

## Repairing Ontologies via Axiom Weakening

Nicolas Troquard, Roberto Confalonieri, Pietro Galliani,  
Rafael Peñaloza, Daniele Porello, Oliver Kutz

Faculty of Computer Science  
Free University of Bozen-Bolzano  
Piazza Domenicani, 3  
I-39100 Bozen-Bolzano BZ, Italy

### Abstract

Ontology engineering is a hard and error-prone task, in which small changes may lead to errors, or even produce an inconsistent ontology. As ontologies grow in size, the need for automated methods for repairing inconsistencies while preserving as much of the original knowledge as possible increases. Most previous approaches to this task are based on removing a few axioms from the ontology to regain consistency. We propose a new method based on *weakening* these axioms to make them less restrictive, employing the use of refinement operators. We introduce the theoretical framework for weakening DL ontologies, propose algorithms to repair ontologies based on the framework, and provide an analysis of the computational complexity. Through an empirical analysis made over real-life ontologies, we show that our approach preserves significantly more of the original knowledge of the ontology than removing axioms.

### Introduction

Ontology engineering is a hard and error-prone task, where even small changes may lead to unforeseen errors, in particular to inconsistency. Ontologies are not only growing in size, they are also increasingly being used in a variety of AI and NLP applications, e.g., (Bateman et al. 2010; Prestes et al. 2013). At the same time, methods to generate ontologies through automated methods gain popularity: e.g., ontology learning (Lehmann and Hitzler 2010; Sazonau, Sattler, and Brown 2015), extraction from web resources such as DBpedia (Auer et al. 2007), or the combination of knowledge from different sources (Stuckenschmidt, Parent, and Spaccapietra 2009).

Such ontology generation methods are all likely to require ontology repair and refinement steps, and trying to repair an ontology containing hundreds, or even thousands of axioms by hand is infeasible. For these reasons, it has become fundamental to develop automated methods for repairing ontologies while preserving as much of the original knowledge as possible.

Most existing ontology repair approaches are based on removing a few axioms to expel the errors (Schlobach and Cornet 2003; Kalyanpur et al. 2005; 2006; Baader, Peñaloza,

and Suntisrivaraporn 2007). While these methods are effective, and have been used in practice, they have the side effect of removing also many potentially wanted implicit consequences. In this paper,<sup>1</sup> we propose a more fine-grained method for ontology repair based on *weakening* axioms, thus making them more general. The idea is that, through this weakening, more of the original knowledge is preserved; that is, our method is less destructive.

We show, both theoretically and empirically, that axiom weakening is a powerful approach for repairing ontologies. On the theoretical side, we prove that the computational complexity of this task is not greater than that of the standard reasoning tasks in description logics. Empirically, we compare the results of weakening axioms against deleting them, over existing ontologies developed in the life sciences. This comparison shows that our approach preserves significantly more of the original ontological knowledge than removing axioms, based on an evaluation measure inspecting the preservation of taxonomic structure (see e.g., (Alani, Brewster, and Shadbolt 2006; Resnik 1999) for related measures).

The main result of this paper is to present a new ontology repair methodology capable of preserving most of the original knowledge, without incurring any additional costs in terms of computational complexity. By thereby preserving more implicit consequences of the ontology, our methodology also provides a contribution to the ontology development cycle (Neuhaus et al. 2013). Indeed, it can be a useful tool for test-driven ontology development, where the preservation of the entailment of competency questions from the weakened ontology can be seen as a measure for the quality of the repair (Grüninger and Fox 1995; Ren et al. 2014).

We begin by outlining formal preliminaries, including the introduction of refinement operators and a basic analysis of properties of both, specialisation and generalisation operators. This is followed by a complexity analysis of the problem of computing weakened axioms in our approach. We then present several variations of repair algorithms, a detailed empirical evaluation of their performance, and a quality analysis of the returned ontologies. We close with a dis-

<sup>1</sup>An extended version of the paper is available at <https://arxiv.org/abs/1711.03430>.

cussion of related work and an outlook to future extensions and refinements of the presented ideas.

## Preliminaries

From a formal point of view, an ontology is a set of formulas in an appropriate logical language with the purpose of describing a particular domain of interest. The precise logic used is in fact not crucial for our approach as most techniques introduced apply to a variety of logics; however, for the sake of clarity we use description logics (DLs) as well-known examples of ontology languages. We briefly introduce the basic DL  $\mathcal{ALC}$ ; for full details see (Baader et al. 2003). The syntax of  $\mathcal{ALC}$  is based on two disjoint sets  $N_C$  and  $N_R$  of *concept names* and *role names*, respectively. The set of  $\mathcal{ALC}$  concepts is generated by the grammar

$$C ::= A \mid \neg C \mid C \sqcap C \mid C \sqcup C \mid \forall R.C \mid \exists R.C ,$$

where  $A \in N_C$  and  $R \in N_R$ . A *TBox* is a finite set of concept inclusions (GCIs) of the form  $C \sqsubseteq D$  where  $C$  and  $D$  are concepts. It is used to store terminological knowledge regarding the relationships between concepts. An *ABox* is a finite set of formulas of the form  $C(a)$  and  $R(a, b)$ , which express knowledge about objects in the knowledge domain.

The semantics of  $\mathcal{ALC}$  is defined through *interpretations*  $I = (\Delta^I, \cdot^I)$ , where  $\Delta^I$  is a non-empty *domain*, and  $\cdot^I$  is a function mapping every individual name to an element of  $\Delta^I$ , each concept name to a subset of the domain, and each role name to a binary relation on the domain. The interpretation  $\mathcal{I}$  is a *model* of the TBox  $\mathcal{T}$  if it satisfies all the GCIs in  $\mathcal{T}$ . Given two concepts  $C$  and  $D$ , we say that  $C$  is *subsumed* by  $D$  w.r.t. the TBox  $\mathcal{T}$  ( $C \sqsubseteq_{\mathcal{T}} D$ ) if  $C^I \subseteq D^I$  for every model  $I$  of  $\mathcal{T}$ . We write  $C \equiv_{\mathcal{T}} D$  when  $C \sqsubseteq_{\mathcal{T}} D$  and  $D \sqsubseteq_{\mathcal{T}} C$ .  $C$  is *strictly subsumed* by  $D$  w.r.t.  $\mathcal{T}$  ( $C \sqsubset_{\mathcal{T}} D$ ) if  $C \sqsubseteq_{\mathcal{T}} D$  and  $C \not\equiv_{\mathcal{T}} D$ .

$\mathcal{EL}$  is the restriction of  $\mathcal{ALC}$  allowing only conjunctions and existential restrictions (Baader, Brandt, and Lutz 2005). It is widely used in biomedical ontologies for describing large terminologies since classification can be computed in polynomial time.<sup>2</sup> In the following,  $\mathcal{DL}$  denotes either  $\mathcal{ALC}$  or  $\mathcal{EL}$ , and  $\mathcal{L}(\mathcal{DL}, N_C, N_R)$  denotes the set of (complex) concepts that can be built over  $N_C$  and  $N_R$  in  $\mathcal{DL}$ .

**Definition 1.** Let  $\mathcal{T}$  be a  $\mathcal{DL}$  TBox with concept names from  $N_C$ . The set of *subconcepts* of  $\mathcal{T}$  is given by

$$\text{sub}(\mathcal{T}) = \{\top, \perp\} \cup \bigcup_{C \sqsubseteq D \in \mathcal{T}} \text{sub}(C) \cup \text{sub}(D) ,$$

where for  $C \in N_C \cup \{\top, \perp\}$ , and  $\text{sub}(C)$  is the set of subconcepts in  $C$ .

The size  $|C|$  of a concept  $C$  is the size of its syntactic tree where for every role  $R$ ,  $\exists R$ . and  $\forall R$ . are individual nodes.

