Revising General Knowledge Bases in Description Logics

Zhe Wang and **Kewen Wang** and **Rodney Topor**
Griffith University, Australia

Abstract

Revising knowledge bases (KBs) in description logics (DLs) in a syntax-independent manner is an important and nontrivial problem for ontology management and DL communities. Several attempts have been made to adapt classical model-based belief revision/update techniques to DLs, but they are restricted in several ways. In particular, they rarely investigate how to deal with revisions of general DL KBs. An important reason for this is that, unlike propositional logic, a DL KB may have infinitely many models with complex (and possibly infinite) structures, making it difficult to define and compute revisions in terms of models. In this paper, we focus on a specific DL in the DL-Lite family, but aim to define and compute revision by general KBs. For this reason, we first develop an alternative semantic characterization for DL-Lite by introducing the concept of a feature (as an alternative to a model) and then define specific revision operators for DL-Lite KBs based on features. We also present an algorithm for computing best approximations for KB revisions in DL-Lite.

Introduction

Description logic (DL) has proved to be the most successful formalism for representing and reasoning about static knowledge in ontology applications. Such application require a knowledge base (KB) consisting of a TBox (of terminological axioms) and an ABox (of data membership assertions). For example, the W3C ontology language framework OWL designed for Semantic Web applications is based on a particular set of description logics.

However, ontologies in Semantic Web applications are not static, but evolve over time. An important and nontrivial problem for such applications is thus how to effectively and efficiently revise/update KBs in a natural way. A typical scenarios is the need for incremental ontology design to satisfy a changing environment.

This problem of revising/updating a knowledge base has been thoroughly studied for classical logic. Here, classical belief revision/update approaches can be classified as *formula-based* (*syntactic*) approaches or *model-based* (*semantic*) approaches, according to the relevance of syntax (Eiter et al. 1992). Syntactic approaches have the disadvantages of being too sensitive to the syntax of the KB and of being unable to modify formulas (they can only add and delete formulas). Semantic approaches hence appear to be preferable for Semantic Web applications.

Recently, there has also been significant interest in revising/updating knowledge bases in description logics. Initial approaches concentrated on adapting classical revision postulates, *i.e.*, AGM postulates, to DLs, but no specific revision operators were provided, *e.g.*, (Flouris et al. 2006). More recently, there have been several attempts to define specific knowledge change operators for DLs, *e.g.*, (Halaschek-Wiener et al. 2006; Ribeiro et al. 2007; Qi et al. 2008). However, these approaches are all syntactic relevant in nature. There have been relatively few attempts to adapt classical model-based revision/update approaches to DLs, *e.g.*, (Liu et al. 2006; Giacomo et al. 2007; Qi et al. 2009). However, these approaches are restricted in several ways. In particular, these revision/update operators for DLs can deal only with ABoxes (Liu et al. 2006; Giacomo et al. 2007), or only with TBoxes (Qi et al. 2009).

In contrast to previous approaches, we focus on a specific DL, but address the problem of defining and computing revisions for general KBs, consisting of TBoxes *and* ABoxes. The DL-Lite family(Calvanese et al. 2007; Artale et al. 2007), which forms the basis of OWL 2 QL (one of the three profiles of OWL 2), is a family of lightweight DLs with efficient KB reasoning and query answering algorithms. We choose DL-Lite$^{\mathcal{N}}_{bool}$ (Artale et al. 2007), one of the most expressive members of the DL-Lite family, and define revision operators for DL-Lite$^{\mathcal{N}}_{bool}$ KBs in a way analogous to the model-based approaches in propositional logic. Although our approach is based on DL-Lite$^{\mathcal{N}}_{bool}$, we note that the definitions and algorithms can easily be adapted to other DLs. We also note that update traditionally addresses changes of the actual state of the world (*e.g.*, that resulted from some action), whereas revision addresses the incorporation of new knowledge about the world. In this paper, we focus on the revision problem.

The key issues in adapting classical model-based approaches to DLs are how to define the distance between models and how to construct the resulting KB (directly or indirectly) from selected models. However, such adaption is difficult for the following reasons. (1) DL models have complex (possibly infinite) structures, which require a complex definition of the distance between two models. (2) Unlike a propositional theory, a DL KB may have infinitely many models, making it impossible to compute the result directly via models. (3) Given a collection $\mathbb{M}$ of models, there may not exist a single KB $\mathcal{K}$ such that $\mathbb{M}$ is exactly the set of models for $\mathcal{K}$. These are also the reasons for the restrictions in previous approaches to DL revision.

In this paper, we first define *features* for DL-Lite$^{\mathcal{N}}_{bool}$, which precisely capture the most important semantic properties of DL-Lite$^{\mathcal{N}}_{bool}$ KBs, and (unlike models) are always finite. We adapt the techniques of model-based revision in propositional logic to the revision of DL-Lite$^{\mathcal{N}}_{bool}$ KBs, and define two specific revision operators based on two definitions of distance between features. We show that both revision operators possess desirable logical properties, and one of them preserve more knowledge from the original KB and thus yields a better result. As a set of features may not correspond exactly to any DL-Lite$^{\mathcal{N}}_{bool}$ KB, we also present syntactic algorithms for approximating the result of revision as a single DL-Lite$^{\mathcal{N}}_{bool}$ KB.

Due to space limitation, detailed proofs are omitted in this paper but a longer version of this paper with complete proofs can be found at `http://hobbit.ict.griffith.edu.au/~kewen/revision_long.pdf`.

The DL-Lite Family

A *signature* is a finite set $\mathcal{S} = \mathcal{S}_C \cup \mathcal{S}_R \cup \mathcal{S}_I \cup \mathcal{S}_N$ where $\mathcal{S}_C$ is the set of atomic concepts, $\mathcal{S}_R$ is the set of atomic roles, $\mathcal{S}_I$ is the set of individual names and $\mathcal{S}_N$ is the set of natural numbers in $\mathcal{S}$. We assume 1 is always in $\mathcal{S}_N$. $\top$ and $\bot$ will not be considered as atomic concepts or atomic roles. Formally, given a signature $\mathcal{S}$, a DL-Lite$^{\mathcal{N}}_{bool}$ language has the following syntax:

$$R \leftarrow P \mid P^-, \qquad B \leftarrow \top \mid \bot \mid A \mid \geqslant n\, R,$$
$$C \leftarrow B \mid \neg C \mid C_1 \sqcap C_2,$$

where $n \in \mathcal{S}_N$, $A \in \mathcal{S}_C$ and $P \in \mathcal{S}_R$. B is called a *basic concept* and C is called a *general concept*. We write $\exists R$ as a shorthand for $\geq 1\ R$.

A *TBox* $\mathcal{T}$ is a finite set of concept *inclusions* of the form $C_1 \sqsubseteq C_2$, where C_1 and C_2 are general concepts. An *ABox* $\mathcal{A}$ is a finite set of membership *assertions* of the form $C(a), R(a,b)$, where a, b are individual names. We call $C(a)$ a concept assertion and $R(a,b)$ a role assertion. A knowledge base (KB) is a pair $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$.

The semantics of a DL-Lite KB is given by interpretations. An interpretation $\mathcal{I}$ is a pair $(\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set called the *domain* and $\cdot^{\mathcal{I}}$ is an interpretation function which associates each atomic concept A with a subset $A^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$, each atomic role P with a binary relation $P^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, and each individual name a with an element $a^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$.

The interpretation function $\cdot^{\mathcal{I}}$ can be extended to general concept descriptions:

$$(P^-)^{\mathcal{I}} = \{(a^{\mathcal{I}}, b^{\mathcal{I}}) \mid (b^{\mathcal{I}}, a^{\mathcal{I}}) \in P^{\mathcal{I}}\},$$
$$(\geqslant n\, R)^{\mathcal{I}} = \{a^{\mathcal{I}} \mid |\{b^{\mathcal{I}} \mid (a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}\}| \geq n\},$$
$$(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}, \quad (C_1 \sqcap C_2)^{\mathcal{I}} = C_1^{\mathcal{I}} \cap C_2^{\mathcal{I}}.$$

An interpretation $\mathcal{I}$ *satisfies* inclusion $C_1 \sqsubseteq C_2$ if $C_1^{\mathcal{I}} \subseteq C_2^{\mathcal{I}}$; $\mathcal{I}$ *satisfies* assertion $C(a)$ if $a^{\mathcal{I}} \in C^{\mathcal{I}}$; $\mathcal{I}$ *satisfies* assertion $R(a,b)$ if $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$. $\mathcal{I}$ *satisfies* TBox $\mathcal{T}$ (or ABox $\mathcal{A}$) if $\mathcal{I}$ satisfies each inclusion in $\mathcal{T}$ (resp., each assertion in $\mathcal{A}$). $\mathcal{I}$ is a *model* of a KB $\langle \mathcal{T}, \mathcal{A} \rangle$, if $\mathcal{I}$ satisfies both $\mathcal{T}$ and $\mathcal{A}$. We use $\mathsf{Mod}(\mathcal{K})$ to denote the set of models of KB $\mathcal{K}$. $\mathsf{Sig}(\mathcal{K})$ is the signature of $\mathcal{K}$.

