Introduction to AI-Enabled Teaching

Facilitated by
Cassie Collier, Assistant Professor
Computer Management and Information Systems

Adapted from Harvard Business School's Teaching with AI Program

Agenda

What is AI?

Guidelines for the use of AI for teaching

Activity: Finding the jagged frontier of AI in education tasks

Retrospective & next steps

A couple disclaimers

Everyone in this room has experience with / knowledge of AI. I'm here as a facilitator to help us all gain value from our collective experience and knowledge; not as a lecturer to purely deliver content

The ethics and values of AI use is outside the scope for this conversation. There are many ethical reasons to be resistant to the use of AI. Do not take my excitement for the topic as a wholehearted endorsement of all things AI.

What is AI?

Artificial intelligence: the science of making machines mimic human intelligence, e.g. reasoning, problem-solving, learning, and decision-making

Machine learning: a subset of AI where machines learn from data instead of being explicitly programmed

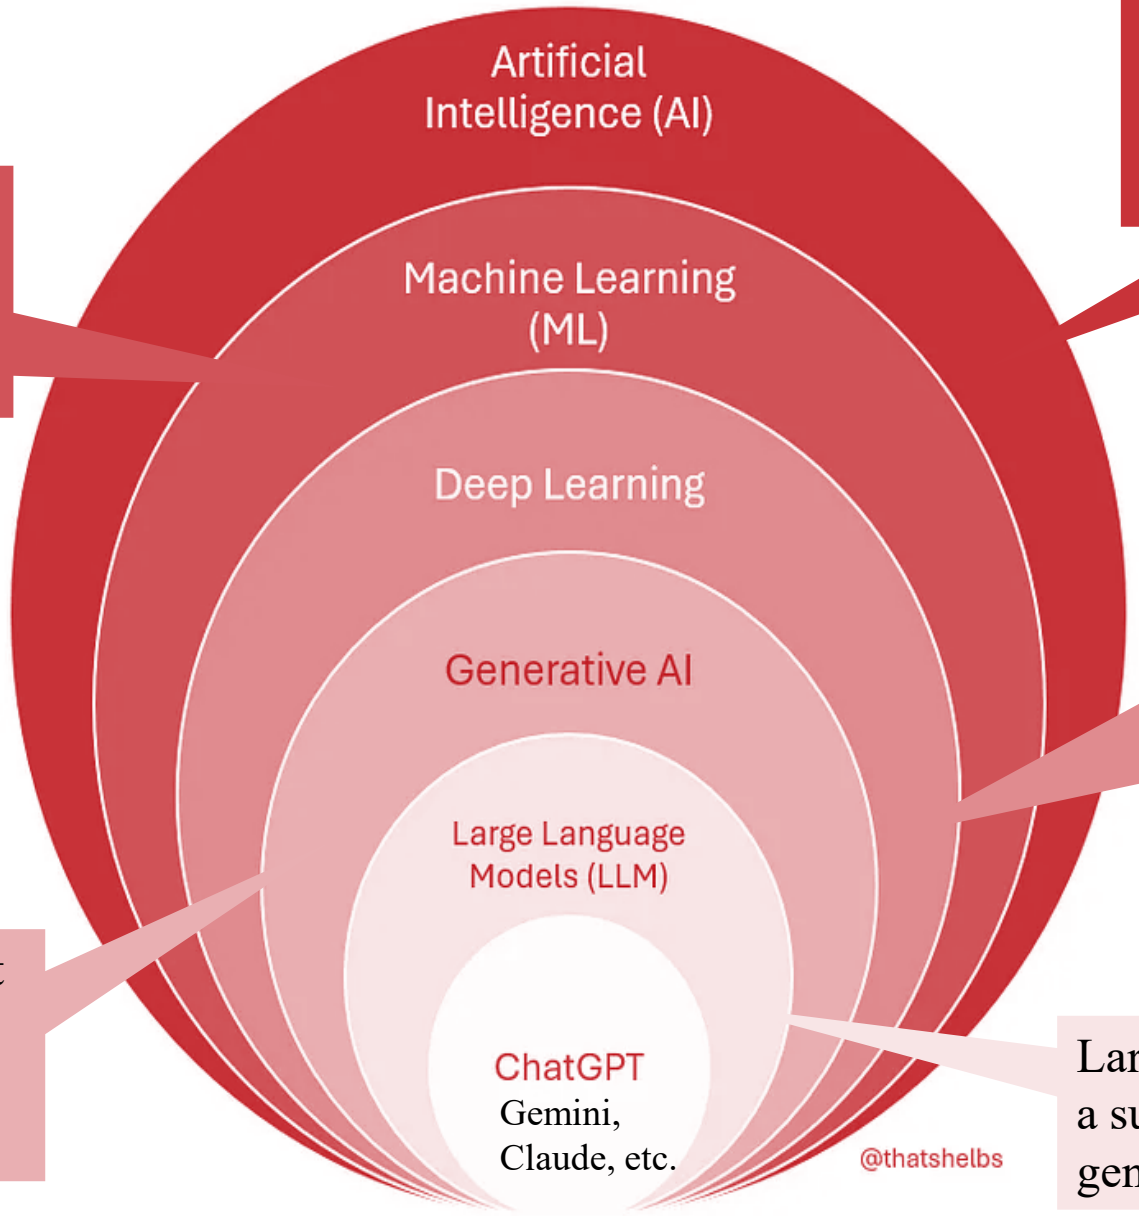

Deep learning: one type of machine learning that uses neural networks with many layers to automatically learn complex features and patterns from massive datasets

Generative AI: a subset of AI that focuses on **generation** of new content, rather than prediction, classification, or recognition

Large Language Models (LLMs): a subset of GenAI that intakes and generates human language

Large Language Models are
Next Token Predictors

Roses are **red**, violets are **blue**

How did you know that?

Four Elements of an LLM

Four Elements of an LLM

Tokenizer | Embedding layer | Transformer | Prediction layer

Input text is broken up into individual **tokens** to be analyzed (words, punctuation, brief phrases)

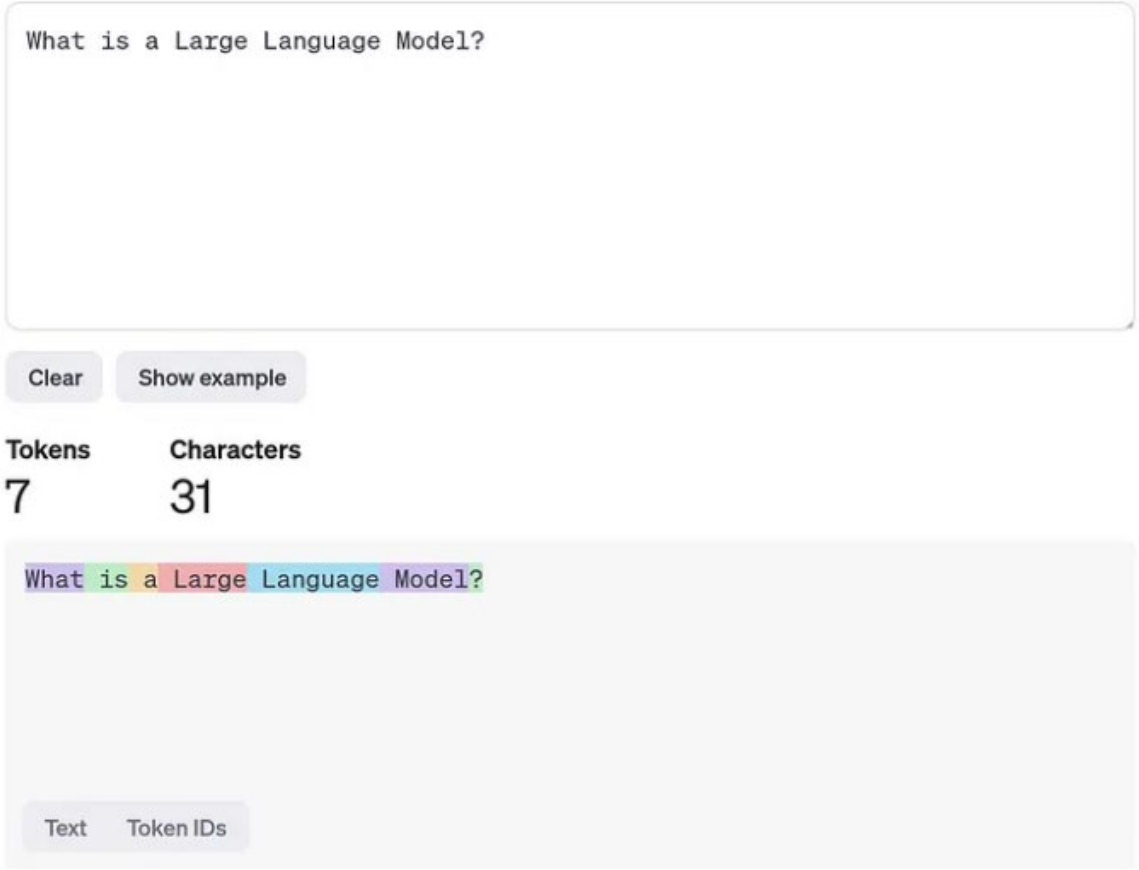

Four Elements of an LLM

Tokenizer | **Embedding layer** | Transformer | Prediction layer

Tokens are evaluated for their **relatedness** based on how often (statistically) they occur together in training data. GPTs are pre-trained in these embeddings based on data from the internet and other sources.

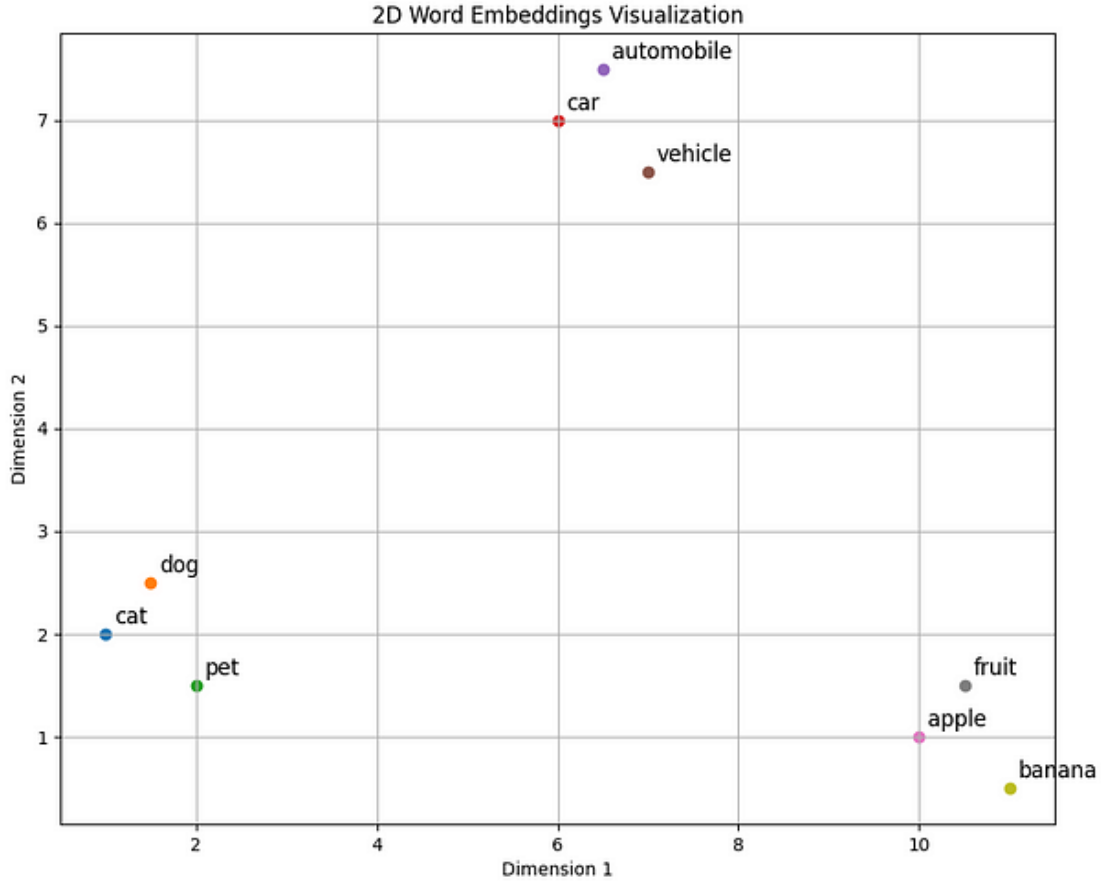

Four Elements of an LLM

Tokenizer | Embedding layer | **Transformer** | Prediction layer

The bat flew out of the cave.

Vs.

He hit the ball with his bat.

Embeddings for each token are mixed together to provide **context** (resulting in better embeddings called contextualized embeddings).

