Which Kind of Module Should I Extract?[*]

Ulrike Sattler[1], Thomas Schneider[1], and Michael Zakharyaschev[2]

[1] University of Manchester, UK, {sattler,schneider}@cs.man.ac.uk
[2] Birkbeck College London, UK, michael@dcs.bbk.ac.uk

Abstract There are various techniques for specifying a module of an ontology that covers all knowledge about a given set of terms. These differ with respect to the size of the module, the complexity of its computation, and certain robustness properties. In this paper, we survey existing logic-based approaches, focus on syntactic approximations, and compare different kinds of modules with respect to their properties. This is intended to give guidelines on how to choose "the right kind of module".

1 Introduction

An ontology provides a common vocabulary (signature) for a domain of interest and describes the relationships between the terms built from that vocabulary. When developing an ontology, it is helpful for the engineer if she can reuse information from external, already existing, ontologies. Ideally, she should be able to extract modules that represent (only) the knowledge she wants to reuse. The problem of module extraction, therefore, can be phrased as follows: given a subset Σ of the vocabulary of an ontology, find a (minimal) subset of that ontology that is relevant for the terms[3] in Σ. This means that we are considering ontologies to be sets of axioms, and their modules to be subsets thereof.

The above requirement of "relevance" for Σ is of course still vague and open to interpretation. We will now provide intuitions for this requirement, using different phrasings. We will later model it precisely, using a set of equivalence relations between ontologies, called *inseparability relations*. There are different approaches to relevance for Σ; they can be grouped into structural ones, e.g. [13, 16], and logic-based ones, e.g. [4, 3, 9]. While the former focus and depend on the syntax of the axioms in the ontology and on the induced concept hierarchy, the latter are concerned with preserving entailments or models over a signature Σ. The latter acts as an interface for communication with the ontology in question: with Σ we specify a set of terms, and we expect to obtain a module—a subset that represents all knowledge about these terms from the original ontology. This is how we understand relevance, and we usually phrase it as "has the same entailments with respect to"—or sometimes as "has the same models modulo", which is a stronger condition. When we say that a logic-based module $\mathcal{M}$ of the ontology $\mathcal{O}$ with respect to a signature Σ "is relevant for" the terms in

---

[3] Terms are concept (class) names and role (property) names.

Σ, we mean that all consequences of $\mathcal{O}$ that can be expressed over Σ are also consequences of $\mathcal{M}$. Then $\mathcal{O}$ is said to be a conservative extension (CE) of $\mathcal{M}$.

The logic-based view seems theoretically sound and elegant and provides a desirable guarantee: reusing only terms from Σ, we will not be able to distinguish between importing $\mathcal{M}$ and importing $\mathcal{O}$ into our ontology. However, deciding CEs is computationally expensive in general: the problem of deciding whether two ontologies entail the same concept inclusions over a given signature is usually harder than standard reasoning tasks. For $DL\text{-}Lite_{horn}$, for instance, complexity grows from polynomial-time to coNP-complete, and for $DL\text{-}Lite_{bool}$ it grows from NP-complete to Π_2^p-complete [11]. For $\mathcal{ALC}$ and extensions, it is usually one exponential higher or even undecidable [5, 12]. These computational obstacles have led to syntactic approximations via locality [3]. Locality-based modules are in general not minimal, but provide the same guarantee, namely the preservation of certain entailments. They have been shown to be useful for economically reusing ontologies [7]. Syntactic locality comes in two variants, "top" and "bottom", and there are several ways of extracting a locality-based module ("top", "bottom", "bottom of top", vice versa, and further nesting).

This paper surveys existing logic-based approaches, focusing on syntactic approximations. We will compare different kinds of locality-based modules with each other and with CE-based modules. In particular, we will explore nested locality-based modules more in-depth than the existing literature. The comparison of module kinds will be based on examining properties relevant for ontology reuse. It will help learning which modules are best suited for which requirements.

The relevant properties have been identified in [8] and are sketched below. For our import scenario, we can consider a module $\mathcal{M}$ of the ontology $\mathcal{O}$ w.r.t. a signature Σ to be a subset of $\mathcal{O}$ that is indistinguishable from $\mathcal{O}$ w.r.t. Σ. The above mentioned inseparability relations generalise the import scenario and thus compare arbitrary ontologies that are not necessarily in the subset relation. They provide a unifying framework for comparing arbitrary definitions of modules.

Robustness under vocabulary restrictions. This property implies that a module of an ontology w.r.t. a signature Σ is also a module of this ontology w.r.t. any subset of Σ. This is important because it means that we do not need to import a different module when we restrict the set of terms that we are interested in.

Robustness under vocabulary extensions. This implies that a module of an ontology $\mathcal{O}$ w.r.t. Σ is also a module of $\mathcal{O}$ w.r.t. any $\Sigma' \supseteq \Sigma$ as long as $\Sigma' \setminus \Sigma$ does not share terms with $\mathcal{O}$. This means that we do not need to import a different module when extending the set of relevant terms with terms not from $\mathcal{O}$.

Robustness under replacement for a logic $\mathcal{L}$. This property implies that, if $\mathcal{M}$ is a module of $\mathcal{O}$ w.r.t. Σ, then the result of importing $\mathcal{M}$ into an $\mathcal{L}$-ontology $\mathcal{O}'$ is a module of the result of importing $\mathcal{O}$ into $\mathcal{O}'$. This is called module coverage in [7]: importing a module does not affect its property of being a module.

Robustness under joins. If two ontologies are indistinguishable w.r.t. Σ and they share only terms from Σ, then each of them is indistinguishable from their union w.r.t. Σ. This property, together with robustness under replacement, implies that it is not necessary to import two indistinguishable versions of the same ontology.

