

Logical Relevance in Ontologies

Chiara Del Vescovo, Bijan Parsia, and Uli Sattler

University of Manchester, UK

Abstract Most ontology development environments (ODEs) are term oriented and take a frame-based view of the information in an ontology about a given term. Even tools, such as Protégé 4, designed for axiom oriented development preserve the frame-based view as the central mode of interaction with the ontology. The frame-based approach has a number of advantages—most prominently that it is comfortable to people familiar with object oriented programming languages. However, in expressive languages the frame-based views suffer from being only sensitive to syntactic relations between axioms and terms, thus possibly missing key logical relations.

In this paper, we first introduce a semantic notion of relevance between a term and axioms in an ontology, and we investigate the relation of this concept with the inseparability relation based on model Conservative Extensions. Unfortunately, we cannot use model conservativity to detect relevance since it is hard, or even impossible, to decide. Hence, we approximate model conservativity using two notions of modules based on locality, that can be efficiently computed, and provide logical guarantees, e.g. they preserve entailments over a given signature. In particular, we define relevance via Atomic Decomposition, that is a dependency graph showing the logical relations enforced by the two notions of modules between the axioms. We define a suitable labelling that allows us to locate axioms that are relevant for a term in the AD dependency structure. Finally, we describe an interesting consequence of such a view in terms of the models of an ontology.

1 Introduction

Most ontology development environments (ODEs) are term oriented and take a frame-based view of the information in an ontology about a given term¹. Even tools, such as Protégé 4, designed for axiom oriented syntaxes (such as the functional syntax of OWL 2) preserve the frame-based view as the central mode of interaction with the ontology. The frame-based approach has a number of advantages—most prominently that it is comfortable to people familiar with object oriented programming languages.

However, in expressive languages the frame-based view is prone to present a misleading view of what is relevant for a given term in an ontology: in particular, axioms which are relevant for the meaning of the term are excluded from such a view, and some extraneous ones may creep in. The main reason for this problem to occur is that frame-based views are generally only sensitive to syntactic relations between an axiom and a term, and thus they can miss key logical relations.

¹ We use *term* to mean any individual, concept, or role name belonging to the signature of the ontology.

Given an axiom α , it is easy to check whether it is logically relevant for a term \mathfrak{t} : if it constrains the meaning of \mathfrak{t} . As an example, let us consider the axiom $\alpha = 'A \sqsubseteq B \sqcap (C \sqcup \neg C)'$. Then, α is clearly relevant for both A and B, because it states a subsumption relation between the two concept names. However, α “does not say anything” about C, since for any interpretation $\mathcal{C}^{\mathcal{I}}$, the expression $C \sqcup \neg C$ is equivalent to \top , and can be discarded from the axiom obtaining an axiom $\alpha' = 'A \sqsubseteq B'$ logically equivalent to α . Hence, axioms irrelevant for \mathfrak{t} can easily sneak into the usage view. Tautologies as $\mathfrak{t} \sqsubseteq \top$, which are automatically generated when a new top class name is entered in an OWL ontology using Protégé 4, are a quite common example.

Another logical relation that we want to preserve concerns the consequences that *sets* of axioms can impose on a term. An issue in the detection of what contributes to the meaning of a term is the fact that a given term \mathfrak{t} does not even need to occur in an axiom’s signature for being constrained by it. As an example let us consider the set of axioms $\{\alpha_i = 'A_{i-1} \sqsubseteq A_i'\}_{i=1..n}$. Then, it is easy to see that, for $n \geq 2$ and $i = 2, \dots, n - 2$, both axioms α_1 and α_n are logically relevant for A_i even though their signatures do not contain it. Note that a complex logical interaction can occur also within ontologies with limited expressivity. However, logical relevance is clearly more interesting for more complex ontologies than taxonomies.

This paper is a preliminary investigation of the notion of relevance of axioms for a term under a model-theoretic perspective. The major aim of our future work consists of identifying an efficiently computable way to reveal the logical interactions between axioms and terms. The main applications of this study can be found in the areas for improving reasoners performance, and in supporting ontology engineers during the modelling process.

2 Preliminaries

We assume the reader to be familiar with Description Logics [1]. As usual in this context, we use \mathcal{O} for ontologies, i.e. finite sets of axioms based on a Description Logic, e.g., *SHIQ*, and $(\Delta^{\mathcal{I}}, \mathcal{I})$ for interpretations of \mathcal{O} over the domain $\Delta^{\mathcal{I}}$. We use $\tilde{\alpha}$ for the signature of an axiom α , i.e., the set of concept, role, and individual names used in α . A generic term \mathfrak{t} is any non logical atomic symbol of the signature $\tilde{\mathcal{O}}$ of the ontology. Given a signature $\Sigma \subseteq \tilde{\mathcal{O}}$, we denote by $\mathcal{I}|_{\Sigma}$ the restriction of the interpretation function \mathcal{I} over the symbols in Σ .

In this section we briefly summarize the key concepts used in the paper, as model conservativity [8], locality-based modules [3], Atomic Decompositions (ADs) [6] and their labelled versions (LADs) [4], plus some notions inherited by the algebraic order theory.

