Big Math and the One-Brain Barrier: The Tetrapod Model of Mathematical Knowledge

JACQUES CARETTE, WILLIAM M. FARMER, MICHAEL KOHLHASE , AND FLORIAN RABE

The Viewpoint column offers readers of the Mathematical Intelligencer the opportunity to write about any issue of interest to the international mathematical community. Disagreement and controversy are welcome. The views and opinions expressed here, however, are exclusively those of the author. The publisher and editors-in-chief do not endorse them or accept responsibility for them. Articles for Viewpoint should be submitted to the editors-in-chief.

In recent decades we have seen mathematics tackle problems whose solutions require increasingly large developments in proofs, computations, data sets, and document collections. This trend has led to intense discussions about the nature of mathematics, raising questions like these:

1. Is a proof that can be verified only with the help of a computer still a valid mathematical proof?
2. Can a collection of mathematical documents that exceeds what can be understood in detail by a single expert be a legitimate justification of a certain mathematical result?
3. Can a collection of mathematical texts—however large—adequately represent a large body of mathematical knowledge?

The first question was first raised by Appel and Haken's proof of the four color theorem [2], which in four hundred pages of regular proof text reduced the problem to approximately two thousand concrete configurations that had to be checked by a computer. It arose again from Thomas Hales's proof of the Kepler conjecture [20], which contained substantial algebraic computations, carried out by computer as part of the proof.

The second question is raised by, for example, the classification of finite simple groups (CFSG), which comprises the work of a large group of mathematicians over decades and which has resisted dedicated efforts even to be written down consistently; see the discussion below.

The third question arises from the ongoing development of digital mathematics libraries (DMLs)—such as the huge collection of papers that constitute the CFSG—that fail to make explicit the abundant interconnections in mathematical knowledge that are needed to find information in these DMLs and reuse it in new contexts.

Let us call such developments Big Math, by analogy with the big data/big everything meme but also alluding to the New Math movement of the 1960s that aimed to extend mathematical education by changing how mathematics was taught in schools. The emerging consensus of the mathematical community seems to be that while the methods necessary for dealing with Big Math are rather problematic, the results so obtained are too important to forgo by rejecting such methods. Thus, we need to take on board Big Math methods and understand the underlying mechanisms and problems.

In what follows, we analyze the problems, survey possible solutions, and propose a unified high-level model that we claim computer support must take into account for scaling mathematics. We believe that suitable and acceptable methods should be developed in a tight collaboration

between mathematicians and computer scientists—indeed such development is already underway but needs to become more comprehensive and integrative.

We propose that all Big Math developments include five main aspects that need to be dealt with at scale:¹

1. *Inference*: deriving statements by deduction (proving), abduction (conjecture formation from best explanations), and induction (conjecture formation from examples).
2. *Computation*: algorithmic manipulation and simplification of mathematical expressions and other representations of mathematical objects.
3. *Concretization*: generating, collecting, maintaining, and accessing sets of concrete objects that serve as examples, suggest patterns and relations, and allow testing of conjectures.
4. *Narration*: bringing the results into a form that can be digested by human beings in natural language but also in diagrams, tables, and simulations, usually in mathematical documents such as papers, books, and webpages.
5. *Organization*: representing, structuring, and interconnecting mathematical knowledge.

Computer support exists for all of these five aspects of Big Math, such as, for example, the following:

Inference:	theorem provers such as Isabelle [50], Coq [14], and Mizar [34].
Computation:	computer algebra systems such as GAP [18], SageMath [44], Maple [29], and Mathematica [53].
Concretization:	mathematical databases such as the L-Functions and Modular Forms Database (LMFDB) [15, 27] and the Online Encyclopedia of Integer Sequences (OEIS) [46].
Narration:	online journals, mathematical information systems such as zbMATH [54] and MathSciNet [32], preprint servers such as arXiv.org, and research-level help systems like MathOverflow [31].
Organization:	libraries of theorem provers and mathematics encyclopedias such as Wolfram MathWorld [33] and PlanetMath [42].

While human beings can easily integrate these five aspects and do so for all mathematical developments (large or otherwise), much research is still necessary into how such an integration can be achieved in software systems. We

want to throw the spotlight on the integration problem to help start off research and development of systems that integrate all five aspects.

Although the correctness of mathematical results is very important, checking correctness is not our primary concern, since there are effective societal mechanisms and, when needed (as for Hales's proof of the Kepler conjecture), inference systems for doing this. Similarly, although the efficiency of mathematical computation is likewise very important, computational efficiency is also not our primary concern, since there are programming languages and computer algebra systems that enable highly efficient mathematical computation. Correspondingly, there are highly scalable systems for working with large databases and collections of narrative documents. We will not even concern ourselves with complexity issues here, since essentially all computational problems involved are either NP-complete or undecidable. Instead, our primary concern is to provide, by integrating these aspects, new capabilities that are needed but currently unavailable for Big Math.

