Local safety of an ontology*

Lukáš Homoľa[1] and Július Štuller[2]

[1] Faculty of Nuclear Sciences and Physical Engineering, Czech Technical University
lukashomola@hotmail.com
[2] Institute of Computer Science, Academy of Sciences of the Czech Republic
stuller@cs.cas.cz

Abstract. *The ability to import ontologies safely, that is, without changing the original meaning of their terms, has been identified as crucial for the collaborative development and the reuse of (OWL) ontologies.*
In this paper, we propose the notion of local safety of an ontology and we identify scenarios in which this notion may be useful in guiding the development of an ontology that is to import other ontologies safely.

1 Introduction

Defined as explicit specifications of conceptualizations of a domain of knowledge (or of a discourse) [1], ontologies are (virtually) always manifestations of a shared understanding of a domain. They typically take the form of a formal (e.g., logical) theory that fixes the vocabulary of a domain and, through constraining possible interpretations and well-formed use of the vocabulary terms, provides meaning for the vocabulary.

Ontologies have been advocated as a tool to support human communication, knowledge sharing and reuse, and interoperability between distributed systems. As such, ontologies have a range of applications in fields like knowledge management, information retrieval and integration, cooperative information systems, bioinformatics, medicine, linguistics, e-commerce, etc. Today, they are perhaps best known as the key technology of the Semantic Web vision.

The construction of a typical ontology is a collaborative process that involves *direct cooperation* among multiple individuals or groups of ontology engineers and domain experts (sometimes from different domains of expertise and different organizations) and/or *indirect cooperation* through the *reuse* of previously published, autonomously developed ontologies.

Most often, each team participating in the development of an ontology focuses on a part of it (a "component ontology") that pertains to the team's domain of expertise/authority and cooperates with other teams to relate the part it is working on with other parts. Performing an upgrade of even only one such a component ontology may require the participation of all the teams as different component ontologies are, when combined together, interrelated, depend on and affect one another (changing one component ontology may thus necessitate changes to the others and might require teams to reconcile their changes).

By reusing an ontology we mean using it as an input to develop a new ontology. In such a process, significant parts of the reused ontology are often extracted, refined, extended or otherwise adapted and then combined with other ontologies to form the final assembly.

One of the prerequisites for efficient collaborative ontology construction and maintenance is the ability to combine ontologies in a controlled way. The interaction among component ontologies should be controlled and well-understood in order to reduce the communication that is needed among different teams and to avoid expensive reconciliation processes. Ideally, controlled interaction should allow different teams to develop, test and upgrade their ontologies independently, to replace a component ontology or extend an ontology with minimal side effects. The issue is also vital for ontology reuse, especially in the case when the reused ontology, rather than being adapted and used as a draft to develop an ontology component, is linked to and remains under the control of its original developers, who may perform changes to it autonomously.

2 Problem definition

The Web Ontology Language (OWL) [2], a widely accepted W3C recommendation for creating and sharing ontologies on the Web, provides only very limited support for combining ontologies.

OWL adopts an importing mechanism, implemented by the `owl:imports`[3] construct, which allows one to include in an OWL ontology all the statements

* The work was supported by the project No. 1M0554 "Advanced Remedial Processes and Technologies" of the Ministry of Education, Youth and Sports of the Czech Republic and partly by the Institutional Research Plan AV0Z10300504 "Computer Science for the Information Society: Models, Algorithms, Applications". The first author also acknowledges the financial support of the Department of Mathematics at his institution.

[3] `http://www.w3.org/2002/07/owl#imports`, to be precise

contained in some other OWL ontology. In the importing ontology, there is no logical difference between the statements that are imported and the proper ones.

A number of recent papers by Grau et al. [3–7] stressed the particular relevance of the ability to import OWL ontologies "safely", that is, in such a way that the imported terms (the terms of the imported ontologies) preserve their original meaning (the meaning these terms have in the imported ontologies) in the importing ontology. This ability is applicable in typical scenarios of OWL ontology development such as in the following one (see the above-mentioned works by Grau et al. for a motivating example):

- an ontology engineer distinguishes between the so-called *external terms* and the so-called *local terms* of the ontology $\mathcal{O}$ he or she is developing;
- the local terms are those whose meaning is assumed to be fully described in the ontology $\mathcal{O}$ itself, possibly with the help of the remaining, external terms;
- the meaning of the external terms is assumed to be only partially described in the ontology $\mathcal{O}$ – in terms of their use in the description of the local terms – and to be further described in some other ontologies (preexistent or concurrently developed) that are to be imported into $\mathcal{O}$;
- the use of the external terms in the statements of the ontology $\mathcal{O}$ is expected not to alter the original meaning these terms have in the ontologies to be imported.

In the paper, we continue in the study, initiated by Grau at al., of the methodology for OWL ontology development in the scenario given above. We propose the notion of *local safety of an ontology* and discuss under which conditions and how this notion can be used to guide the development of OWL ontologies.

3 Preliminaries

In this section we introduce description logics (DLs) [8], a family of logic-based knowledge representation formalisms, which underly modern ontology languages such as OWL. OWL consists of three (sub)languages of increasing expressive power, two of which, namely OWL Lite and OWL DL, roughly correspond to the DLs $\mathcal{SHIF}$ and $\mathcal{SHOIN}$, respectively.

DLs view the world as being populated by objects and allow one to represent the relevant notions of the domain of interest in terms of *concepts*, *roles* and (possibly) *individuals*, representing sets of elements, binary relationships between elements and single elements, respectively. Starting from atomic concepts, atomic roles and individuals, which are denoted simply by a name, complex concepts and complex roles are built using concept and role *constructors*. We assume the sets $\mathbf{C}$, $\mathbf{R}$ and $\mathbf{I}$ of (respectively) *atomic concepts*, *atomic roles* and *individuals* to be countably infinite and mutually disjoint and to be fixed for every DL. An ontology $\mathcal{O}$ formalized in a DL takes the form of a finite set of terminological and role *axioms*, which are used to suitably organize and interrelate multiple concept and role descriptions. DLs are distinguished by constructors and/or types of axioms they provide. We will use the term $\mathcal{L}$-axiom ($\mathcal{L}$-ontology) to emphasize we are talking about an axiom (an ontology) in the DL $\mathcal{L}$.

In the abstract notation we will use the letters A, B to denote atomic concepts, r, s to denote atomic roles, and a, b for individuals (all the letters possibly with a subscript). The letters C, D will be used to denote a concept (atomic or complex), R, S to denote a role, and α, β to denote an axiom.

As the *minimal DLs* of practical interest are usually considered the DLs $\mathcal{EL}$ and $\mathcal{AL}$, which both are fragments of the smallest propositionally closed DL $\mathcal{ALC}$. In $\mathcal{ALC}$, concepts are composed inductively according to the following syntax rule:

$$\begin{aligned}C, D \rightarrow\ & A\ (\text{atomic concept}) \mid \\ & \bot\ (\text{bottom concept}) \mid \top\ (\text{top concept}) \mid \\ & \neg C\ (\text{concept negation}) \mid \\ & C \sqcap D\ (\text{conjunction}) \mid C \sqcup D\ (\text{disjunction}) \mid \\ & \exists R.C\ (\text{existential restriction}) \mid \\ & \forall R.C\ (\text{value restriction}).\end{aligned}$$

Valid constructs for $\mathcal{EL}$ are: $\bot$, $C \sqcap D$ and $\exists R.C$. In $\mathcal{AL}$, the syntax of complex concepts is the following: $\bot$, $\top$, $\neg A$ (atomic concept negation), $C \sqcap D$, $\exists R.\top$ (limited existential restriction), and $\forall R.C$.

