

Mathematical Understanding (work in progress)

Jeremy Avigad

November 9, 2025

We are not trying to meet some abstract production quota of definitions, theorems and proofs. The measure of our success is whether what we do enables *people* to understand and think more clearly and effectively about mathematics.

William Thurston,
“Proof and Progress in Mathematics” [43]

Any fool can know. The point is to understand.

attributed to Albert Einstein,
[36, epigraph to Chapter 1]

Abstract

The philosophy of mathematics has traditionally been concerned with knowledge and justification, but the goals of mathematical activity are not limited to making true mathematical statements. Ultimately, we look to mathematics for *understanding*. This chapter considers the ways we think about mathematical understanding in our assessments of historical and contemporary mathematical results, in the use of computers in mathematics, and in mathematics education. It argues that a theory of mathematical understanding should be a theory of *abilities* to think and reason mathematically, and that such a theory should inform, and be informed by, practical applications. This chapter also situates models of reasoning in the context of symbolic and neural AI, which offer complementary perspectives.

1 Introduction

The distinction between knowledge and true belief is central to Western epistemology. In Plato’s dialogue, *Meno*, Socrates asks a slave a series of questions that lead him to discover a theorem of geometry. Socrates uses this episode to argue that acquisition of knowledge is a matter of recalling

forms that our souls encounter before we are born, a process that serves to secure our beliefs and keep them from running away.

... true opinions, as long as they remain, are a fine thing and all they do are good, but they are not willing to remain long, and they escape from a man's mind, so that they are not worth much until one ties them down by giving an account of the reason why. And that, Meno, my friend, is recollection, as we previously agreed. After they are tied down, in the first place they become knowledge, and then they remain in place. That is why knowledge is prized higher than correct opinion, and knowledge differs from correct opinion in being tied down. [32]

For Plato, mathematical objects and concepts are paradigm instances of forms, geometric reasoning is a paradigm of recollection, and deductive argument is the paradigm form of the *logos* that secures knowledge. The philosophy of mathematics has been concerned with the nature of mathematical objects, inference, and proof ever since. These are the focus, for example, of the *The Oxford Handbook of Philosophy of Mathematics* [35], published in 2005.

Concerns about justification seem less pressing today, however, now that the foundational crises in mathematics are a century behind us and axiomatic foundations provide adequate and uncontroversial accounts of the rules of mathematical inference. Gödel's second incompleteness theorem places strong limits as to what we can do to justify our choices of axioms, and twentieth century philosophy of mathematics encourages us to accept that no philosophical justification is necessary. According to the logical positivists, mathematical and logical axioms are true because they are part of the linguistic framework we use to interpret the world, and according to Quine, they are true because they are part of our holistic web of beliefs. Either way, it is fruitless to look for external justification; mathematics is as mathematics does, and it is not the philosopher's job to question it. What, then, is left for a philosopher of mathematics to do?

Even if our choices of axioms and rules are dictated by mathematical considerations, we can still hope to gain more clarity as to why we do mathematics the way we do. Mathematicians try to do mathematics well, and they are generally thoughtful and reflective about their craft. As our first epigraph suggests, the goal of mathematics is not just to rack up definitions, theorems, and proofs, but to achieve mathematical understanding. As philosophers, we should join mathematicians in thinking about what that means. As a first cut, it may be helpful to take the phrase "mathematical understanding" to stand for the things that we value in our mathematical artifacts that go beyond mere knowledge, that is, what we get from mathematical definitions, theories, and proofs, beyond knowing that a collection of mathematical statements are true. This is a way of picking up where twentieth century philosophy of mathematics left off: if the ultimate justification for our axiomatic frameworks is the role they play in supporting mathematics as a whole, we would do well to come to terms with the goals of mathematics and the way our mathematical practices achieve them.

This chapter makes the case that we can be both philosophical and scientific about mathematical understanding, and that our scientific theories should inform and be informed by our philosophical ones. I will advocate for a specific approach to thinking about mathematical understanding that focuses on studying the *abilities*, or capacities, that underlie mathematical reasoning. I will argue that we can do this in a way that abstracts these abilities away from the particular cognitive architectures that implements them, whether they are students, experts mathematicians, symbolic algorithms, neural networks, or even social and historical processes.

We will also consider two complementary views of mathematical understanding that emerge from computational advances and the mechanization of mathematical reasoning. On the perspective of

symbolic AI and automated reasoning, understanding is a matter of having algorithms, symbolic representations, and heuristics that make it possible to carry out complex mathematical tasks, while on the perspective of machine learning and neural networks, understanding is a matter of having distributed representations, gathered from experience, that make it possible to recognize mathematical patterns and synthesize reasoning strategies. We will explore these perspectives and the way they can inform our discourse on understanding.

