

Just the Right Amount: Extracting Modules from Ontologies*

Bernardo Cuenca Grau, Ian Horrocks, Yevgeny Kazakov and Ulrike Sattler
The University of Manchester
School of Computer Science
Manchester, M13 9PL, UK
{bcg,horrocks,ykazakov,sattler}@cs.man.ac.uk

ABSTRACT

The ability to extract meaningful fragments from an ontology is key for ontology re-use. We propose a definition of a module that guarantees to completely capture the meaning of a given set of terms, i.e., to include all axioms relevant to the meaning of these terms, and study the problem of extracting minimal modules. We show that the problem of determining whether a subset of an ontology is a module for a given vocabulary is undecidable even for rather restricted sub-languages of OWL DL. Hence we propose two “approximations”, i.e., alternative definitions of modules for a vocabulary that still provide the above guarantee, but that are possibly too strict, and that may thus result in larger modules: the first approximation is semantic and can be computed using existing DL reasoners; the second is syntactic, and can be computed in polynomial time. Finally, we report on an empirical evaluation of our syntactic approximation which demonstrates that the modules we extract are surprisingly small.

Categories and Subject Descriptors

I.2.4 [Knowledge Representation Formalisms and Methods]:
Miscellaneous

General Terms

Algorithms

Keywords

Ontologies, Description Logics, OWL, Semantic Web

1. INTRODUCTION

The design, maintenance, reuse, and integration of ontologies are highly complex tasks—especially for ontologies formulated in a logic-based language such as OWL. Like software engineers, “ontology engineers” need to be supported by tools and methodologies that help them to minimise the introduction of errors, i.e., to ensure that ontologies have appropriate consequences. In order to develop this support, important notions from software engineering, such as *module*, *black-box behavior*, and *controlled interaction*, need to be

*This work is supported by the EU Project TONES (Thinking ONtologiES) ref: IST-007603 and by the EPSRC Project REOL (Reasoning in Expressive Ontology Languages) ref:EP/C537211/1. The authors would like to thank Boris Motik for his suggestions and assistance concerning this work.

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.
WWW 2007, May 8–12, 2007, Banff, Alberta, Canada.
ACM 978-1-59593-654-7/07/0005.

adapted so as to take into account the fact that an OWL ontology is, in essence, a logical theory; due to the expressive power of OWL, this turns out to be difficult.

In earlier work [4], we have studied modularity in the context of *collaborative ontology development* and *controlled integration*, and defined what it means for an ontology we are developing to be safely integrated with a “foreign” ontology; roughly speaking, such an integration is safe if it does not change the meaning of the terms in the foreign ontology.

In this paper, we focus on the use of modularity to support the *partial reuse* of ontologies: continuing with the above integration scenario, as a next step, we would like to *extract*, from the foreign ontology, a small fragment that captures the meaning of the terms we use in our ontology. For example, when building an ontology describing research projects, we may use terms such as *Cystic.Fibrosis* and *Genetic.Disorder* in our descriptions of medical research projects. In order to improve the precision of our ontology, we may want to add more detail about the meaning of these terms; for reasons of cost and accuracy, we would prefer to do this by reusing information from a medical ontology. Such ontologies are, however, typically very large, and importing the whole ontology would make the consequences of the additional information costly to compute and difficult for our ontology engineers (who are not medical experts) to understand. Thus, in practice, we need to extract a module that includes just the relevant information. Ideally, this module should be *as small as possible* while still *guaranteeing* to capture the meaning of the terms used; that is, when answering arbitrary queries against our projects ontology, importing the module would give us *exactly the same answers* as if we had imported the whole medical ontology. In this case, importing the module instead of the whole ontology will have no observable effect on our ontology—apart from allowing for more efficient reasoning.

Concerning the efficiency of reasoning, the time needed to process an ontology is often too high for ontology engineering, where fast response under changes in the ontology is required, or for deployment in applications, where fast response to queries is required. The ability to extract modules in the sense described above would address both these problems: it would allow us to identify a (hopefully small) part of the ontology that is affected by a given change or that is sufficient to answer a given query—and then to reason over this part only without losing any consequences.

The contributions of this paper are as follows:

1. We propose a definition of a *module* Q_1 within a given ontology Q for a given vocabulary S .
2. We take the above definition as a starting point, and investigate the problem of computing minimal modules. We show that none of the reasonable variants of this problem is solv-

able in general already for rather restricted sub-languages of OWL DL. In fact, it is even not possible to determine whether a subset \mathcal{Q}_1 of an ontology \mathcal{Q} is a module in \mathcal{Q} for \mathbf{S} .

3. Given these negative results, we propose two “approximations”, i.e., alternative definitions of a module that still guarantee to completely capture the meaning of the terms in \mathbf{S} , but that are possibly too strict, and that may thus result in larger modules; these approximations are based on the notion of *locality* of an ontology with respect to a vocabulary, as first introduced in [4]. The first approximation is semantic, and can be computed using existing OWL reasoners; the second one is a restriction of the first one which can be computed in polynomial time. We propose an algorithm for computing the smallest module for each of these approximations.
4. Finally, we describe our implementation and present our experimental results on a set of real-world ontologies of varying size and complexity. We show that, using our syntactic approximation, we obtain modules that are much smaller than the ones computed using existing techniques, but still sufficient to capture the meaning of the specified vocabulary.

This paper comes with a Technical Report [3], available online, which contains the complete proofs for the results we discuss here.

2. PRELIMINARIES

In this section we introduce description logics (DLs) [2] which underly modern ontology languages, such as OWL DL. The *syntax* of a description logic \mathbf{L} is given by a signature and a set of constructors. A *signature* (or *vocabulary*) \mathbf{S} of a DL is the (disjoint) union of a set \mathbf{C} of *atomic concepts* (A, B, \dots) representing sets of elements, a set \mathbf{R} of *atomic roles* (r, s, \dots) representing binary relations between elements, and a set \mathbf{I} of *individuals* (a, b, c, \dots) representing elements. Every DL provides *constructors* for defining the set $\text{Rol}(\mathbf{S})$ of (general) *roles* (R, S, \dots), the set $\text{Con}(\mathbf{S})$ of (general) *concepts* (C, D, \dots), and the set $\text{Ax}(\mathbf{S})$ of *axioms* (α, β, \dots) for a signature \mathbf{S} which is a union of *role axioms* (RBox), *terminological axioms* (TBox) and *assertions* (ABox).

\mathcal{EL} [1] is a simple description logic which allows one to construct complex concepts using *conjunction* $C_1 \sqcap C_2$ and *existential restriction* $\exists R.C$ starting from atomic concepts A , roles R and the *bottom concept* \perp . \mathcal{EL} provides no role constructors and no role axioms; thus, every role R in \mathcal{EL} is atomic. The TBox axioms of \mathcal{EL} can be either *concept definitions* $A \equiv C$ or *general concept inclusion axioms* (GCIs) $C_1 \sqsubseteq C_2$. \mathcal{EL} assertions are either *concept assertions* $a : C$ or *role assertions* $r(a, b)$.

The *basic description logic* \mathcal{ALC} [14] is obtained from \mathcal{EL} by adding *complement of concepts* $\neg C$. We introduce some additional constructors as abbreviations: the *top concept* \top is a shortcut for $\neg \perp$, the *disjunction of concepts* $C_1 \sqcup C_2$ stands for $\neg(\neg C_1 \sqcap \neg C_2)$, and the *value restriction* $\forall R.C$ stands for $\neg(\exists R.\neg C)$.

\mathcal{S} is an extension of \mathcal{ALC} where, additionally, some atomic roles can be declared to be *transitive* using a role axiom $\text{Trans}(r)$.

Further extensions of description logics include *inverse roles* r^{-} (indicated by appending a letter \mathcal{D}), *role inclusion axioms* (RIs) also called *role hierarchies* $R_1 \sqsubseteq R_2$ ($+\mathcal{H}$), *functional roles* $\text{Funct}(R)$ ($+\mathcal{F}$), *number restrictions* ($\geq n S$) ($+\mathcal{N}$), *qualified number restrictions* ($\geq n S.C$)¹ ($+\mathcal{Q}$), and *nominals* $\{a\}$ ($+\mathcal{O}$). Nominals make it possible to construct a concept representing a singleton set $\{a\}$ (a *nominal concept*) from an individual a . These extensions can be used in different combinations, for example \mathcal{ALCO}

¹the dual constructors ($\leq n S$) and ($\leq n S.C$) are abbreviations for $\neg(\geq n S.\neg C)$ and $\neg(\geq n S.\neg C)$, respectively

is an extension of \mathcal{ALC} with nominals; \mathcal{SHIQ} is an extension of \mathcal{S} with role hierarchies, inverse roles and qualified number restrictions; and \mathcal{SHOIQ} is the DL that uses all the constructors and axiom types we have presented.

Modern ontology languages, such as OWL [12], are based on description logics and, to a certain extent, are syntactic variants thereof. In particular, OWL DL corresponds to \mathcal{SHOIN} [8]. In this paper, we assume an *ontology* \mathcal{O} based on a description logic \mathbf{L} to be a set of axioms in \mathbf{L} . The *signature of an ontology* \mathcal{O} (of an axiom α) is the set $\text{Sig}(\mathcal{O})$ ($\text{Sig}(\alpha)$) of atomic concepts, atomic roles and individuals that occur in \mathcal{O} (respectively in α).

