

# Results for Knowledge Graph Creation Challenge 2025: SDM-RDFizer

Enrique Iglesias<sup>1,2,\*</sup>, Maria-Esther Vidal<sup>1,2,3</sup>

<sup>1</sup>L3S Research Center, Hannover, Germany

<sup>2</sup>Leibniz University of Hannover, Hannover, Germany

<sup>3</sup>TIB Leibniz Information Centre for Science and Technology, Hannover, Germany

## Abstract

In recent years, knowledge graphs (KGs) have grown in popularity. Major companies have incorporated them into their products to enhance the user experience. Consequently, numerous methods have emerged to address KG construction. The RDF Mapping Language (RML) is an extension of the W3C mapping language recommendation R2RML. RML allows for the declarative definition of KG structures based on unified ontologies and data sources in various formats, including CSV, JSON, and XML files. However, RML supports CSV, JSON, and XML files. Furthermore, RML has additional functionalities, such as executing functions, using quoted triples, working with collections, and creating logical views. Thus, RML has become a standalone mapping language separate from R2RML. The KGCW 2025 Challenge defines a dataset comprising a series of test cases that cover all RML functionalities. These test cases evaluate the compliance of state-of-the-art KG creation engines. This paper reports on the conformance evaluation of SDM-RDFizer executing this dataset, highlighting its strengths and areas for improvement to enhance performance.

## Keywords

Knowledge Graph Creation, Data Integration System, RDF Mapping Languages

## 1. Introduction

Knowledge graphs (KGs) have become increasingly commonplace in recent years, driven by the significant growth in daily data generation. Major companies such as Google, Netflix, Amazon, and Microsoft use KGs to create relationships between products, concepts, and themes to enhance the user experience [1]. As a result, multiple methods and mapping languages have been developed to support KG creation. One such mapping language is the RDF Mapping Language (RML) [2]. RML was introduced initially as an extension of R2RML, which focused solely on relational databases. RML expanded support to additional data source formats, including CSV, JSON, and XML. Both RML and R2RML adhere to the rules established by the Resource Description Framework (RDF)<sup>1</sup>. Over time, RML has incorporated additional functionalities such as the definition and execution of value transformation functions (RML+FnO [3]), support for RDF-Star (RML-Star [4]), transformation of collections and containers (RML-CC<sup>2</sup>), and projection and transformation of data sources (RML-LV<sup>3</sup>). The integration of these extensions has allowed RML to evolve into a stand-alone mapping language with its own specification<sup>4</sup>.

The dataset used in the KGCW 2025 Challenge defines an informative set of test cases that cover core functionalities of RML, including the execution of joins, duplicate handling, blank nodes, and empty values. Additionally, the dataset includes test cases for all current RML extensions (e.g., RML-Star, RML-CC, RML-LV). The challenge aims to evaluate the level of compliance of existing KG creation engines with the latest RML specification. This report builds upon the results achieved in the KGCW

Portorož'25: Sixth International Workshop on Knowledge Graph Construction, June 1, 2025, Portorož, SI

\*Corresponding author.

<sup>†</sup>These authors contributed equally.

✉ iglesias@l3s.de (E. Iglesias); maria.vidal@tib.eu (M. Vidal)

ORCID 0000-0002-8734-3123 (E. Iglesias); 0000-0003-1160-8727 (M. Vidal)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

<sup>1</sup><https://www.w3.org/TR/2004/REC-rdf-primer-20040210/>

<sup>2</sup><https://kg-construct.github.io/rml-cc/spec/docs/>

<sup>3</sup><https://kg-construct.github.io/rml-lv/ontology/documentation/index-en.html>

<sup>4</sup><https://kg-construct.github.io/rml-core/spec/docs/>

2024 Challenge [5] by incorporating evaluations of newly introduced test cases (e.g., RML-LV) and addressing cases that were previously incomplete. It presents the modifications applied to SDM-RDFizer to ensure full compliance with the updated specification and provides the outcomes of executing it against the 2025 challenge dataset.

This paper is organized into three additional sections. Section 2 provides an overview of SDM-RDFizer, including its techniques, data structures, and physical operators used to optimize KG creation. Section 3 presents the results of the challenge, including a detailed description of the test cases and the required updates for execution. Finally, Section 4 offers concluding remarks and outlines future directions for SDM-RDFizer.