A KB $\mathcal{K}$ is consistent if it has at least one model. Two KBs $\mathcal{K}_1, \mathcal{K}_2$ that have the same models are said to be equivalent, denoted $\mathcal{K}_1 \equiv \mathcal{K}_2$. A KB $\mathcal{K}$ entails an inclusion or assertion α, denoted $\mathcal{K} \models \alpha$, if all models of $\mathcal{K}$ satisfy α.

Given a set $\mathbb{M}$ of interpretations and a signature $\mathcal{S}$, in most cases there does not exist a KB $\mathcal{K}$ over $\mathcal{S}$ such that the models of $\mathcal{K}$ is exactly $\mathbb{M}$. To tackle this inexpressibility problem, a notion of best approximation is introduced in (Giacomo et al. 2007). A KB $\mathcal{K}$ is said to be a *maximal approximation* of $\mathbb{M}$ over $\mathcal{S}$ if (1) $\mathsf{Sig}(\mathcal{K}) \subseteq \mathcal{S}$, (2) $\mathbb{M} \subseteq \mathsf{Mod}(\mathcal{K})$, and (3) there exists no KB $\mathcal{K}'$ over $\mathcal{S}$ such that $\mathbb{M} \subseteq \mathsf{Mod}(\mathcal{K}') \subset \mathsf{Mod}(\mathcal{K})$. It is shown in (Giacomo et al. 2007) that maximal approximation may not exist for some DLs. However, we can show that maximal approximations always exist in DL-Lite$^{\mathcal{N}}_{bool}$. When it exists, maximal approximation of $\mathbb{M}$ is unique up to KB equivalence.

A *disjunctive knowledge base* (DKB) (Meyer et al. 2005) is a set $\mathbb{K}$ of KBs, defined in a way that $\mathsf{Mod}(\mathbb{K}) = \bigcup_{\mathcal{K} \in \mathbb{K}} \mathsf{Mod}(\mathcal{K})$.

Features in DL-Lite$^{\mathcal{N}}_{bool}$

In this section, we introduce the concept of features in DL-Lite$^{\mathcal{N}}_{bool}$, which provides an alternative semantic characterization for DL-Lite$^{\mathcal{N}}_{bool}$. An advantage of semantic features over models is that the number of all features for a DL-Lite$^{\mathcal{N}}_{bool}$ knowledge base is finite and each feature is finite. These finiteness properties make it possible to recast key approaches to revision for classical propositional logic into DL-Lite$^{\mathcal{N}}_{bool}$.

Features for DL-Lite$^{\mathcal{N}}_{bool}$ are based on the notion of *types* defined in (Kontchakov et al. 2008).

In the following sections, if not specified, we assume $\mathcal{S}$ is a (large enough) fixed signature, *i.e.*, $\mathsf{Sig}(E) \subseteq \mathcal{S}$ for any general concept, inclusion, assertion, or KB E taken into consideration.

An *$\mathcal{S}$-type* τ is a set of basic concepts over $\mathcal{S}$, s.t. $\top \in \tau$, $\bot \notin \tau$, and for any $m, n \in \mathcal{S}_N$ with $m < n$, $\geqslant n\ R \in \tau$ implies $\geqslant m\ R \in \tau$. When the signature $\mathcal{S}$ is clear from context, we will simply call an $\mathcal{S}$-type a *type*. As $\top \in \tau$ for any type τ, we omit it in examples for simplicity.

For example, let $\mathcal{S}_C = \{A, B\}$, $\mathcal{S}_R = \{P\}$, and $\mathcal{S}_N = \{1, 3\}$. Then $\tau = \{\, A, \exists P, \geq 3P, \exists P^- \,\}$ is a type.

Intuitively, if each concept C is viewed as a propositional atom, types correspond to propositional interpretations of C. But a type is different from a propositional interpretation in that an element in a type may be of complex form, *e.g.*, $\geqslant n\ R$ and $\exists P$. Define a type τ *satisfying* a concept in the following way: τ satisfies basic concept B if $B \in \tau$, τ satisfies $\neg C$ if τ does not satisfy C, and τ satisfies $C_1 \sqcap C_2$ if τ satisfies both C_1 and C_2.

We can also define a type τ *satisfies* concept inclusion $C_1 \sqsubseteq C_2$ if τ satisfies concept $\neg C_1 \sqcup C_2$. τ *satisfies* a TBox $\mathcal{T}$ if it satisfies every inclusion in $\mathcal{T}$.

Types are enough to capture the semantics of TBoxes, but insufficient for ABoxes. We need to extend the notion of types and thus define *Herbrand sets* for ABoxes.

Definition 1 *A $\mathcal{S}$-*Herbrand set *(or* Herbrand set *when $\mathcal{S}$ is clear from the context) $\mathcal{H}$ is a finite set of assertions of the form $B(a)$ or $P(a,b)$, where $a, b \in \mathcal{S}_I$, $P \in \mathcal{S}_R$ and B is a basic concept over $\mathcal{S}$, satisfying the following conditions*

1. *For each $a \in \mathcal{S}_I$, $\top(a) \in \mathcal{H}$, $\bot(a) \notin \mathcal{H}$, and $\geqslant n\, R(a) \in \mathcal{H}$ implies $\geqslant m\, R(a) \in \mathcal{H}$ for $m, n \in \mathcal{S}_N$ with $m < n$.*
2. *For each $P \in \mathcal{S}_R$, $P(a, b_i) \in \mathcal{H}$ $(i = 1, \ldots, n)$ implies $\geqslant m\, P(a) \in \mathcal{H}$ for any $m \in \mathcal{S}_N$ such that $m \le n$.*
3. *For each $P \in \mathcal{S}_R$, $P(b_i, a) \in \mathcal{H}$ $(i = 1, \ldots, n)$ implies $\geqslant m\, P^-(a) \in \mathcal{H}$ for any $m \in \mathcal{S}_N$ such that $m \le n$.*

By condition 1 in Definition 1, given all the concept assertions of a in $\mathcal{H}$, $B_1(a), \ldots, B_k(a)$ $(k \ge 1)$, then $\tau = \{B_1, \ldots, B_k\}$ is a type. We call τ the *type of* a in $\mathcal{H}$. Conditions 2 and 3 preserve the consistency of a Herbrand set. Since $\top(a)$ is always in $\mathcal{H}$ for any $\mathcal{H}$ and $a \in \mathcal{S}_I$, for simplicity, we will omit it in examples.

We define a Herbrand set $\mathcal{H}$ *satisfies* concept assertion $C(a)$ if the type of a in $\mathcal{H}$ satisfies concept C. $\mathcal{H}$ *satisfies* role assertion $P(a, b)$ or $P^-(b, a)$ if $P(a, b)$ is in $\mathcal{H}$. $\mathcal{H}$ *satisfies* an ABox $\mathcal{A}$ if $\mathcal{H}$ satisfies every assertion in $\mathcal{A}$.

To provide an alternative characterization for reasoning in a KB, we could use pairs $\langle \tau, \mathcal{H} \rangle$, where τ is a type and $\mathcal{H}$ is a Herbrand set, to replace standard interpretations, such that $\langle \tau, \mathcal{H} \rangle$ satisfies KB $\langle \mathcal{T}, \mathcal{A} \rangle$ if τ satisfies $\mathcal{T}$ and $\mathcal{H}$ satisfies $\mathcal{A}$. The resulting satisfaction relation should guarantee that $\mathcal{K}$ is consistent iff there exists some pair $\langle \tau, \mathcal{H} \rangle$ satisfying $\mathcal{K}$. However, the following example shows that it is not the case.