In the next section we discuss some of the state of the art in computer support for Big Math by way of high-profile mathematical developments, which we use as case studies for Big Math and present the issues and methods involved. Then we present a proposal for the integration of the five aspects into what we call a tetrapod, whose body is organization and whose legs are the other four aspects, arranged like the corners of a tetrahedron. The last section summarizes the main points raised in this position paper.

Computer Support of Mathematics

The classification of finite simple groups (CFSG) is one of the seminal results of twentieth-century mathematics. Its usefulness and mathematical consequences give it a prominent status in group theory, similar to that of the fundamental theorem of arithmetic in number theory. The proof of the CFSG was constructed through the coordinated effort of a large community over a period of at least half a century; the last special cases were completed only in 2004.

The proof itself is spread over many dozens of contributing articles totaling over ten thousand pages. As a consequence, work on collecting and simplifying the proof has been underway since 1985, and it is estimated that the emerging "second-generation proof" can be condensed to five thousand pages.²

It seems clear that the traditional method of doing mathematics, which consists in well-trained, highly creative individuals deriving insights in the small, reporting on them in community meetings, and publishing them in academic journals or monographs, is reaching the natural limits posed by the amount of mathematical knowledge that can be held in a single human brain—we call this the one-brain

¹Finding good names for the five aspects has been a major effort in our discussions about the model. We went through many versions before we arrived at the current set. The names need to be sufficiently general to convey the corresponding activity in all its richness without being arbitrary or unrecognizable. In the end, we decided on nouns derived from active verbs to evoke the sense that they are about "doing mathematics." The most difficult concept was that of "concretization"; it had competitors such as "tabulation" for the act of organizing concrete objects systematically and "data" for collections of concrete objects (which, however, is not an activity like the others).

²See the Wikipedia article on the classification of finite simple groups.

barrier (OBB).³ We posit that transcending the OBB will be a crucial step toward future mathematics. The space of mathematical knowledge on this side of the OBB is bounded by the amount of time and memory capacity that a single individual can devote to learning the scaffolding necessary to understand or build on a specific result. More specifically, the point at which a mathematical domain grows so much that it would take a mathematician more than twenty-five years of work to be able to contribute new ideas would likely signify the end of research in that domain. Indeed, we are seeing a gradual increase in the number of large proofs, which might point to a decrease in important mathematical results inside the OBB. For example, many high-profile open problems like the Riemann hypothesis might be elusive precisely because they are beyond the OBB.

There are two obvious ways around the OBB:

1. breakthroughs in the structural understanding of wide swaths of mathematics, so that the effort of learning about a particular domain can be greatly reduced;
2. large-scale computer support.

These are not mutually exclusive, and computer support may indeed enable such breakthroughs.

Computers Versus Humans in Mathematics

Humans and computers have complementary performance characteristics: Humans excel at vertical tasks, those that involve deep and intricate manipulations, intuitions, and insights but limited amounts of data. In contrast, computers shine where large data volumes, high-speed processing, relentless precision, but shallow inference are required, that is, in horizontal tasks.

In mathematics, vertical tasks include the exploration, conceptualization, and intuitive understanding of mathematical theories and the production of mathematical insights, conjectures, and proofs from existing mathematical knowledge. Horizontal tasks include the verification of proofs; the processing of large lists of examples, counterexamples, and evidence; and information retrieval across large tracts of mathematical knowledge.

Enlisting computers for horizontal tasks has been extremely successful—mathematicians routinely use computer algebra systems, sometimes to perform computations that have pushed the boundary of our knowledge. Other examples include data-driven projects like LMFDB, which tries to facilitate Langland’s program [6] in number theory by collecting and curating objects like elliptic curves and their properties, or OEIS, which similarly collects sequences of integers. These already form respectively Big Computation and Big Concretization (aka Big Data)

approaches, but they do not help in the case of the CFSG, which is more of a Big Inference and Big Narration problem though it involves the other aspects as well. In the sequel, we consider the issues involved in an exemplary Big Inference effort.

A Computer Proof of the Odd-Order Theorem

In 2014, Georges Gonthier’s team presented a machine-checked proof of the Feit–Thompson odd-order theorem (OOT), in the Coq theorem prover [19].⁴ Even though Gonthier characterizes the OOT as the “foothills before the Himalayas constituted by the CFSG,” the Coq proof was a multiperson–multidecade endeavor, and the Coq verification ran multiple hours on a current computer. This proof arguably sits at the edge of the OBB or perhaps already transcends it. In this article, we want to analyze the kind of system we would need for breaking the OBB and pushing the boundaries of mathematical knowledge. But before we do, let us recap how proof verification works.

