Go Big or Go Home: AI-based Mathematics in the Big Science Era

Michael Friedman (University of Bonn, Hausdorff School for Mathematics),
Kati Kish Bar-On (Boston University & Harvard Business School)[1]

Abstract

This paper argues that AI is reshaping mathematics in ways that outstrip the classic Big Science analogy. As in Big Science, today's mathematical work which is classified as Big Science relies on concentrated computation, specialized and hierarchical teams, sociopolitical justification, and large-scale collaboration. But AI also changes who does the epistemic work: systems now generate conjectures, explore proof spaces, and formalize arguments, shifting mathematicians toward roles of curation, evaluation, and orchestration within hybrid human–machine assemblages. This reconfiguration redistributes epistemic authority toward corporate labs and elite hubs that control models, data, compute, and narrative, thereby gating which problems appear tractable and legitimate. Analyzing four dimensions: resource concentration, division of labor, political embedding, and collaboration, we conclude that AI-based mathematics constitutes a novel sociotechnical regime that we call "Big AI Mathematics", in which truth claims are increasingly mediated by opaque pipelines and institutional power, making the future of mathematical knowledge as much a governance problem as an epistemic one.

Keywords: Big Science; Big Mathematics; AI-based Mathematics; Knowledge production; Centralization of Knowledge

1. Introduction

The term Big Science emerged in the 1960s to describe a mode of research that had taken hold in the mid-twentieth century. Coined by Alvin Weinberg, the phrase captured a series of dramatic changes in how science was conducted: large-scale collaborations, dependence on vast infrastructures, government and military funding, and the emergence of bureaucratic managerial structures to coordinate these efforts (Weinberg 1961; Galison & Hevly 1992). The Manhattan Project, CERN, and NASA's space programs are emblematic of this new modality, where science became a national and international enterprise, both politically and economically embedded. Big Science was not simply "more" science, it represented a new social form of knowledge production.

By the late 20th century, several mathematical research projects began to take on more characteristics of Big Science. Supercomputing centers, large-scale modeling in climate science and epidemiology, and international collaborations on formal proof verification all hinted at a nascent form of Big Mathematics (Parshall & Rowe 1994; Kaiser 2005). One can give two examples of such projects of Big Mathematics: first, the classification of finite simple groups, a

[1] The authors contributed equally to this work and are listed in alphabetical order.

project which is explicitly termed by Volker Remmert and Rebecca Waldecker as the first one incorporating characteristics of Big Mathematics[2]. Second, the more distributed projects of formalization of proofs by automated proof-assistants such as *Coq* or *Lean*. Both projects, though obviously of a different character and scale, have increasingly required infrastructure, coordination, and sustained institutional support. Moreover, in the past few years there has been an increasing growth in the use of AI-based technologies in mathematical practice and research, where AI-based tools begin to play a pivotal role. From machine-learning-assisted theorem provers[3] to symbolic reasoning tools integrated with large language models, the role of AI is no longer merely 'symbolical' or computational. Various examples of the usage of AI-based technologies in mathematics can be given: *AlphaTensor* (Fawzi et al. 2022), finding new ways to multiply matrices, *AlphaGeometry* (Trinh et al. 2024) discovering new and unexpected proofs for problems in plane Euclidean geometry, *FunSearch* (Romera-Paredes et al. 2024) which write programs for finding various special sets for the *cap set problem*, or recent discoveries in knot theory (Davies et al., 2021).[4] In this context, we ask: Is this usage of AI-based tools simply more of the same – mathematics adopting the collaborative, tool-driven, and infrastructural characteristics of Big Science – or does it mark a deeper social shift?

What is at stake in this shift is not merely mathematicians' workflow, but the politics of mathematical knowledge itself. AI-based systems in mathematics are developed, maintained, and deployed inside highly asymmetric institutional environments: DeepMind, OpenAI, Google, and a handful of elite academic labs with access to custom hardware, proprietary data, and coordinated engineering labor. These actors are now in a position not only to accelerate mathematical work, but to shape which mathematical problems are considered tractable, valuable, fundable, or even intelligible. This means that mathematical research is beginning to inherit the same political economy that Cathy O'Neil warned about in her book *Weapons of Math Destruction*: computational systems are presented as objective engines of insight, while in practice they encode and amplify existing concentrations of power (O'Neil 2017). We argue that AI-based mathematics is reproducing a quieter, discipline-internal version of this phenomenon.

We hence claim that the analogy to Big Science is no longer enough. Calling AI-driven mathematics "Big Mathematics" may capture the scale of collaboration, but it risks either merely being considered as a continuation of previous constellations or sounding descriptively sociological and ideologically harmless. Our claim is different: we argue that AI-assisted mathematical research is redistributing epistemic authority, from university departments and individual mathematicians toward corporate AI research labs that control computers, models, data, and narrative. This redistribution is already visible in high-profile cases: *AlphaTensor* and *AlphaGeometry* emerge from DeepMind/Google; *FunSearch* is framed as a major advance in

---

[2] https://www2.mathematik.uni-halle.de/waldecker/index_CFSGhistory.html. See also: Gorenstein 1982; Aschbacher 2004.

[3] For example, systems such as GPT-f (Polu and Sutskever 2020) and LEAN-Gym (Polu et al. 2020), as well as *AlphaGeometry* and its successor *AlphaGeometry* 2, which extends the original neural-symbolic architecture with a richer domain language and improved search, enabling near-gold-medalist performance on International Mathematical Olympiad geometry problems (Chervonyi et al. 2025; Trinh et al. 2024; Zhao et al. 2025; also see Nature News 2025 - https://www.nature.com/articles/d41586-025-00406-7).

[4] Some of these examples are presented in (Miao & Wang 2024). Their philosophical importance, especially concerning the notion of mathematical intuition, is discussed in: Friedman and Kish Bar-On, (in review). See also: Davis 2021.

combinatorics by a corporate team with Tensor Processing Unit-scale (TPU) resources;[5] and recent developments such as Terence Tao's ChatGPT-assisted work in analytic number theory show that to do frontier mathematics may increasingly mean to work through corporate models[6]. The danger is that once epistemic authority migrates into these infrastructures, the very standards that define what counts as legitimate, rigorous, or 'frontier' mathematics risk becoming aligned with, and dependent upon, the interests, constraints, and affordances of those corporate systems.

At the same time, it is important to recognize that the enclosure of mathematical knowledge is not a novelty of the AI age, nor are today's emerging gates the first to determine who may participate in mathematical creation. Historically, mathematics has never been a fully open or democratic enterprise. In Plato's *Meno*, the enslaved boy is introduced precisely as someone with no access to mathematical training, and Socrates' demonstration hinges on the assumption that mathematical insight must be elicited from one who has been structurally denied such learning (See in: Cooper & Hutchinson 1997, *Meno* 870-898). For centuries thereafter, formal mathematical education remained restricted to narrow social strata: women, for example, were systematically excluded from universities and scientific academies well into the nineteenth century, with figures such as Émilie du Châtelet and Sophie Germain forced to study outside institutional structures or under male pseudonyms (Grinstein & Campbell 1987; Arianrhod 2012). Even in the twentieth century, access to advanced mathematical training and research positions was tightly bounded by class, gender, race, and institutional affiliation (one of the prominent examples for this is how Jewish mathematicians were treated in Nazi Germany, see: Siegmund-Schultze 2009; or, even more specifically, Emmy Noether is a case in point, see: Rowe & Koreuber 2020). Thus, when we speak about contemporary forms of enclosure we do not describe a rupture from an otherwise egalitarian past, but rather the latest iteration of long-standing social boundaries around mathematical knowledge. AI introduces a new kind of gate, but it does so atop a discipline whose epistemic and institutional architectures have always been somewhat stratified.

Our objective in this paper is twofold: first, we examine which characteristics Big Mathematics in the age of AI shares with Big Science, and second, we explore to what extent AI-based mathematics introduces major transitions in knowledge production and how does that affect Mathematics as a Big Science. To understand what might be distinctive about the usage of AI in mathematics, we must first revisit what Big Science has offered. This is what we do in the next section. We then turn to explore how each of its characteristics manifests in mathematics, and what are their implications. We focus on four aspects of Big Science, following Peter Galison and Bruce Hevly, and analyze how AI-based technologies affect them in a way that represents a potential departure from the expanded notion of Big Mathematics. We conclude by suggesting that this shift marks not merely a quantitative intensification of existing trends, but the beginning of a novel form of distributed mathematical inquiry – one in which cognitive labor, authority, agency and responsibility are co-constructed by human and machine collaborators.

---

[5] Tensor Processing Unit is a neural processing unit application-specific integrated circuit, developed by Google during the 2010s for neural network machine learning.

[6] As mentioned here: https://mathstodon.xyz/@tao/115306424727150237 and here https://mathoverflow.net/questions/501066/is-the-least-common-multiple-sequence-textlcm1-2-dots-n-a-subset-of-t/501125#501125.

2. AI-based Mathematics and Big Science: Characteristics and Implications

As Peter Galison and Bruce Hevly demonstrate in their analyses, Big Science cannot be reduced to a simple matter of scale or expense, but rather constitutes a complex reconfiguration of scientific practice with far-reaching intellectual, institutional, and social dimensions (Galison & Hevly 1992). Their work examines the principal characteristics of Big Science, offering a historical perspective on this transformative development in modern scientific practice in the second half of the 20th century.