Example 1 *Let $\mathcal{K} = \langle \{ \exists P^- \sqsubseteq \bot \}, \{ \exists P(a) \} \rangle$. Obviously, $\mathcal{K}$ is inconsistent and thus has no model. However, let $\mathcal{S} = \{P, a, 1\}$ then $\langle \{\exists P\}, \{\exists\, P(a)\} \rangle$ satisfies $\mathcal{K}$.*

The problem with using pairs $\langle \tau, \mathcal{H} \rangle$ is that they are not sufficient to capture the semantic connection between the TBox and the ABox of a KB. Our investigation shows that it is more plausible to use a set of types (instead of a single type) in such a pair, as an alternative semantic characterization for a DL-Lite$^{\mathcal{N}}_{bool}$ KB. Using sets of types as semantic characterizations of DL-Lite TBoxes is also suggested in (Kontchakov et al. 2008).

Thus, we introduce the definition of a *feature* as follows.

Definition 2 (Features) *Given a signature $\mathcal{S}$, an $\mathcal{S}$-*feature *(or simply* feature *when $\mathcal{S}$ is clear) is defined as a pair $\mathcal{F} = \langle \Xi, \mathcal{H} \rangle$, where Ξ is a non-empty set of $\mathcal{S}$-types and $\mathcal{H}$ a $\mathcal{S}$-Herbrand set, satisfying the following conditions:*

1. *$\exists P \in \bigcup \Xi$ iff $\exists P^- \in \bigcup \Xi$, for each $P \in \mathcal{S}_R$.*
2. *$\tau \in \Xi$, for each $a \in \mathcal{S}_I$ and τ the type of a in $\mathcal{H}$.*

The intuition behind the two conditions in Definition 2 can be easily seen after we explain how features can be used as an alternative for DL-Lite interpretations later.

We define the satisfaction relation of an inclusion or assertion *w.r.t.* a feature $\mathcal{F} = \langle \Xi, \mathcal{H} \rangle$ as follows:

- $\mathcal{F}$ satisfies $C_1 \sqsubseteq C_2$, if τ satisfies $\neg C_1 \sqcup C_2$ for all $\tau \in \Xi$.
- $\mathcal{F}$ satisfies $C(a)$, if the type of a in $\mathcal{H}$ satisfies C.
- $\mathcal{F}$ satisfies $P(a, b)$ or $P^-(b, a)$, if $P(a, b)$ is in $\mathcal{H}$.

We can see that the first condition in Definition 2 guarantees that $\exists P$ is unsatisfiable (*i.e.*, $\exists P \sqsubseteq \bot$ is satisfied) *w.r.t.* the TBox if and only if $\exists P^-$ is also unsatisfiable *w.r.t.* the TBox. The second condition in Definition 2 requires that if $\mathcal{F}$ satisfies assertion $C(a)$, then C must be satisfiable *w.r.t.* the TBox.

We call $\mathcal{F}$ a *model feature* of KB $\mathcal{K}$ if $\mathcal{F}$ satisfies every inclusion and assertion in $\mathcal{K}$. We use $\mathsf{MF}(\mathcal{K})$ to denote the set of all model features of $\mathcal{K}$.

From the definition of features, as we only consider finite signatures, a model feature is always finite in structure, and the number of model features of a KB is also finite.

Example 2 *Consider the KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, where $\mathcal{T} = \{ A \sqsubseteq \exists P,\ B \sqsubseteq \exists P,\ \exists P^- \sqsubseteq B,\ A \sqcap B \sqsubseteq \bot,\ \ge 2\, P^- \sqsubseteq \bot \}$ and $\mathcal{A} = \{ A(a),\ P(a, b) \}$. It is shown in (Calvanese et al. 2006) that $\mathcal{K}$ is a KB having no finite model. In fact, an infinite model $\mathcal{I}$ of $\mathcal{K}$ can be defined as follows:*

$\Delta^{\mathcal{I}} = \{d_a, d_b, d_1, d_2, d_3 \ldots\}$, $a^{\mathcal{I}} = d_a$ and $b^{\mathcal{I}} = d_b$; the concept A is interpreted as a singleton $\{d_a\}$ and B as $\{d_b, d_1, d_2, d_3 \ldots\}$; and role P is interpreted as $\{(d_a, d_b), (d_b, d_1), (d_1, d_2), \ldots, (d_i, d_{i+1}), \ldots\}$.

Take $\mathcal{S} = \mathsf{Sig}(\mathcal{K}) = \{A, B, P, 1, 2, a, b\}$. The (finite) model feature of $\mathcal{K}$ that corresponds to $\mathcal{I}$ is $\mathcal{F} = \langle \Xi, \mathcal{H} \rangle$, where $\Xi = \{\tau_1, \tau_2\}$ with $\tau_1 = \{A, \exists P\}$ and $\tau_2 = \{B, \exists P, \exists P^-\}$, and $\mathcal{H} = \{ A(a),\ \exists P(a),\ B(b),\ \exists P(b),\ \exists P^-(b),\ P(a, b) \}$.

Given an inclusion or assertion α, define $\mathcal{K} \models_f \alpha$ if all features in $\mathsf{MF}(\mathcal{K})$ satisfy α. Given two KBs $\mathcal{K}_1, \mathcal{K}_2$ and $\mathcal{S} = \mathsf{Sig}(\mathcal{K}_1 \cup \mathcal{K}_2)$, define $\mathcal{K}_1 \models_f \mathcal{K}_2$ if $\mathsf{MF}(\mathcal{K}_1) \subseteq \mathsf{MF}(\mathcal{K}_2)$, and $\mathcal{K}_1 \equiv_f \mathcal{K}_2$ if $\mathsf{MF}(\mathcal{K}_1) = \mathsf{MF}(\mathcal{K}_2)$.

The following two results show that model features do capture the semantic properties of DL-Lite KBs.

Proposition 1 *Let $\mathcal{K}$ be a DL-Lite$^{\mathcal{N}}_{bool}$ KB and $\mathcal{S} = \mathsf{Sig}(\mathcal{K})$. Then we have*

- *$\mathcal{K}$ is consistent iff $\mathcal{K}$ has a model feature.*
- *$\mathcal{K} \models (C_1 \sqsubseteq C_2)$ iff $\mathcal{K} \models_f (C_1 \sqsubseteq C_2)$ for any $C_1 \sqsubseteq C_2$ over $\mathcal{S}$.*
- *$\mathcal{K} \models C(a)$ iff $\mathcal{K} \models_f C(a)$ for any $C(a)$ over $\mathcal{S}$.*
- *$\mathcal{K} \models R(a, b)$ iff $\mathcal{K} \models_f R(a, b)$ for any $R(a, b)$ over $\mathcal{S}$.*

Proposition 2 *Let $\mathcal{K}_1, \mathcal{K}_2$ be two DL-Lite$^{\mathcal{N}}_{bool}$ KBs and $\mathcal{S} = \mathsf{Sig}(\mathcal{K}_1 \cup \mathcal{K}_2)$. Then $\mathcal{K}_1 \models \mathcal{K}_2$ iff $\mathcal{K}_1 \models_f \mathcal{K}_2$, and $\mathcal{K}_1 \equiv \mathcal{K}_2$ iff $\mathcal{K}_1 \equiv_f \mathcal{K}_2$.*

Proof (Sketch) Given a signature $\mathcal{S}$, each interpretation $\mathcal{I}$ uniquely determines an $\mathcal{S}$-feature $\mathcal{F}_{\mathcal{I}}$. We also say $\mathcal{I}$ realizes $\mathcal{F}_{\mathcal{I}}$. We can show that $\mathcal{I}$ is a model of $\mathcal{K}$ iff $\mathcal{F}_{\mathcal{I}}$ is a model feature of $\mathcal{K}$. Moreover, $\mathsf{MF}(\mathcal{K}) = \{ \mathcal{F}_{\mathcal{I}} \mid \mathcal{I} \in \mathsf{Mod}(\mathcal{K}) \}$. ■

In the same way as maximal approximation defined for a set of models, we can define the *maximal approximation*

$\mathcal{K}$ of a set $\mathbb{F}$ of $\mathcal{S}$-features: (1) $\mathsf{Sig}(\mathcal{K}) \subseteq \mathcal{S}$, (2) $\mathbb{F} \subseteq \mathsf{MF}(\mathcal{K})$, and (3) there exists no KB $\mathcal{K}'$ over $\mathcal{S}$ such that $\mathbb{F} \subseteq \mathsf{MF}(\mathcal{K}') \subset \mathsf{MF}(\mathcal{K})$.

