

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 OWL DL 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, using our own implementation of semantic locality, which is the first to our knowledge. 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 can either be easily ruled out or their number is relatively small. We classify the axioms that explain the rare differences into four kinds of “culprits” and discuss which of those can be avoided by extending the definition of syntactic locality. Finally, we show that differences between **MEX** and locality-based modules occur for a minority of ontologies from our corpus and largely affect (approximations of) expressive ontologies – this conclusion relies on a much larger and more diverse sample than existing comparisons between **MEX** and syntactic 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

⁴ <http://www.ihtsdo.org/snomed-ct/>

\mathcal{O} but rather only in a limited part of it which is relevant to their application. One recently explored technique for addressing this situation is to use *modules*, i.e., suitably small subsets of \mathcal{O} that behave for specific purposes like the original ontology over a given *signature* Σ , i.e., a set of terms (classes and properties).

Using a module rather than a whole ontology aims at improving performance since only information that is relevant to a restricted vocabulary is processed. However, the correctness of the outcome can be guaranteed only if the used modules satisfy certain well-defined properties. For example, reasoning-based tasks require the modules to *provide coverage* for \mathcal{O} over Σ , i.e., preserve *all* the entailments of \mathcal{O} over Σ (they are called *logical modules* [9,4]). Applications of logical modules include reuse of (a part of) well-established ontologies, ontology integration, and computing justifications to debug ontologies [11]. In these scenarios, though, a stronger notion of logical module is required that satisfies also two additional properties [15,19]: *self-containment* and *depletion*. The former means that the module preserves entailments over all terms that occur in the module (not just those used to extract the module). The latter means that $\mathcal{O} \setminus \mathcal{M}$ does not entail any non-tautological axioms over Σ . In this paper we will analyze only depleting and self-contained logical modules.

Interestingly, a *minimal* depleting and self-contained module for a signature Σ is, under some mild conditions, uniquely determined [15]. Extracting such modules is, unfortunately, computationally hard or even undecidable for expressive ontology languages [10,17,18]. In order to identify notions of modules whose extraction is feasible we can follow two alternative strategies. The first one consists of restricting the expressivity of the ontology language, as in the case of the MEX approach [14]: the MEX system allows for the extraction in polynomial time of the minimal self-contained and depleting module from acyclic \mathcal{ELI} terminologies. The second strategy consists of looking for practical sufficient conditions to guarantee the properties of logical modules without imposing minimality on the module \mathcal{M} , as it is the case for the family of logical modules known as *locality-based modules (LBMs)* [3]; these modules can be extracted from ontologies as expressive as \mathcal{SROIQ} , are self-contained and depleting, but can contain axioms that are not relevant to preserve any entailment over the given Σ .

The family of LBMs consists of module notions that are parameterized according to two features: (1) the technique used for identifying which axioms need to be included in the module (*semantic* or *syntactic*); (2) the kind of placeholder(s) used for those terms not included in the signature (*bottom*, *top*, or *nested*). In the next two paragraphs we provide an intuitive discussion of the meaning of these two features.

The extraction of semantic LBMs requires entailment checks against an empty ontology and thus involve reasoning, which makes the computation as hard as reasoning. Moreover, the kind of reasoning service used is rather unusual for DL reasoners.⁵ Hence, although algorithms for extracting semantic LBMs are known, until now and to the best of our knowledge they had not been

⁵ DL reasoners usually *classify* an ontology: test it for consistency and all concept names for satisfiability/mutual subsumption.

with syntactic bottom modules on a sample of 5000 random signatures and the SNOMED CT ontology. We perform our study on a larger corpus (not restricted to \mathcal{ELT}), compare more notions of logical modules, and also provide rigorous statistical significance results.

The main contributions of this paper are summarized as follows:

- 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 *culprits*, i.e., patterns of axioms that completely explain those rare differences. One includes simple tautologies that can be removed in a straightforward preprocessing step.
- 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 nested-modules) slower than the extraction of syntactic LBMs, and both only take milliseconds to seconds for most ontologies below 10K axioms.
- To obtain these results, we use our own implementation of semantic locality which, to the best of our knowledge, is the first ever to be implemented.
- 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 nested-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} [1,13]), and aim here at fixing the notations and at defining the key notions around module extraction, with a focus on locality-based modules [3] and MEX modules [14].