## 2. SDM-RDFizer

SDM-RDFizer [6, 7] is a KG creation engine capable of transforming structured (i.e., CSV and relational databases) and semi-structured data (i.e., JSON and XML) into RDF triples, following defined mapping rules. It supports the latest RML specification.

SDM-RDFizer adopts a two-fold approach to KG creation, consisting of two main modules: **Triples Maps Planning** (TMP) and **Triples Maps Execution** (TME). Each module plays a distinct role in the KG creation process and employs a set of data structures and operators to handle various aspects such as join execution and duplicate removal. TMP determines the execution order of RML Triples Maps (TMs) with the goal of minimizing memory usage. TME then generates the KG according to the execution plan established by TMP.

To interpret RML TMs, SDM-RDFizer uses a SPARQL-based parser. This parser employs four SPARQL queries: the main query extracts core mapping information, such as the logical source, *rml:subjectMap*, *rml:predicateObjectMap*, join conditions, logical dumps, function calls, etc. Meanwhile, the remaining queries handle nested structures like collections, functions, and fields for logical views.

To support the transformation of different types of TMs, SDM-RDFizer implements several specialized operators. The **Simple Object Map** (SOM) operator executes *rml:template* and *rml:reference*; the **Object Reference Map** (ORM) operator handles parent triples maps; and the **Object Join Map** (OJM) operator executes joins. For duplicate detection, SDM-RDFizer uses hash tables known as **Predicate Tuple Tables** (PTTs). Each generated triple is compared against the corresponding PTT; if it already exists, it is discarded as a duplicate. Otherwise, the triple is added to both the PTT and the KG. A **Dictionary Table** (DT) is used to compress the resources stored in PTTs. For join operations, SDM-RDFizer caches the results in a structure called the **Predicate Join Tuple Table** (PJTT) to avoid redundant join computations.

For the KGCW 2024 Challenge [5], SDM-RDFizer was extended to support the latest RML modules, including RML-FNML, RML-Star, and RML-IO. In the case of RML-IO, SDM-RDFizer was enhanced to accept new input types such as compressed files (e.g., ZIP, TAR), SPARQL endpoints, and remote data sources. It also gained the ability to compress the generated KG into various formats, split triples into multiple files, and export in different RDF serializations (e.g., JSON-LD, RDF/XML, Turtle).

To support RML-FNML, a new operator was introduced to execute functions during data transformation, incorporating strategies from FunMap [8]. FunMap is a TM translator that replaces TMs containing functions with equivalent TMs by executing the functions and transforming the input data accordingly. SDM-RDFizer applies the same principle, enabling real-time data transformation through function execution.

For RML-Star, a new operator was developed to handle the *rml:quotedTriplesMap* construct introduced in this module. Additionally, PJTT was extended to store joins that may occur in either the *rml:subjectMap* or the *rml:objectMap*.

Finally, SDM-RDFizer was further enhanced to support the RML-LV and RML-CC modules, which are discussed later in this report. The version of the SDM-RDFizer used is 4.7.5.12.2. SDM-RDFizer is publicly available on GitHub <sup>5</sup>.

---

<sup>5</sup><https://github.com/SDM-TIB/SDM-RDFizer>

### 3. KGCW 2025 Challenge Test Cases

The KGCW 2025 Challenge <sup>6</sup> aims to assess the compliance of existing KG creation engines with the latest RML specification. The dataset used in this challenge is a refined version of the KGCW 2024 Challenge dataset <sup>7</sup>, with a greater focus on cases that test the core functionality of the updated RML specification, along with new test cases (e.g., RML-LV). The dataset is composed of seven sets of test cases:

