

Big AI-Based Mathematics and the Redistribution of Epistemic Labor

Social Studies of Science

1–25

© The Author(s) 2026

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/03063127261482050

journals.sagepub.com/home/sss

Michael Friedman¹ , and Kati Kish Bar-On^{2,3} 

Abstract

The use of AI-based technologies is reshaping mathematics in ways that extend but also strain the classic Big Science analogy. As in Big Science, the new AI-based mathematical work relies on concentrated computation, specialized and hierarchical teams, sociopolitical justification, and large-scale collaboration. But AI also changes how epistemic labor is distributed and attributed: Systems now generate constructions, explore proof spaces, and formalize arguments, prompting mathematicians to renegotiate how credit, authority, and accountability are assigned across 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), and examining recent cases, we argue that AI-based mathematics constitutes a novel sociotechnical regime 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 AI Mathematics, knowledge production, epistemic authority

¹Mathematical Institute, University of Bonn, Germany

²Boston University, MA, USA

³Harvard Business School, Boston, MA, USA

Corresponding Author:

Kati Kish Bar-On, Department of Philosophy, Boston University, 745 Commonwealth Avenue, 5th Floor, Boston, MA 02215, USA.

Email: kkish@bu.edu

Introduction

The term Big Science emerged in the 1960s to describe a mode of research that had taken hold in the mid-20th 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 (Capshaw & Rader, 1992; Galison & Hevly, 1992; Weinberg, 1961). 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 organization 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 (Aker, 2006; Edwards, 2013). For example, the classification of finite simple groups, a project whose first stage was taking place between the 1950s and the 1980s, was explicitly described by the historian of science Volker Remmert and the mathematician Rebecca Waldecker (Remmert & Waldecker, 2026) as the first project incorporating characteristics of Big Mathematics (see also: Aschbacher, 2004; Gorenstein, 1982; Steingart, 2012). For a second example we might turn to the more distributed projects of formalization of proofs by automated proof-assistants such as *Coq* or *Lean* in the 20th and 21st centuries (Hartnett, 2026). Both projects, though obviously of a different character and scale, have increasingly required infrastructure, coordination, and sustained institutional support. Moreover, in recent years, there has been a rapid growth in the use of AI-based technologies in mathematical practice and research, where these technologies begin to play a pivotal role. From machine-learning-assisted theorem provers¹ to symbolic reasoning tools integrated with large language models (LLMs), the role of AI is no longer merely 'symbolical' or computational but rather more 'generative'. 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* (Chervonyi et al., 2025; Trinh et al., 2024) discovering new and unexpected proofs for problems in plane Euclidean geometry, *FunSearch* (Romera-Paredes et al., 2024) which writes programs for finding various special sets for the cap set problem (and see also its successor *AlphaEvolve*), or recent discoveries in knot theory (Davies et al., 2021).² In addition, two further achievements in 2026 are worth mentioning: First, *Gauss*, an AI-based agent developed by the start-up *Math, Inc.* has completed (or produced completely) an autoformalized 'sorry-free' proof of the sphere packing problem in dimensions 8 and 24 in February 2026, which was proven by human mathematicians in 2016.³ Second, in April 2026, ChatGPT 5.4 Pro has solved Erdős Problem #1196 (a problem in number theory), using a method which human mathematicians did not notice,⁴ 'reveal[ing] a tight connection between [two domains] ... that does not, to my knowledge, have any explicit precursor in the literature', to cite the mathematician Terence Tao.⁵ In this context, we ask: Is this use of AI-based technologies 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?⁶

Before developing our answer, two clarifications are in order. First, since Big Science obviously reshapes labor, we do not claim that Big Science left human epistemic agency intact or that AI alone has disrupted it. Big Science has long produced its own ruptures around agency, authorship, and credit.⁷ Alongside these ruptures a repertoire of institutional and discursive tools has absorbed them: co-authorship conventions, collaboration protocols, long author lists, and formal credit-sharing arrangements with engineering teams (Biagioli & Galison, 2014; on the unfinished character of this absorption, see Sorgner, 2026). However, mathematics, as a discipline long organized around the figure of the individual or small-group mathematicians, lacks comparable models for attributing, delegating, and staging agency across human and non-human contributors (e.g. Heintz, 2000). AI-based mathematics is not introducing the problem of distributed agency for the first time; rather, it is confronting mathematics with a problem that other fields have already begun to institutionalize and conceptualize, under conditions in which the relevant vocabularies and infrastructures remain underdeveloped for mathematics. Second, a word about terminology. We use the term Big AI Mathematics to name the emerging sociotechnical regime we describe. The term signals continuity with Big Science and Big Mathematics, while also marking that AI-based mathematics exhibits a structural tendency toward going big: the specific combination of LLMs, massive datasets of various sorts, and accelerator-scale compute is constitutive of these projects.

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: Google's DeepMind, OpenAI, *Math, Inc.* 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 O'Neil (2017) warns 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. The recent *Leiden Declaration* on Artificial Intelligence and Mathematics, published in June 2026, reacted to this shifting political economy (Leiden Declaration, 2026). We argue that Big AI Mathematics is reproducing a discipline-internal version of this phenomenon.

The analogy to Big Science is no longer enough. Big AI Mathematics redistributes epistemic authority from university departments and individual mathematicians toward corporate AI research labs that control compute, models, data, and narrative. This redistribution is already visible in high-profile cases: *AlphaEvolve*, *AlphaTensor* and *AlphaGeometry* emerge from Google's DeepMind. The formalization of the 8- and 24-dim sphere packing problem was done by *Math, Inc.*, an AI startup dedicated to 'auto-formalization'. *FunSearch* and *AlphaEvolve* show how LLMs could write and 'mutate' small Python programs, and the latter is framed as a major advance in mathematics using Tensor Processing Unit-scale (TPU) resources, being 'Google's custom AI accelerator'.⁸ And recent developments such as Terence Tao's and other mathematicians' ChatGPT-assisted work (as well as other AI-models) in analytic number theory show that doing frontier mathematics may increasingly mean working through corporate models.⁹ 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.¹⁰ For centuries, 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 (Arianrhod, 2012; Grinstein & Campbell, 1987). In the 20th century, access to advanced mathematical training and research positions was tightly bounded by class, gender, race, and institutional affiliation.¹¹ 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 AI Mathematics shares with Big Science. Second, we explore how AI-based technologies are changing how agency, authorship, and trust are attributed and staged within mathematical practice, and what this means for the discipline’s political economy. We conclude by arguing that Big AI Mathematics is not merely a quantitative intensification of existing trends but the formation of a sociotechnical regime whose governance is as important as its epistemology.

AI-Based Mathematics and Big Science: Characteristics and Implications

As Galison and Hevly (1992) show, 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. 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 the physical dimensions of equipment. However, as Hevly argues, ‘Big Science is not simply science carried out with big or expensive instruments’ (p. 355). Instead, the ‘bigness’ of Big Science stems from its extensive reach beyond the traditional boundaries of scientific practice. One may 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 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 (pp. 357–358). Large scientific collaborations have changed the relationship of individual researchers to their experiments and the nature of individual contributions to knowledge creation.

