Towards a Mapping Framework for the Tenders Electronic Daily Standard Forms

Eugeniu **Costetchi**[1], Alexandros **Vassiliades**[1] and Csongor I. **Nyulas**[1]

[1]Meaningfy SARL, 61 route de Fischbach, L-7447, Lintgen, Luxembourg

Abstract

Knowledge graphs are frequently built using declarative rules to bridge diverse data sources to a desired ontology and materialise them as RDF. The materialisation of the full knowledge graph may be a complex task when these data sources are extensive, making it unsuitable for an "on-demand" materialisation. In this paper, we present a methodology on how to map Public Procurement Data from the Tenders Electronic Daily website of the European Union by using RML, based on a innovative idea of mapping partitions. We map the aforementioned data into the eProcurement Ontology, which is a popular ontology when it comes to representing public procurement data. We also provide a method of evaluating the quality of the mapped data by using a mechanism that produces SPARQL queries based on the conceptual mapping of the Tenders Electronic Daily website data into the eProcurement Ontology. We then give an empirical evaluation over the quality of the produced data, and provide a detailed discussion on what the method presented in this paper has to offer.

Keywords

Knowledge Graph Generation, RDF Mapping Language (RML), Tenders Electronic Daily (TED), eProcurement Ontology (ePO), Public Procurement Data, Conceptual Mapping, Technical Mapping.

1. Introduction

A knowledge graph (KG), consisting of a relative simple knowledge organisation and linking of a usually very large number of resources represented in RDF, is a suitable knowledge representation structure for any knowledge-based system. Moreover, the population of the KG is an aspect that is left as an after work and in most cases, procedural languages are used to map data into an ontology. Even more interesting is the fact that the data which is mapped into an ontology comes in the majority of cases from relational databases or other structured formats, such as tables or comma separated values, among others. The aforementioned formats, even though are helpful for mapping data with procedural languages, produce non-scaleable

mapping mechanisms.

The RDF Mapping Language (RML)[1] comes to tackle the problem of creating mapping mechanisms based on procedural languages or that are restricted to a single dataset (e.g. the RDB to RDF Mapping Language(R2RML)[2]), as RML offers a generic method, based on declarative rules, to map data into an ontology and that is not restricted to a single dataset [1, 2]. Moreover, there are additional advantages when mapping data with RML: (i) it decreases significantly the time required for mapping the data into the ontology, (ii) it reduces the maximum peak of memory usage, and (iii) it scales to data sizes that other engines are incapable of processing currently. For the aforementioned reasons and the fact that RML is recommended by W3C as the tool to map heterogeneous data into an ontology, RML has become an off-the-shelf solution when mapping heterogeneous data into an ontology.

In this paper, we present a methodology to map European Union (EU) Public Procurement Data (PPD) published on the public Tenders Electronic Daily (TED) website[3] into the eProcurement Ontology (ePO) [3, 4]. More specifically, we implement our methodology over a specific subset of the TED data called the Standard Forms for Public Procurement[4]. The method is based on mapping the various concepts that appear in each Standard Form into fragments of ePO, a procedure called Conceptual Mapping (CM) of the data. Then, based on this CM, we create our RML mapping rules that convert the data in the XML files, which encode the content of the filled out Standard Forms, into instances of ePO, a procedure that we call Technical Mapping (TM). Finally, we present a validation mechanism that automatically produces SPARQL queries and SHACL shapes rules to check the quality of the produced data from the mapping mechanism.

The motivation behind this paper lies in the fact that mapping heterogeneous data into an ontology is a difficult task which requires a sophisticated analysis of the data to be mapped in order to develop the RML mapping rules. For this reason, we propose a new methodology of mapping heterogeneous data based on the innovative idea of mapping partitions. Groups of mapping rules are referred to as mapping partitions when they produce distinct subsets of the knowledge graph. Our interest is in mapping TED data, which is of high importance and value for the EU citizen, into ePO ontology, which is an emerging semantic standard for PPD. Moreover, our motivation lies in the fact that the data that is produced should constantly be evaluated regarding its quality against the input data and the ontology in which is mapped to.

The key contribution of this paper consists in the methodology that we present for mapping PPD into fragments of ePO, by using a CM and a TM. More specifically, the CM offers: (a) the identification of the Business concepts in both the source and the target representations; (b) it serves as a source to generate validation tests; (c) it manages the complexity of mapping multiple versions of the source to a version of the target; and (d) organises the mapping rules in terms of mapping suites, as they are designed in Standard Forms. Next, the TM offers a generic mapping methodology for mapping heterogeneous PPD into the ePO ontology by using the RML mapping rules. In this mapping methodology we propose how to manage complexity by having the mapping rules being managed as incomplete fragments, some reusable and some specific to a "mapping suite" (i.e., Form number). Another contribution of the paper is the

---

[1]https://rml.io/specs/rml/

[2]http://www.w3.org/TR/r2rml/

[3]https://ted.europa.eu/TED/browse/browseByMap.do

[4]https://simap.ted.europa.eu/web/simap/standard-forms-for-public-procurement

validation mechanism that checks for the quality of the produced data, by automatically creating SHACL shape rules and SPARQL queries. Finally, we provide a set of Command Line Interface (CLI) tools available publicly[5] to anyone that can be used to aggregate what is needed for each mapping suite in a self-sufficient package.

