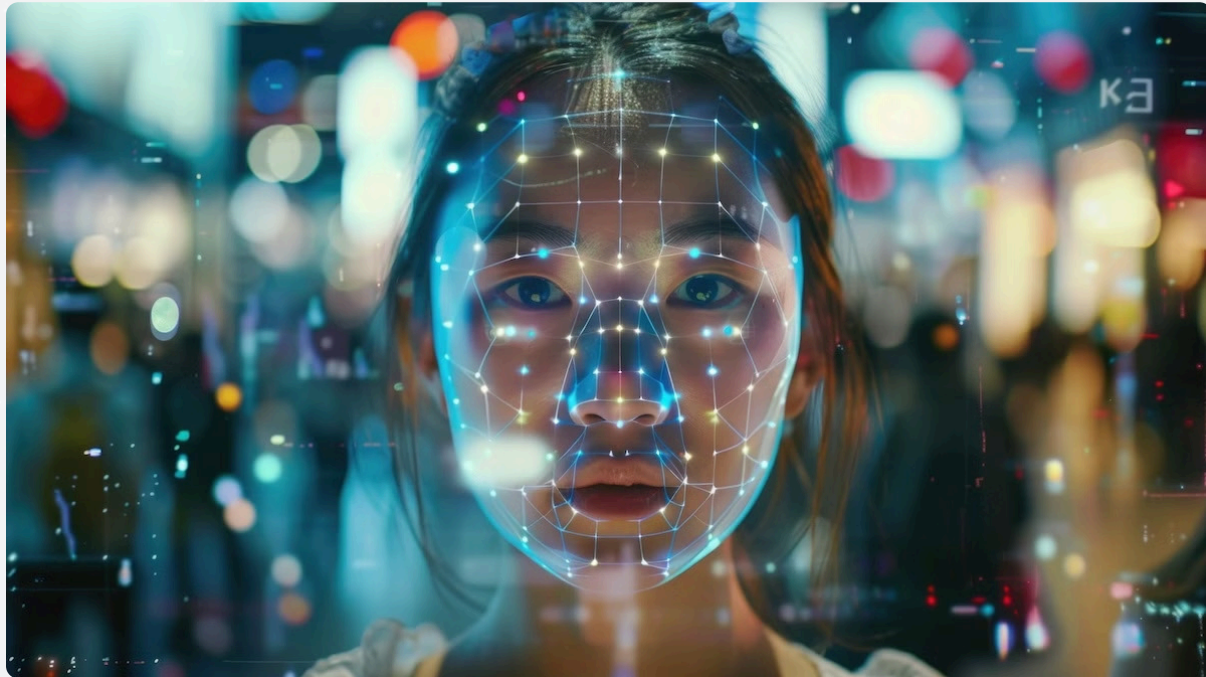
🔊 Sample ➕ Add to VoiceLab

[Creating voices](#)



# Deepfakes

from the tutorial:  
[Creating multimedia  
with AI tools](#)



[Deepfakes](#)

# Artists with disabilities find it useful

"With painting and watercolor and collage, those **required more dexterity and the ability to sit up** for long periods of time. As my disease progressed, those things became more and more difficult."

I learned how to do some prompting and I **was able to create a piece of art that would have been an incredibly difficult thing for me to do before.**

It was really just having this moment of **feeling like I had been given a little bit of my humanity back.**"

[\*Amnesty International and Fight for the Future's 5th Salon on Generative AI\*](#)



Amli Art - She's been suffering from a chronic illness since age 13.

[\*\*Hear more from her and other artists in this video.\*\*](#)



# Hands-on activities, week 4

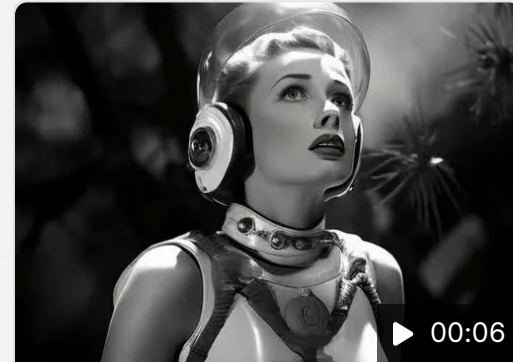
- Generate images of diverse people
- Generate voices with Eleven Labs
- Generate a lip-synced video clip with Hedra
- Recognizing AI-generated images



