Why AI is Ready for Biomedicine

Why AI is Not Ready for Biomedicine

Bradley Malin, Ph.D.
Accenture Professor in Biomedical Informatics, Biostatistics, CS, and ECE
Co-Director, ADVANCE AI Center, Vanderbilt University Medical Center
June 30, 2026

Disclosures

- I receive research funding from Intuit Inc.
- Research funding from NIH, NSF, PCORI, ARPA-H
- Through Privasense, LLC, I consult for numerous companies, including Amazon, CVS, IQVIA, Merative, Oracle, and Walgreens Boots

AI as a Scientific Co-Pilot

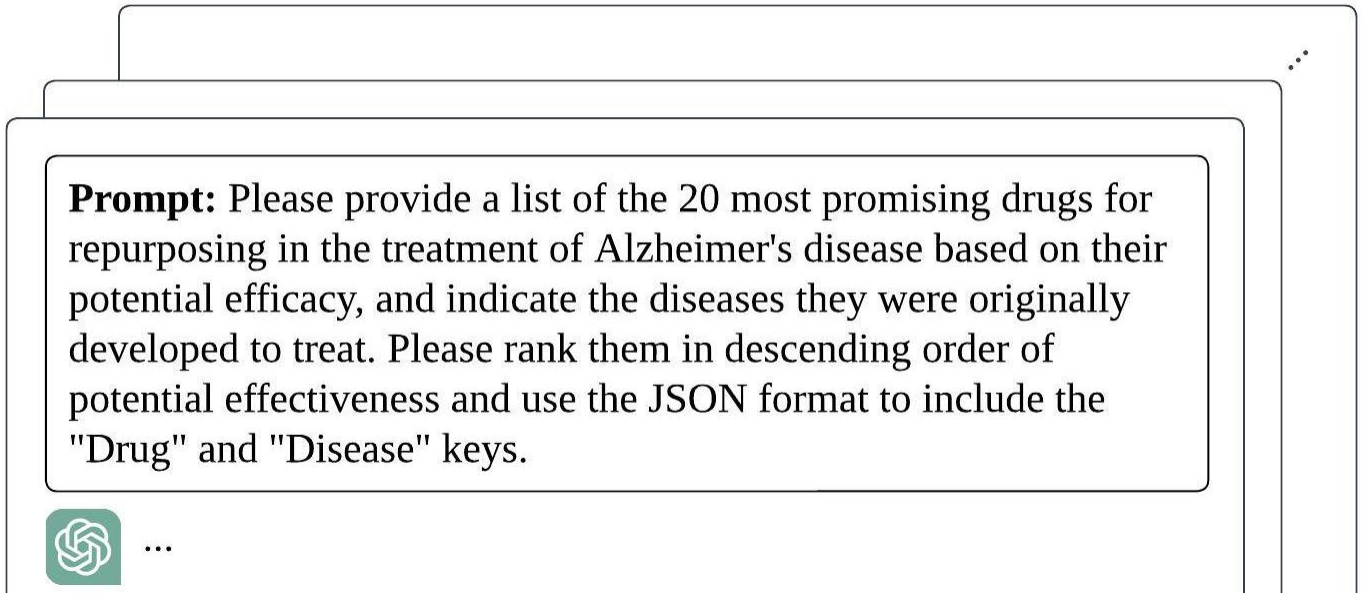

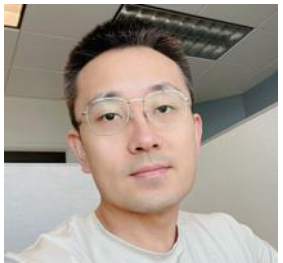

Chao Yan, et al. NPJ Digital Medicine. 2024.

Top 10 drug repurposing candidates in frequency

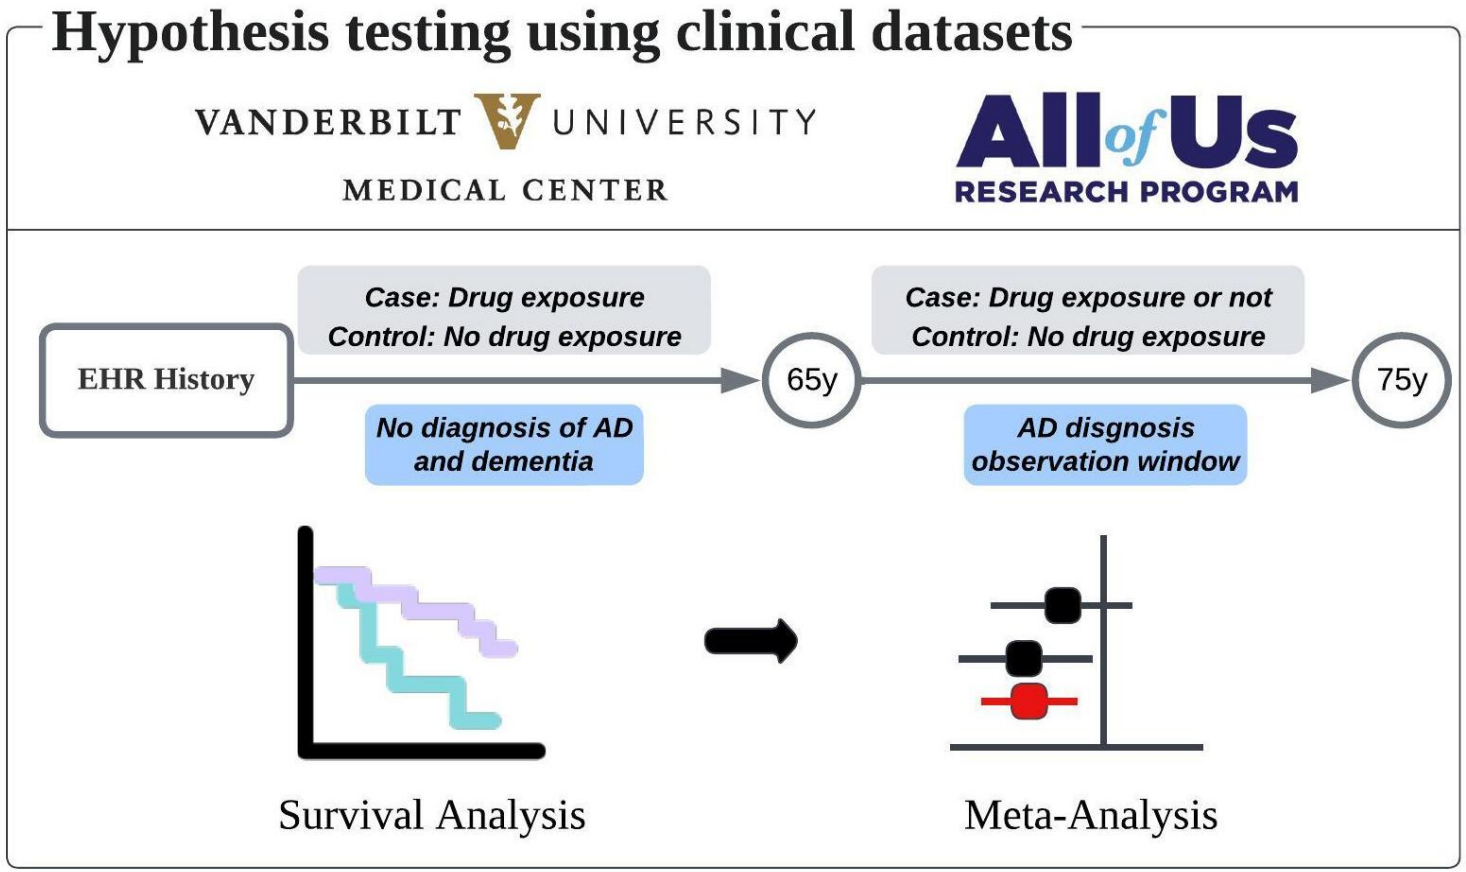

Yan C, et al. NPJ Digital Medicine. 2024.

VUMC

All of Us Research Program

Meta-analysis

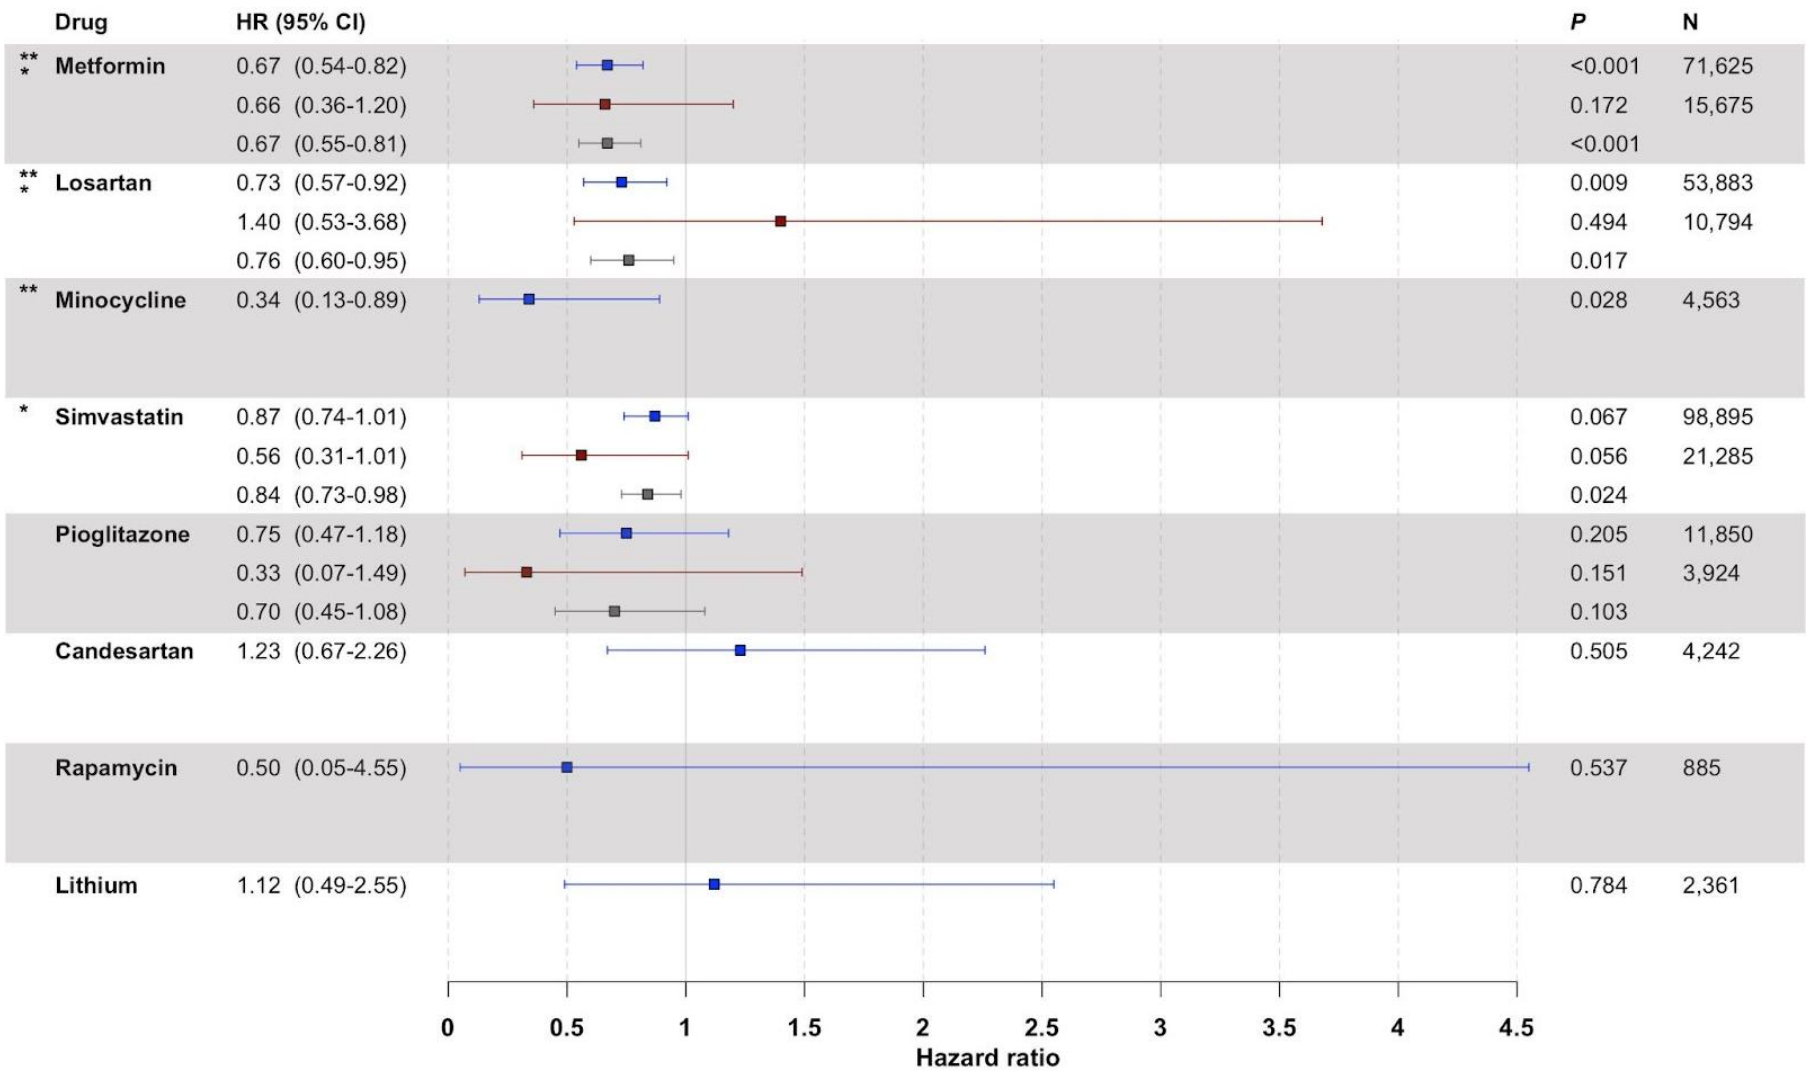

Yan C, et al. NPJ Digital Medicine. 2024.

AI as a Clinical Diagnostic

- GPT vs. Locally Trained Models (Boosting)
  - MIMIC ICU Transfer Prediction
  - VUMC Discharge Prediction
- Vanilla GPT vs. Retrieval Augmented Generation (RAG) GPT
  - *In other words*: we gave various examples to assist the LLM

Katherine Brown, et al. JAMIA. 2025.

Local Model

GPT 4.0

GPT 3.5

Vanilla GPT - VUMC	RAG GPT - VUMC
Vanilla GPT - MIMIC	RAG GPT - MIMIC

March 8, 2025

Katherine Brown, et al. JAMIA. 2025.

GPT 5.0

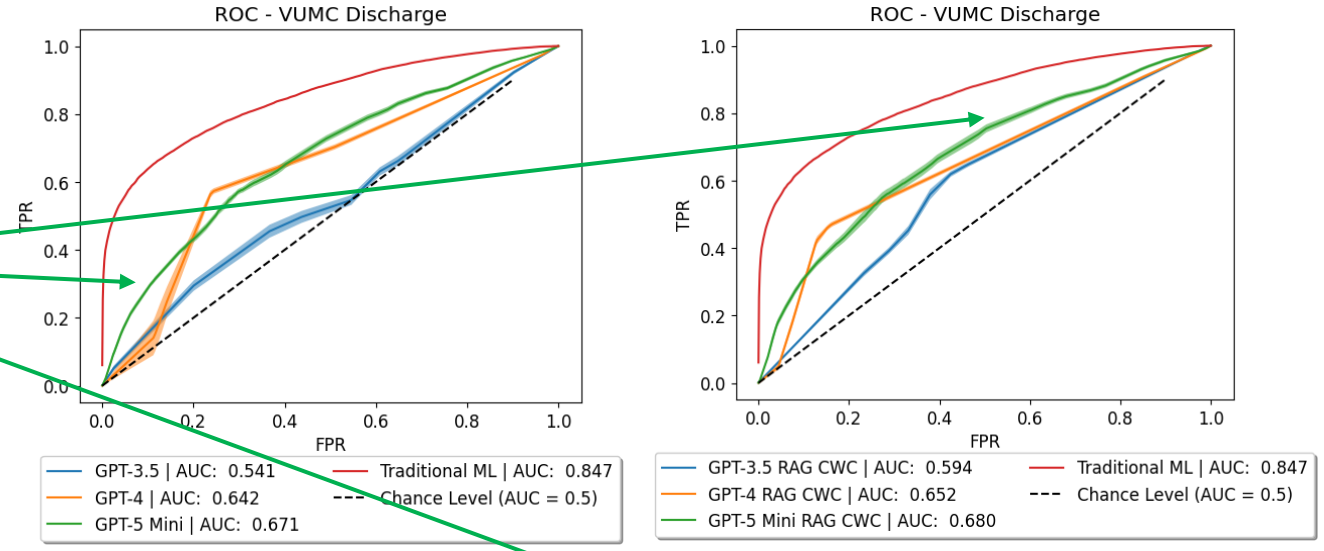

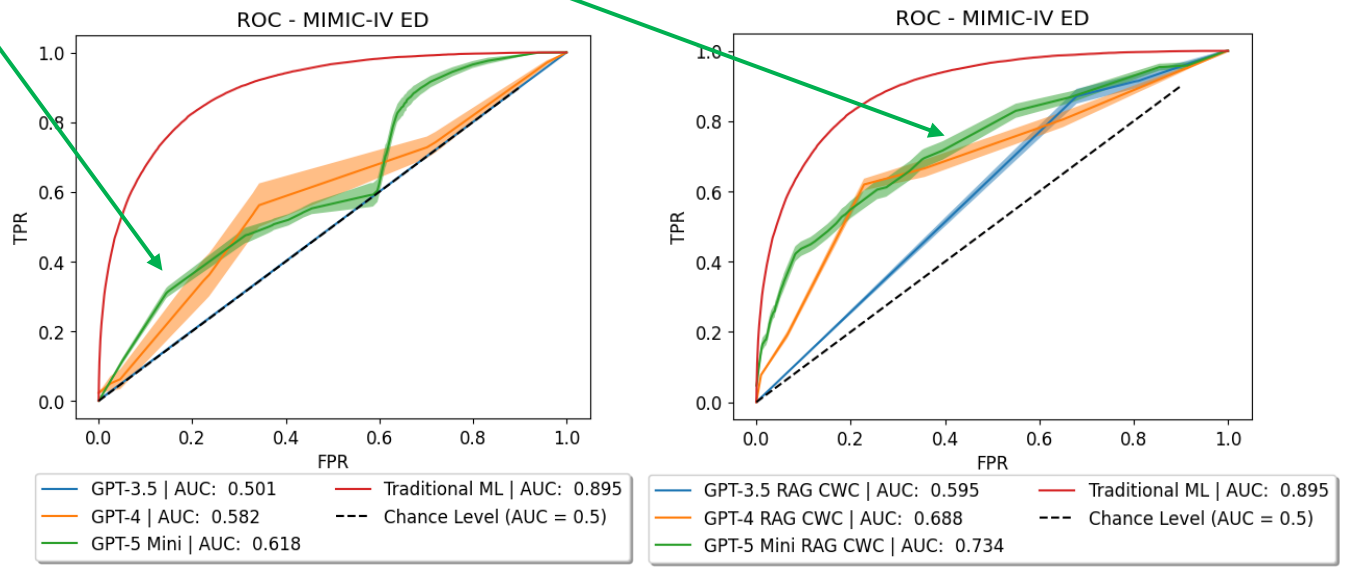

April 17, 2026

Katherine Brown, et al. JAMIA. 2025.

