



# AI is turning research into a scientific monoculture

Cecilie Steenbuch Traberg, Jon Roozenbeek & Sander van der Linden



Generative AI deserves scientific attention. But the rush to study it is producing a feedback loop of topical and methodological convergence, flattening scientific imagination and crowding out the pluralism needed to keep research adaptive, resilient, and intellectually generative.

Conferences, journals, and funding calls in the social and behavioural sciences are increasingly dominated by (generative) AI<sup>1</sup>. Many academics have rebranded themselves as “AI researchers”. Every project finds its “AI angle.” This shift is understandable and important: generative AI is a consequential technological development, and psychologists and behavioural scientists are well-positioned to examine its impacts<sup>2</sup>. But this focus is becoming all-encompassing. *The New Yorker* recently argued that AI is “homogenizing our thoughts”<sup>3</sup>: that by repeatedly surfacing the most probable continuations of human thought, these systems are nudging human reasoning toward conformity. Ironically, scientific culture is drifting toward a meta-version of that claim. While earlier work warned that increasing AI-adoption may lead to a scientific monoculture<sup>4</sup>, empirical evidence now suggests this process is underway<sup>5</sup>. In studying AI, research practices are themselves becoming more uniform - converging not only in what is studied, but in how questions are framed, investigated, and evaluated. Understanding this convergence as a feedback loop rather than an unavoidable trend opens the possibility of targeted interventions to preserve scientific diversity before monocropping becomes fully entrenched.

## The Rush Effect

Across social science disciplines, a race has emerged to show what AI can do<sup>1</sup>. The logic is partly pragmatic: funders, journals, and institutions reward topicality. But it is also cultural. To not work on AI is increasingly perceived as a missed opportunity, or worse - irrelevance. Ironically, AI tools themselves amplify this acceleration. Recent evidence shows AI-augmented papers are cited more, AI-adopting researchers publish more, and career advancement accelerates<sup>5</sup>. As scientists increasingly use LLMs to generate ideas and synthesise literatures<sup>4</sup> - often on the topic of AI itself - the technology feeds its own growth, increasing the pace of production<sup>6</sup>, while narrowing the space for slow, divergent thinking<sup>4,5</sup>.

## The feedback loop

What follows is a self-reinforcing cycle of topical, methodological, conceptual, and linguistic convergence (see Figure 1). The AI-monoculture feedback loop spans multiple levels, from broad techno-cultural salience and institutional incentives to methodological practices and epistemic feedback. Recent large-scale evidence shows that several components of this loop are already in motion: for example, AI adoption is associated with strong individual-level incentives, shared methodological uptake, and system-level

