



Towards Semantic Markup of Mathematical Documents via User Interaction

Luka Vrečar^(✉), Joe Wells, and Fairouz Kamareddine

Heriot-Watt University, Edinburgh, UK
{lv21,F.D.Kamareddine}@hw.ac.uk

Abstract. Mathematical documents written in \LaTeX often contain ambiguities. We can resolve some of them via semantic markup using, e.g., \STeX , which also has other potential benefits, such as interoperability with computer algebra systems, proof systems, and increased accessibility. However, semantic markup is more involved than “regular” typesetting and presents a challenge for authors of mathematical documents. We aim to smooth out the transition from plain \LaTeX to semantic markup by developing semi-automatic tools for authors. In this paper we present an approach to semantic markup of formulas by (semi-)automatically generating grammars from existing \STeX macro definitions and parsing mathematical formulas with them. We also present a GUI-based tool for the disambiguation of parse results and showcase its functionality and potential using a grammar for parsing untyped λ -terms.

1 Introduction

Formulas in mathematical documents written in \LaTeX are often ambiguous. We can disambiguate them using \STeX [7, 8], a package for *semantic markup* of mathematical documents using *semantic macros*. This allows the explicit encoding of the semantics of formulas without changing their typeset appearance.

Example 1 (The ambiguity of $P \times Q$). The formula $P \times Q$ is usually typeset using the ambiguous typesetting `P \times Q` (where the meaning of `\times` depends on the type of P and Q). Instead one can use the semantic macro `\cart{P,Q}` to denote the **Cartesian product** of two sets, or `\matrixtimes[x]{P,Q}` to denote the multiplication of two matrices.

We used \STeX to typeset the example above. Note that this also generates hyperlinks to online resources¹. Semantic markup can facilitate clearer communication with readers, e.g., students or people with disabilities. It can also

¹ Coloured text denotes such hyperlinks (opens a new browser tab or jumps to the relevant definition in the PDF). Blue is used for highlighting notational components of macros, and teal is used for textual references to macros. Magenta colouring is used in the online definitions to highlight the defining occurrence of a symbol.

support connections between human-readable documents and proof systems (for formal verification), computer algebra systems (for computation), etc. Semantically marked up documents can be dynamic² and adapt to user preferences and interaction.

Example 2 (Semantic markup for encoding the structure of a formula).

Consider the formula $\sin x/y - z$ [6], typeset as `\sin x/y-z`. This can be interpreted as $\sin(\frac{x}{y}) - z$, $\sin(\frac{x}{y-z})$, or even $\frac{\sin x}{y-z}$. A reader can disambiguate this with sufficient background knowledge and context, e.g. some of the readings can be nonsensical if $y - z = 0$. Additionally, implicit parentheses can be “filled in” based on prior reader knowledge (in the case of \sin , we usually write it without parentheses when the argument is a monomial).

We could use, e.g., `\realminus{\sine{\realdive{x}{y}},z}`, to reflect the intended structure using $\mathsf{\STeX}$, which would be typeset as $\sin x/y - z$ and can be read equivalently to $\sin(\frac{x}{y}) - z$. This precisely encodes the intended structure of the formula while also offering all of the advantages of $\mathsf{\STeX}$, such as hyperlinks to online resources. In an active document, we can then add ways to show the structure of the formula (e.g., by rendering it as a tree), or automatically insert brackets for disambiguation³.

Semantic markup clearly has advantages over plain $\mathsf{\LaTeX}$, but producing documents via $\mathsf{\STeX}$ is more involved. Our goal is to facilitate a smooth transition to $\mathsf{\STeX}$ for new users, by devising ways to (semi-)automatically add semantic markup to $\mathsf{\LaTeX}$ documents. We will refer to this as $\mathsf{\STeX}$ -ification, or $\mathsf{\STeX}$ -ifying a document.

The initial focus of our work are documents which do not contain semantic annotations. We sampled a number of papers from a number of areas of mathematics and computer science and found a complete absence of macros that record the meaning or structure of formulas being used. Our conclusion was that a significant number of documents does not contain semantic annotations, so it made sense to focus on those for now, and explore handling documents with existing (author-defined) semantic annotations in future work.

Authors of documents without semantic annotations (especially those involved in teaching) likely do not have experience with them, but could benefit from the features built around $\mathsf{\STeX}$ such as the hyperlinks to definitions, definitions on hover⁴ and ALeA [2]. By offloading the actual semantic annotating (i.e., rewriting of formulas to use semantic macros) we aim to provide the benefits of semantic annotations, without authors having to learn $\mathsf{\STeX}$ beyond the very basics (compiling documents with $\mathsf{\STeX}$, and defining new semantic macros/notations). We also aim to provide a way of annotating a document for authors who are not familiar with fast text-editing techniques like Emacs macros.

² Not necessarily in PDF, but other formats like HTML support dynamic documents with ease.

³ Automatic bracketing is already a feature in $\mathsf{\STeX}$, see [7], section 7.4.1.

⁴ In HTML documents, which can be generated from a $\mathsf{\LaTeX}$ source using RusTeX [10].

Table 1. The $\text{\texttt{S}\TeX}$ -ification process

1. Generating all the prerequisites
 - (a) Manually identify all the semantic macros required, and define any new macros that might not be available in existing collections.
 - (b) Create a grammar for parsing all the formulas. This can be done manually, or (semi-)automatically (as described in Section 5).
2. Actually $\text{\texttt{S}\TeX}$ -ifying the document
 - (a) Parse each formula in the $\text{\texttt{L}\TeX}$ source with the grammar from 1b.
 - (b) Disambiguate any parses by prompting the user to select the correct parse tree via a graphical interface we present in Section 4,
 - (c) Convert all the parse trees into $\text{\texttt{S}\TeX}$ macros.
 - (d) Create a copy of the original document, with each formula replaced by its $\text{\texttt{S}\TeX}$ counterpart.

