Comparative Verification of the Digital Library of Mathematical Functions and Computer Algebra Systems

André Greiner-Petter[1]([✉]) , Howard S. Cohl[2] , Abdou Youssef[2,3],
Moritz Schubotz[1,4] , Avi Trost[5], Rajen Dey[6], Akiko Aizawa[7] , and Bela Gipp[1]

[1] University of Wuppertal, Wuppertal, Germany,
`{greinerpetter,schubotz,gipp}@uni-wuppertal.de`
[2] National Institute of Standards and Technology,
Mission Viejo, CA, U.S.A., `howard.cohl@nist.gov`
[3] George Washington University, Washington, D.C., U.S.A, `ayoussef@gwu.edu`
[4] FIZ Karlsruhe, Berlin, Germany, `moritz.schubotz@fiz-karlsruhe.de`
[5] Brown University, Providence, RI, U.S.A., `avitrost@gmail.com`
[6] University of California Berkeley, Berkeley, CA, U.S.A., `rajhataj@gmail.com`
[7] National Institute of Informatics, Tokyo, Japan, `aizawa@nii.ac.jp`

Abstract. Digital mathematical libraries assemble the knowledge of years of mathematical research. Numerous disciplines (e.g., physics, engineering, pure and applied mathematics) rely heavily on compendia gathered findings. Likewise, modern research applications rely more and more on computational solutions, which are often calculated and verified by computer algebra systems. Hence, the correctness, accuracy, and reliability of both digital mathematical libraries and computer algebra systems is a crucial attribute for modern research. In this paper, we present a novel approach to verify a digital mathematical library and two computer algebra systems with one another by converting mathematical expressions from one system to the other. We use our previously developed conversion tool (referred to as L^AC_AST) to translate formulae from the NIST Digital Library of Mathematical Functions to the computer algebra systems `Maple` and `Mathematica`. The contributions of our presented work are as follows: (1) we present the most comprehensive verification of computer algebra systems and digital mathematical libraries with one another; (2) we significantly enhance the performance of the underlying translator in terms of coverage and accuracy; and (3) we provide open access to translations for `Maple` and `Mathematica` of the formulae in the NIST Digital Library of Mathematical Functions.

Keywords: Presentation to Computation, LaCASt, LaTeX, Semantic LaTeX, Computer Algebra Systems, Digital Mathematical Library

1 Introduction

Digital Mathematical Libraries (DML) gather the knowledge and results from thousands of years of mathematical research. Even though pure and applied mathematics are precise disciplines, gathering their knowledge bases over many years results in

issues which every digital library shares: consistency, completeness, and accuracy. Likewise, Computer Algebra Systems (CAS)[8] play a crucial role in the modern era for pure and applied mathematics, and those fields which rely on them. CAS can be used to simplify, manipulate, compute, and visualize mathematical expressions. Accordingly, modern research regularly uses DML and CAS together. Nonetheless, DML [7,14] and CAS [1,20,11] are not exempt from having bugs or errors. Durán et al. [11] even raised the rather dramatic question: *"can we trust in [CAS]?"*

Existing comprehensive DML, such as the Digital Library of Mathematical Functions (DLMF) [10], are consistently updated and frequently corrected with errata[9]. Although each chapter of the DLMF has been carefully written, edited, validated, and proofread over many years, errors still remain. Maintaining a DML, such as the DLMF, is a laborious process. Likewise, CAS are eminently complex systems, and in the case of commercial products, often similar to black boxes in which the magic (i.e., the computations) happens in opaque private code [11]. CAS, especially commercial products, are often exclusively tested internally during development.

An independent examination process can improve testing and increase trust in the systems and libraries. Hence, we want to elaborate on the following research question.

> How can digital mathematical libraries and computer algebra systems be utilized to improve and verify one another?

Our initial approach for answering this question is inspired by our previous studies on translating DLMF equations to CAS [7]. In order to verify a translation tool from a specific LaTeX dialect to `Maple`[10]. , we performed *symbolic* and *numeric* evaluations on equations from the DLMF. Our approach presumes that a proven equation in a DML must be also valid in a CAS. In turn, a disparity in between the DML and CAS would lead to an issue in the translation process. However, assuming a correct translation, a disparity would also indicate an issue either in the DML source or the CAS implementation. In turn, we can take advantage of the same approach to improve and even verify DML with CAS and vice versa. Unfortunately, previous efforts to translate mathematical expressions from various formats, such as LaTeX [8,14,29], MathML [31], or OpenMath [18,30], to CAS syntax have shown that the translation will be the most critical part of this verification approach.

In this paper, we elaborate on the feasibility and limitations of the translation approach from DML to CAS as a possible answer to our research question. We further focus on the DLMF as our DML and the two general-purpose CAS `Maple` and `Mathematica` for this first study. This relatively sharp limitation is necessary in order to analyze the capabilities of the underlying approach to verify commercial CAS

---

[8] In the sequel, the acronyms CAS and DML are used, depending on the context, interchangeably with their plurals.

[9] https://dlmf.nist.gov/errata/ [accessed 09/01/2021]

[10] The mention of specific products, trademarks, or brand names is for purposes of identification only. Such mention is not to be interpreted in any way as an endorsement or certification of such products or brands by the National Institute of Standards and Technology, nor does it imply that the products so identified are necessarily the best available for the purpose. All trademarks mentioned herein belong to their respective owners.

and large DML. The DLMF uses semantic macros internally in order to disambiguate mathematical expressions [27,35]. These macros help to mitigate the open issue of retrieving sufficient semantic information from a context to perform translations to formal languages [31,14]. Further, the DLMF and general-purpose CAS have a relatively large overlap in coverage of special functions and orthogonal polynomials. Since many of those functions play a crucial role in a large variety of different research fields, we focus in this study mainly on these functions. Lastly, we will take our previously developed translation tool LACAST [8,14] as the baseline for translations from the DLMF to `Maple`. In this successor project, we focus on improving LACAST to minimize the negative effect of wrong translations as much as possible for our study. In the future, other DML and CAS can be improved and verified following the same approach by using a different translation approach depending on the data of the DML, e.g., MATHML [31] or OpenMath [18].

In particular, in this paper, we fix the majority of the remaining issues of LACAST [7], which allows our tool to translate twice as many expressions from the DLMF to the CAS as before. Current extensions include the support for the mathematical operators: sum, product, limit, and integral, as well as overcoming semantic hurdles associated with Lagrange (prime) notations commonly used for differentiation. Further, we extend its support to include `Mathematica` using the freely available *Wolfram Engine for Developers* (WED)[11] (hereafter, with `Mathematica`, we refer to the WED). These improvements allow us to cover a larger portion of the DLMF, increase the reliability of the translations via LACAST, and allow for comparisons between two major general-purpose CAS for the first time, namely `Maple` and `Mathematica`. Finally, we provide open access to all the results contained within this paper, including all translations of DLMF formulae, an endpoint to LACAST[12], and the full source code of LACAST[13].

The paper is structured as follows. Section 2 explains the data in the DLMF. Section 3 focus on the improvements of LACAST that had been made to make the translation as comprehensive and reliable as possible for the upcoming evaluation. Section 4 explains the symbolic and numeric evaluation pipeline. Since Cohl et al. [7] only briefly sketched the approach of a numeric evaluation, we will provide an in-depth discussion of that process in Section 4. Subsequently, we analyze the results in Section 5. Finally, we conclude the findings and provide an outlook for upcoming projects in Section 6.