DLs $\mathcal{EL}$, $\mathcal{AL}$ and $\mathcal{ALC}$ provide no role constructors.

The listed $\mathcal{ALC}$ constructors are not all independent ($\top = \neg\bot$, $C \sqcup D = \neg(\neg C \sqcap \neg D)$, $\forall R.C = \neg(\exists R.\neg C)$). In fact, $\mathcal{ALC}$ can be obtained from both $\mathcal{EL}$ and $\mathcal{AL}$ by adding the concept negation constructor.

A *terminological axiom* in $\mathcal{EL}$, $\mathcal{AL}$ and $\mathcal{ALC}$ is an expression of the following forms: $A \equiv C$ (*concept definition*), $A \sqsubseteq C$ (*concept specialization*) or $C \sqsubseteq D$ (*general concept inclusion*, GCI). The abbreviation of the form $C \equiv D$ (*concepts equality*) stands for the two GCIs $C \sqsubseteq D$ and $D \sqsubseteq C$. $\mathcal{EL}$, $\mathcal{AL}$ and $\mathcal{ALC}$ provide no role axioms.

$\mathcal{S}$ is an extension of $\mathcal{ALC}$ in which an atomic role can be declared transitive using the role axiom of the form $\mathsf{Trans}(r)$.

Further extensions of DLs are indicated by appending letters to the DL's name. Advanced concept constructors include *number restrictions* of the form

$\geq nR$ (indicated by appending the letter $\mathcal{N}$), *qualified number restrictions* $\geq nR.C$ (appending $\mathcal{Q}$) and *nominals* $\{a\}$ (appending $\mathcal{O}$). In the case of number restrictions and qualified number restrictions, the dual constructors $\leq nR$ and $\leq nR.C$ are introduced as abbreviations for $\neg(\geq n+1R)$ and $\neg(\geq n+1R.C)$, respectively. Nominals allows to construct a concept representing a singleton set containing one individual. Enumeration $\{a_1, \ldots, a_n\}$ is an abbreviation for $\{a_1\} \sqcup \ldots \sqcup \{a_n\}$.

Yet other extensions include role constructors, of which the *inverse role* constructor r^- (appending $\mathcal{I}$) is the most prominent one. Another important type of role axioms is the *role inclusion* $R \sqsubseteq S$ (appending $\mathcal{H}$).

These extensions can be used in different combinations, for example $\mathcal{ALN}$ is an extension of $\mathcal{AL}$ with number restrictions, and $\mathcal{SHOIN}$ is the DL that uses 5 of the constructors we have presented.

The semantics of DLs is defined via interpretations. An *interpretation* $\mathcal{I}$ is a pair $\mathcal{I} = (\Delta^\mathcal{I}, .^\mathcal{I})$, where $\Delta^\mathcal{I}$ is a non-empty set, called the *domain* of the interpretation, and $.^\mathcal{I}$ is the *interpretation function*, which maps atomic concepts to subsets of $\Delta^\mathcal{I}$, atomic roles to binary relations over $\Delta^\mathcal{I}$ and individuals to elements of $\Delta^\mathcal{I}$. The interpretation function extends to complex concepts as follows:

$$
\begin{aligned}
\bot^\mathcal{I} &= \emptyset, \top^\mathcal{I} = \Delta^\mathcal{I},\\
(C \sqcap D)^\mathcal{I} &= C^\mathcal{I} \cap D^\mathcal{I}, (C \sqcup D)^\mathcal{I} = C^\mathcal{I} \cup D^\mathcal{I},\\
(\neg C)^\mathcal{I} &= \Delta^\mathcal{I} - C^\mathcal{I},\\
(\forall R.C)^\mathcal{I} &= \{x \in \Delta^\mathcal{I};\ \forall y\,((x,y) \in R^\mathcal{I} \to y \in C^\mathcal{I})\},\\
(\exists R.C)^\mathcal{I} &= \{x \in \Delta^\mathcal{I};\ \exists y\,((x,y) \in R^\mathcal{I} \land y \in C^\mathcal{I})\},\\
(\geq nR)^\mathcal{I} &= \{x \in \Delta^\mathcal{I};\ |\{y \in \Delta^\mathcal{I};\ (x,y) \in R^\mathcal{I}\}| \geq n\},\\
(\geq nR.C)^\mathcal{I} &= \{x \in \Delta^\mathcal{I};\\
&\qquad |\{y \in \Delta^\mathcal{I};\ (x,y) \in R^\mathcal{I} \land y \in C^\mathcal{I}\}| \geq n\},\\
\{a\}^\mathcal{I} &= \{a^\mathcal{I}\},\\
(r^-)^\mathcal{I} &= \{(x,y);\ (y,x) \in r^\mathcal{I}\}.
\end{aligned}
$$

The semantics of terminological axioms is defined in terms of a satisfaction relation $\models$, which relates interpretations to the terminological axioms they satisfy. It is defined as follows: $\mathcal{I} \models C \sqsubseteq D$ iff $C^\mathcal{I} \subseteq D^\mathcal{I}$, $\mathcal{I} \models R \sqsubseteq S$ iff $R^\mathcal{I} \subseteq S^\mathcal{I}$, $\mathcal{I} \models C(a)$ iff $a^\mathcal{I} \in C^\mathcal{I}$, $\mathcal{I} \models R(a,b)$ iff $(a^\mathcal{I}, b^\mathcal{I}) \in R^\mathcal{I}$, $\mathcal{I} \models \mathsf{Trans}(r)$ iff the relation $r^\mathcal{I}$ is transitive. Interpretations satisfying an axiom are said to be its models.

An interpretation $\mathcal{I}$ is a *model* of an ontology $\mathcal{O}$ (written $\mathcal{I} \models \mathcal{O}$) iff $\mathcal{I} \models \alpha$ for all $\alpha \in \mathcal{O}$. An ontology is said to be *consistent* if it has at least one model and is said to be *inconsistent* otherwise.

An ontology $\mathcal{O}$ *entails* an axiom α (written $\mathcal{O} \models \alpha$) iff all models of $\mathcal{O}$ satisfy α, especially we will speak about *subsumption* between C and D in the case of $\mathcal{O} \models C \sqsubseteq D$, and *satisfiability of concept* C in the case $\mathcal{O} \not\models C \sqsubseteq \bot$.

Interpretations $\mathcal{I}$ and $\mathcal{J}$ are *isomorphic* (written $\mathcal{I} \cong \mathcal{J}$) iff there is a bijection $\mu : \Delta^\mathcal{I} \to \Delta^\mathcal{J}$ such that for every $x, y \in \Delta^\mathcal{I}, A \in \mathbf{C}, r \in \mathbf{R}, a \in \mathbf{I}$ the following holds: $x \in A^\mathcal{I}$ iff $\mu(x) \in A^\mathcal{J}$, $(x,y) \in r^\mathcal{I}$ iff $(\mu(x), \mu(y)) \in r^\mathcal{J}$, $x = a^\mathcal{I}$ iff $\mu(x) = a^\mathcal{J}$. Isomorphic interpretations are semantically indistinguishable (in particular, they satisfy the same axioms).