Feature Distance and Revision

In this section, we define two notions of distance between features, in the spirit of Hamming distance for propositional models. The first distance is defined as the set of concept and role names interpreted differently in the two features. The second distance is based on a generalized notion of symmetric difference. Based on these two distances, we define two specific revision operators for DL-Lite KBs in an analogous way to Satoh's (Satoh 1988), and show that they have desired properties.

Given a set Σ of concept and role names and $\mathcal{S}$-types τ_1, τ_2, denote $\tau_1 \sim_\Sigma \tau_2$ if for all basic concepts B over $\mathcal{S} - \Sigma$, $B \in \tau_1$ iff $B \in \tau_2$.

Let $\mathcal{F}_1 = \langle \Xi_1, \mathcal{H}_1 \rangle$ and $\mathcal{F}_2 = \langle \Xi_2, \mathcal{H}_2 \rangle$ be two $\mathcal{S}$-features, and $\Sigma \subseteq \mathcal{S}_C \cup \mathcal{S}_R$. Define $\mathcal{F}_1 \leftrightarrow_\Sigma \mathcal{F}_2$ if the following conditions are satisfied:

1. For each $\tau_1 \in \Xi_1$, there exists $\tau_2 \in \Xi_2$ s.t. $\tau_1 \sim_\Sigma \tau_2$; conversely, for each $\tau_2 \in \Xi_2$, there exists $\tau_1 \in \Xi_1$ s.t. $\tau_1 \sim_\Sigma \tau_2$.
2. For each $a \in \mathcal{S}_I$, $\tau_1 \sim_\Sigma \tau_2$, where τ_i ($i = 1, 2$) is the type of a in $\mathcal{H}_i$; and $P(a, b) \in \mathcal{H}_1$ iff $P(a, b) \in \mathcal{H}_2$ for each $P \in \mathcal{S}_R - \Sigma$ and $a, b \in \mathcal{S}_I$.

Intuitively, the minimal sets Σ such that $\mathcal{F}_1 \leftrightarrow_\Sigma \mathcal{F}_2$ are the sets of concept and role names on whose interpretations $\mathcal{F}_1, \mathcal{F}_2$ disagree.

Given two KBs $\mathcal{K}_1, \mathcal{K}_2$ and $\mathcal{S} = \mathsf{Sig}(\mathcal{K}_1 \cup \mathcal{K}_2)$, define the distance between $\mathcal{K}_1$ and $\mathcal{K}_2$ as the set of all minimal distances between model features of $\mathcal{K}_1$ and $\mathcal{K}_2$:

$d_f(\mathcal{K}_1, \mathcal{K}_2) = \min_\subseteq(\{ \Sigma \subseteq \mathcal{S}_C \cup \mathcal{S}_R \mid \exists \mathcal{F}_1 \in \mathsf{MF}(\mathcal{K}_1), \exists \mathcal{F}_2 \in \mathsf{MF}(\mathcal{K}_2) \text{ s.t. } \mathcal{F}_1 \leftrightarrow_\Sigma \mathcal{F}_2 \})$.

To define a revision operator $\mathcal{K} \circ \mathcal{K}'$ in analogy to classical model-based revision, we need to specify the subset of $\mathsf{MF}(\mathcal{K}')$ that is *closest* to $\mathsf{MF}(\mathcal{K})$ (*w.r.t.* feature distance).

Definition 3 (S-Revision) *Let $\mathcal{K}, \mathcal{K}'$ be two DL-Lite$^{\mathcal{N}}_{bool}$ KBs and $\mathcal{S} = \mathsf{Sig}(\mathcal{K} \cup \mathcal{K}')$. Define the* s-revision *of $\mathcal{K}$ by $\mathcal{K}'$, denoted $\mathcal{K} \circ_s \mathcal{K}'$, such that $\mathsf{MF}(\mathcal{K} \circ_s \mathcal{K}') = \mathsf{MF}(\mathcal{K}')$ if $\mathsf{MF}(\mathcal{K}) = \emptyset$, and otherwise,*

$$\mathsf{MF}(\mathcal{K} \circ_s \mathcal{K}') = \{ \mathcal{F}' \in \mathsf{MF}(\mathcal{K}') \mid \exists \mathcal{F} \in \mathsf{MF}(\mathcal{K}) \text{ s.t. } \mathcal{F} \leftrightarrow_\Sigma \mathcal{F}' \text{ and } \Sigma \in d_f(\mathcal{K}_1, \mathcal{K}_2) \}.$$

Example 3 *Consider the following KB,*

$$\mathcal{K} = \langle \{ \mathit{PhD} \sqsubseteq \mathit{Student} \sqcap \mathit{Postgrad}, \mathit{Student} \sqsubseteq \neg\exists \mathit{teaches}, \exists \mathit{teaches}^- \sqsubseteq \mathit{Course}, \mathit{Student} \sqcap \mathit{Course} \sqsubseteq \bot \}, \{ \mathit{PhD}(\mathit{Tom}) \} \rangle.$$

The TBox of $\mathcal{K}$ specifies that PhD students are postgraduate students, and students are not allowed to teach any courses, while the ABox states that Tom is a PhD student. Suppose PhD students are actually allowed to teach, and we want to revise $\mathcal{K}$ with $\mathcal{K}' = \langle \{ \mathit{PhD} \sqsubseteq \exists \mathit{teaches} \}, \emptyset \rangle$.

$\mathcal{F}' = \langle \{\tau_1\}, \{\mathit{Student}(\mathit{Tom}), \mathit{Postgrad}(\mathit{Tom})\} \rangle$, where $\tau_1 = \{\mathit{Student}, \mathit{Postgrad}\}$ is a model feature of $\mathcal{K}'$. From the model features of $\mathcal{K}$, take $\mathcal{F} = \langle \{\tau_1, \tau_2\}, \{\mathit{PhD}(\mathit{Tom}), \mathit{Student}(\mathit{Tom}), \mathit{Postgrad}(\mathit{Tom})\} \rangle$ where $\tau_2 = \{\mathit{PhD}, \mathit{Student}, \mathit{Postgrad}\}$. We can see that $\mathcal{F} \leftrightarrow_{\{\mathit{PhD}\}} \mathcal{F}'$ and $\{\mathit{PhD}\} \in d_f(\mathcal{K}, \mathcal{K}')$. Thus, $\mathcal{F}'$ is a model feature of $\mathcal{K} \circ_s \mathcal{K}'$.

In fact, we can show that $\mathcal{K} \circ_s \mathcal{K}'$ is a DKB $\{\mathcal{K}_1, \mathcal{K}_2\}$ where

$$\mathcal{K}_1 = \langle \{ \mathit{PhD} \sqsubseteq \exists \mathit{teaches}, \mathit{Student} \sqsubseteq \neg\exists \mathit{teaches}, \exists \mathit{teaches}^- \sqsubseteq \mathit{Course}, \mathit{Student} \sqcap \mathit{Course} \sqsubseteq \bot \}, \{ \mathit{Student}(\mathit{Tom}), \mathit{Postgrad}(\mathit{Tom}) \} \rangle, \text{ and}$$

$$\mathcal{K}_2 = \langle \{ \mathit{PhD} \sqsubseteq \exists \mathit{teaches}, \mathit{PhD} \sqsubseteq \mathit{Student} \sqcap \mathit{Postgrad}, \mathit{Student} \sqcap \mathit{Course} \sqsubseteq \bot \}, \{ \mathit{PhD}(\mathit{Tom}) \} \rangle.$$

In (Qi et al. 2009), a similar notion of distance (based on signature interpretation difference) is introduced, but it is defined between models. However, their distance uses only concept names. The distance is defined as the number of concept names interpreted differently by the two models if they interpret all roles the same; and otherwise, the distance is simply the number of all concept names. This definition looks not very fine-grained. First, the distance cannot correctly reflect the difference between the two models on their interpretations on roles. The distance of two models interpreting one role name differently is the same as another pair interpreting all role names differently. Second, the distance they used is Dalal's distance, which works well for propositional logic but not be able to distinguish different concept names. For instance, two models disagreeing on concept PhD have the same distance with another pair of models disagreeing on concept $\mathit{Student}$.

