

Ontology Contraction: Beyond the Propositional Paradise

Bernardo Cuenca Grau¹, Evgeny Kharlamov², and Dmitriy Zheleznyakov²

¹ University of Oxford, bernardo.cuenca.grau@cs.ox.ac.uk

² Free University of Bozen-Bolzano, {kharlamov,zheleznyakov}@inf.unibz.it

Abstract. The dynamic nature of ontology development has motivated the formal study of ontology evolution problem. This paper addresses contraction —the problem of retracting information that should no longer hold in an ontology. We survey existing model and formula based semantics to contraction and investigate their properties for the description logics *DL-Lite* and \mathcal{EL} , which underpin the QL and EL profiles of OWL 2. Our results suggest that these contraction semantics, which are well-understood and well-behaved for propositional logics, are intrinsically problematical in the context of ontology languages. We believe that a starting point for addressing these problems might be the recent semantics proposed in [1].

1 Introduction

Ontologies written in the Web Ontology Language (OWL) [2] and its revision OWL 2 [3] are widely used in applications. The formal underpinning of OWL is based on Description Logics (DLs) – knowledge representation formalisms with well-understood computational properties [4]. A DL ontology \mathcal{K} consists of a TBox \mathcal{T} , describing schema-level domain knowledge, and an ABox \mathcal{A} , providing data about specific individuals.

Ontologies are not static entities, but rather they are frequently modified when new information needs to be incorporated, or existing information is no longer considered valid. The impact of such changes on the semantics of the ontology, however, is difficult to predict and understand. This dynamic nature of ontologies motivates the study of *ontology evolution problems* from both foundational and practical perspectives [5–12, 1].

In this paper, we focus on a particular aspect of evolution, namely *contraction* – the process of “retracting” information that should no longer hold [13, 14]. From a logic-based perspective, the desirable properties of contraction are dictated by the *principle of minimal change* [13], according to which the semantics of the ontology should change “as little as possible”, thus ensuring that the contraction has the least possible impact.

Logic-based semantics derived from the principle of minimal change have been studied in the more general context of ontology evolution and update. These semantics are either *model-based* (MBS) or *formula-based* (FBS). Under both types of semantics, evolution of ontology \mathcal{K} results in another ontology \mathcal{K}_{op} in which the required information is incorporated, retracted, or updated; the difference is in the way \mathcal{K}_{op} is obtained.

Under MBS the models \mathcal{M} of \mathcal{K} (i.e., the set of all first order interpretations satisfying \mathcal{K}) evolves into a set \mathcal{M}' of interpretations that are “as close as possible” to those in \mathcal{M} (w.r.t. some notion of distance between interpretations); then, \mathcal{K}_{op} is the ontology that axiomatises \mathcal{M}' [7–9, 15, 16]. Under FBS, \mathcal{K}_{op} is a finite subset of the *deductive closure*

of \mathcal{K} satisfying the evolution requirements, with FBS differing in their subset selection mechanism. FBS have been less studied in the context of ontologies [8, 17].

Approaches to ontology contraction typically adopted in practice are, however, *syntactic* [6, 10, 18, 19]. For example, to retract an axiom α entailed by \mathcal{K} , it suffices to compute a maximal subset of \mathcal{K} that does not entail α . This solution complies with a syntactical notion of minimal change: retracting α results in the deletion of a minimal set of axioms, and the structure of \mathcal{K} is maximally preserved. By removing axioms from \mathcal{K} , however, we may also be retracting consequences of \mathcal{K} other than α , which are intended. Identifying and recovering such intended consequences is an important issue.

In this paper we compare different syntactic, model-based, and formula-based semantics for ontology contraction, and study their basic properties. We consider two scenarios, which are important for many ontology design and management tasks:

- *TBox contraction*, where the axiom α to be retracted is a TBox axiom; and
- *ABox contraction*, where α is an ABox assertion, and the TBox of the original ontology cannot be modified as a result of the contraction.

OWL TBoxes are extensively used in the clinical sciences, and ontologies such as SNOMED are subject to frequent modifications that involve retracting TBox consequences [20]. ABox contraction is important for applications relying on widely-used *reference TBoxes*. For example, experimental results on gene extraction can be described using an ABox according to standard gene TBoxes. New experiments may imply that facts about specific genes no longer hold, which should be reflected in the ABox; at the same time, TBoxes should clearly not be affected by these manipulations of the data.

In our formal study of ontology contraction problems we will put special emphasis on *expressibility*: given a \mathcal{DL} -ontology \mathcal{K} , an axiom α to be retracted, and a “protected” part \mathcal{P} of \mathcal{K} , we will study whether a \mathcal{DL} -ontology \mathcal{K}_{op} entailed by \mathcal{K} exists – an *optimal contraction* – such that (i) $\mathcal{K}_{op} \not\models \alpha$, (ii) $\mathcal{K}_{op} \models \mathcal{P}$, and (iii) \mathcal{K}_{op} is “as similar as possible” to \mathcal{K} according to the notion of minimal change in the semantics under consideration.

Our framework is general, and applies to arbitrary first-order languages. When studying expressibility problems, however, we provide results for two prominent (families of) DLs: *DL-Lite* [21] and \mathcal{EL} [22], which underpin OWL 2 QL and OWL 2 EL, respectively. We show that existing inexpressibility results obtained for the problem of *update* and *revision* in *DL-Lite* under MBS [8, 15] also hold for contraction with minor modifications; furthermore, we conjecture that these results might hold even if the optimal contraction of a *DL-Lite* ontology is allowed to be expressed in full first-order logic. Concerning FBS, we focus on two semantics: the so-called *bold semantics* [8] and *WIDTIO semantics* [8, 23]; in both cases, we show inexpressibility results that apply to both TBox and ABox contraction in \mathcal{EL} .

Furthermore, we show that existing MBS, FBS, and syntactic approaches to ontology contraction are rather incompatible. Although one would expect FBS approaches to behave better in terms of expressibility, this is not always the case; in particular, we identify simple cases for which inexpressibility can be shown in the FBS approach, but which can be easily captured using a model-based semantics.

