

Insights from an OTTR-centric Ontology Engineering Methodology

Moritz Blum^{1,*,\dagger}, Basil Ell^{1,2,\dagger} and Philipp Cimiano¹

¹Bielefeld University, CITEC, Inspiration 1, 33619, Bielefeld, Germany

²University of Oslo, Problemveien 11, 0313 Oslo, Norway

Abstract

OTTR is a language for representing ontology modeling patterns, which enables to build ontologies or knowledge bases by instantiating templates. Thereby, particularities of the ontological representation language are hidden from the domain experts, and it enables ontology engineers to, to some extent, separate the processes of deciding about what information to model from deciding about how to model the information, e. g., which design patterns to use. Certain decisions can thus be postponed for the benefit of focusing on one of these processes. To date, only few works on ontology engineering where ontology templates are applied are described in the literature.

In this paper, we outline our methodology and report findings from our ontology engineering activities in the domain of Material Science. In these activities, OTTR templates play a key role. Our ontology engineering process is bottom-up, as we begin modeling activities from existing data that is then, via templates, fed into a knowledge graph, and it is top-down, as we first focus on which data to model and postpone the decision of how to model the data.

We find, among other things, that OTTR templates are especially useful as a means of communication with domain experts. Furthermore, we find that because OTTR templates encapsulate modeling decisions, the engineering process becomes flexible, meaning that design decisions can be changed at little cost.

Keywords

Ontology Engineering Methodology, Templates, Modularization

1. Introduction

Knowledge representation through ontologies plays a crucial role across multiple domains. However, creating large ontologies poses notable challenges. In particular, multiple design decisions must be made during the ontology development process, e. g., determining the relevant information to model and deciding on suitable ontologies or design patterns to build on. Additionally, collaborating with domain experts unfamiliar with the technical details of RDF or OWL further compounds the challenges faced in ontology development.

Just as modularization enabled the development of large software systems, this technique has found application in ontology engineering as well (see, e. g., [1]). Ontology templates


*Corresponding author.

^{\dagger}These authors contributed equally.

✉ mblum@techfak.uni-bielefeld.de (M. Blum); bell@techfak.uni-bielefeld.de (B. Ell);

cimiano@techfak.uni-bielefeld.de (P. Cimiano)

ORCID 0000-0003-4924-3903 (M. Blum); 0000-0002-8863-3157 (B. Ell); 0000-0002-4771-441X (P. Cimiano)

 © 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

embrace the concept of modularization, facilitating the construction of ontologies through the instantiation of templates. The type of modularization we realize is not that the content of an ontology is distributed into specific modules. Instead, in our case, templates encapsulate modeling decisions so that these modeling decisions do not need to be specified repeatedly and redundantly, which could lead to inconsistencies when the ontology evolves, but lead to uniformly modeled data. Despite the potential advantages of ontology templates and template instances, so far, no generally applicable ontology engineering methodology has made use of ontology templates as central design building blocks.

We propose an ontology engineering methodology that is centered around Reasonable Ontology Templates (OTTR) [2]. OTTR is a language and framework that introduces ontology templates, allowing to develop ontologies at a higher level of abstraction. The templates can be instantiated to construct ontologies.

OTTR allows to decouple the process of determining the information to be modeled from the process of deciding how to model it, thus allowing a more agile development approach. We propose to leverage this feature to postpone design decisions until the data is pre-structured. Additionally, the abstraction through templates simplifies the communication with domain experts by staying on a higher level without domain experts getting in touch with RDF.

We propose bottom-up ontology engineering and begin with existing data to start building the OTTR template headers. Domain experts actively participate and contribute their knowledge, which serves as a foundation for a discussion.

Subsequently, we adopt a top-down refinement approach, utilizing OTTR template headers as a starting point to develop OTTR template bodies iteratively. Thereby, the ontology content is continuously refined by establishing sub-templates that model certain parts of the data.

The template-based ontology engineering approach is related to the model-driven engineering methodology (MDE), in that a model is generated from a meta-model. In our case, the meta-model is a set of template definitions and instances and the generated model is an ontology.

In the remainder of this paper, we present our methodology by outlining the steps to take, which do not have to be followed strictly linearly. Additionally, we discuss different design decisions, like when to split a template into multiple templates. Moreover, we report our experience gained while developing a material science domain ontology, and we outline how we interact with domain experts and how we integrate real-world data to populate the ontology.

2. Related Work

In this paper, we focus on ontology engineering activities that make use of ontology templates. As we are not proposing a genuinely new ontology engineering methodology, we do not provide a comprehensive overview of existing ontology engineering methodologies.