A *signature* $\mathbf{S}$ is a finite subset of $\mathbf{C} \cup \mathbf{R} \cup \mathbf{I}$. Two interpretations $\mathcal{I}$ and $\mathcal{J}$ *coincide* on a signature $\mathbf{S}$ (written $\mathcal{I}|_\mathbf{S} = \mathcal{J}|_\mathbf{S}$) iff $\Delta^\mathcal{I} = \Delta^\mathcal{J}$ and $X^\mathcal{I} = X^\mathcal{J}$ holds for all $X \in \mathbf{S}$.

We say that $\mathcal{I}$ has been obtained from $\mathcal{J}$ through a *domain expansion* with the set Δ (such $\mathcal{I}$ will by denoted by $\mathcal{J}_{\cup\Delta}$) iff Δ is a non-empty set disjoint with $\Delta^\mathcal{J}$, $\Delta^\mathcal{I} = \Delta^\mathcal{J} \cup \Delta$, and $X^\mathcal{I} = X^\mathcal{J}$ holds for all $X \in \mathbf{C} \cup \mathbf{R} \cup \mathbf{I}$. Note that $\mathcal{J}_{\cup\Delta}$ and $\mathcal{J}$ only differ in that the domain of $\mathcal{J}$ is a proper subset of the domain of $\mathcal{J}_{\cup\Delta}$ (with Δ being the set of additional domain elements).

A DL $\mathcal{L}$ is said to have the *finite model property* (FMP) iff every consistent $\mathcal{L}$-ontology admits a model that is finite (i.e., with a finite domain). One of the most prominent DLs that exhibit the FMP is $\mathcal{SHOQ}$, while $\mathcal{SHIN}$ is an example of a DL that lacks the FMP. For a DL $\mathcal{L}$ with the FMP, $\mathcal{L}$-ontology $\mathcal{O}$ and $\mathcal{L}$-axiom α the following holds: $\mathcal{O} \models \alpha$ iff $\mathcal{I} \models \alpha$ for all finite models $\mathcal{I} \models \mathcal{O}$.

We say that an interpretation $\mathcal{K}$ is a *disjoint union* of interpretations $\mathcal{I}$ and $\mathcal{J}$ (written $\mathcal{K} = \mathcal{I} \uplus \mathcal{J}$) iff there exist some interpretations $\tilde{\mathcal{I}}$ and $\tilde{\mathcal{J}}$ satisfying $\tilde{\mathcal{I}} \cong \mathcal{I}$, $\tilde{\mathcal{J}} \cong \mathcal{J}$ and $\Delta^{\tilde{\mathcal{I}}} \cap \Delta^{\tilde{\mathcal{J}}} = \emptyset$ for which the following holds: $\Delta^\mathcal{K} = \Delta^{\tilde{\mathcal{I}}} \cup \Delta^{\tilde{\mathcal{J}}}$, $X^\mathcal{K} = X^{\tilde{\mathcal{I}}} \cup X^{\tilde{\mathcal{J}}}$ for all $X \in \mathbf{C} \cup \mathbf{R}$ and $a^\mathcal{K} = a^{\tilde{\mathcal{I}}}$ for all $a \in \mathbf{I}$. Intuitively, the interpretation $\mathcal{K}$ for which $\mathcal{K} = \mathcal{I} \uplus \mathcal{J}$ holds is composed of two unrelated parts one being isomorphic to $\mathcal{I}$ and the other to $\mathcal{J}$. Disjoint union of a set of interpretations is defined analogously.

A DL $\mathcal{L}$ is said to have the *disjoint union model property* (DUMP) iff the set of models of arbitrary $\mathcal{L}$-ontology is closed under disjoint unions.

A prominent example of a DL that enjoys the DUMP is $\mathcal{SHIQ}$. DLs that support nominals do not have this property.

In the subsequent sections, will use $\mathbf{C}(\alpha)$ to denote the set of all atomic concepts that occur in the axiom α (the sets $\mathbf{R}(\alpha)$ and $\mathbf{I}(\alpha)$ are defined analogously). We will use $\mathbf{Sig}(\alpha)$ as a shorthand for $\mathbf{C}(\alpha) \cup \mathbf{R}(\alpha) \cup \mathbf{I}(\alpha)$. $\mathbf{C}(\mathcal{O})$ will stand for $\bigcup_{\alpha \in \mathcal{O}} \mathbf{C}(\alpha)$ (the sets $\mathbf{R}(\mathcal{O})$, $\mathbf{I}(\mathcal{O})$ and $\mathbf{Sig}(\mathcal{O})$ are defined analogously).

4 Related work

In the papers by Grau et al. [3–7], safety of ontology import is formulated using the notion of conservative extension, in the context of ontologies first used in [9] and recently further studied in [10, 11].

Definition 1 (Conservative extension). *Let $\mathcal{L}$ be a DL, $\mathcal{O}_1$ and $\mathcal{O}_2$ two ontologies such that $\mathcal{O}_1 \subseteq \mathcal{O}_2$.*

We say that $\mathcal{O}_2$ is a deductive conservative extension of $\mathcal{O}_1$ w.r.t. $\mathcal{L}$, if for every $\mathcal{L}$-axiom α with $\mathbf{Sig}(\alpha) \subseteq \mathbf{Sig}(\mathcal{O}_1)$, we have $\mathcal{O}_2 \models \alpha$ iff $\mathcal{O}_1 \models \alpha$.

An ontology $\mathcal{O}$ into which an ontology $\mathcal{O}'$ can be safely imported is said to be *safe for $\mathcal{O}'$*.

Definition 2 (Safety for an ontology). *Let $\mathcal{L}$ be a DL, $\mathcal{O}$ and $\mathcal{O}'$ two ontologies.*

We say that $\mathcal{O}$ is safe for $\mathcal{O}'$ w.r.t. $\mathcal{L}$, if $\mathcal{O} \cup \mathcal{O}'$ is a conservative extension of $\mathcal{O}'$ w.r.t. $\mathcal{L}$.

Ghilardi at al. [12] studied novel DL reasoning services aimed at supporting developers in customizing their ontology to be safe for a particular ontology.

As regards the scenario we are concerned with, Grau et al. [3–7] argues that in practice it is often convenient, or even necessary, for the developers of an ontology $\mathcal{O}$ to abstract from particular ontologies that are to be imported into it and focus instead only on $\mathcal{O}$ and on its external terms:

- ontologies to be imported might not be available during the development of $\mathcal{O}$ (as it is in the case when these ontologies are developed concurrently with $\mathcal{O}$);
- the developers of $\mathcal{O}$ are usually not willing to commit to particular versions of the ontologies they intend to import (the development of a typical ontology is a never-finished process);
- at a later time, the developers might find ontologies other than those initially considered more suitable for providing the meaning of the external terms of $\mathcal{O}$.

Grau et al. proposed the following condition to be used to guide the development of an ontology $\mathcal{O}$ in such cases.

