

# Mapping Spreadsheets to RDF: Supporting Excel in RML

Markus Schröder, Christian Jilek, and Andreas Dengel

<sup>1</sup> Smart Data & Knowledge Services Dept., DFKI GmbH, Kaiserslautern, Germany

<sup>2</sup> Computer Science Dept., TU Kaiserslautern, Germany

{markus.schroeder, christian.jilek, andreas.dengel}@dfki.de

**Abstract.** The RDF Mapping Language (RML) enables, among other formats, the mapping of tabular data as Comma-Separated Values (CSV) files to RDF graphs. Unfortunately, the widely used spreadsheet format is currently neglected by its specification and well-known implementations. Therefore, we extended one of the tools which is RML Mapper to support Microsoft Excel spreadsheet files and demonstrate its capabilities in an interactive online demo. Our approach allows to access various meta data of spreadsheet cells in typical RML maps. Some experimental features for more specific use cases are also provided. The implementation code is publicly available in a GitHub fork.

**Keywords:** Spreadsheet · Excel · RML · RDF · Knowledge Graph

## 1 Introduction

As soon as knowledge graphs have to be constructed from given data, the use of mapping languages is a common practice. Such languages let users define declarative rules to map various input data formats to a complex interconnected graph. In the Semantic Web community such graphs are usually modeled with the Resource Description Framework (RDF) [11] and (semi-)structured data is mapped with the RDF Mapping Language (RML) [2] – a superset of the W3C-recommended mapping language R2RML [10]. The specification of RML [1] describes how semi-structured data, like XML or JSON, and tables in form of Comma-Separated Values (CSV) can be mapped properly. However, tabular data can also be found in other file types, prominently in spreadsheets.

Since the spreadsheet methodology enables a well understood, easy and fast possibility to enter data, they are widely used by knowledge workers, especially in the industrial sector. In contrast to simply structured CSV files, spreadsheets can model complex workbooks containing multiple sheets with meta data rich cells. Besides its content, a single cell may store additional information like its appearance (colors, styles and borders) or cell comments. How these cells should be filled or styled by users is not predetermined by spreadsheet applications in general. As a consequence, in practice, cells can contain inconsistent and unstructured content which can be arbitrarily arranged in a sheet. That is why a mapping of such data to RDF can become a challenging task.

Although, spreadsheets are frequently used in industry, well-known RML-supporting tools like RML Mapper<sup>3</sup>, CARML<sup>4</sup>, RocketRML [9], SDM-RDFizer [4] or Mapeathor [5] currently do not support them natively. To fix this issue, we extended RML Mapper to support Microsoft Excel spreadsheet files. We have chosen RML Mapper because its code structure let us easy embed our new code. Natively supporting spreadsheets in RML has several advantages. There is no need anymore to preprocess and transform spreadsheets in a format that can be handled by a mapping tool (e.g. CSV). This removes the otherwise additional effort for mapping experts and let them focus on defining proper rules. Having RML rules that directly refer to spreadsheet data makes interpretation of rule definitions straightforward. Thus, it becomes easier to communicate to a data provider how their spreadsheets will be mapped to a knowledge graph. Additionally, practitioners of RML who already learned the language become able to map spreadsheets with almost no extra effort.

In the next section we describe in detail how we realized the Excel support in RML.

## 2 Supporting Excel in RML

Our approach is implemented as an extension to the RML Mapper tool in a separate GitHub fork<sup>5</sup>. For demonstration purpose, we additionally provide an interactive demo page<sup>6</sup> where visitors can try various mapping examples.

At its core, our component utilizes the Apache POI<sup>7</sup> library to read Microsoft Excel spreadsheets. In order to refer to spreadsheet contents as a logical source, a small spreadsheet ontology<sup>8</sup> (usually prefixed with `ss`) was designed. As demonstrated in Listing 1.1, using a spreadsheet reference formulation (Line 2), a workbook source should name a spreadsheet file (Line 5), a referring sheet by name (Line 6) and a range of cells in the sheet (Line 7). A triples map that uses this logical source will iterate over single cells in the given range.

Listing 1.1: Exemplary RML definition of a spreadsheet file as a logical source.

```

1 [ a rml:LogicalSource ;
2   rml:referenceFormulation ql:Spreadsheet ;
3   rml:source [
4     a ss:Workbook;
5     ss:url "workbook.xlsx" ;
6     ss:sheetName "Papers" ;
7     ss:range "A2:A5" ;
8     ss:javascriptFilter "/Know\\w*\\.test(valueString)" # optional
9   ]
10 ]

```

<sup>3</sup> <https://github.com/RMLio/rmlmapper-java>

<sup>4</sup> <https://github.com/carml/carml>

<sup>5</sup> <https://github.com/mschroeder-github/rmlmapper-java/tree/mschroeder-features>

<sup>6</sup> <http://www.dfki.uni-kl.de/~mschroeder/demo/excel-rml>

<sup>7</sup> <https://poi.apache.org/>

<sup>8</sup> <http://www.dfki.uni-kl.de/~mschroeder/ld/ss>

Optionally, a filter can be added that picks certain cells based on a JavaScript program (Line 8). Regarding our example, a regular expression is used to iterate over only those cells which contain the phrase “Know”. Using the expressive script language and variables that represent cell meta data, appropriate filter procedures can be realized.

