BULLETIN (New Series) OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 61, Number 3, July 2024, Pages 395–410
https://doi.org/10.1090/bull/1823
Article electronically published on May 15, 2024

PROOFS FOR A PRICE: TOMORROW'S ULTRA-RIGOROUS MATHEMATICAL CULTURE

SILVIA DE TOFFOLI

Abstract. Computational tools might tempt us to renounce complete certainty. By forgoing of rigorous proof, we could get (very) probable results for a fraction of the cost. But is it really true that proofs (as we know and love them) can lead us to certainty? Maybe not. Proofs do not wear their correctness on their sleeve, and we are not infallible in checking them. This suggests that we need help to check our results. When our fellow mathematicians will be too tired or too busy to scrutinize our putative proofs, computer proof assistants could help. But feeding a mathematical argument to a computer is hard. Still, we might be willing to undertake the endeavor in view of the extra perks that formalization may bring—chiefly among them, an enhanced mathematical understanding.

1. The commandment of rigor

In a provocative essay, "Theorems for a price: tomorrow's semi-rigorous mathematical culture," Doron Zeilberger [Zei94] envisages a future in which mathematicians will renounce "absolute certainty" and content themselves with "almost certainty." Mathematicians of (the day after) tomorrow will rebel against the commandment *thou shalt prove everything rigorously* that regulates today's mathematical practice.[1] They will work with new computational tools to establish results forgoing the deductive method, or so predicts Zeilberger. Why insist on absolute certainty in cases in which we can get very close to it for a fraction of the cost? In his words,

> we might witness many results for which we would know how to find a proof (or refutation), but we would be unable, or unwilling, to pay for finding such proofs, since "almost certainty" can be bought so much cheaper. I can envision an abstract of a paper, c. 2100, that reads: "We show, in a certain precise sense, that the Goldbach conjecture is true with probability larger than 0.99999, and that its complete truth could be determined with a budget of $10B." [Zei94, p. 14]

Received by the editors July 17, 2023.
2020 *Mathematics Subject Classification.* Primary 00A30.
This research was supported by MUR – Ministero dell'Universitá e della Ricerca through PRIN PNRR Missione 4 "Istruzione e Ricerca" – Componente C2 Investimento 1.1, "Fondo per il Programma Nazionale di Ricerca e Progetti di Rilevante Interesse Nazionale (PRIN)" Funded by the European Union Next GenerationEU (Project: "Understanding Scientific Disagreement and its Impact on Society," n. P2022A8F82) – CUP I53D23006880001.

[1]Since rigor can be cashed out in many alternative ways, this commandment is quite vague as stated. But it will do for now. For a precisification, see [BDT22].

Zeilberger gives a few examples of how new technologies could potentially change mathematical practice dramatically. And he was writing in the early nineties! Nowadays, computational methods are commonly used—but not merely in the way envisaged by Zeilberger. Machines can surely help us attain (very) probable results, but they can also help us solve mathematical problems in the good old-fashioned way: through rigorous proofs. As a matter of fact, machines are giving us access to a whole new class of computationally heavy proofs. For example, in the field of discrete geometry, it is common to discover *and prove* results by deploying heterogeneous computational methods such as "optimization techniques, combinatorial enumeration, validated numerical computation, linear programming methods, Monte Carlo simulation, search techniques, propositional satisfiability algorithms, and computer algebra" [Avi22, p. 108]. And artificial intelligence is bringing in a whole new range of possibilities.

It is a fact that machines help us reach new results, keeping or not the commandment of rigor. Mostly keeping it: technology is not (at least not *yet*) urging us to let go of complete certainty. That is why Zeilberger's claims still sound very controversial today. His assumptions do not sound controversial at all, however. "In the future, not all mathematicians will care about absolute certainty" [Zei94, p. 11], says Zeilberger. This claim suggests that now all (or at least the majority of) mathematicians care about absolute certainty. This seems right. *But can mathematicians really get what they care for?*

Although mathematicians can't always get what they want (who can?), it is plausible to think that at times they can. They don't always manage to produce proofs, but when they do, it is reasonable to think that they can be certain of their results. After all, according to the traditional story, it takes a proof to truly *justify* a mathematical claim. And genuine proofs are rigorous. Nothing new under the sun. Already in 1900, before listing his famous twenty-three problems, David Hilbert clarified what would count as a solution:

> It remains to discuss briefly what general requirements may be justly laid down for the solution... This requirement of logical deduction by means of a finite number of processes is simply the requirement of rigor in reasoning. Indeed the requirement of rigor... has become proverbial in mathematics. [Hil02, p. 409]

A mathematical claim is fully justified only through a (rigorous) proof. Since proofs are correct deductive arguments in which the premises entail the conclusion, they can provide absolute certainty. But there is a caveat. Proofs must start somewhere, and foundations can be shaky. This should not make us tremble, however. Even if we cannot be sure of the unqualified truth of the stated conclusions of our theorems, at least we can be certain of the conditional claims that the premises (which often are left implicit) imply the conclusions, or so it seems. That is, if we stay put and resist the temptation of tomorrow's semi-rigorous methods, we may get what we really want: certainty.

But the story is not so simple, and the price to pay might be even higher than the one envisioned by Zeilberger. To see this, let me first say something about the nature of proofs and something about the nature of our grasp of proofs.

