

Conclusion

The years during which this dissertation was completed (2022-2026) coincided with a series of profound changes in academia and research. LLMs were born, and then quickly expanded to dominate every aspect of public, personal, and professional lives. Unprecedented political attacks threatened research funding, independence of academic institutions, and students' abilities to conduct research. Industry labs emerged that hoovered up leading researchers away from universities unlike anytime before. It seems fitting to conclude on some personal thoughts from all this, from the incredible promise that academia holds and perils it faces, both internal and external. While largely relevant to my field of computer science, some may hold more broadly.

Modern academia pioneered the process through which scholars felt emboldened to work on diverse, open-ended, and often “impractical” research problems. It cultivated the peer-review process through which that research was filtered, credentialed, and publicly disseminated, allowing others to build on promising directions and sidestep the pitfalls. It developed the modern Ph.D. degree where bright young minds could find the space, time, and resources to explore novel ideas with guidance and close mentorship. Though imperfect in many ways, it is a phenomenal accomplishment that created and democratized research to a wider populace. These public and social benefits led to the establishment of steady government grants to fund it. This success also led to the emergence of industry labs shaped in a similar mold, from Bell Labs in the latter half of the 20th century, to early instantiations of Google and Microsoft Research in the early internet age. While they were naturally more restrained than academia, by-and-large, they still valued open-ended research and encouraged their members to publicly disseminate their work and engage with peer-review and criticisms. Much of the fundamental research that led to the development of our modern technological society, including LLMs, arose from this foundation; it involved work from both academia and industry labs building upon each other, published openly and scrutinized by their peers.

This classical academic and research environment now faces a string of interconnected challenges: (1) the unsustainability of peer review, (2) uncertain public funding of academia, and (3) the exodus of leading researchers away from academia. The peer review process that underpins the modern research environment is deeply overburdened in computer science. Prestigious venues such as NeurIPS and ICML have seen year-over-year exponential growth for almost a decade, with NeurIPS 2025 receiving over 21,000 submissions; unofficial estimates for NeurIPS 2026 place it well over 30,000 submissions. On the supply side, the expansion of Computer Science as a discipline and its industrial utility has led to “# of papers” at such venues emerging as a simple and easy metric, in both academia and industry. It has now become expected that incoming CS Ph.D. students at top institutions have some publication at such venues. On the other hand, the quantity and quality of reviewers is simply limited. In many machine learning adjacent venues, it is rare to see senior faculty members or PIs reviewing papers: that burden instead falls on graduate students and post-docs, many of whom are new to the area or to research itself. This imbalance has only been exacerbated by the dominance of LLMs, which are extensively used on both sides of the review process. As a result, the modal outcome in Computer Science increasingly threatens to be that we publish more papers than we can read or consume.

More broadly, academia is also facing significant external headwinds. The United States, long the largest contributor and promoter of scientific research, has seen NSF and NIH funding significantly slashed in 2025. The funding that does exist, in the US and beyond, is increasingly targeted and narrowed to specific applied

domains, threatening to upend the premise of academia as a generator of broad and general knowledge. Public trust in academia and its mandate is also dwindling. The causes of these headwinds are many, from inefficient administrative spending by universities, to the rise of a political climate that is increasingly hostile to academia. The end result, however, is a far more hostile environment to conduct research and build a career around it, especially for junior scholars.

Addressing both these challenges requires significant leadership from those within. Computer Science as a discipline needs to re-evaluate what publications mean and how we value them, how to re-imagine research in the age of LLMs, and how to better communicate the broader value of our work to society. But it is precisely here where the most striking headwinds lie. We are seeing many leading academicians, who are the natural leaders to take the mantle and address such issues, transition away from academia into frontier labs, or to launch AI themed start-ups. While this may seem just a larger version of past cycles of academia-to-industry flights, it is more pronounced. First, much of the research work done by many frontier labs is not published. What little is released is never peer-reviewed, stymieing natural critiques and making it difficult to reproduce or build on. So while the developments made by these labs are incredible, the broader benefits to the research community are much more limited. Second, the overall focus of these labs is much narrower: increase AI capabilities to produce a better product offering. They are happy to cede much of the intellectual ground of discovery and exploration to AI (as it strengthens their product story) and instead focus on enhancing their capabilities. In this sense, it is less about exploring the unknown and more about tooling AI to do so. This cannot cover the same intellectual breadth as before (even if the depth is more). Both these stances are completely natural and legitimate from the perspectives of an AI company focused on providing returns to their shareholders¹⁶: but the by-product is an intellectual monoculture where the advances are shaped and understood by a few, with little room to engage in criticisms.

While many leading academics join such enterprises under the noble goal of “AI Safety”, this has become a catch-all buzzword whose meaning is increasingly muddled. Addressing some negative aspects of AI is the interest of these firms: rogue AI models, unintended usage by nefarious parties, etc. But many others are not: the decline of cognitive abilities due to mass adoption of AI, labour market impact, concentration of AI capabilities in the hands of a few, and so on. It is also unclear whether the closed research culture of such firms (and its resulting implications) constitute an ideal environment toward resolving societally important questions around AI Safety. On the flip-side, this exodus deprives us of the talent, skill, and leadership needed to address the challenges faced by academia and the broader enterprise of peer-reviewed research. The next generation of scholars may not find an environment as conducive to exploring their interests as those before them; they may not receive adequate feedback, guidance, and mentorship. Public funding and support of research and academia may continue dwindling and never recover if the best continue leaving.

I have no doubts that academia and peer-reviewed research as an institution will survive. But we are at risk of it being performative and hollow, a Potemkin village. Perhaps this shift is inevitable, and the old perspectives of research and academia no longer have a place in the new world dominated by AI. But at the least, we should be honest about the societal and intellectual ramifications of it. And if that seems unpalatable and we are in a position to improve it, we should be emboldened to do so.

*Come my friends;
Tis not too late to seek a newer world*

— Alfred, Lord Tennyson, *Ulysses*

¹⁶The irony of lamenting firms holding private information and strategically revealing it in a thesis that is about profiting from information asymmetry is not lost on me.