Our results suggest that classical approaches to ontology contraction, which are well-understood and well-behaved for propositional logics, are intrinsically problematical in

\mathcal{DL}	concepts, roles	axioms	
\mathcal{EL}	$C := \perp \mid \top \mid A \mid \exists R.C \mid C_1 \sqcap C_2$	$C_1 \sqsubseteq C_2$	
$DL\text{-Lite}$	$C := \perp \mid \top \mid A \mid \exists R.\top$ $R := P \mid P^-$	$C_1 \sqsubseteq C_2, C_1 \sqsubseteq \neg C_2, C \sqsubseteq C_1 \sqcap C_2,$ $(\text{funct } R), R_1 \sqsubseteq R_2$	
concepts, roles	semantics	axiom	semantics
\perp, \top	$\emptyset, \Delta^{\mathcal{I}}, \text{ resp.}$	$C_1 \sqsubseteq C_2$	$C_1^{\mathcal{I}} \subseteq C_2^{\mathcal{I}}$
$C_1 \sqcap C_2$	$C_1^{\mathcal{I}} \cap C_2^{\mathcal{I}}$	$C_1 \sqsubseteq \neg C_2$	$C_1^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \setminus C_2^{\mathcal{I}}$
$\exists R.C$	$\{a \mid \exists b \text{ s.t. } (a, b) \in R^{\mathcal{I}}, b \in C^{\mathcal{I}}\}$	$(\text{funct } R)$	$R^{\mathcal{I}} \text{ is functional}$
R^-	$\{(a, b) \mid (b, a) \in R^{\mathcal{I}}\}$	$R_1 \sqsubseteq R_2$	$R_1^{\mathcal{I}} \subseteq R_2^{\mathcal{I}}$

Fig. 1. Syntax and semantics of \mathcal{EL} and $DL\text{-Lite}$ concepts, roles and axioms

the context of ontology languages. A starting point for addressing these problems might be the semantics in [1], which unifies and extends FBS and syntactic approaches: the challenge remains to extend this semantics to encompass also MBS contraction.

2 Preliminaries

We discuss contraction in the context of first-order logic (FOL). Our work, however, is motivated by description logic ontologies, so we will use DL terminology throughout the paper. We assume standard definitions of (function-free) FOL signature, predicates, formulae, sentences, interpretations and models, satisfiability, and entailment.

An ontology $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ consists of finite sets of first-order sentences \mathcal{T} (the TBox) and ground atoms \mathcal{A} (the ABox). In the most general setting, a description logic can be defined as a recursive set of ontologies closed under renaming of constants and the subset relation. Predicates in DL signatures are typically restricted to be unary (*atomic concepts*) or binary (*atomic roles*). DLs use a specialised syntax, where variables are omitted, and which provides operators for constructing complex concepts and roles from simpler ones, as well as a set of axioms.

We consider two (families of) DLs: $DL\text{-Lite}$ [21] and \mathcal{EL} [22]. The syntax of (the variants of) $DL\text{-Lite}$ and \mathcal{EL} considered here are given in Figure 1, where A is an atomic concept, C , C_1 , and C_2 are arbitrary concepts, and \top and \perp are special concepts which are mapped by every interpretation to the domain and the empty set, respectively. TBoxes in $DL\text{-Lite}$ are restricted: first, \top cannot occur on the left-hand side of axioms; second, for each R_2 s.t. $R_1 \sqsubseteq R_2 \in \mathcal{T}$, it holds that $\mathcal{T} \not\models (\text{funct } R_2)$ and $\mathcal{T} \not\models (\text{funct } R_2^-)$.

Let \mathcal{DL} and \mathcal{DL}' be DLs with $\mathcal{DL}' \subseteq \mathcal{DL}$. The *closure* $\text{Cl}_{\mathcal{DL}'}(\mathcal{K})$ of $\mathcal{K} \in \mathcal{DL}$ w.r.t. \mathcal{DL}' is the set of all \mathcal{DL}' -axioms α entailed by \mathcal{K} .

The evaluation of $DL\text{-Lite}$ and \mathcal{EL} concepts and axioms under $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ is also given in Figure 1. We adopt the *standard name assumption* and assume that $c^{\mathcal{I}} = c$ for each constant c ; that is, we do not distinguish between constants and domain elements.

3 The Ontology Contraction Problem

As already mentioned, contraction can be seen at a high level as the process of retracting an axiom α that holds in an ontology \mathcal{K} while preserving a *protected part* \mathcal{P} of \mathcal{K} .

The following notion of a *contraction setting* formalises the basic requirements that \mathcal{K} , α and \mathcal{P} must satisfy for the contraction process to make sense.

Definition 1. Let \mathcal{DL} be a DL. A \mathcal{DL} -contraction setting is a triple $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$, with $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ a \mathcal{DL} -ontology, α a \mathcal{DL} -axiom s.t. $\mathcal{K} \models \alpha$, and \mathcal{P} a \mathcal{DL} -ontology s.t. $\mathcal{K} \models \mathcal{P}$ and $\mathcal{P} \not\models \alpha$. We say that \mathcal{C} is a TBox-contraction setting if α is a TBox axiom in \mathcal{DL} ; \mathcal{C} is an ABox contraction setting if α is an ABox assertion in \mathcal{DL} and $\mathcal{P} \models \mathcal{T}$. \square

A contraction for a \mathcal{DL} -setting $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ can now be defined as a \mathcal{DL} -ontology \mathcal{K}_{op} that preserves \mathcal{P} and in which α no longer holds. Furthermore, since \mathcal{K}_{op} should not add new information to \mathcal{K} , we also require \mathcal{K}_{op} to be entailed by \mathcal{K} .

Definition 2. Let $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ be a \mathcal{DL} -contraction setting. We say that $\mathcal{K}_{op} \in \mathcal{DL}$ is a contraction for \mathcal{C} if (i) $\mathcal{K} \models \mathcal{K}_{op}$; (ii) $\mathcal{K}_{op} \models \mathcal{P}$; and (iii) $\mathcal{K}_{op} \not\models \alpha$. \square

The properties that an “optimal” contraction needs to satisfy are dictated by the *principle of minimal change* according to which the semantics of the ontology should be changed “as little as possible”, thus ensuring that modifications have the least possible impact. Hence, the *contraction problem* can be understood at a high level as follows:

[CONTRACT]: Is a given \mathcal{DL} -ontology \mathcal{K}_{op} an *optimal contraction* for a given \mathcal{DL} -contraction setting $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$?; in other words, is \mathcal{K}_{op} a contraction for \mathcal{C} such that no other contraction \mathcal{K}'_{op} for \mathcal{C} is “more similar” to \mathcal{K} than \mathcal{K}_{op} ?

Thus, optimal contractions are those that are “as similar to \mathcal{K} as possible”. The notion of optimal contraction immediately suggests the following *expressibility problem*.

[EXPRESS]: Does an optimal contraction exist for a given \mathcal{DL} -contraction setting?

Contraction semantics essentially differ in their formalisation of *optimality*. These can roughly be divided into three groups, which we shall discuss next: *model-based semantics* (MBS), *formula-based semantics* (FBS), and *syntactic approaches*.

4 Syntactic Contraction

Approaches to ontology contraction typically adopted in practice are essentially *syntactic* [6, 10, 18]. In particular, to retract an axiom α entailed by \mathcal{K} , it suffices to compute a maximal subset \mathcal{K}_{op} of \mathcal{K} that does not entail α . Thus, retracting α results in the deletion of a minimal set of axioms and hence the structure of \mathcal{K} is maximally preserved. In this setting, optimal contractions can be defined as follows.

Definition 3. Let $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ be a \mathcal{DL} -contraction setting. A contraction \mathcal{K}_{op} for \mathcal{C} is syntactically optimal if $\mathcal{K}_{op} \subseteq \mathcal{K}$ and no contraction \mathcal{K}' for \mathcal{C} exists s.t. $\mathcal{K}_{op} \subset \mathcal{K}' \subseteq \mathcal{K}$. \square

In this case, **[CONTRACT]** and **[EXPRESS]** essentially amount to entailment checking. Furthermore, practical algorithms for computing such optimal \mathcal{K}_{op} have been implemented in ontology development platforms. By adopting this approach to contraction, however, we may inadvertently retract consequences of \mathcal{K} that are “intended”. Identifying and recovering such intended consequences is an important issue.

