

Empirical Study of Logic-Based Modules: Cheap Is Cheerful

Chiara Del Vescovo¹, Pavel Klinov², Bijan Parsia¹,
Ulrike Sattler¹, Thomas Schneider³, and Dmitry Tsarkov¹

¹ University of Manchester, UK
{delvescc|bparsia|sattler|tsarkov}@cs.man.ac.uk

² University of Ulm, Germany
pavel.klinov@uni-ulm.de

³ Universität Bremen, Germany
tschneider@informatik.uni-bremen.de

Abstract. For ontology reuse and integration, a number of approaches have been devised that aim at identifying modules, i.e., suitably small sets of “relevant” axioms from ontologies. Here we consider three logically sound notions of modules: MEX modules, only applicable to inexpressive ontologies; modules based on semantic locality, a sound approximation of the first; and modules based on syntactic locality, a sound approximation of the second (and thus the first), widely used since these modules can be extracted from *SRIQ* ontologies in time polynomial in the size of the ontology.

In this paper we investigate the quality of both approximations over a large corpus of ontologies. In particular, we show with statistical significance that, in most cases, there is no difference between the two module notions based on locality; where they differ, the additional axioms are in general unproblematic since either they can be easily ruled out or their number is relatively small. Finally, we show that the same can be said about the relation between MEX and locality-based modules.

1 Introduction

Some notable examples of ontologies describe large and loosely connected domains, as it is the case for SNOMED-CT, the Systematized Nomenclature Of MEDicine, Clinical Terms,⁴ which describes the terminology used in medicine including diseases, drugs, etc. Users often are not interested in a whole ontology \mathcal{O} , rather only in a limited relevant part of it. In this context, the idea has been recently explored to use *modules*, i.e., suitably small subsets of ontologies that behave for specific purposes as the original ontologies over a given *signature* Σ , i.e., a set of terms (non-logical symbols – concept and role names). The notion of *logical module* [9,4] focuses on providing *coverage*, i.e., on preserving all the entailments of \mathcal{O} over Σ .

In [19] the authors comment on the crucial role played by coverage and by two additional properties of modules for ontology reuse and integration. Let \mathcal{M}

⁴ <http://www.ihtsdo.org/snomed-ct/>

be a subset of \mathcal{O} . We say that: (1) \mathcal{M} is *self-contained* if it provides coverage for its signature; (2) \mathcal{M} is *depleting* if the remainder $\mathcal{O} \setminus \mathcal{M}$ of the ontology does not entail any non-tautological axiom η over Σ . Under some mild conditions a minimal depleting and self-contained module is also uniquely determined [15].

Extracting the uniquely determined module for a signature Σ is, however, hard or even impossible for expressive languages [18,10,17]. For identifying notions of modules whose extraction is feasible, we can either restrict the expressivity of the ontology language, or look for feasible sufficient conditions that guarantee \mathcal{M} to be a module for Σ , even though not necessarily the smallest.

For inexpressive logics, one can make use of the module extractor implemented in the MEX system [13]. This module extractor works on acyclic \mathcal{ELI} terminologies and extracts the minimal module in polynomial time.

For expressive logics, module extraction can be implemented making use of the notion of *locality*. The resulting modules, known as locality-based modules (LBMs) [2] are, in general, not minimal. Locality comes in two flavors: semantic and syntactic locality, which have a bottom-, a top-, and a star-variant which is contained in both the top- and the bottom-variants. For any of the three variants, a syntactic LBM contains the corresponding semantic LBM. Algorithms for the extraction of syntactic LBMs are known that run in time polynomial in the size of the ontology (thus much cheaper than reasoning), are implemented in the OWL API,⁵ and are currently used for ontology reuse and integration. In contrast, despite the fact that algorithms for extracting *semantic* LBMs are known, until now and to the best of our knowledge they had not been implemented. They require entailment checks against an empty ontology and thus involve reasoning of a kind that is rather unusual for DL reasoners.⁶

We know that the MEX module for a signature Σ is contained in the star semantic LBM which, in turn, is contained in the star syntactic LBM. Thus, syntactic locality can be seen as an approximation of semantic locality which, in turn, is an approximation of MEX modules. An interesting question arising here is how *good* these approximations are: how much larger are the modules extracted by the approximations, and how much faster is the extraction?

To answer these questions, we present the first implementation of a semantic LBM extractor and report on experiments on a large corpus of ontologies. We compare the performance and results of the semantic LBM extractor with those of a syntactic LBM extractor and with those of MEX. A comparison between MEX- and syntactic bottom-modules only for SNOMED-CT is reported in [13].

The contributions of this paper are as follows.

- We provide the first implementation for the extraction of semantic LBMs, which is embedded in the latest release of the FACT++ reasoner. Our results show that the extraction of semantic LBMs, which is in principle hard, is feasible in practice: on average, it is between 3 times (for top-modules) and 15 times (for bottom- and star-modules) slower than the extraction of syntactic LBMs, and both only take milliseconds to seconds for most ontologies below 10K axioms.