**Definition 2.** The *size*  $|C|$  of a concept  $C$  is inductively defined as follows. For  $C \in N_C \cup \{\top, \perp\}$ ,  $|C| = 1$ . Then,  $|\neg C| = 1 + |C|$ ;  $|C \sqcap D| = |C \sqcup D| = 1 + |C| + |D|$ ; and  $|\exists R.C| = |\forall R.C| = 1 + |C|$ .

<sup>2</sup>The OWL 2 EL profile significantly extends the basic  $\mathcal{EL}$  logic whilst maintaining its desirable polynomial time complexity, see <https://www.w3.org/TR/owl2-profiles/>.

The *size*  $|\mathcal{T}|$  of the TBox  $\mathcal{T}$  is  $\sum_{C \sqsubseteq D \in \mathcal{T}} (|C| + |D|)$ . Clearly, for every  $C$  we have  $\text{card}(\text{sub}(C)) \leq |C|$  and for every TBox  $\mathcal{T}$  we have  $\text{card}(\text{sub}(\mathcal{T})) \leq |\mathcal{T}| + 2$ .

We now define the upward and downward cover sets of concept names. Intuitively, the upward set of the concept  $C$  collects the most specific subconcepts of the TBox  $\mathcal{T}$  that subsume  $C$ ; conversely, the downward set of  $C$  collects the most general subconcepts from  $\mathcal{T}$  subsumed by  $C$ . The concepts in  $\text{sub}(\mathcal{T})$  are *some* concepts that are relevant in the context of  $\mathcal{T}$ , and that are used as building blocks for generalisations and specialisations. The properties of  $\text{sub}(\mathcal{T})$  guarantee that the upward and downward cover sets are finite.

**Definition 3.** Let  $\mathcal{T}$  be a  $\mathcal{DL}$  TBox and  $C$  a concept. The *upward cover* and *downward cover* of  $C$  w.r.t.  $\mathcal{T}$  are:

$$\begin{aligned} \text{UpCov}_{\mathcal{T}}(C) &:= \{D \in \text{sub}(\mathcal{T}) \mid C \sqsubseteq_{\mathcal{T}} D \text{ and} \\ &\quad \nexists D' \in \text{sub}(\mathcal{T}) \text{ with } C \sqsubset_{\mathcal{T}} D' \sqsubset_{\mathcal{T}} D\}, \\ \text{DownCov}_{\mathcal{T}}(C) &:= \{D \in \text{sub}(\mathcal{T}) \mid D \sqsubseteq_{\mathcal{T}} C \text{ and} \\ &\quad \nexists D' \in \text{sub}(\mathcal{T}) \text{ with } D \sqsubset_{\mathcal{T}} D' \sqsubset_{\mathcal{T}} C\}. \end{aligned}$$

Observe that  $\text{UpCov}_{\mathcal{T}}$  and  $\text{DownCov}_{\mathcal{T}}$  miss interesting refinements. Note also that this definition only returns meaningful results when used with a consistent ontology; otherwise it returns the whole set  $\text{sub}(\mathcal{T})$ . Hence, when dealing with the repair problem of an inconsistent ontology  $O$ , we need a derived, consistent ‘reference ontology’  $O^{\text{ref}}$  to steer the repair process; this is outlined in greater detail in the section on repairing ontologies.

**Example 4.** Let  $A, B, C \in N_C$  and  $\mathcal{T} = \{A \sqsubseteq B\}$ . We have  $\text{UpCov}_{\mathcal{T}}(A \sqcap C) = \{A\}$ . Iterating, we get  $\text{UpCov}_{\mathcal{T}}(A) = \{A, B\}$  and  $\text{UpCov}_{\mathcal{T}}(B) = \{B, \top\}$ . We could reasonably expect  $B \sqcap C$  to be also a generalisation of  $A \sqcap C$  w.r.t.  $\mathcal{T}$  but it will be missed by the iterated application of  $\text{UpCov}_{\mathcal{T}}$ . Similarly,  $\text{UpCov}_{\mathcal{T}}(\exists R.A) = \{\top\}$ , while we can expect  $\exists R.B$  to be a generalisation of  $\exists R.A$ .

To take care of these omissions, we introduce a generalisation and specialisation operator. We denote as  $\text{nfn}(C)$  the negation normal form of the concept  $C$ . Let  $\uparrow$  and  $\downarrow$  be two functions from  $\mathcal{L}(\mathcal{DL}, N_C, N_R)$  to the powerset of  $\mathcal{L}(\mathcal{DL}, N_C, N_R)$ . We define  $\zeta_{\uparrow, \downarrow}$ , the *abstract refinement operator*, by induction on the structure of concept descriptions as shown in Table 1. Complying with the previous observation, we define two concrete refinement operators from the abstract operator  $\zeta_{\uparrow, \downarrow}$ .

**Definition 5.** The *generalisation operator* and *specialisation operator* are defined, respectively, as

$$\begin{aligned} \gamma_{\mathcal{T}} &= \zeta_{\text{UpCov}_{\mathcal{T}}, \text{DownCov}_{\mathcal{T}}} , \text{ and} \\ \rho_{\mathcal{T}} &= \zeta_{\text{DownCov}_{\mathcal{T}}, \text{UpCov}_{\mathcal{T}}} . \end{aligned}$$

Returning to our example, notice that for  $\mathcal{T} = \{A \sqsubseteq B\}$ , we now have  $\gamma_{\mathcal{T}}(A \sqcap C) = \{B \sqcap C, A \sqcap \top, A\}$ .

**Definition 6.** Given a  $\mathcal{DL}$  concept  $C$ , its  *$i$ -th refinement iteration* by means of  $\zeta_{\uparrow, \downarrow}$  (viz.,  $\zeta_{\uparrow, \downarrow}^i(C)$ ) is inductively defined as follows:

- $\zeta_{\uparrow, \downarrow}^0(C) = \{C\}$ ;
- $\zeta_{\uparrow, \downarrow}^{j+1}(C) = \zeta_{\uparrow, \downarrow}^j(C) \cup \bigcup_{C' \in \zeta_{\uparrow, \downarrow}^j(C)} \zeta_{\uparrow, \downarrow}(C')$ ,  $j \geq 0$ .

Table 1: Abstract refinement operator

$$\begin{aligned}
\zeta_{\uparrow,\downarrow}(A) &= \uparrow(A) \\
\zeta_{\uparrow,\downarrow}(\neg A) &= \{\text{nnf}(\neg C) \mid C \in \downarrow(A)\} \cup \uparrow(\neg A) \\
\zeta_{\uparrow,\downarrow}(\top) &= \uparrow(\top) \\
\zeta_{\uparrow,\downarrow}(\perp) &= \uparrow(\perp) \\
\zeta_{\uparrow,\downarrow}(C \sqcap D) &= \{C' \sqcap D \mid C' \in \zeta_{\uparrow,\downarrow}(C)\} \cup \\
&\quad \{C \sqcap D' \mid D' \in \zeta_{\uparrow,\downarrow}(D)\} \cup \uparrow(C \sqcap D) \\
\zeta_{\uparrow,\downarrow}(C \sqcup D) &= \{C' \sqcup D \mid C' \in \zeta_{\uparrow,\downarrow}(C)\} \cup \\
&\quad \{C \sqcup D' \mid D' \in \zeta_{\uparrow,\downarrow}(D)\} \cup \uparrow(C \sqcup D) \\
\zeta_{\uparrow,\downarrow}(\forall R.C) &= \{\forall R.C' \mid C' \in \zeta_{\uparrow,\downarrow}(C)\} \cup \uparrow(\forall R.C) \\
\zeta_{\uparrow,\downarrow}(\exists R.C) &= \{\exists R.C' \mid C' \in \zeta_{\uparrow,\downarrow}(C)\} \cup \uparrow(\exists R.C)
\end{aligned}$$

The set of all concepts reachable from  $C$  by means of  $\zeta_{\uparrow,\downarrow}$  in a finite number of steps is  $\zeta_{\uparrow,\downarrow}^*(C) = \bigcup_{i \geq 0} \zeta_{\uparrow,\downarrow}^i(C)$ .