However, it is not difficult to generalize the definitions to take role interpretations into consideration. We generalize Definition 1 of (Qi et al. 2009) and define *signature distance* between two models $\mathcal{I}_1$ and $\mathcal{I}_2$ over a signature $\mathcal{S}$ to be $\mathit{diff}_\mathcal{S}(\mathcal{I}_1, \mathcal{I}_2) = \{ E \in \mathcal{S}_C \cup \mathcal{S}_R \mid E^{\mathcal{I}_1} \neq E^{\mathcal{I}_2} \}$. This definition generalizes Qi *et al.*'s approach in that $\mathit{diff}_\mathcal{S}(\mathcal{I}_1, \mathcal{I}_2)$ can also reflect differences between $\mathcal{I}_1$ and $\mathcal{I}_2$ on the way they interpret roles. In order to obtain a more fine-grained comparison of signature distances, we use set containment rather than set cardinality to define minimal distance, and generalize the revision operator for TBoxes (Definition 3) in (Qi et al. 2009) to KB revision as follows: Given two KBs $\mathcal{K}_1, \mathcal{K}_2$ and $\mathcal{S} = \mathsf{Sig}(\mathcal{K}_1 \cup \mathcal{K}_2)$, $\mathcal{K} \circ_m \mathcal{K}' \equiv \mathcal{K}'$ if $\mathcal{K}$ is inconsistent, and otherwise,

$$\mathsf{Mod}(\mathcal{K} \circ_m \mathcal{K}') = \{ \mathcal{I}' \in \mathsf{Mod}(\mathcal{K}') \mid \exists \mathcal{I} \in \mathsf{Mod}(\mathcal{K}) \text{ s.t. } \mathit{diff}_\mathcal{S}(\mathcal{I}, \mathcal{I}') \in d_m(\mathcal{K}, \mathcal{K}') \},$$

where $d_m(\mathcal{K}, \mathcal{K}') = \min_\subseteq(\{ \mathit{diff}_\mathcal{S}(\mathcal{I}, \mathcal{I}') \mid \mathcal{I} \in \mathsf{Mod}(\mathcal{K}), \mathcal{I}' \in \mathsf{Mod}(\mathcal{K}') \})$.

We note that, although the concept-based definition of model distance can be generalized to distinguish both concepts and roles, it is less clear how to compute $\mathcal{K} \circ_m \mathcal{K}'$ in DL-Lite$^{\mathcal{N}}_{bool}$. For this reason, we introduce a method for approximating $\mathcal{K} \circ_m \mathcal{K}'$ by showing its close connection with $\mathcal{K} \circ_s \mathcal{K}'$.

The following result shows that $\mathcal{K} \circ_s \mathcal{K}'$ and $\mathcal{K} \circ_m \mathcal{K}'$ are equivalent up to maximal approximation.

Proposition 3 *Let $\mathcal{K}, \mathcal{K}'$ be two DL-Lite$^{\mathcal{N}}_{bool}$ KBs and $\mathcal{S} = \mathsf{Sig}(\mathcal{K} \cup \mathcal{K}')$. Then $\mathcal{K} \circ_s \mathcal{K}'$ and $\mathcal{K} \circ_m \mathcal{K}'$ have the same maximal approximation.*

Using the algorithms that will be introduced in the next section, we are also able to compute the maximal approximation of $\mathcal{K} \circ_m \mathcal{K}'$.

Motivated by resolving incoherence, (Qi et al. 2009) embeds two further requirements into their definition of model distance (Definitions 4 and 5). We note that corresponding feature-based revision operators can be defined for all three revision operators in (Qi et al. 2009), and can be used to resolve incoherence, as they are there, but, as resolving incoherence is not the focus of this paper, we omit the details here.

An interesting observation is that $\mathcal{K} \circ_s \mathcal{K}'$ can be computed by query-based forgetting. In particular, let $\mathsf{forget}(\mathcal{K}, \Sigma)$ denote a result of $\mathcal{Q}^u_{\mathcal{L}}$-forgetting about Σ in $\mathcal{K}$ (Wang et al. 2007), we have the following connection between revision and forgetting.

Proposition 4 *Let $\mathcal{K}, \mathcal{K}'$ be two consistent DL-Lite$^{\mathcal{N}}_{bool}$ KBs and $\mathcal{S} = \mathsf{Sig}(\mathcal{K} \cup \mathcal{K}')$. Then*

$$\mathcal{K} \circ_s \mathcal{K}' = \{ \mathsf{forget}(\mathcal{K}, \Sigma) \cup \mathcal{K}' \mid \Sigma \in d_f(\mathcal{K}, \mathcal{K}') \},$$

where $\mathsf{forget}(\mathcal{K}, \Sigma)$ is a result of $\mathcal{Q}^u_{\mathcal{L}}$-forgetting about Σ in $\mathcal{K}$.

Proof (Sketch) It is shown in (Wang et al. 2007) that a KB $\mathcal{K}''$ over $\mathcal{S} - \Sigma$ is a result of $\mathcal{Q}^u_{\mathcal{L}}$-forgetting about Σ in $\mathcal{K}$ iff (1) $\mathcal{K} \models \mathcal{K}''$; and (2) for each model $\mathcal{I}' \in \mathsf{Mod}(\mathcal{K}'')$, there exists $\mathcal{I} \in \mathsf{Mod}(\mathcal{K})$ such that $\mathcal{I}$ and $\mathcal{I}'$ realize the same feature over $\mathcal{S} - \Sigma$. It can be verified that the model features of $\mathsf{forget}(\mathcal{K}, \Sigma)$ over $\mathcal{S}$ are $\mathsf{MF}(\mathsf{forget}(\mathcal{K}, \Sigma)) = \{ \mathcal{F}' \mid \text{there exists } \mathcal{F} \in \mathsf{MF}(\mathcal{K}), \textit{s.t.}\ \mathcal{F} \leftrightarrow_\Sigma \mathcal{F}' \}$. From the definition of $\mathcal{K} \circ_s \mathcal{K}'$, we have that the model features of $\mathcal{K} \circ_s \mathcal{K}'$ are exactly those of $\{ \mathsf{forget}(\mathcal{K}, \Sigma) \cup \mathcal{K}' \mid \Sigma \in d_f(\mathcal{K}, \mathcal{K}') \}$. ■

As shown in (Wang et al. 2007), the result of $\mathcal{Q}^u_{\mathcal{L}}$-forgetting is always expressible in DL-Lite$^u_{bool}$, an simple extension of DL-Lite$^{\mathcal{N}}_{bool}$ (Kontchakov et al. 2008). Thus we have shown that $\mathcal{K} \circ_s \mathcal{K}'$ is always be expressible as a DKB in DL-Lite$^u_{bool}$.

Revision under Approximation

For many applications, it is desirable to have the revision as a single DL-Lite$^{\mathcal{N}}_{bool}$ KB rather than a DKB, *i.e.*, the maximal approximation of the revision is desired. However, in most cases, $\mathcal{K} \circ_s \mathcal{K}'$ and $\mathcal{K} \circ_m \mathcal{K}'$ are too weak in preserving knowledge of the original KB, as shown in the following example.

Example 4 *In Example 3, $\mathcal{K} \circ_s \mathcal{K}'$ is a DKB, whose maximal approximation is the following KB,* $\langle \{ \mathit{PhD} \sqsubseteq \exists \mathit{teaches}, \mathit{Student} \sqcap \mathit{Course} \sqsubseteq \bot \}, \{ (\mathit{Student}(\mathit{Tom}), \mathit{Postgrad}(\mathit{Tom}), (\mathit{PhD} \sqcup (\neg\exists \mathit{teaches} \sqcap \neg\exists \mathit{teaches}^-))(\mathit{Tom}) \} \rangle$.

