On the unreasonable reliability of mathematical inference

Brendan Philip Larvor[1]

Abstract

In (Avigad, 2020), Jeremy Avigad makes a novel and insightful argument, which he presents as part of a defence of the 'Standard View' about the relationship between informal mathematical proofs (that is, the proofs that mathematicians write for each other and publish in mathematics journals, which may in spite of their 'informal' label be rather more formal than other kinds of scientific communication) and their corresponding formal derivations ('formal' in the sense of computer science and mathematical logic). His argument considers the various strategies by means of which mathematicians can write informal proofs that meet mathematical standards of rigour, in spite of the prodigious length, complexity and conceptual difficulty that some proofs exhibit. He takes it that showing that and how such strategies work is a necessary part of any defence of the Standard View.

In this paper, I argue for two claims. The first is that Avigad's list of strategies is no threat to critics of the Standard View. On the contrary, this observational core of heuristic advice in Avigad's paper is agnostic between rival accounts of mathematical correctness. The second is that that Avigad's project of accounting for the relation between formal and informal proofs requires an answer to a prior question: what sort of thing is an informal proof? His paper havers between two answers. One is that informal proofs are ultimately syntactic items that differ from formal derivations only in completeness and use of abbreviations. The other is that informal proofs are not purely syntactic items, and therefore the translation of an informal proof into a derivation is not a routine procedure but rather a creative act. Since the 'syntactic' reading of informal proofs reduces the Standard View to triviality, makes a mystery of the valuable observational core of his paper, and underestimates the value of the achievements of mathematical logic, he should choose some version of the second option.

Keywords Proof · Informal · Derivation · Heuristic · Translation · Avigad · Hamami

---

Extended author information available on the last page of the article

1 Two views of Proof: history of the question

Avigad's paper adds to a debate that can be traced to the origins of modern mathematics in the early twentieth century. The historic division is between those like Hilbert[1] and the Bourbakists[2] who wished to drive the modern tendency towards ever more formal and abstract mathematics to its limit, and those like Poincaré[3] and Brouwer[4] who thought (in different ways) that mathematics ineliminably refers to human experience and activity, and therefore judged that complete formalisation is either unwise or impossible. This broad contrast developed into a debate over (what we now call) the Standard View: that full formalisability is not merely a splendid and valuable feature of mathematical proofs, but is a criterion or norm of rigour. For example, Saunders Mac Lane claimed that, "In practice, a proof is a sketch, in sufficient detail to make possible a routine translation of this sketch into a formal proof" and that, "…the test for the correctness of a proposed proof is by formal criteria and not by reference to the subject matter at issue" (1986 pp. 377-8). Here is a looser statement from Thomas Hales,

> The ultimate standard of proof is a formal proof, which is nothing other than an unbroken chain of logical inferences from an explicit set of axioms. While this may be the mathematical ideal of proof, actual mathematical practice generally deviates significantly from the ideal. (2012, p. x)

Where Mac Lane describes the informal proof as a 'sketch', Hales writes (in the title of his book) of a 'blueprint'. Solomon Feferman, at the end of his (2012), sought to specify the relation between an informal proof and a corresponding formal derivation without using such metaphors:

> …the Formalizability Thesis should be given a very strict reading, namely that (i) every good proof has an underlying logical structure, (ii) that structure is completely analyzed in the derivation that formalizes the proof, and, finally (iii) that derivation assures the correctness of the theorem proved on the basis of the background assumptions expressed by the axioms and rules of the system in which the proof is formalized.

---

[1] See Hilbert (1927), especially pp. 472-4 of van Heijenoort (ed.) for his polemics against Poincaré and Brouwer. Hilbert's philosophy of mathematics is much richer than here presented, and includes a conception of finitary mathematics as contentful, that is, not wholly syntactic. See Zach (2019).

[2] Bourbaki (1949).

[3] Poincaré (1908). On his conception of proof, "*Verification* differs from proof precisely because it is analytical,… It leads to nothing because the conclusion is nothing but the premises translated into another language. A real proof, on the other hand, is fruitful because the conclusion is in a sense more general than the premises." (1902, p. 396 in Benacerraf & Putnam). For a nuanced discussion of Poincaré's views on formal logic, see McLarty (1997).

[4] "The question where mathematical exactness does exist, is answered differently by the two sides; the intuitionist says: in the human intellect, the formalist says: on paper." Brouwer (1913) p. 83.

For a more recent expression of the idea that proof is really about syntax, here is Joel D. Hamkins,

> Proof and truth… lie on opposite sides of the syntax-semantics divide, for at bottom, a proof is a kind of argument, a collection of assertations structured syntactically in some way, while the truth of an assertion is grounded in deeply semantic issues concerning the way things are… Proof… lies solidly on the syntactic side, since ideally one can verify and analyze a proof as a purely syntactic object, without a concept of meaning and without ever interpreting the language in any model. (2020, pp. 157-8)

