Obsession With AI May Be Starving the Research Fields It Needs to Survive

Originally published Weds 21 Oct 2026 at https://www.unite.ai/obsession-with-ai-may-be-starving-the-research-fields-it-needs-to-survive/



I've been writing exclusively about artificial intelligence for ten years. In 2022, there were only a handful of other journalists in the world exclusively dedicated to AI; now we number in the multiple thousands or more, with better-heeled outlets employing more than one staff member to exclusively cover artificial intelligence.

As a species, we now concentrate our attention, efforts, fears and hopes around a once-in-a-century cultural and technological phenomena that, until the advent of *useful* Large Language Models and *effective* diffusion-based image and video generation, had been languishing in the 'curiosities' area, with deepfakes a political football, but world economies apparently unthreatened.

The 'Routine' AI Storm

Obviously, that's all over now. Everything is threatened, everything is disrupted, and, arguably, 2026 has been the year in which predictions *about the following year* were hardest to ground*.

Weekly, sometimes *several* times a week, new closed-source and open-weight models are released, now at such a frequency that some are not even reviewed by the enthusiast or hobbyist community, who are developing a jaundiced outlook in this regard, and an ever-rising standard for attention in the face of the constant stream of releases.

As someone who keenly and professionally observes the research and release scene, I must say that improvements across the various sectors of AI R&D, though they have a great impact at large intervals, tend (and have always tended) to emerge *in small increments*, rather than revolutionary or evolutionary tent-poles, such as Transformers. These days, the shocks are mostly about AI's effect on the economy, rather than its latest capabilities.

Additionally, as I have observed before, the *quality* of ideas represented in the enormous boom in AI research since 2022 does not always match the quantity; many are 'Frankenstein'-style re-strikes, recombinations, or re-assemblies of existing and more novel work – patchwork architectures designed to create *apparent* significant improvements on prior works – but which, on examination, are hamstrung by brittle methodologies, or absurdly high VRAM costs.

The Topical Ghost Towns

When those thousands of journalists migrated to AI coverage, the subjects they had previously covered must inevitably have been left with less direct attention. In a way, it hardly mattered, because the topic and implementation of AI became so ubiquitous in culture and infrastructure that – for instance – an energy correspondent would now be covering AI-related issues more than any other; likewise for those on the economics, employment, politics and policy beats, among others.

It seemed that all these topics could now be filtered through the lens of AI, which may be a partial explanation for the huge increase in 'dedicated' AI coverage, in an ever-diminishing field (albeit one that has rallied a little lately).

Consequently, one has to wonder if *other* topics now receive inadequate coverage due to 'AI obsession' – topics with stories, trends and branches that may eventually affect the 'golden topic' of AI itself?

But what can you do? Resources are limited, and so far AI is the story of the century.

© Martin Anderson and/or commissioning entities

The Great AI Migration

This, according to a new paper that has stirred my blood on the subject, is what may be happening to AI itself – that, besides the risk that retrenchment and re-feeding will lower the quality of future models, the wider scientific ecosystem in which ideas emerge and cross-pollinate across disciplines is becoming so heavily concentrated on AI research that 'serendipitous' moments – when unexpected connections between seemingly unrelated ideas can transform entire fields – may become increasingly rare.

In short, the new work, titled *Stronger artificial intelligence breakthroughs hasten their own exhaustion*, suggests that AI obsession is sucking away the ambient conditions in which new developments, *including developments directly affecting AI*, occur; and that (to grossly paraphrase the authors) there is a certain 'incestuous trend' in this new locus of interest that seems unlikely to arrive at a good end.