- **RML-Core:** This set includes basic test cases originally defined in the RML test suite <sup>8</sup> to evaluate the compliance of KG creation engines. While the 2024 dataset used CSV, JSON, XML files, and relational databases (MySQL and PostgreSQL) as data sources, the 2025 dataset focuses solely on JSON files. The remaining formats have been redistributed across the RML-IO and RML-IO-Registry modules. This revision aims to produce a more representative and concise set of test cases while reducing the requirements for achieving compliance with RML-Core.
- **RML-FNML:** This set includes test cases that apply functions to transform data using a set of predefined operations <sup>9</sup>.
- **RML-Star:** This set contains test cases for RDF-Star <sup>10</sup>, adapted from the RML-Star test suite <sup>11</sup> to conform to the updated specification.
- **RML-IO:** Formerly part of a unified module (which is now divided into RML-IO and RML-IO-Registry) in the 2024 dataset, this set now includes a wide range of remote data sources such as endpoints, compressed files, and JSON and XML files <sup>12</sup>. It also defines output specifications for various formats, including Turtle, RDF/JSON, JSON-LD, and compressed formats like ZIP and TAR.
- **RML-CC:** This set includes test cases that cover collections and containers <sup>13</sup>.
- **RML-LV:** This set features test cases where data sources are generated through projection and joins across multiple sources, including mixtures of different data formats <sup>14</sup>.

This work presents the results of executing all the modules with SDM-RDFizer. Table 1 shows the total number of test cases in the dataset and which cases were passed and failed by SDM-RDFizer. The full results are available on GitHub <sup>15</sup>.

#### 3.1. Results of RML-Core

RML-Core consists of test cases designed to validate fundamental RML functionalities, including the definition of classes, the use of *rml:template* and *rml:reference*, the execution of parent triples maps and joins, and the handling of data types, language tags, and named graphs. In the 2024 challenge, this module utilized a range of data sources such as CSV, JSON, XML, and relational databases. For the 2025 challenge, however, the dataset has been streamlined to focus exclusively on JSON files.

Given this focus, one of the primary challenges lies in correctly navigating nested JSON structures. To address this, SDM-RDFizer implements a recursive traversal mechanism that locates the relevant data by descending through the JSON hierarchy according to the specified iterator.

---

<sup>6</sup><https://zenodo.org/records/14970817>

<sup>7</sup><https://zenodo.org/records/10973433>

<sup>8</sup><https://kg-construct.github.io/rml-core/test-cases/docs/>

<sup>9</sup><https://kg-construct.github.io/rml-fnml/test-cases/docs/>

<sup>10</sup><https://kg-construct.github.io/rml-star/test-cases/docs/>

<sup>11</sup><https://zenodo.org/records/6518802>

<sup>12</sup><https://kg-construct.github.io/rml-io/test-cases/docs/>

<sup>13</sup><https://kg-construct.github.io/rml-cc/test-cases/docs/>

<sup>14</sup><https://kg-construct.github.io/rml-lv/test-cases/docs/>

<sup>15</sup>[https://github.com/SDM-TIB/SDM-RDFizer/tree/master/kgcw\\_2025\\_challenge](https://github.com/SDM-TIB/SDM-RDFizer/tree/master/kgcw_2025_challenge)

Module	Test Cases	Passed Cases	Cases Under Review	Cases Future Release
RML-Core	59	58	1	0
RML-FNML	17	16	1	0
RML-Star	18	18	0	0
RML-IO	73	64	9	0
RML-IO-Registry	103	70	31	2
RML-CC	35	35	0	0
RML-LV	32	32	0	0
Total	337	293	42	2

**Table 1**

Test Cases of the KGCW 2025 Challenge dataset.

SDM-RDFizer employs a parser query to process the input mappings. To comply with the latest RML specification, the parser has been updated to adopt the new *rml* namespace, remove references to the deprecated R2RML namespace, and revise the *rml:logicalSource* definition. This includes support for *rml:path* and *rml:root*, and replaces *rml:query* with *rml:iterator*.

SDM-RDFizer successfully executed 58 of the 59 test cases in the RML-Core module. RMLTC0019b-JSON is deemed an incorrect test case since it presents an empty expected output when it should generate triples. This test case significantly overlaps with RMLTC0019a-JSON, and in this test case, the expected output contains triples. This test case is being reviewed <sup>16</sup>.

### 3.2. Results of RML-FNML

RML-FNML includes test cases that apply functions for value transformations, such as replacing, concatenating, or changing the case of strings, based on the RML+FnO specification. SDM-RDFizer supports these transformations by leveraging FunMap [8], a tool that rewrites TMs containing functions into equivalent ones that reflect the function results. SDM-RDFizer also implements a dedicated operator that executes these functions on the fly, ensuring that the values incorporated into the KG are those produced by the function maps.

