The Space of Mathematical Software Systems — A Survey of Paradigmatic Systems

Katja Berčič[1], Jacques Carette[2], William M. Farmer[2], Michael Kohlhase[1],
Dennis Müller[1], Florian Rabe[1], and Yasmine Sharoda[2]

[1] Computer Science, FAU Erlangen-Nürnberg, `http://kwarc.info`

[2] Computing and Software, McMaster University,
`http://www.cas.mcmaster.ca/research/mathscheme/`

February 13, 2020

Abstract

Mathematical software systems are becoming more and more important in pure and applied mathematics in order to deal with the complexity and scalability issues inherent in mathematics. In the last decades we have seen a cambric explosion of increasingly powerful but also diverging systems.

To give researchers a guide to this space of systems, we devise a novel conceptualization of mathematical software that focuses on five aspects: *inference* covers formal logic and reasoning about mathematical statements via proofs and models, typically with strong emphasis on correctness; *computation* covers algorithms and software libraries for representing and manipulating mathematical objects, typically with strong emphasis on efficiency; *concretization* covers generating and maintaining collections of mathematical objects conforming to a certain pattern, typically with strong emphasis on complete enumeration; *narration* covers describing mathematical contexts and relations, typically with strong emphasis on human readability; finally, *organization* covers representing mathematical contexts and objects in machine-actionable formal languages, typically with strong emphasis on expressivity and system interoperability.

Despite broad agreement that an ideal system would seamlessly integrate all these aspects, research has diversified into families of highly specialized systems focusing on a single aspect and possibly partially integrating others, each with their own communities, challenges, and successes. In this survey, we focus on the commonalities and differences of these systems from the perspective of a future multi-aspect system.

Our goal is to give new researchers, existing researchers from each of these communities, or outsiders like mathematicians a basic overview that enables them to match practical challenges to existing solutions, identify white spots in the software space, and to deepen the integration between systems and paradigms.

CAVEAT: This paper is intended as a living survey that is updated in-place on `http://arXiv.org` from time to time. We publish this early as a pre-preprint to let the community discuss and maybe provide feedback to the authors at `tetrapod@lists.informatik.uni-erlangen.de`.

Contents

1 Introduction

In the last half decade we have seen mathematics tackle problems that lead to increasingly large developments: proofs, computations, data sets, and document collections. This trend has led to intense discussions about the nature of mathematics, ventilating questions like:

i) Is a proof that can only be verified with the help of a computer still a mathematical proof?

ii) Is a mathematical proofscape that exceeds what can be understood in detail by a single expert a legitimate justification of a mathematical result?

iii) Can a collection of mathematics papers — however big — adequately represent a large body of mathematical knowledge?

In [Car+20] we have discussed these questions under the heading of **Big Math** and propose a unified, high-level model. We claim that computer support will be necessary for scaling mathematics, and that suitable and acceptable methods should be developed in a tight collaboration between mathematicians and computer scientists — indeed such method development is already under way, but needs to become more comprehensive and integrative.

We propose that all Big Math developments comprise four main aspects that need to be dealt with at scale:

i) *Inference*: deriving statements by *deduction* (i.e., proving), *abduction* (i.e., conjecture formation from best explanations), and *induction* (i.e., conjecture formation from examples).

ii) *Computation*: algorithmic manipulation and simplification of mathematical expressions and other representations of mathematical objects.

iii) *Concretization*: generating, collecting, maintaining, and accessing collections of examples that suggest patterns and relations and allow testing of conjectures.

iv) *Narration*: bringing the results into a form that can be digested by humans, usually in mathematical documents like articles, books, or preprints, that expose the ideas in natural language but also in diagrams, tables, and simulations.

These aspects — their existence and importance to mathematics — should be rather uncontroversial. Figure 1 may help convey the part which is less discussed, and not less crucial: that they are tightly related. For a convenient representation in three dimensions, we choose to locate the organization aspect at the barycentre of the other four since they are all consumers and producers of mathematical knowledge.
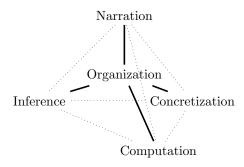

Figure 1: Five Aspects of Big Math Systems, a Tetrapod Structure

Computer support exists for all of these four aspects of Big Math, e.g.,

i) theorem provers like Isabelle, Coq, or Mizar;

ii) computer algebra systems like GAP, SageMath, Maple, or Mathematica; and

iii) mathematical data bases like the L-functions and Modular Forms Data Base (LMFDB) [Cre16; LM] and the Online Encyclopedia of Integer Sequences (OEIS) [Slo03];

iv) online journals, mathematical information systems like zbMATH or MathSciNet, preprint servers like arXiv.org, or research-level help systems like MathOverflow.

While humans can easily integrate these four aspects and do that for all mathematical developments (large or otherwise), much research is still necessary into how such an integration can be

achieved in software systems. [1]

<div align="right">EdN:1</div>

Overview We want to throw the spotlight on the integration problem to help start off research and development of systems that integrate all four aspects. To facilitate this, we give a high-level survey of mathematical software systems from the Tetrapod perspective. Because almost every one of these systems has one primary aspect and because systems with the same primary aspect are often very similar, we group systems by their primary aspects.

In each group we try to further subdivide the systems. We want to stress that this classification is mostly meant for convenience, e.g., by simplifying the description of several similar systems. We do not mean to imply a strict separation between these groups of systems, and often the borders are fluid. In particular, there are some systems that already allow aspect switching in a way that precludes ascribing a primary aspect. For convenience, we still assign these systems into one of the groups and discuss which other aspects they support.[2]

<div align="right">EdN:2</div>

While most systems use one of the four aspects as the primary one, there are also some systems that focus primarily or even exclusively on the ontology. Therefore, use a fifth group[3] for those systems.

<div align="right">EdN:3</div>

In survey in the next sections we use the following general terms for all aspects: *Syntax* is a set of rules for forming objects, and *data* is any piece of well-formed syntax. *Semantics* is a set of rules for interpreting objects, and *knowledge* is a pair of a datum and its semantics.

2   Primary Aspect: Inference

Various methods have been developed to represent and perform inferences. We structure our presentation by how each method relates to computation, the aspect most whose integration with inference has drawn the most attention. In general, the ubiquity of underspecified function symbols and quantified variables means that logical expressions usually do not normalize to unique values. At best, computations like $y := f(x)$ can be represented as open-ended conjectures where different options for y are produced, each together with a proof of the respective equality. Therefore, inference systems usually sacrifice computation or at least its efficiency.

Proof assistants sit at the extreme end of this spectrum. They employ strong logics and high-level declarations to provide a convenient way to formalize domain knowledge and reason about it. The reasoning is usually interactive in order to represent inferences that are too difficult to be fully automated. Most proof assistants integrate at least some of the other methods to overcome this weakness.

Further along the spectrum, *automated theorem provers* use simpler logics than interactive proof assistants. They are fully automatic and much faster, but can handle much fewer theorems, and typically do not check their proofs. *Satisfiability checkers* continue this progression by aiming at decidable automation support, whereas theorem proving is usually an semi-decidable search problem. That limits them to propositional logic or specific theories of more expressive logics (usually of first-order logic) that are complete, i.e., where every formula can be proved or

---

[1]EDNOTE: MK: Possible outline:

1. The mathematics process involves several integrated activities.

2. Contemporary mathematical software systems usually focus on just one of these activities. Moreover, these systems are not designed to work with each other and do not employ a common knowledge base.

3. The limitations of contemporary systems have been exposed by the tremendous growth in the production of mathematical knowledge and "big math" projects like the Classification of Finite Simple Groups.

4. We need a mathematical software system that is holistic in the sense that it is designed to support the entire mathematics process and to be integrated with existing systems.

5. The first step towards the goal of a holistic system is to survey the mathematical software systems that are available today, which is subject of this paper.

[2]EDNOTE: MK: I do not think we should do it that way, but mention them in both
[3]EDNOTE: MK: this is currently the first one discussed, we either need to announce that or move it last.

disproved. In the special cases, where satisfiability checkers are applicable, they come close to verified computation systems.

Orthogonal to the above triplet, there are several methods for realizing Turing-complete computation naturally inside a logic. Here imperative and object-oriented computation are usually avoided in favor of other programming paradigms that are easier to reason about. *Rewriting* aims at optimizing the $f(x) \rightsquigarrow y$ progression, allowing users to mark specific transformations as rewrite steps. *Terminating recursion* is the method of adding recursive functions to a logic in order to make it a pure functional programming language. Finally, *logic programming* restricts attention to theorems of a special form, for which proof search is simple and predictable so that users can represent computations by supplying axioms that guide the proof search.

In the sequel, we describe each method in some more detail.