While there is a growing philosophical literature on scientific understanding, considerably less has been written on the topic of mathematical understanding (e.g., [11, 15, 21, 42]). This chapter is not a survey of the literature, but, rather, revisits themes I have explored in the past, and situates them in the context of recent developments in AI.

2 Motivating Questions

Let us start with some examples of the kinds of questions that a theory of mathematical understanding might help us answer.

2.1 Understanding vs. knowledge

From the traditional point of view, the point of proving a theorem is to establish that the theorem is true, thereby justifying claims to know that it is true. However, it is often the case the multiple proofs of the same theorem appear in the mathematical literature [2, 12, 13].

Proofs convey all sorts of understanding. Consider the following example (discussed at length in [2]). In the *Arithmetic*, Diophantus notes that $5 = 2^2 + 1^2$, $13 = 3^2 + 2^2$, and $5 \times 13 = 65 = 8^2 + 1^2 = 7^2 + 4^2$; in other words, the product of two sums of integers squares is again an integer square. This is a general phenomenon:

Theorem. *If each of x and y is the sum of two integer squares, then so is xy .*

Proof. Suppose $x = a^2 + b^2$ and $y = c^2 + d^2$. Then

$$\begin{aligned} xy &= (a^2 + b^2)(c^2 + d^2) \\ &= a^2c^2 + b^2d^2 + a^2d^2 + b^2c^2 \\ &= a^2c^2 - 2abcd + b^2d^2 + a^2d^2 + 2abcd + b^2c^2 \\ &= (ac - bd)^2 + (ad + bc)^2, \end{aligned}$$

a sum of two squares. □

This proof uses nothing more than high-school algebra, but it renders the result mysterious. For an alternative proof, consider the *Gaussian integers*, which is to say, complex numbers of the form $a + bi$, where a and b are integers. In this system of numbers, -1 has two square roots, i and $-i$. As far as the operations of arithmetic are concerned, the two are indistinguishable; replacing i by $-i$ uniformly in any calculation using addition and multiplication, the calculation remains valid. In particular, if $\alpha = a + bi$, we define the *conjugate* $\bar{\alpha}$ of α to be $a - bi$, and we have that if $\alpha\beta = \gamma$, then $\bar{\alpha}\bar{\beta} = \bar{\gamma}$. This tells us that the conjugate $\bar{\alpha}\bar{\beta}$ of a product is equal to the product $\bar{\alpha}\bar{\beta}$ of the conjugates, for every α and β .

Now, for any Gaussian integer, α , we define the *norm* of α , $N(\alpha)$, by

$$N(\alpha) = \alpha\bar{\alpha} = (a + ib)(a - ib) = a^2 - i^2b^2 = a^2 + b^2.$$

So the norm of $N(a + ib)$ is the sum of squares $a^2 + b^2$, and any sum of squares $a^2 + b^2$ can be expressed as $N(a + ib)$. It is not hard to show that the norm of a product is the product of the norms:

$$N(\alpha\beta) = (\alpha\beta)(\overline{\alpha\beta}) = \alpha \cdot \beta \cdot \bar{\alpha} \cdot \bar{\beta} = \alpha\bar{\alpha} \cdot \beta\bar{\beta} = N(\alpha)N(\beta).$$

With that little bit of theory in hand, our theorem has a short proof.

Proof. Suppose x and y are sums of squares. If $x = N(\alpha)$ and $y = N(\beta)$, then $xy = N(\alpha\beta)$, a sum of two squares. $\square$

Seeing this proof for the first time is typically an aha! moment. The proof elicits a warm feeling inside, even in seasoned number theorists. The ideas—the study of number systems extending the integers and their symmetries—are fundamental to number theory today, and there is a lot we can say about why we like the proof [2]. However, there is nothing like the visceral feeling of having grasped something valuable to make the case that there is a phenomena here that deserves to be explained.

Mark Steiner [39] distinguished between explanatory and non-explanatory proofs, initiating debate on the nature of mathematical explanation [29]. However, not every assessment of proof is cast in terms of explanation; for example, proofs can be assessed on grounds of simplicity, generality, depth, and insight. At least some of these virtues seem to bear on the way the proof contributes to mathematical understanding, and a theory of mathematical understanding should help us make sense of them. This raises a fundamental question: what do we value in mathematical proofs beyond their ability to tell us that a mathematical statement is true?