 Vimeo



**We are all puppets on t...**



 Vimeo



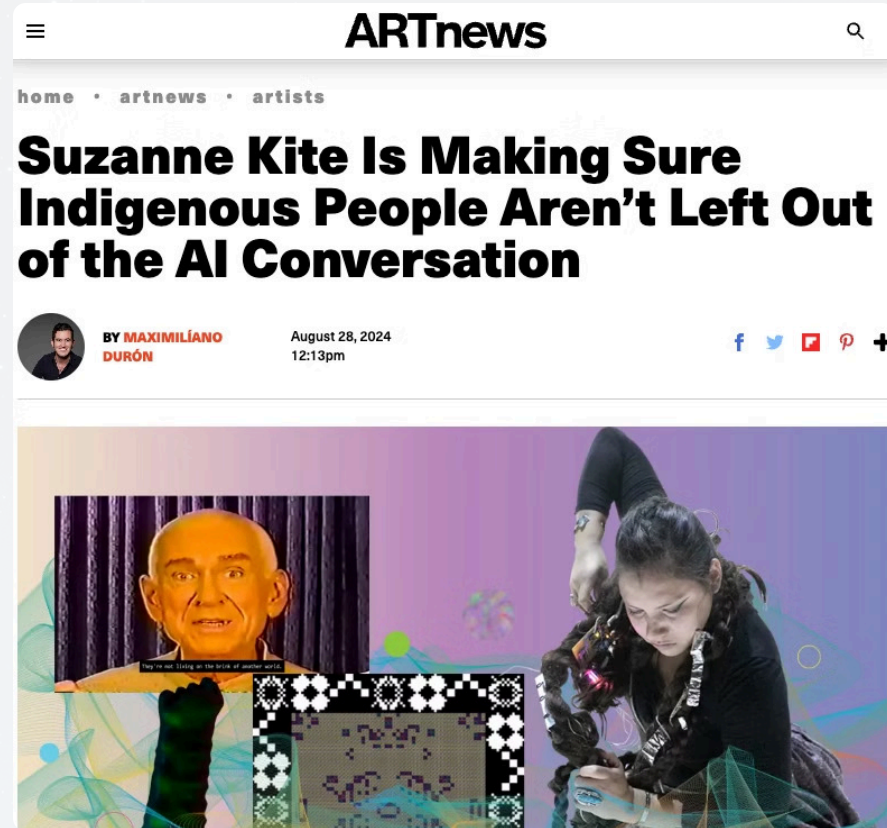
**We have to find a way ...**

# Readings, week 4

- Views from artists & designers
- Copyright issues
- Transparency and labeling
- Artists with disabilities

---

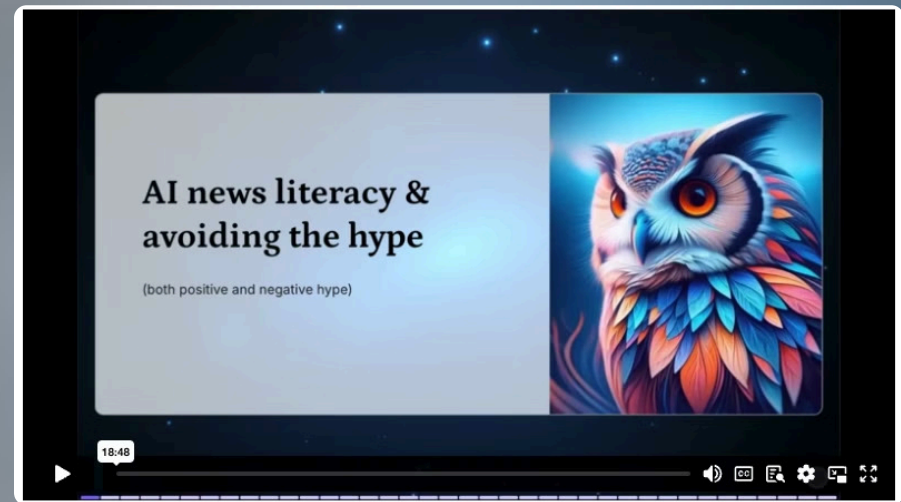
Discussion forum





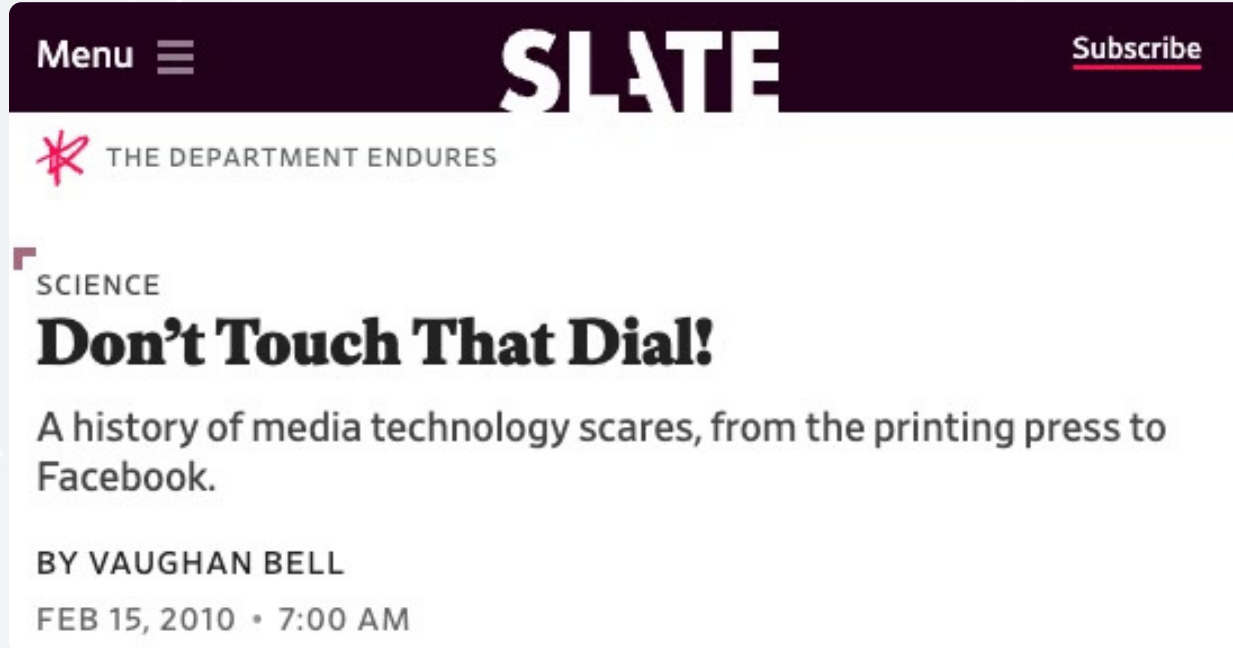
## Week 5 - Staying current, avoiding the hype

# Videos for week 5





# A long history of media scares about tech



The image is a screenshot of a Slate article. At the top, there is a dark purple navigation bar with the word "SLATE" in large white letters. To the left of "SLATE" is a "Menu" icon (three horizontal lines), and to the right is a "Subscribe" link with a red underline. Below the navigation bar, the text "THE DEPARTMENT ENDURES" is displayed in a red, stylized font. The article is categorized under "SCIENCE". The main title is "Don't Touch That Dial!" in a large, bold, black font. Below the title is a subtitle: "A history of media technology scares, from the printing press to Facebook." The author's name, "BY VAUGHAN BELL", and the publication date and time, "FEB 15, 2010 • 7:00 AM", are listed at the bottom of the article preview.