Model inseparability Two ontologies $\mathcal{O}_1, \mathcal{O}_2$ are *model-inseparable* w.r.t. a signature Σ —denoted $\mathcal{O}_1 \equiv_{\Sigma}^{mCE} \mathcal{O}_2$ —if $\{\mathcal{I}|_{\Sigma} \mid \mathcal{I} \models \mathcal{O}_1\} = \{\mathcal{J}|_{\Sigma} \mid \mathcal{J} \models \mathcal{O}_2\}$. We can then define an *mCE-module* w.r.t. a signature Σ to be a minimal set of axioms $\mathcal{M} \subseteq \mathcal{O}$ such that, for each model \mathcal{I} of \mathcal{M} , there is a model \mathcal{J} of \mathcal{O} such that $\mathcal{J}|_{\Sigma} = \mathcal{I}|_{\Sigma}$. Another notion we use is the *\mathfrak{t} -variant* of an interpretation \mathcal{I} , defined as an interpretation \mathcal{J} such that, for each symbol $s \in \tilde{\mathcal{O}} \setminus \mathfrak{t}$, we have $s^{\mathcal{I}} = s^{\mathcal{J}}$.

Locality-based modules Unfortunately, deciding if a set of axioms is an mCE module is hard or even impossible for expressive DLs [8,12]. Thus, efficiently computable approximations have been devised, as those defined via the notion of *syntactic locality*. A locality-based module \mathcal{M} for a signature Σ is an approximation of the mCE-module for Σ , in the sense that it is a (not minimal, but generally small) set of axioms that preserves all models over Σ . Interestingly, \mathcal{M} also preserves all entailments over Σ , even though possibly not only those. Locality-based modules are particularly interesting because the extraction of a module can be performed in polynomial time. We give an intuition of the definition in what follows, and refer the interested reader to [3] for a deeper discussion.

Intuitively, an axiom α is (syntactically) local w.r.t. a signature Σ if there is no axiom over Σ that is entailed by α . Locality is anti-monotonic, that is, if an axiom is non local w.r.t. Σ , then it is non local also w.r.t. any Σ' that contains Σ . So we can define a *minimal seed signature* for an axiom α to be a signature Σ such that α is non local w.r.t. Σ but it is local w.r.t. any proper subset of Σ .

Locality comes in two flavours: \perp and \top . Intuitively, an axiom is \perp -local w.r.t. a term when it fails to constrain it “from above”. As an example, let us consider $\alpha = \text{‘}A \sqsubseteq B\text{’}$; then, α is local w.r.t. B because, for any interpretation \mathcal{I} over $\{A\}$, \mathcal{I} can be extended by interpreting B as $\Delta^{\mathcal{I}}$ and still $\mathcal{I} \models \alpha$. Similarly, α is \top -local w.r.t. A because it fails to constrain A “from below”.

Locality-based modules then inherit a similar intuition: roughly speaking, a \perp -module for Σ (denoted $\perp\text{-mod}(\Sigma, \mathcal{O})$), when non empty, gives a view “from above” because it contains all subconcepts of concept names in Σ ; a \top -module for Σ (denoted $\top\text{-mod}(\Sigma, \mathcal{O})$) gives a view “from below” since it contains all superconcepts of concept names in Σ . Please note that \mathcal{M} is not simply the union of all non-local axioms w.r.t. Σ . The extraction algorithm is described in [3], and a module extractor based on syntactic locality is available in the OWL API.²

(Labelled) Atomic Decomposition The number of modules of an ontology \mathcal{O} can be exponential in the minimum amongst the number of axioms of \mathcal{O} and the size of its signature [16]. However we can focus on *genuine* modules, i.e. modules that are not the union of two “ \subseteq ”-uncomparable modules. Such modules define a base for all modules, and interestingly the size of the family of genuine modules for \mathcal{O} is linearly dependent on its size [6].

Some sets of axioms never split across two modules [6], revealing a strong logical interrelation. The notion of *Atomic Decomposition* provided next is central to our paper.

Definition 1. For $x \in \{\top, \perp\}$, we call x -atom a maximal set $\alpha^x \subseteq \mathcal{O}$ such that, for each x -module \mathcal{M}^x , either $\alpha^x \subseteq \mathcal{M}^x$, or $\alpha^x \cap \mathcal{M}^x = \emptyset$. The family of x -atoms of \mathcal{O} is denoted by $\mathcal{A}(\mathcal{O})$ and is called x -Atomic Decomposition (x -AD).

If the module notion x is clear from the context, or irrelevant, we drop it.

Since every atom is a set of axioms, and atoms are pairwise disjoint, the AD is a partition of the ontology, and its size is at most linear w.r.t. the size of the ontology. In particular, each axiom³ α belongs to one and only one atom, denoted α_α .


² <http://owlapi.sourceforge.net>

³ Syntactic tautologies do not occur in any atom; however, since they do not impose any constraint on the terms of an ontology \mathcal{O} because they are always true, we can safely remove them from \mathcal{O} and only consider the case where \mathcal{O} does not contain any such axioms.