Definition 3 (Safety for a signature). *Let $\mathcal{L}$ be a DL, $\mathcal{O}$ an ontology and $\mathbf{S}$ a signature.*

We say that $\mathcal{O}$ is safe for $\mathbf{S}$ w.r.t. $\mathcal{L}$, if for every $\mathcal{L}$-ontology $\mathcal{O}'$ such that $\mathbf{Sig}(\mathcal{O}) \cap \mathbf{Sig}(\mathcal{O}') \subseteq \mathbf{S}$, $\mathcal{O}$ is safe for $\mathcal{O}'$ w.r.t. $\mathcal{L}$.

Once an ontology $\mathcal{O}$ is safe for the signature $\mathbf{S}$ (which is presumably the set of its external terms) w.r.t. $\mathcal{L}$, one can safely import into $\mathcal{O}$ any ontology $\mathcal{O}'$ written in $\mathcal{L}$ and sharing only terms from $\mathbf{S}$ with $\mathcal{O}$.

As Grau et al. showed, even the problem of checking whether an ontology consisting of a single $\mathcal{ALC}$ axiom is safe for a signature w.r.t. $\mathcal{ALCO}$ is undecidable. It is not yet known whether the safety for a signature is decidable for weaker DLs, such as $\mathcal{EL}$, or for more expressive DLs. Grau et al. proposed several safety classes of ontologies, parametrized by a signature $\mathbf{S}$ and representing sufficient conditions for safety for $\mathbf{S}$, that are decidable and can be checked syntactically in polynomial time.

Several extensions to OWL have been proposed to better support collaborative ontology development and ontology reuse, including P-OWL [13], C-OWL [14], the extension based on $\mathcal{E}$-connections [15] and the extension based on the so-called semantic import [16]. All such extension are, however, still subjects of research and are not included in the current candidate recommendation for OWL 2 [17], an ongoing extension to and revision of OWL.

5 Local safety of an ontology

The notion of safety for a signature, along with the corresponding safety classes, facilitates the construction of an ontology that is safe for any ontology (in a given DL) with which it shares only some prearranged set of terms.

In the scenario we are interested in here, however, an ontology engineer does not always need to have the ontology $\mathcal{O}$ safe for every possible ontology (every possible set of axioms in a certain DL), but often only needs to have it safe for a certain, conveniently chosen class of *candidate* ontologies. This is the case, for instance, when the scenario applies to collaborative ontology development and $\mathcal{O}$ is considered as a component ontology for a larger ontology developed distributively as a set of ontologies importing one another. The development of component ontologies in such a case is typically coordinated to some extent (e.g., some principles on which individual component ontologies will be build are resolved beforehand and made explicit) and the developers can make assumptions about some qualities and characteristics of the ontologies they import (as well as about the way these ontologies may further evolve).

5.1 Local ontologies

Ontologies, like other engineering artifacts, are designed. When we choose how to represent something in an ontology[4], we are making design decisions. The

[4] "There is no one correct way to model a domain – there are always viable alternatives." [18]

best solution to ontology design depends on a number of factors, of which the most important include the intended use of an ontology, and the anticipated extensions and refinements to it.

Generally accepted and widely cited are the five design criteria Gruber [19] proposed for ontologies whose purpose is knowledge sharing and interoperation among programs. They include the following criterion:

> *An ontology should offer a conceptual foundation for a range of anticipated tasks, and the representation should be crafted so that one can extend and specialize the ontology monotonically. Especially, one should be able to define new terms for special uses based on the existing vocabulary, in a way that does not require the revision of the existing definitions.*

To facilitate the design, deployment and reuse of ontologies, Guarino [20] suggested the development of different kinds of ontologies with different levels of generality and dependence on a particular domain, task or point of view, namely top-level ontologies, domain and task ontologies and application ontologies. Terms of ontologies on a lower level are, in some sense, held to be specializations of terms of ontologies on a level above. Top-level ontologies, which describe concepts independent of a particular problem or domain (such as space, time, object, event, action, etc.) are meant to be unifying for a large group of ontologies on lower levels.

Swartout et al. [21] proposed a number of desiderata aimed primarily at domain, task and application ontologies. They include the two following:

> *An ontology should be extensible. [...] Extension should be possible both at a low level, by adding domain-specific subconcepts, or at high level by adding intermediate or upper level concepts that cover new areas.*
>
> *Ontologies should not be "stovepipes." The derisive term "stovepipe system" is used to describe a system that may be vertically integrated but cannot be integrated horizontally with other systems.*

Here we propose a notion of restricting the meaning of the top concept, which, as we believe, provide a partial characterization of ontologies violating the aforementioned criteria.

Definition 4 (Restricting the meaning of the top concept). *We say that the ontology $\mathcal{O}$ restricts the meaning of the top concept, if there are atomic concepts $A_1, \ldots, A_n$, atomic roles $r_1, \ldots, r_m, s_1, \ldots, s_k$ and individuals $a_1, \ldots, a_l$ in $\mathbf{Sig}(\mathcal{O})$ such that:*

$$\mathcal{O} \models \top \sqsubseteq A_1 \sqcup \ldots \sqcup A_n \sqcup \exists r_1.\top \sqcup \ldots \sqcup \exists r_m.\top \sqcup \\ \sqcup \exists s_1^-.\top \sqcup \ldots \sqcup \exists s_k^-.\top \sqcup \{a_1, \ldots, a_l\}.$$

Intuitively, an ontology restricts the meaning of the top concept if it introduces its vocabulary in such a way that the vocabulary can only be further specialized but not otherwise monotonically extended. In the case of domain, task and application ontologies at least, such an ontology can be considered badly-designed:

- it can not be, without previous modification, extended to cover a broader subject area than it already does,
- it is unsuitable for importing into any ontology that touches, even marginally, a subject area disjoint with that already covered by it.

As regards top-level ontologies, we studied the Basic Formal Ontology[5] and also the design of several other top-level ontologies [22], and came to the conclusion that even in this case the developers prefer not to restrict the meaning of the top concept. The only exception we found is the top-level ontology[6] proposed by John Sowa.

The notion of restricting the meaning of the top concept is closely related to the notion of localness of an ontology studied (also under the name safety of an ontology) by Grau et al. [23–26].

Definition 5 (Localness). *An ontology $\mathcal{O}$ is local if the set of its models is closed under domain expansion (i.e., if $\mathcal{I} \models \mathcal{O}$ implies $\mathcal{I}_{\cup\Delta} \models \mathcal{O}$ for every interpretation $\mathcal{I}$ and every non-empty set Δ disjoint with $\Delta^{\mathcal{I}}$).*

In [23], a syntactic characterization of localness for $\mathcal{SHOIQ}$ ontologies is given, which allows one to check localness of an $\mathcal{SHOIQ}$ ontology in polynomial time. We used this characterization[7] to show that $\mathcal{SHOIQ}$ ontologies restricting the meaning of the top concept are exactly those that are not local.

Proposition 1. *A $\mathcal{SHOIQ}$ ontology restricts the meaning of the top concept iff it is not local.*

Grau et al. [25] also reported testing over 700 ontologies available on the Web for localness and finding more than 99% of them local. However, we could not trace any further details on their experimental results.

5.2 Local safety

Regarding the development of an ontology in our scenario, the considerations above suggest that it is often possible to consider only local ontologies as the candidates for importing.