Big Science did more than scale up existing practices; it displaced the individual epistemic subject in favor of the experimental apparatus itself, thereby reconfiguring the epistemic labor involved. Such a reconfiguration in AI-based mathematics takes a particular shape: Machines no longer merely verify results or carry out calculations; they now find new correlations, suggest new proofs and auxiliary constructions, and develop programs that mathematicians only then read, interpret, and sometimes adopt. While several contemporary discussions revolve around the question of whether this involves agency or not (e.g. Dung, 2025; Floridi, 2025), we want to ask how agency is being attributed, delegated, and staged across the pipeline that produces a result. This question goes beyond the abstract aspects of agency (of asking whether a machine 'really' authors or discovers a proof, or co-authors a publication), since we now ask which parts of the work are handed off to which actors, which outputs are treated as finished and which as drafts that require human rework, and how the trustworthiness of a result becomes stabilized across human and non-human contributors. Hence this is a reconfiguration not just of scale or hierarchy, but of the humans' relationship to knowledge itself (Collins, 1998; Ravetz, 1971).

Cross-disciplinary entanglements of this kind characterize other fields, including medicine, applied physics, computational biology, and data-intensive social science, many of which likewise emphasize ideals of autonomy and internal rigor. What makes mathematics particularly revealing, 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 in which novel ideas originate (only) from (human) mathematicians becomes unstable. 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).

Based on such characteristics, AI-based technologies in mathematics can be understood as manifestations of Big Science in the 21st century. However, as we will show, each of the four characteristics acquires a distinctive shape when AI enters the scene, and the sum of these shifts marks what we are calling Big AI Mathematics. We argue that in mathematics, these shifts destabilize traditional notions of expertise and intellectual authority, and force us to rethink what it means for mathematics to be a human activity,

if theoretically the entire body of mathematics (or practically just entire sub-disciplines of it) can be ‘fed’ into and consequently developed by a machine. The recent ‘conversations’ between various mathematicians and ChatGPT or other AI-based technologies that produced (new) proofs of Erdős problems or counterexamples to a conjecture in analytic number theory (i.e. Erdős problem #90) illustrate this on a small scale.¹² The results cannot be termed purely AI-authored proofs, since they are based (also) on other results discovered by humans, but they also cannot be solely attributed to human mathematicians, and the discipline does not yet have settled conventions for describing what exactly takes place or how basic epistemological notions have changed. In the next sections we analyze each of the characteristics of Big Science through an AI-based prism of mathematical practice, using AI-assisted episodes (among others, *AlphaGeometry*, *FunSearch*, knot theory and also how the Leiden Declaration frames such characteristics) that examine labor delegation, the process of discovery and its outcomes, and under which circumstance it can establish trust.

Concentration and Centralization of Computational Resources

The first characteristic of Big Science—the concentration of resources and data in specialized research centers—can be seen in recent AI-mathematical research developments. This characteristic can be seen in two aspects: first, it can be seen in the unequal access to costly, proprietary AI and computational systems, which will be discussed below; and second, it can be seen in the capacity to handle vast amounts of data, which will be discussed here. Contemporary AI-based 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., 2026, p. 84). 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).¹³ A similar pattern emerges with *AlphaGeometry*: The developers ‘extracted 100 million synthetic theorems’ and ‘produce[d] nearly 10 million synthetic proof steps that construct auxiliary points, reaching beyond the scope of pure symbolic deduction’ (Trinh et al., 2024, p. 476). Concretely, once a problem has been given, *AlphaProof*, the engine behind *AlphaGeometry*, generates numerous solution candidates, by computing the ‘deduction closure’, i.e. the set of all deducible theorems given a core set of initial statements. These proofs are ‘used to reinforce *AlphaProof*’s language model’. If a solution is not found, various LLMs add auxiliary constructions, and once again the ‘deduction closure’ is computed; this process is repeated until a complete and valid proof is found (p. 480).

As such, these technologies embody the distinctive contours of Big AI Mathematics—illustrating a shift not only in scale (as Big Science offered) but also in mode: toward a sociotechnical regime in which humans and machine agents collaborate on cognitive work, where proof-generation is automated and problem discovery is mechanized. Another example illustrates this pattern: Knot theory research using AI involved processing ‘more than 2.7 million knots’ through neural networks (Davies et al., 2021, pp. 71–72). Such computational demands necessitate the kind of resource concentration that

characterized the emergence of Big Physics laboratories in the postwar period.¹⁴ These examples represent a fundamental shift from mathematics as an individual or small-group endeavor to big collaborative constellations.

The acceptance of *AlphaGeometry*'s outputs (as well as of outputs of other AI-based technologies) as genuine mathematical results depends on a specific trust arrangement. The symbolic engine's deductive steps are formally verifiable (each inference can be checked) so the logical skeleton of a proof is treated as reliable. But the auxiliary constructions proposed by the language model are not derived from any surveyable rationale, which again results in epistemic opacity; they emerge from training on roughly 100 million synthetic problem instances. Trust in these constructions can be stabilized not through mathematical understanding of why they work, but through the conjunction of formal verification (the symbolic engine confirms that the construction leads to a valid proof) and institutional reputation (DeepMind's track record and the involvement of recognized mathematicians as evaluators). When Chervonyi et al. (2025, p. 14) report that their geometry experts consider several solutions to exhibit 'superhuman creativity', that statement functions as a legitimization device that converts a machine output into a disciplinary achievement. Because the pipeline cannot be rebuilt outside DeepMind, trust in the result defaults to the institution that produced it.

However, as we noted earlier, this is not the first time that there has been a shift from individual to collective work in the history of mathematics. The idea that mathematics is a field in which discovery and proof were primarily associated with individual efforts was challenged by 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). Remmert and Waldecker (2026) 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. Another example of such a collaborative project which did not use any AI-based technology is the Busy Beaver Challenge, an international effort to determine the precise value of $BB(5)$, a notoriously hard-to-compute function rooted in Turing machine theory.¹⁵ 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*).

What these pre-AI projects already establish is that mathematics can sustain large collaboration settings in which no single mathematician has a global view of the proof. The contrast with Big AI Mathematics lies elsewhere. In CFSG and in Busy Beaver Challenge, the constellation is still human-readable: Every step is produced, checked, and signed off by identifiable contributors, and even when formal proof assistants are involved, the verification pipeline is designed to be surveyable by further humans. In the *AlphaGeometry* case, by contrast, the concentration of compute and data is what allows the auxiliary-construction heuristics to work, and that concentration is not easily reproduced outside the institution that built it. The shift is not from small to big collaboration, but from collaboration whose infrastructure is broadly distributed to collaboration whose infrastructure and databases are not only extremely centralized, but also privately held.

Hierarchical Organization and Specialized Division of Labor

The second characteristic of Big Science—hierarchical organization with specialized workforce divisions—appears 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 the division of conceptual, technical, and interpretive labor (Galison & Hevly, 1992, pp. 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. This hierarchy is directly addressed by the *Leiden Declaration*: ‘Mathematicians and academic organizations collaborating with industry often face asymmetries in their bargaining positions, as well as in access to professional support such as legal resources’ whereas ‘[r]esearchers who do not have access to the technologies or decision-making related to them’ are disadvantaged (Leiden Declaration, 2026). As in Big Science, this distribution of labor demands new coordination mechanisms, protocols and agreement concerning quality control and power asymmetries, and evolving norms of authorship and accountability. To see the distribution of labor in a real case study, we turn to *FunSearch*.