1.1 Related Work

Existing verification techniques for CAS often focus on specific subroutines or functions [26,20,5,12,6,25,21,17], such as a specific theorems [23], differential equations [19], or the implementation of the `math.h` library [24]. Most common are verification approaches that rely on intermediate verification languages [5,20,21,19,17], such as *Boogie* [25,2] or *Why3* [21,4], which, in turn, rely on proof assistants and theorem provers, such as *Coq* [5,3], *Isabelle* [19,28], or *HOL Light* [20,16,17]. Kaliszyk and Wiedijk [20] proposed on entire new CAS which is built on top of the proof assistant HOL Light so that each simplification step can be proven by the underlying

---

[11] https://www.wolfram.com/engine/ [accessed 09/01/2021]

[12] https://lacast.wmflabs.org/ [accessed 01/01/2022]

[13] https://github.com/ag-gipp/LaCASt [accessed 04/01/2022]

architecture. Lewis and Wester [26] manually compared the symbolic computations on polynomials and matrices with seven CAS. Aguirregabiria et al. [1] suggested to teach students the known traps and difficulties with evaluations in CAS instead to reduce the overreliance on computational solutions.

Cohl et al. [7] developed the aforementioned translation tool LaCASt, which translates expressions from a semantically enhanced LaTeX dialect to `Maple`. By evaluating the performance and accuracy of the translations, we were able to discover a sign-error in one the DLMF's equations [7]. While the evaluation was not intended to verify the DLMF, the translations by the rule-based translator LaCASt provided sufficient robustness to identify issues in the underlying library. To the best of our knowledge, besides this related evaluation via LaCASt, there are no existing libraries or tools that allow for automatic verification of DML.

2 The DLMF dataset

In the modern era, most mathematical texts (handbooks, journal publications, magazines, monographs, treatises, proceedings, etc.) are written using the document preparation system LaTeX. However, the focus of LaTeX is for precise control of the rendering mechanics rather than for a semantic description of its content. In contrast, CAS syntax is coercively unambiguous in order to interpret the input correctly. Hence, a transformation tool from DML to CAS must disambiguate mathematical expressions. While there is an ongoing effort towards such a process [32,22,34,13,36,33], there is no reliable tool available to disambiguate mathematics sufficiently to date.

The DLMF contains numerous relations between functions and many other properties. It is written in LaTeX but uses specific semantic macros when applicable [35]. These semantic macros represent a unique function or polynomial defined in the DLMF. Hence, the semantic LaTeX used in the DLMF is often unambiguous. For a successful evaluation via CAS, we also need to utilize all requirements of an equation, such as constraints, domains, or substitutions. The DLMF provides this additional data too and generally in a machine-readable form [35]. This data is accessible via the i-boxes (information boxes next to an equation marked with the icon ⓘ). If the information is not given in the attached i-box or the information is incorrect, the translation via LaCASt would fail. The i-boxes, however, do not contain information about branch cuts (see Section B) or constraints. Constraints are accessible if they are directly attached to an equation. If they appear in the text (or even a title), LaCASt cannot utilize them. The test dataset, we are using, was generated from DLMF Version 1.1.3 (2021-09-15) and contained 9,977 formulae with 1,505 defined symbols, 50,590 used symbols, 2,691 constraints, and 2,443 warnings for non-semantic expressions, i.e., expressions without semantic macros [35]. Note that the DLMF does not provide access to the underlying LaTeX source. Therefore, we added the source of every equation to our result dataset.

3 Semantic LaTeX to CAS translation

The aforementioned translator LaCASt was developed by Cohl and Greiner-Petter et al. [8,7,14]. They reported a coverage of 58.8% translations for a manually selected

part of the DLMF to the CAS `Maple`. This version of LᴬCᴀsT serves as a baseline for our improvements. In order to verify their translations, they used symbolic and numeric evaluations and reported a success rate of $\sim 16\%$ for symbolic and $\sim 12\%$ for numeric verifications.

Evaluating the baseline on the entire DLMF result in a coverage of only 31.6%. Hence, we first want to increase the coverage of LᴬCᴀsT on the DLMF. To achieve this goal, we first increasing the number of translatable semantic macros by manually defining more translation patterns for special functions and orthogonal polynomials. For `Maple`, we increased the number from 201 to 261. For `Mathematica`, we define 279 new translation patterns which enables LᴬCᴀsT to perform translations to `Mathematica`. Even though the DLMF uses 675 distinguished semantic macros, we cover $\sim 70\%$ of all DLMF equations with our extended list of translation patterns (see Zipf's law for mathematical notations [15]). In addition, we implemented rules for translations that are applicable in the context of the DLMF, e.g., ignore ellipsis following floating-point values or \choose always refers to a binomial expression. Finally, we tackle the remaining issues outlined by Cohl et al. [7] which can be categorized into three groups: (i) expressions of which the arguments of operators are not clear, namely sums, products, integrals, and limits; (ii) expressions with prime symbols indicating differentiation; and (iii) expressions that contain ellipsis. While we solve some of the cases in Group (iii) by ignoring ellipsis following floating-point values, most of these cases remain unresolved. In the following, we elaborate our solutions for (i) in Section 3.1 and (ii) in Section 3.2.

3.1   Parse sums, products, integrals, and limits

Here we consider common notations for the sum, product, integral, and limit operators. For these operators, one may consider mathematically essential operator metadata (MEOM). For all these operators, the MEOM includes *argument(s)* and *bound variable(s)*. The operators act on the arguments, which are themselves functions of the bound variable(s). For sums and products, the bound variables are referred to as *indices*. The bound variables for integrals[14] are called *integration variables*. For limits, the bound variables are continuous variables (for limits of continuous functions) and indices (for limits of sequences). For integrals, MEOM include precise descriptions of regions of integration (e.g., piecewise continuous paths/intervals/regions). For limits, MEOM include limit points (e.g., points in $\mathbb{R}^n$ or $\mathbb{C}^n$ for $n \in \mathbb{N}$), as well as information related to whether the limit to the limit point is independent or dependent on the direction in which the limit is taken (e.g., one-sided limits).

For a translation of mathematical expressions involving the LATEX commands \sum, \int, \prod, and \lim, we must extract the MEOM. This is achieved by (a) determining the argument of the operator and (b) parsing corresponding subscripts, superscripts, and arguments. For integrals, the MEOM may be complicated, but certainly contains the argument (function which will be integrated), bound (integration) variable(s) and details related to the region of integration. Bound variable extraction is usually straightforward since it is usually contained within a differential expression

---

[14] The notion of integrals includes: antiderivatives (indefinite integrals), definite integrals, contour integrals, multiple (surface, volume, etc.) integrals, Riemannian volume integrals, Riemann integrals, Lebesgue integrals, Cauchy principal value integrals, etc.