The outline of this paper is the following. In Section 2, we present the related work to this paper. Next, in Section 3 we describe the nature of data, we give a high-level analysis of the ePO ontology, and we describe the RML mapping mechanism. We also present our validation mechanism which produces SHACL shape rules and SPARQL queries based on CM of the PPD Standard Forms, and we show the mapping suite dissemination. In Section 4 we present the validation report of our framework. We conclude our paper with a discussion over the resulting methodology by displaying some conclusions and proposing future work directions.

2. Related Work

The related work will be separated into two main subsections, one for generic mapping methodologies that use RML[1] or other mapping languages that exploit declarative rules, and one for methodologies that are concentrated to procurement data. It is worth mentioning that currently RML is perhaps the most commonly used method for mapping knowledge into an ontology, and its popularity is steadily increasing. For instance, the authors in [5] demonstrate how the use of standard declarative mapping rules (i.e., R2RML) guarantees a systematic and sustainable workflow for constructing and maintaining a KG.

Looking at generic mapping methodologies that attempt to map heterogeneous data into an ontology, by using RML or other mapping languages based on declarative rules, we observe that most of these efforts are either treated only at a theoretical level, or are tested only over a handful of context restricted datasets. More specifically, the studies [6, 7] offer a generic method on how to map heterogeneous data into ontologies, but do not test their method over any specific dataset. Next, the studies [8, 9, 10] also offer a generic methodology for mapping data into an ontology, and test their method into specific datasets, but are different from our context of PPD of TED. [8] and [9] use the SDM-Genomic-Dataset [7] and the GTFS-Madrid-Bench [11], while [10] uses the NPD Benchmark [12]. Another interesting aspect, when comparing the aforementioned studies with our study, is that they do not offer a validation mechanism to check the quality of their data. One could also read the thesis of David Chaves-Fraga [13], which gives a more complete view on how to map heterogeneous data into an ontology.

The area of mapping heterogeneous data from public procurement databases is quite rich, as well. We can find numerous studies, such as (i) [14], which uses data from the Public Procurement Pilot Experience and (ii) [15], which focuses on the European railway domain. Similarly to the first category of the related works, these studies fall into a different category of experiments with procurement data than our study. The aforementioned studies though do not provide a conceptual mapping in a commonly used ontology, and they also lack a validation mechanism that evaluates the quality of the produced data.

---

[5]https://github.com/meaningfy-ws/mapping-workbench

3. Methodology

In this section, we will start by presenting the mapping methodology overview, followed by a description of the nature of source data and the high-level structure of eProcurement ontology, which represents the mapping target.

In Figure 1, one can see the architecture of the framework presented in this paper. In the *Conceptual Mapping* layer, the relevant PPD Standard Forms from the TED website are selected (see Subsection 3.1) to be mapped in the CM. A sample dataset is created for the purpose of testing and validating the mapping rules, and, a Conceptual Mapping is created aligning business concepts, XML paths and ontology fragments. Next, in the *Technical Mapping* layer, the CM is being implemented using RML language[1]; this is referred to as *Create Technical Mapping*. Furthermore, the *sample dataset is transformed* with the implemented TM rules to enable quality control. In the third layer (*Validation*), we depict SPARQL and SHACL validation steps which evaluate the quality of the produced data, and if violations and inconsistencies are found the mechanism will point which parts of the CM seem to have an issue. Once the validation is passed successfully, in the fourth layer *Dissemination*, the *Mapping Suite is available* (for a Notice Type) and they are stored in the *Mapping Suite Repository*[6] to be used by the transformation pipeline, when necessary.

For the validation procedure, as well as for transforming data from the XML files of the Standard Forms, we offer a set of CLI tools that one can use in order to access, transform, and validate the data. The command line interaction tools can be found here[7], where a throughout documentation on how to use them is provided.

3.1. Nature of source data

The data we are mapping into ePO refer to the PPD that can be found in the Standard Forms of the TED[4]. These forms exist to help citizens to publish EU PPD in the Official Journal of the EU. The European Commission has created Standard Forms aligned with each of the EU legal bases in place for publishing this data, namely: (i) TED schema forms set out in Regulation (EU) 2015/1986 and (ii) eForms set out in Regulation (EU) 2019/1780. More specifically, currently we mapped forms F03, F06, F13, F20, F21, F22, F23 and F25[8], and we will be progressing with the remaining ones.

The TED Standard Forms that we are currently mapping to ePO are in PDF format, but they also have an XML counterpart[9], for each notice. We work with these XML notices, as it is a more appropriate format to map. By notice, we mean an instance of a completed form type, where form type are the different TED Standard Forms (i.e., form F03, F06, F13, etc).