Let \mathcal{O} denote an ontology, \mathbf{N}_C a set of class names, and \mathbf{N}_R a set of property names. A *signature* is a set $\Sigma \subseteq \mathbf{N}_C \cup \mathbf{N}_R$ of *terms*. Given a class, property, or axiom X , we call the set of terms in X the *signature of X* , denoted \tilde{X} . Given a \mathcal{SROIQ} ontology \mathcal{O} , 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

direction is in general computationally hard, or even undecidable for expressive DLs [10,17,18]. 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 .

Let Σ be a signature and \mathcal{O} be an ontology. Let $x \in \{\text{MEX}, \emptyset, \Delta, \perp, \top\}$ be a notion of module. For each such notion, an oracle “ x -check” can be defined that determines whether an axiom α may be involved in preserving an entailment η of \mathcal{O} over Σ . Then, the x -module $x\text{-mod}(\Sigma, \mathcal{O})$ for Σ in \mathcal{O} can be computed by performing Algorithm 1.

Algorithm 1 Extraction of an x -module for Σ

Input: Ontology \mathcal{O} , seed signature Σ , oracle x -check

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 the x -check for α against $\Sigma \cup \widetilde{\mathcal{M}}$ is positive **then**

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

until changed = **false**

return \mathcal{M}

Algorithm 1 is a special case of the one in [3, Figure 4], and its output \mathcal{M} does not depend on the order in which the axioms α are selected [3].

Due to space limitations, we can just briefly sketch the intuition behind the definition of each oracle and the corresponding results of interest for this paper. We refer the interested reader to [3,14] for further details.

The MEX system. In [14], the notion of a MEX-module is defined for acyclic terminologies, i.e., ontologies that satisfy two conditions: (1) they only contain axioms of the form $\mathbf{A} \equiv \mathbf{C}$ or $\mathbf{A} \sqsubseteq \mathbf{C}$ where \mathbf{A} is a class name and \mathbf{C} is a complex class; (2) for each \mathbf{A} , there is at most one axiom with \mathbf{A} on the left-hand side; if one such axiom α exists, then \mathbf{A} is said to be *defined*, and to be *directly dependent on* all the terms \mathbf{X} that occur on the right-hand side of α (denoted $\mathbf{A} \succ \mathbf{X}$). The MEX method requires to determine for each defined class \mathbf{A} the set $\text{depend}_{\mathcal{O}}(\mathbf{A})$ of all the terms \mathbf{X} in \mathcal{O} such that the pair (\mathbf{A}, \mathbf{X}) belongs to the transitive closure of \succ . Intuitively, then, the MEX-check for an axiom α against a signature Σ tests whether either α defines a class $\mathbf{A} \in \Sigma \cup \widetilde{\mathcal{M}}$ and *uses*⁷ at least one term $\mathbf{X} \in \text{depend}_{\mathcal{O}}(\mathbf{A}) \cap (\Sigma \cup \widetilde{\mathcal{M}})$ in $\mathcal{O} \setminus \mathcal{M}$, or if every term on which \mathbf{A} depends only via \equiv -axioms is *used to define*⁷ some term in $\Sigma \cup \widetilde{\mathcal{M}}$. The authors prove that, if \mathcal{O} is an acyclic \mathcal{ELI} ontology, then using the oracle MEX-check in Algorithm 1 generates the minimal depleting self-contained module for a signature Σ in polynomial time.

Semantic locality In [3], the authors define a family of notions of locality with different parameters, the prominent notions being those where the placeholder x belongs to $\{\emptyset, \Delta\}$. These two notions of locality can be intuitively described

⁷ The expressions *use* and *used to define* are high-level intuitive descriptions of the two conditions given in [14, Fig. 4], to which we refer the reader since a formal definition goes beyond the scope of this paper.

as follows: a *SRIOQ* axiom α is \emptyset -local (resp. Δ -local) w.r.t. signature Σ if α' obtained by replacing all terms in $\tilde{\alpha} \setminus \Sigma$ with \perp (resp. \top) is a tautology, in which case the x -check returns negative. This treatment of α independently of the remaining axioms distinguishes the \emptyset - and Δ -check (as well as the \perp - and \top -check introduced in the next paragraph) from the MEX-check; hence the name *local*. The authors of [3] 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 *SRIOQ* axiom, so standard reasoners need to be extended.