Existing ontology engineering methodologies include the Cyc methodology [3], ONIONS [4], Methontology [5], CommonKADS [6], the On-To-Knowledge Methodology (OTKM) [7], UPON [8], DILIGENT [9], DOGMA [10], and the NeOn methodology [11]. Some of these methodologies are described and compared in a survey by Iqbal et al. [12] and in a book by Maria Keet [13]. A survey by Simperl et al. focuses on collaborative ontology engineering methodologies [14] such as DILIGENT. Complementing these methodologies that introduce pro-

cess models, there exist a wide range of ontology evaluation methodologies (see Vrandečić [15] for an overview), such as the OntoClean [16] and a list of OWL ontology antipatterns [17].

All of these methodologies have been developed before ontology templates were introduced. However, ontology templates can be seen as enabling ontology experts to develop domain specific languages (DSL). Thus, related work exists regarding design guidelines for DSLs [18]. We refer to the book by Martin Fowler [19] for a general introduction to DSLs. Mernik et al. [20] discuss when and how DSLs should be used. In the context of ontology engineering, an expressive and redundancy-free DSL could be developed that is understandable by domain experts and abstracts from the particularities of the ontological representation language (i. e., RDF(S) and OWL). E. g., Semantic Application Design Language (SADL) Crapo and Moitra [21] offers a controlled-English language that translates into OWL or SPARQL. In contrast to a DSL, not every OTTR template needs to be domain specific. Instead, ontological modeling is modularized such that there are user-facing domain-specific templates which are understandable by domain experts, but these templates make use of modeling-encapsulating templates that step-by-step guide the translation of data into an ontology.

The idea behind Ontology Design Patterns (ODP) is to have available components *that will support ontology design at the level that is more natural to domain experts and laymen, i.e. the level at which small, expertise-aware components can be assembled as easy-to-apply, easy-to-customize building blocks* [22]. According to Skjæveland et al. [2], OTTR templates have a similar ambition, but go one step further by enhancing the practicality of applying ODPs and ensuring their consistent implementation. Furthermore, OTTR templates abstract from the particularities of the ontological representation language.

Creating an ontology via OTTR template instantiation can be seen as a model-driven design method. Some works exist on model-driven ontology engineering [23, 24, 25] and make use of the Ontology Definition Metamodel (ODM), developed by the Object Management Group (OMG). ODM enables to translate OWL ontologies into UML and back, but it does not support encapsulation of modeling decisions, as it is the case with OTTR templates.

Lupp et al. [26] describe an ontology engineering methodology in which OTTR template headers are created from highly structured data by creating parameters that correspond to columns in tables. Therefore, this methodology focuses on a scenario where OTTR templates can be derived directly from the data. Moreover, they encounter a low number of data sources and no relations between those.

3. Preliminaries: OTTR

In this section, we provide a very brief introduction to Reasonable Ontology Templates (OTTR) [2], which is *a language with supporting tools for representing and instantiating RDF graph and OWL ontology modeling patterns*.¹

An OTTR template library is a set of OTTR template definitions. See Fig. 1 for an example. Each template definition consists of a header (e. g., `ax:SubClassOf[?sub, ?super]` – `ax:SubClassOf` is the template’s name; and `?sub` and `?super` are its parameters) and a body (enclosed by curly brackets). Template bodies can refer to other templates – thus, templates can

¹See <http://ottr.xyz>

be nested. `ottr:Triple` is the base template that creates RDF triples. A template instantiation denotes binding concrete values to the template parameters and to create triples.

In this paper we distinguish between *template instantiation*, i.e., the process of instantiating a template, and *template instance*, i.e., the result of instantiating a template. Furthermore, we distinguish between templates that are only instantiated from within other templates (such as `ax:SubClassOf` in Fig. 1), and those that are not (such as `pz:Pizza` in Fig. 1). The latter we call user-facing templates. User-facing templates are instantiated by domain experts, or these template instances are generated by some tool.

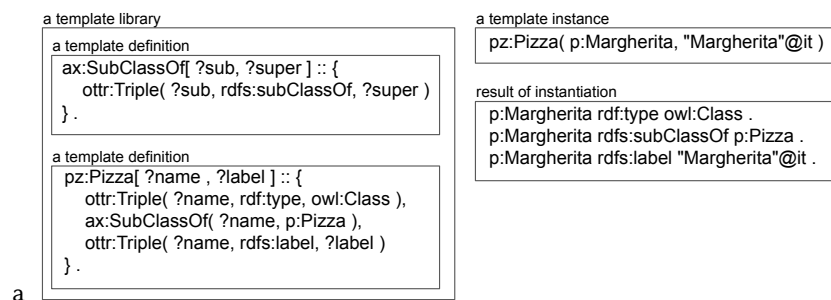


Figure 1: A pizza template library is shown on the left and an instance of the pizza template is shown on the right. The terms `p:Margherita` and `"Margherita"@en` are bound to the template parameters `?name` and `?label` to create an instance of the pizza template, whose triples are shown underneath.