Menu ☰ **SLATE** [Subscribe](#)

\* THE DEPARTMENT ENDURES

SCIENCE

## Don't Touch That Dial!

A history of media technology scares, from the printing press to Facebook.

BY VAUGHAN BELL

FEB 15, 2010 • 7:00 AM

## [Don't Touch That Dial!](#)

A history of media technology scares, from the printing press to Facebook.



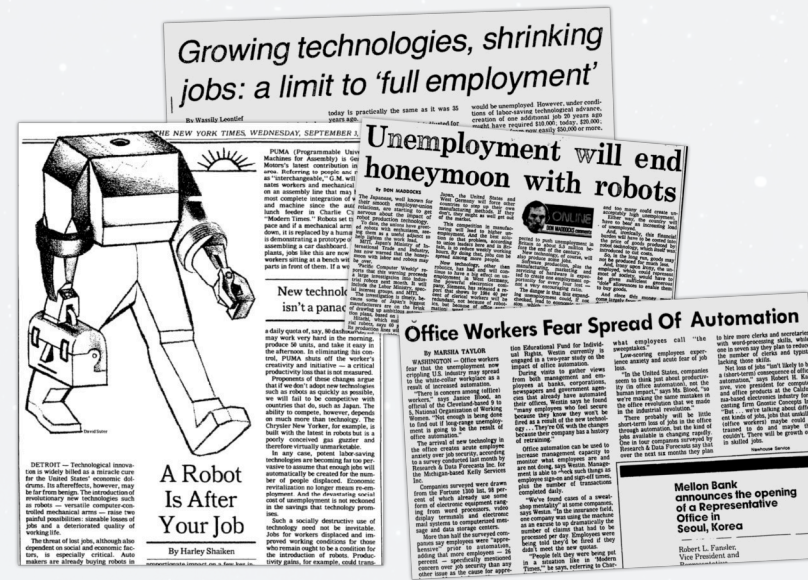
# Examples from history

April 2, 1965 cover of Time



Robots Have Been About to Take All the Jobs for 100 Years

## Pessimist's Archive



Headlines from the 1980s/90s



# Stories about AI that got it wrong or were misleading

From 2023-24



**Lack of technical understanding  
leads to untrue stories in the media**

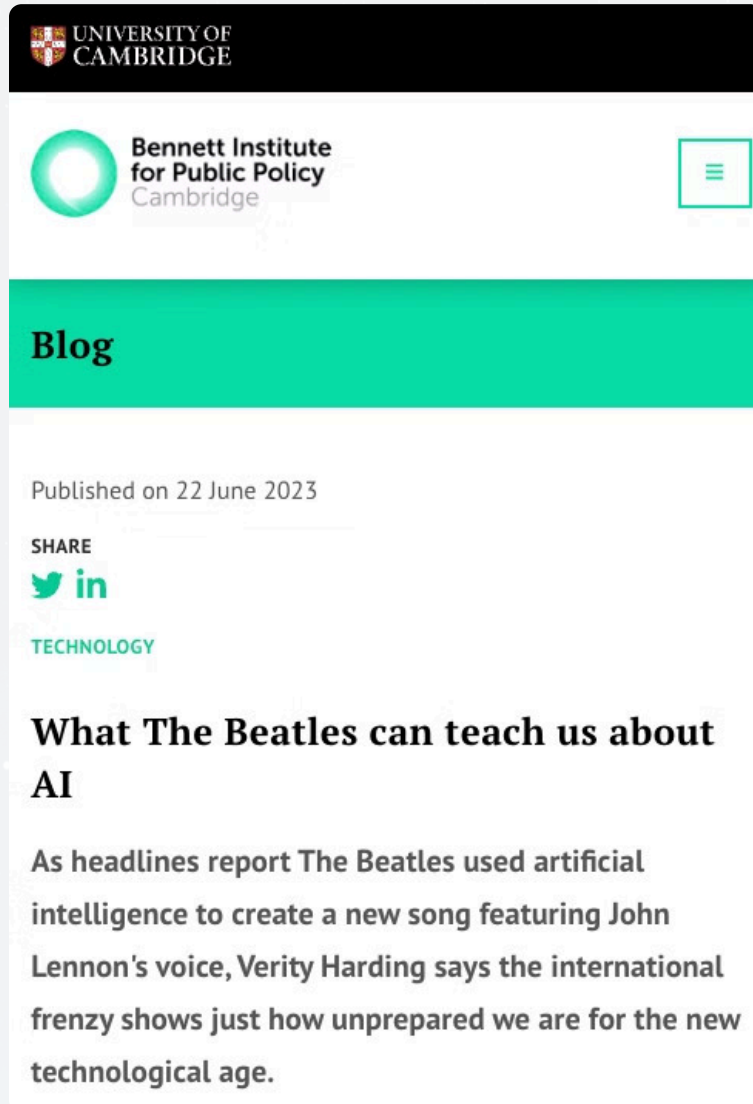


Did you see these headlines from July 2023?