Interestingly, there is a 1-1 correspondence between atoms and genuine modules: for each atom a we denote by \mathcal{M}_a the corresponding genuine module, that is also the smallest module containing a . Then we can define a second logical relation between atoms: an atom a is *dependent* on a distinct atom b (written $a \succ b$) if $\mathcal{M}_b \subseteq \mathcal{M}_a$. Note that this property then holds for all modules containing a . The dependence relation \succ on AD is a poset (i.e., transitive, reflexive, and antisymmetric), thus can be represented by means of a Hasse diagram. Moreover, it is computable in polynomial time [6]. To easy the understanding of what an AD of an ontology is, the Example 1 illustrates a small ontology and its \perp -AD.

Example 1. Consider the following toy ontology and its \perp -AD:

$\mathcal{O} = \{(\alpha_1) \text{Animal} \sqsubseteq \exists \text{hasGender.Thing},$
 $(\alpha_2) \text{Animal} \sqsubseteq \geq \text{lhasHabitat.Thing},$
 $(\alpha_3) \text{Person} \sqsubseteq \text{Animal},$
 $(\alpha_4) \text{Vegan} \equiv \text{Person} \sqcap \forall \text{eats.}(\text{Vegetable} \sqcup \text{Mushroom}),$
 $(\alpha_5) \text{Student} \sqsubseteq \text{Person} \sqcap \exists \text{hasHabitat.University},$
 $(\alpha_6) \text{GraduateStudent} \equiv \text{Student} \sqcap \exists \text{hasDegree.}(\{\text{BA}, \text{BS}\}),$
 $(\alpha_7) \text{Car} \sqsubseteq \text{Vehicle},$
 $(\alpha_8) \text{Truck} \sqsubseteq \text{Vehicle},$
 $(\alpha_9) \text{Car} \sqsubseteq \neg \text{Truck}\}$



Here the \perp -atoms in the AD contain the following axioms respectively:

$\mathbf{a}_1 = \{\alpha_1, \alpha_2\}$, $\mathbf{a}_2 = \{\alpha_3\}$, $\mathbf{a}_3 = \{\alpha_4\}$, $\mathbf{a}_4 = \{\alpha_5\}$, $\mathbf{a}_5 = \{\alpha_6\}$, $\mathbf{a}_6 = \{\alpha_7\}$, $\mathbf{a}_7 = \{\alpha_8\}$, and $\mathbf{a}_8 = \{\alpha_9\}$.

Atoms can be seen as building blocks for modules: for each x -module \mathcal{M} of an ontology \mathcal{O} , there are atoms $\mathbf{a}_1, \dots, \mathbf{a}_\kappa$ in $\mathcal{A}(\mathcal{O})$ such that $\mathcal{M} = \bigcup_{i=1}^{\kappa} \mathbf{a}_i$. The converse does not hold, since not all combinations of atoms are modules. However, in [5] we studied and implemented an algorithm to extract modules of ontologies directly from their ADs, that is without loading the ontology. We use an enriched version of the ADs, called *Labelled Atomic Decomposition* (LAD), where each atom \mathbf{a} is mapped to the set minimal sets Σ of terms that make \mathbf{a} be included in the module for Σ . Depending on the task we want to use LADs for, different labels can be defined. A first investigation of tasks and suitable labels can be found in [4].

Order theory The poset structure induced over the atoms of an ontology allows us to take advantage of some useful algebraic notions. Given an atom \mathbf{a} , we define its *principal ideal* $\downarrow \mathbf{a}$ to be the set union of \mathbf{a} with all atoms \mathbf{b} such that $\mathbf{a} \succ \mathbf{b}$.⁴ Similarly, we can define the *principal filter* $\uparrow \mathbf{a}$ of \mathbf{a} as the set union of all atoms \mathbf{c} such that $\mathbf{c} \succ \mathbf{a}$. More in general, given a set \mathcal{S} of atoms $\{\mathbf{a}_1, \dots, \mathbf{a}_\ell\}$ we can define its ideal (filter) to be the union of the principal ideals (filters) of the atoms in \mathcal{S} . The usefulness of these algebraic notions for ADs is proven by the ease of getting the genuine module of an atom $\mathbf{a} \sqsubseteq \mathcal{O}$ from the AD of \mathcal{O} : we can just extract the principal ideal $\downarrow \mathbf{a}$.

⁴ Slightly abusing the notation, we define ideals as the union over poset elements rather than sets of poset elements. This choice allows ideals to be set of axioms, hence ontologies.

3 Semantic-based Relevance

While the semantic of the terminology of an ontology defines the objects that the ontology deals with, it does not say how these objects are related. The relationship between terms is defined by the *axioms* of the ontology, that constrain which interpretations are allowed, and which are not. For this reason, we are interested in looking for a semantic notion of relevance of an axiom for a term. In particular, the interpretations of two distinct terms can be conflicting only if some axioms are then violated. In this perspective, the natural choice is to investigate the notion of relevance of an axiom for a term, rather than relevance between terms. First, we introduce the following useful notions.

Definition 2. *Given a consistent ontology \mathcal{O} and a signature $\Sigma \subseteq \tilde{\mathcal{O}}$, we define a Σ -model w.r.t. \mathcal{O} to be an interpretation \mathcal{I} over Σ such that, there exists a model \mathcal{J} for \mathcal{O} such that $\mathcal{J}|_{\Sigma} = \mathcal{I}$. In this case, we say that \mathcal{I} is extendable to a model \mathcal{J} for \mathcal{O} , and any such \mathcal{J} is called an \mathcal{O} -extension of \mathcal{I} .*

If \mathcal{O} is clear from the context, we simply drop it and say Σ -model. Please also note that Def. 2 is also valid in the case of \mathcal{O} being a single axiom.