We omit explaining parameter types and how type consistency is checked, optional parameters and default parameter values, lists, list expansion methods, and the different serialisation formats.

4. OTTR-centric Ontology Engineering Methodology

This section presents a list of steps and guidelines to develop an ontology using OTTR templates. It is important to note that this list of steps does not follow a strictly linear process model. In addition, we discuss our practical experience after most steps, gained in the context of the DiProMag project.² One of the ambitions of DiProMag is to model experiments related to materials, specifically magnetocaloric alloys, from the production, over the characterization to the prototypical application in an application ontology.

From the ontology engineering perspective, the project context favors a bottom-up (from data to ontology) ontology engineering approach, as properties of materials are measured by machines automatically, producing digital data that largely adheres to common data formats and standards. As a result, the structured data can be utilized in designing the template headers.

The subsequent steps to be followed are as follows: define the scope of the ontology and screen the available data (Sec. 4.1), design of template headers (Sec. 4.2), template header design verification and documentation (Sec. 4.3), design of template bodies (Sec. 4.4), handling of axiomatic triples (Sec. 4.5), template body documentation (Sec. 4.6), template library documentation (Sec. 4.7), and template instantiation & data integration (Sec. 4.8).

²See <https://www.dipromag.de>

The ontology engineering process involves close collaboration between ontology engineers and domain experts. The ontology engineers primarily manage most activities based on the insights and feedback provided by domain experts during various discussions. However, there are specific tasks that ontology engineers handle independently, such as identifying suitable existing ontologies and ODPs. Conversely, domain experts take the lead in documenting the process for inputting data into the templates, as it heavily relies on domain-specific knowledge.

Although developing OTTR templates requires ontology experts to possess additional skills, in the long run making use of OTTR templates simplifies ontology engineering, population, and maintenance.

4.1. Define the Scope of the Ontology and Screen the Available Data

In the first step, we follow the common practice of ontology engineering: we define the ontology's scope through competency questions (CQs) [27, 28] and let domain experts examine relevant data (e. g., datasets, records, and other information repositories) with those CQs. Based on this foundation, the ontology engineers engage in a discussion with the domain experts so that the ontology engineers understand the domain and the data's meaning, as well as its structure and parameters.

As in other methodologies, there is a risk of overlooking relevant information or the risk of modeling information irrelevant for the CQs. For example, modeling all sensor readings over time might be unnecessary if the CQs only ask for certain values. Therefore, this should be discussed with the domain experts with respect to the CQs which help to find missing source data and help to identify and filter out irrelevant data. Throughout the process, the CQs may evolve, being adapted and refined as domain experts continuously offer insights.

Experiences: Our CQs describe knowledge that would also be published in research papers, like results of experiments, but also details required for their reproduction. In addition, the ontology should be usable by machine-learning components to generate hypotheses about the properties of materials. For this reason, and to make the data easily queryable with SPARQL, it is essential to follow standards as much as possible, e. g., using the International System of Units, and to model the data uniformly. In general, our CQs are not designed specifically for our OTTR template based ontology engineering methodology, and, therefore, do not differ from CQs developed in other methodologies.

During data collection, we noticed a substantial degree of heterogeneity in the data management practices applied to the data. To facilitate future data integration, we not only collected representative data samples but also metadata like where the data is located (i.e., machine, volume, folder, etc.), how it can be accessed (i. e., file, database, service), how often it is updated, and the format of the files.