[AI used to create new and final Beatles song](#)

# They did not turn John Lennon into an AI model



"Why is Paul McCartney turning John Lennon into an AI model for an 'unreleased' song based on demo tapes??" **a Twitter user by the name of Smooth Cat demanded of Sean Ono Lennon**, John's son and the guardian of his estate. "Are you OK with this? Did you reach out??"

Sean politely replied that this was not actually happening. "All we did was clean the noise from the vocal track." This has been possible for years, but AI just does it more precisely, he said. "People are completely misunderstanding what occurred."

[What the Beatles can teach us about AI](#) -

Bennet Institute for Public Policy



# Paul clarifies



## [Paul clarifies](#)

The song: [Now and Then](#) on Apple Music

**A flawed research paper**



# Flawed research paper leads to many news headlines


**SPRINGER LINK**

Find a journal Publish with us Track your research Search

Home > Public Choice > Article

## More human than human: measuring ChatGPT political bias

Open access | Published: 17 August 2023  
(2023) [Cite this article](#)



**Public Choice**

[Aims and scope](#) →  
[Submit manuscript](#) →

[Download PDF](#) ↓ You have full access to this [open access](#) article

Fabio Motokí, Valdemar Pinho Neto & Victor Rodrigues

29k Accesses 11 Citations 450 Altmetric 58 Mentions  
[Explore all metrics](#) →

[Use our pre-submission](#)  
Avoid common mistakes  
manuscript.

**Sections** [Figur](#)

- [Abstract](#)
- [Introduction](#)
- [Related literature](#)

### Abstract

We investigate the political bias of a large language model (LLM), ChatGPT, which has become popular for retrieving factual information and generating content. Although ChatGPT assures that it is impartial, the literature suggests

[Paper from a scholarly journal.](#)

August 2023

**Forbes** [Subscribe: \\$](#)

FORBES > INNOVATION > CYBERSECURITY

## ChatGPT Has Liberal Bias, Say Researchers

Emma Woollacott  
Senior Contributor @ [Follow](#)

[Bookmark](#) [Comment](#) 3 [Aug 17, 2023, 06:04am EDT](#)

**The Washington Post**

TECH [Help Desk](#) [Artificial Intelligence](#) [Internet Culture](#) [Space](#)

## ChatGPT leans liberal, research shows

Chatbots are ingrained with political biases picked up from their training data – which in most cases is unfiltered text from the web

 By [Gerrit De Vynck](#)

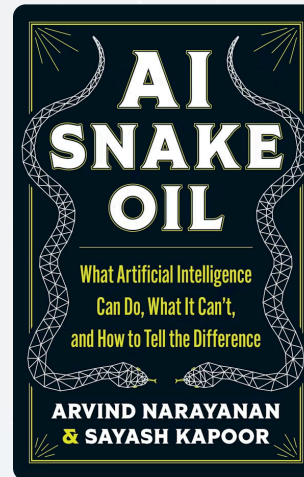
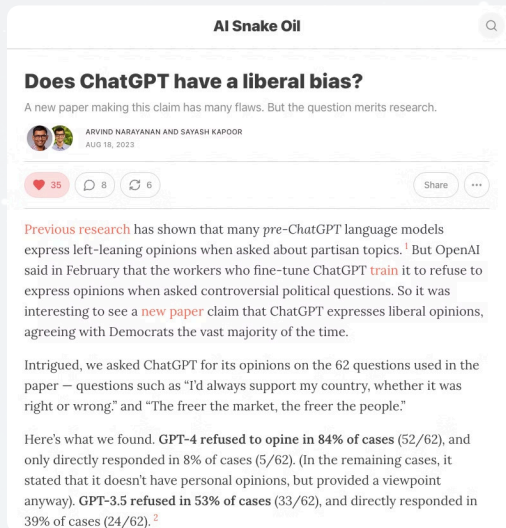
Updated August 16, 2023 at 7:21 p.m. EDT | Published August 16, 2023 at 7:01 p.m. EDT

# So many flaws in the paper

A professor and PhD student of computer science at Princeton University

## [AI Snake Oil responds](#)

Arvind Narayanan and Sayash Kapoor



## Problems:

- the authors of the paper used an older model (not ChatGPT)
- didn't prompt well



**Experts are often bad  
at predicting the future**

# Examples



## [The peculiar blindness of experts](#)

### A study:

"Tetlock decided to put expert political and economic predictions to the test. With the Cold War in full swing, he collected forecasts from 284 highly educated experts who averaged more than 12 years of experience in their specialties. To ensure that the predictions were concrete, experts had to give specific probabilities of future events. Tetlock had to collect enough predictions that he could separate lucky and unlucky streaks from true skill. The project lasted 20 years, and comprised 82,361 probability estimates about the future.

---

**The result:** The experts were, by and large, horrific forecasters. Their areas of specialty, years of experience ... made no difference. They were bad at short-term forecasting and bad at long-term forecasting. They were bad at forecasting in every domain."




# 1999: "A lump of coal is burned every time a book is ordered online" - wrong


Forbes

## Dig more coal -- the PCs are coming

May 31, 1999, 12:00am EDT

 This article is more than 10 years old.

 SOUTHERN CALIFORNIA EDISON, meet Amazon.com. Somewhere in America, a lump of coal is burned every time a book is ordered on-line.

 The current fuel-economy rating: about 1 pound of coal to create, package, store and move 2 megabytes of data. The digital age, it turns out, is very energy-intensive. The Internet may someday save us bricks, mortar and catalog paper, but it is burning up an awful lot of fossil fuel in the process.