The paper states†:

> *'A dominant paradigm removes not only its rivals but the diversity from which its own successor would have been recombined.*
>
> *'The post-2017 narrowing of attention, the migration of more than half of all field-switching authors into AI, the declining vitality of new ideas, and the steepening tail-decline of active idea populations across successive cycles together describe a system whose creative capacity is being eroded by the success of the dominant paradigms it has produced.'*

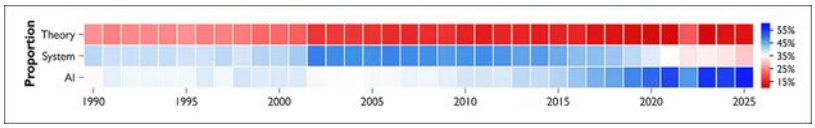
Prior work has observed that AI may be turning research into a monoculture, but the authors of the new work also believe that AI's growing dominance may be starving it of the fresh ideas it needs to keep advancing, comparing AI research to biological sciences, where breakthroughs such as cryo-EM and AlphaFold have triggered similar cycles of innovation – and decline.

To produce empirical evidence for their theories, the authors analyzed 7.18 million OpenAlex computer science papers published between 1990 and 2025, tracking changes in research topics and citations; the movements of 3.7 million researchers between fields; and 73,089 peer reviews to see how AI's growing dominance has affected scientific diversity.

One of the most interesting findings to emerge from the data relates to the recent migration of researchers towards AI and away from the subjects they were already studying:

> *'Migrants are mid-career, moderately established researchers who leave the work that earned them their standing to pursue the gravitational pull of the field, even at the cost of a pivot penalty '*

The study found that more than half of the 918,862 recorded switches between computer science fields were *into AI*, with other areas, particularly systems and theoretical research, steadily losing people.



From the new paper, changes in the share of computer science papers devoted to AI, systems and theory from 1990 to 2025. Each square represents one year, w and red a smaller one. AI's share rose sharply after 2017, while the other fields declined. Source

The researchers most likely to make the jump were experienced academics with established careers. Meanwhile, AI's share of all computer science papers grew from 40% in 2017 to more than 55% in 2025.

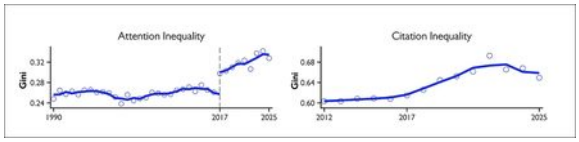


From the Arxiv blog, the rise in Computer Science submissions over a 2-year period 2024-26. Source

Three Signals

The paper identifies three 'converging signatures' that the authors argue have come to define computer science since 2017.

The first is a sharp narrowing of *research attention* after 2017, when the Transformer architecture was introduced. By extracting the ten most prominent keywords from each of 7.18 million computer science papers, the authors tracked how much research was concentrating on the same subjects:



Graphs showing how computer science research has become increasingly concentrated since 2017. Left: researchers are focusing on a narrower range of topics. Right: a growing share of citations is going to a smaller selection of papers. Both trends accelerated around the arrival of the Transformer.

© Martin Anderson and/or commissioning entities

For decades, this concentration had remained relatively stable, even through the rise of the internet, smartphones and deep learning. But between 2017 and 2025, the Gini coefficient†† for research attention (see image above) rose from 0.26 to 0.33, an increase of 27%. Citations became more concentrated too, with their Gini coefficient climbing from 0.60 to a peak of 0.69.

The researchers, who come from Tsinghua University, the University of Chicago, and the Santa Fe Institute, state:

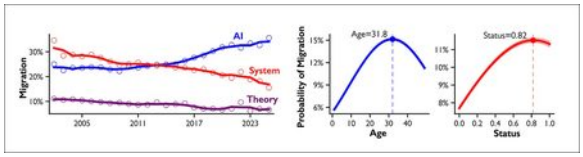
> *'The Transformer disrupted the field on a scale its predecessor did not. The pattern is not an artifact: it survives variation in [keyword count].'*

And continue:

> *'AI's share of computer-science publications rises from 40% to over 55% between 2017 and 2025, while systems research recedes from a long golden age and theory contracts further still.'*

Therefore the findings suggest that, even as computer science expanded, researchers were increasingly *pursuing the same ideas*, and citing a narrower selection of work.