Proof Assistants (C) The most successful proof assistants represent tens or hundreds of person-year investments into

- Define the foundational logic. Usually logics much stronger than textbook first- or higher-order logic are needed in practice. A big question has been the trade-off between flexible, untyped languages that rely on undecidable reasoning and rich type system that are more restrictive but have better computational properties. Typical choices are first-order set theory (e.g., Mizar [Miza]), higher-order logic (e.g., HOL [HOL], HOL Light [Har96], Isabelle/HOL [Pau94], PVS [ORS92]), and constructive type theory (e.g., Coq [Coq15], Matita [Asp+06], Lean [de +15], Agda [Nor05]). Some systems use undecidable type systems (e.g., Mizar, PVS) as compromises.
- Implement the logic. Usually the implementation starts with a kernel that checks proofs and then grows outwards in layers until a human-friendly surface syntax is exposed. Much work has been put into automatically filling in as many gaps left by the user as possible. Tactic languages (e.g., HOL, HOL Light, Coq, Isabelle) and high-level proof languages (e.g., Isabelle, Coq) and the integration of automated provers (e.g., Isabelle, Coq) and decision procedures (e.g., PVS, Isabelle) have been crucial here.
- Build a library of data structures. Usually proof assistants are only valuable if their standard library provides many data basic structures of mathematics (e.g., groups, real function) and computer science (e.g., records, inductive types). This wa most prominently envisioned in the QED manifesto [Qed]. But representing these and proving their characteristic properties has proved very expensive and remains a tough benchmark for the design of logic and system.

The tactic languages are often Turing-complete themselves (e.g., HOL Light, Isabelle, Coq). This is realized by writing tactics in the underlying programming language. Recently systems have tried to represent tactics in the systems itself either declaratively (e.g., Coq, Matita) or programmatically through reflection of the kernel data structures (e.g., Lean).

Most major proof assistants allow interspersed narrative structure, at the very least through comments. Some systems mimic sectioning where the scope of variables is determined not by the logic but by the narrative structure (e.g., Coq). Some systems (e.g., Isabelle, Agda) are narratively strong enough to make it feasible to write narrative documents (most importantly the documentation of the system itself) in the system.

[4]

Automated Theorem Provers (A) Automated provers have been mostly developed for relatively simpler logics, where full automation is feasible. Therefore, they are often used as backend system integrated into, e.g., interactive proof assistants.

Most systems work with variants of first-order logic and compete regularly in the CASC competition. Examples are Vampire [RV02], E [Sch01], and Spass [Wei+02]. An ongoing trend is the gradual extension of first-order logic with additional features such as definitions, primitive numbers, types, and polymorphism.

---

[4]EDNOTE: JC: recent paper "A Survey on Theorem Provers in Formal Methods" - https://arxiv.org/abs/1912.03028 has just appeared. We should at the very least cite all the software that appears in it.

Automated provers for higher-order logics are much harder to develop but are gaining strength. Example systems are Leo [Ben+08] and Satallax [Bro12].

Recently efforts have been made for automated provers to return checkable proofs in order to use them in proof assistants (see [Bla+16] for an overview). A big problem here is the selection of useful axioms to reduce the search space for the automated prover. Recently machine learning has been employed successfully for this purpose (e.g., [KU15]).

Satisfiability Checkers (S)   While satisfiability is decidable for propositional logic and a few first-order theories (most importantly variants of linear arithmetic), the overall scope is limited. Decision problems are typically stated in the form of satisfiability problems and called *SAT solvers* (for propositional logic) or *SMT solvers* (satisfiability modulo theory, for specific theories of, typically, first-order logic). Specific decision procedure-based systems for higher logics include Z3 [dB08] and CVC [**cvc**]. [5]   EdN:5

Satisfiability checkers for propositional logic have become so powerful that it is handle feasible to encode high-level problems as propositional problems by using large amounts of propositional variables.

To increase scope, efforts are made at combining decision procedures or applying them to theories with decidable fragments. While these efforts lose decidability, they may still be highly valuable in practice, e.g., by directly computing a value instead of extracting it from a proof, or to reduce the search space by eliminating unsatisfiable branches.

Terminating Recursions (T)   Many proof assistants include sound and incomplete termination checkers (e.g., PVS, Coq, Agda) that only accept provably terminating functions. These are Turing-complete in the sense that the syntax can represent all computable functions but the termination checker will not accept all of them. Some systems (e.g., Isabelle, Coq) can export those programs in external programming languages after verifying their correctness in the logics.

In some cases it can be hard to draw the line between inference and computation systems. This is the case for systems that use functional programming combined with the representation of proofs-as-programs (e.g., Coq, Agda).

Rewriting (R)   Rewriting implements a directed equality relation between expressions. This can be used both to obtain Turing-complete computation and for reasoning (by rewriting formulas to a boolean value).

Compared to other forms of computation, rewriting is very inefficient, e.g., rewriting polynomials into normal form takes exponentially longer than an algorithm based on plain "arithmetic". But the embedding of rewrite systems in logic permits proving the soundness of each rewrite rule.

Originally most rewrite systems were based on first-order logic due to its decidable unification (e.g., Maude [Cla+96]).[6]   But recently more systems are employing rewriting in higher-order   EdN:6
setting (e.g., Dedukti [BCH12]). [7]   EdN:7

A lot of effort has gone into establishing the confluence and termination of sets of rewrite rules. This is critical to establish that rewriting implements a deterministic computation. Therefore, many systems try to establish confluence and termination of sets of rewrite rules automatically. [8]   EdN:8

Many proof assistants integrate rewrite engines that use some of the proved theorems as rewrite rules that are applied automatically by the system (e.g., PVS, Isabelle).

Logic Programming (L)   If the proof search behavior of a prover is known to the user and predictable in practice, this can be instrumented for computation. This works particularly well with axioms in Horn-form where searching a proof of the conclusion triggers searching proofs of the assumptions. This is the basic idea of logic programming, where a formalization consists of a

---

[5]EDNOTE: FR: cite some SAT and SMT solvers here
[6]EDNOTE: add more citations
[7]EDNOTE: check Cynthia Kop's talk on higher-order rewriting for more citations
[8]EDNOTE: cite examples

set of Horn axioms, which can be seen as both a specification and (via predicatble proof search) as a program. If special predicate symbols are added, whose "proofs" consists of extra-logical operations like I/O side effects, this yields a general purpose programming language.

Logic programming was mostly investigated in untyped first-order logic via various Prolog dialects [CM81]. But higher-order variants exists as well, e.g., such as λProlog [Mil]. [9]

[10]

3 Primary Aspect: Computation

[11] [12]

Computation is a rather broad topic: for example, term rewriting systems and finite state machines are often regarded as performing computations. At the Turing-complete end of the spectrum, we could list all programming languages as performing computations — because, well, they do!

As our primary focus is still centered on mathematics, this helpfully narrows things down. Even though one can indeed use just about any language to do mathematics, it makes sense to instead focus on those languages that have been designed with mathematics in mind.

The systems can be usefully divided according to the kinds of data they were primarily designed to handle:

- Typed

- Symbolic

- Algebraic

- Numeric

  - analytic
  - statistical

By *typed* data, we mean data that can be expressed in "type theory", as originating from Russell and Whitehead [WR10], through Church [Chu40], until today; the historical development until 10 years ago is well documented in [KLN04]. Although we have seen a wide varieties of type theories used for systems that perform inference, for systems that take computations seriously, one family emerges: dependent type theories. By *symbolic* data, we mean data that can be best expressed as *abstract syntax trees*, usually containing "free variables". From a type-theoretical point of view, these are significantly harder to deal with, as these would then form *open terms*, which are notoriously difficult to manipulate correctly. This is why all know symbolic computation systems are untyped. By *algebraic*, we basically mean data that belongs to the mathematical sub-domain of Algebra. What distinguishes these is that, although it is frequently convenient to use "free variables" for the visual display of these objects, they are not fundamentally required for an adequate representation. Under Algebra, we also include systems that do exact computations on natural numbers, integers, Gaussian integers, etc, as these are also algebraic. By *numeric*, we mean systems whose data include "real numbers" in one way or another; it is useful to subdivide this class further, into the systems that specialize in more analytic problems (quadrature, differential equations) from those that deal with statistics. Both kinds excell at computations based on linear algebra.

Some of the systems that we survey below are quite broad, and so implement many features in common: one can indeed do statistics in Maple, and symbolic computation in Matlab. The classification is not meant to "squeeze" any system into a narrow box, but rather to express the

---

[9]EdNote: need more references here
[10]EdNote: Responsible: Florian
[11]EdNote: Responsible: Jacques
[12]EdNote: MK: We need letters for the big table here. I will use "T" for Turing Complete; need that for TEX

7

fundamental organizational system around which the system grew outward to encompass much more.

Typed   Agda [Nor05], Idris [Bra13]
   ATS.

Symbolic   Mathematica
   Maple
   Axiom

Exact Computation Systems   Gap: groups
   Singular: ideals
   Sage: integrated with Python and (via Python) other languages

Scientific Computation   Matlab, scilab, Octave.
   ChebFun, NumPY,

Statistics Packages   R. SPSS, SAS, Minitab.

More categories   Machine learning, probabilistic programming. Term rewriting?

commentary   [13] By *reflection*, we mean the ability of a logic to have access to (some of) the reasoning facilities usually associated with its meta-logic. Computational reflection is similar, in the context of programming languages; this usually proceeds via a representation of some (otherwise opaque) concepts, which can be manipulated, before being *reified*. If the representation is adequate and the manipulations are meaning preserving, then computation can implement reasoning. EdN:13

If programming languages use a sufficiently strong *type system*, they may be able to embed deduction into computation. Types are, via Curry-Howard, (simple) propositions that are considered true if they are inhabited. Then type checking is a form of verification (i.e., checking that a particular proof/term inhabits a particular proposition/type). *Abstract interpretation* moves beyond "simple" types into being able to attach significantly more powerful properties to programs.