This paper explores (semi-)automatic semantic annotating and the problems that must be solved to make our approach feasible for large documents (e.g., lecture notes or papers). We propose a method for automatically generating grammars from $\text{\texttt{S}\TeX}$ macro definitions, parsing formulas with them, and using a graphical user interface (GUI) for disambiguation. In Sect. 2, we describe $\text{\texttt{S}\TeX}$ in more detail, and also introduce the untyped λ -calculus, which we are using as a testing ground for our methods.

We envision $\text{\texttt{S}\TeX}$ -ification as a two-stage process (see Table 1). Section 3 describes the first stage, which involves identifying the macros we need to $\text{\texttt{S}\TeX}$ -ify a given document (stage 1a), and creating a grammar for parsing the formulas within it (stage 1b). The grammar can be manually specified or automatically generated (see Sect. 5) based on the users' selection of relevant $\text{\texttt{S}\TeX}$ modules⁵. We also present the `parglare` parsing library [3] we are using. Section 4 describes the second stage, which involves parsing individual formulas (stage 2a) and using our GUI-based tool for disambiguating parses (stage 2b). Once all the formulas are disambiguated, the parse trees are converted into $\text{\texttt{S}\TeX}$ macros and a copy of the document is created, where formulas are replaced with their $\text{\texttt{S}\TeX}$ -ified counterparts (stage 2d). We demonstrate this with an example grammar (see Example 3). Section 5 showcases the current state of our automatic grammar generation, which was developed to make stage 1 of $\text{\texttt{S}\TeX}$ -ification easier. Section 6 summarizes the work and lists some future improvements.

⁵ Semantic macro definitions in $\text{\texttt{S}\TeX}$ are organized into *archives* (for areas of mathematics, e.g., calculus) and *modules* (for defining macros for a specific mathematical object, e.g., the integral).

2 Background

This section outlines the background needed to contextualize the research. We assume familiarity with context-free grammars (CFGs) [1] and GLR parsers [14]. This section gives a brief introduction to $\mathsf{\S\TeX}$ and the (untyped) λ -calculus⁶.

2.1 $\mathsf{\S\TeX}$

The $\mathsf{\S\TeX}$ package [7] was developed by the KWARC research group. Its main feature is a systematic way to generate *semantic macros*, which allow users to semantically annotate their mathematical documents. This allows for the explicit recording of, e.g., structure in a mathematical expression. The group also started the development of the Semantic Multilingual Glossary of Mathematics (SMGloM) [9], which contains semantic macro definitions for hundreds of concepts from various areas of mathematics. The SMGloM also provides natural language definitions of those concepts in English, and sometimes other languages (German being the second most common). The package is well documented in [7]. Nonetheless we will reference some features which we use in this paper.

By semantically marking up terms using $\mathsf{\S\TeX}$ we can also use the systems surrounding it, like the automatically generated hyperlinks to online definitions.

Previous Attempts at Automatic $\mathsf{\S\TeX}$ -ification. A notable previous attempt at automatic $\mathsf{\S\TeX}$ -ification was by Müller and Kalyszyk [11], who trained a machine learning (ML) model on a corpus of $\mathsf{\LaTeX}$ fragments and their semantically annotated counterparts produced using $\mathsf{\S\TeX}$. They achieved some success with GPT-2, producing correct $\mathsf{\S\TeX}$ -ifications around 50% of the time.

Instead of probability-based predictions, we present an approach that relies on user input for correct $\mathsf{\S\TeX}$ -ification. This also removes the need for training an ML model, while being sufficiently flexible to handle arbitrary semantic macro definitions without large modifications to the system we developed.

2.2 The (Untyped) λ -calculus

To demonstrate the methods we propose in this paper, we need to choose an area of mathematics for which we want to $\mathsf{\S\TeX}$ -ify documents. The overall problem of semantically annotating mathematical formulas is present in all areas of mathematics, but the focus of this paper will be the (untyped) λ -calculus. We do this because we are the teaching team of the Foundations course on the λ -calculus, which is a subject of our experimentation in $\mathsf{\S\TeX}$ -ifying its entire $\mathsf{\LaTeX}$ library. We wish to extend tests to our students across all our campuses (around 250). That being said, here is a short recap of the λ -calculus:

⁶ We will use the phrases “untyped λ -calculus” and “ λ -calculus” interchangeably from this point on.

Definition 1 (Variables). Let \mathcal{V} be the set of variables, defined as $\mathcal{V} = \{v, v', v'', \dots\}$. We will denote the meta-variables that range over \mathcal{V} with lowercase letters (e.g., x, y, z), that can have apostrophes or subscripted index attached (e.g., x', y_1, z_2'').

Definition 2 (The set of λ -terms). Let Λ be the set of all λ -terms. We will denote the meta-terms that range over Λ with uppercase letters (e.g., A, B, C), that can have apostrophes or subscripted index attached (e.g., A', B_1, C_2''). We define Λ inductively as follows:

- If $x \in \mathcal{V}$, then x is in Λ .
- If $A, B \in \Lambda$, then the *application of A to B* , denoted by (AB) , is in Λ .
- If $x \in \mathcal{V}$ and $A \in \Lambda$, then the *abstraction in A over x* , denoted by $(\lambda x.A)$, is in Λ .

Definition 3 (Notational conventions of the λ -calculus). We employ some notational conventions when writing out λ -terms. We follow the conventions from our Foundations course notes:

- We can remove the outermost parentheses in a term: we can write AB instead of (AB) .
- Application is left-associative: we can write $(AB)C$ as ABC .
- The scope of an abstraction extends as far to the right as possible: $\lambda x.xy$ is equivalent to $\lambda x.(xy)$, NOT $(\lambda x.x)y$.
- Multiple consecutive abstractions can be “compressed”: we can write $\lambda x.(\lambda y.(\lambda z.A))$ as $\lambda xyz.A$.

