Rethinking Mathematical Intuition in the Age of AI

Michael Friedman and Kati Kish Bar-On

It goes without saying that mathematics is traditionally viewed as the pinnacle of human intellectual achievement and that it is characterized by rigorous proofs and profound theoretical insights. However, with the emergence of AI-based technologies that solve complicated problems and produce formal proofs, mathematics as a research field has been undergoing a fundamental transformation. The question at hand is whether these technologies can go beyond helping with mathematical tasks to significantly changing the nature of mathematical practice and the process of creation and discovery. Here, we explore the more philosophical implications of these AI-based technologies in mathematics, focusing on how they challenge the concept of *mathematical intuition*. To show why we focus on this notion, a recent example is worth inspecting, since it highlights how such technologies transform mathematics in practice: the recent solution of Erdős problem #90, or better formulated, finding a counterexample to Erdős's conjecture.

Posed by Paul Erdős in 1946, the planar unit distance problem asks a simple geometric question: If you place n points on a flat two-dimensional plane, what is the maximum number of pairs that can be exactly 1 unit of distance apart? Erdős predicted that the maximum number of pairs could grow only slightly faster than linearly, bounded by $n^{1+o(1)}$ (specifically $n^{1+O(1/\log\log n)}$). For eight decades, mathematicians assumed that a standard square grid (where points are lying in a maximal way on circles of radius 1) was the unique best configuration. In May 2026, an internal *OpenAI* model disproved Erdős's conjecture by proving the existence of configurations with $\geq n^{1+\delta}$ (where $\delta > 0$ is a universal constant) points.[1] The model constructed an infinite unramified tower of totally real algebraic number fields, which was a technique that had not been employed in former (human-led) attempts to approach the problem, and used methods usually neither used nor known in this field.[2] By adjoining the imaginary unit i, these fields are used to produce high-dimensional lattices. These lattices are then projected down onto the 2-dimensional plane, and one obtains the configurations of points that disprove the conjecture. However, as noted, all the attempts done before were directed toward proving the conjecture. The reactions to the counterexample ranged from amazement to anxiety [2]; indeed, the model's success came from a combination of persistence, technical breadth, and unconventional research choices. The strength of this counterexample lies in pursuing lines of inquiry demanding both wide-ranging technical knowledge and techniques that human mathematicians might abandon (and have indeed abandoned or ignored, since they were usually considered, if they were even considered, to be unpromising, time-consuming, or too laborious.

Other examples, of course, could be given,[3] but what this example already shows us is that, to follow Jacob Tsimerman, "it is hard to think through the analytic regimes [of this construction] and retain guiding intuition" [2, p. 16]; that is, while all the components of the counterexample were known, their combination and implementation was unexpected. Even if one describes such a solution (or other solutions achieved by AI-based technologies) as a maverick one, or as just relying on perseverance, the mathematicians reacting to this solution—and also to other solutions, as we will see below—explicitly state that their intuition is slipping away; the solution has not (yet) made sense to mathematicians. But what exactly does this assertion mean? How can we describe this shift in intuition compared to other, older, discoveries of "nonintuitive" mathematical objects such as non-Euclidean geometries and imaginary numbers?

A first step toward answering this question is the realization that the development of such intuition comes—if it indeed could be even developed—much more retroactively; that is, it becomes a *reaction*. In an age in which every week several Erdős problems are being solved (or claimed to have been solved) by AI-based technologies, one can claim, as Terence Tao argues (see the section

[1] See the blog post https://openai.com/index/model-disproves-discrete-geometry-conjecture/ (last accessed on June 10, 2026).

[2] For example, the Golod–Shafarevich theorem, stating when class field towers can be infinite.

[3] For example, Erdős Problem #1196 was solved by OpenAI's GPT-5.4 Pro in April 2026 by introducing unexpected mathematical connections that human experts had overlooked for 60 years. Other achievements should also be mentioned: *AlphaTensor* was used to discover new, more efficient algorithms for multiplying 5×5 matrices over $\mathbb{Z}/2\mathbb{Z}$ [19]; the well-known project *First Proof*; the formalization of proofs of the densest sphere packing in 8 and 24 dimensions with a specialized AI program, *Gauss*, developed by *Math, Inc.*; or the list of 67 mathematical problems, most of which were solved (or partially solved) by *AlphaEvolve* [22]. See also https://terrytao.wordpress.com/2025/11/05/mathematical-exploration-and-discovery-at-scale/ (last accessed on June 10, 2026).

https://doi.org/10.1007/s00283-026-10561-y

below on relocating intuition), that we are entering an era of "proof abundance," in which proof generation and verification are cheap but making sense of them is not; our case studies in this paper put this claim under the lens of mathematical intuition. This claim, which we will explore in this paper, leads us to argue that such an intuition—often associated with distinctively *human* reasoning and *actions*—is being transformed by today's AI-based technologies.[4]

