Publishing Math Lecture Notes as Linked Data

Catalin David¹, Michael Kohlhase¹, Christoph Lange¹, Florian Rabe¹,
Nikita Zhiltsov², and Vyacheslav Zholudev¹

¹ Computer Science, Jacobs University Bremen, Germany
{c.david,m.kohlhase,ch.lange,f.rabe,v.zholudev}@jacobs-university.de

² Mathematics, Kazan State University, Russia
nikita.zhiltsov@gmail.com

Abstract. We mark up a corpus of L^AT_EX lecture notes semantically and expose them as Linked Data in XHTML+MathML+RDFa. Our application makes the resulting documents interactively browsable for students. Our ontology helps to answer queries from students and lecturers, and paves the path towards an integration of our corpus with external sites.

1 Application: Computer Science Lecture Notes

Over the last seven years, the second author has accumulated a large corpus of teaching materials, comprising more than 2,000 slides, about 1,000 homework problems, and hundreds of pages of course notes, all written in L^AT_EX. The material covers a general first-year introduction to computer science, graduate lectures on logics, and research talks on mathematical knowledge management. This situation is typical for educators and researchers and represents the state of the art in mathematics, physics, computer science, and engineering: L^AT_EX has proven suitable for writing high-quality lecture notes and publishing them as PDF. However, in our educational setting, we would like to benefit from the much larger degree of interactivity that screen reading and e-books support. For example, while reading notes students want to directly look up the meaning of a symbol (e.g. $\models$) in a formula, or examples for a difficult concept (e.g. structural induction). They may want to select advanced material for self-study from the whole body of lecture notes, based on the topics covered in the lecture. They want to use a search engine to find related material in other universities' online course notes, on mathematical web sites, or Wikipedia. Lecturers want to query their repository for document parts reusable in an upcoming lecture, given the prerequisites students are expected to meet and the material that has already been covered. In a course for a special audience, e.g. mathematics for physicists, they want to draw examples from that domain even though they are less familiar with it. They also want to locate didactic gaps, such as concepts without examples, or unjustified proof steps. These services require semantic annotations in the lecture notes that are understandable for external search engines. Plain L^AT_EX is barely usable for anything *beyond* on-screen reading and printing. Even simple semantic annotations are uncommon, rare exceptions are the `\title` command making its meaning explicit or `\frac{a}{b}` focusing on functional

structure instead of visual layout. This is especially problematic for symbols in formulæ, which are often overloaded with multiple definitions or presentable in different notations. $\binom{n}{k}$ can be a vector or a binomial coefficient, and a French or Russian would write the latter as $\mathcal{C}_n^k$. Therefore, we have developed a semantic representation of mathematical knowledge in $\text{\LaTeX}$ and a presentation process that preserves these semantic structures as Linked Data in the output, exposing them to mashups for interactive exploration, as well as semantic searching and querying. These are based on an ontology for mathematical knowledge so that mathematical content can be linked across different repositories.

2 Research Background and Related Work

$\text{\LaTeX}$'s importance in scientific authoring and its extensibility by macros have led to semantic extensions enabling modern publishing workflows. SALT (semantically annotated $\text{\LaTeX}$ [8]) marks up rhetorical structures and fine-grained citations in scientific documents. Its markup is not sufficiently fine-grained for formulæ, and its vocabulary is limited to rhetorics and citations and not extensible. Our own sTeX offers macros for introducing new mathematical symbols and using arbitrary metadata vocabularies. Some math e-learning systems, such as ActiveMath [1] or MathDox [17], use semantic representations of formulæ and higher-level structures, e.g. proof steps or course module dependencies, in the standard XML languages OpenMath [5] and OMDoc [11]. They utilize semantic structures but do not *publish* them in a standard representation like RDF, which would promote general-purpose queries beyond the built-in services and integration with other systems on the web. The Linking Open Data movement promotes best practices for publishing data on the web [9], as standalone RDF or embedded into HTML documents as RDFa [2]. Applications include Sindice, an engine that crawls and indexes Linked Data [19], and the Sparks O_3 Browser, a mashup that utilizes RDFa annotations in HTML for interactive browsing [20]. Our interactive documents work similarly but additionally support annotations in MathML formulæ. MathML has pioneered embedded annotations long before RDFa, albeit with a more limited scope. Its *parallel markup* interlinks the rendered appearance and the semantic structure of mathematical expressions; the meaning of mathematical symbols is usually defined in lightweight ontologies called OpenMath content dictionaries [4]. HELM (Hypertext Electronic Library of Mathematics [3]) pioneered representing structures of mathematical knowledge in RDF, e.g. what mathematical theory introduces a symbol, what of its properties have been declared or asserted, and how the latter are proved. The HELM ontology has not gained wide acceptance, though. At the time of its development, there was no RDFa-like way of embedding RDF into web documents.