Four Elements of an LLM

Tokenizer → Embedding layer → Transformer → **Prediction layer**

The best thing about AI is its ability to

learn	4.5%
predict	3.5%
make	3.2%
understand	3.1%
do	2.9%

Produces a huge list of probabilities for the next token based on contextualized embeddings and makes a selection.

Four Elements of an LLM

Tokenizer

Input text is broken up into individual **tokens** to be analyzed (words, punctuation, brief phrases)

Tokens	Characters
7	31

What is a Large Language Model?

Embedding layer

Tokens are evaluated for their **relatedness** in an embedding layer. More related = closer together. GPTs are pre-trained in these embeddings

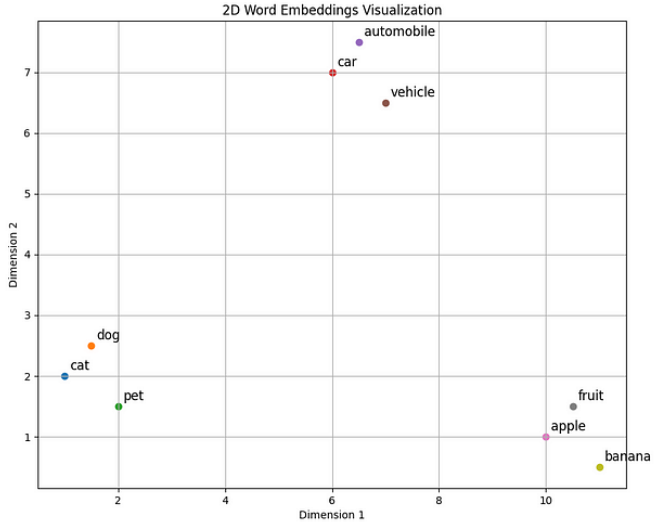

Transformer

Embeddings for each token are mixed together to provide **context** (resulting in better embeddings called contextualized embeddings).

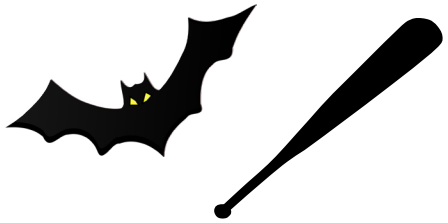

Prediction layer

Produces a huge list of probabilities for the next token based on contextualized embeddings. What we ultimately see produced by an LLM is the token with an acceptably high **probability** of being the "right" next token.

The best thing about AI is its ability to

learn	4.5%
predict	3.5%
make	3.2%
understand	3.1%
do	2.9%

It’s not magic or human thinking – it’s math!

… which is why sometimes it doesn’t work

Exposing The Brittleness Of Generative AI As Exemplified By The Recent Gibberish Meltdown Of ChatGPT

By Lance Eliot, Contributor. ⓘ Dr. Lance B. Eliot is a world-renowned AI scienti... Follow Author

Published Feb 25, 2024, 07:00am EST

- “When one user inquired what a computer was, ChatGPT allegedly replied, “It does this as the good work of a web of art for the country, a mouse of science, an easy draw of a sad few, and finally, the global house of art, just in one job in the total rest.” (Source: “ChatGPT glitches out: Rogue AI responding in nonsensical Spanglish, gibberish” by Ben Cost, *New York Post*, February 20, 2024).

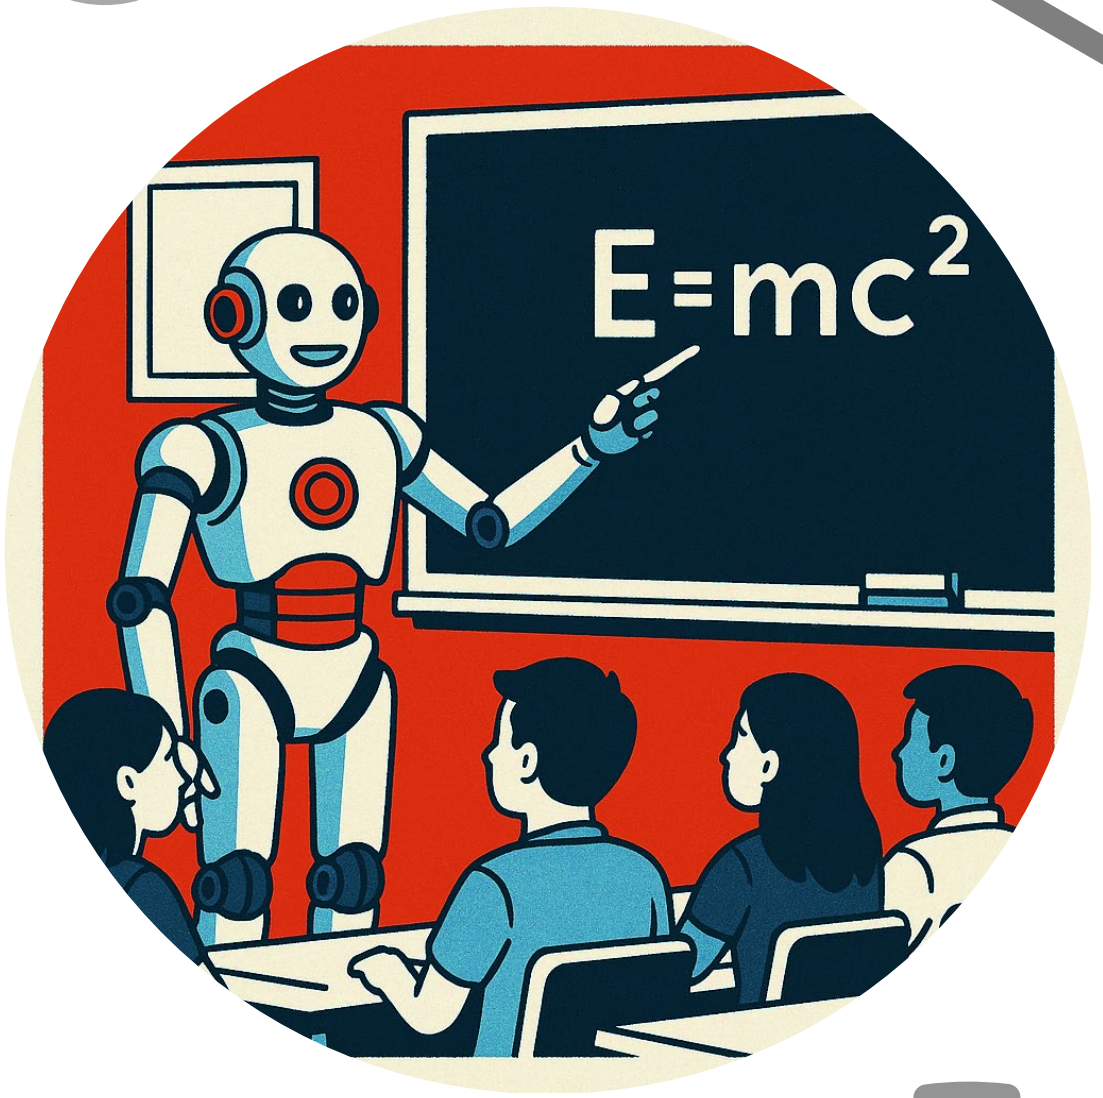

Using AI in Teaching

Photo generated by ChatGPT

If students use AI, will they still be learning?

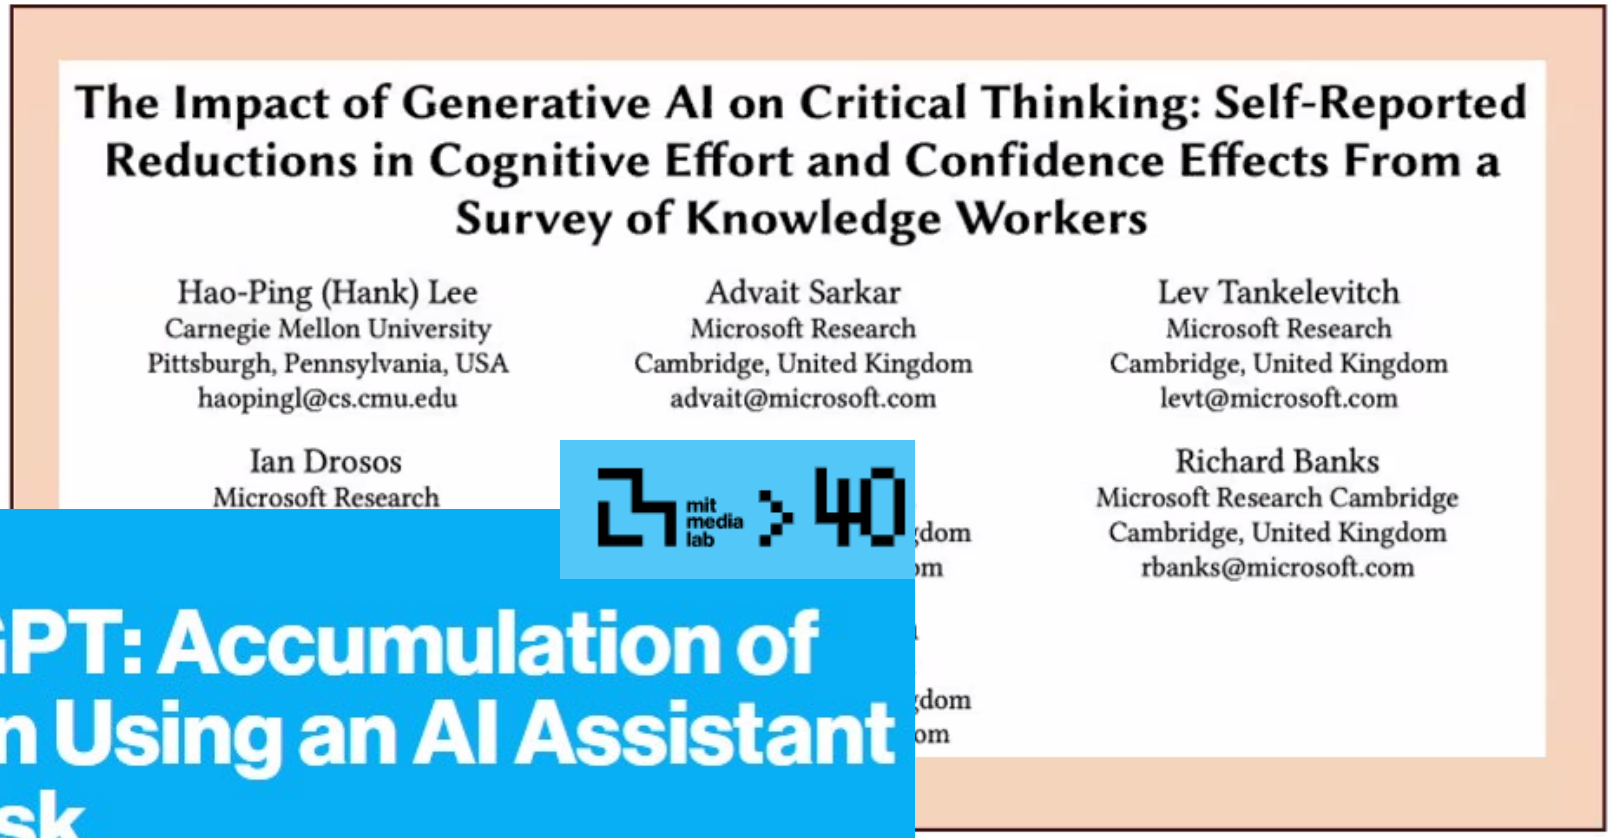

The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers

Hao-Ping (Hank) Lee
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
haopingl@cs.cmu.edu

Advait Sarkar
Microsoft Research
Cambridge, United Kingdom
advait@microsoft.com

Lev Tankelevitch
Microsoft Research
Cambridge, United Kingdom
levt@microsoft.com

Ian Drosos
Microsoft Research

Richard Banks
Microsoft Research Cambridge
Cambridge, United Kingdom
rbanks@microsoft.com

mit media lab > 40

Publication

Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task

Every new technology has brought concerns that it will lead humans to stop using our brains

Writing

Plato: ***They will cease to exercise memory because they rely on that which is written, calling things to remembrance no longer from within themselves, but by means of external marks***.

Cell phones

New York Times, 2005: ***HEY, cellphone user, when was the last time you memorized a phone number?***

Calculators

New York Times, 1985: ***...calculators could lead to "speed and ease, but not to learning."***

The 10x Educator

The goal: Proper use of AI can increase our efficiency and effectiveness by 10 times

What is the “proper” use of AI?