2 Preliminaries

In the following, we will use $\mathcal{L}$ to denote a description logic contained in $\mathcal{SHOIQ}$, whose standard syntax and semantics are defined in [1]. The reason for this choice is that locality-based modules are defined for logics up to $\mathcal{SHOIQ}$. We will consider an ontology as a TBox, i.e., a set of axioms.

Let $\mathsf{N_C}$ be a set of concept names, and $\mathsf{N_R}$ a set of role names. A *signature* $\varSigma$ is a set of terms, i.e., $\varSigma \subseteq \mathsf{N_C} \cup \mathsf{N_R}$. We can think of a signature as specifying a topic of interest. Axioms that only use terms from $\varSigma$ can be thought of as "on-topic", and all other axioms as "off-topic". For instance, if $\varSigma = \{\mathsf{Animal}, \mathsf{Duck}, \mathsf{Grass}, \mathsf{eats}\}$, then $\mathsf{Duck} \sqsubseteq \exists\mathsf{eats}.\mathsf{Grass}$ is on-topic, while $\mathsf{Duck} \sqsubseteq \mathsf{Bird}$ is off-topic.

Any concept name, role name, TBox, or axiom that uses only terms from $\varSigma$ is called a $\varSigma$-*concept*, $\varSigma$-*role*, $\varSigma$-*TBox*, or $\varSigma$-*axiom*. Given any such object X, we call the set of terms in X the *signature of* X and denote it with $\mathrm{sig}(X)$ or $\widetilde{X}$.

Given an interpretation $\mathcal{I}$, we denote its restriction to the terms in a signature $\varSigma$ with $\mathcal{I}|_\varSigma$. Two interpretations $\mathcal{I}$ and $\mathcal{J}$ are said to *coincide on a signature* $\varSigma$, in symbols $\mathcal{I}|_\varSigma = \mathcal{J}|_\varSigma$, if $\Delta^\mathcal{I} = \Delta^\mathcal{J}$ and $X^\mathcal{I} = X^\mathcal{J}$ for all $X \in \varSigma$.

2.1 Conservative extensions and locality

There are a number of variants of the notion of conservative extensions, which capture the desired preservation of knowledge to different degrees. We focus on the following basic ones.

Definition 1. Let $\mathcal{L}$ be a DL, $\mathcal{M} \subseteq \mathcal{T}$ be $\mathcal{L}$-TBoxes, and $\varSigma$ a signature.

(1) $\mathcal{T}$ is a *deductive* $\varSigma$-*conservative extension* ($\varSigma$-dCE) of $\mathcal{M}$ w.r.t. $\mathcal{L}$ if for all GCI axioms α over $\mathcal{L}$ with $\widetilde{\alpha} \subseteq \varSigma$, it holds that $\mathcal{M} \models \alpha$ if and only if $\mathcal{T} \models \alpha$.
(2) $\mathcal{T}$ is a *model* $\varSigma$-*conservative extension* ($\varSigma$-mCE) of $\mathcal{M}$ if $\{\mathcal{I}|_\varSigma \mid \mathcal{I} \models \mathcal{M}\} = \{\mathcal{I}|_\varSigma \mid \mathcal{I} \models \mathcal{T}\}$.
(3) $\mathcal{M}$ is a *dCE-based* (*mCE-based*) *module for* $\varSigma$ *of* $\mathcal{T}$ if $\mathcal{T}$ is a $\varSigma$-dCE ($\varSigma$-mCE) of $\mathcal{M}$ w.r.t. $\mathcal{L}$.

It is clear that $\mathcal{T}$ being a $\varSigma$-mCE of $\mathcal{M}$ implies that $\mathcal{T}$ is a $\varSigma$-dCE of $\mathcal{M}$.

Due to the computational difficulty to decide both kinds of CEs, approximations have been introduced [3]. They are based on *locality* of single axioms, which means that, given $\varSigma$, the axiom can always be satisfied independently of the interpretation of the $\varSigma$-terms, but in a restricted way: by interpreting all non-$\varSigma$ terms either as the empty set ($\emptyset$-locality) or as the full domain[4] (Δ-locality).

Definition 2. An axiom α over a logic $\mathcal{L}$ is called $\emptyset$-*local* (Δ-*local*) w.r.t. *signature* $\varSigma$ if, for each interpretation $\mathcal{I}$, there exists an interpretation $\mathcal{J}$ such that $\mathcal{I}|_\varSigma = \mathcal{J}|_\varSigma$, $\mathcal{J} \models \alpha$, and for each $X \in \widetilde{\alpha} \setminus \varSigma$, $X^\mathcal{J} = \emptyset$ (for each $C \in \widetilde{\alpha} \setminus \varSigma$, $C^\mathcal{J} = \Delta$ and for each $R \in \widetilde{\alpha} \setminus \varSigma$, $R^\mathcal{J} = \Delta \times \Delta$).

---

[4] Or, in the case of roles, the set of all pairs of domain individuals.

It has been shown in [3] that $\mathcal{M} \subseteq \mathcal{O}$ and all axioms in $\mathcal{O} \setminus \mathcal{M}$ being $\emptyset$-local (or all axioms being Δ-local) w.r.t. $\Sigma \cup \widetilde{\mathcal{M}}$ is sufficient for $\mathcal{O}$ to be a Σ-mCE of $\mathcal{M}$. The converse does not hold: e.g., the axiom $A \equiv B$ is neither $\emptyset$- nor Δ-local w.r.t. $\{A\}$, but the ontology $\{A \equiv B\}$ is an $\{A\}$-mCE of the empty ontology.

Furthermore, locality can be tested using available DL-reasoners [3], which makes this problem considerably easier than testing conservativity, see also Section 3.2. However, reasoning in expressive DLs is still complex, e.g. NExpTime-complete for $\mathcal{SHOIQ}$. In order to achieve *tractable* module extraction, the following syntactic approximation of locality has been introduced in [3].