In fairness, it should be noted that this quotation comes from an introductory lecture series in which Hamkins outlines a range of perspectives and leaves the deeper questions that they raise for students to discuss. Nonetheless, these examples are enough to show why this is called 'the Standard View'. It occurs in mathematicians' discussions of proof, often with little or no supporting argument, as if it were an unobjectionable statement of what all mathematicians mean by 'proof'. It is almost always joined with a recognition that reality rarely approaches this ideal (as in Mac Lane and Hales, just quoted). Here is Hamkins, a few pages on: "Most contemporary mathematical proofs are written in prose, essay-style… In mathematical practice, a proof is any sufficiently detailed convincing mathematical argument that logically establishes the truth of a theorem from its premises." (p. 160). Students might wonder how prose, even mathematical prose, can 'lie solidly on the syntactic side'. In most of these presentations, the tension between the ideal of a purely syntactic object and the reality of proofs written in semantically rich prose is either shrugged off or set aside with a claim (as in Mac Lane) that the translation between informal proof and formal derivation is 'routine' or otherwise unproblematic. Of course, some aspects of mathematical argumentation are perfectly well modelled by formal logic. Every proof by contradiction, for example, is a clear, undisguised case of *reductio ad absurdum.*

Some mathematicians do worry about in the relation between an informal proof and its corresponding derivations,[5] such as Hersh (1997) and most notably Rav (1999). Rav asked why we prove theorems, and his answers all turned on a claim that proofs include 'topic-specific moves' (p. 26) that cannot be re-cast in formal logic without loss or violence. This is also a central point for Hersh, "The passage from informal to formalized theory must entail loss of meaning or change of meaning" (1997, p. 160). Mac Lane, in contrast, claims that if an informal proof is rigorous, this translation must be 'routine' and that the test for correctness should made no reference to the subject matter, that is, should not be topic-specific (quoted above).

Thus, we have some mathematicians (Hersh, Rav) insisting that informal proofs are often so different from formal derivations that the latter have little relevance to the business of writing and checking the former. Other mathematicians (Mac Lane, Feferman, perhaps Hales and Hamkins) sail close to claiming that a rigorous informal

---

[5] Typically, there are lots of derivations that might formalise a given proof, and sometimes they vary sufficiently to make trouble for Feferman's claim that *the* corresponding derivation captures *the* logical structure of the informal proof. See Tanswell (2015) for cases.

proof in some sense *is* a formal derivation ('at bottom', under 'complete analysis' or up to 'routine' translation). What, in this circumstance, is a philosopher to do? One of the most widely discussed accounts of the matter is Jody Azzouni's (2004)[6] 'derivation-indicator' view.

Azzouni, seeking to specify a role for formal derivations while doing justice to the rich descriptions of proof-practice that he found in Rav and others, argued that an informal proof is not a formal derivation somehow abbreviated or disguised. Rather, a valid informal proof indicates the existence of a formal derivation and in this way meets the norm of formal correctness. Azzouni later modified this view (2009), giving up the phrase 'derivation-indicator' in favour of talk of topic-specific inference packages that sound rather like Rav's topic-specific moves. In spite of this partial rapprochement, there are some important differences between Azzouni's eventual position and Rav's outlook. Azzouni works hard in his (2009) paper to identify the role that formal logic has in informal proof practice (see his discussion of 'content containment'), and he ties his account to human cognitive processes and capacities, whereas Rav talks about topic-specific mathematical moves that might in principle be carried out by a cephalopod or an alien.

Azzouni's chief motive for giving up the derivation-indicator view seems to be that he found himself explaining observable, measurable social facts—the reliability and stability of mathematical knowledge, the confidence of mathematicians in their results and the high degree of consensus in the mathematical research community—by reference to something that in almost all cases is non-existent:

> I kept falling (against my will) into a view that mathematicians had to be engaged in *something* like sophisticated syntactic pattern-recognition while perusing informal mathematical proofs, so that they would be sensitive (without realizing it) to a background of *nonexistent* formal derivations.[7]

Philosophers of the Rav tendency often rely on some version of this point. Formal derivations corresponding to informal proofs almost never exist in material reality. They might be posited in principle, and sometimes the structure of the informal proof might suggest a map of how a derivation could be worked up. Nevertheless, the posited formal derivations almost always remain materially non-existent and therefore ineligible to be elements in the causal history of a psychological or social fact. Whatever job it is that formal derivations are supposed to do in mathematical practice, they somehow must be able to do it without existing.[8]

---

[6] Azzouni modified his view considerably in his (2009) and (2017). Nevertheless, his phrase 'derivation-indicator' has taken on a life beyond his writing. Since Avigad's most recent work is the focus of the present paper, it is worth noting that Azzouni lists Avigad among the acknowledgements of his (2009) paper.

[7] Azzouni, 2009, p. 25. Italics in original.

[8] Of course, beliefs about non-existent formal derivations can be causally effective, and Azzouni makes moves in that direction in his (2017). There, he attempts to rescue his claim—that in reading an informal proof, mathematicians are unconsciously recognising algorithmic processes—by suggesting that the objects of the algorithms need not be linguistic strings. This thought is already in the literature, e.g., De Toffoli (2017), Giardino (2017), Larvor (2012), and Manders (2008) except that Azzouni re-presents manipulations as algorithmic processes. That may work for Azzouni, but it comes at the cost of separating

There are two more points to make about this debate before turning to Avigad's recent paper. First, practicing mathematicians supply the contending philosophical positions. The fact that mathematicians feel moved to say what they do is part of the data for philosophers of mathematical practice. In particular, philosophers who favour Rav's outlook need to explain why mathematicians so often announce some version of the Standard View. It matters, even if it is strictly false.[9] The philosopher's work would not yet be complete if one side of this debate triumphed over the other, because (whichever way it went), all the data must be accounted for, including the reflections of mathematicians committed to the defeated view.