In a nutshell, the theorem and all the prerequisite background knowledge are expressed in a logic, i.e., a formal language $\mathcal{L}$ equipped with a calculus $\mathcal{C}$. In $\mathcal{L}$, the well-formedness of an expression is rigorously defined, so that it can be decided by running a program in finite time, while $\mathcal{C}$ is a set of rules that transform $\mathcal{L}$ -expressions into other $\mathcal{L}$ -expressions. Finally, a proof of a theorem t (an $\mathcal{L}$ -expression) is a series of applications of $\mathcal{C}$ -rules to $\mathcal{C}$ -axioms that end in t . Essentially, a $\mathcal{C}$ -proof gives us absolute assurance that t is a theorem in $\mathcal{C}$. Crucially, the property of being a $\mathcal{C}$ -proof is also decidable. However, the existence of a proof is usually undecidable, and the cost of producing a proof is significant, since its structure must be made explicit enough that a machine can fill all gaps using both decision procedures and heuristic search. Of course, all lemmas that lead up to the theorem are checked as well, so that—in contrast to informal mathematics—we are sure that all the pieces fit together exactly.

The Engineering Aspect: Inference and Knowledge Organization

But this is only half the story. To cope with the complexity of calculus-level proofs, it is much more convenient to use expressive logics—the calculus of inductive constructions (CIC) [7] in the case of Coq—and programs that support the user in proof construction.⁵ A proof like the one for the OOT can have billions of steps and is generated in memory of the Coq system only during proof checking. Programs like Coq are engineering marvels, optimized to cope with such computational loads.

³It might be argued that much mathematical research is now carried out in small groups instead of by individuals and that this should rather be called the “small group brain-pool barrier” (SGBB), but there are natural limits to collaboration on complex topics, as was epitomized in the seminal 1975 book *The Mythical Man-Month* [8]. The main result of this study is that adding members to a team can even slow down progress, because it induces communication overhead; the main solution proposed is to introduce individuals to the team that achieve “a detailed understanding of the whole project,” making the difference between an OBB and an SGBB a matter of degree rather than kind.

⁴The theorem, whose proof was published in 1963, states that every finite group of odd order is solvable.

⁵Proofs in calculus and other areas of mathematics are traditionally expressed in a natural language: mathematical vernacular [10], a stylized form of English with interspersed mathematical formulas, tables, and diagrams.

Optimization is also needed in the organization of the requisite knowledge if one is going to achieve the scale necessary for the OOT. Without care, we frequently end up re-proving similar lemmas, resulting in an exponential blowup of the quantity of work required. To alleviate this problem, we follow the (informal) mathematical practice of generalizing results and proving each lemma at the most general level possible. However, in the formal-methods setting, we need to extend the logics, calculi, and proof construction machinery involved to take modular development into account and optimize them accordingly. For the OOT, Gonthier and his team developed the method of mathematical components [28] (akin to object-oriented modeling in programming languages, but better suited to mathematics) inside CIC and used that to control the combinatorics of mathematical knowledge representation. Indeed, the development of the library of reusable intermediate results comprised about eighty percent of the development effort in the OOT proof.⁶

Five Aspects of Big Math Systems

We have seen two essential components of computer systems that can scale up to the Big Math level: efficient and expressive theorem proving systems and systems for organizing mathematical knowledge in a modular fashion. We recall the five basic aspects of mathematics mentioned in the introduction to this paper:

Inference:	the acquisition of new knowledge from what is already known.
Computation:	the algorithmic transformation of representations of mathematical objects into more readily comprehensible forms.
Concretization:	the creation of static, concrete data pertaining to mathematical objects and structures that can be readily stored, queried, and shared.
Narration:	the human-oriented description of mathematical developments in natural language.
Organization:	the modular structuring of mathematical knowledge.

These aspects—their existence and importance to mathematics—should be rather uncontroversial. In order to help the reader understand their tight interrelationships, Figure 1 arranges them in a convenient representation in three dimensions: we locate the organizational aspect at the center and the other four aspects at the corners of a tetrahedron, since the latter are all consumers and producers of the mathematical knowledge represented by the former. A four-dimensional representation might be more accurate but less intuitive. We note that the names of the aspects are all derived from verbs describing mathematical activity: (i)

inferring mathematical knowledge, (ii) computing representations of mathematical objects, (iii) concretizing mathematical objects and structures, (iv) narrating how mathematical results are produced, and (v) organizing mathematical knowledge.

Below we look at each aspect in turn using the CFSG and related efforts as guiding case studies and survey existing solutions with respect to the tetrapod structure from Figure 1.

Inference

We have already seen an important form of machine-supported mathematical inference: deduction via machine-verified proof. There are other forms: automated theorem provers can prove simple theorems by systematically searching for calculus-level proofs (usually for variants of first-order logic), model generators construct examples and counterexamples, and satisfiability solvers check for Boolean satisfiability. All of these can be used to systematically explore the space of mathematical knowledge and can thus constitute a horizontal complement to human facilities.