3 Laying the Groundwork

In this section, we present the first stage of our approach to $\text{\texttt{S}\TeX}$ -ification (see Table 1). Throughout this section, we will assume that we wish to $\text{\texttt{S}\TeX}$ -ify a document, whose formulas contain only untyped λ -terms (let us call it `demo-file.tex`). Its contents are shown in Example 4. In Sect. 3.2 we will also present the parsing library we are using.

3.1 Identifying the Required $\text{\texttt{S}\TeX}$ Macros

The goal of this section is to define some macros and a grammar in order to $\text{\texttt{S}\TeX}$ -ify `demo-file.tex`. Table 2 lays out some semantic macros we defined to semantically annotate λ -terms.

Table 2. The macros we define for semantically annotating λ -terms.

- `\var{#1}` - marks up its argument as a variable.
- `\abs{#1}{#2}` - the argument 1 is a *flexary argument*^a representing a sequence of variables, and argument 2 is the body of the abstraction.
- `\app{#1}{#2}` - both arguments are terms and this denotes the application of the first to the second.

^a These are arguments of variable length, supplied as comma-separated values. For more details, we refer the reader to the \LaTeX documentation [7], section 3.1.3.

Table 3 shows examples of source code using the macros, and the typeset results. Our aim is to use our macros and tools to \LaTeX -ify the material for our course at Heriot-Watt. Having defined the macros we need to \LaTeX -ify `demo-file.tex`, we can now create a grammar to parse the formulas within the document.

Table 3. A side by side comparison of \LaTeX source code with \LaTeX macros and the typeset results.

Source code	Typeset result
<code>\var{x}</code>	x
<code>\abs{\var{x},\var{y},\var{z}}{A}</code>	$\lambda xyz. A$
<code>\app{A}{B}</code>	AB

3.2 Parsing Formulas with `parglare`

For our parsing, we are using `parglare` [3], a GLR parser written in Python. GLR parsers are suitable for our parsing needs as they are exhaustive and support ambiguous grammars. Parsing terms with `parglare` produces parse trees, which we can transform into abstract syntax trees (ASTs) using *parse actions*. Parse actions are a `parglare` feature which allows for transforming of parse trees into other forms, either by using the built-in functions provided by `parglare`, or defining custom ones. We use custom parse actions to build ASTs that map easily onto the \LaTeX macros. They are built so that the names of nodes match the names of semantic macros identified in Sect. 3, or native \LaTeX macros like `dobrackets`, which wraps its argument in parentheses.

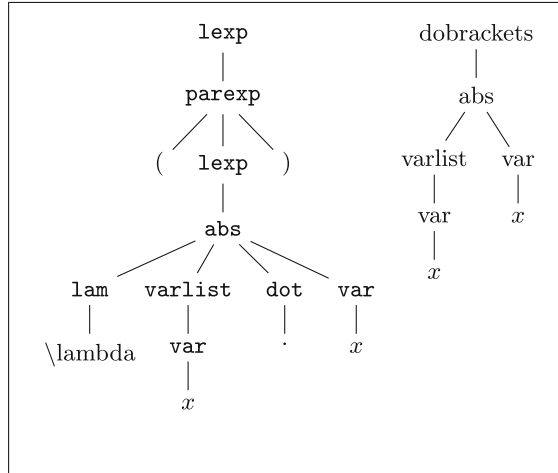


Fig. 1. The parse tree and AST for $(\lambda x.x)$

Once parsing is complete, we can transform the ASTs into a series of macro calls and their arguments to create semantic annotations. This is relatively straightforward, as the nested structure of semantic macros is similar to the ASTs produced during parsing. Figure 1 shows the parse tree created by parsing ($\lambda x.x$) using the grammar in Example 3, and the AST produced by running our parse actions over it. Text in `typewriter` font denotes nonterminals.

Example 3 (The grammar used for parsing λ -terms).

<code>lexp</code> \rightarrow <code>app</code> <code>var</code> <code>abs</code> <code>parexp</code>	<code>lexp</code> matches any λ -term
<code>app</code> \rightarrow <code>lexp</code> <code>lexp</code>	<code>app</code> matches applications
<code>abs</code> \rightarrow <code>lam</code> <code>varlist</code> <code>dot</code> <code>lexp</code>	<code>abs</code> matches abstractions
<code>lam</code> \rightarrow " <code>\</code> " " <code>\lambda</code> "	<code>lam</code> matches either λ or <code>\lambda</code>
<code>varlist</code> \rightarrow <code>var</code> <code>varlist</code> <code>var</code>	<code>varlist</code> matches a list of variables
<code>parexp</code> \rightarrow "(" <code>lexp</code> ")"	<code>parexp</code> matches a <code>lexp</code> in parentheses
<code>var</code> \rightarrow [<code>a-z</code>] ⁺	<code>var</code> is a regex matching single letters
<code>dot</code> \rightarrow "."	<code>dot</code> matches a single "." ⁷

The grammar we define is ambiguous, as it has no notion of associativity of application, or how far the scope of an abstraction extends. That means that it parses terms like xyz , and $\lambda x.yx$ ambiguously⁸.

The grammar was designed by hand, which can be a lengthy process and would not scale to more areas of mathematics easily. We propose a solution in the form of (automatic) grammar generation, described in Sect. 5.

