

Dear researchers: Is AI all you've got?

Austin Z. Henley

Carnegie Mellon University, azhenley@cmu.edu

What important problems have software engineering researchers abandoned so that they could focus on AI instead?

In the early 1900s, bacteriophage therapy was a popular topic in medical and biological research for its potential to target specific bacteria [1]. However, after the discovery of penicillin in 1928, the research community shifted its attention to antibiotics, mostly abandoning bacteriophage research by the 1950s [2]. Interest surged again in the 2000s as antibiotic resistance grew more common and the highly targeted nature of phage therapy became more appealing [3].

Is a similar phenomenon happening in the software engineering community? Are researchers over-indexing on AI? At ICSE '25, one-third of research-track papers and two-thirds of industry-track papers involved AI [4]. In contrast, a study of research trends from 1992 to 2016 in software engineering did not even include AI or ML in the top 10 most popular topics [5]. Furthermore, AI conferences have grown exponentially, with AAAI going from 9000 submissions in 2022 to 29,000 submissions for the 2026 conference [6].

1. What might be missed?

My concern is not that AI doesn't have considerable value (I have worked on AI in both academia and industry) or that it isn't a huge innovation for the field, but rather that other topics are important too. What innovations might we miss? What problems will continue to go unsolved? What daily pain points exist beyond AI?

As software engineers build highly distributed systems, architectural and implementation decisions play a critical role in system resilience. In 2025, we saw several significant outages that had cascading consequences across

the globe. In June 2025, Google Cloud had an outage that impacted Google products for millions of users as well as Spotify, Discord, and ChatGPT [7]. In October 2025, AWS had an outage that impacted Zoom, Venmo, Reddit, Snapchat, Amazon, and banks around the world [8]. Only a few days later, Azure had an outage that affected various Microsoft services, retail stores, and airline systems [9]. In November and December 2025, Cloudflare had outages that caused disruptions to ChatGPT, LinkedIn, Shopify, Square, and many others [10, 11].

Given that approximately 20% of all websites use Cloudflare [12] and that AWS, Azure, and GCP account for 62% of the global cloud infrastructure [13], this is of huge concern. The AWS outage was estimated to have cost businesses upwards of 581 million dollars in just 15 hours [14] and one of the Cloudflare outages to have cost 300 million dollars in less than 4 hours [15]. In fact, I argue that it is a global catastrophe waiting to happen and we have seen repeated warning signs, yet I do not see researchers flocking to solve the problem as if lives depend on it. Has AI dramatically increased the reliability of software?

From my experience in industry, most recently at a startup that went through an acquisition, the problems we faced were fundamentally human problems. For example, convincing business partners, customers, and lawyers that our software does what we say that it does and that it won't have major disruptions was a huge challenge that we spent considerable time on. AI allowed us to scale faster, but it did not help us convince others to trust our system.

Looking over the keynote talks from ICSE and FSE leading up to the release of ChatGPT in late 2023, we can remind ourselves of the topics that were once top of mind for researchers: software safety, reliability, run-time monitoring, testing, industry impact, software engineering education, research rigor and reproducibility, socio-technical coordination, ethics and privacy, environmental impacts, etc. Several of these talks are relevant to the global cloud outages we have been experiencing, such as Marsha Chechik's keynote at FSE 2022 on the safety and reliability of software [16]. Unfortunately, it seems that her call to action was not enough.

2. Follow the incentives

But perhaps there are innovations out there that we aren't even thinking about. Clayton Christensen, in his seminal book, *The Innovator's Dilemma* [17],

argues that as fields mature, they often optimize and sustain existing innovations while also overlooking novel, disruptive innovations. He goes on to state that disruption often begins in places that others dismiss as too small, too immature, or too orthogonal, which are the exact areas that risk being neglected when AI becomes the default answer to every research question.

A straightforward explanation for why researchers would jump on the AI bandwagon lies in incentives. There has been an incredibly strong force pulling everyone into AI, including funding agencies asking for AI-related proposals, conferences adding multiple topics in AI to their calls for papers, and universities posting numerous faculty openings for AI researchers. It makes it incredibly difficult for a researcher to avoid. In fact, the papers I co-author involving AI are being cited at a far, far higher rate than any of my other work. As Zimmermann observes, academic behavior is often shaped less by intrinsic motivation than by the incentive structures and metrics that govern how academics are evaluated [18]. In other words, “incentives, incentives, and incentives”.