Other forms of inference yield plausible conclusions instead of provable facts: abduction (conjecture formation from best explanations) and induction (conjecture formation from examples). Machine-supported abduction and induction have been studied much less than machine-supported deduction, at least for producing formal mathematics. However, there is now a conference series [1] that studies the use of machine learning techniques in theorem proving.

One of the main problems with Big Inference for mathematics is that inference systems are (naturally and legitimately) specialized to a particular logic. For instance, interactive proof assistants like Coq and HOL Light [23] have very expressive languages, whereas automated proof search is possible only for simpler logics in which the combinatorial explosion of the proof space can be controlled. This makes inference systems very difficult to interoperate, and thus all their libraries are silos of formal mathematical knowledge, leading to duplicated work and missed synergies—in analogy to the OBB, we could conceptualize this as a one-system barrier of formal systems. There is some work on logic- and library-level interoperability—we speak of logical pluralism—using metalogical frameworks (logics to represent logics and their relations) [26, 41, 45]. We contend that this is an important prerequisite for organizing mathematical knowledge in Big Math.

Computation

Computer scientists have a very broad view of what computation is. Be it β -reduction (in the case of the lambda calculus), transitions and operations on a tape (for Turing machines), or rewrites in some symbolic language, all of these are somehow quite removed from what a mathematician thinks when “computing.” Here, for the sake of

⁶Georges Gonthier, private communication, 2017.

Figure 1. The five aspects of Big Math systems as a tetrapod.

simplicity and familiarity, we will largely be concerned with symbolic computation, e.g., manipulation of expressions containing symbols that represent mathematical objects. Of course, there are also many flavors of numeric computation such as in scientific computing, simulation/modeling, statistics, and machine learning.

In principle, mathematical computation can be performed by inference, e.g., by building a constructive proof that the sum $2 + 2$ exists. But this is not how humans do it—they are wonderfully flexible in switching between the computational and the inferential aspects. Current inference-based systems in wide use have not achieved this flexibility, although systems like Coq, Agda [38, 39], and Idris [24] are making inroads. In any case, computation via inference is intractable (even basic arithmetic ends up in an unexpected complexity class)—somewhat dual to how inference via computation in decision procedures has only limited success.

Instead, the most powerful computation systems are totally separate from inference: computer algebra systems like Maple, Mathematica, SageMath, and GAP can tackle computations that are many orders of magnitude larger than what humans can do—often in mere milliseconds.

But these systems face the same interoperability problems as inference systems do, open standards for object representation like OpenMath [40] and MathML [3] notwithstanding. Just to name a trivial but symptomatic example, a particular dihedral group is called D_4 in Sage and D_8 in GAP due to differing conventions in the respective communities. More mathematically involved and therefore more difficult to fix is that most of the implementations of special functions in computer algebra systems differ in the underlying branch cuts [12]. Inference during computation would enable some of these problems to be fixed, but this has been sacrificed in favor of computational efficiency. Another source of complexity is that today's most feature-rich symbolic computation systems are both closed-source and commercial, which makes integrating them into a system of trustable tools challenging. That having been said, the kinds of effort devoted to the development of those systems is significantly greater than what can currently be achieved in academia, where code contributions are undervalued compared to the publication of research papers. Furthermore, obtaining funding for sustainable development of large software is difficult.

Lastly, there is the question of acceptability of certain computations in proofs. In part, this derives from the difficulty in determining whether a program written in a mainstream programming language is actually correct, at least to the same level of rigor that other parts of mathematics are subjected to. Some of this problem can be alleviated by the use of more modern programming languages that have well-understood operational and denotational semantics, thus putting them on a sound mathematical footing. Nevertheless, it will remain that any result that requires thousands of lines of code or hours of computation (or more) that cannot be verified by humans is likely to be doubted unless explicit steps are taken to ensure its correctness. The Flyspeck project [22]—a computer verification of Thomas Hales's proof of the Kepler conjecture [20] in the HOL Light and Isabelle proof assistants of comparable magnitude to the OOT proof effort—provides an interesting case study, since it includes the verification of many complex computations.

Concretization

If we look at the CFSG, we see that that result already contains an instance of concretization: the collection of the 26 sporadic groups, which are concrete mathematical objects that can be represented, for example, in terms of matrices of numbers. Even more importantly, many of the insights that led to the CFSG were reached only by constructing particular groups, which were tabulated (as parts of journal articles and lists that were passed around). We also see this in other Big Math projects, e.g., Langland's program of trying to relate Galois groups in algebraic number theory to automorphic forms and the representation theory of algebraic groups over local fields and adeles. This is supported by LMFDB, which contains about 80 tables with representations of mathematical objects ranging from elliptic curves to Hecke fields—almost a terabyte of data in all. These are used, for instance, to find patterns in the objects and their properties and to support or refute conjectures. Other well-known examples are OEIS, with over 300 000 integer sequences, and the Small Groups Library [47], with more than 450 million groups of small order. See [4] for a work-in-progress survey of math databases.

