On the Approximation of Instance Level Update and Erasure in Description Logics

Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, Riccardo Rosati

Dipartimento di Informatica e Sistemistica
Sapienza Università di Roma, Italy
lastname@dis.uniroma1.it

Abstract

A Description Logics knowledge base is constituted by two components, called TBox and ABox, where the former expresses general knowledge about the concepts and their relationships, and the latter describes the properties of instances of concepts. We address the problem of how to deal with changes to a Description Logic knowledge base, when these changes affect only its ABox. We consider two types of changes, namely update and erasure, and we characterize the semantics of these operations on the basis of the approaches proposed by Winslett and by Katsuno and Mendelzon. It is well known that, in general, Description Logics are not closed with respect to updates, in the sense that the set of models corresponding to an update applied to a knowledge base in a Description Logic $\mathcal{L}$ may not be expressible by ABoxes in $\mathcal{L}$. We show that this is true also for erasure. To deal with this problem, we introduce the notion of best approximation of an update (erasure) in a DL $\mathcal{L}$, with the goal of characterizing the $\mathcal{L}$ ABoxes that capture the update (erasure) at best. We then focus on *DL-Lite*$_\mathcal{F}$, a tractable Description Logic, and present polynomial algorithms for computing the best approximation of updates and erasures in this logic, which shows that the nice computational properties of *DL-Lite*$_\mathcal{F}$ are retained in dealing with the evolution of the ABox.

Introduction

Recent years have witnessed a strong interest in Description Logics (Baader *et al.* 2003) (DL). From one hand, DLs are widely considered as the logical basis for representing and reasoning over web ontologies[1]. On the other hand, the use of DL systems such as RacerPro, FaCT++, Pellet, QuOnto[2] is becoming more and more common in various application areas. These systems essentially provide basic reasoning services over DL knowledge bases, such as consistency checking, subsumption, instance checking, and, at least some of them, query answering. However, the kind of support they provide is somehow limited to reasoning over a static knowledge base. This has to be contrasted with the increasing interest in supporting the *evolution* of a DL knowledge base. Indeed, recent papers, in particular, (Haase & Stojanovic 2005; Liu *et al.* 2006; De Giacomo *et al.* 2006; Gutiérrez, Hurtado, & Vaisman 2006) have started to formally address this issue.

In this paper, we study DL knowledge base evolution under three main assumptions. First, as done in (De Giacomo *et al.* 2006), we restrict our attention to the case where the evolution concerns only the *instance level* of the knowledge base (ABox), and therefore the change to the knowledge base does not modify its intensional level, i.e., the general knowledge about the concepts and their relationships, represented by the TBox. We share the view of (De Giacomo *et al.* 2006), that studying instance-level evolution is an interesting starting point of a principled approach to the problem, as it is very common that the ABox will change much more frequently than the TBox in many real-world contexts. Second, we consider two basic evolution operations, called update and erasure, roughly corresponding to the addition and the deletion of a set of facts, respectively. Following (Liu *et al.* 2006; Gutiérrez, Hurtado, & Vaisman 2006; De Giacomo *et al.* 2006), we base our investigation on a formal semantics for these operations proposed by Winslett (Winslett 1988; 1990) and Katsuno and Mendelzon (Katsuno & Mendelzon 1991). Third, since every DL system is tailored to a specific DL, we believe that the notion of evolution should be studied under the assumption that, given a fixed DL $\mathcal{L}$, the result of any evolution operation must be expressed in $\mathcal{L}$.

One of the main contributions of (Liu *et al.* 2006) was to show that there are cases where, for a fixed DL $\mathcal{L}$ used to express the knowledge base, the result of an instance level evolution operation, in particular update interpreted under Winslett's semantics, is not expressible in $\mathcal{L}$. This is one of the fundamental problems we want to address in this paper. More precisely, we present the following contributions.

- We show that the non-expressibility problem arises also for erasure, even in the context of very simple DLs, and very simple forms of ABox assertions.
- To address the non-expressibility problem, we propose to resort to the notion of maximal approximation of update and erasure. Roughly speaking, given a fixed DL $\mathcal{L}$, and a knowledge base $\mathcal{K}$ in $\mathcal{L}$ constituted by a TBox and an ABox, the maximal approximation of an update (resp., erasure) of $\mathcal{K}$ is the knowledge base in $\mathcal{L}$ that has the same TBox as the original knowledge base, and whose set of models captures at best the set of models characterizing the result of the update (resp., erasure). To the best of our knowledge, this is the first proposal of dealing with the non-expressibility problem by means of the notion of

[1] http://www.w3.org/TR/owl-features/

[2] http://www.cs.man.ac.uk/~sattler/reasoners.html

approximation.

- We consider the *DL-Lite*$_{\mathcal{F}}$ (Calvanese *et al.* 2005) DL, where all reasoning tasks can be done in polynomial time with respect to the size of the knowledge base. We show that, in general, the result of both an update and an erasure on a *DL-Lite*$_{\mathcal{F}}$ knowledge base is not expressible in *DL-Lite*$_{\mathcal{F}}$. We then present algorithms for both approximated update and approximated erasure in this logic. We also show that these algorithms have polynomial time complexity, and that instance checking on the knowledge base resulting from an approximated update (resp., erasure) provides exactly the same answers as instance checking over the set of models capturing the exact result of the update (resp., erasure). Note that the same does not hold for conjunctive query answering.

Preliminaries

DL knowledge bases Description Logics (DLs) (Baader *et al.* 2003) are knowledge representation formalisms, tailored for representing the domain of interest in terms of *concepts* (or classes), which denote sets of objects, and *roles* (or relations), which denote binary relations between objects. A DL *knowledge base* (KB) $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ is formed by two distinct parts $\mathcal{T}$ and $\mathcal{A}$, called *TBox* and *ABox*, respectively. The TBox $\mathcal{T}$ is a finite set of assertions representing the *intensional level* of the KB, i.e., providing an intensional description of the domain of interest. The *ABox* $\mathcal{A}$ is a finite set of assertions about the extension of concepts and roles, i.e., providing information on the *instance level* of the KB.