Demonstrating good values / adhering to policy

Technically proficient interaction with AI (crafting input + evaluating output)

(A few) Considerations Regarding AI Use in Education

1. Research & adhere to legal limitations – copyright, FERPA, etc. (*Future FDCI AI seminars will address fair use & intellectual property etc.)*
2. Enact explicit policies regarding acceptable student use of AI for a given program, class, and/or assignment (*CFDI has sample policies available for use online now in their AI Resources)*

GenAI Usage Policy Example

	Level descriptor:	What this looks like in practice:
Level 0	No use of GenAI	Students will create their own, original work without the use of GenAI for any manner.
Level 1	Organizational use of GenAI	Students will create their own, original work without the use of GenAI; however, the use of GenAI for personal efficiency (i.e., summarizing notes/readings, clarifying content) is acceptable.
Level 2	Use of GenAI for brainstorming or idea generation	Students can consult GenAI as a tool for brainstorming or idea generation, but are expected to create their own, original work without the use of GenAI.
Level 3	Use of GenAI for feedback	Students create their own work, then use GenAI as a tool to provide feedback on their work. Students are expected to use feedback from GenAI to conduct their own revisions of their own work, so any work submitted should be GenAI-supported, not Gen-AI created.
Level 4	Use of GenAI to co-create and revise work	Students can use GenAI to develop drafts/outlines of their work but are expected to carefully edit and revise GenAI-created content as appropriate for their learning context. It is expected that any use of GenAI-created content is properly disclosed and attributed.
Level 5	Unrestricted, attributed use of GenAI	Students can freely use GenAI if the use of any GenAI-created content is properly disclosed and attributed.
Level 6	Unrestricted, unattributed use of GenAI	Students can freely use GenAI in any form. Attribution is not necessary. *Note: While it it possible to use GenAI in an unrestricted manner without attribution, any who apply this level of GenAI use in their course should carefully consider ethical and legal implications of such AI use.*

Center for Integrated Professional Development - Illinois State

(A few) Considerations Regarding AI Use in Education

1. Research & adhere to legal limitations – copyright, FERPA, etc. (*Future FDCI AI seminars will address fair use & intellectual property etc.)*
2. Enact explicit policies regarding acceptable student use of AI for a given program or class (*CFDI has sample policies available for use online now in their AI Resources)*
3. Avoid the trap of "policing" AI use – aim instead to develop course content that *embraces* AI use (*main focus of 2/26 & 3/4 FDCI seminars)*
4. **Use AI well, ethically, and transparently, to increase your own effectiveness as an educator and provide a model of "good" AI use for your students.**

… but how?

Technically Proficient AI interaction

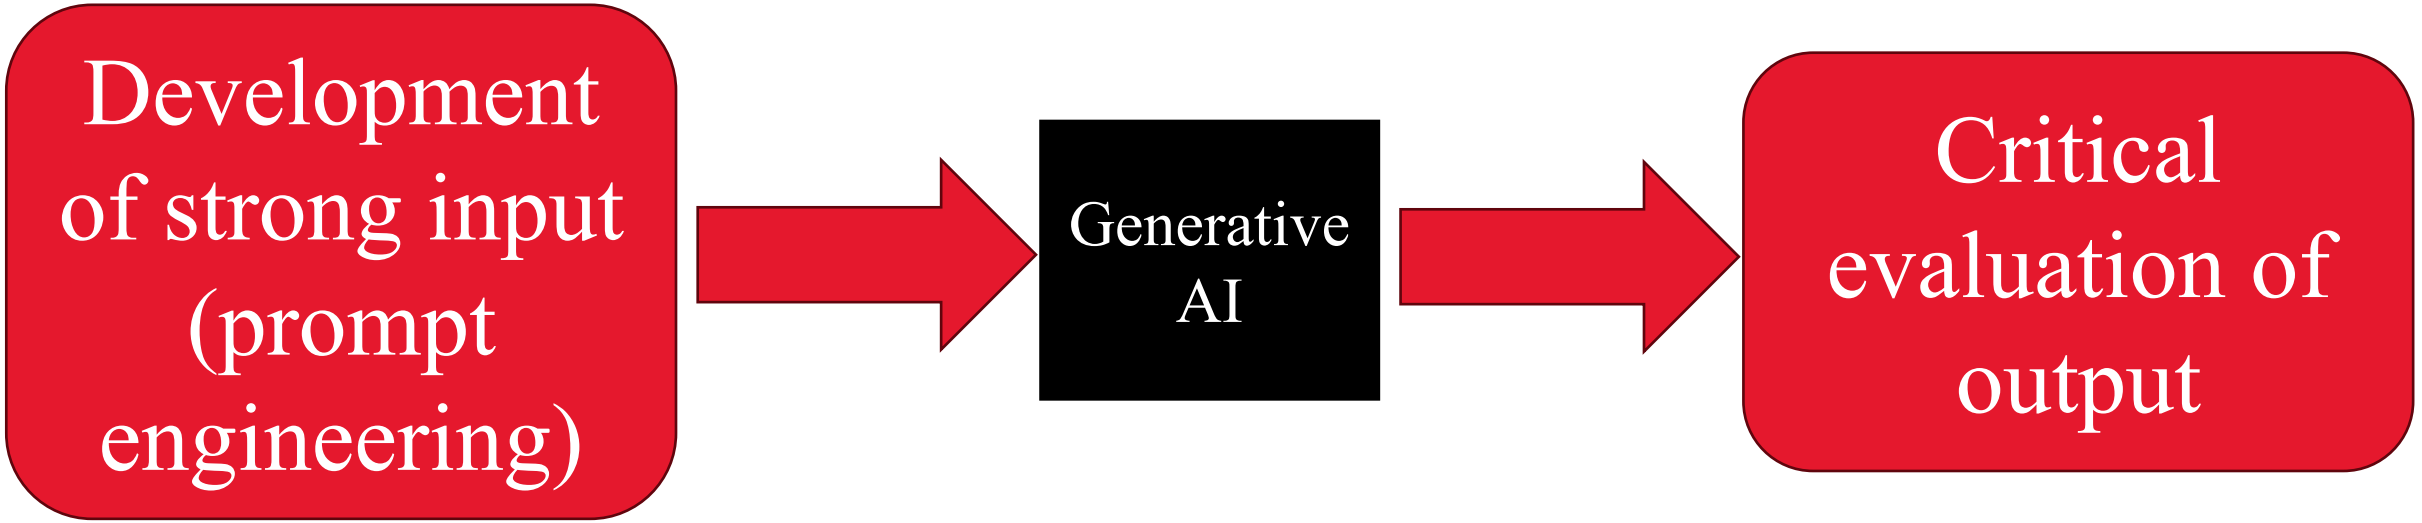

Prompt Engineering

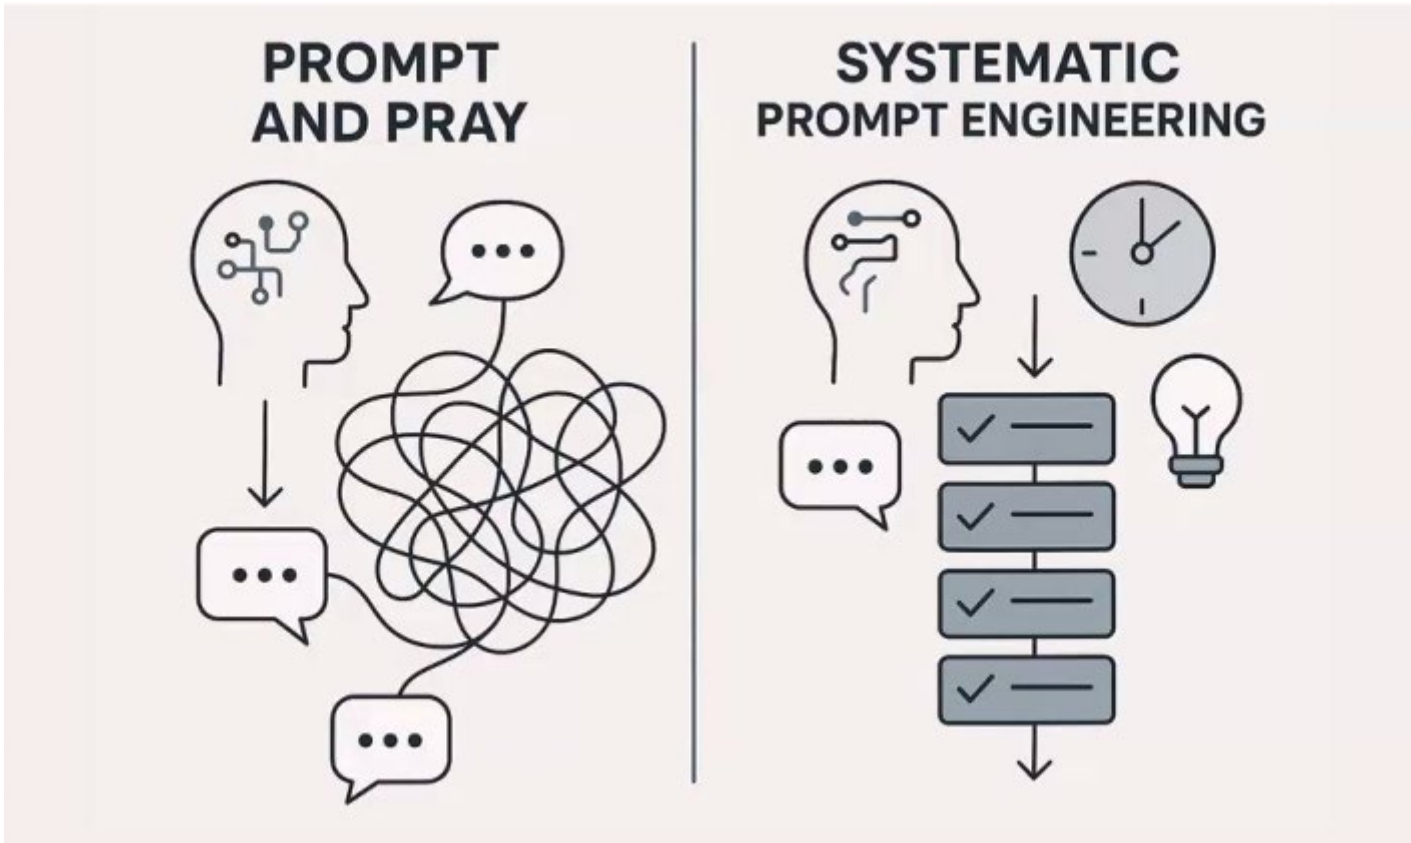

From Harvard Business Impact Webinar: Using Generative AI to Create Engaging, Customized Learning Experiences

Prompting Basics

Be specific

Assign AI a role

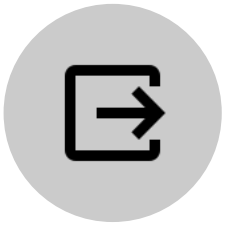

Describe preferred output

Use “do” and “don’t”

Use examples

Identify key sources

Prompting Basics, Cont’d

Consider tone and audience

Build on previous prompts

Correct mistakes and give feedback

Ask AI what else it needs

Or…
“ChatGPT, help me develop an excellent prompt to teach X topic”

Output Evaluation: (A Few) Generative AI Pitfalls & Concerns

- Hallucinations
- Sycophancy
- Vague & verbose output
- Environmental concerns
- Intellectual property concerns
- Potential biases