Unfortunately, such math databases are typically not integrated into systems for mathematical inference, narration, or knowledge organization and are only weakly integrated into systems for computation. Usually these databases supply a human-oriented web interface. If they also offer application programming interface (API) access to the underlying database, they do so only for database-level interaction, where an elliptic curve might be a quadruple of numbers encoded as decimal strings to work around size restrictions of the underlying database engine. What we would need instead for an integration in a Big Math system is an API that supplies access to the mathematical objects, such as the representations of elliptic curves as they are treated in organization, computation, inference, and narration systems; see [51] for a discussion.

As usual, there are exceptions. The GAP Small Groups Library system and LMFDB are notable examples: the former is deeply integrated into the GAP computer algebra system, and the latter includes almost a thousand narrative descriptions of the mathematical concepts underlying the tabulated objects.

Narration

Consider Figure 2, which shows an intermediate result in the OOT in the Coq inference system (foreground) and its corresponding narrative representation (background). Even though great care has been taken to make the Coq text human-readable, i.e., short and suggestive of traditional mathematical notation, there is still a significant language barrier to all but members of the OOT development team.

Indeed, mathematical tradition is completely different from its representation in inference, computation, and concretization systems. Knowledge and proofs are presented in documents like journal articles, preprints, monographs, and talks for human consumption. While rigor and correctness are important concerns, the main emphasis is on the efficient communication of ideas, insights, intuitions, and inherent connections to colleagues and students. As a consequence, more than half of the text of a typical mathematical document consists of introductions, motivations, recaps, remarks, outlooks, conclusions, and references. Even though this packaging of mathematical knowledge into documents leads to extensive overhead and duplication, it seems to be an efficient way of dealing with the OBB and thus a necessary cost for scholarly communication.

In current proof assistants like Coq, the narration aspect is undersupported even though tools like $\text{\LaTeX}$ have revolutionized mathematical writing. Source comments in the input files are possible in virtually all inference and computation systems, but they are not primary objects for the system and are thus used much less than in narrative representations. Knuth's literate programming idea [25] has yet to take root in mathematics, although it is worthwhile noting that one of the earliest inference systems, Automath [9], had extensive features for narration. The main modern exception among inference systems is Isabelle: it supports the inclusion of marked-up text and programs, and it turns the underlying mathematical library into a deeply integrated document management system, which allows recursive nesting of narrative, inference, and computation [49].

Figure 2. Informal and formal representations of mathematical knowledge.

Some computational communities such as parts of statistics, especially users of the software environment R, make use of some features of literate programming. Jupyter notebooks as used with SageMath as well as the document interfaces of Maple and Mathematica are also somewhat literate, although the simple fact that they do not interoperate smoothly with $\text{\LaTeX}$ hampers their adoption as methods of conveying large amounts of knowledge narratively.

In any case, inference and computation systems are notoriously bad for expressing the vague ideas and underspecified concepts that are characteristic of early phases in the development of mathematical theories and proofs, a task at which narration excels. Therefore, the Flyspeck project [22] used a $\text{\LaTeX}$ book [21] that reorganized the original proof to orchestrate and track the formal proof via cross-references to the identifiers of all formal definitions, lemmas, and theorems. Incidentally, the ongoing effort to establish a second-generation proof of the CFSG has a similar book, consisting of seven volumes already published and five additional volumes that will be published in the future.⁷

Organization

In the discussion and survey of the four corners of the tetrapod from Figure 1, we have seen that all four aspects share and are based on the representation of knowledge, which we call the mathematical ontology,⁸ and that they can interoperate effectively through this ontology. And we have seen that a modular, redundancy-minimizing organization of the ontology is crucial for getting a handle on the inherent complexity of mathematical knowledge.

Most inference and computation systems feature some kind of modularity features to organize their libraries. For inference systems, this was pioneered in the IMPS system [17] in the form of theories and theory interpretations and has been continued, for example, in the Isabelle system (type classes and locales). In systems like Coq and Lean [36] that feature dependent record types, theories and their morphisms can be encoded inside the base logic. Computation systems feature similar concepts. Finally, the MMT system [35, 43] systematically combines modular representation with a metalogical framework, in which the logics and logic-morphisms can be represented, yielding a foundation-independent (bring-your-own-logic) framework for mathematical knowledge representation that can be used to establish system interoperability.