The second 'signature' concerns the growing number of researchers abandoning their specializations to work in AI: tracking 3.7 million researchers, the authors identified nearly 919,000 migrations between computer science fields. Of these, 53.81% were *into AI* – a research area which gained researchers every year, while systems and theory steadily lost them:

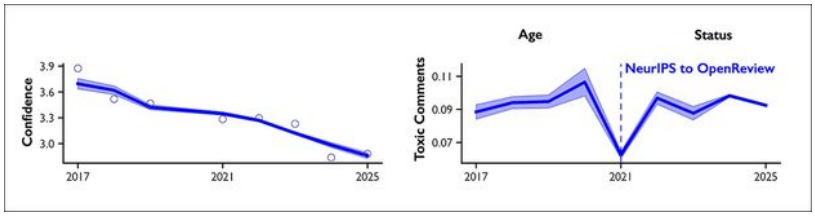


Researcher migration into AI. The left graph shows AI attracting a growing share of researchers from other computer science fields, while systems and theory lose ground. The two graphs on the right show who is most likely to make the move: researchers roughly 32 years into their publishing careers, with stronger publication records than about 82% of their peers. A change of field is counted when a researcher starts publishing most of their papers in a different area.

Those making the switch tended to be established, mid-career academics rather than newcomers or veterans. The likelihood of moving into AI peaked at an academic age of 31.8 years and among researchers at the 82nd percentile for publication standing.

The third of the three trends that the paper notes is a *growing strain on peer review* – the process in which other researchers assess a scientific paper's quality and reliability before publication.

Examining 73,089 reviews of 26,731 papers submitted to 45 conferences and other venues between 2017 and 2025, the authors found that reviewers' confidence in their ability to judge the papers fell from 3.9 to 2.9, on a five-point scale:



Graphs of changes in scientific peer review between 2017 and 2025. The left graph shows reviewers becoming less confident in their ability to assess submitted p 3.9 to 2.9 out of five. The right graph tracks hostile review comments, which briefly declined in 2021 when the NeurIPS conference moved to the more transparent rising again. Both lines show three-year rolling averages.

The authors suggest that this decline may reflect AI's rapid expansion, which has made it harder to match papers with suitably qualified reviewers:

> *'Reviewer pools have become flooded with newcomers drawn in by AI's gravity. Reviewers increasingly receive papers laced with AI methods that fall outside their core expertise. Matching submissions to genuinely expert reviewers has become structurally difficult.*
>
> *'Acceptance, in this regime, increasingly resembles a coin flip.'*

After the Gold Rush(es)

The paper also suggests that AI's greatest breakthroughs may ultimately hasten their own exhaustion, since each major advance is followed by a surge of new ideas – but this initial flurry of innovation then settles into torpid repetition, as research becomes increasingly concentrated around the breakthroughs themselves.†:

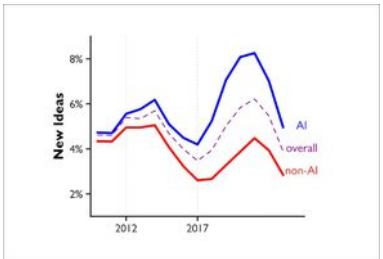
> *'Successive cycles in modern AI saturate earlier and exhaust more sharply. The cycle is compressing.*
>
> *'Across 36 years of computer science development, spanning the three major technological advances of AlexNet, the Transformer, and LLMs, we repeatedly observe the same spiral [pattern].*
>
> *'Notably, this pattern is not confined to computer science. In biology, we observe the same cycle of exhaustion following major non-computational (e.g., cryo-EM) and computational (e.g., AlphaFold) advancements.'*

This 'exhaustion' certainly happened with autoencoders, Generative Adversarial Networks (GANs), and autoencoder approaches such as deepfakes. All of these were mercilessly flogged by researchers for years beyond their clear exhaustion of potential, until new frameworks and approaches finally arrived to take the pressure off them.

© Martin Anderson and/or commissioning entities

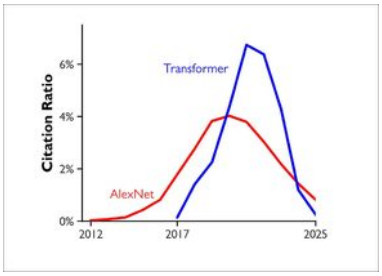
In any case, to track these cycles, the authors searched scientific abstracts for distinctive nouns and phrases, following their appearance, disappearance and occasional return as indicators of changing research ideas.