Definition 3. An axiom α is called *syntactically $\bot$-local ($\top$-local) w.r.t. signature Σ* if it is of the form $C^\bot \sqsubseteq C$, $C \sqsubseteq C^\top$, $C^\bot \equiv C^\bot$, $C^\top \equiv C^\top$, $R^\bot \sqsubseteq R$ ($R \sqsubseteq R^\top$), or $\mathsf{Trans}(R^\bot)$ ($\mathsf{Trans}(R^\top)$), where C is an arbitrary concept, R is an arbitrary role name, $R^\bot \notin \Sigma$ ($R^\top \notin \Sigma$), and $C^\bot$ and $C^\top$ are from $\mathrm{Bot}(\Sigma)$ and $\mathrm{Top}(\Sigma)$ as defined in Figure 1 (a) (Figure 1 (b)).

(a) $\bot$-Locality   Let $A^\bot, R^\bot \notin \Sigma$, $C^\bot \in \mathrm{Bot}(\Sigma)$, $C^\top_{(i)} \in \mathrm{Top}(\Sigma)$, $\bar{n} \in \mathbb{N} \setminus \{0\}$

$\mathrm{Bot}(\Sigma) ::= A^\bot \mid \bot \mid \neg C^\top \mid C \sqcap C^\bot \mid C^\bot \sqcap C \mid \exists R.C^\bot \mid \geqslant \bar{n}\, R.C^\bot \mid \exists R^\bot.C \mid \geqslant \bar{n}\, R^\bot.C$
$\mathrm{Top}(\Sigma) ::= \top \mid \neg C^\bot \mid C^\top_1 \sqcap C^\top_2 \mid \geqslant 0\, R.C$

(b) $\top$-Locality   Let $A^\top, R^\top \notin \Sigma$, $C^\bot \in \mathrm{Bot}(\Sigma)$, $C^\top_{(i)} \in \mathrm{Top}(\Sigma)$, $\bar{n} \in \mathbb{N} \setminus \{0\}$

$\mathrm{Bot}(\Sigma) ::= \bot \mid \neg C^\top \mid C \sqcap C^\bot \mid C^\bot \sqcap C \mid \exists R.C^\bot \mid \geqslant \bar{n}\, R.C^\bot$
$\mathrm{Top}(\Sigma) ::= A^\top \mid \top \mid \neg C^\bot \mid C^\top_1 \sqcap C^\top_2 \mid \exists R^\top.C^\top \mid \geqslant \bar{n}\, R^\top.C^\top \mid \geqslant 0\, R.C$

Figure 1. Syntactic locality conditions

It has been shown in [3] that $\bot$-locality ($\top$-locality) of an axiom α w.r.t. Σ implies $\emptyset$-locality (Δ-locality) of α w.r.t. Σ. Therefore, all axioms in $\mathcal{O} \setminus \mathcal{M}$ being $\bot$-local (or all axioms being $\top$-local) w.r.t. $\Sigma \cup \widetilde{\mathcal{M}}$ is sufficient for $\mathcal{O}$ to be a Σ-mCE of $\mathcal{M}$. The converse does not hold; examples can be found in [3].

Modules of $\mathcal{O}$ for each of the four locality notions are obtained by starting with an empty set of axioms and subsequently adding axioms from $\mathcal{O}$ that are non-local. In order for this procedure to be correct, the signature against which locality is checked has to be extended with the terms in the axiom that is added in each step. Definition 4 (1) introduces locality-based modules, which are always mCE-based (and therefore dCE-based) modules [3], although not necessarily minimal ones. Modules based on syntactic locality can be made smaller by nesting $\top$-extraction into $\bot$-extraction and vice versa, and the result is still an mCE-based module. These so-called $\top\bot$-modules and $\bot\top$-modules are introduced in Definition 4 (2). Finally, we will see in Section 3.2 that iterated nesting of the latter can lead to even smaller (still CE-based) modules. These $\top\bot^*$-modules and $\bot\top^*$-modules are introduced in Definition 4 (3).

Definition 4. Let $x \in \{\emptyset, \Delta, \bot, \top\}$; let $y, z \in \{\bot, \top\}$ with $y \neq z$; let $\mathcal{T}$ be a TBox and Σ a signature.

(1) A TBox $\mathcal{M}$ is the *x-module of $\mathcal{T}$ w.r.t. Σ* if it is the output of Algorithm 1. We write $\mathcal{M} = x\text{-mod}(\Sigma, \mathcal{T})$.

(2) A TBox $\mathcal{M}$ is the *yz-module of $\mathcal{T}$ w.r.t. Σ*, written $\mathcal{M} = yz\text{-mod}(\Sigma, \mathcal{T})$, if $\mathcal{M} = y\text{-mod}(\Sigma, z\text{-mod}(\Sigma, \mathcal{T}))$.

(3) Let $(\mathcal{M}_i)_{i \geqslant 0}$ be a sequence of TBoxes such that $\mathcal{M}_0 = \mathcal{T}$ and $\mathcal{M}_{i+1} = yz\text{-mod}(\Sigma, \mathcal{M}_i)$ for every $i \geqslant 0$. For the smallest $n \geqslant 0$ with $\mathcal{M}_n = \mathcal{M}_{n+1}$, we call $\mathcal{M}_n$ the *yz^*-module of $\mathcal{T}$ w.r.t. Σ*, written $\mathcal{M} = yz^*\text{-mod}(\Sigma, \mathcal{T})$.