The parser query is extended to recognize function maps. An additional parser query is used to isolate functional maps in the case of nested functions. This enables proper handling of nested structures, as each function map is extracted and executed individually. SDM-RDFizer dynamically resolves function dependencies when one function receives the output of another.

SDM-RDFizer successfully completes 16 out of 17 test cases in this set. RMLFNMLTC0001-CSV involves a function that generates a random UUID, making it highly unlikely to match the expected output exactly. Nevertheless, the test is considered correct if a UUID is generated. In contrast, RMLFNMLTC0041-CSV is regarded as an incorrect case, as the function used in its triples map cannot produce the value presented in the expected output.

### 3.3. Results of RML-Star

RML-Star includes test cases based on RDF-Star, an RDF extension that introduces quoted triples, which allow a triple to be used as the subject or object of another triple. This capability supports the expression of metadata about statements—such as provenance, certainty, or attribution. To accommodate this feature, RML-Star introduces the *rml:quotedTriplesMap* construct for defining quoted triples within a KG. SDM-RDFizer supports this by implementing a dedicated operator capable of generating quoted triples, including recursively nested ones.

A particular challenge addressed in this module is the execution of joins in the *rml:subjectMap*, in addition to the more common use in the *rml:objectMap*. SDM-RDFizer handles both scenarios consistently using its OJM operator for join processing and PJTT for managing intermediate results. The parser was also extended to recognize and process *rml:quotedTriplesMap*.

<sup>16</sup><https://github.com/kg-construct/rml-core/issues/214>

RML-Star introduces a new type of TM, *rml:NonAssertedTriplesMap*. This type of TM is designed solely for generating quoted triples and does not contribute asserted triples to the KG. SDM-RDFizer treats these mappings as auxiliary, invoking them only when needed to produce quoted triples, and ignoring them otherwise.

SDM-RDFizer successfully executes all 18 test cases in this module.

### 3.4. Results of RML-IO and RML-IO-Registry

RML-IO and RML-IO-Registry modules comprise test cases that cover a wide range of input data source formats, including compressed files, JSON and XML documents, and data extracted from SPARQL endpoints. These modules also test the ability to generate the KG in various RDF serialization formats, such as Turtle, RDF/XML, and JSON-LD, and to compress outputs into formats like ZIP and TAR. Some cases introduce the concept of directing specific triples to designated output files. These modules aim to assess the capacity of KG creation engines to manage diverse input types and output requirements. Initially, the 2024 dataset included such test cases under RML-Core, but in the 2025 dataset, they are handled by these two modules.

RML-IO-Registry focuses on specialized source specifications and reference formulations, including datatype mappings, different encodings, and handling inputs containing comments.

To extract data from diverse sources, SDM-RDFizer utilizes several Python libraries: *csv* for CSV files, *json* for JSON, *xml* for XML, *mysql-connector* for MySQL, *psycopg2* for PostgreSQL, and *pyodbc* for SQL Server.

SDM-RDFizer uses the *requests* library to retrieve remote data. It employs the *SPARQLWrapper* library for SPARQL endpoints to execute queries and format results similarly to CSV. Compressed input files are downloaded locally and decompressed using appropriate libraries (e.g., *zip* for ZIP files). Output files are serialized into RDF formats using the *rdflib* library.

Some test cases explore the use of alternative output destinations for specific triples. These outputs can be defined within various mapping components, such as *rml:subjectMap*, *rml:predicateMap*, or *rml:objectMap*. Depending on the component, SDM-RDFizer directs the relevant triples to a specified output file. For example, if an alternate output is declared in the *subjectMap*, all triples from that mapping are sent to the alternative file. When defined in the *predicateObjectMap*, only the relevant triples to that predicate are redirected. Triples without an alternate destination go to the default output file set at execution. SDM-RDFizer prioritizes the creation of these auxiliary files and can compress them when needed.

SDM-RDFizer successfully executed 64 out of 73 test cases in the RML-IO module.

Test cases RMLSTC0006a and RMLSTC0006b contained blank spaces in the table headers used as data sources, leading to mismatches between the column names expected in the TMs and those present in the input<sup>17</sup>. Once these blank spaces were removed, SDM-RDFizer successfully executed both cases. They are therefore considered passed.