The main reasoning task for ontologies is *query answering*: given an ontology \mathcal{O} and an axiom α , check if \mathcal{O} implies α .

The logical entailment \models is defined using the *usual Tarski-style set-theoretic semantics* for description logics as follows. Given a signature $\mathbf{S} = \mathbf{R} \cup \mathbf{C} \cup \mathbf{I}$, an *\mathbf{S} -interpretation* \mathcal{I} is a pair $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set, called the *domain* of the interpretation, and $\cdot^{\mathcal{I}}$ is the *interpretation function* that assigns: to every $A \in \mathbf{C}$ a subset $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, to every $r \in \mathbf{R}$ a binary relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, and to every $a \in \mathbf{I}$ an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$.

The interpretation function $\cdot^{\mathcal{I}}$ is extended to complex roles and concepts via DL-constructors in the standard way (see [2, 3] for details). The *satisfaction* relation $\mathcal{I} \models \alpha$ between an interpretation \mathcal{I} and a DL axiom α (read as \mathcal{I} satisfies α) is also standard and can be found in [2, 3]. An interpretation \mathcal{I} is a *model* of an ontology \mathcal{O} if \mathcal{I} satisfies all axioms in \mathcal{O} . An ontology \mathcal{O} implies an axiom α (written $\mathcal{O} \models \alpha$) if $\mathcal{I} \models \alpha$ for every model \mathcal{I} of \mathcal{O} . An axiom α is a *tautology* if it is implied by the empty ontology.

Let \mathbf{S}_1, \mathbf{S} be signatures such that $\mathbf{S}_1 \subseteq \mathbf{S}$. The *restriction of an \mathbf{S} -interpretation* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ to \mathbf{S}_1 is an interpretation $\mathcal{I}|_{\mathbf{S}_1} = (\Delta^{\mathcal{I}_1}, \cdot^{\mathcal{I}_1})$ over \mathbf{S}_1 such that $\Delta^{\mathcal{I}_1} = \Delta^{\mathcal{I}}$ and $X^{\mathcal{I}_1} = X^{\mathcal{I}}$ for every $X \in \mathbf{S}_1$. An *expansion of an \mathbf{S}_1 -interpretation* \mathcal{I}_1 to \mathbf{S} is an \mathbf{S} -interpretation \mathcal{I} such that $\mathcal{I}|_{\mathbf{S}_1} = \mathcal{I}_1$. A *trivial expansion of an \mathbf{S}_1 -interpretation* \mathcal{I}_1 to \mathbf{S} is an expansion of \mathcal{I}_1 to \mathbf{S} such that $X^{\mathcal{I}} = \emptyset$ for every atomic concept and atomic role $X \in \mathbf{S} \setminus \mathbf{S}_1$.

3. MODULES FOR KNOWLEDGE REUSE

For exposition, suppose that an ontology engineer wants to build an ontology about research projects. The ontology defines different types of projects according to the research topics they focus on. Suppose that the ontology engineer defines two concepts `Genetic.Disorder.Project` and `Cystic.Fibrosis.EUProject` in his ontology \mathcal{P} . The first one describes projects about genetic disorders; the second one describes European projects about cystic fibrosis, as given by the axioms P1 and P2 in Figure 1.

The ontology engineer is supposed to be an expert on research projects: he knows, for example, that a `EUProject` is a `Project` (axiom P3). He is unfamiliar, however, with most of the topics the projects cover and, in particular, with the terms `Cystic.Fibrosis` and `Genetic.Disorder` mentioned in P1 and P2. In this case, he decides to reuse the knowledge about these subjects from a well-established and widely-used medical ontology

The most straightforward way to reuse these concepts is to import the medical ontology. This may be, however, a large ontology, which deals with other matters in which the ontology engineer is not interested, such as genes, anatomy, surgical techniques, etc. Ideally, one would like to extract a (hopefully small) fragment of the medical ontology—a *module*—that describes in detail the concepts we are reusing in our ontology. Intuitively, importing the module \mathcal{Q}_1 into \mathcal{P} instead of the full ontology \mathcal{Q} should have no impact on the modeling of the ontology \mathcal{P} .

Continuing with the example, suppose that `Cystic.Fibrosis` and `Genetic.Disorder` are described in an ontology \mathcal{Q} containing ax-

Ontology of medical research projects \mathcal{P} :	
P1	Genetic_Disorder_Project \equiv Project \sqcap $\sqcap \exists \text{has.Focus.Genetic_Disorder}$
P2	Cystic_Fibrosis_EUProject \equiv EUProject \sqcap $\sqcap \exists \text{has.Focus.Cystic_Fibrosis}$
P3	EUProject \sqsubseteq Project
Ontology of medical terms \mathcal{Q} :	
M1	Cystic_Fibrosis \equiv Fibrosis $\sqcap \exists \text{located.In.Pancreas}$ \sqcap $\sqcap \exists \text{has.Origin.Genetic_Origin}$
M2	Genetic_Fibrosis \equiv Fibrosis \sqcap $\sqcap \exists \text{has.Origin.Genetic_Origin}$
M3	Fibrosis $\sqcap \exists \text{located.In.Pancreas}$ \sqsubseteq Genetic_Fibrosis
M4	Genetic_Fibrosis \sqsubseteq Genetic_Disorder
M5	DEFBI_Gene \sqsubseteq Immuno_Protein_Gene \sqcap $\sqcap \exists \text{associated.With.Cystic_Fibrosis}$

Figure 1: Reusing medical terminology for an ontology on medical research projects

ioms M1-M5 in Figure 1. If we include in the module \mathcal{Q}_1 just the axioms that mention Cystic_Fibrosis or Genetic_Disorder, namely M1, M4 and M5, we lose the following dependency:

$$\text{Cystic_Fibrosis} \sqsubseteq \text{Genetic_Disorder} \quad (1)$$

The concept inclusions Cystic_Fibrosis \sqsubseteq Genetic_Fibrosis \sqsubseteq Genetic_Disorder follow from M1-M5, but not from M1, M4, M5, since the dependency Cystic_Fibrosis \sqsubseteq Genetic_Fibrosis does not hold after removing M2 and M3. The dependency (1), however, is crucial for our ontology \mathcal{P} as it (together with axiom P3) implies the following axiom:

$$\text{Cystic_Fibrosis_EUProject} \sqsubseteq \text{Genetic_Disorder_Project} \quad (2)$$

This means, in particular, that all the projects annotated with Cystic_Fibrosis_EUProject must be included in the answer for a query on Genetic_Disorder_Project. Consequently, importing a part of \mathcal{Q} containing only axioms that mention the terms used in \mathcal{P} instead of \mathcal{Q} results in an underspecified ontology. We stress that the ontology engineer might be unaware of dependency (2), even though it concerns the concepts of his primary scope.

The example above suggests that the central requirement for a module $\mathcal{Q}_1 \subseteq \mathcal{Q}$ to be reused in our ontology \mathcal{P} is that $\mathcal{P} \cup \mathcal{Q}_1$ should yield the *same* logical consequences in the vocabulary of \mathcal{P} as $\mathcal{P} \cup \mathcal{Q}$ does. Note that, as seen in the example, this requirement does not force us to include in \mathcal{Q}_1 all the axioms in \mathcal{Q} that mention the vocabulary to be reused, nor does it imply that the axioms in \mathcal{Q} that do not mention this vocabulary should be omitted.

Based on the discussion above, we formalize our first notion of a *module* as follows:

Definition 1 [Module]. Let $\mathcal{Q}_1 \subseteq \mathcal{Q}$ be two ontologies and \mathbf{S} a signature. We say that \mathcal{Q}_1 is an *S-module* in \mathcal{Q} w.r.t. a language L , if for every ontology \mathcal{P} and every axiom α expressed in L with $\text{Sig}(\mathcal{P} \cup \{\alpha\}) \cap \text{Sig}(\mathcal{Q}) \subseteq \mathbf{S}$, we have $\mathcal{P} \cup \mathcal{Q} \models \alpha$ iff $\mathcal{P} \cup \mathcal{Q}_1 \models \alpha$.

In Definition 1 the signature \mathbf{S} acts as the *interface* signature between \mathcal{P} and \mathcal{Q} in the sense that it contains the symbols that \mathcal{P} and α may share with \mathcal{Q} . It is also important to realize that there are two free parameters in Definition 1, namely the ontology \mathcal{P} and the axiom α . Both \mathcal{P} and α are formulated in some ontology language L , which might not necessarily be a sub-language of OWL DL.

Fixing the language L in which \mathcal{P} and α can be expressed is essential in Definition 1 since it may well be the case that \mathcal{Q}_1 is a module in \mathcal{Q} w.r.t. a language L_1 , but not w.r.t. L_2 . Fixing L , however, is not always reasonable. If \mathcal{Q}_1 is an *S-module* in \mathcal{Q} , it should always be possible to replace \mathcal{Q} with \mathcal{Q}_1 regardless of the particular language in which \mathcal{P} and α are expressed. In fact, we may extend our ontology \mathcal{P} with a set of Horn rules, or extend our query language to support arbitrary conjunctive queries. In any case, extending the ontology language for \mathcal{P} and the query language for α should not prevent \mathcal{Q}_1 from being a module in \mathcal{Q} .