Our starting point is simple: as AI-based technologies continue to encroach on areas once thought to be exclusively human, one must examine such developments critically.[5] We suggest that AI-based technologies' ability to explore vast problem spaces and find connections or proofs that humans may overlook (or cannot even find) suggests that these technologies do more than offer merely "brute-force" solutions or passive calculating power. These programs are *actively* engaged in the process, affecting it not only quantitatively but also qualitatively, calling on us to rethink mathematical intuition. While various papers have already suggested rethinking mathematical practices in light of these developments,[6] the notion of mathematical intuition has received little attention; our contribution is to show that the *temporal placement* and *targets* of intuition are changing in systematically different ways across AI workflows, and that this forces us to revise standard accounts of intuition as anticipatory and purely human. We begin with tracing the historical plasticity of mathematical intuition, situating contemporary usage of the term as a plural, practice-sensitive concept that is nevertheless often treated as structurally prior to proof and discovery; moreover, we claim that even when nonintuitive discoveries have been made, they were always embedded in the human context of mathematical research and practice, enabling them to be used and eventually accepted as "making sense." We then present a brief review of the history of mathematical machines to foreground earlier tensions between mechanical procedure and intuition, before examining three recent AI-assisted mathematical episodes (*AlphaGeometry*, knot theory, and *FunSearch* on cap sets) to show how different human–machine configurations reorganize mathematical intuition. We conclude by arguing that AI does not abolish intuition but relocates it, expanding it from a human and mainly anticipatory faculty into a guided, retroactive, or co-constitutive faculty, which is also at the risk of disappearing due to the possible future abundance of true but seemingly meaningless results.

Historical Review: Mathematical Intuition—a Flexible and Evolving Concept

First a word of clarification is in order: we are not going to present here a full-blown historical survey of the notion of intuition in mathematics, nor all the narratives related to it. We intentionally omit these histories not because they are not important, but because our goal here is to focus on the historical plasticity of the notion, in order to understand its potential future flexibilities with the rise of AI. Given this caveat, one can begin with Immanuel Kant, since it was he who (re)shaped such a notion of intuition in the late eighteenth century—in the sense of *Anschauung*—in the Western world during the nineteenth and twentieth centuries. Kant's views on the pure form of intuition (*Anschauung*) and its relation to mathematics, and specifically regarding pure space as a formal condition for any sensory intuitions, made an enormous impact on the mathematical community. Such a conception necessarily tied geometry as well as geometrical concepts with the spatiotemporal ability of human beings to recognize objects with their senses. However, the discovery of non-Euclidean geometries during the 1830s complicated this understanding of intuition, especially since for Kant, space was the three-dimensional *Euclidean* space. Hence, not only that "the" space as such no longer "naturally" existed or was given, but also establishing the consistency of these newly discovered (as well as Euclidean) geometries became unclear and debatable. This led to a rethinking of the status of spatial mathematical intuition [18].

The discovery of projective geometry, higher-dimensional spaces, and "monstrous" functions also undermined the status of mathematical intuition [50, 51], leading directly to the early-twentieth-century debates about how intuition should be constrained and redefined, and culminating in Hans Hahn's 1932 lecture "crisis of intuition" [25]. These discoveries and crises show already how the notion of intuition was flexible and how it radically changed during the late nineteenth century. One approach to deal with the changes of the status of intuition was to turn to methods aligned with mathematical formalism (seen in its extreme form, considering "mathematics as a game of symbols"); another approach emphasized intuition as referring only to palpable visuality (*Anschaulichkeit*), highlighting the plasticity of mathematical intuition: Felix Klein, for example, introduced a distinction between "naïve" and "refined" intuition and even between various forms of intuition [30]. The late twentieth century repositioned intuition as distinctly human, particularly following the reactions to

[4]To emphasize: not every attempt to use AI-based technologies or LLMs leads necessarily to a breakthrough; as we will see, *AlphaGeometry* (to be discussed below) cannot solve every problem; some of the problems in the list of 67 mathematical problems mentioned in the former footnote were not solved or improved.

[5]Indeed, Feng et al., presenting 13 Erdős problems co-solved with AI-based technologies, emphasize that "we caution against over-excitement about its mathematical significance. Any of the open questions answered here could have been easily dispatched by the right expert" [20, p. 8].

[6]For example: [4, 7, 13, 14, 21, 23, 29, 31, 36, 37, 46, 52]. See also the various contributions in the two special issues of the *Bulletin of the AMS* Num. 61, vols. 2 and 3 from 2024: https://www.ams.org/journals/bull/2024-61-02/ and https://www.ams.org/journals/bull/2024-61-03/ (last accessed on June 10, 2026).

computer-assisted proofs and visualizations. Following the objections raised against such proofs during the 1970s and later,[7] intuition was considered "as a lived, embodied, and sensory capacity which distinguishes humans from machines" [39, pp. 197–198], that is, as a form of experience akin to sensory perception, diagrammatic and visual reasoning, and at the same time as what may be directed toward abstract objects rather than concrete, material ones.[8]