The semantics of a DL KB is given in terms of interpretations over a fixed infinite domain Δ of objects. We assume to have one constant for each object in Δ denoting exactly that object. In this way, we blur the distinction between constants and objects, so that we can use them interchangeably (with a little abuse of notation). An interpretation $\mathcal{I} = \langle \Delta, \cdot^{\mathcal{I}} \rangle$ consists of a first order structure over Δ, where $\cdot^{\mathcal{I}}$ is the interpretation function, i.e., a function mapping each concept to a subset of Δ and each role to a subset of $\Delta \times \Delta$. We say that $\mathcal{I}$ is a *model of an assertion* (i.e., either a TBox or an ABox assertion) α if $\mathcal{I}$ *satisfies* α, i.e., α is true in $\mathcal{I}$. We say that $\mathcal{I}$ is a *model of the KB* $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, if $\mathcal{I}$ is a model of all the assertions in $\mathcal{T}$ and $\mathcal{A}$.

Given a set of assertions $\mathcal{S}$, we denote as $Mod(\mathcal{S})$ the set of models of all assertions in $\mathcal{S}$. In particular, the set of *models of a KB* $\mathcal{K}$, denoted as $Mod(\mathcal{K})$, is the set of models of all assertions in $\mathcal{T}$ and $\mathcal{A}$, i.e. $Mod(\mathcal{K}) = Mod(\langle \mathcal{T}, \mathcal{A} \rangle) = Mod(\mathcal{T} \cup \mathcal{A})$. A KB $\mathcal{K}$ is *consistent* if $Mod(\mathcal{K}) \neq \emptyset$, i.e. it has at least one model. A set of models $\mathcal{M}$ *logically implies* an assertion α, written $\mathcal{M} \models \alpha$, if for every interpretation $\mathcal{I} \in \mathcal{M}$, we have $\mathcal{I} \in Mod(\alpha)$, i.e. all the models in $\mathcal{M}$ are also models of α. We say that a KB $\mathcal{K}$ *logically implies* an assertion α, written $\mathcal{K} \models \alpha$, if $Mod(\mathcal{K}) \models \alpha$.

***DL-Lite*$_{\mathcal{F}}$** In this paper, we focus on a particular DL, called *DL-Lite*$_{\mathcal{F}}$, belonging to the *DL-Lite* family (Calvanese *et al.* 2005; 2006). DLs in this family are tailored towards capturing conceptual modeling constructs (such as those typical of ontology languages, UML Class Diagrams and Entity-Relationship Diagrams), while keeping reasoning, including conjunctive query answering, tractable and first-order reducible (i.e, LOGSPACE in data complexity). In *DL-Lite*$_{\mathcal{F}}$ that is the logic originating the whole *DL-Lite* family in (Calvanese *et al.* 2005), concepts are defined as follows:

$$\begin{array}{rcl} B & ::= & A \mid \exists R \\ C & ::= & B \mid \neg B \\ R & ::= & P \mid P^- \end{array}$$

where A denotes an atomic concept, P an atomic role, B a basic concept, and C a general concept. A basic concept can be either an atomic concept, a concept of the form $\exists P$, i.e. the standard DL construct of unqualified existential quantification on roles, or a concept of the form $\exists P^-$, which involves *inverse roles*. A *DL-Lite*$_{\mathcal{F}}$ TBox is a finite set of universal assertions of the form

$$\begin{array}{ll} B_1 \sqsubseteq B_2 & \text{inclusion assertion} \\ B_1 \sqsubseteq \neg B_2 & \text{disjointness assertion} \\ (\mathsf{funct}\ R) & \text{functionality assertion} \end{array}$$

An inclusion assertion specifies that each instance of the basic concept B_1 is also an instance of the basic concept B_2, i.e., B_1 is subsumed by B_2. A disjointness assertion specifies that the set of instances of B_1 and B_2 are disjoint. Finally, a functionality assertion expresses the functionality of an atomic role, or of the inverse of an atomic role. Note that disjunction is disallowed, and that negation is used in a restricted way, in particular for asserting disjointness of concepts. Notably, if we remove either of these limitations, reasoning becomes intractable – see (Calvanese *et al.* 2006).

A *DL-Lite*$_{\mathcal{F}}$ ABox is a finite set of membership assertions of the form

$$B(a), \quad R(a, b)$$

stating, respectively, that the object a is an instance of the basic concept B, and that the pair of objects (a, b) is an instance of the role R.

Given an interpretation $\mathcal{I} = \langle \Delta, \cdot^{\mathcal{I}} \rangle$ the interpretation function $\cdot^{\mathcal{I}}$ interprets the constructs of *DL-Lite*$_{\mathcal{F}}$ as follows:

$$\begin{array}{ll} A^{\mathcal{I}} \subseteq \Delta & (P^-)^{\mathcal{I}} = \{(c', c) \mid (c, c') \in P^{\mathcal{I}}\} \\ P^{\mathcal{I}} \subseteq \Delta \times \Delta & (\exists R)^{\mathcal{I}} = \{c \mid \exists c'.(c, c') \in R^{\mathcal{I}}\} \\ (\neg B)^{\mathcal{I}} = \Delta \setminus B^{\mathcal{I}} & \end{array}$$

An interpretation $\mathcal{I}$ is a *model* of an inclusion assertion $B_1 \sqsubseteq B_2$ if $B_1^{\mathcal{I}} \subseteq B_2^{\mathcal{I}}$. $\mathcal{I}$ is a model of a disjointness assertion $B_1 \sqsubseteq \neg B_2$ if $B_1^{\mathcal{I}} \cap B_2^{\mathcal{I}} = \emptyset$. $\mathcal{I}$ is a model of a functionality assertion $(\mathsf{funct}\ R)$ if for all c, c', c'', $(c, c') \in R^{\mathcal{I}}$ and $(c, c'') \in R^{\mathcal{I}}$ implies $c' = c''$. $\mathcal{I}$ is a model of a membership assertion $C(a)$ (resp., $R(a, b)$) iff $a \in C^{\mathcal{I}}$ (resp., $(a, b) \in R^{\mathcal{I}}$).

We observe that the form of the *DL-Lite*$_{\mathcal{F}}$ ABox is that of the original proposal in (Calvanese *et al.* 2005), and is a restriction with respect to the one studied in (De Giacomo *et al.* 2006), where instance-level updates for the *DL-Lite* family were first introduced. Specifically, here we do not allow for negation or "variables" in the membership assertions. We will see that, without these extensions, *DL-Lite*$_{\mathcal{F}}$ is akin to the vast majority of DLs, see (Liu *et al.* 2006), in that the result of the update (or the erasure for the matter) is not expressible as a new *DL-Lite*$_{\mathcal{F}}$ ABox, thus motivating the need for approximation.