Some basic properties about  $\gamma_{\mathcal{T}}$  and  $\rho_{\mathcal{T}}$  follow.

**Lemma 7.** *For every TBox  $\mathcal{T}$ :*

1. **generalisation:** if  $X \in \gamma_{\mathcal{T}}(C)$  then  $C \sqsubseteq_{\mathcal{T}} X$   
**specialisation:** if  $X \in \rho_{\mathcal{T}}(C)$  then  $X \sqsubseteq_{\mathcal{T}} C$
2. **reflexivity:** if  $C \in \text{sub}(\mathcal{T})$  then  $C \in \text{UpCov}_{\mathcal{T}}(C)$  and  $C \in \text{DownCov}_{\mathcal{T}}(C)$
3. **semantic stability of cover:** if  $C_1 \equiv_{\mathcal{T}} C_2$  then  $C_1 \in \text{UpCov}_{\mathcal{T}}(C)$  iff  $C_2 \in \text{UpCov}_{\mathcal{T}}(C)$  and  $C_1 \in \text{DownCov}_{\mathcal{T}}(C)$  iff  $C_2 \in \text{DownCov}_{\mathcal{T}}(C)$
4. **relevant completeness:**  $\text{UpCov}_{\mathcal{T}}(C) \subseteq \gamma_{\mathcal{T}}(C)$  and  $\text{DownCov}_{\mathcal{T}}(C) \subseteq \rho_{\mathcal{T}}(C)$
5. **generalisability:** if  $C, D \in \text{sub}(\mathcal{T})$  and  $C \sqsubseteq_{\mathcal{T}} D$  then  $D \in \gamma_{\mathcal{T}}^*(C)$   
**specialisability:** if  $C, D \in \text{sub}(\mathcal{T})$  and  $D \sqsubseteq_{\mathcal{T}} C$  then  $D \in \rho_{\mathcal{T}}^*(C)$
6. **trivial generalisability:**  $\top \in \gamma_{\mathcal{T}}^*(C)$   
**falsehood specialisability:**  $\perp \in \rho_{\mathcal{T}}^*(C)$
7. **generalisation finiteness:**  $\gamma_{\mathcal{T}}(C)$  is finite  
**specialisation finiteness:**  $\rho_{\mathcal{T}}(C)$  is finite

Although  $\gamma_{\mathcal{T}}(C)$  and  $\rho_{\mathcal{T}}(C)$  are always finite (see Lemma 7.7), this is not the case for  $\gamma_{\mathcal{T}}^*(C)$  and  $\rho_{\mathcal{T}}^*(C)$ . Indeed, their iterated application can produce an infinite chain of refinements.

**Example 8.** If  $\mathcal{T} = \{A \sqsubseteq \exists r.A\}$ , then  $\gamma_{\mathcal{T}}(A) = \{A, \exists r.A\}$ . Thus  $\gamma_{\mathcal{T}}(\exists r.A) = \{\exists r.A, \exists r.\exists r.A\} \cup \{\top\}$  (notice that  $\top \in \gamma_{\mathcal{T}}^2(A)$ ). Continuing the iteration of  $\gamma_{\mathcal{T}}$  on  $A$ , we get  $(\exists r.)^k A \in \gamma_{\mathcal{T}}^k(A)$  for every  $k \geq 0$ .

This is not a feature caused by the existential quantification alone. Similar examples exist that involve universal quantification, disjunction, and conjunction.<sup>3</sup> Notice that although the covers of two provably equivalent concepts are the same (Lemma 7.3), it is not the case that

<sup>3</sup>From the perspective of ontology repair, infinite refinement chains are not an issue since there are always finite chains (Lemma 7.6). If needed, it can be simply circumvented by imposing a bound on the size of the considered refinements.

$\gamma_{\mathcal{T}}(C_1) = \gamma_{\mathcal{T}}(C_2)$  whenever  $C_1 \equiv_{\mathcal{T}} C_2$ . For example, with the TBox  $\mathcal{T} = \{A \sqsubseteq B\}$ , we have  $\gamma_{\mathcal{T}}(A) = \{A, B\}$  and  $\gamma_{\mathcal{T}}(\top \sqcap A) = \{\top \sqcap A, \top \sqcap B, A, B\}$ .

## Complexity

We now analyse the computational aspects of the refinement operators.

**Definition 9.** Given a TBox  $\mathcal{T}$  and concepts  $C, D$ , the problems  $\gamma_{\mathcal{T}}$ -MEMBERSHIP and  $\rho_{\mathcal{T}}$ -MEMBERSHIP ask whether  $D \in \gamma_{\mathcal{T}}(C)$  and  $D \in \rho_{\mathcal{T}}(C)$ , respectively.

We show that  $\gamma_{\mathcal{T}}$  and  $\rho_{\mathcal{T}}$  are efficient refinement operators, in the sense that deciding  $\gamma_{\mathcal{T}}$ -MEMBERSHIP and  $\rho_{\mathcal{T}}$ -MEMBERSHIP is not harder than deciding (atomic) concept subsumption in the underlying logic. Recall that subsumption is ExpTime-complete in  $\mathcal{ALC}$  and PTime-complete in  $\mathcal{EL}$ . We show that the same complexity bounds hold for  $\gamma_{\mathcal{T}}$ -MEMBERSHIP.

For proving hardness, we first show that deciding whether  $C' \in \text{UpCover}_{\mathcal{T}}(C)$  is as hard as atomic concept subsumption (Theorem 13). Then we show that  $\gamma_{\mathcal{T}}$ -MEMBERSHIP is just as hard (Theorem 17). For the upper bounds, we first establish the complexity of computing the set  $\text{UpCover}_{\mathcal{T}}(C)$  (Theorem 15). We then show that we can decide  $\gamma_{\mathcal{T}}$ -MEMBERSHIP resorting to at most a linear number of computations  $\text{UpCover}_{\mathcal{T}}(C')$  (Theorem 18). Combining Theorem 17 and Theorem 18, we obtain the result.

**Theorem 10.**  $\gamma_{\mathcal{T}}$ -MEMBERSHIP is ExpTime-complete for  $\mathcal{ALC}$  and PTime-complete for  $\mathcal{EL}$ .

Similar arguments can be used to establish the same complexities for  $\rho_{\mathcal{T}}$ -MEMBERSHIP.

**Corollary 11.**  $\rho_{\mathcal{T}}$ -MEMBERSHIP is ExpTime-complete for  $\mathcal{ALC}$  and PTime-complete for  $\mathcal{EL}$ .

The remainder of this section provides the details. We first prove a technical lemma used in the reduction from concept subsumption to deciding whether  $C' \in \text{UpCover}_{\mathcal{T}}(C)$ .

**Lemma 12.** Let  $\mathcal{T}$  be a  $\mathcal{DL}$  TBox and  $X \notin \text{sub}(\mathcal{T})$ . Then, for every model  $\mathcal{I}$  of  $\mathcal{T}' := \mathcal{T} \cup \{X \sqcap B \sqsubseteq \top\}$  there is a model  $\mathcal{J}$  of  $\mathcal{T}'$  such that

1.  $X^{\mathcal{J}} = \emptyset$ , and
2. for every  $C \in \text{sub}(\mathcal{T})$ ,  $C^{\mathcal{I}} = C^{\mathcal{J}}$ .