The canonical understanding of Big Science has often emphasized quantitative measures – budgets, personnel counts, and physical dimensions of equipment. However, as Hevly argues, "Big Science is not simply science carried out with big or expensive instruments", and terms it an "elusive term" (ibid, p. 355). Instead, what truly constitutes the 'bigness' of Big Science is its extensive reach beyond the traditional boundaries of scientific practice. One may hence offer four key characteristics: (1) *Concentration and Specialization of Resources*: Big Science involves not merely an increase in resources devoted to scientific research, but their increasing concentration on a scale unmatched, for example, by leading pre-second world war academic centers (p. 356). (2) *Hierarchical Organization and Specialized Workforce*: Within these centralized institutions, the scientific workforce has become increasingly specialized and hierarchically organized (p. 357). Laboratories are divided not only into groups of theoreticians, experimenters, and instrument builders, but also into hierarchies of group leaders, laboratory managers, and business coordinators. This complex management structure is evident in projects like multinational facilities like CERN. (3) *Social and Political Significance*: Big Science depends fundamentally on the attachment of social and political significance to scientific projects. Whether justified by contributions to national health, military power, industrial potential, or the new interlacements of science and technology, this continuous process of justification has influenced researchers' understanding of their work and ultimately its intellectual content (p. 357). (4) *Collaborative Research*: The products of large, long-term experimental collaborations have undergone complex social processes before emerging from laboratories (p. 357-8). Large scientific collaborations have changed the relationship of individual researchers to their experiments and the nature of individual contributions to knowledge creation.

Based on such characteristics of Big Science, various distinct ways emerge through which AI-based technologies in mathematics can be understood as a manifestation of Big Science in the 21st century. However, as we will show, terming AI-based mathematics as just another instantiation of Big Science would be misleading, since the rise of AI-based technologies in mathematics has far wider implications, specifically regarding the role of the mathematician and the production of mathematical knowledge.

One might argue that in a material sense, Big Science changed the scale and structure of science: it enabled the tackling of problems of previously unimaginable complexity, new forms of interdisciplinary collaboration, and a more bureaucratized, hierarchical organization of knowledge. But AI does not simply extend scale – it reconfigures the very *epistemic labor* involved. In traditional Big Science, humans remained the central epistemic agents, directing instruments, interpreting results, and managing theory. In AI-based mathematics, machines are not only verifying results, but increasingly suggesting novel conjectures, discovering patterns, and formalizing arguments. In some cases AI systems have discovered new and even unexpected results[7], thereby shifting the role of the mathematician from proving and discovering to evaluating, curating, and collaborating with machine agents.

This change raises several questions about authorship, authority, and expertise. Who is the epistemic subject when a machine produces a proof? How do we assign credit, responsibility, or trust in such systems? The mathematician is no longer a solitary genius, nor even a cog in a collaborative team, but becomes also more like a supervisor or a collaborator with semi-autonomous agents. This is a reconfiguration not just of scale or hierarchy, but of the human's relationship to mathematical knowledge itself (Ravetz 1971; Collins 1998). Cross-disciplinary entanglements of this kind also characterize other fields, including applied physics, computational biology, and data-intensive social science, many of which likewise emphasize ideals of autonomy and internal rigor. What makes mathematics a particularly revealing case, however, is that these ideals have long been woven into how the discipline understands its own identity, especially in relation to abstraction and proof-based standards of justification. As mathematicians increasingly collaborate with computer scientists, engineers, and AI researchers, the identity of mathematics as a self-contained field becomes tangible. The infrastructure of AI (training datasets, formal large languages models, neural network architectures) becomes part of the epistemic scaffolding of mathematical knowledge, which echoes arguments in the sociology of science that stress how material and social arrangements condition what counts as knowledge (Knorr Cetina 1999; Latour 1987).

It is important to note that we do not claim that AI-based mathematics does not share any characteristics with Big Science; it definitely does – elements such as scale, infrastructure, coordination are a crucial part of AI-based mathematics. We also do not argue that this shift is unique to mathematics – introducing AI-based technologies into the production of knowledge might have similar repercussions for physics or biology. What we do argue is that in mathematics, this shift is a qualitative one, as it distributes cognition across humans and machines, destabilizes traditional notions of expertise and intellectual authority, and forces us to rethink what it means for mathematics to be a human activity, if the entire body of mathematics (or just entire sub-disciplines of it, as plane Euclidean geometry) can be 'fed' into and consequently developed by a machine. A recent example of AI-based tools actively collaborating with the mathematician is seen with the conversations taking place during the beginning of October 2025 of the mathematician Terence Tao with ChatGPT, together providing a counterexample to a conjecture in analytic number theory.[8] Such an example, even though it does not consist of an entire proof found by AI-based technology, shows already that the integration of AI-based tools into mathematics combined with the shift to Big Mathematics reshapes basic notions found at the epistemological basis of the discipline. Such notions as mathematical agency, intuition, creativity, responsibility, objectivity, authorship and collaboration – all need to be redefined in light of Big Science and how

[7] For example, beyond formal theorem-proving systems, AI-based methods such as *AlphaTensor* (discussed above) and symbolic regression models have been used to uncover novel and unexpected mathematical structures. Symbolic regression systems aim to discover explicit mathematical expressions directly from data, often producing interpretable equations that were not previously hypothesized by human researchers, as demonstrated in applications to equation discovery in physics (Cranmer et al. 2020).

[8] See: https://mathstodon.xyz/@tao/115306424727150237 and https://mathoverflow.net/questions/501066/is-the-least-common-multiple-sequence-textlcm1-2-dots-n-a-subset-of-t/501125#501125. (last accessed on 31 January 2026). As Tao also notes, this "argument has now been formalized in Lean [...]." (ibid.)

it is interlaced with 'AI revolution'. In the next sections we analyze each of the characteristics of Big Science through an AI-based prism of mathematical practice, suggesting that we may be witnessing not just the 'biggification' of mathematics, but the emergence of a new, hybrid sociotechnical regime of mathematical knowledge production.

2.1 Concentration and Centralization of Computational Resources

The first characteristic of Big Science – the concentration of resources and data in specialized research centers – finds strong support in recent AI-mathematical research developments. Contemporary automated theorem proving systems like *TongGeometry* demonstrate this pattern, establishing "the most extensive repository of geometry theorems to date" by discovering "6.7 billion geometry theorems requiring auxiliary constructions" (Zhang et al., 2024). This scale of discovery is not only unprecedented for mathematical research, but it necessitates centralized, high-performance computing infrastructure that far exceeds traditional mathematical resources (even if some later testing can be run on smaller machines, building a repository of this size requires substantial centralized compute). But the significance of *TongGeometry* goes beyond infrastructure and scale: it exemplifies the distinctive epistemic roles AI-based tools can play, pushing mathematics into a novel mode. At this point it helps to clarify how *TongGeometry* and *AlphaGeometry* differ: AlphaGeometry is designed mainly to solve a fixed set of olympiad-style problems using a neural model and a symbolic prover. TongGeometry builds on related ideas but is presented as aiming at something broader: expanding the space of known theorems and proposing new problems, not only solving given ones. Unlike earlier projects, where human agents generated, communicated, and evaluated proofs, the authors of *TongGeometry* report that the system can autonomously explore large spaces of geometric relations and propose conjectures, some of which have been selected as olympiad-level problems in real competitions, while acting as a "geometry coach" rather than a mere solver (Zhang et al., 2024, Sections 4 & 5)[9].

As such, it embodies the distinctive contours of AI-based Big Mathematics – illustrating a shift not only in scale (as Big Science offered) but in *mode*: toward a sociotechnical regime where humans and machine agents collaborate in cognitive work, where proof-generation is automated and problem discovery becomes mechanized. Other examples illustrate this pattern: *AlphaGeometry*'s development required "large-scale synthetic data" (Trinh, 2024: 476) and sophisticated neural language models, while knot theory research using AI involved processing "more than 2.7 million knots" through neural networks (Davies et al., 2021: 71–2). Such computational demands necessitate the kind of resource concentration that characterized the emergence of Big Physics laboratories in the postwar period.[10] These examples represent a fundamental shift from mathematics as an individual or small-group endeavor to big collaborative constellations.

---

[9] Since *TongGeometry* is developed by a China-based team, it is also a reminder that the concentration of compute and data in AI-based mathematics has different national and political settings. The Big-Science pattern here is shaped not only by research labs and companies in general, but also by how different countries organize and support large-scale AI research. In practice, this can affect who has routine access to high-performance compute, what kinds of datasets and benchmarks are prioritized, and which institutions shape evaluation standards (for example, through funding structures, education and competition pipelines, and publication incentives). For the broader political implications of these differences see section 2.3.

[10] Large-scale AI training and inference relies on energy-intensive data centers and accelerated computing. The International Energy Agency estimates that data centres consumed about 1.5% of global electricity use and projects that

However, as was noted at the introduction, this is not the first time that there has been a shift from individual work to collective work in the history of mathematics. The idea that mathematics is a field where the processes of discovery and proof were primarily associated with individual efforts was challenged with the rise of mathematical groups and collaborative efforts, such as the works of Bourbaki and the Classification project of finite simple groups (Gorenstein 1982; Steingart 2012). Waldecker and Remmert underline that the Classification of Finite Simple Groups (CFSG) may represent one of the first examples of what could be called Big Mathematics – a large-scale, collaborative, and institutionally mediated mode of mathematical research that departs significantly from the traditional image of the lone mathematician pursuing isolated problems[11]. Another example of such a collaborative project is the Busy Beaver Challenge, an international effort to determine the precise value of BB(5), a famously extremely hard-to-compute function rooted in Turing machine theory[12]. Like the CFSG, the project exemplifies Big Mathematics in both scale and structure: it mobilizes a large, decentralized network of contributors, relies on shared open-source infrastructure, coordinates complex computational workflows, and integrates formal proof assistants (like *Coq* and *Lean*).

These works – both of pre-AI-based mathematics as well as those using AI-based technologies – suggest that other social entities, such as scientific communities, groups of mathematicians, and collectives, can produce knowledge together, as a mutual and collaborative act, and this knowledge cannot be pinned down or reduced to being originated from one sole individual mathematician. These collective efforts suggested that a mathematical proof is not necessarily the product of one mathematician, but can be a collective effort where the participating actors are *also* machine based, i.e. non-human actors, in which no single mathematician has a global view of the proof or can attest to its veracity. In this sense, in a similar fashion that the CFSG and, in a smaller scale, the Bourbaki group, are to be considered as examples of 'Big' collaborative constellations, so do corporations of mathematicians, programmers and non-human actors.

