

Automated Reasoning Support for Ontology Development

Megan Katsumi and Michael Grüninger

Department of Mechanical and Industrial Engineering,
University of Toronto Toronto, Ontario, Canada M5S 3G8

Abstract. The design and evaluation of ontologies in first-order logic poses many challenges, many of which focus on the specification of the intended models for the ontology's concepts and the relationship between these models and the models of the ontology's axioms. In this paper we present a methodology for the verification of first-order logic ontologies, and provide a lifecycle in which it may be implemented to develop a correct ontology. Automated reasoning plays a critical role in the specification of requirements, design, and verification of the ontology. The application of automated reasoning in the lifecycle is illustrated by examples from the PSL Ontology.

1 Introduction

The design of ontologies is a complicated process, and there has been much work in the literature devoted to the development of methodologies to assist the ontology engineer in this respect. In this paper we present an approach to the semiautomatic verification of first-order logic ontologies, and describe a lifecycle in which it may be implemented to develop a correct ontology. Our objective is to use automated reasoning as the basis for formally defining the steps within a development methodology that focuses on the model-theoretic properties of the ontology. Although the ontology lifecycle described in this paper was developed to address issues that arise as a result of the semidecidable nature of first-order logic, the methodology and lifecycle is applicable to automated reasoning with less expressive languages.

We will discuss some existing methodologies for ontology development, and provide motivation for the methodology presented here. This work is a result of experiences in the development of an extension of the PSL ontology [1,7] that is sufficient to represent flow modelling notations such as UML, BPMN, and IDEF3. The focus of the paper will be a discussion of the role that theorem proving plays in each phase of the lifecycle, with an emphasis on pragmatic guidance for the semiautomatic verification of ontologies.

2 Ontology Development in Literature

High-level methodologies for ontology development tend to cover the breadth of the development process, however they do not provide techniques at the more detailed level.

The On-To-Knowledge Management (OTKM) [11] methodology covers the full lifecycle of ontology development. It provides useful insights into the steps required both pre- and post- application, but it does not provide exact details on how activities like testing should be performed, or what is to be done if an ontology fails to satisfy a requirement. METHONTOLOGY [2] covers areas of the lifecycle similar to what is presented in the OTKM methodology, with a focus on the explanation of the concepts at each stage in development. DILIGENT [10] provides guidance to support the distributed (geographically or organizationally) development of ontologies; in particular, the authors present a method to enable the adaptation of an ontology from different parties.

Low-level ontology design methodologies provide detailed instruction, concentrating on means to accomplish some necessary step in ontology development. The methodology that was used in the design of the TOVE Ontology for enterprise modelling [5] introduced the notion of competency questions to define the requirements of an ontology and to guide the formal definition of a set of axioms for the ontology. The Enterprise Ontology [13], also arising from work with enterprise ontologies, presents an approach for ontology capture that emphasizes the notion of “middle-out” design, giving equal weight to top-down issues (such as requirements) and bottom-up issues (such as reuse).

The evaluation of an ontology is generally accepted as a key process in its development. The main efforts in this area focus on taxonomy evaluation [3]; to the best of our knowledge, there have been no efforts towards evaluation methodologies that are semantically deeper than this. This is possibly due to the issue of semidecidability of first-order logic and the inherent intractability of theorem provers that presents a challenge for the evaluation of test results. High level methodologies typically do not specify the requirements in a verifiable form. In any case, the existing methodologies do not sufficiently address the issue of ontology verification, and our goal is to address this hole, in particular, the challenges encountered in first-order logic development.

3 The Ontology Lifecycle

The lifecycle presented in Figure 1 is intended to serve as a structured framework that provides guidance during the ontology development process. This lifecycle addresses the limitations of existing verification techniques while providing a structure within which existing techniques for ontology design and requirements identification may be applied. The feedback and interactions between the various phases illustrates the tight integration between the design and the verification of the ontology. Throughout this section, we illustrate the efficacy of this lifecycle framework using examples from the design and maintenance of the PSL Ontology¹.

¹ The Process Specification Language (PSL) [1,7] has been designed to facilitate correct and complete exchange of process information. PSL has been published as an International Standard (ISO 18629) within the International Organisation of Standardisation and the full set of axioms (which we call T_{psl}) in the Common Logic Interchange Format is available at <http://www.mel.nist.gov/psl/ontology.html>. This paper only uses the PSL Ontology as the context of a case study in ontology development; the focus is not on the content of the PSL Ontology, but rather on how various axioms were proposed and revised at different stages of the ontology lifecycle.




Fig. 1. The Ontology Lifecycle

We define five phases in the ontology lifecycle:

- The Requirements Phase produces a specification of the intended models of the ontology.
- The Design Phase produces an ontology to axiomatize the class of models that captures the requirements. The feedback loop that is shown in Figure 1 between the Design Phase and the Requirements Phase occurs as the ontology develops. In the Requirements Phase the intended models must initially be specified informally, until the design of the ontology has matured such that its requirements may be specified using its vocabulary.
- The Verification Phase guarantees that the intended models of the ontology which are specified in the Requirements Phase are equivalent to the models of the axioms which are produced in the Design Phase.
- The Tuning Phase addresses the pragmatic issue of dealing with cases in which the theorem provers used in the Verification Phase fail to return a definitive answer.
- The Application Phase covers the different ways in which the ontology is used, such as decision support systems, semantic integration, and search.

Each phase in the ontology lifecycle is associated with a set of reasoning problems that are defined with respect to the axioms of the ontology T_{onto} , a domain theory Σ_{domain} that uses the ontology, and a query Φ that formalizes the competency questions, and/or intended models produced in the Requirements Phase. Since these reasoning problems are entailment problems of the form:

$$T_{onto} \cup \Sigma_{domain} \models \Phi$$

we can utilize theorem provers to verify that the axiomatization of the ontology satisfies its requirements. The relationships between the stages of the ontology lifecycle and the different aspects of the reasoning problems are shown in Figure 2.