Example 4. Let $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$, where $\mathcal{K} = \{A \sqsubseteq B \sqcap C, A \sqsubseteq \exists R.A \sqcap C\}$, $\mathcal{P} = \emptyset$, and $\alpha = A \sqsubseteq C$. Clearly, $\mathcal{K}_{op} = \emptyset$ is syntactically optimal. By computing \mathcal{K}_{op} , however, we have also retracted consequences of \mathcal{K} unrelated to α , i.e., $A \sqsubseteq B$ and $A \sqsubseteq \exists R.A$. \square




Fig. 2. Model-based contraction semantics: example and notation.

5 Model-Based Contraction

Intuitively, under an MBS a contraction \mathcal{K}_{op} for a setting $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ is optimal if the set of models $Mod(\mathcal{K}_{op})$ of \mathcal{K}_{op} is precisely the union of the models $Mod(\mathcal{K})$ of \mathcal{K} and the set of interpretations \mathcal{I} such that (i) $\mathcal{I} \models \mathcal{P}$, (ii) $\mathcal{I} \not\models \alpha$; and (iii) \mathcal{I} is “minimally distant” from the models of \mathcal{K} [13, 24, 25].

Calvanese et al. [8] considered two notions of “minimal distance”, which they called *local* and *global*. We next define these semantics in the context of our framework.

Definition 5. Let $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ be a \mathcal{DL} -contraction setting, and let $dist(\cdot, \cdot)$ be a distance function between interpretations. For \mathcal{I} an interpretation, let $loc_min(\mathcal{I}, \mathcal{P}, \alpha)$ be the set of interpretations \mathcal{J} s.t., $\mathcal{J} \models \mathcal{P}$, $\mathcal{J} \not\models \alpha$, and on \mathcal{J} the value of $dist(\mathcal{I}, \mathcal{J})$ is minimal for the given \mathcal{I} . Then, $\mathcal{K}_{op} \in \mathcal{DL}$ is **L**-optimal for \mathcal{C} if the following holds:

$$Mod(\mathcal{K}_{op}) = Mod(\mathcal{K}) \cup \bigcup_{\mathcal{I} \models \mathcal{K}} loc_min(\mathcal{I}, \mathcal{P}, \alpha). \quad \square$$

Thus, under local semantics, the models of \mathcal{I} of \mathcal{K} are considered one by one and $Mod(\mathcal{K})$ is extended with those models \mathcal{J} of \mathcal{P} such that $\mathcal{J} \not\models \alpha$ and \mathcal{J} is minimally distant from \mathcal{I} . The specific distance under consideration can vary, but it typically maps each pair of interpretations to either a number or a subset of some fixed set.

To get a better intuition of local semantics, consider the left hand side of Figure 2, which depicts two models \mathcal{I}_0 and \mathcal{I}_1 of \mathcal{K} and four interpretations $\mathcal{J}_0, \dots, \mathcal{J}_3$ which satisfy \mathcal{P} but not α . The distance between \mathcal{I}_i and \mathcal{J}_j is represented by the length of the line connecting them; solid lines correspond to minimal distances, and dashed lines to distances that are not minimal. In this case, \mathcal{J}_0 must be included in $Mod(\mathcal{K}_{op})$ because of \mathcal{I}_0 , and $\mathcal{J}_2, \mathcal{J}_3$ must be included in $Mod(\mathcal{K}_{op})$ because of \mathcal{I}_1 .

Definition 6. Let $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ be a \mathcal{DL} -contraction setting, and $dist(\cdot, \cdot)$ a distance function between interpretations. For an interpretation \mathcal{J} , let $dist(Mod(\mathcal{K}), \mathcal{J}) = \min_{\mathcal{I} \in Mod(\mathcal{K})} dist(\mathcal{I}, \mathcal{J})$. Furthermore, let $glob_min(\mathcal{K}, \mathcal{P}, \alpha)$ be the set of interpretations \mathcal{J} s.t. $\mathcal{J} \models \mathcal{P}$, $\mathcal{J} \not\models \alpha$, and for each interpretation \mathcal{J}' such that $\mathcal{J}' \models \mathcal{P}$ and $\mathcal{J}' \not\models \alpha$ it does not hold that $dist(Mod(\mathcal{K}), \mathcal{J}') < dist(Mod(\mathcal{K}), \mathcal{J})$.

Then, $\mathcal{K}_{op} \in \mathcal{DL}$ is **G**-optimal for \mathcal{C} if the following condition holds:

$$Mod(\mathcal{K}_{op}) = Mod(\mathcal{K}) \cup glob_min(\mathcal{K}, \mathcal{P}, \alpha). \quad \square$$

If we consider again Figure 2 and assume that the distance between \mathcal{I}_0 and \mathcal{J}_0 is the global minimum, then \mathcal{J}_2 and \mathcal{J}_3 are not included in $Mod(\mathcal{K}_{op})$.

Finally, note that if \mathcal{K}_{op} is optimal (either locally or globally), then \mathcal{K}_{op} is a contraction for \mathcal{C} , as in Definition 2. Furthermore, **L**-optimal and **G**-optimal contractions are also clearly unique modulo logical equivalence.

5.1 Measuring Distance Between Interpretations.

Classical MBS semantics were originally developed for propositional theories [23]. In this setting, interpretations can be seen as finite sets of propositional symbols, and the symmetric difference “ \ominus ” between such sets can be used to define specific distance functions. More precisely, we can define $\mathcal{I} \ominus \mathcal{J}$ as the set $(\mathcal{I} \setminus \mathcal{J}) \cup (\mathcal{J} \setminus \mathcal{I})$ and introduce the following two distance functions, where $|\mathcal{I} \ominus \mathcal{J}|$ is the cardinality of the set $\mathcal{I} \ominus \mathcal{J}$:

$$\text{dist}_{\subseteq}(\mathcal{I}, \mathcal{J}) = \mathcal{I} \ominus \mathcal{J} \quad \text{and} \quad \text{dist}_{\#}(\mathcal{I}, \mathcal{J}) = |\mathcal{I} \ominus \mathcal{J}|. \quad (1)$$

Distances under dist_{\subseteq} are compared using set inclusion: $\text{dist}_{\subseteq}(\mathcal{I}_1, \mathcal{J}_1) \leq \text{dist}_{\subseteq}(\mathcal{I}_2, \mathcal{J}_2)$ iff $\text{dist}_{\subseteq}(\mathcal{I}_1, \mathcal{J}_1) \subseteq \text{dist}_{\subseteq}(\mathcal{I}_2, \mathcal{J}_2)$. Finite distances under $\text{dist}_{\#}$ are natural numbers and are compared in the standard way.