*Proof.* We define the interpretation  $\mathcal{J}$  where all role names are interpreted as in  $\mathcal{I}$ , and for every concept name  $A \in N_C$

$$A^{\mathcal{J}} := \begin{cases} \emptyset & \text{if } A = X \\ A^{\mathcal{I}} & \text{otherwise.} \end{cases}$$

Since  $X$  only appears in a tautology,  $\mathcal{J}$  is also a model of  $\mathcal{T}'$ . Using induction on the structure of the concepts, it is easy to show that the second condition of the lemma holds.  $\square$

The following theorem is instrumental in the proof of Theorem 17.

**Theorem 13.** Let  $\mathcal{T}$  be a  $\mathcal{DL}$  TBox and let  $C$  be an arbitrary  $\mathcal{DL}$  concept. Deciding whether  $D \in \text{UpCov}_{\mathcal{T}}(C)$  is as hard as deciding atomic subsumption w.r.t. a TBox over  $\mathcal{DL}$ .

*Proof.* We propose a reduction from the problem of deciding atomic subsumption w.r.t. a TBox. Let  $\mathcal{T}$  be a  $\mathcal{DL}$  TBox, and  $A, B$  be two concept names. We assume w.l.o.g. that  $\{A, B\} \subseteq \text{sub}(\mathcal{T})$ . Define the new TBox

$$\mathcal{T}' := \mathcal{T} \cup \{X \sqcap B \sqsubseteq \top\},$$

where  $X$  is a new concept name (not appearing in  $\mathcal{T}$ ).<sup>4</sup> We show that  $A \sqsubseteq_{\mathcal{T}} B$  iff  $X \sqcap B \in \text{UpCov}_{\mathcal{T}'}(X \sqcap A)$ .

[ $\Rightarrow$ ] If  $A \sqsubseteq_{\mathcal{T}} B$ , then  $X \sqcap A \sqsubseteq_{\mathcal{T}} X \sqcap B$ , and hence it also holds  $X \sqcap A \sqsubseteq_{\mathcal{T}'} X \sqcap B$ . Assume that there is some  $E \in \text{sub}(\mathcal{T}')$  with  $X \sqcap A \sqsubset_{\mathcal{T}'} E \sqsubset_{\mathcal{T}'} X \sqcap B$ . Then  $E$  cannot be  $X \sqcap B$ , nor  $X$ . Hence  $E \in \text{sub}(\mathcal{T})$ . Let  $\mathcal{I}$  be an arbitrary model of  $\mathcal{T}'$ . By Lemma 12, there is a model  $\mathcal{J}$  with  $E^{\mathcal{I}} = E^{\mathcal{J}}$  and  $X^{\mathcal{J}} = \emptyset$ . But since by assumption  $E \sqsubseteq_{\mathcal{T}'} X \sqcap B$ , it must be that  $E^{\mathcal{J}} = \emptyset$ , and hence  $E^{\mathcal{I}} = \emptyset$ . It then follows that for every model  $\mathcal{I}$  of  $\mathcal{T}'$ , we have  $E^{\mathcal{I}} = \emptyset$ , which is a contradiction with the assumption  $X \sqcap A \sqsubset_{\mathcal{T}'} E$ . We conclude that  $X \sqcap B \in \text{UpCov}_{\mathcal{T}'}(X \sqcap A)$ .

[ $\Leftarrow$ ] If  $A \not\sqsubseteq_{\mathcal{T}} B$ , there is a model  $\mathcal{I}$  of  $\mathcal{T}$  with  $A^{\mathcal{I}} \not\subseteq B^{\mathcal{I}}$ . We can extend this interpretation to a model  $\mathcal{J}$  of  $\mathcal{T}'$  by setting  $X^{\mathcal{J}} = \Delta^{\mathcal{I}}$ , and  $A^{\mathcal{J}} = A^{\mathcal{I}}$  for all other concept names; and  $r^{\mathcal{J}} = r^{\mathcal{I}}$  for all role names. Then  $(X \sqcap A)^{\mathcal{J}} \not\subseteq (X \sqcap B)^{\mathcal{J}}$ , and hence  $X \sqcap B \notin \text{UpCov}_{\mathcal{T}'}(X \sqcap A)$ .  $\square$

**Theorem 14.** *Let  $\mathcal{T}$  be a  $\mathcal{DL}$  TBox and let  $C$  be an arbitrary  $\mathcal{DL}$  concept. Deciding whether  $D \in \text{UpCov}_{\mathcal{T}}(C)$  can be done in exponential time when  $\mathcal{DL} = \mathcal{ALC}$  and in polynomial time when  $\mathcal{DL} = \mathcal{EL}$ .*

*Proof.* An algorithm goes as follows. If  $D \notin \text{sub}(\mathcal{T})$  or  $C \not\sqsubseteq_{\mathcal{T}} D$ , return **false**. Then, for every  $E \in \text{sub}(\mathcal{T})$ , check whether: (1)  $C \sqsubseteq_{\mathcal{T}} E$ , (2)  $E \sqsubseteq_{\mathcal{T}} D$ , (3)  $E \not\sqsubseteq_{\mathcal{T}} C$ , and (4)  $D \not\sqsubseteq_{\mathcal{T}} E$ . If conditions (1)–(4) are all satisfied, return **false**. Return **true** after trying all  $E \in \text{sub}(\mathcal{T})$ . The routine requires at most  $1+4 \times \text{card}(\text{sub}(\mathcal{T}))$  calls to the subroutine for  $\mathcal{DL}$  concept subsumption. Since  $\text{card}(\text{sub}(\mathcal{T}))$  is linear in  $|\mathcal{T}|$ , the overall routine runs in exponential time when  $\mathcal{DL} = \mathcal{ALC}$  and in polynomial time when  $\mathcal{DL} = \mathcal{EL}$ .  $\square$

The following theorem is instrumental in the proof of Theorem 18.

**Theorem 15.** *Let  $\mathcal{T}$  be a  $\mathcal{DL}$  TBox and let  $C$  be a  $\mathcal{DL}$  concept.  $\text{UpCov}_{\mathcal{T}}(C)$  is computable in exponential time when  $\mathcal{DL} = \mathcal{ALC}$  and in polynomial time when  $\mathcal{DL} = \mathcal{EL}$ .*

*Proof.* It suffices to check for every  $D \in \text{sub}(\mathcal{T})$  whether  $D \in \text{UpCov}_{\mathcal{T}}(C)$  and collect those concepts for which the answer is positive. Since  $\text{card}(\text{sub}(\mathcal{T}))$  is linear in the size of  $\mathcal{T}$ , the result holds.  $\square$

**Lemma 16.** *Let  $\mathcal{T}$  be a  $\mathcal{DL}$  TBox,  $C$  a  $\mathcal{DL}$  concept, and  $X \notin \text{sub}(\mathcal{T})$ . Define  $\mathcal{T}' := \mathcal{T} \cup \{X \equiv C\}$ . If  $D \in \text{sub}(\mathcal{T})$  then  $D \in \text{UpCov}_{\mathcal{T}}(C)$  iff  $D \in \text{UpCov}_{\mathcal{T}'}(C)$ .*

*Proof.* We have  $\text{sub}(\mathcal{T}') = \text{sub}(\mathcal{T}) \cup \{X\}$ . Let  $D \in \text{sub}(\mathcal{T})$ . Suppose  $D \in \text{UpCov}_{\mathcal{T}}(C)$ . Then  $C \sqsubseteq_{\mathcal{T}'} D$  and there is no  $E \in \text{sub}(\mathcal{T}')$  such that  $C \sqsubset_{\mathcal{T}'} E \sqsubset_{\mathcal{T}'} D$ . We thus have

<sup>4</sup>We use this tautology only to ensure that  $X \sqcap B \in \text{sub}(\mathcal{T}')$ , to satisfy the restriction on the definition of the upward cover.

