Logical Formalizations of Commonsense Reasoning: Papers from the 2015 AAAI Spring Symposium

A New Look at Ontology Correctness

Bahar Aameri
Department of Computer Science
University of Toronto
bahar@cs.toronto.edu

Michael Grüninger
Department of Mechanical and Industrial Engineering
University of Toronto
gruninger@mie.utoronto.ca

Abstract

The design of ontologies for new commonsense domains continues to pose challenges, particularly in cases where multiple potential axiomatizations satisfy the requirements for the ontology. One approach is to specify the requirements with respect to the intended semantics of the terminology; from a mathematical perspective the requirements may be characterized by the class of structures (referred to as the required models) which capture the intended semantics. This approach leads to a natural notion of the correctness as a relationship between the models of the axiomatization of the ontology and the required models for the ontology. In this paper, we consider three possible generalizations of the notion of the correctness of an ontology in the case in which the ontology and the required models have different signatures. We show that these notions of correctness lead to different approaches for ontology evaluation and discuss the benefits and drawbacks of each approach.

Introduction

When developing or selecting an axiomatization of an ontology[1] for an application domain, the knowledge engineer typically has some requirements in mind. These requirements for the ontology are specified with respect to the intended semantics of the terminology; from a mathematical perspective the requirements may be characterized by the class of structures which capture the intended semantics, and such structures can be referred to as the required structures for the ontology. Previous work in this approach (Katsumi and Gruninger 2010; Guarino, Oberle, and Staab 2009) has focused on the case in which the required models and the axioms of the ontology have the same signature. In this case, the correctness of the ontology is defined with respect to the relationship between the required structures for the ontology (which we will denote by $\mathfrak{M}^{req}$) and the models of its axiomatization T_{ont} (see Figure 1). If the ontology is too weak,

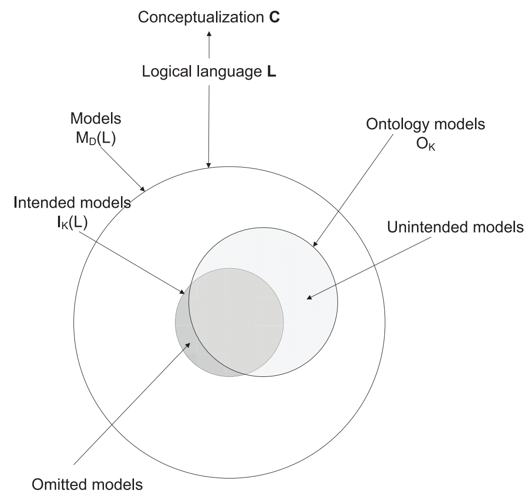

Figure 1: The possible relationships between the models of an ontology and the required structures (taken from (Guarino, Oberle, and Staab 2009). Note that Guarino's notions of intended and unintended models correspond to our notions of required and spurious structures.

then there exist models which are spurious (that is, they are not isomorphic to any required structures):

$$\mathcal{M} \in Mod(T_{ont}) \ \text{ and } \ \mathcal{M} \notin \mathfrak{M}^{req}$$

If the ontology is too strong, then there exist required models which are omitted:

$$\mathcal{M} \in \mathfrak{M}^{req} \ \text{ and } \ \mathcal{M} \notin Mod(T_{ont})$$

In other words, an axiomatization is correct if and only if it does not include any spurious models, *and* it does not omit any required models:

$$\mathcal{M} \in Mod(T_{ont}) \ \text{iff} \ \mathcal{M} \in \mathfrak{M}^{req} \qquad \textbf{(C)}$$

Suppose there are multiple proposed ontologies that satisfy the entire set of semantic requirements (that is, each of the ontologies satisfies the above notion of correctness); how do we select the right ontology? Since the ontologies have the same signature, it is straightforward to evaluate the alternatives. In fact, the ontologies must be logically equivalent,

[1]By an ontology, we mean an axiomatic theory in the language of first-order logic. We, therefore, use the words 'ontology' and 'theory' interchangeably. We consider a theory to be a set of first-order sentences closed under logical entailment, and a subtheory to be a subset of the corresponding theory. For a theory T, $Mod(T)$ denotes the class of all models of T.

otherwise there would exist a model of one which is not a model of the others, so that one ontology would either have an spurious model or an omitted required model. Therefore, since the ontologies are logically equivalent and they all satisfy the requirements, any one of them can be chosen.