2.2 Understanding proofs

There is a reciprocal relationship between understanding and proof: a good proof conveys understanding, while reading or discovering a proof *requires* understanding. Understanding a proof involves, in part, being able to explain or justify the correctness of each individual step, but that is not all: it is not unusual for a mathematician to feel that they have been able to verify a proof line by line without fully understanding it. Poincaré sums it up nicely:

Does understanding the demonstration of a theorem consist in examining each of the syllogisms of which it is composed in succession, and being convinced that it is correct and conforms to the rules of the game? In the same way, does understanding a definition consist simply in recognizing that the meaning of all the terms employed is already known, and being convinced that it involves no contradiction?

... Almost all are more exacting; they want to know not only whether all the syllogisms of a demonstration are correct, but why they are linked together in one order rather than in another. As long as they appear to them engendered by caprice, and not by an intelligence constantly conscious of the end to be attained, they do not think they have understood. [33]


The logic of discovery and the logic of justification are thereby linked: understanding a proof involves seeing a process under which the proof could have been discovered. Rebecca Morris has argued that we consider a proof to be *motivated* when the presentation helps us see how each step could have been anticipated, given the context and goal [31]. Yacin Hamami and Morris have argued that understanding a proof requires seeing it as an instance of a rational plan, and being able to reconstruct a rational process by which it is generated [20].

Many authors have noted that formal proofs do not convey the same sort of understanding as an informal proof, and there is a substantial philosophical literature investigating the gap between the formal and the informal. (See [18, 38] for overviews.) A standard view holds that formal proof is an idealized standard of correctness, in the sense that an informal proof is correct if and only if it appropriately warrants the existence of a formal derivation. However, formal derivations can be inordinarily long, and even a single error invalidates the derivation. This poses a clear objection to the standard view: how can our judgements of the correctness of an informal proof possibly warrant the existence of something so complex and fragile? I have described one approach to answering this question [7], which can be described in a nutshell: we don't gain confidence in the correctness of proofs by checking them formally but, rather, we *understand* them.

2.3 Diagrammatic reasoning

Proposition 16 of Book I of Euclid's *Elements* reads "In any triangle, if one of the sides be produced, the exterior angle is greater than either of the interior and opposite angles." Euclid proves the claim for one of the interior angles and notes that the proof for the other is similar. The proof below is essentially his.

Theorem. *Let ABC be any triangle, and extend the segment BC to D . Then the external angle ACD is greater than interior angle BAC .*



Proof. Let E bisect segment AC , and extend the segment BE to F so that BE is congruent to EF . By side-angle-side, triangle ABE is congruent to CFE : AE is congruent to EC by the choice of E , BE is congruent to EF by construction, and angle AEB is congruent to angle CEF because they are vertical angles. As a result, angle BAC is congruent to the corresponding angle ACF . Since angle ACD is clearly greater than angle ACF , we have that ACD is greater than BAC , as required. $\square$

The propositions of the *Elements* consist of *demonstrations*, which establish that a property holds of a given diagram, and *constructions*, which establish that one can construct geometric objects with certain properties with a ruler and compass. Each construction and inference in the proof above is justified explicitly by an appeal to a prior proposition or one of Euclid's axioms. There is one exception, namely, the inference marked "clearly." What justifies that?

Diagrammatic inferences like that are ubiquitous in the *Elements*. Edward Dean, John Mumma, and I provide an analysis [9], building on Ken Manders’ observation [30] that diagrammatic inferences track topological relationships in the geometric data rather than exact metric relationships. In this case, the construction guarantees that E and F are on the same side of line CD and that F and D are on the same side of line AC , which is what we mean when we say that F is inside angle ACD . This implies that angle ACF is smaller than angle ACD . We are somehow able to track these relationships with some measure of reliability as we read the proof and verify the correctness of the inference. The goal of our analysis was not to explain how we do this; rather, it was to characterize the competences that are implicit in the practice. The point is that reading and writing Euclidean proofs requires a certain geometric understanding, and we wanted to describe that understanding. With such an analysis in hand, one can then address the cognitive question as to how we are able to carry out such inferences; see, for example, [22].

The fact that the formal axiomatic tradition fails to account for visual and diagrammatic reasoning has spawned a considerable philosophical literature (for overviews, see [16, 44]). In each particular use case, we can ask about the understanding that a diagram conveys and the understanding that is required to use the diagram effectively. There is nothing special about diagrams; the use of text-based representations in mathematics is equally complex and mysterious [8]. Questions about how we understand diagrammatic proofs and representations are therefore just a part of the larger question of how we understand mathematical proofs and representations generally.

2.4 Historical progress

Mathematicians often claim that some mathematical advance was made possible by a prior conceptual development. For example, number theorists will tell you that Riemann’s introduction of the complex zeta function and the use of complex analysis made it possible for Hadamard and de la Vallée Poussin to prove the prime number theorem in 1896. What is the sense of “possibility” here? Riemann wrote down the definition of the zeta function and derived some of its properties, which was something that required considerable ingenuity, expertise, and insight. We need to make sense of a conceptual space in which writing down a definition can be a difficult and important achievement, one that can, in turn, open up new avenues of reasoning.