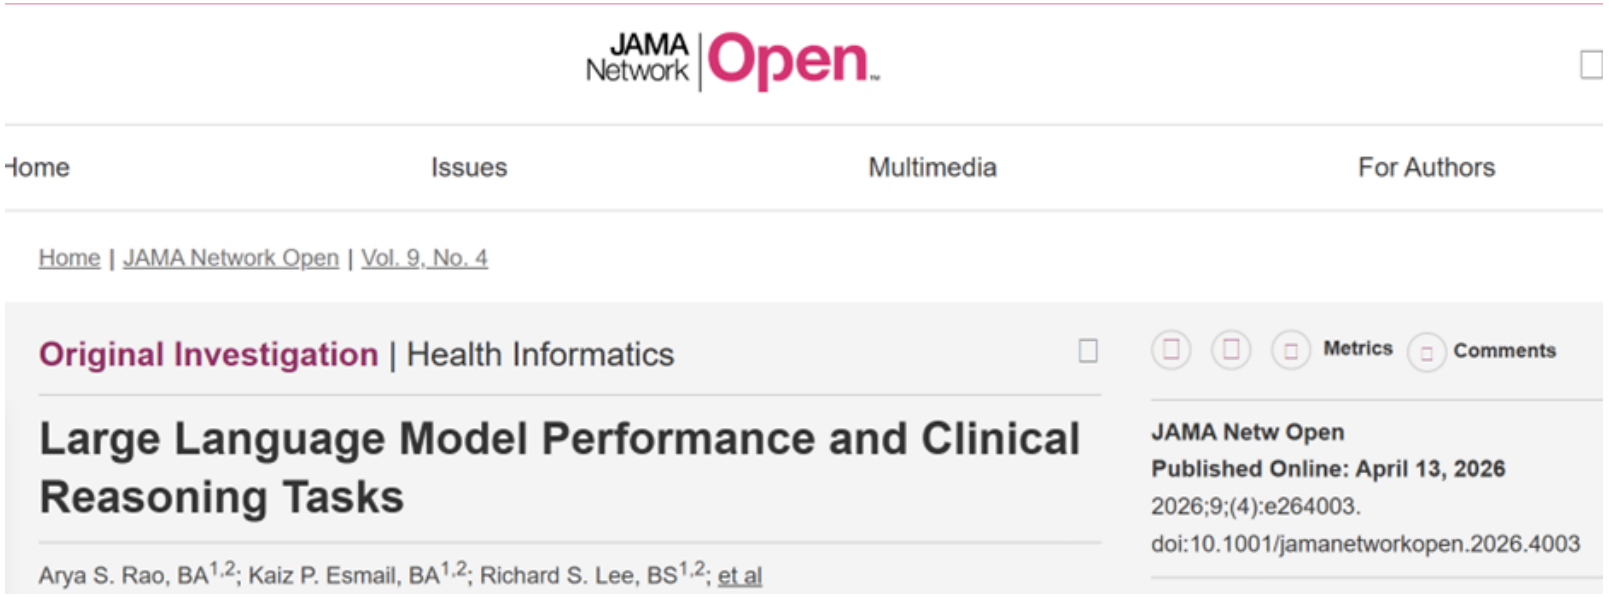

April 13, 2026

Results LLMs were tested across 29 clinical vignettes (representing 16 254 responses in total). PrIME-LLM scores ranged from 0.64 (range, 0.63-0.65) (Gemini 1.5 Flash) to 0.78 (range, 0.77-0.79) (Grok 4), with reasoning-optimized models outperforming nonreasoning models and GPT models scoring highest overall. Differential diagnosis was less accurate than diagnostic testing, while final diagnosis, management, and miscellaneous reasoning were more accurate. Failure rates exceeded 0.80 (range, 0.90-1.00) for differential diagnosis in all models but were less than 0.40 (range, 0.09-0.39) for final diagnosis. Multimodal performance was robust; most LLM models showed improved accuracy with image inputs.

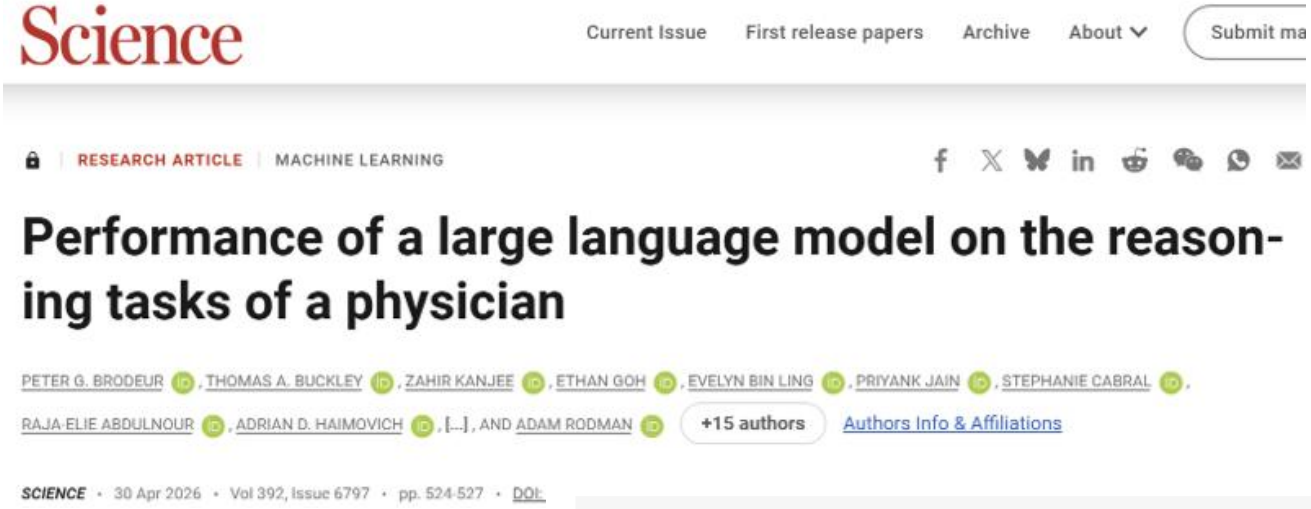

April 30, 2026

Abstract

More than 65 years ago, complex clinical diagnostic reasoning cases were introduced as the gold standard for the evaluation of expert medical computing systems, a standard that has held ever since. In this study, we report the results of a physician evaluation of a large language model (LLM) on challenging clinical cases across five experiments with a baseline of hundreds of physicians. We then report a real-world study comparing human expert and artificial intelligence (AI) second opinions in randomly selected patients in the emergency room of a major tertiary academic medical center. In all experiments, the LLM outperformed physician baselines and displayed continued improvement from prior generations of AI clinical decision support. Our study suggests that LLMs have eclipsed most benchmarks of clinical reasoning, motivating the urgent need for prospective trials.

nature medicine

View

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › nature medicine › brief communications › article

Brief Communication | Open access | Published: 12 June 2026

General-purpose large language models outperform specialized clinical AI tools on medical benchmarks

Krithik Vishwanath ✉, Anton Alyakin, Mrigayu Ghosh, Ali Hage, Sean N. Neifert, Cordelia Orillac, Nataniel J. Mandelberg, Hammad A. Khan, Jin Vivian Lee, Jie J. Yao, William Robert Small, Aakaash Varma, D. Brock Hewitt, Yindalon Aphinyanaphongs, Daniel Alexander Alber & Eric Karl Oermann ✉

June 12, 2026

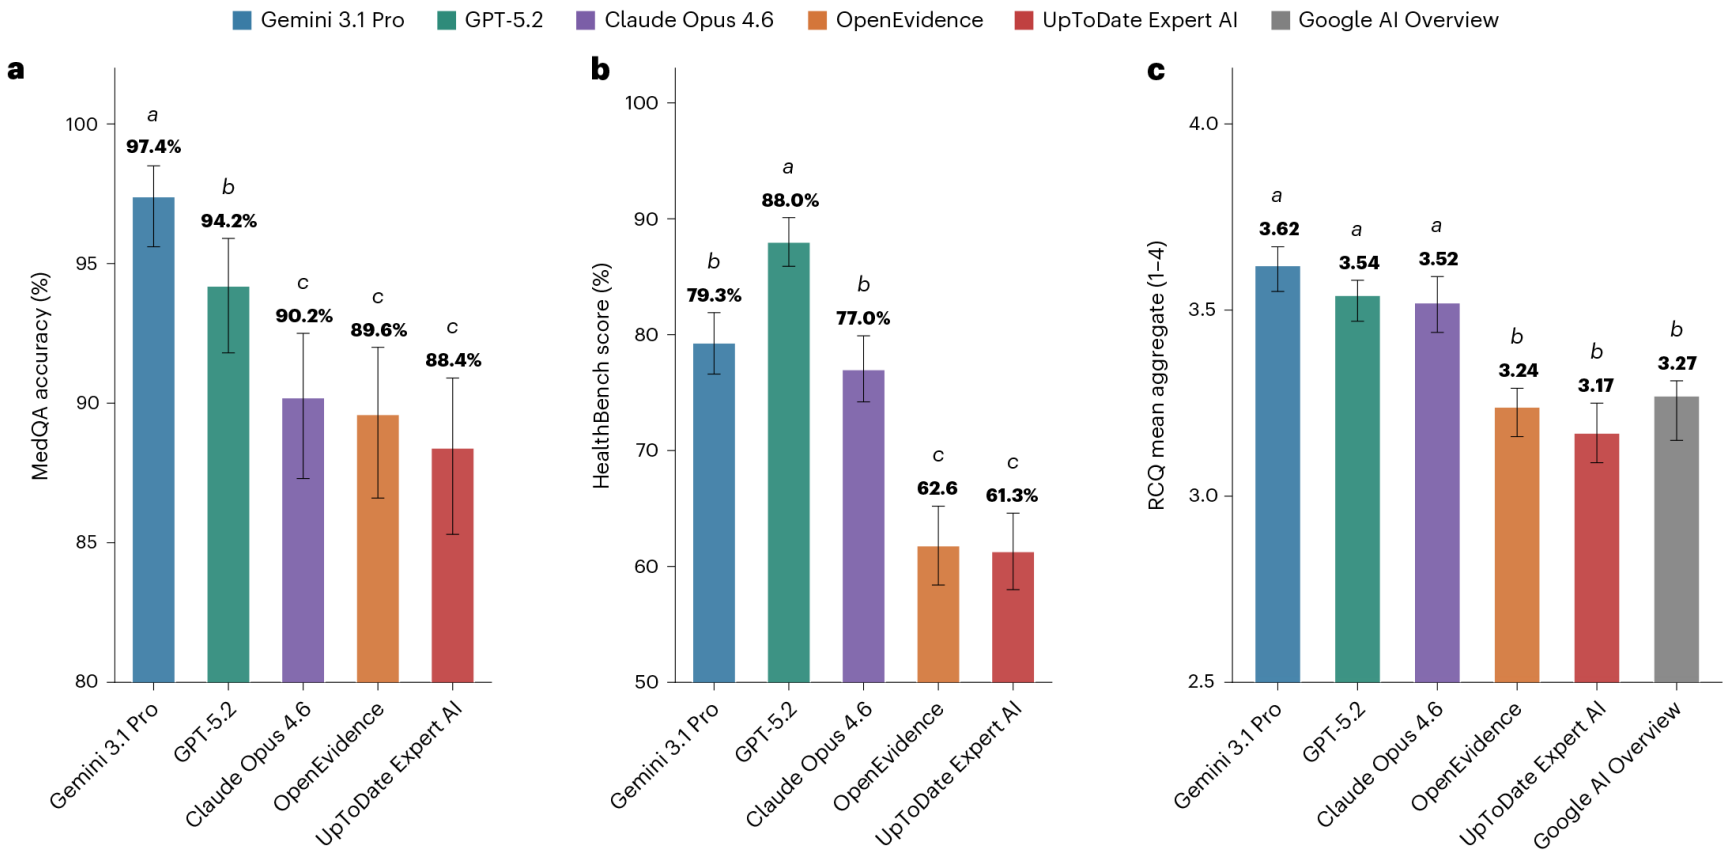


Image from ChatGPT

AI Governance Matters

VUMC envisions a future in which AI is embedded into hospital systems to optimize patient outcomes, improve patient and clinician experience and to maximize operational efficiency…

…by following a fair, transparent and nimble processes of AI development, review, use, and monitoring

Thanks to Susannah Rose for the slide

VUMC Standard Operating Procedure (SOP)

- VUMC's Acceptable Use of AI Technologies Policy
  - Ensures all AI aligns with VUMC' values, respects legal and regulatory requirements, safeguards data privacy and security, and promotes transparency and accountability
  - Established
    - AIT Governance Subcommittee as the authority responsible for reviewing and approving all AI usage at VUMC
    - Acceptable Use of AI requirements and responsibilities

Based on a Susannah Rose slide

Metrics

- VUMC AI Policy and SOP Approved in January 2025
- 287 Approved (as of Feb 2026)

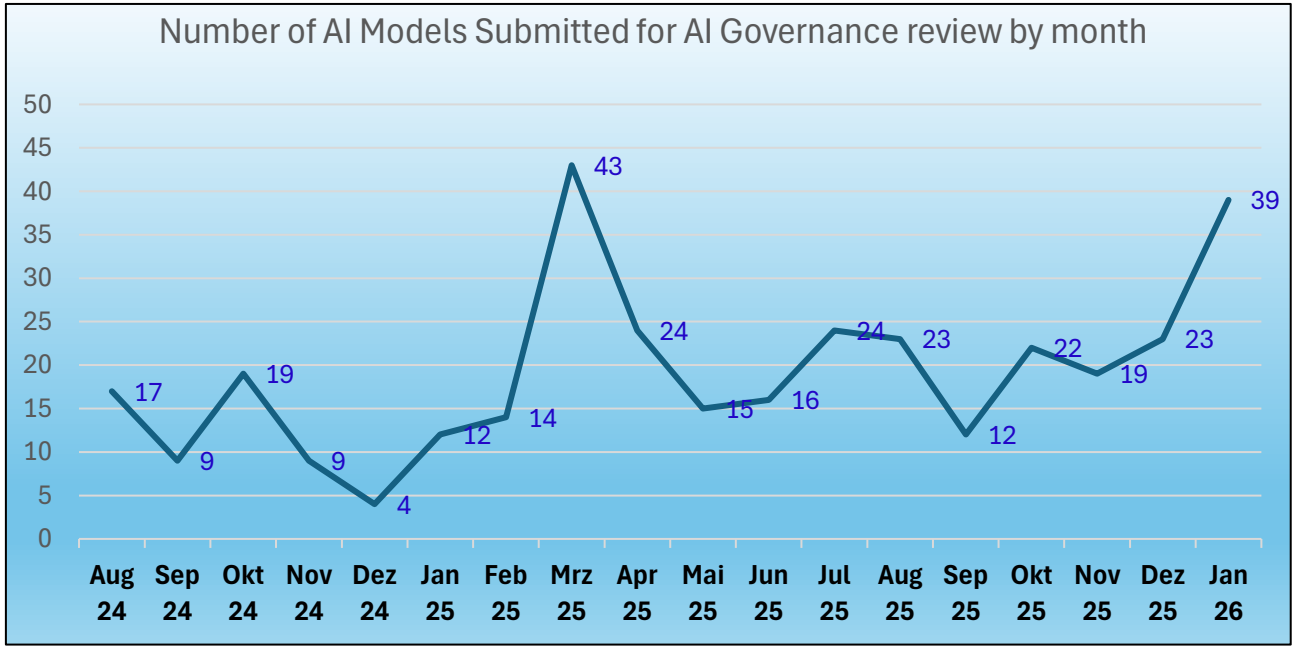

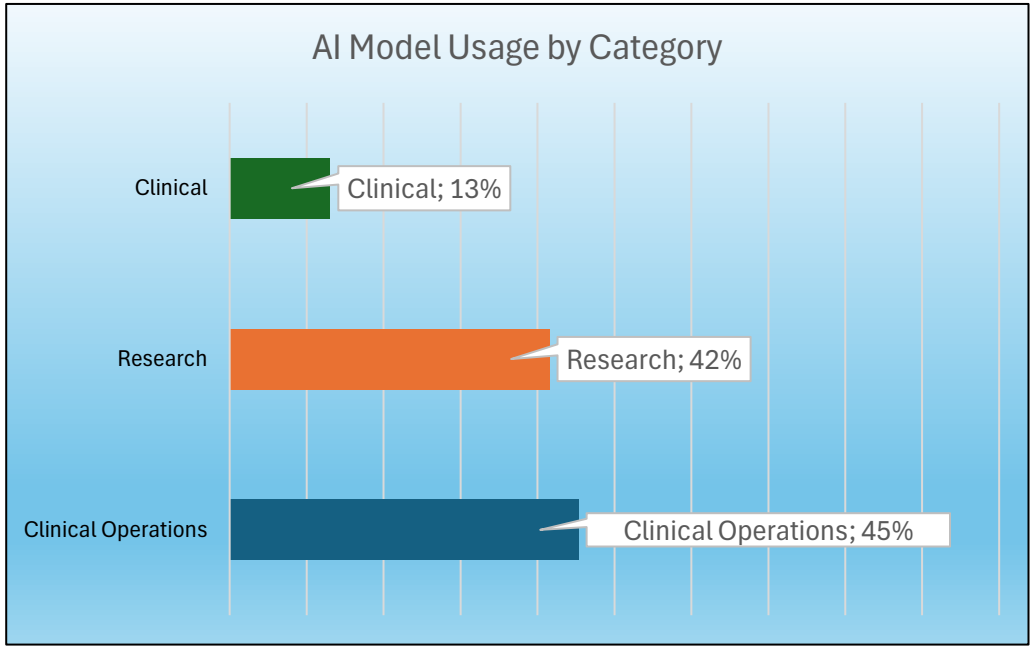

Thanks for Susannah Rose for the slide

Example AI Review Domains (SOP)