(infinitesimal, pushforward, differential 1-form, exterior derivative, measure, etc.), e.g., dx. Argument extraction is less straightforward since even though differential expressions are often given at the end of the argument, sometimes the differential expression appears in the numerator of a fraction (e.g., $\int \frac{f(x)dx}{g(x)}$). In which case, the argument is everything to the right of the \int (neglecting its subscripts and superscripts) up to and including the fraction involving the differential expression (which may be replaced with 1). In cases where the differential expression is fully to the right of the argument, then it is a *termination symbol*. Note that some scientists use an alternate notation for integrals where the differential expression appears immediately to the right of the integral, e.g., $\int dx f(x)$. However, this notation does not appear in the DLMF. If such notations are encountered, we follow the same approach that we used for sums, products, and limits (see Section 3.1).

Extraction of variables and corresponding MEOM The subscripts and superscripts of sums, products, limits, and integrals may be different for different notations and are therefore challenging to parse. For integrals, we extract the bound (integration) variable from the differential expression. For sums and products, the upper and lower bounds may appear in the subscript or superscript. Parsing subscripts is comparable with the problem of parsing constraints [7] (which are often not consistently formulated). We overcame this complexity by manually defining patterns of common constraints and refer to them as blueprints. This blueprint pattern approach allows LACAST to identify the MEOM in the sub- and superscripts. A more detailed explanations with examples about the blueprints is available in the Appendix A[15].

Identification of operator arguments Once we have extracted the bound variable for sums, products, and limits, we need to determine the end of the argument. We analyzed all sums in the DLMF and developed a heuristic that covers all the formulae in the DLMF and potentially a large portion of general mathematics. Let x be the extracted bound variable. For sums, we consider a summand as a part of the argument if (I) it is the very first summand after the operation; or (II) x is an element of the current summand; or (III) x is an element of the following summand (subsequent to the current summand) and there is no termination symbol between the current summand and the summand which contains x with an equal or lower depth according to the parse tree (i.e., closer to the root). We consider a summand as a single logical construct since addition and subtraction are granted a lower operator precedence than multiplication in mathematical expressions. Similarly, parentheses are granted higher precedence and, thus, a sequence wrapped in parentheses is part of the argument if it obeys the rules (I-III). Summands, and such sequences, are always entirely part of sums, products, and limits or entirely not.

A termination symbol always marks the end of the argument list. Termination symbols are relation symbols, e.g., $=$, $\neq$, $\leq$, closing parentheses or brackets, e.g.,),], or $>$, and other operators with MEOMs, if and only if, they define the same bound variable. If x is part of a subsequent operation, then the following operator

---

is considered as part of the argument (as in (II)). However, a special condition for termination symbols is that it is only a termination symbol for the current chain of arguments. Consider a sum over a fraction of sums. In that case, we may reach a termination symbol within the fraction. However, the termination symbol would be deeper inside the parse tree as compared to the current list of arguments. Hence, we used the depth to determine if a termination symbol should be recognized or not. Consider an unusual notation with the binomial coefficient as an example

$$\sum_{k=0}^{n}\binom{n}{k} = \sum_{k=0}^{n}\frac{\prod_{m=1}^{n}m}{\prod_{m=1}^{k}m\prod_{m=1}^{n-k}m}. \tag{1}$$

This equation contains two termination symbols, marked red and green. The red termination symbol $=$ is obviously for the first sum on the left-hand side of the equation. The green termination symbol $\prod$ terminates the product to the left because both products run over the same bound variable m. In addition, none of the other $=$ signs are termination symbols for the sum on the right-hand side of the equation because they are deeper in the parse tree and thus do not terminate the sum.

Note that varN in the blueprints also matches multiple bound variable, e.g., $\sum_{m,k\in A}$. In such cases, x from above is a list of bound variables and a summand is part of the argument if one of the elements of x is within this summand. Due to the translation, the operation will be split into two preceding operations, i.e., $\sum_{m,k\in A}$ becomes $\sum_{m\in A}\sum_{k\in A}$. Figure 1 shows the extracted arguments for some example sums. The same rules apply for extraction of arguments for products and limits.

$\boxed{\sum_{n=1}^{N}c}+2$

$\boxed{\sum_{n=1}^{N}c+\frac{c}{n}}$

$\boxed{\sum_{n=1}^{N}c+n^2}+N$

$\boxed{\sum_{n=1}^{N}n}+\boxed{\sum_{k=1}^{N}k}$

$\boxed{\sum_{n=1}^{N}n+\boxed{\sum_{k=1}^{n}k}}$

$\boxed{\sum_{n=1}^{N}c+\boxed{\sum_{k=1}^{N}k}+n}$

Fig. 1: Example argument identifications for sums.

3.2 Lagrange's notation for differentiation and derivatives

Another remaining issue is the Lagrange (prime) notation for differentiation, since it does not outwardly provide sufficient semantic information. This notation presents two challenges. First, we do not know with respect to which variable the differentiation should be performed. Consider for example the Hurwitz zeta function $\zeta(s,a)$ [10, §25.11]. In the case of a differentiation $\zeta'(s,a)$, it is not clear if the function should be differentiated with respect to s or a. To remedy this issue, we analyzed all formulae in the DLMF which use prime notations and determined which variables (slots) for which functions represent the variables of the differentiation. Based on our analysis, we extended the translation patterns by meta information for semantic macros according to the slot of differentiation. For instance, in the case of the Hurwitz zeta function, the first slot is the slot for prime differentiation, i.e., $\zeta'(s,a)=\frac{d}{ds}\zeta(s,a)$. The identified variables of differentiations for the special functions in the DLMF can be considered to be the standard slots of differentiations, e.g., in other DML, $\zeta'(s,a)$ most likely refers to $\frac{d}{ds}\zeta(s,a)$.

The second challenge occurs if the slot of differentiation contains complex expressions rather than single symbols, e.g., $\zeta'(s^2,a)$. In this case, $\zeta'(s^2,a) = \frac{\mathrm{d}}{\mathrm{d}(s^2)}\zeta(s^2,a)$ instead of $\frac{\mathrm{d}}{\mathrm{d}s}\zeta(s^2,a)$. Since CAS often do not support derivatives with respect to complex expressions, we use the inbuilt substitution functions[16] in the CAS to overcome this issue. To do so, we use a temporary variable `temp` for the substitution. CAS perform substitutions from the inside to the outside. Hence, we can use the same temporary variable `temp` even for nested substitutions. Table 1 shows the translation performed for $\zeta'(s^2,a)$. CAS may provide optional arguments to calculate the derivatives for certain special functions, e.g., `Zeta(n,z,a)` in `Maple` for the n-th derivative of the Hurwitz zeta function. However, this shorthand notation is generally not supported (e.g., `Mathematica` does not define such an optional parameter). Our substitution approach is more lengthy but also more reliable. Unfortunately, lengthy expressions generally harm the performance of CAS, especially for symbolic manipulations. Hence, we have a genuine interest in keeping translations short, straightforward and readable. Thus, the substitution translation pattern is only triggered if the variable of differentiation is not a single identifier. Note that this substitution only triggers on semantic macros. Generic functions, including prime notations, are still skipped.

A related problem to MEOM of sums, products, integrals, limits, and differentiations are the notations of derivatives. The semantic macro for derivatives `\deriv{w}{x}` (rendered as $\frac{\mathrm{d}w}{\mathrm{d}x}$) is often used with an empty first argument to render the function behind the derivative notation, e.g., `\deriv{}{x}\sin@{x}` for $\frac{\mathrm{d}}{\mathrm{d}x}\sin x$. This leads to the same problem we faced above for identifying MEOMs. In this case, we use the same heuris-

Table 1: Example translations for the prime derivative of the Hurwitz zeta function with respect to s^2.