Test case RMLSTC0009a is expected to raise an error because the CSV input file contains column names enclosed in quotes, which is syntactically incorrect. However, the engine successfully produces an output since SDM-RDFizer is implemented in Python, and Python interprets these values without issue. It remains unclear whether this case should be marked as passed or failed. For now, it is considered passed until an official clarification is provided.

Test cases RMLTTC0004f and RMLTTC0004g specify incorrect file extensions for their logical target dumps<sup>18</sup>. Although SDM-RDFizer produces the correct output format, the wrong file extensions could cause confusion. These cases are also considered passed.

Finally, RMLTTC0002f, RMLTTC0002g, RMLTTC0002h, RMLTTC0002i, RMLTTC0002k, RMLTTC0002l, RMLTTC0002m, RMLTTC0002n, and RMLTTC0002r are considered incorrect and are being review<sup>19</sup>

<sup>17</sup><https://github.com/kg-construct/rml-io/issues/134>

<sup>18</sup><https://github.com/kg-construct/rml-io/issues/136>

<sup>19</sup><https://github.com/kg-construct/rml-io/issues/129>, <https://github.com/kg-construct/rml-io/issues/130>, <https://github.com/kg-construct/rml-io/issues/131>



due to incorrect graph assignments in the expected output.

SDM-RDFizer successfully executed 70 out of 103 test cases in the RML-IO-Registry module.

Test case RMLIOREGTC0003e failed due to a missing "@" symbol at the beginning of the triples map, resulting in a syntax error. Once corrected, SDM-RDFizer executes the case successfully, and it is therefore considered passed.

RMLIOREGTC0003d is considered incorrect and is being reviewed <sup>20</sup> since the defined namespace is not reflected correctly in the XML source. Thus, the iterator of the triples map is wrong.

RMLIOREGTC0004k, RMLIOREGTC0004l, RMLIOREGTC0005k, RMLIOREGTC0005l, RMLIOREGTC0006k, and RMLIOREGTC0006l are considered incorrect since the source query does not extract all the values needed for the triples map. Additionally, the triples require the "IDs" column, which doesn't exist in the database tables. These test cases are under review <sup>21</sup>.

RMLIOREGTC0004o, RMLIOREGTC0004t, RMLIOREGTC0004w, RMLIOREGTC0005o, RMLIOREGTC0005t, RMLIOREGTC0005w, RMLIOREGTC0006o, RMLIOREGTC0006t, and RMLIOREGTC0006w are considered incorrect and are under review <sup>22</sup> since the expected outputs express the values of the property "amount" in scientific notation, while the values in the database are in the standard decimal form. For example, "20.0" is expressed as "2.0E1" in the expected output. It is unclear if the SDM-RDFizer should convert the values from the database into their corresponding scientific notation or leave them as they are, since both are equivalent.

Test cases RMLIOREGTC0004y and RMLIOREGTC0006y expect the property "paid" to be represented as a boolean. However, when the source data is stored in MySQL or SQL Server, boolean values are expressed as integers "1" or "0". As a result, the extracted values are interpreted as integers, not booleans, causing a mismatch.

RMLIOREGTC0004z, RMLIOREGTC0005z, and RMLIOREGTC0006z are considered incorrect and are currently under review <sup>23</sup>. These test cases expect byte string outputs; however, the values are encoded differently when uploaded to the various databases, making it impossible to reproduce the expected results exactly.

Test case RMLIOREGTC0011a is considered incorrect because its SPARQL query is incomplete; it lacks the necessary prefix definitions. Additionally, the prefix declarations in the mapping end with ";", instead of "." as Turtle syntax requires. This case is currently under review <sup>24</sup>.

RMLIOREGTC0012b, RMLIOREGTC0012c, RMLIOREGTC0012d, RMLIOREGTC0012e, RMLIOREGTC0012f, RMLIOREGTC0012g, RMLIOREGTC0012h, and RMLIOREGTC0012i are incorrect and are under review <sup>25</sup> since the mappings of all test cases are in RML's old formulation.

Test cases RMLIOREGTC0009a and RMLIOREGTC0010a require access to Kafka and MQTT streams, respectively, to extract JSON input. A new version of SDM-RDFizer is currently being developed to support these streaming sources, but it is not publicly available.