Recognizers. Another powerful feature of `parglare` are its *recognizers*. A recognizer is a Python function the user can supply to the parser for tokenizing portions of the input. For example, matching natural numbers without leading 0s would require a complex regular expression or several grammar rules. It might be more convenient to write a recognizer for such a task. We can improve the grammar we defined in Example 3, by replacing the rule `var` \rightarrow [`a-z`]⁺ with a recognizer for variables with apostrophes or subscripted indexes (to better match the definition of variables in Definition 1). Implementing this with CFG rules is difficult due to \TeX 's peculiarities (e.g., optional braces around arguments), but a recognizer can do it with a Python function. This allows us to use libraries for parsing \LaTeX , (e.g., `pylatexenc` [4]), that make the task more manageable.

Limitations of `parglare`. The main limitation of `parglare` that we encountered is its inability to handle cycles in grammars (see Sect. 5.5 for more detail). A

⁷ This is due to a limitation of the parsing library used, see Sect. 3.2.

⁸ The aim of this paper isn't to create an unambiguous grammar for parsing λ -terms. Those already exist in the literature, but an unambiguous grammar would not showcase all the features of our other work, such as the GUI (see Sect. 4).

different, rather small drawback is that a period cannot appear at the start of a string literal in the grammar. This can be easily resolved by adding a production of the form `dot` \rightarrow "." to each grammar (see Example 3).

4 A Disambiguation GUI for Parsing During \LaTeX -ification

This section describes the second stage of our approach, the disambiguation GUI and its functionality. It was developed using the Python bindings for GTK [13]. By developing the GUI in Python, it easily integrates with `parglare`.

Once we have identified the macros and created a `parglare` grammar, we can open our `.tex` source file within the GUI. All the formulas within the file are extracted and parsed using the grammar. Any ambiguous parses will show up in the GUI's *formula view* (section 3 in Fig. 2). The user can then go through all the formulas and select the correct parses (see Sect. 4.1). Once all the formulas have been disambiguated in this way, a copy of the original `.tex` file is created, with all the formulas replaced by their \LaTeX equivalents.

We have successfully used the GUI to semantically annotate some slides for the Foundations course. It did not fully \LaTeX -ify every formula (the grammar did not cover all of them), but it sped up the process significantly⁹.

Example 4 (The contents of `demo-file.tex`).

```
\documentclass{article}
\begin{document}
Here we have some terms that are ambiguous (according to the grammar
 $\hookrightarrow$  we defined).
 $\lambda xyz.xy$ 
Here, the parser cannot decide whether it is an application of
 $\hookrightarrow \lambda xy.x$  to  $y$  or an abstraction with an application  $xy$ 
 $\hookrightarrow$  inside the body.
 $xyzw$  trips up the parser because it is not aware of the left
 $\hookrightarrow$  associativity of application.
The issue with  $(\lambda xy.xy)$  is that the parser could read this
 $\hookrightarrow$  as an application of  $\lambda xy.x$  to  $y$  wrapped in
 $\hookrightarrow$  parentheses, rather than an abstraction (with an application  $xy$ 
 $\hookrightarrow$  inside the body) that is wrapped in parentheses.
\end{document}
```

4.1 Overview

The main window of the GUI has 8 sections. Figure 2 shows the GUI with no file opened, and Fig. 3 shows the GUI after `demo-file.tex` has been opened. The sections of the GUI, as shown in Fig. 2, are: (1) the toolbar, (2) a dropdown for selecting the parsed formula, (3) the *formula view*, showing all the ambiguous

⁹ Based on the author's subjective experience with manually \LaTeX -ifying mathematical documents in the past.

formulas as Unicode strings¹⁰, (4) the name of the document being processed, (5) a button to confirm the user’s choices, (6) a label displaying the currently selected formula, (7) a grid layout for displaying ASTs, and (8) a “quit” button.




Fig. 2. The main GUI window, with enumerated sections

Figure 3 shows the GUI at work by parsing all the formulas found in `demo-file.tex` using the grammar described in Example 3. We select the file we wish to \LaTeX -ify by navigating to File \rightarrow Open in the toolbar. The formula view shows 3 formulas that parsed ambiguously.

Once we select a

formula, the ASTs for all the possible parses are shown in the grid layout. The currently selected tree has a magenta outline, and the selection is also shown in the “selected option” column of the formula view. To disambiguate a single formula, we need to either click the AST corresponding to the correct parse, or use the dropdown menu to select the correct option. We then need to repeat this

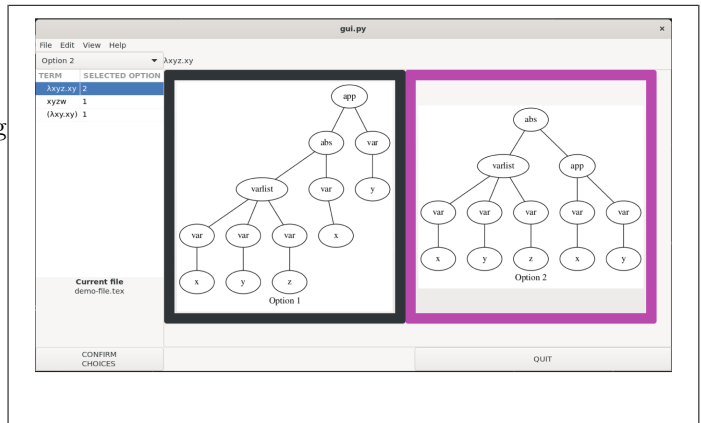


Fig. 3. The GUI once a file has been opened

¹⁰ This is done to improve readability, as it replaces all \LaTeX macros with something more closely resembling their typeset results.

for all the formulas in the formula view, by clicking on each formula and then selecting the correct AST.