Algorithm 1 Extract a locality-based module

Input: TBox $\mathcal{T}$, signature Σ, $x \in \{\emptyset, \Delta, \bot, \top\}$
Output: x-module $\mathcal{M}$ of $\mathcal{O}$ w.r.t. Σ

$M \leftarrow \emptyset$; $\mathcal{T}' \leftarrow \mathcal{T}$
repeat
  changed $\leftarrow$ **false**
  for all $\alpha \in \mathcal{T}'$ **do**
    if α not x-local w.r.t. $\Sigma \cup \widetilde{\mathcal{M}}$ **then**
      $\mathcal{M} \leftarrow \mathcal{M} \cup \{\alpha\}$
      $\mathcal{T}' \leftarrow \mathcal{T}' \setminus \{\alpha\}$
      changed $\leftarrow$ **true**
    end if
  end for
until changed = **false**
return $\mathcal{M}$

As for (1), it has been shown in [3] that the output $\mathcal{M}$ of Algorithm 1 does not depend on the order in which the axioms α are selected.[5] Furthermore, the integer n in (3) exists because the sequence $(\mathcal{M}_i)_{i \geqslant 0}$ is decreasing. (We even have $\mathcal{M}_0 \supset \cdots \supset \mathcal{M}_n = \mathcal{M}_{n+1} = \dots$)

Modulo the locality check, Algorithm 1 runs in time cubic in $|\mathcal{T}| + |\Sigma|$ [3]. Modules based on $\bot/\top$-locality are therefore a feasible approximation for modules based on $\emptyset/\Delta$-locality. In both cases, modules are extracted axiom by axiom but, as said above, the $\emptyset/\Delta$-locality check is more complex.

2.2 Inseparability relations

The notions of modules defined so far were induced by CEs and different notions of locality. We will now put them into a more general context of modules generated by inseparability relations. For a given logic $\mathcal{L}$, an *inseparability relation* is a family $S = \{\equiv^S_\Sigma \mid \Sigma \text{ is a signature}\}$ of equivalence relations on the set of $\mathcal{L}$-TBoxes. The intuition behind this notion is as follows: $\mathcal{T}_1 \equiv^S_\Sigma \mathcal{T}_2$ means that $\mathcal{T}_1$ and $\mathcal{T}_2$ are indistinguishable w.r.t. Σ, i.e., they represent the same knowledge about the topic represented by Σ. The exact meaning of the terms "indistinguishable" and "the same knowledge" depends on the precise definition of the

---

[5] Our algorithm is a special case of the one in [3, Figure 4].

inseparability relation. $\mathcal{M}$ being a module for Σ of $\mathcal{T}$ should be equivalent to $\mathcal{M} \subseteq \mathcal{T}$ and $\mathcal{M}$ being inseparable w.r.t. Σ from $\mathcal{T}$.

The requirement to preserve entailments or models leads to the following inseparability relations, which have been examined in [8].

- $\mathcal{T}_1$ and $\mathcal{T}_2$ are Σ-*concept name inseparable*, written $\mathcal{T}_1 \equiv_\Sigma^c \mathcal{T}_2$, if for all Σ-concept names C, D, it holds that $\mathcal{T}_1 \models C \sqsubseteq D$ if and only if $\mathcal{T}_2 \models C \sqsubseteq D$.
- $\mathcal{T}_1$ and $\mathcal{T}_2$ are Σ-*subsumption inseparable* w.r.t. a logic $\mathcal{L}$, written $\mathcal{T}_1 \equiv_\Sigma^s \mathcal{T}_2$, if for all terms X and Y that are concept expressions over Σ or role names from Σ, it holds that $\mathcal{T}_1 \models X \sqsubseteq Y$ if and only if $\mathcal{T}_2 \models X \sqsubseteq Y$.
- $\mathcal{T}_1$ and $\mathcal{T}_2$ are Σ-*query inseparable*, written $\mathcal{T}_1 \equiv_\Sigma^q \mathcal{T}_2$, if for all Σ-ABoxes $\mathcal{A}$, Σ-queries $q(\mathbf{x})$ and tuples $\mathbf{a}$ of object names from $\mathcal{A}$, it holds that $(\mathcal{T}_1, \mathcal{A}) \models q(\mathbf{a})$ if and only if $(\mathcal{T}_2, \mathcal{A}) \models q(\mathbf{a})$.
- $\mathcal{T}_1$ and $\mathcal{T}_2$ are Σ-*model inseparable*, written $\mathcal{T}_1 \equiv_\Sigma^{\mathrm{sem}} \mathcal{T}_2$, if $\{\mathcal{I}|_\Sigma \mid \mathcal{I} \models \mathcal{T}_1\} = \{\mathcal{I}|_\Sigma \mid \mathcal{I} \models \mathcal{T}_2\}$.

We denote the respective sets of inseparability relations with S^c, S^s, and S^{sem}. It is easy to see that, for each signature Σ, it holds that $\equiv_\Sigma^{\mathrm{sem}} \subseteq \equiv_\Sigma^s \subseteq \equiv_\Sigma^c$.

Inseparability relations induce modules as follows.

Definition 5. Let S be an inseparability relation, $\mathcal{T}$ a TBox, $\mathcal{M} \subseteq \mathcal{T}$, and Σ a signature. We call $\mathcal{M}$

(1) an S_Σ-*module of* $\mathcal{T}$ if $\mathcal{M} \equiv_\Sigma^S \mathcal{T}$;
(2) a *self-contained* S_Σ-*module of* $\mathcal{T}$ if $\mathcal{M} \equiv_{\Sigma \cup \widetilde{M}}^S \mathcal{T}$;
(3) a *depleting* S_Σ-*module of* $\mathcal{T}$ if $\emptyset \equiv_{\Sigma \cup \widetilde{M}}^S \mathcal{T} \setminus \mathcal{M}$.