Syntactic locality In order to achieve *tractable* module extraction, the two syntactic notions of x -locality for $x \in \{\perp, \top\}$ have been defined in [3]. Similarly to semantic locality, the x -check for an axiom α against a signature Σ operates on the transformed axiom α' obtained by replacing all terms not in Σ with the placeholder x . However, rather than invoking a reasoner, the x -check of α against Σ makes use of a simple syntactic test [3, Sec. 5.5]. For example, $\perp \sqsubseteq \mathcal{C}$ is clearly a tautology for each class \mathcal{C} . If the x -check is negative, α is said to be \perp - or \top -local w.r.t. Σ . The x -check used in syntactic LBMs is sound in identifying non-tautological axioms, but it may fail to spot a tautology, i.e., 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. Thus, 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 is performed in time polynomial in the size of the ontology.

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 [3,19], called $\top\perp^*$ - and $\Delta\emptyset^*$ -modules.

Algorithm 1 guarantees that the module notions considered here are self-contained and depleting: self-containment holds because of the iteration until the signature of \mathcal{M} remains unchanged; depletion holds because the axioms left out of \mathcal{M} are those whose x -check against the enlarged signature is negative.

3 Research questions and experimental design

A natural question arising is whether syntactic and semantic LBMs differ in practice, and, if yes, by how much. A second question is whether semantic module extraction is *noticeably* more costly: the x -check has to be carried out often—once per axiom and signature that the algorithm goes through—and it is hard to predict the feasibility of semantic LBM extraction. Altogether, we want to know whether syntactic LBMs are a *good* approximation of semantic LBMs, and how much they differ in cost. Similarly, for acyclic \mathcal{ELI} ontologies the analogous question arises: how good an approximation of MEX modules are LBMs?

An answer to these questions will allow for a more informed choice of which module extraction technique to select. 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 we 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 could have exceeded 300 hours. However, to include at least one case of a huge ontology, we have kept in the corpus NCI, an $\mathcal{SH}(\mathcal{D})$ ontology with 123,270 axioms.

This selection results in a corpus of 242 ontologies, which even beside NCI greatly vary in expressivity (from \mathcal{AL} to $\mathcal{SROIQ}(\mathcal{D})$) and in size (10–16,066 axioms, 10–16,068 terms) [12]. For a full list of the corpus, please refer to [5].

As mentioned above, it is not possible for some ontologies 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 use of the \top -role as required by the semantic locality tests.

Since MEX handles only acyclic \mathcal{ELT} ontologies, we created an \mathcal{ELT} version $\mathcal{ELT}(\mathcal{O})$ of each ontology \mathcal{O} in our corpus by filtering unsupported axioms and breaking terminological cycles. A principled way of doing this is beyond the scope of this paper, and we have used the heuristic described in [5]. The resulting corpus contains 239 ontologies since 3 were left empty after the \mathcal{ELT} -fication.

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, the semantic, and the MEX modules for Σ ; 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 \emptyset -local (Δ -local, resp.) w.r.t. Σ but not \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

⁸ <http://bioportal.bioontology.org>

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

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

- $\top\perp^*$ -mod(Σ, \mathcal{O}) \neq $\Delta\emptyset^*$ -mod(Σ, \mathcal{O}), or
 - $\Delta\emptyset^*$ -mod(Σ, \mathcal{O}) \neq MEX-mod(Σ, \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 $m = \#\tilde{\mathcal{O}}$, 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 $\Sigma_2 = \{\text{Mineral}\}$ [7].

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. Interestingly, a module \mathcal{M} is genuine iff there exists an axiom α such that $\mathcal{M} = x\text{-mod}(\tilde{\alpha}, \mathcal{O})$. As a consequence, there are only linearly many genuine modules in the size of \mathcal{O} , and extracting one module per axiom is enough for obtaining all of them. Moreover, all modules of \mathcal{O} are composed from genuine modules [7]. Thus, genuine modules are of special interest, and we can investigate *all* of them, together with the corresponding *genuine signatures*.

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.

With this in mind, we will analyze the modules obtained by random signatures with $p = 1/2$, and we will see in Section 4 that the module sizes obtained do allow for a reliable statement about the differences observed.