Once we select the correct parses for all formulas, we click the “Confirm choices” button, and the ASTs are converted into \LaTeX macros (see Sect. 4.1). A new file called `demo-file-stex.tex` is produced, where each formula from `demo-file.tex` is replaced with its \LaTeX -ified counterpart. Table 4 shows all the formulas in `demo-file.tex` and their corresponding \LaTeX -ified versions.

Visualising ASTs. To show the ASTs of all possible parses to users of the GUI, we need to visualize them into something easier to read. We currently do this in two ways. We can render them as SVGs using `graphviz` [5], by treating the nodes of the trees as nodes in a directed graph. This allows them to be resized without loss of quality.

We can also use indentation to represent the hierarchy of a tree. All the child nodes of a given parent have the same indentation. For readability, we can also colour different “levels” in the tree with customizable colours (like some IDEs which apply different colours to nested parentheses, braces, etc.). Figure 4 shows the different visualisations of the AST for the term $\lambda xy.xy$.




Fig. 4. Visualising the AST for $\lambda xy.xy$.

Converting ASTs to \LaTeX Macros Each AST can easily be converted into \LaTeX macros. The trees have been constructed with this in mind, so it is only a matter of “flattening” them into strings. Special treatment (see Example 10) is needed for flexary arguments.

The \LaTeX -ified formulas are inserted into a copy of the original document. For each of them, the original formula is preserved in a comment right next to it in the source code. If a formula could not

Table 4. Original \LaTeX source code next to its \LaTeX counterparts. A * indicates an ambiguous formula.

demo-file.tex	demo-file-stex.tex
<code>*\lambda xyz.xy</code>	<code>\abs{\var{x}, \var{y}, \var{z}}{\app{\var{x}}{\var{y}}}</code>
<code>\lambda xy.x</code>	<code>\abs{\var{x}, \var{y}}{\var{x}}</code>
<code>y</code>	<code>\var{y}</code>
<code>xy</code>	<code>\app{\var{x}}{\var{y}}</code>
<code>*xyzw</code>	<code>\app{\app{\app{\var{x}}{\var{y}}}{\var{z}}}{\var{w}}</code>
<code>*(\lambda xy.xy)</code>	<code>(\abs{\var{x}, \var{y}}{\app{\var{x}}{\var{y}}})</code>
<code>\lambda xy.x</code>	<code>\abs{\var{x}, \var{y}}{\var{x}}</code>
<code>y</code>	<code>\var{y}</code>

be parsed, it is highlighted in red in the copied document. Some “boilerplate” is also added, to enable \LaTeX functionality. This includes a `\usepackage` command to include the \LaTeX package, a `\smodule` environment to wrap around the contents of the document¹¹, and a `\usemodule` command which imports the λ -calculus macros we defined. Table 4 shows the formulas in `demo-file.tex` next to their \LaTeX -ified counterparts.

Possible Improvements. In the future, we plan to add more options for displaying parse trees, such as displaying formulas with full parentheses (in this way, the nesting of parentheses reflects the structure of a given parse tree), and by “joining” the parse trees into forests similar to those produced by GLR parsers. We will also investigate ways to handle large formulas which might not easily fit on the screen of the GUI, by possibly breaking them up into subformulas. We can also improve the visual aspects of the GUI and the accessibility features, as currently everything is done by clicking with the mouse. Support for customizable keyboard interactions will make using the GUI more user friendly, accessible, and potentially faster to use for those who prefer keyboard shortcuts.

5 Grammar Generation

To use the GUI to parse mathematical formulas, we need to supply a grammar to the parser. Creating a single grammar to parse arbitrary formulas would be incredibly difficult and time consuming. Even if we managed to create one, a large grammar is slower to parse with an exhaustive parsing algorithm like GLR, as the parser has to check a lot of possibilities that eventually get discarded. As such, we believe that a modular approach is needed, to minimize wasted computation.

Even designing small grammars for specific portions of mathematical language is not easy. It takes knowledge of grammar design, the relevant area of mathematics, and time to create such a grammar. Designing grammars by hand would not scale well, so we started investigating automatic grammar generation.

The modularity of \LaTeX archives and our intention to use \LaTeX macros for semantic annotations inspired us to try generating grammars directly from \LaTeX macro definitions. The SMGloM [9] already contains lots of definitions, so developing a way to systematically generate grammars from them could yield a highly modular and flexible tool for parsing and semantically annotating mathematical formulas. This section describes the current state of our approach.

5.1 Initial Approach

First, let us consider an approach one might initially take to generating grammars from \LaTeX macro definitions. We could extract all the definitions provided by `\symdef`, `\symdecl`, and `\notation` macros¹² from a given \LaTeX module. Then,

¹¹ See Section 7.1 of the \LaTeX manual [7].

¹² See sections 7.2 and 7.4 of the \LaTeX manual [7].

we could turn them into grammar rules by replacing all argument placeholders with a special nonterminal, `arg`, and turning all the \LaTeX code (which is just typesetting instructions) into terminals.

If a semantic macro was defined as, e.g., `\symdef{abs}[args=2]{\lambda #1 . #2}`, a rule of the form `abs \rightarrow "\lambda #1 . #2"` `arg` would be created. Then, we can add a “master rule”, with all the other nonterminals and a simple regular expression to match single letters¹³ on the RHS: `arg \rightarrow \dots | abs | \dots | [a-z]+?`. This approach quickly proves to be flawed. It makes the assumption that anything is a valid argument to a semantic macro¹⁴, which leads to *over-generating* grammars and a large number of nonsensical trees. It is evident that the approach has to be refined.

Example 5 (An example grammar).