$\mathcal{M}$ is called a *minimal* (self-contained, depleting) $\equiv_\Sigma$-module of $\mathcal{T}$ if $\mathcal{M}$, but no proper subset of $\mathcal{M}$, is a (self-contained, depleting) $\equiv_\Sigma$-module of $\mathcal{T}$.

Due to the shift from Σ to $\Sigma \cup \widetilde{M}$, it is not necessarily the case that every self-contained (or depleting) S_Σ-module of $\mathcal{T}$ is an S_Σ-*module of* $\mathcal{T}$. However, under certain robustness properties, this implication holds.

While self-contained and depleting S_Σ-modules tend to be bigger than S_Σ-modules, they have important applications. One of those is the computation of all justifications for an entailment η of the ontology $\mathcal{T}$ [6], where depleting modules are essential. For appropriate inseparability relations S, Definition 5 (1) ensures that each $S_{\widetilde{\eta}}$-module of $\mathcal{T}$ contains at least one justification for η, but not necessarily all. We have good reasons to believe that each self-contained depleting $S_{\widetilde{\eta}}$-module of $\mathcal{T}$ contains all justifications for η.

Another application of depleting modules is the import scenario [3, 10]: if $\mathcal{M}$ is a depleting S_Σ-module of $\mathcal{T}$ and S satisfies certain robustness properties, then, for every module $\mathcal{M}'$ with $\widetilde{\mathcal{M}'} \cup \widetilde{\mathcal{T}} \subseteq \Sigma \cap \widetilde{M}$, we can import $\mathcal{M}'$ into $\mathcal{T} \setminus \mathcal{M}$ because they do not interfere with each other: $(\mathcal{T} \setminus \mathcal{M}) \cup \mathcal{M}' \equiv_{\Sigma'}^S \mathcal{M}'$.

Finally, while there can be exponentially many minimal S_Σ-modules, minimal depleting modules are uniquely determined—under mild conditions involving inseparability relations.

3 Modules and their properties

In order to find natural candidates for inseparability relations for locality-based modules, we proceed by analogy to inseparability relations for conservativity: ontologies $\mathcal{T}_1$ and $\mathcal{T}_2$ are inseparable if they have the same modules, i.e., if the module extraction algorithm returns the same output for each of them. We have replaced the semantic criterion "the same entailments/models" in conservativity-based inseparability with a syntactic criterion "the same extraction result" simply because locality-based modules are defined algorithmically. Furthermore, we will see that two of the thus obtained inseparability relations have almost all desired properties—which makes them even superior to dCE-based modules.

We consider the following inseparability relations for locality-based modules, where x stands for one of the locality notions $\emptyset$, Δ, $\bot$, and $\top$, and yz stands for one of the combinations $\bot\top$ and $\top\bot$.

Relation	$\mathcal{T}_1$ and $\mathcal{T}_2$ are in relation if . . .
$\equiv_\Sigma^x$	$x\text{-mod}(\Sigma, \mathcal{T}_1) = x\text{-mod}(\Sigma, \mathcal{T}_2)$
$\equiv_\Sigma^{yz}$	$yz\text{-mod}(\Sigma, \mathcal{T}_1) = yz\text{-mod}(\Sigma, \mathcal{T}_2)$
$\equiv_\Sigma^{yz^*}$	$yz^*\text{-mod}(\Sigma, \mathcal{T}_1) = yz^*\text{-mod}(\Sigma, \mathcal{T}_2)$

Evidently, they are all equivalence relations.

3.1 Robustness properties of inseparability relations

In the introduction, we have already sketched four important properties of inseparability relations and have seen why they are of interest for applications of modules. We will now define them.

Definition 6. Let $\mathcal{L}$ be a DL. The inseparability relation S is called

(1) *robust under vocabulary restrictions* if, for all $\mathcal{L}$-TBoxes $\mathcal{T}_1, \mathcal{T}_2$ and all signatures Σ, Σ' with $\Sigma \subseteq \Sigma'$, the following holds: if $\mathcal{T}_1 \equiv_{\Sigma'}^S \mathcal{T}_2$, then $\mathcal{T}_1 \equiv_\Sigma^S \mathcal{T}_2$.
(2) *robust under vocabulary extensions* if, for all $\mathcal{L}$-TBoxes $\mathcal{T}_1, \mathcal{T}_2$ and all signatures Σ, Σ' with $\Sigma' \cap (\widetilde{\mathcal{T}_1} \cup \widetilde{\mathcal{T}_2}) \subseteq \Sigma$: if $\mathcal{T}_1 \equiv_\Sigma^S \mathcal{T}_2$, then $\mathcal{T}_1 \equiv_{\Sigma'}^S \mathcal{T}_2$.
(3) *robust under replacement* if, for all $\mathcal{L}$-TBoxes $\mathcal{T}_1, \mathcal{T}_2$, all signatures Σ and every $\mathcal{L}$-TBox $\mathcal{T}$ with $\widetilde{\mathcal{T}} \cap (\widetilde{\mathcal{T}_1} \cup \widetilde{\mathcal{T}_2}) \subseteq \Sigma$, the following holds: if $\mathcal{T}_1 \equiv_\Sigma^S \mathcal{T}_2$, then $\mathcal{T}_1 \cup \mathcal{T} \equiv_\Sigma^S \mathcal{T}_2 \cup \mathcal{T}$.
(4) *robust under joins* if, for all $\mathcal{L}$-TBoxes $\mathcal{T}_1, \mathcal{T}_2$ and all signatures Σ with $\widetilde{\mathcal{T}_1} \cap \widetilde{\mathcal{T}_2} \subseteq \Sigma$ and every $i = 1, 2$, the following holds: if $\mathcal{T}_1 \equiv_\Sigma^S \mathcal{T}_2$, then $\mathcal{T}_i \equiv_\Sigma^S \mathcal{T}_1 \cup \mathcal{T}_2$.