⁵ <http://owlapi.sourceforge.net/>

⁶ DL reasoners usually *classify* an ontology: test it for consistency and all concept names for satisfiability/mutual subsumption.

- We show with statistical significance that, for almost all members of a large corpus of existing ontologies, there is no difference between any syntactic LBM and its semantic counterpart. In the few cases where differences occur, those are extremely modest so that it is questionable whether extracting semantic LBMs is worth the increased computational cost.
- We isolate four patterns of axioms that completely explain those rare differences. One includes simple tautologies that can be removed in a straightforward preprocessing step.
- We modify the original corpus to obtain for each ontology an acyclic \mathcal{EL} version suitable for the use with the MEX system. We then compare MEX-modules and the star-variants of LBMs, and find differences in only $\sim 27\%$ of the corpus. We explain one reason for the largest differences observed.

2 Preliminaries

We assume the reader to be familiar with Description Logic languages (e.g. \mathcal{SROIQ} [1,12]), and aim here at fixing the notations and at defining the key notions around module extraction, with a focus on locality-based modules [2] and MEX modules [14].

Let \mathcal{O} denote a \mathcal{SROIQ} ontology, \mathbf{N}_C a set of concept names, and \mathbf{N}_R a set of role names. We do not consider individual names as they do not play any role in the extraction of the modules here analyzed. We refer to either a concept or a role name by using the word “*term*”. A *signature* is a set $\Sigma \subseteq \mathbf{N}_C \cup \mathbf{N}_R$. Given a concept, role, or axiom X , we call the set of terms in X the *signature of X* , denoted \tilde{X} . Given a set $\mathcal{M} \subseteq \mathcal{O}$ of axioms from \mathcal{O} , and a signature Σ , we say that \mathcal{O} is a *deductive Σ -conservative extension (Σ -dCE)* of \mathcal{M} if, for all \mathcal{SROIQ} -axioms α with $\tilde{\alpha} \subseteq \Sigma$, it holds that $\mathcal{O} \models \alpha$ if and only if $\mathcal{M} \models \alpha$. \mathcal{O} is a *model Σ -conservative extension (Σ -mCE)* of \mathcal{M} if $\{\mathcal{I}|_{\Sigma} \mid \mathcal{I} \models \mathcal{O}\} = \{\mathcal{I}|_{\Sigma} \mid \mathcal{I} \models \mathcal{M}\}$. Dually, \mathcal{M} is a *dCE-based module* of \mathcal{O} for Σ if \mathcal{O} is a Σ -dCE of \mathcal{M} , and it is an *mCE-based module for Σ* if \mathcal{O} is a Σ -mCE of \mathcal{M} . All dCE-based modules are also mCE-based modules, whilst the converse is not always true. A module $\mathcal{M} \subseteq \mathcal{O}$ for Σ is called *depleting* if there is no non trivial entailment η over Σ such that $\mathcal{O} \setminus \mathcal{M} \models \eta$; \mathcal{M} is called *self-contained* if \mathcal{M} is a module for $\Sigma = \tilde{\mathcal{M}}$.

Since $\mathcal{M} \subseteq \mathcal{O}$ the monotonicity of \mathcal{SROIQ} implies that every entailment η over Σ derivable from \mathcal{M} is also derivable from \mathcal{O} . Deciding the converse directions is in general computationally hard, or even undecidable for expressive DLs [18,10,17]. Since we do not need to find *all* the subsets of \mathcal{O} that are a module for Σ , we can use easier conditions which guarantee that a set of axioms $\mathcal{M} \subseteq \mathcal{O}$ is a module for Σ .

One strategy to extract a module $\mathcal{M} \subseteq \mathcal{O}$ for a *seed signature* Σ consists of defining a suitable oracle x to decide whether a single axiom $\alpha \in \mathcal{O}$ is *x -relevant* for preserving the non-tautological entailments over Σ : clearly, if the x -check is positive for α , then by depletion \mathcal{M} needs to contain α . To guarantee also self-containment, the fixpoint procedure described in Algorithm 1 needs to be performed. The module extracted will be called an *x -module*, parameterized by the notion x of relevance. Algorithm 1 is a special case of the one in [2,

Algorithm 1 Extraction of an x -module for Σ

Input: Ontology \mathcal{O} , seed signature Σ , notion of relevance x

Output: x -module \mathcal{M} of \mathcal{O} w.r.t. Σ

$\mathcal{M} \leftarrow \emptyset$; $\mathcal{O}' \leftarrow \mathcal{O}$

repeat

 changed \leftarrow **false**

for all $\alpha \in \mathcal{O}'$ **do**

if α is x -relevant w.r.t. $\Sigma \cup \widetilde{\mathcal{M}}$ **then**

$\mathcal{M} \leftarrow \mathcal{M} \cup \{\alpha\}$; $\mathcal{O}' \leftarrow \mathcal{O}' \setminus \{\alpha\}$; changed \leftarrow **true**

until changed = **false**

return \mathcal{M}

Figure 4], and its output \mathcal{M} does not depend on the order in which the axioms α are selected [2].