4.2. Design of Template Headers

After collecting data relevant to the design of the ontology, OTTR template headers can be developed. An OTTR template header mainly specifies a template by giving it a name and listing its parameters together with their data types. The abstraction through template headers

simplifies the communication with domain experts by discussing on the level of template names and parameters while abstracting from how the data is modeled in RDF. When designing template headers, trade-offs need to be made, taking the following principles into account:

- **Complexity of templates vs. complexity of template relations** If a template (header) has many parameters, then it can become difficult to use or update the template (header). Thus, it can be advisable to split a template header into multiple template headers while at the same time not distributing parameters that form a group over multiple templates. However, distributing information across multiple templates can also increase complexity for end users, as they need to understand how these templates are related so that they can be used correctly. The so-called workflows, that we describe in Section 4.3, are intended to help here.
- **Conceptual similarity** Data that is conceptually similar should be taken care of by the same template. In principle, one can follow standard practices from relational modeling. A template instance can be seen as a tuple in a relational model. Thus, design principles for entities and relationships from entity-relationship-modeling can be used when designing template headers [29].
- **Multiple similar but specific templates vs. few general templates** For a set of similar situations, either for each situation another template is created such that each of these templates almost models the same data. Thus, each template is specific to one situation. Alternatively, one template is created that fits multiple situations, where in each situation not all situation-specific parameters need to be used. For domain experts, situation-specific templates might be easier to understand and use. For ontology engineers, sets of situation-specific templates require more care to ensure uniform modeling across situations.
- **Avoid duplication wisely, but not at all costs** This principle from software engineering applies to template design, too. Although similar situations should be modeled similarly, which could mean with the same template, it can be that the similarity might be likely to disappear in the future when the ontology evolves. Here, the single-responsibility principle needs to be taken into account.
- **Avoiding data redundancies and data inconsistencies** Data redundancy can occur at the instantiation level and at the output level. The former means, that the same information is used in multiple template instances, e. g., in instances of different templates or of the same template. This can lead to inconsistencies, if the template instances need to reflect externally available data that can change, as changes in the external data require instances to be updated and in that case, one needs to identify all instances that need to be updated.

Data redundancy at the output level means that a triple or a set of triples is generated by more than one template instance. As an RDF graph is a set of triples, these redundant triples do not require additional memory. However, if an ontology engineer wishes to make changes such that a specific triple does not anymore occur in the output, then identifying all template instances that need to be adapted might be challenging.

- **Additional Metadata** Beyond modeling parameters for existing data, e. g., from files or databases, the ontology engineer can add further parameters to the template headers so that provenance information or any other meta data about the data generating process can be captured.

To maintain clarity and consistency throughout the template design, we document dependencies between parameters within one template and between templates. This practice will ensure appropriate modeling of the template bodies.

The result is an initial collection of template headers. It might be useful to improve them iteratively in collaboration with domain experts.

Experiences: As potential template candidates, we used one template per process step, i. e., one template per material type, one template per synthesis method, and one template per measurement method. The logical relation between those steps served as a basis for connecting the templates.

Next, we derived the template parameters directly from the available data. In cases where there is no data available to answer a certain CQ, the domain experts provided us with some artificially created data similar to the expected real data. This process directly involved domain experts that supported us with details about the different parameters and their data types.

The resulting template headers were iteratively improved to follow the introduced design principles. For example, one critical aspect of splitting data across multiple templates concerns the correlations and dependencies between data points. Although such connections may not be explicitly present in the raw data, modeling them explicitly in the ontology might be essential. E. g., two machines are used to measure a material's physical properties and the experiment's environmental conditions. Simply considering the data files, both appear to be separate and unrelated. However, the physical properties might only be meaningful in the context of the particular environment in which they were measured.

4.3. Template Header Design Verification and Documentation

Verifying the developed template headers is essential at this stage, as issues will propagate further to the development of the template bodies. To verify the headers of user-facing templates, domain experts can instantiate templates with sample data by informally filling out forms or tables where they do not need to focus on syntax and specific data types. As a result, design decisions from the header design (Sec. 4.2) might be revisited, because, e. g., some data structures (like complex JSON objects or tuples) are not represented correctly by the template parameters.

To simplify the instantiation of templates with (multiple) relationships to other templates, we propose to define workflows that describe how templates are related to each other. These workflows provide instructions for the instantiation order to build one single connected ontology instead of multiple unconnected ontologies which model certain parts of the information separately. Having such a proposed instantiation order has two benefits. First, domain experts have a clearly-defined sequence of steps to follow, which reduces the chance of forgetting a step. Second, with an instantiation order that ensures that the ontology is connected after each instantiation, inspecting and querying the ontology while working through a sequence of instantiations can become easier.

After creating the template headers, domain experts document user-facing templates to ensure easy understanding and usability by other domain experts. Ontology engineers only provide support on topics like the data types or the naming conventions.

The resulting documentation includes a general description about the purpose of the template (what data can be modelled; answering questions about when and how the template is applied), modeling limitations (explaining what is not modeled or what can not be modelled), and specific details like data types, the data a parameter name refers to, default values, and examples of parameter values. Moreover, the documentation provides valuable information for the development of the template bodies, e. g., the information about default values and example parameter values can serve as information to build a first taxonomy, and will ensure that domain knowledge is modelled correctly.