This brief and selective historical overview already shows that mathematical intuition has never been a fixed notion with a stable referent. It evolved to function as an umbrella term for several epistemically and methodologically distinct roles within mathematical practice, and many of these roles preserve an anticipatory dimension, even when intuition is treated as trainable and historically variable rather than as a fixed psychological type. One of the reasons for intuition's historical plasticity is that mathematicians have been reacting for centuries to non- or counterintuitive results—from the irrationality of the diagonal of a square and imaginary numbers to Torricelli's horn, non-Euclidean geometries, and the above-mentioned "monstrous" functions. To recall: imaginary numbers were termed by Gerolamo Cardano in 1545 a "mental torture" and described pejoratively by various mathematicians (e.g., by Descartes as "imaginary" and by Lambert in 1770 as "unthinkable non-thing," in his letter to Kant [28, p. 86]). But nevertheless, their usage and conception emerged from human practices and practical problems (solving cubic equations), and they were eventually absorbed into mathematical understanding. With AI-based technologies, a change is happening on a different level: a nonhuman actor participates generatively in the mathematical practice, producing results that mathematicians must then interpret or assess for their fit with existing intuitions, if these actually exist or can be developed. Since the interpretation of such results may become opaque, such an interpretation is, as we will emphasize in the examples below, *always* a retroactive interpretation of results that are obtained not only in a *nonhuman* context but in masses that can hardly be digested. The emergence of such tools therefore raises the question whether intuition should still be understood as anticipatory and as a tool for discovery. We will address the question in the next section through a careful examination of three examples of AI-generated mathematics.

Mathematical Intuition in AI-Pervaded Mathematics

The idea that mathematics might be aided by or relegated to an automated machine appeared long before Turing's 1948 paper "Intelligent Machinery" [48]; this idea had been mentioned already by Leibniz, who spoke of "an arithmetical instrument, which carries over all the work of the mind to wheels" [33, pp. 487–488].[9] Leibniz's stepped reckoner was one of the first examples of an automated machine, though he may have been referring to Blaise Pascal's Pascaline. Both machines aimed to fulfill the dream of a powerful calculator that would save mathematicians from their "slavery"[10] by removing the burden of calculation. In the nineteenth century, Charles Babbage's *Analytical Engine*, developed in the 1830s, represented a conceptual leap toward the ideal of universal mechanical computation. Another step in such mechanization was introduced by William Stanley Jevons's *Logical Piano*, developed during the 1860s, which attempted to mechanize logical reasoning itself. These new technologies, however, were not accepted without criticism. Henri Poincaré criticized such purely mechanical approaches to mathematics by comparing Jevons's piano to Chicago sausage-processing machines [41, p. 147], thereby highlighting the tension between formal mechanical procedures and mathematical intuition, and insisting that "only intuition could enable us to distinguish among" logical rules [42, p. 473]. Such a critique, though it may seem banal, already indicates the place reserved for the mathematician's intuition.

In the 1950s, computer-based verifications of a finite number of cases for general theorems were taken to provide a positive indication that the theorem in question was correct, but they were not considered a proof.[11] Moreover, a new epistemic situation emerged in the 1970s, when the computer-assisted proof of the four-color theorem was presented. The four-color theorem was the first major theorem to be proved with extensive computer assistance, and as such it generated extensive discussion, raising objections regarding whether it should even count as a proof [38]. Eventually it was accepted by the community,[12] yet the proof remained unverifiable, in the sense that human mathematicians could not have physically and time-wise done the same computer calculations. One may argue that this concern still applies to the use of AI-based technologies,[13] since they produce proofs that sometimes cannot be

[7] See the discussion below on the four-color theorem.

[8] See [9, 10, 16, 24].

[9] "… in Arithmetico instrumento, quod omnem animi laborem in rotas transfert."

[10] Leibniz describes in the 1685 *Machina arithmetica* the machine's various uses for geometers, astronomers, and opticians: "For it is unworthy of excellent men to lose hours like *slaves* in the labor of calculation, which could be safely relegated to anyone else if the machine were used" (cited from: [44, p. 181]). See also [27].

[11] This was indeed the case for the various verifications of Fermat's last theorem: in 1954 it was verified by computer that for the first 2,000 integers n for $n > 2$, the equation $a^n + b^n = c^n$ does not have a solution in positive integers [32].

[12] See [3].

[13] Concerning the non-surveyability problem, see [7].

verified. Against this background, what is the difference between the four-color theorem proof and AI-generated proofs? One striking difference is that in contrast to Leibniz's machine, AI-based technologies are no longer merely *passive* devices that can only calculate or automate manual procedures. Such technologies are *generative* in the sense that they provide new outcomes that mathematicians are not necessarily able to come up with on their own. The mathematician herself is hence no longer considered to be the sole author of the mathematical proof or text, since other, nonhuman, actors take part in the proof of such theorems.[14] This raises two practical and philosophical issues. First, a general issue concerning most, if not all, of the results obtained by AI-based technologies is the issue of computational opacity: in contrast to Poincaré's sausage machine, are we even able to understand how these AI-based technologies reached their output? The second issue is more specific to mathematics, concerning the possible un-explainability of the results. Who can make sense of AI-produced proofs? Can the mathematician's intuition be guided by the results provided by the machine? Is she re- or co-writing the proof, partially following the machine's instructions? Will mathematical intuition always come retroactively? Below we illustrate possible answers with three episodes.