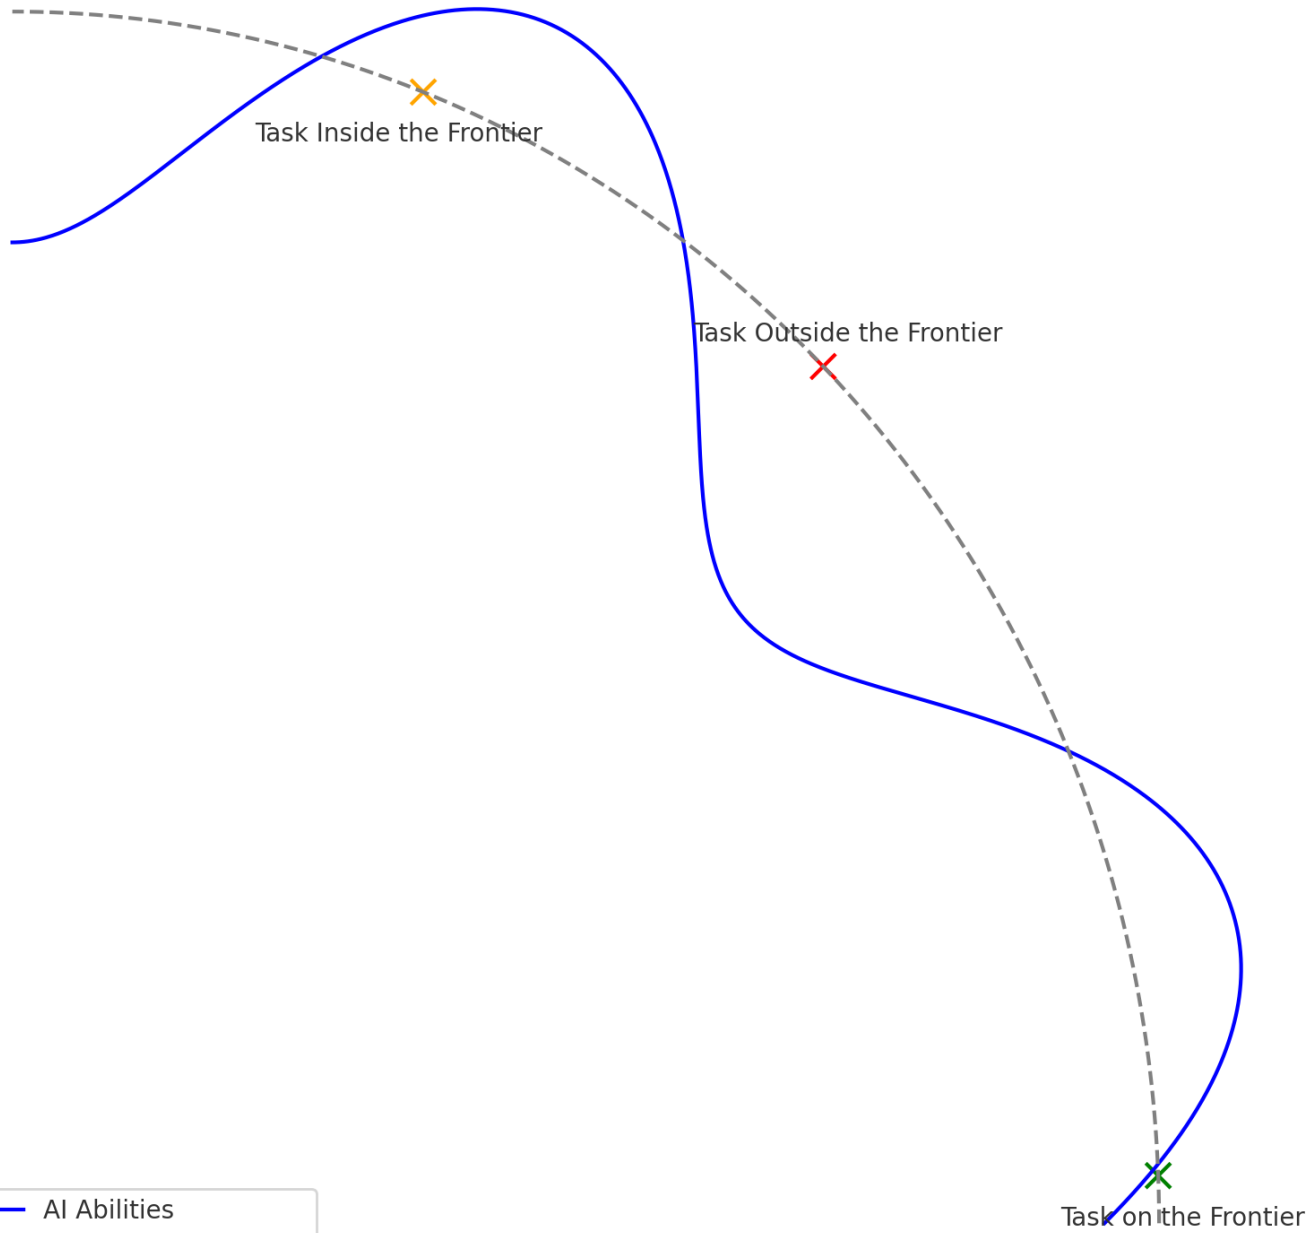

Identifying the Jagged Frontier of AI in Education

Dell'Acqua et al. 2023

AI Capabilities are Not Intuitive

AI Capabilities are Not Intuitive…

(July 25, 2025)

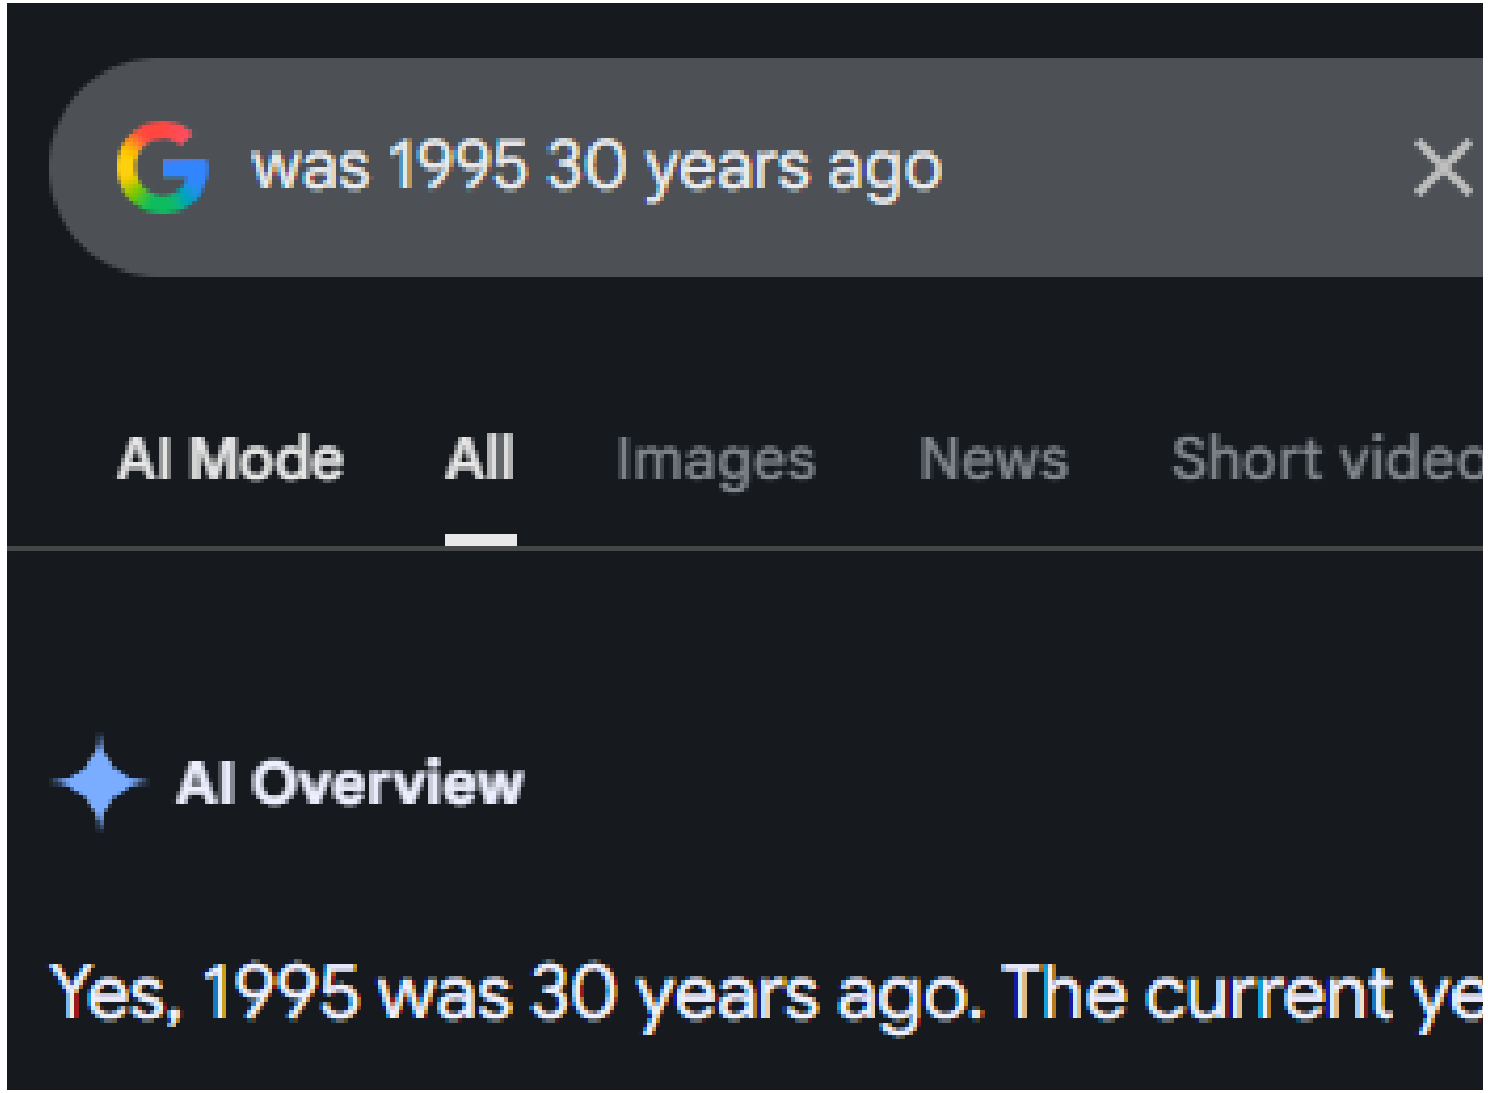

Nor are they permanent

(September 3, 2025)