Summation

Using the classification of finite simple groups as an example of Big Math, we diagnosed the one-brain barrier as a major impediment to large-scale developments and results in mathematics. We saw that the seemingly obvious

answer to this problem—employ computer support—is not without problems and can be a barrier itself. We proposed that computer-based mathematical assistants should have a tetrapodal structure, integrating inference, computation, concretization, and narration centered around a shared knowledge organization feature. We claim that only with special consideration of all these five aspects will mathematical software systems be able to render the support that is necessary for Big Math projects like the CFSG to go beyond the proof certification service rendered by big formal proofs like OOT and Flyspeck.

While the holistic conception of Big Math envisioned by our tetrapod is new, the general sentiment that mathematical assistant systems must escape their native corner is implicitly understood in the mathematical software community. Indeed, many of the systems we mentioned above, while focusing on a particular aspect and excelling at it, also integrate features of some other aspects. We have briefly surveyed current efforts toward such integrations. A thorough review of the state of the art, which would more clearly delineate the progress on the roadmap implicitly given by the tetrapod proposal, is beyond the scope of this position paper but is under active development by the authors for future publication; see [11] for a preliminary version.

We have observed the central place of the ontology in the proposed system functionality architecture, and we claim that such systems are best served by a global digital mathematical library (GDML) [37], which serves as a pivotal point for integrating systems and system functionalities. Again, a discussion of this important resource is beyond the scope of this paper, but see [16] for a motivation and [26] for an avenue on how it could be realized by marshaling existing resources like the libraries of proof assistants.

A GDML would constitute a critical resource for mathematics; it should thus be unsurprising to find organization at the center of the tetrapod. Arguably, the comprehensive ontology of mathematics that would constitute Big Organization cannot be created by a small set of individuals—it is both subject to and a way around the one-brain barrier—but has to be a collaborative effort of the whole community. A prerequisite for this is that the ontology be FAIR (Findable, Accessible, Interoperable, and Reusable) [52], open, and not encumbered by commercial or personal interests that are insurmountable for researchers in developing countries. This remains true even though revenue streams to fund the maintenance of such an ontology are necessary; thus, suitable licensing schemes and business models that reconcile the openness and funding considerations will have to be found.

Another trade-off to consider is that current commercial software development for mathematics has been most effective in regard to large-scale integration projects—the various offerings by Wolfram Inc. (Mathematica and its notebooks, Wolfram Alpha, and the Wolfram Language) together constitute one of the most tetrapodal systems in

⁷See the Wikipedia article on the classification of finite simple groups.

⁸Note that we will use the word “ontology” in the original wider sense of “a set of concepts and categories in a subject area or domain that shows their properties and the relations between them,” not just for specific technologies of the Semantic Web [48].

existence. Noncommercial funding for a similar open effort simply does not seem to exist at the moment. The downside of commercial software and accompanying resources is that it often cannot be adapted for research purposes without considerable legal and monetary effort. No matter how useful it is, closed-source projects ought to have the same impact as omitted proofs: greatly inspiring doubt.

Nevertheless, the authors believe that a confederated, well-organized community effort to develop open-source software and open resources can succeed. This is well supported by the current practices of the mathematical community, who by and large embrace open communication, open-source software, and open-access publication of results and resources. This is not anticommercialization, since there are several very successful companies thriving on open-source software; and even closed-source behemoths like Microsoft and Google have increasingly released part of their source code to the community.

We acknowledge the fact that our tetrapod proposal does not incorporate the fact that mathematics is a social process and that for Big Math problems we will need mathematical social machines, i.e., “an environment comprising humans and technology interacting and producing outputs or action which would not be possible without both parties present” [5]. We conjecture that the social machine aspect is one that will live quite comfortably on top of tetrapod-shaped mathematical software systems and can indeed not fully function without all of its five aspects interacting well; see [13, 30] for a discussion and further pointers to the literature.

Finally, we remark that of course the OBB is not particular to mathematics but affects all scientific and engineering disciplines, where we conjecture that tetrapodal paradigms similar to ours apply. Here, mathematics is a very good test case for the design of knowledge-based systems, since the knowledge structures and algorithms are so overt.

ACKNOWLEDGMENTS

The authors gratefully acknowledge fruitful discussions with many of our colleagues in the mechanized mathematics community. The work reported here was supported by the Horizon 2020 European Research Infrastructures project OpenDreamKit (#676541), the DFG project RA-18723-1 OAF, and the NSERC Grants RGPIN-2018-05812 and RGPIN-5100-2015.