System	$\zeta'(s^2,a)$
DLMF	`\Hurwitzzeta'@{s^2}{a}`
Maple	`subs(temp=(s)^(2),diff(` `Zeta(0,temp,a),temp$(1)))`
Mathe-matica	`D[HurwitzZeta[temp,a],` `{temp,1}]/.temp->(s)^(2)`

tic as we did for sums, products, and limits. Note that derivatives may be written following the function argument, e.g., $\sin(x)\frac{\mathrm{d}}{\mathrm{d}x}$. If we are unable to identify any following summand that contains the variable of differentiation before we reach a termination symbol, we look for arguments prior to the derivative according to the heuristic (I-III).

Wronskians With the support of prime differentiation described above, we are also able to translate the Wronskian [10, (1.13.4)] to `Maple` and `Mathematica`. A translation requires one to identify the variable of differentiation from the elements of the Wronskian, e.g., z for $\mathscr{W}\{\mathrm{Ai}(z),\mathrm{Bi}(z)\}$ from [10, (9.2.7)]. We analyzed all Wronskians in the DLMF and discovered that most Wronskians have a special

---

[16] Note that `Maple` also support an evaluation substitution via the two-argument `eval` function. Since our substitution only triggers on semantic macros, we only use `subs` if the function is defined in `Maple`. In turn, as far as we know, there is no practical difference between `subs` and the two-argument `eval` in our case.

function in its argument—such as the example above. Hence, we can use our previously inserted metadata information about the slots of differentiation to extract the variable of differentiation from the semantic macros. If the semantic macro argument is a complex expression, we search for the identifier in the arguments that appear in both elements of the Wronskian. For example, in $\mathscr{W}\{\mathrm{Ai}(z^a),\zeta(z^2,a)\}$, we extract z as the variable since it is the only identifier that appears in the arguments z^a and z^2 of the elements. This approach is also used when there is no semantic macro involved, i.e., from $\mathscr{W}\{z^a,z^2\}$ we extract z as well. If LACAST extracts multiple candidates or none, it throws a translation exception.

4 Evaluation of the DLMF using CAS
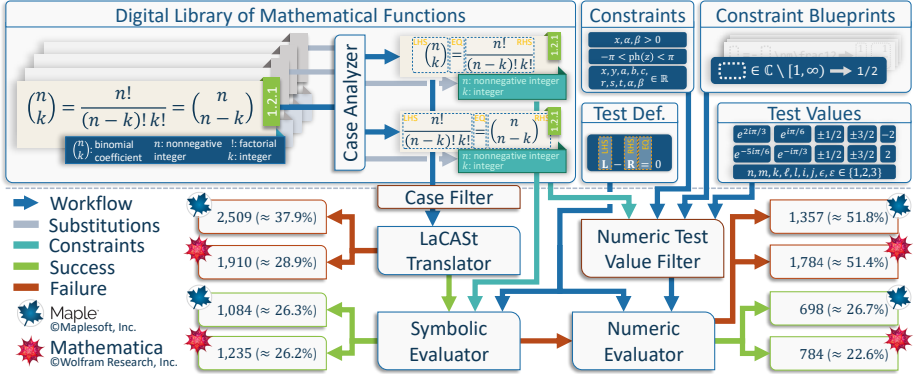
Fig. 2: The workflow of the evaluation engine and the overall results. Errors and abortions are not included. The generated dataset contains 9,977 equations. In total, the case analyzer splits the data into 10,930 cases of which 4,307 cases were filtered. This sums up to a set of 6,623 test cases in total.

For evaluating the DLMF with `Maple` and `Mathematica`, we follow the same approach as demonstrated in [7], i.e., we symbolically and numerically verify the equations in the DLMF with CAS. If a verification fails, symbolically and numerically, we identified an issue either in the DLMF, the CAS, or the verification pipeline. Note that an issue does not necessarily represent errors/bugs in the DLMF, CAS, or LACAST (see the discussion about branch cuts in Section B). Figure 2 illustrates the pipeline of the evaluation engine. First, we analyze every equation in the DLMF (hereafter referred to as test cases). A case analyzer splits multiple relations in a single line into multiple test cases. Note that only the adjacent relations are considered, i.e., with $f(z)=g(z)=h(z)$, we generate two test cases $f(z)=g(z)$ and $g(z)=h(z)$ but not $f(z)=h(z)$. In addition, expressions with $\pm$ and $\mp$ are split accordingly, e.g., $i^{\pm i}=\mathrm{e}^{\mp\pi/2}$ [10, (4.4.12)] is split into $i^{+i}=\mathrm{e}^{-\pi/2}$ and $i^{-i}=\mathrm{e}^{+\pi/2}$. The analyzer utilizes the attached additional information in each line, i.e., the URL in the DLMF, the used and defined symbols, and the constraints. If a used symbol is defined elsewhere in the DLMF, it performs substitutions. For example, the multi-equation [10, (9.6.2)] is split into six test cases and every ζ is replaced by $\frac{2}{3}z^{3/2}$ as defined in [10, (9.6.1)]. The substitution is performed on the parse tree of expressions [14]. A definition is

only considered as such, if the defining symbol is identical to the equation's left-hand side. That means, $z = (\frac{3}{2}\zeta)^{3/2}$ [10, (9.6.10)] is not considered as a definition for ζ. Further, semantic macros are never substituted by their definitions. Translations for semantic macros are exclusively defined by the authors. For example, the equation [10, (11.5.2)] contains the Struve $\mathbf{K}_\nu(z)$ function. Since `Mathematica` does not contain this function, we defined an alternative translation to its definition $\mathbf{H}_\nu(z) - Y_\nu(z)$ in [10, (11.2.5)] with the Struve function $\mathbf{H}_\nu(z)$ and the Bessel function of the second kind $Y_\nu(z)$, because both of these functions are supported by `Mathematica`. The second entry in Table 3 in the Appendix D shows the translation for this test case.

Next, the analyzer checks for additional constraints defined by the used symbols recursively. The mentioned Struve $\mathbf{K}_\nu(z)$ test case [10, (11.5.2)] contains the Gamma function. Since the definition of the Gamma function [10, (5.2.1)] has a constraint $\Re z > 0$, the numeric evaluation must respect this constraint too. For this purpose, the case analyzer first tries to link the variables in constraints to the arguments of the functions. For example, the constraint $\Re z > 0$ sets a constraint for the first argument z of the Gamma function. Next, we check all arguments in the actual test case at the same position. The test case contains $\Gamma(\nu + 1/2)$. In turn, the variable z in the constraint of the definition of the Gamma function $\Re z > 0$ is replaced by the actual argument used in the test case. This adds the constraint $\Re(\nu + 1/2) > 0$ to the test case. This process is performed recursively. If a constraint does not contain any variable that is used in the final test case, the constraint is dropped.