First remark: *proof is a success term*. This means that there is no such thing as a fallacious proof: if we find a substantial error in a putative proof, then we establish that our putative proof was no proof at all. It wasn't a *real* proof, and

things that aren't real proofs, even if they *look like* proofs, are no proofs. It is true that in mathematical practice the term *proof* is used loosely, but on close scrutiny it makes sense to say that *proof* is a success term. Here is a qualification. Saying that fallacious putative proofs aren't proofs at all cannot mean that putative proofs containing minor mistakes are no proofs because those are ubiquitous in mathematics; banning them would lead to excluding too many of the things that mathematicians are happy to call proofs. That is why *essentially correct* putative proofs are proofs. A putative proof that is essentially correct may contain minor errors, but may neither contain errors that would invalidate its conclusion, nor errors that would be too hard to fix.

Here is a case that will help motivate this first remark. What Andrew Wiles had when he first announced to have proven Fermat's last theorem in 1993 was *not* a proof. He thought he had a proof, but he did not—this much seems uncontroversial. With his usual pungency, André Weil observed:

> to some extent, proving Fermat's theorem is like climbing Everest. If a man wants to climb Everest and falls short of it by 100 yards, he has not climbed Everest. (Reported in [Hor94, p. 33].)

But surely, we can fall short of 1 yard (in fact, it should not count as falling short at all)—that is why essentially correct putative proofs are proofs. About a year after his first announcement, collaborating with his former student Richard Taylor, Wiles finally managed to come up with a proof (or at least we have a lot of evidence to think that his final argument is a proof). Most likely, this proof contains minor errors and imperfections. Again, having a proof does *not* require complete formal precision.

To be sure, it is hard to spell out what *essential correctness* amounts to, but for now, a rough idea will do. For the sake of ease of expression, I will simply talk of proofs to mean *essentially correct* putative proofs.[2] As an aside: this problem is linked to that of articulating criteria of identity for proofs—a very hard problem that plausibly does not admit of a context-independent solution.[3]

Second remark: *proofs don't wear their correctness on their sleeve.* Checking the correctness of a putative proof can be tough, and mathematicians are not infallible in discerning genuine proofs from erroneous putative proofs. And this matters. A lot. This seemingly obvious fact puts pressure on the idea that we can actually get what we care for, certainty.

Although Weil's mountaineering analogy is evocative, proving a theorem *in practice* is *not* like climbing a mountain. It is instead like *simulating* climbing a mountain. Let me explain. When we climb a mountain, if the road is blocked or we meet a gorge, we cannot continue in the way we envisaged. We simply cannot. If we are good mountaineers, we might find alternative routes. Still, we cannot pretend that the obstructions we encounter are not there. Moreover, mistakes in climbing can lead to death, and there is an unflinching, all-knowing "adversary," the mountain, which determines if one truly reached the peak or not.

[2]This is quite commonsensical among mathematicians. Avigad puts it as follows: "an informal proof can have small mistakes and yet still reasonably lead us to believe in the correctness of its conclusion" [Avi21, p. 7385].

[3]This is an issue that also invests other topics in the philosophy of mathematics, such as the epistemology of diagrams and visualization [DT23].

When we simulate climbing a mountain, and the road is blocked, most likely we realize that it is blocked and that we cannot proceed. However, sometimes we misperceive, we do not see the obstruction, and we keep going as if it wasn't there. Wiles at first kept going right through a gorge, and only when his putative proof underwent scrutiny from other experts, did the gap became visible. And this is by no means an isolated case. It is precisely through the self-checking activity that the community of mathematicians operates on itself that the simulations are usually so good.

Nothing surprising so far. Nobody doubts that we make mistakes. This is just a fact of life. However, taking this fact seriously obliges us to think hard about what really justifies a mathematical claim. Do philosophers (or mathematicians) still want to hold on to the claim that only genuine proofs can provide us with mathematical justification? To answer this question, we should first get a sense of what *epistemic justification* is in general.

2. Mathematical justification

One way to think about epistemic justification is to link it to (theoretical) rationality. I am justified in believing a certain proposition if and only if it is rational for me to believe it.

By way of example, I am justified in believing that the optimal way to pack spheres is the face-centered cubic packing that is commonly used in fruit stands to display oranges, (i.e., that the Kepler conjecture is true) because Tom Hales proved it (and later formalized with a team of collaborators). More mundanely, I am justified in believing that in my fruit basket there are four oranges since I bought four this morning and ate none. Both beliefs are rational since they are based on good reasons. But I could go wrong, like in the following case. Unbeknown to me, my friend Agnese sneakily took an orange an hour ago, and so now there are only three oranges in my fruit basket. In this case, I have a false but justified belief.

This implies that justification in general is not *factive*. That is, it does not entail truth. On the other hand, knowledge is factive. That is why justification alone is not enough for knowledge. Still, justification is an essential ingredient of knowledge. Consider a second variation on the oranges' example. I am not justified in thinking that there are eight oranges in my fruit basket just because I am absent-minded and a bit of a wishful thinker. Suppose however I believe it, irrationally. Further suppose that my friend Agense, instead of taking an orange, put in four extra oranges as a kind gesture. In this case, I have a true belief, but it is not knowledge because it is not justified—it is not based on a good reason.

We saw that one way to think about epistemic justification is to link it to having good reasons. For instance, if I believe that p, a certain mathematical proposition, is true because I have a good argument for it, then I am justified in believing it. Another good source of justification is testimony or authority. If I read in the *Annals of Mathematics* that a certain theorem has been proven, then I acquire a good reason to believe that the theorem holds.