These distances can be extended to DL interpretations in two ways. First, one can consider interpretations \mathcal{I} and \mathcal{J} as sets of atoms, in which case the symmetric difference $\mathcal{I} \ominus \mathcal{J}$ and the corresponding distances are defined as in the propositional case. In contrast to the propositional case, however, $\mathcal{I} \ominus \mathcal{J}$ (and hence also distances) can be infinite. We denote the distances in Equation (1) as $\text{dist}_{\subseteq}^a(\mathcal{I}, \mathcal{J})$ and $\text{dist}_{\#}^a(\mathcal{I}, \mathcal{J})$, respectively.

Finally, one can also define distances at the level of the concept and role *symbols* in the signature Σ underlying the interpretations:

$$\text{dist}_{\subseteq}^s(\mathcal{I}, \mathcal{J}) = \{S \in \Sigma \mid S^{\mathcal{I}} \neq S^{\mathcal{J}}\}, \quad \text{and} \quad \text{dist}_{\#}^s(\mathcal{I}, \mathcal{J}) = |\{S \in \Sigma \mid S^{\mathcal{I}} \neq S^{\mathcal{J}}\}|.$$

To sum up, we can classify MBS semantics along three dimensions, which in turn lead to eight different MBS semantics for contraction: (1) *local* vs. *global* approach, (2) *atom*-based vs. *symbol*-based distances; and (3) *set inclusion* vs. *cardinality* to compare symmetric differences.

These three dimensions are depicted on the right-hand side of Figure 2. We denote each of the resulting eight semantics by using a combination of three symbols, indicating the choice in each dimension, e.g., $\mathbf{L}_{\#}^a$ denotes the local semantics where the distances are expressed in terms of cardinality of sets of atoms. We can then define \mathbf{L}_x^y -optimality (respectively, \mathbf{G}_x^y -optimality) as in Definition 5 (respectively, as in Definition 6) by using the specific distances determined by the values of x and y .

Since in the propositional case there is no distinction between atom and symbol-based semantics, we use our notation without superscripts for propositional MBS. The classical local MBS by Winslett [26] and Forbus [27] correspond to \mathbf{L}_{\subseteq} , and $\mathbf{L}_{\#}$, respectively. Borgida’s semantics [28] is a variant of \mathbf{L}_{\subseteq} . The classical global MBS proposed by Satoh [29] and Dalal [30] correspond to \mathbf{G}_{\subseteq} , and $\mathbf{G}_{\#}$, respectively.

Definition 7. Let \mathcal{DL} be a DL and let \mathbf{M}_x^y with $\mathbf{M} \in \{\mathbf{L}, \mathbf{G}\}$, $x \in \{\subseteq, \#\}$, and $y \in \{s, a\}$ be an MBS contraction semantics. We say \mathcal{DL} is closed under \mathbf{M}_x^y -contraction (or \mathbf{M}_x^y -contraction is expressible in \mathcal{DL}) if for each \mathcal{DL} -contraction setting \mathcal{C} , an ontology $\mathcal{K}_{op} \in \mathcal{DL}$ exists such that \mathcal{K}_{op} is \mathbf{M}_x^y -optimal.

Analogously, \mathcal{DL} is closed under TBox (resp. ABox) \mathbf{M}_x^y -contraction if for each TBox (resp. ABox) \mathcal{DL} -contraction setting \mathcal{C} , an \mathbf{M}_x^y -optimal $\mathcal{K}_{op} \in \mathcal{DL}$ exists. \square

5.2 Challenges in Capturing MBS Contractions.

MBS contraction is not always expressible in lightweight DLs, such as *DL-Lite* and \mathcal{EL} . Inexpressibility problems mainly originate from the inability of these logics to express disjunction, as shown in the next proposition [8].

Proposition 8. *Let \mathcal{K} be either a *DL-Lite* or an \mathcal{EL} ontology. If $\mathcal{K} \models A(a) \vee B(b)$, with A, B atomic concepts, then either $\mathcal{K} \models A(a)$ or $\mathcal{K} \models B(b)$.*

We next illustrate inexpressibility issues for *DL-Lite* and \mathcal{EL} under each of the eight MBS contraction semantics introduced in this paper. We start with TBox contraction.

Example 9. Let $\mathcal{DL} \in \{\text{DL-Lite}, \mathcal{EL}\}$ and let $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$, where $\mathcal{K} = \{A \sqsubseteq B \sqcap C, A(a)\}$, $\mathcal{P} = \{A(a)\}$, and $\alpha = A \sqsubseteq B \sqcap C$. Pick an MBS semantics \mathbf{M}_x^y and assume $\mathcal{K}_{op} \in \mathcal{DL}$ exists s.t. \mathcal{K}_{op} is \mathbf{M}_x^y -optimal for \mathcal{C} .

First, observe that the following conditions hold for each $\mathcal{I} \in \text{Mod}(\mathcal{K})$ and each $\mathcal{J} \models \mathcal{P}$ s.t. $\mathcal{J} \not\models \alpha$ (both seen as sets of atoms): (i) $B(a), C(a) \in \mathcal{I}$; (ii) either $\{B(a)\} \subseteq \text{dist}_{\subseteq}(\mathcal{I}, \mathcal{J})$ or $\{C(a)\} \subseteq \text{dist}_{\subseteq}(\mathcal{I}, \mathcal{J})$ for $\mathcal{J} \in \text{Mod}(\mathcal{K}_{op})$; (iii) $1 \leq \text{dist}_{\#}(\mathcal{I}, \mathcal{J})$ for $\mathcal{J} \in \text{Mod}(\mathcal{K}_{op})$; (iv) if \mathcal{J} is such that $B(a) \notin \mathcal{J}$ and $C(a) \notin \mathcal{J}$, then $\mathcal{J} \notin \text{Mod}(\mathcal{K}_{op})$.

Now, consider the following models: $\mathcal{J}_1 = \{A(a), B(a)\}$, and $\mathcal{J}_2 = \{A(a), C(a)\}$. From Items (ii) and (iii) we conclude that $\mathcal{J}_i \in \text{loc_min}(\{A(a), B(a), C(a)\}, \mathcal{P}, \alpha)$. Then, we can use items (i) and (iv) to conclude that $\mathcal{K}_{op} \models B(a) \vee C(a)$. By Proposition 8, we must have either $\mathcal{K}_{op} \models B(a)$ or $\mathcal{K}_{op} \models C(a)$. But then, since items (ii) and (iii) ensure that $\mathcal{K}_{op} \not\models B(a)$ and $\mathcal{K}_{op} \not\models C(a)$, we obtain a contradiction. \square

Observe next that similar issues arise for ABox contraction under local MBS.