Due to space limitations, we cannot provide a full description of the notions of modules here investigated, and refer the interested reader to [2,14] for further details. We briefly sketch the intuition behind each notion of relevance and the corresponding results of interest for this paper.

The MEX system. In [14], a procedure called MEX for extracting minimal mCE-based modules from acyclic \mathcal{ELI} ontologies is given. The notion of MEX-relevance is based on a relation $\text{depend}_{\mathcal{O}}$, which associates each concept name A with the set $\text{depend}_{\mathcal{O}}(A)$ of all the symbols X that are used in the definition⁷ of A in \mathcal{O} . Then, an axiom $\alpha \in \mathcal{O}$ is MEX-relevant for a signature Σ if it defines a term in Σ . The authors prove that, if \mathcal{O} is an acyclic \mathcal{ELI} ontology, then using the MEX-relevance in Algorithm 1 generates the minimal depleting self-contained module for each signature Σ , and that the extraction runs in polynomial time.

Semantic locality In [2], the authors define the notion of locality and distinguish two flavors, here called \emptyset - and Δ -locality. Intuitively, a $SR\mathcal{OIQ}$ axiom α is \emptyset -local (resp. Δ -local) w.r.t. signature Σ if α' obtained by replacing all terms in $\widetilde{\alpha} \setminus \Sigma$ with \perp (resp. \top) is a tautology. The authors prove that, if all axioms in $\mathcal{O} \setminus \mathcal{M}$ are \emptyset -local (or all axioms are Δ -local) w.r.t. $\Sigma \cup \widetilde{\mathcal{M}}$, then \mathcal{M} is an mCE-based (and hence dCE-based) module of \mathcal{O} for Σ . Since deciding \emptyset - or Δ -locality requires tautology checks, this problem is as hard as standard reasoning. In some cases, α' is not a $SR\mathcal{OIQ}$ axiom, so standard reasoners need to be extended.

Syntactic locality In order to achieve *tractable* module extraction, [2] define the two syntactic notions of \perp - and \top -locality. Those approximate the two notions of semantic \emptyset - and Δ -locality. Intuitively, the syntactic rules provided describe axioms α for which α' is obviously a tautology—in which case α is said to be \perp - or \top -local w.r.t. Σ . Thus, every \emptyset -local (Δ -local, resp.) axiom w.r.t. Σ is also \perp -local (\top -local, resp.) w.r.t. Σ , but not vice versa. As a consequence, also \perp - and \top -modules are mCE- and dCE-based modules for Σ . Applying the syntactic rules requires polynomial time, hence the extraction of this kind of modules can be performed in time polynomial in the size of the ontology.

⁷ This notion of “definition” is specific for the MEX-modules

Modules based on syntactic (semantic) locality can be made smaller by iteratively nesting \top - and \perp -extraction (Δ - and \emptyset -extraction), again obtaining mCE- and dCE-based modules [2,19], called $\top\perp^*$ - and $\Delta\emptyset^*$ -modules. The $\Delta\emptyset^*$ -module for Σ is always contained in the corresponding $\top\perp^*$ -module. Moreover, for acyclic \mathcal{ELI} ontologies, the MEX-module for Σ is always contained in the corresponding $\Delta\emptyset^*$ -module.

3 Research questions and experimental design

A natural question is whether syntactic modules are likely to be much larger than the semantic ones, which are, in theory, computationally more costly. Hence, a second question is whether semantic module extraction is *noticeably* more costly: the tautology test has to be carried out often—once per axiom and signature that the algorithm goes through—and it is thus hard to predict the feasibility of semantic LBM extraction. Altogether, we want to know whether syntactic LBMs are a *necessary* approximation and *how good* an approximation of semantic LBMs they are. Similarly, for acyclic \mathcal{ELI} terminologies the analogous question arises: how good an approximation of MEX modules are LBMs?

One can always construct ontologies with huge differences in size and time between syntactic and semantic LBMs and between LBMs and MEX modules. Here, we are interested in these differences in currently available ontologies, and thus need to design, run, and analyse suitable experiments.

Selection of the corpus. For our experiments, we have built a corpus containing: (1) all the ontologies from the NCBO BioPortal ontology repository,⁸ version of November 2012; (2) ontologies from the TONES repository⁹ which have already been studied in previous work on modularity [6]: Koala, Mereology, University, People, miniTambis, OWL-S, Tambis, Galen. From this corpus, we have removed ontologies that cannot be downloaded, whose .owl file is corrupted or impossible to parse, or which are inconsistent. Furthermore, we have excluded those large ontologies (exceeding 10K axioms) where the extraction of a semantic LBM repeatedly took more than 2 minutes: for each such ontology, the estimated time needed to perform our experiments would have exceeded 300 hours.