The philosopher of mathematics, Ken Manders, was an advocate of the method of historical case studies as a means to philosophical analysis. Start with any episode in the history of mathematics that is commonly viewed as an important advance, like the invention of calculus or analytic geometry. Calling it an advance implies that something about our state of knowledge, or understanding, changed over the course of the development. If we can figure out how to characterize the relevant change in epistemic state and why that change was valuable, we will have learned something about what is important to mathematics. Historical analysis can thus contribute to our philosophical understanding even if we don’t care about history per se.

When a mathematician finishes writing a proof, they often revise it substantially, breaking out definition and lemmas, and simplifying and streamlining the argument. In software engineering, the corresponding practice is known as *refactoring*. In the nineteenth century, Richard Dedekind published four versions of his proof of unique factorization of ideals in the ring of integers in any algebraic number field; considering what changed over the course of the development and why is illuminating [3]. Rebecca Morris and I studied the history of proofs of Dirichlet’s celebrated theorem on primes in arithmetic progressions from Dirichlet’s first publication in 1837 to a version by Landau in 1927, and we similarly found that the history taught us a lot about the nature of

algebraic and set-theoretic abstraction [10].

Studying the history of mathematics can also raise new questions, such as the following. The value of algebraic reasoning is often attributed to its generality; for example, the axiomatization of *groups* in the nineteenth century unified instances in Galois theory, number theory, and geometry. However, although Dedekind and Weber generalized the notion of an ideal from rings of algebraic integers to rings of functions in 1882, Dedekind clearly felt that the abstraction provided a better understanding of the theory of algebraic integers even before the generalization. That raises the question as to what it is about abstraction that contributes to understanding of a particular body of results and, simultaneously, enables generalization. Historical examples often provide insights that make it possible to answer such questions [3, 6, 10].

2.5 Computers and understanding

Computer-assisted proofs have become increasingly common in contemporary mathematics, and there are several notable examples of results that have been established only with the help of computation. Kenneth Appel and Wolfgang Haken used extensive computation to prove the four-color theorem in 1976 [1], and Thomas Hales announced a proof of the Kepler conjecture in 1998, again using extensive computation [17]. Propositional satisfiability solvers are being used to solve combinatorial problems, in some cases, producing proofs that are terabytes long, or even longer [24]. What are we to make of them? Thurston raises some of the issues nicely:

The question is not even “How do mathematicians make progress in mathematics?” Rather, I prefer, “How do mathematicians advance human understanding of mathematics?” This question brings to the fore something that is fundamental and pervasive: that what we are doing is finding ways for people to understand and think about mathematics.

The rapid advance of computers has helped dramatize this point, because computers and people are very different. For instance, when Appel and Haken completed a proof of the 4-color map theorem using a massive automatic computation, it evoked much controversy. I interpret the controversy as having little to do with doubt people had as to the veracity of the theorem or the correctness of the proof. Rather, it reflected a continuing desire for *human understanding* of a proof, in addition to knowledge that the theorem is true.

On a more everyday level, it is common for people first starting to grapple with computers to make large-scale computations of things they might have done on a smaller scale by hand. They might print out a table of the first 10,000 primes, only to find that their printout isn’t something they really wanted after all. They discover by this kind of experience that what they really want is usually not some collection of “answers”—what they want is understanding. [43]

We should ask about the extent to which computational methods contribute to mathematical understanding and the extent to which they fall short. The answers will bear on the way we use computers in mathematics, and the way we assess the results of computer-assisted proofs.

2.6 Education

As a society, we devote considerable resources to teaching mathematics to our children. From early childhood through high school, students are subjected to countless hours of mathematics instruction, and have been evaluated and assessed on countless homework assignments, quizzes, and exams. If one goal of our educational practices is to convey mathematical understanding, we owe it to our students to think about what that means.

More than a quarter of a century ago, Liping Ma wrote a dissertation title that was subsequently published, *Knowing and Teaching Elementary Mathematics: Teachers' Understanding of Fundamental Mathematics in China and the United States* [28]. The work was based on interviews in which Chinese and American teachers were asked to solve mathematical problems and explain how they would teach their students to do the same. The book argued that Chinese teachers had a deeper and more flexible understanding of elementary mathematics and tried to determine the root causes. The study has been widely discussed and debated, and has influenced educational policy in the United States. The word “understanding” occurs hundreds of times in the text, often seven or eight times on a single page. The work distinguishes between different kinds of understanding, such as conceptual and procedural understanding; it assesses depth of understanding demonstrated in teacher interviews; and it recommends instructional procedures that are designed to convey understanding.