Experiences: We developed *OTTR Extension*,³ an extension for Semantic MediaWiki (SMW),⁴ which enables to use OTTR functionality within SMW. The *OTTR Extension* enables to define and instantiate OTTR templates through SMW's web interface and can generate convenient input forms for template headers automatically. We used this SMW to validate the user-facing template headers through a manual instantiation by the domain experts. As a result, we recognized minor issues with the developed template headers, e. g., a parameter was defined as a single value, whereas a list was necessary.

For the documentation of template headers, we encouraged documenting the parameters in tables (i.e., parameter name, data type, default values, example values, and description).

Furthermore, workflows recommending an instantiation order were developed based on the template dependencies to avoid mistakes when connecting templates through URIs. E. g., a material should be specified conceptually before modeling its measured properties.

4.4. Design of Template Bodies

Based on the template headers that define the shape of the input data and based on the provided documentation, the corresponding template bodies are developed by the ontology engineers.

As by other ontology engineering methodologies propose, we encourage the reuse of existing ontologies with approved modeling decisions inside the template bodies to ensure interoperability with other ontologies [30]. Ontology engineers make use of the template header documentation when identifying relevant ontologies and ODPs.

The development of the template bodies is straightforward, as a pre-structuring is done through developing template definitions and through the ontologies that will be reused. Here, we refer to established ontology engineering methodologies that can be used to develop the ontology inside the template bodies.

The major decisions of the ontology engineer with respect to OTTR are to create appropriate sub-templates. As our approach works top-down, templates with many parameters are the starting point to iteratively develop sub-templates.

As we consider our proposed methodology to be non-linear, steps back are always possible. Nevertheless, the template headers usually do not require significant adjustment caused by

³See <https://www.mediawiki.org/wiki/Extension:OtrParser>

⁴See https://www.semantic-mediawiki.org/wiki/Semantic_MediaWiki

issues encountered during the development of the bodies. The only conflict might be the incompatibility of the developed template library to the ontologies selected for reuse. An example is that a reused mid-level ontology defines a process that can not be modelled in a single template. Instead, the various steps are spread across multiple templates. In those cases, it must be discussed with the domain experts whether it makes sense to adjust the template headers or to use different ontologies and ODPs.

Where for the design of template headers the avoidance of redundancy, specifically data redundancy (see Section 4.2), is an important design principle, for the design of template bodies a relevant design principle is the avoidance of modeling redundancy. With modeling redundancy, we mean that multiple template bodies implement similar design decisions. This is problematic, as this can lead to modeling inconsistencies. Therefore, modeling redundancy should be avoided wisely, but not at all costs, too, and appropriate sub-templates should be used instead.

It is important to note that domain experts do not need to understand the template bodies. Instead, we propose to show them graph visualizations of the template relations and seek their feedback on their sufficiency. The main requirement is that the ontology can be queried to answer the competency questions, which can be done in ontology validation similar to [31].

Experiences: Drawing from mid-level ontologies and ODPs, we decided to build on the PMD core ontology⁵ and EMMO.⁶ The development of the template bodies was straightforward and the use of sub-templates lead to a faster development process as similar data occurs frequently across our data which can be covered by the same sub-templates. A lot of the pre-structuring was done in previous steps, and the header documentation served as a good source of information. We observed that a collection of similar physical experiments often share the same data about the environment the experiments take place, like temperature and humidity. This hints at the fact that a sub-template should be introduced to model the (shared) data about the environment.

4.5. Inclusion of Axiomatic Triples into the Ontology

When developing an ontology by designing and instantiating OTTR templates, an RDF file is the result of instantiating the templates. If the resulting ontology should contain axioms, for example that a certain property is symmetric, then the question is where to define these axioms so that they are contained in the resulting output.

We propose the following method. If a template body makes use of a term (i.e., the template body mentions it) about which we want to specify axioms, then the template's body instantiates another template that encapsulates all axiomatic triples about that term. Thus, that template is the only place where the axiomatic triples are stored, and an ontology engineer knows where to find these axioms. E. g., the domain and range axioms for the property *pz:hasTopping* would be added through a *pz:AxiomHasTopping* template call. Moreover, these axioms can be defined through template calls, e. g., *o-owl-ma:DomainRange(pz:hasTopping, pz:Pizza, pz:Topping)*.⁷

If a term is passed to a template for which no template exists that makes use of that term and could thus instantiate the template that generates the axioms about that term, then the

⁵See <https://github.com/materialdigital/core-ontology/>

⁶See <https://github.com/emmo-repo/EMMO>

⁷See <http://tpl.ottr.xyz/owl/macro/0.1/DomainRange>

above solution does not work. Instead, the solution is to develop templates that can then be instantiated about that term that then generate the axioms.