4   Primary Aspect: Concretization

[14] EdN:14

Overview   The naming of this aspect of knowledge section has proved surprisingly difficult. We mean to include any practice of representing mathematical objects in terms of concrete data structures. Mathematics has a long tradition of such efforts, going back to, e.g., clay tablets of Pythagorean triples [Abd10], lists of decimal digits of π, or logarithm tables. More modern incarnations include computer-supported practices like large prime numbers, the database of integer sequences, and the enumeration of isomorphism classes of simple finite groups. But this practice does not a standard name.

By *concrete data structures*, we mean any data formed using primitive objects such as integers and strings via constructors like lists, records, and tables. These objects have in common that they have an objective physical reality that is beyond doubt: for example, any particular finite list of integers exists absolutely, whereas the existence of, e.g., proofs or programs may be relative to philosophical assumptions (e.g., impredicativity, classical reasoning, axiom of choice)

---

[13]EDNOTE: while I understand why this was put here, I think this should just be deleted. In many ways, this is anti-tetrapodian thinking, as it priviledges inference over everything else.
   [14]EDNOTE: responsible: use summary of Katja's survey

or mathematical conditions (e.g., soundness of an argument, termination of an algorithm). Thus, our concrete objects are tangible and material in a way that is opposite to the Platonic objects that deduction and narration and to some extent computation are concerned with. Thus, we can also think of them as the shadows of Platonic philosophy or as Aristotelian objects in the sense of being empirical, observable, and practical.

Most of the time, concrete objects are aggregated in the form of tables such as logarithm tables or the many large relational databases described below. We considered *tabulation* as an alternative name for this aspect but opted against it to avoid excluding other representation languages for concrete data such as arrays and JSON.

Encodings   There are two ways to define the semantics of concrete objects. Firstly, we can characterize the concrete objects as the subset of those mathematical objects that are self-denoting: e.g., any natural number is interpreted as itself. This is in contrast to the other three aspects, where interpretations are needed to map objects to their denotation. Secondly, many non-concrete mathematical objects can be represented in terms of concrete ones, a process that we call *encoding*. Encodings are commonly used in databases of algebraic structures such as elliptic curves or isomorphism classes of graphs. Seen mathematically, any encoding is based on a representation theorem, which states the encoding $e(o)$ of an object o fully characterizes o (up to isomorphism ideally). However, such representation theorems do not always exist because sets and functions, which are the foundation of most mathematics, are inherently hard to represent concretely.

Conceptually, any effective representation of mathematical objects requires some encoding as concrete objects because only those can be acted on by computers. Therefore, in general, the semantics of concrete objects consists in applying the dual *decoding* operation. We speak of *codecs* for a pair of encoding or decoding that represent a class of mathematical objects as concrete ones.

Data and Knowledge   Until the advent of computer-supported mathematics, the creation and sharing of data only received attention in passing. The logarithm tables are a good example of this, as are the examples collected in the community effort initiated by Gordon Royle on MathOverflow [MOc].

Billey and Tenner introduced the concept of a fingerprint database of theorems in 2013 [BT13]. The primary example of this is the OEIS. They stress the importance of the following aspects: searchability, collaborativeness, citability of the contents, and indexing by small, language independent and canonical data.

Mathematical datasets and databases today are highly varied and range from small to large in several aspects. The datasets can easily reach the Gigabyte range: LMFDB (1TB data in number theory), or a lattice dataset by Kohonen (uncompressed about 1.5TB of lattices), to name some of the largest. The can range from short lists of objects that are extremely hard to compute, to gigantic lists of millions of objects, such as the GAP Small Groups Library (about 450 million finite groups) and the previously mentioned dataset of $17 \cdot 10^9$ lattices. Similarly, the authorship varies from single-author datasets, to community efforts such as the OEIS, with thousands of contributors. The structure of mathematical datasets can be as simple as having a list of objects, a list of records that can contain information like mathematical invariants in addition to the object, to complex databases of related tables such as the LMFDB.

Commonly occurring themes in mathematical data are ad-hoc implementations of codecs, lack of community guidelines for data, and similar.[15]   EdN:15

We will use the following facets for classifying systems in Table **??**: We start out with three classes of representations of concrete objects:

Record Data (R)   where datasets are sets of records conforming to the same schema. Record data and querying is very well-standardized by the relational (SQL) model. However, if encodings are used, SQL can never answer queries about the semantics of the original object.

---

[15]EDNOTE: Mention FAIRMat?

Array Data (A)   consists of very large, multidimensional arrays that require optimized management. Array data tends to come up in settings with large but simply-structured datasets such as simulation time series, while record data is often needed to represent complex objects, especially those from pure mathematics. Array data bases, which offer efficient access to contiguous — possibly lower-dimensional — sub-arrays of datasets (voxels), are less standardized, but OPenNDAP [ODAP] is becoming increasingly recognized even outside the GeoData community, where it originated.

Linked Data (L)   introduces identifiers for objects and then treats them as blackboxes, only representing the identi- fier and not the original object. The internal structure and the semantics of the object remain unspecified except for maintaining a set of named relations and attributions for these identifiers. The named relations allow forming large networks of objects, and the attributions of concrete values provide limited information about each one. Linked data can be subdivided into knowledge graphs and metadata, e.g., as used in publication indexing services.

Isomorphism Classes (I)   In many cases, all mathematical properties of an class of objectsare invariant under the natural isomorphisms of the class. In this case, we are less interested in the objects themselves, but only in their isomorphism classes. Groups, graphs, and elliptic curves are prime examples of this. Some math data systems have sustematic representing isomoprhismclasses, which can be considerably more difficult to represent than their representatives. To achieve this e.g. for graphs we need to generate a canonical labeling for the encoding, i.e. where two graphs are isomorphic, whenever the canonical labeling is identical. And then we need to port all the persistence layer, the algorithms, and the UI to the enhanced graph encoding.

Redundant Information/Computed Properties (P)   Often, a small set of properties of an object suffice to characterize if fully (up to isomorphism); we will call such properties constitutive. Some concretization systems[16] support the storage of "redundant information", i.e. information about of the objects that can in principle be computed from the constitutive properties, but may be too costly, or might be needed as keys for object selection queries.

EdN:16

Complete Enumeration (E)   Contrary to most other forms of represented mathematical knowledge, concrete mathematical data provides the chance to completely enumerate a set (of isomoprhism classes) of objects. E.g. all groups up to a given order. In all other collections of represented mathematical knowledge – think theorem prover libraries, preprint collections, or computer algebra systems – objects are curated because they are "interesting" because they have special properties.

Named Objects/Symbolic Terms (S)   [17]

EdN:17

Integration with Computation System (C)   Most mathematical software systems that store and manage represented mathematical knowledge are based on a conventional persistence layer, i.e. files with object encodings, a relational, array, or graph data base together with a user interface. The knowledge itself is usually created by users utilizing the system itself – e.g. proofs created in the system, papers written in LaTeX, or programs developed developed in the integrated IDE – and are curated by a community or a commercial entity. Alternatively – especially in the cases of complete enumeration (see above) – the content can be created programmatically by a computational system. In this case, there is usually a tight integration with a dedicated computational system.

BOP:18

Standalone Databases   OEIS [OEIS]: some ontology and narration

LMFDB [LM]: content produced by computation; filled by computation; computation integrated into frontend via Sage; narration and ontology for background knowledge

FindStat [BSa14]:

ATLAS of Finite Group Representations (`http://for.mat.bham.ac.uk/atlas/`):

Database of Ring Theory (`http://ringtheory.herokuapp.com/`):

Math Counterexamples (`http://www.mathcounterexamples.net/`):

Manifold Atlas: (`http://www.map.mpim-bonn.mpg.de`)

Distributome (`http://www.distributome.org/`)

SuiteSparse (`https://sparse.tamu.edu/`)

House of Graphs (`http://hog.grinvin.org`)

Digital Library of Mathematical Functions (`https://dlmf.nist.gov/`)

Data Sets within Computation Systems   Mathematics (embedded and external data sources)

Sage

GAP data libraries (`http://www.gap-system.org/Datalib/datalib.html`): (e.g., table of transitive groups)

5   Primary Aspect: Narration

[19] The "narration" aspect of the tetrapod is concerned with mathematical knowledge in a form that can be digested by humans. Narratively represented mathematical knowledge usually exists in mathematical documents[1] like articles, books, or preprints, that expose the ideas in natural language but also in diagrams, tables, and simulations. While rigour and correctness are important concerns in narration, the main emphasis is on communicating ideas, insights, intuitions, and inherent connections efficiently to colleagues well-versed in the particular topic or students who want to become that. 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 the "packaging" of mathematical knowledge into documents leads to some duplication in the mathematical literature, it seems to be an efficient way of dealing with communication and knowledge preservation and can thus be seen as a necessary overhead in scholarly communication.

The primary aim of this survey is to explore holistic mathematical software systems – here software support for mathematical documents. This in turn depends on the depth of explicit structural and semantic markup in documents. Currently, there are four (plus one) levels of representation of mathematical documents:

RL0. **written up**: for communication (e.g. chalk on blackboards) or archival purposes – e.g. on papyrus scrolls $\mathrel{\widehat{=}} {\sim}90\%$[2] of the mathematical documents.