It is therefore convenient to formulate a more general notion of a module which abstracts from the particular language under consideration; that is, we say that \mathcal{Q}_1 is an *S-module* in \mathcal{Q} iff it is an *S-module* in \mathcal{Q} , according to Definition 1 for *every* language L with Tarski-style set-theoretic semantics. The modules we obtain in this paper will be modules in precisely this stronger sense.

In our knowledge reuse scenario, small modules are preferred over large modules. Therefore, it makes sense to focus only on minimal modules. We say that \mathcal{Q}_1 is a *minimal S-module* in \mathcal{Q} if there is no $\mathcal{Q}_2 \subsetneq \mathcal{Q}_1$ that is also an *S-module* in \mathcal{Q} . In our example from Figure 1, there are two minimal *S-modules* $\mathcal{Q}_1 = \{M1, M2, M4\}$ and $\mathcal{Q}_2 = \{M1, M3, M4\}$: if we remove any axiom from them, the dependency (1) will no longer hold. Hence minimal modules are not necessarily unique. While in some cases it is reasonable to extract all minimal modules, in others it may suffice to extract just one. Thus, given \mathcal{Q} and \mathbf{S} , the following tasks are of interest:

- T1. compute *all* minimal *S-modules* in \mathcal{Q}
- T2. compute *some* minimal *S-module* in \mathcal{Q}

Surprisingly, we can show (see [3] for detail) that these tasks are inter-reducible; that is, an algorithm that solves T1 can be used to solve T2 and vice-versa. Let us now consider the axioms M1–M4. These axioms occur in both minimal *S-modules* \mathcal{Q}_1 and \mathcal{Q}_2 ; thus, they are, in a certain sense, essential for dependency (1). In certain situations, one can be interested in computing just the set \mathcal{Q}_e of such essential axioms, instead of computing all minimal modules. This is the case, for example, if the ontology engineer wants to compute a module that is “safe” under removal of axioms: if we remove M2 from \mathcal{Q} , then $\mathcal{Q}'_1 = \mathcal{Q}_1 \setminus M2 = \{M1, M4\}$ is no longer an *S-module* for the updated ontology $\mathcal{Q}' := \mathcal{Q} \setminus \{M2\}$ since the dependency (1) is lost, but $\mathcal{Q}'_e := \mathcal{Q}_e \setminus \{M2\}$ is still a module in \mathcal{Q} . This example suggests the following definition:

Definition 2 [Essential Axiom]. Given a signature \mathbf{S} and an ontology \mathcal{Q} , we say that an axiom $\alpha \in \mathcal{Q}$ is *S-essential* in \mathcal{Q} w.r.t. L if α belongs to some minimal *S-module* in \mathcal{Q} w.r.t. L .

Hence, the following task may also be of interest:

- T3. compute *the union* of all minimal *S-modules* in \mathcal{Q} , which is the set of all *S-essential* axioms in \mathcal{Q}

Obviously, task T3 is not harder than task T1: a procedure for computing all minimal modules can be used in a straightforward way to compute the union of these minimal modules.

PROPOSITION 1. *Tasks T1 and T2 are reducible to task T3; that is, any procedure for T1 or T2 can be used for solving T3.*

In the last few years, numerous techniques for extracting fragments of ontologies for knowledge reuse purposes have been developed. Most of these techniques rely on syntactically traversing the axioms in the ontology and employ various heuristics for determining which axioms are relevant and which are not.

An example of such a procedure is the algorithm implemented in the PROMPT-FACTOR tool [11]. Given a signature \mathbf{S} and an

ontology \mathcal{Q} , the algorithm retrieves a fragment $\mathcal{Q}_1 \subseteq \mathcal{Q}$ as follows: first, the axioms in \mathcal{Q} that mention any of the symbols in \mathbf{S} are added to \mathcal{Q}_1 ; second, \mathbf{S} is expanded with the symbols in $\text{Sig}(\mathcal{Q}_1)$. These steps are repeated until a fixpoint is reached. In our example, the axioms M1-M5 would be retrieved.

Another example is the algorithm in [15], which was used for segmentation of the medical ontology GALEN [13]. Given a signature \mathbf{S} and an ontology \mathcal{Q} , the algorithm adds to \mathcal{Q}_1 all definitions $A \equiv C$ for symbols in \mathbf{S} , expands \mathbf{S} with symbols in $\text{Sig}(\mathcal{Q}_1)$, and then repeats these steps again until a fixpoint is reached. The main idea of this algorithm is to prune irrelevant axioms by traversing the class hierarchy only “upwards” and across existential restrictions. Unfortunately, this algorithm does not detect other dependencies, in particular those expressed by GCIs. In our example, when initialized with Cystic.Fibrosis and Genetic.Disorder, the algorithm retrieves only the axiom M1 and the dependency (1) is lost.

Therefore, none of these algorithms is appropriate for extracting modules according to Definition 1. On the one hand, the PROMPT-FACTOR algorithm extracts many unnecessary axioms (such as M5 in our case) whereas, on the other hand, the segmentation algorithm from [15] misses essential axioms (like M2, M3 and M4).

In our example, the PROMPT-FACTOR algorithm would extract a module (though not a minimal one). In general, however, this is also not the case. For example, consider an ontology $\mathcal{Q} = \{\top \sqsubseteq \{a\}, A \sqsubseteq B\}$, $\alpha = (A \sqsubseteq \forall r.A)$, and $\mathbf{S} = \{A\}$. It is easy to see that \mathcal{Q} admits only single element models and α is satisfied in every such a model; that is, $\mathcal{Q} \models \alpha$. The PROMPT-FACTOR algorithm extracts in this case $\mathcal{Q}_1 = \{A \sqsubseteq B\}$, which does not imply α .

The main problem with these algorithms is that they ignore the semantics of the ontologies. As a consequence, they may, on the one hand, extract irrelevant axioms and, on the other hand, miss essential axioms. These algorithms, however, were not intended to extract modules in accordance to a formal collection of requirements; instead, they were intended to extract “relevant parts” of ontologies which are “likely to be related” to the given signature, and they do not guarantee the correctness of the results. Correctness, however, is the primary requirement for the procedures we present in this paper.

3.1 Modules and Conservative Extensions

The notion of a module is closely related to the notion of a conservative extension which has been used to characterize formal requirements in ontology integration tasks [7, 5, 4, 10]. In the literature we can find at least two different notions of conservative extensions in the context of ontologies [10]:

Definition 3 [Conservative Extensions].

Let $\mathcal{Q}_1 \subseteq \mathcal{Q}$ be two ontologies, \mathbf{S} a signature and L a logic.

We say that \mathcal{Q} is a *deductive \mathbf{S} -conservative extension* of \mathcal{Q}_1 w.r.t. L , if for every axiom α over L with $\text{Sig}(\alpha) \subseteq \mathbf{S}$, we have $\mathcal{Q} \models \alpha$ iff $\mathcal{Q}_1 \models \alpha$.

We say that \mathcal{Q} is a *model \mathbf{S} -conservative extension* of \mathcal{Q}_1 if, for every model \mathcal{I}_1 of \mathcal{Q}_1 , there exists a model \mathcal{I} of \mathcal{Q} such that $\mathcal{I}|_{\mathbf{S}} = \mathcal{I}_1|_{\mathbf{S}}$.

Intuitively, an ontology \mathcal{Q} is a deductive conservative extension of an ontology $\mathcal{Q}_1 \subseteq \mathcal{Q}$ for a signature \mathbf{S} iff every logical consequence α of \mathcal{Q} constructed using only symbols from \mathbf{S} is already a consequence of \mathcal{Q}_1 ; that is, the additional axioms in \mathcal{Q} do not add new logical consequences over the vocabulary \mathbf{S} . Analogously to modules, the notion of a deductive conservative extension depends on the ontology language L in which \mathcal{Q} and α are expressed. In contrast, model conservative extensions are not defined in terms

of logical entailment, but using the models directly. Intuitively, an ontology \mathcal{Q} is a model conservative extension of $\mathcal{Q}_1 \subseteq \mathcal{Q}$ if every model of \mathcal{Q}_1 can be expanded to a model of \mathcal{Q} by interpreting new symbols and leaving the interpretations of the old symbols unchanged.

The notion of semantic conservative extension is strictly stronger than the syntactic one [10] since it does not depend on expressivity of the ontology language. That is, if \mathcal{Q} is a model \mathbf{S} -conservative extension of \mathcal{Q}_1 , it is also a deductive \mathbf{S} -conservative extension of \mathcal{Q}_1 , but not necessarily vice versa.

Example 1. Let \mathcal{Q} be the ontology consisting of axioms M1 – M5 in Figure 1. Let \mathcal{Q}_1 consist of the axioms M1 – M4 and let $\mathbf{S} = \{\text{Cystic.Fibrosis}, \text{Genetic.Disorder}\}$. We show that \mathcal{Q} is a model \mathbf{S} -conservative extension of \mathcal{Q}_1 and, hence, also a deductive conservative extension of \mathcal{Q}_1 .