These episodes will show, unlike the historical casting of intuition as usually preceding formal reasoning or what reframes it inside the human mathematical fabric, how AI-based tools challenge the anticipatory orientation of mathematical intuition by producing results that lack prior human anticipation and human context, forcing intuition—*if* and *when* the mathematician claims to have one concerning the obtained results—to operate retroactively and in a co-constitutive manner. It is important to note that this retroactive stance does not appear uniformly; the various cases reviewed below reveal different (temporal) modes of co-constitution of intuition between human and machine, each with a distinct temporal and epistemic structure. But each case shows a major shift in temporality, though different from case to case, and indicates a decoupling between the process of proving and the process of "digestion," making sense and developing intuition.[15] We hence focus on these specific examples, some coming from domains that traditionally rely heavily on geometric intuition or visual reasoning (such as Euclidean geometry and knot theory). Such examples may be framed as presenting only maverick solutions or statistical correlations, but we claim that they point toward a more radical change in how mathematics is practiced and understood.

Retroactive and Co-constitutive: The Case of AlphaGeometry

We begin with one of the most widely discussed recent examples: solving problems from the International Mathematical Olympiad (IMO). Trinh et al. [47] report that these AI-based systems reach what they describe as a silver-medal standard, eventually attaining gold-medal standard [1, 8]. The problems in question were solved using *AlphaProof*, *AlphaGeometry*, and *AlphaGeometry* 2, or more recently, with the AI-based system *Aristotle*. These systems can generate the auxiliary constructions needed for a solution on their own, rather than relying on human-provided demonstrations [47, p. 478]. Concretely, once a problem is given, *AlphaGeometry* transforms it into formalized expressions. Then, a "symbolic engine" computes the "deduction closure," i.e., the set of all deducible facts given a core set of initial facts. If this does not solve the problem, *AlphaGeometry* operates various large language models (LLMs) to add auxiliary points or lines to the problem, then activates the "engine" again, and this process repeats: the LLMs suggest constructions, the symbolic engine attempts to prove the theorem with them, until, if possible, a complete and valid proof is found.[16]

To give a concrete example, for one Euclidean-plane geometry problem, the system produced a solution within 19 seconds of receiving the formalized problem (an extremely short time by human standards). This was done by introducing an auxiliary construction that was described by Timothy Gowers as a "non-obvious construction,"[17] suggesting that these tools may alter not only the pace but also the character of mathematical problem-solving. Such descriptions are not to be found only in Gowers's statement. Other actors elevate the "non-obviousness" character of *AlphaGeometry*'s solutions to be a "non" or

[14]Such a shift in the role of the mathematician is seen more clearly with programs such as *Lean* and other automated proof assistants, such as *Coq*, which were developed during the 2010s [6]; for the history of the attempts to mechanize proofs in the twentieth century, see also [34]. Yet one cannot call these technologies "generative," but merely "symbolic AI" [5]. To recall: Doron Zeilberger's paper, coauthored with *Shalosh B. Ehad* (his UNIX workstation computer) and "Euclid" [26], was initially rejected for having unacceptable authorship; in the age of *Lean* and AI-based technologies, such mixed authorship is no longer a playful provocation but a practical reality that forces us to rethink how intuition and authorship align.

[15]The different stances regarding intuition not only reflect the flexibility of this notion (see the historical review above) but also stem inherently from the different technologies used. Indeed, it is almost banal to claim that every technological change may lead to a change in our physical capacities and mental faculties, but in the case of AI-based technologies, one must differentiate between the various technologies being employed and not just call them "AI."

[16]See also [35, pp. 32–33].

[17]See https://deepmind.google/discover/blog/ai-solves-imo-problems-at-silver-medal-level/ (last accessed on: September 15, 2025).

"counterintuitive" one. Similarly to how the counterexample of Erdős problem #90 was considered, Trinh et al. describe some solutions as having a "much less intuitive character" [8], and while presenting *AlphaGeometry* 2 and one of the solutions proposed by it, they note that "at first, these constructions seem very *counter*-intuitive, since most humans would not construct these points …" [8, p. 27]; they also claim that their geometry experts and IMO medalists "consider many AlphaGeometry solutions to exhibit *superhuman creativity*" [8, p. 14].

Nevertheless, it is clear that one cannot take these statements at face value, and indeed, as Terence Tao [45] correctly points out, the comparison between human abilities and AI-based methods is inadequate. With *AlphaGeometry*, claims Tao, one does not (and cannot) measure a mathematical capability comparable to those of humans, but rather something categorically different, namely formal optimization with almost unlimited (spatial and temporal) resources, trained on millions of proofs, for which the question of understanding *why* constructions work, let alone having intuition about them (and not merely verifying them), becomes irrelevant.[18]

What, then, is the role of the human mathematician in this constellation? As Tao implies, there is a potential "flood" of such geometric proofs, due to the availability of almost unlimited resources. How should one engage with and make sense of thousands of proofs not generated by humans? One may suggest that mathematicians will be needed to provide a retroactive account of the sense or the "intuitions behind the auxiliary construction."[19] In this picture, intuition is no longer primarily an anticipatory insight that precedes and enables problem-solving and formalization; it becomes something that can be articulated after the fact, though not necessarily and not always. The AI-proposed solution can hence be considered to demand a form of retroactive intuition or of "sense-making" of unexpected constructions, in order to render them intelligible. This mode of operation is different from the retroactive intuition provided to former counter- or nonintuitive results. In our case, *AlphaGeometry* works only with texts (formalized statements) and not, for example, with geometric diagrams. Visual reasoning and diagrammatic thinking are not abandoned or rejected (as they were in the case of "monstrous" functions); rather, they are simply nonexistent, though they can be produced retroactively. In this sense, geometric or visual intuition is not first guiding discovery or proof, but reentering only *later*, after the proof is presented and the diagram is retroactively drawn by *AlphaGeometry*, to reconstruct meaning and motivation for the proposed construction. Such retroactive intuition is dependent essentially on *AlphaGeometry*'s technology and hence is co-constitutive, since the latter provides a solution basically without the intervention of the mathematician. This mode exemplifies interpretive co-constitution in retrospect: the machine generates first; the human interprets later.