According to the received view in the philosophy of mathematics, good first-hand (that is, nontestimonial) reasons for mathematical propositions that are not axioms are genuine proofs. To be sure, we can be justified in believing that the Riemann hypothesis is true by virtue of nondeductive arguments. That is all fine and good, but the Riemann hypothesis does not deserve the status of theorem because it is

not established through a proof. This suggests that genuine first-hand *mathematical* justification, the one aspired by mathematicians, is only given by proofs.

But I can believe mathematical propositions for a host of different reasons. For example, if I believe p because of wishful thinking, p might be true, but I would not be justified in believing it. Wishing that p is true is *not* a good reason to believe that p is true. Moreover, the good reasons must be *epistemic*. That is, they must be related to the truth of the proposition in question. If someone points a gun at my temple and tells me either to believe p or to prepare to die, I might have very good *prudential* reasons to believe p (and should believe it, if I can somehow find a way to force myself to), but still no *epistemic* reason. Epistemic justification (differently than practical justification) is by its nature *truth-conducive*. That is, justified beliefs tend to be true.

Another influential way of thinking about (epistemic) justification of beliefs is considering the processes that formed such beliefs rather than the reasons one might have for them. If these processes are reliable, then the belief is justified, otherwise it is not; see [Gol79].

Justification is important for epistemologists because, as I mentioned above, it is thought to be the central component of knowledge. A belief that is not justified cannot constitute knowledge. If I form a true mathematical belief by flipping a coin, I might get it right, but I won't have knowledge since I did not have a good reason (and in fact my belief-forming process was utterly unreliable).

Thinking that knowledge requires justification goes all the way back to Plato. It is from the *Theaetetus* [Pla92] that we get the account of knowledge as Justified True Belief (JTB). According to the JTB story, a subject S knows proposition p if and only if S is justified in believing p, p is true, and S believes p.

Intuitively, knowledge requires belief: I cannot know something I don't believe. It requires truth as well: I cannot know something false. Finally, it requires justification: there must be some sort of connection between a true belief and its truth for it to constitute knowledge. While justification is truth-conducive, it is not factive: I can be justified in believing false propositions. Philosophers accepted the JTB account basically from Plato until 1964—more on this soon. Let us first consider some additional examples.

I am justified in believing that tomorrow will be sunny because my weather app, which is quite reliable, predicts so. However, it turns out that tomorrow will rain. This is a mundane example of a false but justified belief. Weather forecasts are pretty good, but far from being infallible. What about mathematics? Differently than in other domains, in mathematics (as well as in other areas usually considered to be within the domain of the a priori),[4] first-hand justification has often been considered by philosophers to be factive. This is because it has been associated with proofs. If a subject S is mathematically justified in believing p, then S has a proof for p, and this means that p is true—at least if we take p modulo the starting point and the logical principles.

This is remarkable and would set mathematical justification apart from other types of justification. As we just saw, I could be wrong in believing empirical

[4] *A priori* propositions are propositions that can be justified independently of experience. This is a general definition that can be understood in many different ways—but it will do for the present context.

propositions about oranges, the weather, or other natural phenomena. Yet in mathematics, if I am rational, it seems I cannot end up believing something false, at least for beliefs based on putative proofs. If I go wrong, it means I did not think hard enough, and thus, plausibly, I was not rational. This is why, in mathematics, justification and knowledge have traditionally been conflated. If there is no gap between truth and justification and the belief condition is in place, then justification entails knowledge.

This is connected with the a priori nature of mathematics. There is no external experience that can mess up my justification. Think about a different case, perception. Perceptual beliefs are very common and (assuming we are not in a skeptical scenario—that is, that we are not in *The Matrix* or similar weird places) are reliable. Still, I could misperceive. For instance, I could very well be justified in believing that there is a sheep in the field while in fact there is none. Perhaps what I am seeing is a dog, say a Cockapoo—from afar, it really looks like a sheep! I have a good reason supporting my false belief.

In mathematics, it looks like these cases cannot happen—accordingly, misperceiving in mathematics would not confer good reasons on us. If the only good first-hand reasons to believe a mathematical claim are proofs, then only true beliefs can be justified. But this is in stark contrast with mathematical practice! From the perspective of the working mathematician, it seems that we can also be justified in the absence of a proof. And this is not only in cases in which we openly use semi-rigorous methods. This also holds when we think we are dealing with rigorous proofs.

2.1. **Simil-proofs.** Here is a case in point. Vladimir Voevodsky was awarded the Fields medal in 2002. In two different cases, he found errors in results he had previously published. First, there were some issues in his work for which he was awarded the Fields medal—the theorems held, but a particular Lemma needed to be replaced by a more complicated one.[5] Second and more interesting, he found a significant error in a work in a different area. In this case, a result he thought he had established was outright false. He discovered the problem with his putative proof much later than when he had published it, to be precise, more than ten years later, in 2013.[6]

It is plausible to think that before he found the errors, Voevodsky was justified in believing his results. After all, on pain of skepticism, we must accept that mathematicians providing careful arguments for their results are justified, especially if their putative proofs have been checked and vouched for by their fellow practitioners—even more so if said mathematicians are awarded prestigious prices. But this implies that contrary to the received view, mathematical justification, like other types of epistemic justification, is not factive. That is, it is possible to hold a justified false belief in mathematics, one that is based on something that looks like a proof but it is not.