4.6. Template Body Documentation

To ensure the reuse of templates and the maintenance of their bodies, we propose a template body documentation that is achieved through comments in the OTTR code. This documentation is technical and intended for the ontology engineers (not the domain experts) to explain the functionality and purpose of the statements inside the template bodies.

Additionally, OTTR offers the DocTTR template library, designed to document OTTR templates using templates. The DocTTR⁸ template library includes various documentation elements, including change notes, examples, parameters, provenance information, and versioning [32].

Overall, the template body documentation enhances collaboration and knowledge sharing.

4.7. Template Library Documentation

The final documentation step has the goal to create an overview over the template library, including the structure, relationships, and functionality of the templates in the template library to the ontology engineers and domain experts.

The documentation process encompasses the creation of the following parts:

1. **List of all Templates:** The documentation should list all templates in the library, showing which templates are user-facing and which ones are sub-templates and thus play an internal and supporting role within the library.
2. **Template Inclusion/Call Hierarchy:** Understanding the template call hierarchy is crucial for understanding the dependencies among templates. This hierarchy can be derived automatically from the template library, but it is beneficial for both, ontology engineers and domain experts, to visualize and comment on it.
3. **Instantiation Order and Naming Guidelines:** Documenting the instantiation order is important to avoid mistakes when connecting templates. These instantiation orders must be recorded throughout the development process to provide a logical sequence for domain experts to follow. Additionally, guidelines for naming URIs are important to ensure consistency throughout the ontology.
4. **Individual Template Documentation:** This denotes the documentation already provided in Section 4.3 and Section 4.6. The documentation ensures that new users (domain experts & ontology engineers) understand, and existing users can look up, each template's functionality, ensuring a correct utilization of the template library.

By documenting the template library in the outlined way, ontology engineers and domain experts can efficiently find out which templates to instantiate for a given need and ontology engineers can efficiently maintain the template library in a structured and organized manner throughout its lifecycle.

For the publication of the template library, we propose to use either a Wiki or GitHub.

⁸See <http://tpl.ottr.xyz/p/docttr/0.1/>

Experiences: In our project, we use a SMW to organize and store the template library documentation along with the templates. We create pages to address the aspects mentioned before, and we make use of the semantic functionalities of the SMW to make the template library and the documentation queryable, e. g., to generate a list of all templates tagged to belong to a certain topic, project, or user.

We introduced the naming conventions into the documentation, as we observed that the naming of URIs and templates held significant importance for the domain experts. It allowed them to easily identify information within the ontology and ensured reusability in other projects.

4.8. Template Instantiation & Data Integration

Based on the template library, the ontology is generated by instantiating the templates with data to populate the A-Box, the T-Box, or both. Templates allow to build the T-Box step by step in contrast to traditional ontology engineering methods, which encourage collecting many details of the T-Box, e. g., to build the subclass hierarchy, before starting to populate the A-Box.

We have not yet explained how to integrate real-world data to instantiate templates and how to maintain the ontology. Therefore, a pipeline can automate the process of extracting, formatting, and feeding the data into OTTR templates (the information about the data sources collected, as described in Sec. 4.1, can be used) – thus, generating template instances. Then, the Lutra engine⁹ takes care of expanding the templates and to generate an ontology which can be stored in a triple store. This step is less about ontology engineering but more about data integration. Nevertheless, the data integration activities lead to templates instantiations being generated and templates being instantiated. New data can then reside in databases, as files, in the form of sensor readings etc. This process might involve aggregations of data, stream processing, complex event processing, synchronization, archiving of original data, dealing with data heterogeneity and quality issues, etc. Furthermore, it needs to be decided where the template instances should reside and where to store the results of the instantiations (i.e., in a triple store or as a file).

Experiences: In our case, the templates are instantiated either from data, or manually by domain experts, e. g., through an interface in SMW or in eLabFTW.¹⁰

In our case, eLabFTW generates template instances and sends these to our SMW (i.e., creates new wiki pages that then contain template instances), which returns unique identifiers for each template instance to allow changing the data later. If template bodies are changed by an ontology engineer, then the template bodies can be easily updated in SMW, and the triples resulting from instantiations will automatically update. Therefore, SMW serves as a platform for managing templates, their instances, and as a triple store.