RL1. **digital** usually digitized from print – e.g. as TIFFs $\mathrel{\widehat{=}} {\sim}50\%$

RL2. **presentational**: encoded text interspersed with presentation markup – e.g. PDF, Word, TEX or presentation MathML $\mathrel{\widehat{=}} {\sim}20\%$

RL3. **semantic**: encoded text with functional markup for the meaning, e.g. LATEX, STEX, Mathematica Notebooks $\mathrel{\widehat{=}} {\sim}1\%$

RL4. **formal**: The meaning of the document is fully specified and thus machine-actionable at all levels. $\leq 0.1\%$

We remark that the delineation of levels is somewhat fuzzy and that the "levels" themselves are far from uniform. Nevertheless they constitute a useful categorization for our discussion. Especially

---

[19]EDNOTE: responsible: Michael

[1]While we take a rather inclusive view on "documents", we limit ourselves to written documents, excluding audio recordings and videos. Our rationale is that this does not constitute a loss of generality since the latter could be transcribed into written documents without significant loss in meaning.

[2]The percentages in this classification are rough estimates that should be taken as qualitative indications of the relative size rather than actual quantities.

in the higher levels, the presentational and semantic markup is usually restricted to particular aspects of the document functionality. We have ordered the examples of representation formats by increasing opportunities for structural and functional markup.

Note transforming from higher levels to lower levels is usually simple via largely context-free **styling** rules, whereas the opposite direction – **semantics extraction** – involves non-trivial context-dependent heuristic choices. For instance in the transformation between the levels RL1. and RL2. we have the difference between printing (down transformation) or OCR (optical character recognition; up transformation).

We finally note that computation and thus machine support needs explicitly represented structures and thus higher representation levels lead to more opportunities for software support. Therefore we will structure our survey bottom-up in the hierarchy above, starting with level RL2. since levels RL0. and RL1. have no discernable math-specific aspects.

At the Presentational Level (P): Word Processors & Document Preparation Systems
At this level, we have any kind of software system and for document preparation that can deal with **mathematical vernacular**, the peculiar mixture of natural language, mathematical formulae, and diagrams digitally. In contrast to the image-based formats at level RL1., text is encoded as sequences of characters, whereas formulae and diagrams are in some kind of presentational markup. This is sufficient to e.g. make documents at the presentation level searchable in conventional bag-of-words-based search engines like Google or Bing. The systems and representation formats mainly differ in their

1. **authoring model**: WYSIWYG or formatted/programmable text,
2. **target media**: paginated or flexible page size,
3. treatment of mathematical formulae.

Word processors like MS Word or LibreOffice Writer implement a WYSIWYG – "what you see is what you get" – authoring model and target paginated media. Typesetting systems like TeX/LaTeX, the document preparation system preferred in mathematics, let authors encode documents as Unicode strings with executable – often user-definable – control sequences, which a formatter expands into a primitive page description format. For publication, the standard target page description format in both cases is usually PDF (Adobe's standardized Portable Document Format). This fixes page layouts down to the character position. Diagrams and formula components that do not come from one of the available fonts are represented as vector graphics. In particular, mathematical formulae lose all structural information during the PDF transformation, so that higher-level services like mathematical search or screen readers have nothing to go on, and would OCR-like facilities to function. In essence, mathematical formulae and diagrams are still at level RL1 (image-like) in PDF, even though the "source" (Office Open XML [OOXML06] for MS Word and Open Document Format [DB12] for LibreOffice Writer) may still have had the necessary structures.

Most document preparation systems also allow the export of HTML5 [Fau+17], a web markup format for interactive multimedia documents, that can encode mathematical formulae via MathML (the Mathematical Markup Language; see [MML310]) and diagrams via SVG (Scalable Vector Graphics; see [Dah+11]). For MS Word, MathML export is a native feature, for LibreOffice Writer via a plugin Writer2xhtml [W2X], TeX/LaTeX can be exported to HTML5 (and the eBook format ePUB3 [Con+11] based on it) via the LaTeXML engine or TeX4HT. Note that MathML has two sub-languages: **presentation MathML** specifies the visual layout of formulae (i.e. level RL2), and **content MathML** for the meaning (the associated operator trees; i.e. RL3). The exports above all restrict themselves to presentation MathML (though LaTeXML [LTX] does a best-effort attempt at inferring content MathML). But even presentation MathML has enough structure to support formula screen readers like MathPlayer [Mat] or mathematical search engines; see [GSC15; Aiz+16] for pointers to the state of the art and [ZBF] for an online example.

Formal, Narrative Documents (F)   To understand the semantic representation level (RL3) of narrative mathematical documents, which mixes language and with formal aspects, it is good

to think about fully formal narrative documents. By definition, these could consist of a sequence of logical propositions[3], possibly extended by a formalization of sectioning, discourse, and rhetorical structures. To the best of the knowledge of the authors, there are no fully formal, narrative representations of mathematical knowledge. All fully formal representations (e.g. proof assistant libraries cf. Section 2) are organized with respect to the inherent structures of the knowledge space, with little concern to a given narrative for human readers. Where attempts of a narrative are made – e.g. in comments – these are informal natural language, not logic. A notable example is van Benthem Jutting' s formalization [BJ77] of Landau's "*Grundlagen der Analysis* [Lan30], where the formalization closely follows the structure of the original. Arguably the narrative structure of the narrative structure of the Grundlagen is very limited, and the formalization excluded the explanatory introduction. The NaProChe [Cra+10; CKS11] project develops a controlled natural language for mathematics, i.e. a formal language that is – syntactically – a subset of mathematical vernacular, with the aim of verifying mathematics in the Isabelle proof assistant. Again, introductions, motivations, recaps, remarks, outlooks, conclusions, and references are not part of the language.

A similar but dual example – verbalization instead of formalization – is the case of Mizar articles, from which human-oriented presentations for the Journal of Formalized Mathematics [JFM] are generated. Again, abstracts and introductions have to supplied by (human) authors. The remaining content is generated from the Mizar theorems and proofs. The latter are sequences of statements and justifications which can be verified by the Mizar prover. In fact, formal proofs given as proof step sequences may be the only fully formal narrative documents currently available. In contrast to proof objects – λ terms in an expressive type theory – they combine full formality with a narrative (step-by-step with justifications) structure conducive to human understanding. Generally, proof presentation – i.e. transforming proofs to mathematical vernacular – has been studied in various contexts; see [ABR01; Hua96; Hor00] for details and pointers.

A particularly influential design has been the ISAR[4] (Intelligible semi-automated reasoning) proof document language [Wen07] of the Isabelle proof assistant. ISAR tries to bridge the semantic gap between prover-internal – proof objects generated by tactic scripts – and an appropriate level of abstraction for user-level work. ISAR proof texts consist document constructors, atopic steps via high-level tactic invocations, and library references. Thus they admit a purely static reading, thus being intelligible later without requiring dynamic replay that is so typical for traditional proof scripts. This is a very important characteristic of narrative representations.

Semi-Formal Systems (S): Documents at the Semantic Level   The semantic level relaxes the requirement of full formality and allows informal elements interspersed with formal ones. For instance, Isabelle provides native syntax for LaTeX-like commands as well as raw LaTeX, and all Isabelle documents can be turned into LaTeX for documentation. Similarly, Agda [Nor05] can read two kinds of files: documentation files with interspersed Agda code or Agda code files with interspersed documentation.

Thus semiformal systems can choose which aspects to formalize and focus on services using those, leaving the informal ones to humans; [Koh13] introduces the concept of **flexiformality** (flexible formality) and discusses the issues involved. A good example is weak type theory [KN04], a λ-calculus with a linguistically inspired type system, which is intended as an intermediate step in the formalization of mathematical developments. The MathLang system [Kam+14] based on ideas from weak type theory allows to annotate (i.e. flexiformalize) various aspects of a mathematical document, up to a point where enough semantic information to drive verification of the document in a proof assistant [Ret09].

The OMDoc (Open Mathematical Documents) format [Koh06] is possibly the most complete framework for flexiformal mathematics. It subsumes all the other representation formats and can serve as an interoperability layer. It specifies markup for mathematical documents and knowledge

---

[3]In fact, the resolution of the Grundlagen Crisis of Mathematics in the last century is that any mathematical document can – in principle – be formalized in first-order logic with axiomatic set theory.

[4]inspired by the Mizar language, hence the name

in a system-independent general framework and relates the two aspects. OMDoc provides representational infrastructure at three levels: the object level for mathematical formulae, based on OpenMath[20] and content MathML, the statement level for definitions, theorems, and proofs, and the theory level for document sectioning and knowledge grouping. OMDoc documents can be created by writing them in sTeX [Koh08; sTeX], a variant of LaTeX that allows "semantic preloading", i.e. invisible OMDoc markup in the LaTeX sources and then converting to OMDoc/XML via the LaTeXML system. Conversely, OMDoc documents can be transformed into **active documents**, which use the semantic information for embedded services – the more semantic preloading, the more services – see [Koh+11] for a discussion.