3 Architecture and Demo

Our architecture publishes semantically enriched $\text{\LaTeX}$ lecture notes as XHTML+MathML+RDFa Linked Data. We kept $\text{\LaTeX}$ as an input language, as

it is familiar to authors and well supported by editors, and as high-quality PDF can be obtained from it. With sTeX (semantically enhanced TeX), we have introduced L^ATeX macros for marking up the semantic structure of formulæ and documents [12] and manually annotated our complete corpus using the sTeX plugin for the Emacs editor. One can, e.g., declare a symbol *union*, formally define it, and make its semantic representation `\union{A,B,C}` expand to `A\cup B\cup C` for human-readable rendering. There are environments for mathematical statements and theories, e.g. `\begin{example}[for=union]`. L^ATeXXML transforms this into a semantically equivalent intermediate XML representation, using the standard XML languages OpenMath for formulæ [5] and OMDoc for higher-level structures [11]. Finally, our JOMDoc rendering library [10] generates human-readable output from this XML – an output that still contains the full semantic structure as annotations. A custom Java implementation renders formulæ as parallel markup of Presentation MathML annotated with OpenMath¹; rendering higher-level structures as XHTML+RDFa [2] is implemented in XSLT. RDF is extracted from XML by our Krexor XML→RDF library [15], which generates URIs for all mathematical objects in a document. It uses our OMDoc ontology (cf. [14]) as a vocabulary for representing all mathematical structures (e.g. “*d* is a definition, *e* is an example for *d*”) plus full text, inspired by HELM and designed as a more expressive counterpart of the OMDoc XML schema.

The whole transformation process is integrated into our versioned XML database TNTBase [22]; see <http://kwarc.info/LinkedLectures>. TNTBase has a Subversion-compatible interface making it suitable as a lecture notes repository. The TeX→XML and XML→RDF transformations are automatically triggered by a hook upon committing a new revision of an sTeX lecture module. If the generated OMDoc+OpenMath is not schema-valid, the commit is rejected. On the other hand, it follows Linked Data best practices and, depending on the MIME type an HTTP client requests, serves a document as OMDoc, as RDF (only a structural outline, not the full text and formulæ), or as XHTML+MathML+RDFa. The latter contains JavaScript code from our JOBAD library for interactive documents [13,7], which operationalizes the annotations – Linked Data and other – in the rendered documents. JOBAD’s definition lookup determines the OpenMath annotation of the Presentation MathML symbol the user clicked on, from that obtains the URI of the symbol, and then requests XHTML from that URI (resulting in the symbol’s declaration and definition), which is then displayed in a popup. The RDFa annotations are used for making parts of a document (e.g. steps of a structured proof) foldable, and for displaying the local neighborhood in the RDF graph (e.g. related examples) in popups; this is implemented using the rdfQuery library [18], relying on the Linked Data structure in the latter case. Further third-party services can be integrated in a mashup style; we have demonstrated this for a unit conversion service [13,7]. Besides

¹ A proposal for fully representing formulæ in RDF [16] has not gained wide acceptance. RDF-based reasoners are often limited to decidable first order logic subsets, which is insufficient for mathematical applications, and XML has a straightforward notion of order (e.g. of the arguments of an operator or of a set constructor).

enabling JOBAD’s services, we have implemented machinery to load the extracted RDF into a triple store and query it using SPARQL. We also provide a widget for formulating queries without knowing SPARQL and the OMDoc ontology. It allows to ask some non-trivial queries, e.g. “find examples for all concepts from graph theory (about which I’m planning a lecture), assuming as prerequisites the concepts from formal languages (and their prerequisites)”. This would yield the parse tree of a context-free language as an example for the concept “tree” – as operating systems were not among the prerequisites.

Our demo shows the complete pipeline in action: (i) annotating a document with our sTeX Emacs mode, (ii) committing it to TNTBase, (iii) automatic translation to OMDoc, schema validation, and RDF extraction, (iv) loading the extracted RDF data into a triple store, (v) retrieving the document in different representations, (vi) browsing the XHTML+MathML+RDFa rendering, (vii) interacting with the Linked Data in it, (viii) and querying a triple store. Additionally, we will demonstrate the generation of PDF from the sTeX sources.