This selection results in a corpus of 242 ontologies, which greatly vary in expressivity (from \mathcal{AL} to $\mathcal{SROIQ}(\mathcal{D})$) and in size (10–16,066 axioms, 10–16,068 terms) [11]. For a full list of the corpus, please refer to [5].

As mentioned before, for some ontologies is not possible to test Δ -locality (and thus for extracting Δ - and $\Delta\emptyset^*$ -modules) using standard DL reasoners., see [5] for details. To cover these cases, we have extended the reasoner FACT++ to cover the uses of the \top -role specific to the semantic locality tests.

Since MEX handles only acyclic \mathcal{ELI} ontologies, we created an \mathcal{ELI} version $\mathcal{ELI}(\mathcal{O})$ of each ontology \mathcal{O} in our corpus by filtering unsupported axioms and breaking terminological cycles. While a principled way of doing this is beyond the scope of this paper, we have used a heuristic, which is described in [5].

⁸ <http://bioportal.bioontology.org>

⁹ <http://owl.cs.manchester.ac.uk/repository/>

Comparing modules and locality. In order to compare syntactic and semantic locality, as well as LBMs and MEX modules, we want to understand (1) whether, for a given seed signature Σ , it is likely that there is a difference between the syntactic and the semantic Σ -module or the MEX module for Σ and the latter and, if so, the size of the difference;¹⁰ and (2) how feasible the extraction of semantic LBMs is. For this purpose, we compare (a) \emptyset -semantic and \perp -syntactic locality, Δ -semantic and \top -syntactic locality, (b) \emptyset - and \perp -modules, Δ - and \top -modules, $\Delta\emptyset^*$ - and $\top\perp^*$ -modules, (c) MEX modules and $\Delta\emptyset^*$ -modules.

Due to the recursive nature of Algorithm 1, our investigation is both on a

per-axiom-basis: given axiom α and signature Σ , is it likely that α is semantically \emptyset -local (Δ -local, resp.) w.r.t. Σ but not syntactically \perp -local (\top -local, resp.) w.r.t. Σ ?

per-module basis: given a signature Σ , is it likely that

- $\perp\text{-mod}(\Sigma, \mathcal{O}) \neq \emptyset\text{-mod}(\Sigma, \mathcal{O})$, or
- $\top\text{-mod}(\Sigma, \mathcal{O}) \neq \Delta\text{-mod}(\Sigma, \mathcal{O})$, or
- $\top\perp^*\text{-mod}(\Sigma, \mathcal{O}) \neq \Delta\emptyset^*\text{-mod}(\Sigma, \mathcal{O})$, or
- $\Delta\emptyset^*\text{-mod}(\Sigma, \mathcal{O}) \neq \text{MEX-mod}(\Sigma, \mathcal{O})$?

If yes, is it likely that the difference is large?

Clearly we need to pick, for each ontology in our corpus, a suitable set of signatures, and this poses a significant problem. A full investigation is infeasible: if $m = \#\mathcal{O}$, there are 2^m possible seed signatures, so that testing axioms for locality against *all* the signatures is already impossible for $m \sim 100$. One could assume that comparing modules is easier since many signatures can lead to the same module. However, previous work [6,8] has shown that the number of modules in ontologies is, in general, exponential w.r.t. the size of the ontology. Still, different seed signatures can lead to the same module, which makes it hard to extract enough *different* modules.

We will consider seed signatures of two kinds: genuine seed signatures and random seed signatures.

Genuine seed signatures. A module does not necessarily show an internal coherence: e.g., if we had an ontology \mathcal{O} about the domains of geology and philosophy, we could extract the module for the signature $\Sigma = \{\text{Epistemology}, \text{Mineral}\}$. That module is likely to be the union of the two disjoint modules for $\Sigma_1 = \{\text{Epistemology}\}$ and $\Sigma_2 = \{\text{Mineral}\}$.

In contrast, *genuine modules* can be said to be coherent: they are those modules that cannot be decomposed into the union of two “ \subseteq ”-uncomparable modules. Notably, there are only linearly many genuine modules in the size of \mathcal{O} since each genuine x -module equals $x\text{-mod}(\tilde{\alpha}, \mathcal{O})$, for some axiom $\alpha \in \mathcal{O}$. Moreover, all modules of \mathcal{O} are composed from genuine modules [7]. Thus, genuine modules are of special interest, and we can investigate them, and the corresponding *genuine signatures*, in full.

¹⁰ Recall: the MEX module is always a subset of the semantic Σ -module, which is always a subset of the syntactic Σ -module.