Finally, another example of flexiformal – here computational – documents are Wolfram Notebooks [MNB] or Jupyter [JN] notebooks. Here, computational "cells" with executable code (Mathematica for Wolfram notebooks and a wide variety of computational systems for Jupyter notebooks) are interleaved with text cells, which provide an (informal) narrative. Computational cells show the results of computations, the code can be arbitrarily edited and re-executed, giving a very flexible way of exploring the mathematical contents. Computational cells can also drive "widgets", which pipe computation results into special-purpose interaction forms.

6 Primary Aspect: Organization

[21] Every mathematical system organizes the body of mathematical knowledge (MK) relevant to the system as a structure consisting of *units of mathematical knowledge*, *means for combining the units*, and an underlying *semantics* for understanding what the units and their combinations mean.

Compound Units of Mathematical Knowledge (C)   There are several kinds of MK units, both **atomic** and **compound**. Atomic MK units are not composed of smaller MK units, while compound MK units contain components that are MK units themselves. Examples of atomic units include equations, tables, algorithms, definitions, theorems, and proofs, while examples of compound units include theorem-proof pairs, term-definition pairs, question-answer pairs, articles, and axiomatic theories.

Indexed Knowledge Collections (I)   There are various ways that mathematical knowledge units can be combined in the structure. The default case is that they are ordered in some kind of a **collection** and indexed by some salient features (e.g. concept names or sizes) either as an organization principle or for a retrieval-based user interface. A mere *list* is a special case of this, there items are indexed by their position, (mathematical) encyclopaedias like the Encyclopedia of Mathematics [EM], PlanetMath [PM], Wikipedia [Wik], and Wolfram MathWorld [Wei] are prime examples; they are indexed by concept names. Other examples are mathematical data sets like the *Digital Library of Mathematical Functions* (DLMF) [Nat10] – indexed by function; the *Online Encyclopedia of Integer Sequences* (OEIS) [OEIS]), and triangle centers (as in the Encyclopedia of Triangle Centers (ETC) [ETC].

An interesting case in this class is the MathOverflow [MOa], where mathematical knowledge is organized (i.e.g indexed) by questions it answers. Of course this poses interesting problems e.g. about when two questions are "equal"; MathOverFlow has developed community-organized solutions here.

Graph-Structured/Semantic Organization (G)   In this more elaborate organizational principle, knowledge items are interconnected by semantical relations, e.g. concepts in a taxonomy. Other examples of the latter include a collection of articles connected by hyperlinks and a collection of axiomatic theories interconnected by theory morphisms. The main mechanism of these

---

[20] EdNote: MK: introduce OpenMath somewhere, probably in the inference aspect
[21] EdNote: responsible: Bill, Yasmine

graph-shape organization forms (knowledge graphs) is that they support some kind of knowledge inheritance mechanism, that enhances space-efficiency, consistency, and maintainability of the knowledge collection.

Arguably the most general concept in this space is that of **theory graphs**: A directed graph whose nodes are axiomatic theories and whose edges are theory morphisms. The latter are meaning-preserving mappings that enable information to flow from abstract theories to more concrete theories or equally abstract theories. This organizational structure is found in many proof assistants, specification systems, and logical frameworks.

Collections of numeric or symbolic algorithms constitute an interesting instance of this class of collections. They are usually organized by code/library dependencies, and thus inherit functionality. Algorithm collections (we think of them as **algorithmic theories**) are embodied in computer algebra systems such as Maple [MA] and Mathematica [WM].

Heterogeneous Organization (H) The units in an organization structure can be *homogeneous*, i.e., all of one kind, or *heterogeneous*, i.e. of several kinds.[22] For algorithmic theories, we think of a collection to be heterogeneous, if it involves multiple programming languages or paradigms. For instance the OEIS [OEIS] collects implementation of integer sequences in many computational stystems.

Theorem Prover libraries, such as those of such as HOL Light [HLL] and Mizar [MizLib] are heterogeneous: they are (essentially) lists of the axioms, definitions, theorems; but they also compose a graph of deductive developments in which the root of the tree is foundational axiomatic theory.

Formal Organization (F) The underlying semantics can be based on traditional informal mathematical practice; formal logics like first-order logic, simple type theory, set theory, and dependent type theory.

Organization by Mathematical Practice (P) In contrast to the above, the knowledge units can be organized by mathematical practice. An example is the Math Subject Classification [Ame09], an

Explicitly represented mathematical knowledge Most mathematical system represents mathematical knowledge. Either implicitly inscribed into the source code used in the implementation of the system or explicitly represented. Many systems manage quite a lot of represented knowledge and data.[24]

Figure 2[25] gives an overview over the scale of mathematical knowledge explicitly represented in state of the art systems. We indicate the tetrapod aspects involved in the last column to give a preview o the discussions in Sections 2 to **??**. Table 1 below gives a more detailed classification of the aspects in terms of specific facets that will be introduced below.

7 The Big Table of Systems and their Aspects

[26]   [27]

[28] [29] [30] [31] [32] [33] [34] [35] [36] [37] [38] [39]

EdN:22

BNP:23

EdN:24
EdN:25

ENP:23

EdN:26
EdN:27
EdN:28
EdN:29
EdN:30

---

[22]EDNOTE: MK@BF: we need to define this: do you mean Florian's "homogenous vs. heterogeneous methods?
[23]NEW PART: MK: moved over from concretization, makes more sense here.
[24]EDNOTE: MK: we need to talk about FAIR and what it means for math somewhere.
[25]EDNOTE: MK: we should probably extend this by some more systems
[26]EDNOTE: MK: I had problems finding or citing the inverse symbolic calculator. There does not seem to be a working version online. Do we care?
[27]EDNOTE: Conjecturing, Theory Exploration
[28]EDNOTE: not Python
[29]EDNOTE: Matlab, Wolfram Alpha? Geogebra?
[30]EDNOTE: Geogebra represents Graphing Calc

Data/Knowledge in	Description	Aspects
Theorem prover libraries [OAF]	≈ 5 proof libraries, $\approx 10^5$ theorems each, ≈ 200 GB	Inf
Computer algebra systems [Deva]	e.g., SageMath distribution bundles ≈ 4 GB of various tools and libraries	Comp
Modelica libraries [MOb]	> 10 official, > 100 open-source, ≈ 50 commercial, > 5.000 classes in the Standard Library, industrial models can reach .5M equations	Comp
Integer Sequences [OEIS]	≈ 330K sequences, ≈ 1 TB	Conc, Inf, Nar
Sequence Identities [KDH]	$\approx .3$M sequence identities, ≈ 2.5 TB	Comp, Inf
Highly symmetric graphs, maps, polytopes [COa; AP; AAP; COb; TG; EET]	≈ 30 datasets, $\approx 2 \cdot 10^6$ objects, ≈ 1 TB	Conc
Finite lattices [FL; AG; UL]	7 datasets, $\approx 17 \cdot 10^9$ objects, ≈ 1.5 TB	Conc
Combinatorial statistics and maps [BSa14]	≈ 1.500 objects	Conc
SageMath databases [SDB]	12 datasets	Comp, Comp
L-functions and modular forms [LM]	≈ 80 datasets, $\approx 10^9$ objects, ≈ 1 TB	Comp, Conc, Nar
zbMATH [ZBM]	≈ 4M publication records with semantic data, ≈ 30M reference data, > 1M disambig. authors, $\approx 2,7$M full text links: ≈ 1M OA	Nar, Link,
swMATH [SWM]	≈ 25K software records with > 300K links to > 180K publications	Nar,Link
EuDML [EUD]	≈ 260K open full-text publications	Nar
Wikidata [WD]	34 GB linked data, thereof about 4K formula entities, interlinked, e.g., with named theorems, persons, and/or publications	Link, Nar
arXiv.org	≈ 300K math preprints (of ≈ 1.6M) most with LaTeX sources	Nar,link
MathOverFlow	$\approx 1,1$M questions/answers, ≥ 11K answer authors	Nar
Stacks project	≥ 6000 pages, semantically annotated, curated, searchable textbook	Nar
nLab	≥ 13K pages on category theory and applications	Nar

Figure 2: Represented Knowledge/Data in Mathematical Software Systems

	System	Reference	Orga.	Inference	Comp.	Concr.	Narration
Deduction	Coq	[Coq15]		C A T P R			
	Isabelle	[ISA]		P			S
	Mizar	[Mizb]		P			
	Otter	[McC03]		A			
	* Hammer			A			
	CVC/Z3						
	Prolog	[CM81]		L			
	FOIL	[Qui90]		L			
Computation	Mathematica	[WM]					S
	SageMath	[Sage]					
	GAP	[GAP]					
	GeoGebra	[GG]					
	CoCalc	[CC]					
	Octave	[GU]					
	Simulink	[MS]					
	R	[R]					
	Stan	[Stan]					
Concretization	Math Gene. Proj.	[MGP]				L	
	WikiData	[WD]	K			R L	
	OEIS	[OEIS]	E			R	P
	LMFDB	[Cre16; LM]				R	
	Small Groups Lib.	[EBO]				R C^1 E	
	DLMF	[Nat10; Loz03]				**?TS?**	S
	Inv. Symb. Calc.					**?TS?**	
Narration	arXiv	[ArX]					P
	zbMath	[ZBM]				L	P
	TeX/LaTeX	[Knu84; Lam94]			?^{2}		P S^2
	pMathML	[MML310]					P
	sTeX	[Koh08]					P S F
	SIUnitsX	[Wri]					F
	OpenMath/ cMathML	[Bus+04; MML310]					F^3
	Wikipedia	[Wik]	E^4				P S^5
Orga.	MathOverflow	[MOa]	I				P
	Polymath	[Pol]					P
	AFP	[AFP]	G H				
	MSC	[MSC10]	P				
	MathHub	[MH]	G G				

Table 1: Tetrapod Systems and their Aspects

We give a overview of mathematical software systems from a tetrapod perspective in Table 1. We list systems in the first row, and specify which aspects they support in the last five using the letter codes specified in the "sub-aspects" in Sections 6 to 5. Note that the letter codes are only unique per column.[40] Where necessary, we mark the codes with comments which can be referenced in the list below.