DL instance-level updates and erasure Several approaches for changing a knowledge base have been considered in the literature, see, e.g., (Herzig & Rifi 1999), for a survey. Following the work in (De Giacomo *et al.* 2006), we adopt the Winslett's notion of update (Winslett 1988; 1990) and its counterpart, defined in (Katsuno & Mendelzon 1991), as notion of erasure. However, as in (De Giacomo *et al.* 2006) we refine such notions to take into account that we are interested in studying changes at the instance-level, while we insist that the intensional level of the knowledge base is considered stable and hence remains invariant. Intuitively the result of updating (erasing) a KB $\mathcal{K}$ with a finite set of formulas $\mathcal{F}$, should be any KB that logically implies (does not logically imply) all formulas in $\mathcal{F}$, and whose set of models minimally differs from the set of models of $\mathcal{K}$. Taking into account that one of our basic assumption is to maintain the TBox unchanged, the definitions of update and erasure in our context are as follows.

Definition 1 Let $\mathcal{I} = \langle \Delta, \cdot^{\mathcal{I}} \rangle$ and $\mathcal{I}' = \langle \Delta, \cdot^{\mathcal{I}'} \rangle$ be two interpretations over the same alphabet. We say that $\mathcal{I}$ *is contained in* $\mathcal{I}'$, written $\mathcal{I} \subseteq \mathcal{I}'$, if $\mathcal{I}, \mathcal{I}'$ are such that (*i*) $A^{\mathcal{I}} \subseteq A^{\mathcal{I}'}$, for every atomic concept A, and (*ii*) $R^{\mathcal{I}} \subseteq R^{\mathcal{I}'}$, for every atomic role R. We say that $\mathcal{I}$ *is properly contained* $\mathcal{I}'$, written $\mathcal{I} \subset \mathcal{I}'$, if $\mathcal{I} \subseteq \mathcal{I}'$ and $\mathcal{I}' \not\subseteq \mathcal{I}$.

Definition 2 Let $\mathcal{I} = \langle \Delta, \cdot^{\mathcal{I}} \rangle$ and $\mathcal{I}' = \langle \Delta, \cdot^{\mathcal{I}'} \rangle$ be two interpretations (over the same alphabet). We define the *difference between* $\mathcal{I}$ *and* $\mathcal{I}'$, written $\mathcal{I} \ominus \mathcal{I}'$, as the interpretation $(\Delta, \cdot^{\mathcal{I} \ominus \mathcal{I}'})$ such that:

- $A^{\mathcal{I} \ominus \mathcal{I}'} = A^{\mathcal{I}} \ominus A^{\mathcal{I}'}$, for every atomic concept A
- $R^{\mathcal{I} \ominus \mathcal{I}'} = R^{\mathcal{I}} \ominus R^{\mathcal{I}'}$, for every atomic role R,

where $S \ominus S'$ denotes the symmetric difference between sets S and S', i.e. $S \ominus S' = (S \cup S') \setminus (S \cap S')$.

Definition 3 Let $\mathcal{T}$ be a TBox in a DL $\mathcal{L}$, $\mathcal{I}$ a model of $\mathcal{T}$, and $\mathcal{F}$ a finite set of membership assertions expressed in $\mathcal{L}$ such that $Mod(\mathcal{T} \cup \mathcal{F}) \neq \emptyset$. The *update of* $\mathcal{I}$ *with* $\mathcal{F}$, denoted $U^{\mathcal{T}}(\mathcal{I}, \mathcal{F})$, is defined as follows:

$$U^{\mathcal{T}}(\mathcal{I}, \mathcal{F}) = \{\mathcal{I}' \mid \mathcal{I}' \in Mod(\mathcal{T} \cup \mathcal{F}) \text{ and there exists no } \mathcal{I}'' \in Mod(\mathcal{T} \cup \mathcal{F}) \text{ s.t. } \mathcal{I} \ominus \mathcal{I}'' \subset \mathcal{I} \ominus \mathcal{I}'\}$$

Observe that $U^{\mathcal{T}}(\mathcal{I}, \mathcal{F})$ is the set of models of both $\mathcal{T}$ and $\mathcal{F}$ whose difference with respect to $\mathcal{I}$ is $\subseteq$-minimal. With these notions in place we can define instance-level update and erasure.

Definition 4 (**Update**) Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a KB expressed in a DL $\mathcal{L}$, and $\mathcal{F}$ a finite set of membership assertions expressed in $\mathcal{L}$ such that $Mod(\mathcal{T} \cup \mathcal{F}) \neq \emptyset$. The *update of* $\mathcal{K}$ *with* $\mathcal{F}$, denoted $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$, is defined as follows:

$$\mathcal{K} \circ_{\mathcal{T}} \mathcal{F} = \bigcup_{\mathcal{I} \in Mod(\mathcal{K})} U^{\mathcal{T}}(\mathcal{I}, \mathcal{F}).$$

Definition 5 (**Erasure**) Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a KB expressed in a DL $\mathcal{L}$, and $\mathcal{F}$ a finite set of membership assertions expressed in $\mathcal{L}$ such that $Mod(\mathcal{T} \cup \neg\mathcal{F}) \neq \emptyset$, where $\neg\mathcal{F}$ denotes the set of membership assertions $\{\neg F_i \mid F_i \in \mathcal{F}\}$. The *erasure of* $\mathcal{K}$ *with* $\mathcal{F}$, denoted $\mathcal{K} \bullet_{\mathcal{T}} \mathcal{F}$, is defined as follows:

$$\mathcal{K} \bullet_{\mathcal{T}} \mathcal{F} = Mod(\mathcal{K}) \cup (\bigcup_{\mathcal{I} \in Mod(\mathcal{K})} U^{\mathcal{T}}(\mathcal{I}, \neg\mathcal{F})).$$

Observe that, from the above definition, erasing a set of assertions means simultaneously erasing each assertion in the set. Note also that assertions in the set $\neg F_i$ might not be in the language $\mathcal{L}$, in particular when $\mathcal{L}$ is not closed with respect to negation of membership assertions. However, even in this case, the above semantics of erasure is well-defined, provided that we refer to the traditional first-order semantics for negated membership assertions.