The existing TED Standard Forms data needs to be transformed from XML into RDF format. Then, a set of XML transformation scripts can be developed. Once they are ready, they need to

---

[6]https://docs.ted.europa.eu/rdf-mapping/repository-structure.html

[7]https://github.com/OP-TED/ted-rdf-mapping

[8]see Standard Forms for Public Procurement (set out in Regulation (EU) 2015/1986) on SIMAP website: https://simap.ted.europa.eu/standard-forms-for-public-procurement

[9]see TED XML schemas (R2.0.9 & R2.0.8) for Standard Forms on EU Vocabularies website: https://op.europa.eu/en/web/eu-vocabularies/e-procurement/tedschemas
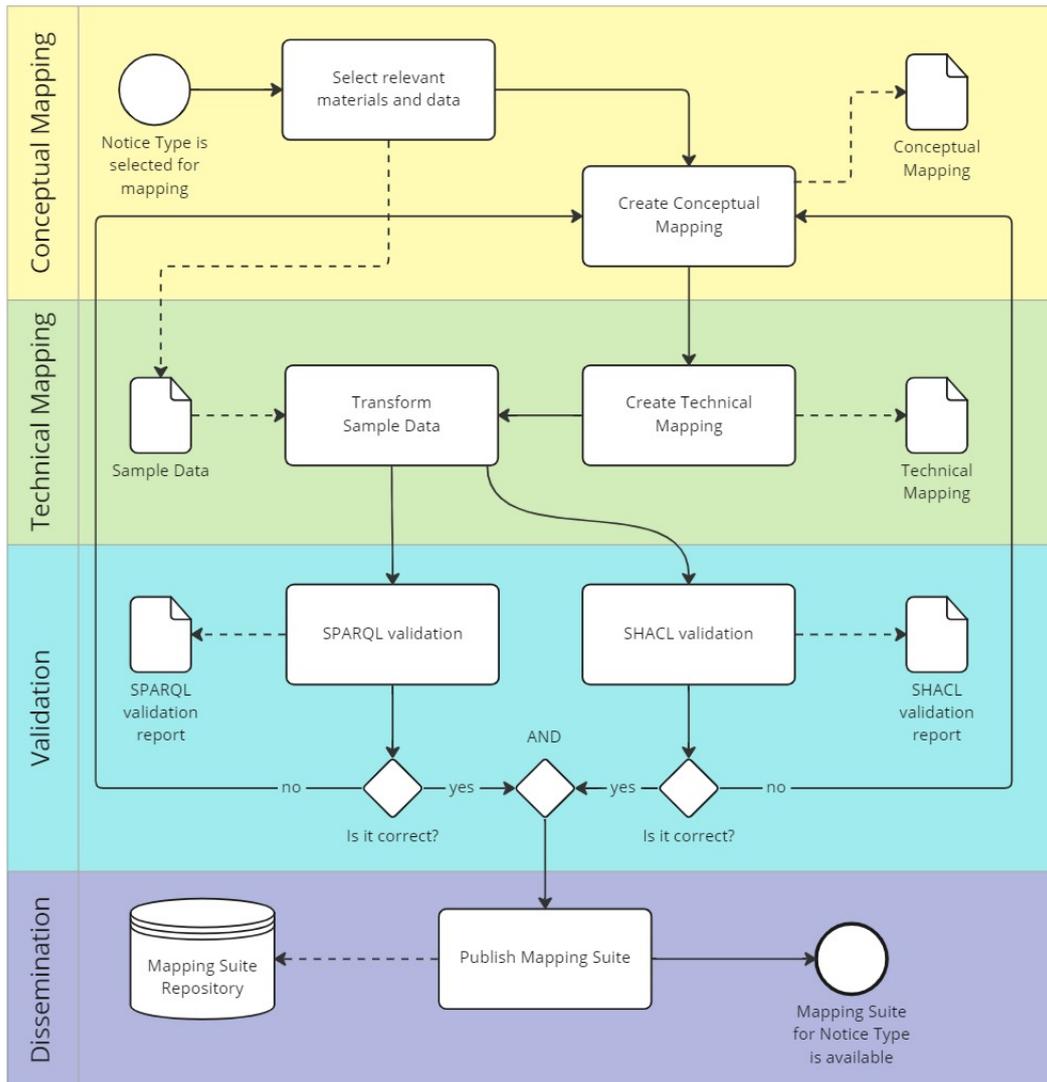
Figure 1: Mapping methodology workflow for the EU public procurement data

be executed on previously selected datasets, to convert them into RDF data instantiating the formal ontology.

3.2. Target ontology - eProcurement ontology

eProcurement ontology (ePO) [3, 4] is a semantic data model that conceptualises and formally encodes the knowledge representation of the public procurement domain. Its primary purpose is to bridge the interoperability gap in the European public procurement data space and can be used for data exchange, access, and reuse. The ePO ontology was created because vocabularies and the semantics that they are introduced through PPD, the phases of public procurement

that they are covering, and the technologies that they are using all differ. These differences hamper data interoperability and thus its reuse by the wider public. This creates the need for a common data standard for publishing public procurement data, hence allowing data from different sources to be easily accessed and linked, and consequently reused. ePO facilitates encoding procurement data and making it available in an open, structured and machine-readable format.

The ultimate objective of the eProcurement ontology project is to put forth a commonly agreed OWL ontology (and other necessary artefacts such as SHACL [16] data shapes and additional reasoning axioms) that will conceptualise, formally encode and make available in an open, structured and machine-readable format data about public procurement, covering end-to-end procurement, i.e., from notification, through tendering to awarding, ordering, invoicing and payment [17].

ePO offers a UML [18] representation[10] with which one could interact and get familiar with the ontology schema and the various object/data-type properties that it has. The ontology consists of about 140 classes, nearly 300 object properties, about 220 data-type properties, and uses more than 50 controlled vocabularies.