1 The Small Groups library is deeply integrated into the GAP computer algebra system and uses the GAP-internal representations to represent the groups extremely space-efficiently (a handful of bits per group). In particular, extracting the data out of GAP leads to a dataset that is multiple orders of magnitude less space-efficient.

---

[31]EdNote: Numeric/Scientific computation, Optimisation, Statistics: refer to surveys?
[32]EdNote: Keep in mind: IDRIS, F*
[33]EdNote: Otter representative of Prover 9, Mace
[34]EdNote: probabilistic proof? is it a computation thing?
[35]EdNote: zbMath representative for MathSciNet, has swMath, Google Scholar, Scopus, arXiv
[36]EdNote: MathTutor History of Mathematics (bibliographies) archive http://www-history.mcs.st-and.ac.uk/
[37]EdNote: not Oracle
[38]EdNote: not InDesign, Markdown
[39]EdNote: MathWorld
[40]EdNote: Make an example where this happens

Primary Aspect	Secondary Aspect				
	Org	Inf	Comp	Tab	Narr
Org	category of theories module systems	type inference			
Inf	theories	meta-theorems verification proof-checking	recursion rewriting logic programming tactics ATP decision procedures		documentation semi-formal proofs
Comp	specifications	verification	preprocessing code generation profiling	memoization package repositories	documentation
Tab	schemas		querying built-in functions		
Narr			active documents literate programming macros		documentation

Figure 3: Paradigms for Supporting Secondary Aspects

2 TEX [Knu84] pairs a set of layout primitives with a Turing-complete macro expansion facility, which is use d by a large community to define libraries of macros. LATEX [Lam94] is the widely used one; it establishes semantic markup for sectioning, crossreferences, bibliographic references, and statements.

3 OpenMath and – by reference – presentation MathML (whose semantics is given in terms fo OpenMath) fully describe the structure of mathematical formulae and give the meaning of the symbols in terms of OpenMath Content Dictionaries, which are semantic mathematical documents themselves.

4 Wikipedia uses a restricted subset of TEXto create presentation MathML.

5 Some of the links and concepts are classified with Linked Open Data annotations.

8 Realizing Secondary Aspects

Systems usually use additional aspects, which we call secondary. The secondary aspect can have multiple roles:

- It might enhance knowledge written in the primary aspects, e.g., narrative documentation of programs.
- It may substitute for knowledge written in the primary aspects, e.g., a narrative snippet describing an omitted proof step in a semi-formal proof.
- It may be used to talk about knowledge written in the primary aspect, e.g., a computational tactic that produces proofs or the verification of a computational system.

Note that some of these roles allow for the primary and secondary aspect to be the same. For example, we can use a computational preprocessor to generate programs before compilation. Or we can proof a meta-theorem that states the admissibility of an additional inference rule.

9 Conclusion

In this living survey we survey the state of the art in mathematical software systems from the perspective of the tetrapodal model of mathematical knowledge introduced in [Car+20]. Other than

the survey itself, we also contribute a set of facets of the main five aspects *i*) inference, *ii*) computation, *iii*) concretization, *iv*) narration, and *v*) organization of the tetrapod (see Figure 1). Each paradigmatic mathematical software system is analyzed with these facets in mind— with the results in Table 1. The development of the facets has been an iterative process of recognizing patterns in system functionality, using these for classification, and verifying the soundness of the results on the systems. The table shows that current systems are still predominantly single-aspect, but some trans-aspect facets are creeping in. We hope that this trend continues, and that we will see true tetrapodal systems in the future.

We will monitor the situation in later versions of the survey.

References

[AAP]   Dimitri Leemans. *Atlases of Abstract Polytopes*. URL: `http://homepages.ulb.ac.be/~dleemans/resources.html` (visited on 01/23/2019).

[Abd10]   Abdulrahman Ali Abdulaziz. *The Plimpton 322 Tablet and the Babylonian Method of Generating Pythagorean Triples*. 2010. arXiv: `1004.0025`.

[ABR01]   Ahmed Amerkad, Yves Bertot, and Laurence Rideau. "Mathematics and Proof Presentation in Pcoq". In: *Proceedings of the Workshop on Proof Transformation, Proof Presentations and Complexity of Proofs (PTP-01)*. Ed. by Uwe Egly et al. Universitá degli studi di Siena, 2001, pp. 51–60.

[AFP]   AFP. *Archive of Formal Proofs*. URL: `https://www.isa-afp.org` (visited on 09/20/2019).

[AG]   Dimitri Leemans. *An atlas of subgroup lattices of finite almost simple groups*. URL: `http://homepages.ulb.ac.be/~dleemans/atlaslat/` (visited on 01/25/2019).

[Aiz+16]   Akiko Aizawa et al. "NTCIR-12 MathIR Task Overview". In: *Proceedings of the 12th NTCIR Conference on Evaluation of Information Access Technologies*. Ed. by Noriko Kando, Tetsuya Sakai, and Mark Sanderson. Tokyo, Japan: NII, Tokyo, 2016, pp. 299–308. URL: `http://research.nii.ac.jp/ntcir/workshop/OnlineProceedings12/pdf/ntcir/OVERVIEW/01-NTCIR12-OV-MathIR-ZanibbiR.pdf`.

[AP]   Michael Hartley. *Abstract Polytopes*. URL: `http://www.abstract-polytopes.com/atlas/index.html` (visited on 01/23/2019).

[ArX]   *arxiv.org e-Print archive*. URL: `http://www.arxiv.org` (visited on 06/12/2012).

[Asp+06]   A. Asperti et al. "Crafting a Proof Assistant". In: *TYPES*. Ed. by T. Altenkirch and C. McBride. Springer, 2006, pp. 18–32.

[BCH12]   M. Boespflug, Q. Carbonneaux, and O. Hermant. "The λΠ-calculus modulo as a universal proof language". In: *Proceedings of PxTP2012: Proof Exchange for Theorem Proving*. Ed. by D. Pichardie and T. Weber. 2012, pp. 28–43.

[Ben+08]   C. Benzmüller et al. "LEO-II - A Cooperative Automatic Theorem Prover for Classical Higher-Order Logic (System Description)". In: *Automated Reasoning*. Ed. by A. Armando, P. Baumgartner, and G. Dowek. Springer, 2008, pp. 162–170.

[BJ77]   L.S. van Benthem Jutting. "Checking Landaus "Grundlagen" in the AUTOMATH System". PhD thesis. Eindhoven University of Technology, 1977.

[Bla+16]   J. Blanchette et al. "Hammering towards QED". In: *Journal of Formalized Reasoning* 9.1 (2016), pp. 101–148.

[Bra13]   E. Brady. "Idris, a general-purpose dependently typed programming language: Design and implementation". In: *Journal of Functional Programming* 23.5 (2013), pp. 552–593.

[Bro12]   C. Brown. "Satallax: An Automatic Higher-Order Prover". In: *Automated Reasoning*. Ed. by B. Gramlich, D. Miller, and U. Sattler. Springer, 2012, pp. 111–117.

[BSa14]   C. Berg, C. Stump, and al. *FindStat: The Combinatorial Statistic Finder.* `http://www.FindStat.org`. [Online; accessed 31 August 2016]. 2014.

[BT13]   Sara C. Billey and Bridget E. Tenner. "Fingerprint databases for theorems". In: *Notices Amer. Math. Soc.* 60.8 (2013), pp. 1034–1039. ISSN: 0002-9920. DOI: `10.1090/noti1029`.

[Bus+04]   Stephen Buswell et al. *The Open Math Standard, Version 2.0.* Tech. rep. The OpenMath Society, 2004. URL: `http://www.openmath.org/standard/om20`.

[Car+20]   Jacques Carette et al. "Big Math and the One-Brain Barrier – The Tetrapod Model of Mathematical Knowledge". In: *Mathematical Intelligencer* (2020). in press. URL: `https://arxiv.org/abs/1904.10405`.

[CC]   *CoCalc: Collaborative Calculation in the Cloud.* URL: `https://cocalc.com` (visited on 01/28/2019).

[Chu40]   Alonzo Church. "A Formulation of the Simple Theory of Types". In: *Journal of Symbolic Logic* 5 (1940), pp. 56–68.

[CKS11]   Marcos Cramer, Peter Koepke, and Bernhard Schröder. "Parsing and Disambiguation of Symbolic Mathematics in the Naproche System". In: *Intelligent Computer Mathematics.* Ed. by James Davenport et al. LNAI 6824. Springer Verlag, 2011, pp. 180–195. ISBN: 978-3-642-22672-4.

