

AI and Research

(...and society)

(.....and your backyard)

Brady M. Chisholm & Juraj Mesik

GAMMs Club 5/20/2026



MINNESOTATM



UNIVERSITY OF MINNESOTA

Driven to Discover[®]

Some Disclaimers

- We're not experts in AI
- Have interest in AI, both professionally and personally
- May have blind spots to things AI “pros” would mention
- Purpose is to raise awareness about various aspects of AI. Do not mean to prescribe certain tools or AI use as right or wrong
- Meeting inspired by Matthias Nau, Assistant Professor of Computational Cognitive Neuroscience at the Vrije Universiteit Amsterdam



Today's Agenda

- Survey AI use and attitudes from the group
- Brief summary of AI from 2022 to now
- Implications of AI today
- Modern AI in Research
 - Example use cases of AI in research
 - Questions for the group about AI

My Goals

1. Raise awareness about AI
2. Get people thinking about AI at their own institutions
(tech offerings, policies, training we (researchers) need)
3. Mention some of the larger implications of AI
4. Discuss how we do and don't use AI for research



Why Discuss AI?

- AI is developing rapidly
 - OpenAI’s flagship GPT 3.5 model released Nov. 2022 (approx. 1,275 days ago)
- Seemingly AI is here to stay; knowledge, familiarity, and rate of adoption vary widely across the world
- University of Minnesota
 - Psychology Department Committees on AI
 - AI Hub Announced March 4th, 2026
 - Feb. 2026 delegation from MN (including UMN AI Hub reps) visits Nvidia HQ in CA
- Other BIG10 institutions already “ahead”? Michigan Institute for Data and AI in Society (MIDAS) was founded in 2015



Questions for the Audience

- Within the last month, who has interacted with...
 - An LLM? (Gemini, ChatGPT, Anthropic, etc.)
 - A CoPilot or CLI coding tool? (VS Code Line Assist, GitHub CoPilot)
 - An Agentic AI tool/harness? (Antigravity, OpenClaw, Codex, etc.)

Any experienced AI users? Anyone *never* used AI?



Chat GPT (2022, early)

- Novel and impressive, fun to explore
 - e.g. “How many S’s in STRAWBERRY?”
“Summarize the key points of macroeconomics”
- Not too practically useful
- Seemed highly prone to pitfalls
 - Hallucination
 - Logical contradiction

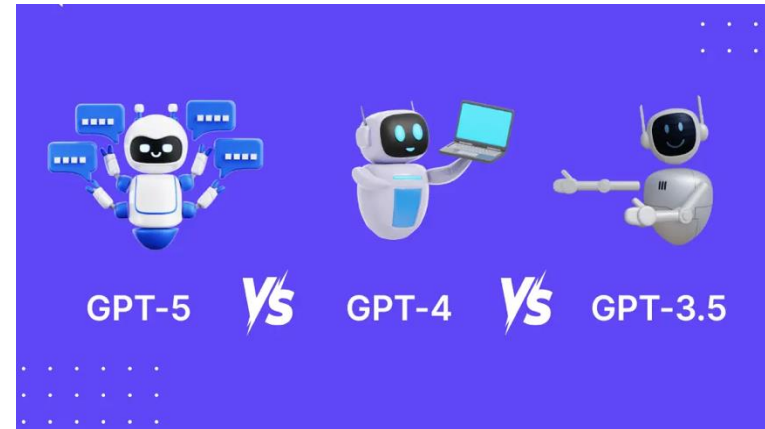



IMAGE: <https://www.kommunicate.io/blog/chatgpt-4-vs-chatgpt-3-5-key-differences/>

What's a GPT, Anyway?