A first very basic notion of relevance is introduced in the following example: let α be the axiom $A \sqsubseteq B \sqcap (C \sqcup \neg C)$. Then, in order for an interpretation \mathcal{I} to be a model of α , it needs to satisfy the relation $A^{\mathcal{I}} \subseteq B^{\mathcal{I}}$. In other words, the acceptable interpretations of both A and of B are constrained. On the contrary, the interpretation of C is not constrained by α , and we can even rewrite the axiom into the equivalent $A \sqsubseteq B$ that does not even mention C. This case of relevance provided by a single axiom is described in the next definition.

Definition 3. *An axiom α is directly relevant to a term \mathfrak{t} if there exists a model \mathcal{I} of α and a $\{\mathfrak{t}\}$ -variant \mathcal{I}' of \mathcal{I} such that $\mathcal{I}' \not\models \alpha$.*

Note that if an axiom α does not contain a term \mathfrak{t} in its signature, then it does not directly constrain it. However, the inverse implication does not hold as we saw before.

An interesting strong relation between the notion of direct relevance and the notion of model conservativity is discussed in Prop. 1.

Proposition 1. *Let α be a consistent axiom and $\mathfrak{t} \in \tilde{\alpha}$ a term. If there exists a signature $\Sigma \subseteq \tilde{\alpha}$ such that $\Sigma \ni \mathfrak{t}$, $\alpha \not\equiv_{\Sigma}^{mCE} \emptyset$, and $\alpha \equiv_{\Sigma \setminus \{\mathfrak{t}\}}^{mCE} \emptyset$, then α is directly relevant for \mathfrak{t} .*

Proof. Let α be an axiom and \mathfrak{t} be a term such that there exists a signature $\Sigma \subseteq \tilde{\alpha}$ satisfying the hypothesis. Since $\alpha \not\equiv_{\Sigma}^{mCE} \emptyset$, we know that there exists an interpretation \mathcal{J} over Σ that cannot be extended to a model for α . In contrast, since $\alpha \equiv_{\Sigma \setminus \{\mathfrak{t}\}}^{mCE} \emptyset$, we have that $\mathcal{J}|_{\Sigma \setminus \{\mathfrak{t}\}}$ can be extended to a model \mathcal{I} for α . Let us now consider the interpretation \mathcal{J}' that interprets all symbols in $\alpha \setminus \{\mathfrak{t}\}$ as \mathcal{I} does, whilst it interprets \mathfrak{t} as \mathcal{J} does. Then, \mathcal{J}' is not a model for α , and it is a \mathfrak{t} -variant for \mathcal{I} . Hence, α is directly relevant for \mathfrak{t} . \square

The inverse implication in Prop. 1 does not hold in general. For example, let us consider the axiom $\{a\} \sqsubseteq A$. Then, there is no model that interprets the symbol A as the empty set. In particular, $\alpha \not\equiv_{\emptyset}^{mCE} \emptyset$. However, the inverse implication can still hold

for less expressive DL than *SHROIQ*. An investigation on the characterization of the languages for which the inverse of Prop. 1 is part of our future work.

Let us now recall the trivial example in the introduction, where $\mathcal{O} = \{\alpha_i\}_{i=1\dots n}$ and the i -th axiom is $A_{i-1} \sqsubseteq A_i$. If $n \geq 2$ then the axiom α_n does not contain A_0 . However, α_n does *indirectly* constrain A_0 because in any model \mathcal{I} of \mathcal{O} where $A_{n-1}^{\mathcal{I}} = \emptyset$ (and α_n is then satisfied for any interpretation of A_n) we have that $A_0^{\mathcal{I}}$ is forced to be empty. In other words, in order to define irrelevance between a term and an axiom we need then to look at their interpretations *in the context of the ontology*. Intuitively, an axiom α is irrelevant for \mathfrak{t} w.r.t. the ontology \mathcal{O} if the interpretation of \mathfrak{t} and the interpretation of symbols in $\tilde{\alpha} \setminus \mathfrak{t}$ can be chosen independently from each other, even though we still have to take into account the constraints provided by \mathcal{O} . This idea is formalised in Def. 4.