$C \sqsubseteq_{\mathcal{T}} D$ . Since  $\text{sub}(\mathcal{T}) \subset \text{sub}(\mathcal{T}')$  there is no  $E \in \text{sub}(\mathcal{T})$  such that  $C \sqsubset_{\mathcal{T}} E \sqsubset_{\mathcal{T}} D$ .

Let  $D \in \text{UpCov}_{\mathcal{T}}(C)$ . Then  $C \sqsubseteq_{\mathcal{T}} D$  and  $C \sqsubseteq_{\mathcal{T}'} D$ . Moreover, there is no  $E \in \text{sub}(\mathcal{T})$  with  $C \sqsubset_{\mathcal{T}} E \sqsubset_{\mathcal{T}} D$ . So there is no  $E \in \text{sub}(\mathcal{T})$  such that  $C \sqsubset_{\mathcal{T}'} E \sqsubset_{\mathcal{T}'} D$ .

Since  $X \equiv_{\mathcal{T}'} C$ , it is not the case that  $C \sqsubset_{\mathcal{T}'} X \sqsubset_{\mathcal{T}'} D$ . Since  $\text{sub}(\mathcal{T}') = \text{sub}(\mathcal{T}) \cup \{X\}$ , there is no  $E \in \text{sub}(\mathcal{T}')$  such that  $C \sqsubset_{\mathcal{T}'} E \sqsubset_{\mathcal{T}'} D$ . Then  $D \in \text{UpCov}_{\mathcal{T}'}(C)$ .  $\square$

**Theorem 17.** *Deciding  $\gamma_{\mathcal{T}}$ -MEMBERSHIP is as hard as deciding whether  $D \in \text{UpCov}_{\mathcal{T}}(C)$ .*

*Proof.* Let  $\mathcal{T}$  be a  $\mathcal{DL}$  TBox,  $C$  a concept, and  $X \notin \text{sub}(\mathcal{T})$ . Define  $\mathcal{T}' := \mathcal{T} \cup \{X \equiv C\}$ . For every concept  $D \neq X$ , we show that  $D \in \text{UpCov}_{\mathcal{T}}(C)$  iff  $D \in \gamma_{\mathcal{T}'}(X)$ .

By Lemma 16,  $D \in \text{UpCov}_{\mathcal{T}}(C)$  iff  $D \in \text{UpCov}_{\mathcal{T}'}(C)$ . Since  $X \equiv_{\mathcal{T}'} C$ , Lemma 7.3 yields  $D \in \text{UpCov}_{\mathcal{T}'}(C)$  iff  $D \in \text{UpCov}_{\mathcal{T}'}(X)$ . As  $X$  is a concept name, by definition of  $\gamma$  we have  $D \in \text{UpCov}_{\mathcal{T}'}(X)$  iff  $D \in \gamma_{\mathcal{T}'}(X)$ .  $\square$

**Theorem 18.** *Let  $\mathcal{T}$  be a  $\mathcal{DL}$  TBox and  $C$  a concept.  $\gamma_{\mathcal{T}}$ -MEMBERSHIP can be decided in exponential time when  $\mathcal{DL} = \mathcal{ALC}$  and in polynomial time when  $\mathcal{DL} = \mathcal{EL}$ .*

*Proof.* We can decide whether  $\gamma_{\mathcal{T}}(C)$  contains a particular concept by computing only a linear number of times  $\text{UpCov}_{\mathcal{T}}(C')$ , where  $|C'|$  is linearly bounded by  $|C'| + |\mathcal{T}|$ . Theorem 15 tells us that each of these computations can be done in exponential time when  $\mathcal{DL} = \mathcal{ALC}$  and in polynomial time when  $\mathcal{DL} = \mathcal{EL}$ . This yields an exponential time procedure when  $\mathcal{DL} = \mathcal{ALC}$  and a polynomial time procedure when  $\mathcal{DL} = \mathcal{EL}$ .  $\square$

## Repairing Ontologies

Our refinement operators can be used as components of a method for repairing inconsistent ontologies by weakening, instead of removing, problematic axioms.

Given an inconsistent ontology  $O$ , we proceed as described in Algorithm 1. Briefly, we first need to find a consistent subontology  $O^{\text{ref}}$  of  $O$  to serve as *reference ontology* in order to be able to compute a non-trivial upcover and downcover. The *brave* approach (which we use in our evaluation) picks a random maximally consistent subset of  $O$  and chooses it as reference ontology  $O^{\text{ref}}$ . The *cautious* approach takes as  $O^{\text{ref}}$  the intersection of all maximally consistent subsets (Ludwig and Peñaloza 2014; Lembo et al. 2010). While the brave approach is faster to compute and still guarantees to find solutions, the cautious approach has the advantage of not excluding certain repairs a priori. However, it also returns, e.g., a much impoverished upcover.

Once a reference ontology  $O^{\text{ref}}$  has been chosen, and as long as  $O$  is inconsistent, we select a “bad axiom” and replace it with a random weakening of it with respect to  $O^{\text{ref}}$ . In view of evaluation, we consider two variants of the subprocedure  $\text{FindBadAxiom}(O)$ . The first variant (“mis”) randomly samples a number of minimally inconsistent subsets  $I_1, I_2, \dots, I_k \subseteq O$  and returns one axiom from the ones occurring the most often, i.e., an axiom from the set

---

**Algorithm 1** RepairOntologyWeaken( $O$ )

---

$O^{\text{ref}} \leftarrow \text{MaximallyConsistent}(O)$   
**while**  $O$  is inconsistent **do**  
   $\text{BadAx} \leftarrow \text{FindBadAxiom}(O)$   
   $\text{WeakerAx} \leftarrow \text{WeakenAxiom}(\text{BadAx}, O^{\text{ref}})$   
   $O \leftarrow O \setminus \{\text{BadAx}\} \cup \{\text{WeakerAx}\}$   
**end while**  
Return  $O$

---

$\text{argmax}_{\phi \in O}(\text{card}(\{j \mid \phi \in I_j \text{ and } 1 \leq j \leq k\}))$ . The second variant ('rand') of FindBadAxiom( $O$ ) merely returns an axiom in  $O$  at random.

The set of all weakenings of an axiom with respect to a reference ontology is defined as follows:

**Definition 19** (Axiom weakening). Given a subsumption axiom  $C \sqsubseteq D$  of  $O$ , the set of (least) *weakenings* of  $C \sqsubseteq D$  w.r.t.  $O$ , denoted by  $g_O(C \sqsubseteq D)$  is the set of all axioms  $C' \sqsubseteq D'$  such that

$$C' \in \rho_O(C) \text{ and } D' \in \gamma_O(D) .$$

Given an assertional axiom  $C(a)$  of  $O$ , the set of (least) *weakenings* of  $C(a)$ , denoted  $g_O(C(a))$  is the set of all axioms  $C'(a)$  such that

$$C' \in \gamma_O(C) .$$

The subprocedure WeakenAxiom( $\phi, O^{\text{ref}}$ ) randomly returns one axiom in  $g_O(\phi)$ . For every subsumption or assertional axiom  $\phi$ , the axioms in the set  $g_O(\phi)$  are indeed weaker than  $\phi$ .

**Lemma 20.** For every subsumption or assertional axiom  $\phi$ , if  $\phi' \in g_O(\phi)$ , then  $\phi \models_O \phi'$ .