Second, notice how Azzouni got drawn (against his will) into making a claim in philosophy of mind. The thought that informal proofs are really ('at bottom', after 'routine' translation) syntactic items leads naturally to the idea that in reading them, mathematicians are doing some sort of syntactic processing. This is a special case of the idea that all human thinking is syntactic processing, in spite of appearances. That general idea owes some of its grip on the contemporary philosophical imagination to the omnipresence of digital devices, which now carry out tasks that look as if they belong solidly on the semantic side, but which are at bottom (at the level of machine-code), syntactic processes. In addition, the brain-as-computer idea has a role in cognitive science because it helps researchers to ask precise questions (about, for example, memory size and structure). It may be also true, but results in experimental psychology suggest otherwise (Kahneman, 2012). None of the authors cited in this paper have committed themselves to computationalism in the philosophy of mind. Nevertheless, that is the direction in which these arguments point.

To summarise: there is a tension in the Standard View arising from the differences between the ideal of a wholly syntactic proof and the reality of informal proofs written essay-style in prose. Attempts to resolve this tension pull philosophers and mathematicians towards claiming that informal proofs are, in spite of appearances, not really informal. To a philosopher feeling this tension, informal proofs can come to look like formal derivations in disguise, their underlying syntactic nature waiting to be revealed by translation or analysis. From here, it is a short step to claiming that

---

his thesis (reading mathematical proofs is algorithm-recognition) from the idea that informal proofs are always-already formed in something like first-order logic with ZFC. He knows this, and makes moves to reconnect the two claims in the later sections of his paper.

[9] This point has been raised in a helpfully clear way by Zoe Ashton (PhD thesis, forthcoming). It is not a topic of the present paper, but it is easy to see where one might start. Formalisation does all sorts of work in mathematics aside from securing proofs, such as sharpening and deepening ideas and preparing them for communication (see Hamkins pp. 162-3). It is distinctive of modern mathematics (roughly, Hilbert and after). Moreover, proof theory (the branch of formal logic) supplies mathematical models of mathematical reasoning. One might say that proof theory is what mathematics has to say about proof. It is not then surprising that it is also what many mathematicians have to say about proof. This is Hersh's line (1997 p. 154): the mathematician who insists on the Standard View is like the economist who forgets that his model of the economy is only a model. Hersh calls for the model to be tested. Perhaps current work on digital proof assistants is an answer to that call.A philosopher in the Rav tendency might argue thus: mathematicians formalise, but the art of it is tacit knowledge. It is not a summatively assessed part of mathematical training (see Tatton-Brown (2020) for an account of how it is learned). Mathematicians who have entirely internalised the art of formalisation may feel that they're doing almost nothing, something 'routine', rather as a virtuoso jazz musician might insist that they're not thinking about music theory or technique while improvising.

when mathematicians read informal proofs they are really engaging in syntactic processing. With that step taken, the philosopher of mathematics who began by trying to articulate the Standard View then stands on the brink of making a commitment in the long-standing debate about whether the human brain works like a digital computer. There may be an independent argument for the brain-as-computer thesis, but it lies outside the philosophy of mathematical practice.

2 Avigad, J. [2020] reliability of Mathematical Inference

The abstract of Avigad's paper reads as follows:

> Of all the demands that mathematics imposes on its practitioners, one of the most fundamental is that proofs ought to be correct. It has been common since the turn of the twentieth century to take correctness to be underwritten by the existence of formal derivations in a suitable axiomatic foundation, but then it is hard to see how this normative standard can be met, given the differences between informal proofs and formal derivations, and given the inherent fragility and complexity of the latter. This essay describes some of the ways that mathematical practice makes it possible to reliably and robustly meet the formal standard, preserving the standard normative account while doing justice to epistemically important features of informal mathematical justification.

The first thing to note is that what we might call the 'observational core' of his argument—his partial natural history of strategies for constructing sound informal proofs (Sects. 3 and 4 of his paper)—is a significant and insightful contribution to the philosophy of mathematical practice, and will remain so regardless of the fate of his larger argument.