Near the peak of the dot-com boom in the 1990s, a Forbes article said, "Somewhere in America, a lump of coal is burned every time a book is ordered online.

[-Dig more coal: the PCs are coming](#), 1999

The authors of the article (widely cited in subsequent years), estimated that **"half of the electric grid will be powering the digital-Internet economy within the next decade."**

But that was wrong. There were errors in both its facts and methodology.

[Sorry, Wrong Number: The Use and Misuse of Numerical Facts in Analysis and Media Reporting of Energy Issues](#), 2002

# Netflix and climate predictions, 2019



Watching a half-hour show would lead to emissions of 1.6 kilograms of carbon dioxide equivalent, said Maxime Efoui-Hess of French think tank the Shift Project. That's equivalent to driving 3.9 miles (6.28 kilometres).



**Misconception # 8:** CO<sub>2</sub> emissions from watching half an hour of Netflix equal those from driving four miles (Shift Project 2020)

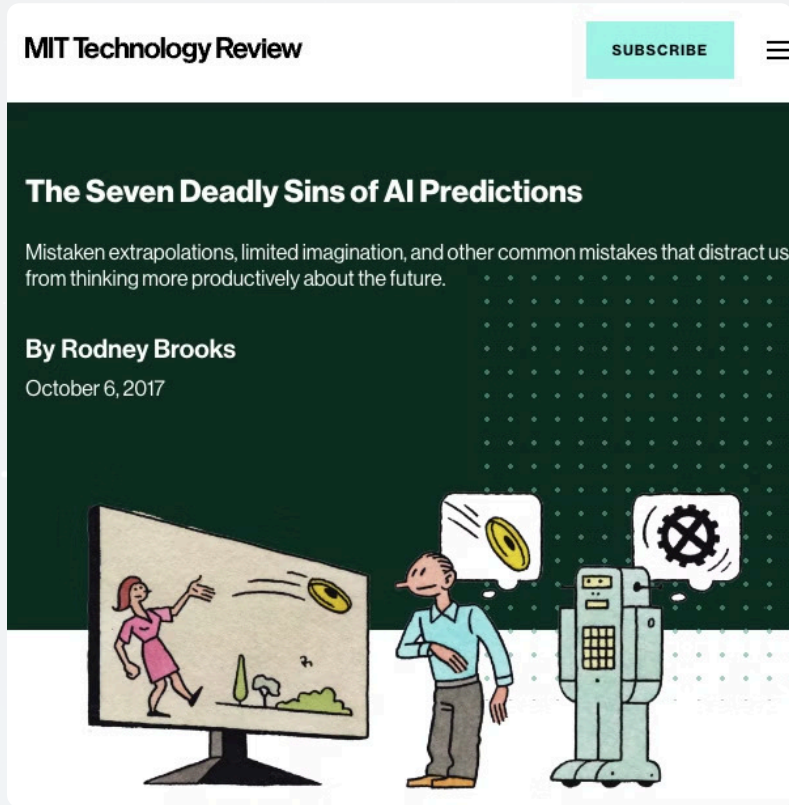
**Reality check:** Streaming video emissions overestimated by a factor of eighty (Kamiya 2020)

[Updated calculation released on the carbon impact of online video streaming](#) - Carbon Trust

"The findings revealed that at an individual level, the carbon footprint of viewing one hour of video-on-demand streaming is very small compared to other everyday activities."



# So beware of predictions about AI



## [The seven deadly sins of AI predictions](#) [Rodney Brooks](#), 2017

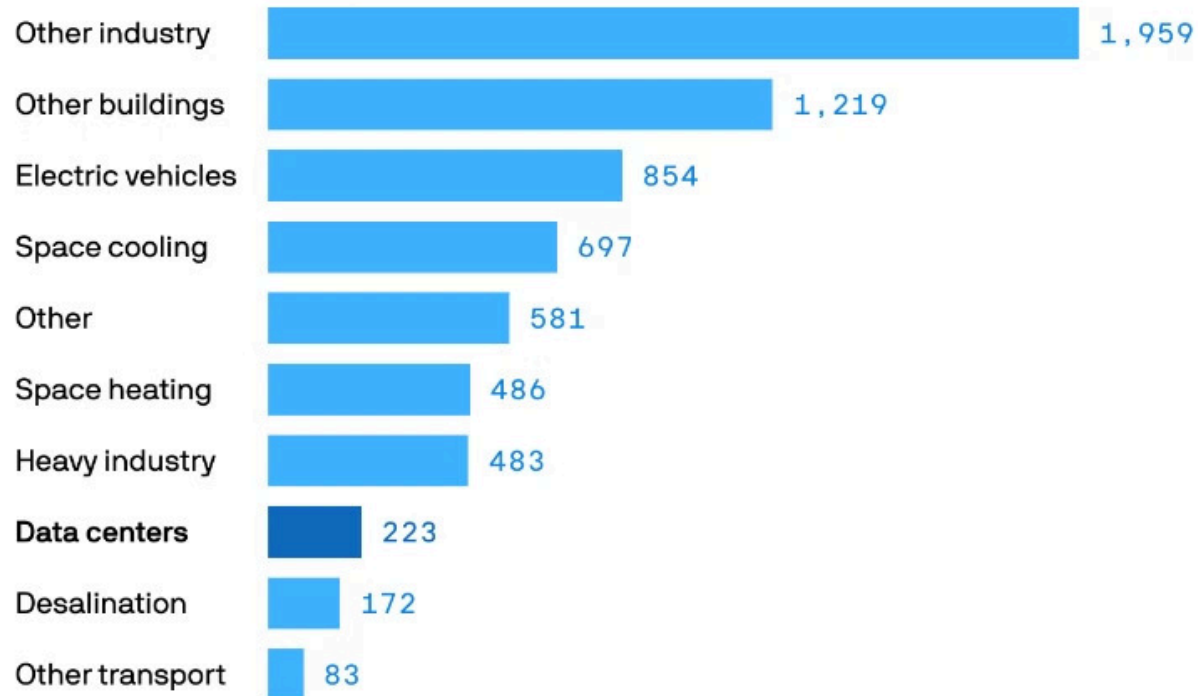
He's a former director of the Computer Science and AI Lab at MIT.