As mentioned in Section 2.2, some of these robustness properties have an effect on the relation between the three different kinds of induced S_Σ-modules from Definition 5, see [15].

Let us now examine the properties of the inseparability relations S^{sem}, S^s, S^c, S^x, S^{yz} and S^{yz^*}, where $x \in \{\emptyset, \Delta, \bot, \top\}$ and $yz \in \{\bot\top, \top\bot\}$. First, in addition to the inclusions $\equiv_\Sigma^{\text{sem}} \subseteq \equiv_\Sigma^s \subseteq \equiv_\Sigma^c$ from Section 2.2, the following inclusions hold, see [15] for a proof.

Theorem 7. *Let Σ be a signature, $x \in \{\emptyset, \Delta, \bot, \top\}$ and $yz \in \{\top\bot, \bot\top\}$. Then the following properties hold:* $\equiv_\Sigma^x \subseteq \equiv_\Sigma^{\mathrm{sem}}$; $\equiv_\Sigma^{yz} \subseteq \equiv_\Sigma^{\mathrm{sem}}$; $\equiv_\Sigma^{yz^*} \subseteq \equiv_\Sigma^{\mathrm{sem}}$.

property relation	S^{sem}	S^s	S^c	S^x	S^{yz}	S^{yz^*}
corresponding module notion	mCE	dCE	—	x-mod	yz-mod	yz^*-mod
(min.) modules	✓⁻	✓⁻	—	✓	✗	✓
(min.) self-contained mod.s	✗	✗	—	✓	✓	✓
(min.) depleting modules	✗	✗	—	✓	✓	✓
robustness voc. restr.	✓	✓	✓	✓	✗	✓
robustness voc. ext.	✓	(✗)	(✓)	✓	✗	✓
robustness replacement	✓	(✗)	(✗)	✓	✓	✓
robustness joins	✓	(✗)	(✗)	✓	✓	✓

Symbols:

S^x	stands for $S^\emptyset$, S^Δ, $S^\bot$, $S^\top$
S^{yz}, S^{yz^*}	stand for $S^{\top\bot}$, $S^{\bot\top}$, $S^{\bot\top^*}$, $S^{\top\bot^*}$
✓, ✗	property holds/fails
(✓), (✗)	property holds/fails for many standard description logics
✓⁻	property holds except for minimality
—	property not considered: no corresponding module notion

Figure 2. Properties of inseparability relations for different module notions.

The properties of the equivalence relations defined in Subsection 2.2 are summarised in Figure 2. The first four lines give the module notion which is generalised by this inseparability relation, and indicate whether each such module for Σ is a (minimal) S_Σ-module, or a (minimal) self-contained S_Σ-module, or a (minimal) depleting S_Σ-module. For S^c, this question is meaningless because there is no corresponding standard module notion, apart from redefining Σ-dCEs to take only Σ-concept inclusions into account. For S^{sem} and S^s, it is clear that Definition 1 (3) leads to S_Σ-modules, but not to minimal, self-contained, or depleting ones. This can, however, easily be achieved by adopting stronger module notions as in [9, 10]. Hence, the negative entries in this part of the table are not problematic, and we can say that locality-based and conservativity-based modules are equally "good"—except for yz-modules: they are not always S_Σ^{yz}-modules, which is critical. However, they are always minimal self-contained (depleting) S_Σ^{yz}-modules.

The remaining four lines indicate whether the respective inseparability relation satisfies the four robustness properties from Definition 6. The results for S^{sem}, S^s and S^c are taken from [3, 8]; those for the locality-based relations will be proven in the following. This part of the table reveals the following insights: while it is not surprising and known that S^c lacks some of the important robustness properties, it is interesting (but known as well) that this is also the case for S^s. As for the locality-based inseparability relations, it is truly surprising that the S^{yz} lack two of four robustness properties, while the S^x and S^{yz^*} appear to be flawless and as good as S^{sem}.

The following theorem states the results in Figure 2 for locality-based inseparability relations. Their proofs can be found in [15].

Theorem 8. *Let $x \in \{\emptyset, \Delta, \bot, \top\}$ and $yz \in \{\bot\top, \top\bot\}$.*

(1) *The inseparability relation S_Σ^x is robust under replacement, vocabulary restrictions, vocabulary extensions and joins.*
(2) *The inseparability relation S_Σ^{yz} is not robust under vocabulary restrictions or extensions, but under replacement and joins.*
(3) *The inseparability relation $S_\Sigma^{yz^*}$ is robust under replacement, vocabulary restrictions, vocabulary extensions and joins.*
(4) *Let $\mathcal{T}$ be a TBox, Σ a signature, and $\mathcal{M} = x\text{-}mod(\Sigma, \mathcal{T})$. Then $\mathcal{M}$ is a minimal S_Σ^x-module, a minimal self-contained S_Σ^x-module, and a minimal depleting S_Σ^x-module of $\mathcal{T}$.*
(5) *Let $\mathcal{T}$ be a TBox, Σ a signature, and $\mathcal{M} = yz\text{-}mod(\Sigma, \mathcal{T})$. Then $\mathcal{M}$ is not generally a minimal S_Σ^{yz}-module, but it is a minimal self-contained S_Σ^{yz}-module and a minimal depleting S_Σ^{yz}-module of $\mathcal{T}$.*
(6) *Let $\mathcal{T}$ be a TBox, Σ a signature, and $\mathcal{M} = yz^*\text{-}mod(\Sigma, \mathcal{T})$. Then $\mathcal{M}$ is a minimal $S_\Sigma^{yz^*}$-module, a minimal self-contained $S_\Sigma^{yz^*}$-module, and a minimal depleting $S_\Sigma^{yz^*}$-module of $\mathcal{T}$.*

3.2 Minimality versus efficient computability

Theoretical results. For $\mathcal{ALC}$ and $\mathcal{ALCQI}$, the problem of deciding whether two ontologies entail the same concept inclusions over a given signature Σ is 2EXP-TIME-complete [5, 12], which is one exponential harder than standard reasoning tasks. For $\mathcal{ALCQIO}$, and hence for the more expressive OWL 1 and 2, this problem is undecidable [12]. The related problem of deciding whether two ontologies have the same models with respect to Σ is undecidable already for $\mathcal{ALC}$ [12].