2.2 Hierarchical Organization and Specialized Division of Labor

The second characteristic of Big Science – hierarchical organization with specialized workforce divisions – manifests clearly in AI-mathematical research through the emergence of new professional roles and management structures. As Galison and Hevly emphasize, Big Science brought together scientists, engineers, technicians, and administrators into coordinated, large-scale enterprises, often structured around lab-based hierarchies and project bureaucracies that managed

global data-centre electricity consumption could roughly double that by 2030, with accelerated computing for AI as a key driver (IEA, Energy and AI, 2025). For the United States, the Department of Energy notes that data-center electricity demand has grown rapidly and is projected to double or triple by 2028, underscoring that compute centralization is also an energy-and-infrastructure issue (Shehabi et al., 2024; Strubell, Ganesh, and McCallum 2019)

[11] See: https://www2.mathematik.uni-halle.de/waldecker/index_CFSGhistory.html and : https://gepris.dfg.de/gepris/projekt/528127386?language=en

[12] BB(5) refers to the 5-state, 2-symbol Busy Beaver value: informally, the maximum number of steps taken (starting from a blank tape) by any 5-state Turing machine that eventually halts. The Busy Beaver function is noncomputable in general (there is no general algorithm that takes n and outputs BB(n)), but specific small-n values can sometimes be established through extensive enumeration plus mathematical and formal-verification work. For BB(5), a long-standing lower bound of 47,176,870 steps was found in 1989 (Marxen & Buntrock 1990), and in 2024 the bbchallenge.org massively collaborative project proved BB(5) = 47,176,870, including a Coq/Rocq-verified proof ("Coq-BB5"). For more details see: https://bbchallenge.org

the division of conceptual, technical, and interpretive labor (Galison & Hevly 1992: 2, 11-12). Similarly, in AI-based mathematics, the work is no longer conducted solely by individual mathematicians but by distributed teams that include AI engineers, formal methods specialists, proof assistants, interface designers, and mathematical theorists. This layered collaboration reflects a "multidimensional framework" of intellectual authorship, with distinct contributions across axes such as content generation, structural assistance, creative input, and analytical reasoning (Hutson 2025). These differentiated roles are not merely functional but often hierarchical: senior mathematicians may define high-level goals, while AI experts and tool developers implement, test, and refine system outputs, and others validate results or translate them back into human-readable mathematical arguments. As in Big Science, this distribution of labor demands new coordination mechanisms, quality control protocols, and evolving norms of authorship and accountability.

Recent work on automated theorem proving reveals additional layers of specialization, where researchers must develop expertise in "premise selection, proof guidance in several settings, AI systems and feedback loops iterating between reasoning and learning" (Blaauwbroek et al. 2024). This division and restructuring of labor creates distinct hierarchical levels reminiscent of the complex management structures described in Big Physics. The example of the AI-based technology *FunSearch* and its combinatorial discoveries of various sets (Romera-Paredes et al. 2024) demonstrate how this integration creates new forms of mathematical knowledge production that require both human interpretation and machine computation.[13] *FunSearch* does not go over various sets, but rather combined "several programs sampled from the programs database […]. The prompt is then fed to the pretrained LLM and new programs are created." (ibid., p. 469) More explicitly, the *FunSearch* project exemplifies this specialization: mathematicians guide prompt engineering, while designing effective prompts is crucial for obtaining the desired results from LLMs and is described as "nontrivial" (ibid, "Methods" section). This mirrors the novel labor structures described by Galison and Hevly, where different groups specialized in theoretical work, experimental design, and instrumentation. Similarly, in learning-guided reasoning, we see coordination between human-designed heuristics, machine learning models that guide inference, and systems architects who manage feedback loops between proving and learning (Blaauwbroek et al. 2024, 19). These efforts demand integrated team structures, parallel computational infrastructure, and shared conditions that are associated with Big Science.

But beyond changing hierarchies and creating complex management structures, the collaborative dimension of AI-based mathematics transforms mathematical practice itself, as the 'working mathematician' is no longer just an individual or a group of mathematicians; rather, mathematical labor is now distributed between humans and machines. Originally, the term 'working mathematician' was sought to characterize the mathematician as a practitioner actively engaged in research, focusing on applicable methods rather than purely abstract theorization or philosophical

[13] Such an example comes from the area of combinatorics. A *cap set* is defined as a subset of the n-dimensional affine space over the three-element field, where no three elements sum to the zero vector. The *cap set problem* is the problem of finding the size of the largest possible cap set, as a function of n. Determining whether a cap set exists from a collection of given elements is known to be NP-complete (yet it is known how to quickly obtain cap sets of size 2^n). In 1970, Giuseppe Pellegrino (Pellegrino 1970) proved that four-dimensional cap sets have a maximum size 20 (note that $2^4 = 16$). This result led to finding lower bounds larger than 2^n for any higher dimension. Such a lower bound was improved to be 2.218^n by Tyrrell (Tyrrell 2022). However, Romera-Paredes et al. (Romera-Paredes et al. 2024) managed to improve the lower bound by using an LLM termed *FunSearch* with an evaluator, to 2.2202^n.

considerations (Bourbaki 1949; Mac Lane 1971; Corry 2004). However, as AI-based programs increasingly undertake tasks traditionally performed by mathematicians, the distinction between human and machine labor in mathematics becomes increasingly tenuous. The question now is not just whether these programs contribute to mathematical discovery, but whether they can meaningfully be considered 'working mathematicians' in the way this category has historically been understood.

This new perspective of the 'working mathematician' as referring to a collaborative act between humans and machines parallels the way Big Science transformed the nature of experimental practice through human–machine collaborations. However, in contrast to experimental physics, where machines tend to produce data for humans to interpret, the logic of AI-assisted mathematical discovery includes a deeper entanglement: machines can now guide, propose, verify, and refine the very inferential steps that constitute mathematical reasoning. Learning-guided automated reasoning systems, for instance, like *AlphaGeometry* (Trinh et al., 2024: 478–480), incorporate feedback loops where deductive logic not only supports machine learning but is itself shaped by it, suggesting a hybrid epistemic agent composed of both human and machine contributors. This new notion of hybrid mathematical agency requires rethinking traditional concepts of authorship and authority in mathematics, as well as embracing a new image of the mathematician: as a participant in a *sociotechnical* network of humans and machines rather than a collaborator in a *social* network comprised only of other humans, in which computer programs are considered as a mere aid for various calculations but *not* as generative agents. In this sense, the integration of AI-based tools into mathematical practice not only prompts constellations which can be termed Big Mathematics (as a parallel to Big Science) but also transformed the way we think about the role of the mathematician in discovering, creating, and producing new knowledge.

2.3 Political and Social Significance

The third characteristic: attachment of social and political significance to scientific projects, emerges powerfully in contemporary AI-based mathematical research through concerns about institutional power, accessibility, and epistemic authority. The question arises whether AI-based mathematical research risks becoming concentrated within elite institutions that have privileged access to computational resources, echoing other political dimensions of historical Big Science projects. One of the most striking indicators of the political economy at work in AI-based mathematics is compute scale and resource concentration. For example, DeepMind reports that to train *AlphaGeometry* it produced 100 million unique synthetic geometry problem instances to train its neural-symbolic pipeline[14]. Meanwhile, in the case of *AlphaTensor*, DeepMind's technical reporting notes that the search space for relevant tensor decompositions was "more than 10^{33} possible moves at each step"[15]. While such concentration of resources per se does not have to necessarily imply influence of political power, one aspect should be underlined: these systems ran on Google's custom TPU pods, which are racks of thousands of specialized chips, representing millions of core-hours of computation. Such resources, including the engineering teams that 'feed' these problems and maintain these TPUs, remain far beyond the reach of typical mathematics departments in state universities and colleges. In this sense, AI-mathematical discovery has

---

[14] https://deepmind.google/discover/blog/alphageometry-an-olympiad-level-ai-system-for-geometry/
[15] https://spectrum.ieee.org/matrix-multiplication-deepmind

become an infrastructural game: access to accelerators, bespoke data pipelines, and massive human-machine engineering teams is itself a gate to participation.

To emphasize: The question concerning the existence of political influence is justified, however, since several of the known discoveries by AI-based technologies, as *AlphaGeometry* or *AlphaTensor*, were developed *by* Google. The recent discovery of Terrance Tao, mentioned above, was done due to a conversation with ChatGPT, a technology owned by *OpenAI*. The involvement of companies in the research follows patterns similar to historical Big Science[16]. Their agenda is, almost banal to note, not at all neutral or free of financial, reputational, and infrastructural motives. Cathy O'Neil's *Weapons of Math Destruction* (2017) already warned us that while algorithmic systems are presented as objective tools for advancing knowledge, they often reinforce asymmetries of power and accountability, a concern now mirrored in the concentration of AI-based mathematical discovery within private tech laboratories. Although such agendas emphasize the potential of their technologies to advance mathematics, discover new knowledge, and push the boundaries of human understanding, they simultaneously shape which forms of mathematical inquiry receive attention and resources, thereby embedding commercial and institutional priorities within the epistemic fabric of the field (Jaton 2021). This agenda is also promoted with the new forms of science-technology integration that fundamentally reshape research practice,[17] an aspect that represents perhaps the most profound parallel between AI-based mathematics and historical Big Science.