- Generative Pre-Trained Transformer
 - Trains on an existing corpus of data; limited knowledge window
 - Prone to biases; myriad of other issues

GPT training typically involves the following steps:

1. **Data preparation:** The first step is to gather and prepare a massive dataset of text and code. This dataset is carefully curated to be as diverse and representative as possible, covering a wide range of topics, writing styles, and languages.
2. **Tokenization:** The text data is then divided into smaller units called "tokens." These can be individual words, parts of words, or even characters, depending on the specific GPT model and the desired level of granularity.
3. **Model initialization:** The GPT model is initialized with random parameters. These parameters will be adjusted during the training process as the model learns from the data.
4. **Self-supervised learning:** The model is then fed the tokenized text data and tasked with predicting the next token in a sequence. For example, given the input "The cat sat on the", the model might predict "mat."
5. **Backpropagation and optimization:** The model's predictions are compared to the actual next tokens in the training data, and the difference between them is used to calculate a "loss" value. This loss represents how far off the model's predictions are from the truth. The model then uses backpropagation to adjust its internal parameters to minimize this loss. This iterative process of prediction, loss calculation, and parameter adjustment continues over many epochs, with the model gradually improving its ability to predict the next token in a sequence accurately.

Audio from 4/30/26
META meeting

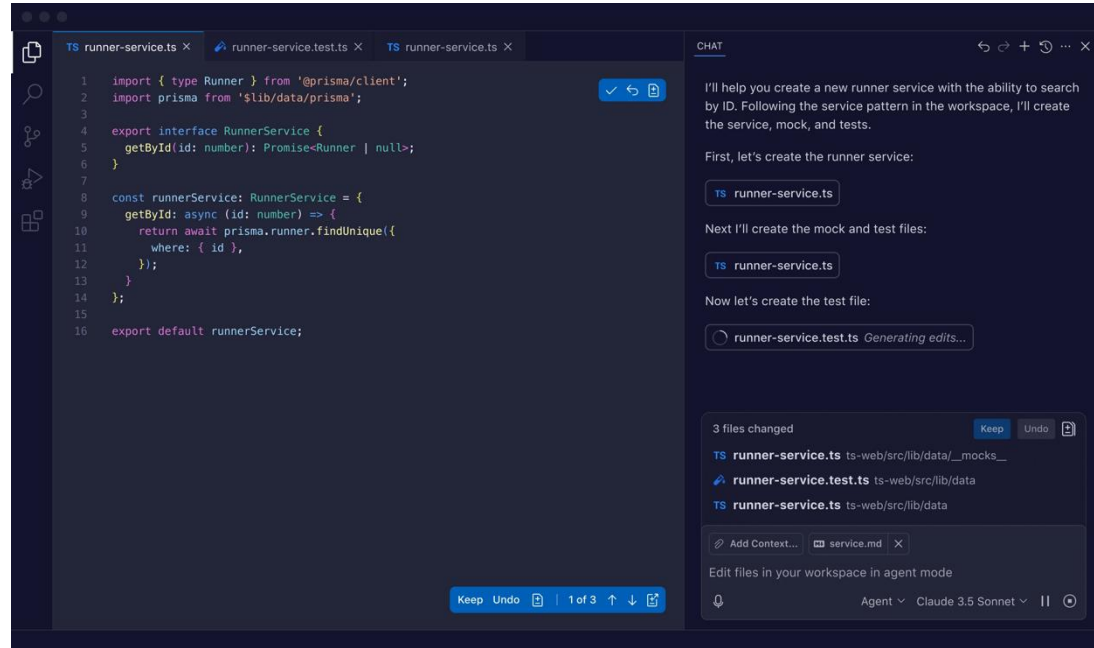
[Zuckerberg] “The AI models learn from **watching really smart people do things...**the average intelligence of the people at this company [META] is significantly higher than the average set of people that you can get to do tasks.

...if we're trying to teach the models coding, for example, then **having people internally build tools or solve tasks that help teach the model how to code, we think is going to dramatically increase our model's coding ability faster than what others in the industry have the capability to do**, who don't have thousands and thousands of extremely strong engineers at their company.”



CoPilot and IDE Assistants (reasoning)

- These tools help accelerate coding / debugging by parsing your code and helping predict what you want to write. Almost like an LLM wrapper (removes resistance from chatbot-style tools)
- Visual Studio Code CoPilot integration throughout 2025
- Most IDEs offer some type of CLI assistant now...