Listing 1.2: Demonstration of how cell meta data can be accessed in RML maps.

```

1 rr:subjectMap [
2   rr:template "http://example.org/{address}"
3 ] ;
4 rr:predicateObjectMap [
5   rr:predicateMap [
6     rr:template "http://example.org/{[2,0].valueString}"
7   ] ;
8   rr:objectMap [
9     rml:reference "(2,0).valueNumeric"
10  ]
11 ]

```

In various RML maps, meta data of the currently iterated cell can be retrieved, as demonstrated in Listing 1.2. For example in Line 2, a template expression inserts the cell’s address (like “A2”) to form a unique URI for a subject resource. However, often it is required to refer to nearby cells in order to construct appropriate statements. To refer to other cells relatively from the current cell, we introduce a parenthesis notation, like `(column,row)`. A column shift (x-axis of the sheet) and a row shift (y-axis of the sheet) allows to reach any cell relative to the cell representing the subject resource. Regarding Line 9 of our example, numeric values which are located two columns away on the right are used as objects in the mapped statements. In a similar way, cells can be referenced absolutely by using square brackets (Line 6).

As already mentioned, spreadsheet cells have several meta data values that need to be retrievable in a mapping through variables (written in the following with monospaced font). A cell’s location in a sheet can be accessed either as a usual spreadsheet **address** or as **column** and **row** indices. Since a cell (if not empty) stores either a string value or a numeric value, various possibilities are given to access its content: **valueNumeric** retrieves a floating point value (**valueInt** an integer value), **valueBoolean** a boolean value, **valueFormula** if a cell contains a formula (**valueError** to get its possible error code) and **valueString** retrieves its text content. Besides content, one could also be interested in the appearance of a cell in a sheet: since it can be colored, **backgroundColor** and **foregroundColor** can be queried as a hexadecimal RGB value. Regarding the used font, **fontColor**, **fontName** and **fontSize** are available. On a more fine-grained level, cells can also store formatted text which is returned by **valueRichText**. The formatted text is represented in an HTML-like syntax, for example,

“`<b><i><font face='Arial' color='#ff0000'>red, italic and bold</font></i></b>`”.

If spreadsheets were completed in an inconsistent manner, it happens that cells were unintentionally filled with different data types. For such cases, one can use the **value** variable to always obtain a string representation of a cell’s content regardless of its cell type. However, if this is not sufficient because more details are needed, the **json** variable could be used to retrieve a JavaScript Object

Notation (JSON) representation of a cell. The JSON object contains, besides the cell type, various data types mentioned above.

## 2.1 Experimental Features

We also would like to propose some experimental features that we found useful in our use cases. However, the official introduction of these extensions would require to change the RML specification on some points.

**Multiple Different Properties in a Cell.** In our use cases, we frequently experienced that users record multiple information in one cell. Each piece of information  $i_j \in I$  then corresponded to a different property  $p_j \in P$ . Following the current RML specification, one has to define for each  $p_j \in P$  a separate predicate-object map. Instead, we propose a shortcut such that only one predicate-object map is needed. This is done by allowing that an RDF list of properties  $p_j \in P$  can be passed in a predicate map. By implementing a suitable Function Ontology (F<sub>n</sub>O) [7] procedure, we extract information pieces  $i_j \in I$  from a cell's content and return them as an object list  $o_j \in O$ . Usually, RML provides that a Cartesian product is formed between predicates and objects which is  $P \times O$ . However, in our case we need to zip predicate and object lists such that the following set is made:  $\{(p_j, o_j)\}$ . Thus, instead of all possible pairs of the  $P \times O$  matrix, only the diagonal ones are selected. This new behavior is activated by adding a `ss:zip true` statement to a predicate-object map.

**Multiple Complex Entities in a Cell.** Similarly to the previous observation, we often discovered that users mention several entities in a single cell. A prominent example is a list of persons having first and last names, for instance book authors. Represented as an RDF graph, such complex entities potentially require several statements to be fully expressed. An F<sub>n</sub>O function that is able to perform entity extraction would need to return an arbitrary large RDF graph in a certain serialization format (e.g. Turtle) instead of a single value. In order to integrate such a return value in an object map, we define a new term type `ss:Graph`. Once this term type is chosen, the returning RDF graph is parsed and added to the emerging knowledge graph. By using a special `ss:SelectedObjects` resource together with `rr:object` statements, only selected objects will be mapped in an object map.

## 3 Related Work

In the past, several similar approaches were implemented that map spreadsheets to RDF using a language. Domain specific languages (DSL) other than RML are provided to let knowledge engineers express how input data shall be mapped to a graph structure.

Spread2RDF<sup>9</sup> uses a Ruby-internal language while *M*<sup>2</sup> (Mapping Master) [8] builds upon a compact syntax for OWL ontologies (Manchester syntax).