Recall the semantic macros defined in Sect. 3.1. For this example, we simplified the abstraction macro by replacing the flexary argument with a regular one. Our initial approach generates the grammar in Fig. 5. It is clear that the grammar will over-generate, as an abstraction or application can appear as the first argument of an abstraction (which should only be a variable).

```
arg $\rightarrow$ var | abs | app | text
var $\rightarrow$ arg
abs $\rightarrow$ "\lambda" arg "." arg
app $\rightarrow$ arg arg
text $\rightarrow$ [a - z]+?
```

Fig. 5. A grammar generated from the macros for the λ -calculus

5.2 Adding Types

An option to provide a *type* to a macro is available in \gTeX . We can extract types from macro definitions and use them to restrict the possible arguments of a macro.

Example 6 (Macro with a type and the corresponding grammar rules).

Suppose that we have a macro definition of the form `\symdef{app}[name=application, args=2, type=\funspace{\setOfLambdas, \setOfLambdas}{\setOfLambdas}]{#1 #2}`. Its type is $A \times A \rightarrow A$, it takes two λ -terms as input and returns a λ -term. We can use this information to restrict the arguments in the grammar rule corresponding to `\app`. The rules we generate from it are `app \rightarrow arg.setOfLambdas "+" arg.setOfLambdas` and `arg.setOfLambdas \rightarrow app`.

For each type present in the grammar, we also need to provide a default to match if all else fails (like the letter-matching regular expression in Fig. 5). We can use `parglare`’s *recognizers* for this, but they would have to be manually specified. For the `arg.setOfLambdas` rule, the recognizer can be a simple function that finds uppercase letters possibly followed by apostrophes or subscripted indexes.

¹³ This is the last to be checked, only if none of the other nonterminals matched.

¹⁴ In practice, this is not the case. For example, a semantic macro for addition of natural numbers should not take sets as arguments.

There are two issues to solve before we can scale this to more areas of mathematics. First, not all macro definitions in the SMGloM provide a type signature. Adding them for all macros is a difficult task without automation of type checking, which requires a particular choice of type system to serve as a foundation. Additionally, it is not always obvious what the type of a mathematical object (and thus the semantic macro representing it) should be.

Types in a More “Loose” Sense. We could add types in a less strict sense, akin to some programming languages. When defining a function in, e.g., Java, we need to specify the types for all its arguments, and the type of any value that is returned. We could take a similar approach to typing when generating grammars, by specifying what a valid argument for each macro is. A grammar could be generated using any available types. Then, for rules which have no type restrictions for their arguments, we can prompt the user to select all the nonterminals or recognizers of the correct “type”.

Example 7 (Adding types to an “untyped” rule). Recall the grammar rule for abstractions from the grammar in Fig. 5. We can restrict its first argument to only accept variables by changing the rule for `abs` to `abs → "\lambda" abs_arg_1 "." abs_arg_2`, and adding the rules `abs_arg_1 → var` and `abs_arg_2 → arg`¹⁶.

5.3 Using Precedence to Restrict Grammars Further

We can use *precedence* to restrict grammars even further. If rules have different priorities assigned, `parglare` uses them to disambiguate parses by considering only the rule with the highest priority. In `SiTeX`, precedence provides information about how strong a macro binds its arguments, which is used e.g., for automated bracketing. Lower numbers¹⁶ indicate stronger binding and thus higher precedence. We can use `SiTeX`’s precedence to assign priority to grammar rules, but we must first map them to non-negative integers used by `parglare`.

5.4 Standardising Grammars and Parse Actions

Besides improving the initial approach to grammar generation using types and precedence, we also standardised parse actions and the naming of nonterminals to follow a specific pattern. Each semantic macro definition (via `\symdef` or `\symdec1`) provides the name of a macro, e.g. `macroname` (abbreviated as `mn`). We use that to create a nonterminal with the name `main_macroname_rule`. For each

¹⁵ The second rule is purely illustrative. In practice, we would restrict this grammar further by e.g., adding a recognizer for variables and possibly specifying a general nonterminal (and semantic macro) for a general λ -term (like the `lexp` in the grammar shown in Example 3).

¹⁶ Precedences in `SiTeX` can be negative and range from -2^{32} to 2^{32} .

notation for a given `mn` (defined via `\symdef` or `\notation`), we create a *notation rule* of the form `mn_name_rule` \rightarrow `RHS`, where `name` is the name of the notation (this is an empty string if the notation is defined using `\symdef`). We generate `RHS` by replacing argument placeholders with nonterminals representing types, as per Sect. 5.2. Then, a rule of the form `main_mn_rule` \rightarrow `mn_n1_rule` | ... | `mn_nN_rule` is created, where each `mn_nK_rule` is the LHS of a notation rule.

Example 8 (Automatic grammar generation from a sample archive). After including precedences and types, our macros for annotating λ -terms from Table 2 now look like this:

```
\symdef{var}[name=variable, args=1, type=\varSet]{#1}
\symdef{abs}[name=abstraction, args=ai, prec=51;\infprec x\infprec,
  \type=\funspace{\varSet, \setOfLambdas}{\setOfLambdas}]
  \maincomp{\lambda}\argsep{#1}{}\comp{.}#2}
\symdef{app}[name=application, args=2, prec=50;50x49,
  \type=\funspace{\setOfLambdas, \setOfLambdas}{\setOfLambdas}]#1
  #2}
```