Using AI for “good fit” tasks increases quality of the output

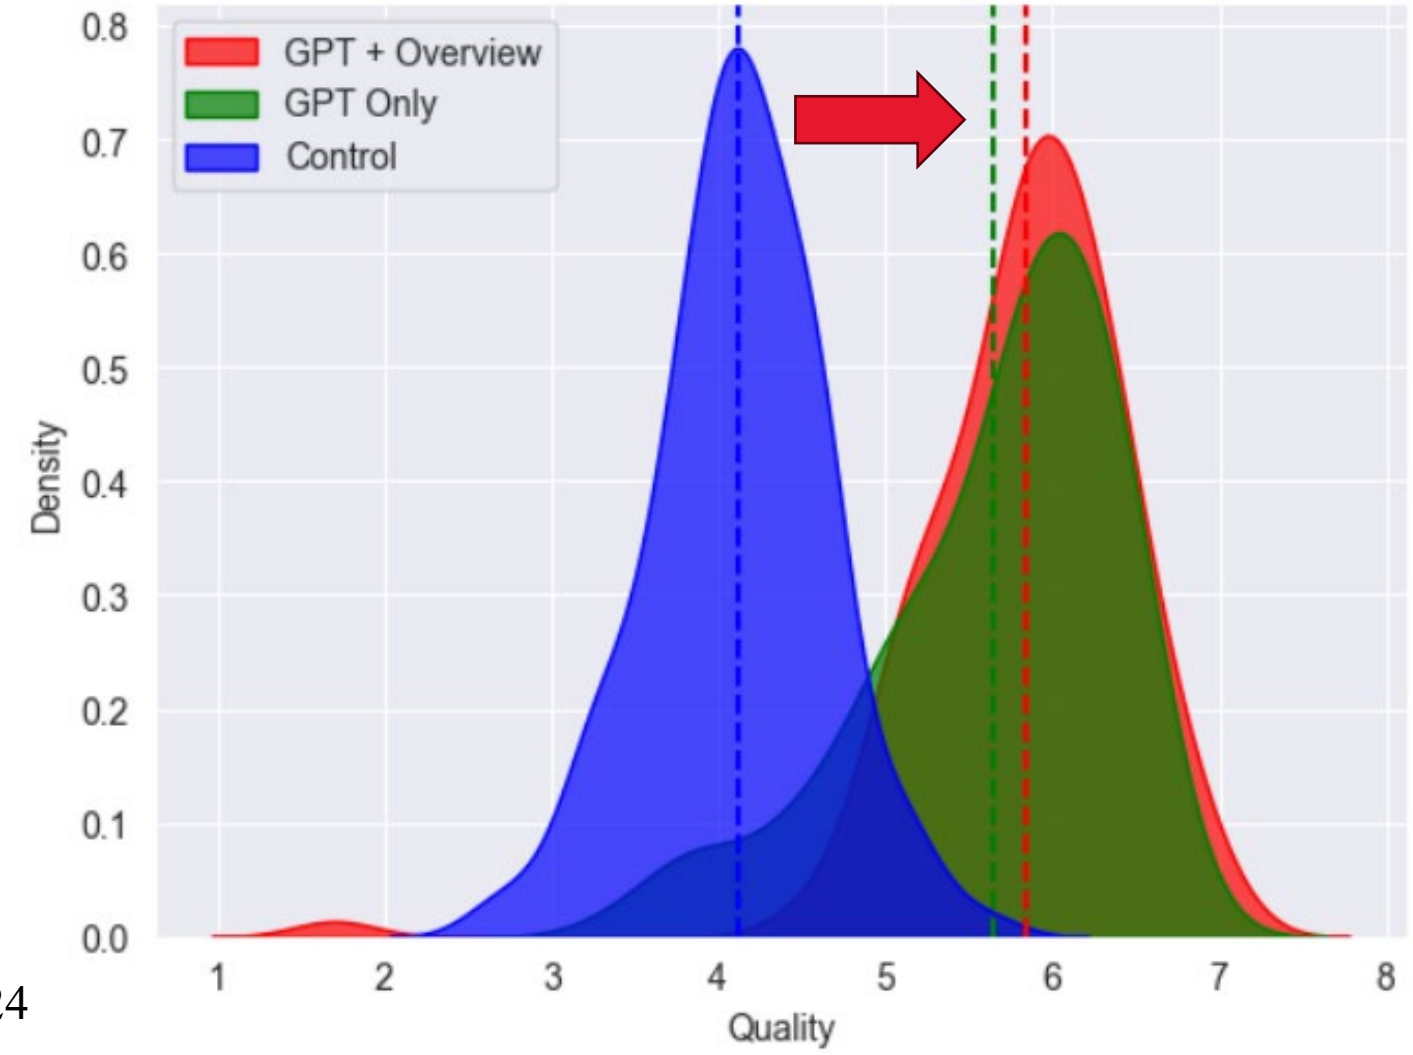

Dell’Acqua et al. 2024

Using AI for “bad fit” tasks decreases quality of the output

Figure 7: **Performance - Outside the Frontier**

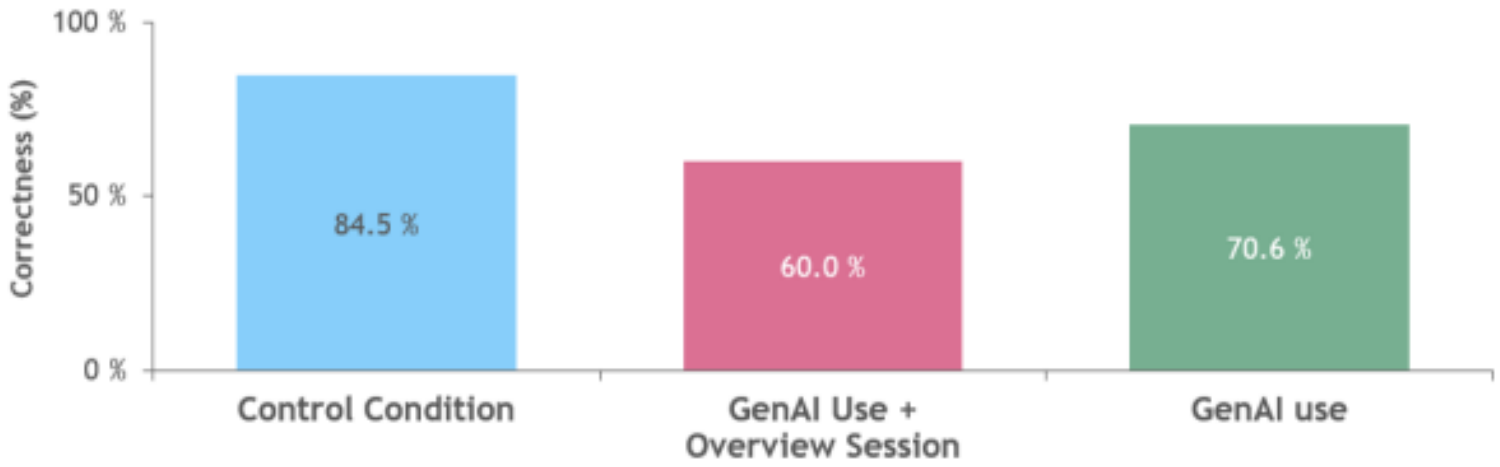

Dell’Acqua et al. 2024

In summary…

The experiments show that the shape and position of the frontier are vital to understanding the impact of AI on work. On tasks within the frontier, AI significantly improved human performance. Outside of it, humans relied too much on the AI and were more likely to make mistakes. Not all users navigated the jagged frontier with equal adeptness. While some completed their task incorrectly, others showcased a remarkable ability to harness the power of AI effectively. We conducted further analyses of the

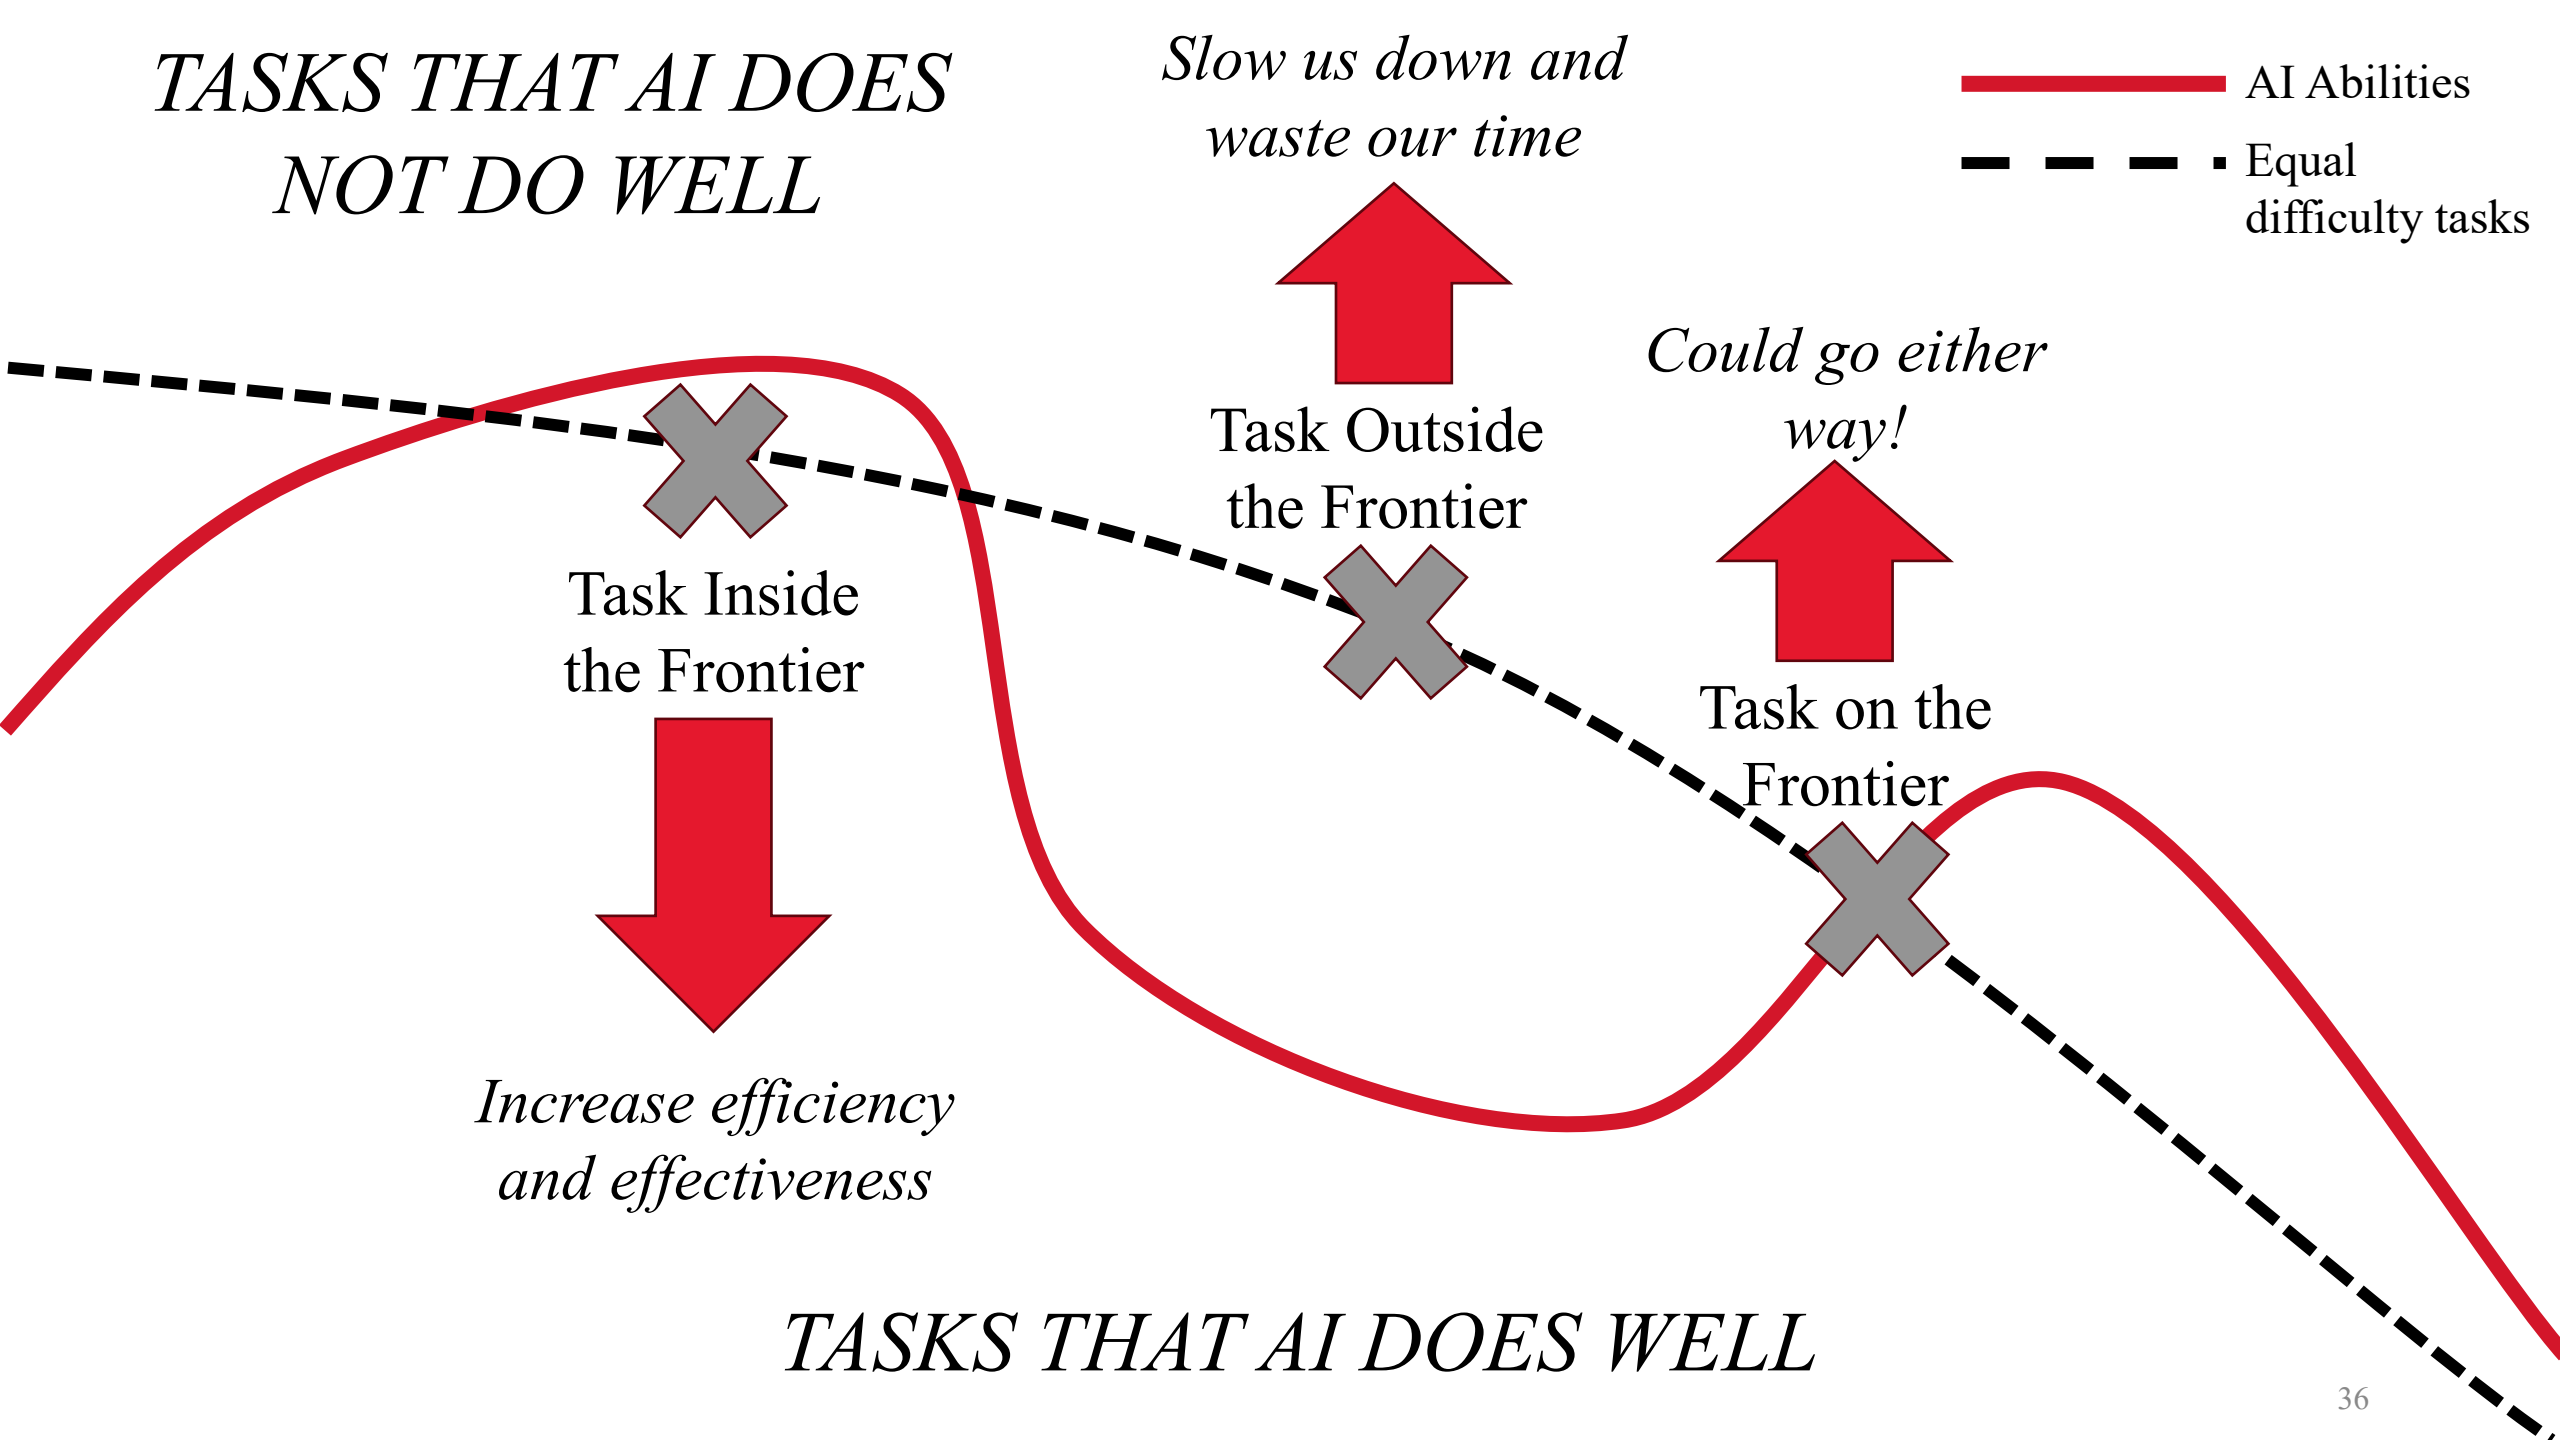
TASKS THAT AI DOES NOT DO WELL
Slow us down and waste our time
AI Abilities
Equal difficulty tasks
Could go either way!
Task Outside the Frontier
Task Inside the Frontier
Task on the Frontier
Increase efficiency and effectiveness
TASKS THAT AI DOES WELL

