

The *Interactions* website (interactions.acm.org) hosts a stable of bloggers who share insights and observations on HCI, often challenging current practices. Each issue we'll publish selected posts from some of the leading and emerging voices in the field.

What Happens When CHI Reaches 30,000 Submissions in 2030?

Preparing for the Tsunami of AI-Accelerated Papers

Johannes Schöning, University of St. Gallen

An irony is unfolding in computer science. The AI communities that built large language models are now struggling to cope with what those models have unleashed: a rapid, accelerating increase in paper submissions that threatens to overwhelm the peer review systems on which research quality depends.

In her recent keynote at AlpCHI 2026, Elisabeth André from the University of Augsburg gave a striking analogy that might apply to those of us interested in HCI. She highlighted that it is worth looking at AAAI, a flagship AI conference. In 2023, AAAI received around 9,000 submissions. By 2025, that figure had climbed to roughly 13,000. Then, in a single year, it nearly doubled: AAAI 2026 received close to 29,000 submissions to its main technical track, of which about 23,000 entered the review process. More than 75,000 unique authors were listed. The organizers had to recruit more than 28,000 program committee members, nearly three times the number from the previous year.^a In its own words, the conference was pushing “the absolute limit” of its reviewing systems with respect to storage, compute, bandwidth, and qualified reviewer time. The scale of correspondence alone was staggering; the team received more than five times the total volume of the entire previous year.

Our flagship venue, ACM CHI, is not solely focused on AI, and its growth curve may differ. CHI 2023 received roughly 3,200 paper submissions. By 2026, that number had risen to 6,730.^b Submissions have more than doubled, and the growth rate is not slowing; the

most recent year-over-year increase was 34 percent. But LLM-assisted writing affects every research discipline, and the tools that make paper production faster are available to everyone. If CHI were to experience even a modest version of AAAI's trajectory, we could be looking at 15,000 to 20,000 submissions by 2030. A more dramatic scenario, one that mirrors AAAI's doubling pattern, would place us somewhere between 25,000 and 30,000. At that scale, with two reviews and two associate chairs per submission, we would need 100,000 to 120,000 individual reviews. The current system cannot deliver that: It's not designed for this type of volume, and it doesn't have the infrastructure that would be needed to support it.

A surge in submissions might suggest that research activity and funding are booming, but the reality is less exciting. Government budget allocations for R&D across the Organization for Economic Cooperation and Development (OECD) area grew by only 0.6 percent in 2023, down sharply from 5.5 percent the year before. The increase in 2024 was 0.5 percent, well below GDP growth. According to Eurostat, R&D spending as a share of GDP has stagnated at 2.24 percent in the EU, far short of the long-standing 3 percent target. R&D intensity across

the OECD has remained flat at 2.7 percent since 2020.^c Proportional expansion of research capacity is not driving growth in submissions; more papers are being produced from roughly the same pool of resources.

I have served on program committees long enough to recognize the early symptoms of this strain. Recruiting qualified reviewers is harder each cycle; the pool of willing experts does not grow at anywhere near the rate submissions do; review quality has become uneven, not because reviewers lack good intentions, but because the match between expertise, experience, and assigned papers loosens as numbers climb. AAAI 2026 acknowledged this development, noting that some reviewers would assess papers “adjacent to their core research fields rather than papers that perfectly match their expertise.” Reviewer fatigue is real, it is growing, and it corrodes the care that good reviewing demands [1]. Additionally, many researchers carelessly use AI tools. I stumbled over many hallucinated references when reviewing for a large SIGCHI conference, which also has seen a 50 percent increase in submissions.

And then there is the feedback loop: AI-assisted papers are increasingly reviewed by AI-assisted reviewers. Research [2] suggests that between 6.5 and 16.9 percent of peer review text at some AI conferences may have been substantially generated or modified by LLMs. arXiv has responded by tightening its rules for computer science, now requiring prior peer review for review articles and position papers, explicitly citing the flood of content that generative AI makes “fast and easy to write.”

Using AI to polish text is

Our community should start treating the possibility of 30,000 CHI submissions by 2030 as a realistic planning scenario, not a distant hypothetical.

a <https://bit.ly/40OxPZU>

b <https://bit.ly/3PteavY>

c <https://bit.ly/4daQkPD>

understandable, and for many non-native English speakers it represents a genuine equalizer (a reason why I am also using such tools for this blog post). The deeper concern is not that individuals are using tools; it is that the structural workload is driving systemic reliance on automated shortcuts at both ends of the evaluation process. When the person writing and the person reviewing are both partially off-loading judgment to the same class of technology, something in the scholarly exchange thins. The question of whether a contribution advances knowledge becomes harder to answer as fewer humans genuinely engage with the work.