Example 10. Let $\mathcal{DL} = \mathcal{EL}$ and let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be the following \mathcal{DL} -ontology:

$$\mathcal{T} = \{A \sqsubseteq \exists R.B, B \sqcap C \sqsubseteq \perp\}; \quad \mathcal{A} = \{A(a), R(a, b), B(b), C(c), C(d)\}.$$

Finally, let \mathcal{C} be the \mathcal{DL} -contraction setting defined by \mathcal{K} , $\mathcal{P} = \mathcal{T} \cup \{A(a)\}$, and $\alpha = R(a, b)$. Pick a local MBS semantics \mathbf{L}_x^y and assume $\mathcal{K}_{op} \in \mathcal{DL}$ exists s.t. \mathcal{K}_{op} is \mathbf{L}_x^y -optimal for \mathcal{C} . Consider the following interpretations:

$$\begin{aligned} \mathcal{I}_0 : \quad & A^{\mathcal{I}} = \{a\}, \quad B^{\mathcal{I}} = \{b\}, \quad C^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus \{b\}, \quad R^{\mathcal{I}} = \{(a, b)\}; \\ \mathcal{J}_1 : \quad & A^{\mathcal{J}} = \{a\}, \quad B^{\mathcal{J}} = \{b, c\}, \quad C^{\mathcal{J}} = \Delta^{\mathcal{J}} \setminus \{b, c\}, \quad R^{\mathcal{J}} = \{(a, c)\}; \\ \mathcal{J}_2 : \quad & A^{\mathcal{J}} = \{a\}, \quad B^{\mathcal{J}} = \{b, d\}, \quad C^{\mathcal{J}} = \Delta^{\mathcal{J}} \setminus \{b, d\}, \quad R^{\mathcal{J}} = \{(a, d)\}. \end{aligned}$$

Clearly, $\mathcal{I}_0 \models \mathcal{K}$ and one can check that $\mathcal{J}_1, \mathcal{J}_2 \in \text{loc_min}(\mathcal{I}_0, \mathcal{P}, \alpha)$. Moreover, since $\mathcal{J}_1 \not\models C(c)$ and $\mathcal{J}_2 \not\models C(d)$, we have $\mathcal{K}_{op} \not\models C(c)$ and $\mathcal{K}_{op} \not\models C(d)$. At the same time, one can show that $\mathcal{K}_{op} \models C(c) \vee C(d)$. Proposition 8 then leads to a contradiction.

Inexpressibility of *DL-Lite* ABox contraction can be shown analogously by taking \mathcal{A}, \mathcal{P} , and α as before and using a TBox $\mathcal{T} = \{A \sqsubseteq \exists R.\top, \exists R^-. \top \sqsubseteq B, B \sqsubseteq \neg C\}$, which ‘‘mimics’’ the behaviour of the \mathcal{EL} -TBox considered before. \square

Theorem 11. *Let $\mathcal{DL} \in \{\text{DL-Lite}, \mathcal{EL}\}$ and let $\mathbf{M} \in \{\mathbf{L}, \mathbf{G}\}$, $x \in \{\subseteq, \#\}$ and $y \in \{s, a\}$. Then, \mathcal{DL} is not closed under TBox \mathbf{M}_x^y -contraction. Furthermore, \mathcal{DL} is not closed under ABox \mathbf{L}_x^y -contraction.*

These inexpressibility results can be overcome by allowing fewer models in optimal contractions. To this effect, one can define a partial order \preceq on models where $\mathcal{I}_1 \preceq \mathcal{I}_2$ if \mathcal{I}_1 changes certain aspects of $\text{Mod}(\mathcal{K})$ less than \mathcal{I}_2 ; then, \mathcal{K}_{op} can only have models that are \preceq -minimal. For example, in [7] changes in concepts rather than roles were preferred. Formally, $\mathcal{I}_1 \preceq_{\mathcal{K}}^R \mathcal{I}_2$ if there is $\mathcal{I} \in \text{Mod}(\mathcal{K})$ such that $\text{dist}_{\subseteq}^s(\mathcal{I}, \mathcal{I}_2)$ contains a role and $\text{dist}_{\subseteq}^s(\mathcal{I}, \mathcal{I}_1)$ contains only concepts. Another example of \preceq is to order the signature of \mathcal{K} , e.g., $A \preceq B$ if B is a “more important concept” than A ; that is, if $\text{dist}_{\subseteq}^a(\mathcal{I}, \mathcal{I}_2) = \{A(a)\}$ and $\text{dist}_{\subseteq}^a(\mathcal{I}, \mathcal{I}_1) = \{B(a), B(b)\}$, then $\mathcal{I}_1 \preceq \mathcal{I}_2$ since \mathcal{I}_2 changes a more important concept A .

Definition 12. *Let \preceq be a partial order on the class of all first-order interpretations and min_{\preceq} the class of interpretations minimal w.r.t. \preceq . Let $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ be a \mathcal{DL} -contraction setting. We say that $\mathcal{K}_{op} \in \mathcal{DL}$ is $[\preceq, \mathbf{L}]$ -optimal for \mathcal{C} if it holds that:*

$$\text{Mod}(\mathcal{K}_{op}) = \text{Mod}(\mathcal{K}) \cup \bigcup_{\mathcal{I} \models \mathcal{K}} (\text{loc_min}(\mathcal{I}, \mathcal{P}, \alpha) \cap \text{min}_{\preceq}).$$

Finally, $\mathcal{K}_{op} \in \mathcal{DL}$ is $[\preceq, \mathbf{G}]$ -optimal for \mathcal{C} if the following condition holds:

$$\text{Mod}(\mathcal{K}_{op}) = \text{Mod}(\mathcal{K}) \cup (\text{glob_min}(\mathcal{K}, \mathcal{P}, \alpha) \cap \text{min}_{\preceq}). \quad \square$$

We next show that this solution can lead to even more severe inexpressibility problems; in particular, optimal contractions might not even be expressible with FOL ontologies.

Example 13. Let $\mathcal{DL} = \mathcal{ELFI}$, which extends \mathcal{EL} with functionality and inverses. Let \mathcal{C} be defined by $\alpha = A(a_1)$, $\mathcal{P} = \mathcal{T}$, and the following \mathcal{DL} ontology $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$:

$$\mathcal{T} = \{A \sqsubseteq \exists R.B, B \sqsubseteq \exists R.A, A \sqcap B \sqsubseteq \perp, (\text{funct } R), (\text{funct } R^{-}), \mathcal{A} = \{A(a_1)\}.$$

Each model \mathcal{I} of \mathcal{K} is of one of the following two types (see Figure 3): (i) \mathcal{I} contains a cycle of the form $R(a_1, a_2), \dots, R(a_{2n}, a_1)$ for some $n \in \mathbb{N}$ such that for each $1 \leq k \leq 2n$ we have $A(a_k) \in \mathcal{I}$ if k is odd and $B(a_k) \in \mathcal{I}$ if k is even; (ii) \mathcal{I} contains an infinite R -chain $R(a_1, a_2), \dots, R(a_k, a_{k+1}), \dots$ such that for each $k \geq 1$ we have $A(a_k) \in \mathcal{I}$ if k is odd and $B(a_k) \in \mathcal{I}$ if k is even.