Notably, the book does not attempt to give a general account of what it means to understand mathematics. The meaning of the phrase has to be inferred from what Ma has to say about the case studies. This does not detract from the value of the work, but it makes it hard to assess the book's arguments rigorously and critically, and one might hope that a general theory of mathematical understanding could sharpen the arguments.

Contemporary philosophy of mathematics has only begun to address questions of mathematical understanding in education. Some educational researchers have begun to reflect on the nature of understanding (e.g. [14, 23, 45]); philosophers of mathematics have begun publishing in the educational literature (e.g. [11, 19, 27]); and philosophers and educators have begun to collaborate (e.g. [41]). Given that our educational practices aim to instill and assess mathematical understanding in students, we need to ask, “what does it mean to understand mathematics?”

3 Toward a theory of mathematical understanding

The questions raised in Section 2 have several things in common. First, they all have a generally epistemological character; they have to do with our ability to discover, verify, and communicate mathematical knowledge. Second, they have a generally normative character: they are questions about how we *should* go about our mathematical activities. Third, they don't seem all that mysterious; if you ask a mathematician why a certain theorem or historical development is important, or ask a teacher how to convey understanding to students, they will have plenty to say. Finally, we care about the answers; we devote considerable resources to mathematical research and education, and the questions bear on the way we carry out and evaluate our work. These facts don't guarantee the existence of a philosophical theory of mathematical understanding that can help us answer the questions, but it should at least encourage us to look.

3.1 The objects of understanding

The objects of mathematical understanding are varied. To start with, we can understand a definition, a theorem, or a proof. These sorts of syntactic objects are well represented by the conventional methods of mathematical logic. The same holds for problems, conjectures, examples, and theories; for example, we can think of a theory as a collection of definitions, theorems, and proofs. It is not as straightforward to come up with a general definition of a diagram [16, 44], however, toward explaining what it means to understand one. We can often characterize the information conveyed by particular diagrams, such as diagrams in Euclidean geometry, knot diagrams, or diagrams in homological algebra, but it is less straightforward to make sense of the different ways that diagrams may be presented to us, and how that bears on understanding.

We also speak of understanding more abstract entities such as methods, concepts, ideas, strategies, motivations, and goals. The philosophy of mathematics doesn't yet give us the means to talk about them with any rigor and precision. Sometimes what we understand are processes rather than objects; we can understand how to prove a theorem or how to solve a problem. And just as we can understand all these sorts of things, they can, in turn, convey understanding. A diagram can help us understand a proof, and we can come to understand a definition by seeing the role it plays in proving a theorem.

3.2 Understanding as ability

The fact that we value mathematical understanding in so many different contexts suggests that we should not expect a simple, homogeneous account of what it means to understand mathematics. A common thread running through the examples above, however, is that understanding is closely tied to our mathematical *capacities*, that is, what we can do with our mathematical knowledge. Understanding a proof involves not only the ability to justify each inference but also the ability to discover the proof, and others like it. Understanding Euclidean geometry involves the ability to read and write Euclidean proofs, which, in turn, relies on the ability to recognize legitimate inferences and reject illegitimate ones. Understanding Riemann's definition of the zeta function contributed to mathematicians' ability to prove the prime number theorem, and understanding elementary mathematics involves the ability to solve elementary mathematical problems across a range of variations.

These capacities are, moreover, fundamentally capacities for *reasoning*. The conventional, twentieth-century view sees mathematical knowledge as a static body of definitions, theorems, and proofs, whereas reasoning is a dynamic process of passing through a sequence of epistemic states in pursuit of a goal. We step through the inferences in a proof in order to verify the conclusion; we search for a proof in order to establish a new theorem; we manipulate mathematical representations of mathematical objects in order to solve problems; we look for patterns and formulate conjectures in the pursuit of new knowledge. Knowledge is both an outcome and a component of understanding, but at any given point in time, knowledge is static: it is the network of definitions, theorems, and proofs that we have at our disposal. Understanding is what enables us to use that knowledge effectively and to discover more.

Thinking first about what it means to reason will help us get a grip on understanding. Mathematics is hard; it calls on us to solve complex tasks, like solving problems and proving theorems, that require long sequences of steps that need to be carried out in precise ways. Each step we take changes our epistemic state: we introduce new objects, recall relevant definitions and facts, and draw conclusions. The action space is typically huge, and the combinatorial explosion of options

threatens to overwhelm. Understanding is what gets us through it, taming complexity and guiding us toward our goals.