Guided and Co-constitutive: The Case of Knot Theory

Our second example moves from Euclidean geometry to knot theory. In one of the first mathematical papers employing AI-based technologies, with the highly promising title "Advancing mathematics by guiding human intuition with AI" [11], Alex Davies et al. address the connection between two branches of knot theory by finding relations between their invariants, that is, between the *signature* of a knot (a classical topological-algebraic invariant) and a geometric invariant in hyperbolic knot theory (the hyperbolic *volume* of a knot). To do this, they used a neural network to predict the signature based on hyperbolic invariants.

The technology behind this discovery is based on the choice of a dozen geometric quantities as candidate invariants, generating a dataset of around 2.7 million knots. Then a fully connected neural network is trained to predict the signature. Next, by analyzing the neural network itself, the mathematicians used gradient-based attribution methods to inspect which inputs most influenced its predictions. This revealed that most of the predictive power was concentrated in three invariants. Retraining the model using only those three invariants indicated that they capture the essential geometric invariants determining the signature. The mathematicians then found that the fact that the prediction depends primarily on those three invariants is related to the geometric feature of the knot called the cusp. This observation led them to define a new quantity called the natural slope and to make an initial conjecture. After further investigation, they revised their conjecture, leading them to a proof of a relationship between signature, volume, and another geometric invariant, the injectivity radius [11, p. 72].

Whether these results should be attributed primarily to AI-based guidance of human intuition is not straightforward, and this question has been debated on the grounds that similar discoveries might, in principle, have been reached with earlier and less-sophisticated technologies [12]. The mathematician remains an agent in the process, since she is the one who "guides the selection of conjectures that not only fit the data but also look interesting." Even so, given the sheer volume of data that must be processed, it is hard to view the technology as merely a passive instrument. In this case, the AI-based technology detects

[18] It is important to emphasize that *AlphaGeometry* is not able to solve every problem in plane Euclidean geometry. Trinh et al. note that geometry-specific languages are narrowly defined and therefore cannot express many human proofs that draw on tools outside standard Euclidean geometry, such as complex numbers, which appear in at least one human solution to an IMO problem [47, p. 476].

[19] See https://storage.googleapis.com/deepmind-media/DeepMind.com/Blog/imo-2024-solutions/P4/index.html (last accessed on June 10, 2026).

correlations across massive datasets, while the mathematician determines which conjectures are worth exploring. While the mathematician's role is of importance here, we still see a shift in the concept of intuition—similar to the *AlphaGeometry* example[20]—since various knot-theoretic arguments were historically based on human visual reasoning [17] and are now absent from the suggestions made by AI-based programs. Indeed, research on various theorems in knot theory was also based on visual and material arguments (see [15]). Still, in contrast to the *AlphaGeometry* case, intuition here is guided and has an anticipatory role: mathematicians continue to shape the space of relevance and decide which machine-generated proposals are mathematically promising. In that sense, the interaction can be described as co-constitutional: the machine provides a palette of patterns, and the human's intuition curates what becomes mathematically meaningful.

What makes this case significant is a disruption of the order of the inquiry process, turning the process into a "back-and-forth" one, a process that can also be observed in other AI-based discoveries; such a process distinguishes it from earlier uses of statistical computation in mathematics [22, p. 13]. Here, the researchers' own trajectory of inquiry was redirected by the model's outputs, since it was an AI-based tool that first detected the relevance of the cusp-related invariants, and only then did the mathematicians follow that lead toward a conjecture they had not anticipated, using a retrained LLM. Thus intuition, in this case, was partly constituted by the machine, as opposed to being only retroactively reconstructed.

Retroactive and Object-Shifted: The Case of Cap Sets

The third example turns to the area of combinatorics. A *cap set* is defined as a subset of 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 obtain cap sets of size 2^n quickly). In 1970, Giuseppe Pellegrino [40] proved that four-dimensional cap sets have a maximum size of 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 2.218^n by Tyrrell [49]. However, Romera-Paredes et al. [43] managed to improve the lower bound, using an LLM with an evaluator, to 2.2202^n.

To obtain their result, Romera-Paredes et al. employed a procedure called *FunSearch*, short for "searching in the function space," which pairs a pretrained large language model with a systematic evaluator [43, p. 468]. The technical idea is to perform the search not in the space of sets, but in the space of Python programs that generate them: starting with a program and then using an LLM to generate numerous similar but slightly different programs (termed mutations). Such programs are then scored by evaluating the sets they produce. As Georgiev et al. [22, p. 3] point out, the search in program space favors simplicity and structure, which assists in differentiating between the less optimal ("ugly") solutions and the elegant and optimal ones.