Some template instances contain data obtained via experiments by the material science researchers involved in our project. It is important to distinguish between data that has already been reported in scientific publications and data that has not. In the latter case, this data should be kept internal. We used parameters to indicate the publication status. Our template library does not include sensitive data such as the template instances and can, therefore, be published.

⁹See <https://gitlab.com/ottr/lutra/lutra>

¹⁰See <https://www.elabftw.net>

5. Discussion

We would like to use the discussion to address two open questions regarding the use of ODPs in combination with OTTR and the extraction of the ontology from the templates definitions.

ODPs in combination with OTTR There are two ways of using ODPs in combination with our OTTR-centric ontology engineering methodology. The first method is to have a template that only instantiates a single ODP. This way, the ODP can be used in other templates easily. Then, the OTTR features like, e. g., list expansion, dealing with optional/default parameters, and checking (datatype and type) constraints, make using ODPs much easier than applying them manually. However, depending on the complexity of the ODP and the template library, this method is not always feasible. Therefore, the second method is to distribute an ODP across multiple templates. In this case, it is necessary to introduce new relations between templates such that the ODP is still followed.

Extraction of the ontology After the development of the OTTR template library, the question raises how to extract the ontology from the template definitions only. We recognize that the ontology is generated from a combination of the template library and the template instances. Therefore, extracting the complete T-Box from the template library alone is not a straightforward task, and it might be the case that the entire T-Box can only be obtained when also incorporating the template instances.

6. Limitations & Conclusion

In this paper, we outline an ontology engineering methodology that relies on OTTR templates for representing ontology modeling patterns. We present the key steps that include defining the ontology's scope, designing template headers, verifying and documenting them, incorporating axiomatic triples, creating and documenting template bodies, as well as documenting the entire template library, and instantiating templates and integrating data.

Additionally, we share experiences from a project, demonstrating how OTTR templates naturally support ontology engineering. We observed benefits such as flexibility in the case where modeling decisions are re-engineered and enhanced communication with domain experts.

What we propose in this paper is not a comprehensive ontology engineering methodology, because it should be combined with existing ontology engineering practices, e. g., for the development of competency questions, the creation of the template bodies, and for the evaluation of ontologies. The direction we are exploring is towards incorporating principles of modularization into the design of ontologies so that design decisions are not represented redundantly. Much research has been carried out on abstraction and modularization in the domain of software engineering, and we believe that there are multiple principles and techniques, from model-driven engineering and beyond, that could be adapted and applied to the field of ontology engineering.

Acknowledgments

This work was done in the context of DiProMag, a BMBF (German Federal Ministry of Education and Research) funded research project under Grant No. 13XP5120B, and Sirius with the Norwegian Research Council project No. 237898.

References

- [1] F.-Z. Hannou, V. Charpenay, M. Lefrançois, C. Roussey, A. Zimmermann, F. Gandon, The ACIMOV Methodology: Agile and Continuous Integration for Modular Ontologies and Vocabularies, in: 2nd Workshop on Modular Knowledge, 2023, pp. 1–13.
- [2] M. G. Skjæveland, D. P. Lupp, L. H. Karlsen, H. Forssell, Practical ontology pattern instantiation, discovery, and maintenance with reasonable ontology templates, in: The Semantic Web–ISWC 2018: 17th International Semantic Web Conference, Monterey, CA, USA, October 8–12, 2018, Proceedings, Part I 17, Springer, 2018, pp. 477–494.
- [3] C. Elkan, R. Greiner, Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project: D.B. Lenat and R.V. Guha, *Artificial Intelligence* 61 (1993) 41–52.
- [4] A. Gangemi, G. Steve, F. Giacomelli, ONIONS: An ontological methodology for taxonomic knowledge integration, in: *ECAI Workshop on Ontological Engineering*, volume 95, 1996, pp. 95–106.
- [5] M. Fernández-López, A. Gómez-Pérez, N. Juristo, Methontology: from ontological art towards ontological engineering, in: 1997 AAAI Spring Symposium, 1997, pp. 33–40.
- [6] G. Schreiber, *Knowledge engineering and management: the CommonKADS methodology*, MIT press, 2000.
- [7] Y. Sure, S. Staab, R. Studer, *On-To-Knowledge Methodology (OTKM)*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, pp. 117–132.
- [8] A. De Nicola, M. Missikoff, R. Navigli, A proposal for a unified process for ontology building: UPON, in: *DEXA Proceedings*, 2005, pp. 655–664.
- [9] H. S. Pinto, S. Staab, C. Tempich, Y. Sure, *Distributed Engineering of Ontologies (DILIGENT)*, Springer Berlin Heidelberg, 2006, pp. 303–322.
- [10] M. Jarrar, R. Meersman, *Ontology Engineering – The DOGMA Approach*, Springer Berlin Heidelberg, 2009, pp. 7–34.
- [11] M. C. Suárez-Figueroa, A. Gómez-Pérez, M. Fernández-López, *The NeOn Methodology for Ontology Engineering*, Springer Berlin Heidelberg, 2012, pp. 9–34.
- [12] R. Iqbal, M. A. A. Murad, A. Mustapha, N. M. Sharef, et al., An analysis of ontology engineering methodologies: A literature review, *Research journal of applied sciences, engineering and technology* 6 (2013) 2993–3000.
- [13] C. M. Keet, *An Introduction to Ontology Engineering*, University of Cape Town, 2018.
- [14] E. Simperl, M. Luczak-Rösch, Collaborative ontology engineering: a survey, *The Knowledge Engineering Review* 29 (2014) 101–131.
- [15] Z. Vrandečić, *Ontology Evaluation*, Ph.D. thesis, Karlsruhe Institute of Technology, 2010.