If this general view of understanding is correct, the challenge is to figure out how to talk about understanding and reasoning in a precise and rigorous way. Initially, I thought about the relevant capacities as *methods*, that is, quasi-heuristic, fallible reasoning procedures, perhaps described as a kind of cognitive-mathematical pseudocode [2]. However, thinking in procedural terms seems to be too fine-grained. For example, to establish the correctness of a statement in number theory, it may be sufficient to recognize that the statement follows from a general result in algebra. In that case, we may say that the requisite understanding lies in the ability to recognize the number-theoretic statement as an instance of the algebraic one, independent of the process that leads to that recognition. As a result, I later argued in favor of thinking about understanding in terms of the *abilities* to carry out the requisite reasoning steps [4, 5]. That will not get us away from procedural thinking entirely, since we often view some mathematical abilities as constitutive of others. For example, the ability to add multi-digit numbers presupposes the ability to add one-digit numbers and the ability to maintain place-value relationships. The natural way to explain that dependence is to view abilities as specifications of methods, or families of methods, that invoke other methods as subroutines.

Although there are challenges to developing rigorous ways of talking about reasoning abilities, I am convinced it is the way to go. Other philosophers of mathematics have also argued in favor of thinking about mathematical reasoning in procedural terms [15, 19, 40, 41], something that is clearly related to Ryle’s distinction between knowing that and knowing how [34, 37]. Views of understanding in terms of abilities or intellectual know-how can also be found in the general epistemological literature [25] and in the epistemology of science [26]. Thinking about understanding in terms of abilities, however, is especially well-suited to mathematics and mathematical reasoning. In the next section I will argue that such an approach is methodologically sound and that there are plenty of opportunities to make progress.

3.3 Methodology

We want to make sense of the various ways we talk about understanding, in a manner that is clear, precise, and internally coherent, and that accords with our intuitions, fits the data, and can inform, and be informed by, related sciences. I have argued that we should think of mathematical reasoning as a process of passage through epistemic states en route to a mathematical goal, and that we should think of mathematical understanding as possession of the abilities needed to carry out such reasoning successfully. Some of these abilities are specific and easy to characterize in behavioral terms, such as the ability to draw an appropriate inference or carry out a calculation. Others function at a higher level, for example, the ability to see an analogy between two mathematical domains or the ability to implement a high-level proof strategy. Sections 4 and 5 will explore some of the vocabulary that has been used in cognitive science and AI to talk about such abilities, but we should also draw inspiration and guidance from all the domains of application enumerated in Section 2.

Such an approach accords well with our intuitions. If you tell me that a student in your class understands harmonic analysis and I ask you to explain what you mean, you will likely describe the kinds of problems the student can solve, the types of reasoning they can carry out, and the standard theorems that they can state and apply correctly. If you tell me that a particular historical result marked an advance in our understanding of an important problem and I ask you what you mean,

you will probably tell me what people have been able to do with the result. The good news that there is no shortage of such data: we can study assessments of understanding in several domains and characterize them in terms of the associated abilities. From that data, we should aim to develop a general account of our reasoning abilities and how they contribute to the observable results.

The main methodological claim of this section is that this is all a theory of mathematical understanding needs to do [4, 5]. In other words, a theory of mathematical understanding ought to be a theory of abilities that sets aside the question of how these abilities are implemented in human cognition, algorithms, neural networks, or historical and social processes. The latter are tasks for cognitive science, computer science, machine learning, and the history and sociology of mathematics. A theory of understanding that accounts for the historical and contemporary literature and that explains the practices and value judgments we see there can provide a philosophical foundation that clarifies discourse in all those domains.

We should proceed from the bottom up, answering specific instances of the kinds of questions raised in Section 2 with an eye toward seeing what the answers have in common. What are the diagrammatic inferences that are allowed in Euclid’s *Elements*? How did the introduction of the zeta function lead to the solution of the prime number theorem? What kinds of understanding can we get from a brute-force computation? In what ways are the practices of Chinese mathematics instructors better at conveying mathematical understanding than those of their American counterparts? We will only make progress by addressing such questions thoughtfully and then reflecting on what the answers tell us about mathematical understanding more generally.

4 Insights from symbolic AI

Following the successes of mathematical logic and the theory of computation in the early twentieth century, and the advent of the digital computer in the 1940s, a logical, symbolic, rule-based view of reasoning became the dominant paradigm in artificial intelligence. The associated views of language and thought shaped research in linguistics, cognitive science, philosophy of language, and philosophy of mind. This section explores mathematical reasoning and understanding from this perspective, while the next section explores complementary perspectives from machine learning and neural AI.