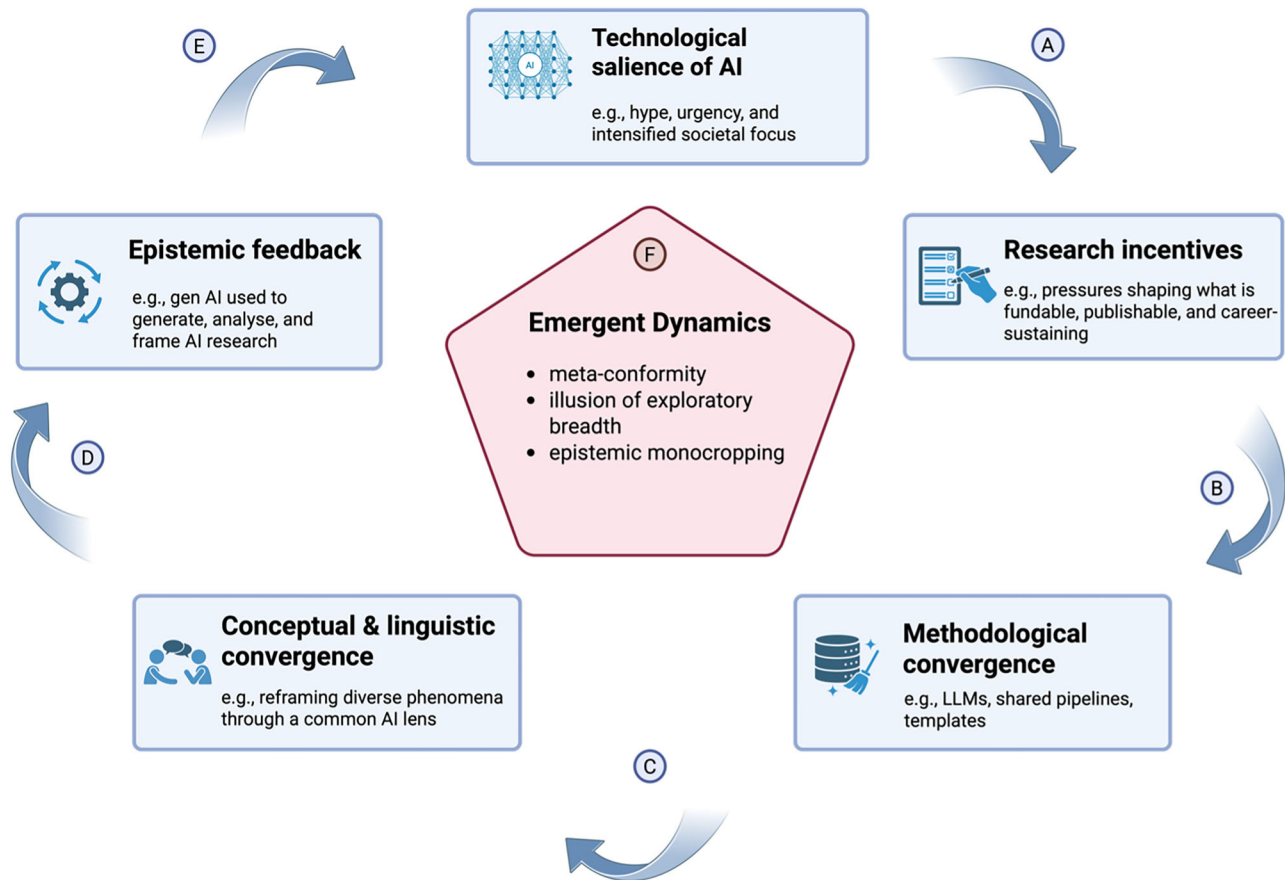
outcomes such as reduced topical breadth and scientific engagement<sup>5</sup>. Together, these findings indicate that the dynamics underpinning scientific monocropping are no longer hypothetical but actively unfolding.

At the same time, viewing these patterns through a feedback-loop lens highlights additional dynamics that warrant systematic investigation. The cultural salience of AI is typically treated as background rather than measured directly; conceptual and linguistic convergence is often inferred indirectly rather than analysed in its own right; and epistemic feedback - how AI tools shape idea generation, framing, and evaluation - remains poorly understood. These elements may be critical for understanding how monocropping accelerates and where intervention is most effective.

**Topical convergence.** Across disciplines, diverse research agendas are reframed through the AI lens. Questions that once centred on cognition, communication or institutions now become questions about *AI and* cognition, *AI and* communication, and *AI and* institutions. This shift is visible in publication trends: the AI Index 2025 Report<sup>1</sup> documents a sharp rise in AI-related publications across scientific fields. The expansion of AI research into disciplines that historically had little engagement with AI raises questions about how research agendas are being reframed. This reframing reflects heightened technological salience and institutional signals about what counts as timely and consequential research, alongside intellectual curiosity.

**Methodological convergence.** Analytic pipelines, toolkits, and benchmark datasets are becoming increasingly shared across the behavioural and social sciences. Large language models now serve as general-purpose analytic tools, used for classification, text analysis, content generation, and behavioural modelling. This mirrors an earlier wave of methodological convergence that accompanied the rise of computational social science, when diverse research traditions began adopting common machine-learning and data-mining workflows. What were once distinct methodological communities - experimentalists, qualitative researchers, ethnographers - are again being pulled toward a common computational paradigm. As these tools become the default templates for inquiry, methodological pluralism risks being eclipsed by increasingly standardised analytic pipelines. Embedded as default research infrastructure, they also shape how ideas are generated, filtered, and evaluated, feeding back into which questions appear tractable or worth pursuing.

**Linguistic convergence.** Even the surface texture of science begins to sound alike. Proposals and papers repeat familiar phrases: “trustworthy AI,” “human-AI collaboration,” “ethical deployment.” Researchers echo what appears credible, relevant or fundable, but in so doing, such informational conformity may inadvertently create a narrow terminology and, thereby, limit variation in research questions. The result is a flattening of discourse - the homogenisation of scientific language itself. Furthermore, over time, this narrowing of language may reinforce meta-conformity, as