Needless to say, such discussions have implicit but highly clear political dimensions, as recent research on AI literacy in higher education emphasizes that "the potential for continued inequities, exclusion and divides must not be ignored" (Van Wyk 2024). The concentration of AI-mathematical capabilities raises questions about institutional power, research accessibility, and research policy that parallel the political concerns surrounding historical Big Science projects. Just as Cold War-era physics was dominated by well-funded laboratories backed by government and military institutions, contemporary AI-mathematics research tends to cluster in elite academic departments, corporate AI labs, and a small number of globally dominant research hubs with access to high-performance computing and proprietary models. Such a concentration is obviously embedded in the AI research in general, which is heavily funded by e.g. the US government (Taylor & Stan 2024)[18]. This raises a few more fundamental questions: Who sets the research agenda in AI-based mathematics? Who has the computational capacity to participate meaningfully in large-scale formalization or discovery efforts? Which languages or formal systems are encoded into these tools, and whose mathematical traditions are left out? These are not only technical

[16] As recent analyses of corporate AI research have shown, these developments are embedded in broader political economies of data, computation, and influence (Crawford 2021; Birhane 2021)

[17] In the bigger social framework, this agenda is already noted in Sam Altman's manifesto from 2021: https://moores.samaltman.com/, calling for a "new social contract".

[18] For example, an analysis of federal AI and IT R&D spending estimates that the federal government allocated about US$3.316 billion for core and cross-cut AI funding in Fiscal Year 2025, across agencies including NIH and NSF showing significant public funding flows to US elite institutions (https://federalbudgetiq.com/insights/federal-ai-and-it-research-and-development-spending-analysis/). In addition, in the FY 2025 budget supplement for the National Artificial Intelligence Research Institutes program it is noted that by 2023 and 2024 federal agencies had invested approximately US$118.5 million and US$69 million, respectively, in these AI research institutes (as can be seen in this appendix: https://www.nitrd.gov/pubs/FY25-National-AI-Institutes.pdf). Moreover, as documented in the National Security Commission on Artificial Intelligence (NSCAI) report, the commission recommends tripling the number of federally funded national AI research institutes (https://reports.nscai.gov/final-report/chapter-11)

concerns, but political ones, tied to how research capacity, prestige, and epistemic authority are distributed across institutions and regions.

For instance, initiatives like *AlphaTensor* or *FunSearch* are developed by teams with privileged access to custom-designed TPUs and massive training datasets, supported by proprietary platforms and extensive engineering teams[19] (Fawzi et al. 2022; Romera-Paredes et al. 2024) – resources that are not necessarily accessible to every mathematics department or research institute (Urman, Smirnov & Lasser 2024; Besiroglu et al. 2024). This asymmetry may lead to mathematical innovation being increasingly channelled through a narrow circle of well-resourced institutions, while others risk positioning themselves as downstream consumers of tools and outputs they cannot fully shape or interrogate[20] (Ahmed & Wahed, 2020). Such dynamics can reinforce knowledge monopolies and undermine the public accountability of scientific innovation (Wink et al., 2024).

Research accessibility is shaped not only by infrastructure but by epistemic entry points, since without adequate training, resources, and localized expertise, entire regions may be excluded from meaningful participation in AI-mediated fields – including mathematical research (Wall, Saxena, and Brown 2021). This makes AI-based mathematics not just a question of technological progress, but one of educational justice and global epistemic equity. From this viewpoint, questions of research policy emerge: Should AI-based mathematical tools function as public infrastructure or proprietary assets? While one can claim that Terence Tao's 'conversations' with ChatGPT are open to anyone having access to this technology, this is hardly true concerning the technology underlying, for example, *AlphaTensor*. That is, the question arises, should funders prioritize access and interoperability? Do the results of the mathematical research led by such corporations truly belong to the general public?[21] These structural questions extend the Big Science analogy but also go beyond it, suggesting that AI-based mathematics increasingly shapes political priorities and future research trajectories.

2.3.1 Agency

Related to questions of social and political significance but somehow differently is the concept of agency. As is clear, integrating AI into mathematical practice crosses the boundary between human and machine cognition (Bressan et al., 2024), a crossing which by itself has major political implications. Here, the political stakes are not only about access to funding or infrastructure but about control over the epistemic norms themselves: who (or what) gets to reason, validate, and explain mathematical truth and objectivity, leading to the ultimate question about what is

[19] As documented in DeepMind's own publications: https://deepmind.google/discover/blog/discovering-novel-algorithms-with-alphatensor/ and https://deepmind.google/discover/blog/funsearch-making-new-discoveries-in-mathematical-sciences-using-large-language-models/

[20] For example, Carnegie Mellon University has been awarded ~US $225 million in NSF grants for AI-related research over 14 years (more than any other U.S. university in that dataset - https://sciencebusiness.net/news/carnegie-mellon-uc-san-diego-top-us-grant-winners-ai-research) and institutions like MIT, Harvard, USC, Johns Hopkins, and Stanford University are making substantial pledges (in billions) for AI initiatives (as can be seen in this overview: https://profalexreid.com/2025/09/22/major-university-led-ai-initiatives-and-their-focus-2019-2025/).

[21] Such a question is not trivial; for example *AlphaTensor* developed by Google has discovered new, more efficient ways to multiply matrices, being one of the main operations at the basis of everyday computer calculations.

mathematical agency in AI-based Big Mathematics. It is clear that the usage of AI-based technologies has prompted a lively (philosophical) discussion regarding whether such technologies are agents, as well as a shaft in how one should consider the notion of agency in the 21st century (see e.g. Floridi 2025). However, in mathematics, such a discussion has hardly taken place.

Most discussions on agency in mathematical practice refer to mathematical activity as human activity, focusing on what *people* do when they do mathematics (Hamami 2023; Mancosu et al. 2005). In these discussions, mathematical agents must possess specific knowledge, skills, and comprehension to carry out the relevant mathematical activities. They are primarily characterized as human agents with typical cognitive capabilities and constraints (Ferreirós 2016). Mathematical agents are portrayed as planning entities with the ability to use manipulative imagination, execute inferential steps, and comprehend the underlying reasoning behind mathematical processes (De Toffoli and Giardino 2014; Larvor 2019; Chemla 2015; Weber and Tanswell 2022). While some acknowledge that computers can play a role within mathematical agency, there has been little substantive exploration of how AI systems that generate original and creative proofs (and not only perform calculations) fit into this framework, or the implications of such an integration.

Indeed, we claim that the problematic issue or rather the blind spot with the above accounts is that they mostly view mathematics as a human creation, deeply rooted in the human community (even if using computers for calculations or for verifying proofs, as with *Lean* or *Coq* eventually leading to the improvement of proofs). However, the increasing reliance on AI challenges this view. The challenge is not only theoretical or conceptual; it also has methodological consequences. Our fundamental notion of mathematical agency partially shapes how we analyze specific mathematical activities and practices, which, in turn, can generate new empirical data that may refine and reshape our theoretical perspectives. In practice, the mathematician as the sole active agency is being replaced by a more collaborative model, where human and machine contributions intertwine. This shift forces us to reconsider the essence of mathematical activity and its connection to human cognition and creativity, challenging the idea that mathematics and mathematical knowledge are rooted or grounded into elementary human activities (Kitcher 1984; Ferreirós 2016). At the same time, if we acknowledge machines as active participants in mathematical discovery or even as agents, we must develop a deeper understanding of the nature of the knowledge they generate. This represents a new, hybrid form of understanding that cannot be regarded as purely non-creative, yet also cannot always be fully verified or validated.

Interactions with AI-based programs as producers of new knowledge requires mathematicians to develop new skills and adapt to new modes of working. In this new kind of mathematical practice, mathematicians may focus on developing and refining AI systems, or specialize in interpreting and applying their outputs. On the one hand, it creates more roles for mathematicians in collaborative work with AI systems, such as creating and configuring prompts (Romera-Paredes et al., 2024). On the other hand, it may take away from the mathematician the role of discovery. The "working mathematician" is no longer just an individual or a group of mathematicians; rather, mathematical labor is now distributed between humans and machines.

These shifts in agency, both conceptual and practical, bring us back to the broader picture of AI-based Big Mathematics: as machines increasingly contribute to various mathematical tasks once reserved for humans, we see not just a redistribution of cognitive labor, but a transformation in the very structure of mathematical agency and knowledge production. This new mode of practice exemplifies the convergence of Big Science dynamics with a fundamentally novel, hybrid sociotechnical regime, reshaping what it means to do mathematics, who is entitled to participate, and how mathematical truth is constructed and recognized.

2.4 Collaborative Research: Creativity, Authority, Responsibility

Collaboration is deeply embedded in all the characteristics discussed above and plays a central role in both institutional and Big Science. Large-scale scientific projects have reshaped how individual researchers relate to their work and have transformed the meaning and visibility of individual contributions within the collective process of knowledge production. In Big Mathematics, collaborations can take two forms: large-scale efforts involving individuals, groups, institutions, and collectives - all human-based - and hybrid collaborations between humans and AI-based technologies.

Hybrid human-AI collaborations in mathematics bring to light critical limitations in current AI systems. Studies examining how AI-based technologies are approaching mathematical problems suggest that, although these tools may introduce new questions about the nature of proof and reasoning, they often operate through brute-force search or statistical pattern recognition, typically restricted to statements of low logical complexity (Dean and Naibo 2024)[22]. For example, recent research on reinforcement learning (RL) in combination with large language models (LLMs) shows that, despite some progress, these models still struggle with compositional reasoning and the integration of multiple strategies. As one study notes, "RL's impact on compositional tasks remains modest: models still struggle to integrate multiple learned strategies coherently. […] RL struggles to induce genuinely new reasoning patterns, showing negligible progress on transformative generalization that requires shifting to novel solution paradigms" (Sun et al., 2025). These findings underscore a fundamental tension in hybrid collaboration: while AI can extend the range and scale of mathematical exploration, it does not yet replicate the kinds of creative, conceptual shifts that characterize deep mathematical insight. However, even if incapable (yet) of creating 'creative leaps' as the creation of new concepts, the above research does challenge traditional notions of mathematical creativity and raises questions about whether machine-generated proofs can be considered 'creative ideas' in the traditional sense.[23]

Though we do not want to open a full-length debate on mathematical creativity, such creativity is a complex notion, typically described as the capacity to generate original work that meaningfully expands the field of knowledge, introduces new lines of inquiry, and leads to unconventional, innovative, or insightful solutions to existing problems (Liljedahl & Sriraman, 2006; Nadjafikhah, Yaftian & Bakhshalizadeh 2012; Boden 2004; Hadamard 1945). Some of the literature suggests

---

[22] Dean and Naibo base their claim on three "AI-inspired" families of technologies they examine as case studies: (i) traditional automated theorem proving (e.g., resolution/unification-based provers) as in the computer-assisted resolution of the Robbins problem (pp. 17–20); (ii) SAT-solving methods (including modern SAT-solvers) as in SAT-based resolutions of finite combinatorial conjectures, their key example being the Boolean Pythagorean triples problem (pp. 21–23); and (iii) large language models used alongside search/optimization procedures in recent work that improves bounds in additive combinatorics (their example is the cap-set lower-bound result obtained with an LLM in concert with other techniques, pp. 24–28).