Mathematical justification should therefore not be connected with proofs, but with what in previous work I called *simil-proofs* [DT21]. These are arguments that look like proofs to the relevant subjects but may contain significant errors.

[5]He found an error in 1999–2000 (before the Fields Medal), which he corrected in a paper published in 2006.

[6]See [Voe14]. The story is further complicated by the presence of what some thought to be a counterexample to his results, but the details are beyond the scope of this paper.

Crucially, not any argument can be a simil-proof, however. Arguments containing blatant mistakes are excluded. Here is a working definition:

> SIMIL-PROOF: A mathematical argument that is taken to be a genuine proof by at least one (or a group of) appropriately trained subject(s), but that might not be. Moreover, it satisfies the standard of acceptability of the mathematical community to which it is addressed, and it has not been the object of serious criticism.

To be sure, this is a loose definition, but it will do for the present context. Linking mathematical justification to simil-proofs instead of genuine proofs goes along the lines of thinking of proofs as convincing mathematical arguments. Akshay Venkatesh characterizes proofs thus:

> (proofs) which are defined by the fact that they should induce uniform agreement about their validity, without any need for replication. [Ven24, p. 205]

And indeed, in mathematical practice, it is common to call "proof" also arguments that are not (genuine) proofs. For example, Leslie Lamport introduces his influential guidelines on how to write (simil-)proofs as follows:

> A method of writing proofs is described *that makes it harder to prove things that are not true*. [Lam11, p. i, emphasis added]

In my view, the ambiguity of the term *proof*, which sometimes is used to refer to proofs and sometimes to refer to simil-proofs (which may or may not be proofs), has created much confusion in the epistemology of mathematics. That is why it is helpful to disambiguate the term. It might be fine to use the word *proof* loosely in mathematics, but in order to pursue an epistemological inquiry into mathematics, disambiguation is called for.

To recap, at a given time, simil-proofs are phenomenologically indistinguishable from proofs, but they might contain substantial mistakes. This means that not all simil-proofs are proofs. However, if a simil-proof contains a substantial mistake, this must be a subtle mistake that the mathematical community is blind to—at least for some time. The idea is that by performing a self-monitoring activity on itself, the mathematical community gradually filters out all erroneous simil-proofs so that only correct simil-proofs (which, therefore, are proofs) remain.

We saw that not all simil-proofs are proofs. The reverse holds as well: not all proofs are simil-proofs. An example is given by an argument that is a proof, but, for some reason, no mathematician recognizes it as a proof. This could be because of the fact that the argument looks fallacious even if it is not, or, if we endorse an abstract definition of proof, that it has not been considered by any mathematician at all.

At the cost of being pedantic, from now on I will stick to the new term *simil-proof* to refer to accepted mathematical arguments put forward as proofs and generally referred to simply as *proofs* by mathematicians—bear with me.

The philosophical moral of the story is that, like in the case of perception in which I can justifiably believe that there is a sheep in the field while I am looking at something that looks like a sheep (a Cockapoo) but is not, so in mathematics, I can justifiably believe that a result holds in virtue of a something that looks like a proof (a simil-proof containing a major error), but it is not. That is, as Voevodsky's case suggests, beliefs based on fallacious simil-proof are justified.

2.2. **Gettier.** Let's get back to the general epistemological story. Philosophers went along with the JTB account from Plato all the way until the 1960s, when Edmund Gettier unhinged the tradition. Offering a couple of compelling counterexamples, he showed that the three conditions (justification, truth, and belief) are not jointly sufficient for knowledge.

Actually, a *Gettier case* (as they are called nowadays) was already concocted by Bertrand Russell [Rus09, p. 91]. It goes along the following lines. It is four o'clock in the morning, and you wake up and look at your analog clock. You thus form the true justified belief that it is four in the morning. However, your clock's batteries are dead—as a matter of fact, the clock stopped working yesterday precisely at four in the afternoon! So, you have a belief that is true and justified (you do not have any reason to think the clock is not working), but intuitively you don't know that it is four o'clock in the morning because your justification is severed from the truth of your belief. Some sort of epistemic luck—incompatible with knowledge—is involved.

Here is another case. As before, you are looking at a dog in a field, and you believe it is a sheep. You form the justified belief that there is a sheep in the field. Unbeknownst to you, there is indeed a sheep in the field, but it is outside your field of vision. Again, you have a justified true belief that does not constitute knowledge.[7]

In the wake of such counterexamples, epistemologists embarked on the quest of finding a fourth condition for knowledge (which they did *not* find—eventually abandoning the enterprise). What matters for us is that they generally did not question that the traditional analysis (the JTB story) would be perfectly fine for the case of mathematics. Here is Alvin Goldman, a very influential epistemologist, on the matter:

> My concern will be with knowledge of empirical propositions only, since I think that the traditional analysis is adequate for knowledge of nonempirical truths. [Gol67, p. 357]

But in fact, if there is a gap between justification and knowledge—and Voevodsky's case suggests that there is—then counterexamples to the JBT account are possible in mathematics as well. Here is one. It has to do with the 4-color theorem. Alfred B. Kempe was the first to publish a simil-proof for it in 1869. Kempe's was a careful argument by induction that was accepted by the mathematical community. However, it was not a genuine proof! After eleven years, a significant gap was unearthed. The story is well known; it took more than a century before Apple and Haken came up with a simil-proof that is still accepted today (which is often considered to be the very first computer-assisted (simil-)proof). So, it is plausible to think that Kempe had a justified true belief that did not amount to knowledge since, contrary to what he thought, his simil-proof was not a proof.