**Fig. 1 | The AI monoculture feedback loop.** A As AI becomes culturally prominent, characterised by hype, urgency, and intensified societal focus, it attracts attention and legitimacy. This salience is translated into institutional signals about what counts as relevant, timely, and consequential research. B These signals are formalised through research incentives, including funding priorities, journal norms, and career pressures, which reward alignment with AI-centred topics and approaches. As a result, pursuing non-AI work may increasingly carry reputational and career risk. C Incentive structures increasingly favour the use of AI as a general-purpose research instrument—used to generate data, conduct analyses, synthesise findings, and scale scholarly output. As these tools become normalised as default research infrastructure, analytic workflows and methodological practices converge around shared AI-driven pipelines and

templates. D As methods converge, so too do the linguistic and conceptual frames used to define and explain research problems. Diverse phenomena are increasingly reframed through a common AI lens and language, as methods shape what is easy to articulate, justify, and evaluate. E AI systems are then used to generate ideas, synthesise literatures, frame research questions, and provide feedback on ongoing work—often about AI itself. This epistemic feedback accelerates the production and visibility of AI-related research, making AI appear increasingly indispensable and self-evidently central to scientific inquiry, thereby reinforcing its technological salience. F Together, these interactions give rise to emergent system-level dynamics, including meta-conformity, an illusion of exploratory breadth, and epistemic monocropping, even when individual research decisions are locally rational.

familiar framings become taken-for-granted markers of relevance and rigor.

The above process may be explained through a behavioural science lens. Individuals often align with perceived group norms, driven by both informational and reputational pressures<sup>7</sup>. Scientists are not immune to these pressures; they too respond to cues about what their communities value, reward or attend to.

The prevalence of AI in research culture illustrates a recursive form of conformity: meta-conformity. The object of study - AI - begins to shape the form of inquiry itself. It is an epistemic mirror: science studying intelligence while gradually adopting its logic of standardisation and efficiency. And, while LLM outputs may resemble human judgement, they rely on fundamentally different generative processes based on statistical pattern

recognition rather than social or contextual understanding; a phenomenon termed *epistemia*<sup>8</sup>.

### What is at stake

Building on prior warnings about scientific monocultures and epistemic narrowing in AI-driven research, the dynamics described here carry three system-level consequences.

First, loss of intellectual diversity. Topics that do not fit easily within the AI narrative may struggle for legitimacy or resources. Important social, behavioural, and methodological questions risk marginalisation simply because they are not computationally tractable. Second, methodological pluralism: the ability to triangulate complex phenomena through multiple epistemic lenses. An overreliance on a single class of AI-based tools may

leave science blind to what those tools cannot capture<sup>4</sup>. Third, field-level optionality: the collective capacity of a research ecosystem to pivot when the world changes. Just as biodiversity protects ecosystems from collapse, intellectual heterogeneity protects science from paradigm shocks. But the current mass pivot toward AI may itself erode the diversity that enables future pivots. Without a broad portfolio of approaches and expertise, fields become fragile, less able to adapt when the next scientific challenge or technological shift emerges.

### A path forward

AI must be studied, including by social scientists. Yet the *way* we study AI must not narrow the scientific imagination or erase the intellectual diversity that preceded it. The goal is not to retreat from AI, but to resist allowing its urgency to collapse the range of methods and perspectives that make scientific ecosystems resilient.

To counteract epistemic monocropping, the scientific community can cultivate structural and cultural safeguards, which we propose below. If – as we suggest – scientific monocropping emerges from reinforcing feedback loops rather than isolated choices, then effective responses must target multiple points in that loop. No single intervention will suffice. Instead, preserving scientific diversity requires a portfolio of safeguards that slow, counterbalance, or redirect the dynamics outlined in Figure 1.

**Funding diversification.** As techno-cultural salience intensifies, institutional incentives increasingly align around AI-centred topics, reinforcing topical convergence. Funding priorities play a pivotal role in translating salience into scientific momentum, shaping which questions appear timely, consequential, and worthy of sustained investment. Allocating a protected share of research resources to non-AI projects can counteract this amplification effect at its point of entry. Such diversification can be implemented at multiple levels – from national funding bodies and international programmes, to universities, departments, and individual grant calls – helping to decouple scientific value from immediate AI relevance.

**Methodological rotation.** As AI tools become embedded as default research infrastructure, methodological convergence risks hardening into epistemic lock-in. Supporting experimental, qualitative, ethnographic, and design-based approaches alongside computational ones helps prevent shared AI pipelines from becoming the unquestioned template for inquiry. Such rotation preserves the capacity to triangulate complex phenomena across complementary ways of knowing, ensuring that what is easy to compute does not become synonymous with what is worth studying.