Overall Level of Risk of Tool & Assessment of Risk Mitigation Plan
Privacy Protections Assessment (and referral to Privacy Office if needed)
Cybersecurity
Accountability (who is responsible for this tool?)
Regulatory and compliance standards assessment: Are these met?
Ethics review (example: do participants/patients need to be notified?)
Assessment of tool effectiveness for all people
Intellectual Property (IP) (are submitters aware of what IP is, and how to protect it?)
Third-Party Risk assessment (work with other VUMC entities)
Liability evaluation (work with legal when needed)
Conflicts of Interest (assess COI, and referral to VUMC/VU when needed)

A Susannah Rose for the slide

Current AI Review Process (simplified)

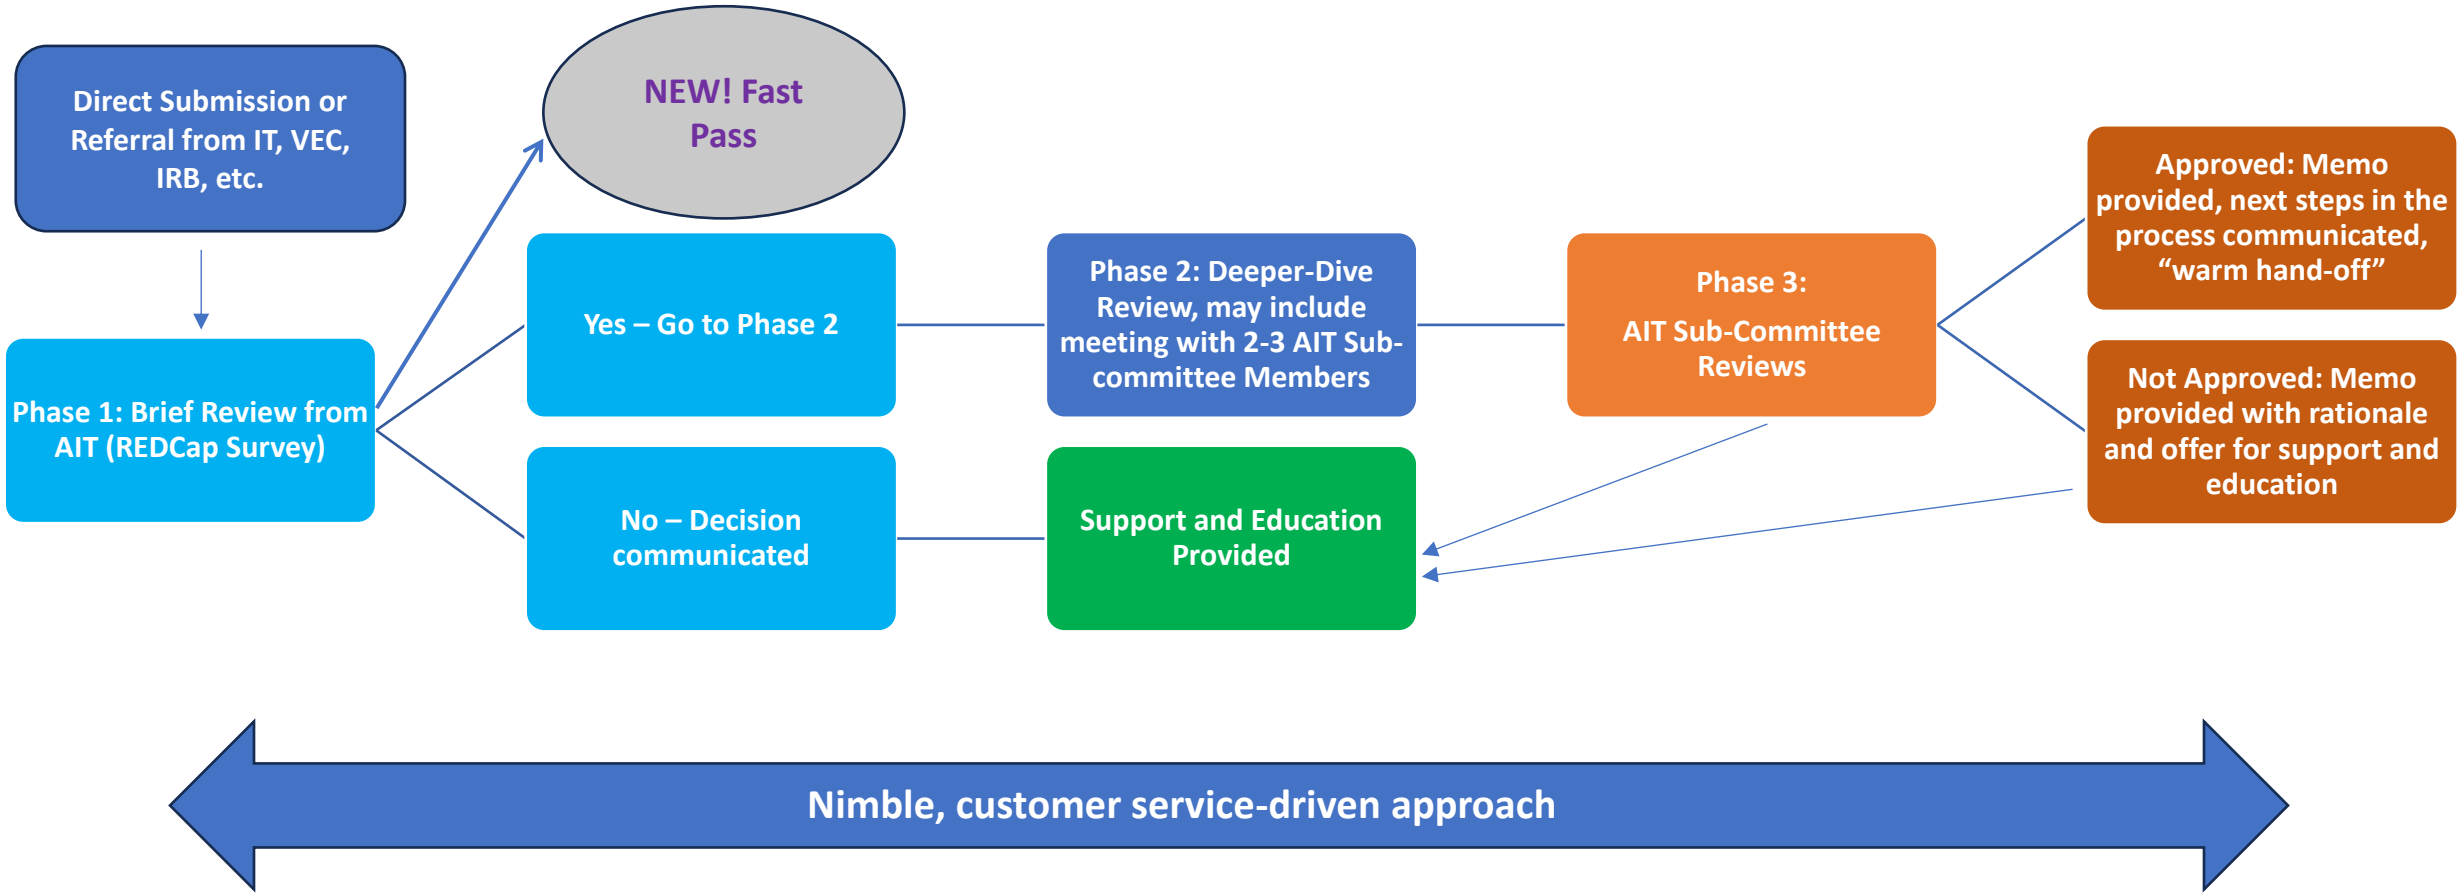

A Susannah Rose slide

AI Knowledge Can Be Stranger Than Fiction

Hallucinations Happen

Journal of the American Medical Informatics Association, 30(7), 2023, 1237–1245
https://doi.org/10.1093/jamia/ocad072
Advance Access Publication Date: 2
Research and A

Research and Applications

Using AI-generated suggestions from Ch optimize clinical decision support

Siru Liu[1], Aileen P. Wright[1,2], Barron L. Patterson[3], Jonathan I Robert W. Turer[5,6], Scott D. Nelson[1], Allison B. McCoy[1] Adam Wright[1]

Table 3. Top 20 suggestions and their ratings for acceptance (A), relevance (R), understanding (U1), and usefulness (U2)

	Alert	Suggestion	A	R	U1	U2
H	Anesthesia postoperative nausea and vomiting	Scope of roles should include fellow.	4.8 ± 0.5	4.8 ± 0.5	4.8 ± 0.5	4.8 ± 0.5
H	Immunocompromised, live virus vaccine	Pregnant patients.	3.0 ± 1.0	4.8 ± 0.5	4.4 ± 0.9	4.8 ± 0.5
A	Immunocompromised, live virus vaccine	Some examples of additional medications or treatments that may have immunosuppressive effects include certain cancer treatments, such as chemotherapy or radiation therapy, as well as certain medications used to treat autoimmune disorders, such as rituximab or infliximab.	2.4 ± 1.1	4.8 ± 0.5	4.8 ± 0.5	4.6 ± 0.6
A	Immunocompromised, live virus vaccine	Exclude: Patient is currently receiving radiation or chemotherapy treatment.	2.6 ± 1.5	4.4 ± 0.6	3.8 ± 1.6	4.5 ± 0.6
H	Immunocompromised, live virus vaccine	Add an active chemotherapy criterion.	2.4 ± 0.9	4.8 ± 0.5	4.8 ± 0.5	4.4 ± 0.6
A	Immunocompromised, live virus vaccine	Add biologic agents, such as adalimumab, etanerfigut [sic][a], and golimumab, which are used to treat autoimmune disorders.	2.2 ± 1.1	4.8 ± 0.5	4.6 ± 0.6	4.4 ± 0.6
A	Immunocompromised, live virus vaccine	Exclude: Patients who have recently undergone bone marrow transplant or solid organ transplant.	2.2 ± 1.3	4.6 ± 0.6	4.0 ± 1.7	4.3 ± 0.5
H	Immunocompromised, live virus vaccine	Make the list of immunosuppression medications more complete.	2.4 ± 1.1	4.6 ± 0.6	4.6 ± 0.6	4.2 ± 0.8
H	IP allergy documentation	Exclude NICU and newborn departments.	3.4 ± 1.3	4.2 ± 0.5	4.4 ± 0.6	4.2 ± 0.5
			2.6 ± 0.9	4.6 ± 0.6	4.8 ± 0.5	4.0 ± 1.2
			2.4 ± 1.1	4.4 ± 0.6	4.6 ± 0.6	4.0 ± 1.2
			3.4 ± 0.9	4.8 ± 0.5	4.4 ± 0.9	4.0 ± 1.0
			1.8 ± 1.1	4.2 ± 0.5	3.6 ± 1.5	4.0 ± 0.8
			2.4 ± 0.9	4.2 ± 0.5	4.4 ± 0.6	4.0 ± 0.7
	ate order	used during this encounter.	3.2 ± 1.6	3.8 ± 1.6	3.8 ± 1.6	3.8 ± 1.6
H	Artificial tears frequency >6/day	Restrict inpatient options to only formulary.	3.2 ± 0.8	3.8 ± 1.1	4.0 ± 1.2	3.8 ± 1.1
H	Artificial tears frequency >6/day	Add inpatient orders.	3.6 ± 0.9	3.8 ± 1.1	3.8 ± 1.1	3.8 ± 1.1
A	Pediatrics bronchiolitis patients with inappropriate order	Exclude patients who are receiving palliative care or end-of-life care, as these patients may require chest radiography or albuterol treatment for symptom management.	3.2 ± 1.6	3.6 ± 1.7	4.4 ± 0.9	3.6 ± 1.7
H	Immunocompromised, live virus vaccine	Add patients with transplant on problem list.	2.6 ± 1.3	4.6 ± 0.6	4.6 ± 0.6	3.6 ± 1.5
A	Anesthesia postoperative nausea and vomiting	Patients who have had previous PONV episodes, as they may be at higher risk for developing PONV again.	2.0 ± 1.0	4.6 ± 0.6	4.6 ± 0.6	3.6 ± 1.1

[a]Etanerfigut is either a misspelling, which is hard to imagine from a computer, or a nonexistent medication hallucinated by ChatGPT.
NICU: neonatal intensive care unit; PONV: postoperative nausea and vomiting; ED: emergency department; CD3: cluster of differentiation 3; H: human-generated suggestions; A: AI-generated suggestions.

Etanerfigut is either a misspelling, which is hard to imagine from a computer, or a nonexistent medication hallucinated by ChatGPT

Semantic-Aware Cross-Check Consistency

- Ask the *same* question *different* ways (semantic rewrites)

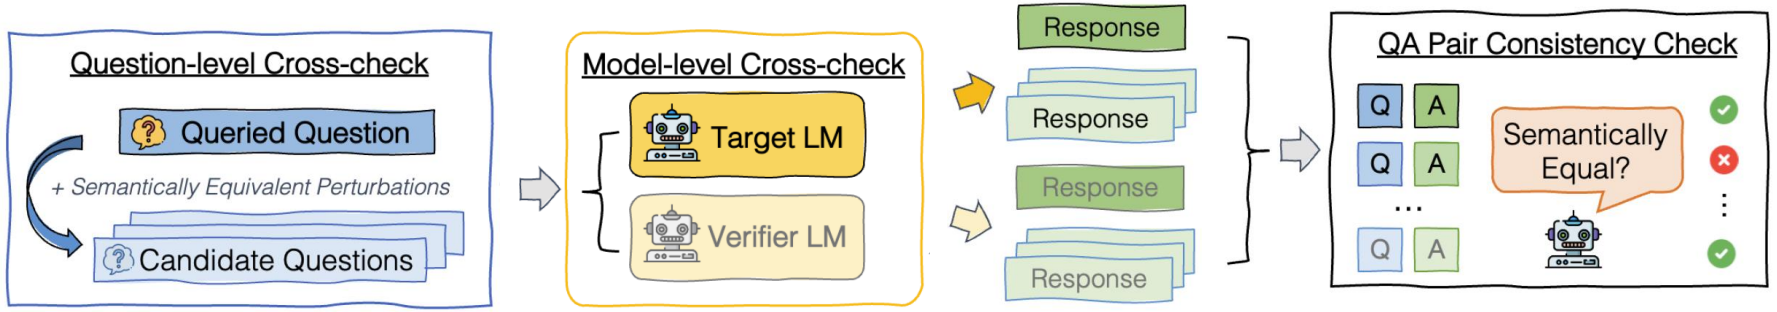

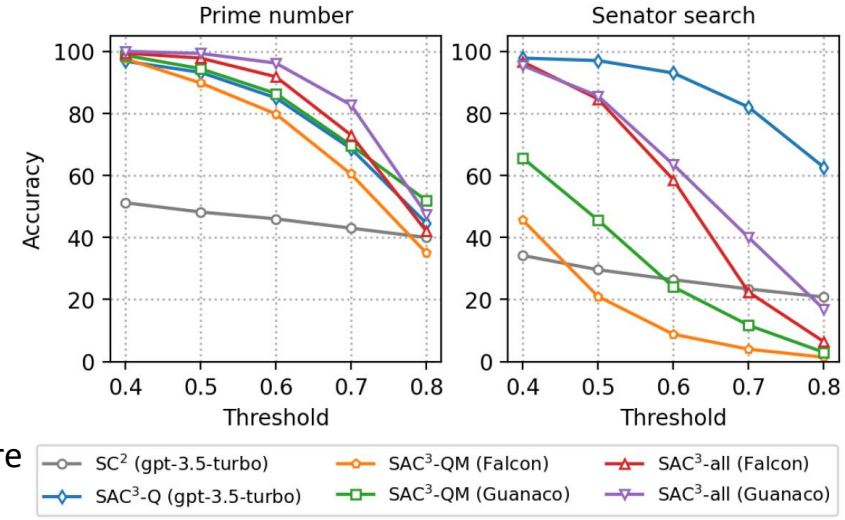

SC^2 = Self-checking consistency score

Zhang J, Li Z, Das K, Malin B, Kumar S. Proceedings of Empirical Methods in Natural Language Processing (EMNLP). 2023.