The Design Phase provides the axioms $T_{onto} \cup \Sigma_{domain}$ that form the antecedent of the reasoning problem. In the Verification Phase, we use theorem provers to determine whether or not the sentences that capture the requirements are indeed entailed by the axioms of the ontology.




Fig. 2. Reasoning problems and stages of the ontology lifecycle

3.1 Requirements

Requirements for an ontology are specified with respect to the intended semantics of the terminology, and the challenge of the initial phase of the ontology lifecycle is to cast these requirements in a testable form. From a mathematical perspective the requirements may be characterized by their intended models; to allow for semiautomatic verification we can specify these semantic requirements as entailment problems with the use of competency questions. From a reasoning problem perspective the output of the Requirements Phase is a set of consequents for the entailment problems described above.

For more mature ontologies, we may be able to use the representation theorem, to completely characterize the requirements with the use of a relative interpretation. In this case, the output of the Requirements Phase is the set of resulting entailment problems stating the necessary and sufficient conditions for the relative interpretation; these entailment problems differ slightly from Figure 2 and will be presented later in more detail.

Intended Models. In current ontology research, the languages for formal ontologies (such as RDFS, OWL, and Common Logic) are closely related to mathematical logic, in which the semantics are based on the notion of an interpretation. If a sentence is true in the interpretation, we say that the sentence is satisfied by the interpretation. If every axiom in the ontology is satisfied by the interpretation, then the interpretation is called a model of the ontology. With a formal ontology, the content of the ontology is specified as a theory, so that a sentence is consistent with that theory if there exists a model of the theory that satisfies the sentence; a sentence can be deduced if it is satisfied by all models of the theory. Therefore, the semantics of the ontology's terminology can be characterized by this implicit set of models, which we refer to as the set of intended models.

We aim to define the requirements for a *semantically correct* ontology with the relationship between the intended models for the ontology, and the actual models of its axiomatization, (see Figure 3). An axiomatization is semantically correct if and only if it does not include any unintended models, *and* it does not omit any intended models. Formally, we define these potential *semantic errors* as follows:

Definition 1. *An error of unintended models is present in the ontology if and only if there exists any model \mathcal{M} such that*

$$\mathcal{M} \in Mod(T_{onto}) \Rightarrow \mathcal{M} \notin \mathfrak{M}^{intended}$$

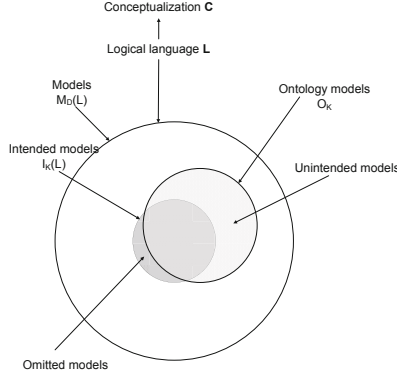


Fig. 3. The relationship between intended models for an ontology and the ontology’s axioms (from [9]). In this paper we refer to the models of the ontology O_k as $Mod(T_{onto})$, and we refer to the intended models $I_k(L)$ as $\mathfrak{M}^{intended}$.

Definition 2. An error of model omission is present in the ontology if and only if there exists any model \mathcal{M} such that

$$\mathcal{M} \in \mathfrak{M}^{intended} \Rightarrow \mathcal{M} \notin Mod(T_{onto})$$

Intended models are specified with respect to some well-understood class of mathematical structures (such as partial orderings, graph theory, and geometry). The extensions of the relations in the model are then specified with respect to properties of these mathematical structures.

For example, the intended models for the subactivity occurrence ordering extension to the PSL Ontology are intuitively specified by two properties:

1. the partial ordering over the subactivity occurrences;
2. the mapping that embeds the partial ordering into the activity tree.

Formally, the intended models for the subactivity occurrence ordering extension are defined by

Definition 3. Let \mathfrak{M}^{soo} be the following class of structures such that for any $\mathcal{M} \in \mathfrak{M}^{soo}$,

1. \mathcal{M} is an extension of a model of T_{psl} (i.e. the PSL Ontology);
2. for each activity tree τ_i , there exists a unique partial ordering $\varrho_i = (P_i, \prec)$ and a mapping $\theta : \tau_i \rightarrow \varrho_i$ such that
 - (a) $\langle s_1, s_2, a \rangle \in \mathbf{min_precedes} \Rightarrow \theta(s_2) \not\prec \theta(s_1)$
 - (b) $\langle \theta(s), s, a \rangle \in \mathbf{mono}$;
 - (c) comparable elements in ϱ are the image of comparable elements in τ .
3. $\langle s, a \rangle \in \mathbf{soo}$ iff $s \in P_i$;
4. $\langle s_1, s_2, a \rangle \in \mathbf{soo_precedes}$ iff $s_1 \prec s_2$.

One can take this approach to explicitly specify the intended models as a class of mathematical structures- however, early in development we often lack the knowledge to do so. Therefore, in practice initially the specification of the intended models is based on use cases for the application of the ontology. The two primary application areas for the PSL Ontology have been in semantic integration and decision support systems. In this approach, the intended models are defined implicitly with respect either to the set of sentences that are satisfied by all of the intended models or to sets of sentences that should be satisfied by some model. If we recall that sentences satisfied by all models are entailed by the axiomatization, then this leads us to the idea of implicitly defining the intended models of the ontology with respect to reasoning problems. The reasoning problems that are associated with the Requirements Phase are competency questions (which arise primarily from decision support use cases), and relative interpretations (derived from the specification of intended models).