A study published in *Science* in 2025 [3] analyzed more than 2 million preprints and found that LLM use accelerates manuscript output, reduces barriers for non-native English speakers, and diversifies the discovery of prior literature. These are real benefits. But the same study found that the traditional correlation between writing quality and research quality is breaking down: Papers likely written with LLM assistance showed a weaker link between linguistic sophistication and eventual journal acceptance. At the same time, however, a recent paper in *Nature* [4] shows that the adoption of AI in science creates an apparent paradox, where individual researchers can achieve greater impact, while the overall scope of collective scientific activity becomes more limited. In other words, the signal that reviewers and editors have long relied on to sort strong work from weak work is becoming unreliable, precisely now when the volume of work to sort is exploding.

The research ecosystem was designed for slower knowledge production. Its filtering mechanisms assumed that writing a paper was itself a form of intellectual labor, one that took time and reflected sustained engagement with a problem. When LLMs compress the writing phase, the bottleneck shifts to ideation and empirical rigor, which are also increasingly accelerated by AI. That shift is not inherently bad, but our institutions and conferences have not adapted to it.



The underlying tension predates LLMs. The publish-or-perish culture and the emphasis on the h-index have always rewarded volume over depth. LLMs have not created this problem; they have amplified it. The efforts and cost of producing a paper have dropped, but the incentive to produce has not changed. The predictable result is more output for the same or smaller input of careful thought.

If submission growth continues unmanaged, the consequences are foreseeable. Review quality will decline further as overloaded reviewers produce shallower assessments; program committees will burn out; desk rejections and automated filtering will become de facto gatekeepers, displacing expert judgment; valuable work will become invisible inside a flood of adequate-but-unremarkable papers. Communities may fragment into narrower sub-venues, losing the cross-pollination that a generalist conference like CHI has always enabled. NeurIPS 2025 already provided an interesting episode: Even after expanding to two venues in two cities, the conference reportedly had to reject roughly 400 papers that had already been accepted, simply because physical capacity ran out. When logistics override scholarly merit, something fundamental is broken.

I don't have a solution to offer, and I don't know if anyone else does yet. But our community should start treating the possibility of 30,000 CHI submissions by 2030 as a realistic planning scenario, not a distant hypothetical. That means

beginning to discuss alternative review models, new conference structures, filtering mechanisms with transparent criteria, and incentive reforms that reward depth over volume. It means revisiting whether a single annual deadline and a single physical event remain the right architecture for a field of this size. It means reconsidering what we measure when evaluating one another's contributions. And it means being honest about the fact that the tools we helped build are now reshaping the conditions under which we do our work.

If CHI were to reach more than 25,000 submissions by 2030, the challenge would extend far beyond logistics. It would indicate that the rate of paper production has exceeded the capacity of peer review systems designed for a much smaller community. Addressing the coming surge of AI-accelerated research will demand substantial collective effort. Discussions about how to adapt our review processes, conference structures, and evaluation practices must begin now. In a world where producing papers has become easy, how can we safeguard the value of the critical thinking that should precede the writing?

REFERENCES

1. Chidambaram, V. The role of LLMs in academic reviewing. SIGOPS blog, July 30, 2025; <https://bit.ly/4dDUWh6>
2. Liang, W. et al. Monitoring AI-modified content at scale: A case study on the impact of ChatGPT on AI conference peer reviews. *Proc. of the 41st International Conference on Machine Learning*. ACM, 2024, Article 1192, 29575–29620.
3. Kusumegi, K. et al. Scientific production in the era of large language models. *Science* 390, 6779 (2025), 1240–1243; doi:10.1126/science.adw3000
4. Hao, Q., Xu, F., Li, Y., and Evans, J. Artificial intelligence tools expand scientists' impact but contract science's focus. *Nature* 649 (2026), 1237–1243.

👤 **Johannes Schöning** is a professor of HCI at the University of St. Gallen and an affiliated professor at the Mohamed bin Zayed University of Artificial Intelligence (MBZUAI). He aims to empower individuals and communities with the information they need to make better data-driven decisions by developing novel user interfaces with them.
→ johannes.schoening@unisg.ch