---

[5] `http://www.ifomis.org/bfo`

[6] `http://www.jfsowa.com/ontology/toplevel.htm`

[7] Relevant points of the characterization are reproduced at the beginning of the Appendix.

Definition 6 (Local safety for a signature). *Let $\mathcal{L}$ be a DL, $\mathbf{S}$ a signature and $\mathcal{O}$ an ontology.*

We say that $\mathcal{O}$ is locally safe for $\mathbf{S}$ w.r.t. $\mathcal{L}$, if for every local $\mathcal{L}$-ontology $\mathcal{O}'$ with $\mathbf{Sig}(\mathcal{O}) \cap \mathbf{Sig}(\mathcal{O}') \subseteq \mathbf{S}$, $\mathcal{O}$ is safe for $\mathcal{O}'$ w.r.t. $\mathcal{L}$.

The following proposition provides sufficient condition for local safety w.r.t. $\mathcal{SHOIQ}$.

Proposition 2. *Let $\mathbf{S}$ be a signature and $\mathcal{O}$ an ontology such that for every interpretation $\mathcal{J}$ there exists a model $\mathcal{I}$ of $\mathcal{O}$ such that*

- $\Delta^{\mathcal{I}} = \Delta^{\mathcal{J}} \cup \Delta$ *for some* Δ, $\Delta \cap \Delta^{\mathcal{J}} = \emptyset$,
- $X^{\mathcal{I}} = X^{\mathcal{J}}$ *for all* $X \in \mathbf{S}$.

Then $\mathcal{O}$ is locally safe for $\mathbf{S}$ w.r.t. $\mathcal{SHOIQ}$.

The following example demonstrates that, in comparison with safety condition proposed by Grau et al., the condition of local safety is less restrictive and may allow for a more convenient use of the external terms.

Example 1. Let us consider building an OWL ontology intended to provide a reference terminology for the annotation of films. Let us assume we intend to safely import into our ontology $\mathcal{O}$ some well-designed (and thus local) ontology that defines the categorization of films by genre. Suppose that the atomic concept Film is expected to be the only term our ontology will share with the imported ontology. Whereas the ontology

$$\mathcal{O} = \{\mathsf{Director} \sqsubseteq \neg\mathsf{Film}, \mathsf{Film} \sqsubseteq \exists \mathsf{hasDirector}.\mathsf{Director}\}$$

is not safe for $\{\mathsf{Film}\}$ even w.r.t. $\mathcal{AL}$ (take the non-local ontology $\mathcal{O}' = \{\top \sqsubseteq \mathsf{Film}\}$ as a counterexample), it is, according to Proposition 2, locally safe for $\{\mathsf{Film}\}$ w.r.t. $\mathcal{SHOIQ}$.

When we are concerned with local safety w.r.t. $\mathcal{SHOQ}$ (which has the FMP) we can use the following sufficient condition.

Proposition 3. *Let $\mathbf{S}$ be a signature and $\mathcal{O}$ an ontology such that for every finite interpretation $\mathcal{J}$ there exists a model $\mathcal{I}$ of $\mathcal{O}$ such that*

- $\Delta^{\mathcal{I}} = \Delta^{\mathcal{J}} \cup \Delta$ *for some* Δ, $\Delta \cap \Delta^{\mathcal{J}} = \emptyset$,
- $X^{\mathcal{I}} = X^{\mathcal{J}}$ *for all* $X \in \mathbf{S}$.

Then $\mathcal{O}$ is locally safe for $\mathbf{S}$ w.r.t. $\mathcal{SHOQ}$.

The two following propositions give a recipe for deciding local safety of an ontology written in $\mathcal{SHIQ}$ (which has the DUMP) w.r.t. $\mathcal{SHOQ}$ in the case when all external terms are atomic concepts.

Proposition 4. *Let $\mathbf{S}$ be a signature such that $\mathbf{S} \subseteq \mathbf{C}$, $\mathcal{O}$ a $\mathcal{SHIQ}$ ontology and a an individual. Assume that for every subset $\tilde{\mathbf{S}} \subseteq \mathbf{S}$ the ontology*

$$\mathcal{O} \cup \{A \equiv \{a\}\}_{A \in \tilde{\mathbf{S}}} \cup \{A \equiv \bot\}_{A \in \mathbf{S} - \tilde{\mathbf{S}}}$$

is consistent.

Then $\mathcal{O}$ is locally safe for $\mathbf{S}$ w.r.t. $\mathcal{SHOQ}$.

Proposition 5. *Let $\mathbf{S}$ be a signature such that $\mathbf{S} \subseteq \mathbf{C}$, $\mathcal{O}$ a $\mathcal{SHIQ}$ ontology and a an individual. Assume that there exists a subset $\tilde{\mathbf{S}} \subseteq \mathbf{S}$ such that the ontology*

$$\mathcal{O} \cup \{A \equiv \{a\}\}_{A \in \tilde{\mathbf{S}}} \cup \{A \equiv \bot\}_{A \in \mathbf{S} - \tilde{\mathbf{S}}}$$

is inconsistent.

Then $\mathcal{O}$ is not locally safe for $\mathbf{S}$ w.r.t. $\mathcal{ALO}$ ($\mathcal{ELO}$).

Corollary 1. *Let $\mathbf{S}$ be a signature such that $\mathbf{S} \subseteq \mathbf{C}$, $\mathcal{O}$ a $\mathcal{SHIQ}$ ontology.*

Then $\mathcal{O}$ is locally safe for $\mathbf{S}$ w.r.t. $\mathcal{SHOQ}$ iff $\mathcal{O}$ is locally safe for $\mathbf{S}$ w.r.t. $\mathcal{ALO}$ ($\mathcal{ELO}$).

Corollary 2. *Let $\mathcal{L}$ be a DL that is in between $\mathcal{ALO}$ ($\mathcal{ELO}$) and $\mathcal{SHOQ}$.*

The problem of deciding whether a $\mathcal{SHIQ}$ ontology is locally safe for a signature $\mathbf{S}$, $\mathbf{S} \subseteq \mathbf{C}$, w.r.t. $\mathcal{L}$ is reducible to the problem of checking consistency of a finite set of $\mathcal{SHOIQ}$ ontologies, and thus decidable.

Corollary 3. *Let $\mathcal{O}$ be a $\mathcal{SHIQ}$ ontology locally safe for $\mathbf{S} \subseteq \mathbf{C}$ w.r.t. $\mathcal{SHOQ}$, $\mathcal{O}'$ a local $\mathcal{SHOQ}$ ontology such that $\mathbf{Sig}(\mathcal{O}) \cap \mathbf{Sig}(\mathcal{O}') \subseteq \mathbf{S}$. Then $\mathcal{O} \cup \mathcal{O}'$ is locally safe for $\mathbf{S} \setminus \mathbf{Sig}(\mathcal{O}')$ w.r.t. $\mathcal{SHOQ}$.*

6 Conclusion and outlook

This paper contributes to the framework for ontology development presented by Grau et al. We proposed the notion of local safety of an ontology and showed its applicability in the development of real-world ontologies. We showed that local safety for a signature consisting solely of atomic concepts is decidable for an interesting group of description logics.