In total, the case analyzer would identify four additional constraints for the test case [10, (11.5.2)]. Table 3 in the Appendix D shows the applied constraints (including the directly attached constraint $\Re z > 0$ and the manually defined global constraints from Figure 3). Note that the constraints may contain variables that do not appear in the actual test case, such as $\Re \nu + k + 1 > 0$. Such constraints do not have any effect on the evaluation because if a constraint cannot be computed to true or false, the constraint is ignored. Unfortunately, this recursive loading of additional constraints may generate impossible conditions in certain cases, such as $|\Gamma(iy)|$ [10, (5.4.3)]. There are no valid real values of y such that $\Re(iy) > 0$. In turn, every test value would be filtered out, and the numeric evaluation would not verify the equation. However, such cases are the minority and we were able to increase the number of correct evaluations with this feature.

To avoid a large portion of incorrect calculations, the analyzer filters the dataset before translating the test cases. We apply two filter rules to the case analyzer. First, we filter expressions that do not contain any semantic macros. Due to the limitations of LᴬCᴀsT, these expressions most likely result in wrong translations. Further, it filters out several meaningless expressions that are not verifiable, such as $z = x$ in [10, (4.2.4)]. The result dataset flag these cases with '*Skipped - no semantic math*'. Note that the result dataset still contains the translations for these cases to provide a complete picture of the DLMF. Second, we filter expressions that contain ellipsis[17] (e.g., `\cdots`), approximations, and asymptotics (e.g., $\mathcal{O}(z^2)$) since those expressions cannot be evaluated with the proposed approach. Further, a definition is skipped if it is not a definition of a semantic macro, such as [10, (2.3.13)], because definitions without

---

[17] Note that we filter out ellipsis (e.g., `\cdots`) but not single dots (e.g., `\cdot`).

an appropriate counterpart in the CAS are meaningless to evaluate. Definitions of semantic macros, on the other hand, are of special interest and remain in the test set since they allow us to test if a function in the CAS obeys the actual mathematical definition in the DLMF. If the case analyzer (see Figure 2) is unable to detect a relation, i.e., split an expression on $<, \leq, \geq, >, =$, or $\neq$, the line in the dataset is also skipped because the evaluation approach relies on relations to test. After splitting multi-equations (e.g., $\pm, \mp, a = b = c$), filtering out all non-semantic expressions, non-semantic macro definitions, ellipsis, approximations, and asymptotics, we end up with 6,623 test cases in total from the entire DLMF.

After generating the test case with all constraints, we translate the expression to the CAS representation. Every successfully translated test case is then symbolically verified, i.e., the CAS tries to simplify the difference of an equation to zero. Non-equation relations simplifies to Booleans. Non-simplified expressions are verified numerically for manually defined test values, i.e., we calculate actual numeric values for both sides of an equation and check their equivalence.

4.1 Symbolic Evaluation

The symbolic evaluation was performed for `Maple` as in [7]. However, we use the newer version `Maple` 2020. Another feature we added to LACAST is the support of packages in `Maple`. Some functions are only available in modules (packages) that must be preloaded, such as `QPochhammer` in the package `QDifferenceEquations`[18]. The general `simplify` method in `Maple` does not cover q-hypergeometric functions. Hence, whenever LACAST loads functions from the q-hyper-geometric package, the better performing `QSimplify` method is used. With the WED and the new support for `Mathematica` in LACAST, we perform the symbolic and numeric tests for `Mathematica` as well. The symbolic evaluation in `Mathematica` relies on the full simplification[19]. For `Maple` and `Mathematica`, we defined the global assumptions $x, y \in \mathbb{R}$ and $k, n, m \in \mathbb{N}$. Constraints of test cases are added to their assumptions to support simplification. Adding more global assumptions for symbolic computation generally harms the performance since CAS internally uses assumptions for simplifications. It turned out that by adding more custom assumptions, the number of successfully simplified expressions decreases.

4.2 Numerical Evaluation

Defining an accurate test set of values to analyze an equivalence can be an arbitrarily complex process. It would make sense that every expression is tested on specific values according to the containing functions. However, this laborious process is not suitable for evaluating the entire DML and CAS. It makes more sense to develop a general set of test values that (i) generally covers interesting domains and (ii) avoid singularities, branch cuts, and similar problematic regions. Considering these two attributes, we come up with the ten test points illustrated in Figure 3. It contains four complex values on the unit circle and six points on the real axis. The test values cover the

---

general area of interest (complex values in all four quadrants, negative and positive real values) and avoid the typical singularities at $\{0,\pm 1,\pm i\}$. In addition, several variables are tied to specific values for entire sections. Hence, we applied additional global constraints to the test cases.

The numeric evaluation engine heavily relies on the performance of extracting free variables from an expression. Unfortunately, the inbuilt functions in CAS, if available, are not very reliable. As the authors explained in [7], a custom algorithm within `Maple` was necessary to extract identifiers. `Mathematica` has the undocumented function `Reduce'FreeVariables` for this purpose. However, both systems, the custom solution in `Maple` and the inbuilt `Mathematica` function,

Test Values $\Im$ **Special Test Values**
$n,m,k,\ell,l,i,j,\epsilon,\varepsilon \in \{1,2,3\}$

Global Constraints
$x,\alpha,\beta > 0$
$-\pi < \mathrm{ph}(z) < \pi$
$x,y,a,b,c,r,s,t,\alpha,\beta \in \mathbb{R}$

Fig. 3: The ten numeric test values in the complex plane for general variables. The dashed line represents the unit circle $|z| = 1$. At the right, we show the set of values for special variable values and general global constraints. On the right, i is referring to a generic variable and not to the imaginary unit.

have problems distinguishing free variables of entire expressions from the bound variables in MEOMs, e.g., integration and continuous variables. `Mathematica` sometimes does not extract a variable but returns the unevaluated input instead. We regularly faced this issue for integrals. However, we discovered one example without integrals. For `EulerE[n,0]` from [10, (24.4.26)], we expected to extract $\{n\}$ as the set of free variables but instead received a set of the unevaluated expression itself $\{$`EulerE[n,0]`$\}$[20]. Since the extended version of LᴬCᴀsT handles operators, including bound variables of MEOMs, we drop the use of internal methods in CAS and extend LᴬCᴀsT to extract identifiers from an expression. During a translation process, LᴬCᴀsT tags every single identifier as a variable, as long as it is not an element of a MEOM. This simple approach proves to be very efficient since it is implemented alongside the translation process itself and is already more powerful as compared to the existing inbuilt CAS solutions. We defined subscripts of identifiers as a part of the identifier, e.g., z_1 and z_2 are extracted as variables from $z_1 + z_2$ rather than z.

The general pipeline for a numeric evaluation works as follows. First, we replace all substitutions and extract the variables from the left- and right-hand sides of the test expression via LᴬCᴀsT. For the previously mentioned example of the Struve function [10, (11.5.2)], LᴬCᴀsT identifies two variables in the expression, ν and z. According to the values in Figure 3, ν and z are set to the general ten values. A numeric test contains every combination of test values for all variables. Hence, we generate 100 test calculations for [10, (11.5.2)]. Afterward, we filter the test values that violate the attached constraints. In the case of the Struve function, we end up with 25 test cases.