The *FunSearch* AI-based technology, developed to solve the problem of cap set, coming from the area of combinatorics in order to find a specific set in an n -dimensional space,¹⁶ offers a concrete view of how this division of labor works in practice, and of how delegation shapes what counts as a result. *FunSearch* searches not in the space of mathematical objects (in this case, sets) directly but in the space of computer programs written in Python that generate them: Starting with a seed program, a pretrained language model produces ‘mutations’ (that are basically slightly altered programs) which are then scored by an automated evaluator. The programs that score well are fed back into the model to generate further mutations (Romera-Paredes et al., 2024, p. 469). In the case of the cap set problem, this loop produced a program that generated a 512-element cap set in dimension 8, improving on the previously known bound of 496.

The labor behind this result is distributed across at least four distinct roles. Mathematicians formulate the problem and define the scoring function that determines what counts as a good output. AI engineers design and maintain the search architecture and the feedback loop between model and evaluator. The pretrained language model generates candidate programs, which is a generative role that is neither passive computation nor independent authorship. Mathematicians then re-enter at the end to inspect the obtained programs, simplify their codes, and extract from it a humanly intelligible construction (Romera-Paredes et al., 2024, p. 471). The prompt engineering that mediates between mathematicians and the model is itself described as ‘nontrivial’, forming a further layer of specialized work.¹⁷

What makes this division significant for our purposes is that the output crosses several hands before it is accepted as a mathematical result: The Python program produced by the language model is not yet mathematics. The LLM's search space is the space of *all* computer programs, not only mathematical objects. It may become a mathematical object only after human mathematicians inspect the code, identify the construction it encodes, and reconstruct that construction in terms the discipline recognizes; that is, human interpretation supplies the mathematical meaning that the code alone does not carry. The result is therefore assembled across actors and platforms, not by a single actor (or a group of human actors) or tool alone.

This new human-machine redistribution and reconfiguration also transforms mathematical practice itself and its grounding concepts, such as the 'working mathematician'. The term 'working mathematician' has traditionally referred to a practitioner who is engaged in research and proves theorems (Bourbaki, 1949; Corry, 2004; Mac Lane, 1971). Yet the *FunSearch* case suggests that this category is becoming strained, as the labor that produces such results becomes distributed across human and machine actors, whereas the technology underlying the latter may be proprietary; indeed, this is not only true for *FunSearch* (and its successor *AlphaEvolve*) but also for the counter example of Erdős problem #90 that was obtained by an *internal* reasoning model developed by OpenAI. Another example of the shifting division of labor is the solution of Erdős Problem #1196 by ChatGPT 5.4 Pro (mentioned above); while the solution in this case was unexpected, human mathematicians still had to interpret the result and conceptually connect the methods employed.¹⁸ In these constellations, no single participant can be called a 'working mathematician' in the classical sense; instead, what emerges is a socio-technical division of labor in which the mathematician's role shifts toward problem formulation, evaluation design, and the interpretive work of converting machine outputs into disciplinary knowledge.

Political and Social Significance

Infrastructure and Control. 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.¹⁹ 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'.²⁰ 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-based mathematical discovery has become an infrastructural game: access to accelerators, bespoke data pipelines, and massive human-machine engineering teams are themselves gates to participation.

The institutions producing these results are not universities but corporations. *AlphaGeometry* and *AlphaTensor* were developed by Google DeepMind, and proofs of Erdős Problems and counterexamples emerged from use of and conversations with AI systems as ChatGPT, owned by OpenAI. These actors' agendas are not neutral, as they carry financial, reputational, and infrastructural motives that shape which problems are deemed important, which questions receive attention and how the results are framed. The Leiden Declaration identifies exactly such unequal access to proprietary AI-based technologies and computational systems, underlining that this 'disadvantages researchers who do not have access to the technologies or decision-making related to them' (Leiden Declaration, 2026). O'Neil (2017) warns us that while algorithmic systems are presented as objective tools for advancing knowledge, they often reinforce asymmetries of power and accountability, and this warning becomes even more relevant in the case of AI-based mathematical discovery in private laboratories (Jaton, 2021).

The discovery of a new relation in knot theory described by Davies et al. (2021) is exemplary, since it is described as one of the first examples of AI-based mathematical discovery. It illustrates a subtler dimension of this political embedding: not only who controls the infrastructure, but how the resulting discovery is narrated in ways that render that dependence invisible. Davies et al. investigated the relationship between two branches of knot theory by looking for connections between their invariants: the signature of a knot (a topological-algebraic invariant) and the hyperbolic volume (a geometric invariant). A neural network was then trained on a huge dataset (of roughly 2.7 million hyperbolic knots), and that neural network detected a correlation that no mathematician had anticipated by processing a dataset whose scale was beyond any mathematician's capacity to survey (Davies et al., 2021, p. 72). But then the mathematicians took the lead and formulated conjectures on the basis of it, revised those conjectures, and eventually proved a theorem. The paper's title 'Advancing mathematics by guiding human intuition with AI' describes the relationship between humans and AI tools as one of guidance rather than co-authorship. But beyond the mathematical discoveries, what is notable is the explicit remark that 'This research was funded by DeepMind.' That is, the corporate infrastructure—DeepMind's involvement and finance as well as the technology required to train the network on 2.7 million knots—was a precondition of the discovery, even though the final product was certified by a human proof. Already the narrative underlining one of the first papers advocating AI-based mathematical discovery normalizes a dependency that most mathematics departments could not replicate.

The concentration of AI-mathematical capabilities raises questions about research accessibility and policy that parallel earlier Big Science concerns. Just as Cold War-era physics clustered around well-funded laboratories backed by government and military institutions, contemporary AI-mathematics research tends to cluster in corporate AI labs and elite academic departments with access to high-performance computing and proprietary models (Taylor & Stan, 2024).²¹ 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 concerns, but political ones, tied to how research capacity, prestige, and epistemic authority are distributed across institutions and regions.

Possible answers to some of these concerns can already be found in recent institutional arrangements. In late 2025, Imperial College London, the Institute for Advanced Study, IHES, the Simons Institute at Berkeley, and TIFR entered a funded partnership with Google DeepMind, receiving financial support and access to Gemini Deep Think, AlphaEvolve, and AlphaProof, with the declared aim of identifying which mathematical problems are ripe for AI-driven insight.²² In this case, one might argue, agenda-setting is not imposed on the discipline from outside but organized jointly with it, by institutions that hold considerable epistemic authority and whose terms of access are nonetheless set by the company. Nevertheless, in June 2026 the Leiden declaration notes that ‘Current events illustrate the need for public alternatives to proprietary technologies, from basic services for online collaboration, to computer clusters for mathematical modeling and machine learning applications’ (Leiden Declaration, 2026). Hence, it is not entirely clear who gets to make final decisions in their ‘powerful feedback loop’, and according to whose standards ‘ripe’ is being evaluated.