However, requirements for an ontology are not always specified in the same language as the ontology's: to facilitate the verification process, requirements are often characterized in the language of existing verified ontologies (e.g. standard mathematical structures like graphs and lattices). Moreover, in some cases multiple ontologies are available for reuse, all of them satisfy the semantic requirements, but they have different signatures. Such cases pose more of a challenge since we cannot simply compare the ontologies by their entailments. We cannot even say that a model of the ontology is isomorphic to a required model, because the notion of isomorphism is only defined for structures with the same signature. Furthermore, the above notion of the correctness of ontologies no longer applies.

In this paper, we consider three possible generalizations of the notion of the correctness of an ontology with respect to its required models to the case in which the ontology and the required models have different signatures. One approach which is widely used is that an ontology is correct if it does not have any spurious models. In the second approach, a correct ontology has no spurious models and it does not omit any required models. With the third approach, an ontology is correct if there is a one-to-one correspondence between the set of required models and the set of models of the ontology.

We will show that these notions of correctness lead to different methodologies for selecting the right ontology from a set of ontologies that satisfy the semantic requirements. Note that when we talk about "*ontology selection*" it does not necessary mean that the alternative ontologies already exist. Rather, the ontology developer can choose to design the ontology in a way that it satisfies one of the notions of correctness. In the first approach, the selected ontology is the strongest one which can interpret the theories that axiomatize the required models. The third approach generalizes the situation with ontologies that have the same signature as the required models – theories which satisfy the requirements must be equivalent.

We begin by a review of the different metalogical relationships among theories with different signatures. We use these relationships to formalize the three notions of correctness for ontologies and then describe the ontology selection approaches that are based on these notions. Finally, we compare the benefits and drawbacks of each selection approach.

Relationships between Theories

The differences between the different notions of ontology correctness are rooted in the way that they exploit the metalogical relationships among ontologies. We therefore begin by reviewing the different metalogical relationships between ontologies that we use throughout the paper and then give a generalized definition for the notion of strength of theories that applies to theories with different signatures.

Among theories with a common signature, a theory T_2 is considered to be stronger than another theory T_1 if T_2 entails all sentences in T_1, i.e., $T_2 \models T_1$.

In that sense, being stronger than a theory is equivalent to being an extension of that theory.

Definition 1 *Let T_1, T_2 be two first-order theories such that $\Sigma(T_1) \subseteq \Sigma(T_2)$.*[2]

T_2 is an extension of T_1 iff for any sentence $\Phi \in \mathcal{L}(T_1)$,

$$T_1 \models \Phi \Rightarrow T_2 \models \Phi.$$

T_2 is a conservative extension of T_1 iff for any sentence $\Phi \in \mathcal{L}(T_1)$,

$$T_1 \models \Phi \Leftrightarrow T_2 \models \Phi.$$

We are also interested in comparing theories regardless of their signature. Therefore, we exploit the notion of relative interpretation (Enderton 1972) which generalizes the notion of extension between theories with distinct signatures.

Definition 2 *An interpretation π of the theory T_1 into a theory T_2 is a function on the set of non-logical symbols of $\Sigma(T_1)$ and formulae in $\mathcal{L}(T_1)$ such that*

1. *π assigns to $\forall$ a formula $\pi_\forall$ of $\Sigma(T_2)$ in which at most the variable v_1 occurs free, such that*
$$T_2 \models (\exists v_1)\, \pi_\forall$$
2. *π assigns to each n-place function symbol f a formula π_f of L_1 in which at most the variables $v_1, ..., v_n, v_{n+1}$ occur free, such that*
$$T_2 \models (\forall v_1, ..., v_n)\, \pi_\forall(v_1) \wedge ... \wedge \pi_\forall(v_n) \supset$$
$$(\exists x)(\pi_\forall(x) \wedge ((\forall v_{n+1})(\pi_f(v_1, ..., v_{n+1}) \equiv (v_{n+1} = x))))$$
3. *π assigns to each n-place relation symbol P a formula π_P of $\Sigma(T_2)$ in which at most n variables occur free.*
4. *for any sentence $\Phi, \Psi \in \Sigma(T_1)$,*
   - *if Φ is an atomic sentence with relation symbol P, $\pi(\Phi) = \pi(P)$;*
   - *$\pi(\neg\Phi) = \neg\pi(\Phi)$;*
   - *$\pi(\Phi \supset \Psi) = \pi(\Phi) \supset \pi(\Psi)$;*
   - *$\pi(\exists x\, \Phi) = \exists x\, \pi_\forall(x) \wedge \pi(\Phi)$;*
   - *$\pi(\forall x\, \Phi) = \forall x\, \pi_\forall(x) \supset \pi(\Phi)$;*
5. *For any sentence $\Phi \in \Sigma(T_1)$,*
$$T_1 \models \Phi \Rightarrow T_2 \models \pi(\Phi).$$