Definition 4. *Let \mathcal{O} be an ontology, $\mathfrak{t} \in \tilde{\mathcal{O}}$ a term, and $\alpha \in \mathcal{O}$ an axiom. We say that α is \mathcal{O} -irrelevant for \mathfrak{t} if, for any $\{\mathfrak{t}\}$ -model $(\Delta^{\mathcal{I}_1}, \mathcal{I}_1)$ w.r.t. $\mathcal{O} \setminus \alpha$ and any $\{\tilde{\alpha} \setminus \mathfrak{t}\}$ -model $(\Delta^{\mathcal{I}_2}, \mathcal{I}_2)$ w.r.t. \mathcal{O} , there exists a model \mathcal{J} which is an \mathcal{O} -extension of both \mathcal{I}_1 and \mathcal{I}_2 . We say that α is \mathcal{O} -relevant for \mathfrak{t} if it is not \mathcal{O} -irrelevant.*

In many cases, if an axiom $\alpha \in \mathcal{O}$ is directly relevant for a term \mathfrak{t} , then α is also \mathcal{O} -relevant for \mathfrak{t} . However, this condition fails to hold when both α and $\mathcal{O} \setminus \alpha$ imply that there is only one valid $\{\mathfrak{t}\}$ -model w.r.t. \mathcal{O} . Such a peculiar case is described in the following example: let us consider the ontology $\mathcal{O} = \{\mathfrak{t} \sqsubseteq \perp, \mathfrak{t} \sqsubseteq A \sqcap \neg A\}$. Now, both axioms are clearly directly relevant for \mathfrak{t} . However, both of them are not \mathcal{O} -relevant for \mathfrak{t} , because we cannot find any $\{\mathfrak{t}\}$ -model w.r.t. \mathcal{O} where \mathfrak{t} is interpreted differently. Hence, we need the following unifying notion of relevance.

Definition 5. *An axiom $\alpha \in \mathcal{O}$ is said to be relevant for a term $\mathfrak{t} \in \tilde{\mathcal{O}}$ if α is either directly relevant or \mathcal{O} -relevant for \mathfrak{t} . Otherwise, α is said to be irrelevant for \mathfrak{t} .*

Note that this notion of relevance is still defined in the context of the ontology \mathcal{O} .

In the following, we denote:

1. the set of (semantically) directly relevant axioms w.r.t. a term \mathfrak{t} by $semDR_{\mathcal{O}}(\mathfrak{t}) = \{\alpha \in \mathcal{O} \mid \alpha \text{ is directly relevant for } \mathfrak{t}\}$
2. the set of (semantically) relevant axioms w.r.t. \mathfrak{t} by $semRel_{\mathcal{O}}(\mathfrak{t}) = \{\alpha \in \mathcal{O} \mid \alpha \text{ is relevant for } \mathfrak{t}\}$.

In this paper we are not interested in investigating the complexity for deciding relevance of an axiom to a term. Our aim is to use modules based on syntactic locality to efficiently get the two approximations $\mathbf{DC}_{\mathcal{O}}(\mathfrak{t})$ for the set $semRel_{\mathcal{O}}(\mathfrak{t})$, and $\mathbf{C}_{\mathcal{O}}(\mathfrak{t})$ for $semDR_{\mathcal{O}}(\mathfrak{t})$. By approximation we mean that *all* relevant axioms are preserved, even if some irrelevant axiom can sneak into such sets. This is still ongoing research, and from this point on the reader will find only definitions, examples, and conjectures. Proving these results is included in our future work.

4 Locality-based Relevance

In [7] we have analysed several forms of modularity to detect logically coherent subsets of an ontology (and more in general of a logical theory). The idea behind that paper

is that each kind of module determines a granular structure in the ontology, and this identifies clusters of axioms that stick together and axioms that can be separated from those clusters. Some of these notions of modularity relate also each cluster of axioms to a subset of the vocabulary used—hence relating axioms to terms. However, all these kinds of modularity suffer from inducing a coarse notion of internal coherence, and in some notable examples the ontology cannot be decomposed into smaller bits, even if it seems to be well structured.

In the same paper, we also analyse the partitioning of an ontology provided by one of its ADs. Atoms are generally very small bits, as discussed in [5], hence in principle such bits do not suffer from aggregating together (too many) unrelated axioms. However, ADs did not come with a semantically-based notion of relevance between the atoms and the terms of an ontology.

In what follows we define two labelling functions, the first mapping each axiom to relevant terms, and the second mapping each atom to the set of relevant terms. Then, we describe and conjecture the relations of the resulting LAD with the two sets $semRel_{\mathcal{O}}(\tau)$ and $semDR_{\mathcal{O}}(\tau)$, defined in the previous paragraph, that contain the relevant and the directly relevant axioms to a term.

Definition 6. *Let α be an axiom, and let τ be a term such that $\tau \in \tilde{\alpha}$. We say that:*

1. α constrains τ from above if there exists a minimal seed signature containing τ that makes α non \perp -local
2. α constrains τ from below if there exists a minimal seed signature containing τ that makes α non \top -local
3. α constrains τ if α constrains τ either from above, or from below.

We denote:

- by $C_{\mathcal{O}}^{\perp}(\tau)$ the set of axioms in \mathcal{O} constraining a term τ from above
- by $C_{\mathcal{O}}^{\top}(\tau)$ the set of axioms in \mathcal{O} constraining a term τ from below
- by $C_{\mathcal{O}}(\tau)$ the set of axioms in \mathcal{O} constraining a term τ .

We conjecture that the notion of constraining a term is an approximation of the notion of direct relevance.

Conjecture 1. Let τ be a term in the signature of the ontology \mathcal{O} . Then,

$$semDR_{\mathcal{O}}(\tau) \subseteq C_{\mathcal{O}}(\tau).$$

5 LADs and the Double Cones of Relevance

The next step is to relate the set $semRel_{\mathcal{O}}(\tau)$ to a suitable efficiently computable approximation. The idea is to identify in the ADs of the ontology the axioms constraining a term, and then follow along the ADs how the consequences logically “span” across the whole ontology. Specifically, we define the following labelling functions, and we conjecture that the resulting LADs are able to keep track of the logical relations between the axioms of an ontology \mathcal{O} and the terms in $\tilde{\mathcal{O}}$.