Start-ups like *Math, Inc.* or AI-based systems like *AlphaTensor* or *FunSearch* are developed by teams with privileged access to specific technologies and infrastructure; for example, *AlphaTensor* and *FunSearch* employ custom-designed TPUs and massive training datasets²³ (Fawzi et al., 2022; Romera-Paredes et al., 2024)—resources that are not necessarily accessible to every mathematics department or research institute (Besiroglu et al., 2024; Urman et al., 2024). This asymmetry may lead to mathematical innovation being increasingly channeled 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 (Ahmed & Wahed, 2020; Wink et al., 2024).²⁴ Another concern is for that dependency to run upstream, for example in the case of the FrontierMath benchmark, that was written by more than sixty mathematicians who were not informed that OpenAI had funded it and retained exclusive access to the problems and their solutions (Glazer et al., 2024; Harris, 2025). Here mathematicians supplied the expert labor that made an evaluative standard credible, without visibility into the arrangement that standard served.

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 ‘conversations’ with ChatGPT are open to anyone having access to this technology, this is hardly true concerning the technology underlying, for example, *AlphaTensor* or *Math, Inc.* These structural questions extend the Big Science analogy but also go beyond it, suggesting that Big AI Mathematics increasingly shapes political priorities and future research trajectories.

Agency. The political stakes of Big AI Mathematics extend beyond questions of funding and access to the question of how agency itself is attributed within mathematical work. The issue of agency and authorship is highlighted in the Leiden declaration, as its authors

note that ‘Results are attributable to specific authors who take credit for their discovery and assume responsibility for their correctness’, stating unequivocally that ‘credit and responsibility continue to belong to humans within the mathematical community and should not be given to automated systems’ (Leiden Declaration, 2026). The concept of agency has been discussed across several disciplines and underwent considerable rethinking over the years. In history and sociology of science, agency has been treated as something that is attributed, delegated, staged, and denied through specific material and institutional arrangements, and not as an attribute of individuals (Callon, 1984; Latour, 1987; Pickering, 1995). In philosophy of mathematical practice, agency refers to a distinctly human capacity (Hamami, 2023; Mancosu et al., 2005) according to which 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 (Chemla, 2015; De Toffoli & Giardino, 2014; Ferreirós, 2016; Larvor, 2019; Weber & Tanswell 2022). Computers, where they appear in these accounts, are cast as passive tools or aids for calculation or verification, whose contribution does not alter the agential structure of the practice. This framing worked well enough for proof assistants like *Lean* or *Coq*, where the tool checks human-authored proofs and the mathematician retains clear authorial control. It fits less comfortably when a large language model proposes an auxiliary construction that no mathematician anticipated, or when it produces an object or a counterexample that mathematicians must reverse-engineer to understand. The Leiden declaration’s statements are set against this backdrop of a shift in the author’s status.

In such cases, the machine’s contribution cannot be classified as either passive execution or independent authorship; the counterexamples and constructions would not exist without the model, even though all of the components of these constructions were already existing in the (human written) literature and even though the model had no understanding of what it proposed.²⁵ Asking where agency is located or whether the machine ‘really’ possesses agency seems less productive; the more productive question is how agency is distributed across this human–machine combination that produces a result, and who ends up being treated as responsible for what.

The examples discussed above illustrate different arrangements of delegated agency. In *AlphaGeometry*, in Erdős Problem #1196 or in Erdős Problem #90, the generative work (proposing auxiliary constructions or suggesting using unexpected methods) is delegated to the language model, and the mathematician’s agency is exercised retroactively, in the interpretive and evaluative work of making sense of the output. In the knot theory case, the network detects a correlation, but the mathematician retains anticipatory agency: She decides which correlations are worth pursuing, formulates conjectures, and provides the proof. In the *FunSearch* case, agency circulates through a loop: Mathematicians set the scoring function, the model generates programs, and mathematicians re-enter to ‘make sense’ of the result. The (semi-)autoformalization of *Math, Inc.*, even though not prompting or producing new results, presents a completely different arrangement of delegated agency, in which *Gauss*, the autoformalization agent, has just overtaken the entire project of *Lean* formalization of the existing proofs. None of these arrangements fits the classical picture of a sole human agent directing a passive instrument, but they also do not fit a picture in which the machine has ‘become’ a sole agent

(in any robust sense). What they do show is that the attribution of agency has become a practical problem that each project approaches differently, depending on its technical architecture and the nature of the problem.

This practical instability has a consequence that extends beyond the philosophy of mathematics. As long as agency is treated as a property that inheres in individuals (or of groups of humans), the question of who gets credit for a result can be settled by asking who did the intellectual work. Once agency is distributed across a pipeline of engineers, language models, evaluators, and interpreting mathematicians, credit becomes a matter of convention. Resolving this becomes even more pressing since in mathematics, unlike in high-energy physics or biomedicine, the conventions for distributing credit across heterogeneous contributors are still rudimentary (Biagioli & Galison, 2014). In the absence of such conventions, attribution becomes a variable that the infrastructure owner sets. This claim can be claimed, as when *AlphaGeometry* and *FunSearch* results are publicly framed as achievements of *DeepMind* or when the counterexample of Erdős Problem #90 as achievements of OpenAI rather than of the mathematicians who formulated the problems or interpreted the outputs, or when *Math, Inc.* announces that it has completed the formalization of the strong Prime Number Theorem, a challenge set by Tao and Kontorovich.²⁶ It can equally be effaced, as when a general-purpose model is presented as a neutral utility available to anyone and the resulting mathematics is credited to the mathematician who prompted it, while the infrastructure that made the exchange possible recedes from view. These are opposite choices, and each is selected according to whether a company seeks recognition as a producer of mathematics or normalization as its background infrastructure. In this sense, the agency aspect is inseparable from the political-economy question raised at the start of this section: The actor who controls the pipeline also controls the attribution of intellectual contribution.

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. The latter surfaces a tension that runs through much of what follows: while AI can extend the range and scale of mathematical exploration, it does not yet replicate the kinds of creative, conceptual shifts characteristic of profound mathematical insight.

This tension bears directly on mathematical creativity. Creativity in mathematics has traditionally been understood as the capacity to generate original work that expands the field, introduces new lines of inquiry, or leads to solutions that were not previously envisioned (Boden, 2004; Hadamard, 1945; Liljedahl & Sriraman, 2006). The underlying assumption behind these accounts is that creativity results from the conscious and intentional efforts of the human mind. AI-based technologies create a dissonance here: Systems like *AlphaGeometry* produce solutions described by their developers as

exhibiting ‘superhuman creativity’ (Chervonyi et al., 2025, p. 14), but instead of humans leading the creative process, the AI-based technologies present solutions that mathematicians did not think of or that were unexpected. 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 as well as agency 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.

AI-based collaboration poses a challenge for the social constitution of mathematical knowledge. The view that mathematical truths are socially constituted—that their acceptance depends on collective scrutiny, peer review, and shared standards of justification—has a long history (Bloor, 1978, 1991; Bos & Mehrkens, 1977; Ernest, 1998). That view presupposed that the actors doing the scrutinizing were human, and that the objects of scrutiny were, in principle, surveyable. However, in Big AI Mathematics, both presuppositions are under pressure. When AI systems generate proofs whose internal workings are opaque (i.e., ‘the exact way these LLMs are operating is not disclosed to users’, Srdarov & Leaver, 2024), the classical model of peer review as a check on knowledge claims becomes limited. The problems of verification and surveyability that surfaced with earlier computer-assisted proofs (Parshina, 2023) return in a sharper form, because the outputs of AI systems are not merely long but structurally opaque: produced by processes that even their designers cannot fully explain. This is clearly evident with the latest achievement of *Math, Inc.*, automatically formalizing the human proofs of the 8- and 24-dimensional sphere packing. Not only is reviewing such formalized code practically impossible, as it comprises approximately 200,000 lines of code, but the process of auto-formalization is opaque to end users (or readers).