It is particularly compelling to think that Kempe was justified because, like Voevodsky, not only he had a careful argument of which he was convinced, but he also had additional evidence that his argument was correct given by the acceptance of it by his mathematical community. It is because we are fallible and know that we are, that this additional evidence is particularly important.

[7] A similar example was proposed in [Chi66].

2.3. **Shareability.** I hinted at the fact that part of the success of mathematics rests on the possibility that the mathematical community performs a self-monitoring activity. Without invoking mathematical arguments that can be shared among our fellow mathematicians, we cannot overcome our individual shortcomings. It is because it underwent scrutiny from other mathematicians that Wiles's original proof was found wanting. But as Kempe and Voevodsky's cases make it clear, errors are not always spotted so quickly—still, ideally, sooner or later they are indeed spotted.

It is because of this reason that if simil-proofs are to provide justification, they might contain errors, but they cannot be idiosyncratic arguments that nobody except a single subject could, in principle, grasp. They must have the potential to be understood by multiple mathematicians. That is, they must be *shareable*. And if they are arguments that satisfy the standards of acceptability of a legitimate mathematical community, they are indeed shareable. Note that a shareable argument might not be the kind of thing that a single mathematician can grasp; excluding large proofs like the ones involving automated computations or large collaborations would be too restrictive.

However, if proofs are defined simply as valid deductive arguments (from some accepted starting points), there might be proofs that are not shareable at all. This is counterintuitive. A valid deductive argument could have as many inferential steps as there are atoms in the universe. Suppose an extraterrestrial creature sufficiently similar to us (but having a far greater ability to process and keep track of inferential steps) can quickly grasp such an argument. At least in some cases, such a creature will not be able to share its grasp of the argument with us. Consequently, we would not be able to form justified beliefs based on such an argument. It is for this reason that, in my view, not all valid deductive arguments are proofs. This reveals that, like simil-proofs, proofs must be shareable as well.[8]

Requiring proofs and simil-proofs to be shareable restricts their domain to humanly accessible arguments. Shareability is naturally a graded notion that presents several distinct dimensions of evaluation. It measures not only how good of a reason a simil-proof is for our belief but also what it would take to share such a reason with other appropriately trained subjects. It is natural that different simil-proofs present different degrees of shareability.

One worry brought about by the development of ever more sophisticated computational tools is that in the future, there will be more and more simil-proofs that are hard to share, that is, that are hard to grasp from human practitioners. In these cases, we might have to resort to extra help for checking our results.

3. A price we might want to pay

When Voevodsky found out that his results were erroneous, he started to worry. Those were important and widely circulated results. He thought that mathematics (or some of it anyway) was extending into an arduous territory that was just not suited to the human mind. Michael Harris explains:

> Voevodsky obtained his prestigious position at the IAS and his Fields Medal for his work in a field in which "too-long [simil-]proofs" are common and in which the relatively small number of competent

[8]For more details, see [DT21].

> potential referees typically spend much of their time writing "too-long [simil-]proofs" of their own, so he might understandably be concerned that [simil-]proofs are not being read as carefully as they should. [Har15, p. 58]

Too-long simil-proofs are more and more customary in mathematics. Until now, we have mostly relied on each other to check the correctness of our simil-proofs. But in the future, we might need to be obliged to ask for help elsewhere. Computer proof assistants[9] are tools that allow us to create formal counterparts of our informal simil-proof and formally verify them if they are indeed correct. These tools could therefore help us implement a more thorough check on our simil-proofs.[10] And Voevodsky thought that in the future, these tools would be used as a matter of routine:

> Voevodsky predicted it would soon be possible to design proof checkers based on univalent foundations that could effectively verify correctness of mathematical [simil-]proofs written in the appropriate machine-readable language. In a few years, he added, journals will only accept articles accompanied by their machine-verifiable equivalents. [Har15, p. 60]

Although the few years might end up being substantially more, computer proof assistants are indeed gaining terrain in mathematics, especially in some parts of it. Obvious candidates for formal verification are computer-assisted simil-proofs. This is because they are sometimes accepted only with reservation by the mathematical community (precisely because they tend to score low on shareability). A famous case is the one of the 4-color theorem mentioned above. A formally verified simil-proof was produced in 2005 by Georges Gonthier.

Another case involves Tom Hales's 2005 simil-proof of the Kepler conjecture (remember the oranges?). After years of work for the twelve referees, the simil-proof was finally published in the *Annals of Mathematics*. However, the simil-proof involved computer-assisted calculations, and the referees admitted they were only 99% sure of its correctness. The ambiguous status of his simil-proof led Hales and collaborators to embark on the *Flyspeck project*, the massive enterprise of formally verifying the simil-proof, which was achieved in 2014.[11]

Even traditional simil-proofs from core areas of mathematics might call for formal verification. A recent case involves Peter Scholze and Dustin Clausen's work on *condensed mathematics* [Sch19]—their goal is to propose a new framework in which

[9] In the current context, a more appropriate, albeit more pedantic, name for these tools would be *computer simil-proof assistants*. However, for the sake of simplicity, I will keep adopting the usual terminology.