*Proof.* Suppose  $C' \sqsubseteq D' \in g_O(C \sqsubseteq D)$ . Then, by definition of  $g_O$  and Lemma 7.1,  $C' \sqsubseteq C$  and  $D \sqsubseteq D'$  are inferred from  $O$ . Thus, by transitivity of subsumption, we obtain that  $C \sqsubseteq D \models_O C' \sqsubseteq D'$ . For the weakening of assertions, the result follows immediately from Lemma 7.1 again.  $\square$

Clearly, substituting an axiom  $\phi$  with one axiom from  $g_O(\phi)$  cannot diminish the set of interpretations of an ontology. By Lemma 7.6, any subsumption axiom is a finite number of refinement steps away from the trivial axiom  $\perp \sqsubseteq \top$ . Any assertional axiom  $C(a)$  is also a finite number of generalisations away from the trivial assertion  $\top(a)$ . It follows that by repeatedly replacing an axiom with one of its weakenings, the weakening procedure will eventually obtain an ontology with some interpretations. Hence, the algorithm will eventually terminate.

In the next section, we compare Algorithm 1 with Algorithm 2, which merely removes bad axioms until an ontology becomes consistent. We do so for both variants 'mis' and 'rand' of FindBadAxiom( $O$ ). As we will see, Algorithm 1 generally allows us to obtain consistent ontologies which retain significantly more of the informational content of the axioms of the original (and inconsistent) ontology than the ones obtained through Algorithm 2. This is most significant with the 'mis' variant of FindBadAxiom( $O$ ) which reliably pinpoints the problematic axioms.

---

**Algorithm 2** RepairOntologyRemove( $O$ )

---

**while**  $O$  is inconsistent **do**  
   $\text{BadAx} \leftarrow \text{FindBadAxiom}(O)$   
   $O \leftarrow O \setminus \{\text{BadAx}\}$   
**end while**  
Return  $O$

---

---

Abbreviation	Name
bctt	Behaviour Change Technique Taxonomy
co-wheat	Wheat Trait Ontology
elig	Eligibility Feature Hierarchy
hom	Homology and Related Concepts in Biology
icd11	Body System Terms from ICD11
ofsmr	Open Food Safety Model Repository
ogr	Ontology of Geographical Region
pe	Pulmonary Embolism Ontology
taxrank	Taxonomic Rank Vocabulary
xeo	XEML Environment Ontology

---

Table 2: BioPortal ontologies considered for experimental validation

## Evaluation

The question of which one of two consistent repairs  $O_1$  and  $O_2$  of a given inconsistent ontology  $O$  is preferable is not, in general, well-defined. In this work, we compare two such repairs by taking into account the corresponding *inferred class hierarchies*. To this end, we define:

$$\text{Inf}(O_i) = \{A \sqsubseteq B : A, B \in N_C, O_i \models A \sqsubseteq B\} .$$

The intuition behind the choice of measure is that if  $\text{card}(\text{Inf}(O_1) \setminus \text{Inf}(O_2)) > \text{card}(\text{Inf}(O_2) \setminus \text{Inf}(O_1))$  (that is, if there exist more subsumptions between classes which can be inferred in  $O_1$  but not in  $O_2$  than vice versa) then  $O_1$  is to be preferred to  $O_2$ . Furthermore, class subsumptions, which can be inferred from both  $O_1$  or  $O_2$ , should be of no consequence to determine which repaired ontology is preferable. That is, whenever  $\text{Inf}(O_1) \subseteq \text{Inf}(O'_1)$ ,  $\text{Inf}(O_2) \subseteq \text{Inf}(O'_2)$  and  $\text{Inf}(O'_1) \setminus \text{Inf}(O_1) = \text{Inf}(O'_2) \setminus \text{Inf}(O_2)$  it should hold that the quality of  $O_1$  with respect to  $O_2$  is the same as the quality of  $O'_1$  with respect to  $O'_2$ . Thus, we define the following measure to compare the inferable information content of two ontologies.

**Definition 21.** Let  $O_1$  and  $O_2$  be two consistent ontologies. If  $\text{Inf}(O_1) \neq \text{Inf}(O_2)$ , we define the *inferable information content*  $\text{IIC}(O_1, O_2)$  of  $O_1$  w.r.t.  $O_2$  as  $\text{IIC}(O_1, O_2) =$

$$\frac{\text{card}(\text{Inf}(O_1) \setminus \text{Inf}(O_2))}{\text{card}(\text{Inf}(O_1) \setminus \text{Inf}(O_2)) + \text{card}(\text{Inf}(O_2) \setminus \text{Inf}(O_1))} ;$$

if instead  $\text{Inf}(O_1) = \text{Inf}(O_2)$ , we set  $\text{IIC}(O_1, O_2) = 0.5$ .

It is readily seen that this definition satisfies the two conditions mentioned above. Furthermore, the following properties hold:

1.  $\text{IIC}(O_1, O_2) \in [0, 1]$ ;
2.  $\text{IIC}(O_1, O_2) = 1 - \text{IIC}(O_2, O_1)$ ;




Figure 1: Comparing weakening-based ontology repair with removal-based ontology repair. Mean IIC of weakening-based against removal-based repair for each ontology, when choosing axioms at random (left) or by sampling minimally inconsistent sets (right).

3.  $IIC(O_1, O_2) = 0.5$  if and only if  $\text{card}(\text{Inf}(O_1)) = \text{card}(\text{Inf}(O_2))$ ;
4.  $IIC(O_1, O_2) = 1$  if and only if  $\text{Inf}(O_2) \subset \text{Inf}(O_1)$ ;
5.  $IIC(O_1, O_2) > 0.5$  if and only if  $\text{card}(\text{Inf}(O_1) \setminus \text{Inf}(O_2)) > \text{card}(\text{Inf}(O_2) \setminus \text{Inf}(O_1))$ .

Although this is by no means the only possible measure for comparing two ontologies (Tartir et al. 2005; Alani, Brewster, and Shadbolt 2006; Vrandečić and Sure 2007; Vrandečić 2009), these properties suggest that our definition captures a notion of “quality” that is meaningful for our intended application: in particular, if for two proposed repairs  $O_1, O_2$  of an inconsistent ontology  $O$  we have  $IIC(O_1, O_2) > 0.5$ , then there are more class subsumptions which can be inferred in  $O_1$  but not in  $O_2$  than vice versa, and hence—all other things being equal— $O_1$  is a better repair of  $O$  than  $O_2$ .

One possible criticism of our definition of  $IIC(O_1, O_2)$  is that its value depends only on  $\text{Inf}(O_1)$  and  $\text{Inf}(O_2)$ : if  $O_1$  and  $O_2$  differ only w.r.t. subsumptions between complex concepts, then  $IIC(O_1, O_2) = 0.5$  (even though the implications of  $O_1$  might still be considerably richer than those of  $O_2$ ). On the other hand, focusing on atomic subsumptions makes also conceptual sense, as these are the ones that our inconsistent ontology—as well as the proposed repairs—discuss about. It is, in any case, certainly true that our measure is fairly *coarse*: if  $IIC(O_1, O_2)$  is significantly greater than 0.5 there are good grounds to claim that  $O_1$  is a better repair of  $O$  than  $O_2$  is, but it may easily be that repair candidates between which our measure cannot discriminate are nonetheless of different quality.