To begin, consider Avigad's gloss on the Standard View:

> According to the standard view, a mathematical statement is a theorem if and only if there is a formal derivation of that statement, or, more precisely, a suitable formal rendering thereof. When a mathematical referee certifies a mathematical result, then, whether or not the referee recognizes it, the correctness of the judgement stands or falls with the existence of such a formal derivation. (2020, p. 4)

It matters to what follows to get the modalities of this statement right. Avigad is not simply saying that wherever there is a sound informal proof, there is as a matter of fact the possibility-in-principle of writing a corresponding formal derivation. If that were the whole of the Standard View, there would be little discussion because, as he points out, long-running programmes in the foundations of mathematics give us compelling reasons to think that every sound informal proof can in principle be supplied with a correct formal correlate, and what is more there is a growing library of formal derivations of significant theorems. Avigad's claim is something stronger: correspondence with a formal derivation is the very meaning of correctness for an informal

proof. Throughout his paper, he uses Azzouni's formulation that the existence of a derivation is the 'normative standard' for the correctness of an informal proof. He rejects Detlefsen's (2008) suggestion that '[informal] rigor and formalization are independent concerns', on the grounds that it 'requires us to provide an independent characterization of rigor and provide some other explanation of what it means for an informal proof to be correct.' (2020, p. 6). It is evident from the dialectical context that he does not expect this challenge to be met. On the other hand, he says nothing here about the relation between the informal proof and its corresponding derivations (unlike Mac Lane's insistence that the translation between them must be 'routine' or Feferman's requirement that the derivation is the logical analysis of the informal proof).

The chief obstacle for the Standard View, according to Avigad, is to explain how informal proofs—long, complex and difficult as they may be—can be the sort of things that entail the existence of formal derivations. To meet this challenge, he deploys the observational core of his argument. Part of his answer is that informal proofs (unlike formal derivations) do not have to be wholly error-free—an informal proof can play its epistemic and logical roles even if it has small, easily fixed errors, mistakes in book-keeping, mislabelling and the like (2020, p. 7). This robustness is in part due to the stabilising role of the skilled mathematical reader, who can see the mistakes and figure out what ought to have been printed. Avigad proceeds to develop the beginnings of a natural history of desirable proof-features. He observes that proofs are modular—they divide into chunks that prove lemmas, and the chunks can be passed around from proof to proof, so a new proof may rely on some previously established modular parts that have already shown their reliability, and new modular parts can be checked in isolation from the rest of the proof. Inferences in proofs are motivated—a move in a proof is intelligible because it plays some part in the overall proof strategy, which makes them easier for the reader to grasp. Avigad suggests three general strategies for writing robust proofs:

- Isolate and minimize critical information.
- Maximize exposure to error detection.
- Leverage redundancy.

The working-through of a proof may involve a lot of standard moves, and a proof written for robustness will downplay this familiar material in order to foreground whatever is novel or unique, so that readers can focus on it and interrogate it. The second point means building in checks so that an error will show up vividly. Finally, the robustness of informal proofs depends in part on the fact that there are often many different ways of proving the same lemma or result. Avigad then suggests a non-exhaustive list of some more specific heuristics:

1. Reason by analogy.
2. Modularize.
3. Generalize.
4. Use algebraic abstraction.
5. Collect examples.

6. Classify.
7. Develop complementary approaches.
8. Visualize.

The analogies Avigad has in mind can be between structures (say between the integers and the Gaussian integers), or arguments, or construction heuristics like embedding a class of objects into a bigger structure with more tractable properties. By 'algebraic abstraction', Avigad means capturing the mathematically important features of a structure in axioms, such as the group axioms. These heuristics can support each other. For example, collecting examples is an obvious precursor to classification, but a classification can be more powerful if it is ordered according to some structural feature, and this feature may be a candidate for capture in axioms, and so become a case of algebraic abstraction.[10] Avigad supplies examples and discussion that I will not reproduce here because I have nothing critical or novel to say about the detail of what I am calling the 'observational core' of his paper.[11] Note, though, that these strategies all depend to some degree on the mathematics being meaningful to human mathematicians. Some (such as 2 and 4) carry over to formal derivations relatively smoothly. Others, such as 8, do not. Curiously, Avigad does not explore the extent to which these strategies for informal proofs do or do not tend to produce proofs that map easily ('routinely') across to formal derivations. On the contrary, he presents these strategies as means for managing logical relations internal to the informal proof. They are not about building connections between the informal proof and a corresponding formal derivation.[12] They are about making sure that the informal proof works in its own terms.

This is why the observational core of Avigad's paper could just as well have been written by someone building on the views of Yehuda Rav or Michael Detlefsen, that is to say, the claim that the rigour of informal proofs can (on the whole, in most cases) be accounted for without reference to formal derivations. Avigad's analysis does not directly explain how informal proofs meet the standard of correctness set out in the Standard View. His analysis cannot do that because it does not mention formal derivations at all. Rather, it explains (or at least, begins to explain) how it is that informal proofs can be truth-preserving. In doing so, it indirectly helps us to understand how they meet the formalisability norm set out in Avigad's reading of the Standard View—but only because results in the foundations of mathematics tell us that a sound informal proof will be formalisable. Informal mathematical proofs are highly truth-preserving in part because they isolate critical information, expose themselves to error detection, leverage redundancy, and so on through the rest of Avigad's good heuristic advice. The observational core of Avigad's paper can be adopted wholesale by critics of the Standard View, because it says nothing about the alleged relation

---

[10] Some anecdotal corroboration for some of these points may be found here: https://mathoverflow.net/questions/338607/why-doesnt-mathematics-collapse-even-though-humans-quite-often-make-mistakes-in.

[11] I will indulge in a joke about the creative tension between 3 and 5. There are two sorts of mathematicians: the ones who, when presented with a new mathematical structure ask, 'What is an example of this?' and the ones who ask, 'What is this an example of?'