Random seed signatures. Since a full investigation of all the signatures is impossible, we compare locality—both on a per-axiom and per-module basis—as well as LBMs and MEX modules on a *random* signature Σ , which we select by setting each named entity E in the ontology to have probability $p = 1/2$ of being included in Σ . This ensures that each Σ will have the same probability to be chosen. This approach has a clear setback: the random variable “size of the seed signature generated” follows a binomial distribution, so a random seed signature is highly likely to be rather large and to contain half the terms of the ontology. However, we do not yet have enough insight into what *typical* seed signatures are for module extraction, so biasing the selection of signatures to, for example, those of a certain size has no *rationale*. In contrast, selecting random seed signatures avoids the introduction of any bias. Moreover, this choice is complementary to the selection of *all* the genuine signatures, which are in general small.

Whilst the selection of genuine signatures is complete, we can only aim at selecting a number of random signatures to obtain statistically significant statements about modules. To reach a confidence level of 95% that the true proportion of differences between modules lies in the confidence interval ($\pm 5\%$) of the observed proportion, we have to sample at least 385 seed signatures (see the detailed explanations in [5]). For ontologies with at least 9 elements in the signature, we will therefore draw a sample of size 400. For all other ontologies, we will look at *all* of the ≤ 400 signatures.

Summary. We compare, for every ontology \mathcal{O} in our corpus,

- (T1) for random seed signatures Σ from \mathcal{O} ,
 - (a) for each axiom α in \mathcal{O} , is α
 - \emptyset -local w.r.t. Σ but not \perp -local w.r.t. Σ ?
 - Δ -local w.r.t. Σ but not \top -local w.r.t. Σ ?
 - (b) is
 - \perp -mod(Σ, \mathcal{O}) \neq \emptyset -mod(Σ, \mathcal{O})?
 - \top -mod(Σ, \mathcal{O}) \neq Δ -mod(Σ, \mathcal{O})?
 - $\top\perp^*$ -mod(Σ, \mathcal{O}) \neq $\Delta\emptyset^*$ -mod(Σ, \mathcal{O})?
 - $\Delta\emptyset^*$ -mod($\Sigma, \mathcal{ELI}(\mathcal{O})$) \neq MEX-mod($\Sigma, \mathcal{ELI}(\mathcal{O})$)?
- (T2) the same, Σ ranging over *all* the genuine signatures $\tilde{\beta}$ for $\beta \in \mathcal{O}$.

4 Results of the Experiments

No differences in locality. The main result of the experiment is that, for the vast majority of the ontologies in our corpus, no difference between syntactic and semantic locality is observed, for all three variants \perp vs. \emptyset , \top vs. Δ , and $\top\perp^*$ vs. $\Delta\emptyset^*$. More precisely, for 210 out of 242 ontologies, we obtain that:

- (T1) for random seed signatures, there is no statistically significant difference
 - (a) between semantic and syntactic locality of any kind,
 - (b) between semantic and syntactic LBMs of any kind;
- (T2) given *any* genuine signature, there is no such difference.

More specifically, for all randomly generated seed signatures and *all* genuine signatures, the corresponding bottom-modules (and the corresponding top- and star-modules, respectively) agree, and every axiom is either \perp - and \emptyset -local, or none of both (and either \top - and Δ -local, or none of both).

The 210 ontologies include *Galen* and *People*, which are renowned for having unusually large \perp -modules [2,8]. In most cases, extracting a semantic and syntactic LBM each took only a few milliseconds; hence, a performance comparison is not meaningful. For some ontologies, the semantic LBM took considerably longer to extract than the syntactic: up to 5 times for star-modules in *Molecule Role*, and up to 34 times in *Galen*.

Differences in locality. We have observed differences between syntactic and semantic locality for 32 ontologies in our corpus. We call the axioms that cause these differences *culprits* – patterns of axioms which are not \perp -local (\top -local, respectively) w.r.t. some signature Σ , but which are \emptyset -local (Δ -local, respectively) w.r.t. Σ . We have identified four types of patterns, *a–d*, and we describe them in the following. Sometimes, culprit axioms *pull* additional axioms into the syntactic LBM, due to signature extension during module extraction.

We denote concept *names* by *A, B*, complex concepts by *C, D*, roles by *r, s, . . .*, nominals by *a*, non-empty data ranges (e.g., *int* or *int*^{0..9}) by *R*, possibly with indices. Σ denotes a signature for which a module is extracted or against which an axiom is checked for locality. Terms outside Σ are overlined; we further use notation \mathbf{C}^\perp and \mathbf{C}^\top to denote concepts that are bottom- or top-equivalent due to the grammar defining syntactic locality in [3, Def. 6] and the analogous grammar for semantic locality.