A basic question that arises when dealing with instance-level update and erasure is whether the result of updating/erasing a KB in a given DL $\mathcal{L}$ can be expressed as a new KB again expressed in $\mathcal{L}$ (with the same TBox, since the TBox is invariant). Unfortunately, as shown in (Liu *et al.* 2006; De Giacomo *et al.* 2006), in general the result of update cannot be expressed in the same language as the original KB. In particular, we will show in the next sections that neither update nor erasure can always be expressed in *DL-Lite*$_{\mathcal{F}}$. This fundamental problem leads us to study *approximated* instance-level update and erasure.

Approximation

To study approximated instance-level update and erasure, we first define a general notion of approximation (Schaerf & Cadoli 1995) that is natural in our setting. What we aim at is to approximate in the best possible way the actual update and erasure, while still remaining in the language of the original KB. Since we cannot change the TBox, this means that the approximated update and erasure must be encoded in the ABox.

Definition 6 (**Sound** $(\mathcal{L}, \mathcal{T})$**-Approximation**) Let $\mathcal{M}$ be a set of models, and $\mathcal{T}$ a TBox in a DL $\mathcal{L}$ such that $\mathcal{M} \subseteq Mod(\mathcal{T})$. We say that a DL KB $\mathcal{K}$ is a *sound* $(\mathcal{L}, \mathcal{T})$*-approximation* of $\mathcal{M}$ in $\mathcal{L}$, if (*i*) $\mathcal{K}$ is in $\mathcal{L}$, (*ii*) $\mathcal{K}$ is of the form $\langle \mathcal{T}, \mathcal{A} \rangle$, and (*iii*) $\mathcal{M} \subseteq Mod(\mathcal{K})$.

Definition 7 (**Maximal** $(\mathcal{L}, \mathcal{T})$**-Approximation**) Let $\mathcal{M}$ be a set of models, and $\mathcal{T}$ a TBox in a DL $\mathcal{L}$ such that $\mathcal{M} \subseteq Mod(\mathcal{T})$. We say that a DL KB $\mathcal{K}$ is a *maximal* $(\mathcal{L}, \mathcal{T})$*-approximation* of $\mathcal{M}$ if (*i*) $\mathcal{K}$ is a sound $(\mathcal{L}, \mathcal{T})$-approximation of $\mathcal{M}$, and (*ii*) there exists no KB $\mathcal{K}'$ that is a sound $(\mathcal{L}, \mathcal{T})$-approximation of $\mathcal{M}$, and is such that $Mod(\mathcal{K}') \subset Mod(\mathcal{K})$.

Such a notion of approximation intuitively aims at the best set of models which can be selected through a KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, with $\mathcal{T}$ fixed, expressed in the DL $\mathcal{L}$. The set of models of $\mathcal{K}$ must contain the set $\mathcal{M}$ to be approximated (e.g., the one obtained through update or erasure), and it should be as close as possible to $\mathcal{M}$.

Interestingly, when such a best approximation exists, it is unique up to logical equivalence, as the following theorem shows.

Theorem 1 *Let* $\mathcal{M}$ *be a set of models, and* $\mathcal{T}$ *a TBox in a DL* $\mathcal{L}$ *such that* $\mathcal{M} \subseteq Mod(\mathcal{T})$. *If a KB* $\mathcal{K}$ *exists that is a maximal* $(\mathcal{L}, \mathcal{T})$*-approximation of* $\mathcal{M}$, *then all maximal* $(\mathcal{L}, \mathcal{T})$*-approximations of* $\mathcal{M}$ *are equivalent to* $\mathcal{K}$.

Proof. Suppose that $\mathcal{K}' = \langle \mathcal{T}, \mathcal{A}' \rangle$ is a maximal $(\mathcal{L}, \mathcal{T})$-approximation of $\mathcal{M}$ that is not equivalent to $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$. By definition of maximal $(\mathcal{L}, \mathcal{T})$-approximation, we have that both $\mathcal{K}$ and $\mathcal{K}'$ are sound $(\mathcal{L}, \mathcal{T})$-approximations of $\mathcal{M}$, which implies that $\mathcal{M} \subseteq Mod(\mathcal{K})$, and $\mathcal{M} \subseteq Mod(\mathcal{K}')$. Hence, $\mathcal{M} \subseteq Mod(\mathcal{K}) \cap Mod(\mathcal{K}')$. Now, let $\mathcal{K}''$ be $\langle \mathcal{T}, \mathcal{A} \cup$

$\mathcal{A}'\rangle$. It is easy to see that $Mod(\mathcal{K}'') = Mod(\mathcal{K}) \cap Mod(\mathcal{K}')$, and therefore $\mathcal{K}''$ is also a sound $(\mathcal{L}, \mathcal{T})$-approximation of $\mathcal{M}$. But then, since $Mod(\mathcal{K}) \neq Mod(\mathcal{K}')$, we obtain that $Mod(\mathcal{K}'') \subset Mod(\mathcal{K})$, which contradicts the fact that $\mathcal{K}$ is a maximal $(\mathcal{L}, \mathcal{T})$-approximation. $\square$

The next theorem characterizes the kind of approximation we have defined in terms of logical implication: a (or, with a slight abuse of terminology, "the") maximal $(\mathcal{L}, \mathcal{T})$-approximation captures exactly all membership assertions in the DL $\mathcal{L}$ that are logically implied by $\mathcal{M}$.

Theorem 2 *Let $\mathcal{M}$ be a set of models, and $\mathcal{T}$ a TBox in a DL $\mathcal{L}$ such that $\mathcal{M} \subseteq Mod(\mathcal{T})$. If $\mathcal{K}$ is a maximal $(\mathcal{L}, \mathcal{T})$-approximation of $\mathcal{M}$, then for every membership assertion α in $\mathcal{L}$ it holds that $\mathcal{M} \models \alpha$ iff $\mathcal{K} \models \alpha$.*

Proof. The if-direction is obvious. As for the only-if direction, let us assume that there is a model $\mathcal{I} \in \mathcal{K}$ such that $\mathcal{I} \models \mathcal{K}$ but $\mathcal{I} \not\models \alpha$. But then $\mathcal{K} \cup \alpha$ would be also a sound $(\mathcal{L}, \mathcal{T})$-approximation of $\mathcal{M}$, which contradicts the fact that $\mathcal{K}$ is a maximal $(\mathcal{L}, \mathcal{T})$-approximation of $\mathcal{M}$. $\square$

Unfortunately, maximal $(\mathcal{L}, \mathcal{T})$-approximations may not exist, as the following example shows.