An interpretation π of a theory T_1 into a theory T_2 is faithful iff for any sentence $\Phi \in \Sigma(T_1)$,

$$T_1 \not\models \Phi \Rightarrow T_2 \not\models \pi(\Phi).$$

An interpretation of T_1 in T_2 can be axiomatized by a set of translation definitions from the signature of T_1 into the language of T_2 (Szczerba 1977):

Definition 3 *Let T_1 and T_2 be two theories such that $\Sigma(T_1) \cap \Sigma(T_2) = \varnothing$.*

Translation definitions for T_1 into T_2 are conservative definitions of the form

$$(\forall \overline{x})\, p_i(\overline{x}) \equiv \Phi(\overline{x}),$$

where p_i is a symbol in $\Sigma(T_1)$ and Φ is a formula in $\mathcal{L}(T_2)$.

[2]For a theory T, $\Sigma(T)$ denotes the signature of T, which is the set of non-logical symbols used in sentences of T, and $\mathcal{L}(T)$ denotes the language of T, which is the set of all first-order formulae generated by symbols in $\Sigma(T)$.

According to (Gruninger et al. 2012), if T_2 interprets T_1, there exist translation definitions Δ for T_1 into T_2 such that

$$T_2 \cup \Delta \models T_1.$$

Consequently, we can say that a relative interpretation of a theory entails (with the help of the corresponding translation definitions) all sentences in the theory, and therefore is stronger than it. This is also compatible with the (Visser 2006)'s view of stronger theories.

Definition 4 *A theory T_2 is stronger than another theory T_1 iff T_1 is interpretable in T_2.*

Two theories have the same strength if they are mutually interpretable.

Beside interpretability strength, there are two other properties that can impact ontology selection: decidability and preservation of models. Although mutual interpretability is an equivalence relation, it does not preserve either of these properties. Faithful interpretability preserves decidability, but model preservation requires stronger notions such as mutual faithful interpretability and logical synonymy.

Definition 5 *Theories T_1 and T_2 are definably equivalent iff T_1 faithfully interprets T_2 and T_2 faithfully interprets T_1.*

Definable equivalence preserves properties of models up to elementary equivalence, but many applications require a one-to-one correspondence between the required models of the ontology and the models of the selected axiomatization. (Pearce and Valverde 2012) show that this can be achieved through the notion of logical synonymy.

Definition 6 *Two theories T_1 and T_2 are synonymous iff there exist two sets of translation definitions Δ and Π, respectively from T_1 to T_2 and from T_2 to T_1, such that $T_1 \cup \Pi$ is logically equivalent with $T_2 \cup \Delta$.*

Equivalently, T_1 and T_2 are synonymous iff they have a common definitional extension. It is also easy to see that logical synonymy implies definable equivalence.

Since we are also interested in model-preserving reductions, we borrow the notion of reducibility from (Gruninger et al. 2010).

Definition 7 *A theory T is reducible to a set of theories $T_1, ..., T_n$ iff*

1. *T faithfully interprets each theory T_i;*
2. *T is synonymous with $T_1 \cup ... \cup T_n$.*

It is easy to see that two synonymous theories are reducible to each other.

Mappings between Models

Since the notion of correctness for an ontology is concerned with the relationship between the required models of the ontology and the models of the axiomatization of the ontology, it is important to understand how each of the metalogical relationships (interpretation, faithful interpretation, definable equivalence, and synonymy) induces a mapping on the sets of models of the theories.

We can use results from (Gruninger et al. 2010) and (Pearce and Valverde 2012) to prove the following:

Theorem 1 *Let T_1 and T_2 be two first-order theories.*

1. *If T_2 interprets T_1, then there exists a mapping $\mu : Mod(T_1) \rightarrow Mod(T_2)$.*
2. *If T_2 faithfully interprets T_1, then the mapping $\mu : Mod(T_1) \rightarrow Mod(T_2)$ is surjective.*
3. *If T_2 and T_1 are definably equivalent, then both mappings $\mu_1 : Mod(T_1) \rightarrow Mod(T_2)$ and $\mu_2 : Mod(T_2) \rightarrow Mod(T_1)$ are surjective.*
4. *If T_2 is synonymous with T_1, then $\mu : Mod(T_1) \rightarrow Mod(T_2)$ is bijective.*

In the next section, we will use the different metalogical relationships (on both theories and sets of models) to define different notions of ontology correctness.