I’m certainly not the first to voice concerns about the fixation on AI, in fact, several others have argued about the potential downsides of AI. Johnson and Menzies stated that “AI overhype” is dangerous and that it is “the ethical duty of software professionals to rally against such remarks” [19]. Early studies of AI use in the classroom have shown evidence that it may disrupt learning (e.g., [20, 21, 22]). This is further complicated by the disruption in the job market for software engineers, especially those that are early in their careers, caused by AI [23].

3. Right now is the opportunity

Given that virtually every researcher is focused on AI, now is the opportunity to make progress on any other topic. Even Yan LeCun, the 2018 Turing Award winner for his foundational work on deep neural networks, said that “an LLM is basically an off-ramp, a distraction, a dead end” [24], in regards to his belief that LLMs will not continue to scale, and “don’t work on LLMs” [25]. He also warned, “right now, they are sucking the air out of the room anywhere they go, and so there’s basically no resources for anything else. And so for the next revolution, we need to take a step back and figure out what’s missing from the current approaches” [26].

Perhaps we can use machine learning as inspiration to solve this concern of over-indexing on AI. Machine learning algorithms have mechanisms

for escaping local maxima. Simulated annealing and stochastic gradient descent introduce noise, evolutionary algorithms create variation with random starting points and mutations, and optimization methods use random perturbations and restarts to explore new areas of a search space.

Can researchers apply these concepts to how they decide their next research topics? I worry that researchers are stuck in a local maximum. Worse yet, they are intentionally reducing the randomness in their exploration of research topics with the goal of joining (or staying on) the AI bandwagon.

While everyone is researching AI, the software systems that quietly run our society continue to fail in ways that are neither rare nor surprising. If the global cloud can go down repeatedly, affecting banks, airlines, retailers, and governments, then resilience and trust are not solved problems. The unfortunate part is not that these problems are hard, but rather that they increasingly feel unfashionable.

In the 2021 book *Think Again*, organizational psychologist Adam Grant claimed, “Thinking like a scientist involves more than just reacting with an open mind. It means being actively open-minded. It requires searching for reasons why we might be wrong, not for reasons why we must be right, and revising our views based on what we learn.” [27] Whether it is through this scientific approach or through random perturbations, we can’t get stuck in local maxima.

So software and systems researchers, is AI all you’ve got?

4. Acknowledgements

I would like to thank Olaf Zimmermann for the many rounds of feedback on this article.

References

- [1] G. P. C. Salmond, P. C. Fineran, A century of the phage: past, present and future, *Nature Reviews Microbiology* 13 (12) (2015) 777–786.
- [2] X. Wittebole, S. D. Roock, S. M. Opal, A historical overview of bacteriophage therapy as an alternative to antibiotics for the treatment of bacterial pathogens, *Virulence* 5 (1) (2014) 226–235. doi:10.4161/viru.25991. URL <https://doi.org/10.4161/viru.25991>

- [3] F. L. G. Altamirano, J. J. Barr, Phage therapy in the postantibiotic era, *Clinical Microbiology Reviews* 32 (2) (2019) 10.1128/cmr.00066-18. arXiv:<https://journals.asm.org/doi/pdf/10.1128/cmr.00066-18>, doi:10.1128/cmr.00066-18.
URL <https://journals.asm.org/doi/abs/10.1128/cmr.00066-18>
- [4] Icse 2025 research track - event overview (2025).
URL <https://conf.researchr.org/track/icse-2025/icse-2025-research-track>
- [5] G. Mathew, A. Agrawal, T. Menzies, Finding trends in software research, *IEEE Transactions on Software Engineering* 49 (4) (2023) 1397–1410. doi:10.1109/TSE.2018.2870388.
- [6] Aaai-26 review process update (2026).
URL <https://aaai.org/conference/aaai/aaai-26/review-process-update/>
- [7] Google cloud down for thousands of users, downdetector shows (2025).
URL <https://www.reuters.com/business/google-cloud-down-thousands-users-downdetector-shows-2025-06-12/>
- [8] Amazon’s cloud unit reports outage, several websites down (2025).
URL <https://www.reuters.com/business/retail-consumer/amazon-cloud-unit-reports-outage-several-websites-down-2025-10-20/>
- [9] Microsoft hit with azure, 365 outage ahead of quarterly earnings (2025).
URL <https://www.cnn.com/2025/10/29/microsoft-hit-with-azure-365-outage-ahead-of-quarterly-earnings.html>
- [10] Cloudflare restores services after outage impacts thousands of internet users (2025).
URL <https://www.reuters.com/business/elon-musks-x-down-thousands-us-users-downdetector-shows-2025-11-18/>
- [11] Cloudflare is investigating issues with dashboard and related apps (2025).
URL <https://www.cnn.com/2025/12/05/cloudflare-is-investigating-issues-with-dashboard-and-related-apps.html>