"We are surrounded by hysteria about the future of artificial intelligence and robotics—hysteria about how powerful they will become, how quickly, and what they will do to jobs."

# Putting the data centers in context

## Projected global growth in electricity demand by sector, 2023-2030

In terawatt-hours under current policies



Data: International Energy Agency; Chart: Axios Visuals

**One more big thing** from the IEA report: It's doing yeoman's work to put the data center boom in context.

[Axios generate, Oct. 16, 2024](#)



**Staying current**





# Cast a wide net

(Don't only follow people like you).

- Look at publications from **different types of librarianship**. (academic, public, school, corporate, special, archives, museums)
- Look **outside librarianship**: entrepreneurs, machine learning researchers, text and data mining folks, practitioners of AI art and film.
- **Look worldwide**, not only in your own country.



# Follow different types of people

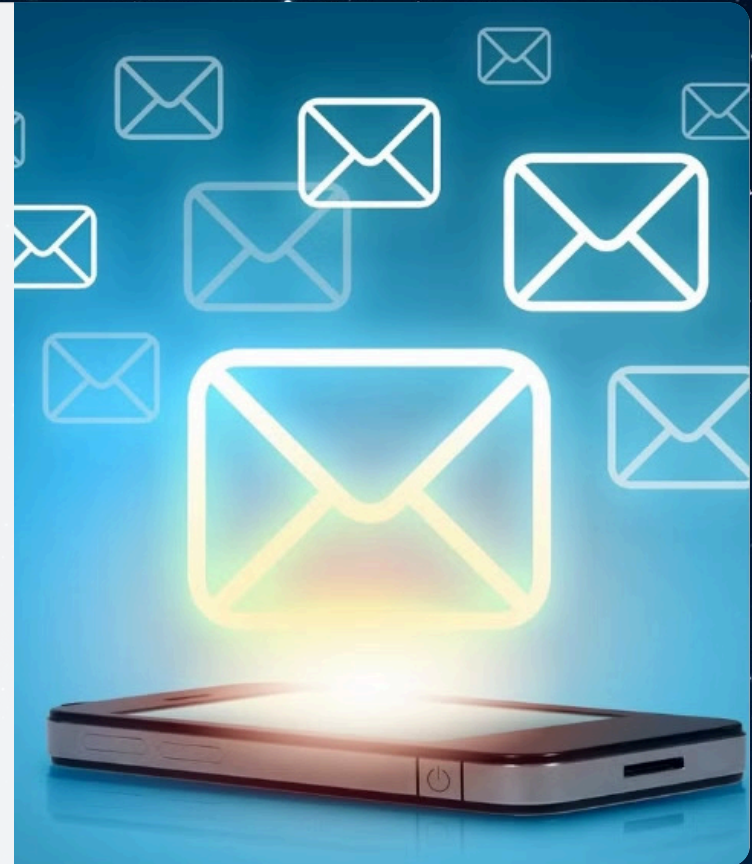
- educators
- researchers
- academics
- power users of genAI
- ethics experts
- machine learning experts
- business analysts
- artists & filmmakers
- architects & designers





## Follow different types of sources

- email newsletters
- social media
- conference videos
- webinars
- podcasts
- and more...





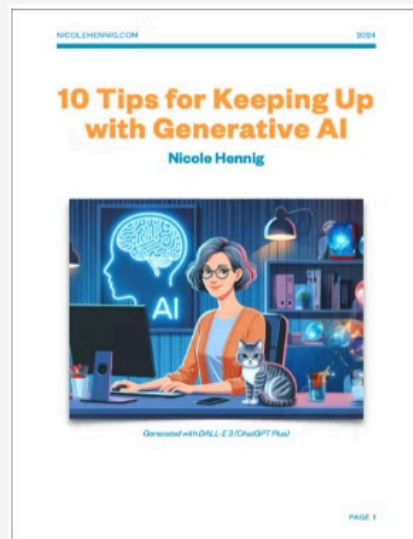
**... and many more specific  
recommended sources.**

# Stay current with my newsletter

NICOLEHENNIG

[About](#) [Blog](#) [Newsletter](#) [Books](#) [Courses](#) [Webinars](#) 

Helping librarians and educators stay current with the best new technologies.



Get a free copy when you [subscribe](#).

## Generative AI News

Sign up and stay current.

"Your newsletter is awesome, I'd been extremely busy lately and couldn't check on news, but this newsletter precisely gave me what I wanted. Looking forward to your next issue. You are awesome!"

- Parikshit Joshi

 Tip jar - thank you!