[23] Note also that Timothy Gowers has recently noted that a solution to a problem in Euclidean geometry proposed by *AlphaGeometry 2* involved a "non-obvious construction". See: https://deepmind.google/discover/blog/ai-solves-imo-problems-at-silver-medal-level/

that the foundation of creativity lies in the ability to think about a problem or situation through analogy (Ghiselin 1952). Other studies argue that employing specific cognitive tools such as imagination, empathy, and embodiment can foster creative outcomes (Csikszentmihalyi 1996; Liljedahl and Allan 2013). In all these instances, the underlying theory maintains that the final articulation of a creative idea results from the conscious and intentional efforts of the human mind. The integration of AI-based programs and technologies into the mathematical creative process creates a dissonance with respect to ownership and leadership (as well as a conceptual and definition problem): instead of humans leading the creative process, the AI-based technologies present solutions that mathematicians did not think of or that were unexpected.[24] To what extent, if at all, can these machine-generated proofs be regarded as creative ideas? If they are deemed creative, then a redefinition of mathematical creativity (and creativity in general) becomes necessary. If they are not considered creative, the mathematical community must determine how to address machine-generated solutions that undeniably contribute to the expansion of knowledge. Are they deemed non-creative solely because they originate from a non-human agency?

Another point that becomes acute within the collaborative space between human and AI-based reasoning is the question of mathematical authorship. Traditional notions of authorship in mathematics are grounded in the belief that mathematicians alone are responsible for the creative act of producing mathematical knowledge and expressing mathematical intuition, when the tools employed are merely passive ones and play no active role. While the view that individuals and collectives have social responsibility for the fundamental truths of mathematics was once considered a controversial and provocative view, it was based on the idea that this responsibility is collective in the sense that it builds on social critique (Bloor 1978; Bos and Mehrtens 1976; Restivo 2001; Ernest 1998). However, as human-machine collaboration becomes more prominent, social critique becomes increasingly dispensable, particularly because human mathematicians often cannot independently verify machine-generated proofs as was mentioned above.[25] As AI-based programs become increasingly capable of generating such proofs, they challenge not only the long-standing assumption that the mathematician is the sole author of these intellectual products, but also the idea that collaborative work and peer reviews are essential for discovering mathematical truths. Instead, what emerges is a more complex and hybrid notion of authorship that integrates human intuition,[26] computational power, and machine intelligence into a shared process of discovery and creation. Such a shared process of knowledge creation challenges the common discussions about the social constitution of knowledge, as there is now a new, non-human player, that creates knowledge that can not necessarily be verified. The problems of verification and surveyability that had already surfaced in earlier computer-assisted proofs (Parshina 2023) reemerge in AI-based Big Mathematics, as contemporary research highlights tensions surrounding the 'black box' nature of AI-mathematical systems, in which "the exact way these LLMs are operating is not disclosed to users," as stated by the authors of a recent paper on "Generative AI Glitches" (Srdarov and Leaver 2024).[27]

[24] Floridi (2025) argues that AI-based technologies should (still) be considered as lacking intelligence.

[25] This is the non-surveyability problem, see: Berg & Tanswell, 2025.

[26] See: Friedman & Kish Bar-On, (in review).

[27] Such a critique on the 'black-box'-ness is not unique to AI, but was already emphasized with respect to algorithms themselves, as O'Neil (2017) notes.

This shift toward human-machine collaboration prompts not only questions of authorship, but also pressing concerns about epistemic authority. As AI-based technologies assume the role of co-authors, they simultaneously become candidates for epistemic authority in mathematical discovery. 'Epistemic authority' is the capacity of an individual, group, or system to be recognized as a legitimate and trustworthy source of knowledge, such that others are justified in deferring to its claims or judgments within a specific domain (Jäger 2025; Zagzebski 2015; Hauswald 2021). To better understand the implications of Big AI-based mathematics on knowledge creation, we must explore how epistemic authority should be reconceptualized in hybrid human-machine networks by (a) investigating which types of technologies, individuals, or institutions justifiably hold such authority in AI-mediated mathematics[28], (b) considering how epistemic authority is socially constructed and institutionally legitimized (following Gieryn's (1999) extension of Weber's concept of domination), and (c) asking what kind of epistemic goods can be transmitted from AI-based systems to human researchers. Is it knowledge, understanding, or merely instrumental reliability? Understanding (and maybe believing too) requires grasping systematic connections, which may be difficult when dealing with such black-box models. Yet if mathematicians develop new ways of working with, interpreting, and contextualizing AI outputs, we may begin to speak of a new form of 'distributed epistemic understanding', one that blends machine-generated insight with human conceptual framing. Clearly, these questions are not only epistemological but also political, and determining who or what can serve as an epistemic authority shapes access, trust, and legitimacy in the emerging field of AI-based Big Mathematics, making it all the more important to examine these forces right now.

* * *

As this section suggests, collaboration in mathematics – whether big, small, AI-based or human based – have many faces, and therefore many implications on other domains in mathematical practice and research. After surveying these four central characteristics of Big Science in light of developments in AI-based mathematics, two things become clear: (1) each of the four is present in mathematical research (turning it into Big Mathematics), and (2) all of the four characteristics are affected by the integration of AI-based tools into mathematics, to various degrees.

3. Discussion: The implications of AI-based tools on mathematics as Big Science and Beyond

As a reminder, our objectives in this paper were to show that (a) AI-driven mathematics inherits the large-scale infrastructure, organizational complexity, and coordinated effort characteristic of Big Science, and (b) in doing so, it introduces distinct forms of task distribution that fundamentally alter the epistemic structure of mathematical practices, making it something different than just big scale mathematics. While the first objective has been addressed through the detailed analysis in the preceding sections, the second objective requires further elaboration to fully capture its implications beyond mere scale.

[28] Jäger's distinction between objective and subjective accounts of authority, as well as his concept of "recognized epistemic authority," may help differentiate between authority grounded in actual epistemic superiority (e.g., an AI's performance record) and authority based on human users' trust or perceived expertise (Jäger 2025, 6).

The deeper implication is political. If AI-based mathematical practice inherits the infrastructures of Big Science, it also inherits the politics of infrastructure. The fact that a small number of institutions now possess the infrastructure and the technology which enables them to generate conjectures, propose candidate theorems, test them at scale, and narrate them as advances in mathematics, is not philosophically neutral; it is the construction of a gate. It has been argued that large-scale computational systems are rarely merely technical solutions, but are governance devices that decide who participates and on what terms. We caution that some mathematical practices may be moving toward a model in which access to discovery and its underlying technology becomes increasingly gated[29], shaped not simply by talent or insight, but by the availability of proprietary large language models, high-performance accelerators, massive curated proof datasets, and the specialized labor required to orchestrate them.

This has two consequences. First, mathematical agency becomes stratified. Some actors, such as well-funded labs, elite collaborations, and researchers embedded in industry–academic hybrids, gain the capacity to direct discovery and frame it as discipline-defining. Others, including many mathematics departments without access to TPU-scale infrastructure, are effectively repositioned as interpreters, validators, or pedagogical downstreamers. Second, legitimacy itself begins to migrate. Results that emerge from corporate AI systems are increasingly framed not only as technically impressive, but as socially inevitable (suggesting that "this is where mathematics is going"). That rhetorical move matters, as it positions corporate infrastructures as the natural home of the future of mathematics, and everyone else as lagging behind. This suggests that in AI-based mathematics, the decisive instrument is no longer a particle accelerator or a space telescope, but an opaque cognitive pipeline which we don't seem to understand all the way through[30].

Each of the four characteristics discussed in the previous sections highlighted a distinctive aspect of AI-based mathematics in terms of traditional Big Science; yet each of them also introduces changes that transform Big Mathematics into something qualitatively different from traditional Big Science applied to mathematics. *Concentration and Centralization* point to a shift in the *mode* of doing mathematics, representing a new hybrid sociotechnical regime where humans and AI-based technologies collaborate in creating new proofs and theories in non-traditional ways. This hybrid ecology of humans and machines working together to develop theories and verify proofs (as seen in projects like *FunSearch* and *TongGeometry*) signals a major change in how we understand and define *mathematical creativity*, as it is no longer the sole product of the human mind. If creativity can be simulated or mechanized within distributed systems, then mathematics is no longer the privileged domain of human reason alone, which raises deeper concerns about the ontological status of mathematical objects and truths. This shift suggests a move away from mathematics as an endeavor rooted in human intentionality and creativity toward one in which

[29] That is not to say that mathematics was not gated before, but in different ways. As we mentioned in the introduction, the social gates of mathematics were not open to women until the mid twentieth century, nor to other minorities or outcasts of society even before that (Kenschaft 2005; Henrion 1997). AI-based technologies form other kinds of gates, but this is clearly not the first encounter of enclosure in mathematics.