Culprits of type *a* are simple tautologies that accidentally entered the “inferred view” (closure under certain entailments) of an ontology. These axioms do not occur in the original “asserted” versions and could, in principle, be detected in a simple preprocessing step. Type-*a* culprits occur in 10 ontologies of the above 32, and are of the following kinds: $A \sqsubseteq A$, $r \equiv (r^-)^-$, and $A \sqcap C \sqcap D \sqsubseteq A \sqcap C$. Each such tautology is trivially \emptyset -local and Δ -local w.r.t. any Σ , but not always \perp - or \top -local: if Σ contains all terms in α , then both sides of the subsumption (equivalence) are neither \perp - nor \top -equivalent.

Differences caused not solely by culprits of type *a* have been observed for 26 ontologies. In only 6 of these cases, the differences affect modules; in the remaining 20, they only affect locality of single axioms (tests T1 a and T2 a). We will focus on the former 6, listed in Table 1, and refer to [5] for details on all 26.

Ontology	Abbreviation	DL expressivity	#axioms	#terms
MiniTambis-repaired	MiniT	<i>ALCN</i>	170	226
Tambis-full	Tambis	<i>SHIN(D)</i>	592	496
Bleeding History P...	BHO	<i>ALCIF(D)</i>	1,925	581
Neuro Behavior O...	NBO	<i>AL</i>	1,314	970
Pharmacogenomic...	PhaRe	<i>ALCIF(D)</i>	459	311
Terminological and...	TOK	<i>SRIQ(D)</i>	466	330

Table 1. Ontologies that exhibit differences in modules

According to Table 1, differences between modules occur for ontologies of medium to large size and medium to high expressivity. Differences in locality alone additionally affect small ontologies such as Koala (42 axioms) and Pilot Ontology (85 axioms), as well as large ontologies such as Galen (4,735 axioms) and Experimental Factor Ontology (7,156 axioms). The number of axioms *causing* these differences (i.e., matching the culprit patterns) in the affected ontologies is small except for Galen, and most of the observed differences are relatively small.

Table 2 gives a representative selection of the differences in *modules* observed. For a complete overview, including differences in locality of single axioms, consult the table in [5].

Ontology	Types affected	#diffs	size of diffs		culprit type	freq.
			#axs	(rel.)		
miniT	bot, star	14-25%	1-7	0-600% ^b	<i>c</i>	3
Tambis	bot, star	32-57%	2-41 ^c	1-62% ^c	<i>c</i>	8
BHO ^a	star	17%	1-12	0-300%	<i>b</i>	31
NBO ^a	star	3%	2	0-200%	<i>d</i>	3
PhaRe ^a	top, star	1-8%	1-326 ^d	0-6,520% ^d	<i>d</i>	10
TOK	top, star	49-100%	1-7	0-9%	<i>d</i>	3

^adifferences only for genuine modules

^bdifferences > 5% only for genuine modules

^cdifferences > 11 axioms (> 2%) only for genuine modules

^ddifferences > 13 axioms (> 1,300%) only for top-modules

The columns show: ontology name (abbreviations: see Table 1); type of modules affected; relative number of module pairs with differences; number of axioms in the differences (absolute and relative to the \emptyset - or Δ - or $\Delta\emptyset^*$ -case); type of culprit present and number of axioms of this type involved in differences.

Table 2. Overview of observed differences between modules

Table 2 shows small absolute differences for miniT, BHO, NBO, and TOK. In Tambis, large differences occur only for genuine modules, which suggests that they are unlikely to occur in practical cases with usually larger seed signatures. Finally, in PhaRe, large differences occur only for top-modules. For all these ontologies, a single syntactic or semantic module was extracted within only a few milliseconds, making module extraction times roughly equal.

Culprits of type *b* are axioms with an \exists -restriction on a set of nominals or a non-empty data range on the right-hand side, such as $A \sqsubseteq \exists \bar{r}. \{a_1, \dots, a_n\}$ or $A \sqsubseteq \exists \bar{r}. R$. These axioms are Δ -local w.r.t. any signature Σ that does not contain r because they become tautologies if r is replaced by \top . However, they can only be \top -local when A is a \perp -equivalent concept w.r.t. Σ .

Culprit-*b* axioms affect genuine modules of BHO, and (only) locality of single axioms for 4 more ontologies. We observed a slightly more sophisticated variant of the form $A \equiv C^\top \sqcap \exists r. R$.

Culprits of type *c* are axioms α that contain a concept description C such that (a) C becomes equivalent to \perp (or \top) if all terms outside Σ are replaced by \perp (or \top); (b) this causes α to be semantically \perp -local (or \top -local); but (c)

the grammars for syntactic locality do not “detect” \mathbf{C} to be a \mathbf{C}^\perp (or \mathbf{C}^\top). For example, $\mathbf{C} = \forall r. \bar{A} \sqcap \exists r. \top$ becomes \perp -equivalent if A is replaced by \perp ; the same holds with cardinality restrictions in place of “ \exists ”. Consequently, axioms such as $A^\perp \equiv B \sqcap \forall r. \mathbf{C}^\perp \sqcap \forall s. \{a\} \sqcap = 3 r. \top$, (taken from *Koala*) are \emptyset -local but not \perp -local.