3.3. Conceptual Mapping

Let us begin a small example on how the CM works, in order to have an intuitive understanding first. Consider, the notice 113175-2023[11] in Section II and subsection II.1.1 has a title.

Then, for each notice that is similar to notice 113175-2023 (i.e., is of the same form) it is expected for the mapping rule to map the information about the title in the same section of the ontology.

The purpose of the CM is to map sections, subsections and fields of the PPD Standard Form into ePO ontology fragments, which are carefully chosen sequences of properties and classes to represent well the instantiation context. In Table 1, one can see a line (here reverted into a column) of the one concept from the forms translated into fragments of the ontology. Also, notice that some information was omitted here due to space restrictions.

- The **Form Number** row indicates in which form the concept from the **Section** is found.
- The **Standard Form Field Number** row indicates the name of the **Section** as found in the form.
- The **Field XPath** row indicates the XPath of the concept in the XML counterpart of the forms. This usually is generic for each form, as the XPath is the same for each notice in a form.
- The **Class Path** row indicates how the concept is being mapped into ePO class, in this case for the *Title* concept, a instance of the *epo:Procedure* is being created which is associated though a property with the class *xsd:string*.
- The **Property Path** row indicates how the concept is being mapped into ePO properties, in this case for the instance *Title* of the class *epo:Procedure* that is associated with the *xsd:string* class this is performed through the property *epo:hasTitle*.

---

[10]https://docs.ted.europa.eu/EPO/latest/_attachments/html_reports/ePO/index.html?goto=1:1:7:142
[11]https://ted.europa.eu/udl?uri=TED:NOTICE:113175-2023:TEXT:EN:HTML&src=0

Table 1
Information contained in a CM rule (concise form)

Form Number	3,6,13,20,21,22,23,25
Section	II.1.1
Standard Form Field Number	Title
Field XPath	*OBJECT_CONTRACT/TITLE*
Class Path	*epo:Procedure/xsd:string*
Property Path	*?this epo:hasTitle ?value*

3.4. Technical Mapping

The RML mapping mechanism refers to the declarative rules that convert the data from the XML files of the Standard Forms into RDF triples. The development of the mapping rules was more natural due to the pre-processing that we have done on our data, as the CM helped us understand to which class and property, we should map each element in the XML files.

We provide an example in this section to give an intuitive understanding of how the transition from the CM to the RDF occurred. Considering the information in Table 1, the idea behind this RML mapping rule is simple, as an instance of the class *epo:Procedure* will be created with a unique naming based on the URI of its XPath. This instance will be associated with a title that exists in the XPath, `ancestor::STANDARD_FORM_NUMBER/@LG`. `STANDARD_FORM_NUMBER` varies according to the Standard Form, also notice that *epo:* is the namespace prefix of the ePO ontology.

```
1    tedm:Procedure a rr:TriplesMap ;
2      rr:subjectMap
3          [
4              rr:template "ID" ;
5              rr:class epo:Procedure
6          ] ;
7      rml:logicalSource
8          [
9              rml:source "data/source.xml" ;
10             rml:iterator "/TED_EXPORT/FORM_SECTION/STANDARD_FORM_NUMBER/OBJECT_CONTRACT" ;
11             rml:referenceFormulation ql:XPath
12         ] ;
13     rr:predicateObjectMap
14         [
15             rr:predicate epo:hasTitle ;
16             rr:objectMap
17                 [
18                     rml:reference "TITLE";
19                     rml:languageMap [
20                         rml:reference "lower-case(ancestor::STANDARD_FORM_NUMBER/@LG)"
21                     ]
22                 ] ;
23         ] .
```

The URI creation, provided in subject map template, is based on a hashing function. This functionality is accessed through a REST call to a digest API. This guarantees unique reference

to the element. This mechanism of generating a unique deterministic URI is useful in both cases: (a) when generating the URI of an instance (in *rml:subjectMap*), and (b) when referring to the URI of an instance (in *rml:objectMap*).

Reflecting on the mapping rules, in most cases we managed to create generic rules that will apply over all Standard Forms, but that was not always the case as there were exceptions to that. Meaning that some mapping rules were restricted to a specific Standard Form, this usually occurred because some Standard Forms contain sections or subsections that were found only in a specific form.

In order to handle the complexity of mapping the Standard Forms into ePO we had to consider some baselines for the RML mapping rules to be more customizable. We have applied the following solutions:

- *Sectioning within a form*, meaning that we have mappings for each form section in order to increase maintainability. When any changes apply to a section, rules for other sections will not be affected.
- *Segregation of rules* (generic and form specific), meaning that there are generic files and a file per mapping suite.
- *Apply relative paths* in the mapping rules for handling versioning in the XML files
- *Reuse of rules across Standard Forms* and packages of Standard Forms means that there is a set of general source files where all the rules as a kept as single source of truth. There is a selection and packaging process that picks the necessary modules to form a unified, self-sufficient package for each Standard Form (see Section 3.6).
- *Management of* rml:TripleMap *parts*, meaning that we had to separate the statements of *rml:subjectMap* and *rml:logicalSource* in form-specific modules, whereas the statements of *rml:predicateObject* were contained in modules reused across forms. Only after an assembly of parts (and packaging) process, the mapping rules are integrated and executable (see Section 3.6).