[30] One might object that no individual scientist fully "understands" a particle accelerator or a space telescope either. This is surely right at the level of individual expertise. However, in these cases there typically exists a well-defined community of specialists who can, in principle and often in practice, give a relatively complete, mechanistic account of how the instrument functions and why it yields the data it does. This locus of distributed but collectively accessible understanding is precisely what seems to be under strain in the case of contemporary AI-based models, where even the engineers and scientists who develop them may lack a clear explanation of how particular outputs are generated or why certain solution paths are favored (O'Neil gives several examples in her book (2017)).

knowledge emerges from complex, often inscrutable interactions between formal systems, statistical learning, and human interpretation.[31]

From a sociological perspective, this shift involves more than a transition from individual to collective work, or from small-group to large-scale collaboration; it introduces new forms of *sociotechnical interactions* among individuals, research communities, and machines. This is however not surprising, as every integration of new technology introduces these new forms. In AI-based mathematics, these interactions unsettle established norms of authorship, responsibility, and trust: when no individual can verify a result or fully explain how it was produced, authority shifts from the human mathematician to the credibility of the entire system - to its infrastructure, design, and institutional embedding. However, these sociotechnical interactions also provide a space for collaborative ideas and knowledge exchange between humans and machines unlike any kind of interaction seen before. Changes of this kind might also disrupt the conventional narrative of mathematics as a cumulative, human-centered endeavor and instead suggest that future histories of mathematics may need to document not only the development of theories, but also the architectures of AI-based computation, the training regimes of models and the enormous data sets, and the sociotechnical contexts in which discovery occurs. Seen this way, AI-based Big Mathematics is not just a story about 'smarter' or more efficient tools; it is a story about enclosure of knowledge, since such knowledge risks being partially relocated into privately governed infrastructures whose incentives are largely reputational, competitive, and commercial.

Concerning *Hierarchical Organization* AI-based mathematical labor may include AI engineers, interface designers, mathematical theorists as well as proof assistants as *non-human actors*, and this reflects a new *division of labor* of various contributors, all part of a layered collaborative effort to produce new knowledge. The intriguing element of this change is its effect on the very notion of the 'working mathematician' which undergoes a conceptual as well as practical transformation: from an individual or a group of mathematicians doing the 'work' to a collaborative act of humans and machines that not only work together, but also take turns in shaping the project. In the framework of AI-based mathematics, the definition of the working mathematician, as it was understood historically and sociologically, is slowly becoming irrelevant; there is now a need for a new definition since the work is being done by humans and AI-based technologies, collaboratively. This shift destabilizes the assumed unity of the mathematical subject since the working mathematician becomes less a discrete figure and more a *relational position* within an evolving epistemic network, defined as much by interface and orchestration as by proof and reasoning. In this new regime, 'doing mathematics' no longer means exercising internal conceptual mastery alone; it also means learning how to collaborate with non-human agents, interpret their outputs, and shape inquiry around tools that increasingly influence the trajectory of the discipline itself.

The *social and political significance* of AI-based mathematics comes with the prices of asymmetrical research resources, rich closed research circles, and political identities of those who get to 'do' mathematics and their social and institutional affiliations. But embedded within this shift is a deeper epistemological transformation in the very concept of *mathematical agency*: it is no longer attributed solely to the individual mathematician, but is instead distributed across

[31] Note that we do not claim that such a shift has happened in mathematics only with the rise of AI-based technologies: the proof of the 4-color theorem in the 1970s already promoted similar discussions (Parshina 2023).

human–AI-based systems, which means that the identity of the mathematical agent becomes layered and contingent. This raises a tension for the philosophy of mathematical practice: can agency still be the organizing principle of mathematical reasoning if that agency is no longer wholly human, intentional, or interpretable? What becomes of mathematical explanation or understanding when the generation of mathematical results is mediated by tools that do not themselves possess these qualities, yet shape their conditions? The change in the concept of agency from human to hybrid also means shifts in what it means to know, to understand, and to act within mathematics, restructuring not only who gets to do mathematics, but what it means to do mathematics at all. Here it should be emphasized that before the rise of AI-based technologies, also human mathematicians employed hybrid reasoning, but the agency did not change due to this hybridity.[32] Here the shift is more substantial: it is due to the hybrid nature of agency that hybrid types of reasoning emerge.

Collaborative Research, which is woven into each aspect of AI-driven Big Mathematics, introduces changes in authorship and epistemic authority. These changes represent a fundamental shift in the architecture of knowledge production and undermine a central tenet of methodological rationalism: that epistemic authority depends on the capacity for justification and explanation (Wilholt 2013; Hauswald 2025; Croce & Baghramian 2024). If AI-generated outputs cannot be explained in the same terms as traditional proofs, yet they produce results accepted as valid, we must ask whether comprehensibility is still an epistemic requirement, or whether new norms of trust, delegation, and reproducibility are displacing older ideals of rational understanding. The emergence of *AI-based co-authors* requires us to historicize (future) mathematical knowledge as increasingly dependent on non-human agents whose contributions are not communicative in the traditional sense, but are still original and creative. Collaborative human-AI authorship also forces a revaluation of legitimacy and accountability in the distribution of epistemic roles: if systems gain authority through performance rather than intelligibility, then we face a shift from consensus-driven validation (where human experts evaluate results) to algorithmic closure (where the credibility of a result rests not on understanding or debate, but on its verifiability and reproducibility through automated procedures). The shift in authorship and authority leads also to a hybrid notion of human-machine collaboration in mathematics, which in its turn reinforces a different approach to mathematical practices, as something that: (1) are not necessarily rooted in human conventions, (2) can be detached from their (human) contributor, and (3) refer to a form of knowledge we might not be certain about, as noted above. That is, it is not clear whether an AI-generated proof can be considered a mathematical object in the same sense that human-generated proof can, and to what extent it is socially constituted, if at all.

* * *

To conclude we would like to highlight the accelerating redefinition of what it means to "collaborate" in mathematical research, which emphasizes why this discussion is both timely and important. To do so, we review three talks held in the recent 2025 *Big proof: formalizing mathematics at scale* workshop held at the Isaac Newton Institute at Cambridge (9 to 13 June 2025),[33] all of them exemplify the above discussed shifts. First, Mateja Jamnik in her talk "AI for

[32] For example, see the research of the braid group between 1925 and 1950 in: Friedman 2019.
[33] See https://www.newton.ac.uk/event/bprw03/

Math: The Future of Collaborative Discovery" described a vision of human-AI collaboration grounded in interactivity, where AI systems serve as conjecture generators, proof assistants, and exploratory partners in a dynamic loop of feedback and revision[34]. Her empirical study with professional mathematicians emphasized that even flawed or incorrect AI outputs can be perceived as "helpful" when they inspire novel directions or heuristics, suggesting a shift from correctness to utility as a measure of epistemic contribution. This aligns with the sociological and philosophical shifts discussed above: collaboration here is a co-evolutionary process in which human and machine reasoning mutually shape the trajectory of mathematical discovery.

This was not the only model proposed for human-AI collaboration. In contrast, Terence Tao's talk described a different mode of large-scale collaboration, one that relies on human–machine symbiosis at the infrastructural and organizational level rather than cognitive partnership[35]. The *Equational Theorist* Project he led did not depend on deep reasoning from AI systems, but instead on a division of labor between humans, automated theorem provers, visualization tools, and a modular proof-verification architecture. AI tools played a largely instrumental role: they generated counterexamples, verified lemmas, or converted outputs between formal systems. The collaboration was massive in scale as it spanned thousands of implications and dozens of contributors, but it was not hybrid in cognition. Rather than a co-author, AI functioned more like an environment: structuring workflow, accelerating tasks, and coordinating proof infrastructure. This acceleration is also to be seen in the sort of human-AI collaboration grounded in interactivity: conversations with ChatGPT, resulting in a counterexample to a conjecture in analytic number theory. It must be emphasized, however, that while such conversations may be seen as accessible to anyone, this is in fact an illusion: the funding, founding and having access to the infrastructure itself is certainly not open to anyone.

Even though both Jamnik and Tao envision collaboration, they address collaboration differently: Tao's *Equational Theorist* describes it as a scalable protocol for distributing formal tasks, while it is Jamnik's view that pushes us to revise our conception of what it means to think with a machine. But not every large-scale formalization effort involves machine reasoning or creative input. Floris van Doorn's *Carlson Project*, which formalizes a certain proof in *Lean*, exemplifies a deeply human collaborative model, where AI is absent from the reasoning process altogether[36]. What van Doorn's project reveals is that scale and infrastructure alone do not define the sociological and philosophical novelty of AI-based Big Mathematics. It is only when collaborative labor includes AI agents contributing original insights – conjectures, proof paths, mathematical objects – that we enter a new epistemic terrain. This differentiation matters because it marks a point of divergence between AI-based Big Mathematics and the more 'traditional' Big Science. In particle physics or genomics, or even in CFSG or in *Lean*-based projects, large-scale collaboration expanded the institutional and technological capacity of science without altering its conceptual structure. But in AI-based mathematics, collaboration cuts deeper: it reconfigures the very act of reasoning. The challenge ahead is to discern when AI-supported work reflects a distributed human effort mediated

---

[34] See https://www.youtube.com/watch?v=ogJSGYs_OCE

[35] See https://www.youtube.com/watch?v=T4DE27uk0jw.

[36] See https://www.youtube.com/watch?v=n1v7q7fxylI . To explain, Lennart Carleson published in 1966 a proof of a theorem in harmonic analysis, that states that the Fourier series converges pointwise to the original function under weak conditions. The *Carlson Project* aims to formalize and verify all the details of a generalized Carleson theorem in *Lean*.

by machines, and when it constitutes a genuinely new mode of joint cognition, one that calls for new frameworks of agency, authorship, objectivity, and creativity.