In this example, *FunSearch* discovered a larger cap set in dimension ($n = 8$) than previously known: instead of a cap set of size 496 in $n = 8$, the authors discovered a cap set of size 512 in $n = 8$. By iteratively recombining several programs and then feeding them the pretrained LLM to generate new ones, not only the set of 512 eight-dimensional vectors was discovered, but also and *first of all* a program that generates it [43, p. 469]. The first problem that emerges is how human mathematicians can even examine all the thousands of programs generated. Nevertheless, it is only by "inspecting the code, we [the authors] obtain a degree of understanding of what this set is: specifically, manual simplification of [the code] provides the construction" of this set of 512 elements, which "we were able to manually construct thanks to having discovered the cap set by searching in function space" [43, p. 471]. However, there is also a downside to this process, namely the potential loss of interpretability in the search process [22, p. 3]. This is the major point we wish to emphasize: the mathematician's understanding of the result emerged only *after FunSearch* produced a successful program; by inspecting and simplifying the discovered code, they were able to extract a humanly intelligible construction and then manually reconstruct the 512-element cap set [43, p. 471]. This case therefore exemplifies a distinctly temporal pattern: the AI-based technology produces a generating procedure first, only then followed by an object, entailing a possible "loss of interpretability," and mathematical intuition enters in a retroactive mode, as interpretation and reconstruction of what the machine has found.

FunSearch thus presents a stronger form of human–machine hybridity than the cases of knot theory or Euclidean geometry. Intuition and "sense-making" are purely retroactive, also because the space of researched objects has shifted; moreover, Jordan Ellenberg [53], while describing their own collaborations with similar AI-based technologies to *FunSearch*, admit the possibility that human mathematicians may *never* be able to understand—even in retrospect—what was discovered by these technologies. They note that sometimes the "improvements [done by such technologies] are difficult to interpret … [For a specific solution] we have not been able to concoct a satisfying story of 'why,'" and "it is possible that no such story exists." One can read an AI-generated proof as true or take note of an AI-discovered mathematical correlation, but one may not understand the logic leading to this output.

Along with this potential loss, the significance here lies also in the change in the *target* of intuition. One has to develop intuition regarding the *programs* produced

[20]There is also a difference to be noted when comparing this example to *AlphaGeometry*. In the knot theory case, there is no final solution offered, while in the previous case study, the mathematician has to *retroactively* explain the possible intuitions or rationale leading to the solution once a solution is given.

to generate the "final object" (in this case, the cap set), a process that is essentially different from the previous examples. Moreover, the model was trained on the space of *all* computer programs, that is, not on a specific data set of mathematical objects. In this case, AI-based technology does not merely search for the object itself or its properties (as in the case of knot theory) or a possible proof or a counterexample (as in *AlphaGeometry* or Erdős problem #90); it searches in the space of Python programs that generate such objects. The transition from specific objects and proofs to these new spaces prompts a fundamental shift concerning what the mathematician's intuition should focus on.

Relocating Intuition

The above three cases as well as Erdős problem #90 show that mathematical intuition is not necessarily disappearing; it is being pushed into new places. The striking feature of current AI-based practice is not that machines imitate a familiar "flash of insight," but that they flood us with correct or plausible outputs that no human anticipated, forcing intuition either to move downstream into interpretation, sense-making, and conceptual integration or to completely refocus its target objects. What our three cases add to existing discussions of AI and mathematical practice is that this shift is not uniform: it takes three distinguishable shapes—retroactive interpretation (*AlphaGeometry*), guided co-constitution (knot theory), and a radical object shift (*FunSearch*). Even if one casts such advancements as mere combinatorics of already known results or as maverick proofs no one expected, the sheer number of results and their interpretative opacity force a change in how we conceptualize mathematical intuition. What is new compared to earlier "crises" of intuition is not just the strangeness of objects, examples, or proofs, but the decoupling between proof generation, verification, and digestion; intuition has moved to the digestion stage. Historically, intuition could be cast as what comes first—a human capacity that orients inquiry, suggests constructions, and singles out promising directions—or, in a more refined form, as what is distilled after long engagement with problems, examples, and concepts, but still as an internal human faculty. The decades-long attempts on Erdős #90 illustrate the older picture vividly: intuition guided the community and guided them in the wrong direction. In our case studies, by contrast, the crucial intellectual work often begins *after* the system has already generated the proof, pattern, or program.

The three cases reveal not only new mathematical workflows, but also the emergence of "too many proofs," i.e., proof abundance, in which AI-based systems generate and check more arguments than experts can possibly digest. In that setting, mathematics is no longer a "proof discipline," and the rare resource is now "proof digestion," to follow Terence Tao,[21] to connect AI-produced material so that it can actually be integrated into *human* mathematical practice. Intuition, in this view, besides retroactively making sense of a single result, might also be employed to decide which of a thousand formally correct outputs deserves a place in the canon. In *AlphaGeometry*, the system may propose various auxiliary constructions that no human would naturally draw first, and intuition is pushed into retroactive reconstruction and explanation. In the knot theory discovery, a neural network, processing millions of knot invariants, exposes which geometric features actually drive a classical invariant; human anticipatory intuition endures, but as a co-constitutive filter on patterns discovered by the model. With *FunSearch*, what must be understood is not merely a set but its generating program; intuition is retroactively exercised in reexpressing and making sense—if it can even do so—of the code itself.