This opacity presses on the concept of epistemic authority, which refers to the capacity of an individual, group, or system to be recognized as a legitimate and trustworthy source of knowledge within a domain (Hauswald, 2021, 2025; Jäger, 2025; Zagzebski, 2015). In traditional mathematics, epistemic authority rested on the capacity for justification and explanation: a mathematician could claim authority over a result because she could walk others through the proof (Croce & Baghrarian, 2024; Wilholt, 2013). AI-generated outputs complicate this picture. If a system produces results accepted as valid but cannot explain them in terms the discipline recognizes, authority shifts from justification to performance: the system is trusted because it works, not because anyone understands why it works or why the reasoning ‘makes sense’. This raises also the issue of ‘epistemic’ or ‘interpretative’ opacity – i.e. the mathematicians’ difficulty of understanding AI-produced results (Beisbart, 2021, 2026). To better understand what this shift entails, we need to ask (a) which types of technologies, individuals, or institutions justifiably hold such authority in AI-mediated mathematics,²⁷ (b) following Gieryn’s (1999) extension of Weber’s concept of domination, how such authority is socially constructed and institutionally legitimized, and (c) what kind of epistemic goods can be transmitted from AI-based systems to human researchers. Is it knowledge, understanding, or merely instrumental reliability? Understanding 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. 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 Big AI Mathematics, making it all the more important to examine these forces right now.

Discussion and Concluding Remarks

Let us start with an obvious yet non-trivial point this paper makes: if AI-based mathematical practice reproduces the infrastructures of Big Science, it also reproduces the politics of infrastructure. The fact that a small number of institutions, start-ups and private companies now possess the technology that enables them to generate solutions, propose candidate theorems, test them at scale, and narrate them as advances in mathematics, is not philosophically or socially neutral; it is the construction of a gate.

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 with neither access to TPU-scale infrastructure nor connections to private start-ups or to *Google*, 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 ‘mathematics is cooked’ popularized by the physicist Alexander Wissner-Gross, or that ‘The AI Revolution in Math Has Arrived’²⁸). That rhetorical move matters, as it positions corporate infrastructures as a possible home of the future of mathematics, and everyone else as lagging behind. In Big AI 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.²⁹

Seen this way, Big AI Mathematics is not only a story about more efficient tools; it is also a story about the enclosure of knowledge, as mathematical discovery risks being partially relocated into privately governed infrastructures whose incentives are largely reputational, competitive, and commercial. The cases examined in this paper each illustrate various facets of this enclosure. In *Math, Inc.* and *AlphaGeometry*, the entire proof pipeline is proprietary. In the knot theory case or in the solution of the Erdős Problems, the internal models, or the corporate compute or paid access to AI systems was a precondition; however, the final theorems or proofs were certified and interpreted by human mathematicians, preserving a degree of disciplinary autonomy. In *FunSearch*, the result—which was first achieved by searching in the space of all python programs, a task no human could do—was assembled across a chain of human and machine handoffs in which no single actor could claim full authorship. While there are various modes of enclosure of knowledge, what these cases do share is that mathematical knowledge was stabilized not through the classical mechanism of peer-reviewable proof alone, but through varying combinations of formal verification, institutional reputation, and expert endorsement.

The cases also suggest that the figure of the ‘working mathematician’ is undergoing a broader socio-practical transformation, echoing current discussions of authorship. Across

all these episodes, the mathematician's role shifts from that of a sole author toward a relational position within a larger epistemic network, defined as much by the capacity to formulate problems for machines, interpret their outputs, and orchestrate hybrid workflows as by the ability to prove theorems. 'Doing mathematics' in this regime means learning how to collaborate with non-human agents, and that collaboration introduces a form of knowledge production whose epistemic status is genuinely new and unsettled. Mathematical practices in Big AI Mathematics are becoming (1) not necessarily rooted in human conventions alone, (2) partially detachable from their human contributors, and (3) bound up with a form of knowledge that the discipline has no settled vocabulary to describe. Whether an AI-generated proof can be considered a mathematical object in the same sense as a human-generated one, and to what extent it is socially constituted, remain open questions.

This discussion of mathematical practices shows the transformation of mathematical agency. It should be emphasized that before the rise of AI-based technologies, human mathematicians also employed hybrid reasoning—for example, using different methods in the course of the same proof: algebraic, geometric, topological—but hybridity did not change the concept of agency. That is, hybrid reasoning was a tool choice made by a human agent, leaving agency untouched. What is different now is the direction of the dependency: It is the hybrid nature of agency that produces hybrid types of reasoning, rather than the other way around. This change matters because it means that the shift from consensus-driven validation (where human experts evaluate results through debate and scrutiny) toward what we might call algorithmic closure—where the credibility of a result rests on its formal verifiability and reproducibility through automated procedures rather than on human comprehension—is not a temporary accommodation but a structural feature of the emerging regime. 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 sociotechnical contexts in which discovery occurs.

Therefore, 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 O'Neil's (2017) 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. To be clear, our claim is that unlike earlier phases of Big Science, Big AI Mathematics does not just demand public funding and centralized laboratories; it demands access to privately owned computational infrastructures and proprietary models. Hence, 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 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. But Big AI Mathematics also differs from Big Science in blurring the boundary between tool and collaborator. When AI systems propose conjectures, fabricate counter examples, generate auxiliary constructions, or produce programs that expand the boundaries of known mathematics, their role in knowledge production becomes both more active and more opaque than anything Big

Science instrumentation involved—CERN’s particle accelerator did not suggest which particles to look for. However, 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 the coming years) but they remain prospective rather than present realities.

What should not be obscured by the hype around AI’s current capabilities is how its integration is already reorganizing epistemic labor and reshaping the norms that hold mathematics together as a discipline. The question going forward is not only what AI can discover but who will be in a position to claim ownership over mathematical discovery, which institutions will set the agenda for what mathematics becomes, and whether the discipline can develop governance structures, such as conventions for credit, standards for transparency, and norms for reproducibility, that keep its future open rather than enclosed. These are not questions that mathematics has had to answer before, and the answers it arrives at will shape the discipline for a long time to come.

ORCID iDs

Michael Friedman  <https://orcid.org/0000-0002-2905-0170>

Kati Kish Bar-On  <https://orcid.org/0000-0001-6867-3701>

Author Contributions

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

Funding

The authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Kati Kish Bar-On was supported by the Alfred P. Sloan Foundation [grant number G-2025-25317].