Let’s find the jagged frontier of AI in education

Course Design and Content

Classroom Prep and Facilitation

Student Assessment and Support

Instructions

1. Think of a task that you do for each of the three areas: **course design and content, class prep and facilitation, student assessment & support.**
2. Prompt an AI tool of your choice (ChatGPT, Copilot, Gemini, Claude, Anthropic, etc.) to help you with that task
3. Evaluate output and iterate
4. Keep notes and share with others

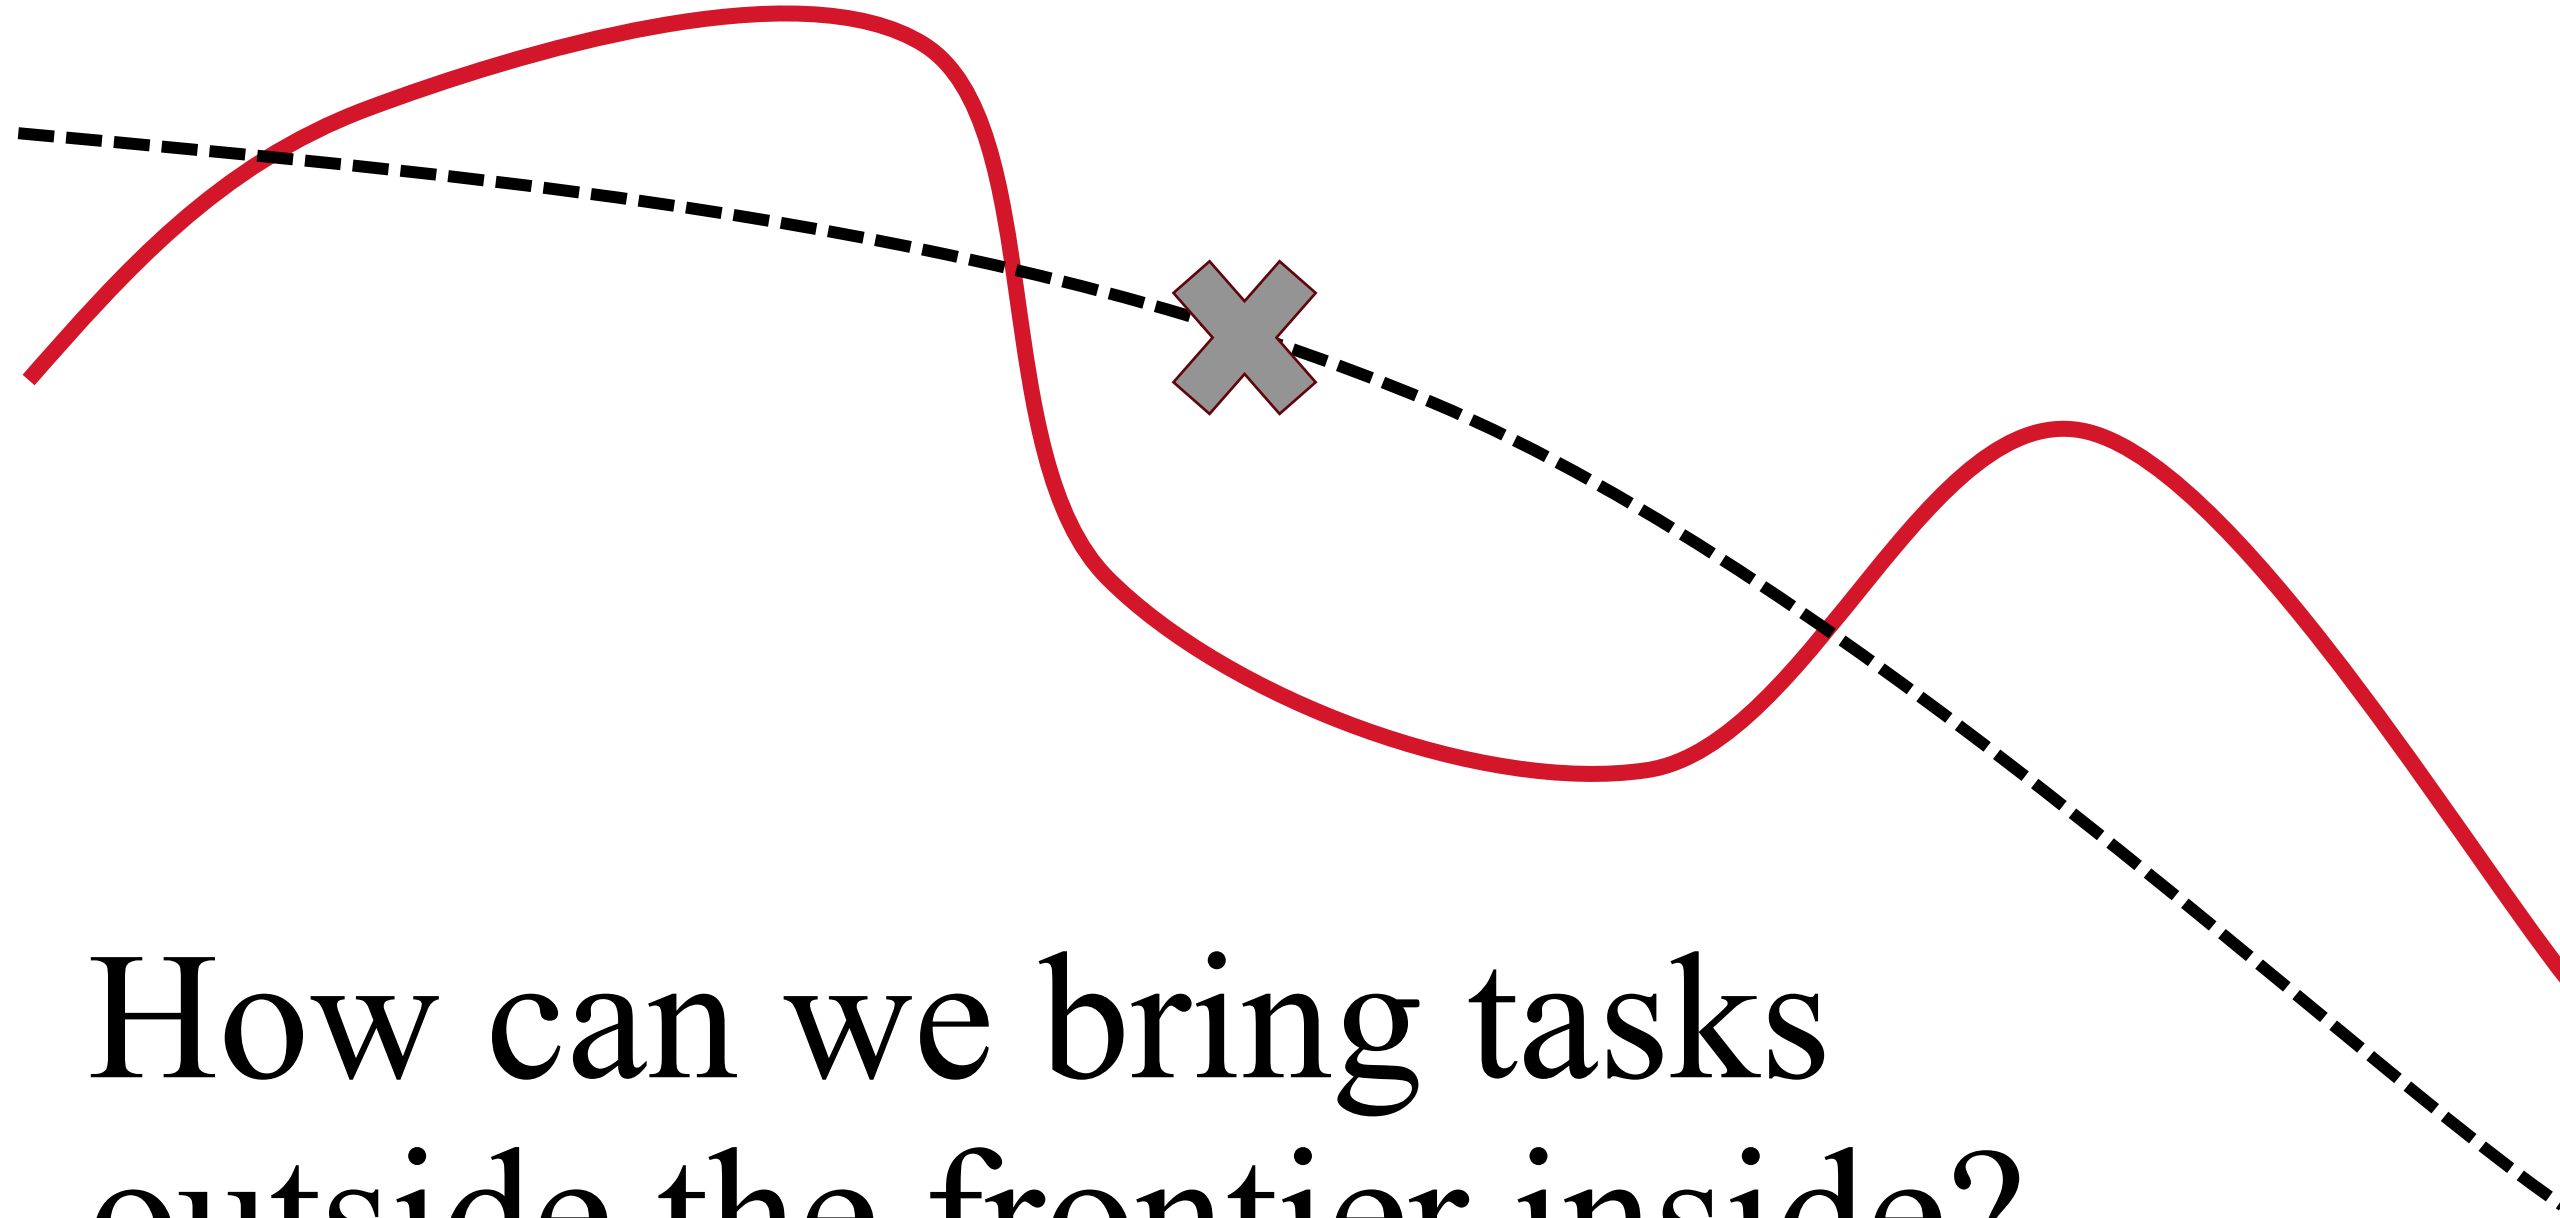

How can we bring tasks outside the frontier inside?