FUNDING

Open Access funding provided by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit

line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Jacques Carette
Department of Computing and Software
McMaster University
1280 Main Street
West Hamilton, Ontario L8S 4L7
Canada
e-mail: carette@mcmaster.ca

William M. Farmer
Department of Computing and Software
McMaster University
1280 Main Street
West Hamilton, Ontario L8S 4L7
Canada
e-mail: wmfarmer@mcmaster.ca

Michael Kohlhase
Computer Science
University of Erlangen-Nuremberg
Martensstrasse 3, Room 11.139 91058 Erlangen
Germany
e-mail: michael.kohlhase@fau.de

Florian Rabe
Computer Science
University of Erlangen-Nuremberg
Martensstrasse 3, Room 11.137 91058 Erlangen
Germany
e-mail: florian.rabe@gmail.com

REFERENCES

- [1] AITP: Conference on artificial intelligence and theorem proving. Available at <http://aitp-conference.org/>.
- [2] Kenneth Appel and Wolfgang Haken. *Every Planar Map Is Four Colorable*, Contemporary Mathematics 98. American Mathematical Society, 1989.
- [3] Ron Ausbrooks, Stephen Buswell, David Carlisle, et al. Mathematical Markup Language (MathML) version 3.0. W3C Recommendation, World Wide Web Consortium (W3C), 2010.
- [4] Katja Berčič. Math databases wiki. Available at <https://github.com/MathHubInfo/Documentation/wiki/Math-Databases>.
- [5] Tim Berners-Lee and Mark Fischetti. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor*. Harper, 1999.
- [6] Steve Bernstein and Joseph Gelbart, editors. *An Introduction to the Langlands Program*. Birkhäuser, 2003.
- [7] Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development—Coq'Art: The Calculus of Inductive*