For the future work, we would like to study decidability and computational properties of (sufficient conditions for) local safety for a signature that contains atomic roles as well. The results obtained in the paper are also directly applicable to the problem of extracting reusable ontology parts, or ontology modules, as conceived by Grau et al. in the cited works.

References

1. T.R. Gruber: *A translation approach to portable ontology specifications.* Knowledge Acquisition, 5, 2, 1993, 199–220.
2. P. Patel-Schneider, P. Hayes, and I. Horrocks: *OWL Web Ontology Language: Semantics and Abstract Syntax.* W3C recommendation, W3C, 2004. Available at `http://www.w3.org/TR/owl-semantics/`.
3. B.C. Grau, I. Horrocks, Y. Kazakov, and U. Sattler: *Ontology Reuse: Better Safe than Sorry.* In Description Logics, Bozen/Bolzano University Press, 2007, 41–52.

4. B.C. Grau, I. Horrocks, Y. Kazakov, and U. Sattler: *Just the right amount: extracting modules from ontologies.* In Proc. of the 16th Int. World Wide Web Conf., 2007, 717–727.
5. B.C. Grau, I. Horrocks, Y. Kazakov, and U. Sattler: *A logical framework for modularity of ontologies.* In Proc. of the 20th Int. Joint Conf. on Artificial Intelligence, 2007.
6. B.C. Grau, I. Horrocks, Y. Kazakov, and U. Sattler: *Modular reuse of ontologies: theory and practice.* Journal of Artificial Intelligence Research, 31, 2008, 273–318.
7. E. Jimenez-Ruiz, B.C. Grau, T. Schneider, U. Sattler, and R. Berlanga: *Safe and economic re-use of ontologies: a logic-based methodology and tool support.* In Proc. of the 5th European Semantic Web Conf., Springer LNCS, 2008.
8. 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.
9. G. Antoniou and A. Kehagas: *A note on the the refinement of ontologies. Int. Journal of Intelligent Systems*, 15, 7, 2000, 623–632.
10. C. Lutz and F. Wolter: *Conservative extensions in the lightweight description logic EL.* In Proc. of the 21th Conf. on Automated Deduction, Springer-Verlag, 2007, volume 4603, 84–99.
11. C. Lutz, D. Walther, and F. Wolter: *Conservative extensions in expressive description logics.* In M. Veloso, editor, Proc. of the 20th Int. Joint Conf. on Artificial Intelligence, AAAI Press, 2007, 453–458.
12. S. Ghilardi, C. Lutz, and F. Wolter: *Did I damage my ontology? A case for conservative extensions in description logics.* In P. Doherty, J. Mylopoulos, and C. Welty, editors, Proc. of the 10th Int. Conf. on Principles of Knowledge Representation and Reasoning, AAAI Press, 2006, 187–197.
13. J. Bao and V. Honavar: *Ontology language extensions to support localized semantics, modular reasoning, and collaborative ontology design and ontology reuse.* Technical report, Iowa State University, 2004.
14. P. Bouquet, F. Giunchiglia, F. van Harmelen, L. Serafini, and H. Stuckenschmidt: *C-OWL: Contextualizing Ontologies.* In Proc. of the Second Int. Semantic Web Conf., Springer, 2003, 164–179.
15. B.C. Grau, B. Parsia, and E. Sirin: *Combining OWL ontologies using $\mathcal{E}$-Connections.* Web Semantics: Science, Services and Agents on the World Wide Web, 4, 1, 2006, 40–59.
16. J.Z. Pan, L. Serafini, and Y. Zhao: *Semantic import: an approach for partial ontology reuse.* In P. Haase, V. Honavar, O. Kutz, Y. Sure, and A. Tamilin, editors, Proc. of the 1st Int. Workshop on Modular Ontologies, volume 232, CEUR-WS.org, 2006.
17. M. Boris, B.C. Grau, I. Horrocks, Z. Wu, A. Fokoue, and C. Lutz: *OWL 2 Web Ontology Language: Profiles.* W3C draft, W3C, 2009. Available at `http://www.w3.org/TR/2009/CR-owl2-profiles-20090611/`.
18. N.F. Noy and D.L. Mcguinness: *Ontology development 101: a guide to creating your first ontology.* Technical report, Stanford Knowledge Systems Laboratory, 2001. Available at `http://protege.stanford.edu/publications/ontology_development/ontology101-noy-mcguinness.html`.
19. T.R. Gruber: *Toward principles for the design of ontologies used for knowledge sharing* . Int. Journal of Human-Computer Studies, 43, 5-6, 1995, 907–928.
20. N. Guarino: *Formal ontology and information systems.* In N. Guarino, editor, Proc. of the 1st Int. Conf. on Formal Ontologies in Information Systems, IOS Press, 1998, 3–15.
21. B. Swartout, R. Patil, K. Knight, and T. Russ: Toward distributed use of large-scale ontologies. In Proc. of the 10th Knowledge Acquisition for Knowledge-Based Systems Workshop, 1996.
22. N.F. Noy and C.D. Hafner: *The state of the art in ontology design: A survey and comparative review. AI Magazine*, 18, 1997, 53–74.
23. B.C. Grau, B. Parsia, E. Sirin, and A. Kalyanpur: *Modularity and web ontologies.* In P. Doherty, J. Mylopoulos, and C. A. Welty, editors, Proc. the 20th Int. Conf. on Principles of Knowledge Representation and Reasoning, AAAI Press, 2006, 198–209.
24. B.C. Grau, B. Parsia, E. Sirin, and A. Kalyanpur: *Automatic partitioning of OWL ontologies using $\mathcal{E}$-connections.* In Proc. of the 18th Int. Workshop on Description Logics, 2005.
25. B.C. Grau, I. Horrocks, O. Kutz, and U. Sattler: *Will my ontologies fit together?* In Proc. of the 19th Int. Workshop on Description Logics, 2006.
26. B.C. Grau and O. Kutz: *Modular ontology languages revisited.* In Workshop on Semantic Web for Collaborative Knowledge Acquisition at the 20th Int. Joint Conf. on Artificial Intelligence, 2007.

7 Appendix

As shown in [23]:

- for each $\mathcal{SHOIQ}$ concept C, one of the following holds:
  - $C^{\mathcal{I}_{\cup\Delta}} = C^{\mathcal{I}}$ for all $\mathcal{I}$ and Δ, $\Delta \cap \Delta^{\mathcal{I}} = \emptyset$ (such C is said to be *local*);
  - $C^{\mathcal{I}_{\cup\Delta}} = C^{\mathcal{I}} \cup \Delta$ for all $\mathcal{I}$ and Δ, $\Delta \cap \Delta^{\mathcal{I}} = \emptyset$ (C is *non-local*).
- If R is a $\mathcal{SHOIQ}$ role, then $R^{\mathcal{I}_{\cup\Delta}} = R^{\mathcal{I}}$ for all $\mathcal{I}$ and Δ, $\Delta \cap \Delta^{\mathcal{I}} = \emptyset$.
- A $\mathcal{SHOIQ}$ ontology is not local iff it explicitly contain a GCI $D \sqsubseteq C$ such that C is local and D is non-local.