4.1 Concepts

4.2 Representations

4.3 Cognitive difficulty

4.4 Modularity and abstraction

5 Insights from machine learning

Although both logic-based and machine-learning-based approaches were represented at the 1956 Dartmouth workshop that gave the field of artificial intelligence its name, the latter have become more prominent over the last two decades, which brought striking successes in the use of deep learning, large language models, and reinforcement learning. Symbolic and neural approaches embody different perspectives on reasoning and understanding: whereas symbolic AI views reasoning

as a matter of carrying out logical inferences, neural AI see reasoning as a matter of drawing inferences from experience and data, and whereas symbolic AI makes use of hand-crafted, explicit rules and representations, neural AI draws on rules and representations that are implicit and distributed across the parameters of a neural network. Despite these differences, research on neural AI for mathematics draws on many of the same paradigms for mathematical reasoning that were developed in the context of symbolic AI. This section explores how these paradigms play out in the neural setting.

5.1 Machine learning in mathematics

5.2 Neural perspectives on reasoning

6 Conclusions

I hope that I have convinced you that there is important work to be done in systematizing and clarifying the ways we think about mathematical understanding, which bear on the way we teach mathematics to our children and the way we use mathematics in our daily lives. Coming to terms with our mathematical values requires us to reflect on the nature and purpose of mathematics itself, which becomes even more pressing now as we witness the rise of mechanized systems that can perform mathematical reasoning tasks that we once thought were uniquely human. Technological developments compel us to ask what role *we* should play in the mathematical process, why *our* mathematical understanding matters, and how we ought to engage with the emerging technologies. Throughout history—from ancient Greece to the scientific revolution, to the foundational crises of the early twentieth century—philosophy of mathematics has been there to guide us. Today, as new technologies promise both to extend and disrupt our mathematical practice, philosophy once again has a vital role to play in helping us navigate the changes ahead.

References

- [1] K. Appel and W. Haken. Every planar map is four colorable. *Bull. Amer. Math. Soc.*, 82(5):711–712, 1976.
- [2] Jeremy Avigad. Mathematical method and proof. *Synthese*, 153(1):105–159, 2006.
- [3] Jeremy Avigad. Methodology and metaphysics in the development of Dedekind’s theory of ideals. In *The architecture of modern mathematics*, pages 159–186. Oxford Univ. Press, Oxford, 2006.
- [4] Jeremy Avigad. Understanding proofs. In Paolo Mancosu, editor, *The Philosophy of Mathematical Practice*, pages 317–353. Oxford University Press, Oxford, 2008.
- [5] Jeremy Avigad. Understanding, formal verification, and the philosophy of mathematics. *Journal of the Indian Council of Philosophical Research*, 27:161–197, 2010.
- [6] Jeremy Avigad. Modularity in mathematics. *Rev. Symb. Log.*, 13(1):47–79, 2020.
- [7] Jeremy Avigad. Reliability of mathematical inference. *Synthese*, 198(8):7377–7399, 2021.

- [8] Jeremy Avigad. The design of mathematical language. In Bharath Sriraman, editor, *Handbook of the History and Philosophy of Mathematical Practice*, pages 3151–3189. Springer International Publishing, Cham, 2024.
- [9] Jeremy Avigad, Edward Dean, and John Mumma. A formal system for Euclid’s *Elements*. *Review of Symbolic Logic*, 2:700–768, 2009.
- [10] Jeremy Avigad and Rebecca Morris. Character and object. *Rev. Symb. Log.*, 9(3):480–510, 2016.
- [11] Jessica Carter. Mathematical understanding – common themes in philosophy and mathematics education. *The Journal of Mathematical Behavior*, 76:101202, 2024.
- [12] John W. Dawson, Jr. Why do mathematicians re-prove theorems? *Philos. Math. (3)*, 14(3):269–286, 2006.
- [13] John W. Dawson, Jr. *Why prove it again?* Springer, Cham, 2015. Alternative proofs in mathematical practice, With the assistance of Bruce S. Babcock and with a chapter by Steven H. Weintraub.
- [14] Elizabeth Fennema and Thomas A. Romberg, editors. *Mathematics Classrooms That Promote Understanding*. Routledge, 1999.
- [15] Janet Folina. Towards a better understanding of mathematical understanding. In *Truth, existence and explanation*, volume 334 of *Boston Stud. Philos. Hist. Sci.*, pages 121–146. Springer, Cham, 2018.
- [16] Marcus Giaquinto. *Visual thinking in mathematics: an epistemological study*. Oxford University Press, Oxford, 2007.
- [17] Thomas C. Hales. A proof of the Kepler conjecture. *Ann. of Math. (2)*, 162(3):1065–1185, 2005.
- [18] Yacin Hamami. Mathematical rigor and proof. *Rev. Symb. Log.*, 15(2):409–449, 2022.
- [19] Yacin Hamami and Rebecca Lea Morris. Philosophy of mathematical practice: a primer for mathematics educators. *ZDM Mathematics Education*, 52, 2020.
- [20] Yacin Hamami and Rebecca Lea Morris. Plans and planning in mathematical proofs. *Rev. Symb. Log.*, 14(4):1030–1065, 2021.
- [21] Yacin Hamami and Rebecca Lea Morris. Understanding in mathematics: The case of mathematical proofs. *Noûs*, 58(4):1073–1106, 2024.
- [22] Yacin Hamami, John Mumma, and Marie Amalric. Counterexample search in diagram-based geometric reasoning. *Cognitive Science*, 45(4), 2021.
- [23] Guershon Harel. Students’ understanding of proofs: a historical analysis and implications for the teaching of geometry and linear algebra. *Linear Algebra Appl.*, 302/303:601–613, 1999. Special issue dedicated to Hans Schneider (Madison, WI, 1998).