Assume there exists a FOL ontology \mathcal{K}_{op} that is $[\preceq_{\mathcal{K}}^R, \mathbf{G}_{\subseteq}^s]$ -optimal for \mathcal{C} . By Definition 12, $\text{Mod}(\mathcal{K}_{op})$ consists of following two disjoint sets of models: $\text{Mod}(\mathcal{K})$ and $\mathcal{S} = \text{glob_min}(\mathcal{K}, \mathcal{P}, \alpha) \cap \text{min}_{\preceq}$. Observe that \mathcal{S} consists of the models obtained from $\text{Mod}(\mathcal{K})$ as described next (see r.h.s. of Figure 3). If \mathcal{I} is of Type (i), then one has to drop all atoms of the form $A(a_k)$ and $B(a_k)$ for each a_k involved in the cycle, which gives us again a cyclic model; let \mathcal{S}_1 be the set of all such models. If \mathcal{I} is of Type (ii), then one must drop only $A(a_1)$, which yields a model with an infinite chain; let \mathcal{S}_2 be the set of all such models. Clearly, \mathcal{S}_1 , \mathcal{S}_2 , and $\text{Mod}(\mathcal{K})$ are mutually disjoint. The following ontology has exactly \mathcal{S}_2 as models: $\mathcal{K}_2 = \mathcal{T} \cup \{\exists R.B(a_1), \neg A(a_1)\}$. We next use the locality property of FOL to show that no FOL ontology \mathcal{K}_1 exists s.t. $\text{Mod}(\mathcal{K}_1) = \mathcal{S}_1$. Assume such \mathcal{K}_1 exists and consider \mathcal{I}_1 containing an R -cycle of even length, and \mathcal{I}_2 containing an R -cycle of odd length. If the cycle’s length is big enough, \mathcal{K}_1 cannot distinguish between \mathcal{I}_1 and \mathcal{I}_2 ; thus, $\mathcal{I}_2 \in \text{Mod}(\mathcal{K}_1)$, a contradiction. Finally, $\text{Mod}(\mathcal{K}_{op} \wedge \neg \mathcal{K} \wedge \neg \mathcal{K}_2) = \mathcal{S}_1$, which contradicts the fact that \mathcal{S}_1 cannot be captured by a FOL ontology. Hence, no optimal \mathcal{K}_{op} exists.³ \square

³ We denote with $\neg \mathcal{K}_{(i)}$ the formula obtained by negating the conjunction of all formulas in $\mathcal{K}_{(i)}$.




Fig. 3. Models illustrating issues with FOL expressibility of contraction

6 Formula-Based Contraction

Given a contraction setting \mathcal{C} , the *bold semantics* [8] selects a maximal subset of the closure of the corresponding ontology that is a contraction for \mathcal{C} .

Definition 14. Let $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ be a \mathcal{DL} -contraction setting and $\mathcal{M}(\mathcal{C})$ the class of maximal subsets \mathcal{S} of $\text{Cl}_{\mathcal{DL}}(\mathcal{K})$ s.t. $\mathcal{S} \models \mathcal{P}$ and $\mathcal{S} \not\models \alpha$. Then, $\mathcal{K}_{op} \in \mathcal{DL}$ is **BS**-optimal for \mathcal{C} if there exists $\mathcal{S} \in \mathcal{M}(\mathcal{C})$ such that $\mathcal{K}_{op} \equiv \mathcal{S}$. \square

Note that under Bold semantics \mathcal{K}_{op} is not unique in general (even modulo equivalence). There have been several proposals for combining all elements of $\mathcal{M}(\mathcal{C})$ into a single set of formulas [8, 23, 26]. Under *Cross-Product* (**CP**) semantics, an optimal evolution is equivalent to the “disjunction” of all relevant maximal subsets of the closure, whereas under *When In Doubt Throw It Out* (**WIDTIO**) semantics, one takes the “intersection”.

Definition 15. Let $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ be a \mathcal{DL} -contraction setting. Then, $\mathcal{K}_{op} \in \mathcal{DL}$ is **CP**-optimal for \mathcal{C} if $\mathcal{K}_{op} \equiv \{\bigvee_{\mathcal{S} \in \mathcal{M}(\mathcal{C})} (\bigwedge_{\beta \in \mathcal{S}} \beta)\}$, **WIDTIO**-optimal if $\mathcal{K}_{op} \equiv \bigcap_{\mathcal{S} \in \mathcal{M}(\mathcal{C})} \mathcal{S}$. \square

Expressibility of optimal contractions can now be defined in the obvious way.

Definition 16. Let \mathcal{DL} be a DL and let $\mathbf{S} \in \{\mathbf{BS}, \mathbf{CP}, \mathbf{WIDTIO}\}$ be an FBS contraction semantics. We say \mathcal{DL} is closed under \mathbf{S} -contraction (or \mathbf{S} -contraction is expressible in \mathcal{DL}) if for each \mathcal{DL} -contraction setting \mathcal{C} , an \mathbf{S} -optimal $\mathcal{K}_{op} \in \mathcal{DL}$ exists. \square

Intuitively, **CP** has the advantage of not “losing information”; however, **CP**-optimal contractions can be exponentially larger than the original ontology, even if its closure is a finite set. In addition, even if we consider a *DL-Lite* contraction setting \mathcal{C} , the corresponding **CP**-optimal contraction for \mathcal{C} may not be expressible in *DL-Lite*, since a language with disjunction may be required. In contrast, **WIDTIO**-optimal contractions for *DL-Lite* are always expressible, but important information may be lost.

Thus, both **CP** and **WIDTIO** semantics are somewhat problematic, even for languages such as *DL-Lite*, where the deductive closure is always a finite set. In contrast, **BS** semantics is well-behaved for *DL-Lite*: optimal contractions always exist; furthermore, in the case of ABox contraction they are also unique and computable in polynomial time [8]. In general, however, **BS**-optimal contractions are non-unique for TBox contractions.

6.1 Expressibility issues for \mathcal{EL} .

Unfortunately, as soon as we consider logics such as \mathcal{EL} , for which the closure of an ontology can be infinite, **WIDTIO** and **BS** semantics lead to inexpressibility problems. This is illustrated by the following example (adapted from Lemma 19 in [1]).

Example 17. Let $\mathcal{C} = (\mathcal{T}, \mathcal{P}, \alpha)$ be the \mathcal{EL} contraction setting defined as follows:

$$\mathcal{T} = \{Z \sqsubseteq \exists R.A, A \sqsubseteq \exists R.A, \exists R.B \sqsubseteq B, A \sqsubseteq B\}; \quad \alpha = A \sqsubseteq B; \quad \mathcal{P} = \mathcal{T} \setminus \{\alpha\}$$

Furthermore, for each $k \in \mathbb{N}$, let

$$\gamma_k = Z \sqsubseteq \exists R^k.(A \sqcap B), \quad \beta_k = Z \sqsubseteq \exists R^k.B, \quad \Lambda_k = \{\gamma_i \mid 1 \leq i \leq k\}.$$

For each $k \in \mathbb{N}$, observe that $\gamma_k \models \beta_k$, $\Lambda_k \not\models \beta_{k+1}$, and $\gamma_k \in \text{Cl}_{\mathcal{EL}}(\mathcal{T})$; also, $\mathcal{P} \cup \Lambda_k \not\models \alpha$. Let \mathcal{T}_{op} be **BS**-optimal for \mathcal{C} ; we then have $\mathcal{T}_{op} \models \beta_k$ for each $k \in \mathbb{N}$. Furthermore, for each finite $\mathcal{S} \subseteq \text{Cl}_{\mathcal{EL}}(\mathcal{T})$ there is n such that $\mathcal{P} \cup \Lambda_n \models \mathcal{S}$. Thus, such $n = n_0$ exists for \mathcal{T}_{op} and hence $\Lambda_{n_0} \not\models \beta_{n_0+1}$ and by monotonicity of FOL, $\mathcal{T}_{op} \not\models \beta_{n_0+1}$ and we obtain a contradiction to the maximality of \mathcal{T}_{op} ; hence, \mathcal{T}_{op} cannot be **BS**-optimal. \square

We formalize the intuition from the example in the next theorem (which is an adaptation of Theorem 20 from [1]).