Generalizing Ontology Correctness

As we discussed in the introduction, the semantic requirements for a domain are specified through defining a class of structures, which we refer to them as *required structures* or *required models*. We have also seen that the major challenge for ontology selection appears in the case in which multiple candidate ontologies have different signatures than the required models. We therefore need a generalized notion of ontology correctness that allows evaluating axiomatizations specified in different languages.

In this section, we apply the metalogical relationships reviewed in the preceding section, and formally define three notions of correctness that have been (often implicitly) employed as the bases for ontology evaluation and selection.

The first notion considers an ontology to be correct iff none of the models of the ontology correspond to a spurious model. In other words, an ontology is correct if each model of the ontology corresponds to a model in $\mathfrak{M}^{req}$.

Definition 8 *An ontology T_{ont} is interpretably correct with respect to $\mathfrak{M}^{req}$ iff there exists a mapping $\mu : Mod(T_{ont}) \rightarrow \mathfrak{M}^{req}$.*

Note that an ontology is interpretably correct even if it omits a subset of required models. By Theorem 1, ontology T_{ont} is interpretably correct with respect to $\mathfrak{M}^{req}$ if T_{ont} interprets $Th(\mathfrak{M}^{req})$ ($Th(\mathfrak{M}^{req})$ denotes the theory of $\mathfrak{M}^{req}$).

The second notion of correctness states that an ontology T_{ont} is correct iff its models do not correspond to spurious models and it does not omit any required models. In that case, the mapping between $Mod(T_{ont})$ and $\mathfrak{M}^{req}$ is surjective.

Definition 9 *An ontology T_{ont} is faithfully correct with respect to $\mathfrak{M}^{req}$ iff the mapping $\mu : Mod(T_{ont}) \rightarrow \mathfrak{M}^{req}$ is a surjection.*

By Theorem 1, ontology T_{ont} is faithfully correct with respect to $\mathfrak{M}^{req}$ if T_{ont} faithfully interprets $Th(\mathfrak{M}^{req})$.

The third notion of correctness, which is also the strongest one, states that an ontology is correct iff there is a one-to-one correspondence between the models of the ontology and the required models.

Definition 10 *An ontology T_{ont} is verifiably correct with respect to $\mathfrak{M}^{req}$ iff the mapping $\mu : Mod(T_{ont}) \rightarrow \mathfrak{M}^{req}$ is a bijection.*

The notion of verifiable correctness generalizes the correctness condition (**C**) for ontologies that are not in the same language as required models. By Theorem 1, ontology T_{ont} is verifiably correct with respect to $\mathfrak{M}^{req}$ if T_{ont} is synonymous with $Th(\mathfrak{M}^{req})$.

Ontology Selection

Suppose we are given a set of semantic requirements and we want to select (or design) an appropriate axiomatization for these requirements. From the point of view of an ontology designer, a good axiomatization must correctly define concepts that are needed for representing the semantic requirements, and cover *all* required concepts and their relationships, but with making minimal ontological commitment (Vrandecic 2009). If the semantic requirements are specified by a class of required models $\mathfrak{M}^{req}$, these criteria can be achieved by specifying the axiomatization of $Th(\mathfrak{M}^{req})$.

In practice, however, the language of potential ontologies and the required models are distinct, and so the selection process is different. In particular, the required 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 models are then specified with respect to the properties of these mathematical structures. For instance, relationships between timepoints are specified with respect to linear orderings, or parthood relations between components of a spatial domain are specified by Boolean lattices. Potential ontologies are theories that capture the corresponding mathematical structures, e.g., two potential ontologies for timepoints are the theory of linear orderings (with only one symbol $\leq$ in its signature) and the theory of real numbers (in which the ordering relation $\leq$ is definable by addition and multiplication).

In this section, we review ontology selection approaches that are common within the knowledge representation and ontology communities. The approaches are based on the three notions of ontology correctness we described in the preceding section. We argue that only ontologies that are definably equivalent or synonymous with $Th(\mathfrak{M}^{req})$ satisfy all the above-mentioned criteria, while ontologies that are stronger than $Th(\mathfrak{M}^{req})$ violate at least one criterion.