4. Concluding Remarks

This paper has examined how the integration of AI-based technologies into mathematics intersects with and diverges from the established framework of Big Science. While many of the infrastructural, organizational, and collaborative features of AI-based mathematics align with the historical trajectory of Big Science, we have argued that these continuities also point to an epistemic rupture. What distinguishes AI-based mathematics is not simply its reliance on scale or coordination, but its transformation of core categories in mathematical practice: authorship and creativity are no longer confined to human intellect, whereas agency becomes distributed across human-machine systems, pointing much more to the philosophical problem of non-surveyability of the obtained results that than relying on intersubjective justification. These shifts are not merely conceptual; they alter the very structure of mathematical labor and reshape how mathematical knowledge is generated, validated, and institutionalized. The emergence of AI as a co-creative force compels us to revisit foundational questions about what it means to do mathematics, and who or what is entitled to participate in its unfolding. What follows from this is that AI-based mathematics needs to be treated at least as a partial site of political struggle, not only as a site of epistemic transformation. Following Cathy O'Neil's argument, we may be watching the emergence of what might be called "mathematical infrastructure capitalism": the progressive relocation of high-end mathematical discovery into institutional contexts that are unequally accessible, privately governed, and computationally intensive.

In this light, our analysis also points toward a broader need to interrogate the contours of what might be called Big AI. Discussions of artificial intelligence often invoke scale: large models, massive datasets, and increasingly expansive computational infrastructure. But scale alone does not define the epistemic or political significance of AI systems. Like Big Science, Big AI is embedded in institutions, shaped by global asymmetries, and entangled with state, market, and academic interests. Unlike earlier phases of Big Science, however, Big AI in mathematics does not just demand public funding and centralized laboratories; it demands access to privately owned computational infrastructures and, in many cases, to proprietary models and software stacks. This is not because mathematics itself (or its theories) is "protected data," but because reproducible access often hinges on compute, model weights, and curated training/evaluation corpora that are not always openly available. In other words, that means the preconditions for participating in frontier mathematics are themselves politically distributed. In practice, this creates a class structure inside mathematical research: those with access to AI infrastructure (e.g., the compute to run and fine-tune models, the technical capacity to modify code and pipelines, or direct collaboration channels with the engineers who build and maintain these systems) can pose certain questions and be recognized as moving mathematics forward, while those without access become structurally dependent on systems they neither own nor audit. At the same time, it is crucial to note that we are not yet at a point where AI has generated new mathematical theories or conceptual frameworks that fundamentally transform the discipline. While AI-assisted results are accelerating discovery and hinting at future possibilities, a breakthrough that could only have been achieved by AI has not yet occurred. Such developments may come (perhaps in years or decades) but they remain prospective rather than present realities.

Yet Big AI also differs from Big Science in another important respect: it blurs the boundaries between tool and collaborator, automaton and author. As AI systems increasingly generate not only outputs but also insights by proposing conjectures, reconfiguring workflows, and challenging human intuitions, their role in knowledge production becomes both more active and more opaque. The case of AI-based Big Mathematics offers an intriguing example of this shift, making visible the philosophical and sociological stakes of treating AI systems not merely as instruments of analysis but as co-constructors of knowledge. While Big Mathematics asks from us, among others, to understand how large-scale technologies, infrastructure and enterprises are incorporated into the practices of doing mathematics, AI-based Big Mathematics sheds light on an additional, more contemporary shift: understanding how the human factor is being reshaped in the age of mathematical collaboration with non-human agents. In this sense, the hype around AI's capabilities should not obscure the more important question: not what AI can supposedly 'discover' today, but how its integration is already reorganizing epistemic labor and reshaping norms of authorship, credit, and legitimacy. To put it bluntly: AI-based Big Mathematics is not only about whether machines can generate proofs. It reveals greater philosophical as well as social challenges to today's mathematical practice: who will be allowed to claim ownership over mathematical discovery, which institutions will be empowered to set the agenda for what mathematics becomes, and how resource concentration will redraw the social map of the discipline.

References:

Ahmed, N. and Wahed, M., 2020. The De-democratization of AI: Deep learning and the compute divide in artificial intelligence research. arXiv preprint https://doi.org/10.48550/arXiv.2010.15581

Aschbacher, M., 2004. The status of the classification of the finite simple groups. *Notices of the AMS*, 51(7), pp.736-740.

Asprey, W. "High Technology and Mathematics: A Historical Perspective." In: *Mathematics and Modern Technology*, edited by William Aspray and Philip Kitcher, Oxford University Press, 2019.

Arianrhod, Robyn. *Seduced by logic. Emilie Du Chatelet, Mary Somerville and the Newtonian Revolution*, Oxford University Press, Oxford, 2012.

Besiroglu, T., Bergerson, S.A., Michael, A., Heim, L., Luo, X. and Thompson, N., 2024. The compute divide in machine learning: A threat to academic contribution and scrutiny?. arXiv preprint arXiv:2401.02452.

Birhane, A., 2021. Algorithmic injustice: a relational ethics approach. Patterns, 2(2).

Blaauwbroek, L., Cerna, D.M., Gauthier, T., Jakubův, J., Kaliszyk, C., Suda, M. and Urban, J., 2024. Learning guided automated reasoning: A brief survey. *Logics and Type Systems in Theory and Practice: Essays Dedicated to Herman Geuvers on The Occasion of His 60th Birthday*, pp.54-83.

Bloor, David. (1978) Polyhedra and the abominations of Leviticus. *Br J Hist Sci* 11(3):245–272.

Bloor, David. (1991). *Knowledge and Social Imagery*. University of Chicago Press.

Boden, M. (2004) (2nd ed) *The creative mind. Myths and Mechanisms*. Routledge: London.

Bos H. J. M., Mehrtens H. (1977). The interactions of mathematics and society in history. Some exploratory remarks. *Hist Math* 4:7–30.

Bressan, B., Kocsis, A., Liyanage, S. and Tello, P.G., (2024) 'Big science as a complex human enterprise' In: *Big science, innovation, and societal contributions: The organisations and collaborations in big science experiments*, Shantha Liyanage (ed.) et al., Oxford University Press, pp. 249-263.

Bourbaki, N. 1949. "Foundations of Mathematics for the Working Mathematician." The Journal of Symbolic Logic 14 (1): 1–8. https://doi.org/10.2307/2268971

Chemla K. (2015). Proof, generality and the prescription of mathematical action: a nanohistorical approach to communication. *Centaurus* 57(4):278–300

Chervonyi, Y., Trinh, T.H., Olšák, M., Yang, X., Nguyen, H., Menegali, M., Jung, J., Verma, V., Le, Q.V. and Luong, T., 2025. Gold-medalist performance in solving olympiad geometry with alphageometry2. arXiv preprint arXiv:2502.03544

Chun, Wendy Hui Kyong. *Control and Freedom: Power and Paranoia in the Age of Fiber Optics*. MIT Press, 2008.

Cole J. C. (2008) Mathematical domains: social constructs? In: Gold B, Simons R (eds) *Proof and other dilemmas: mathematics and philosophy*. Mathematics Association of America, Washington, DC, pp 109–128.

Cole J. C. (2015) Social construction, mathematics, and the collective imposition of function onto reality. *Erkenntnis* 80(6): 1101–1124

Collins, Harry. "The Meaning of Data: Open and Closed Evidential Cultures in the Search for Gravitational Waves." *American Journal of Sociology* 104, no. 2 (1998): 293–338.

Cooper, J.M. and Hutchinson, D.S. eds., 1997. Plato: complete works. Hackett Publishing.

Corry, Leo. 2004. *Modern Algebra and the Rise of Mathematical Structures*. Basel: Springer.

Cranmer, Miles, et al. "Discovering symbolic models from deep learning with inductive biases." *Nature Communications* 11.1 (2020): 1–10.

Crawford, K., 2021. The atlas of AI: Power, politics, and the planetary costs of artificial intelligence. Yale University Press

Croce, M., & Baghramian, M. (2024). *Experts—Part II: The sources of epistemic authority*. Philosophy Compass, e70005. https://doi.org/10.1111/phc3.70005

Csikszentmihalyi C. 1996. Creativity: flow and the psychology of discovery and invention. New York: HarperCollins.

Davies, Alex, et al. "Advancing mathematics by guiding human intuition with AI." *Nature* 600 (2021): 70–74.

De Toffoli, S., & Giardino, V. (2014). An inquiry into the practice of proving in low-dimensional topology. *Boston Studies in the History and the Philosophy of Science*, 308, 315–336.

Dean, W., & Naibo, A. (2024). Artificial intelligence and inherent mathematical difficulty. ArXiv. https://arxiv.org/abs/2408.03345

Ernest P. (1998). *Social constructivism as a philosophy of mathematics*. SUNY, Albany

Fawzi, A., Balog, M., Huang, A., Hubert, T., Romera-Paredes, B., Barekatain, M., Novikov, A., R. Ruiz, F.J., Schrittwieser, J., Swirszcz, G. and Silver, D., 2022. Discovering faster matrix multiplication algorithms with reinforcement learning. Nature, 610(7930), pp.47-53.

Feferman S. (2011) Is the continuum hypothesis a definite mathematical problem? http://logic.harvard.edu/EFI_Feferman_IsCHdefinite.pdf

Ferreirós J. (2016). *Mathematical knowledge and the interplay of practices*. Princeton University Press, Princeton.

Floridi, Luciano (2025). AI as Agency without Intelligence: On Artificial Intelligence as a New Form of Artificial Agency and the Multiple Realisability of Agency Thesis. Philos. Technol. 38, 30. https://doi.org/10.1007/s13347-025-00858-9

Friedman, M. (2019). Mathematical formalization and diagrammatic reasoning: the case study of the braid group between 1925 and 1950. *British Journal for the History of Mathematics*, *34*(1), 43–59.

Friedman, M, and Kish Bar-On, K. "A Tool or A Collaborator? Rethinking Mathematical Intuition, Authorship, and Practice in the Age of AI." *Under Review*.

Galison, P. *Image and Logic: A Material Culture of Microphysics*. University of Chicago Press, 1997.