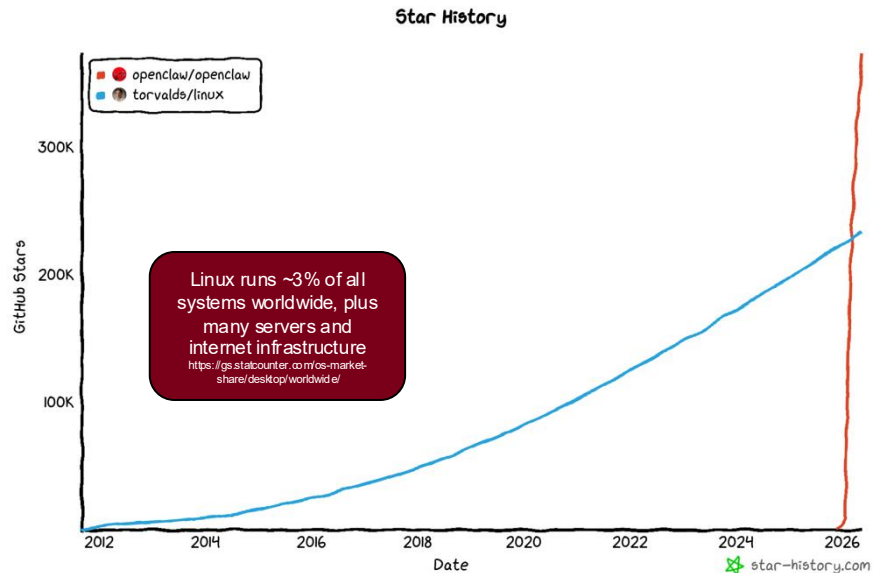


<https://github.com/features/copilot>



Agentic AI (latest tech)

- Agentic AI takes CLI-style tools further by not only generating ideas, but also acting upon the system independently via a variety of tools.
- Different workflow from “chatbot” LLMs
 - Copy-pasting code from a Gemini window vs asking agents to develop and test code for you
- Have become popular in the last ~6 months and rapidly gaining momentum
 - Predominantly in SWE world
 - Unclear how adoption looks like in academia / science
- Autonomous, “agentic” control over your system
 - e.g., writing code in an IDE, running it, interpreting error messages, searching the web for packages to solve issues, downloads, etc
 - Don’t worry, you set boundaries 😊
 - e.g., “don’t read from /lab/experiment3/data/anatomical_fMRI/subjP01/...”
- Iterates for you
 - Can ask for you to review outlines or plans



HERMES-AGENT

<https://github.com/nousresearch/hermes-agent>



Codex
(OpenAI)



Antigravity
(Google)

OPENCLAW

<https://github.com/openclaw/openclaw>



AI Today

Pros

- Can save time
 - Quickly review or summarize papers
 - Parse large datasets
- Accessibility (empowering the user with skills they don't otherwise have)
- Ability for automation, task offloading, personal “agentic” assistants
- Review code and help debug

Cons

- Malicious actors using AI (security concerns)
 - Prompt injections
 - Hacks, data breaches
- Biases (intended or otherwise)
- Skill atrophy, cognitive offloading
- Context window decay
- Intellectual property, patent, and copyright concerns




AI Today (cont.)

- Articles discussing AI

- Gutiérrez, J. B. (2025). *The Flaws of Others: An LLM-driven Framework for Scientific Knowledge Production* (arXiv:2507.06565; Version 1). arXiv. <https://doi.org/10.48550/arXiv.2507.06565>
- Liu, G., Christian, B., Dumbalska, T., Bakker, M. A., & Dubey, R. (2026). *AI Assistance Reduces Persistence and Hurts Independent Performance* (arXiv:2604.04721). arXiv. <https://doi.org/10.48550/arXiv.2604.04721>
- Shen, J. H., & Tamkin, A. (2026). *How AI Impacts Skill Formation* (arXiv:2601.20245). arXiv. <https://doi.org/10.48550/arXiv.2601.20245>
- Dolan, Eric W. (2026) ChatGPT acts as a “cognitive crutch” that weakens memory, new research suggests (<https://www.psypost.org/chatgpt-acts-as-a-cognitive-crutch-that-weakens-memory-new-research-suggests/>)
- Buyl, M., Rogiers, A., Noels, S., Bied, G., Dominguez-Catena, I., Heiter, E., Johary, I., Mara, A.-C., Romero, R., Lijffijt, J., & De Bie, T. (2026). *Large language models reflect the ideology of their creators*. *Npj Artificial Intelligence*, 2(1), 7. <https://doi.org/10.1038/s44387-025-00048-0>
- Bai, H., Voelkel, J. G., Muldowney, S., Eichstaedt, J. C., & Willer, R. (2025). *LLM-generated messages can persuade humans on policy issues*. *Nature Communications*, 16(1), 6037. <https://doi.org/10.1038/s41467-025-61345-5>
- Waight, H., Yang, E., Yuan, Y., Messing, S., Roberts, M. E., Stewart, B. M., & Tucker, J. A. (2026). *State media control influences large language models*. *Nature*. <https://doi.org/10.1038/s41586-026-10506-7>