Let \mathcal{I}_1 be an arbitrary model of \mathcal{Q}_1 . We demonstrate that we can always construct a model \mathcal{I} of \mathcal{Q} which interprets the symbols from \mathbf{S} in the same way as \mathcal{I}_1 does, i.e. $\mathcal{I}|_{\mathbf{S}} = \mathcal{I}_1|_{\mathbf{S}}$.

Indeed, let \mathcal{I} be identical to \mathcal{I}_1 except for the interpretation of the atomic concepts DEFBL.Gene and Immuno.Protein.Gene, and the atomic role associatedWith, all of which we interpret in \mathcal{I} as the empty set. Note that these atomic concepts and this atomic role do not occur in \mathcal{Q}_1 . Hence, \mathcal{I} interprets the concepts in \mathcal{Q}_1 exactly like \mathcal{I}_1 , and so \mathcal{I} is a model of \mathcal{Q}_1 . Furthermore, \mathcal{I} is a model of M5 since the concepts on the left-hand-side and the right-hand-side of this axiom are both interpreted as the empty set. Thus, \mathcal{Q} is a model \mathbf{S} -conservative extension of \mathcal{Q}_1 .

Although Definition 1 is close to the notion of deductive conservative extension, there are two important differences. First, in the definition of deductive conservative extension, the logical consequences are considered only w.r.t. the ontologies \mathcal{Q} and \mathcal{Q}_1 of interest whereas, in our definition of module, all the possible ontologies \mathcal{P} in which the module can be used are taken into account. Second, in the definition of deductive conservative extension, the signature of α is required to be a subset of \mathbf{S} whereas, in our definition of module, only the common part of $\alpha \cup \mathcal{P}$ and \mathcal{Q} is required to be a subset of \mathbf{S} . Despite these differences, the two notions of conservative extensions are related to our notion of module:

PROPOSITION 2. *Let $\mathcal{Q}_1 \subseteq \mathcal{Q}$ be two ontologies. Then:*

1. *If \mathcal{Q}_1 is an \mathbf{S} -module in \mathcal{Q} w.r.t. L then \mathcal{Q} is a deductive \mathbf{S} -conservative extension of \mathcal{Q}_1 w.r.t. L ;*
2. *If \mathcal{Q} is a model \mathbf{S} -conservative extension of \mathcal{Q}_1 then \mathcal{Q}_1 is an \mathbf{S} -module in \mathcal{Q} for every ontology language L with Tarski-style set-theoretic semantics.*

PROOF. 1. Let α be an axiom with $\text{Sig}(\alpha) \in \mathbf{S}$ such that $\mathcal{Q} \models \alpha$. We have to show that $\mathcal{Q}_1 \models \alpha$ (\star). Take $\mathcal{P} := \emptyset$ (the empty ontology). Since \mathcal{Q}_1 is a module in \mathcal{Q} , $\text{Sig}(\mathcal{P} \cup \{\alpha\}) \cap \text{Sig}(\mathcal{Q}) \subseteq \mathbf{S}$, and $\mathcal{P} \cup \mathcal{Q} = \mathcal{Q} \models \alpha$, by Definition 1, we have $\mathcal{Q}_1 = \mathcal{P} \cup \mathcal{Q}_1 \models \alpha$.

2. Assume that \mathcal{Q} is a model \mathbf{S} -conservative extension of \mathcal{Q}_1 , but \mathcal{Q}_1 is not an \mathbf{S} -module in \mathcal{Q} w.r.t. some logic L . According to Definition 1, this means that there exists an ontology \mathcal{P} and an axiom α over L with $\text{Sig}(\mathcal{P} \cup \{\alpha\}) \cap \text{Sig}(\mathcal{Q}) \subseteq \mathbf{S}$, such that $\mathcal{P} \cup \mathcal{Q} \models \alpha$ but $\mathcal{P} \cup \mathcal{Q}_1 \not\models \alpha$. The last implies that for some interpretation \mathcal{I}_1 , we have $\mathcal{I}_1 \models \mathcal{P} \cup \mathcal{Q}_1$, but $\mathcal{I}_1 \not\models \alpha$. Let $\mathcal{I}'_1 := \mathcal{I}_1|_{\text{Sig}(\mathcal{Q})}$. Obviously, $\mathcal{I}'_1 \models \mathcal{Q}_1$. By Definition 3, since \mathcal{Q} is a model \mathbf{S} -conservative extension of \mathcal{Q}_1 , there exists an interpretation \mathcal{I}' such that $\mathcal{I}' \models \mathcal{Q}$ and $\mathcal{I}'|_{\mathbf{S}} = \mathcal{I}'_1|_{\mathbf{S}}$. Let \mathcal{I} be the expansion of $\mathcal{I}'|_{\text{Sig}(\mathcal{Q})}$ to $\text{Sig}(\mathcal{P} \cup \{\alpha\})$ by setting $X^{\mathcal{I}} := X^{\mathcal{I}'_1}$ for every $X \in \text{Sig}(\mathcal{P} \cup \{\alpha\}) \setminus \mathbf{S}$. Note that we also have $\mathcal{I}|_{\mathbf{S}} =$

$\mathcal{I}'|_S = \mathcal{I}'_1|_S = \mathcal{I}_1|_S$, hence $\mathcal{I}|_{\text{Sig}(\mathcal{P} \cup \{\alpha\})} = \mathcal{I}_1|_{\text{Sig}(\mathcal{P} \cup \{\alpha\})}$, and so $\mathcal{I} \models \mathcal{P}$ and $\mathcal{I} \not\models \alpha$. Since $\mathcal{I}|_{\text{SUSig}(\mathcal{Q})} = \mathcal{I}'|_{\text{SUSig}(\mathcal{Q})}$ and $\mathcal{I}' \models \mathcal{Q}$, we have $\mathcal{I} \models \mathcal{Q}$, which yields a contradiction. \square

Proposition 2 shows that our notion of module stays “in between” the two notions of conservative extensions. In particular, by applying Property 2 in Proposition 2 to Example 1, we can show that the axioms M1-M4 in Figure 1 constitute a module in the ontology \mathcal{Q} , consisting of M1-M5. The converse of Property 1 in Proposition 2, however, does not hold in general:

Example 2. Let $\mathcal{Q}_1 = \{\}$, $\mathcal{Q} = \{\top \sqsubseteq \exists R.A\}$ and $\mathbf{S} = \{A\}$. The ontology \mathcal{Q} is a deductive \mathbf{S} -conservative extension of \mathcal{Q}_1 w.r.t. \mathcal{ALC} . Indeed, every \mathcal{ALC} -axiom $\alpha = (C_1 \sqsubseteq C_2)$ over $\mathbf{S} = \{A\}$, is equivalent in \mathcal{ALC} to either $\top \sqsubseteq \top$, $\top \sqsubseteq \perp$, $\top \sqsubseteq A$ or $A \sqsubseteq \perp$, which are indistinguishable by \mathcal{Q}_1 and \mathcal{Q} —that is, the axiom is implied by \mathcal{Q}_1 iff it is implied by \mathcal{Q} . \mathcal{Q}_1 , however, is not an \mathbf{S} -module in \mathcal{Q} . Consider an \mathcal{ALC} -ontology $\mathcal{P} = \{A \sqsubseteq \perp\}$, which is constructed over \mathbf{S} . It is easy to see that $\mathcal{P} \cup \mathcal{Q} \models \top \sqsubseteq \perp$, but $\mathcal{P} \cup \mathcal{Q}_1 \not\models \top \sqsubseteq \perp$.

Given the relationships between our definition of module and conservative extensions, it is worth examining the computational complexity of the associated problems. The problem of deciding whether \mathcal{Q} is an \mathbf{S} -conservative extension of \mathcal{Q}_1 has been studied in [10], where it is proved to be 2NEXPTIME-complete for \mathcal{ALCQ} (roughly OWL-Lite) and undecidable for OWL DL. For model conservative extensions, the problem is highly undecidable (non recursively enumerable), even for \mathcal{ALC} [10].

The decidability result from [10] for deductive conservative extensions, however, does not transfer to our problem since an ontology \mathcal{Q} may well be an \mathbf{S} -deductive conservative extension of \mathcal{Q}_1 , but still \mathcal{Q}_1 might not be an \mathbf{S} -module in \mathcal{Q} . In fact, we show that our problem is already undecidable for \mathcal{ALC} ontologies when the language allows for nominals:

THEOREM 1. *Given a signature \mathbf{S} , an \mathcal{ALC} -ontology \mathcal{Q} and an axiom $\alpha \in \mathcal{Q}$, it is undecidable whether α is \mathbf{S} -essential in \mathcal{Q} w.r.t. $L = \mathcal{ALCO}$.*

The proof of Theorem 1 is a variation of the proof from [10] for undecidability of deductive conservative extensions in \mathcal{ALCQTO} , which is based on a reduction to domino tiling problems. The proof is rather technical and we refer the reader to [3] for details.