Example 1 Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be the $\mathcal{ALCQIO}$ KB (Baader *et al.* 2003) defined as follows:

$$\begin{array}{lllll} \mathcal{T}: & \{\exists R_1 \sqcap \neg A \sqsubseteq \exists R_2, & \exists R_2^- \sqsubseteq \exists R_1, & \exists R_1 \sqsubseteq \neg \exists R_1^-, \\ & (\leq 1\, R_2), & \top \sqsubseteq \forall R_1.\{b\}, & A \sqsubseteq \forall R_2.A\} \\ \mathcal{A}: & \{R_1(a,b), & \forall R_1^-.\neg A(b)\} & \end{array}$$

Now, consider the update of $\mathcal{K}$ with $\mathcal{F} = \{A(a)\}$, and let $\mathcal{M} = \mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$. We show that there exists no finite maximal $(\mathcal{ALCQIO}, \mathcal{T})$-approximation of $\mathcal{K}$ with $\mathcal{F}$. Clearly, each

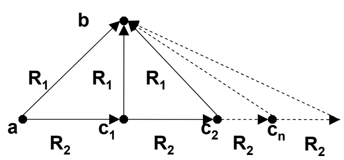

Figure 1: Example of model of $\mathcal{K}$

model $\mathcal{I}$ of $\mathcal{K}$ is such that (*i*) the pair of objects (a, b) is in $R_1^{\mathcal{I}}$, and (*ii*) each object x such that (x, b) is in $R_1^{\mathcal{I}}$, is also in $\neg A^{\mathcal{I}}$. Figure 1 shows a graphical representation of a possible model $\mathcal{I}$ of $\mathcal{K}$, where $a \in \neg A^{\mathcal{I}}$, and, for all i, $c_i \in \neg A^{\mathcal{I}}$. It can be shown that $U^{\mathcal{T}}(\mathcal{I}, \mathcal{F})$ is non-empty, and, for every $\mathcal{J} \in U^{\mathcal{T}}(\mathcal{I}, \mathcal{F})$, we have that $a \in A^{\mathcal{J}}$ and either the pair (a, c_1) does not belong to $R_2^{\mathcal{J}}$, or all c_i's do not belong to $\neg A^{\mathcal{J}}$. Therefore, the update changes the interpretation of $\neg A$ by removing from $\neg A$ only those objects x for which there exists a path from a to x through edges labeled R_2. Now let us define, for $n \geq 0$:

$$\mathcal{A}^n = \{A(a), \forall R_1^-.(\neg A \sqcup \bigsqcup_{0 \leq i \leq n} (\exists R_2^-)^i.\{a\})(b)\}$$

From the above observations it can be shown by induction that for each n: (*i*) $\mathcal{K}^n = \langle \mathcal{T}, \mathcal{A}^n \rangle$ is a sound $(\mathcal{ALCQIO}, \mathcal{T})$-approximation of $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$; and (*ii*) $Mod(\mathcal{K}^n) \subset Mod(\mathcal{K}^{n+1})$. Therefore, no finite set of membership assertions is a maximal $(\mathcal{ALCQIO}, \mathcal{T})$-approximation of $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$.

With the notion of maximal $(\mathcal{L}, \mathcal{T})$-approximation of $\mathcal{M}$ in place, we can come back to the issue of approximating instance-level update and erasure.

Definition 8 (**$(\mathcal{L}, \mathcal{T})$-Update**) Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, $\mathcal{K}^a = \langle \mathcal{T}, \mathcal{A}^a \rangle$ be two KBs in a DL $\mathcal{L}$ and $\mathcal{F}$ a finite set of membership assertions expressed in $\mathcal{L}$ such that $Mod(\mathcal{T}) \cap Mod(\mathcal{F}) \neq \emptyset$. We say that $\mathcal{K}^a$ is a *$(\mathcal{L}, \mathcal{T})$-update of $\mathcal{K}$ with* $\mathcal{F}$ if $\mathcal{K}^a$ is a maximal $(\mathcal{L}, \mathcal{T})$-approximation of $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$.

Definition 9 (**$(\mathcal{L}, \mathcal{T})$-Erasure**) Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, $\mathcal{K}^a = \langle \mathcal{T}, \mathcal{A}^a \rangle$ be two KBs in a DL $\mathcal{L}$ and $\mathcal{F}$ a finite set of membership assertions expressed in $\mathcal{L}$ such that $Mod(\mathcal{T}) \cap Mod(\neg\mathcal{F}) \neq \emptyset$. We say that $\mathcal{K}^a$ is a *$(\mathcal{L}, \mathcal{T})$-erasure of $\mathcal{K}$ with* $\mathcal{F}$ if $\mathcal{K}^a$ is a maximal $(\mathcal{L}, \mathcal{T})$-approximation of $\mathcal{K} \bullet_{\mathcal{T}} \mathcal{F}$.

From Theorem 1 we know that if an $(\mathcal{L}, \mathcal{T})$-update (resp. $(\mathcal{L}, \mathcal{T})$-erasure) of $\mathcal{K}$ with $\mathcal{F}$ exists, it is unique up to logical equivalence. Moreover, by Theorem 2 we know that $(\mathcal{L}, \mathcal{T})$-update (resp. $(\mathcal{L}, \mathcal{T})$-erasure) captures exactly the logical implication of the membership assertions of the "exact" update (resp. erasure). Also the example above shows that in general there are cases for which $(\mathcal{L}, \mathcal{T})$-updates do not exist, and, in fact, the example above can be adapted to show that also $(\mathcal{L}, \mathcal{T})$-erasures do not always exist.