Here, `\varSet` and `\setOfLambdas` are used to represent \mathcal{V} and \mathcal{A} , respectively. In this way, we limit the first argument of an abstraction to only be a variable, while the second argument can be any λ -term. Similarly, both arguments of an application are now necessarily λ -terms. We can then automatically generate a grammar (in `parglare` syntax) from these macro definitions. To use the grammar to parse `demo-file.tex` with `parglare`, some manual tweaks are required¹⁸. Removals from the grammar are indicated in red, while additions are indicated in green. Numbers in braces indicate precedence, and expressions of the form `foo=bar` allow us to pass the value of `bar` to parse actions as a *keyword argument* with name `foo`.

```
main_grammar: main_var_rule {10} | main_abs_rule {11} | main_app_rule
  \{12} | text | par_exp;
var__arg_1: varSet_type_rule;
varSet_type_rule: var__rule | varSet_type_rule_recognizer |
  \variable_recognizer;
var__rule: arg1=var__arg_1 ;
main_var_rule: var__rule {10};
abs__arglist_1: abs__arg_1 abs__arglist_1 | abs__arg_1;
abs__arg_1: setOfLambdas_type_rule varSet_type_rule;
setOfLambdas_type_rule: abs_rule | app_rule |
  \setOfLambdas_type_rule_recognizer | variable_recognizer;
setOfLambdas_type_rule: main_var_rule {10} | main_abs_rule {11} |
  \main_app_rule {12} | setOfLambdas_type_rule_recognizer;
abs__arg_2: main_grammar setOfLambdas_type_rule;
abs_rule: bs "lambda" arg1=abs__arglist_1 dot arg2=abs__arg_2 ;
main_abs_rule: abs_rule {11};
app__arg_1: main_grammar setOfLambdas_type_rule;
app__arg_2: main_grammar setOfLambdas_type_rule;
app_rule: arg1=app__arg_1 arg2=app__arg_2 ;
main_app_rule: app_rule {12};
```

```

par_exp: "(" main_grammar ");
any_type_rule: main_grammar | variable_recognizer;

terminals
bs: "\\ ";
dot: ".";
text: /[A-Z]?/;
varSet_type_rule_recognizer: /[a-z]?/;
setOfLambdas_type_rule_recognizer: /[A-Z]?/;
variable_recognizer: ;

```

Some rules (such as the rules for `par_exp` and `dot`) are hardcoded into this grammar, as they would be needed in most grammars we generate. For brevity, we use regexes instead of recognizers for variables and general λ -terms, since variables in `demo-file.tex` do not contain any apostrophes or subscripts. This grammar can now be used to parse `demo-file.tex`, and ambiguities can be resolved using the GUI. The ambiguities that occur are the same as if we used the manually specified grammar from Example 3. The formulas and their \LaTeX -ified counterparts are laid out in Table 4.

Standardising Parse Pctions. By standardising the grammar structure and naming of nonterminals, we only need a handful of parse actions for transforming the parse results into ASTs, which we assign to certain kinds of nonterminals.

Nonterminals of the form `main_mn_rule` do not contribute to the structure of the AST, so we can define a parse action that returns the subtree without changing it. Nonterminals of the form `mnn_name` (`mnn` is used as a shorthand for `macroname_notation_name`) return an AST with a node named `mnn_rule` as the root, and the arguments of that macro as children.

Example 9 (AST returned by matching a notation).

If we wish to annotate `\lambda x.x`, a portion of the parse tree is given by matching

the rule `abs_rule` \rightarrow `bs "lambda" abs_arglist_1 dot abs_arg_2`. The resulting AST (see Fig. 6) should not include `bs"lambda"` or `dot`, as they are irrelevant to the overall structure of the term. We can use a parse action to build the AST, with the LHS of the rule as the root, and the nonterminals on the RHS that represent arguments, as children. This closely resembles the structure of the semantic macro that we should use to annotate the input string, `\abs{x}{x}`.

¹⁷ Most of these can be made obsolete by improving the grammar generation code, in particular when it comes to using types. This will be addressed in future work.

Special treatment is required for flexary arguments. Grammar rules of the form $\text{mnn_arglist_N} \rightarrow \text{mnn_arg_N sep mnn_arglist_N} \mid \text{mnn_arg_N}$ are used for flexary arguments, and produce distinctly shaped parse trees. We transform them into an AST with a single parent node and one “layer” of children (see Example 10).



Fig. 6. Example AST

Example 10 (A parsed flexary argument and the corresponding AST).

Recall the `\abs` macro defined in Table 2. Its first argument is a flexary sequence of variables that can be parsed by a grammar rule of the form $\text{varlist} \rightarrow \text{var varlist} \mid \text{var}$ (see Example 3). For brevity we are omitting the rules for parsing `var`. When parsing terms like `\lambda xyz.A`, we obtain very “right-leaning” parse trees, which we have to transform into ASTs with one “layer” of children (see Fig. 7).

Standardising nonterminal names and actions yields ASTs which can easily be converted to `TeX` macros. The names of nonterminals in an AST contain the names of relevant semantic macros and notations. Standardising nonterminal names will also ensure fewer conflicts when we merge multiple grammars to parse a document. The only issue to solve is two `TeX` archives defining a macro with the same name. This can be avoided by attaching metadata to nonterminals¹⁸.




Fig. 7. The parse tree (left) and the AST we wish to obtain from it (right)

5.5 Current State and Issues to Address

The grammars we produce can be used to parse small formulas and produce semantically annotated equivalents. We have only done limited testing, because the generated grammars can contain *left recursion* or *cycles*, e.g., in Example 8, $\text{var_rule} \rightarrow \text{var_arg_1} \rightarrow \text{varSet_type_rule} \rightarrow \text{var_rule}$. For clarity, we only show the nonterminals that form this cycle. In the example, we had to manually modify the grammar to remove such cycles. We can potentially eliminate some cycles by improving the grammar generation code, especially when it comes to providing rules for a given type. In Example 8, the rule for `setOfLambdas_type_rule` did not have the “main” nonterminals for a given macro on the RHS, but rather nonterminals associated with a specific notation for a macro.

¹⁸ This is another feature of `paglare`.