Representation Theorems. We use representation theorems to formalize the definition of semantic correctness for our ontology. They are proven in two parts – we first prove every structure in the class is a model of the ontology and then prove that every model of the ontology is elementary equivalent to some structure in the class.

For the new extension T_{soo} to the PSL Ontology, the representation theorem is stated as follows:

Theorem 1. *Any structure $\mathcal{M} \in \mathfrak{M}^{soo}$ is isomorphic to a model of $T_{soo} \cup T_{psl}$.
Any model of $T_{soo} \cup T_{psl}$ is isomorphic to a structure in \mathfrak{M}^{soo} .*

The characterization up to isomorphism of the models of an ontology through a representation theorem has several distinct advantages. First, unintended models are more easily identified, since the representation theorems characterize *all* models of the ontology. We also gain insight into any implicit assumptions within the axiomatization which may actually eliminate models that were intended. Second, any decidability and complexity results that have been established for the classes of mathematical structures in the representation theorems can be extended to the ontology itself. Finally, the characterization of models supports the specification of semantic mappings to other ontologies, since such mappings between ontologies preserve substructures of their models.

Representation theorems are distinct from the notion of the completeness of an ontology. A logical theory T is complete if and only if for any sentence Φ , either $T \models \Phi$ or $T \not\models \Phi$. The ontologies that we develop are almost never complete in this sense. Nevertheless, we can consider representation theorems to be demonstration that the ontology T_{onto} is complete with respect to its requirements (i.e. set of intended models \mathfrak{M}^{onto}). This allows us to say that $T_{onto} \models \Phi$ if and only if $\mathfrak{M}^{onto} \models \Phi$ (that is, the ontology entails a sentence if and only if the class of intended models entails the sentence).

The typical way to prove the Representation Theorem for an ontology is to explicitly construct the models of the ontology in the metatheory and then show that these models are equivalent to the specification of the intended models of the ontology using classes of mathematical structures. An alternative to the mathematical specification of the representation theorem is to employ a theorem prover to verify that the models of the ontology are equivalent to the intended models. Depending on the maturity of the

ontology, this can be accomplished in two ways - with the use of competency questions or a relative interpretation.

Competency Questions. Following [5,8,12], competency questions are queries that impose demands on the expressiveness of the underlying ontology. Intuitively, the ontology must be able to represent these questions and characterize the answers using the terminology. Examples of competency questions for the subactivity occurrence ordering extension include the following:

Which subactivities can possibly occur next after an occurrence of the activity a_1 ?

$$\begin{aligned} & (\forall o, s_1) \text{occurrence}(s_1, a_1) \wedge \text{occurrence}(o, a) \\ & \quad \wedge \text{subactivity_occurrence}(s_1, o) \\ \supset & (\exists a_2, s_2) \text{occurrence}(s_2, a_2) \wedge \text{next_subocc}(s_1, s_2, a) \end{aligned} \quad (1)$$

Does there exist a point in an activity tree for a after which the same subactivities occur?

$$\begin{aligned} & (\exists a, a_1, s_1) \text{subactivity}(a_1, a) \wedge \text{occurrence_of}(s_1, a_1) \\ \wedge & ((\forall o_1, o_2) \text{occurrence_of}(o_1, a) \wedge \text{occurrence_of}(o_2, a) \\ & \quad \wedge \text{subactivity_occurrence}(s_1, o_1) \\ & \quad \wedge \text{subactivity_occurrence}(s_1, o_2) \\ & \quad \wedge \text{min_precedes}(s_1, s_2, a) \\ \supset & (\exists s_3) \text{subactivity_occurrence}(s_3, o_2) \\ \wedge & \text{min_precedes}(s_1, s_3, a) \wedge \text{mono}(s_2, s_3, a) \end{aligned} \quad (2)$$

Recall that sentences such as these constitute the consequent Φ of a reasoning problem (see Figure 2), and that they are supposed to be entailed by the axioms of the ontology T_{onto} together with a domain theory T_{domain} (which in this case is a process description that formalizes a specific UML activity diagram). It is in this sense that competency questions are requirements for the ontology – there must be sufficient axioms in $T_{\text{onto}} \cup T_{\text{domain}}$ to entail the sentences that formalize the competency questions.

In the area of decision support, the verification of an ontology allows us to make the claim that any inferences drawn by a reasoning engine using the ontology are actually entailed by the ontology's intended models. If an ontology's axiomatization has unintended models, then it is possible to find sentences that are entailed by the intended models, but which are not provable from the axioms of the ontology.

To specify requirements, competency questions can be developed from use cases (as above), or with the goal of approximating the Representation Theorem. The relationship between competency questions and the requirements is that the associated query must be provable from the axioms of the ontology alone. Since a sentence is provable if and only if it is satisfied by all models, competency questions implicitly specify the intended models of the ontology. Rather than proving the Representation Theorem directly, we can utilize competency questions for the specification of the requirements for the ontology. We can do this by proving that the extensions of the relations in the models of

the ontology have properties that are equivalent to those satisfied by the extensions of the relations in the intended models. For example, the subactivity occurrence ordering ϱ introduced in the definition of the intended models \mathfrak{M}^{soo} is a partial ordering, and this corresponds to the competency questions which asserts that the *soo-precedes* relation is also a partial ordering, and hence is a transitive relation:

$$\begin{aligned} & (\forall s_1, s_2, s_3, a) \text{soo_precedes}(s_1, s_2, a) \\ & \wedge \text{soo_precedes}(s_2, s_3, a) \supset \text{soo_precedes}(s_1, s_3, a) \end{aligned} \quad (3)$$

This illustrates the technique of identifying competency questions directly from the mathematical definition of the intended models.