The Synonymy Approach

The Synonymy Approach adopts the notion of verifiable correctness as its basis for selecting (or designing) ontologies. Accordingly, the approach only selects ontologies that are synonymous with $Th(\mathfrak{M}^{req})$, so that there exists a bijection between the models in $\mathfrak{M}^{req}$ and the models of the selected ontology. Note that in this case the selected ontology is as strong as $Th(\mathfrak{M}^{req})$.

Suppose T_{ont} is selected by the Synonymy Approach for axiomatizing the class of required models $\mathfrak{M}^{req}$. Since there is a one-to-one correspondence between $Mod(T_{ont})$ and $\mathfrak{M}^{req}$, the axiomatization T_{onto} captures semantic requirements correctly and completely. Moreover, since T_{ont} is the precise axiomatization of the requirements, it cannot be weakened, and hence makes the minimal possible set of ontological commitments.

In many domains, the semantic requirements are specified by a combination of multiple mathematical structures and axiomatized by the union of the respective mathematical theories. Suppose $\mathfrak{M}^{req}$ is axiomatized by the union of some mathematical theories $T_1, ..., T_n$, $n \geq 1$. According to Theorem 1 and the definition of reducible theories, $T_1 \cup ... \cup T_n$ is verifiably correct with respect to $\mathfrak{M}^{req}$ (and so is selected by the Synonymy Approach) iff $Th(\mathfrak{M}^{req})$ is reducible to $T_1, ..., T_n$.

To illustrate the Synonymy Approach more clearly, we will apply it to a simple example. Suppose we want to select an axiomatization for qualities that are unique for each object, i.e., each object in a domain is assigned to exactly one instance of such qualities. This is similar to the DOLCE direct quality (Masolo et al. 2003), and can be used to represent properties like spatial and temporal positions which are unique for a specific object.

The class of required models, denoted by $\mathfrak{M}^{quality}$, can be defined as the following: $\mathcal{M} \in \mathfrak{M}^{quality}$ iff

- $\mathcal{M}$ includes two disjoint sets $\mathbf{obj}, \mathbf{qt} \subset M$, where M is the domain of $\mathcal{M}$;
- $\mathcal{M}$ includes a relation $\mathbf{assoc} : \mathbf{obj} \times \mathbf{qt}$ which maps each member of $\mathbf{obj}$ to exactly one member of $\mathbf{qt}$.

On the other hand, suppose that the nominated ontology is the theory T_{tpar_inc} which axiomatizes a class of mathematical structures, $\mathfrak{M}^{tpar\_inc}$, known as bipartite incident structures (Buekenhout 1995). A bipartite incidence structure is a tuple $\mathcal{I} = (\mathbf{point}, \mathbf{line}, \mathbf{in})$, where $\mathbf{point}, \mathbf{line}$ are sets with $\mathbf{point} \cap \mathbf{line} = \varnothing$, and

$$\mathbf{in} \subseteq ((\mathbf{point} \times \mathbf{line}) \cup (\mathbf{line} \times \mathbf{point})).$$

It is easy to see that there are models in $\mathfrak{M}^{tpar\_inc}$ which have no equivalent in $\mathfrak{M}^{quality}$: a point $\mathbf{p}$ in a model $\mathcal{N} \in \mathfrak{M}^{tpar\_inc}$ can be incident with more than one line. If the sets $\mathbf{point}, \mathbf{line}, \mathbf{in}$ in $\mathcal{N}$ are respectively mapped to the sets $\mathbf{obj}, \mathbf{qt}, \mathbf{assoc}$ in a model $\mathcal{M}$, then $\mathcal{M}$ will not be a model in $\mathfrak{M}^{quality}$. Consequently, T_{tpar_inc} is not strong enough to axiomatize $\mathfrak{M}^{quality}$.

Instead of T_{tpar_inc}, we can use an extension of it $T_{par\_ln\_inc}$ which axiomatizes the class $\mathfrak{M}^{par\_ln\_inc}$ of parallel-lines incident structures. A parallel-lines incident structure is a bipartite incidence structure

$$\mathcal{G} = (\mathbf{point}, \mathbf{line}, \mathbf{in^G})$$

such that all elements in $\mathbf{point}$ are incident with exactly one element in $\mathbf{line}$.