Fortunately, extracting mCE-based modules *can* be tractable, for instance for acyclic $\mathcal{EL}$ ontologies [9]. We will examine how much locality-based modules differ from minimal modules in size and desired properties in this case.

For the DL-Lite family [2], deciding dCEs is Π_2^p-complete for *DL-Lite$_{bool}$* and coNP-complete for *DL-Lite$_{horn}$* [11], i.e., it is most likely intractable in both cases. However, there is experimental evidence [10] that minimal modules of DL-Lite ontologies can be extracted quite efficiently via locality-based modules as a preprocessing step and using QBF solvers.

Experimental results. Since CE-based modules are so hard to extract, it is difficult to perform experiments comparing modules of different notions extracted from the same ontology for the same signature. However, in two cases this has been possible. We will briefly describe the two cases and show the difference between conservativity-based and locality-based modules.

SNOMED CT, the *Systematized Nomenclature of Medicine, Clinical Terms*, is an ontology that consists of approximately 360,000 concepts and 1.4 million relationships, and defines the medical terminology for health care systems in the

US, the UK, and other countries. It is an acyclic $\mathcal{EL}$ TBox; therefore extracting mCE-based modules is tractable: the system MEX [9] extracts mCE-based modules from acyclic $\mathcal{EL}$ ontologies in polynomial time. Three signatures containing some 4,000, 16,000, and 24,000 concept names from SNOMED CT have been provided by the Intensive Care Unit of the Royal Prince Alfred Hospital in Sydney, Australia, from a corpus of 60 million tokens complied from 6 years of notes. (For the process of converting the text to SNOMED CT codes, see [14]). For the given signatures, we have extracted mCE-based modules with MEX, and locality-based modules via the OWL API[6]. In these cases, $\top$-modules comprised almost the whole ontology, and all yz-modules and yz^*-modules coincided. Furthermore, yz^*-modules were only about 1.5 times as big as mCE-modules for signatures greater than 15,000 symbols. The runtime was 1–4 seconds on an average PC in all cases. The sizes of mCE-modules, $\bot$-modules and $\top\bot^*$-modules are given in the following two tables containing exact sizes and percentages, respectively.

$|\Sigma|$	$\bot$	$\top\bot^*$	MEX
2,687	15,351	15,011	6,374
15,747	38,912	38,534	26,183
23,859	55,504	54,958	37,129

$|\Sigma|$	$\bot$	$\top\bot^*$	MEX
0.7%	4.0%	4.0%	1.7%
4.1%	10.2%	10.1%	6.9%
6.3%	14.6%	14.5%	9.8%

The version of SNOMED CT underlying these figures was dated 9 February 2005. A version of 30 December 2006 led to similar results.

In [10], dCE-based modules have been extracted from two ontologies in $DL\text{-}Lite_{bool}$, Core and Umbrella containing 1283 and 1247 axioms, for systematically and randomly generated signatures of size 1 and 10. These modules are MCM, MQM and MDQM, which are minimal S_Σ^c-modules, minimal S_Σ^q-modules (where S_Σ^q is based on query-CEs), and minimal depleting S_Σ^q-modules. The table below gives the average sizes for modules extracted for (a) all singleton signatures, (b) 30 randomly generated signatures consisting of 10 concept names, and (c) 30 random signatures consisting of 5 concept and 5 role names. The entry "—" means that no data is available due to performance reasons.

	(a) $|\Sigma| = 1$		(b) $|\Sigma| = 10$		(c) $|\Sigma| = 5 + 5$	
	Core	Umbrella	Core	Umbrella	Core	Umbrella
MCM	2	2	34	39	81	96
MQM	5	2	54	66	—	—
MDQM	80	57	294	319	497	485
$\top\bot^*$-mod	226	69	465	351	671	526

The results show that, on average, locality-based modules can be considerably bigger than dCE-based modules. On the other hand, since locality-based modules are always depleting (see Section 3.1), it is only fair to compare them with MDQM, and in only one of six cases are $\top\bot^*$-modules more than one and a half times as big as MDQMs on average. In [10], the time needed to extract the modules can be found in addition to their sizes.

---

[6] http://owlapi.sourceforge.net

On the nesting of locality-based modules. We have stated in Section 2.1 that the size of locality-based modules can be reduced by nesting the extraction of $\top$- and $\bot$-modules. In the first instance, this led to the notion of $\top\bot$- and $\bot\top$-modules. To the best of our knowledge, this has not been done before. As shown in the following example, it indeed leads to smaller modules.

Example 9. Let $\Sigma = \{A, D\}$ and $\mathcal{T} = \{A \sqsubseteq B, \ C \sqsubseteq D\}$.
Then we have $\top$-mod$(\Sigma, \mathcal{T}) = \{C \sqsubseteq D\}$ and $\bot$-mod$(\Sigma, \mathcal{T}) = \{A \sqsubseteq B\}$, but $\top\bot$-mod$(\Sigma, \mathcal{T}) = \bot\top$-mod$(\Sigma, \mathcal{T}) = \emptyset$.

There are cases where repeated nesting of $\top\bot$-modules ($\bot\top$-modules) decreases module size even further, see Example 10.