3.5. Mapping Validation

The validation mechanism starts with transformation of the sample XML data with the RML mapping rules provided in the mapping suite. Then, from the conceptual mapping file, automatically is generated a test suite of SPARQL query assertions. These assertions are derived from the ontology fragments to which each XPath was mapped to. These assertions are then used for checking, for a mapping rule in the CM, if a specific ontology fragment is instantiated in the output file.

Moreover, the sample dataset is indexed for unique XPaths found in each sample XML file. This index is used for checking, for a mapping rule in the CM, if a specific input is present in the sample file.

The SPARQL-based validation of the transformed sample dataset includes, for each RDF file, first execution of all SPARQL query assertions, and second, asserting the present XPaths mentioned in the CM. The SHACL validation is streamlined to standard application of data shape files to each RDF output. The result is a set of reports that reflect the quality of the data which was produced by the RML mapping mechanism. In more detail, the validation mechanism

will do two things: (i) it will create for each line of the CM (see Table 1) a SPARQL query that checks if the data corresponding to the XPath mentioned in that line has been translated to an appropriate RDF triple, and (ii) based on the SHACL shape provided in the context of ePO repository[12] checks if the ontology is correctly instantiated.

Additionally, to the SHACL shapes and SPARQL queries we offer another form of evaluating the quality of the data, but this time the evaluation is performed on the input data, i.e., data to be mapped, in our case the XML files that represent the Standard Forms. This last form of evaluation refers to the XPaths of the concepts that exist in the CM and are about to be mapped in ePO. Basically, what the XPath "checker" does is to see if there exist or not an XPath for the concept in the XML file and if so, whether it is unique or not. The XPath checker serves a greater cause than just checking the existence or plurality of XPaths in the data, as it allows us to interpret violations of the SPARQL evaluator, i.e. the *unverifiable* assertions (when they fail on the output and no input for the rule is available either), the *warning* assertions (when they succeed in the output but no input for the rule is available). This it is possible to understand if the issue lies in the output data, in the input data or in the mapping rules (technical or conceptual). Section 4 provides a more detailed description of assertion severities.

3.6. Mapping suite dissemination

The mappings are part of a larger ecosystem where they are used for systematically transforming the TED notices. In this context, the mapping rules are being prepared as self-sufficient mapping packages called *mapping suites*. There is a governance procedure for how they are maintained, consumed and disseminated. In this section we focus mainly on how they are structured.

The mapping suites are maintained and published in a GitHub repository[13]. The maintenance is supported by a custom-built toolchain[14]. The repository from the mapping suites are ingested by the transformation pipeline is organised as follows.

- `/docs` folder contains this documentation. It is written in AsciiDoc format and compiled with Antora system [15].
- `/mappings` folder contains mapping suite packages organised based on the Standard Forms numbers. Their name is formed based on the form number (e.g. F03, F06) prefixed with package_ for readability.
- `/src/mappings` folder holds all the RML mappings files for all Standard Forms in a "single source of truth". Then because of the modularisation and reuse method adopted in this project, they are selected and packaged appropriately.
- `/test_data` folder contains sample TED notices selected with advanced search methods. These methods yielded the minimum number of the most representative notices given a form number and an XSD version. We omit to describe those selection mechanisms in this paper for brevity.

---

[12]see the eProcurement Ontology official GitHub repository https://github.com/OP-TED/ePO
[13]see the TED RDF mappings repository in GitHub https://github.com/OP-TED/ted-rdf-mapping
[14]see the mapping workbench toolchain repository in GitHub https://github.com/meaningfy-ws/mapping-workbench
[15]read more about AsciiDoc and Antora on https://antora.org/

- `/sampling_XX` subfolder contains the forms produced in the time frame XX, for example /sampling_2014-2021 refers to notices produced in from 2014 until 2021.

If we zoom into a mapping suite, for example */package_F03*, it will be composed of several elements assuring its completeness and self-sufficiency in ingestion, eligibility checking, transformation, validation and reporting processes undertaken by the transformation pipeline. Such a package also covers the needs in the development and testing of a given "mapping suite" (aimed for transformation of notices conforming to eligibility constraints e.g. of XSD version, form number, etc.).