COROLLARY 1. *There exists no algorithm for performing any of the tasks T1-T3 from (3), and (4) for \mathcal{ALC} .*

PROOF. Theorem 1 implies directly that there is no algorithm for task T3 from (4), because otherwise, one can check if an axiom α is \mathbf{S} -essential in \mathcal{Q} by simply computing the set of all essential axioms by this algorithm for T3 and then checking if α is contained in this set. The remaining tasks from (3) are unsolvable since they are reducible to T3 by Proposition 1. \square

COROLLARY 2. *Given a signature \mathbf{S} , an \mathcal{ALC} -ontology \mathcal{Q} and an ontology $\mathcal{Q}_1 \subseteq \mathcal{Q}$, it is undecidable whether \mathcal{Q}_1 is an \mathbf{S} -module in \mathcal{Q} w.r.t. $L = \mathcal{ALCO}$.*

PROOF. The procedure for deciding if \mathcal{Q}_1 is an \mathbf{S} -module in \mathcal{Q} can be used for solving task T1, which is not possible by Corollary 1. Indeed, by enumerating the subsets of \mathcal{Q} and checking if they are modules, one can compute all subsets \mathcal{M} of \mathcal{Q} that are \mathbf{S} -modules in \mathcal{Q} . The set of all minimal modules in \mathcal{Q} can be then computed from \mathcal{M} by filtering out those sets in \mathcal{M} that are proper subsets of some other sets in \mathcal{M} . \square

Corollary 2 has a strong impact on the problem of knowledge reuse and forces us to revisit the original problem we aim at solving. As the problem of extracting minimal modules cannot be computationally solved for OWL DL in none of the forms T1-T3, we propose to relax some of the requirements in these tasks. We cannot drop the requirements that extracted fragments should be modules since, in this case, we have no guarantee for the correctness of the result. We can sacrifice, however, the minimality requirements for the computed modules and consider the following weakened version of the task T2:

T2w. compute *some* small enough \mathbf{S} -module in \mathcal{Q} (5)

Although it is always possible to extract an \mathbf{S} -module in \mathcal{Q} (one can simply return \mathcal{Q} which is always an \mathbf{S} -module in \mathcal{Q}), it still makes sense to develop, compare, and practically apply procedures that compute reasonably small modules. In the rest of the paper we describe two procedures of this form, based on the notions of locality, which we first introduced in [4]. The modules we obtain might be larger than the minimal modules and therefore we need to show that, in practice, they are still reasonably small.

4. MODULES BASED ON LOCALITY

In this section, we formulate the notion of locality, first introduced in [4] which will constitute the basis of our algorithm for extracting modules.

4.1 Locality

As a consequence of Case 2 in Proposition 2, model conservative extensions can be used as a sufficient condition for the notion of module. It is not possible, however, to design a procedure that extracts modules based on this condition since the problem of deciding model conservative extensions is highly undecidable [10]. The idea underlying this notion, however, can be used to establish sufficient conditions for the notion of module which are decidable and can be used in practice.

Consider Example 1, where we show that the set \mathcal{Q} of axioms M1-M5 in Figure 1 is a model \mathbf{S} -conservative extension of $\mathcal{Q}_1 = \{M1, \dots, M4\}$, for $\mathbf{S} = \{\text{Cystic_Fibrosis}, \text{Genetic_Disorder}\}$. In this example, the model conservative extension property was shown by finding expansions of $\text{Sig}(\mathcal{Q}_1)$ -interpretations to models of \mathcal{Q} in which all concept and atomic roles not in $\text{Sig}(\mathcal{Q}_1)$ were interpreted as the empty set. One could consider the cases where conservative extensions (and hence modules) can be determined in this manner. This idea can be formalized using the notion of locality:

Definition 4 [Locality [4]]. Let \mathbf{S} be a signature. We say that an axiom α is *local* w.r.t. \mathbf{S} if every trivial expansion of any \mathbf{S} -interpretation to $\mathbf{S} \cup \text{Sig}(\alpha)$ is a model of α . We denote by $\text{local}(\mathbf{S})$ the collection of all axioms that are local w.r.t. \mathbf{S} . An ontology \mathcal{O} is local w.r.t. \mathbf{S} if $\mathcal{O} \subseteq \text{local}(\mathbf{S})$.

Intuitively, an ontology \mathcal{O} is *local* w.r.t. a signature \mathbf{S} if we can take any interpretation for the symbols in \mathbf{S} and extend it to a model of \mathcal{O} that interprets the additional symbols as the empty set.

Example 3. Consider axiom M5 from Figure 1. This axiom is local w.r.t. $\mathbf{S} = \{\text{Cystic_Fibrosis}, \text{Genetic_Disorder}\}$. Indeed, as shown in Example 1, for every trivial expansion \mathcal{I} of an \mathbf{S} -interpretation to $\mathbf{S} \cup \text{Sig}(\alpha)$, the atomic concept DEFBI_Gene is interpreted as the empty set, and so, \mathcal{I} satisfies M5.

On the other hand, M5 is not local w.r.t. $\mathbf{S} = \{\text{DEFBI_Gene}\}$. Indeed, take any \mathbf{S} -interpretation \mathcal{I}_1 in which DEFBI_Gene is interpreted as a non-empty set. Then, for every trivial expansion \mathcal{I} of \mathcal{I}_1 , the concept on the left-hand-side of M5 is always interpreted

as a non-empty set, whereas the concept on the right-hand-side is always interpreted as the empty set. So \mathcal{I} does not satisfy α .

Locality can be used to formulate a sufficient condition for an ontology to be a model conservative extension of another ontology:

PROPOSITION 3. *Let $\mathcal{O}_1, \mathcal{O}_2$ be two ontologies and \mathbf{S} a signature such that \mathcal{O}_2 is local w.r.t. $\mathbf{S} \cup \text{Sig}(\mathcal{O}_1)$. Then $\mathcal{O}_1 \cup \mathcal{O}_2$ is an \mathbf{S} -model conservative extension of \mathcal{O}_1 .*

PROOF. Let \mathcal{I}_1 be a model of \mathcal{O}_1 . We show that there exists a model \mathcal{I} of $\mathcal{O}_1 \cup \mathcal{O}_2$ such that $\mathcal{I}|_{\mathbf{S}} = \mathcal{I}_1|_{\mathbf{S}}$.

Let \mathcal{I} be a trivial expansion of $\mathcal{I}_1|_{\mathbf{S} \cup \text{Sig}(\mathcal{O}_1)}$ to $\mathbf{S} \cup \text{Sig}(\mathcal{O}_1) \cup \text{Sig}(\mathcal{O}_2)$, thus, in particular, $\mathcal{I}|_{\mathbf{S} \cup \text{Sig}(\mathcal{O}_1)} = \mathcal{I}_1|_{\mathbf{S} \cup \text{Sig}(\mathcal{O}_1)}$. We need to show that \mathcal{I} is a model of $\mathcal{O}_1 \cup \mathcal{O}_2$. Since \mathcal{O}_2 is local w.r.t. $\mathbf{S} \cup \text{Sig}(\mathcal{O}_1)$, by Definition 4, \mathcal{I} is a model of \mathcal{O}_2 . Moreover, since $\mathcal{I}|_{\text{Sig}(\mathcal{O}_1)} = \mathcal{I}_1|_{\text{Sig}(\mathcal{O}_1)}$ and $\mathcal{I}_1 \models \mathcal{O}_1$, we have $\mathcal{I} \models \mathcal{O}_1$. Hence, $\mathcal{I} \models \mathcal{O}_1 \cup \mathcal{O}_2$ what was required to show. \square

Using Proposition 3 and Property 2 of Proposition 2 we obtain:

COROLLARY 3. *Let $\mathcal{O}_1, \mathcal{O}_2$ and \mathbf{S} be as given in Proposition 3. Then \mathcal{O}_1 is an \mathbf{S} -module in $\mathcal{O}_1 \cup \mathcal{O}_2$.*

Next, we introduce our first restricted class of modules:

Definition 5 [Modules based on Locality Condition].

Given an ontology \mathcal{Q} and a signature \mathbf{S} , we say that $\mathcal{Q}_1 \subseteq \mathcal{Q}$ is a *locality-based \mathbf{S} -module in \mathcal{Q}* if $\mathcal{Q} \setminus \mathcal{Q}_1$ is local w.r.t. $\mathbf{S} \cup \text{Sig}(\mathcal{Q}_1)$.

Example 4 [Example 3, continued]. We have seen in Example 3 that axiom M5 is local w.r.t. every \mathbf{S} that does not contain the atomic concept DEFBl.Gene. In particular, for \mathcal{Q}_1 consisting of axioms M1-M4 from Figure 1, M5 is local w.r.t. $\text{Sig}(\mathcal{Q}_1)$. Hence, according to Definition 5, \mathcal{Q}_1 is a locality-based \mathbf{S} -module in $\mathcal{Q} = \{\text{M1}, \dots, \text{M5}\}$ for every $\mathbf{S} \subseteq \text{Sig}(\mathcal{Q}_1)$.

4.2 Computing Locality-Based Modules

As demonstrated in Example 3, for testing locality of an axiom α w.r.t. \mathbf{S} , it is sufficient to interpret every atomic concept and atomic role not in \mathbf{S} with the empty set and then check if α is satisfied for all interpretations of the remaining symbols. This observation suggests that locality can be tested by first simplifying the ontology by eliminating atomic roles and concepts that are not in \mathbf{S} , and then checking if the resulting axioms are satisfied in every interpretation for the remaining symbols. This idea is formalized as follows:

PROPOSITION 4 (TESTING LOCALITY).