Invalid Outputs



Some (simple) Examples Using Gemini


Translating Code

Taking a function written in MATLAB that locates data files for an inputted list of subject IDs over to R. Provided context about file directory structure etc. to prevent errors



Creating a Sound Spectrum

Generated a vector graphic (i.e., scalable, non-pixelated image) of sound spectrum depicting different harmonic complex tones.



Other Considerations with AI

Hot take: We need to replace Central Park with the world's largest data center

 **Nick Lindquist** 
@nick_lindquist

Follow  

Central Park is great, but it takes up a lot of space and isn't utilized to its full potential.

That's why I worked with McKinsey on a plan to make it a state of the art data center, complemented by rooftop parking and nuclear power.

We can still build beautiful things.

Theoretical capability and observed usage by occupational category

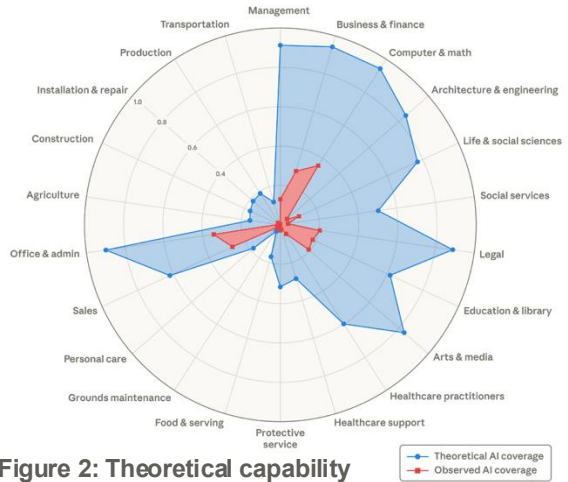


Figure 2: Theoretical capability and observed exposure by occupational category
Share of job tasks that LLMs could theoretically perform (blue area) and our own job coverage measure derived from usage data (red area).

Source: Anthropic
(<https://www.anthropic.com/research/lab-or-market-impacts>.)

(major conflict of interest...)



...aerial view, an Amazon Web Services data center is shown situated near single-family homes on July 17, 2024 in Stone Ridge, Virginia.
<https://www.mprnews.org/episode/2025/05/01/minnesotas-natural-resources-fuel-development-of-large-data-centers>
May 1, 2026, MPR News

Only 2% backed Kevin O'Leary's \$100 billion Utah data center — now residents want a referendum to kill it

As public outcry in Utah's Box Elder County continues in the fight against Kevin O'Leary's \$100 billion data center proposal,...

2 days ago
May 18, 2026

40,000 acres = 62sq. miles = ~24,000 soccer (or) ~33,000 football fields



Some Questions for the Group...

- Other examples of AI usage the group has?
- Who sets the boundaries for AI use in research?
Individual lab PIs? The Office of Research Integrity? The organization/agency which distributes your funding? The APA?
- If we disclose AI use in research, how?
- Is it better for an AI or an undergrad to do repeatable tasks (experiment scoring, recruitment emails, etc.)?
- What do we do about primarily relying on models released by “Big Tech”?
 - What biases? e.g. checking for primarily male authors cited? Only American publications?
- Institution policies on Agentic AI? Any existing policies at all?
- Is GAMMs Club interested in further in-depth discussions about AI? (Tutorials, tools, examples)



Thanks for listening!

Have a great summer GAMMs Club!



Lloyd

Poppy