- `metadata.json` automatically generated from Metadata sheet of `conceptual_mapping.xlsx` describing the parameters for selecting the notices that the mappings can be applied to, and various version information.
- `/transformation/conceptual_mappings.xlsx` is a CM specific to a form number.
- `/transformation/resources` contains additional resources necessary to apply the transformation rules. The content of this folder is automatically generated by the mapping package processor (part of the mapping workbench toolchain), based on the *Resources* sheet of the `conceptual_mappings.xlsx`.
- `/transformation/mappings/*.rml.ttl` the relevant RML transformation rules, organized in module files (copied from the "single source of truth" mappings folder) according to the specification in the "RML Modules" sheet of the `conceptual_mappings.xlsx`. **Note:** to streamline the transformation process in these rules the source XML is always pointing to `data/source.xml` file that will be copied (and renamed) from the `/test_data` folder at the time of the execution of the mapping test data.
- `/test_data` automatically selected test data (possibly grouped in suborders) that contain as few files as possible but which are the most representative and complete specimens in the entire data population.
- `/output` is a placeholder folder created at runtime to store outputs of the sample data transformation. It serves only when the mapping suite is being tested, or executed by some script. For each sample file we create a folder that will contain all the generated artefacts for that sample file, i.e. transformation output, validation reports, XML unique XPath index, etc.
- `/validation/sparql` contains all the SPARQL test suites, used in validation and development process.
- `/validation/shacl` contains all the SHACL test suites, used in validation and development process.
- `/validation/sparql_cm_assertions` contains the SPARQL assertion queries automatically generated from the conceptual mapping.

4. Empirical evaluation

In this section, we analyse briefly the results that the SHACL and SPARQL validators return and how we can interpret the results in order to optimise our mapping rules. Starting from the

Table 2
Sample SHACL validator results

SHACL Violation	Property of Application	Form	Notice
Value does not have class epo:LotSpecificTerm	epo:isSubjectToLotSpecificTerm	3,6,13 20,21,22 23,25	002705-2021 ... 654902-2021
Value does not have class epo:ProcedureSpecificTerm	epo:isSubjectToProcedureSpecificTerm	3,6,21 22,23,25	013921-2021 ... 359962-2021
Value does not have class epo:Duration	epo:hasQualificationSystemDuration	22	012969-2017 309175-2020
More than 1 values	epo:hasMainActivityDescription ... epo:indicatesAwardOfLotToWinner	3,6,13 20,21,22 23,25	135100-2021 ... 344048-2021
Less than 1 values	epo:hasProcurementScopeDividedIntoLot ... epo:unitType	3,6,13 20,21,22 23,25	000163-2021 ... 654902-2021

SHACL shape validator we can see that currently three types of violations that refer to missing class relations (i.e., an instance is not classified correctly), cardinality constraint for more than one value, and cardinality constraint for less than one value (see Table 2). Notice we display just a sample of violations due to space restrictions.

The second column indicates the property to which the violation applies to, the third the form in which it can be found, and the fourth the notice, i.e., Standard Form that in which it can be found. For the third column we can comment that due to clustering for display purposes the cardinality violations appear in all packages. The interpretation of errors is much easier with an analysis like this, as for example for the cardinality issues we can check if the constraint in ePO is perhaps too strict and needs to be relaxed, or the mapping rule needs to be modified.

Moving to the SPARQL evaluation we follow a similar analysis where we summarise the types of SPARQL inconsistencies, i.e., warnings, unverifiable queries, invalid queries, and errors.

- *error:* Refer to SPARQL queries which failed with an error, most likely because of incorrect SPARQL syntax or other technical issue.
- *invalid:* Refer to SPARQL queries which concern data that can be found in the input but not in the output. In other words some concepts from the input were not transformed into RDF triples.
- *unverifiable:* Refer to SPARQL queries which concern data that cannot be found either in the input or the output data.
- *warning:* Refer to SPARQL queries which concern data that cannot be found in the input but can be found, for mysterious reasons, in the output. In other words, some data was translated into RDF triples that it did not exist, but in reality there is always an explanation for such cases, for example two rules taking different inputs are mapped to the same ontology fragment.

The validation reports contain five result statuses: *Valid*, *Unverifiable*, *Warning*, *Invalid* and *Error.* Most of the results are Valid. Some are unverifiable because no input data in the sample triggers a mapping rule. Some Warnings are signalled in cases when the field is found in the output but not detected in the input. Invalid results are generated in cases when the data was found in the input but missing (or not detected by the current reporting tool) in the output.

Errors occur when the query is wrong or cannot be executed. No errors are acceptable, and none are reported in the current reports. Some Invalid results are found in the validation reports. Based on our analysis, they are not reflecting incorrect mapping rules or final data.

Out of a total of 1466 SPARQL queries, automatically generated from the CM, we find 0 errors, 43 invalid queries, 113 unverifiable, and 328 warnings. Table 3 shows a percentage of coverage for each type of inconsistency over the total number of queries. The aforementioned queries were distributed over 8 different Standard Form types and 650 notices. More specifically, for F03, F06, F13, F20, F21, F22, F23, F25 a set of 200, 195, 122, 146, 231, 231, 194, and 147 SPARQL queries were addressed, respectively.

Table 3
SPARQL Validator Result

Type of Inconsistency	Coverage
Error	0%
Invalid	2.9%
Unverifiable	7.7%
Warning	22.3%
Total	32.9%

The warnings and the unverifiable ones are not so relevant, as the first one might be the result of some aspects in the format we require the SPARQL query to have in order for our CLI tools to work correctly, such as title or a description in comments inside the query. For the unverifiable ones, in most cases the issue is a missing XPath in the input data. Nevertheless, we report the warning and unverifiable violations to have a complete view of the coverage of each violation. On the other hand, those that should be analysed and be taken more seriously are the invalid ones because in this case, (a) either the SPARQL query was not correctly generated by the SPARQL validator, (b) the ontology fragment in the CM is not correctly specified, or (c) there is an issue in the selected sample data.