Relative Interpretation. With more mature ontologies, we may have a better understanding of both the ontology, and its intended models. If we are able to identify and axiomatize the class of intended models, a theorem about the relationship between the class of the ontology's models and the class of intended models can be replaced by a theorem about the relationship between the ontology (a theory) and the theory axiomatizing the intended models (assuming that such axiomatization is known). We can use automated reasoners to prove the latter relationship and thus verify an ontology in a (semi-)automated way. The relationship between two theories, T_A and T_B , is the notion of interpretation, which is a mapping from the language of T_A to the language of T_B that preserves the theorems of T_A . We will say that two theories T_A and T_B are definably equivalent iff they are mutually interpretable, i.e. T_A is interpretable in T_B and T_B is interpretable in T_A . The key to representing the requirements for a relative interpretation as entailment problems is the following theorem of reducibility from [6]:

Theorem 2. *A theory T is definably equivalent with a set of theories T_1, \dots, T_n iff the class of models $Mod(T)$ can be represented by $Mod(T_1) \cup \dots \cup Mod(T_n)$.*

The necessary direction of a representation theorem (i.e. if a structure is intended, then it is a model of the ontology's axiomatization) can be stated as

$$\mathcal{M} \in \mathfrak{M}^{intended} \Rightarrow \mathcal{M} \in Mod(T_{onto})$$

If we suppose that the theory that axiomatizes $\mathfrak{M}^{intended}$ is the union of some previously known theories T_1, \dots, T_n , then by Theorem 2 we need to show that T_{onto} interprets $T_1 \cup \dots \cup T_n$. If Δ is the set of translation definitions for this relative interpretation, then the necessary direction of the representation theorem is equivalent to the following reasoning task:

$$T_{onto} \cup \Delta \models T_1 \cup \dots \cup T_n \quad (\text{Rep-1})$$

The sufficient direction of a representation theorem (any model of the ontology's axiomatization is also an intended model) can be stated as

$$\mathcal{M} \in Mod(T_{onto}) \Rightarrow \mathcal{M} \in \mathfrak{M}^{intended}$$

In this case, we need to show that $T_1 \cup \dots \cup T_n$ interprets T_{onto} . If Π is the set of translation definitions for this relative interpretation, the sufficient direction of the representation theorem is equivalent to the following reasoning task:

$$T_1 \cup \dots \cup T_n \cup \Pi \models T_{onto} \quad (\text{Rep-2})$$

By Theorem 2, $Mod(T_{onto})$ is representable by $\mathfrak{M}^{intended}$ iff $T_1 \cup \dots \cup T_n$ is definably equivalent to T_{onto} , which we can show by proving both of the above reasoning tasks. Note that the requirement **Rep-2** varies from Figure 2 in that the axioms being designed for the ontology form the consequent of the reasoning problem. Also both **Rep-1** and **Rep-2** differ in that they include a set of translation definitions as opposed (or in addition) to a domain theory. Regardless, all requirements specified as entailment problems are verified in the same way with the use of an automated theorem prover.

3.2 Design

Recall that the Requirements Phase specified intended models either implicitly or explicitly. The task of the Design Phase is to produce a set of axioms that will entail exactly these intended models. The reasoning problems associated with the Design Phase focus on the notion of the design rationale for axioms. The task in this case is to characterize why specific axioms are required, tracing each axiom back to the original set of requirements for the ontology. While we do not prescribe a specific methodology for the design of the axioms, a key task to assist the process of axiom design in early stages of development is model exploration.

In early phases of the ontology lifecycle, we often lack a complete understanding of the models of the ontology; in fact, we may not even possess sufficient knowledge to formally express our requirements. In model exploration, we can use an automated model builder to generate models of the axioms in an effort to develop an understanding of both the intended structures and the models of the ontology. We attempt to generate models with specific sizes and/or properties. These specific properties are specified by asserting the existence of elements satisfying certain relations in the ontology.

For example, if T_{onto} is a satisfiable axiomatization of an ontology with some n -ary predicate $P(a_1, a_2, \dots, a_n)$ with finite $n \geq 1$ in the language of T_{onto} . Then, we might generate a model for T_{\exists} , where:

$$T_{\exists} = T \cup \{ \exists x_1, x_2, \dots, x_n [\text{all } x_k \text{ distinct} \wedge P(x_1, x_2, \dots, x_n)] \}$$

Models generated in this way serve to illustrate the semantics of the ontology's axioms.

3.3 Verification

As we saw earlier, existing approaches to ontology verification focus on taxonomic relationships and obvious logical criteria for ontology evaluation such as consistency. Strictly speaking, we only need to show that a model exists in order to demonstrate that a first-order theory is consistent. Constructing a single model, however, runs the risk of having demonstrated satisfiability for a restricted case; for example, one could construct a model of a process ontology in which no processes occur, without realizing that the axiomatization might mistakenly be inconsistent with any theory in which processes do occur. We therefore need a complete characterization of all models of the ontology up to isomorphism.

In general, verification is concerned with the relationship between the intended models of an ontology and the models of the axiomatization of the ontology. From a mathematical perspective this is formalized by the notion of representation theorems. From a

reasoning problem perspective this amounts to evaluating the entailment problems, and we are able to perform this verification semiautomatically with the use of a theorem prover.

Outcomes of Verification. As we have just seen, the reasoning problems associated with Verification support the proofs of the Representation Theorem. When verifying any semantic requirement there are three possible outcomes, as illustrated in Figure 4. Each outcome is discussed in more detail below.