[Cla+96]   M. Clavel et al. "Principles of Maude". In: *Proceedings of the First International Workshop on Rewriting Logic.* Ed. by J. Meseguer. Vol. 4. 1996, pp. 65–89.

[CM81]   W. F. Clocksin and C. S. Mellish. *Programming in Prolog.* Springer Verlag, 1981.

[COa]   Marston Conder. *Lists of regular maps, hypermaps and polytopes, trivalent symmetric graphs, and surface actions.* URL: `https://www.math.auckland.ac.nz/~conder/` (visited on 01/23/2019).

[COb]   Primož Potočnik. *Lists of graphs of a prescribed symmetry type and valence, and some other combinatorial and algebraic structures.* URL: `https://www.fmf.uni-lj.si/~potocnik/work.htm` (visited on 01/23/2019).

[Con+11]   Garth Conboy et al. *EPUB 3 Overview.* Recommended Specification. International Digital Publishing Forum (IDPF), Oct. 11, 2011. URL: `http://idpf.org/epub/30/spec/epub30-overview.html`.

[Cra+10]   Marcos Cramer et al. "The Naproche Project Controlled Natural Language Proof Checking of Mathematical Texts". In: *Controlled Natural Language, Workshop on Controlled Natural Language, CNL 2009. Revised Papers.* Ed. by Norbert E. Fuchs. LNCS 5972. Springer, 2010, pp. 170–186. ISBN: 978-3-642-14417-2. DOI: `10.1007/978-3-642-14418-9_11`.

[Cre16]   John Cremona. "The L-Functions and Modular Forms Database Project". In: *Foundations of Computational Mathematics* 16.6 (2016), pp. 1541–1553. ISSN: 1615-3383. DOI: `10.1007/s10208-016-9306-z`.

[Dah+11]   *Scalable Vector Graphics (SVG) 1.1 (Second Edition).* W3C Recommendation. World Wide Web Consortium (W3C), Apr. 2011. URL: `http://www.w3.org/TR/SVG11`.

[dB08]   L. de Moura and N. Bjørner. "Z3: An Efficient SMT Solver". In: *Tools and Algorithms for the Construction and Analysis of Systems.* Ed. by C. Ramakrishnan and J. Rehof. Springer, 2008, pp. 337–340.

[DB12]   *Open Document Format for Office Applications (OpenDocument) v1.2.* OASIS Standard. 2012. URL: `http://docs.oasis-open.org/office/v1.2/os/OpenDocument-v1.2-os.html`.

[Deva]   The Sage Developers. *SageMath, the Sage Mathematics Software System.* `http://www.sagemath.org`. [Online; accessed 30 August 2016]. URL: `http://www.sagemath.org`.

[Devb]   The Sage Developers. *SageMath, the Sage Mathematics Software System.* URL: `http://www.sagemath.org`.

[EBO]   Bettina Eick, Hans Ulrich Besche, and Eamonn O'Brien. *SmallGrp – The GAP Small Groups Library.* URL: `https://www.gap-system.org/Manuals/pkg/SmallGrp-1.3/doc/chap1.html` (visited on 10/13/2018).

[EET]   Steve Wilson and Primož Potočnik. *A Census of edge-transitive tetravalent graphs.* URL: `https://jan.ucc.nau.edu/~swilson/C4FullSite/index.html` (visited on 01/23/2019).

[EM]   *Encyclopedia of Mathematics.* URL: `http://www.encyclopediaofmath.org` (visited on 02/01/2014).

[ETC]   *Encyclopedia of Triangle Centers – ETC.* URL: `https://faculty.evansville.edu/ck6/encyclopedia/ETC.html` (visited on 02/07/2020).

[EUD]   *EuDML – The European Digital Mathematics Library.* URL: `http://eudml.eu` (visited on 08/02/2011).

[Fau+17]   Steve Faulkner et al. *HTML 5.2.* W3C Recommmentation. World Wide Web Consortium (W3C), Oct. 28, 2017. URL: `http://w3.org/TR/html52/`.

[FL]   Jukka Kohonen. *Lists of finite lattices (modular, semimodular, graded and geometric).* URL: `https://www.shsu.edu/mem037/Lattices.html` (visited on 01/25/2019).

[GAP]   The GAP Group. *GAP – Groups, Algorithms, and Programming.* URL: `http://www.gap-system.org` (visited on 08/30/2016).

[GG]   Markus Hohenwarter. *Graphing Calculator – GeoGebra.* URL: `http://www.geogebra.org` (visited on 07/23/2019).

[GSC15]   Ferruccio Guidi and Claudio Sacerdoti Coen. "A Survey on Retrieval of Mathematical Knowledge". In: *Intelligent Computer Mathematics 2015.* Ed. by Manfred Kerber et al. LNCS 9150. Springer, 2015, pp. 296–315. ISBN: 978-3-319-20615-8. DOI: `10.1007/978-3-319-20615-8_20`.

[GU]   *GNU Octave.* URL: `https://gnu.org/software/octave/` (visited on 01/20/2020).

[Har96]   J. Harrison. "HOL Light: A Tutorial Introduction". In: *Proceedings of the First International Conference on Formal Methods in Computer-Aided Design.* Springer, 1996, pp. 265–269.

[HLL]   *HOL Light Library.* URL: `https://github.com/jrh13/hol-light/tree/master/Library` (visited on 02/07/2020).

[Hor00]   Helmut Horacek. "Tailoring Inference-Rich Descriptions Through Making Compromises Between Conflicting Cooperation Principles." In: *Int. J. Human-Computer Studies* 53 (2000), pp. 1117–1146.

[Hua96]   Xiaorong Huang. *Human Oriented Proof Presentation: A Reconstructive Approach.* DISKI 112. Sankt Augustin, Germany: Infix, 1996.

[ISA]   *Isabelle.* URL: `http://isabelle.in.tum.de` (visited on 03/27/2013).

[JFM]   *Journal of Formalized Mathematics.* URL: `http://www.mizar.org/JFM` (visited on 09/27/2012).

[JN]   *Jupyter Notebook.* URL: `http://jupyter-notebook.readthedocs.org/en/latest/notebook.html#notebook-documents` (visited on 08/22/2017).

[Kam+14]   Fairouz Kamareddine et al. "Computerising Mathematical Text". In: *Computational Logic*. Ed. by Dov M. Gabbay, Jrg Siekmann, and John Woods. North Holland, 2014, pp. 343–396. ISBN: 9780444516244.

[KDH]   *KWARC Data Host*. URL: `https://datahost.kwarc.info/` (visited on 01/25/2019).

[KLN04]   Fairouz Kamareddine, Twan Laan, and Rob Nederpelt. *A Modern Perspective on Type Theory: From its Origins until Today*. Applied Logic 29. Kluwer, 2004.

[KN04]   Fairouz Kamareddine and Rob Nederpelt. "A refinement of de Bruijn's formal language of mathematics". In: *Logic, Language and Information* 13.3 (2004), pp. 287–340. URL: `http://www.macs.hw.ac.uk/~fairouz/forest/papers/journals-publications/kjour.pdf`.

[Knu84]   Donald E. Knuth. *The TEXbook*. Addison Wesley, 1984.

[Koh+11]   Michael Kohlhase et al. "The Planetary System: Web 3.0 & Active Documents for STEM". In: *Procedia Computer Science* 4 (2011): *Special issue: Proceedings of the International Conference on Computational Science (ICCS)*. Ed. by Mitsuhisa Sato et al. Finalist at the Executable Paper Grand Challenge, pp. 598–607. DOI: `10.1016/j.procs.2011.04.063`.

[Koh06]   Michael Kohlhase. *OMDoc – An open markup format for mathematical documents [Version 1.2]*. LNAI 4180. Springer Verlag, Aug. 2006. URL: `http://omdoc.org/pubs/omdoc1.2.pdf`.

[Koh08]   Michael Kohlhase. "Using LATEX as a Semantic Markup Format". In: *Mathematics in Computer Science* 2.2 (2008), pp. 279–304. URL: `https://kwarc.info/kohlhase/papers/mcs08-stex.pdf`.

[Koh13]   Michael Kohlhase. "The Flexiformalist Manifesto". In: *14th International Workshop on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC 2012)*. Ed. by Andrei Voronkov et al. Timisoara, Romania: IEEE Press, 2013, pp. 30–36. ISBN: 978-1-4673-5026-6. URL: `http://kwarc.info/kohlhase/papers/synasc13.pdf`.

[KU15]   C. Kaliszyk and J. Urban. "HOL(y)Hammer: Online ATP Service for HOL Light". In: *Mathematics in Computer Science* 9.1 (2015), pp. 5–22.

[Lam94]   Leslie Lamport. *LaTeX: A Document Preparation System, 2/e*. Addison Wesley, 1994.

[Lan30]   Edmund Landau. *Grundlagen der Analysis*. Reprint of the edition, Leipzig, 1970. Darmstadt, Germany; second edition: Wissenschaftliche Buchgesellschaft, 1930.

[LM]   *The L-functions and Modular Forms Database*. URL: `http://www.lmfdb.org` (visited on 02/01/2016).