In addition, we apply a limit of 300 calculations for each test case and abort a computation after 30 seconds due to computational limitations. If the test case generates more than 300 test values, only the first 300 are used. Finally, we calculate

---

[20] The bug was reported to and confirmed by Wolfram Research Version 12.0.

the result for every remaining test value, i.e., we replace every variable by their value and calculate the result. The replacement is done by `Mathematica`'s `ReplaceAll` method because the more appropriate method `With`, for unknown reasons, does not always replace all variables by their values. We wrap test expressions in `Normal` for numeric evaluations to avoid conditional expressions, which may cause incorrect calculations (see Section 5.1 for a more detailed discussion of conditional outputs). After replacing variables by their values, we trigger numeric computation. If the absolute value of the result (i.e., the difference between left- and right-hand side of the equation) is below the defined threshold of 0.001 or true (in the case of inequalities), the test calculation is considered successful. A numeric test case is only considered successful if and only if every test calculation was successful. If a numeric test case fails, we store the information on which values it failed and how many of these were successful.

5   Results

The translations to `Maple` and `Mathematica`, the symbolic results, the numeric computations, and an overview PDF of the reported bugs to `Mathematica` are available online on our demopage. In the following, we mainly focus on `Mathematica` because of page limitations and because `Maple` has been investigated more closely by [7]. The results for `Maple` are also available online. Compared to the baseline ($\approx 31\%$), our improvements doubled the amount translations ($\approx 62\%$) for `Maple` and reach $\approx 71\%$ for `Mathematica`. The majority of expressions that cannot be translated contain macros that have no adequate translation pattern to the CAS, such as the macros for interval Weierstrass lattice roots [10, §23.3(i)] and the multivariate hypergeometric function [10, (19.16.9)]. Other errors (6% for `Maple` and `Mathematica`) occur for several reasons. For example, out of the 418 errors in translations to `Mathematica`, 130 caused an error because the MEOM of an operator could not be extracted, 86 contained prime notations that do not refer to differentiations, 92 failed because of the underlying LaTeX parser [34], and in 46 cases, the arguments of a DLMF macro could not be extracted.

Out of 4,713 translated expressions, 1,235 (26.2%) were successfully simplified by `Mathematica` (1,084 of 4,114 or 26.3% in `Maple`). For `Mathematica`, we also count results that are equal to 0 under certain conditions as successful (called `ConditionalExpression`). We identified 65 of these conditional results: 15 of the conditions are equal to constraints that were provided in the surrounding text but not in the info box of the DLMF equation; 30 were produced due to branch cut issues (see Section B in the Appendix); and 20 were the same as attached in the DLMF but reformulated, e.g., $z \in \mathbb{C} \backslash (1, \infty)$ from [10, (25.12.2)] was reformulated to $\Im z \neq 0 \vee \Re z < 1$. The remaining translated but not symbolically verified expressions were numerically evaluated for the test values in Figure 3. For the 3,474 cases, 784 (22.6%) were successfully verified numerically by `Mathematica` (698 of 2,618 or 26.7% by `Maple`[21]). For 1,784 the numeric evaluation failed. In the evaluation process, 655

---

[21] Due to computational issues, 120 cases must have been skipped manually. 292 cases resulted in an error during symbolic verification and, therefore, were skipped also for numeric evaluations. Considering these skipped cases as failures, decreases the numerically verified cases to 23% in `Maple`.

computations timed out and 180 failed due to errors in `Mathematica`. Of the 1,784 failed cases, 691 failed partially, i.e., there was at least one successful calculation among the tested values. For 1,091 all test values failed. Table 3 in the Appendix D shows the results for three sample test cases. The first case is a false positive evaluation because of a wrong translation. The second case is valid, but the numeric evaluation failed due to a bug in Mathematica (see next subsection). The last example is valid and was verified numerically but was too complex for symbolic verifications.

5.1   Error Analysis

The numeric tests' performance strongly depends on the correct attached and utilized information. The first example in Table 3 in the Appendix D illustrates the difficulty of the task on a relatively easy case. Here, the argument of f was not explicitly given, such as in $f(x)$. Hence, LACAST translated f as a variable. Unfortunately, this resulted in a false verification symbolically and numerically. This type of error mostly appears in the first three chapters of the DLMF because they use generic functions frequently. We hoped to skip such cases by filtering expressions without semantic macros. Unfortunately, this derivative notation uses the semantic macro `deriv`. In the future, we filter expressions that contain semantic macros that are not linked to a special function or orthogonal polynomial.

As an attempt to investigate the reliability of the numeric test pipeline, we can run numeric evaluations on symbolically verified test cases. Since `Mathematica` already approved a translation symbolically, the numeric test should be successful if the pipeline is reliable. Of the 1,235 symbolically successful tests, only 94 (7.6%) failed numerically. None of the failed test cases failed entirely, i.e., for every test case, at least one test value was verified. Manually investigating the failed cases reveal 74 cases that failed due to an `Indeterminate` response from `Mathematica` and 5 returned `infinity`, which clearly indicates that the tested numeric values were invalid, e.g., due to testing on singularities. Of the remaining 15 cases, two were identical: [10, (15.9.2)] and [10, (18.5.9)]. This reduces the remaining failed cases to 14. We evaluated invalid values for 12 of these because the constraints for the values were given in the surrounding text but not in the info boxes. The remaining 2 cases revealed a bug in `Mathematica` regarding conditional outputs (see below). The results indicate that the numeric test pipeline is reliable, at least for relatively simple cases that were previously symbolically verified. The main reason for the high number of failed numerical cases in the entire DLMF (1,784) are due to missing constraints in the i-boxes and branch cut issues (see Section B in the Appendix), i.e., we evaluated expressions on invalid values.

Bug reports `Mathematica` has trouble with certain integrals, which, by default, generate conditional outputs if applicable. With the method `Normal`, we can suppress conditional outputs. However, it only hides the condition rather than evaluating the expression to a non-conditional output. For example, integral expressions in [10, (10.9.1)] are automatically evaluated to the Bessel function $J_0(|z|)$ for the condition[22] $z \in \mathbb{R}$ rather than $J_0(z)$ for all $z \in \mathbb{C}$. Setting the `GenerateConditions`[23] option

---

[22] $J_0(x)$ with $x \in \mathbb{R}$ is even. Hence, $J_0(|z|)$ is correct under the given condition.

[23] https://reference.wolfram.com/language/ref/GenerateConditions.html [accessed 09/01/2021]

to `None` does not change the output. `Normal` only hides $z \in \mathbb{R}$ but still returns $J_0(|z|)$. To fix this issue, for example in (10.9.1) and (10.9.4), we are forced to set `GenerateConditions` to false.

Setting `GenerateConditions` to false, on the other hand, reveals severe errors in several other cases. Consider $\int_z^\infty t^{-1}\mathrm{e}^{-t}\mathrm{d}t$ [10, (8.4.4)], which gets evaluated to $\Gamma(0,z)$ but (condition) for $\Re z > 0 \wedge \Im z = 0$. With `GenerateConditions` set to false, the integral incorrectly evaluates to $\Gamma(0,z) + \ln(z)$. This happened with the 2 cases mentioned above. With the same setting, the difference of the left- and right-hand sides of [10, (10.43.8)] is evaluated to 0.398942 for $x, \nu = 1.5$. If we evaluate the same expression on $x, \nu = \frac{3}{2}$ the result is `Indeterminate` due to `infinity`. For this issue, one may use `NIntegrate` rather than `Integrate` to compute the integral. However, evaluating via `NIntegrate` decreases the number of successful numeric evaluations in general. We have revealed errors with conditional outputs in (8.4.4), (10.22.39), (10.43.8-10), and (11.5.2) (in [10]). In addition, we identified one critical error in `Mathematica`. For [10, (18.17.47)], WED (`Mathematica`'s kernel) ran into a *segmentation fault (core dumped)* for $n > 1$. The kernel of the full version of `Mathematica` gracefully died without returning an output[24].