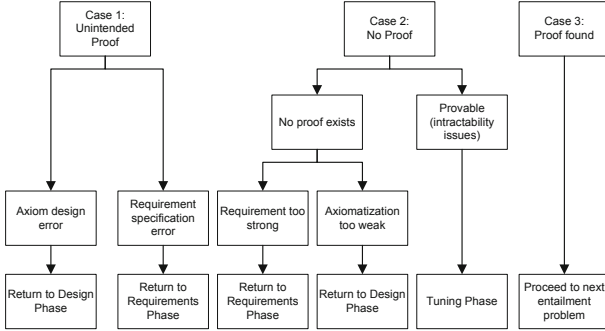


Fig. 4. Outcomes of Semantic Requirement Verification

Case 1: Unintended Proof Found. In this case, a proof was found of a sentence that contradicts the intended semantics of the ontology. This is often encountered when the theorem prover finds a proof without using all, or any clauses from the proposition or query. Given this possibility, a thorough inspection of all proofs found must be performed; if this case is detected, it is indicative of at least one of two possible errors with the ontology:

1. An examination of the proof may lead to the identification of some axiom in the ontology which is not entailed by the intended models; in this case we must return to the Design Phase.

One such result with the design of the subactivity occurrence ordering extension T_{soo} to the PSL ontology arose when testing its consistency with the first addition of T_{soo} . A proof was found, and normally this would indicate an error in the set of definitions that was being tested. However, upon examination we realized that the definition of the *same_tree* relation made it inconsistent for any model of the ontology to have an occurrence in the same activity tree as the root of the tree:

$$T_{soo} \cup T_{psl} \models (\forall s_1, a) \text{root}(s_1, a) \supset \neg(\exists s_2) \text{same_tree}(s_1, s_2, a)$$

This sentence should not be entailed by the intended models of the PSL Ontology itself, yet it was only identified when using the axioms of T_{soo} ; it was a hidden consequence of the PSL Ontology. In particular, it was the axiom below from T_{soo} that played the critical role in deriving the unintended sentence:

$$\begin{aligned}
& (\forall s_1, a) \text{root}(s_1, a) \supset (\exists s_2) \text{same_tree}(s_1, s_2, a) \\
& \quad \wedge \text{mono}(s_1, s_2, a) \wedge \text{soo}(s_2, a)
\end{aligned} \tag{4}$$

As a result of this discovery, the axiom for *same_tree* was modified.

It is interesting to see the relationship between this case and the failure of a potential representation theorem for the ontology. Part of the representation theorem shows that a sentence that is entailed by axioms is also entailed by the set of intended models. Hidden consequences such as we have just considered are counterexamples to this part of the representation, since it is a sentence that is provable from the axioms, yet it is not entailed by the intended models. One can either weaken the axioms (so that the sentence is no longer provable), or one can strengthen the requirements by restricting the class of intended models (so that the sentence is entailed by all intended models).

2. An examination of the proof may lead to the detection of some error in the definition of the requirements; in this case we must return to the Requirements Phase. It is important to devote considerable attention to the detection of this possibility so that the ontology is not revised to satisfy incorrect requirements. We did not encounter an example of this in our experiences, however it would be possible for an error in the requirements specification to lead to an unintended proof. As discussed above, this type of error could be addressed by strengthening the requirements.

Case 2: No Proof Found. As a result of the semi-decidability of first-order logic and the inherent intractability of automated theorem proving, if no proof is found when testing for a particular requirement then a heuristic decision regarding the state of the lifecycle must be made. It could be the case that no proof is found because the sentence really is not provable; there may be a mistake in the definition of the requirement that we are testing, or there may be an error in the axiomatization of the ontology, (i.e. we cannot prove that the requirement is met because it is not met). However, it could be the case that due to the intractability of automated theorem provers, a proof exists but the theorem prover is unable to find it (at least within some time limit). To avoid unnecessary work, some effort must be made to ensure that we are as informed as possible when making this decision; in particular, previously encountered errors and the nature of the requirement that we are attempting to prove is satisfied must be taken into account. In some cases, we may be able to use a model building tool to search for a counter-example, potentially resolving the uncertainty in this case. We discuss this possibility further in the following section; regardless, if no proof is found we must choose between the following options:

1. If we believe there may be some error in the requirements or the design of the ontology, then we must revisit the Requirements Phase or the Design Phase, respectively. It is recommended that the Requirements Phase is revisited first, as generally a much smaller investment is required to investigate the correctness of a requirement, rather than that of the ontology.

In the course of proving the representation theorem for T_{soo} , the theorem prover failed to entail a property regarding the *mono* relation, namely, that the *mono* relation should only hold between different occurrences of the same sub-activity. This requirement was initially expressed as the following proposition:

$$\begin{aligned}
& (\forall s_1, s_2, o, a) \text{mono}(s_1, s_2, a) \wedge \text{occurrence_of}(o, a) \\
& \quad \wedge \text{subactivity_occurrence}(s_1, o) \\
& \quad \supset \neg \text{subactivity_occurrence}(s_2, o)
\end{aligned} \tag{5}$$

It seemed clear that the axiomatizations of the *mono* relation should have restricted all satisfying models to instances where s_1 and s_2 were not in the same activity tree. Returning to the Requirements Phase, an examination of the above sentence led to the realization that the axiomatization of the proposition had been incorrect. We had neglected to specify the condition that s_1 was not the same occurrence as s_2 , (in which case the *subactivity_occurrence* relation clearly holds for s_2 if it holds for s_1). Once this issue was addressed, we continued to the Verification Phase (no revisions were required to the ontology's design) and the theorem prover was able to show that the ontology entailed the corrected property, shown below.

$$\begin{aligned}
& (\forall s_1, s_2, o, a) \text{mono}(s_1, s_2, a) \wedge \text{occurrence_of}(o, a) \\
& \quad \wedge \text{subactivity_occurrence}(s_1, o) \wedge (s_1 \neq s_2) \\
& \quad \supset \neg \text{subactivity_occurrence}(s_2, o)
\end{aligned} \tag{6}$$