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Notes

1. For example, systems such as GPT-f (Polu & 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>).
2. 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 (2026). See also: Davis (2021).
3. See <https://www.math.inc/sphere-packing> (last accessed on 20 April 2026). Autoformalization is the automated process of translating natural language mathematics (in form of informal

proofs or theorems) into machine-verifiable code, which can be read by formal proof assistants (e.g. *Lean*, *Cog*). A ‘sorry-free’ proof refers to a completed, fully verified mathematical proof in a formal proof assistant (like *Lean*) that contains no unproven parts.

4. See: <https://x.com/jdlichtman/status/2044298382852927894> (last accessed on 20 April 2026): ‘The proof produced by GPT5.4 Pro was quite surprising, since it rejected the ‘gambit’ that was implicit in all works on the subject since Erdős’ original 1935 paper. The [initial] idea [...] was so natural & tempting from a human-conceptual point of view, that it obscured a technical possibility to retain (efficient, yet counter-intuitive) analytic terminology throughout’, which GPT5.4 Pro did retain.
5. <https://www.erdosproblems.com/forum/thread/1196> (last accessed on 20 April 2026).
6. It must be noted that we deal here with different technologies, i.e. Gauss’ *Math, Inc.* which auto-formalizes (given) proofs, is essentially different from *AlphaEvolve*, which search the space of Python computer programs, or from *AlphaGeometry* whose aim is to solve geometrical problems. That is, one must differentiate between the various technologies being employed and not just call them ‘AI’.
7. For example, the work of technicians and instrument-builders, the opacity of precision apparatus, and the negotiated invisibility of whole layers of laboratory labor have all been documented in the canonic STS literature (Galison, 1997; Knorr Cetina, 1999; Shapin, 1989; Shapin & Schaffer, 2011).
8. Tensor Processing Unit is a neural processing unit application-specific integrated circuit, developed by Google during the 2010s for neural network machine learning. See: <https://deepmind.google/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/>
9. 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>. See also here the list of Erdos problems co-solved with AI: <https://www.erdosproblems.com/> and <https://github.com/teorth/erdosproblems/wiki/AI-contributions-to-Erd%C5%91s-problems> (the latter website is no longer updated as of Jun 30, 2026).
10. 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, pp. 870–898).
11. One of the prominent examples for this is how Jewish mathematicians were treated in Nazi Germany (see Siegmund-Schultze, 2009; or, more specifically, Emmy Noether is a case in point, (see: Rowe & Koreuber, 2020).
12. See: <https://github.com/teorth/erdosproblems/wiki/AI-contributions-to-Erd%C5%91s-problems>, <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 20 April 2026). As Tao also notes, this ‘argument has now been formalized in *Lean*’.
13. While a discussion on *TongGeometry* is outside the scope of our paper, we note that 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).
14. 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 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 et al., 2019).
 15. 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>
 16. 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, 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 (2022). However, Romera-Paredes et al. (2024) managed to improve the lower bound by using an LLM termed *FunSearch* with an evaluator, to 2.2202^n .
 17. This mirrors the novel labor structures described by Galison and Hevly, where different groups specialized in theoretical work, experimental design, and instrumentation. Another example comes from learning-guided reasoning, where 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, p. 19).
 18. See the 'quiver diagram' of Terence Tao: <https://tinyurl.com/fbztrdzb> (accessed on 20 April 2026)
 19. <https://deepmind.google/discover/blog/alphageometry-an-olympiad-level-ai-system-for-geometry/>
 20. <https://spectrum.ieee.org/matrix-multiplication-deepmind>
 21. 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: <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>)

22. See this report from Google declaring the partnership: <https://blog.google/innovation-and-ai/models-and-research/google-deepmind/ai-for-math/>
23. 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/>
24. 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://pro-falexreid.com/2025/09/22/major-university-led-ai-initiatives-and-their-focus-2019-2025/>).
25. See also the blog posts of Michael Harris from August 2026 dealing, among others, with this issue (Harris, 2026a, 2026b).
26. <https://www.math.inc/gauss>
27. 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, p. 6).
28. We refer to the title of *Quanta's* article from April 13, 2026: <https://www.quantamagazine.org/the-ai-revolution-in-math-has-arrived-20260413/> (last accessed on 20 April 2026).
29. 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 (2017) gives several examples in her book).

References

- Ahmed, N., & Wahed, M. (2020). *The De-democratization of AI: Deep learning and the compute divide in artificial intelligence research*. arXiv. <https://arxiv.org/abs/2010.15581>
- Akera, A. (2006). *Calculating a natural world: Scientists, engineers, and computers during the rise of US cold war research*. MIT Press.
- Arianrhod, R. (2012). *Seduced by logic: Emilie Du Chatelet, Mary Somerville and the Newtonian Revolution*. Oxford University Press.
- Aschbacher, M. (2004). The status of the classification of the finite simple groups. *Notices of the American Mathematical Society*, 51(7), 736–740.
- Beisbart, C. (2021). Opacity thought through: On the intransparency of computer simulations. *Synthese*, 199, 11643–11666. <https://doi.org/10.1007/s11229-021-03305-2>
- Beisbart, C. (2026). In which ways is machine learning opaque? In J. M. Durán & G. Pozzi (Eds.), *Philosophy of science for machine learning* (pp. 3–24). Springer.
- Besiroglu, T., Bergerson, S. A., Michael, A., Heim, L., Luo, X., & Thompson, N. (2024). *The compute divide in machine learning: A threat to academic contribution and scrutiny?* arXiv <https://doi.org/10.48550/arXiv.2401.02452>
- Biagioli, M., & Galison, P. (2014). *Scientific authorship: Credit and intellectual property in science*. Routledge.