Theorem 18. *\mathcal{EL} is not closed under TBox **BS**-contraction.*

A similar inexpressibility result can be obtained for ABox contraction.

Example 19. Let $\mathcal{C} = (\mathcal{A}, \emptyset, \alpha)$ be the \mathcal{EL} -ABox contraction setting defined by $\mathcal{A} = \{A(a), R(a, a)\}$, and $\alpha = R(a, a)$. Then, for each $k \in \mathbb{N}$, we have $\alpha_k = \exists R^k.A(a)$ is in $\text{Cl}_{\mathcal{EL}}(\mathcal{A})$. Clearly, $\bigcup_k \{\alpha_k\} \not\models \alpha$. Thus, if an \mathcal{EL} **BS**-contraction \mathcal{A}_{op} for \mathcal{C} exists, then $\mathcal{A}_{op} \models \alpha_k$ for each k . One can show that no such finite \mathcal{A}_{op} exists. \square

Theorem 20. *\mathcal{EL} is not closed under ABox **BS**-contraction.*

Examples 17 and 19 can also be used to illustrate inexpressibility of **WIDTIO**-optimal contractions. Indeed, observe that in Example 19 every maximal subset \mathcal{A}_{op} of $\text{Cl}_{\mathcal{EL}}(\mathcal{A})$ contains the infinite set $\bigcup_k \{\alpha_k\}$.

Theorem 21. *\mathcal{EL} is not closed under ABox and TBox **WIDTIO**-contraction.*

6.2 FBS vs MBS

We next show that some contraction settings for which no MBS-optimal contraction exists can be easily captured using FBS (and vice versa). Thus, there seem to be certain key points of ‘‘incompatibility’’ between model-based and formula-based contraction semantics, which require further investigation.

Example 22. Consider the setting \mathcal{C} in Example 9. Let $\mathcal{A}' = \{A(a), B(a), C(a)\}$. Under **BS** semantics we have two optimal contractions, namely $\mathcal{K}_{op}^1 = \mathcal{A}' \cup \{A \sqsubseteq B\}$ and $\mathcal{K}_{op}^2 = \mathcal{A}' \cup \{A \sqsubseteq C\}$, which are both in *DL-Lite* and \mathcal{EL} .

Next, recall the \mathcal{EL} -setting \mathcal{C} from Example 10. There is a unique (modulo equivalence) **BS**-optimal contraction for \mathcal{C} , namely $\mathcal{K}_{op} = \mathcal{K} \setminus \{\alpha\}$.

Finally, recall Example 19. Under all model-based contraction semantics considered in this paper, the optimal contraction for \mathcal{C} can be shown to be $\mathcal{A}_{op} = \{A(a)\}$. \square

6.3 Extending the Formula Based Approach

As already discussed, formula-based semantics are well-behaved for DLs such as *DL-Lite*, where the closure of an ontology is always finite [8, 15]. FBS are, however, problematic for DLs like \mathcal{EL} , which do not provide such guarantee. Inexpressibility issues for **BS**-semantics originate from the requirement that optimal contractions for $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ must be equivalent to a maximal subset \mathcal{S} of $\text{Cl}_{\mathcal{DL}}(\mathcal{K})$ satisfying \mathcal{P} but not α ; if $\text{Cl}_{\mathcal{DL}}(\mathcal{K})$ is infinite, it might be that no such \mathcal{S} is equivalent to a finite set of \mathcal{DL} formulas.

FBS are also problematic from a modeling point of view: in contrast to syntactic approaches, they do not distinguish between the axioms in the closure that are explicit in \mathcal{K} , and those that are implied. Ontologies, however, are the result of a time-consuming modeling process, and thus contractions should also aim at preserving the structure of \mathcal{K} .

In [1], these limitations of FBS approaches were addressed (at least partly). On the one hand, the semantics in [1] provides a “bridge” between syntactic and FBS approaches; on the other hand, this semantics provides a distinction between the languages \mathcal{DL} in which both \mathcal{K} and the resulting contraction are expressed, and the language \mathcal{LP} (the *preservation language*), which expresses the entailments of \mathcal{K} that must be maximally preserved. The principle of minimal change is reflected along two dimensions:

- (i) *structural*, where the explicit axioms in \mathcal{K} are maximally preserved;
- (ii) *deductive*, where the consequences in $\text{Cl}_{\mathcal{LP}}(\mathcal{K})$ are maximally preserved.

Definition 23. Let $\mathcal{C} = (\mathcal{K}, \mathcal{P}, \alpha)$ be a \mathcal{DL} -contraction setting and let $\mathcal{LP} \subseteq \mathcal{DL}$. A contraction \mathcal{K}_{op} for \mathcal{C} is **SD-optimal**⁴ for \mathcal{C} and \mathcal{LP} if (i) $\mathcal{K}_{op} \cup \{\beta\} \models \alpha$, for each $\beta \in \mathcal{K} \setminus \mathcal{K}_{op}$; and (ii) $\mathcal{K}_{op} \cup \{\gamma\} \models \alpha$, for each $\gamma \in \text{Cl}_{\mathcal{LP}}(\mathcal{K}) \setminus \text{Cl}_{\mathcal{LP}}(\mathcal{K}_{op})$. \square

The notion of expressibility can be formalised as in Definition 16 in the obvious way. Note that the preservation language \mathcal{LP} provides “control” over the consequences to be preserved. Furthermore, syntactic contractions can be easily captured by taking \mathcal{LP} to be the empty set. We next illustrate **SD**-contractions with an example.

Example 24. Let $\mathcal{DL} = \mathcal{EL}$ and let \mathcal{LP} be the DL consisting of all atomic subsumptions of the form $A \sqsubseteq B$. Let $\mathcal{K} = \{A \sqsubseteq B, B \sqsubseteq C\}$ and let $\alpha = A \sqsubseteq B$. Clearly, $\mathcal{K}_{op} = \{B \sqsubseteq C, A \sqsubseteq C\}$ is an **SD-optimal** contraction for \mathcal{LP} and $\mathcal{C} = (\mathcal{K}, \emptyset, \alpha)$. \square

As shown in [1], there exist practically relevant DLs \mathcal{DL} and \mathcal{LP} such that, on the one hand, \mathcal{DL} is closed under **SD**-contraction for \mathcal{LP} and, on the other hand, $\text{Cl}_{\mathcal{LP}}(\mathcal{K})$ is in general an infinite set for $\mathcal{K} \in \mathcal{DL}$. For example, one may consider \mathcal{DL} to be the set of *acyclic* \mathcal{EL} ontologies – roughly speaking, those \mathcal{EL} ontologies that do not entail cyclic axioms involving existential quantifiers on the right hand side of the axiom (e.g., axioms such as $A \sqsubseteq \exists R.A$). Many \mathcal{EL} ontologies such as SNOMED, satisfy such acyclicity condition. In this setting, restrictions to \mathcal{LP} also apply (see [1] for details).