This also bears on the maintenance of expertise. Researchers trained in non-computational traditions play a critical role in sustaining methodological diversity, while those trained primarily in computational methods face growing pressure to outsource key stages of inquiry to LLMs. Treating AI as a complement rather than a substitute for methodological judgement helps limit epistemic feedback through which AI-optimised workflows increasingly define what questions appear tractable or legitimate.

**Editorial and review practices.** Conceptual and linguistic convergence is reinforced at the point of evaluation, where shared framings, terminologies, and methodological assumptions become markers of rigour and relevance. When AI-centred work is assessed primarily by reviewers with similar computational backgrounds, familiar questions and vocabularies are more readily legitimised, while alternative framings are treated as peripheral.

Ensuring that AI-related work is evaluated by reviewers with diverse methodological and theoretical backgrounds broadens what counts as contribution and rigour. By legitimising multiple ways of defining problems, interpreting results, and articulating significance, editorial practices can slow linguistic homogenisation and preserve conceptual diversity.

**Institutional incentives.** As AI tools become embedded across the research pipeline, institutions increasingly reward outputs that align with speed, scale, and technological familiarity. Under such conditions, researchers may feel compelled to integrate AI across idea generation, analysis, and writing, as AI-optimised workflows yield greater visibility and career advancement. This intensifies epistemic feedback, reinforcing perceptions of relevance and indispensability.

Incentive structures that reward originality, theoretical depth, and long-term contribution can moderate this self-reinforcing dynamic. By protecting heterodox research trajectories and valuing sustained engagement with difficult problems, institutions can preserve field-level optionality and reduce the likelihood that locally rational choices accumulate into system-level monocropping.

At the individual level, scholars might reflect on their own epistemic flexibility: Which problems still matter even if they are not computationally solvable? And which problems should we resist squeezing into an AI lens simply because that lens is currently ascendant?

Whether AI ultimately narrows or expands scientific imagination will depend on how the feedback loops it triggers are shaped and governed across the scientific ecosystem. These dynamics are neither fixed nor inevitable, but contingent on institutional choices about what is valued, rewarded, and sustained.

In warning that AI might make us think alike, we may have begun to think alike about AI. This symmetry should give us pause. The task for social science is to ensure that, in navigating this moment, we do not become artificial ourselves.

Cecilie Steenbuch Traberg <sup>1</sup> ✉, Jon Roosenbeek <sup>2</sup> & Sander van der Linden <sup>2</sup>

<sup>1</sup>Copenhagen Business School, Copenhagen, Denmark. <sup>2</sup>University of Cambridge, Cambridge, UK. ✉e-mail: [ces.digi@cbs.dk](mailto:ces.digi@cbs.dk)

Received: 22 December 2025; Accepted: 11 February 2026;

Published online: 23 February 2026

### References

1. AI Index Steering Committee. *The AI Index 2025 Annual Report*. (2025).
2. Rahwan, I. et al. Machine behaviour. *Nature* **568**, 477–486 (2019).
3. Chayka, K. A.I. Is Homogenizing Our Thoughts. *The New Yorker* (2025).
4. Messeri, L. & Crockett, M. J. Artificial intelligence and illusions of understanding in scientific research. *Nature* **627**, 49–58 (2024).
5. Hao, Q., Xu, F., Li, Y. & Evans, J. Artificial intelligence tools expand scientists' impact but contract science's focus. *Nature* (2026) <https://doi.org/10.1038/s41586-025-09922-y>.
6. Noy, S. & Zhang, W. *Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence*. <https://www.science.org>.
7. Cialdini, R. B. & Goldstein, N. J. Social Influence: Compliance and Conformity. *Annu. Rev. Psychol.* **55**, 591–621 (2004).
8. Loru, E. I. et al. The simulation of judgment in LLMs. 122, (2025).

### Author contributions

Conceptualisation: C.S.T., Methodology: C.S.T., J.R., S.V.D.L., Investigation: C.S.T., Visualisation: C.S.T., Supervision: C.S.T., J.R., S.V.D.L., Writing – original draft: C.S.T., Writing – review & editing: C.S.T., J.R., S.V.D.L.

### Competing interests

The authors declare no competing interests.

---

**Additional information**

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s44271-026-00428-5>.

**Correspondence** and requests for materials should be addressed to Cecilie Steenbuch Traberg.

**Peer review information** The manuscript was considered suitable for publication without further review at Communications Psychology. Primary Handling Editor: Marike Schiffer. A peer review file is available

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026