Approximated updates in *DL-Lite*$_{\mathcal{F}}$

In this section we focus our attention to updates to *DL-Lite*$_{\mathcal{F}}$ KBs. This problem has been studied in the context of the *DL-Lite* family in (De Giacomo *et al.* 2006), where it is shown that, for each $\mathcal{T}$ in *DL-Lite*$_{FS}$, the perfect (*DL-Lite*$_{FS}$, $\mathcal{T}$)-update exists, where *DL-Lite*$_{FS}$ differs from *DL-Lite*$_{\mathcal{F}}$ because it allows for the presence of variables and negation in membership assertions.

In (De Giacomo *et al.* 2006) a sound and complete algorithm, called $ComputeUpdate(\mathcal{T}, \mathcal{A}, \mathcal{F})$ was presented, for computing the update of a *DL-Lite*$_{FS}$ KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ with a set of *DL-Lite*$_{FS}$ membership assertions $\mathcal{F}$. Specifically we have the following theorem.

Theorem 3 ((De Giacomo *et al.* 2006)) *Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a DL-Lite$_{FS}$ KB, $\mathcal{F}$ a finite set of DL-Lite$_{FS}$ membership assertions such that $Mod(\mathcal{T} \cup \mathcal{F}) \neq \emptyset$, and let $\mathcal{K}^p = \langle \mathcal{T}, \mathcal{A}^p \rangle$, where $\mathcal{A}^p$ is the ABox returned by $ComputeUpdate(\mathcal{T}, \mathcal{A}, \mathcal{F})$. Then, $Mod(\mathcal{K}^p) = \mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$.*

In other words for *DL-Lite*$_{FS}$ update and maximal approximation of the update coincide.

Observe that *DL-Lite*$_{FS}$ TBoxes are also *DL-Lite*$_{\mathcal{F}}$ TBoxes, while *DL-Lite*$_{FS}$ ABoxes have a more general form than *DL-Lite*$_{\mathcal{F}}$ ones. As a result of this restriction we have that the nice property above does not hold in general for *DL-Lite*$_{\mathcal{F}}$ since to express the updated KB we have to resort to variable and negation in the membership assertions of the new ABox, which is not allowed in *DL-Lite*$_{\mathcal{F}}$. The following example illustrates this point.

Example 2 Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be the *DL-Lite*$_{\mathcal{F}}$ KB such that $\mathcal{T} = \{A_1 \sqsubseteq \neg \exists P, \exists P \sqsubseteq \neg A_2, \exists P^- \sqsubseteq A_3\}$ and $\mathcal{A} = \{\exists P(a)\}$. We want now to compute the (*DL-Lite*$_{\mathcal{F}}$, $\mathcal{T}$)-update of $\mathcal{K}$ with $\mathcal{F} = \{A_1(a)\}$. Clearly, the set of models that satisfy $A_1(a)$ and minimally differs from $Mod(\mathcal{K})$ is obtained by modifying the interpretation of P in all models of $\mathcal{K}$ so that there exists no couple of objects (a, x) that belongs to the interpretation of P. In particular, this means that for each model $\mathcal{I}$ of $\mathcal{K}$, there is no need to change neither the interpretation of A_2, nor the interpretation of A_3. Thus, we have that (*i*) a should belong to $\neg A_2$, i.e. $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F} \models \neg A_2(a)$, and (*ii*) for each couple (a, x) that belongs to $P^{\mathcal{I}}$, x should be still interpreted as belonging to $A_3^{\mathcal{I}}$, and hence, A_3 cannot be interpreted as the empty set, i.e. $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F} \models \exists x A_3(x)$. It can be easily shown that there is no way to express in *DL-Lite*$_{\mathcal{F}}$ neither (*i*), nor (*ii*) through a set of *DL-Lite*$_{\mathcal{F}}$ membership assertions.

From the previous example, it follows that:

Theorem 4 *The result of an update to a DL-Lite*$_{\mathcal{F}}$ *KB may not be expressible in DL-Lite*$_{\mathcal{F}}$ *itself.*

Thus, in general we have to look for approximate representations of the result of an update in *DL-Lite*$_{\mathcal{F}}$. To this aim, we define an algorithm named $ComputeUpdate^{app}$ that takes as input a TBox $\mathcal{T}$, an ABox $\mathcal{A}$ and a set of membership assertions $\mathcal{F}$, where both $\mathcal{T}$, $\mathcal{A}$ and $\mathcal{F}$ are expressed in *DL-Lite*$_{\mathcal{F}}$. Moreover, we assume that $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ is consistent and $Mod(\mathcal{T} \cup \mathcal{F}) \neq \emptyset$. Then $ComputeUpdate^{app}$ returns the ABox $\mathcal{A}^a$ obtained as follows:

1. compute the ABox $\mathcal{A}^p = ComputeUpdate(\mathcal{T}, \mathcal{A}, \mathcal{F})$;
2. return the projection of A^p to *DL-Lite*$_{\mathcal{F}}$, i.e. the ABox obtained by deleting from $\mathcal{A}^p$ all the assertions that are not *DL-Lite*$_{\mathcal{F}}$ membership assertions.

Since the algorithm $ComputeUpdate(\mathcal{T}, \mathcal{A}, \mathcal{F})$ runs in polynomial time (De Giacomo *et al.* 2006), it follows that the algorithm $ComputeUpdate^{app}$ also terminates and runs in time polynomial with respect to the size of its input.

Example 3 Consider the *DL-Lite*$_{\mathcal{F}}$ KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ introduced of Example 2. Now suppose to compute the (*DL-Lite*$_{\mathcal{F}}, \mathcal{T}$)-update of $\mathcal{K}$ with $\mathcal{F} = \{A_1(a)\}$. First, we apply the update algorithm $ComputeUpdate$ of (De Giacomo *et al.* 2006). This returns a *DL-Lite*$_{FS}$ ABox $\mathcal{A}^p$ that is obtained from $\mathcal{A}$ by removing the assertion $\exists P(a)$, and introducing, besides $A_1(a)$, the assertions $\neg A_2(a)$ and $A_3(z)$, where z is a new variable. Second, we perform the projection of $\mathcal{A}^p$ in *DL-Lite*$_{\mathcal{F}}$, and obtain the *DL-Lite*$_{\mathcal{F}}$ ABox $\mathcal{A}^a = \{A_1(a)\}$.

To prove the correctness of the above algorithm we need the following preliminary lemma.