[10] One thorny issue I will gloss over is what exactly computer proof assistants are supposed to do. Are they used to check the correctness of our simil-proofs (and thus to find out whether they are proofs or not) or are they used to check whether the conclusions of our simil-proofs are indeed implied by their premises? Surely, these are related questions. They are not, however, the same question. This issue has to do with the (hard) problem already mentioned of simil-proofs individuation. Without dwelling on this problem, I endorse the seemingly uncontroversial assumption that at least in some cases, computer proof assistants are indeed used to check the correctness of our simil-proofs (one such case is Scholze's—see below).

[11] [HAB+17]. See [Avi18] for a survey of recent developments in the domain of formal mathematics.

topological spaces are replaced by *condensed sets*, which would make it possible to apply techniques from homological algebra to algebraic geometry. The material is online and has been widely circulated. But Scholze was unsure of some of the key details, and apparently his fellow mathematicians did not help him put his worries to rest. Kevin Buzzard recalls his personal interaction with Scholze:

> At the end of 2020, Scholze approached me and asked if we had had a study group on the work at Imperial; I answered that we had. Scholze then asked whether we had looked through all the details of the [simil-]proof of Theorem 9.1 of [Schb]; I answered that we had not. Scholze then remarked that he had had the same response from other mathematicians, and raised the possibility that perhaps nobody other than himself and Clausen had ever read the [simil-]proof carefully. Furthermore he suggested that perhaps this might remain true even after the refereeing process. The reason he was concerned about this was that, for Scholze, this was the theorem that the entire theory stood upon. [Buz21, p. 12]

At the time, the simil-proof of Theorem 9.1 was a result that, although in principle shareable, nobody except its authors took the time to check in detail. It is for this reason, and for its foundational role within the condensed sets framework, that Scholze decided to address the community of mathematicians that had been playing with computer proof assistants. Buzzard and other members of the Lean community took on the challenge and embarked on what they called the *Liquid Tensor Experiment*.[12] It is extraordinary that the team members (among which Johan Commelin and Patrick Massot played major roles) managed to formalize the bulk of the simil-proof in just six months! This shows that, contrary to prior expectations, computer proof assistants can indeed tackle cutting-edge mathematics in a reasonable amount of time.

In this case, Scholze could not find human checkers, so to speak. But even when other mathematicians are willing to verify our results, they are fallible and are likely to share the same cognitive shortcomings. Computer proof assistant might offer us a whole new level of confidence because, with them, we can subject our simil-proofs to a more rigid scrutiny, one that is more reliable in detecting errors. Or, at least, this is what some mathematicians involved in the computer proof assistant community think. One of them is Massot. He puts it as follows:

> The most obvious benefit of formalized mathematics is super-human certainty that a [simil-]proof is correct when it has been checked by a computer. [Mas21, p. 1]

It is plausible to think, along with Massot, that computer proof assistants can greatly amplify our confidence in the correctness of our simil-proofs. However, we might still fall short of getting a 100% guarantee. Why is complete certainty still eluding us? Because two types of problems lurk in the background. The first is the potential presence of bugs in the kernel or compiler of the proof assistant—this is a remote (but still existent) possibility due to the limited size and simplicity of the software. The second type of problem is trickier. It has to do with the faithfulness of the translation of informal results into formal statements. How can we be sure

[12]Lean is a computer proof assistant that is gaining momentum in the mathematical community.

that what we formalize is indeed what we started with in the first place? This suggests that we can verify our results using computer proof assistants to reach "super-human certainty" of their truth, but that super-human is still less than absolute.

Still, verification, although a major benefit of ultra-rigorous mathematics, is not the only one. Before winding up, I will briefly discuss how shareability could be enhanced through formalizations. This might sound surprising since, as I hinted at before, sharability is often lowered by the use of computational tools. In particular, computer-assisted simil-proofs tend to be rather opaque and thus hard to grasp, thus scoring low on shareability.

3.1. **Shareability on steroids.** Another promise of computer proof assistants is that they will make available a new type of mathematical writing. As we discussed, we want our simil-proofs to be shareable because we want our fellow mathematicians to be able to check them. But shareability is not only in place to check which simil-proofs actually amount to proofs. It is also crucial for providing *mathematical understanding*.

To be sure, it is notoriously hard to articulate a sharp characterization of what understanding is.[13] In the case of mathematics, this is partly due to the fact that understanding is a multi-dimensional phenomenon involving diverse cognitive abilities spanning from symbolic to visual. But this is also due to the fluid, ever-changing nature of mathematical understanding. For example, in a recent contribution, Jeremy Avigad [Avi22] points out how computers in mathematics opened up the possibility of new forms of understanding.

At any rate, the importance of understanding in mathematics is unquestionable. According to a popular article by Bill Thurston, in their activity of finding new results, mathematicians "discover... that what they really want is usually not some collection of "answers"—what they want is understanding" [Thu94, p. 162], and goes as far as to suggests that understanding is the ultimate goal of the mathematical activity.

In usual mathematical practice, simil-proofs are presented at specific levels of detail to facilitate a particular audience's understanding. However, it is often hard to choose just the right level, and the possibility of expanding and hiding the details on the fly would be a great feature, one that could help hit the target of understanding. This is another respect in which computer proof assistants could help:

> I think the most the promising application of formalized mathematics is the dream of producing mathematical documents allowing readers to dynamically choose the detail level and access background knowledge on demand. [Mas21, p. 3]

This type of technology could have a great impact on the way we understand simil-proofs. Moreover, Buzzard explains that computer proof assistants can be coupled with tools that produce dynamic web pages, enhancing even more the dynamicity of mathematical texts:

[13]For the present context, it is important to note that understanding does not have to be factive—that is, that we can understand things are are not true. A non-factive account of understanding has been championed by Catherine Elgin [Elg07].