Using the following translation definitions, we can show that there is a one-to-one correspondence between models in $\mathfrak{M}^{par\_ln\_inc}$ and $\mathfrak{M}^{quality}$:

$$\begin{aligned}
\Delta : (\forall x, y)\, in(x, y) &\equiv assoc(x, y) \vee assoc(y, x) \vee x = y,\\
(\forall x)\, point(x) &\equiv obj(x),\\
(\forall x)\, line(x) &\equiv qt(x).\\
\Pi : (\forall x, y)\, assoc(x, y) &\equiv in(x, y) \wedge point(x) \wedge line(y),\\
(\forall x)\, obj(x) &\equiv point(x),\\
(\forall x)\, qt(x) &\equiv line(x).
\end{aligned}$$

Note that theories that are stronger than $T_{par\_ln\_inc}$ omit a subset of $\mathfrak{M}^{quality}$, and consequently are not selected by the Synonymy Approach.

The Faithful Theory Approach

The Faithful Theory Approach is based on the notion of faithful correctness, and so selects ontologies that faithfully interpret $Th(\mathfrak{M}^{req})$. Although the ontologies that are selected by this approach capture all semantic requirements correctly, they do not always make minimal ontological commitments; the only case that they do is when the selected ontology is definably equivalent with $Th(\mathfrak{M}^{req})$.

Notice that in some application domains, the actual models of the ontology, and not just the sentences that they entail, are of central importance. For example, in computer vision, an interpretation of an image is a model of the ontology, and in scheduling, a schedule is a model of the resource and temporal constraints that are expressed using the ontology. For such domains, definable equivalence is insufficient as it only preserves entailment and satisfiability and does not induce a bijection between the models of the corresponding theories.

The Interpreted Theory Approach

The Interpreted Theory Approach implicitly applies the notion of interpretable correctness: the assumption in this approach is that interpretability (not necessarily faithful) in one direction is sufficient, and so selects any ontology that is stronger than $Th(\mathfrak{M}^{req})$. Examples of applying this approach can be found in (Kautz and Ladkin 1991; Knight and Ma 1994) which use real numbers and related theories to represent and reason about time and duration, or (Borgo and Masolo 2010) which apply real coordinate spaces to compare different mereological and mereotopological theories.

Although the axiomatizations selected by this approach are interpretably correct, they are too strong that omit a subset of required models. Moreover, in some cases they pledge to ontological commitments that are not necessary. We investigate the drawbacks of these issues further in the next section.

Why not Use a Stronger Ontology?

A theory T can be extended to a stronger theory T', with a common signature, in three ways: (i) adding sentences in $\mathcal{L}(T)$ that are not in T (ii) adding sentences that use symbols that are not in $\Sigma(T)$ (iii) or both (i) and (ii). If an extension T' is achieved by (i), then it is a non-conservative extension of T, and every model of T' is also a model of T, i.e.,

$$Mod(T') \subset Mod(T).$$

Extensions achieved by (ii) and (iii) are more expressive than T as their languages are expanded, and consequently more relations and concepts are definable. Note that theories attained by (ii) are conservative extensions of T while extensions by (iii) are non-conservative.

This argument can be extended to relative interpretation; in that sense a theory T' that is stronger than T satisfies one of the following cases:

1. T' (non-faithfully) interprets T, and all subtheories of T' that are stronger than T have the same signature as T'.
2. T' (non-faithfully) interprets T, and has a subtheory T'' such that $\Sigma(T'') \subset \Sigma(T')$ and T'' is stronger than T.
3. T' faithfully interprets T.[3]

When employing an ontology that is stronger than the requirements, i.e. stronger than $Th(\mathfrak{M}^{req})$, the selected axiomatization belongs to one of the above categories. In this section, we apply the properties of these categories to discuss why stronger theories are not appropriate for ontology development and selection.

Relations Without Intended Semantics

Consider the case where a theory T that is stronger than $Th(\mathfrak{M}^{req})$ is used to axiomatize the class $\mathfrak{M}^{req}$. Suppose T has a subtheory T' that interprets $Th(\mathfrak{M}^{req})$ and $\Sigma(T') \subset \Sigma(T)$. This means that there are symbols in $\Sigma(T)$ that are not needed in translations of $Th(\mathfrak{M}^{req})$, and consequently have no meaningful semantics in models in $\mathfrak{M}^{req}$.

We might also encounter a similar problem when the description of a function or relation symbol in T is so strong that cannot be assigned to any concept in the underlying structures. Suppose, for example, that we need to select an ontology for reasoning about time durations. Many approaches, like (Rescher and Urquhart 1971; Kautz and Ladkin 1991; Knight and Ma 1994; Navarrete et al. 2002), suggest using the ontology of a particular algebraic field such as real numbers. However, it does not make sense to multiply two time durations. Even if we assign an artificial semantics to multiplication, the product of two time durations cannot be a time duration, and so the underlying required structures do not form a field.