Besides `Mathematica`, we also identified several issues in the DLMF. None of the newly identified issues were critical, such as the reported sign error from the previous project [7], but generally refer to missing or wrong attached semantic information. With the generated results, we can effectively fix these errors and further semantically enhance the DLMF. For example, some definitions are not marked as such, e.g., $Q(z) = \int_0^\infty \mathrm{e}^{-zt} q(t)\mathrm{d}t$ [10, (2.4.2)]. In [10, (10.24.4)], ν must be a real value but was linked to a *complex parameter* and x should be positive real. An entire group of cases [10, (10.19.10-11)] also discovered the incorrect use of semantic macros. In these formulae, $P_k(a)$ and $Q_k(a)$ are defined but had been incorrectly marked up as Legendre functions going all the way back to DLMF Version 1.0.0 (May 7, 2010). In some cases, equations are mistakenly marked as definitions, e.g., [10, (9.10.10)] and [10, (9.13.1)] are annotated as local definitions of n. We also identified an error in LᴬCᴀsT, which incorrectly translated the exponential integrals $E_1(z)$, $\mathrm{Ei}(x)$ and $\mathrm{Ein}(z)$ (defined in [10, §6.2(i)]). A more explanatory overview of discovered, reported, and fixed issues in the DLMF, `Mathematica`, and `Maple` is provided in the Appendix C.

6   Conclusion

We have presented a novel approach to verify the theoretical digital mathematical library DLMF with the power of two major general-purpose computer algebra systems `Maple` and `Mathematica`. With LᴬCᴀsT, we transformed the semantically enhanced LᴬTᴇX expressions from the DLMF to each CAS. Afterward, we symbolically and numerically evaluated the DLMF expressions in each CAS. Our results are auspicious and provide useful information to maintain and extend the DLMF efficiently. We further identified several errors in `Mathematica`, `Maple` [7], the DLMF, and the transformation tool LᴬCᴀsT, proving the profit of the presented verification approach.

---

[24] All errors were reported to and partially confirmed by Wolfram Research. See Appendix C for more information.

Further, we provide open access to all results, including translations and evaluations[25]. and to the source code of LaCASt[26].

The presented results show a promising step towards an answer for our initial research question. By translating an equation from a DML to a CAS, automatic verifications of that equation in the CAS allows us to detect issues in either the DML source or the CAS implementation. Each analyzed failed verification successively improves the DML or the CAS. Further, analyzing a large number of equations from the DML may be used to finally verify a CAS. In addition, the approach can be extended to cover other DML and CAS by exploiting different translation approaches, e.g., via MATHML [31] or OpenMath [18].

Nonetheless, the analysis of the results, especially for an entire DML, is cumbersome. Minor missing semantic information, e.g., a missing constraint or not respected branch cut positions, leads to a relatively large number of false positives, i.e., unverified expressions correct in the DML and the CAS. This makes a generalization of the approach challenging because all semantics of an equation must be taken into account for a trustworthy evaluation. Furthermore, evaluating equations on a small number of discrete values will never provide sufficient confidence to verify a formula, which leads to an unpredictable number of true negatives, i.e., erroneous equations that pass all tests. A more sophisticated selection of critical values or other numeric tools with automatic results verification (such as variants of Newton's interval method) potentially mitigates this issue in the future. After all, we conclude that the approach provides valuable information to complement, improve, and maintain the DLMF, `Maple`, and `Mathematica`. A trustworthy verification, on the other hand, might be out of reach.

6.1 Future Work

The resulting dataset provides valuable information about the differences between CAS and the DLMF. These differences had not been largely studied in the past and are worthy of analysis. Especially a comprehensive and machine-readable list of branch cut positioning in different systems is a desired goal [9]. Hence, we will continue to work closely together with the editors of the DLMF to improve further and expand the available information on the DLMF. Finally, the numeric evaluation approach would benefit from test values dependent on the actual functions involved. For example, the current layout of the test values was designed to avoid problematic regions, such as branch cuts. However, for identifying differences in the DLMF and CAS, especially for analyzing the positioning of branch cuts, an automatic evaluation of these particular values would be very beneficial and can be used to collect a comprehensive, inter-system library of branch cuts. Therefore, we will further study the possibility of linking semantic macros with numeric regions of interest.

---

[25] https://lacast.wmflabs.org [accessed 01/01/2022]
[26] https://github.com/ag-gipp/LaCASt [accessed 04/01/2022]

References

1. Aguirregabiria, J.M., Hernández, A.M., Rivas, M.: Are we careful enough when using computer algebra? Computers in Physics **8**(1), 56–61 (1994). https://doi.org/10.1063/1.4823260

2. Barnett, M., Chang, B.Y.E., DeLine, R., Jacobs, B., Leino, K.R.M.: Boogie: A modular reusable verifier for object-oriented programs. In: Formal Methods for Components and Objects, pp. 364–387. Springer Berlin Heidelberg (2006). https://doi.org/10.1007/11804192_17

3. Bertot, Y., Castéran, P.: Interactive Theorem Proving and Program Development - Coq´Art: The Calculus of Inductive Constructions. Texts in Theoretical Computer Science. An EATCS Series, Springer Berlin Heidelberg (2004)

4. Bobot, F., Filliâtre, J.C., Marché, C., Paskevich, A.: Why3: Shepherd your herd of provers. Boogie 2011: First International Workshop on Intermediate Verification Languages pp. 53–64 (5 2011), https://hal.inria.fr/hal-00790310/document

5. Boulmé, S., Hardin, T., Hirschkoff, D., Ménissier-Morain, V., Rioboo, R.: On the way to certify computer algebra systems. Electronic Notes in Theoretical Computer Science **23**(3), 370–385 (1999). https://doi.org/10.1016/S1571-0661(05)80609-7, cALCULEMUS 99, Systems for Integrated Computation and Deduction (associated to FLoC'99, the 1999 Federated Logic Conference)

6. Carette, J., Kucera, M.: Partial evaluation of Maple. Science of Computer Programming **76**(6), 469–491 (6 2011). https://doi.org/10.1016/j.scico.2010.12.001

7. Cohl, H.S., Greiner-Petter, A., Schubotz, M.: Automated symbolic and numerical testing of DLMF formulae using computer algebra systems. In: Intelligent Computer Mathematics CICM. vol. 11006, pp. 39–52. Springer (2018). https://doi.org/10.1007/978-3-319-96812-4_4

8. Cohl, H.S., Schubotz, M., Youssef, A., Greiner-Petter, A., Gerhard, J., Saunders, B.V., McClain, M.A., Bang, J., Chen, K.: Semantic preserving bijective mappings of mathematical formulae between document preparation systems and computer algebra systems. In: Intelligent Computer Mathematics CICM. pp. 115–131. Springer (2017). https://doi.org/10.1007/978-3-319-62075-6_9

9. Corless, R.M., Jeffrey, D.J., Watt, S.M., Davenport, J.H.: "According to Abramowitz and Stegun" or arccoth needn't be uncouth. SIGSAM Bulletin **34**(2), 58–65 (2000). https://doi.org/10.1145/362001.362023

10. DLMF: *NIST Digital Library of Mathematical Functions*. https://dlmf.nist.gov/, Release 1.1.4 of 2022-01-15, F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds.