Ways to shift the frontier

- Better prompt engineering
- Stay up to date on functionality and circle back regularly
- Try different models (ChatGPT vs. Gemini vs. Claude) for the same prompt
- Combine multiple tools
- Augment models with more data (RAG)
- Iterate, iterate, iterate
- Human Q&A and testing

just keep trying

You are the worst at AI you will ever be right now. The models and tools you are using are the worst they will ever be right now. Anticipate a learning curve that will delay the positive impacts of AI use on your work.

What's next?

- Provide your feedback on this session
- Future seminars:
  - Architecting Assignments for an AI World, 2/26 & 3/04
  - Leveling up the research process with generative AI, 3/19 & 3/26
  - Who Owns the Output? Fair Use and IP in Generative AI, 3/26 & 4/01

Resources

- Resources specific to each slide are available in the Notes of the slide
- *Co-Intelligence: Living and Working with AI*, Ethan Mollick
- One Useful Thing – Ethan Mollick's blog
- Everyday AI – podcast and newsletter
- *Ethical Machines*, Reid Blackman
- Harvard Business Impact: Teaching with AI (requires creation of a free HBI Educator account)
- CFDI AI Resources Page
- Harvard University Generative AI Resources

Backup Slides

It’s not magic or human thinking – it’s math!

… and why sometimes it “hallucinates” – if there is not enough data to suggest an appropriately high probability, it will *generate* a fictional response that *looks* like an appropriate response.


The New York Times

A.I. Is Getting More Powerful, but Its Hallucinations Are Getting Worse

A new wave of “reasoning” systems from companies like OpenAI is producing incorrect information more often. Even the companies don’t know why.

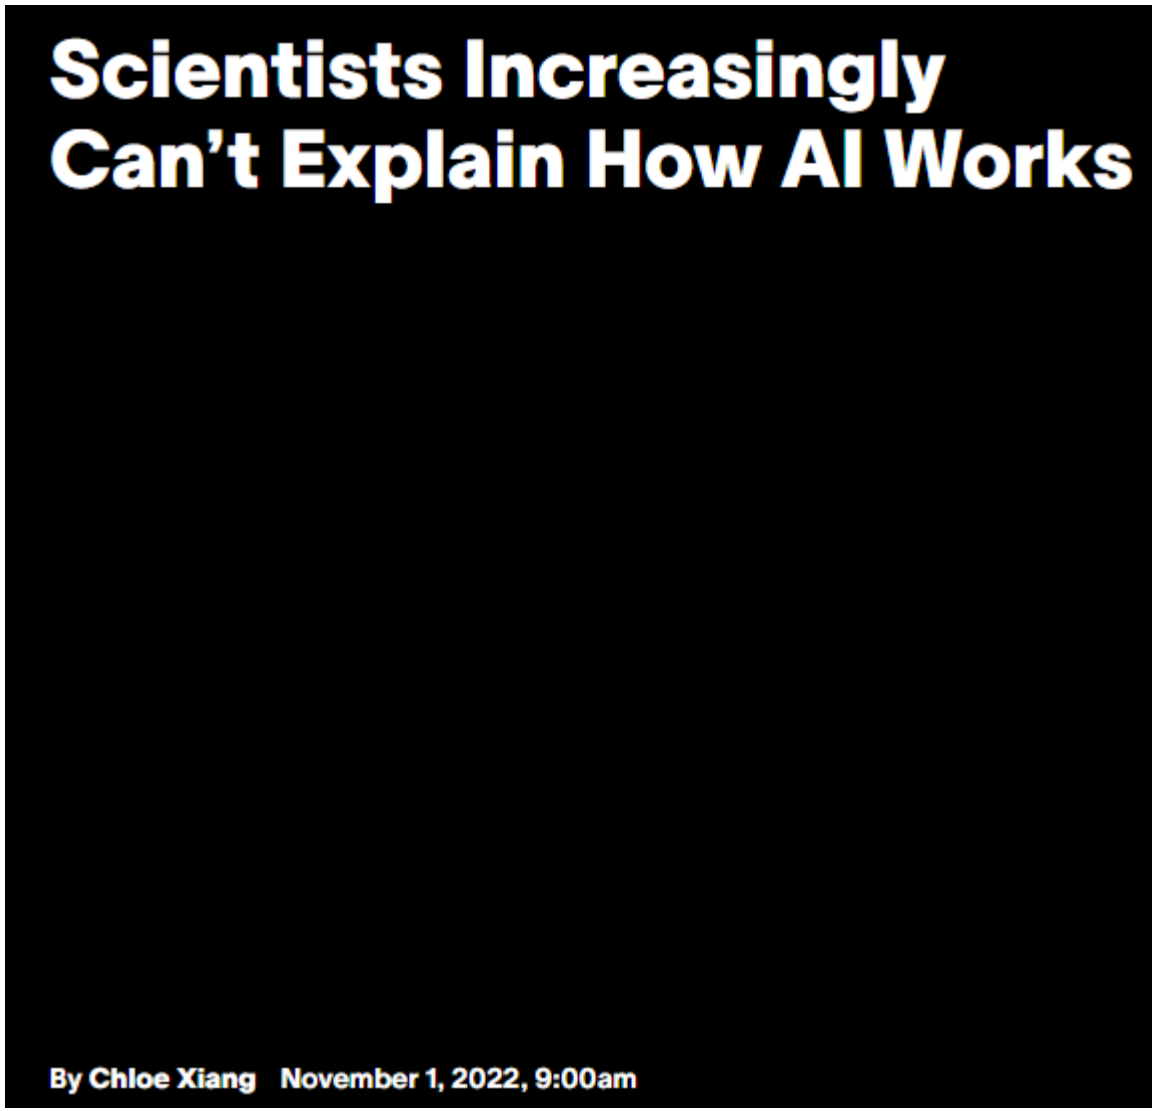

The specific engineering of ChatGPT has made it quite compelling. But ultimately (at least until it can use outside tools) ChatGPT is "merely" pulling out some "coherent thread of text" from the "statistics of conventional wisdom" that it's accumulated. But it's amazing how human-like the results are. And as I've discussed, this suggests something that's at least scientifically very important: that human language (and the patterns of thinking behind it) are somehow simpler and more "law like" in their structure than we thought. ChatGPT has implicitly discovered it. But we can potentially explicitly expose it, with semantic grammar, computational language, etc.

What ChatGPT does in generating text is very impressive—and the results are usually very much like what we humans would produce. So does this mean ChatGPT is working like a brain? Its underlying artificial-neura
on an idealizati
that when we h
of what's going on are quite similar.

Stephen Wolfram, 2023

Better prompting

Ask	Ask it to do (better, newer, more specific)
Explain	Ask it to explain itself (it won't – but could still be illuminating)
Invite	Invite it to take a specific perspective or persona
Organize	Ask it to organize its response
Decompose	Break your question into parts
Reframe	Reframe or rephrase the question
Constraints	Give constraints
Examples	Give examples of good answers, guidelines for success, rubrics
Provide	Provide data and ask it to hypothesize
Apply	Ask it to apply methods of frameworks you use

AI is replacing junior (but not senior) employees

Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data*

Seyed M. Hosseini† Guy Lichtinger+

Preliminary

August 2025

Abstract

We study whether generative artificial intelligence (AI) constitute *niority-biased technological change*, disproportionately affecting junior re workers. Using U.S. résumé and job posting data covering nearly 62 ers in 285,000 firms (2015–2025), we track within-firm employment d niority. We identify AI adoption through a text-analysis approach ings for dedicated "AI integrator" roles, signaling active implementa tive AI. Difference-in-differences and triple-difference estimates show in 2023Q1, junior employment in adopting firms declined sharply r adopters, while senior employment continued to rise. The junior d primarily by slower hiring rather than increased separations, with th