Note that in the above example, knowledge in $\mathcal{K}$ about concept *PhD* and about role *teaches* are totally lost after revision and approximation. In particular, $\mathit{PhD} \sqsubseteq \mathit{Postgrad}$ and $\exists \mathit{teaches}^- \sqsubseteq \mathit{Course}$ are eliminated, though they have nothing to do with the inconsistency. Although assigning preference over concepts (and roles), as in (Qi et al. 2009), may help strengthen the result of revision by eliminating some (less preferred) members from the DKB, it may still suffer from loss of useful information. Recall Example 3, let $\mathcal{T}$ and $\mathcal{T}'$ be the TBoxes of $\mathcal{K}$ and $\mathcal{K}'$, respectively. Under the best revision operator (Definition 5) in (Qi et al. 2009), the revision of $\mathcal{T}$ by $\mathcal{T}'$ is the (single) TBox of $\mathcal{K}_1$, where $\mathit{PhD} \sqsubseteq \mathit{Postgrad}$ is eliminated as an extra cost of restoring coherence.

We argue that, the reason why revision operator $\circ_s$ performs not very well under approximation is that the distance defined on concept and role names is too simple to reflect differences between model features, which can be seen from the following example. Let $\mathcal{F} = \langle \{\tau_1, \tau_3\}, \{A(a), B(a)\} \rangle$, $\mathcal{F}' = \langle \{\tau_1, \tau_2, \tau_3\}, \{A(a), A(b)\} \rangle$, and $\mathcal{F}'' = \langle \{\tau_2\}, \{A(a), A(b), A(c), A(d)\} \rangle$, where $\tau_1 = \{A, B\}$, $\tau_2 = \{A\}$, and $\tau_3 = \emptyset$. Obviously, $\mathcal{F}$ is closer to $\mathcal{F}'$ than to $\mathcal{F}''$. However, such difference cannot be measured using only concept names, as $\mathcal{F} \leftrightarrow_{\{A,B\}} \mathcal{F}'$ and $\mathcal{F} \leftrightarrow_{\{A,B\}} \mathcal{F}''$.

From the above discussion, we can see that it is insufficient if we measure the distance of two features/models in terms of only a set of concepts (and roles). To obtain a better definition of KB revision, we need to introduce a more complex notion of feature distance, which extends the definition of symmetric difference $\triangle$.

Recall that $S_1 \triangle S_2 = (S_1 - S_2) \cup (S_2 - S_1)$ for any two sets S_1 and S_2. Given two $\mathcal{S}$-features $\mathcal{F}_1 = \langle \Xi_1, \mathcal{H}_1 \rangle$ and $\mathcal{F}_2 = \langle \Xi_2, \mathcal{H}_2 \rangle$, we define the *distance* between $\mathcal{F}_1$ and $\mathcal{F}_2$, denoted $\mathcal{F}_1 \triangle \mathcal{F}_2$, as a pair $\langle \Xi_1 \triangle \Xi_2, \mathcal{H}_1 \triangle \mathcal{H}_2 \rangle$. Note that we do not require $\mathcal{H}_1 \triangle \mathcal{H}_2$ to be a Herbrand set.

To compare two distances, given $\mathcal{F}_i = \langle \Xi_i, \mathcal{H}_i \rangle$ for $i = 1, 2, 3, 4$, we could define $\mathcal{F}_1 \triangle \mathcal{F}_2 \subseteq_f \mathcal{F}_3 \triangle \mathcal{F}_4$ if $\Xi_1 \triangle \Xi_2 \subseteq \Xi_3 \triangle \Xi_4$ and $\mathcal{H}_1 \triangle \mathcal{H}_2 \subseteq \mathcal{H}_3 \triangle \mathcal{H}_4$; and $\mathcal{F}_1 \triangle \mathcal{F}_2 \subset_f \mathcal{F}_3 \triangle \mathcal{F}_4$ if $\mathcal{F}_1 \triangle \mathcal{F}_2 \subseteq \mathcal{F}_3 \triangle \mathcal{F}_4$ and $\mathcal{F}_3 \triangle \mathcal{F}_4 \not\subseteq \mathcal{F}_1 \triangle \mathcal{F}_2$. However, our effort shows that such a measure is too weak to preserve enough knowledge of the original KB, as many features are still incomparable under such measure. Instead, we set a preference on Herbrand sets over type sets: $\mathcal{F}_1 \triangle \mathcal{F}_2 \subset_f \mathcal{F}_3 \triangle \mathcal{F}_4$ iff

- $\mathcal{H}_1 \triangle \mathcal{H}_2 \subset \mathcal{H}_3 \triangle \mathcal{H}_4$, or
- $\mathcal{H}_1 \triangle \mathcal{H}_2 = \mathcal{H}_3 \triangle \mathcal{H}_4$ and $\Xi_1 \triangle \Xi_2 \subset \Xi_3 \triangle \Xi_4$.

The intuition behind such a preference on Herbrand set over type set is that, when both TBox inclusions and ABox assertions take part in causing the inconsistency, assertions in the original ABox have priority to be preserved whereas TBox inclusions are candidates for revision. The consideration for this is two-fold: First, due to the nature of revision and update, revision is more suitable for DL TBox change whereas update more suitable for ABoxes. This intuition can be justified by previous approaches in the literature (Liu et al. 2006; Giacomo et al. 2007; Qi et al. 2009). Second, if inconsistency is caused by incoherence of the combination of the TBoxes (as in Example 3), only modifying ABox assertions can help restoring consistency but will leave the resulting TBox incoherent. Giving TBox inclusions priority for change has the merit that inconsistency and incoherence can be resolved at the same time, so that no extra mechanism for is needed for restoring coherence. We will demonstrate this point with an example. But before that, we first provide the definition of our revision operator in terms of the new distance.

Definition 4 (F-Revision) *Let* $\mathcal{K}, \mathcal{K}'$ *be two DL-Lite*$^{\mathcal{N}}_{bool}$ *KBs and* $\mathcal{S} = \mathsf{Sig}(\mathcal{K} \cup \mathcal{K}')$*. Define the* f-revision *of* $\mathcal{K}$ *by* $\mathcal{K}'$*, denoted* $\mathcal{K} \circ_f \mathcal{K}'$*, such that* $\mathsf{MF}(\mathcal{K} \circ_f \mathcal{K}') = \mathsf{MF}(\mathcal{K}')$ *if* $\mathsf{MF}(\mathcal{K}) = \emptyset$*, and otherwise*

$$\mathsf{MF}(\mathcal{K} \circ_f \mathcal{K}') = \{\ \mathcal{F}' \in \mathsf{MF}(\mathcal{K}') \mid \exists \mathcal{F} \in \mathsf{MF}(\mathcal{K}) \text{ s.t. } \forall \mathcal{F}_i \in \mathsf{MF}(\mathcal{K}), \forall \mathcal{F}'_j \in \mathsf{MF}(\mathcal{K}'),\ (\mathcal{F}_i \triangle \mathcal{F}'_j) \not\subset_f (\mathcal{F} \triangle \mathcal{F}')\ \}.$$

The next example shows that $\circ_f$ performs better under maximal approximation.

Example 5 *Consider the KBs* $\mathcal{K}, \mathcal{K}'$ *in Example 3. We can show that the maximal approximation of* $\mathcal{K} \circ_f \mathcal{K}'$ *is*

$$\langle\ \{\ PhD \sqsubseteq Student \sqcap Postgrad,\ PhD \sqsubseteq \exists teaches,\ Student \sqcap \exists teaches \sqsubseteq PhD,\ \exists teaches^- \sqsubseteq Course,\ Student \sqcap Course \sqsubseteq \bot\ \},\ \{\ Student(Tom),\ Postgrad(Tom)\ \}\ \rangle.$$

Note that $Student \sqsubseteq \neg\exists teaches$ is revised (and weakened) to $Student \sqcap \exists teaches \sqsubseteq PhD$. In this way, consistency is restored, as well as coherence. Also, the knowledge in $\mathcal{K}$ is well preserved.

The well-known AGM postulates (R1) – (R6) for propositional belief revision have been adapted to DLs in (Qi et al. 2006). However, the authors present the postulates in terms of models of KBs. In the following, we reformulate them using KB combinations and entailments, in a manner analogous to classical AGM postulates.