Elimination of Required Models

Beside the problem that we just described, using non-faithful interpretations has another disadvantage: they are too strong because they eliminate some of the required models.

Consider again the time duration example. Let $\mathfrak{M}^{duration}$ denotes the class of required structures for time durations, and T_{reals} denote the theory of real numbers. (Gruninger 2010) showed that $Th(\mathfrak{M}^{duration})$ is interpretable in T_{reals}. So, there exists a set of translation definitions Δ such that

$$T_{reals} \cup \Delta \models T,$$

where $Th(\mathfrak{M}^{duration}) \subset T$. Since the interpretation is not faithful we have $\Sigma(Th(\mathfrak{M}^{duration})) = \Sigma(T)$, and so

$$Mod(T) \subset \mathfrak{M}^{duration}.$$

Thus, T_{reals} does not axiomatize *all* semantic requirements for time durations as it omits a subset of required models.

A similar argument can be made to show that $\mathbb{R}^n$ (real coordinate space of size n) cannot capture all semantic requirements of mereotopologies.

[3]Note that T' might have a subtheory T'', with $\Sigma(T'') \subset \Sigma(T')$, which faithfully interprets T. However, since this has no impact on our argument we do not distinguish the two cases.

Someone might argue that by eliminating sentences in $T_{reals} \cup \Delta$ that entail sentence in $T \setminus Th(\mathfrak{M}^{duration})$ we can get a subtheory of T_{reals} that faithfully interprets $Th(\mathfrak{M}^{duration})$. While having a one-direction faithful interpretation has its own problems (which will be discussed in the next section), in some cases it is not possible to reduce a strong theory to a faithful interpretation without losing fundamental properties of the theory. For instance, the theory of linear orderings is interpretable by T_{reals} using the following translation definition

$$(\forall x, y)\ (x \leq y) \equiv ((\exists z)\ (y = x + (z \times z))).$$

Since T_{reals} cannot interpret *discrete* linear orderings the interpretation is not faithful. To get a faithful interpretation we need to eliminate the sentences

$$(\forall x)\ (x \neq 0) \supset (\exists y)\ (x \times y = 1),$$
$$(\forall x)(\exists y)\ \ (x = y^2) \vee (-x = y^2),$$

and their consequences, which makes the models of the new theory something other than closed fields.

Restricting Shareability

Unlike non-faithful interpretations, stronger theories that faithfully interpret $Th(\mathfrak{M}^{req})$ can capture all semantic requirements; yet they are not as desirable as definably equivalent or synonymous theories. The reason is that one-direction faithful interpretations compel additional ontological commitments which are not necessary, and so impose restriction on shareability.

For instance, the theory of dense linear orderings is faithfully interpretable by T_{reals}, but it would be unusual to use T_{reals} to axiomatize such ordering relation. The reason is that some of the restrictions specified for addition and multiplication are not necessary for ordering relation.

As another example, suppose we want to select an ontology for representing *betweenness* (e.g., we want to extract which cities we must pass to go from point A to point B, or which streets are between two given streets). One suggestion might be to use Hilbert's axioms for geometry (Hilbert 1971); however, the full geometry is too strong for reasoning only about betweenness: $T_{hilbert}$ comprises three subtheories, T_{incid}, that axiomatizes properties of the incidence relation, $T_{between}$, which is a theory of betweenness relations, and T_{cong} that describes congruence relationships; for representing betweenness, $T_{between}$ is sufficient.

Despite all of the arguments regarding the benefits of making minimal ontological commitments to support shareability, one might still prefer one-direction faithful interpretations over definable equivalence and synonymy, since in many cases faithfully interpreting theories provide more definable relations, and so are more expressive. Our argument is that this additional expressiveness is not required in the application domain because the necessary degree of expressivity is already captured through specifying required models. As such, the appropriate axiomatization should be as expressive as is needed for representing the specified requirements.