How many seed signatures do we have to sample from a given ontology \mathcal{O} in order to obtain statistically significant statements about modules determined

by the real population of *all* signatures from \mathcal{O} ? We apply the usual statistical model of confidence intervals [20], aiming at a confidence level of 95% that the true proportion of differences between modules – i.e., the proportion of seed signatures that lead to different modules – lies in the confidence interval ($\pm 5\%$) of the observed proportion. Then we can generalize the conclusions for the random sample to the full population because the probability that the proportion of differences among modules for all seed signatures differs by no more than 5% from the proportion observed in the sample (and reported in Section 4) is 95%. In order to reach this confidence level, we need a sample size of at least 385 elements, independently of the size of the full population: for a two-sided test to detect a change in the proportion defective of size δ in either direction, the minimum sample size is

$$N \geq \frac{p(1-p)}{\delta^2} z_{1-\alpha/2}^2,$$

where p is the observed proportion, α the significance level, and $z_{1-\alpha/2}$ the critical value of the underlying distribution [2]. Here, we use the normal distribution as an approximation of the binomial distribution which is usually assumed for proportions in random sampling; hence the significance level of $\alpha = 0.05$ leads to $z_{1-\alpha/2} \approx 1.96$. Furthermore, although we do not know the value p in advance, it is clear that $p(1-p) \leq 0.25$ because $0 \leq p \leq 1$. The confidence interval of $\pm 5\%$ determines the error of $\delta = 0.05$. Therefore, we obtain

$$N \geq \frac{0.25}{0.05^2} \cdot 1.96^2 \approx 384.16,$$

that is, a representative sample for these parameters needs at least 385 elements, and this number is independent of the population size. 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 questions (a) and (b), with Σ ranging over *all* the genuine signatures $\tilde{\beta}$ for $\beta \in \mathcal{O}$.

Our sample selection includes large as well as small seed signatures: the random seed signatures created to answer T1 will tend to contain around half the terms in the ontology, while the signatures used to answer T2 will range over *all* signatures of *single axioms* and therefore tend to be small.

4 Results of the Experiments

4.1 Semantic versus syntactic locality

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 209 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 nested-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 209 ontologies include Galen and People, which are renowned for having unusually large \perp -modules [3,8].

In most cases, extracting a semantic and syntactic LBM each took only a few milliseconds, so 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 nested-modules in Molecule Role, and up to 34 times in Galen.

Differences in locality. We have observed differences between syntactic and semantic locality for 33 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 class *names* by A, B , complex classes by C, D , properties by r, s, \dots , nominals by a , non-empty data ranges (e.g., int or $\text{int}^{0..9}$) 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 C^\perp and C^\top to denote classes that are bottom- or top-equivalent due to the grammar defining syntactic locality in [3, Fig. 3] 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 33 and are of the kinds $A \sqsubseteq A$ or $r \equiv (r^-)^-$. 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 that tautology, 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 27 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 27.

Ontology	Abbreviation	DL expressivity	#axioms	#terms
MiniTambis-repaired	MiniT	\mathcal{ALCN}	170	226
Tambis-full	Tambis	$\mathcal{SHIN}(\mathcal{D})$	592	496
Bleeding History Phenotype	BHO	$\mathcal{ALCIF}(\mathcal{D})$	1,925	581
Neuro Behavior Ontology	NBO	\mathcal{AL}	1,314	970
Pharmacogenomic Relationsh...	PhaRe	$\mathcal{ALCHIF}(\mathcal{D})$	459	311
Terminological and Ontological...	TOK	$\mathcal{SRIQ}(\mathcal{D})$	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, plus the relative sizes of modules extracted for (T1) and (T2). For a complete overview, including differences in locality of single axioms, see the table in [5].

Ontol.	Types affected	#diffs	size of diffs		size of $\Delta\emptyset^*$ -modules				culprit type
			#axs	(rel.)	T1	(%)	T2		
				range	avg.	range	avg.	+ freq.	
miniT	bot, nested	14–25%	1–7	0–600% ^b	48–79	66	0–8	2	c 3
Tambis	bot, nested	32–57%	2–41 ^c	1–62% ^c	75–88	82	0–34	9	c 8
BHO ^a	nested	17%	1–12	0–300%	55–72	65	0–31	4	b 31
NBO ^a	nested	3%	2	0–200%	64–78	71	0–3	0	d 3
PhaRe ^a	top, nested	1–8%	1–326 ^d	0–6,520% ^d	50–70	60	0–8	1	d 10
TOK	top, nested	49–100%	1–7	0–9%	48–68	59	9–17	10	d 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. 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. a signature that does not contain r because they become tautologies if r is replaced by \top . However, they can never be \top -local unless A is replaced by some C^\perp .