5. Discussion

Our intuition, when mapping heterogeneous data into an ontology is that the existence of a pre-processing methodology before developing RML rules, is mandatory. In most cases, a CM seems to be a direction that eases significantly the task of developing RML rules, and also gives assurance of the quality of the produced data. The intuition behind the CM is to map concepts of the data into fragments of the ontology. The reason for which a pre-processing, such as a CM, is important, is because it allows us, on the one hand, to better understand how the mapping rules should be developed, and on the other hand, to check if the data was indeed mapped to the correct property and class after the mapping procedure took place.

Moving to the mapping rules, we consider that when we want to construct a generic mapping mechanism based on RML, we should keep in mind the following key points: (i) sectioning, meaning that the data, if possible, should be split into sections, as this will increase the maintainability of mapping rules, for example changes that are applied to the rules of one section should not affect the rules in other sections, (ii) segregation of generic and form specific rules,

meaning that there are generic files and a file per mapping suite, (iii) use of relative paths in the mapping rules for easier handling of versioning in the XML files, and (iv) reuse of rules across the data, by having a general source file where one would keep all the rules as single source of truth and then package, whatever is needed, for each data instance. The aforementioned key points can help us tackle the complexity of heterogeneous aspects in our data.

Another important aspect when creating a generic mechanism for mapping heterogeneous data into an ontology is the evaluation of the produced data. This procedure in the ideal scenario should be twofold. One should check, on the one hand, the quality of the data in the produced `*.ttl` files, and on the other hand, how the produced data fulfils the constraints posed by the ontology that we mapped to. For the first part, a combination of the XPaths along with a SPARQL-based evaluator, i.e., an evaluator that checks using SPARQL queries if each XPath from the input data has been mapped to a fragment of the ontology, seems ideal. Based on a mechanism like this, we can interpret the types of violations that each SPARQL query returns in order to correct our mapping rules, or maybe clean some noise from the input data. For the violations of the SPARQL queries we can comment that the Warnings and the Invalid are the most important ones, as the Warning violations indicate that the SPARQL query spotted something in the produced data that was not found in the input data, while the Invalid violations indicate that for something in the input data there was no corresponding output data found.

For the second part of the evaluation, i.e., the one checking if the produced data respect the constraints posed by the ontology, a SHACL shape validator that automatically extracts all the conditions from the ontology, and checks the SHACL shapes against the produced data, seems o be a suitable option. Such a SHACL validator is very helpful, as it indicates the types of errors in detail, and one could immediately change the mapping rule or the CM if necessary. Based on our experience, many SHACL shape violations are generated for instances not having as their type the class that they were supposed to. Besides being an obvious error in the mappings, this can also happen because the mapping mechanisms do not generate statements to describe the schema of the ontology, e.g. there are no subclass relations present in the produced files. Meaning that the instances which are shown to have missing classes might in fact be instances of the class that is indicated, but as an instance of some subclass of the indicated one. Taking into consideration also the ontology itself and enabling (basic) reasoning during the validation process would eliminate such false violation reports. Another big group of SHACL shape violations are due to cardinality constraints. Besides erroneous mappings, these kind of violations can happen also due to invalid input data, but most often they are due to over- or under-constrained properties in the ontology.

Concerning some potential limitations that are presented in our methodology, we can point out the following. Firstly, the CM is not automatically aligned to the versioning of the ePO ontology, that means that each time ePO updates if properties or classes are changed-renamed-deleted then we need to reflect this in the CM by hand. Similarly, in the TM, the mapping rules do not support versioning of ePO. Moreover, another limitation is the mandatory use of absolute paths in our TM, this is due to the fact that many paths are not unique and that results in using absolute paths in numerous instances in the iterators or join conditions of the mapping rules, but this also reduces the scalability of our TM as it may not be able to map all the existing Standard Forms. Finally, concerning the SHACL and SPARQL validators, we could say that a beautification to the summariser would be welcome.

As for future work, we plan to start mapping eForms[16] into ePO, as eForms will gradually replace Standard Forms for storing PPD in EU TED. We are also interested in supporting versioning of ePO, meaning that if changes apply to ePO we should be able to easily update our CM and TM to maintain the high quality of the generated data. Finally, we plan to further improve the quality of the mapped data, by analysing the various SHACL and SPARQL violations.

6. Conclusion

In this paper, we presented a mapping methodology that maps Public Procurement Data from the EU TED website into eProcurement Ontology. More specifically, we implemented our methodology over a specific subset of the TED data called the Standard Forms for Public Procurement. The method is based on mapping the various concepts of each Standard Form into fragments of ePO, a procedure called the Conceptual Mapping (CM) of the data, then based on this CM we developed our RML mapping rules which convert the data from the XML files that the Standard Forms are represented in, into instances of ePO classes, we call this procedure the Technical Mapping (TM). Finally, we presented a validation mechanism that automatically produces SPARQL queries and SHACL shapes to check the quality of the produced data.