Let \mathcal{O} be a SHOIQ ontology and \mathbf{S} a signature. Let $\mathcal{O}_{\mathbf{S}}$ be obtained from \mathcal{O} by applying the transformations below, where every A is an atomic concept, every r is an atomic role with $A, r \notin \mathbf{S}$, and every R is a role r or r^- with $r \notin \mathbf{S}$: (1) replace all concepts of form $A, \exists R.C$ or $(\geq n R.C)$ with \perp ; (2) remove every transitivity axiom $\text{Trans}(r)$; (3) replace every assertion $a : A$ and $r(a, b)$ with the contradiction axiom $\top \sqsubseteq \perp$.

Then \mathcal{O} is local w.r.t. \mathbf{S} iff every axiom in $\mathcal{O}_{\mathbf{S}}$ is a tautology.

PROOF. It is easy to check that the transformation above preserves the satisfaction of axioms under every trivial expansion \mathcal{I} of every \mathbf{S} -interpretation to $\mathbf{S} \cup \text{Sig}(\mathcal{O})$. Hence, the resulting ontology $\mathcal{O}_{\mathbf{S}}$ is local w.r.t. \mathbf{S} iff the original ontology \mathcal{O} was local w.r.t. \mathbf{S} . Moreover, it is easy to see that there are no atomic concepts and atomic roles outside \mathbf{S} left in $\mathcal{O}_{\mathbf{S}}$ after the transformation. Hence, every axiom α from $\mathcal{O}_{\mathbf{S}}$ is a tautology iff \mathcal{Q} is local w.r.t. \mathbf{S} . \square

Note that according to Definition 4, assertions $a : A$ and $r(a, b)$ can never be local since they can only be satisfied by interpretations that interpret A and r as non-empty sets. Hence, assertions must be included in every locality-based module, which is reflected by the step (3) of the transformation in Proposition 4.

An important conclusion of Proposition 4 is that one can use the standard capabilities of available DL-reasoners² for testing locality since these reasoners can test for DL-tautologies. Checking for tautologies in description logics is, theoretically, a difficult problem (e.g. for DL SHOIQ is NEXPTIME-complete). There are, however, several reasons to believe that the locality test would perform well in practice. First, and most importantly, the size of the axioms in an ontology is usually small compared to the size of the ontology. Second, DL reasoners are highly optimized for standard reasoning tasks and behave well for most realistic ontologies.

In case this is too costly, it is possible to formulate a tractable approximation to the locality conditions for SHOIQ:

Definition 6 [Syntactic Locality for SHOIQ].

Let \mathbf{S} be a signature. The following grammar recursively defines two sets of concepts $\mathcal{C}_{\mathbf{S}}^{\perp}$ and $\mathcal{C}_{\mathbf{S}}^{\top}$ for a signature \mathbf{S} :

$$\begin{aligned} \mathcal{C}_{\mathbf{S}}^{\perp} &::= A^{\perp} \mid (\neg C^{\top}) \mid (C \sqcap C^{\perp}) \mid (\exists R^{\perp}.C) \\ &\quad \mid (\exists R.C^{\perp}) \mid (\geq n R^{\perp}.C) \mid (\geq n R.C^{\perp}). \\ \mathcal{C}_{\mathbf{S}}^{\top} &::= (\neg C^{\perp}) \mid (C_1^{\top} \sqcap C_2^{\top}). \end{aligned}$$

where $A^{\perp} \notin \mathbf{S}$ is an atomic concept, R is a role, and C is a concept, $C^{\perp} \in \mathcal{C}_{\mathbf{S}}^{\perp}, C_i^{\top} \in \mathcal{C}_{\mathbf{S}}^{\top}, i = 1, 2$, and $R^{\perp} \notin \text{Rol}(\mathbf{S})$ is a role.

An axiom α is *syntactically local w.r.t. \mathbf{S}* if it is of one of the following forms: (1) $R^{\perp} \sqsubseteq R$, or (2) $\text{Trans}(R^{\perp})$, or (3) $C^{\perp} \sqsubseteq C$ or (4) $C \sqsubseteq C^{\top}$. We denote by $\text{s.local}(\mathbf{S})$ the set of all SHOIQ-axioms that are syntactically local w.r.t. \mathbf{S} . A SHOIQ-ontology \mathcal{O} is *syntactically local w.r.t. \mathbf{S}* if $\mathcal{O} \subseteq \text{s.local}(\mathbf{S})$.

Intuitively, every concept in $\mathcal{C}_{\mathbf{S}}^{\perp}$ becomes equivalent to \perp if we replace every symbol A^{\perp} or R^{\perp} not in \mathbf{S} with the bottom concept \perp and the empty role respectively, which are both interpreted as the empty set under every interpretation. Similarly, the concepts from $\mathcal{C}_{\mathbf{S}}^{\top}$ are equivalent to \top under this replacement. Syntactically local axioms become tautologies after these replacements.

For example, the axiom M2 from Figure 1 is local w.r.t. $\mathbf{S} = \{\text{Fibrosis, has-Origin}\}$: if we replace the remaining symbols in this axiom with \perp , we obtain a tautology $\perp \equiv \perp$:

$$\underbrace{\perp}_{\text{Genetic.Fibrosis}} \equiv \text{Fibrosis} \sqcap \underbrace{\exists \text{has-Origin. Genetic-Origin}}_{\perp}$$

Syntactic locality is an approximation for (semantic) locality:

PROPOSITION 5. *Let \mathbf{S} be a signature. Then $\text{s.local}(\mathbf{S}) \subseteq \text{local}(\mathbf{S})$.*

PROOF. Let α be an axiom that is syntactically local w.r.t. \mathbf{S} and let $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$ be a trivial expansion of some \mathbf{S} -interpretation to $\mathbf{S} \cup \text{Sig}(\alpha)$. We have to demonstrate that \mathcal{I} is a model of α . By induction over the definitions of $\mathcal{C}_{\mathbf{S}}^{\perp}$ and $\mathcal{C}_{\mathbf{S}}^{\top}$ from Definition 6, it is easy to show that: (i) every role $R \notin \text{Rol}(\mathbf{S})$ and every every concept from $\mathcal{C}_{\mathbf{S}}^{\perp}$ is interpreted in \mathcal{I} with the empty set, and (ii) every concept from $\mathcal{C}_{\mathbf{S}}^{\top}$ is interpreted in \mathcal{I} with Δ . By checking all the possible cases for a syntactically local axiom α in Definition 5, it is easy to see that in every of these cases \mathcal{I} is a model of α . \square

²See <http://www.cs.man.ac.uk/~sattler/reasoners.html> for a list of currently available reasoners.

Algorithm 1 $\text{extract_module}(\mathcal{Q}, \mathbf{S})$ **Input:**

\mathcal{Q} : ontology
 \mathbf{S} : signature

Output:

\mathcal{Q}_1 : a locality-based \mathbf{S} -module in \mathcal{Q}

```

1:  $\mathcal{Q}_1 \leftarrow \emptyset$   $\mathcal{Q}_2 \leftarrow \mathcal{Q}$ 
2: while not empty( $\mathcal{Q}_2$ ) do
3:    $\alpha \leftarrow \text{select\_axiom}(\mathcal{Q}_2)$ 
4:   if locality_test( $\alpha, \mathbf{S} \cup \text{Sig}(\mathcal{Q}_1)$ ) then
5:      $\mathcal{Q}_2 \leftarrow \mathcal{Q}_2 \setminus \{\alpha\}$  ▷  $\alpha$  is processed
6:   else
7:      $\mathcal{Q}_1 \leftarrow \mathcal{Q}_1 \cup \{\alpha\}$  ▷ move  $\alpha$  into  $\mathcal{Q}_1$ 
8:      $\mathcal{Q}_2 \leftarrow \mathcal{Q} \setminus \mathcal{Q}_1$  ▷ reset  $\mathcal{Q}_2$  to the complement of  $\mathcal{Q}_1$ 
9:   end if
10: end while
11: return  $\mathcal{Q}_1$ 

```




Figure 2: Summary for the main theoretical results of the paper

The converse of Proposition 5 does not hold in general since there are semantically local axioms that are not syntactically local. For example, the axiom $\alpha = (A \sqsubseteq A \sqcup B)$ is a tautology and thus is local w.r.t. every \mathbf{S} . This axiom, however, is not syntactically local w.r.t. $\mathbf{S} = \{A, B\}$ since it involves symbols in \mathbf{S} only. Another example, which is not a tautology, is the GCI $\alpha = (\exists R. \neg A \sqsubseteq \exists R. \neg B)$, which is semantically local w.r.t. $\mathbf{S} = \{R\}$ ($\exists R. \top \sqsubseteq \exists R. \top$ is a tautology), but not syntactically local. Thus, the limitation of syntactic locality is its inability to perform reasoning elements from \mathbf{S} .

We distinguish the notion of modules based on these two locality conditions as *semantic locality-based modules* and *syntactic locality-based modules*.

COROLLARY 4. *If \mathcal{Q}_1 is a syntactic locality-based \mathbf{S} -module in \mathcal{Q} , then \mathcal{Q}_1 is a semantic locality-based \mathbf{S} -module in \mathcal{Q} .*

For the reference and for the convenience of the reader, we illustrate in Figure 2 the relationships between the key theoretical results of this paper.