Galison, Peter, and Bruce Hevly (eds.). *Big Science: The Growth of Large-Scale Research*. Stanford University Press, 1992.

Ghiselin B. (1952) The creative process: reflections on invention in the arts and sciences. Berkeley: University of California Press.

Gieryn, Thomas F. (1999). *Cultural Boundaries of Science: Credibility on the Line*. Chicago/London: University of Chicago Press.

Gorenstein, D. 1982. *Finite Simple Groups: An Introduction to Their Classification*. New York: Plenum Press.

Grinstein, Louise S., and Paul J. Campbell (eds.). *Women of Mathematics: A Biobibliographic Sourcebook*. New York: Greenwood Press, 1987.

Hadamard, J. (1945), The Psychology of Invention in the Mathematical Field. Princeton University Press.

Hamami, Y. (2023). Agency in Mathematical Practice. In: Sriraman, B. (eds) *Handbook of the History and Philosophy of Mathematical Practice*. Springer, Cham. https://doi.org/10.1007/978-3-030-19071-2_48-1

Hauswald, Rico. (2021) *Epistemische Autoritäten*. Individuelle und plurale. Dresden.

Hauswald, R. (2025). Artificial Epistemic Authorities. Social Epistemology, 39(6), 716–725. https://doi.org/10.1080/02691728.2025.2449602

Henrion, C., 1997. Women in mathematics: The addition of difference. Indiana University Press.

Hutson, J., 2025. Human-ai collaboration in writing: A multidimensional framework for creative and intellectual authorship. *International Journal of Changes in Education*.

IEA (2025), Energy and AI, IEA, Paris https://www.iea.org/reports/energy-and-ai, Licence: CC BY 4.0

Jäger, Christoph. (2025). Epistemic Authority. In: *Oxford Handbook of Social Epistemology*, hrsg. v. Jennifer Lackey und Aidan McGlynn, New York.

Jaton, F., 2021. The constitution of algorithms: Ground-truthing, programming, formulating. MIT Press.

Joseph, J., (2025). The algorithmic self: how AI is reshaping human identity, introspection, and agency. *Frontiers in Psychology*, 16, p.1645795.

Kaiser, David. (2005). *Drawing Theories Apart: The Dispersion of Feynman Diagrams in Postwar Physics*. University of Chicago Press.

Kenschaft, Patricia C. *Change is possible: Stories of women and minorities in mathematics*. American Mathematical Soc., 2005.

Kitcher P. (1984). *The nature of mathematical knowledge*. Oxford University Press, New York

Knorr Cetina, Karin. *Epistemic Cultures: How the Sciences Make Knowledge*. Harvard University Press, 1999.

Larvor B. (2019) From Euclidean geometry to knots and nets. *Synthese* 196(7): 2715–2736

Latour, Bruno. *Science in Action: How to Follow Scientists and Engineers through Society*. Harvard University Press, 1987.

Liljedahl, P., & Sriraman, B. (2006). Musings on mathematical creativity. *For The Learning of Mathematics*, 26(1), 17-19.

Liljedahl, P., Allan, D. (2013). Mathematical Discovery. In: Carayannis, E.G. (eds) *Encyclopedia of Creativity, Invention, Innovation and Entrepreneurship*. Springer, New York, NY. https://doi.org/10.1007/978-1-4614-3858-8_376

Mac Lane, Saunders. 1971. *Categories for the Working Mathematician*. New York: Springer.

Mancosu P., Jørgensen K. F., Pedersen S. A. (eds) (2005) *Visualization, explanation and reasoning styles in mathematics*. Springer, Dordrecht.

Marxen, H., & Buntrock, J. (1990). Attacking the Busy Beaver 5. Bull. EATCS, 40, 247-251.

Nadjafikhah, M., Yaftian, N., & Bakhshalizadeh, S. (2012). Mathematical Creativity: Some Definitions and Characteristics. *Procedia-Social and Behavioral Sciences*, 31, 285-291. http://dx.doi.org/10.1016/j.sbspro.2011.12.056

O'Neil, C., 2017. Weapons of math destruction: How big data increases inequality and threatens democracy. Crown.

Parshall, Karen Hunger, and David E. Rowe. 1994. *The Emergence of the American Mathematical Research Community, 1876–1900*. AMS & London Math Society.

Parshina, Katia. (2023). "Philosophical Assumptions Behind the Rejection of Computer-Based Proofs" *KRITERION – Journal of Philosophy*, vol. 37, no. 2-4,, pp. 105-122.

Polu, S., Han, J.M., Zheng, K., Baksys, M., Babuschkin, I. and Sutskever, I., 2022. Formal mathematics statement curriculum learning. arXiv preprint arXiv:2202.01344

Polu, S. and Sutskever, I., 2020. Generative language modeling for automated theorem proving. arXiv preprint arXiv:2009.03393.

Ravetz, Jerome R. *Scientific Knowledge and Its Social Problems*. Oxford University Press, 1971.

Restivo S. (2001). *Mathematics in society and history: sociological inquiries*. Springer Science + business Media, Dordrecht.

Romera-Paredes, B., Barekatain, M., Novikov, A., Balog, M., Kumar, M.P., Dupont, E., Ruiz, F.J., Ellenberg, J.S., Wang, P., Fawzi, O. and Kohli, P., (2024). Mathematical discoveries from program search with large language models. Nature, 625(7995), pp.468-475.

Rowe, D. E., Koreuber, M. 2020. *Proving It Her Way. Emmy Noether, a Life in Mathematics*. Cham: Springer.

Searle, J. R. (1995) *The construction of social reality*. Free Press, New York.

Siegmund-Schultze, R. 2009. *Mathematicians Fleeing from Nazi Germany: Individual Fates and Global Impact*. Princeton: Princeton University Press.

Shehabi, Arman, Sarah J. Smith, et al. 2024. 2024 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory (LBNL-2001637). December 2024 https://doi.org/10.71468/P1WC7Q

Shrum, Wesley, Joel Genuth, and Ivan Chompalov. *Structures of Scientific Collaboration*. MIT Press, 2007.

Si, C., Hashimoto, T., & Yang, D. (2025). The Ideation-Execution Gap: Execution Outcomes of LLM-Generated versus Human Research Ideas. ArXiv. https://arxiv.org/abs/2506.20803

Srdarov, S., & Leaver, T. (2024). Generative AI Glitches: The Artificial Everything. *M/C Journal*, 27(6). https://doi.org/10.5204/mcj.3123

Steingart, Alma. 2012. "A Group Theory of Group Theory: Collaborative Mathematics and the 'uninvention' of a 1000-Page Proof." *Social Studies of Science* 42 (2): 185–213.

Strubell, Emma, Ananya Ganesh, and Andrew McCallum. 2019. Energy and Policy Considerations for Deep Learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 3645–3650, Florence, Italy. Association for Computational Linguistics.

Subercaseaux, B., Nawrocki, W., Gallicchio, J., Codel, C., Carneiro, M., & Heule, M. J. (2024). Formal Verification of the Empty Hexagon Number. ArXiv. https://arxiv.org/abs/2403.17370

Sun, Y., Hu, S., Zhou, G., Zheng, K., Hajishirzi, H., Dziri, N., & Song, D. (2025). OMEGA: Can LLMs Reason Outside the Box in Math? Evaluating Exploratory, Compositional, and Transformative Generalization. ArXiv. https://arxiv.org/abs/2506.18880

Taylor, Zachary W., and Kayla Stan. 2024. "Exploring the Stratified Nature of Artificial Intelligence Research Funding in United States Educational Systems: A Bibliometric and Network Analysis" *Education Sciences* 14, no. 11: 1248. https://doi.org/10.3390/educsci14111248

Trinh, T.H., Wu, Y., Le, Q.V. et al. Solving olympiad geometry without human demonstrations. Nature 625, 476–482 (2024). https://doi.org/10.1038/s41586-023-06747-5

Urman, A., Smirnov, I. and Lasser, J., 2024. The right to audit and power asymmetries in algorithm auditing. EPJ Data Science, 13(1), p.19. https://epjds.epj.org/articles/epjdata/abs/2024/01/13688_2024_Article_454/13688_2024_Article_454.html

Van Wyk, B., 2024. Exploring the philosophy and practice of AI literacy in higher education in the Global South: a scoping review. *Cybrarians Journal*, (73), pp.1-21.

Wall, P.J., Saxena, D. and Brown, S., (2021). Artificial intelligence in the Global South (AI4D): Potential and risks. arXiv preprint https://doi.org/10.48550/arXiv.2108.10093

Weber, Keith, and Fenner Stanley Tanswell. 2022. "Instructions and recipes in mathematical proofs." *Educ Stud Math* 111 (1): 73–87

Weinberg, Alvin. "Impact of Large-Scale Science on the United States." *Science* 134, no. 3473 (1961): 161–164.

Wilholt, T., 2013. Epistemic trust in science. *The British Journal for the Philosophy of Science* 64 (2): 233-253.

Wink, Ruediger, Alberto Di Meglio, Marilena Streit-Bianchi, and Shantha Liyanage, (2024). 'Big Science and Social Responsibility of the Digital World', in Shantha Liyanage, Markus Nordberg, and Marilena Streit-Bianchi (eds), *Big Science, Innovation, and Societal Contributions: The Organisations and Collaborations in Big Science Experiments,* Oxford Academic https://doi.org/10.1093/oso/9780198881193.003.0012

Zagzebski, Linda. (2015). *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford, Oxford University Press.

Zhao, J., Zhang, T., Sun, J., Tian, M. and Huang, H., 2025. Pi-gps: Enhancing geometry problem solving by unleashing the power of diagrammatic information. arXiv preprint arXiv:2503.05543.

Zhang, C., Song, J., Li, S., Liang, Y., Ma, Y., Wang, W., Zhu, Y., & Zhu, S. (2024). Proposing and solving olympiad geometry with guided tree search. ArXiv. https://arxiv.org/abs/2412.10673