We found this pattern in 8 ontologies. Only in *miniT* and *Tambis*, it affects a large proportion of bottom- and star-modules, with additional axioms “pulled in”. Still, the size of the differences is modest, as argued above. Some of the remaining 6 ontologies contain different kinds of complex concepts that cause differences in top-locality of single axioms.

Culprits of type d are axioms where a concept (or role) name from the left-hand side occurs on the right-hand side together with a top-equivalent role (or concept), causing differences in top-modules. The simplest kind of axiom of this type is $A \sqsubseteq \exists r. A$, which is Δ -local because replacing r with \top makes it a tautology. The axiom is only \top -local if Σ contains neither r nor A . We have found further examples of increasing complexity in *Adverse Event Reporting Ontology* and *Galen*; see [5].

We have observed culprits of type d in 17 ontologies, see the detailed overview in [5]. Only in 3 cases (*NBO*, *PhaRe*, and *TOK*) are modules affected.

Galen contains 121 culprit- D axioms, but they only affect locality of single axioms. In addition, the time differences for *Galen* are remarkable: checking all axioms for Δ -locality takes up to 70 times longer than checking them for \top -locality.

Summary. All culprits hardly ever cause significant differences in modules. Only for *PhaRe* are differences between semantic and syntactic modules not negligible, but we were able to relativize them.

Table 1 may suggest that culprits occur only in expressive ontologies. However, patterns a , c , d can, in principle, already occur in simple terminologies in \mathcal{EL} and \mathcal{ALC} , respectively. Evidently, type- a culprits can easily be filtered out in a preprocessing step. For types c and d , there is no hope for an exhaustive extension to locality because they can (and do) occur in arbitrarily complex shapes and contexts.

Patterns of type b rely on nominals or datatypes – but they are *repairable* by a straightforward extension to the definition of syntactic locality: one can extend the locality definition to distinguish \perp - and \top -*distinct* concepts, by adding appropriate grammars to the definition of syntactic locality, and adding more cases of \perp - and \top -equivalent concepts to the existing grammars. However, from the small numbers of differences observed, we doubt that such an extension of syntactic locality will have any significant effects in practice.

LBM vs MEX results. The results of the experimental comparison of syntactic/semantic LBMs and MEX modules are summarized in Table 3. They show that MEX modules smaller than the corresponding LBMs can be found in $\sim 27\%$ of the preprocessed ontologies, for either random or axiom-based seed signatures. At the same time, unsurprisingly, syntactic and semantic LBMs do not differ at all for these simple \mathcal{ELI} ontologies.

In experiments with random seed signatures, it can be seen that for those ontologies where there *are* differences (most notably, *Galen*), they occur in many

Experiment	#ontol. with diffs.	% tests with diffs.	avg size of diffs #axs	rel.
Random signatures	66	84%	0–26	0–13%
Axiom signatures	61	12%	0–13	0–80%

The results from the third column on are averaged over all ontologies with differences LBM–MEX in at least one module. For example, the last two columns show the average min and max absolute (resp. relative) difference between LBMs and MEX modules.

Table 3. Differences between MEX and LBMs ($\top\perp^*$, $\Delta\emptyset^*$)

tests. Thus, the difference appears to be caused by features of the ontology, not some particular seed signatures. Also, the difference sometimes comes out large in certain tests, also for genuine modules. For example, for the signature of the following axiom in *Galen*, both $\Delta\emptyset^*$ -mod and $\top\perp^*$ -mod contain 127 axioms while the MEX-module only contains the axiom itself:¹¹ $\text{RICF} \equiv \text{ICF} \sqcap \exists \text{ISFO.RSH}$. The likely reason is the proliferation of concept equivalence axioms in *Galen*. For example $A \equiv B$ will end up in the $\Delta\emptyset^*$ -mod for any seed signature containing either A or B. It is, however, an mCE of \emptyset w.r.t. to either $\{A\}$ or $\{B\}$.

5 Conclusion and outlook

Summary. We obtain three main observations from our experiments. (1) In general, there is no or little difference between semantic and syntactic locality. Hence, the computationally cheaper syntactic locality is a good approximation of semantic locality. For the ontologies *Galen* and *People*, which are “renowned” for having disproportionately large modules, syntactic and semantic LBMs do not differ. (2) In most cases, there is no or little difference between LBMs and MEX modules. Only for *Galen* are MEX modules considerably smaller than LBMs. (3) Though in principle hard to compute, semantic LBMs can be extracted rather fast in practice. Still, their extraction often takes considerably longer than for syntactic LBMs. We cannot make any statement about MEX module extraction times because we use the original MEX implementation, which combines loading and module extraction. Due to (1), hardly any benefit can be expected from preferring the potentially smaller semantic LBMs to the cheaper syntactic LBMs. From (2) we can say that semantic LBMs can be seen as the best available approximation of MEX modules for ontologies in highly expressive languages.