- [16] N. Guarino, C. Welty, Evaluating ontological decisions with ontoclean, *Commun. ACM* 45 (2002) 61–65.
- [17] C. Roussey, O. Corcho, L. M. Vilches-Blázquez, A catalogue of OWL ontology antipatterns, in: *Proceedings of the Fifth International Conference on Knowledge Capture, K-CAP '09*, Association for Computing Machinery, 2009, p. 205–206.
- [18] G. Karsai, H. Krahn, C. Pinkernell, B. Rumpe, M. Schneider, S. Völkel, Design Guidelines for Domain Specific Languages, in: *Proceedings of the 9th OOPSLA Workshop on Domain-Specific Modeling (DSM '09)*, volume B-108, 2009, pp. 7–13.
- [19] M. Fowler, *Domain-specific languages*, Pearson Education, 2010.
- [20] M. Mernik, J. Heering, A. M. Sloane, When and how to develop domain-specific languages, *ACM Comput. Surv.* 37 (2005) 316–344.
- [21] A. Crapo, A. Moitra, Toward a unified english-like representation of semantic models, data, and graph patterns for subject matter experts, *International Journal of Semantic Computing* 7 (2013) 215–236.
- [22] A. Gangemi, V. Presutti, *Ontology Design Patterns*, in: *Handbook on Ontologies*, Springer, 2009, pp. 221–243.
- [23] D. Gašević, D. Djurić, V. Devedžić, *Model driven architecture and ontology development*, volume 10, Springer, 2006.
- [24] Y. Pan, G. Xie, L. Ma, Y. Yang, Z. Qiu, J. Lee, *Model-Driven Ontology Engineering*, in: *Journal on Data Semantics VII*, Springer, 2006, pp. 57–78.
- [25] S. Cranefield, J. Pan, Bridging the gap between the model-driven architecture and ontology engineering, *International Journal of Human-Computer Studies* 65 (2007) 595–609.
- [26] D. P. Lupp, M. Hodkiewicz, M. G. Skjæveland, Template libraries for industrial asset maintenance: A methodology for scalable and maintainable ontologies, in: *CEUR Workshop Proceedings*, volume 2757, Technical University of Aachen, 2020, pp. 49–64.
- [27] M. Gruninger, Methodology for the design and evaluation of ontologies, in: *Proceedings IJCAI'95, Workshop on Basic Ontological Issues in Knowledge Sharing*, 1995, pp. 1–10.
- [28] V. Presutti, E. Daga, A. Gangemi, E. Blomqvist, eXtreme design with content ontology design patterns, in: *Proceedings Workshop on Ontology Patterns*, 2009, pp. 83–97.
- [29] P. P.-S. Chen, The entity-relationship model—toward a unified view of data, *ACM Trans. Database Syst.* 1 (1976) 9–36.
- [30] N. F. Noy, D. L. McGuinness, et al., *Ontology development 101: A guide to creating your first ontology*, 2001.
- [31] M. Booshehri, L. Emele, S. Flügel, et. al, Introducing the Open Energy Ontology: Enhancing data interpretation and interfacing in energy systems analysis, *Energy and AI* 5 (2021).
- [32] M. G. Skjæveland, D. P. Lupp, L. H. Karlsen, H. Forssell, Practical Ontology Pattern Instantiation, Discovery, and Maintenance with Reasonable Ontology Templates, in: *The Semantic Web – ISWC 2018*, 2018, pp. 477–494.