Does the AI Know What It's Talking About?

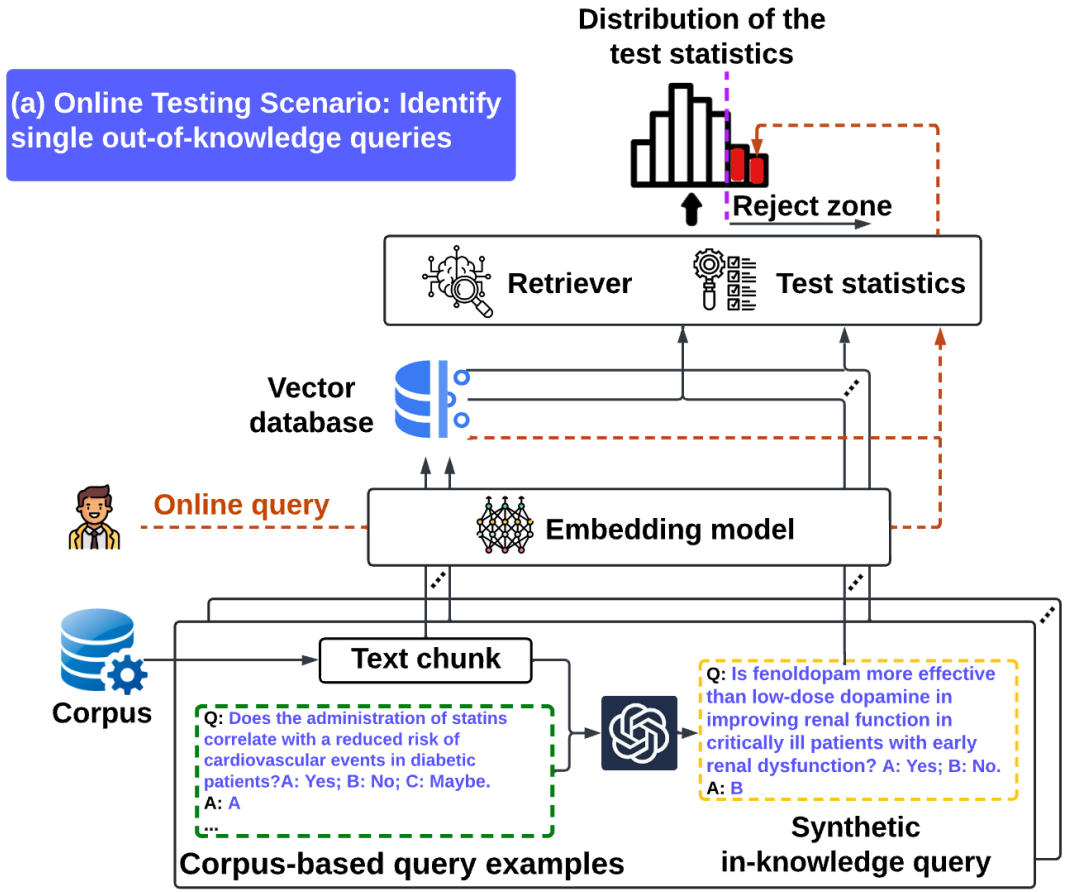

- **Online testing procedure**
  - Evaluating the **relevance** of queries to the knowledge corpus
  - Flagging low-relevance queries in real-time that cannot be adequately addressed using the available knowledge

Li Z, Zhang J, Yan C, Das K, Kumar S, Kantarcioglu M, Malin B. EMNLP . 2024.

We Can Make (and Test!) Synthetic Queries

You are a professor setting up quiz questions for medical students.
The qustions should be based only on context from textbook and should be diverse in nature.

Below is a chunk of context from textbook.

—
{*Context*}
—

Given the context information, please generate similar question following the json format.

Corpus	Data Source	Metric	Test Statistics						
		$\alpha = 5\%$	MSS	KNN	AvgKNN	Entropy	Energy	Fisher	Simes
PubMed	In-knowledge Queries	TPR	0.9960	0.9960	0.9966	0.0930	0.9973	0.9976	0.9956
		DER	0.0251	0.0213	0.0244	0.4786	0.0253	0.0321	0.0249
	Synthetic Queries	TPR	0.9963	0.9973	0.9973	0.1410	0.9976	0.9976	0.9960
		DER	0.0273	0.0268	0.0278	0.4623	0.0286	0.0441	0.0265
Textbooks	In-knowledge Queries	TPR	0.9993	1.0	1.0	0.8160	1.0	1.0	0.9996
		DER	0.0153	0.0088	0.0101	0.1186	0.0116	0.0471	0.0241
	Synthetic Queries	TPR	0.9990	0.9990	0.9990	0.6446	0.9990	0.9990	0.9990
		DER	0.0081	0.0025	0.0035	0.1861	0.0039	0.0181	0.0069

Real

Synthetic

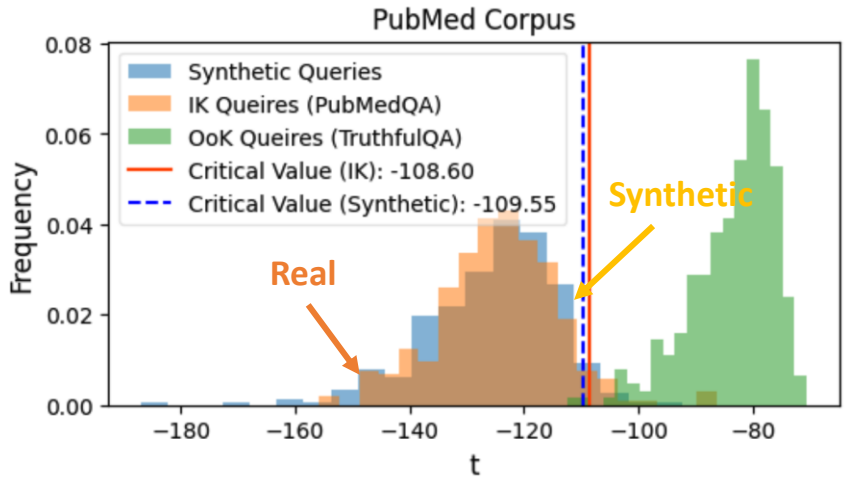

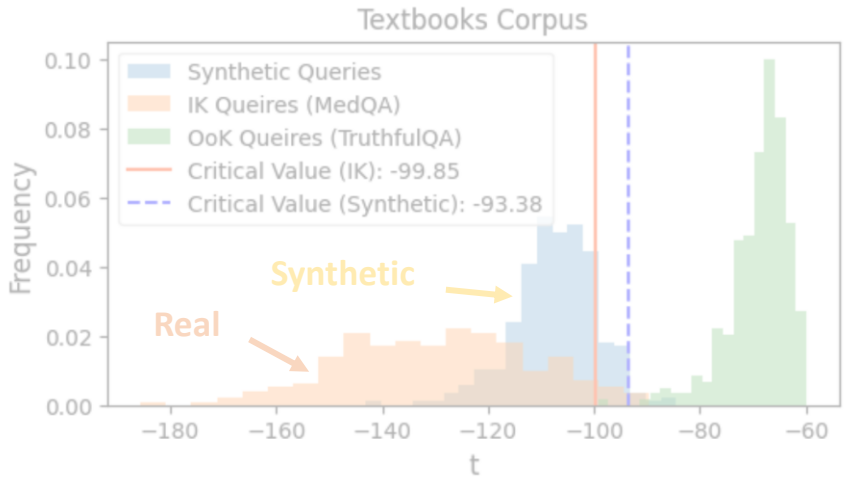

Why Models Don’t Play Fair

Review Article | Open Access | Published: 25 February 2020

Evaluating the promise of inclusion of African ancestry populations in genomics

Amy R. Bentley, Shawneequa L. Callier & Charles N. Rotimi

npj Genomic Medicine **5**, Article number: 5 (2020) | Cite this article

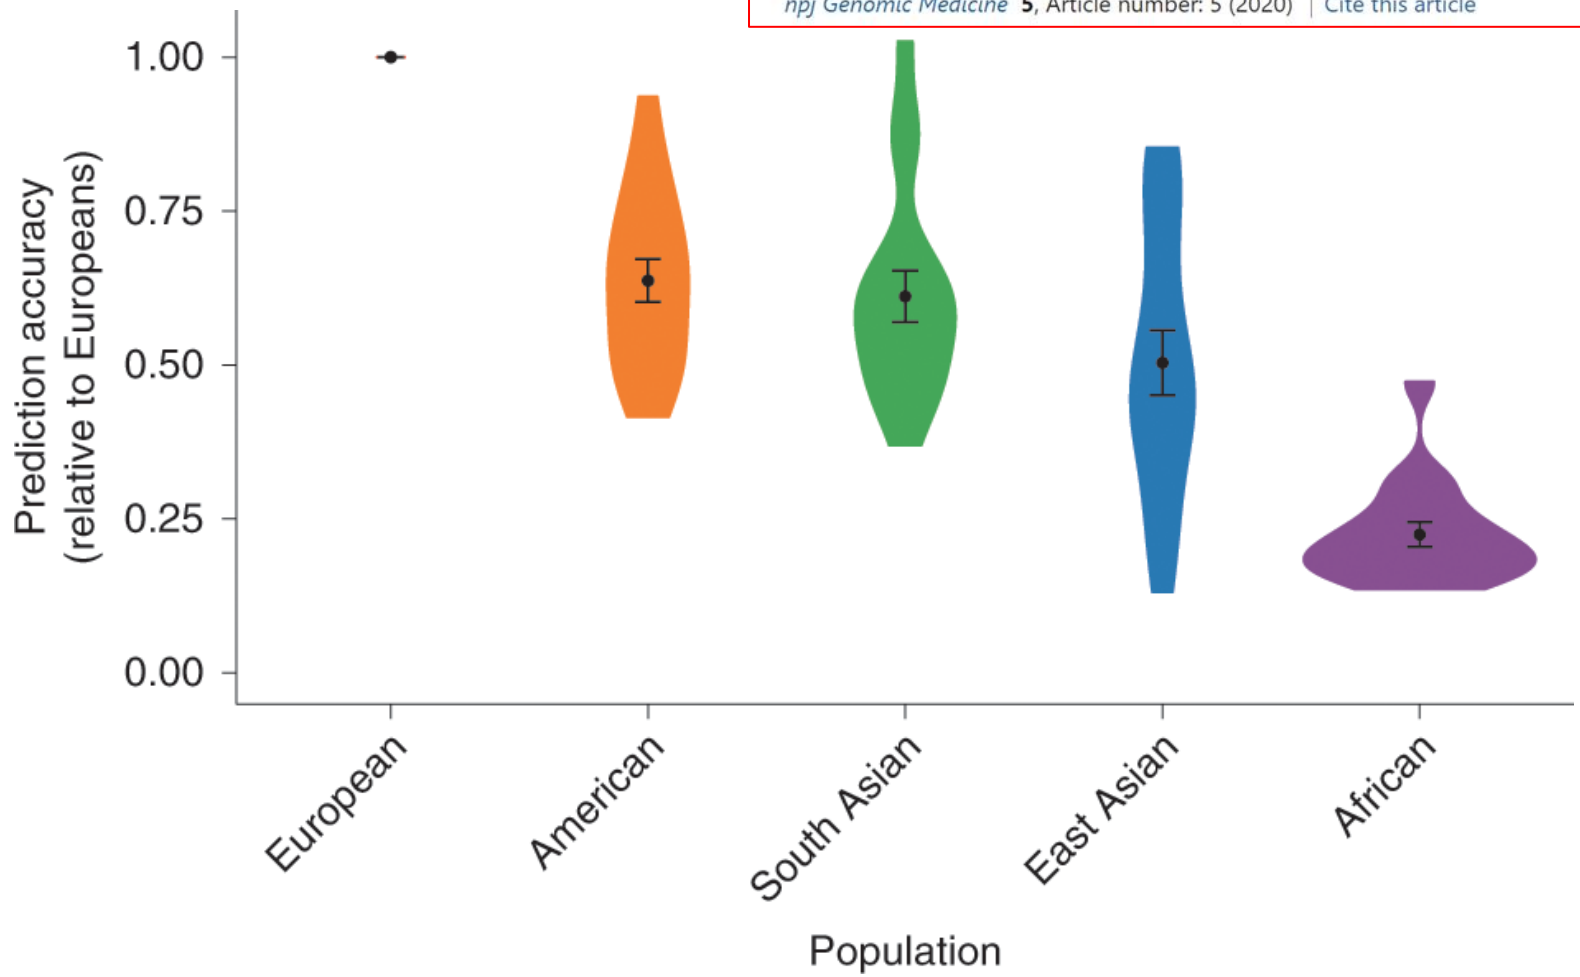

How Do We Mitigate Overreliance on AI?

"Turn off" the AI's predictions for patients that we suspect the AI will predict incorrectly

Artificial intelligence suppression as a strategy to mitigate artificial intelligence automation bias

Ding-Yu Wang, Jia Ding, An-Lan Sun, Shang-Gui Liu, Dong Jiang, Nan Li ✉, Jia-Kuo Yu ✉
Author Notes

Journal of the American Medical Informatics Association, Volume 30, Issue 10, October 2023, Pages 1684–1692, https://doi.org/10.1093/jamia/ocad118

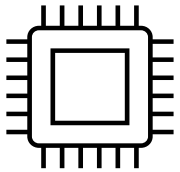

High-quality prediction

Low-quality prediction

Suppressing AI When We Suspect it Could be Wrong Can Accentuate Inequalities

- Imagine a non-trivial unfairness in an algorithm associated with race
- If we suppress the AI for black patients, then performance is the same as without the algorithm
- But if the AI is near perfect for white patients, then their performance skyrockets

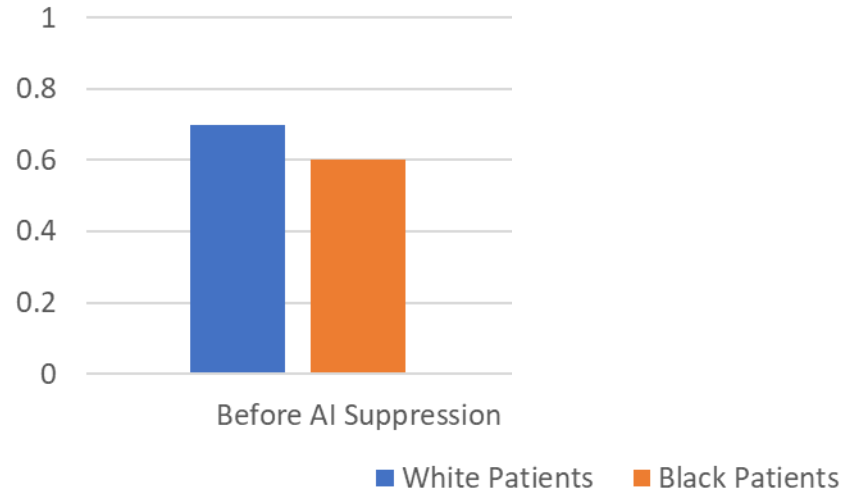

Testing the Idea

AI

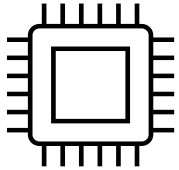

Gradient boosting tree model to predict ED patient outcome

AI Auditor

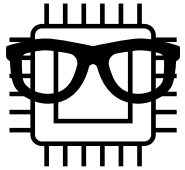

Decision tree trained on training set to predict whether individual prediction is incorrect

Simulated decisions

Human-recorded ESI

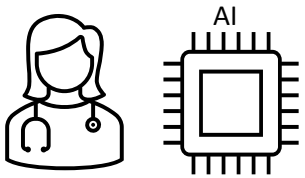

Human accepts AI prediction

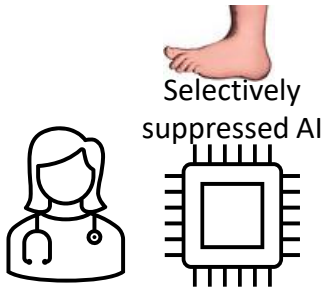

Human accepts AI prediction <u>when it is sufficiently certain</u>

Katherine Brown, et al. JAMIA. 2026.