11. Durán, A.J., Pérez, M., Varona, J.L.: The misfortunes of a trio of mathematicians using computer algebra systems. Can we trust in them? Notices of the AMS **61**(10), 1249–1252 (2014)

12. Elphick, D., Leuschel, M., Cox, S.: Partial evaluation of MATLAB. In: Gen. Prog. and Component Eng., pp. 344–363. Springer (2003). https://doi.org/10.1007/978-3-540-39815-8_21

13. Greiner-Petter, A., Schubotz, M., Aizawa, A., Gipp, B.: Making presentation math computable: Proposing a context sensitive approach for translating LaTeX to computer algebra systems. In: International Congress of Mathematical Software (ICMS). Lecture Notes in Computer Science, vol. 12097, pp. 335–341. Springer (2020). https://doi.org/10.1007/978-3-030-52200-1_33

14. Greiner-Petter, A., Schubotz, M., Cohl, H.S., Gipp, B.: Semantic preserving bijective mappings for expressions involving special functions between computer algebra systems and document preparation systems. Aslib Journal of Information Management **71**(3), 415–439 (2019). https://doi.org/10.1108/AJIM-08-2018-0185

15. Greiner-Petter, A., Schubotz, M., Müller, F., Breitinger, C., Cohl, H.S., Aizawa, A., Gipp, B.: Discovering mathematical objects of interest - A study of mathematical notations. In: WWW. pp. 1445–1456. ACM / IW3C2 (2020). https://doi.org/10.1145/3366423.3380218

16. Harrison, J.: HOL Light: A tutorial introduction. In: Srivas, M., Camilleri, A. (eds.) Formal Methods in Computer-Aided Design (FMCAD). Lecture Notes in Computer Science, vol. 1166, pp. 265–269. Springer Berlin Heidelberg. https://doi.org/10.1007/BFb0031814

17. Harrison, J.R., Théry, L.: A skeptic's approach to combining HOL and Maple **21**(3), 279–294. https://doi.org/10.1023/A:1006023127567

18. Heras, J., Pascual, V., Rubio, J.: Using open mathematical documents to interface computer algebra and proof assistant systems. In: Intelligent Computer Mathematics MKM at CICM. Lecture Notes in Computer Science, vol. 5625, pp. 467–473. Springer (2009). https://doi.org/10.1007/978-3-642-02614-0_37

19. Hickman, T., Laursen, C.P., Foster, S.: Certifying differential equation solutions from computer algebra systems in Isabelle/HOL http://arxiv.org/abs/2102.02679

20. Kaliszyk, C., Wiedijk, F.: Certified computer algebra on top of an interactive theorem prover. In: Towards Mechanized Math. Assist., pp. 94–105. Springer (2007). https://doi.org/10.1007/978-3-540-73086-6_8

21. Khan, M.T.: Formal Specification and Verification of Computer Algebra Software. phdthesis, Johannes Kepler University Linz (Apr 2014)

22. Kristianto, G.Y., Topić, G., Aizawa, A.: Utilizing dependency relationships between math expressions in math IR. Information Retrieval Journal **20**(2), 132–167 (3 2017). https://doi.org/10.1007/s10791-017-9296-8

23. Lambán, L., Rubio, J., Martín-Mateos, F.J., Ruiz-Reina, J.L.: Verifying the bridge between simplicial topology and algebra: the Eilenberg-Zilber algorithm. Logic Journal of IGPL **22**(1), 39–65 (8 2013). https://doi.org/10.1093/jigpal/jzt034

24. Lee, W., Sharma, R., Aiken, A.: On automatically proving the correctness of math.h implementations. Proc. ACM on Prog. Lang. (POPL) **2**(47), 1–32 (2018). https://doi.org/10.1145/3158135

25. Leino, K.R.M.: Program proving using intermediate verification languages (IVLs) like Boogie and Why3. ACM SIGAda Ada Letters **32**(3), 25–26 (11 2012). https://doi.org/10.1145/2402709.2402689

26. Lewis, R.H., Wester, M.: Comparison of polynomial-oriented computer algebra systems. SIGSAM Bull. **33**(4), 5–13 (12 1999). https://doi.org/10.1145/500457.500459

27. Miller, B.R., Youssef, A.: Technical aspects of the digital library of mathematical functions. Ann. Math. Artif. Intell. **38**(1-3), 121–136 (2003). https://doi.org/10.1023/A:1022967814992

28. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL - A Proof Assistant for Higher-Order Logic, Lecture Notes in Computer Science, vol. 2283. Springer Berlin Heidelberg (2002). https://doi.org/10.1007/3-540-45949-9

29. Parisse, B.: Compiling LATEX to computer algebra-enabled HTML5 http://arxiv.org/abs/1707.01271

30. Prieto, H., Dalmas, S., Papegay, Y.: Mathematica as an OpenMath application **34**(2), 22–26. https://doi.org/10.1145/362001.362016

31. Schubotz, M., Greiner-Petter, A., Scharpf, P., Meuschke, N., Cohl, H.S., Gipp, B.: Improving the representation and conversion of mathematical formulae by considering their textual context. In: ACM/IEEE JCDL. pp. 233–242. ACM (2018). https://doi.org/10.1145/3197026.3197058

32. Schubotz, M., Grigorev, A., Leich, M., Cohl, H.S., Meuschke, N., Gipp, B., Youssef, A.S., Markl, V.: Semantification of identifiers in mathematics for better math information retrieval. In: ACM SIGIR'16. pp. 135–144. ACM Press (2016). https://doi.org/10.1145/2911451.2911503

33. Shan, R., Youssef, A.: Towards math terms disambiguation using machine learning. In: Kamareddine, F., Sacerdoti Coen, C. (eds.) Proceedings of the International Conference on Intelligent Computer Mathematics (CICM). Lecture Notes in Computer Science, vol. 12833, pp. 90–106. Springer. https://doi.org/10.1007/978-3-030-81097-9_7

34. Youssef, A.: Part-of-math tagging and applications. In: Intelligent Computer Mathematics CICM. Lecture Notes in Computer Science, vol. 10383, pp. 356–374. Springer (2017). https://doi.org/10.1007/978-3-319-62075-6_25

35. Youssef, A., Miller, B.R.: A contextual and labeled math-dataset derived from NIST's DLMF. In: Intelligent Computer Mathematics CICM. Lecture Notes in Computer Science, vol. 12236, pp. 324–330. Springer (2020). https://doi.org/10.1007/978-3-030-53518-6_25

36. Zanibbi, R., Oard, D.W., Agarwal, A., Mansouri, B.: Overview of ARQMath 2020: CLEF lab on answer retrieval for questions on math. In: CLEF. Lecture Notes in Computer Science, vol. 12260, pp. 169–193. Springer (2020). https://doi.org/10.1007/978-3-030-58219-7_15