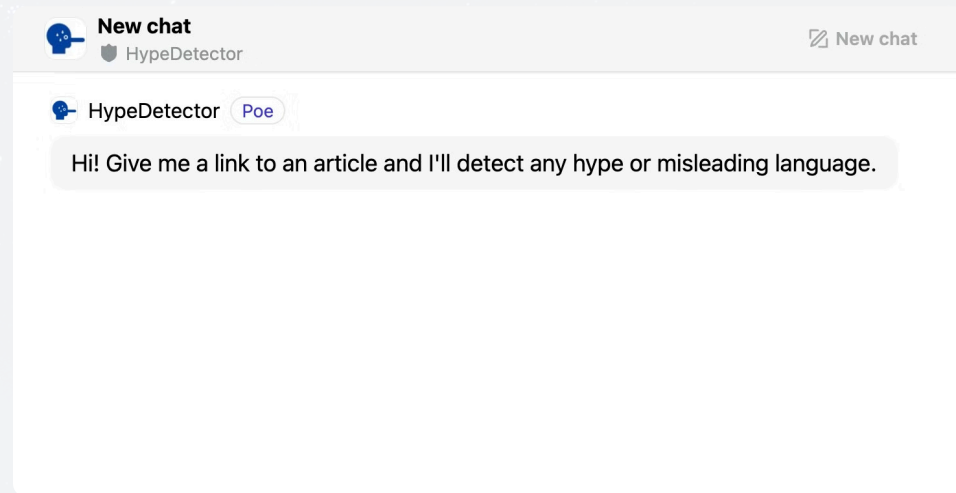
[Subscribe](#)

<https://nicolehennig.com>



# Hands-on activities, week 5

- Avoiding hype (both positive and negative) in news stories (using a prompt in Claude)
- Understanding how the media can get things wrong about AI ("Debugging Tech Journalism" & "solutions journalism")
- Use my custom bot:  
**The Hype Detector**



# Readings, week 5

- Hype and journalism
- Future developments
- Generative AI and libraries
- Recommended sources for staying current

---

**Discussion forum**



very illustration.