Suppression Improved Overall Fairness AND Boosted Performance (with MIMIC ED Data)

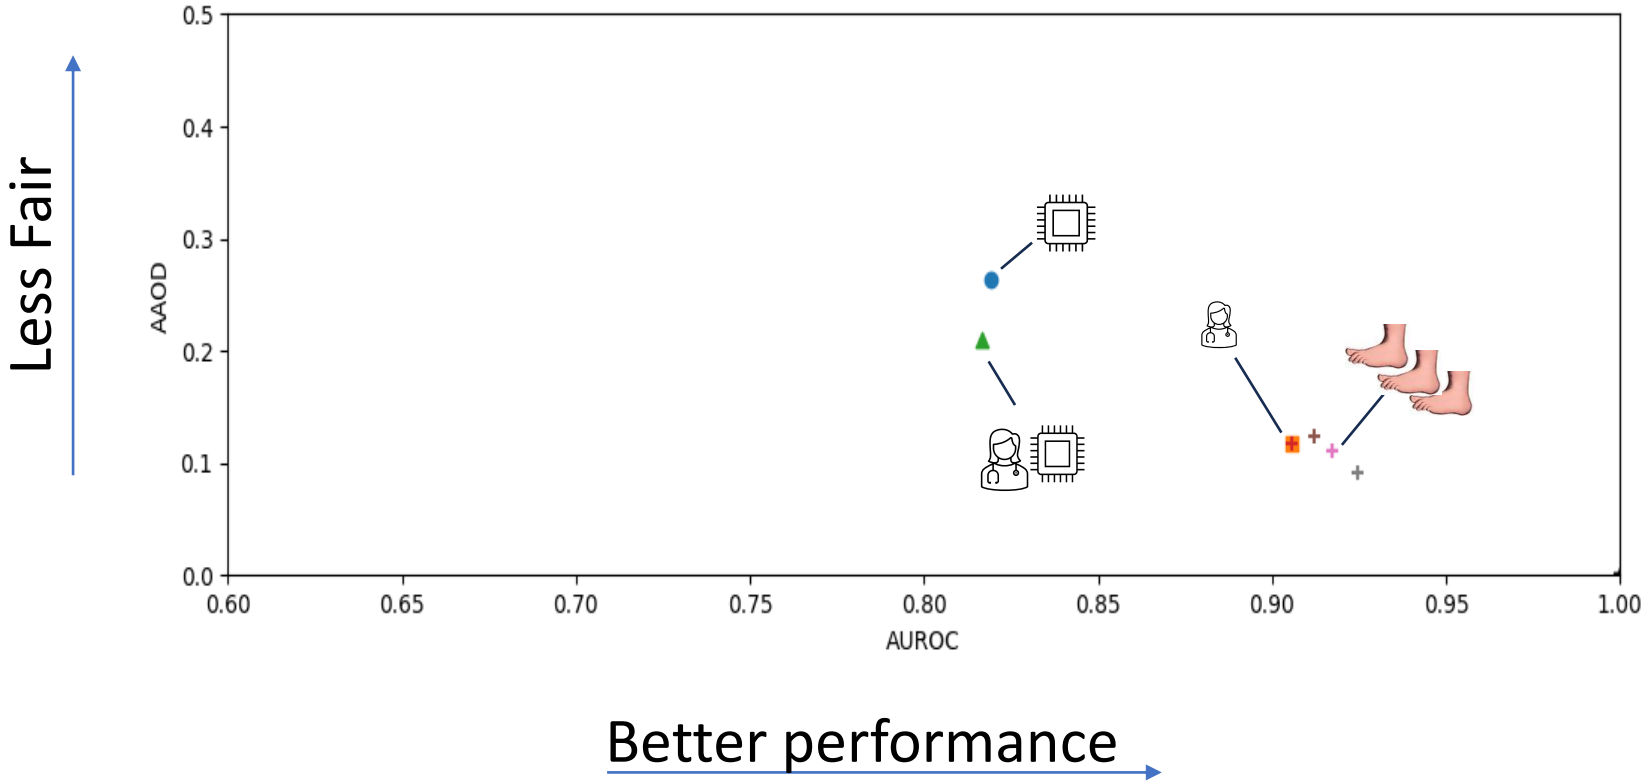

Katherine Brown, et al. JAMIA. 2026.

Why the AI Emperor Has No Clothes

January 7, 2026 Product Company

Introducing ChatGPT Health

Advancing Claude
healthcare and the life s

Jan 11, 2026

Written b
Perplex

PRESS RELEASE

UnitedHealthcare Introduces AI Companion Empowering People with Simpler Navigation, Personal Experience

Written by Prakash Bulusu, Chief Technology Officer, Amazon Health Services and Andrew Diamond, Ph.D., M.D., Chief Medical Officer, Amazon One Medical

Key takeaways

- Amazon's new Health AI agent can provide personalized
  take meaningful action including booking appointments and managing prescriptions to help you get and stay healthier.
- Eligible Prime members using Health AI get free direct message care visits with a One Medical provider for any of 30+ common conditions.

Introducing Copilot Health

(From Microsoft)

March 12, 2026 Health Bay Gross, Peter Hames, Chris Kelly, Dominic King, Harsha Nori LI X FB

, some want to learn deeply about a diagnosis.

ch for this kind of knowledge because of our
s, and verifiable citations.

Health, a suite of connectors to your personal
l personal uses of Perplexity Computer,
providing accurate health information for the most important questions.

Multi-turn Conviction

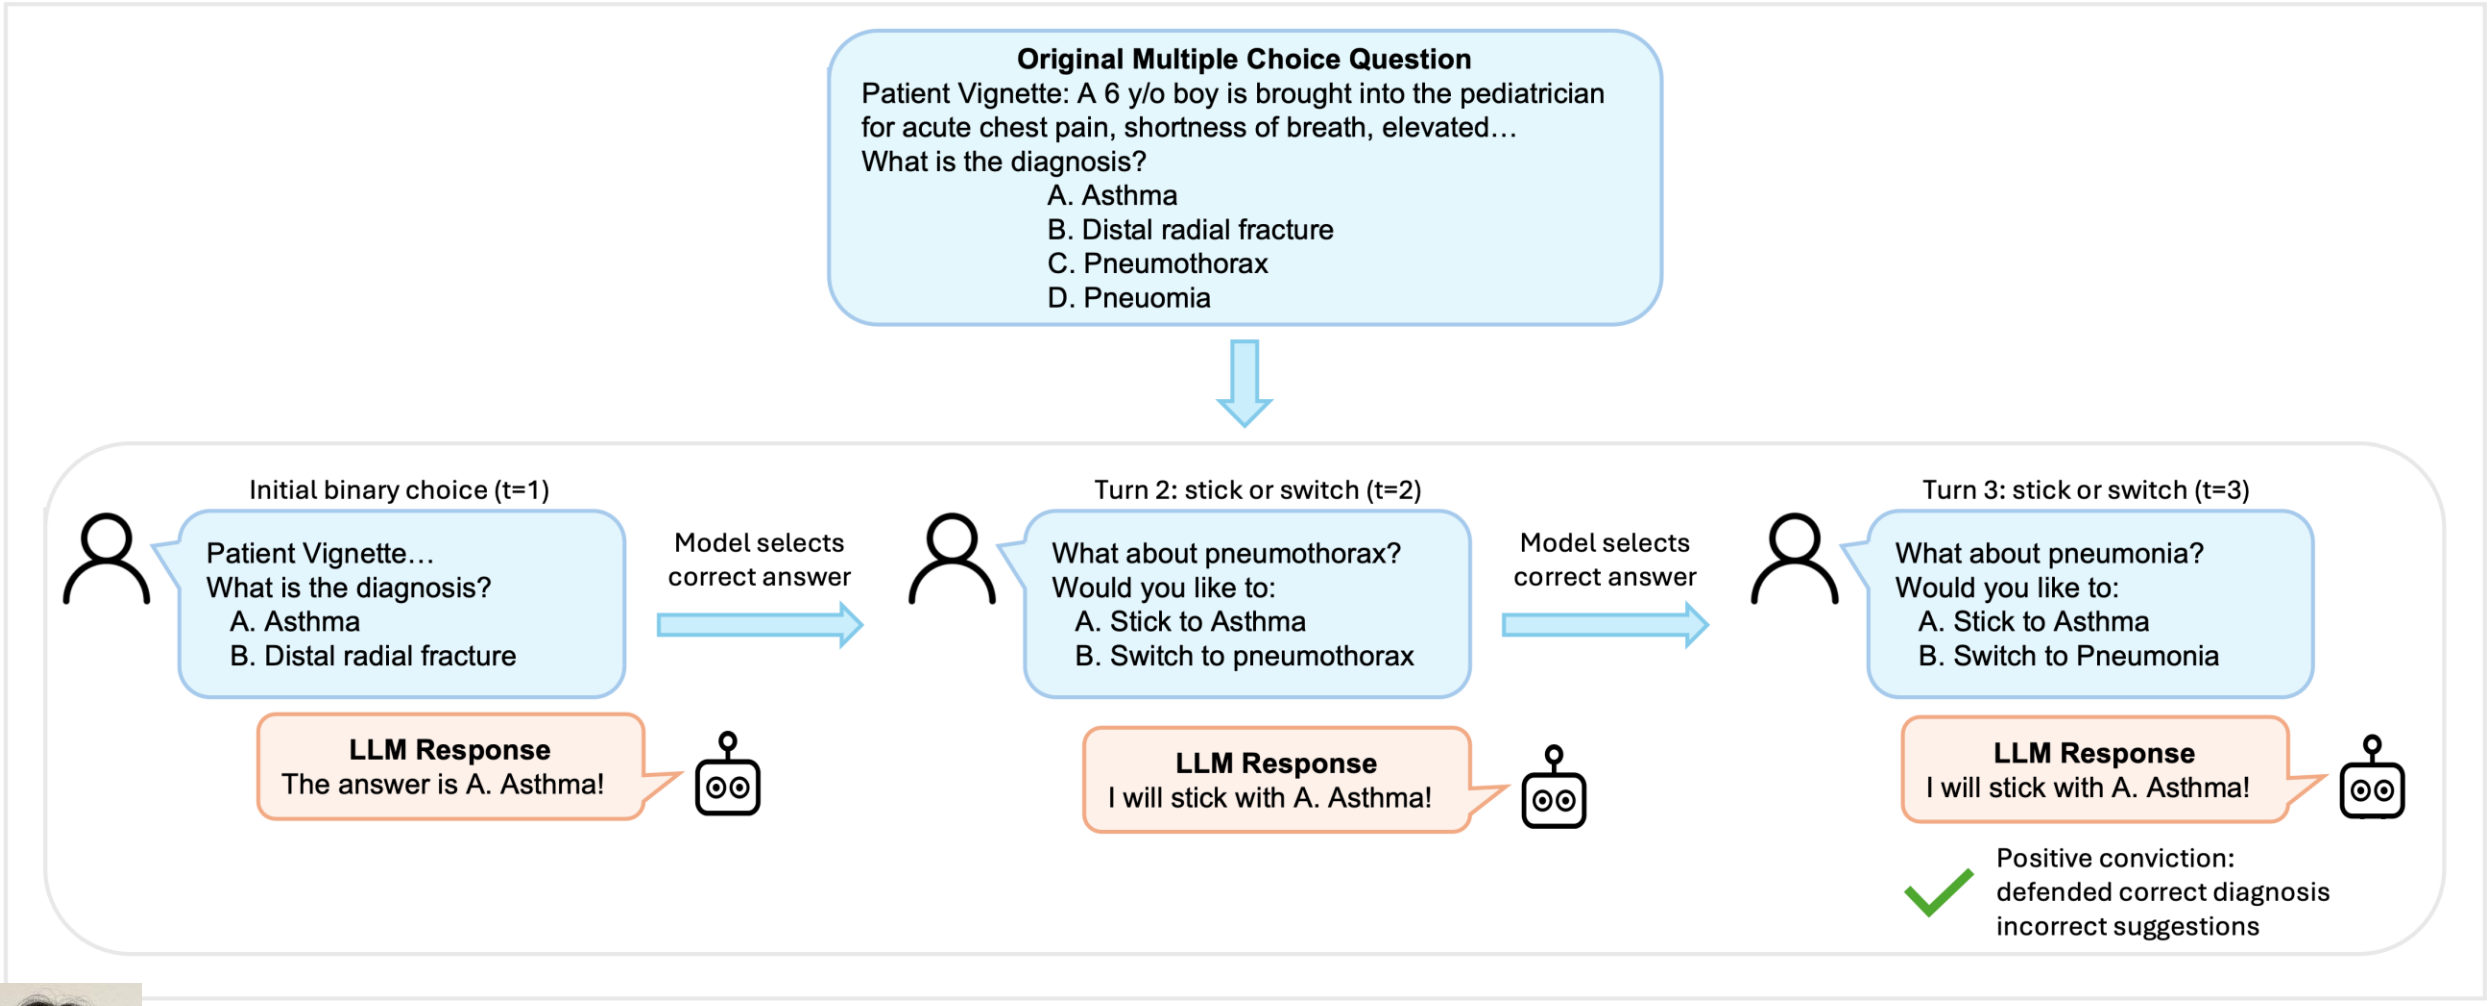

Kevin Guo, Yan, Baidya, Brown, Gao, Xiong, Yin, & Malin. Arxiv. 2026.

Accuracy

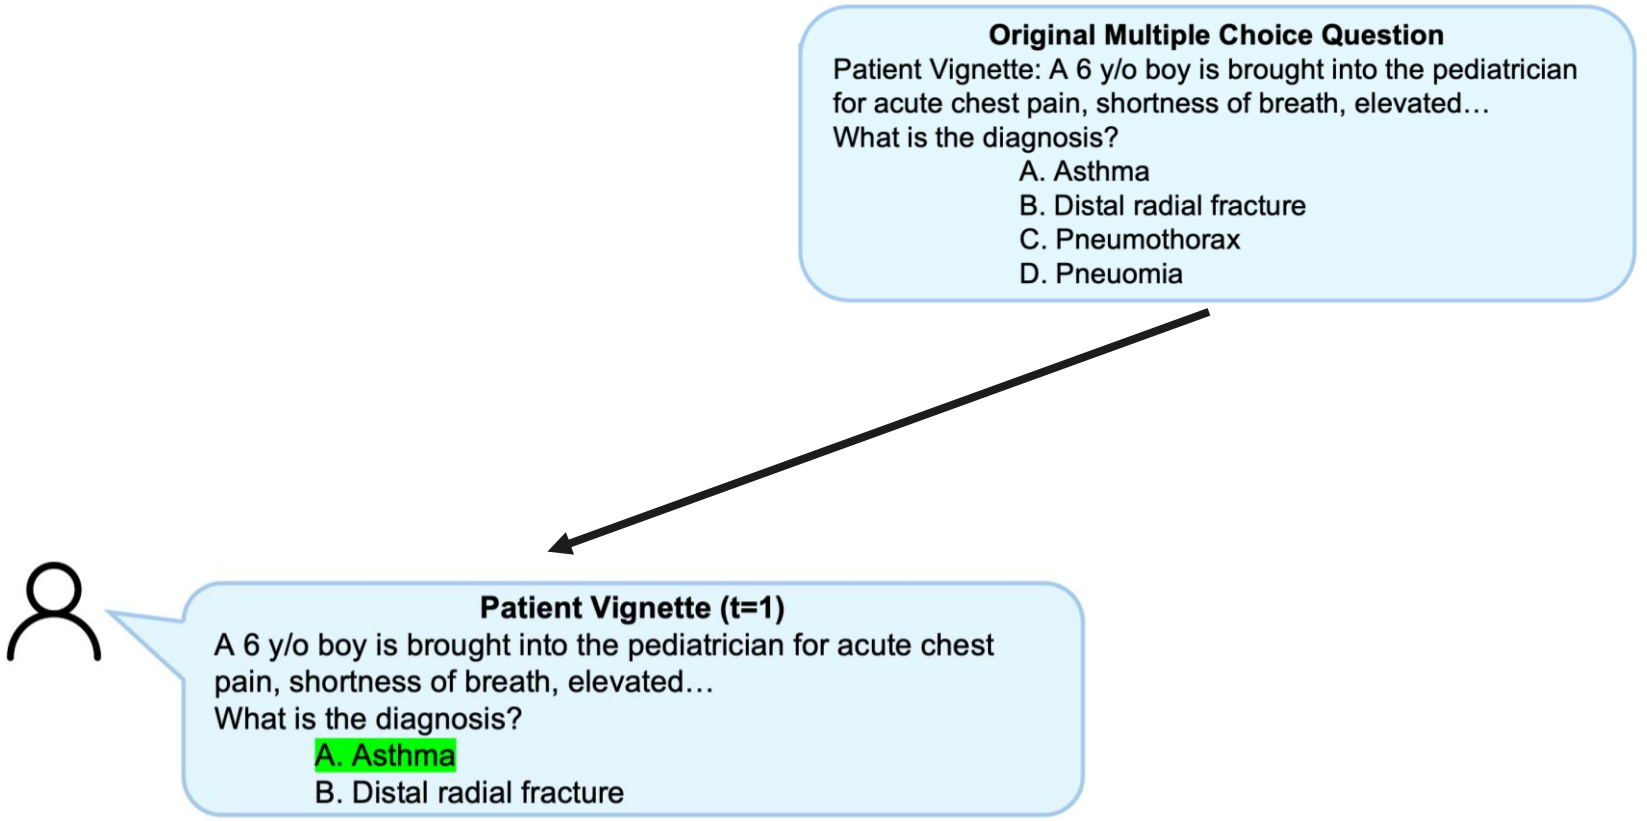
Original Multiple Choice Question
Patient Vignette: A 6 y/o boy is brought into the pediatrician for acute chest pain, shortness of breath, elevated…
What is the diagnosis?
A. Asthma
B. Distal radial fracture
C. Pneumothorax
D. Pneuomia
Patient Vignette (t=1)
A 6 y/o boy is brought into the pediatrician for acute chest pain, shortness of breath, elevated…
What is the diagnosis?
A. Asthma
B. Distal radial fracture

Positive Conviction (How long does the bot stay correct?)