### 3.5. Results of RML-CC

The RML-CC module contains test cases designed to generate RDF collections and containers. It introduces the terms *rml:gather* and *rml:gatherAs* into RML. The *rml:gather* term specifies which data elements should be grouped into a collection, while *rml:gatherAs* defines the type of RDF container or collection to be used (e.g., *rdf:Alt*, *rdf:Bag*, *rdf:List*, or *rdf:Seq*). The gathered data may consist of literals, IRIs, or results from join operations. In some cases, nested collections are formed—for example, an *rdf:Bag* might contain multiple *rdf:List* instances. The *rml:gather* term can be used within both the *rml:subjectMap* and *rml:objectMap*. Properly handling blank nodes is essential in this module, as they are intermediate nodes linking elements within a collection.

<sup>20</sup><https://github.com/kg-construct/rml-io-registry/issues/38>

<sup>21</sup><https://github.com/kg-construct/rml-io-registry/issues/36>

<sup>22</sup><https://github.com/kg-construct/rml-io-registry/issues/37>

<sup>23</sup><https://github.com/kg-construct/rml-io-registry/issues/39>

<sup>24</sup><https://github.com/kg-construct/rml-io-registry/issues/40>, <https://github.com/kg-construct/rml-io-registry/issues/43>

<sup>25</sup><https://github.com/kg-construct/rml-io-registry/issues/41>

SDM-RDFizer extends its parser query to support *rml:gather* and *rml:gatherAs*, and introduces additional logic to retrieve and manage nested collections. A new operator is implemented to collect data and generate the corresponding triples based on the specified collection type. This operator supports recursive application, enabling the construction of nested structures. When the collected data originates from join operations, SDM-RDFizer utilizes the PJTT data structure to store the intermediate results, which are accessed as needed during collection generation.

SDM-RDFizer successfully executed all 35 test cases from this module.

### 3.6. Results of RML-LV

RML-LV is a set of test cases that apply RML logical views to input data. This module is distinctive in that it focuses on generating new input data by combining and projecting existing sources and using various types of joins (e.g., inner and left joins). The output of a logical view is expected to be flat, meaning that the projected data must not retain any nested structures, even when the original input is a JSON file. RML-LV extends the definition of a TM's *logicalSource* by introducing the terms *rml:viewOn* and *rml:field*. The *rml:viewOn* term specifies the source of data for the logical view, while *rml:field* defines which values are extracted and their corresponding names. The fields of a Logical view may also be nested. RML-LV supports data format intermixing; for example, a JSON structure may contain values resembling CSV content.

SDM-RDFizer extends its parser to support logical views and implements a separate mechanism to extract nested views. A new operator is introduced and executed at the TMP module level, unlike the other operators executed at the TME module level; it produces raw values instead of RDF triples. Following the same design philosophy as with other data source formats, SDM-RDFizer ensures that each logical view is generated only once per execution. This operator can also be executed recursively to process nested logical views. When executing joins, SDM-RDFizer uses a modified version of the PJTT data structure, which stores raw values instead of RDF entities. In the case of data format intermixing, these values are extracted and loaded into memory as independent files of the corresponding type (e.g., JSON structures within a CSV file). Finally, all projected data is flattened to eliminate any nested structures.

SDM-RDFizer successfully executed all 32 test cases of this module.

## 4. Conclusions

The KGCW 2025 Challenge dataset evaluates the compliance of state-of-the-art engines with the new RML formulation. It comprises 337 test cases across seven modules: RML-Core, RML-IO, RML-IO-Registry, RML-FNML, RML-Star, RML-CC, and RML-LV. SDM-RDFizer successfully executed 293 of 337 test cases, covering all proposed modules. The remaining test cases are currently under review, given that they have errors that need addressing and two test cases will be addressed in a future release. The SDM-RDFizer is fully RML compliant.

To achieve this, SDM-RDFizer introduced a new parsing query and separate queries for the extraction of nested functions, collections, and views, the extension of existing data structures for the proper handling of the new operators, an operator for the execution of functions on the fly, an operator for generating quoted triples, an operator for the generation of RDF collection, and an operator for the creation of logical views. Moving forward, the authors will extensively test the operators implemented to comply with RML-CC and RML-LV (especially the flattening logical views) to determine that all border cases are appropriately covered. Additionally, the operator for the execution of RML-Star is currently being redesigned to improve the handling of intermediate results and remove redundant code. Furthermore, a new module is being developed to extract data from Kafka and MQTT streams. Finally, an improved mapping parser is planned to remove the reliance on ever more complicated parsing queries.