However, eliminating all cycles during grammar generation is not a guarantee¹⁹ and `parglare` cannot handle cyclic grammars. This leaves us with two options: (1) modify the grammars to remove cycles, or (2) use a parser that handles cyclic grammars. Algorithms for removing cycles exist, but they change the structure (and size) of the grammar. We wish to avoid that, as “cyclic grammars can be the most compact way to describe a language” [12]. Thus we will address this issue in the parser itself. This will require extensive research of the state of the art, or adapting existing parsing algorithms.

Adding Types to Grammars. Not many semantic macro definitions have types added to them. As such, we might have to consider manually adding “types” of some sort to grammars that we generate from “untyped” semantic macro definitions. This could be as simple as just editing the grammar files, or we could develop a special interface for the task. Each grammar we create could then be shared in an open-source repository similar to the SMGloM, to avoid a reduplication of effort from other users trying to $\mathcal{S}\mathcal{T}\mathcal{E}\mathcal{X}$ -ify similar documents.

5.6 Alternatives

We could search for existing grammars and manually adapt them to fit with our work, and connect them with $\mathcal{S}\mathcal{T}\mathcal{E}\mathcal{X}$ archives based on whether the parse results can be represented with macros from a given archive. We could, however, lose the one-to-one correspondence between grammars and $\mathcal{S}\mathcal{T}\mathcal{E}\mathcal{X}$ archives that we have by generating grammars using the approach we presented in this paper.

We could also “crowdsource” the grammars by starting a community-developed collection of small grammars based on $\mathcal{S}\mathcal{T}\mathcal{E}\mathcal{X}$ archives. This would require precise instructions and rigorous checks for contributors, to avoid e.g., name conflicts in nonterminals, and ensure the grammars can seamlessly be used simultaneously.

6 Conclusion and Future Work

In this paper we presented the exploration of a novel approach to $\mathcal{S}\mathcal{T}\mathcal{E}\mathcal{X}$ -ification by generating grammars from $\mathcal{S}\mathcal{T}\mathcal{E}\mathcal{X}$ macro definitions to parse a document and using a GUI to disambiguate parses. We have shown how it can help $\mathcal{S}\mathcal{T}\mathcal{E}\mathcal{X}$ -ify mathematical documents on a practical example. We also discuss a method to generate grammars for use with the GUI, and elaborated on the issues we face.

In the future, we plan on doing at least the following: (1) as mentioned in Sect. 4.1, we will add more features to the GUI; (2) as mentioned in Sect. 5.5, we will address the issue of cyclic grammars, and develop tools to speed up the grammar creation process; (3) we will develop an interface with which users

¹⁹ Even after modifying the grammar in Example 8, we still have the cycle `abs_arg_2` \rightarrow `setOfLambdas_type_rule` \rightarrow `main_abs_rule` \rightarrow `abs_rule` \rightarrow `abs_arg_2`, for example.

can select the macros/archives they wish to use in their document, from which a grammar is to be generated (4) we will look into ways to share grammars between users, to avoid the reduplication of effort; (5) we will develop a process for merging the smaller grammars we generate into grammars that can handle larger, less specific documents than `demo-file.tex`; (6) we will generalise the “boilerplate” added to an \LaTeX -ified document to add all modules whose macros are being used; (7) we will add support for author-defined macros (this might take the form of a preprocessing step that selectively expands them); (8) we will investigate ways to (semi-)automatically annotate other parts of a document, e.g., definitions, examples, and textual references to mathematical objects.

References

1. Aho, A.V., Lam, M.S., Sethi, R., Ullman, J.D.: Compilers: Principles, Techniques, and Tools, 2 edn. Addison-Wesley
2. Berges, M., Betzendahl, J., Chugh, A., Kohlhase, M., Lohr, D., Müller, D.: Learning support systems based on mathematical knowledge management. In: Dubois, C., Kerber, M. (eds.) CICM 2023. LNCS, vol. 14101, pp. 84–97. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-42753-4_6
3. Dejanović, I.: Parglare: a LR/GLR parser for python. Sci. Comput. Program. **214**, 102734. <https://doi.org/10.1016/j.scico.2021.102734>. <https://www.sciencedirect.com/science/article/pii/S0167642321001271>
4. Faist, P.: Welcome to pylatexenc’s documentation! - pylatexenc 3.0alpha000021 documentation. <https://pylatexenc.readthedocs.io/en/latest/>
5. Graphviz: Graphviz. <https://graphviz.org/>
6. Kohlhase, M., Iancu, M.: Co-representing structure and meaning of mathematical documents. Sprache und Datenverarbeitung. Int. J. Lang. Data Process. **38**, 49–80
7. Kohlhase, M., Müller, D.: The STEX3 Package Collection. <https://anorien.csc.warwick.ac.uk/mirrors/CTAN/macros/latex/contrib/stex/doc/stex-doc.pdf>
8. Kohlhase, M., Müller, D.: System description STEX3 - a LATEX-based ecosystem for semantic/active mathematical documents. In: Buzzard, K., Kutsia, T. (eds.) CICM 2022. LNCS, vol. 13467, pp. 184–188. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16681-5_13
9. KWARC: SMGloM - GitLab. <https://gl.mathhub.info/smgloom>
10. Müller, D.: Slatex/RusTeX. <https://github.com/slatex/RusTeX>
11. Müller, D., Kaliszyk, C.: Disambiguating symbolic expressions in informal documents. In: International Conference on Learning Representations. <https://openreview.net/forum?id=K5j7D81ABvt>
12. Rekers, J.: Parser generation for interactive environments. Ph.D. thesis. <https://ir.cwi.nl/pub/29930/29930D.pdf>
13. The GTK Team: The GTK Project - A free and open-source cross-platform widget toolkit. <https://www.gtk.org/>
14. Tomita, M., Ng, S.K.: The generalized LR parsing algorithm. In: Tomita, M. (ed.) Generalized LR Parsing, pp. 1–16. Springer, Boston (1991)