To empirically test whether weakening axioms is a better approach to ontology repair than removing them, we tested our approach on ten ontologies from BioPortal (Matentzoglou and Parsia 2017), expressed in *ALC* (see Table 2). On aver-

	Random	MIS
bctt	0.55 (0.35)	<b>0.72 (0.36)</b>
co-wheat	<b>0.69 (0.29)</b>	<b>0.76 (0.31)</b>
elig	<b>0.61 (0.30)</b>	<b>0.72 (0.27)</b>
hom	<b>0.68 (0.26)</b>	<b>0.71 (0.31)</b>
icd11	0.60 (0.30)	<b>0.71 (0.40)</b>
ofsmr	<b>0.65 (0.31)</b>	<b>0.76 (0.29)</b>
ogr	0.56 (0.32)	<b>0.70 (0.35)</b>
pe	0.56 (0.33)	<b>0.67 (0.41)</b>
taxrank	0.56 (0.31)	<b>0.82 (0.36)</b>
xeo	<b>0.67 (0.29)</b>	<b>0.67 (0.34)</b>

Table 3: Mean and standard deviation (in parentheses) of IIC between RepairOntologyWeaken and RepairOntologyRemove, both when choosing axioms at random (left column) and by sampling minimally inconsistent sets (right). Bolded values are significant ( $p < 0.05$ ) with respect to both Wilcoxon and T-test with Holm-Bonferroni correction; non-bolded values were not significant for either.

age the ontologies have 105 logical axioms and 90 classes. We compared the performance of RepairOntologyWeaken (Algorithm 1) with the one of the non weakening-based RepairOntologyRemove (Algorithm 2) by first making the ontologies inconsistent through the addition of random axioms, then attempting to repair them through the two algorithms (using the original ontology as the reference), and then computing IIC. This procedure has the following rationale: one may think that the axioms added constitute some new claims made concerning the relationships between the classes of the ontology, which however unfortunately made it inconsistent. It is thus desirable to fix this inconsistency while preserving as much as possible of the informational content of these axioms and of the other axioms in the ontology.

The procedure was repeated one hundred times per ontology, selecting the axioms to weaken or remove by sampling minimally inconsistent sets, and one further hundred times selecting the axioms to remove or weaken completely randomly. We tested the significance of our results through both Wilcoxon signed-rank tests and T-tests, applying the Holm-Bonferroni correction for multiple comparison, with a p-value threshold of 0.05.

Figure 1 and Table 3 summarise the results of our experiments. When choosing the axioms to weaken or remove through sampling minimally inconsistent sets, the means (in the case of the T-test) and medians (in the case of the Wilcoxon test) of the IIC for RepairOntologyWeaken against RepairOntologyRemove were all significantly greater than 0.5 for all ontologies. This confirms that our repair-by-weakening technique is able to preserve more of the informational content of axioms than repair-by-removal techniques. When selecting the axioms to repair randomly, on the other hand, this was not always the case, as shown in Table 3. This illustrates how our repair-by-weakening approach constitutes a genuine improvement over removal-based ontology repair only when problematic axioms can be

reliably pinpointed.

Figure 1 highlights the effect of choosing the axioms to repair or remove randomly rather than through sampling inconsistent sets. While the difference is not statistically significant for all ontologies,<sup>5</sup> we observe that the quality of the repair compared to the corresponding removal is always improved by choosing the axioms to repair via sampling. The natural next step in this line of investigation would consist in evaluating the effect of varying the number of minimally inconsistent sets sampled by FindBadAxiom, which for these experiments was set to one tenth of the ontology size.

To summarise, the main conclusion of our experiments is that, when problematic axioms can be reliably identified, our approach is better able to preserve the informational content of inconsistent ontologies than the corresponding repair-by-removal method.

## Related Work

Refinements operators were also discussed in the context of inductive logic programming (van der Laag and Nienhuys-Cheng 1998), and formalised in description logics for concept learning (Lehmann and Hitzler 2010). The refinement operators used by our weakening approach were introduced in (Confalonieri et al. 2016), and further analysed in (Confalonieri et al. 2017) in relation to incoherence detection. They were not previously applied to ontology repair.

The problem of identifying and repairing inconsistencies in ontologies has received much attention in recent years. Our approach differs from many other works in the area, see for instance (Schlobach and Cornet 2003; Kalyanpur et al. 2005; 2006; Baader, Peñaloza, and Suntisrivaraporn 2007; Haase and Qi 2007), in that—rather than removing the problematic axioms altogether—we attempt to repair the ontology by replacing them with weakened rewritings. On the one hand, our method requires the choice of a (consistent) *reference ontology* with respect to which one can compute the weakenings; on the other hand, it allows us to perform a softer, more fine-grained form of ontology repair.

A different approach for repairing ontologies through weakening was discussed in (Lam et al. 2008). Our own approach is, however, quite different from it: while the repair algorithm of (Lam et al. 2008) operates by pinpointing (and subsequently removing) the subcomponents of the axioms responsible for the contradiction, ours is based on a refinement operator, which combines both semantic (via the cover operators) and syntactic (via the compositional definitions of generalisations and specialisations of complex formulas) information in order to identify candidates for the replacement of the offending axiom(s). In particular, this implies—using the terminology of (Ji et al. 2014)—that our repair algorithm, in contrast to (Lam et al. 2008), is ‘black box’ in that it treats the reasoner as an oracle, and can thus be more easily combined with different choices of reasoner (or, with slightly more effort, applied to different logics).

<sup>5</sup>For instance, w.r.t. the Wilcoxon test it is statistically significant for *bctt*, *elig*, *ogr*, *pe* and *taxrank*, but not for the other five ontologies.

Another influential approach to ontology repair is discussed in (Qi, Liu, and Bell 2006a) and in (Qi, Liu, and Bell 2006b). That approach, like ours, attempts to weaken problematic axioms; but it does so by adding *exceptions* to value restrictions  $\forall R.C(a)$ ,<sup>6</sup> rather than by means of a more general-purpose transformation.

We leave to future work the evaluation of our approach in comparison to other state-of-the-art ontology repair frameworks. As already stated, this is not an entirely well-posed problem; but if, as in this work, we accept that a suggested repair  $O_1$  is preferable to another suggested repair  $O_2$  whenever  $\text{card}(\text{Inf}(O_1) \setminus \text{Inf}(O_2)) > \text{card}(\text{Inf}(O_2) \setminus \text{Inf}(O_1))$  then the question becomes amenable to analysis. Possibly, complementary metrics for further evaluations can be chosen from (Alani, Brewster, and Shadbolt 2006). Experiments involving user evaluation could be also considered in this context.

## Conclusions

We have proposed a new strategy for repairing ontologies based on the idea of weakening terminological and assertional axioms. Axiom weakening is a way to improve the balance between regaining consistency and keeping as much information from the original ontology as possible.

We have investigated the theoretical properties of the refinement operators that are required in the definition of axiom weakening and analysed the computational complexity of employing them. Furthermore, the empirical evaluation shows that our weakening-based approach to repairing ontologies performs significantly better, in terms of preservation of information, than the removal-based approach.

Future work will concentrate on a few directions. Our experiments show that weakening-based repairs allow one to preserve more information than removal-based repairs. In practice, it will be important to involve humans in the loop to help make decisions, that are not only logically sensible, but that also make sense from a domain perspective. As initiated in (Porello et al. 2017), we plan to integrate the repair procedure into social choice mechanisms. In such mechanisms, experts can express their opinions as votes and preferences over the axioms of an ontology, that can then be used to steer the repair. This will lead us to investigate other evaluation measures that reflect the experts’ opinions. Other measures of the quality of repairs could reflect the preservation of entailments of competency questions, or the enabling of particular tasks (McNeill and Bundy 2007). Finally, we plan to extend the presented approach to axiom weakening to more expressive DL languages, including *SROIQ* underlying OWL 2 DL (Horrocks, Kutz, and Sattler 2006) and to full first-order logic, for which debugging is a particularly challenging problem (Kutz and Mossakowski 2011). We expect that, for more complex languages, the weakening-based strategy will likewise significantly improve on the removal-based strategy, and indeed be even more appropriate by exploiting the higher syntactic complexity.

<sup>6</sup>Another difference is that we are also interested in repairing TBoxes, whereas the approach of (Qi, Liu, and Bell 2006b) operates only over ABoxes.