AlexNet in 2012 and Transformers in 2017 each unleashed a wave of new ideas that eventually subsided. AI pulled further ahead with each breakthrough, while other computer science fields fell behind, eventually generating new ideas at a lower rate than before 2012:



Comparison of the share of new research ideas in AI (blue), non-AI computer science (red), and both combined (dashed purple). New ideas surged following AlexNet in 2012 and the Transformer in 2017, before declining. AI's advantage widened after the second breakthrough. Lines show three-year rolling averages.

The pattern extends to ideas falling out of use, with new breakthroughs sweeping aside older approaches, sometimes years after their arrival. Yet the breakthroughs themselves are caught in the same cycle, in that each attracts a surge of research attention that eventually subsides, leaving the field waiting for *another discovery* to revive its momentum (as occurred with Gaussian Splatting, after a 31-year interval):



Comparison of citations to the original AlexNet (red) and Transformer (blue) papers from 2012 to 2025. Both attracted growing shares of citations before declining sharply, illustrating how attention to major breakthroughs eventually fades. Lines show three-year rolling averages.

AI's success concentrates researchers and ideas around a handful of dominant approaches, while the surrounding fields that might help produce their successors continue to lose ground the authors argue. And while they don't predict outright that this will prevent future breakthroughs, they do raise the possibility that AI is gradually consuming the diversity on which its next leap may depend.

The number of ideas still in active use was also tracked; across successive waves of AI development, this number fell progressively further below expectations, with a larger shortfall in the wake of Transformers than after AlexNet:

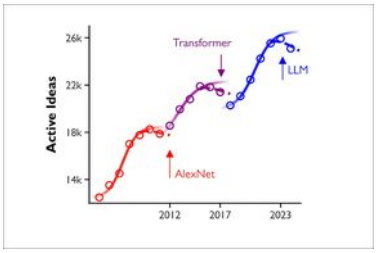


Chart showing the number of active research ideas across successive AI breakthroughs. Solid lines show modeled growth, while dashed sections show actual counts falling below projected levels. The shortfall is greater after the Transformer than after AlexNet, with an early decline also appearing during the LLM era.

In summary, the authors conclude that the concentration of effort and people on an increasingly narrow range of challenges related to AI threatens to impoverish adjacent fields that would once have naturally fed ideas into the AI sector – and diverse other sectors.

Referring to John Hassler's comment regarding the award of the 2025 Nobel Prize in Economic Sciences, they conclude:

> *'The Nobel committee warned that the mechanisms of creative destruction must be upheld if stagnation is to be avoided. Our results identify one such mechanism, intellectual diversity, and show that current institutions are eroding it.*
>
> *'The institutions that allocate attention, including reviewing systems, submission fees and rules, reward structures, hiring, promotion, and retirement conventions 12, all shape whether the next paradigm has room to emerge.*
>
> *'Designing them well means resisting the impulse to double down on whichever paradigm is currently winning, taxing concentration rather than exploration, and shifting evaluative authority away from the prolific and toward the novel.*

© Martin Anderson and/or commissioning entities

'Creative destruction sustains growth only as long as it leaves something behind to recombine. A singularity has no successor.'

** There should be a word for this: 'forecast-resistant'? Well, that's not a discrete word, but a term, and also sounds like it refers to 2026 instead of the view of 2027 from 2026. 'Unfuturable'? That got a red line in the spell check, so let's settle on it.*

† *My conversion of the authors' inline citations to hyperlinks.*

†† *The Gini coefficient indicates how unevenly something is distributed, on a scale from 0 (perfect equality) to 1 (maximum concentration). Here, a higher score means that researchers are focusing on fewer topics.*

First published Wednesday, October 21, 2026

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.
Website: martinanderson.ai
Contact: martin@martinanderson.ai