The previous example was a case where the error that was corrected resulted from a misrepresentation of the intended semantics of the requirements. Another interesting situation was encountered where the requirements' semantics were redefined following a series of inconclusive test results. Originally, T_{soo} was to be a conservative extension of PSL². In part, this meant that the axiomatization of T_{soo} had to account for all of the kinds of activity trees that were represented in PSL. One particular class of activity trees was represented by the zigzag relation, defined below.

$$\begin{aligned}
& (\forall s_1, s_3, a) \text{zigzag}(s_1, s_3, a) \\
& \quad \equiv (\exists s_2) \text{preserve}(s_1, s_2, a) \\
& \quad \wedge \text{preserve}(s_2, s_3, a) \wedge \neg \text{preserve}(s_1, s_3, a)
\end{aligned} \tag{7}$$

With the inclusion of the zigzag class in T_{soo} we were unable to entail the transitivity of the preserve relation. Review of the inconclusive test results led to the belief that with the inclusion of the zigzag class of activity trees there was in fact, no proof of the transitivity of the preserve relation:

$$\begin{aligned}
& T_{soo} \cup T_{psl} \not\models (\forall s_1, s_2, s_3, a) \text{preserve}(s_1, s_2, a) \\
& \quad \wedge \text{preserve}(s_2, s_3, a) \supset \text{preserve}(s_1, s_3, a)
\end{aligned} \tag{8}$$

Careful consideration of the situation led to the decision that the ability of T_{soo} to entail the transitivity of the preserve relation was more important than developing it as a non-conservative extension of PSL. This decision resulted in a

² A detailed discussion of conservative extensions and the role that they play in ontology development can be found in [4].

change in the axiomatization of T_{soo} , however this change represented a change in the requirements. We were no longer considering the zigzag class, because we had revised the requirements such that T_{soo} did not have to be a conservative extension of PSL. After this change was implemented, we were able to successfully prove that the axioms of the ontology entailed the transitivity of the preserve relation.

2. If we are strongly confident about the correctness of both the requirements and the design of the ontology, then we consider the possibility that it is the intractable nature of the theorem prover that is preventing a proof from being found. In this case, we proceed to the Tuning Phase and attempt to adapt the ontology so that the theorem prover performance is improved to a level where the requirements can be verified.

In another case, we were unsuccessful in proving a particular property about the *min_precedes* and the *preserve* relations. Based on the intended semantics of the two relations, it would appear straightforward that in any model where *min_precedes* holds for two occurrences, *preserve* must hold as well. We attempted to verify this by proving that we could entail the following proposition from the ontology with the theorem prover, however the results were inconclusive:

$$\begin{aligned} (\forall s_1, s_2, a) \text{min_precedes}(s_1, s_2, a) \\ \supset \text{preserve}(s_1, s_2, a) \end{aligned} \quad (9)$$

Being fairly certain about the correctness of the ontology and specification of the property, we moved to the Tuning Phase. In consideration of the definition of the *preserve* relation and the proposition we were attempting to verify, the reflexivity of the *preserve* relation was an intuitively important property. Two lemmas regarding the reflexivity of the *mono* relation that had already been shown to satisfy the models of the ontology were:

$$\begin{aligned} (\forall s_1, s_2, a) \text{min_precedes}(s_1, s_2, a) \\ \supset \text{mono}(s_2, s_2, a) \end{aligned} \quad (10)$$

and

$$\begin{aligned} (\forall s, o, a) \text{subactivity_occurrence}(s, o) \wedge \text{legal}(s) \\ \supset \text{occurrence_of}(o, a) \supset \text{mono}(s, s, a) \end{aligned} \quad (11)$$

The addition of these lemmas to the original reasoning problem aided the theorem prover sufficiently so that the property was proved and we could continue testing the other requirements.

This example illustrates a phenomenon that distinguishes theorem proving with ontologies from more traditional theorem proving – we are not certain that a particular sentence is actually provable. Effectively, every theorem proving task with an ontology is an open problem.

Case 3: All Requirements Met. If we obtain a proof that a requirement is satisfied, and it is consistent with the intended semantics of the ontology, then we may proceed with testing the remaining requirements. Once we obtain such proofs for each requirement, the ontology we have developed satisfies our requirements and we can proceed to the Application Phase.

The results of the Verification Phase may lead to revisions of either the design or the requirements of the ontology. The Design Phase is revisited in the case that an error is detected or suspected in the Verification Phase. If an error in the design is identified, it is noted that the developer must be cautious and consider the entire design when making the correction so that further errors are not created. Additionally, when a correction is made to the design, all previous test results should be reviewed, and any tests (proofs) that were related to the error should be rerun; an error in the design has the potential to positively or negatively impact the results of any tests run prior to its identification and correction. The Requirements Phase is revisited in the case that an error in the requirements is found or suspected during the Verification Phase. It may also be revisited if revisions to the requirements are necessary because of corrections to the ontology that were implemented in the Design Phase.

Note that due to the structure of the entailment problem Rep-2, verification results may require a slightly different interpretation. When attempting to diagnose an error, we should take into account that we are likely reasoning with a well-established axiomatization of some mathematical theory.