[12] I owe this observation to a comment by Dirk Schlimm at the 2020 meeting of the Association for the Philosophy of Mathematical Practice, during the Q&A of Avigad's presentation of his paper.

between informal proofs and formal derivations. Apart from a few references to a paper by Hamami (2019), Avigad's heuristic observations are entirely about informal proofs alone, not about their relation to derivations. The indifference of Avigad's heuristics to formal derivations rather supports the thought that the practically effective norm for mathematical proofs is that they should be sound.

That said, the challenge that Avigad presents to the Rav tendency remains, namely, to provide some alternative account of the correctness of informal proofs. After all, most of Avigad's heuristics would be good advice for legal or historical argument (slot in sub-arguments that have already been well tested, find multiple arguments for the same conclusion, etc.) and almost all of it is good advice for philosophers (collect examples, classify, identify abstract principles, etc.). The challenge for the Rav tendency is to explain how informal proofs attain mathematical levels of reliability, and this cannot be achieved by only pointing to general heuristics of this sort. Some in the Rav tendency have proposals, but testing them is not the aim of the present paper.[13] Proponents of the strongest version of the Standard View—that sound informal proofs are formal derivations in disguise—can turn to formal logic to explain how proofs work, but do have to make their core thesis plausible in the face of Avigad's worry about fragility, Azzouni's worry about how mathematicians engage with derivations and fact that informal proofs and derivations look different and seem to work differently. Someone who wishes to defend the Standard View without claiming that informal proofs are only superficially informal would face the same challenge as the Rav tendency. If we allow that informal proofs are not merely syntactic objects disguised in human language, then the challenge that Avigad presents to the Rav tendency is the very problem that he addresses in his paper. I will argue in the next section that Avigad equivocates between these two versions of the Standard View.

This completes the case for the first claim of this paper: that the observational core of heuristic advice in Avigad's paper is entirely agnostic between rival accounts of mathematical correctness.

3 The syntactic side of the tension

I now move to the second claim of this paper, that the tension at the heart of the Standard View is still there in Avigad's version and sometimes pushes him towards radical theses in philosophy of mind. Recall how that temptation goes: informal proofs and formal derivations often look quite different, and in any case the formal derivation corresponding to a proof is almost never produced. Nevertheless, according to the Standard View, the existence of formal derivations is somehow normative for informal proofs. A tempting resolution is to argue that in spite of appearances, the informal proof is really a derivation in disguise. Then, two problems for the Standard View—the non-existence of derivations and their apparent difference in kind from informal proofs—both disappear. In that case, 'routine' translation and faithful logical analysis do not transform a proof but merely reveal its essence. Once this line has been taken, it is natural to adopt a computational view of the mathematical mind.

---

[13] See Larvor (2012, 2019).

Since mathematicians are humans, a general computationalism in the philosophy of mind beckons.

As we saw, Avigad's characterisation of the Standard View does not describe the relation between informal proofs and formal derivations, but he does offer some clues in the body of his paper. Near the start of his paper, Avigad adopts the familiar image of informal proofs as sketches, and Azzouni's language of derivation indication:

> On the standard view, [informal proofs] are only high-level sketches that are intended to indicate the existence of formal derivations. But providing less information only exacerbates the problem… (2020, p. 4)

Avigad does not unpack the sketch analogy, but he does recommend Yacin Hamami's (2019) which does exactly that.[14] Echoing Mac Lane, Hamami writes: 'a mathematical proof is rigorous if and only if it can be *routinely translated* into a formal proof.'[15] The central part of Hamami's paper is a careful specification what he thinks 'routine translation' has to mean for the Standard View to be cogent:

> …the notion of translation in the standard view is quite different from the one of linguistic translation. …a better analogy would be with the process of compilation, that is, with the 'translation' of a computer program written in a high-level programming language into machine language.[16]

Hamami distinguishes four levels of 'code' linked by three stages of 'compilation'. The essential point is that on his view (which Avigad sometimes suggests he shares), the 'routine translation' of informal proof into formal derivation that the Standard View requires is not like translating a meaningful text from one human language to another. It is, rather, like the use of syntactic substitutions to move from a high-level computer language where a programmer can call an arbitrarily large and complex subroutine with a single command down to the level just above the firmware. 'Routine' here does not mean 'easy for humans' or even 'feasible for humans'. It means that a digital machine could do it. It also means that, layers of compilation and syntactic abbreviation notwithstanding, informal proofs and formal derivations are essentially the same sort of thing. In other words, in Hamami (though not directly in Avigad), we find the strong version of the Standard View, that sound informal proofs are syntactic items in disguise.

Avigad offers some other analogies in the same vein:

> …an informal proof is a form of data compression. Coding schemes for data represent the most common patterns concisely, reserving extra bits for those that are unusual and unexpected. To invoke another analogy, software developers use version control software to store a project's history in a shared reposi-

---

[14] Avigad says that, 'Hamami's model is essentially correct…' (2020, p. 15). 'Essentially' here warns us not to hold Avigad responsible for every nuance of Hamami's elaboration of the Standard View.

[15] From the abstract. Italics in original.