Lemma 1 (Auxiliary). *Let α be a $\mathcal{SHOIQ}$ axiom, $\mathcal{I}$ an interpretation and Δ a set disjoint with $\Delta^{\mathcal{I}}$. Then $\mathcal{I} \not\models \alpha$ implies $\mathcal{I}_{\cup\Delta} \not\models \alpha$.*

Proof. Suppose that $\mathcal{I} \not\models \alpha$.
If α is of the form $C \sqsubseteq D$ that means $C^{\mathcal{I}} \not\sqsubseteq D^{\mathcal{I}}$ ($\star$). For C, D local, we have $C^{\mathcal{I}_{\cup\Delta}} = C^{\mathcal{I}}$ and $D^{\mathcal{I}_{\cup\Delta}} = D^{\mathcal{I}}$. By

($\star$), we get $C^{\mathcal{I}_{\cup\Delta}} \not\sqsubseteq D^{\mathcal{I}_{\cup\Delta}}$. For C local, D non-local, we have $C^{\mathcal{I}_{\cup\Delta}} = C^{\mathcal{I}}$ and $D^{\mathcal{I}_{\cup\Delta}} = D^{\mathcal{I}} \cup \Delta$. By ($\star$) and by the fact that $\Delta \cap C^{\mathcal{I}} = \emptyset$ (because $\Delta \cap \Delta^{\mathcal{I}} = \emptyset$), we get $C^{\mathcal{I}_{\cup\Delta}} \not\sqsubseteq D^{\mathcal{I}_{\cup\Delta}}$. For C non-local, D local, we have $C^{\mathcal{I}_{\cup\Delta}} = C^{\mathcal{I}} \cup \Delta$ and $D^{\mathcal{I}_{\cup\Delta}} = D^{\mathcal{I}}$. By ($\star$), we get $C^{\mathcal{I}_{\cup\Delta}} \not\sqsubseteq D^{\mathcal{I}_{\cup\Delta}}$. For C, D are non-local, we have $C^{\mathcal{I}_{\cup\Delta}} = C^{\mathcal{I}} \cup \Delta$ and $D^{\mathcal{I}_{\cup\Delta}} = D^{\mathcal{I}} \cup \Delta$. By ($\star$) and by the fact that $\Delta \cap C^{\mathcal{I}} = \emptyset$, $\Delta \cap D^{\mathcal{I}} = \emptyset$ (because $\Delta \cap \Delta^{\mathcal{I}} = \emptyset$), we get $C^{\mathcal{I}_{\cup\Delta}} \not\sqsubseteq D^{\mathcal{I}_{\cup\Delta}}$. For each of the four possible cases we showed that $\mathcal{I}_{\cup\Delta} \not\models C \sqsubseteq D$.

The remaining types of $\mathcal{SHOIQ}$ axioms ($C \equiv D$, $\mathsf{Trans}(r)$, $R \sqsubseteq S$) can be treated in the same way. $\square$

Proof (of Proposition 1.). The proposition is obviously true for inconsistent ontologies.
Assume that a consistent ontology $\mathcal{O}$ restricts the meaning of the top concept. Then, by Definition 4, $\mathcal{O} \models \top \sqsubseteq C$ for some C of the form $A_1 \sqcup \ldots \sqcup A_n \sqcup \exists r_1.\top \sqcup \ldots \sqcup \exists r_m.\top \sqcup \exists s_1^-.\top \sqcup \ldots \sqcup \exists s_k^-.\top \sqcup \{a_1, \ldots, a_l\}$. Take any model $\mathcal{I}$ of $\mathcal{O}$ and any $x \notin \Delta^{\mathcal{I}}$. As $\top^{\mathcal{I}_{\cup\{x\}}} = \Delta^{\mathcal{I}} \cup \{x\}$ and $C^{\mathcal{I}_{\cup\{x\}}} = \Delta^{\mathcal{I}}$, $\mathcal{I}_{\cup\{x\}} \not\models \top \sqsubseteq C$, and thus $\mathcal{I}$ is not a model of $\mathcal{O}$. This shows $\mathcal{O}$ is not local.

Assume a consistent $\mathcal{SHOIQ}$ ontology $\mathcal{O}$ does not restrict the meaning of the top concept. Then, by Definition 4, $\mathcal{O} \not\models \top \sqsubseteq C$ for C of the form $A_1 \sqcup \ldots \sqcup A_n \sqcup \exists r_1.\top \sqcup \ldots \sqcup \exists r_m.\top \sqcup \exists r_1^-.\top \sqcup \ldots \sqcup \exists r_m^-.\top \sqcup \{a_1, \ldots, a_l\}$, where $A_1, \ldots, A_n$, $r_1, \ldots, r_m$ and $a_1, \ldots, a_l$ are all atomic concepts, atomic roles and individuals in $\mathbf{Sig}(\mathcal{O})$. Since $\mathcal{O} \not\models \top \sqsubseteq C$, there exists a model $\mathcal{I}$ of $\mathcal{O}$ such that $\mathcal{I} \not\models \top \sqsubseteq C$. Pick any such model $\mathcal{I}$. Since $\mathcal{I} \not\models \top \sqsubseteq C$, there exists an object $x \in \Delta^{\mathcal{I}}$ such that x do not participate in the interpretation $X^{\mathcal{I}}$ of any atomic concept, atomic role and individual X in $\mathbf{Sig}(\mathcal{O})$. Observe that: for all local $\mathcal{SHOIQ}$ concepts C_1 composed of the symbols from $\mathbf{Sig}(\mathcal{O})$, $x \notin C_1^{\mathcal{I}}$ holds; for all non-local $\mathcal{SHOIQ}$ concepts C_2 composed of the symbols from $\mathbf{Sig}(\mathcal{O})$, $x \in C_2^{\mathcal{I}}$ holds. Therefore, $\mathcal{O}$ does not contain a GCI of the form $C_2 \sqsubseteq C_1$ (otherwise $\mathcal{I}$ were not its model) and thus is local. $\square$

Proof (of Proposition 2.). Let $\mathcal{O}'$ be an arbitrary local $\mathcal{SHOIQ}$ ontology with $\mathbf{Sig}(\mathcal{O}) \cap \mathbf{Sig}(\mathcal{O}') \subseteq \mathbf{S}$. We need to show that $\mathcal{O} \cup \mathcal{O}'$ is a conservative extension of $\mathcal{O}'$ w.r.t. $\mathcal{SHOIQ}$.

Assume (for contradiction) that there exists a $\mathcal{SHOIQ}$ axiom α with $\mathbf{Sig}(\alpha) \subseteq \mathbf{Sig}(\mathcal{O}')$ for which both $\mathcal{O}' \not\models \alpha$ ($\star$) and $\mathcal{O} \cup \mathcal{O}' \models \alpha$ ($*$) hold.
By ($\star$), there is a model $\mathcal{J}$ of $\mathcal{O}'$ such that $\mathcal{J} \not\models \alpha$. ($\diamond$)

The conditions of the proposition imply the existence of a model $\mathcal{I}$ of $\mathcal{O}$ such that $\Delta^{\mathcal{I}} = \Delta^{\mathcal{J}} \cup \Delta$, $\Delta \cap \Delta^{\mathcal{J}} = \emptyset$ and $X^{\mathcal{I}} = X^{\mathcal{J}}$ for all $X \in \mathbf{S}$. Pick any interpretation $\mathcal{K}$ satisfying: $\Delta^{\mathcal{K}} = \Delta^{\mathcal{I}}$, $X^{\mathcal{K}} = X^{\mathcal{I}}$ for all $X \in \mathbf{Sig}(\mathcal{O})$, $X^{\mathcal{K}} = X^{\mathcal{J}}$ for all $X \in \mathbf{Sig}(\mathcal{O}')$. Such an interpretation exists as $X^{\mathcal{I}} = X^{\mathcal{J}}$ for $X \in \mathbf{Sig}(\mathcal{O}) \cap \mathbf{Sig}(\mathcal{O}')$, and symbols not occurring in $\mathbf{Sig}(\mathcal{O}) \cap \mathbf{Sig}(\mathcal{O}')$ can be interpreted arbitrarily.