We believe that based on the methodology presented in this paper, one could benefit from significantly when mapping heterogeneous data. Firstly, the existence of a CM allows for better "controlling" where the data will be mapped, it enables quality control for the produced data, and it also makes it easier to develop mapping rules. Next, the bullet points presented for the TM in subsection 3.4 show how we can better partition the data that we have to map, over mapping rules. Finally, the SPARQL and SHACL evaluators ensure to a great extent the quality of the produced data, by indicating where we need to fix or adjust a mapping rule, or change the mapping we have in the CM.

References

[1] D. Van Assche, T. Delva, G. Haesendonck, P. Heyvaert, B. De Meester, A. Dimou, Declarative rdf graph generation from heterogeneous (semi-) structured data: A systematic literature review, Journal of Web Semantics (2022) 100753.

[2] J. F. Sequeda, F. Priyatna, B. Villazón-Terrazas, Relational database to rdf mapping patterns., in: WOP, 2012, p. 1.

[3] M. Dekkers, E. Stani, F. Barthelemy, D02.02 - Project Charter proposal, Deliverable SC378DI07171, Publications Office of the European Union, 2017.

[4] M. Dekkers, E. Stani, B. Wyns, F. Barthelemy, D02.01 - Specification of the process and methodology to develop the eProcurement ontology with initial draft of the eProcurement Ontology for 3 use cases, Deliverable SC378DI07171, Publications Office of the European Union, 2017.

[5] D. Chaves-Fraga, O. Corcho, F. Yedro, R. Moreno, J. Olías, A. De La Azuela, Systematic construction of knowledge graphs for research-performing organizations, Information 13 (2022) 562.

---

[16]https://single-market-economy.ec.europa.eu/single-market/public-procurement/digital-procurement/eforms_en

[6] S. Jozashoori, D. Chaves-Fraga, E. Iglesias, M.-E. Vidal, O. Corcho, Funmap: Efficient execution of functional mappings for knowledge graph creation, in: The Semantic Web–ISWC 2020: 19th International Semantic Web Conference, Athens, Greece, November 2–6, 2020, Proceedings, Part I 19, Springer, 2020, pp. 276–293.

[7] 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: Proceedings of the 29th ACM International Conference on Information & Knowledge Management, 2020, pp. 3039–3046.

[8] J. Arenas-Guerrero, D. Chaves-Fraga, J. Toledo, M. S. Pérez, O. Corcho, Morph-kgc: Scalable knowledge graph materialization with mapping partitions, Semantic Web (2022) 1–20.

[9] J. Arenas-Guerrero, M. Scrocca, A. Iglesias Molina, J. Toledo, L. Pozo-Gilo, D. Dona, O. Corcho, D. Chaves-Fraga, Knowledge graph construction with r2rml and rml: an etl system-based overview, in: CEUR workshop proceedings., volume 2873, CEUR Workshop Proceedings, 2021, p. 1.

[10] D. Calvanese, B. Cogrel, S. Komla-Ebri, R. Kontchakov, D. Lanti, M. Rezk, M. Rodriguez-Muro, G. Xiao, Ontop: Answering sparql queries over relational databases, Semantic Web 8 (2017) 471–487.

[11] D. Chaves-Fraga, F. Priyatna, A. Cimmino, J. Toledo, E. Ruckhaus, O. Corcho, Gtfs-madrid-bench: A benchmark for virtual knowledge graph access in the transport domain, Journal of Web Semantics 65 (2020) 100596.

[12] D. Lanti, M. I. Rezk, G. Xiao, D. Calvanese, The npd benchmark: Reality check for obda systems, in: Advances in database technology-EDBT 2015: 18th International Conference on Extending Database Technology, Brussels, Belgium, March 23-27, 2015, proceedings, University of Konstanz, University Library, 2015, pp. 617–628.

[13] D. Chaves Fraga, Knowledge Graph Construction from Heterogeneous Data Sources exploiting Declarative Mapping Rules, Ph.D. thesis, ETSI_Informatica, 2021.

[14] C. Guasch, G. Lodi, S. V. Dooren, Semantic knowledge graphs for distributed data spaces: The public procurement pilot experience, in: The Semantic Web–ISWC 2022: 21st International Semantic Web Conference, Virtual Event, October 23–27, 2022, Proceedings, Springer, 2022, pp. 753–769.

[15] J. A. Rojas, M. Aguado, P. Vasilopoulou, I. Velitchkov, D. Van Assche, P. Colpaert, R. Verborgh, Leveraging semantic technologies for digital interoperability in the european railway domain, in: The Semantic Web–ISWC 2021: 20th International Semantic Web Conference, ISWC 2021, Virtual Event, October 24–28, 2021, Proceedings 20, Springer, 2021, pp. 648–664.

[16] H. Knublauch, D. Kontokostas, Shapes Constraint Language (SHACL), W3C Recommendation, W3C, 2017. Https://www.w3.org/TR/2017/REC-shacl-20170720/.

[17] N. Loutas, S. Kotoglou, D. Hytiroglou, D04.07 - Report on policy support for eProcurement, Deliverable SC245DI07171, ISA programme of the European Commission, 2016.

[18] S. Cook, C. Bock, P. Rivett, T. Rutt, E. Seidewitz, B. Selic, D. Tolbert, Unified Modeling Language (UML) Version 2.5.1, Standard formal/2017-12-05, Object Management Group (OMG), 2017. URL: https://www.omg.org/spec/UML/2.5.1.