[16] Hamami credits this analogy to Gil Kalai (Kalai 2008). (Hamami 2019, p. 30).

tory. To save space, the repository only has to store the difference between successive versions, rather than a new copy.

As in Hamami, the difference between the informal proof and the formal derivation is really just a process of selective abbreviation. They are both syntactic objects—it's just that the formal derivation is much longer because it is expressed in a foundational system (usually taken to be a combination of set theory and classical logic). In these data compression and software development analogies, Avigad seems to commit himself to the strong version of the Standard View, that informal proofs are syntactic objects in disguise. However, Avigad does not always follow this line.

4 The semantic side of the Tension

The claim of the present paper is that the tension in the Standard View is there in Avigad's article as it was in Mac Lane, Azzouni, and the other authors we canvassed above. The Standard View looks implausible because most informal proofs do not seem to work like purely syntactic objects, even though mathematical notation may make them look like syntactic objects. Like other kinds of reason-giving and argument-making, informal mathematical proofs depend on communities of practitioners sharing a common understanding of what moves are permitted within the practice, in virtue of their grasp of the nature of the subject-matter. Mathematicians have some advantages over other communities of reasoners. They can offload some of their thinking to highly efficient notational systems that partially encode the rigour of the practice. Ordinary high-school algebraic notation can serve an example of this: its introduction in the early seventeenth century massively extended the range, depth and reliability of mathematics. Following the implicit rules for forming and manipulating algebraic expressions makes it much easier to avoid making a mistake—but it's still easy to divide by zero without noticing, unless you pay attention to the meaning of what you're doing as well as the syntax. Mathematicians can gain clarity and rigour by exploiting the human capacity for visual and kinetic imagination, either by inward imagining or by using diagrams or physical models (this is Avigad's eighth heuristic). They can exploit Peircean iconicity, because mathematics is about structure, so we can use representations that share structure with the objects under investigation.[17] Increasingly, as the variety of mathematical domains grows, mathematicians can create arguments by moving between different parts of mathematics. Ken Manders calls this the 'conceptual agility' of contemporary mathematics:

> …hopping a functor to another category any time superfluous details are sensed (homology groups in topology); bringing in details seemingly extraneous to one's question by a representing functor (group representations in the theory of abstract groups); explaining families of simple number-theoretic facts any-

---

[17] See, e.g., De Toffoli (2017), Giardino (2017), and Manders (2008).

one can see one-by-one, by some hard to master 'underlying' abstract structure (Fermat's problem).[18]

Here Kevin Buzzard, a mathematician at the heart of the *Lean* proof formalisation project, makes a similar point:

> Mathematicians [are] so good at instantly switching between the various 'obviously equivalent' ways that a mathematician looks at a complicated algebraic object ('It's an equivalence relation! Now it's a partition! Now it's an equivalence relation again! Let your mental model jump freely to the point of view which makes what I'm saying in this particular paragraph obvious!', or 'Matrices are obviously associative under multiplication because functions are associative under composition.'... Some of the proofs we're writing [in *Lean*] are simply proofs that humans are behaving correctly when using mathematical objects. (Buzzard 2020)

Proving that the moves mathematicians make are correct is not the same activity as making those same moves more slowly in a less rich notation.

These jumps are helpful precisely because they allow us to bring domain-specific techniques from one area to bear on a problem first found in another. A detour through complex analysis in order to solve a problem in real analysis, or a visit to real analysis to help prove a result in number theory, does not look much like re-coding of syntactic data. It is more like moving from one natural language to another in search of the *mot juste* or looking for legal arguments in a different area of law from the case in hand. Mathematicians translate problems and proof-ideas between different parts of mathematics all the time, and translation of a proof into a foundational idiom is just a special case of this. Rather than choosing a metaphor to describe translation between informal proofs and formal derivations with the aim of preserving the Standard View, it might be less question-begging to work up a general account of intra-mathematical translation on a broad evidence-base, and then see what it says about the special case of translations of proofs into formal derivations.

This view of mathematics as a meaningful human reason-giving activity appears in places in Avigad's article. He notes that mathematical proofs written for human use are robust in the sense that that they can survive errors of syntax and abuses of notation—in fact, abusing notation tactically is an important mathematical skill, but this thought makes no sense if proofs are syntactic items. They're also insensitive to minor mistakes in calculations or localised reasoning, because readers can correct the errors (as Avigad observes). This can only be because readers understand what the inferences are about and are not simply processing symbol strings syntactically. Some of Avigad's heuristics—analogy, generalisation, abstraction, exemplification, classification and visualisation—only make sense as strategies for handling meaningful mathematical content. An analogy is not, at the point where it is useful, a syntactic relation. Strategic generalisation is not something a machine can do, not least

---

[18] Unpublished paper. Though this paper is unpublished, it is a developed piece of work and benefits from two decades of refinement. We may therefore treat it as a reasonably reliable expression of Manders' view.

because very often more than one generalisation is available (for example, if you know something about equilateral triangles, you might generalise to other planar regular polygons or you might generalise to simplexes in higher dimensions). Similarly, noticing that there is a structure in common that is open to axiomatisation, identifying examples that will be both tractable and informative, and classifying in a way that carves the domain at the joints all depend on familiarity with and understanding of mathematical content. It's true that these activities are preparations for proving rather than parts of proofs. However, it is implausible to suppose that the same mathematical content suddenly switches from being richly significant to become a series of syntactic operations when it is tidied up and arranged into a proof.[19] One of Avigad's metaphors makes the very point:

> Comparing proof to a narrative allows us to draw on intuitions regarding the distinction between the plot and its syntactic presentation. After reading *Pride and Prejudice*, we can reliably make the claims about the characters and their motives, but we cannot reliably make claims about the number of occurrences of the letter 'p.' If the judgment as to the correctness of a proof is more like the former, we have a chance; if it is more like the latter, it is hopeless. (2020, p. 8)

Evidently, Avigad favours the former option: thinking about the correctness of a proof is more like making judgments about the plot and characters of a novel, than it is like guessing how many times a letter occurs in the manuscript. This thought, however, undermines the metaphors in his paper that encourage syntactic readings of informal proofs. The relation between the plot of a novel and the typescript is not like data-compression or code-compiling. Turning a plot-and-character idea into a fully written out novel is not a routine translation, in any sense of 'routine', let alone in the sense that Hamami develops in his paper.

To fix ideas, consider this example: prove that given any two potatoes, there is a closed loop of the same size and shape on both skins. Like many mathematical proofs in philosophy articles, this is very short and untechnical. It is in these respects quite unlike typical proofs in research mathematics, but it is nevertheless a mathematical proof so whatever we say about mathematical proofs in general must be true of it. The solution is in this[20] footnote. It is, of course, possible to translate this proof into the language of contemporary geometry and thence into a suitable foundational language, so that it no longer relies on spatial intuition and the inferences are all applications of inferential rules of the general logic of the chosen foundational idiom. The issue is the nature of this translation. According to Hamami, and to Avigad when he

---

[19] Though it can sometimes feel that way. Here is Alain Connes, '…when a mathematician works, he is in fact reflecting up on a certain field, in which he encounters mathematical beings, and ends up playing with them, until they become familiar to him… After a time… either we get nowhere,… or else we manage to get hold of some result. Then begins the onerous task—the obligation to write up for the benefit of the mathematical community a polished article that is as compelling as possible.' (Connes et al., 2001, p. 24f). Connes insists that these are two wholly distinct types of activity. Nevertheless, one supplies materials for the other—untranslated.

[20] Imagine pushing one potato surface into the other, so that they intersect. The intersection is the required common loop. Proof taken from Tim Gowers' Twitterfeed.

follows Hamami, the relation between the proof as offered here and its corresponding derivation is one of data-compression. For this to be true, the proof as written here must already be a syntactic object, and the reasoning in it must be the application of wholly general syntactic rules. It is, of course, possible to conceive of it thus, but in doing so, we would be re-reading it in order to save a philosophical theory from refutation. Alternatively, we could take it to be what it seems to be, namely, a piece of spatial reasoning that works not by manipulating meaningless symbol strings but rather by manipulating imagined potato surfaces.

This brings me to another difficulty Hamami's computational picture of the translation between informal proof and formal derivation: it undersells the achievement of the foundational programmes in mathematics. Avigad sometimes wonders aloud[21] why critics of the Standard View seem to want to disregard the great advances in rigour achieved by twentieth-century mathematics. It is wonderful that it is now possible in principle to check any piece of mathematical reasoning mechanically, that mathematicians can pursue piecemeal formalisation as far as they need to in any given case, and that there is a growing library of fully formal proofs of significant theorems.[22] However, if we think of mathematics as carried out by humans as syntactic processing, as the data-compression and code-compilation metaphors require, then this becomes less of an achievement. If all mathematical thinking is and always was syntactic processing, then the great foundational programmes did nothing more than set up the compilation tables. The real wonder, the deeply significant achievement, is that it is now possible, with enough work, to check even the most richly human mathematics by machine. It is as if we could judge the coherence of the plot of a novel by asking a computer to count the 'p's in the manuscript. Hamami's view of translation undervalues this achievement and misdescribes the work involved in preparing a proof for mechanical checking. In Avigad's paper, there is already a hint about how the formalisation of meaningful human mathematics is possible: many of the heuristics described in the observational core of Avigad's paper simultaneously support the rigour of the informal proof *and* prepare it for translation into a more formal idiom. Following this clue may help us to understand how machine proof assistants are useful to mathematicians, but only if we resist any temptation to begin by positing that human mathematicians were machines all along.

5 Conclusions

I have argued for two claims in this paper. The first is that the observational core of Avigad (2020) is independent of the Standard View that formalisability is the criterion of correctness for mathematical proofs. In doing so, I distinguished between 'strong' versions of the Standard View that claim that sound informal proofs are, in spite of appearances, really syntactic objects and that when mathematicians read them, they're doing some sort of syntactic processing, and 'weaker' versions that don't have that commitment. My second claim is that Avigad, in his (2020), havers

---

[21] Personal communication at the APMP conference in Zurich, 2020.

[22] For a spectacular recent example, see Castelvecchi (2021).

between these two versions of the Standard View. He is not alone in this—the tension between these two options is present in the debate from Mac Lane onwards. For the discussion to advance, proponents of the Standard View need to specify which version they endorse.