Lemma 1 *Let* $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ *be a DL-Lite*$_{FS}$ *KB, and* α *be a DL-Lite*$_{\mathcal{F}}$ *assertion. If* $\mathcal{K} \models \alpha$*, then there exists a DL-Lite*$_{\mathcal{F}}$ *membership assertion* α' *in* $\mathcal{A}$ *such that* $\langle \mathcal{T}, \alpha' \rangle \models \alpha$*.*

Proof. The proof is based on the nice property of *DL-Lite*$_{FS}$ KBs of having a minimal model that is built by using a *chase-like technique* (Poggi 2006). □

Theorem 5 *Let* $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ *be a DL-Lite*$_{\mathcal{F}}$ *KB,* $\mathcal{F}$ *a finite set of DL-Lite*$_{\mathcal{F}}$ *membership assertions such that* $Mod(\mathcal{T} \cup \mathcal{F}) \neq \emptyset$*, and let* $\mathcal{K}^a = \langle \mathcal{T}, \mathcal{A}^a \rangle$*, where* $\mathcal{A}^a$ *is the ABox returned by* $ComputeUpdate^{app}(\mathcal{T}, \mathcal{A}, \mathcal{F})$*. Then,* $\mathcal{K}^a$ *is a* (*DL-Lite*$_{\mathcal{F}}, \mathcal{T}$)*-update of* $\mathcal{K}$ *with* $\mathcal{F}$*.*

Proof. Clearly the algorithm $ComputeUpdate^{app}$ terminates, since so does $ComputeUpdate$. Now, let $\mathcal{A}^a = ComputeUpdate^{app}(\mathcal{T}, \mathcal{A}, \mathcal{F})$. We first show that $\langle \mathcal{T}, \mathcal{A}^a \rangle$ is a sound (*DL-Lite*$_{\mathcal{F}}, \mathcal{T}$)-approximation of $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$, then we show that it is a maximal one. Let $\mathcal{A}^p = ComputeUpdate(\mathcal{T}, \mathcal{A}, \mathcal{F})$. By construction, $\mathcal{A}^a \subseteq \mathcal{A}^p$ and therefore, since *DL-Lite*$_{\mathcal{F}}$ is monotone, $Mod(\mathcal{A}^p) \subseteq Mod(\mathcal{A}^a)$. Hence $Mod(\langle \mathcal{T}, \mathcal{A}^p \rangle) \subseteq Mod(\langle \mathcal{T}, \mathcal{A}^a \rangle)$. Moreover, by Theorem 3, $Mod(\langle \mathcal{T}, \mathcal{A}^p \rangle) = \mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$. It follows that $\langle \mathcal{T}, \mathcal{A}^a \rangle$ is a sound (*DL-Lite*$_{\mathcal{F}}, \mathcal{T}$)-approximation of $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$. Now, let us show that $\mathcal{K}^a = \langle \mathcal{T}, \mathcal{A}^a \rangle$ is the maximal (*DL-Lite*$_{\mathcal{F}}, \mathcal{T}$)-approximation of $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$. By contradiction, let $\mathcal{K}' = \langle \mathcal{T}, \mathcal{A}' \rangle$ be a sound (*DL-Lite*$_{\mathcal{F}}, \mathcal{T}$)-approximation of $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$ such that $Mod(\mathcal{K}') \subset Mod(\mathcal{K}^a)$. Since $Mod(\mathcal{K}' \cup \mathcal{K}^a) = Mod(\mathcal{K}') \cap Mod(\mathcal{K}^a)$, we have that $Mod(\mathcal{K}' \cup \mathcal{K}^a) = Mod(\mathcal{K}')$, which implies that that $\mathcal{K}^a \subset \mathcal{K}'$, and thus that there exists a *DL-Lite*$_{\mathcal{F}}$ membership assertion α such that $\alpha \in \mathcal{A}'$, $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F} \models \alpha$ and $\mathcal{K}^a \not\models \alpha$. Let $\mathcal{A}^p = ComputeUpdate(\mathcal{T}, \mathcal{A}, \mathcal{F})$ and $\mathcal{K}^p = \langle \mathcal{T}, \mathcal{A}^p \rangle$. By Theorem 3, $Mod(\mathcal{K}^p) = \mathcal{K} \circ_{\mathcal{T}} \mathcal{F}$. Then we have that $\mathcal{K}^p \models \alpha$, where α is a membership assertion in *DL-Lite*$_{\mathcal{F}}$. By Lemma 1, there must exist a *DL-Lite*$_{\mathcal{F}}$ membership assertion α' in $\mathcal{A}^p$ such that $\langle \mathcal{T}, \alpha' \rangle \models \alpha$. But then, by construction, α' belongs to $\mathcal{A}^a$, contradicting $\mathcal{K}^a \not\models \alpha$. □

From Theorem 5 and Theorem 2, it follows that $ComputeUpdate^{app}$ captures, in a sound and complete way, logical implication of *DL-Lite*$_{\mathcal{F}}$ assertions after update.

Theorem 6 *Let* $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ *be a DL-Lite*$_{\mathcal{F}}$ *KB, let* $\mathcal{F}$ *be a finite set of DL-Lite*$_{\mathcal{F}}$ *membership assertions such that* $Mod(\mathcal{T} \cup \mathcal{F}) \neq \emptyset$*, and let* $\mathcal{K}^a = \langle \mathcal{T}, \mathcal{A}^a \rangle$*, where* $\mathcal{A}^a$ *is the ABox returned by* $ComputeUpdate^{app}(\mathcal{T}, \mathcal{A}, \mathcal{F})$*. Then, for every membership assertion* α *in DL-Lite*$_{\mathcal{F}}$*, we have that* $\mathcal{K} \circ_{\mathcal{T}} \mathcal{F} \models \alpha$ *iff* $\mathcal{K}^a \models \alpha$*.*

Finally, we point out that the above results do not contradict what is reported on (De Giacomo *et al.* 2006), where it is shown that the result of an update to a KB expressed in *DL-Lite*$_{FS}$ is always expressible in this logic. Indeed, *DL-Lite*$_{FS}$ is strictly more expressive than *DL-Lite*$_{\mathcal{F}}$, and our result shows indeed that this extra expressive power is crucial for expressing updates.