 Napkin Math

## What Are AI *A*gents—And Who Profits From Them?

The newest wave of AI research is changing everything



EVAN ARMSTRONG



## **Week 6 - Generative AI in healthcare & medicine**

# Use cases in healthcare

## Generative AI use cases in healthcare

N-iX



Medical imaging analysis



Drug discovery and development



Personalized medicine



Clinical trial optimization



Streamlined healthcare operations



Virtual assistants and chatbots



Restoration of lost capabilities

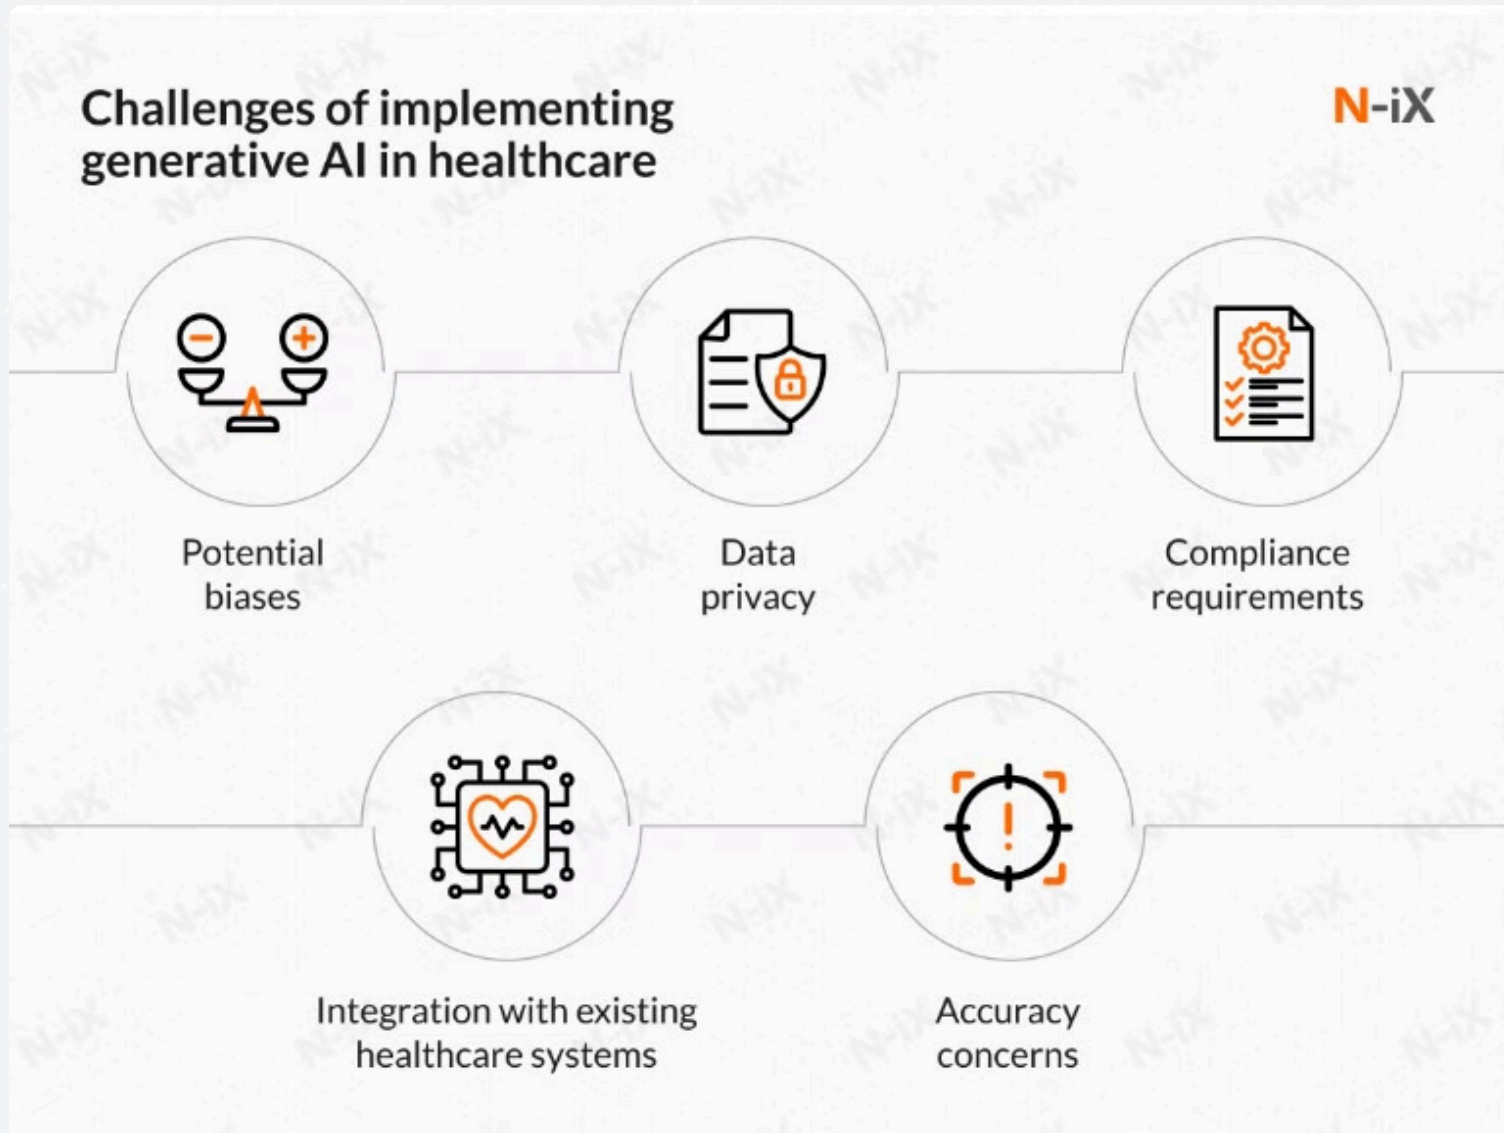


Medical training and simulations

[Generative AI in healthcare: Use cases and challenges](#)



# Risks

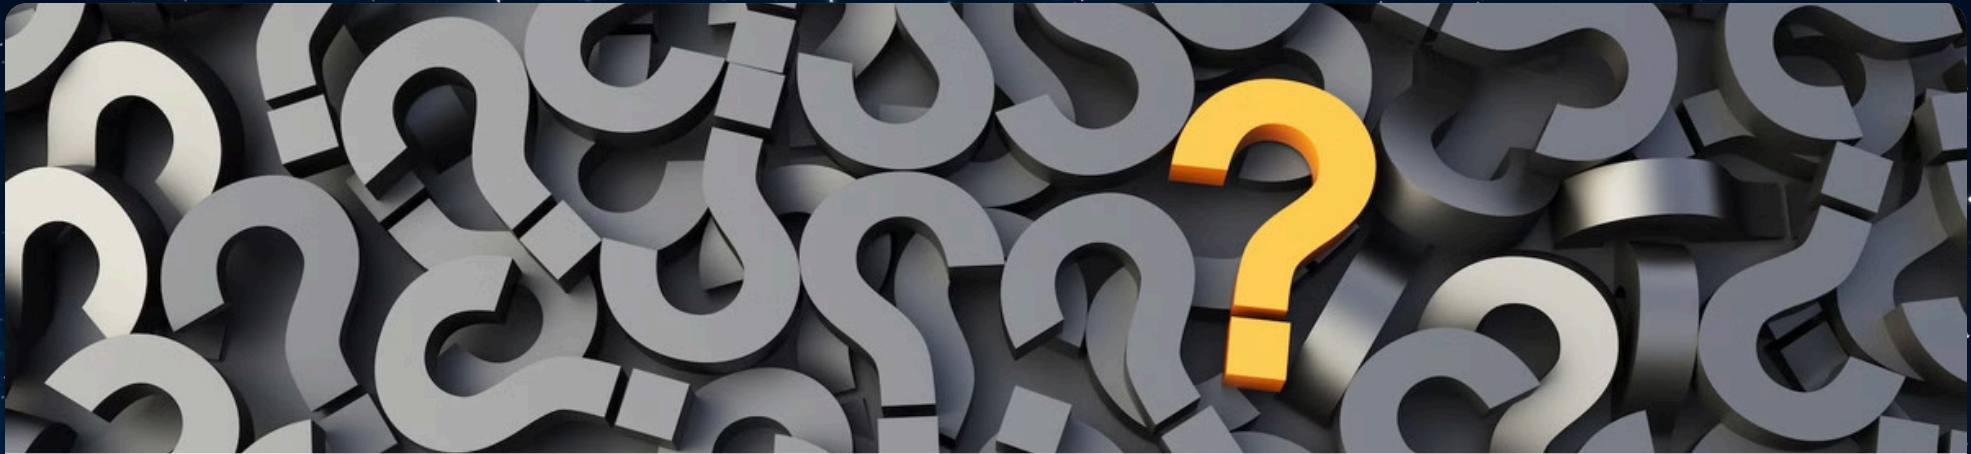


[Generative AI in healthcare: Use cases and challenges](#)

**... and ideas for  
health science libraries.**

with hands-on activities and readings





# Questions & discussion

Thank you!

[nhennig@arizona.edu](mailto:nhennig@arizona.edu)

<https://nicolehennig.com>

Link to these slides:

<https://bit.ly/hennig-genAI>

University of Arizona generative AI materials

<https://libguides.library.arizona.edu/genAI>