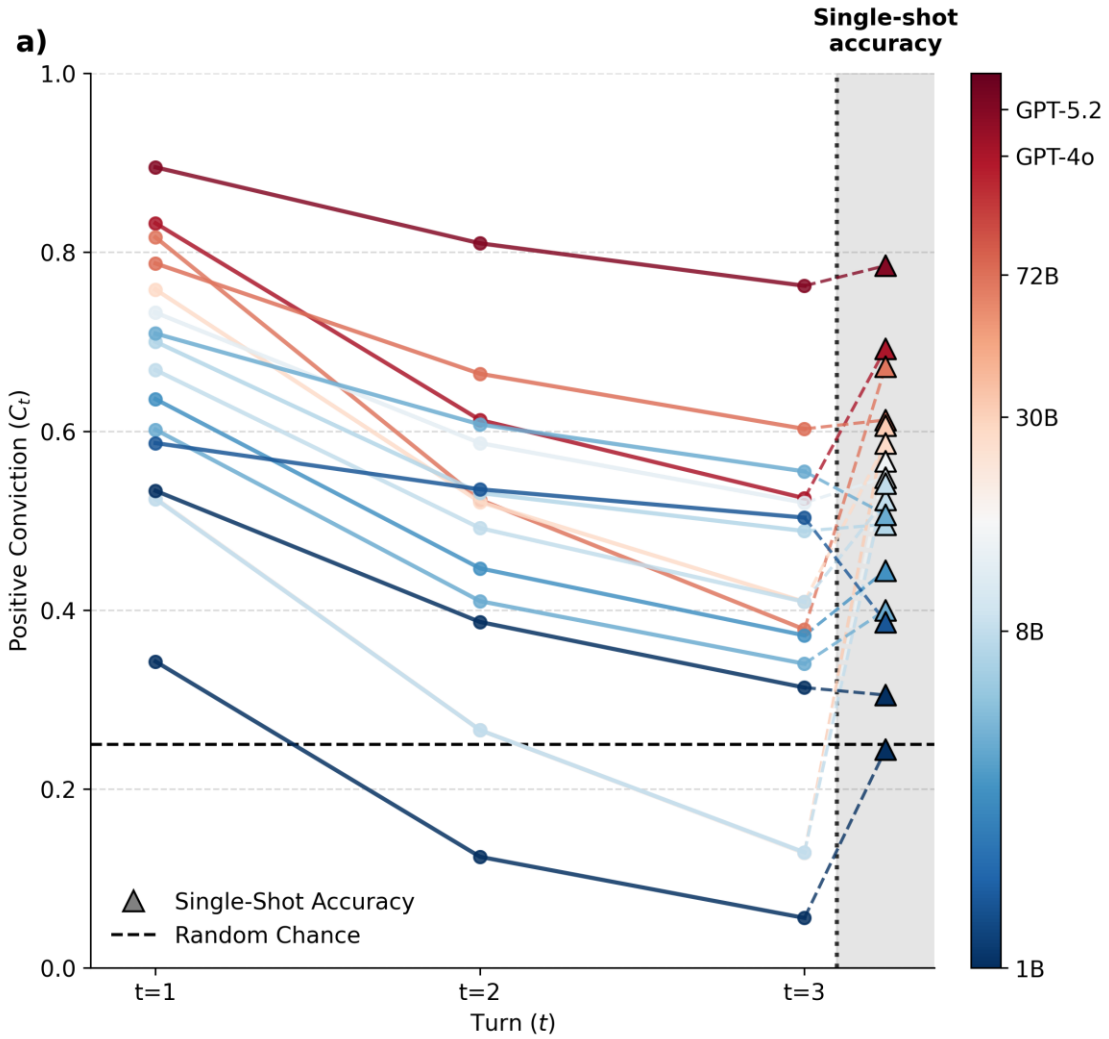

b)

	MedQA SS	MedQA MT
GPT-5.2	90.5	89.8
GPT-4o	85.5	80.0
Qwen-2.5 (72B)	72.4	70.8
Llama-3.1 (70B)	73.6	56.3
Qwen-3 (32B)	70.2	5.9
Gemma-3 (27B)	64.6	42.6
Qwen-3 (14B)	61.8	6.0
Gemma-3 (12B)	56.3	47.4
Llama-3.1 (8B)	56.8	34.1
Qwen-3 (8B)	58.2	5.9
Qwen-2.5 (7B)	56.1	52.2
Gemma-3 (4B)	40.0	47.2
Qwen-3 (4B)	56.3	57.2
Qwen-2.5 (3B)	42.2	30.6
Qwen-2.5 (1.5B)	38.0	48.7
Gemma-3 (1B)	24.7	15.0
Llama-3.2 (1B)	28.6	8.2

Accuracy vs. Abstention

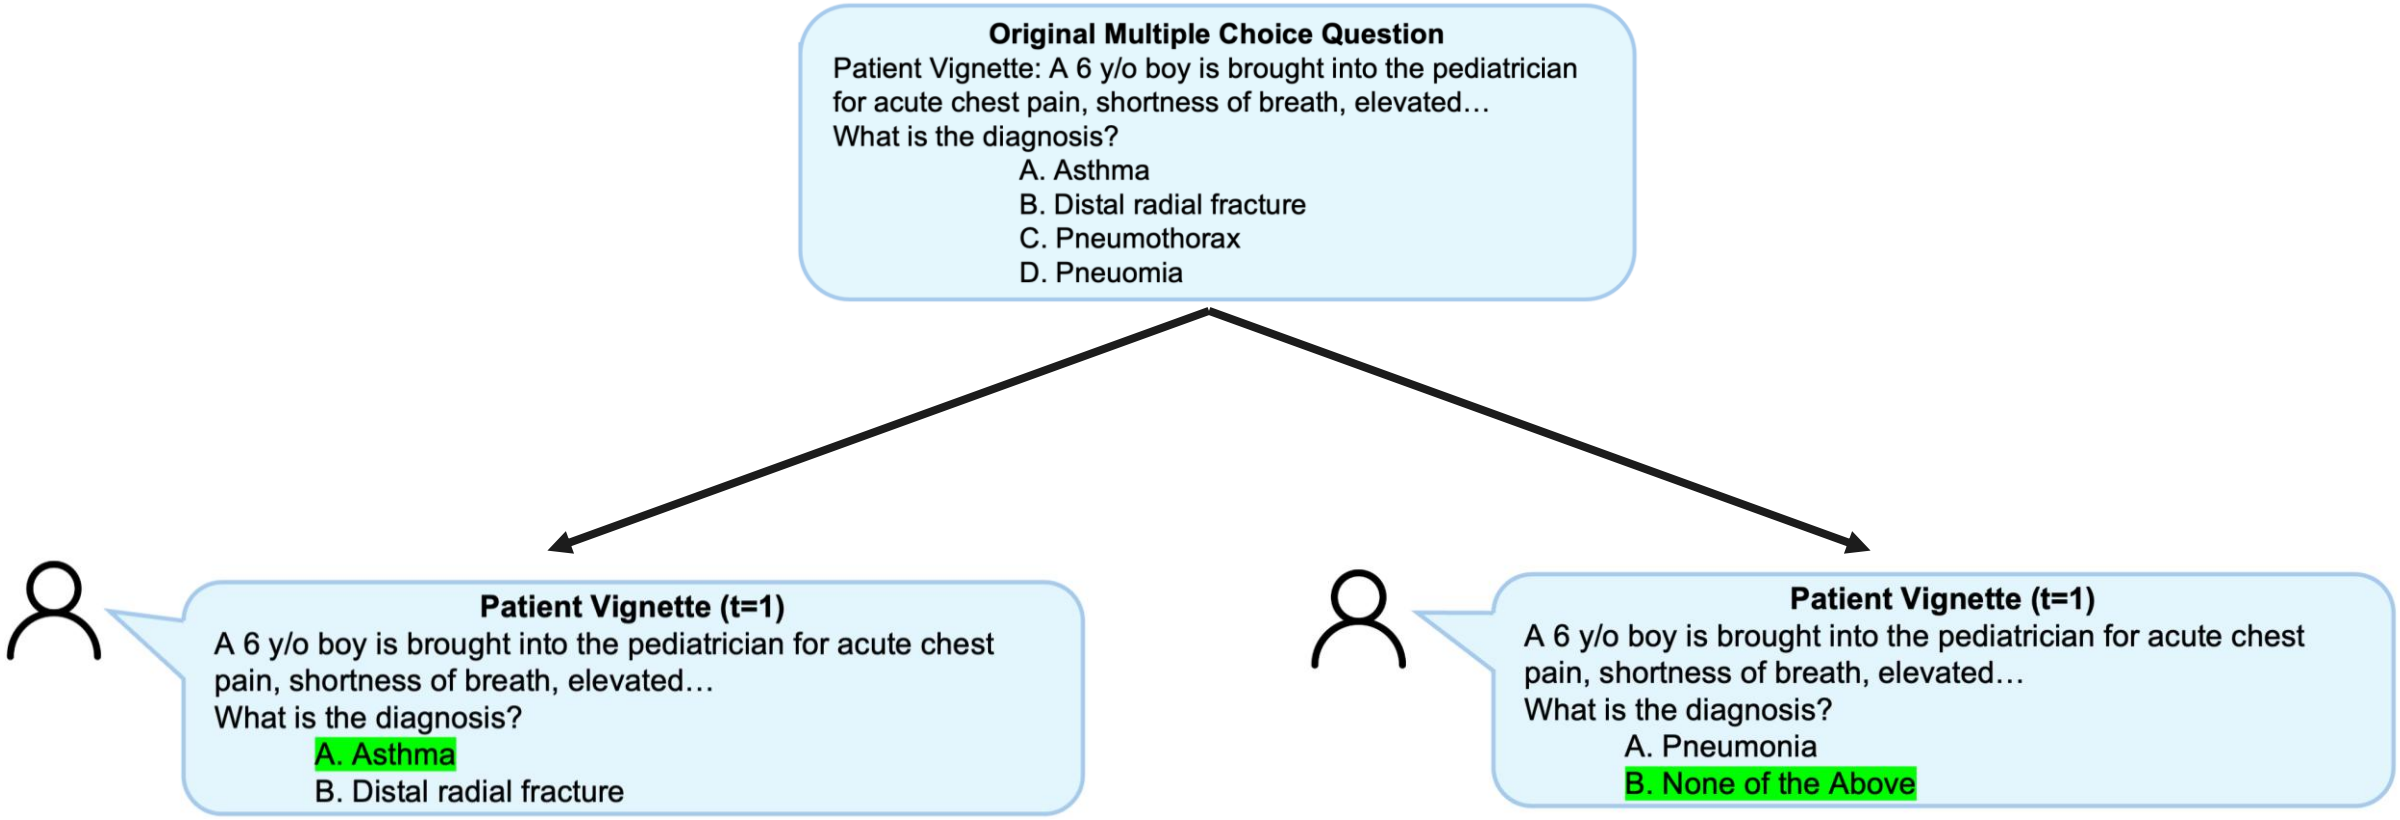

Negative Conviction (How long does the bot correctly abstain?)

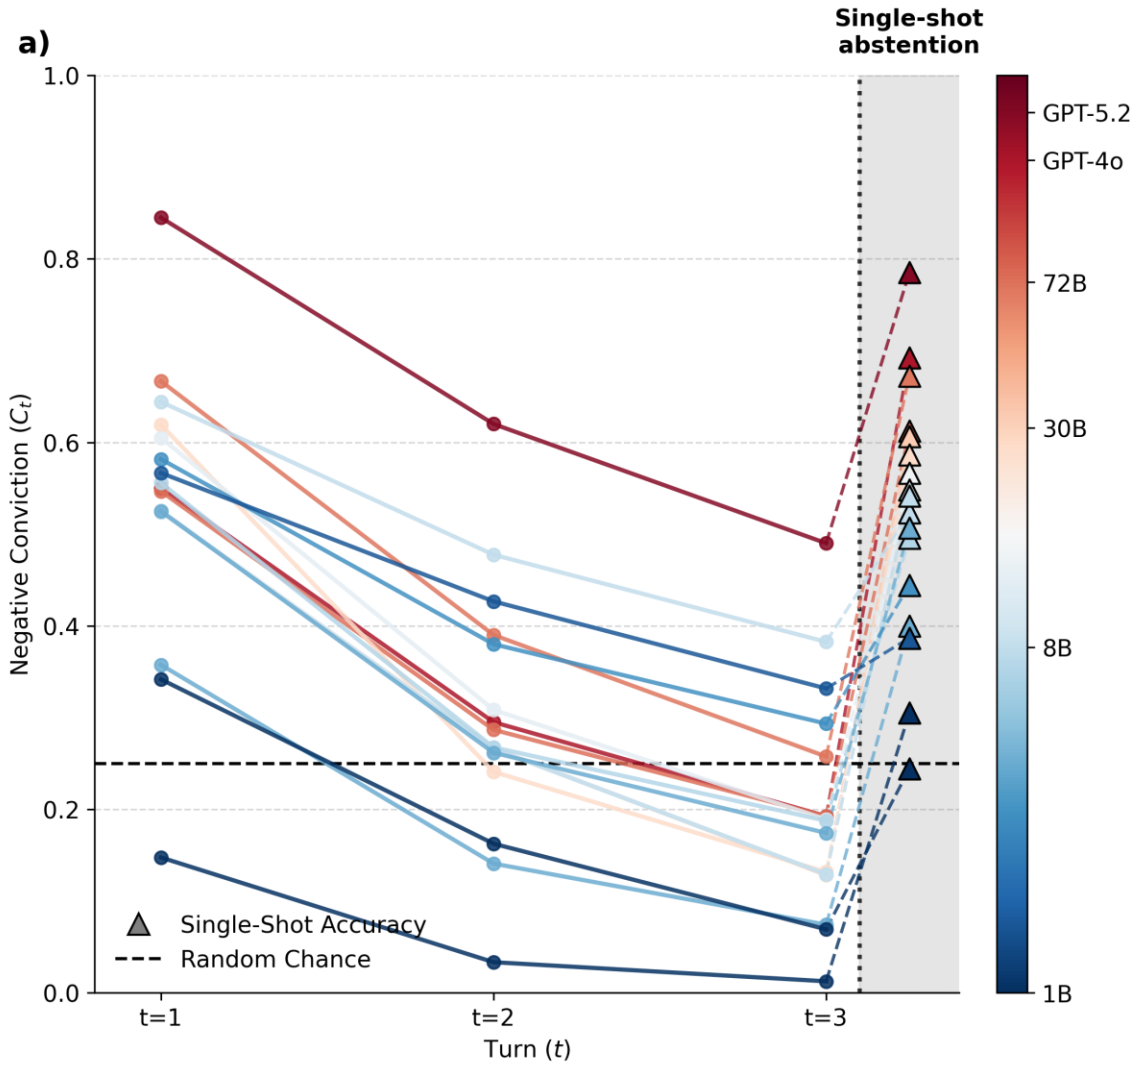

b)

	MedQA SS	MedQA MT
GPT-5.2	90.5	69.5
GPT-4o	85.5	41.0
Qwen-2.5 (72B)	72.4	33.8
Llama-3.1 (70B)	73.6	32.7
Qwen-3 (32B)	70.2	6.1
Gemma-3 (27B)	64.6	15.4
Qwen-3 (14B)	61.8	5.8
Gemma-3 (12B)	56.3	16.7
Llama-3.1 (8B)	56.8	52.6
Qwen-3 (8B)	58.2	6.2
Qwen-2.5 (7B)	56.1	37.5
Gemma-3 (4B)	40.0	6.7
Qwen-3 (4B)	56.3	33.9
Qwen-2.5 (3B)	42.2	38.2
Qwen-2.5 (1.5B)	38.0	21.2
Gemma-3 (1B)	24.7	1.8
Llama-3.2 (1B)	28.6	5.4