- [12] Historical trends in the usage statistics of reverse proxy services for websites (2026).
URL https://w3techs.com/technologies/history_overview/proxy/all
- [13] Cloud market share q2 2025: Microsoft dips, aws still kingpin (2025).
URL https://www.crn.com/news/cloud/2025/cloud-market-share-q2-2025-microsoft-dips-aws-still-kingpin?utm_source=chatgpt.com
- [14] Cybercube shares insurance loss estimate for aws “amazonk” outage (2025).
URL <https://www.cybcube.com/news/insurance-loss-estimate-for-aws-amazonk-outage>
- [15] Cloudflare’s outage: Another wake-up call for cloud resilience (2025).
URL <https://www.forrester.com/blogs/cloudflares-outage-another-wake-up-call-for-cloud-resilience/>
- [16] On safety assurance and reliability: A software engineering perspective (keynote) (2022).
URL <https://2022.esec-fse.org/details/fse-2022-keynotes/1/On-Safety-Assurance-and-Reliability-A-Software-Engineering-Perspective-Keynote->
- [17] C. M. Christensen, The Innovator’s Dilemma: When New Technologies Cause Great Firms to Fail, Harvard Business Review Press, 2015.
- [18] O. Zimmermann, Overcoming the research-practice gap: Root cause analysis and topics of practical relevance in software architecture and distributed systems, *Journal of Systems and Software* 230 (2025) 112539. doi:<https://doi.org/10.1016/j.jss.2025.112539>.
URL <https://www.sciencedirect.com/science/article/pii/S0164121225002080>
- [19] B. Johnson, T. Menzies, Ai over-hype: A dangerous threat (and how to fix it), *IEEE Software* 41 (6) (2024) 131–138. doi:[10.1109/MS.2024.3439138](https://doi.org/10.1109/MS.2024.3439138).

- [20] M. Kazemitabaar, J. Williams, I. Drosos, T. Grossman, A. Z. Henley, C. Negreanu, A. Sarkar, Improving steering and verification in ai-assisted data analysis with interactive task decomposition, in: Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology, UIST '24, Association for Computing Machinery, New York, NY, USA, 2024. doi:10.1145/3654777.3676345. URL <https://doi.org/10.1145/3654777.3676345>
- [21] H. Bastani, O. Bastani, A. Sungu, H. Ge, Ö. Kabakcı, R. Mariman, Generative ai can harm learning, The Wharton School Research Paper (2024).
- [22] J. Prather, B. N. Reeves, J. Leinonen, S. MacNeil, A. S. Randrianasolo, B. A. Becker, B. Kimmel, J. Wright, B. Briggs, The widening gap: The benefits and harms of generative ai for novice programmers, in: Proceedings of the 2024 ACM Conference on International Computing Education Research - Volume 1, ICER '24, Association for Computing Machinery, New York, NY, USA, 2024, p. 469–486. doi:10.1145/3632620.3671116. URL <https://doi.org/10.1145/3632620.3671116>
- [23] How badly is ai cutting early-career employment? (2025). URL <https://spectrum.ieee.org/ai-impact-on-job-market>
- [24] 'imagine a cube floating in the air': The new ai dream allegedly driving yann lecun away from meta (2025). URL <https://gizmodo.com/yann-lecun-world-models-2000685265>
- [25] Yann lecun, pioneer of ai, thinks today's llm's are nearly obsolete (2025). URL <https://www.newsweek.com/nw-ai/ai-impact-interview-yann-lecun-artificial-intelligence-2054237>
- [26] Meta ai: Yann lecun says llms won't get us to true intelligence, pushes "world models" (2025). URL <https://www.businessinsider.com/meta-ai-yann-lecun-llm-world-model-intelligence-criticism-2025-11>
- [27] A. Grant, Think Again: The Power of Knowing What You Don't Know, Penguin, 2023.