Generation of Counter-Examples. Although the theorem prover's results are inconclusive in Case 2, we have the potential to use an automated model building tool to determine whether or not a proof exists. If an ontology's axiomatization has unintended models, then it is possible to find sentences φ that are entailed by the intended structures, but which are not provable from the axioms of the ontology. Therefore we can address the ambiguity of Case 2 by evaluating the satisfiability³, of

$$T_{onto} \cup \neg\varphi$$

where φ is the requirement we are attempting to prove. In other words, if we are unable to find a proof for a particular requirement, we search for a model that is a counter-example of the requirement. The (non-)existence of such a model will tell us the cause of Case 2 with certainty. If a counter-example is found, it can then be examined (and evaluated) to determine the course of corrective action required. In other words if the counter-example is a desired model then the requirements are too strong should be corrected accordingly, otherwise this indicates that the axiomatization is too weak and must be corrected to exclude the counter-example. In the latter case, we may attempt to generate multiple counter-examples to improve our understanding of the error. In this way, model generation may be used not only to resolve the uncertainty of Case 2, but to assist in the decision-making process required if a proof does not exist.

³ We would be evaluating the satisfiability of $T_1 \cup \dots \cup T_n \cup \neg\varphi$ if the requirement in question was from Rep-2.

3.4 Tuning

The Tuning Phase focuses on the mitigation of theorem prover intractability. Similar to the Design Phase, the Tuning Phase develops a set of axioms, however the input and the function of each phase makes the two distinct. In contrast to the Design Phase, the input of the Tuning Phase is a version of the ontology that we believe to be correct, but have not been able to conclusively test. Automated theorem provers sometimes have difficulty finding a proof, though one exists; the aim of this phase is to apply techniques to streamline the ontology in an effort to mitigate theorem prover intractability. This can be accomplished with the development of subsets and the use of lemmas, discussed in further detail below.

Subsets. To develop a subset requires that we remove some of the axioms of the ontology that are not relevant (in the axiomatization of the ontology) to the particular reasoning problem we are considering. The idea is that by excluding these axioms from the reasoning problem, we can increase the efficiency of the theorem prover, as it will not be considering axioms that would not be used for the proof. By reducing the number of clauses that must be considered by the theorem prover, there is the potential to reduce the time required to find a solution, if one exists. The selection of a subset could be, but is not necessarily based on the modules of an ontology. We have currently not explored all aspects of this technique in depth, however it appears that subset selection introduces a new complication that must be considered. When selecting a subset of the ontology, we must ensure that all of the necessary axioms have been included. If any related axiom is excluded, this could allow for an inconclusive test result (we might not prove a sentence because some of the axioms required to prove it have been excluded). At this point, we do not provide any methodology to address this challenge; a clear and complete understanding of the ontology is required to be certain of what axioms must be selected for a subset for a particular reasoning problem.

For example, when working with the PSL ontology, we could exclude axioms from PSL-Core that had to do with timepoints when testing reasoning problems that were related to the composition of activities. This was possible because of our understanding of the concepts of the ontology. The size of the PSL ontology makes the use of subsets necessary for most reasoning problems, however we have observed that these subsets are often successfully reused for other related reasoning problems. While testing the ontology, we were able to use the same subset⁴ to successfully entail 14 of 16 propositions related to the ordering of subactivity occurrences.

Lemmas. We use the term lemma in its traditional, mathematical sense. Their use to improve theorem prover performance is a commonly accepted technique. In the Tuning Phase, we can apply this technique to assist in obtaining conclusive test results. Lemmas may be used (in conjunction or independent of subsets) to improve performance as a means of reducing the number of steps required to obtain a proof. By adding a lemma, we hope to provide the theorem prover with a sentence that will be used in the proof; in this case the number of steps required to obtain the proof is reduced (since the theorem

⁴ This subset can be found at <http://stl.mie.utoronto.ca/colore/process/psl-subset-519.clif>.

prover no longer needs to prove the lemma before using it). We also speculate that the addition of such a lemma provides a sort of guidance for the search, in the direction of the desired proof. Lemmas should be developed intelligently, with some idea of how the addition of the sentence to the ontology will assist in finding a proof. Another point to consider is that of reusability. Some effort should be made to design a lemma that is general enough to be applied for other reasoning problems. In the event that we have already developed a lemma for one reasoning problem, we should consider its potential use in the Tuning Phase for a related reasoning problem.

An example of the use of a lemma while testing the PSL ontology is discussed in Case 2 in the previous section.

4 Discussion

The description of the Verification Phase highlighted the heuristic decisions in the methodology that result from the semidecidability of first-order logic, and the intractability of theorem proving in this case. As discussed, if we do not obtain a proof when testing a requirement, then we cannot always be certain if this is because a proof does not exist (if this is the case, then we know that our current ontology does not satisfy the requirement, unless it was incorrectly specified) or if a proof does exist but the theorem prover reaches a specified time limit before it is able to find it. We suspect that it is this issue of uncertain paths resulting from Case 2 that will stimulate some criticism of our verification methodology and the lifecycle as it is proposed here. We address this concern with the following remarks:

- In two of the three possible cases that we have identified, we can be certain of the direction we must proceed in (the cases when a proof is found).
- In Case 2, when a proof is not found and there is uncertainty about the cause, we can be certain that the requirements for the ontology's application have not been met. Applications of the ontology that utilize a theorem prover must be able to answer a query (competency questions) or entail a proposition (infer a property). In other words a theorem prover must be able to find a proof of the requirements in some reasonable amount of time. Therefore, we can say that in all cases the verification methodology is capable of testing if the requirements are met; the uncertainty exists in how to proceed in development when a requirement is not met.
- The uncertainty of which path should be followed when a requirement is not met may be mitigated. If thorough documentation practices are followed through development, we may seek out trends to indicate the most likely source of the error. We also presented an application of model generation to search for counter-examples; this has the potential to completely resolve any uncertainty for this case.
- Furthermore, we have demonstrated the feasibility and effectiveness of our methodology in practice with examples of each possible case in the Verification Phase.

5 Future Work

One direction of future work that arose from the development of the subactivity occurrence ordering extension of PSL is in the area of development environments.