- [24] Marijn J. H. Heule and Oliver Kullmann. The science of brute force. *Commun. ACM*, 60(8):70–79, 2017.
- [25] Allison Hill. Understanding why. *Noûs*, 50(4), 2016.
- [26] Anna Elisabeth Höhl. *Scientific Understanding: What It Is and How It Is Achieved*. Philosophie - Aufklärung - Kritik, Vol. 2. transcript Verlag, Bielefeld, Germany, 2024.
- [27] Marc Lange. What are explanatory proofs in mathematics and how can they contribute to teaching and learning? *The Journal of Mathematical Behavior*, 76:101191, 2024.
- [28] Liping Ma. *Knowing and Teaching Elementary Mathematics: Teachers’ Understanding of Fundamental Mathematics in China and the United States*. Routledge, New York, 1999.
- [29] Paolo Mancosu, Francesca Poggioli, and Christopher Pincock. Mathematical Explanation. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2023 edition, 2023.
- [30] Kenneth Manders. The Euclidean diagram. In Paolo Mancosu, editor, *The philosophy of mathematical practice*, pages 80–133. Oxford University Press, Oxford, 2008. MS first circulated in 1995.
- [31] Rebecca Lea Morris. Motivated proofs: what they are, why they matter and how to write them. *Rev. Symb. Log.*, 13(1):23–46, 2020.
- [32] Plato. *Plato: Complete Works*. Hackett Publishing Company, Indianapolis, 1997.
- [33] Henri Poincaré. *Science et Méthode*. Flammarion, Paris, 1908. Translated by Francis Maitland as *Science and method*, Dover Publications, New York, 1952.
- [34] Gilbert Ryle. *The Concept of Mind*. University of Chicago Press, Chicago, New Univer edition, 2000 edition, 1949.
- [35] Stewart Shapiro, editor. *The Oxford Handbook of Philosophy of Mathematics and Logic*. Oxford University Press, Oxford, New York, 2005.
- [36] George F. Simmons. *Precalculus Mathematics in a Nutshell: Geometry, Algebra, Trigonometry*. Janson Publications, 1987.
- [37] Jason Stanley. Knowing (how). *Noûs*, 45(2):207–238, 2011.
- [38] Fenner Stanley Tanswell. *Mathematical Rigour and Informal Proof*. Elements in the Philosophy of Mathematics. Cambridge University Press, 2024.
- [39] Mark Steiner. Mathematical Explanation. *Philosophical Studies*, 34:133–151, 1978.
- [40] Fenner Stanley Tanswell and Matthew Inglis. The language of proofs: a philosophical corpus linguistics study of instructions and imperatives in mathematical texts. In *Handbook of the history and philosophy of mathematical practice*, pages 2925–2952. Springer, Cham, 2024.
- [41] Fenner Stanley Tanswell and Keith Weber. Instructions and recipes in mathematical proofs. *Educational Studies in Mathematics*, 111(1).

- [42] Jamie Tappenden. Proof style and understanding in mathematics i: Visualization, unification and axiom choice. In Paolo Mancosu, Klaus Froyen Jørgensen, and Stig Andur Pedersen, editors, *Visualization, Explanation and Reasoning Styles in Mathematics*, pages 147–214. Springer Netherlands, Dordrecht, 2005.
- [43] William P. Thurston. On proof and progress in mathematics. *Bulletin of the American Mathematical Society*, 30(2):161–177, 1994.
- [44] Silvia De Toffoli. *How to Prove Things with Diagrams*. Oxford University Press, to appear.
- [45] Keith Weber and Matthew Inglis. Mathematics education research on mathematical practice. In *Handbook of the history and philosophy of mathematical practice*, pages 2637–2663. Springer, Cham, 2024.