This reconfiguration has consequences for what it means to make sense of mathematical results. Mathematicians have always cared about more than knowing *that* a theorem holds; they want to see *why*. AI-assisted mathematics exacerbates this gap by uncoupling proving from sense making, delivering correct but possibly opaque arguments that must be worked backward into intelligible form. A retroactive account of mathematical intuition is therefore not an optional philosophical luxury or refinement; it is a practical necessity for mathematicians working alongside these systems. And the stakes are not small: in some future cases, there may be no human context to recover, no story to tell about why. If mathematicians renounce that role, one risks "ruining" mathematics by outsourcing not just proofs, but understanding itself.

Acknowledgments

Kati Kish Bar-On was supported by the Alfred P. Sloan Foundation [grant number G-2025-25317]. The authors contributed equally to this work and are listed in alphabetical order.

Funding Open Access funding enabled and organized by Projekt DEAL.

[21]Following the talk by Terence Tao, "New Mathematical Workflows," May 2, 2026, at the conference *Future of Mathematics Symposium*, Stanford University.

References

[1] T. Achim et al. Aristotle: IMO-level automated theorem proving. arXiv:2510.01346, 2025.

[2] N. Alon, T. Bloom, et al. Remarks on the disproof of the unit distance conjecture. arXiv:2605.20695, 2026.

[3] K. I. Appel and W. Haken. *Every Planar Map is Four Colorable*. American Mathematical Society, 1989.

[4] J, Avigad. *Mathematical Understanding*. Available at https://philsci-archive.pitt.edu/27708/, 2026.

[5] J. Avigad. Mathematics and the formal turn. *Bulletin of the American Mathematical Society* 61:2 (2024), 225–240.

[6] J. Avigad. Automated reasoning for mathematics. In *Automated Reasoning, IJCAR 2024*, edited by C. Benzmüller, M. J. Heule, and R. A. Schmidt. Lecture Notes in Computer Science vol. 14739. Springer, 2024.

[7] Á. Berg and F. Tanswell. The philosophical prospects of large language models in the future of mathematics. Forthcoming in the *Annals of Mathematics and Philosophy*, 2025.

[8] Y. Chervonyi, H. T. Trieu, et al. Gold-medalist performance in solving olympiad geometry with AlphaGeometry2. *Journal of Machine Learning Research* 26:241 (2025), 1–39.

[9] E. Chudnoff. Intuition in Mathematics. In *Rational Intuition*, edited by Linda Osbeck and Barbara Held. Cambridge University Press, 2014.

[10] E. Chudnoff. The nature of intuitive justification. *Philosophical Studies* 153:2 (2011), 313– 333.

[11] A. Davies, P. Veličković, et al. Advancing mathematics by guiding human intuition with AI. *Nature* 600:7887 (2021), 70–74.

[12] E. Davis. Deep learning and mathematical intuition. Review of [11]. Available at https://arxiv.org/abs/2112.04324, 2021.

[13] E. Davis. Mathematics, word problems, common sense, and artificial intelligence. Available at https://arxiv.org/abs/2301.09723, 2023.

[14] W. Dean and A. Naibo. Artificial intelligence and inherent mathematical difficulty. *Philosophia Mathematica* 33:3 (2025), 283–329.

[15] S. De Toffoli and V. Giardino. Forms and roles of diagrams in knot theory. *Erkenntnis* 79:3 (2014), 829–842.

[16] S. De Toffoli. Reconciling rigor and intuition. *Erkenntnis* 86:6 (2021), 1783–1802.

[17] M. Epple. *Die Entstehung der Knotentheorie: Kontexte und Konstruktionen einer modernen mathematischen Theorie*. Springer, 2013.

[18] M. Epple. "Analogies," "Interpretations," "images," "systems," and "models": some remarks on the history of abstract representation in the sciences since the nineteenth century. In *Model and Mathematics: From the 19th to the 21st Century*, edited by M. Friedman and K. Krauthausen, Trends in the History of Science. Birkhäuser, 2022. Available at https://doi.org/10.1007/978-3-030-97833-4_9.

[19] A. Fawzi, M. Balog, et al. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* 610:7930 (2022), 47–53.

[20] T. Feng. Semi-autonomous mathematics discovery with gemini: a case study on the Erdős problems. Available at https://arxiv.org/abs/2601.22401, 2026.

[21] J. L. Gastaldi. How to do maths with words: neural machine learning applications to mathematics and their philosophical significance. In *Handbook of the History and Philosophy of Mathematical Practice*, edited by B. Sriraman. Springer, 2024. Available at https://doi.org/10.1007/978-3-031-40846-5_142.

[22] B. Georgiev, J. Gómez-Serrano, et al. Mathematical exploration and discovery at scale. arXiv:2511.02864, 2025.

[23] W. T. Gowers. Rough structure and classification. In *Visions in Mathematics: GAFA 2000* Special Volume, Part I, pp. 79–117. Birkhäuser, 2000.