- Constructions*, Texts in Theoretical Computer Science. Springer, 2004.
- [8] Frederick Brooks. *The Mythical Man-Month: Essays on Software Engineering*. Addison-Wesley, 1975.
 - [9] Nicolaas Govert de Bruijn. The mathematical language AUTOMATH, its usage and some of its extensions. In *Symposium on Automatic Demonstration*, Lecture Notes in Mathematics 125, pp. 29–61. Springer, 1970.
 - [10] Nicolaas Govert de Bruijn. The mathematical vernacular, a language for mathematics with typed sets. In R. P. Nederpelt, J. H. Geuvers, and R. C. de Vrijer, editors, *Selected Papers on Automath, Studies in Logic and the Foundations of Mathematics* 133, pp. 865–935. Elsevier, 1994.
 - [11] Jacques Carette, William M. Farmer, Yasmine Sharoda, Katja Berčič, Michael Kohlhase, Dennis Müller, and Florian Rabe. The space of mathematical software systems—a survey of paradigmatic systems. Preprint available at <http://arxiv.org/abs/2002.04955>, 2020.
 - [12] R. M. Corless, J. H. Davenport, D. J. Jeffrey, and S. M. Watt. According to Abramowitz and Stegun. *SIGSAM Bulletin* 2 34 (2000), 58–65.
 - [13] Joseph Corneli, Ursula Martin, Dave Murray-Rust, Alison Pease, Raymond Puzio, and Gabriela Rino Nesin. Modelling the way mathematics is actually done. In *Proceedings of the 5th ACM SIGPLAN International Workshop on Functional Art, Music, Modeling, and Design*, FARM 2017, pp. 10–19. ACM, 2017.
 - [14] Coq Development Team. *The Coq Proof Assistant: Reference Manual*. HAL-INRIA. Available at <https://hal.inria.fr/inria-00069968>.
 - [15] John Cremona. The L -functions and modular forms database project. *Foundations of Computational Mathematics* 16:6 (2016), 1541–1553.
 - [16] William M. Farmer. MKM: A new interdisciplinary field of research. *Bulletin of the ACM Special Interest Group on Symbolic and Automated Mathematics (SIGSAM)* 38:2 (2004), 47–52.
 - [17] William M. Farmer, Joshua D. Guttman, and F. Javier Thayer. IMPS: An Interactive Mathematical Proof System. *Journal of Automated Reasoning* 11:2 (October 1993), 213–248.
 - [18] The GAP Group. GAP—Groups, Algorithms, and Programming. Available at <http://www.gap-system.org>, 2016.
 - [19] G. Gonthier, A. Asperti, J. Avigad, Y. Bertot, C. Cohen, F. Garillot, S. Le Roux, A. Mahboubi, R. O’Connor, et al. A machine-checked proof of the odd order theorem. In Sandrine Blazy, Christine Paulin-Mohring, and David Pichardie, editors, *Interactive Theorem Proving, Lecture Notes in Computer Science* 7998, pp. 163–179. Springer, 2013.
 - [20] Thomas C. Hales. A Proof of the Kepler Conjecture. *Annals of Mathematics, Second Series* 162:3 (2005), 1065–1185.
 - [21] Thomas Hales. *Dense Sphere Packings*, London Mathematical Society Lecture Note Series 400. Cambridge University Press, 2012.
 - [22] Thomas Hales, Mark Adams, Gertrud Bauer, et al. A formal proof of the Kepler conjecture. *Forum of Mathematics*, Pi 5, E2, <https://doi.org/10.1017/fmp.2017.1>, 2017.
 - [23] John Harrison. HOL Light: a tutorial introduction. In *Proceedings of the First International Conference on Formal Methods in Computer-Aided Design*, pp. 265–269. Springer, 1996.
 - [24] Idris language documentation. Available at <http://docs.idris-lang.org/en/latest/>.
 - [25] Donald E. Knuth. *Literate Programming*. University of Chicago Press, 1992.
 - [26] Michael Kohlhase and Florian Rabe. QED reloaded: towards a pluralistic formal library of mathematical knowledge. *Journal of Formalized Reasoning* 9:1 (2016), 201–234.
 - [27] The L -functions and modular forms database. Available at <http://www.lmfdb.org>.
 - [28] Assia Mahboubi and Enrico Tassi. Mathematical components. Available at <https://math-comp.github.io/mcb/book.pdf>.
 - [29] Maple: The essential tool for mathematics. Available at <https://www.maplesoft.com>.
 - [30] Ursula Martin. Computational logic and the social. *Journal of Logic and Computation* 26 (2016), 467–477.
 - [31] MathOverflow. Available at <https://mathoverflow.net>.
 - [32] MathSciNet Mathematical Reviews on the Net. American Mathematical Society. Available at <http://www.ams.org/mathscinet/>.
 - [33] Mathworld. Available at <http://mathworld.wolfram.com>.
 - [34] Mizar. Available at <http://xbpapp://www.mizar.org>.
 - [35] MMT: language and system for the uniform representation of knowledge. Project web site at <https://uniformal.github.io/>.
 - [36] Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. The lean theorem prover. In Amy P. Felty and Aart Middeldorp, editors, *Automated Deduction – CADE-25, Lecture Notes in Computer Science* 9195, pp. 378–388. Springer, 2015.
 - [37] National Research Council. *Developing a 21st Century Global Library for Mathematics Research*. The National Academies Press, 2014.
 - [38] Ulf Norell. The Agda Wiki. Available at <http://wiki.portal.chalmers.se/agda>.
 - [39] Ulf Norell. *Towards a Practical Programming Language Based on Dependent Type Theory*. PhD thesis, Chalmers University of Technology and Göteborg University, 2007.
 - [40] Openmath, 2009. Available at <http://wiki.openmath.org>.
 - [41] Frank Pfenning. Logical frameworks. In Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, vols. I and II. Elsevier Science and MIT Press, 2001.
 - [42] PlanetMath.org—math for the people, by the people. Available at <http://www.planetmath.org>.
 - [43] Florian Rabe and Michael Kohlhase. A scalable module system. *Information & Computation* 0.230 (2013), 1–54.
 - [44] Sagemath, the Sage Mathematics Software System. Available at <http://www.sagemath.org>.
 - [45] Ronan Saillard. *Typechecking in the $\lambda\Pi$ -Calculus Modulo: Theory and Practice*. PhD thesis, École nationale supérieure des mines de Paris, 2015.
 - [46] Neil J. A. Sloane. The on-line encyclopedia of integer sequences. *Notices of the AMS* 50:8 (2003), 912.
 - [47] The small groups library. Available at http://www.icm.tu-bs.de/ag_algebra/software/small/small.html.
 - [48] W3 consortium. webpage at <http://www.w3.org>. seen February 2007.
 - [49] Makarius Wenzel. Isabelle/jEdit as IDE for domain-specific formal languages and informal text documents. In Paolo Masci,

Rosemary Monahan, and Virgile Prevosto, editors, *F-IDE 2018 – 4th Workshop on Formal Integrated Development Environment*, 2018.

- [50] Makarius Wenzel, Lawrence C. Paulson, and Tobias Nipkow. The Isabelle framework. In Ait Mohamed, Munoz, and Tahar, editors, *Theorem Proving in Higher Order Logics (TPHOLs 2008)*, Lecture Notes in Computer Science 5170, pp. 33–38. Springer, 2008.
- [51] Tom Wiesing, Michael Kohlhase, and Florian Rabe. Virtual theories—a uniform interface to mathematical knowledge bases. In Johannes Blömer, Temur Kutsia, and Dimitris Simos, editors, *MACIS 2017: Seventh International Conference on Mathematical Aspects of Computer and Information Sciences*, Lecture Notes in Computer Science 10693, pp. 243–257. Springer, 2017.

[52] Mark D. Wilkinson, Michel Dumontier, IJsbrand Aalbersberg, et al. The fair guiding principles for scientific data management and stewardship. *Scientific Data*, 3 (2016), 160018.

[53] Wolfram Mathematica. Available at <https://www.wolfram.com/mathematica/>.

[54] zbMATH the first resource in mathematics. Available at <http://zbmath.org>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.