Not only does our study evaluate how good the cheap syntactic locality approximates semantic locality and model conservativity, it also enabled us to fix bugs in the implementation of syntactic modularity.

Future work. Two questions are interesting for future work: (1) How can we redesign the experiments so that we can include the very large ontologies? (2) How do LBMs compare to other types of conservativity-based modules?

Concerning (2), one could include, for example, the technique based on reduction to QBF for the OWL 2 QL profile [16] when an off-the-shelf implementation becomes available.

¹¹ The acronyms denote `RightIneffectiveCardiacFunction`, `IneffectiveCardiacFunction`, `isSpecificFunctionOf`, `RightSideOfHeart`.

References

1. Franz Baader, Deborah Calvanese, Diego McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.
2. Bernardo Cuenca Grau, Ian Horrocks, Yevgeny Kazakov, and Ulrike Sattler. Modular reuse of ontologies: Theory and practice. *J. of Artif. Intell. Research*, 31(1):273–318, 2008.
3. Bernardo Cuenca Grau, Ian Horrocks, Yevgeny Kazakov, and Ulrike Sattler. Extracting modules from ontologies: A logic-based approach. In Stuckenschmidt et al. [20], pages 159–186.
4. Bernardo Cuenca Grau, Bijan Parsia, Evren Sirin, and Aditya Kalyanpur. Modularity and Web ontologies. In *Proc. of KR-06*. AAAI Press/The MIT Press, 2006.
5. Chiara Del Vescovo, Pavel Klinov, Bijan Parsia, Ulrike Sattler, Thomas Schneider, and Dmitry Tsarkov. Empirical study of logic-based modules: Cheap is cheerful. Technical report, 2013. <https://sites.google.com/site/cheapischeerful/>.
6. Chiara Del Vescovo, Bijan Parsia, Ulrike Sattler, and Thomas Schneider. The modular structure of an ontology: an empirical study. volume 573 of *ceur-ws.org*, 2010.
7. Chiara Del Vescovo, Bijan Parsia, Ulrike Sattler, and Thomas Schneider. The modular structure of an ontology: Atomic decomposition. In *Proc. of IJCAI-11*, pages 2232–2237, 2011.
8. Chiara Del Vescovo, Bijan Parsia, Ulrike Sattler, and Thomas Schneider. The modular structure of an ontology: Atomic decomposition and module count. volume 230 of *FAIA*, pages 25–39, 2011.
9. James Garson. Modularity and relevant logic. 30(2):207–223, 1989.
10. Silvio Ghilardi, Carsten Lutz, and Frank Wolter. Did I damage my ontology? A case for conservative extensions in Description Logics. In *Proc. of KR-06*, pages 187–197. AAAI Press/The MIT Press, 2006.
11. Matthew Horridge, Bijan Parsia, and Ulrike Sattler. The state of bio-medical ontologies. 2011.
12. Ian Horrocks, Oliver Kutz, and Ulrike Sattler. The even more irresistible *SR0IQ*. In *Proc. of KR-06*, pages 57–67, 2006.
13. Boris Konev, Carsten Lutz, Dirk Walther, and Frank Wolter. Semantic modularity and module extraction in description logics. In *Proc. of ECAI-08*, pages 55–59, 2008.
14. Boris Konev, Carsten Lutz, Dirk Walther, and Frank Wolter. Formal properties of modularization. In Stuckenschmidt et al. [20], pages 25–66.
15. Roman Kontchakov, Luca Pulina, Ulrike Sattler, Thomas Schneider, Petra Selmer, Frank Wolter, and Michael Zakharyashev. Minimal module extraction from DL-Lite ontologies using QBF solvers. In *Proc. of IJCAI-09*, pages 836–841, 2009.
16. Roman Kontchakov, Frank Wolter, and Michael Zakharyashev. Logic-based ontology comparison and module extraction, with an application to DL-Lite. *Artificial Intelligence*, 174(15):1093–1141, 2010.
17. Carsten Lutz, Dirk Walther, and Frank Wolter. Conservative extensions in expressive Description Logics. In *Proc. of IJCAI-07*, pages 453–458, 2007.
18. Carsten Lutz and Frank Wolter. Deciding inseparability and conservative extensions in the description logic \mathcal{EL} . 45(2):194–228, 2010.
19. Ulrike Sattler, Thomas Schneider, and Michael Zakharyashev. Which kind of module should I extract? volume 477 of *ceur-ws.org*, 2009.
20. Heiner Stuckenschmidt, Christine Parent, and Stefano Spaccapietra, editors. volume 5445 of *LNCS*. Springer-Verlag, 2009.