[24] V. Giardino. Manipulative imagination: how to move things around in mathematics. *Theoria: An International Journal for Theory, History and Foundations of Science* 33:2 (2018), 345–360.

[25] H. Hahn. The crisis in intuition. In *Empiricism, Logic and Mathematics: Philosophical papers*, pp. 73–102. Springer, 1980.

[26] Z. Hu, Y. Cui, et al. Shalosh B. Ekhad: a computer credit for mathematicians. *Scientometrics* 122 (2020), 71–97.

[27] M. L. Jones. *Reckoning with Matter: Calculating Machines, Innovation, and Thinking About Thinking from Pascal to Babbage*. University of Chicago Press, 2016.

[28] I. Kant. *Immanuel Kants Werke. Band 9: Briefe von und an Kant; 1. 1749–1789*. Cassirer, 1918.

[29] S. Kantamneni and M. Tegmark. Language models use trigonometry to do addition. Available at https://doi.org/10.48550/arXiv.2502.00873, 2025.

[30] F. Klein. *The Evanston Colloquium*. Lectures on Mathematics. American Mathematical Society, 1911.

[31] T. Klowden and T. Tao. Mathematical methods and human thought in the age of AI. arXiv:2603.26524v1, 2026.

[32] D. H. Lehmer, E. Lehmer, and H. S. Vandiver. An application of high-speed computing to Fermat's last theorem. *Proceedings of the National Academy of Sciences* 40:1 (1954), 25–33.

[33] G. W. Leibniz. *Sämtliche Schriften und Briefe*. Berlin-Brandenburgische Akademie der Wissenschaften, Akademie der Wissenschaften zu Göttingen, 1923.

[34] D. MacKenzie. *Mechanizing Proof: Computing, Risk, and Trust*. MIT Press, 2004.

[35] Q. Miao and F. Y. Wang. *Artificial Intelligence for Science (AI4S)*. SpringerBriefs in Service Science. Springer, 2024.

[36] M. Pantsar. Two approaches to developing human-like artificial mathematical intelligence. In *Philosophy of Artificial Intelligence: Optimistic and Pessimistic Views* pp. 165–182. Springer, 2025.

[37] M. Pantsar. Theorem proving in artificial neural networks: new frontiers in mathematical AI. *Euro Journal Phil. Sci*. 14:4 (2024). Available at https://doi.org/10.1007/s13194-024-00569-6.

[38] K. Parshina. Philosophical assumptions behind the rejection of computer-based proofs. *KRITERION – Journal of Philosophy* 37:2-4 (2023), 105–122.

[39] C. Pedwell. Speculative machines and us: more-than-human intuition and the algorithmic condition. *Cultural Studies* 38:2 (2024), 188–218.

[40] G. Pellegrino. Sul massimo ordine delle calotte in S4,3. *Matematiche (Catania)* 25 (1970), 149–157.

[41] H. Poincaré. *Science and Method*. Translated by F. Maitland. Thomas Nelson and Sons, 1914.

[42] H. Poincaré. *The Foundations of Science*. Science Press, 1913.

[43] B. Romera-Paredes, M. Barekatain, et al. Mathematical discoveries from program search with large language models. *Nature* 625 (2024), 468–475. Available at https://doi.org/10.1038/s41586-023-06924-6.

[44] D. E. Smith, editor. *A Source Book in Mathematics*. Dover, 1959.

[45] T. Tao. It is tempting to view the capability of current AI technology.... Mastodon, 2025. Available at https://mathstodon.xyz/@tao/114881418225852441.

[46] A. Testolin. can neural networks do arithmetic? a survey on the elementary numerical skills of state-of-the-art deep learning models. *Applied Sciences* 14:2 (2024), 744. Available at https://doi.org/10.3390/app14020744.

[47] T. H. Trinh, Y. Wu, et al. Solving olympiad geometry without human demonstrations. *Nature* 625 (2024), 476–482. Available at https://doi.org/10.1038/s41586-023-06747-5.

[48] A. Turing. Intelligent machinery (1948). In *The Essential Turing*, pp. 395–432. Oxford University Press, 2004.

[49] F. Tyrrell. New lower bounds for cap sets. arXiv:2209.10045, 2022.

[50] K. Volkert. *Die Krise der Anschauung: eine Studie zu formalen und heuristischen Verfahren in der Mathematik seit 1850*. Vandenhoeck & Ruprecht, 1986.

[51] K. Volkert. Die Geschichte der pathologischen Funktionen – Ein Beitrag zur Entstehung der mathematischen Methodologie. *Archive for History of Exact Sciences* 37/3 (1987), 193–232.

[52] J. O. Weatherall and J. Wolfson. Correctness, artificial intelligence, and the epistemic value of mathematical proof. https://arxiv.org/html/2602.12463v1, 2026.

[53] J. Ellenberg, N. Libedinsky, et al. Bruhat intervals that are large hypercubes. arXiv:2601.01235, 2026.

Michael Friedman, University of Bonn, Bonn, Germany. E-mail: friedman@math.uni-bonn.de
Kati Kish Bar-On, Boston University, Boston, Massachusetts, USA. E-mail: kkish@bu.edu