(R1) $\mathcal{K} \circ \mathcal{K}' \models_f \mathcal{K}'$;

(R2) if $\mathcal{K} \cup \mathcal{K}'$ is consistent, then $\mathcal{K} \circ \mathcal{K}' = \mathcal{K} \cup \mathcal{K}'$;

(R3) if $\mathcal{K}'$ is consistent, then $\mathsf{MF}(\mathcal{K} \circ \mathcal{K}') \neq \emptyset$;

(R4) if $\mathcal{K}_1 \equiv \mathcal{K}_2$ and $\mathcal{K}'_1 \equiv \mathcal{K}'_2$, then $\mathcal{K}_1 \circ \mathcal{K}'_1 \equiv_f \mathcal{K}_2 \circ \mathcal{K}'_2$;

(R5) $(\mathcal{K} \circ \mathcal{K}') \cup \mathcal{K}'' \models_f \mathcal{K} \circ (\mathcal{K}' \cup \mathcal{K}'')$;

(R6) if $(\mathcal{K} \circ \mathcal{K}') \cup \mathcal{K}''$ is consistent, then $\mathcal{K} \circ (\mathcal{K}' \cup \mathcal{K}'') \models_f (\mathcal{K} \circ \mathcal{K}') \cup \mathcal{K}''$.

As with Satoh's (Satoh 1988) revision operator, both $\circ_s$ and $\circ_f$ satisfy the first five postulates. However, $(R6)$ is not always satisfied by Satoh's, which is also the case for our revision operators.

Theorem 1 *The revision operators defined in Definitions 3 and 4 both satisfy postulates* $(R1)$ – $(R5)$.

If using its maximal approximation to replace the revision, postulates $(R1)$ – $(R4)$ *are always satisfied.*

Algorithm for Computing Revision

In this section, we introduce an algorithm for computing the maximal approximation of revision syntactically. As the two revision operators are based on model features of KBs, we first introduce a method for computing all the model features of a KB, and then show how the maximal approximation of revision can be constructed via model features.

We first introduce a method that computes $\mathsf{MF}(\mathcal{K})$ from $\mathcal{K}$ through syntactic transformations (ref. Figure 1).

Note that each pair $\langle \Xi, \mathcal{H} \rangle$ added to $\mathbb{F}$ is a feature, and satisfies both $\mathcal{T}$ and $\mathcal{A}$. Algorithm 1 returns $\emptyset$ if and only if $\mathcal{K}$ is inconsistent.

Algorithm 1
Input: A DL-Lite$^{\mathcal{N}}_{bool}$ KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ and a signature $\mathcal{S}$.
Output: $\mathsf{MF}(\mathcal{K})$.
Method: Initially, let $\mathbb{F} = \emptyset$.
Step 1. Compute the set $\Xi_{\mathcal{T}}$ of all $\mathcal{S}$-types satisfying $\mathcal{T}$.
Let $\mathcal{P} = \{\ P(a,b) \mid P(a,b) \text{ or } P^-(b,a) \in \mathcal{A}\ \}$.
Step 2. Add into $\mathbb{F}$ all pairs $\langle \Xi, \mathcal{H} \rangle$ such that:
1. $\Xi \subseteq \Xi_{\mathcal{T}}$, and $\exists P \in \bigcup \Xi$ iff $\exists P^- \in \bigcup \Xi$ for all $P \in \mathcal{S}_R$.
2. $\mathcal{P} \subseteq \mathcal{H}$, and for each $a \in \mathcal{S}_I$, the type τ of a in $\mathcal{H}$ satisfies the following conditions:
(1) $\tau \in \Xi$, and τ satisfies C for every $C(a) \in \mathcal{A}$.
(2) $\geqslant m\ P \in \tau$, for each $P \in \mathcal{S}_R$ *s.t.* $P(a, b_i) \in \mathcal{H}$ with $i = 1, \ldots, n$, and $m \in \mathcal{S}_N$ with $m \leq n$.
(3) $\geqslant m\ P^- \in \tau$, for each $P \in \mathcal{S}_R$ *s.t.* $P(b_i, a) \in \mathcal{H}$ with $i = 1, \ldots, n$, and $m \in \mathcal{S}_N$ with $m \leq n$.
Step 3. Return $\mathbb{F}$ as $\mathsf{MF}(\mathcal{K})$.

Figure 1: Compute model features.

Theorem 2 *Given a DL-Lite*$^{\mathcal{N}}_{bool}$ *KB* $\mathcal{K}$ *and a signature* $\mathcal{S}$*, Algorithm 1 always returns* $\mathsf{MF}(\mathcal{K})$*.*

Given that $\mathsf{MF}(\mathcal{K})$ and $\mathsf{MF}(\mathcal{K}')$ can be computed by Algorithm 1, and they are finite, it is straightforward to obtain $\mathsf{MF}(\mathcal{K} \circ_s \mathcal{K}')$ and $\mathsf{MF}(\mathcal{K} \circ_f \mathcal{K}')$ through Definitions 3 and 4. Now we show that the maximal approximation of $\mathcal{K} \circ \mathcal{K}'$ can be constructed from $\mathsf{MF}(\mathcal{K} \circ \mathcal{K}')$. Indeed, the maximal approximation of $\mathbb{F}$ can be constructed for any set $\mathbb{F}$ of model features in the same way. Given a $\mathcal{S}$-type τ, we denote the concept $C_\tau = \bigsqcap_{B \in \tau} B \sqcap \bigsqcap_{B \notin \tau} \neg B$, where B is a basic concept over $\mathcal{S}$. In what follows, we present an algorithm for DL-Lite$^{\mathcal{N}}_{bool}$ KB revision (ref. Figure 2).

Algorithm 2
Input: Two DL-Lite$^{\mathcal{N}}_{bool}$ KBs $\mathcal{K}$ and $\mathcal{K}'$, $\mathcal{S} = \mathsf{Sig}(\mathcal{K} \cup \mathcal{K}')$.
Output: $\mathcal{K} \circ_f \mathcal{K}'$.
Method: Initially, let $\mathcal{T} = \emptyset$ and $\mathcal{A} = \emptyset$.
Step 1. Use Algorithm 1 to compute $\mathsf{MF}(\mathcal{K})$ and $\mathsf{MF}(\mathcal{K}')$.
Step 2. Obtain $\mathsf{MF}(\mathcal{K} \circ_f \mathcal{K}')$ from $\mathsf{MF}(\mathcal{K})$ and $\mathsf{MF}(\mathcal{K}')$ by Definition 4.
Step 3. For each $\mathcal{S}$-type τ not occurring in any type set in $\mathsf{MF}(\mathcal{K} \circ_f \mathcal{K}')$, add inclusion $C_\tau \sqsubseteq \bot$ into $\mathcal{T}$.
Step 4. For each individual $a \in \mathcal{S}_I$, add concept assertion $(\bigsqcup_{\tau \in \Xi_a} C_\tau)(a)$ into $\mathcal{A}$, where $\Xi_a = \{\ \tau \mid \exists \langle \Xi, \mathcal{H} \rangle \in \mathsf{MF}(\mathcal{K} \circ_f \mathcal{K}')$ *s.t.* τ is the type of a in $\mathcal{H}\ \}$.
Step 5. For each role assertion $P(a,b)$ occurring in every Herbrand set in $\mathsf{MF}(\mathcal{K} \circ_f \mathcal{K}')$, add $P(a,b)$ into $\mathcal{A}$.
Step 6. Return $\langle \mathcal{T}, \mathcal{A} \rangle$ as $\mathcal{K} \circ_f \mathcal{K}'$.

Figure 2: Compute f-revision.

Note that Algorithm 2 does not rely on the definition of a specific revision operator, as long as it is defined in terms of feature distance. Hence, Algorithm 2 can be used to compute $\mathcal{K} \circ_s \mathcal{K}'$ if Definition 3 is adopted in Step 2. Such an algorithm based on features gives us the flexibility to compute other possible revisions proposed according different application needs.

Theorem 3 *Given two consistent DL-Lite$^{\mathcal{N}}_{bool}$ KBs $\mathcal{K}$ and $\mathcal{K}'$, Algorithm 2 always returns the maximal approximation of $\mathcal{K} \circ_f \mathcal{K}'$.*

Proof (Sketch) Let $\mathcal{K}^*$ be the output KB in Algorithm 2. We can show $\mathsf{MF}(\mathcal{K} \circ_f \mathcal{K}') \subseteq \mathsf{MF}(\mathcal{K}^*)$.