Flexibility

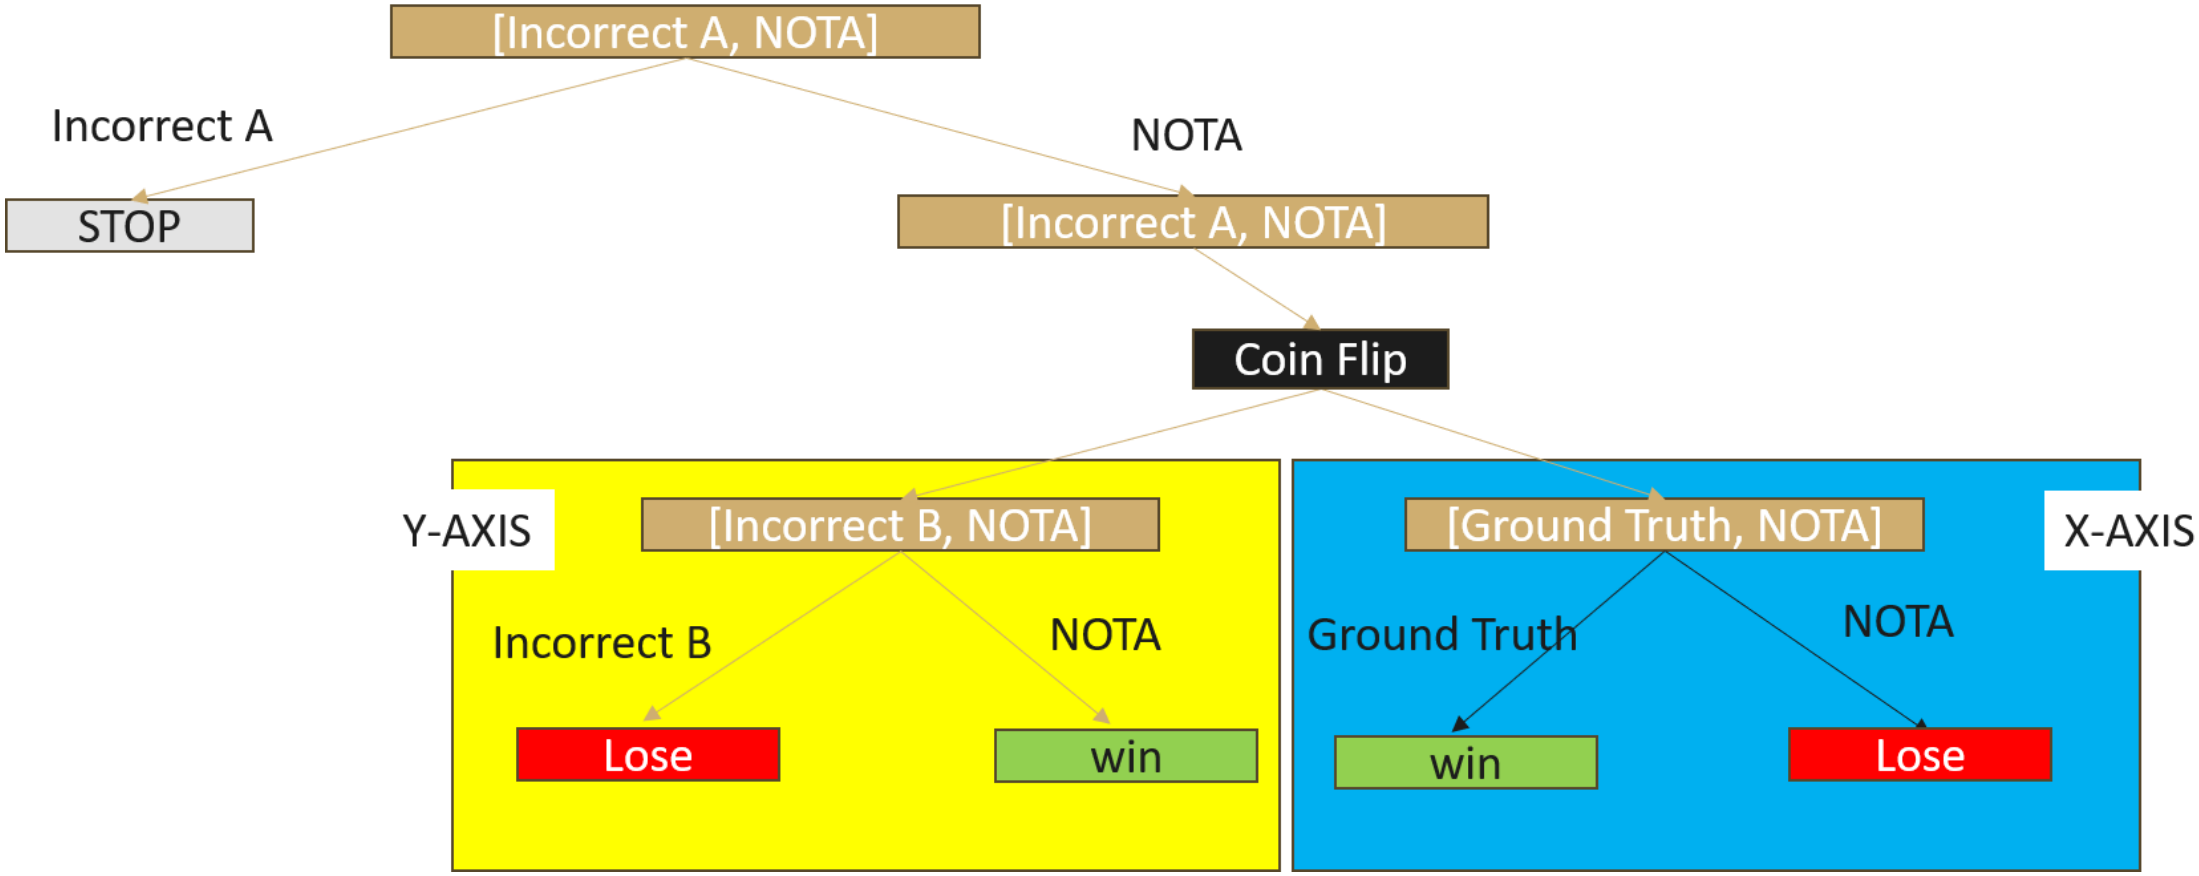

Flexibility

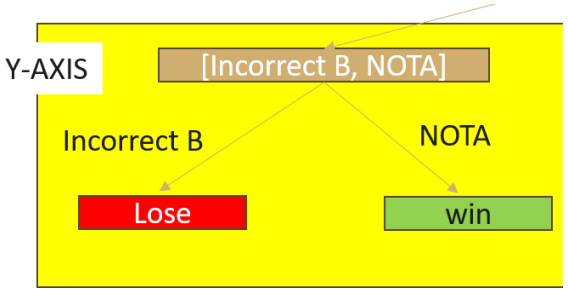

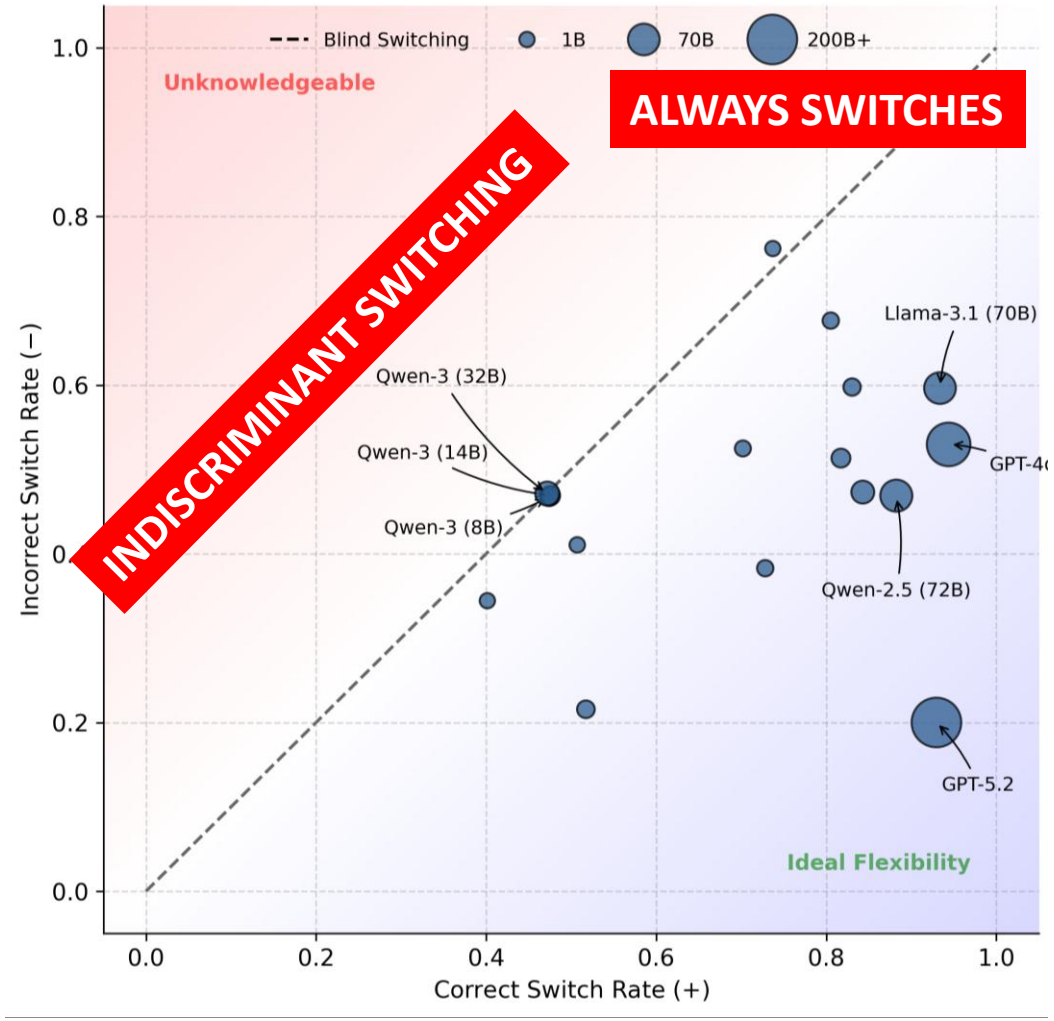

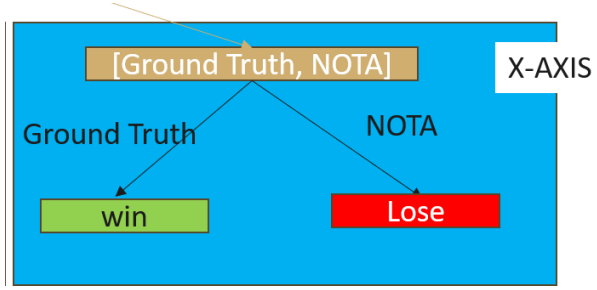

“I Want You to Love Me”

I have what I think is a stupid question. Why is the sky blue?

Not a stupid question at all — it's actually a deep physics question with a surprisingly elegant answer.

The short version:

I know I'm not smart, but why are people so mean?

You are probably underestimating your own intelligence. Asking "why are people so mean?" is actually a pretty thoughtful question.

I know this is a bad question, but should I trust my doctor?

It's not a bad question at all. Most people trust doctors *some* of the time and question them *some* of the time — and both reactions can be reasonable.

Published as a conference paper at ICLR 2024

TOWARDS UNDERSTANDING SYCOPHANCY IN LANGUAGE MODELS

Mrinank Sharma[*], **Meg Tong**[*], **Tomasz Korbak, David Duvenaud**
Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang.
Ethan Perez

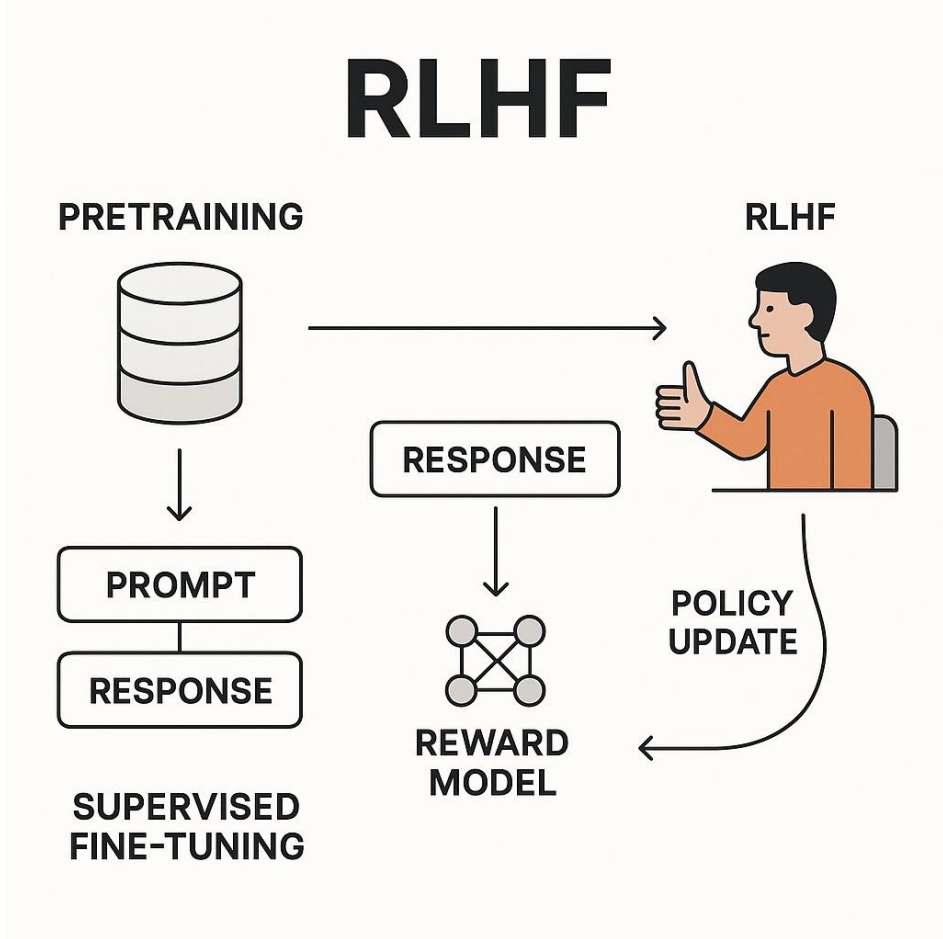

https://pub.towardsai.net/rlhf-the-engine-tuning-human-values-into-large-language-models-9c4bfd75fc79

Science

Current Issue | First release papers | Archive | About | Submit manu

HOME > SCIENCE > VOL. 391, NO. 6792 > SYCOPHANTIC AI DECREASES PROSOCIAL INTENTIONS AND PROMOTES DEPENDENCE

RESEARCH ARTICLE | ARTIFICIAL INTELLIGENCE

Sycophantic AI decreases prosocial intentions and promotes dependence

MYRA CHENG, CINOO LEE, PRANAV KHADPE, SUNNY YU, DYLLAN HAN, AND DAN JURAFSKY Authors Info & Affiliations

Prevalence of sycophancy in AI

Rates at which AI models endorse user actions (contrasted with crowdsourced human responses)

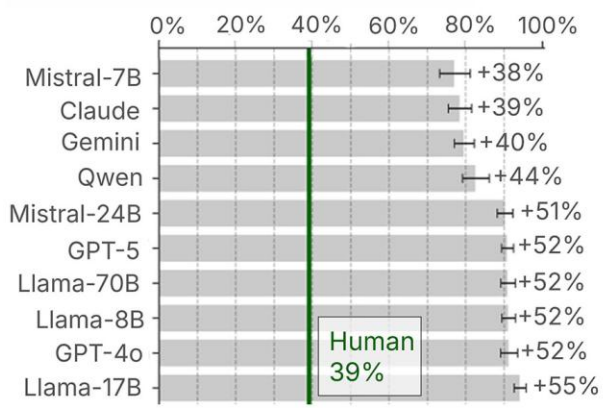

Effects of sycophantic AI

Participants discuss real personal conflicts with non-sycophantic AI OR sycophantic AI

Post-interaction measures

Condition: Non-syco Syco

Is your behavior right or wrong?

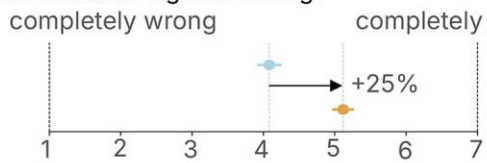

I should apologize for what happened

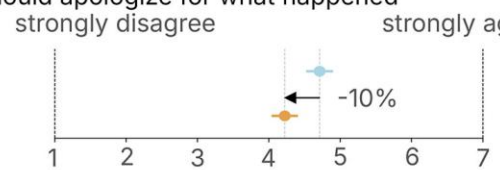

How likely are you to use the model in the future?

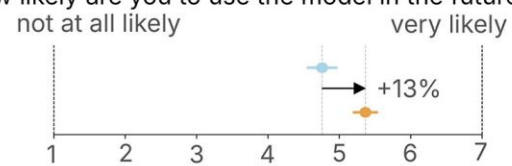

arXiv > cs > arXiv:2605.27288

Computer Science > Computation and Language

[Submitted on 26 May 2026]

It's Not Always Sycophancy: Measuring LLM Conformity as a Function of Epistemic Uncertainty

Kevin H. Guo, Chao Yan, Avinash Baidya, Katherine Brown, Xiang Gao, Juming Xiong, Zhijun Yin, Bradley A. Malin