Definition 7. *Let α be an axiom in the ontology \mathcal{O} . We define the following labelling functions:*

1. $lab^\perp : \mathcal{O} \rightarrow \wp(\tilde{\mathcal{O}})$ that maps each axiom to the set of all terms that are constrained by α from above
2. $lab^\top : \mathcal{O} \rightarrow \wp(\tilde{\mathcal{O}})$ that maps each axiom to the set of all terms that are constrained by α from below
3. for any notion of module $x \in \{\perp, \top\}$, $Lab^x : \mathcal{A}^x(\mathcal{O}) \rightarrow \wp(\tilde{\mathcal{O}})$ that maps each atom α to the set $\bigcup_{\alpha \in \mathfrak{a}} lab^x(\alpha)$.

By Def. 7.3 we can define a LAD that is able to keep track of the logical relevance throughout the whole ontology. Before proceeding further, we want to include an example to support the understanding of what follows.

Example 2. Consider the toy ontology as in Example 1. Then, the labelling function Lab^\perp is defined as follows:

- $\mathfrak{a}_1 \mapsto \{\text{Animal}\}$,
- $\mathfrak{a}_2 \mapsto \{\text{Person}\}$,
- $\mathfrak{a}_3 \mapsto \{\text{Vegan, Person, eats}\}$,
- $\mathfrak{a}_4 \mapsto \{\text{Student}\}$,
- $\mathfrak{a}_5 \mapsto \{\text{GraduateStudent, Student, hasHabitat}\}$,
- $\mathfrak{a}_6 \mapsto \{\text{Car}\}$,
- $\mathfrak{a}_7 \mapsto \{\text{Truck}\}$,
- $\mathfrak{a}_8 \mapsto \{\text{Car, Truck}\}$.

The corresponding \perp -LAD is represented in Fig. 1. The terms `hasGender`, `hasHabitat`,

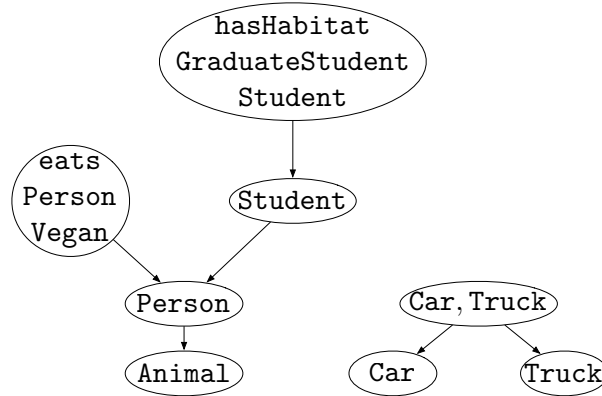


Figure 1. \perp -LAD of our example ontology

Vegetable, Mushroom, University, BA, and BS are not shown in this LAD, since they are not constrained from the above in this ontology.

In order to obtain functions that map each term to the set of logically related *atoms*, rather than axioms, we can invert the labelling functions just defined.

Definition 8. Given a notion $x \in \{\perp, \top\}$, we define the set of home atoms of \mathfrak{t} to be $h_{\mathcal{O}}^x(\mathfrak{t}) = \{\mathfrak{a} \mid \exists \alpha \in \mathfrak{a}, \mathfrak{t} \in lab^x(\alpha)\}$.

As an example, we then have that in our toy ontology $\mathbf{h}_{\mathcal{O}}^{\perp}(\text{Student}) = \{\mathbf{a}_5, \mathbf{a}_6\}$.

Definition 9. Given an ontology \mathcal{O} , a notion of locality $x \in \{\perp, \top\}$, the x -LAD of \mathcal{O} ($\mathcal{A}(\mathcal{O}), \succ, \text{Lab}^x$), and a term $\mathfrak{t} \in \tilde{\mathcal{O}}$, we define the double cone of x -relevance for \mathfrak{t} to be the set

$$\mathbf{DC}_{\mathcal{O}}^x(\mathfrak{t}) = \bigcup_{\mathbf{a} \in \mathbf{h}_{\tilde{\mathcal{O}}}(\mathfrak{t})} (\downarrow \mathbf{a} \cup \uparrow \mathbf{a}).$$

We conjecture that the LADs described in this paper allows us to identify which axioms are logically relevant for a term in an ontology.

Conjecture 2. If an axiom α is relevant for a term \mathfrak{t} , then $\alpha \in \mathbf{DC}_{\mathcal{O}}^{\perp}(\mathfrak{t}) \cup \mathbf{DC}_{\mathcal{O}}^{\top}(\mathfrak{t})$.