[Loz03]   D. W. Lozier. "NIST Digital Library of Mathematical Function". In: *Annals of Mathematics and Artificial Intelligence — Special Issue on Mathematical Knowledge Management* 38 (2003), pp. 105–119. URL: `http://dlmf.nist.gov/about/publications/MKM-Lozier.pdf`.

[LTX]   Bruce Miller. *LaTeXML: A LATEX to XML Converter*. URL: `http://dlmf.nist.gov/LaTeXML/` (visited on 03/12/2013).

[MA]   *Maple: The Essential Tool for Mathematics*. URL: `https://www.maplesoft.com/products/maple/` (visited on 01/07/2019).

[Mat]   *MathPlayer. ¡Display MathML in your browser¿*. URL: `http://www.dessci.com/en/products/mathplayer` (visited on 03/19/2012).

[McC03]   William McCune. *Otter 3.3 Reference Manual*. Technical Report ANL/MCS-TM-263. Argonne, IL, USA: Mathematics and Computer Science Division, Argonne National Laboratory, 2003.

[MGP]   *The Mathematics Genealogy Project.* URL: https://genealogy.math.ndsu.nodak.edu/ (visited on 01/20/2020).

[MH]   *MathHub.info: Active Mathematics.* URL: http://mathhub.info (visited on 01/28/2014).

[Mil]   Dale Miller. *λProlog.* URL: http://www.lix.polytechnique.fr/Labo/Dale.Miller/lProlog/.

[Miza]   *Mizar.* httpxbpap://www.mizar.org. seen 2013-02-27. URL: http://www.mizar.org.

[Mizb]   *Mizar.* URL: http://www.mizar.org (visited on 02/27/2013).

[Mizc]   *Mizar.* URL: http://www.mizar.org.

[MizLib] *Mizar Mathematical Library.* URL: http://www.mizar.org/library (visited on 09/27/2012).

[MML310] Ron Ausbrooks et al. *Mathematical Markup Language (MathML) Version 3.0.* Ed. by David Carlisle, Patrick Ion, and Robert Miner. 2010. URL: http://www.w3.org/TR/MathML3.

[MNB]   *Wolfram Notebooks.* URL: http://www.wolfram.com/notebooks/ (visited on 03/29/2019).

[MOa]   *MathOverflow.* URL: http://mathoverflow.net (visited on 09/08/2011).

[MOb]   *Modelica and the Modelica Association.* URL: https://modelica.org/ (visited on 01/22/2019).

[MOc]   Gordon Royle. *What are some early examples of creation of lists / catalogues of (particularly) combinatorial objects? – MathOverflow.* URL: https://mathoverflow.net/questions/47044/what-are-some-early-examples-of-creation-of-lists-catalogues-of-particularly (visited on 01/09/2020).

[MS]   *Simulation and Model-Based Design – MatLab & Simulink.* URL: https://www.mathworks.com/products/simulink.html (visited on 01/20/2020).

[MSC10]   *Mathematics Subject Classification MSC2010.* 2010. URL: http://msc2010.org (visited on 11/16/2011).

[Nor05]   U. Norell. *The Agda WiKi.* http://wiki.portal.chalmers.se/agda. 2005.

[OAF]   *OAF: An Open Archive for Formalizations.* URL: https://kwarc.info/projects/oaf/ (visited on 04/26/2018).

[ODAP]   *OPenNDAP - Advanced Software for Remote Data Retrieval.* URL: https://opendap.org/ (visited on 01/25/2019).

[OEIS]   *The On-Line Encyclopedia of Integer Sequences.* URL: http://oeis.org (visited on 05/28/2017).

[OOXML06] *Standard ECMA-376 - Office Open XML File Formats.* Tech. rep. 2006. URL: http://www.ecma-international.org/publications/standards/Ecma-376.htm.

[ORS92]   S. Owre, J. Rushby, and N. Shankar. "PVS: A Prototype Verification System". In: *11th International Conference on Automated Deduction (CADE).* Ed. by D. Kapur. Springer, 1992, pp. 748–752.

[Pau94]   L. Paulson. *Isabelle: A Generic Theorem Prover.* Vol. 828. Lecture Notes in Computer Science. Springer, 1994.

[PM]   *PlanetMath.org – Math for the people, by the people.* URL: http://planetmath.org (visited on 11/11/2012).

[Pol]   *The polymath blog.* URL: http://polymathprojects.org/ (visited on 07/12/2019).

[Qed]   http://www-unix.mcs.anl.gov/qed/.

[Qui90]   J. Ross Quinlan. "Learning Logical Definitions from Relations". In: *Machine Learning* 5.3 (1990), pp. 239–266.

[R]   *The R Project for Statistical Computing*. URL: https://r-project.org (visited on 01/20/2020).

[Ret09]   Krzysztof Retel. "Gradual Computerisation and Verification of Mathematics. MathLang's Path into Mizar". PhD thesis. Edinburgh: Heriot-Watt University, Apr. 2009.

[RV02]   A. Riazanov and A. Voronkov. "The design and implementation of Vampire". In: *AI Communications* 15 (2002), pp. 91–110.

[Sage]   The Sage Developers. *SageMath, the Sage Mathematics Software System*. URL: http://www.sagemath.org (visited on 09/30/2016).

[Sch01]   S. Schulz. "System Abstract: E 0.61". In: *International Joint Conference on Automated Reasoning*. Ed. by R. Goré, A. Leitsch, and T. Nipkow. Springer, 2001, pp. 370–375.

[SDB]   *SageMath Databases*. URL: http://doc.sagemath.org/html/en/reference/databases/index.html (visited on 01/25/2019).

[Slo03]   Neil J. A. Sloane. "The On-Line Encyclopedia of Integer Sequences". In: *Notices of the AMS* 50.8 (2003), p. 912.

[Stan]   *Stan*. URL: https://mc-stan.org/ (visited on 01/20/2020).

[sTeX]   *sTeX: A semantic Extension of TeX/LaTeX*. URL: https://github.com/sLaTeX/sTeX (visited on 05/15/2015).

[SWM]   *Mathematical Software – swMATH*. URL: http://swmath.org (visited on 09/07/2017).

[TG]   Brendan McKay, Gordon Royle, and Alexander Hulpke. *Transitive Graphs*. URL: http://staffhome.ecm.uwa.edu.au/~00013890/remote/trans/index.html (visited on 01/23/2019).

[UL]   Martin Malandro. *Unlabeled lattices on <= 15 nodes*. URL: https://b2share.eudat.eu/records/dbb096da4e364b5e9e37b982431f41de (visited on 01/25/2019).

[W2X]   *Writer2xhtml – Apache OpenOffice Extensions*. URL: https://extensions.openoffice.org/fr/project/writer2xhtml (visited on 12/01/2019).

[WD]   *Wikidata:Introduction*. URL: https://wikidata.org/wiki/Wikidata:Introduction (visited on 01/25/2015).

[Wei]   Eric W. Weisstein, ed. *Wolfram MathWorld. the web's most extensive mathematics resource*. Wolfram Research. URL: http://mathworld.wolfram.com (visited on 12/02/2009).

[Wei+02]   C. Weidenbach et al. "SPASS Version 2.0". In: *Conference on Automated Deduction*. Ed. by A. Voronkov. Springer, 2002, pp. 275–279.

[Wen07]   Makarius Wenzel. "Isabelle/Isar — a generic framework for human-readable proof documents." In: *From Insight to Proof: Festschrift in Honour of Andrzej Trybulec*. Ed. by R. Matuszewski and A. Zalewska. Vol. 10:23. Studies in Logic, Grammar and Rhetoric. University of Białystok, 2007, pp. 277–298. URL: http://mizar.org/trybulec65/.

[WM]   *WOLFRAM MATHEMATICA The world's definitive system for modern technical computing*. URL: https://www.wolfram.com/mathematica/ (visited on 01/24/2019).

[WR10]   Alfred North Whitehead and Bertrand Russell. *Principia Mathematica*. 2nd ed. Vol. I. Cambridge, UK: Cambridge University Press, 1910.

[Wri]   Joseph Wright. `siunitx` *– A comprehensive (SI) units package*. URL: http://ctan.org/pkg/siunitx (visited on 12/02/2019).

[ZBF]   *zbMATH – Formula Search*. URL: https://zbmath.org/formulae/ (visited on 03/28/2019).

[ZBM]   *zbMATH the first resource in mathematics*. URL: http://zbmath.org (visited on 01/29/2019).

[Ame09]   American Mathematical Society. *Mathematics Subject Classification MSC2010*. http://www.ams.org/mathscinet/msc/. 2009.

[Coq15]   Coq Development Team. *The Coq Proof Assistant: Reference Manual*. Tech. rep. INRIA, 2015.

[HOL]   HOL4 development team. http://hol.sourceforge.net/.

[Nat10]   National Institute of Standards and Technology (NIST), ed. *Digital Library of Mathematical Functions*. May 7, 2010. URL: http://dlmf.nist.gov.

[Wik]   Wikimedia Foundation, ed. *Wikipedia, the free encyclopedia*. URL: http://www.wikipedia.org (visited on 08/11/2010).

[de +15]   L. de Moura et al. "The Lean Theorem Prover (System Description)". In: *Automated Deduction*. Ed. by A. Felty and A. Middeldorp. Springer, 2015, pp. 378–388.