Specifically, an environment capable of capturing the development history of an ontology - including test results, changes made to axioms and requirements, and different versions of the ontology developed in the Tuning Phase. This would be valuable not only from a documentation perspective, but as a potential aid to heuristic decisions. A summary of test results could indicate trends (design errors, or lemmas required) that would assist the decision process in the case that no proof is found. Also, if a particular test remains inconclusive after considerable effort, an environment capable of tracking test details could allow the developer to temporarily leave the test unresolved and return to it after additional tests have been performed. With the necessary test history information available, the developer could easily and efficiently apply any corrections, subsets, or lemmas discovered in subsequent tests to the “problem” test. An environment dedicated to this type of documentation has the potential to be theorem prover independent. This would be beneficial since different theorem provers may be more useful for different applications or with different ontologies, so committing to one theorem prover restricts the potential application of the methodology presented here.

Additional areas for future work that we have identified are briefly described below:

- Perform experiments to identify the possibility of problem-general heuristics or techniques to reduce uncertainty in the heuristic decision following a Case 2 of the Verification Phase.
- We acknowledge that techniques involving model generation are not possible with all ontologies, i.e. when all models are infinite. An interesting question to investigate in this case is if it is possible to identify a “finite version” of the ontology that is able to preserve the semantics, without forcing infinite models.
- Include maintenance considerations in the lifecycle phases. Different types of maintenance activities (bug fixes, changes in the ontology’s domain) should be performed differently within the lifecycle.
- Investigate the role of non-functional requirements in the ontology lifecycle. In keeping with the analogy typically drawn between software and ontology development, we identify the specification of intended models as the functional requirements of an ontology. This leaves the design and evaluation of non-functional requirements to be explored: how can these qualities be identified and measured? and how can they be integrated in the development lifecycle of ontologies?

6 Conclusions

Existing ontology development methodologies do not provide an account of ontology evaluation that is adequate for verifying ontologies with respect to their model-theoretic properties. In this paper, we have provided an approach to the ontology lifecycle that focuses on support for semiautomatic verification of ontologies, including a methodology that takes into account the pragmatic issues of semi-decidability in first-order logic. The effective use of such a methodology addresses the challenges posed by ontologies that use more expressive languages, such as first-order logic.

Our presentation of the ontology lifecycle rests on the connection between the mathematical definition of the intended models of an ontology and the reasoning problems

that are equivalent to the verification of these intended models with respect to the axiomatization of the ontology. It is this connection which allows theorem provers to play a pivotal role in ontology design, analysis, and evaluation.

Nothing in the methodology presented here is specific to the PSL ontology, or to first-order logic. The lifecycle accounts for the difficulties of development with first-order logic; however, since the semiautomatic verification of requirements satisfaction could be beneficial in any application of ontology development, we close with an invitation for the techniques presented here to be applied with other ontologies.

References

1. Bock, C., Grüninger, M.: PSL: A semantic domain for flow models. *Software and Systems Modeling* 4, 209–231 (2004)
2. Fernández, M., Gómez-Pérez, A., Juristo, N.: Methontology from ontological art towards ontological engineering. In: *Symposium on Ontological Engineering of AAAI* (1997)
3. Gómez-Pérez, A., Corcho, O., Fernández-Lopez, M.: *Ontological Engineering: with examples from the areas of Knowledge Management*. In: *e-Commerce and the Semantic Web*, 1st edn. *Advanced Information and Knowledge Processing*. Springer, Heidelberg (2004), <http://www.worldcat.org/isbn/1852335513>
4. Cuenca Grau, B., Parsia, B., Sirin, E.: *Ontology Integration using ϵ -Connections*. In: *Stuckenschmidt, H., Parent, C., Spaccapietra, S. (eds.) Modular Ontologies*. LNCS, vol. 5445, pp. 293–320. Springer, Heidelberg (2009)
5. Grüninger, M., Fox, M.S.: *Methodology for the design and evaluation of ontologies*. In: *International Joint Conference on Artificial Intelligence (IJCAI 1995), Workshop on Basic Ontological Issues in Knowledge Sharing* (1995)
6. Grüninger, M., Hahmann, T., Hashemi, A., Ong, D.: *Ontology verification with repositories*. In: *Conference on Formal Ontology in Information Systems (FOIS 2010)*, pp. 317–333. IOS Press (2010)
7. Grüninger, M.: *Using the PSL ontology*. In: *Handbook of Ontologies*, pp. 419–431. Springer, Berlin (2009)
8. Gruninger, M., Fox, M.S.: *The role of competency questions in enterprise engineering*. In: *Proceedings of the IFIP WG5.7 Workshop on Benchmarking – Theory and Practice* (1994)
9. Guarino, N., Oberle, D., Staab, S.: *What is an ontology?* In: *Handbook of Ontologies*, pp. 1–17. Springer, Berlin (2009)
10. Pinto, H.S., Tempich, C., Staab, S.: *Ontology engineering and evolution in a distributed world using diligent*. In: *Handbook on Ontologies, International Handbooks on Information Systems*, pp. 153–176. Springer, Heidelberg (2003)
11. Sure, Y., Staab, S., Studer, R., Gmbh, O.: *On-to-knowledge methodology (otkm)*. In: *Handbook on Ontologies, International Handbooks on Information Systems*, pp. 117–132. Springer, Heidelberg (2003)
12. Uschold, M., Grüninger, M.: *Ontologies: Principles, methods and applications*. *Knowledge Engineering Review* 11, 93–136 (1996)
13. Uschold, M., King, M.: *Towards a methodology for building ontologies*. In: *Workshop on Basic Ontological Issues in Knowledge Sharing, held in conjunction with IJCAI 1995* (1995)