- Blaauwbroek, L., Cerna, D. M., Gauthier, T., Jakubův, J., Kaliszyk, C., Suda, M., & Urban, J. (2024). Learning guided automated reasoning: A brief survey. In V. Capretta, R. Krebbers, & F. Wiedijk (Eds.), *Logics and type systems in theory and practice: Essays dedicated to Herman Geuvers on the occasion of his 60th birthday* (pp. 54–83). Springer.
- Bloor, D. (1978). Polyhedra and the abominations of Leviticus. *The British Journal for the History of Science*, 11(3), 245–272. <https://doi.org/10.1017/S000708740004379X>
- Bloor, D. (1991). *Knowledge and social imagery*. University of Chicago Press.
- Boden, M. (2004). *The creative mind: Myths and mechanisms* (2nd ed.). Routledge.
- Bos, H. J. M., & Mehrtens, H. (1977). The interactions of mathematics and society in history some exploratory remarks. *Historia Mathematica*, 4, 7–30. [https://doi.org/10.1016/0315-0860\(77\)90031-3](https://doi.org/10.1016/0315-0860(77)90031-3)
- 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>
- Callon, M. (1984). Some elements of a sociology of translation: Domestication of the scallops and the fishermen of St Briec Bay. *The Sociological Review*, 32, 196–233. <https://doi.org/10.1111/j.1467-954X.1984.tb00113.x>
- Capshew, J. H., & Rader, K. A. (1992). Big science: Price to the present. *Osiris*, 7, 3–25. <https://doi.org/10.1086/368703>
- Chemla, K. (2015). Proof, generality and the prescription of mathematical action: A nanohistorical approach to communication. *Centaurus*, 57(4), 278–300. <https://doi.org/10.1111/1600-0498.12111>
- Chervonyi, Y., Trinh, T. H., Olšák, M., Yang, X., Nguyen, H., Menegali, M., Jung, J., Verma, V., Le, Q. V., & Luong, T. (2025). *Gold-medalist performance in solving olympiad geometry with alphageometry2*. arXiv. <https://doi.org/10.48550/arXiv.2502.03544>
- Collins, H. (1998). The meaning of data: Open and closed evidential cultures in the search for gravitational waves. *American Journal of Sociology*, 104(2), 293–338. <https://doi.org/10.1086/210040>
- Cooper, J. M., & Hutchinson, D. S. (Eds.). (1997). *Plato: Complete works*. Hackett Publishing.
- Corry, L. (2004). *Modern algebra and the rise of mathematical structures*. Springer.
- Croce, M., & Baghramian, M. (2024). Experts—Part II: The sources of epistemic authority. *Philosophy Compass*, 19(9–10), Article e70005. <https://doi.org/10.1111/phc3.70005>
- Davies, A., Veličković, P., Buesing, L., Blackwell, S., Zheng, D., Tomašev, N., Tanburn, R., Battaglia, P., Blundell, C., Juhász, A., Lackenby, M., Williamson, G., Hassabis, D., & Kohli, P. (2021). Advancing mathematics by guiding human intuition with AI. *Nature*, 600, 70–74. <https://doi.org/10.1038/s41586-021-04086-x>
- Davis, E. (2021). *Deep learning and mathematical intuition: A review of (Davies et al. 2021)*. <https://arxiv.org/abs/2112.04324>
- De Toffoli, S., & Giardino, V. (2014). An inquiry into the practice of proving in low-dimensional topology. In G. Lolli, M. Panza, & G. Venturi (Eds.), *From logic to practice: Italian studies in the philosophy of mathematics*. Boston Studies in the History and the Philosophy of Science (Vol. 308, pp. 315–336). Springer.
- Dung, L. (2025). Understanding artificial agency. *The Philosophical Quarterly*, 75(2), 450–472. <https://doi.org/10.1093/pq/pqae010>
- Edwards, P. N. (2013). *A vast machine: Computer models, climate data, and the politics of global warming*. MIT Press.
- Ernest, P. (1998). *Social constructivism as a philosophy of mathematics*. State University of New York.
- Fawzi, A., Balog, M., Huang, A., Hubert, T., Romera-Paredes, B., Barekatin, M., Novikov, A., R. Ruiz, F. J., Schrittwieser, J., Swirszcz, G., & Silver, D. (2022). Discovering faster matrix

- multiplication algorithms with reinforcement learning. *Nature*, 610(7930), 47–53. <https://doi.org/10.1038/s41586-022-05172-4>
- Ferreirós, J. (2016). *Mathematical knowledge and the interplay of practices*. Princeton University Press.
- Floridi, L. (2025). AI as Agency without Intelligence: On artificial intelligence as a new form of artificial agency and the multiple realisability of agency thesis. *Philosophy & Technology*, 38, Article 30. <https://doi.org/10.1007/s13347-025-00858-9>
- Friedman, M., & Kish Bar-On, K. (2026). Rethinking mathematical intuition in the age of AI. *The Mathematical Intelligencer*. <https://doi.org/10.1007/s00283-026-10561-y>
- Galison, P. (1997). *Image and logic: A material culture of microphysics*. University of Chicago Press.
- Galison, P., & Hevly, B. (Eds.). (1992). *Big science: The growth of large-scale research*. Stanford University Press.
- Gieryn, T. F. (1999). *Cultural boundaries of science: Credibility on the line*. University of Chicago Press.
- Glazer, E., Erdil, E., Besiroglu, T., Chicharro, D., Chen, E., Gunning, A., Olsson, C. F., Denain, J.-S., Ho, A., de Oliveira Santos, E., Järvinen, O., Barnett, M., Sandler, R., Vrzala, M., Sevilla, J., Ren, Q., Pratt, E., Levine, L., Barkley, G., ... Wildon, M. (2024). *Frontiermath: A benchmark for evaluating advanced mathematical reasoning in AI*. arXiv. <https://doi.org/10.48550/arXiv.2411.04872>
- Gorenstein, D. (1982). *Finite simple groups: An introduction to their classification*. Plenum Press.
- Grinstein, L. S., & Campbell, P. J. (Eds.). (1987). *Women of mathematics: A biobibliographic sourcebook*. Greenwood Press.
- Hadamard, J. (1945). *The psychology of invention in the mathematical field*. Princeton University Press.
- Hamami, Y. (2023). Agency in mathematical practice. In B. Sriraman (Ed.), *Handbook of the history and philosophy of mathematical practice* (pp. 1–19). Springer. https://doi.org/10.1007/978-3-030-19071-2_48-1
- Harris, M. (2025). The FrontierMath scandal. *Silicon Reckoner*. <https://siliconreckoner.substack.com/p/the-frontier-math-scandal>
- Harris, M. (2026a). Quick clarification on authorship. *Silicon Reckoner*. <https://siliconreckoner.substack.com/p/quick-clarification-on-authorship>
- Harris, M. (2026b). Mathematics is not rational, really! *Silicon Reckoner*. <https://siliconreckoner.substack.com/p/mathematics-is-not-rational-really>
- Hartnett, K. (2026). *The proof in the code: How a truth machine is transforming math and AI*. Quanta Books.
- Hauswald, R. (2021). *Epistemische Autoritäten: Individuelle und plurale [Epistemic Authorities: Individual and Plural]*. Springer.
- Hauswald, R. (2025). Artificial epistemic authorities. *Social Epistemology*, 39(6), 716–725. <https://doi.org/10.1080/02691728.2025.2449602>
- Heintz, B. (2000). *Die Innenwelt der Mathematik. Zur Kultur und Praxis einer beweisenden Disziplin [The Inner World of Mathematics: On the Culture and Practice of a Discipline of Proof]*. Springer.
- Hutson, J. (2025). Human-AI collaboration in writing: A multidimensional framework for creative and intellectual authorship. *International Journal of Changes in Education*. Advance online publication. <https://doi.org/10.47852/bonviewIJCE52024908>
- IEA. (2025). *Energy and AI*. Author. <https://www.iea.org/reports/energy-and-ai>
- Jäger, C. (2025). Epistemic authority. In J. Lackey & A. McGlynn (Eds.), *Oxford handbook of social epistemology* (pp. 63–84). Oxford University Press.