> Tools like Alectryon will enable us to make documents which will allow links to dynamic web pages displaying anything from mathematical details to interactive pictures, in a human-readable form, and which will allow one to keep digging right down to the axioms, although of course it is unlikely that anyone would like to go down this far. [Buz21, p. 21]

Such dynamicity could have a terrific pedagogical impact as well:

> One could also imagine error-free undergraduate textbooks also written in this way, where statements which a student cannot understand (perhaps because they are ambiguous) can be inspected in more details until difficulties are resolved. [Buz21, p. 21]

All these considerations show that computer proof assistants like Lean could help us come up with better, more shareable simil-proofs that can be accessed in a personalized way.

This is a reward that goes beyond the particular type of mathematical understanding that is gained in the process of formalizing. As a matter of fact, the very process of feeding a simil-proof to a computer requires thinking hard about the structure and the details of the original simil-proofs, thus helping us gain additional understanding. This became clear in the Liquid Tensor Experiment. Scholze observed that during the process of formalizing his simil-proof, Commelin, one of the leaders of the project, came up with a better (more explicit and more elementary) version of the original argument. More generally, the formalization process gave him a clearer picture of

> [w]hat actually makes the [simil-]proof work! When I wrote the blog post half a year ago, I did not understand why the argument worked, and why we had to move from the reals to a certain ring of arithmetic Laurent series. But during the formalization, a significant amount of convex geometry had to be formalized (in order to prove a well-known lemma known as Gordan's lemma), and this made me realize that actually the key thing happening is a reduction from a non-convex problem over the reals to a convex problem over the integers.[14]

Computer proof assistants could also help us better access already available results by making new searchable databases possible. Moreover, with the development of AI, they will become better and better at suggesting proof strategies, thus aiding us in finding new simil-proofs. This means that they will likely also be tools for the discovery of new mathematics, and not just tools for the verification of old mathematics.

These are great promises. But it is not all puppies and rainbows. Computer proof assistants are only in their infancy, and they are still too clunky to be used by most mathematicians. The learning curve is very steep. Moreover, it is true that computer proof assistants might help increase the shareability of our simil-proofs, but only along specific dimensions. They might also lower it along other dimensions. For example, current proof assistants are not very good at handling diagrams and other types of visual information, which is another way practicing mathematicians

[14]From Buzzard's *Xena* Blog of June 5, 2021: `https://xenaproject.wordpress.com/2021/06/05/half-a-year-of-the-liquid-tensor-experiment-amazing-developments/`

use to amplify their understanding and thus increase the shareability of their simil-proofs. Another problem is that there are many competing systems (e.g., Coq, HOL Light, Isabelle, and Lean), which are often incompatible, and importing results from one to the other is no easy task—it is like having a bunch of different, incompatible operating systems and having to develop separate apps for all of them.

Still, this should not dishearten us. Nobody said that transitioning to a new ultra-rigorous formalized mathematics would be easy or that it would be quick. We might just have to wait some time until these tools become more flexible and more user-friendly.

4. Conclusion

Philosophers have traditionally thought that first-hand mathematical justification of propositions (that are not axioms) is provided by proofs exclusively—but this is a mistake. Looking at the practice of mathematics, we soon realize that there are compelling cases in which a subject can be justified in believing a false mathematical proposition (as in Voevodsky's case) or a true mathematical proposition in virtue of a fallacious argument (as in Kempe's case). Being associated with simil-proofs rather than proofs, mathematical justification is fallible.

But our individual fallibility can be partially overcome by working together and checking each other's results. Problems arise, however, when our simil-proofs are too long or too technical: our fellow mathematicians might be too busy or just unable to scrutinize them.

Weil compared proving a theorem to climbing a mountain. With the terminology introduced, we can now say that in practice we do not always produce proofs; what we produce are simil-proofs (that, to be sure, often are indeed proofs). That is why rather than climbing a mountain, we *simulate* climbing a mountain. According to some, however, computer proof assistant might get us the real thing:

> Having the ability to check partial progress with absolute certainty can be extremely useful to increase confidence and determination. [Mas21, p. 6]

That is, with computer proof assistants, we cannot proceed unless the terrain is really cleared. So, they might get us what we really wanted in the first place, certainty. Or at least, given that there is still the possibility of bugs or mismatches between informal and formal statements, to something closer to it compared to what we generally get relying exclusively on our human abilities.

After all, computers may indeed change the way in which we do mathematics. But it looks like they might lead us in the antipodal direction compared to the one indicated by Zeilberger. A counterpart of his prediction would go like this:

> We might witness many results for which we would have found a simil-proof, but we would be unable, or unwilling, to pay for formalizing it, since "almost certainty" will suffice. I can envision an abstract of a paper, c. 2050, that reads: "We show, in a certain precise sense, that the Goldbach conjecture is true (we checked our simil-proof thoroughly), and that formally verifying our simil-proof could be done with a budget of $1M."

Notice that the time is closer, and the price is lower than in the original quotation. Indeed, the main hurdle still consists in finding a simil-proof in the first place.