<sup>9</sup> <https://github.com/marcelotto/spread2rdf>

Sheet2RDF [3] uses a special ProjEction of Annotations Rule Language (PEARL), whereas XLWrap [6] utilizes template graphs together with a special expression language to refer to contents of sheets. A different approach is followed by TabLinker<sup>10</sup> which requires that input spreadsheets are annotated in advance with certain styles.

With our implementation, we combine the feature of mapping spreadsheets with the advantages of the well specified RML approach. Those who already used the language and need to process spreadsheets do not have to look for an alternative anymore.

## 4 Conclusion and Outlook

In this paper we discussed that spreadsheets, despite being widely used, are still not supported by the RDF Mapping Language (RML). We therefore proposed a first solution which extends the already existing RML Mapper tool with necessary components. Additionally, experimental features were introduced too. A publicly available web page<sup>11</sup> demonstrates the new capabilities.

In the future, we plan to integrate our proposed solution in the RML specification [1] so that other RML tools may support spreadsheets too. Further, we intend to run several performance and compliance tests with our current reference implementation. This could be done in the R2RML implementation report<sup>12</sup>.

**Acknowledgements** This work was funded by the BMBF project SensAI (grant no. 01IW20007).

## References

1. Dimou, A., Sande, M.V.: RDF Mapping Language (RML). <https://rml.io/specs/rml/> (2020), accessed: 2021-02-01
2. Dimou, A., Vander Sande, M., Colpaert, P., Verborgh, R., Mannens, E., Van de Walle, R.: RML: a generic language for integrated RDF mappings of heterogeneous data. In: Proceedings of the 7th Workshop on Linked Data on the Web. CEUR Workshop Proceedings, vol. 1184 (Apr 2014)
3. Fiorelli, M., Lorenzetti, T., Pazienza, M.T., Stellato, A., Turbati, A.: Sheet2rdf: a flexible and dynamic spreadsheet import&lifting framework for RDF. In: Current Approaches in Applied Artificial Intelligence - 28th International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems, IEA/AIE 2015, Seoul, South Korea, June 10-12, 2015, Proceedings. Lecture Notes in Computer Science, vol. 9101, pp. 131–140. Springer (2015)
4. Iglesias, E., Jozashoori, S., Chaves-Fraga, D., Collarana, D., Vidal, M.: Sdm-rdfizer: An RML interpreter for the efficient creation of RDF knowledge graphs. In: CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020. pp. 3039–3046. ACM (2020)

<sup>10</sup> <https://github.com/Data2Semantics/TabLinker>

<sup>11</sup> <http://www.dfki.uni-kl.de/~mschroeder/demo/excel-rml>

<sup>12</sup> <https://github.com/kg-construct/r2rml-implementation-report>

5. Iglesias-Molina, A., Pozo-Gilo, L., Doña, D., Ruckhaus, E., Chaves-Fraga, D., Corcho, Ó.: Mapeathor: Simplifying the specification of declarative rules for knowledge graph construction. In: Proceedings of the ISWC 2020 Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 19th International Semantic Web Conference (ISWC 2020), Globally online, November 1-6, 2020 (UTC). CEUR Workshop Proceedings, vol. 2721, pp. 25–30. CEUR-WS.org (2020), <http://ceur-ws.org/Vol-2721/paper488.pdf>
6. Langegger, A., Wöß, W.: Xlwrap - querying and integrating arbitrary spreadsheets with SPARQL. In: The Semantic Web - ISWC 2009, 8th International Semantic Web Conference, ISWC 2009, Chantilly, VA, USA, October 25-29, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5823, pp. 359–374. Springer (2009)
7. Meester, B.D., Dimou, A., Verborgh, R., Mannens, E.: An ontology to semantically declare and describe functions. In: The Semantic Web - ESWC 2016 Satellite Events, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers. Lecture Notes in Computer Science, vol. 9989, pp. 46–49 (2016). [https://doi.org/10.1007/978-3-319-47602-5\\_10](https://doi.org/10.1007/978-3-319-47602-5_10), [https://doi.org/10.1007/978-3-319-47602-5\\_10](https://doi.org/10.1007/978-3-319-47602-5_10)
8. O’Connor, M.J., Halaschek-Wiener, C., Musen, M.A.: M2: A language for mapping spreadsheets to OWL. OWL: Experiences and Directions (OWLED), Sixth International Workshop (2010)
9. Simsek, U., Kärle, E., Fensel, D.: Rocketrml - A nodejs implementation of a use case specific RML mapper. In: Joint Proceedings of the 1st International Workshop on Knowledge Graph Building and 1st International Workshop on Large Scale RDF Analytics co-located with 16th Extended Semantic Web Conference (ESWC 2019), Portorož, Slovenia, June 3, 2019. CEUR Workshop Proceedings, vol. 2489, pp. 46–53. CEUR-WS.org (2019), <http://ceur-ws.org/Vol-2489/paper5.pdf>
10. World Wide Web Consortium: R2RML: RDB to RDF mapping language. <https://www.w3.org/TR/r2rml/> (2012), accessed: 2021-02-01
11. World Wide Web Consortium: RDF 1.1 Primer. <https://www.w3.org/TR/rdf11-primer/> (2014), accessed: 2021-02-01