## Acknowledgments

This work was supported by the “Leibniz Best Minds: Programme for Women Professors”, through funding of the “TrustKG-Transforming Data in Trustable Insights” project (Grant P99/2020), and by the Lower Saxony Ministry of Science and Culture (MWK) with funds from the Volkswagen Foundation’s zukunft.niedersachsen program (CAIMed - Lower Saxony Center for AI and Causal Methods in Medicine; GA No. ZN4257).

## Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

## References

- [1] N. F. Noy, Y. Gao, A. Jain, A. Narayanan, A. Patterson, J. Taylor, Industry-scale knowledge graphs: lessons and challenges, *Commun. ACM* 62 (2019) 36–43. URL: <https://doi.org/10.1145/3331166>. doi:10.1145/3331166.
- [2] A. Dimou, M. Vander Sande, P. Colpaert, R. Verborgh, E. Mannens, R. Van de Walle, RML: A Generic Language for Integrated RDF Mappings of Heterogeneous Data, in: *Workshop on Linked Data on the Web*, 2014.
- [3] B. De Meester, A. Dimou, R. Verborgh, E. Mannens, An ontology to semantically declare and describe functions, in: *European Semantic Web Conference*, Springer, 2016, pp. 46–49.
- [4] T. Delva, J. Arenas-Guerrero, A. Iglesias-Molina, Ó. Corcho, D. Chaves-Fraga, A. Dimou, Rml-star: A declarative mapping language for rdf-star generation, in: O. Seneviratne, C. Pesquita, J. Sequeda, L. Etcheverry (Eds.), *Proceedings of the ISWC 2021 Posters, Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 20th International Semantic Web Conference (ISWC 2021)*, Virtual Conference, October 24-28, 2021, volume 2980 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2021. URL: <https://ceur-ws.org/Vol-2980/paper374.pdf>.
- [5] E. A. Iglesias, M. Vidal, Results for knowledge graph creation challenge 2024: Sdm-rdfizer, in: D. Chaves-Fraga, A. Dimou, A. Iglesias-Molina, U. Serles, D. V. Assche (Eds.), *Proceedings of the 5th International Workshop on Knowledge Graph Construction co-located with 21th Extended Semantic Web Conference (ESWC 2024)*, Hersonissos, Greece, May 27, 2024, volume 3718 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024. URL: <https://ceur-ws.org/Vol-3718/paper12.pdf>.
- [6] E. Iglesias, S. Jozashoori, D. Chaves-Fraga, D. Collarana, M.-E. Vidal, SDM-RDFizer: An RML Interpreter for the Efficient Creation of RDF Knowledge Graphs, in: *CIKM*, 2020. doi:10.1145/3340531.3412881.
- [7] E. Iglesias, M.-E. Vidal, D. Collarana, D. Chaves-Fraga, Empowering the sdm-rdfizer tool for scaling up to complex knowledge graph creation pipelines1, *Semantic Web* 16 (2025) SW-243580. URL: <https://journals.sagepub.com/doi/abs/10.3233/SW-243580>. doi:10.3233/SW-243580. arXiv:<https://journals.sagepub.com/doi/pdf/10.3233/SW-243580>.
- [8] S. Jozashoori, D. Chaves-Fraga, E. Iglesias, M. Vidal, Ó. Corcho, Funmap: Efficient execution of functional mappings for knowledge graph creation, in: J. Z. Pan, V. A. M. Tamma, C. d’Amato, K. Janowicz, B. Fu, A. Polleres, O. Seneviratne, L. Kagal (Eds.), *The Semantic Web - ISWC 2020 - 19th International Semantic Web Conference*, Athens, Greece, November 2-6, 2020, *Proceedings, Part I*, volume 12506 of *Lecture Notes in Computer Science*, Springer, 2020, pp. 276–293. URL: [https://doi.org/10.1007/978-3-030-62419-4\\_16](https://doi.org/10.1007/978-3-030-62419-4_16). doi:10.1007/978-3-030-62419-4\_16.