- Jaton, F. (2021). *The constitution of algorithms: Ground-truthing, programming, formulating*. MIT Press.
- Knorr Cetina, K. (1999). *Epistemic cultures: How the sciences make knowledge*. Harvard University Press.
- Larvor, B. (2019). From Euclidean geometry to knots and nets. *Synthese*, 196(7), 2715–2736. <https://doi.org/10.1007/s11229-017-1558-x>
- Latour, B. (1987). *Science in action: How to follow scientists and engineers through society*. Harvard University Press.
- Leiden Declaration. (2026). *Leiden Declaration on Artificial Intelligence and Mathematics*. <https://leidendeclaration.ai/>
- Liljedahl, P., & Sriraman, B. (2006). Musings on mathematical creativity. *For The Learning of Mathematics*, 26(1), 17–19.
- Mac Lane, S. (1971). *Categories for the working mathematician*. Springer.
- Mancosu P., Jørgensen K. F., & Pedersen S. A. (Eds.). (2005). *Visualization, explanation and reasoning styles in mathematics*. Springer.
- Marxen, H., & Buntrock, J. (1990). Attacking the busy beaver 5. *Bulletin of the European Association for Theoretical Computer Science*, 40, 247–251.
- Miao, Q., & Wang, F. Y. (2024). *Artificial intelligence for science (AI4S)*. Springer.
- O’Neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
- Parshina, K. (2023). Philosophical assumptions behind the rejection of computer-based proofs. *KRITERION – Journal of Philosophy*, 37(2–4), 105–122. <https://doi.org/10.1515/krt-2022-0015>
- Pellegrino, G. (1970). Sul massimo ordine delle calotte in $S_{4,3}$. *Matematiche (Catania)*, 25, 149–157.
- Pickering, A. (1995). *The mangle of practice: Time, agency, and science*. University of Chicago Press.
- Polu, S., Han, J. M., Zheng, K., Baksys, M., Babuschkin, I., & Sutskever, I. (2022). *Formal mathematics statement curriculum learning*. arXiv. <https://doi.org/10.48550/arXiv.2202.01344>
- Polu, S., & Sutskever, I. (2020). *Generative language modeling for automated theorem proving*. arXiv. <https://doi.org/10.48550/arXiv.2009.03393>.
- Ravetz, J. R. (1971). *Scientific knowledge and its social problems*. Oxford University Press.
- Remmert, V., & Waldecker, R. (2026). Big mathematics – reflections on the history of the classification of finite simple groups, 1950s to 1980s. *European Mathematical Society Magazine*, 140, 5–14. <https://doi.org/10.4171/mag/296>
- Romera-Paredes, B., Barekatin, M., Novikov, A., Balog, M., Kumar, M. P., Dupont, E., R. Ruiz, F. J., Ellenberg, J. S., Wang, P., Fawzi, O., & Kohli, P. (2024). Mathematical discoveries from program search with large language models. *Nature*, 625(7995), 468–475. <https://doi.org/10.1038/s41586-023-06924-6>
- Rowe, D. E., & Koreuber, M. (2020). *Proving it her way: Emmy noether, a life in mathematics*. Springer.
- Siegmund-Schultze, R. (2009). *Mathematicians fleeing from Nazi Germany: Individual fates and global impact*. Princeton University Press.
- Shapin, S. (1989). The invisible technician. *American Scientist*, 77(6), 554–563.
- Shapin, S., & Schaffer, S. (2011). *Leviathan and the air-pump: Hobbes, Boyle, and the experimental life*. Princeton University Press.
- Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M. A., Holeccek, B., Koomey, J. G., Masanet, E. R., & Sartor, D. A. (2024). *2024 United States data center energy usage*

- report. Lawrence Berkeley National Laboratory (LBNL-2001637). <https://doi.org/10.71468/P1WC7Q>
- Sorgner, H. (2026). Symbolic capital or social glue? Recognition in collaborative high-energy physics research. *Social Studies of Science*. Advance online publication. <https://doi.org/10.1177/03063127261474439>
- Srdarov, S., & Leaver, T. (2024). Generative AI glitches: The artificial everything. *M/C Journal*, 27(6). <https://doi.org/10.5204/mcj.3123>
- Steingart, A. (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. <https://doi.org/10.1177/0306312712436547>
- Strubell, E., Ganesh, A., & McCallum, A. (2019). *Energy and policy considerations for deep learning in NLP* [Conference session]. Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy (pp. 3645–3650). Association for Computational Linguistics.
- Taylor, Z. W., & Stan, K. (2024). Exploring the stratified nature of artificial intelligence research funding in United States educational systems: A bibliometric and network analysis. *Education Sciences*, 14(11), Article 1248. <https://doi.org/10.3390/educsci14111248>
- Trinh, T. H., Wu, Y., Le, Q. V., He, H., & Luong, T. (2024). Solving olympiad geometry without human demonstrations. *Nature*, 625, 476–482. <https://doi.org/10.1038/s41586-023-06747-5>
- Tyrrell, F. (2022). *New lower bounds for cap sets*. arXiv. <https://arxiv.org/abs/2209.10045>
- Urman, A., Smirnov, I., & Lasser, J. (2024). The right to audit and power asymmetries in algorithm auditing. *EPJ Data Science*, 13(1), Article 19. https://epjds.epj.org/articles/epjdata/abs/2024/01/13688_2024_Article_454/13688_2024_Article_454.html
- Weber, K., & Tanswell, F. S. (2022). Instructions and recipes in mathematical proofs. *Educational Studies in Mathematics*, 111(1), 73–87. <https://doi.org/10.1007/s10649-022-10156-2>
- Weinberg, A. (1961). Impact of large-scale science on the United States. *Science*, 134(3473), 161–164. <https://doi.org/10.1126/science.134.3473.161>
- Wilholt, T. (2013). Epistemic trust in science. *The British Journal for the Philosophy of Science*, 64(2), 233–253. <https://doi.org/10.1093/bjps/axs007>
- Wink, R., Di Meglio, A., Streit-Bianchi, M., & Liyanage, S. (2024). Big science and social responsibility of the digital world. In S. Liyanage, M. Nordberg, & M. Streit-Bianchi (Eds.), *Big science, innovation, and societal contributions: The organisations and collaborations in big science experiments* (pp. 264–286). Oxford Academic. <https://doi.org/10.1093/oso/9780198881193.003.0012>
- Zagzebski, L. (2015). *Epistemic authority: A theory of trust, authority, and autonomy in belief*. Oxford University Press.
- Zhao, J., Zhang, T., Sun, J., Tian, M., & Huang, H. (2025). *Pi-GPS: Enhancing geometry problem solving by unleashing the power of diagrammatic information*. arXiv. <https://doi.org/10.48550/arXiv.2503.05543>
- Zhang, C., Song, J., Li, S., Liang, Y., Ma, Y., Wang, W., Zhu, Y., & Zhu, S.-C. (2026). Proposing and solving olympiad geometry with guided tree search. *Nature Machine Intelligence*, 8, 84–95. <https://doi.org/10.1038/s42256-025-01164-x>

Author Biographies

Michael Friedman is a historian of mathematics and science whose research focuses on conceptions of materiality and visual practices from the 17th to the 21st century. He is currently an associate researcher at Bonn University, and has held positions at Tel Aviv University, Humboldt

University Berlin, the Fourier Institute in Grenoble, and the Max Planck Institute for Mathematics. His current project investigates how computer-generated visualizations and AI-generated diagrams are transforming mathematical visualization, intuition, and authority in contemporary practice.

Kati Kish Bar-On is a philosopher of mathematics and social science working at the intersection of AI, cognitive science, and social behavior. She is an Alfred P. Sloan Meta-Science and AI post-doctoral fellow at Boston University, a visiting scholar at Harvard Business School, and a lecturer at Northeastern University. Previously a Marie Skłodowska-Curie Global Fellow at MIT's Science, Technology, and Society program, her work examines how AI-based technologies alongside social and cognitive structures jointly shape mathematical knowledge and reasoning.