Recall that, according to Definition 5, in order to construct a locality-based \mathbf{S} -module in an ontology \mathcal{Q} , it suffices to partition the ontology \mathcal{Q} as $\mathcal{Q} = \mathcal{Q}_1 \cup \mathcal{Q}_2$ such that \mathcal{Q}_2 is local w.r.t. $\mathbf{S} \cup \text{Sig}(\mathcal{Q}_1)$. Algorithm 1 outlines a simple procedure which performs this task. Assuming there is an effective locality test $\text{locality_test}(\alpha, \mathbf{S})$ (either using a reasoner or the syntactical approximation) that returns true only for axioms α that are local w.r.t. \mathbf{S} , the algorithm first initializes the partition to the trivial one: $\mathcal{Q}_1 =$

#	\mathcal{Q}_1	\mathcal{Q}_2	New elements in $\mathbf{S} \cup \text{Sig}(\mathcal{Q}_1)$	α	loc.?
1	\emptyset	M1–M5	Cystic_Fibrosis, Genetic_Disorder	M1	No
2	M1	M2–M5	Fibrosis, located_In, Pancreas, has_Origin, Genetic_Origin	M2	No
3	M1, M2	M3–M5	Genetic_Fibrosis	M3	No
4	M1–M3	M4, M5	–	M4	No
5	M1–M4	M5	–	M5	Yes
6	M1–M4	–	–	–	–

Table 1: A trace of Algorithm 1 for $\mathcal{Q} = \{M1, \dots, M5\}$ and $\mathbf{S} = \{\text{Cystic_Fibrosis}, \text{Genetic_Disorder}\}$

\emptyset and $\mathcal{Q}_2 = \mathcal{Q}$, and then repeatedly moves to \mathcal{Q}_1 those axioms from \mathcal{Q}_2 that are not local w.r.t. $\mathbf{S} \cup \text{Sig}(\mathcal{Q}_1)$ until no such axioms are left in \mathcal{Q}_2 .

In Table 1 we provide a trace of Algorithm 1 for the input $(\mathcal{Q}, \mathbf{S})$, where \mathcal{Q} consists of the axioms M1–M5 from Figure 1 and $\mathbf{S} = \{\text{Cystic_Fibrosis}, \text{Genetic_Disorder}\}$. Each row in the table corresponds to an iteration of the while loop in Algorithm 1. The last column of the table provides the result of the locality test in line 4. Note that the syntactic locality condition was sufficient in all tests: all axioms that were semantically non-local were also syntactically non-local.

PROPOSITION 6 (CORRECTNESS OF ALGORITHM 1).

For every input \mathcal{Q} and \mathbf{S} , Algorithm 1 computes a locality-based \mathbf{S} -module in \mathcal{Q} .

PROOF. We have to show that (1) Algorithm 1 terminates for every input \mathcal{T} and \mathbf{S} , and (2) the output $\text{extract_module}(\mathbf{S}, \mathcal{Q})$ is a locality-based \mathbf{S} -module in \mathcal{Q} .

(1) Termination of the algorithm follows from the fact that in every iteration of the while loop either the size of \mathcal{Q}_1 increases, or the size of \mathcal{Q}_1 remains the same but the size of \mathcal{Q}_2 decreases. Note that this means that Algorithm 1 terminates in quadratic time in the number of axioms in \mathcal{Q} , assuming constant time locality test.

(2) It is easy to observe that every axiom α that is neither in \mathcal{Q}_1 nor in \mathcal{Q}_2 is local w.r.t. $\mathbf{S} \cup \text{Sig}(\mathcal{Q}_1)$, since the only way such an α can appear is at the line 3 of the algorithm, and α remains in $\mathcal{Q} \setminus (\mathcal{Q}_1 \cup \mathcal{Q}_2)$ only if $\mathbf{S} \cup \text{Sig}(\mathcal{Q}_1)$ does not change. \square

Note that there is an implicit non-determinism in Algorithm 1, namely, in line 3 in which an axiom from \mathcal{Q}_2 is selected. It might well be the case that several choices for α are possible at this moment. For example, in Table 1 at step 2 we might have selected axiom M3 instead of M2. It is possible to show (see [3] for detail) that the output of Algorithm 1 is uniquely determined by its input \mathcal{Q} and \mathbf{S} , and, moreover, is a subset of every locality-based module:

PROPOSITION 7. *The output of Algorithm 1 for \mathcal{Q} and \mathbf{S} is the smallest (syntactic) locality-based \mathbf{S} -module in \mathcal{Q} .*

4.3 Properties of Locality-based Modules

In this section, we outline some interesting properties of locality-based modules which make it possible to use them for applications other than knowledge reuse.

Let \mathcal{Q}_S^{loc} be the smallest locality-based \mathbf{S} -module in \mathcal{Q} , which is unique by Proposition 6 and is the output of Algorithm 1 for \mathcal{Q} and \mathbf{S} . The first property is a consequence of Proposition 6:

PROPOSITION 8. *\mathcal{Q}_S^{loc} contains all \mathbf{S} -essential axioms in \mathcal{Q} w.r.t. every logic L with Tarski-style set-theoretic semantics.*

Ontology	Language	# Atomic Concepts	A1: Prompt-Factor [11]		A2: Modularization from [6]		A3: Locality-based mod.	
			Max. Size (%)	Avg. Size (%)	Max. Size (%)	Avg. Size (%)	Max. Size(%)	Avg. Size(%)
NCI	\mathcal{EL}	27772	24342 (87.6)	21045 (75.8)	15254 (55)	8565 (30.8)	226 (0.8)	22 (0.08)
SNOMED	\mathcal{EL}	255318	255318 (100)	255318 (100)	255318 (100)	255318 (100)	136 (0.5)	12.8 (0.05)
GO	\mathcal{EL}	22357	226 (1)	22 (0.1)	226 (1)	22 (0.1)	92 (0.4)	13 (0.05)
SUMO	\mathcal{EL}	869	869 (100)	869 (100)	869 (100)	869 (100)	18 (2)	8 (0.09)
GALEN-Small	\mathcal{SHF}	2749	2748 (100)	2748 (100)	2748 (100)	2748 (100)	297 (10)	47.7 (1.7)
GALEN-Full	\mathcal{SHIF}	24089	24089 (100)	24089 (100)	24089 (100)	24089 (100)	7379 (29.8)	865.5 (3.5)
SWEET	\mathcal{SHOIF}	1816	1750 (96.4)	1610 (88.7)	1512 (83.3)	935 (51.5)	34 (1.9)	1.7 (0.1)
DOLCE-Lite	\mathcal{SHOIN}	499	498 (100)	497.9 (100)	498 (100)	497.9 (100)	186 (37.3)	123.4 (24.6)

Table 2: Comparison of Different Modularization Algorithms

As shown in Table 1, the minimal locality-based \mathbf{S} -module extracted from \mathcal{Q} contains all \mathbf{S} -essential axioms M1–M4. In our case, the module contains only essential axioms; in general, however, locality-based modules might contain non-essential axioms; otherwise, they would provide a solution for our task T3 in (4).

PROPOSITION 9. *Let \mathcal{Q} be ontology, A and B atomic concepts and $\mathbf{S}_{(i)}$ a signature. Then:*

1. $\mathbf{S}_1 \subseteq \mathbf{S}_2$ implies $\mathcal{Q}_{\mathbf{S}_1}^{loc} \subseteq \mathcal{Q}_{\mathbf{S}_2}^{loc}$ (monotonicity);
2. $\mathcal{Q} \models (A \sqsubseteq B)$ iff $\mathcal{Q}_{\{A\}}^{loc} \models (A \sqsubseteq B)$.

Proposition 9 (see [3] for a proof) gives two interesting properties of locality-based modules. The first one states that such modules may only grow if the input signature extends. The second one implies that the module for a single atomic concept A provides complete information about all the super-classes of A . This property can be used for optimizing classification: in order to *classify an ontology* \mathcal{Q} , i.e. to compute all *subsumption relations* $A \sqsubseteq B$ between pairs A, B of atomic concepts in \mathcal{Q} , it is sufficient to (1) extract all modules $\mathcal{Q}_{\{A\}}^{loc}$ of \mathcal{Q} for each atomic concept A (2) classify each of these modules *independently* (possibly *in parallel*), and (3) merge the results of the individual classifications. By Property 2, if the subsumption $A \sqsubseteq B$ is implied by the ontology \mathcal{Q} then it is implied by the module $\mathcal{Q}_{\{A\}}^{loc}$ and, hence, it will be obtained in step (2).

5. RELATED WORK

The problem of extracting modular fragments of ontologies has recently been addressed in [16], [11] and [15].

In [16], the authors have proposed an algorithm for partitioning the concepts in an ontology. The intended application is to facilitate the visualization of and navigation through the ontology. The algorithm uses a set of heuristics for measuring the degree of dependency between the concepts in the ontology and outputs a graphical representation of these dependencies. The algorithm is intended as a visualization technique, and does not establish a correspondence between the nodes of the graph and sets of axioms in the ontology.

The algorithms in [11] and [15], which we have briefly outlined in Section 3, use structural traversal to extract modules of ontologies for a given signature. None of these approaches provides a characterization of the logical properties of the extracted modules, nor do they establish a notion of correctness of the modularization.