## References

- Alani, H.; Brewster, C.; and Shadbolt, N. 2006. Ranking ontologies with AKTiveRank. In *Proc. of ISWC'06*, volume 4273 of *LNCS*, 1–15. Springer-Verlag.
- Auer, S.; Bizer, C.; Kobilarov, G.; Lehmann, J.; Cyganiak, R.; and Ives, Z. 2007. DBpedia: A Nucleus for a Web of Open Data. In *Proc. of ISWC'07/ASWC'07*, volume 4825 of *LNCS*, 722–735. Springer-Verlag.
- Baader, F.; Calvanese, D.; McGuinness, D. L.; Nardi, D.; and Patel-Schneider, P. F., eds. 2003. *The Description Logic Handbook: Theory, Implementation, and Applications*. New York, NY, USA: Cambridge University Press.
- Baader, F.; Brandt, S.; and Lutz, C. 2005. Pushing the  $\mathcal{EL}$  envelope. In *Proc. of IJCAI-05*, 364–369.
- Baader, F.; Peñaloza, R.; and Suntisrivaraporn, B. 2007. Pinpointing in the description logic  $\mathcal{EL}^+$ . In *Proc. of KI 2007*, volume 4667 of *LNCS*, 52–67. Springer.
- Bateman, J.; Hois, J.; Ross, R.; and Tenbrink, T. 2010. A Linguistic Ontology of Space for Natural Language Processing. *Artificial Intelligence* 174(14):1027–1071.
- Confalonieri, R.; Eppe, M.; Schorlemmer, M.; Kutz, O.; Peñaloza, R.; and Plaza, E. 2016. Upward refinement operators for conceptual blending in the description logic  $\mathcal{EL}^{++}$ . *Annals of Mathematics and Artificial Intelligence*.
- Confalonieri, R.; Kutz, O.; Galliani, P.; Peñaloza, R.; Porello, D.; Schorlemmer, M.; and Troquard, N. 2017. Coherence, Similarity, and Concept Generalisation. In *Proc. of DL 2017*, volume 1879 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Grüninger, M., and Fox, M. S. 1995. The role of competency questions in enterprise engineering. In *Benchmarking—Theory and Practice*. Springer. 22–31.
- Haase, P., and Qi, G. 2007. An analysis of approaches to resolving inconsistencies in DL-based ontologies. In *Proc. of IWOD-07*, 97–109.
- Horrocks, I.; Kutz, O.; and Sattler, U. 2006. The Even More Irresistible *SRCTQ*. In *Proc. of KR'06*, 57–67. AAAI Press.
- Ji, Q.; Gao, Z.; Huang, Z.; and Zhu, M. 2014. Measuring effectiveness of ontology debugging systems. *Knowledge-Based Systems* 71:169–186.
- Kalyanpur, A.; Parsia, B.; Sirin, E.; and Hendler, J. 2005. Debugging unsatisfiable classes in OWL ontologies. *Web Semantics: Science, Services and Agents on the World Wide Web* 3(4):268–293.
- Kalyanpur, A.; Parsia, B.; Sirin, E.; and Grau, B. C. 2006. Repairing unsatisfiable concepts in OWL ontologies. In *ESWC*, volume 6, 170–184. Springer.
- Kutz, O., and Mossakowski, T. 2011. A Modular Consistency Proof for Dolce. In *Proc. of AAAI-11*. AAAI Press.
- Lam, J. S. C.; Sleeman, D.; Pan, J. Z.; and Vasconcelos, W. 2008. A fine-grained approach to resolving unsatisfiable ontologies. In *Journal on Data Semantics X*. Springer. 62–95.
- Lehmann, J., and Hitzler, P. 2010. Concept learning in description logics using refinement operators. *Machine Learning* 78(1-2):203–250.
- Lembo, D.; Lenzerini, M.; Rosati, R.; Ruzzi, M.; and Savo, D. F. 2010. Inconsistency-tolerant semantics for description logics. In *Proc. of RR 2010*, volume 6333 of *LNCS*, 103–117. Springer.
- Ludwig, M., and Peñaloza, R. 2014. Error-tolerant reasoning in the description logic  $\mathcal{el}$ . In *Proc. of JELIA 2014*, volume 8761 of *LNCS*, 107–121. Springer.
- Matentzoglou, N., and Parsia, B. 2017. BioPortal snapshot 30.03.2017. last accessed, 2017/08/04.
- McNeill, F., and Bundy, A. 2007. Dynamic, Automatic, First-Order Ontology repair by Diagnosis of Failed Plan Execution. *International Journal on Semantic Web and Information Systems* 3(3):1–35.
- Neuhaus, F.; Vizedom, A.; Baclawski, K.; Bennett, M.; Dean, M.; Denny, M.; Grüninger, M.; Hashemi, A.; Longstreth, T.; Obrst, L.; et al. 2013. Towards ontology evaluation across the life cycle – the ontology summit 2013. *Applied Ontology* 8(3):179–194.
- Porello, D.; Troquard, N.; Confalonieri, R.; Galliani, P.; Kutz, O.; and Peñaloza, R. 2017. Repairing socially aggregated ontologies using axiom weakening. In *Proc. of PRIMA 2017*, volume 10621 of *LNCS*, 441–449. Springer.
- Prestes, E.; Carbonera, J. L.; Fiorini, S. R.; Jorge, V. A. M.; Abel, M.; Madhavan, R.; Locoro, A.; Goncalves, P.; Barreto, M. E.; Habib, M.; Chibani, A.; Gérard, S.; Amirat, Y.; and Schlenoff, C. 2013. Towards a core ontology for robotics and automation. *Robotics and Autonomous Systems* 61(11):1193 – 1204. Ubiquitous Robotics.
- Qi, G.; Liu, W.; and Bell, D. 2006a. A revision-based approach to handling inconsistency in description logics. *Artificial Intelligence Review* 26(1):115–128.
- Qi, G.; Liu, W.; and Bell, D. A. 2006b. Knowledge base revision in description logics. In *European Workshop on Logics in Artificial Intelligence*, 386–398. Springer.
- Ren, Y.; Parvizi, A.; Mellish, C.; Pan, J. Z.; van Deemter, K.; and Stevens, R. 2014. Towards competency question-driven ontology authoring. In *The Semantic Web: Trends and Challenges*. Springer. 752–767.
- Resnik, P. 1999. Semantic similarity in a taxonomy: An information-based measure and its application to problems of ambiguity in natural language. *Journal of Artificial Intelligence Research* 11:95–130.
- Sazonau, V.; Sattler, U.; and Brown, G. 2015. General terminology induction in OWL. In *Proc. ISWC 2015, Part I*, volume 9366 of *LNCS*, 533–550. Springer.
- Schlobach, S., and Cornet, R. 2003. Non-standard reasoning services for the debugging of description logic terminologies. In *Proc. of IJCAI-03*, 355–362. Morgan Kaufmann.
- Stuckenschmidt, H.; Parent, C.; and Spaccapietra, S., eds. 2009. *Modular Ontologies—Concepts, Theories and Techniques for Knowledge Modularization*, volume 5445 of *LNCS*. Springer.
- Tartir, S.; Arpinar, I. B.; Moore, M.; Sheth, A. P.; and Aleman-Meza, B. 2005. OntoQA: Metric-based ontology quality analysis. In *Proc. of 2005 IEEE ICDM Workshop on KADASH*, 45–53.
- van der Laag, P. R., and Nienhuys-Cheng, S.-H. 1998. Completeness and properness of refinement operators in inductive logic programming. *The Journal of Logic Programming* 34(3):201 – 225.
- Vrandečić, D. 2009. Ontology evaluation. In *Handbook on Ontologies*. Springer. 293–313.
- Vrandečić, D., and Sure, Y. 2007. How to design better ontology metrics. In *Proc. of ESWC '07*, volume 4519 of *LNCS*, 311–325. Springer-Verlag.