7 Conclusion and Future Work

In propositional logic, contraction is always expressible under both MBS and FBS. In the case of propositional MBS, contraction results are finite sets of finite models, which

⁴ **SD** stands for *Structural-Deductive*

can always be axiomatised as a set of propositional formulas. The problem of interest in the propositional case is the complexity of axiomatisation [23]. For FBS the situation is similar: deductive closure of any propositional theory is a finite set (modulo equivalence); thus, **BS**, **CP**, **WIDTIO**, or **SD**-optimal contractions always exist. Again, the challenge here is the complexity of optimal contraction computation [23].

However, as soon as we move from propositional to first-order logic (and even to computationally well-behaved fragments such as Description Logics), we are also forced to leave propositional paradise. As we have shown, inexpressibility issues arise in rather simple cases for both FBS and MBS. These results suggest that classical approaches to contraction are intrinsically problematical in the context of ontologies. A starting point for addressing these problems might be the semantics in [1], which unifies and extends FBS and syntactic approaches; the challenge remains to extend this semantics to encompass also MBS contraction—an inspiring problem for future work.

Interesting further steps include: *(i)* understanding the impact of the standard-names assumption on expressibility and computation of contraction, *(ii)* better understanding FOL inexpressibility (e.g., see Example 13) *(iii)* isolating fragments of *DL-Lite* and \mathcal{EL} for which contraction is well-behaved (see preliminary results in [1, 15]).

Acknowledgements. We thank Balder Ten Cate for helpful discussions and anonymous reviewers for the inspiring feedback. B. Cuenca Grau is supported by a Royal Society Fellowship. E. Kharlamov is supported by the EPSRC EP/G004021/1, EP/H017690/1, and ERC FP7 grant Webdam (n. 226513). E. Kharlamov and D. Zheleznyakov are supported by the EU project ACSI (FP7-ICT-257593).

References

1. Cuenca Grau, B., Jimenez-Ruiz, E., Kharlamov, E., Zheleznyakov, D.: Ontology evolution under semantic constraints. In: KR, Rome, Italy (2012)
2. Horrocks, I., Patel-Schneider, P.F., van Harmelen, F.: From \mathcal{SHIQ} and RDF to OWL: the making of a web ontology language. *Journal of Web Semantics* **1**(1) (2003) 7–26
3. Cuenca Grau, B., Horrocks, I., Motik, B., Parsia, B., Patel-Schneider, P., Sattler, U.: OWL 2: The next step for OWL. *Journal of Web Semantics* **6**(4) (2008) 309–322
4. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F.: The description logic handbook: Theory, implementation, and applications. In: *Description Logic Handbook*, Cambridge Uni. Press (2003)
5. Flouris, G., Manakanatas, D., Kondylakis, H., Plexousakis, D., Antoniou, G.: Ontology change: classification and survey. *Knowledge Eng. Review* **23**(2) (2008) 117–152
6. Haase, P., Stojanovic, L.: Consistent evolution of OWL ontologies. In: *ESWC*. (2005)
7. Qi, G., Du, J.: Model-based revision operators for terminologies in Description Logics. In: *IJCAI*. (2009) 891–897
8. Calvanese, D., Kharlamov, E., Nutt, W., Zheleznyakov, D.: Evolution of DL-Lite Knowledge Bases. In: *ISWC*, Shanghai, China (2010)
9. Wang, Z., Wang, K., Topor, R.W.: Revising general knowledge bases in Description Logics. In: *KR*. (2010)
10. Jimenez-Ruiz, E., Cuenca Grau, B., Horrocks, I., Berlanga, R.: Supporting concurrent ontology development: Framework, algorithms and tool. *Data and Know. Eng.* **70**:1 (2011)
11. Konev, B., Walther, D., Wolter, F.: The logical difference problem for description logic terminologies. In: *IJCAR*. (2008)

12. Gonçalves, R.S., Parsia, B., Sattler, U.: Analysing the evolution of the NCI thesaurus. In: CBMS. (2011) 1–6
13. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.* **50**(2) (1985) 510–530
14. Peppas, P.: Belief revision. In: Handbook of Knowledge Representation. (2007)
15. Kharlamov, E., Zheleznyakov, D.: Capturing instance level ontology evolution for DL-Lite. In: ISWC. (2011)
16. Wang, Z., Wang, K., Topor, R.W.: A new approach to knowledge base revision in DL-Lite. In: AAAI. (2010)
17. Lenzerini, M., Savo, D.F.: On the evolution of the instance level of DL-Lite knowledge bases. In: DL. (2011)
18. Kalyanpur, A., Parsia, B., Sirin, E., Grau, B.C.: Repairing unsatisfiable concepts in OWL ontologies. In: ESWC. (2006) 170–184
19. Schlobach, S., Huang, Z., Cornet, R., van Harmelen, F.: Debugging incoherent terminologies. *J. Automated Reasoning* **39**(3) (2007) 317–349
20. Hartung, M., Kirsten, T., Rahm, E.: Analyzing the evolution of life science ontologies and mappings. In: DILS. (2008) 11–27
21. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Tractable reasoning and efficient query answering in Description Logics: The DL-Lite family. *Journal of Automated Reasoning* **39**(3) (2007) 385–429
22. Baader, F., Brandt, S., Lutz, C.: Pushing the EL envelope. In: IJCAI. (2005) 364–369
23. Eiter, T., Gottlob, G.: On the complexity of propositional knowledge base revision, updates, and counterfactuals. In: PODS. (1992) 261–273
24. Katsuno, H., Mendelzon, A.O.: On the difference between updating a knowledge base and revising it. In: KR. (1991)
25. Giacomo, G.D., Lenzerini, M., Poggi, A., Rosati, R.: On instance-level update and erasure in description logic ontologies. *Journal Logic and Computation* **19**(5) (2009) 745–770
26. Winslett, M.: Updating Logical Databases. Cambridge Tracts in Theor. Comp. Sci. (1990)
27. Forbus, K.D.: Introducing actions into qualitative simulation. In: IJCAI. (1989)
28. Borgida, A.: Language features for flexible handling of exceptions in information systems. *ACM Transactions on Database Systems* **10**(4) (1985) 565–603
29. Satoh, K.: Nonmonotonic reasoning by minimal belief revision. In: FGCS. (1988)
30. Dalal, M.: Investigations into a theory of knowledge base revision. In: AAAI. (1988) 475–479