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

Culprits of type c are axioms α that contain a class 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” C to be a C^\perp (or C^\top). For example, $C = \forall r. \bar{A} \sqcap \exists r. \top$ becomes \perp -equivalent if A is replaced by \perp ; the same holds with cardinality restrictions in place of “ \exists ”. Consequently, axioms such as $A^\perp \equiv B \sqcap \forall r. C^\perp \sqcap \forall s. \{a\} \sqcap = 3 r. \top$, (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 nested-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 classes that cause differences in top-locality of single axioms.

Culprits of type d are axioms where a class (or property) name from the left-hand side occurs on the right-hand side together with a top-equivalent property (or class), causing differences in top-modules. The simplest such axiom is $A \sqsubseteq \exists \bar{r}. A$, 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, more complex, examples 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. 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.

Module sizes. The selection of the signatures for the experiment was designed to allow for the analysis of two, complementary, kinds of modules: 1) genuine modules, which constitute a base of all modules, extracted from generally small axiom signatures; 2) a statistically significant amount of random modules, obtained from random, unbiased signatures which are likely to contain half the terms of the ontology. We argue in what follows that it is neither the case that genuine modules are so small to be almost irrelevant sets to investigate, nor that random modules are so big to leave no space for differences to be observed. We will focus on syntactic modules which contain the other kinds of modules.

During the experiment we have computed and analyzed a high number of genuine modules: more than 380K for the \perp -notion, more than 40K for the \top -notion, and more than 440K for the $\top\perp^*$ -notion of locality. As we mentioned above, these modules tend to be quite small. However, they are not of irrelevant size: $\sim 8\%$ of the genuine \perp -modules, $\sim 11\%$ of the genuine \top -modules, and $\sim 5\%$ of the genuine $\top\perp^*$ -modules contain more than 20% of the axioms of the

corresponding ontology. So the low number of differences observed is not due to checking only against very small modules.

With a similar and complementary discussion, we argue that the modules obtained through random, “big” signatures do not necessarily contain almost all of the ontology: e.g., 39% of all random $\top\perp^*$ -modules, and 28% of all random \perp -modules, contain less than 60% of the axioms of the corresponding ontology.

To sum up, the lack of differences between the modules is not due to too small or to too big sizes of the modules selected.

Discussion. 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, see [5].

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. For this reason, the identification of culprits can only be done “on demand”, i.e., by observing the differences in the modules of given ontologies.

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* classes, by adding appropriate grammars to the definition of syntactic locality, and adding more cases of \perp - and \top -equivalent classes 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.

4.2 LBMs 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.

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 ($\top\perp^*$, $\Delta\emptyset^*$)

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 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 $\text{RICF} \equiv \text{ICF} \sqcap \exists\text{ISFO.RSH}$.

We analyzed whether the differences observed correlate with the size of the original ontology, its expressivity or the extent of the modification done in the \mathcal{ELT} -fication. There is no correlation with size but, as is to be expected, with the other two features, which are closely connected to each other. Table 4 illustrates the observations by dividing the 239 ontologies tested into four groups. The ontologies in Group 1 are in a format MEX can handle, so they have not been modified. The others required more or less heavy modifications (Groups 2–4). Differences between MEX and LBMs as described above occur only for ontologies that required heavy modifications (Group 4).

Group	#axioms removed	#ontologies	ontology size (avg.)
1 unchanged ontologies <i>no diff. $\Delta\emptyset^* \setminus \text{MEX}$</i>	0	33 (14%)	19–16,066 (2,176)
2 little-changed ontologies <i>no diff. $\Delta\emptyset^* \setminus \text{MEX}$</i>	1–28	36 (15%)	13– 6,587 (466)
3 largely-changed ontologies <i>no diff. $\Delta\emptyset^* \setminus \text{MEX}$</i>	31–7,836 (avg. 884)	104 (44%)	51–13,153 (2,373)
4 largely-changed ontologies <i>with diff. $\Delta\emptyset^* \setminus \text{MEX}$</i>	30–12,185 (avg. 1,001)	66 (27%)	42–12,344 (1,843)

Table 4. Overview of MEX experiment

As expected, the expressivity among Groups 1 and 2 is generally low: only 21 ontologies in Group 2 use expressivity above $\mathcal{AL}\mathcal{E}$ (up to $\mathcal{SHIF}(\mathcal{D})$, which is an outlier). However, the size of some ontologies in Group 1 is already considerable: 22 out of 33 have > 100 axioms; 10 have $> 1,000$ axioms. In contrast, the ontologies in Group 4 have almost always high expressivity, for example 27 out of 66 contain nominals.