Example 10. Consider $\Sigma = \{A\}$ and $\mathcal{T} = \{A \sqsubseteq B \sqcup C, \ B \sqsubseteq A\}$. Then we have $\mathcal{T} \supset \top\bot$-mod$(\Sigma, \mathcal{T}) \supset \top\bot$-mod$(\Sigma, \top\bot$-mod$(\Sigma, \mathcal{T}))$ because:

$$
\begin{aligned}
\mathcal{N}_1 &= \bot\text{-mod}(\Sigma, \mathcal{T}) &&= \{A \sqsubseteq B \sqcup C, \ B \sqsubseteq A\} \\
\mathcal{M}_1 &= \top\text{-mod}(\Sigma, \mathcal{N}_1) &&= \{B \sqsubseteq A\} \\
\mathcal{N}_2 &= \bot\text{-mod}(\Sigma, \mathcal{M}_1) &&= \emptyset \\
\mathcal{M}_2 &= \top\text{-mod}(\Sigma, \mathcal{N}_2) &&= \emptyset
\end{aligned}
$$

An analogous example for $\bot\top$-modules is $\Sigma = \{A\}$; $\mathcal{T} = \{B \sqsubseteq A \sqcup \neg C, \ A \sqsubseteq B\}$.

It therefore makes sense to continue this nesting up to a fixpoint. The following lemma shows that the number of steps until this fixpoint is reached can be asymptotically as big as the number of axioms in $\mathcal{T}$. The proof is given in [15].

Lemma 11. *For every integer $n \geqslant 1$ and $yz \in \{\top\bot, \bot\top\}$, there exist an $\mathcal{ALC}$-TBox $\mathcal{M}_0$ of size $\mathcal{O}(n)$ and a signature Σ of size $\mathcal{O}(n)$ such that $\mathcal{M}_{i+1} = yz\text{-}mod(\Sigma, \mathcal{M}_i)$, for each $i = 0, \ldots, n-1$, and $\mathcal{M}_0 \supset \cdots \supset \mathcal{M}_n$.*

4 Conclusion

We have compared important properties of conservativity-based and locality-based modules via the more general notion of inseparability relations. It has turned out that, while modules based on locality are in general larger than conservativity-based ones, they are very robust. Two out of three versions of locality-based modules enjoy the same robustness properties as mCE-based modules and are therefore more robust than modules based on dCEs. In addition, they are self-contained and depleting. However, their robustness does not simply follow from the fact that they are a special case of mCE-based modules: yz-modules do not "inherit" all robustness properties.

Furthermore, modules based on syntactic locality can be extracted efficiently for all logics up to $\mathcal{SHOIQ}$. They can thus also serve as an intermediate step for extracting conservativity-based modules [10]. Except for the few cases where mCE-based modules can be extracted efficiently, locality-based modules seem best suited to module extraction scenarios because they combine desirable properties with computational feasibility.

References

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

[2] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. DL-Lite: Tractable description logics for ontologies. In *Proc. of AAAI-05*, pages 602–607, 2005.

[3] B. Cuenca Grau, I. Horrocks, Y. Kazakov, and U. Sattler. Modular reuse of ontologies: Theory and practice. *J. of Artificial Intelligence Research*, 31:273–318, 2008.

[4] B. Cuenca Grau, B. Parsia, E. Sirin, and A. Kalyanpur. Modularity and web ontologies. In *Proc. of KR-06*, pages 198–209, 2006.

[5] S. Ghilardi, C. Lutz, and F. Wolter. Did I damage my ontology? A case for conservative extensions in description logics. In *Proc. of KR-06*, pages 187–197, 2006.

[6] M. Horridge, B. Parsia, and U. Sattler. Laconic and precise justifications in OWL. In *Proc. of ISWC-08*, volume 5318 of *LNCS*, pages 323–338, 2008.

[7] 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*, volume 5021 of *LNCS*, pages 185–199, 2008.

[8] B. Konev, C. Lutz, D. Walther, and F. Wolter. Formal properties of modularization. In H. Stuckenschmidt, S. Spacciapietra, and C. Parent, editors, *Ontology modularization*, volume 5445 of *LNCS*. To be published.

[9] B. Konev, C. Lutz, D. Walther, and F. Wolter. Semantic modularity and module extraction in description logics. In *Proc. of ECAI-08*, pages 55–59, 2008.

[10] R. Kontchakov, L. Pulina, U. Sattler, T. Schneider, P. Selmer, F. Wolter, and M. Zakharyaschev. Minimal module extraction from DL-Lite ontologies using QBF solvers. In *Proc. of IJCAI-09*, 2009. To appear.

[11] R. Kontchakov, F. Wolter, and M. Zakharyaschev. Can you tell the difference between DL-Lite ontologies? In *Proc. of KR-08*, pages 285–295, 2008.

[12] C. Lutz, D. Walther, and F. Wolter. Conservative extensions in expressive description logics. In *Proc. of IJCAI-07*, pages 453–458, 2007.

[13] N. F. Noy and M. A. Musen. Specifying ontology views by traversal. In *Proc. of ISWC-04*, volume 3298 of *LNCS*, pages 713–725, 2004.

[14] J. Patrick, Y. Wang, and P. Budd. An automated system for conversion of clinical notes into SNOMED clinical terminology. In *Proc. of HKMD-07*, volume 68 of *CRPIT*, pages 219–226, 2007.

[15] U. Sattler, T. Schneider, and M. Zakharyaschev. Which kind of module should I extract? Technical report, University of Manchester, 2009. `http://www.cs.man.ac.uk/%7Eschneidt/publ/ssz09_insep_report.pdf`.

[16] J. Seidenberg and A. L. Rector. Web ontology segmentation: analysis, classification and use. In *Proc. of WWW-06*, pages 13–22, 2006.