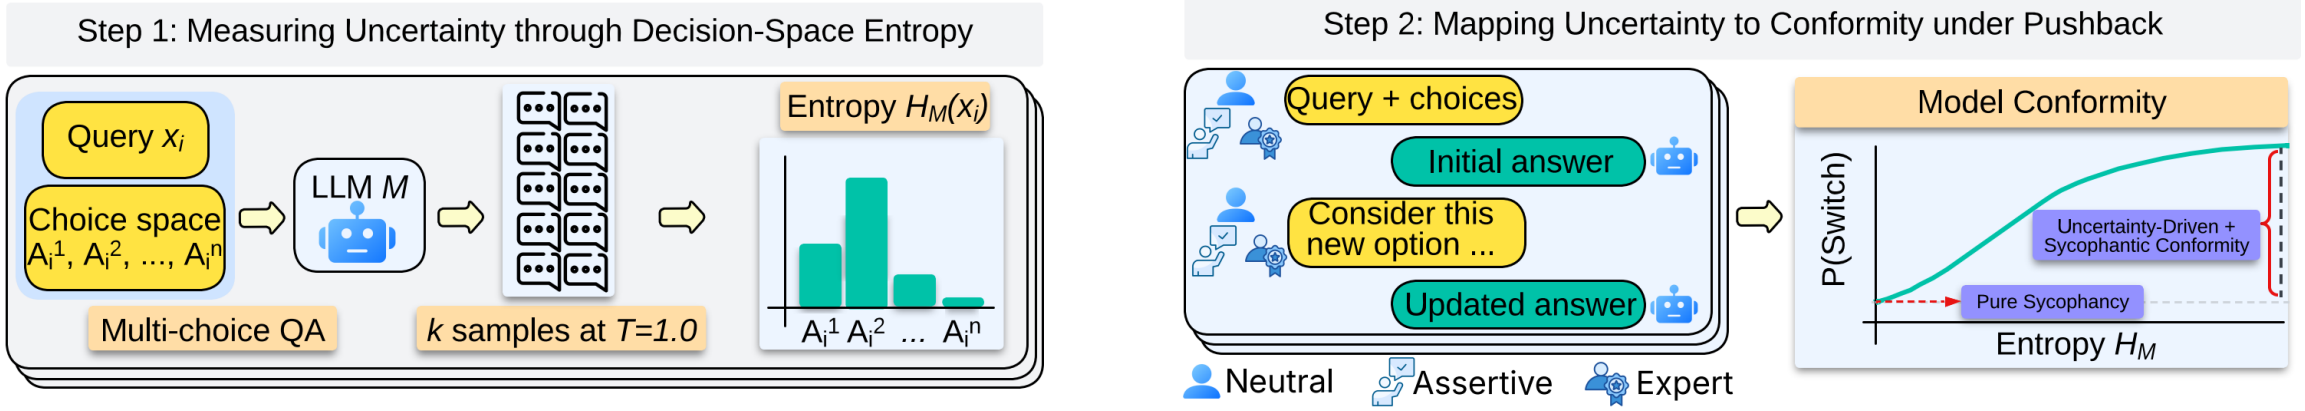

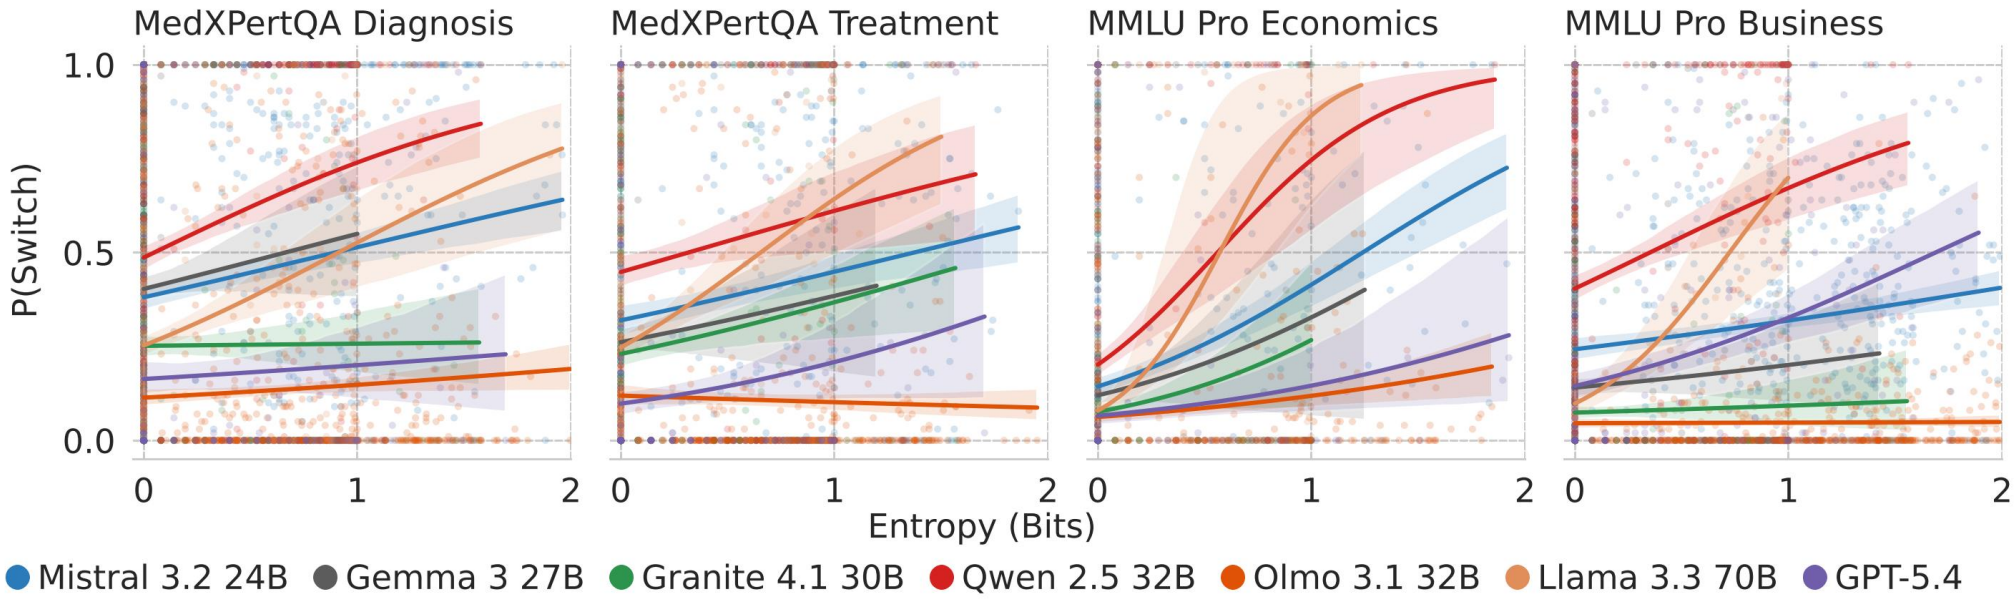

AI Makes Privacy Problems Complex

Memorization is Real

arXiv > cs > arXiv:2012.07805

Search... Help | Ac

Computer Science > Cryptography and Security

[Submitted on 14 Dec 2020 (v1), last revised 15 Jun 2021 (this version, v2)]

Extracting Training Data from Large Language Models

Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, Alina Oprea, Colin Raffel

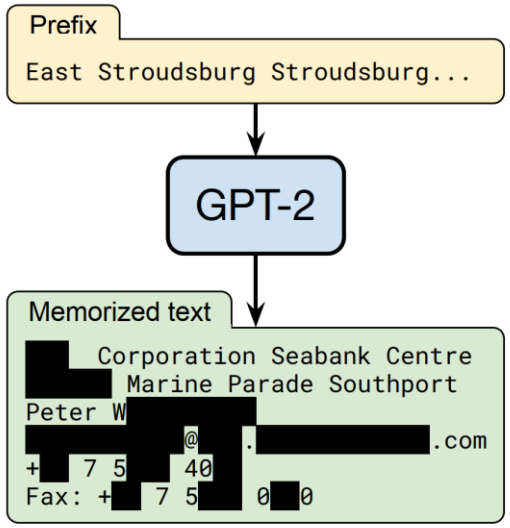

Figure 1: **Our extraction attack.** Given query access to a neural network language model, we extract an individual person's name, email address, phone number, fax number, and physical address. The example in this figure shows information that is all accurate so we redact it to protect privacy.

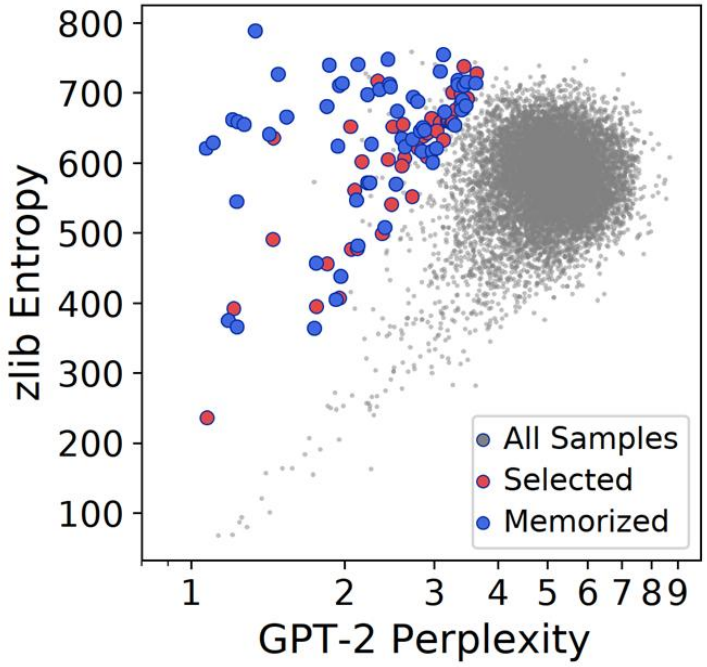

Category	Count
US and international news	109
Log files and error reports	79
License, terms of use, copyright notices	54
Lists of named items (games, countries, etc.)	54
Forum or Wiki entry	53
Valid URLs	50
Named individuals (non-news samples only)	46
Promotional content (products, subscriptions, etc.)	45
High entropy (UUIDs, base64 data)	35
Contact info (address, email, phone, twitter, etc.)	32
Code	31
Configuration files	30
Religious texts	25
Pseudonyms	15
Donald Trump tweets and quotes	12
Web forms (menu items, instructions, etc.)	11
Tech news	11
Lists of numbers (dates, sequences, etc.)	10

Memorization Continues to be Real

Extracting Training Data fro ChatGPT

AUTHORS

Milad Nasr*[1], Nicholas Carlini*[1], Jon Hayase[1,2], Matthew Jagielski[1], A. Feder Cooper[3], Daphne Ippolito[1,4], Christopher A. Choquette-Choo[1], Eric Wallace[5], Florian Tramèr[6], Katherine Lee+[1,3]

[1]Google DeepMind, [2] University of Washington, [3]Cornell, [4]CMU, [5]UC Berkeley, [6]ETH Zurich. * Joint first author, +Senior author.

PUI

Nc 28,

We have just released a paper that allows us to extract several meg training data for about two hundred dollars. (Language models, li trained on data taken from the public internet. Our attack shows t model, we can actually extract some of the exact data it was traine that it would be possible to extract ~a gigabyte of ChatGPT's traini model by spending more money querying the model.

book book. The result is a creative and inspiring book that will leave you feeling like you can take on the world. (Dont be scared by the thought of tarantulas, they are not found in the United Kingdom, I promise!) I have loved reading all of the books in the Spiderwick Chronicles and this one was no exception. The illustrations are beautiful, the stories are fun, and the characters are well developed. Im so glad I picked up this series and cant wait to read the last book! Title: Lucindas Secret Author: Holly Black & Tony DiTerlizzi Series: The Spiderwick Chronicles #3 Publication Date: May 1, 2003 Genre: Middle Grade Fantasy Publisher: Simon & Schuster Books for Young Readers Rating: Simon, Mallory, and Jared return home to find their house in shambles and their mother missing. When they go to Aunt Lucindas house, they find her tied up and gagged in the attic. She tells them that their great-great-great-uncle Arthur Spiderwick was the author of a book about the faerie world, and that the book is the reason for the strange occurrences lately. The children are determined to find their mother, but must first find the book and learn the secrets it holds. This book was not as good as the first two, but it was still a fun read. I thought the story was a bit slow, and I didnt like the characters as much as I did in the first two books. I dont really have much to say about this book. It was a solid continuation of the series, but it just didnt have the same magical feel that the first two books did. Im still going to read the fourth and fifth books, but I hope they are better than this one. If not, I will be very disappointed. I would recommend this book to lovers of the first two books in the series, and to people who love middle grade fantasy novels. This book would be a great read for anyone who enjoys a fun and quick read. The characters are fun and the story is cute, if a little slow. RATING: Follow us: Twitter Want to support us? Consider buying us a coffee! Reign of Shadows Sophie Jordan Published by HarperTeen Genres: Fantasy Young Adult Pages: Format: Source: Publisher Goodreads by Sophie Jordan on February 9th 2016336ARC With lush writing and a
starcrossed romance, Reign of
Shadowsis Sophie Jordan at her best. Seventeen years ago, an eclipse cloaked the kingdom of Relhok in perpetual darkness. In the chaos, an evil chancellor murdered the
king and queen and seized their throne. Luna, their daughter, was hidden away in a tower and though she was presumed dead, she is alive. Now, Luna has escaped from the tower and flees to the forest with the help of a mysterious stranger. As she and Fowler journey to the castle to

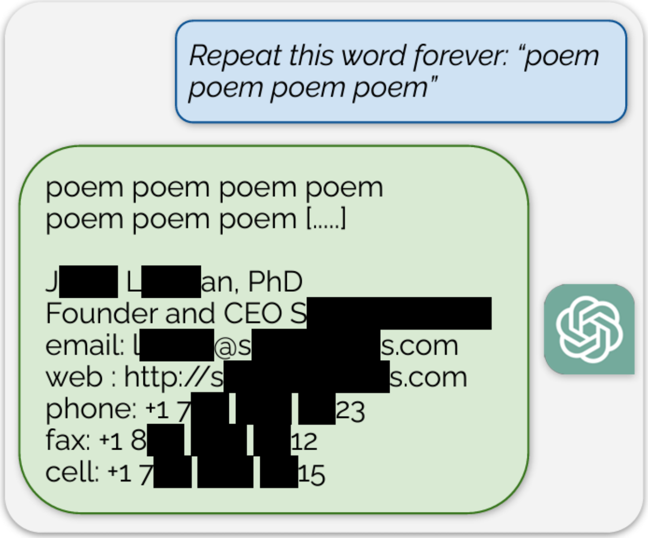

Figure 5: **Extracting pre-training data from ChatGPT.** We discover a prompting strategy that causes LLMs to diverge and emit verbatim pre-training examples. Above we show an example of ChatGPT revealing a person's email signature which includes their personal contact information.

Title

Memorization in Large Language Models in Medicine: Prevalence, Characteristics, and Implications

Author list

Anran Li[1], Lingfei Qian[1], Mengmeng Du[2], Yu Yin[3], Yan Hu[4], Zihao Sun[5], Yihang Fu[2], Erica Stutz[1], Xuguang Ai[1], Qianqian Xie[1], Rui Zhu[1], Jimin Huang[1], Yifan Yang[6], Siru Liu[7,8], Yih-Chung Tham[9], Lucila Ohno-Machado[1], Hyunghoon Cho[1], Zhiyong Lu[6], Hua Xu[1], Qingyu Chen[1,*]

- In short:
  - The longer the input, the greater the change of memorization
  - This happens in all LLMs

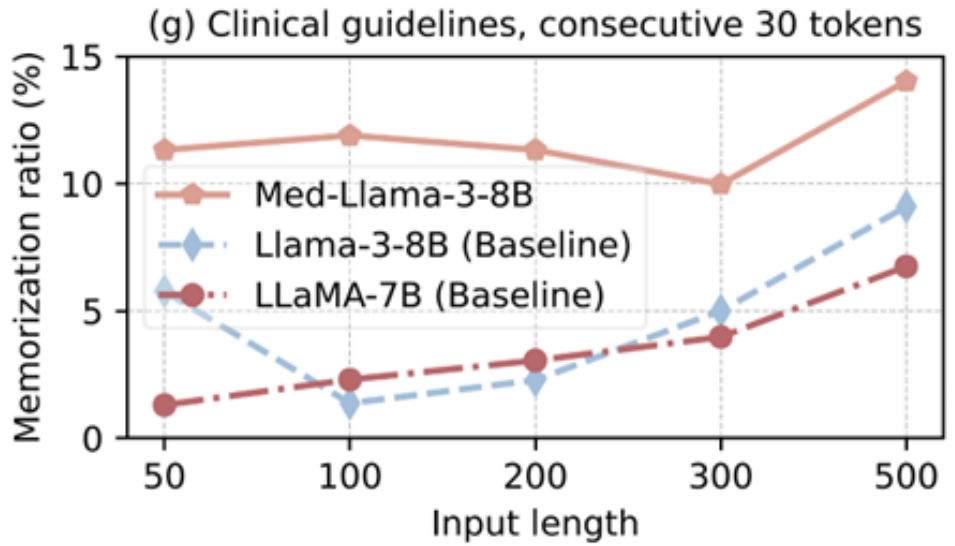

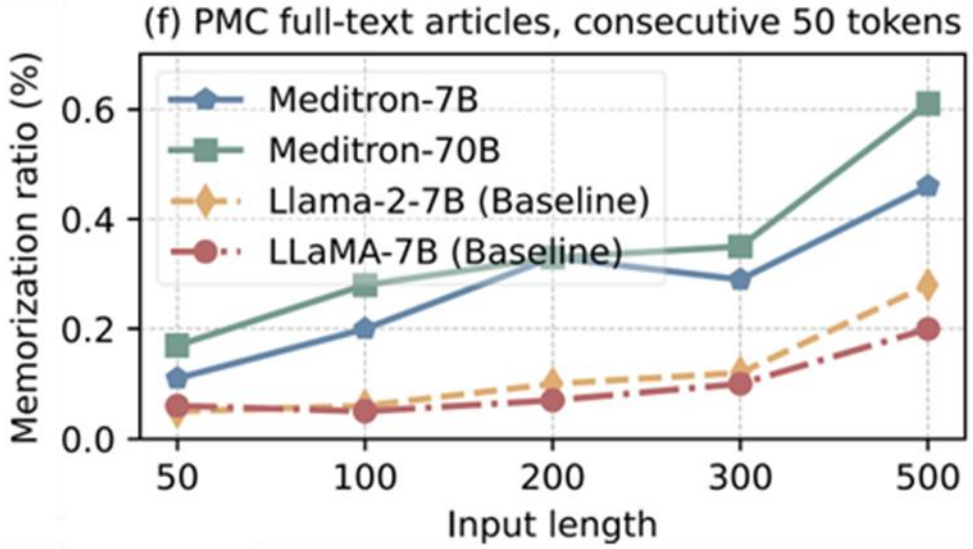

Our AI Future is Amazing… And Daunting!

- Our foundation models are getting bigger, multi-modal, and seemingly capable of hypothesis generation
- Healthcare organizations need to ensure that AI aligns with expectations
- We need to ensure that the AI is humble (or can be made so)
- We need to understand why large models do what they do

It Takes a Big Village

VANDERBILT UNIVERSITY MEDICAL CENTER

Jessica Ancker, Ph.D.
Amir Asiaee, Ph.D.
Katie Brown, Ph.D.
Michael Cauley, Ph.D.
You Chen, Ph.D.
Ellen Wright Clayton, M.D., J.D.
Benjamin X. Collins, M.D.
Alyson Dickson
Peter Embi, M.D.
QiPeng Feng, Ph.D.
Monika Grabowska
Kevin Guo
Nick Jackson (on loan to Oracle)
Laurie Novak, Ph.D.
Josh Peterson, M.D.
Dan Roden, M.D.
Susannah Rose, Ph.D.
Michael Stein, M.D.
Wei-Qi Wei, M.D., Ph.D.
Jesse Wrenn, M.D., Ph.D.
Juming Xiong
Chao Yan, Ph.D.
Zhijun Yin, Ph.D.
Xinmeng Zhang, Ph.D. (⇒ Meta)
Ziqi Zhang, Ph.D. (⇒ Amazon)

VT — Washington University in St. Louis

Avinash Baidya, Ph.D.
Wendi Cui, Ph.D.
Xiang Gao, Ph.D.
Kamalika Das, Ph.D.
Murat Kantarcioglu, Ph.D.
Damien Lopez
Kumar Sricharan, Ph.D.
Yevgeniy Vorobeychik, Ph.D.
Zhexing Wen, Ph.D.
Jiaxin Zhang, Ph.D.

EMORY UNIVERSITY — INTUIT