Despite the correlation between the impact of the \mathcal{ELT} -fication and the differences observed between MEX- and $\Delta\emptyset^*$ -modules, we cannot claim that there is a causation between the two events. Indeed, we have investigated the reasons for the differences observed between the two kinds of modules, and we have noticed that in all the cases the culprit is the proliferation of equivalence axioms. 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 $\{B\}$.

The experimental results in view of this insight are summarized as follows:

Random-modules experiment: the 66 ontologies where differences between random MEX- and $\Delta\emptyset^*$ -modules were observed, coincide exactly with those where equivalences occur in the \mathcal{ELT} -TBox.

Genuine-modules experiment: all 61 ontologies where differences between genuine MEX- and $\Delta\emptyset^*$ -modules were observed contain equivalence axioms.

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

We conjecture that the low expressivity of the \mathcal{ELT} -language reduce the possibility of MEX- and $\Delta\emptyset^*$ -modules to differ only to the presence of equivalences. In addition to the empirical evidence for such a claim, we plan to investigate further this aspect in future work.

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. (2) In most cases, there is no or little difference between LBMs and MEX modules. (3) Though in principle hard to compute, semantic LBMs can be extracted rather fast in practice. Still, their extraction often takes considerably longer than that of 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 results (1) and (2), hardly any benefit can be expected from preferring potentially smaller modules (MEX or semantic LBMs) to cheaper syntactic LBMs. For the ontologies *Galen* and *People*, which are “renowned” for having disproportionately large modules, syntactic and semantic LBMs do not differ. Only for *Galen* are MEX modules considerably smaller than LBMs.

Not only does our study evaluate how good the cheap syntactic locality approximates semantic locality and model conservativity, it also required us to provide the first implementation for extracting modules based on semantic locality. Furthermore, we have been able to fix bugs in the existing implementation of syntactic modularity. A complete report of bugfixes is beyond the scope of the paper; as an example, early runs of the experiment led us to correcting the treatment of reflexivity axioms by the locality checker in the OWL API.

Future work. Two issues are interesting for future work: (1) Sampling seed signatures so that all sizes of signatures are equally likely to be sampled; (2) Comparing LBMs to other types of conservativity-based modules.

As for (1), the current sampling causes small and large signatures to be underrepresented. One might argue that, for big ontologies, the typical module extraction scenario does not require large seed signatures – but it does sometimes require relatively small seed signatures, for example, when a module is extracted to efficiently answer a certain entailment query of typically small size. We therefore plan to conduct a similar experiment using other sampling methods. 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.

References

1. Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.

2. Carroll Croarkin and Paul Tobias, editors. *NIST/SEMATECH e-Handbook of Statistical Methods*. NIST/SEMATECH, 2012. <http://www.itl.nist.gov/div898/handbook>.
3. 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.
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. In *Proc. of DL 2010*, 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. In *Proc. of WoMO-11*, volume 230 of *FAIA*, pages 25–39, 2011.
9. James Garson. Modularity and relevant logic. *Notre Dame J. of Formal 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. Laconic and precise justifications in OWL. In *Proc. of ISWC-08*, volume 5318 of *LNCS*, pages 323–338, 2008.
12. Matthew Horridge, Bijan Parsia, and Ulrike Sattler. The state of bio-medical ontologies. In *Proc. of ISMB-11*, 2011.
13. Ian Horrocks, Oliver Kutz, and Ulrike Sattler. The even more irresistible *SROIQ*. In *Proc. of KR-06*, pages 57–67, 2006.
14. 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.
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} . *J. of Symbolic Computation*, 45(2):194–228, 2010.
19. Ulrike Sattler, Thomas Schneider, and Michael Zakharyashev. Which kind of module should I extract? In *Proc. of DL 2009*, volume 477 of *ceur-ws.org*, 2009.
20. Michael Smithson. *Confidence Intervals*. Quantitative Applications in the Social Sciences. Sage Publications, 2003.