In [6], the authors propose a definition of a module and an algorithm for extracting modules based on that definition. The notion

of a module in an ontology \mathcal{Q} for a signature \mathbf{S} is also based on conservative extensions: if $\mathcal{Q}_1 \subseteq \mathcal{Q}$ is an \mathbf{S} -module in \mathcal{Q} as in [6], then it can be shown that \mathcal{Q} is a model \mathbf{S} -conservative extension of \mathcal{Q}_1 . The definition in [6], however, makes use of additional requirements which lead, in many cases, to the extraction of modules which are larger than one may wish. The reason is that, for every atomic concept $A \in \mathbf{S}$, the module \mathcal{Q}_1 for A in \mathcal{Q} must be a module for all its sub-classes and super-classes.

It is worth pointing out that, given \mathcal{Q} and \mathbf{S} , the fragment obtained using the algorithm in [6] is an \mathbf{S} -module according to Definition 1. This is not the case, however, for the fragment extracted using [15], as we have illustrated in Section 3.

6. IMPLEMENTATION AND EVALUATION

Given an input ontology and an input signature, locality-based modules are not the only possible modules we can obtain. It remains to be shown that the locality-based modules obtained in realistic ontologies are *small enough* to be useful in practice.

For evaluation and comparison, we have implemented the following algorithms using Manchester's OWL API:³

A1: The PROMPT-FACTOR algorithm, as described in [11];

A2: The algorithm for extracting modules described in [6];

A3: Our algorithm for extracting modules (Algorithm 1), based on syntactic locality.

As a test suite, we have collected a set of well-known ontologies available on the Web, which can be divided into two groups:

Simple. In this group, we have included the National Cancer Institute (NCI) Ontology,⁴ the SUMO Upper Ontology,⁵ the Gene Ontology (GO),⁶ and the SNOMED Ontology⁷. These ontologies use a simple ontology language and are of a simple structure; in particular, they do not contain GCIs, but only definitions.

Complex. This group contains the well-known GALEN ontology (GALEN-Full),⁸ the DOLCE upper ontology (DOLCE-Lite),⁹ and

³<http://sourceforge.net/projects/owlapi>

⁴<http://www.mindswap.org/2003/CancerOntology/nciOncology.owl>

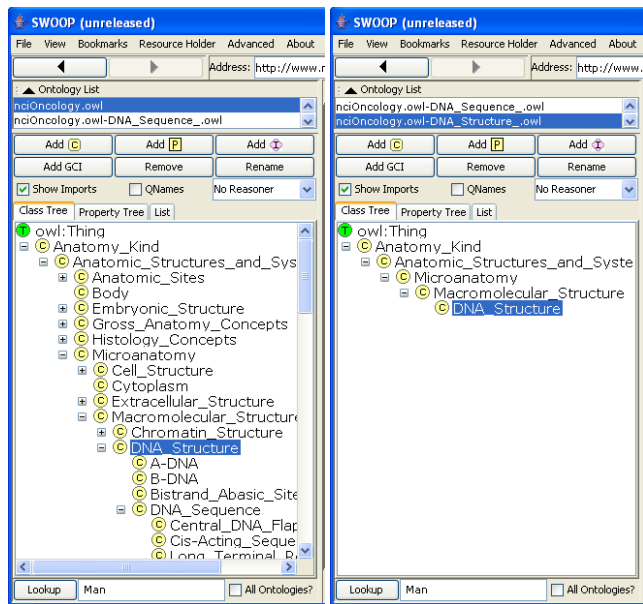
⁵<http://ontology.teknowledge.com/>

⁶<http://www.geneontology.org>

⁷<http://www.snomed.org>

⁸<http://www.openclinical.org/prj-galen.html>

⁹<http://www.loa-cnr.it/DOLCE.html>



(a) DNA_Structure in NCI (b) Synt. Locality-based Module for DNA_Structure in NCI

Figure 4: The Module Extraction Functionality in Swoop

ules and the Y-axis the number of modules whose size is within the given range. The plots thus give an idea of the distribution for the sizes of the different modules.

For SNOMED, NCI and GALEN-Small, we can observe that the size of the modules follows a smooth distribution. In contrast, for GALEN-Full, we have obtained a large number of small modules and a significant number of very big ones, but no medium-sized modules in-between. This abrupt distribution indicates the presence of a big cycle of dependencies in the ontology, which involves all the concepts with large modules. The presence of this cycle can be spotted more clearly in Figure 3(f); the figure shows that there is a large number of modules of size in between 6515 and 6535 concepts. This cycle does not occur in the simplified version of GALEN and thus we have the smooth distribution for that case. In contrast, in Figure 3(e) we can see that the distribution for the “small” modules in GALEN-Full is smooth and much more similar to the one for the simplified version of GALEN.

In order to explore the use of our results for ontology design and analysis, we have integrated our algorithm for extracting modules in the ontology editor SWOOP [9]. The user interface of SWOOP allows for the selection of an input signature and the retrieval of the corresponding module.

As an illustration, consider in Figure 4 the locality-based module for the atomic concept DNA_Structure in the NCI ontology, as obtained in SWOOP. Recall that, according to Case 2 of Proposition 9, the locality-based module $\mathcal{O}_{\{A\}}^{loc}$ for every atomic concept $A \in \text{Sig}(\mathcal{O})$ contains all necessary axioms for, at least, all the (entailed) super-concepts of A in \mathcal{O} . Thus $\mathcal{O}_{\{A\}}^{loc}$ can be seen as the “upper ontology” for A . In fact, Figure 4 shows that the locality-based module for DNA_Structure contains only the concepts in the “path” from DNA_Structure to the top level concept Anatomy_Kind. This suggests that the knowledge in NCI about the particular concept DNA_Structure is very shallow in the sense that NCI only “knows” that a DNA_Structure is a macromolecular structure, etc. which, in the end, is an anatomic structure.

7. CONCLUSION

In this paper, we have proposed a definition of a module for a given vocabulary within an ontology to be reused. Based on this definition, we have formulated three reasoning problems concerning the extraction of minimal modules and shown that none of them is algorithmically solvable, even for simple fragments of OWL DL. We have introduced locality-based modules as an approximation to minimal modules and have empirically demonstrated that such modules are reasonably small for many real-world ontologies.

For the future work, we would like to study other approximations which can produce small modules in complex ontologies like GALEN, and exploit modules for optimizing ontology reasoning.

8. REFERENCES

- [1] F. Baader, S. Brandt, and C. Lutz. Pushing the \mathcal{EL} envelope. In *Proc. IJCAI-2005*, pages 364–370, 2005.
- [2] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.
- [3] B. Cuenca Grau, I. Horrocks, Y. Kazakov, and U. Sattler. Extracting modules from ontologies: Theory and practice. Technical report, University of Manchester, 2007. Available from: <http://www.cs.man.ac.uk/~bcg/Publications.html>.
- [4] B. Cuenca Grau, I. Horrocks, Y. Kazakov, and U. Sattler. A logical framework for modularity of ontologies. In *Proc. IJCAI-2007*, pages 298–304, 2007.
- [5] B. Cuenca Grau, I. Horrocks, O. Kutz, and U. Sattler. Will my Ontologies Fit Together? In *Proc. DL-2006*, 2006.
- [6] B. Cuenca Grau, B. Parsia, E. Sirin, and A. Kalyanpur. Modularity and Web Ontologies. In *Proc. KR-2006*, pages 198–209, 2006.
- [7] S. Ghilardi, C. Lutz, and F. Wolter. Did I Damage my Ontology? A Case for Conservative Extensions in Description Logics. In *Proc. KR-2006*, pages 187–197, 2006.
- [8] I. Horrocks, P. F. Patel-Schneider, and F. van Harmelen. From *SHIQ* and RDF to OWL: The making of a web ontology language. *J. of Web Semantics*, 1(1):7–26, 2003.
- [9] A. Kalyanpur, B. Parsia, E. Sirin, B. Cuenca Grau, and J. Hendler. SWOOP: A web editing browser. *Elsevier’s Journal Of Web Semantics*, 4(2):144–153, 2006.
- [10] C. Lutz, D. Walther, and F. Wolter. Conservative extensions in expressive description logics. In *Proc. of IJCAI-2007*, pages 453–459, 2007.
- [11] N. Noy and M. Musen. The PROMPT suite: Interactive tools for ontology mapping and merging. *Int. Journal of Human-Computer Studies*, 6(59), 2003.
- [12] P. Patel-Schneider, P. Hayes, and I. Horrocks. Web ontology language OWL Abstract Syntax and Semantics. *W3C Recommendation*, 2004.
- [13] A. Rector and J. Rogers. Ontological issues in using a description logic to represent medical concepts: Experience from GALEN. In *Proc. of IMIA WG6 Workshop*, 1999.
- [14] M. Schmidt-Schauß and G. Smolka. Attributive concept descriptions with complements. *Artif. Intell.*, 48(1):1–26, 1991.
- [15] J. Seidenberg and A. Rector. Web ontology segmentation: Analysis, classification and use. In *Proc. WWW-2006*, 2006.
- [16] H. Stuckenschmidt and M. Klein. Structure-based partitioning of large class hierarchies. In *Proc. ISWC-2004*, pages 289–303, 2004.