Besides, the more computer assistants are going to be used, the easier it will become to formalize new mathematics—they will become more flexible, and the libraries will grow. It might become, at some point, a price we will be willing to pay. This is more so because the dream is to use these tools not only as verification devices but also as amplifiers of human understanding—after all, mathematical understanding is what we might really need.

References

[Avi18] Jeremy Avigad, *The mechanization of mathematics*, Notices Amer. Math. Soc. **65** (2018), no. 6, 681–690. MR3792862

[Avi21] Jeremy Avigad, *Reliability of mathematical inference*, Synthese **198** (2021), no. 8, 7377–7399, DOI 10.1007/s11229-019-02524-y. MR4292724

[Avi22] Jeremy Avigad, *Varieties of mathematical understanding*, Bull. Amer. Math. Soc. (N.S.) **59** (2022), no. 1, 99–117, DOI 10.1090/bull/1726. MR4340829

[BDT22] John P. Burgess and Silvia De Toffoli, *What is mathematical rigor?*, Aphex **25** (2022), 1–17.

[Buz21] Kevin Buzzard, *What is the point of computers? A question for pure mathematicians*, Preprint, `arXiv:2112.11598`, (2021).

[Chi66] Roderick M. Chisholm, *Theory of knowledge*, Prentice-Hall, 1966.

[DT21] Silvia De Toffoli, *Groundwork for a fallibilist account of mathematics*, The Philosophical Quarterly **71** (2021), no. 4, 823–844.

[DT23] Silvia De Toffoli, *Who's afraid of mathematical diagrams?*, Philosophers' Imprint **23** (2023), no. 1, 1–20.

[Elg07] Catherine Elgin, *Understanding and the facts*, Philosophical Studies **132** (2007), no. 1, 33–42.

[Gol67] Alvin Goldman, *A causal theory of knowing*, Journal of Philosophy **64** (1967), no. 12, 357–372.

[Gol79] Alvin Goldman, *What is justified belief?*, George S. Pappas (ed.), pp. 1–23, D. Reidel Publishing Company, 1979.

[HAB+17] Thomas Hales, Mark Adams, Gertrud Bauer, Tat Dat Dang, John Harrison, Le Truong Hoang, Cezary Kaliszyk, Victor Magron, Sean McLaughlin, Tat Thang Nguyen, Quang Truong Nguyen, Tobias Nipkow, Steven Obua, Joseph Pleso, Jason Rute, Alexey Solovyev, Thi Hoai An Ta, Nam Trung Tran, Thi Diep Trieu, Josef Urban, Ky Vu, and Roland Zumkeller, *A formal proof of the Kepler conjecture*, Forum Math. Pi **5** (2017), e2, 29, DOI 10.1017/fmp.2017.1. MR3659768

[Har15] Michael Harris, *Mathematics without apologies*, Princeton University Press, Princeton, NJ, 2015. MR3289987

[Hil02] David Hilbert, *Mathematical problems*, Bull. Amer. Math. Soc. (N.S.) **37** (2000), no. 4, 407–436, DOI 10.1090/S0273-0979-00-00881-8. Reprinted from Bull. Amer. Math. Soc. **8** (1902), 437–479. MR1779412

[Hor94] John Horgan, *The last universal mathematician*, Scientific American (1994), 33–34.

[Lam11] Leslie Lamport, *How to write a 21st century proof*, J. Fixed Point Theory Appl. **11** (2012), no. 1, 43–63, DOI 10.1007/s11784-012-0071-6. MR2955006

[Mas21] Patrick Massot, *Why formalize mathematics?*, `https://www.imo.universite-paris-saclay.fr/~patrick.massot/files/exposition/why_formalize.pdf`, 2021.

[Pla92] Plato, *Theaetetus*, Hackett Publishing Company, 1992.

[Rus09] Bertrand Russell, *Human Knowledge: Its Scope and Limits*, New York: Taylor & Francis Routledge, 1st ed. 1948, (2009).

[Sch19] Peter Scholze, *Lectures on analytic geometry*, `https://www.math.uni-bonn.de/people/scholze/Analytic.pdf`, 2019.

[Thu94] William P. Thurston, *On proof and progress in mathematics*, Bull. Amer. Math. Soc. (N.S.) **30** (1994), no. 2, 161–177, DOI 10.1090/S0273-0979-1994-00502-6. MR1249357

[Ven24] Akshay Venkatesh, *Some thoughts on automation and mathematical research*, Bull. Amer. Math. Soc. (N.S.) **61** (2024), no. 2, 203–210, DOI 10.1090/bull/1834. MR4726987

[Voe14] Vladimir Voevodsky, *The origins and motivations of univalent foundations*, The Institute Letter (The Institute for Advanced Studies), (2014), `https://www.ias.edu/ideas/2014/voevodsky-origins`.

[Zei94] Doron Zeilberger, *Theorems for a price: tomorrow's semi-rigorous mathematical culture*, Math. Intelligencer **16** (1994), no. 4, 11–14, 76, DOI 10.1007/BF03024696. Reprinted from Notices Amer. Math. Soc. **40** (1993), no. 8, 978–981 [MR1239765 (94i:00007)]. MR1294994

Department of Humanities and Life Sciences, University School for Advanced Studies IUSS Pavia, 27100, Pavia, Italy

Email address: `silvia.detoffoli@iusspavia.it`