First, because $\mathcal{I} \models \mathcal{O}$ and $\mathcal{K}|_{\mathbf{Sig}(\mathcal{O})} = \mathcal{I}|_{\mathbf{Sig}(\mathcal{O})}$, $\mathcal{K} \models \mathcal{O}$ holds. Second, because $\mathcal{J} \models \mathcal{O}'$ and $\mathcal{O}'$ is local, we have $\mathcal{J}_{\cup\Delta} \models \mathcal{O}'$ and consequently, since $\mathcal{K}|_{\mathbf{Sig}(\mathcal{O}')} = \mathcal{J}_{\cup\Delta}|_{\mathbf{Sig}(\mathcal{O}')}$, $\mathcal{K} \models \mathcal{O}'$. Therefore $\mathcal{K} \models \mathcal{O} \cup \mathcal{O}'$.

Since $\mathcal{J} \not\models \alpha$, we have, using Lemma 1, that $\mathcal{J}_{\cup\Delta} \not\models \alpha$. Furthermore, since $\mathcal{K}|_{\mathbf{Sig}(\mathcal{O}')} = \mathcal{J}_{\cup\Delta}|_{\mathbf{Sig}(\mathcal{O}')}$ and $\mathbf{Sig}(\alpha) \subseteq \mathbf{Sig}(\mathcal{O}')$, $\mathcal{K} \not\models \alpha$.

We showed there exists a model of $\mathcal{O} \cup \mathcal{O}'$ that is not a model of α, which yields a contradiction with the assumption ($*$). $\square$

Proof (of Proposition 3.). Same proof as of Proposition 2 goes through - we only need to replace the sentence labeled with ($\diamond$) with the following: Because ($\star$) and because $\mathcal{SHOQ}$ has the FMP, there exists a finite model $\mathcal{J}$ of $\mathcal{O}'$ such that $\mathcal{J} \not\models \alpha$. $\square$

Proof (of Proposition 4.). Let $\mathcal{J}$ be a finite interpretation.

Let us associate with every $x \in \Delta^{\mathcal{J}}$ ($\Delta^{\mathcal{J}}$ is finite) an unique $a_x \in \mathbf{I}$, $a_x \notin \mathbf{Sig}(\mathcal{O})$ (i.e., different elements are associated with different individuals). For every $x \in \Delta^{\mathcal{J}}$, let us set $\mathbf{S}_x = \{A \in \mathbf{S};\ x \in A^{\mathcal{J}}\}$.

The conditions of the proposition imply that for every $x \in \Delta^{\mathcal{J}}$ there exists a model $\mathcal{I}_x$ of

$$\mathcal{O} \cup \{A \equiv \{a_x\}\}_{A \in \mathbf{S}_x} \cup \{A \equiv \bot\}_{A \in \mathbf{S} - \mathbf{S}_x}.$$

Pick some interpretation $\tilde{\mathcal{I}}$ such that $\tilde{\mathcal{I}} = \biguplus_{x \in \Delta^{\mathcal{J}}} \mathcal{I}_x$ and some interpretation $\mathcal{I}$ isomorphic with $\tilde{\mathcal{I}}$ such that ${a_x}^{\mathcal{I}} = x$ for all $x \in \Delta^{\mathcal{J}}$ (to get such interpretation $\mathcal{I}$ it is enough to "rename" finitely many elements in $\Delta^{\tilde{\mathcal{I}}}$).

Since $\mathcal{I}_x \models \mathcal{O}$ for all $x \in \Delta^{\mathcal{J}}$, $\mathcal{O}$ is a $\mathcal{SHIQ}$ ontology, $\mathcal{SHIQ}$ has the DUMP, $\Delta^{\mathcal{J}}$ is finite, we have $\tilde{\mathcal{I}} \models \mathcal{O}$ and consequently, since $\mathcal{I} \cong \tilde{\mathcal{I}}$, $\mathcal{I} \models \mathcal{O}$ ($\star$).

As ${a_x}^{\mathcal{I}} = x$ for all $x \in \Delta^{\mathcal{J}}$, we have $\Delta^{\mathcal{J}} \subseteq \Delta^{\mathcal{I}}$ ($*$).

It is easy to see that for $A \in \mathbf{S}$ and $x \in \Delta^{\mathcal{J}}$ the following holds: $A^{\mathcal{I}_x} = \{{a_x}^{\mathcal{I}_x}\}$ if $x \in A^{\mathcal{J}}$; $A^{\mathcal{I}_x} = \emptyset$ if $x \notin A^{\mathcal{J}}$. Thus, for $A \in \mathbf{S}$ we have $A^{\tilde{\mathcal{I}}} = \bigcup_{x \in A^{\mathcal{J}}} \{{a_x}^{\tilde{\mathcal{I}}}\}$ and, consequently, $A^{\mathcal{I}} = \bigcup_{x \in A^{\mathcal{J}}} \{{a_x}^{\mathcal{I}}\} = \bigcup_{x \in A^{\mathcal{J}}} \{x\}$, and thus $A^{\mathcal{I}} = A^{\mathcal{J}}$ ($\diamond$).

We showed that, for any finite interpretation $\mathcal{J}$, there exists an interpretation $\mathcal{I}$ satisfying ($\star, *, \diamond$). Using Proposition 3 we have that $\mathcal{O}$ is locally safe for $\mathbf{S}$ w.r.t. $\mathcal{SHOQ}$. $\square$

Proof (of Proposition 5.). Consider an $\mathcal{ALO}$ ($\mathcal{ELO}$) ontology $\mathcal{O}' = \{A \equiv \{a\}\}_{A \in \tilde{\mathbf{S}}} \cup \{A \equiv \bot\}_{A \in \mathbf{S} - \tilde{\mathbf{S}}}$, which evidently is local, satisfies $\mathbf{Sig}(\mathcal{O}) \cap \mathbf{Sig}(\mathcal{O}') \subseteq \mathbf{S}$ and is consistent ($\mathcal{O}' \not\models \top \sqsubseteq \bot$). The conditions of the proposition say that the ontology $\mathcal{O} \cup \mathcal{O}'$ is inconsistent ($\mathcal{O} \cup \mathcal{O}' \models \top \sqsubseteq \bot$). We showed that there exists a local $\mathcal{ALO}$ ($\mathcal{ELO}$) ontology $\mathcal{O}'$ satisfying $\mathbf{Sig}(\mathcal{O}) \cap \mathbf{Sig}(\mathcal{O}') \subseteq \mathbf{S}$ for which $\mathcal{O}$ is not safe. $\square$