Suppose $\mathsf{MF}(\mathcal{K} \circ_f \mathcal{K}') = \{\mathcal{F}_1, \dots, \mathcal{F}_n\}$ ($n \geq 1$) with $\mathcal{F}_i = \langle \Xi_i, \mathcal{H}_i \rangle$ for $i = 1, \dots, n$. The following conditions hold for each model feature $\mathcal{F} = \langle \Xi, \mathcal{H} \rangle$ of $\mathcal{K}^*$: (1) $\Xi \subseteq \bigcup_{1 \leq i \leq n} \Xi_i$. (2) For each individual $a \in \mathcal{S}_I$, the type of a in $\mathcal{H}$ is the type of a in some $\mathcal{H}_i$ ($1 \leq i \leq n$). (3) $P(a, b) \in \mathcal{H}$ if $P(a, b) \in \bigcap_{1 \leq i \leq n} \mathcal{H}_i$ for $P \in \mathcal{S}_I$ and $a, b \in \mathcal{S}_I$. Thus, $\mathcal{F}$ also satisfies $\mathcal{K}''$ for any $\mathcal{K}''$ over $\mathcal{S}$ such that $\mathsf{MF}(\mathcal{K} \circ_f \mathcal{K}') \subseteq \mathsf{MF}(\mathcal{K}')$. That is, $\mathsf{MF}(\mathcal{K}^*) \subseteq \mathsf{MF}(\mathcal{K}')$. ■

The problem of computing the revision is decidable although it may be exponential in the worst case.

Indeed, although in Definitions 3 and 4, $\mathcal{S}$ equals $\mathsf{Sig}(\mathcal{K} \cup \mathcal{K}')$, we can show that only a subset of $\mathcal{K}$, which is *relevant* to $\mathcal{K}'$, needs to be revised, whereas the other irrelevant inclusions and assertions in $\mathcal{K}$ can be added directly into the result of revision. That is, suppose $\mathcal{K}'' \subset \mathcal{K}$ is the subset of $\mathcal{K}$ that is relevant to $\mathcal{K}'$, then we can use $\mathcal{K}''$ and $\mathcal{K}'$ to be the inputs of Algorithm 2, with $\mathcal{S} = \mathsf{Sig}(\mathcal{K}'' \cup \mathcal{K}')$, and union the output with $\mathcal{K} - \mathcal{K}''$.

A notion of (signature) relevance is formally defined in propositional logic via *language splitting* (Parikh 1996). In what follows, we adapt this notion to DLs and show its applications in KB revision. Given a KB $\mathcal{K}$, we say that a set of signatures $\mathbb{S} = \{\mathcal{S}_1 \dots, \mathcal{S}_n\}$ ($n \geq 1$) is a *splitting* of $\mathcal{K}$ if $\mathcal{S}_i \cap \mathcal{S}_j \subseteq \mathcal{S}_N$ for $1 \leq i < j \leq n$, where $\mathcal{S}_N$ consists of all the numbers in $\mathbb{S}$, and there exist KBs $\mathcal{K}_1 \dots, \mathcal{K}_n$ such that $\mathsf{Sig}(\mathcal{K}_i) \subseteq \mathcal{S}_i$ ($1 \leq i \leq n$) and $\mathcal{K} \equiv \bigcup_{1 \leq i \leq n} \mathcal{K}_i$.

The following result states that our revision operators enjoy a decomposition property regarding signature relevance. We use $\circ$ to denote either one of the revision operators defined in Definitions 3 and 4.

Proposition 5 *Let $\{\mathcal{S}_1, \mathcal{S}_2\}$ be a splitting of $\mathcal{K}$, with $\mathcal{K} \equiv \mathcal{K}_1 \cup \mathcal{K}_2$ and $\mathsf{Sig}(\mathcal{K}_i) \subseteq \mathcal{S}_i$ ($i = 1, 2$). Suppose $\mathsf{Sig}(\mathcal{K}') \cap \mathcal{S}_2 \subseteq \mathcal{S}_N$ with $\mathcal{S}_N$ consisting of all the numbers in $\{\mathcal{S}_1, \mathcal{S}_2\}$, then $\mathcal{K} \circ \mathcal{K}' = (\mathcal{K}_1 \circ \mathcal{K}') \cup \mathcal{K}_2$.*

In applications, the new knowledge $\mathcal{K}'$ is often small compared to the large existing KB $\mathcal{K}$. The subset $\mathcal{K}_1$ of $\mathcal{K}$ that is relevant to $\mathcal{K}'$ is usually also small. Thus, it is reasonable to expect the signature $\mathcal{S} = \mathsf{Sig}(\mathcal{K}_1 \cup \mathcal{K}')$ to be relatively small in practical revision algorithms, which grants efficient revision computation.

Conclusion

We have developed a formal framework for revising general KBs (with no specific restriction) in DL-Lite$^{\mathcal{N}}_{bool}$, based on the notion of features. We have defined two specific revision operators, which are natural adaption of model-based revision approaches from propositional logic. We have shown that the second operator performs better *w.r.t.* maximal approximation. We have also developed algorithms for computing maximal approximations of KB revisions in DL-Lite$^{\mathcal{N}}_{bool}$. We note that other propositional revision operators, e. g., Dalal's revision, belief contraction and update can also be easily defined in our framework.

There are several interesting issues remaining for future work. One is to extend our approach to KB revision in other DLs. However, for more expressive DLs, the structure of features need to be more complex. Also, it is interesting to develop more efficient algorithms for each specific revision operator. Another problem is to look at applications of our revision operator in nonmonotonic reasoning problems in DLs.

References

Artale, A.; Calvanese, D.; Kontchakov, R.; and Zakharyaschev, M. 2007. DL-Lite in the light of first-order logic. In *Proc. AAAI*, 361–366.

Calvanese, D.; Giacomo, G. D.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2006. Data complexity of query answering in description logics. In *Proc. KR*, 260–270.

Calvanese, D.; Giacomo, G. D.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2007. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *J. Autom. Reasoning* 39(3):385–429.

Eiter, T., and Gottlob, G. 1992. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence* 57:227–270.

Flouris, G.; Huang, Z.; Pan, J. Z.; Plexousakis, D.; and Wache, H. 2006. Inconsistencies, negations and changes in ontologies. In *Proc. AAAI*.

Giacomo, G. D.; Lenzerini, M.; Poggi, A.; and Rosati, R. 2007. On the approximation of instance level update and erasure in description logics. In *Proc. AAAI*, 403–408.

Halaschek-Wiener, C.; Parsia, B.; Sirin, E.; and Kalyanpur, A. 2006. Description logic reasoning for dynamic ABoxes. In *Proc. DL*.

Kontchakov, R.; Wolter, F.; and Zakharyaschev, M. 2008. Can you tell the difference between DL-Lite ontologies? In *Proc. KR*, 285–295. AAAI Press.

Liu, H.; Lutz, C.; Milicic, M.; and Wolter, F. 2006. Updating description logic ABoxes. In *Proc. KR*, 46–56.

Meyer, T.; Lee, K.; and Booth, R. 2005. Knowledge integration for description logics. In *Proc. AAAI*, 645–650.

Parikh, R. 1996. Beliefs, belief revision, and splitting languages. Logic, Language and Computation, 266–268. CSLI Publications.

Qi, G., and Du, J. 2009. Model-based revision operators for terminologies in description logics. In *Proc. IJCAI*.

Qi, G.; Haase, P.; Huang, Z.; Ji, Q.; Pan, J. Z.; and Völker, J. 2008. A kernel revision operator for terminologies - algorithms and evaluation. In *Proc. ISWC*, 419–434.

Qi, G.; Liu, W.; and Bell, D. A. 2006. Knowledge base revision in description logics. In *Proc. JELIA*.

Ribeiro, M., and Wassermann, R. 2007. Base revision in description logics - preliminary results. In *Proc. International Workshop on Ontology Dynamics*.

Satoh, K. 1988. Nonmonotonic reasoning by minimal belief revision. In *Proc. International Conference on Fifth Generation Computer Systems*, 455–462. Springer-Verlag.

Wang, Z., and Wang, K., and Topor, R., and Pan, J. Z. 2009. Forgetting for knowledge bases in DL-Lite. Technical Report. http://hobbit.ict.griffith.edu.au/ kewen/forgetting.pdf.