4 Conclusion and Outlook

Our architecture makes legacy L^AT_EX lecture notes available as Linked Data. We expose these data to external clients but have also implemented services for interactively exploring the XHTML+MathML+RDFa presentation of our data. We are also working on preserving some of the semantics in the PDF output, as SALT does. Evaluation of our enriched lecture notes by the student end users is planned for the next semester. To the best of our knowledge, we are the first provider of RDF-based Linked Data in the domain of mathematics and among the first to operationalize the Linked Data structures of formula markup. Having successfully transformed more than 300,000 normal, non-semantic L^AT_EX documents from arxiv.org to XHTML+Presentation MathML [21] and working on machinery for automatically annotating them using natural language processing, we will soon be able to expose even more mathematical knowledge as Linked

Open Data; however, due to the inherent complexity of mathematical knowledge, with a less formal semantics than manually annotated documents. Our lecture notes are self-contained so far, but we are now starting to reap the benefits of Linked Data by linking them to other data sets, e.g. DBpedia [6], whose mathematical knowledge does not have a semantics as strong as ours, but which provides abundant informal background knowledge, e.g. about the originators of mathematical theories. On the other hand, hardly any well-known mathematical site (e.g. planetmath.org and mathworld.wolfram.com) currently exposes machine-understandable metadata. We promote our technology, starting with lightweight RDFa annotation using the OMDoc ontology, as a migration path towards their integration into a true mathematical Semantic Web.

References

1. ActiveMath, <http://www.activemath.org>
2. RDFa in XHTML: Syntax and processing. Recommendation, W3C (2008)
3. Asperti, A., Padovani, L., Sacerdoti Coen, C., Guidi, F., Schena, I.: Mathematical knowledge management in HELM. *Annals of Mathematics and Artificial Intelligence*, Special Issue on Mathematical Knowledge Management 38(1-3), 27–46 (2003)
4. MathML 3.0. Candidate Recommendation, W3C (2009)
5. The Open Math standard, version 2.0. Technical report, Open Math Society (2004)
6. DBpedia, <http://www.dbpedia.org>
7. Giceva, J., Lange, C., Rabe, F.: Integrating web services into active mathematical documents. In: Carette, J., Dixon, L., Coen, C.S., Watt, S.M. (eds.) *Calculemus 2009*. LNCS (LNAI), vol. 5625, pp. 279–293. Springer, Heidelberg (2009)
8. Groza, T., Handschuh, S., Möller, K., Decker, S.: SALT – semantically annotated $\text{\LaTeX}$ for scientific publications. In: Franconi, E., Kifer, M., May, W. (eds.) *ESWC 2007*. LNCS, vol. 4519, pp. 518–532. Springer, Heidelberg (2007)
9. Linked data guides, <http://linkeddata.org/guides-and-tutorials>
10. JOMDoc — Java library for OMDoc documents, <http://jomdoc.omdoc.org>
11. Kohlhase, M.: OMDoc – An open markup format for mathematical documents [Version 1.2]. LNCS (LNAI), vol. 4180. Springer, Heidelberg (2006)
12. Kohlhase, M.: Using $\text{\LaTeX}$ as a semantic markup format. *Mathematics in Computer Science* 2(2), 279–304 (2008)
13. Kohlhase, M., Giceva, J., Lange, C., Zholudev, V.: JOBAD – interactive mathematical documents. In: *AI Mashup Challenge* (2009)
14. Lange, C.: SWiM – a semantic wiki for mathematical knowledge management. In: Bechhofer, S., Hauswirth, M., Hoffmann, J., Koubarakis, M. (eds.) *ESWC 2008*. LNCS, vol. 5021, pp. 832–837. Springer, Heidelberg (2008)
15. Lange, C.: Krexitor – an extensible XML→RDF extraction framework. In: *Scripting and Development for the Semantic Web*, SFSW 2009 (2009)
16. Marchiori, M.: The mathematical semantic web. In: Asperti, A., Buchberger, B., Davenport, J.H. (eds.) *MKM 2003*. LNCS, vol. 2594, pp. 216–223. Springer, Heidelberg (2003)
17. MathDox – interactive mathematics, <http://www.mathdox.org>

18. rdfQuery – RDF processing in browser, <http://code.google.com/p/rdfquery/>
19. Sindice – the semantic web index, <http://sindice.com>
20. Sparks *O₃* browser: Enlighten the web, <http://oak.dcs.shef.ac.uk/sparks/>
21. Stamerjohanns, H., Kohlhase, M., Ginev, D., David, C., Miller, B.: Transforming large collections of scientific publications to XML. *Mathematics in Computer Science* (2010)
22. Zholudev, V., Kohlhase, M., Rabe, F.: A (insert xml format) database for (insert cool application). In: *Proceedings of XML Prague* (2010)