Conclusion

There is a common consensus within the knowledge representation community that ontologies must be evaluated with respect to the intended semantic of their terminologies. Still, multiple non-equivalent axiomatizations have been proposed and used for the same sets of semantic requirements. One reason is that ontology designers usually adopt different notions of ontology correctness without explicitly stating them. In this paper, we have formalized three notions of correctness for ontologies and showed how these different notions impact ontology evaluation. In particular, we discussed three design approaches: the first one selects any ontology that is interpretably stronger than the theory $Th(\mathfrak{M}^{req})$ that axiomatizes the requirements. The second approach selects ontologies that faithfully interpret $Th(\mathfrak{M}^{req})$, while the third one only chooses ontologies that are logically synonymous with $Th(\mathfrak{M}^{req})$. We argued that the first two approaches are not desirable in many applications because they may select ontologies that omit some required models or make unnecessary ontological commitments that restrict shareability.

The viewpoint in favor of stronger ontologies (i.e., ontologies that faithfully or non-faithfully interpret $Th(\mathfrak{M}^{req})$) argues that reuse is easier with such ontologies as a small, already-existing collection of strong theories, like full geometry and ordered real closed fields, can axiomatize a wide range of domains– whenever a weaker representation is needed subtheories of the stronger theories can be used. While this idea is effective for applications which are only concerned about preserving entailment and satisfiability, it cannot be applied when the actual models have to be preserved (recall that definable equivalence preserves entailment and satisfiability while logical synonymy and stronger metalogical relationships preserve models). The reason is that any ontology that is weaker than a theory T is definably equivalent with a subtheory of T. Consequently, there are always subtheories in the strong theories collection that are definably equivalent with $Th(\mathfrak{M}^{req})$. However, in order to ensure that the collection always includes a theory synonymous with $Th(\mathfrak{M}^{req})$, the following question must be answered:

Open question: Suppose a theory T_1 interprets another theory T_2. Does T_1 always include a subtheory that is synonymous with T_2?

References

Borgo, S., and Masolo, C. 2010. Full mereogeometries. *The review of symbolic logic* 3(04):521–567.

Buekenhout, F. 1995. *Handbook of incidence geometry: buildings and foundations*. North Holland.

Enderton, H. 1972. *Mathematical Introduction to Logic*. Academic Press.

Gruninger, M.; Hahmann, T.; Hashemi, A.; and Ong, D. 2010. Ontology verification with repositories. In *Proc. of the 6th Int. Conference on Formal Ontologies in Information Systems (FOIS2010)*, 317–330. IOS Press.

Gruninger, M.; Hahmann, T.; Hashemi, A.; Ong, D.; and

Ozgovde, A. 2012. Modular first-order ontologies via repositories. *Applied Ontology* 7(2):169–209.

Gruninger, M. 2010. Ontologies for dates and duration. In *Twelfth International Conference on Principles of Knowledge Representation and Reasoning (KR2010)*. AAAI Publications.

Guarino, N.; Oberle, D.; and Staab, S. 2009. What is an Ontology? In Staab, S., and Studer, R., eds., *Handbook on Ontologies*, 1–17. Berlin: Springer-Verlag.

Hilbert, D. 1971. *Foundations of Geometry*. Open Court Publishing Company.

Katsumi, M., and Gruninger, M. 2010. Theorem Proving in the Ontology Lifecycle. In *Knowledge Engineering and Ontology Design (KEOD 2010)*, 74–86.

Kautz, H. A., and Ladkin, P. B. 1991. Integrating metric and qualitative temporal reasoning. In *Proc. of AAAI91*, 241–246.

Knight, B., and Ma, J. 1994. A temporal database model supporting relative and absolute time. *The Computer Journal* 37(7):588–597.

Masolo, C.; Borgo, S.; Gangemi, A.; Guarino, N.; and Oltramari, A. 2003. Wonderweb deliverable d18–ontology library (final). *Technical Report, Laboratory for Applied Ontology, ISTC-CNR, Trento, Italy*.

Navarrete, I.; Sattar, A.; Wetprasit, R.; and Marin, R. 2002. On point-duration networks for temporal reasoning. *Artificial Intelligence* 140(1):39–70.

Pearce, D., and Valverde, A. 2012. Synonymous theories and knowledge representations in answer set programming. *Journal of Computer and System Sciences* 78:86–104.

Rescher, N., and Urquhart, A. 1971. Temporal logic. *Springer–Verlag, Heidelberg* 42.

Szczerba, L. 1977. Interpretability of Elementary Theories. In *Logic, Foundations of Mathematics and Computability Theory*, 129–145.

Visser, A. 2006. Categories of theories and interpretations. In *Logic in Tehran. Lecture Notes in Logic, vol. 26*, 284–341. Association for Symbolic Logic.

Vrandecic, D. 2009. Ontology evaluation. In *Handbook on Ontologies*. Springer. 293–313.