6 A Consequence of ADs on Models

Finally, we want to look closer at the inseparability relation between two ontologies $\mathcal{O}_1 \subsetneq \mathcal{O}_2$ such that $\mathcal{O}_1 \equiv_{\mathcal{O}}^{mCE} \mathcal{O}_2$. Then, we have by definition that:

$$\{\mathcal{I} \mid \mathcal{I} \models \mathcal{O}_1\} = \{\mathcal{J} \mid_{\mathcal{O}_1} \mathcal{J} \models \mathcal{O}_2\}.$$

Intuitively, we can think of \mathcal{O}_2 as extending \mathcal{O}_1 *without spoiling the models* already identified for \mathcal{O}_1 . This notion can be formalised as in what follows.

Definition 10. A chain of mCEs in \mathcal{O} is a family of ontologies $\mathcal{O}_1 \subsetneq \dots \subsetneq \mathcal{O}_{\ell} = \mathcal{O}$ such that $\mathcal{O}_i \equiv_{\mathcal{O}_i}^{mCE} \mathcal{O}_{i+1}$ for $i \in \{1, \dots, \ell\}$.

Our aim is to identify a chain of mCEs in an ontology \mathcal{O} by using a x -AD. In the following proposition we are going to use the notion of a *join* $\vee(\mathbf{a}_1, \dots, \mathbf{a}_{\kappa})$ of κ atoms, defined as the minimal module that contains all the atoms in $\{\mathbf{a}_1, \dots, \mathbf{a}_{\kappa}\}$.

Conjecture 3. Let \mathcal{O} be an ontology, $(\mathcal{A}(\mathcal{O}), \succ)$ be its x -AD, with $x \in \{\perp, \top\}$. Then, each chain of ontologies $\mathcal{O}_1 \subsetneq \dots \subsetneq \mathcal{O}_{\ell} = \mathcal{O}$ that respects the following criteria is a chain of mCEs in \mathcal{O} .

1. if $\mathcal{O}_j \supseteq \mathbf{a}$ and $\mathbf{a} \succ \mathbf{b}$, then $\mathcal{O}_j \supseteq \mathbf{b}$
2. if $\mathcal{O}_j \supseteq \mathbf{a} \cup \mathbf{b}$, then $\mathcal{O}_j \supseteq \vee(\mathbf{a}, \mathbf{b})$.


Definition 11. Let $x \in \{\perp, \top\}$ be a notion of module and \mathcal{O} be an ontology. Then, a chain $\mathcal{O}_1 \subsetneq \dots \subsetneq \mathcal{O}_{\ell} = \mathcal{O}$ of mCEs in \mathcal{O} defined via the x -LAD of \mathcal{O} is called x -chain of mCEs in \mathcal{O} .

To make the discussion clearer, let us consider the \perp -LAD as in Fig. 1 and refer to Example 1 for the axioms in each atom. Set $\mathcal{O}_1 = \{\alpha_7\}$, $\mathcal{O}_2 = \{\alpha_7, \alpha_8\}$, and $\mathcal{O}_3 = \{\alpha_7, \alpha_8, \alpha_9\}$. Then, $\mathcal{O}_1 \subsetneq \mathcal{O}_3 \subsetneq \mathcal{O}$ is a chain of mCEs, whilst $\mathcal{O}_1 \subsetneq \mathcal{O}_2 \subsetneq \mathcal{O}$ is not. Moreover, if we want to preserve everything that constrains the term *Person* from above, we see that we have to consider the principal ideal $\downarrow \mathbf{a}_3$.

Please note that for different ontologies we can still have that including the join of some set of atoms is not necessary to have a chain of mCEs, as described in the following example.

Example 3. Consider the following ontology \mathcal{O}' and its \perp -LAD:

$(\beta_1) \text{Bicycle} \sqsubseteq \text{NonMotorVehicle} \sqcap \text{TwoWheelsVehicle}$,
 $(\beta_2) \text{NonMotorVehicle} \sqsubseteq \neg \exists \text{hasPart.Engine}$,
 $(\beta_3) \text{TwoWheelsVehicle} \sqsubseteq = 2\text{hasWheel.Wheel}$.



We see that the \perp -AD of this ontology consists of 3 atoms: $\mathfrak{b}_i = \{\beta_i\}$ for $i = 1, 2, 3$, and the inherited poset structure is $\mathfrak{b}_1 \succ \mathfrak{b}_2$, $\mathfrak{b}_1 \succ \mathfrak{b}_3$. In this case, the following is a chain of mCEs: $\mathcal{O}'_1 = \{\beta_2\} \subsetneq \mathcal{O}'_2 = \{\beta_2, \beta_3\} \subsetneq \mathcal{O}'$, even if \mathcal{O}'_2 does not contain the join $\vee(\mathfrak{b}_2, \mathfrak{b}_3)$.

7 Conclusion and Future Work

In this paper we have introduced a notion of logical relevance in ontologies. Moreover, we have shown a promising way to reveal the relevance relations between the axioms and the terms of an ontology with a suitable LAD. Finally, we have described an interesting conjecture relating the models of an ontology and the LAD of an ontology.