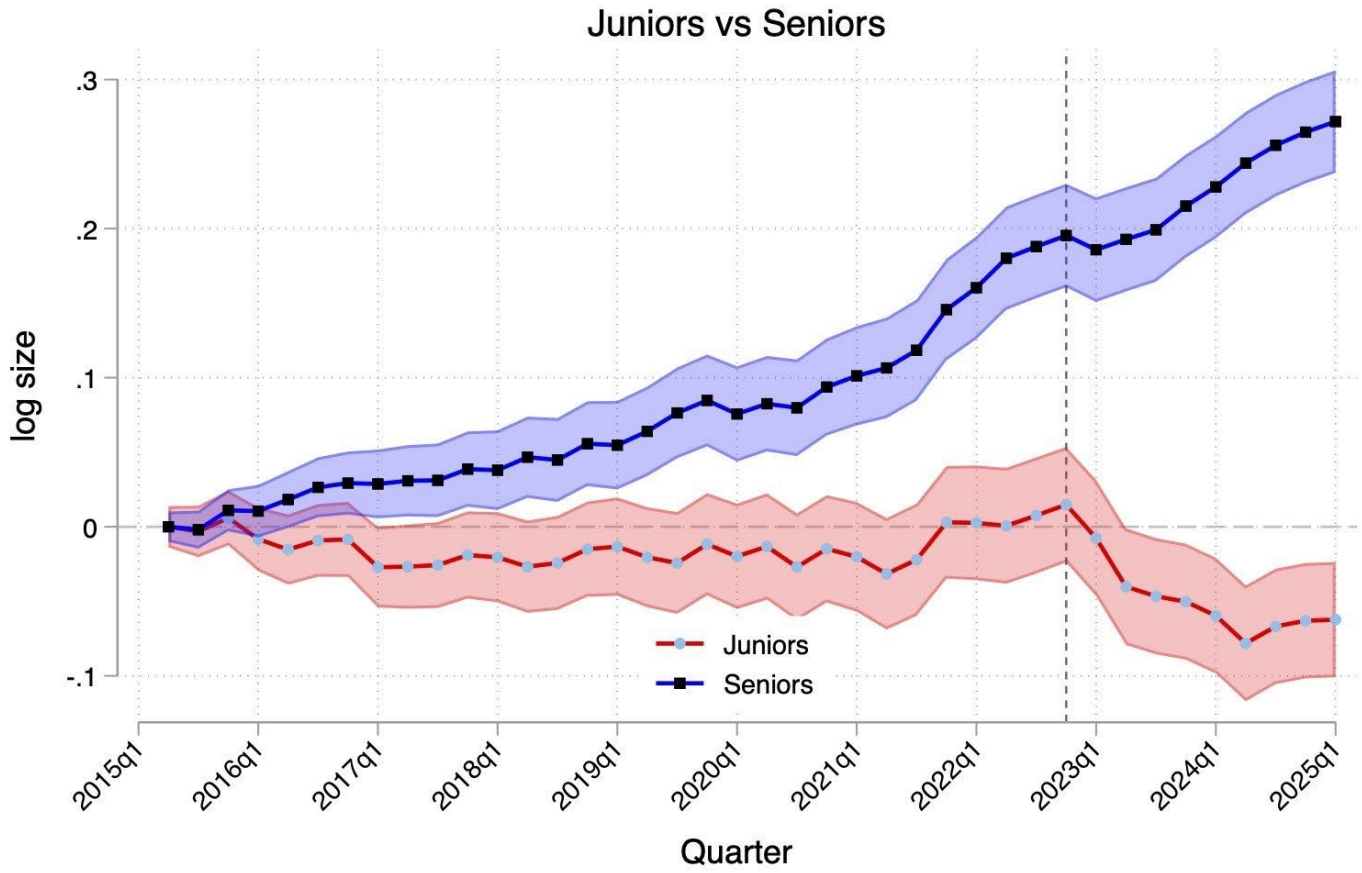

A final note on prompts…

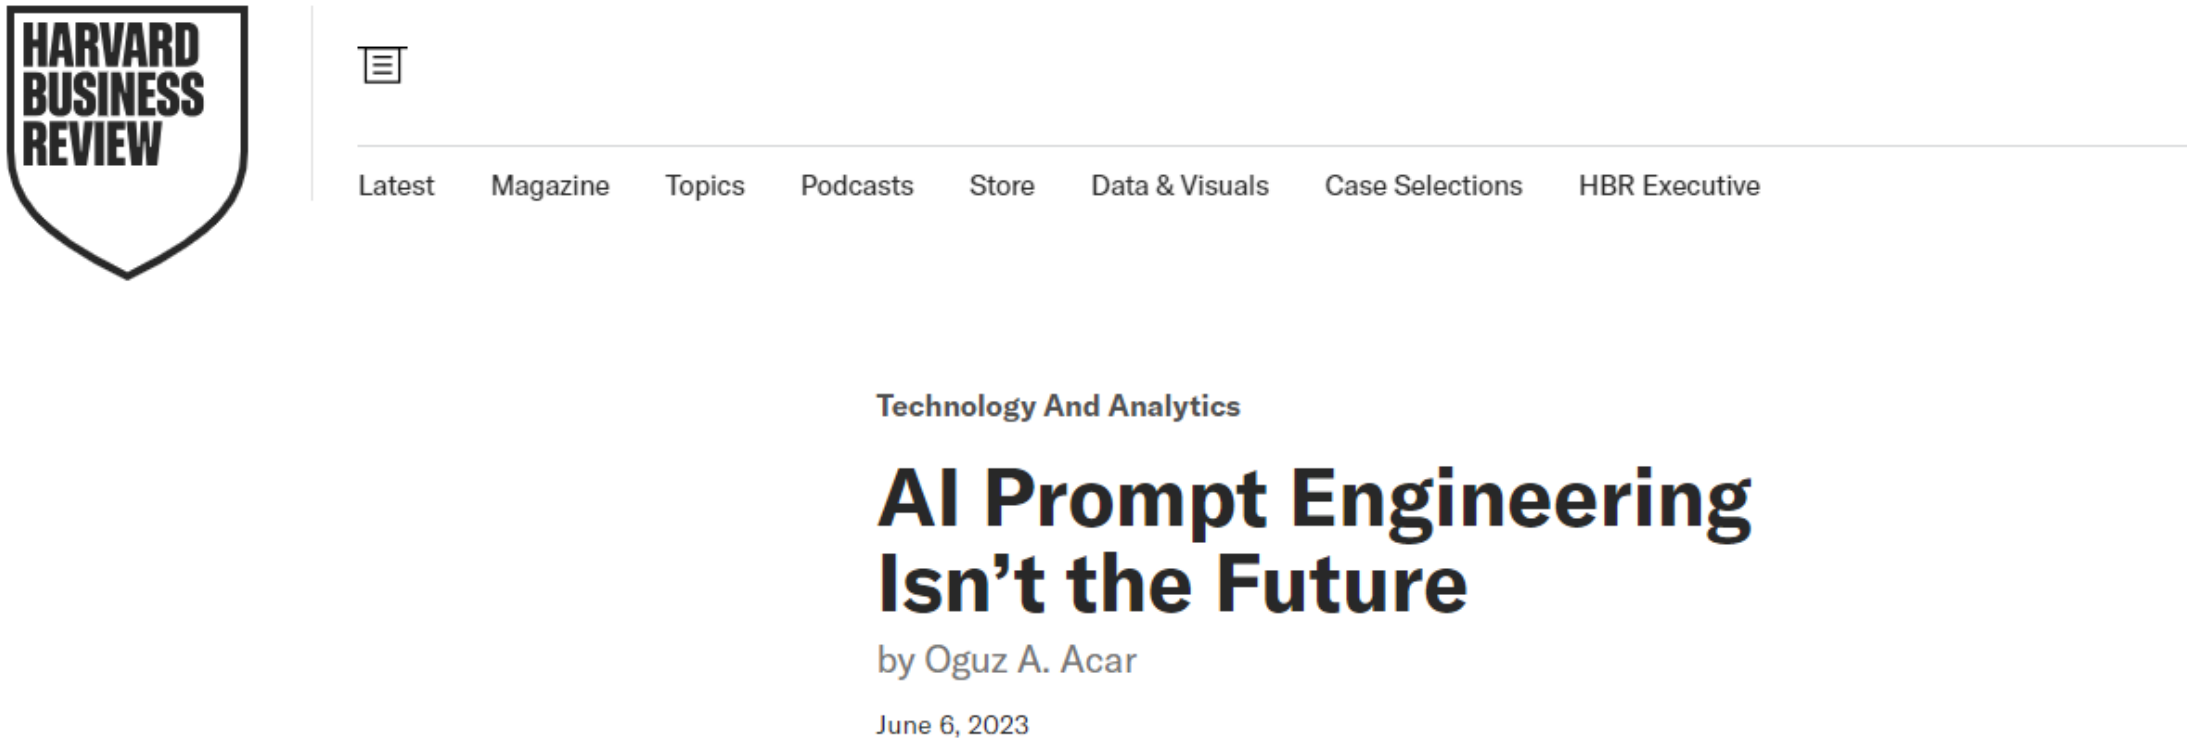

Environmental Impact of AI

What Uses More?

Free to share or adapt

(see **About this app** for credits)

Compare the environmental footprint of digital tasks

▼ **About this app**

It's hard to know what to believe about the ecological impact of generative AI when the Internet seems evenly divided between "Stop worrying" and "It's the apocalypse." This nonpartisan calculator is an attempt to provide a more objective comparison of the energy and water usage of various AI tasks versus other common digital activities, like watching Netflix or searching Google. The calculations are based on academic papers, industry reports, and expert estimates.

Paragraph of AI text x 1: 3 lightbulb-minutes (Wh) [0.02 liters]

versus

Storing 5 GB in the cloud for 1 month x 1: 8 lightbulb-minutes (Wh) [0.06 liters]

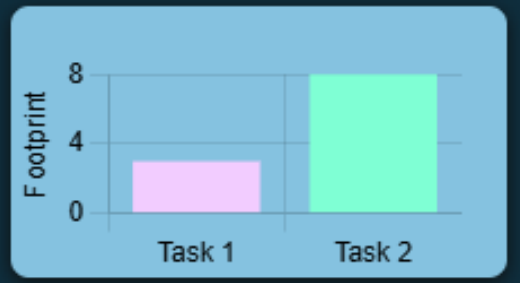

Environmental Impact of Blue Books

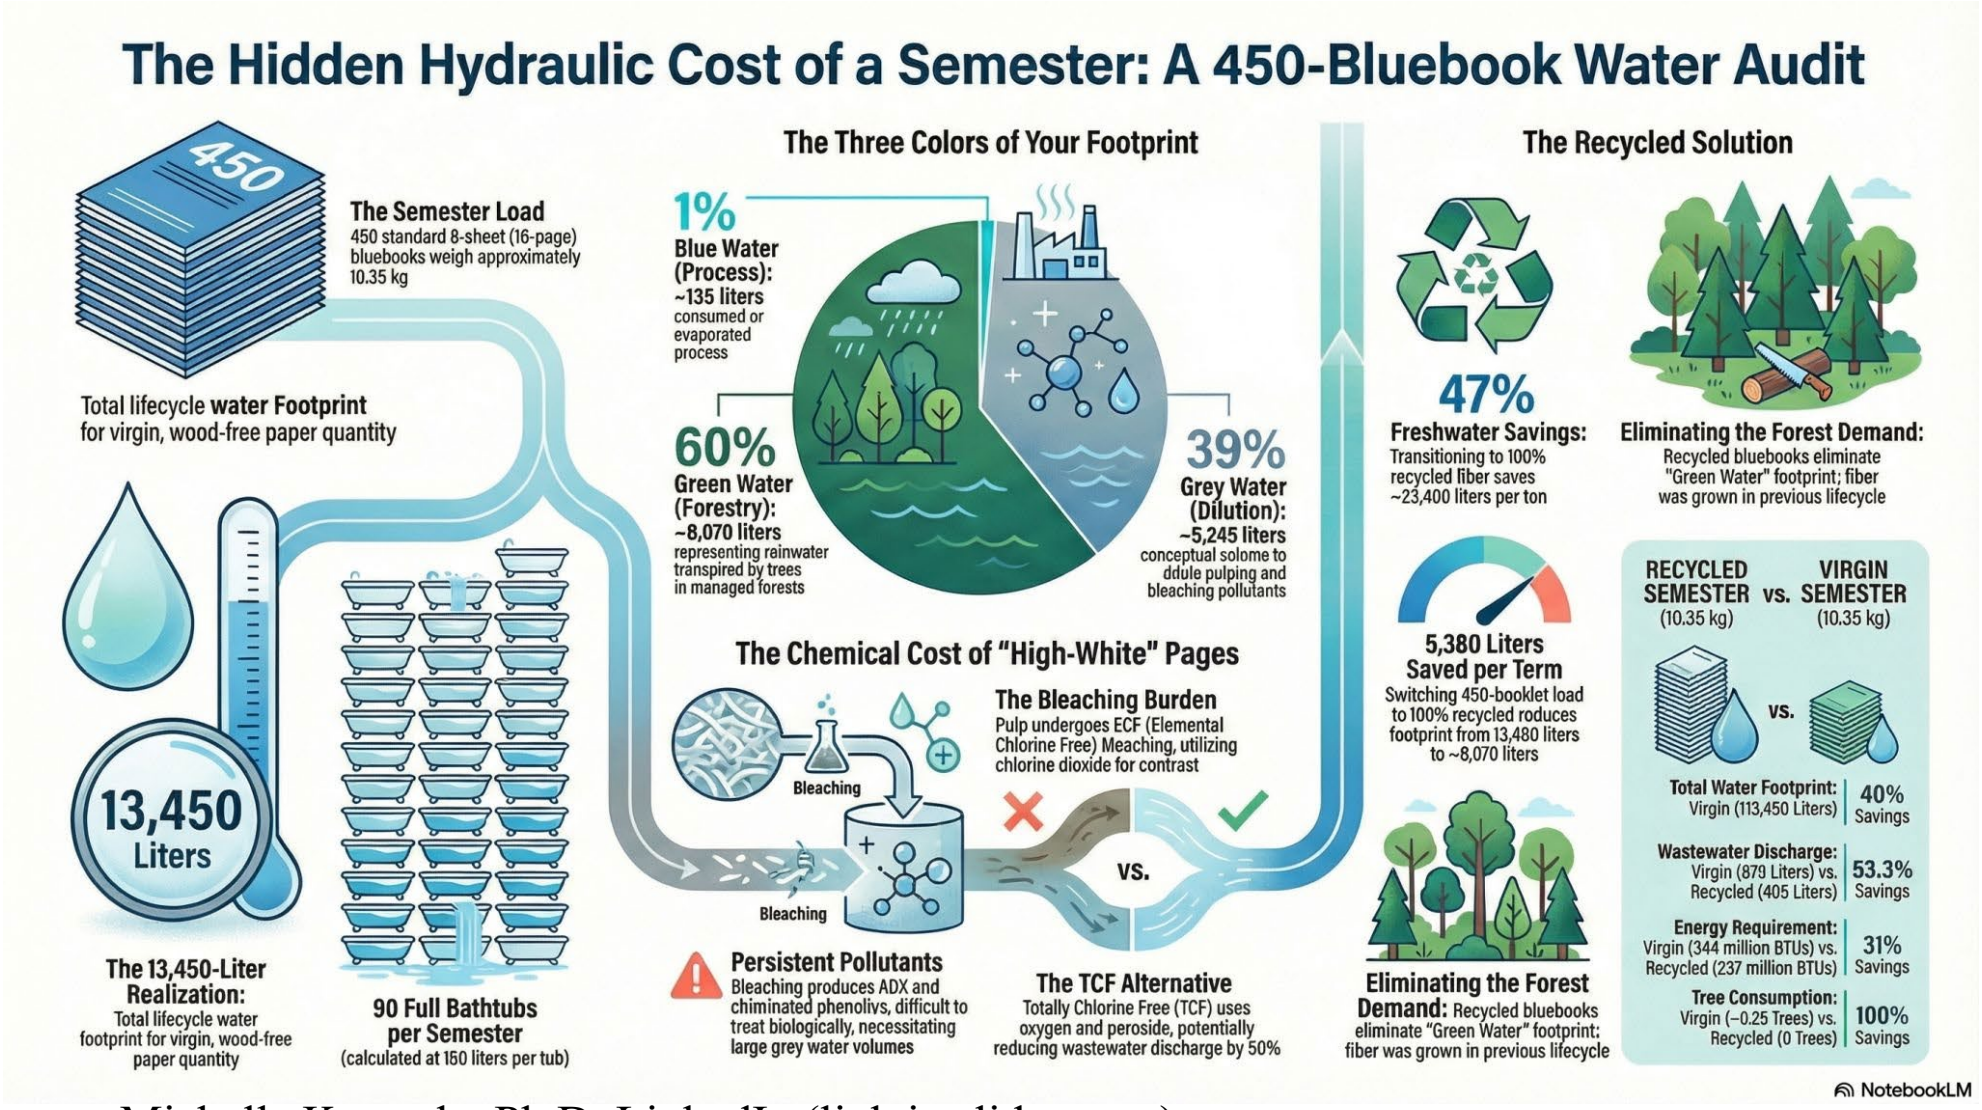

Source: Michelle Kassorla, Ph.D. LinkedIn (link in slide notes)