Declarations

References

Avigad, J. (2020). Reliability of mathematical inference. *Synthese*https://doi.org/10.1007/s11229-019-02524-y

Azzouni, J. (2004). The derivation-indicator view of mathematical practice. *Philosophia Mathematica*, *12*(3), 81–105

Azzouni, J. (2009). Why do informal proofs conform to formal norms? *Foundations of Science*, *14*(12), 9–26

Azzouni, J. (2017). Does reason evolve?(Does the reasoning in mathematics evolve?). *Humanizing mathematics and its philosophy* (pp. 253–289). Cham: Birkhäuser

Bourbaki, N. (1949). Foundations of mathematics for the working mathematician. *The Journal of Symbolic Logic*, *14*(1), 1–8

Brouwer, L. E. J. (1913). Intuitionism and formalism. *Bulletin of the American Mathematical Society*, *20*(2), 81–96

Buzzard, K. (2020). Two types of universe for two types of mathematician Entry in blog *Xenaproject*. Posted on July 23, 2020. https://xenaproject.wordpress.com/2020/07/23/two-types-of-universe-for-two-types-of-mathematician/ Accessed: 23/07/2020

Castelvecchi, D. (2021). Mathematicians welcome computer-assisted proof *Nature*. Vol. 595, pp. 1819

Connes, A., Lichnerowicz, A., & Schützenberger, M. P. (2001). *Triangle of thoughts*. Providence, RI: American Mathematical Society. French Original: *Triangle de pensées*. Paris: Odile Jacob, 2000

Detlefsen, M. (2008). Proof: Its nature and significance. In Bonnie Gold and Roger A. Simons, editors, *Proof and Other Dilemmas: Mathematics and Philosophy*, pages 3–32. Mathematical Association of America

De Toffoli, S. (2017). 'Chasing' the diagram—the use of visualizations in algebraic reasoning. *The Review of Symbolic Logic*, *10*(1), 158–186

Feferman, S. (2012). "And so on… reasoning with infinite diagrams. *Synthese*, *186*, 371. doi:https://doi.org/10.1007/s11229-011-9985-6

Giardino, V. (2017). «L'imagination manipulatoire en mathématique», *Bulletin d'Analyse Phénoménologique* [En ligne], Volume 13, Numéro 2: L'acte d'imagination: Approches phénoménologiques (Actes n°10)

Hales, T. C. (2012). *Dense Sphere Packings: a Blueprint for Formal Proofs*. Cambridge University Press

Hamami, Y. (2019). 'Mathematical Rigor and Proof,'. *The Review of Symbolic Logic* (pp. 1–41). Cambridge University Press

Hamkins, J. D. (2020). *Lectures on the philosophy of mathematics*. MIT Press

Hersh, R. (1997). Prove—Once More and Again *Philosophia Mathematica* (III) 5(2):153–165

Hilbert, D. (1927). The Foundations of Mathematics. Translated in van Heijenoort (ed) *From Frege to Gödel: a source book in mathematical logic, 1879–1931*. Harvard, 1967, 1999

Kahneman, D. (2012). *Thinking Fast and Slow*. Penguin

Kalai, G. (2008). Can Category Theory Serve as the Foundation of Mathematics? Entry in blog *Combinatorics and More*, https://gilkalai.wordpress.com/2008/12/04/can-category-theory-serve-as-the-foundation-of-mathematics/

Larvor, B. (2012). "How to think about informal proofs. *Synthese*, *187*(2), 715–730

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

Mac Lane, S. (1986). *Mathematics: Form and Function*. New York: Springer

Manders, K. (2008). The Euclidean diagram. Mancosu, P., *The Philosophy of Mathematical Practice*, Oxford University Press 2008, pages 80–133, 1995

McLarty, C. (1997). Poincaré: Mathematics, Logic and Intuition. *Philosophia Mathematica*, *5*(2), 97–115

Poincaré, H. (1902). On the nature of mathematical reasoning. Excerpt from *Science and Hypothesis*, reprinted in Benacerraf & Putnam (eds) *Philosophy of Mathematics: Selected Readings.* 2nd Edition, Cambridge University Press, 1983

Poincaré, H. (1908). *Science et Méthode*. Ernest Flammarion, Paris. Translated into English by Francis Maitland as *Science and Method*, Thomas Nelson & Sons, London, 1914

Rav, Y. (1999). Why do we prove theorems? *Philosophia Mathematica*, *7*(3), 5–41

Tanswell, F. (2015). A Problem with the Dependence of Informal Proofs on Formal Proofs. *Philosophia Mathematica (III)*, *23*(3), 295–310

Tatton-Brown, O. (2020). Rigour and Proof. *The Review of Symbolic Logic*, 1–29. doi:https://doi.org/10.1017/S1755020320000398

Zach, R. (2019). "Hilbert's Program", *The Stanford Encyclopedia of Philosophy* (Fall 2019 Edition), Edward N. Zalta (ed.)

Authors and Affiliations

Brendan Philip Larvor[1]

✉   Brendan Philip Larvor
     b.p.larvor@herts.ac.uk

[1]   University of Hertfordshire, Hertfordshire, UK