Future work are 4-fold, and include:

1. the completion of the theoretical investigation, that misses the proofs for many conjectured results.
2. an experimental analysis of how on average the double cone of relevance of a term spans across an ontology. The experiment will take as an input some large datasets of diverse ontologies, for example the NCBO BioPortal ontology repository.
3. an experimental comparison between the notions $\mathbf{C}_{\mathcal{O}}^{\perp}$ and $\mathbf{C}_{\mathcal{O}}^{\top}$ of direct relevance based on LADs, and some frame-based notions, as the description and the usage views in Protégé 4. Chances are that these sets will not differ much. However, such a result would provide the frame-based views with a semantic foundation. Moreover, we will be able to analyze which kinds of logically related axioms are missed, or logically unrelated axioms are included, when we rely on a syntax-based approach to logical relevance.
4. an investigation on possible applications of the notions introduced in this paper. A preliminary idea consists of considering x -chains of mCEs when it comes to reason over an ontology. In fact, we know that, for any non-empty \perp -module \mathcal{M} , and for any concept $A \in \widetilde{\mathcal{M}}$, then $\widetilde{\mathcal{M}}$ contains also all the subsumees of A . This means that we can use the \perp -LAD to predict which concepts can be a subsumee of A .

Acknowledgements We would like to thank the anonymous referees and our colleagues I. Palmisano, E. Mikroyannidi, D. Tsarkov, and N. Matentzoglou for their comments and suggestions, which helped in improving this paper considerably.

References

1. F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. F. Patel-Schneider (eds.) *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press. (2003)

2. F. Baader, M. Bienvenu, C. Lutz, and F. Wolter: Query Answering over DL ABoxes: How to Pick the Relevant Symbols. In B. Cuenca Grau, I. Horrocks, B. Motik, and U. Sattler (eds.): Proc. of DL-09. *ceur-ws.org*, vol. 477. (2009)
3. B. Cuenca Grau, I. Horrocks, Y. Kazakov, and U. Sattler: Modular reuse of ontologies: Theory and practice. In: *J. of Artif. Intell. Research*, vol. 31, pp. 273–318. (2008)
4. C. Del Vescovo: The Modular Structure of an Ontology: Atomic Decomposition towards Applications. In R. Rosati, S. Rudolph, and M. Zakharyashev (eds.): Proc. of DL-11. *ceur-ws.org*, vol. 745. (2011)
5. C. Del Vescovo, D. Gessler, P. Klinov, B. Parsia, U. Sattler, T. Schneider, and A. Winget: Decomposition and Modular Structure of BioPortal Ontologies. In L. Aroyo, C. Welty, H. Alani, J. Taylor, A. Bernstein, L. Kagal, N. Noy, and E. Blomqvist. (eds.): Proc. of ISWC-11. (2011)
6. C. Del Vescovo, B. Parsia, U. Sattler, and T. Schneider: The Modular Structure of an Ontology: Atomic Decomposition. T. Walsh (eds.): Proc. of IJCAI-11. (2011)
7. C. Del Vescovo, B. Parsia, and U. Sattler: Topicality in Logic-Based Ontologies. S. Andrews, S. Polovina, R. Hill, and B. Akhgar (eds.): Proc. of ICCS-11. (2011)
8. S. Ghilardi, C. Lutz, and F. Wolter: Did I damage my ontology? A case for conservative extensions in description logics. In P. Doherty, J. Mylopoulos, and C. Welty (eds.): Proc. of KR-06 pp. 187–197. AAAI Press. (2006)
9. E. Jiménez-Ruiz, B. Cuenca Grau, U. Sattler, T. Schneider, and R. Berlanga Llavori: Safe and economic re-use of ontologies: A logic-based methodology and tool support. In: Proc. of ESWC-08. LNCS, vol. 5021, pp. 185–199. (2008)
10. A. Jimeno, E. Jiménez-Ruiz, R. Berlanga Llavori, and D. Rebolz-Schuhmann: Use of shared lexical resources for efficient ontological engineering. In: SWAT4LS-08. *ceur-ws.org*. (2008)
11. B. Konev, C. Lutz, D. Ponomaryov, and F. Wolter: Decomposing description logic ontologies. In: Proc. of KR-10. pp. 236–246. (2010)
12. B. Konev, C. Lutz, D. Walther, and F. Wolter: Formal properties of modularization. In: *Modular Ontologies*, LNCS, vol. 5445, pp. 25–66. Springer-Verlag. (2009)
13. G. Kourousias, and D. Makinson: Parallel interpolation, splitting, and relevance in belief change. In: *J. Symb. Log.*, vol. 72, n. 3, pp. 994–1002. (2007)
14. R. Kontchakov, L. Pulina, U. Sattler, T. Schneider, P. Selmer, F. Wolter, and M. Zakharyashev: Minimal module extraction from DL-Lite ontologies using QBF solvers. In: Proc. of IJCAI-09. pp. 836–841. (2009)
15. R. Parikh: Beliefs, belief revision, and splitting languages. In: L. S. Moss, J.G., and de Rijke, M. (eds.): *Logic, language and computation*, vol. 2, pp. 266–278. (1999)
16. B. Parsia, and T. Schneider: The modular structure of an ontology: an empirical study. In: Proc. of KR-10. pp. 584–586. AAAI Press (2010)
17. U. Sattler, T. Schneider, and M. Zakharyashev: Which kind of module should I extract? In B. Cuenca Grau, I. Horrocks, B. Motik, and U. Sattler (eds.): Proc. of DL-09. *ceur-ws.org*, vol. 477. (2009)