Approximated erasure in *DL-Lite*$_{\mathcal{F}}$

We now study erasure in *DL-Lite*$_{\mathcal{F}}$ KB. We start by showing that, in *DL-Lite*$_{\mathcal{F}}$, the result of the erasure cannot be always expressed in terms of a *DL-Lite*$_{\mathcal{F}}$ KB. To this aim, we present the following example.

Example 4 Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be the *DL-Lite*$_{\mathcal{F}}$ KB such that $\mathcal{T} = \{A \sqsubseteq B, A \sqsubseteq C\}$ and $\mathcal{A} = \{A(a)\}$. Let $\mathcal{F} = \{C(a)\}$. By definition, $\mathcal{K} \bullet_{\mathcal{T}} \mathcal{F} = Mod(\mathcal{K}) \cup (\mathcal{K} \circ_{\mathcal{T}} \{\neg C(a)\})$. Thus, each model $\mathcal{I}$ in $\mathcal{K} \bullet_{\mathcal{T}} \mathcal{F}$ is obtained from a model $\mathcal{I}$ of $\mathcal{K}$ by either not modifying anything, or by modifying the interpretation of a so that a does not belong to $A^{\mathcal{I}}$. Hence, for each $\mathcal{I}$, either $a \in A^{\mathcal{I}}, a \in C^{\mathcal{I}}$ or $a \notin A^{\mathcal{I}}, a \notin C^{\mathcal{I}}$. Clearly there is no way to express this set of models through a set of *DL-Lite*$_{\mathcal{F}}$ membership assertions.

As an immediate consequence of the above example, we get the following property.

Theorem 7 *The result of an erasure to a DL-Lite*$_{\mathcal{F}}$ *KB may not be expressible in DL-Lite*$_{\mathcal{F}}$ *itself.*

Therefore, like in the case of update, in *DL-Lite*$_{\mathcal{F}}$ it is interesting to look at maximal approximations of the erasure. For this purpose, we define an algorithm named $ComputeErasure^{app}$, which takes as input a TBox $\mathcal{T}$, an ABox $\mathcal{A}$ and a set of membership assertions $\mathcal{F}$, where both

$\mathcal{T}$, $\mathcal{A}$ and $\mathcal{F}$ are expressed in *DL-Lite*$_\mathcal{F}$. Moreover, we assume that $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ is consistent and that $Mod(\mathcal{T} \cup \neg\mathcal{F}) \neq \emptyset$. Then $ComputeErasure^{app}$ returns the ABox $\mathcal{A}^a$ obtained as follows:

1. compute the ABox $\mathcal{A}^p = ComputeUpdate(\mathcal{T}, \mathcal{A}, \neg\mathcal{F})$;
2. compute the projection of A^p to *DL-Lite*$_\mathcal{F}$, i.e. delete from $\mathcal{A}^p$ all the assertions that are not *DL-Lite*$_\mathcal{F}$ membership assertions.

As mentioned in the previous section, the algorithm $ComputeUpdate(\mathcal{T}, \mathcal{A}, \mathcal{F})$ runs in polynomial time, therefore the algorithm $ComputeErasure^{app}$ terminates and runs in time polynomial with respect to the size of its input.

Example 5 Consider the *DL-Lite*$_\mathcal{F}$ KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ introduced in Example 4. Now suppose to compute the $($*DL-Lite*$_\mathcal{F}, \mathcal{T})$-erasure of $\mathcal{K}$ with $\mathcal{F} = \{C(a)\}$. First, we apply the update algorithm $ComputeUpdate$ of (De Giacomo *et al.* 2006) and compute the ABox $\mathcal{A}^p = ComputeUpdate(\mathcal{T}, \mathcal{A}, \{\neg C(a)\})$. This returns a *DL-Lite*$_{FS}$ ABox that is obtained from $\mathcal{A}$ by removing the assertion $A(a)$, and introducing the assertions $\neg C(a)$ and $B(a)$. Second, we perform the projection of $\mathcal{A}^p$ in *DL-Lite*$_\mathcal{F}$, and obtain the *DL-Lite*$_\mathcal{F}$ ABox $\mathcal{A}^a = \{B(a)\}$.

The following theorem shows that the maximal approximation of instance-level erasure in a *DL-Lite*$_\mathcal{F}$ KB always exists and is computed by the algorithm $ComputeErasure^{app}$.

Theorem 8 *Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a DL-Lite$_\mathcal{F}$ KB, $\mathcal{F}$ a finite set of DL-Lite$_\mathcal{F}$ membership assertions such that $Mod(\mathcal{T} \cup \neg\mathcal{F}) \neq \emptyset$, and let $\mathcal{K}^a = \langle \mathcal{T}, \mathcal{A}^a \rangle$, where $\mathcal{A}^a$ is the ABox returned by $ComputeErasure^{app}(\mathcal{T}, \mathcal{A}, \mathcal{F})$. Then, $\mathcal{K}^a$ is a (DL-Lite$_\mathcal{F}, \mathcal{T}$)-erasure of $\mathcal{K}$ with $\mathcal{F}$.*

Proof. First, from definition of erasure, $\mathcal{K} \bullet_\mathcal{T} \mathcal{F} = Mod(\mathcal{K}) \cup \mathcal{K} \circ_\mathcal{T} \neg\mathcal{F}$. Then, by definition of the algorithm $ComputeErasure^{app}$, it follows that for every membership assertion $\alpha \in \mathcal{A}^a - \mathcal{A}$, we have $\mathcal{K} \models \alpha$, which immediately implies that (i) $Mod(\mathcal{K}) \subseteq Mod(\mathcal{K}^a)$. Moreover, let $\mathcal{K}^p = \langle \mathcal{T}, \mathcal{A}^p \rangle$ where $\mathcal{A}^p$ is the ABox returned by $ComputeUpdate^{app}(\mathcal{T}, \mathcal{A}, \neg\mathcal{F})$: by definition, $\mathcal{A}^a \subseteq \mathcal{A}^p$, consequently every model of $\mathcal{K}^p$ is also a model of $\mathcal{K}^a$, and since by Theorem 3 $Mod(\mathcal{K}^p) = \mathcal{K} \circ_\mathcal{T} \neg\mathcal{F}$, it follows that (ii) $\mathcal{K} \circ_\mathcal{T} \neg\mathcal{F} \subseteq Mod(\mathcal{K}^a)$. Hence, from (i) and (ii) it follows that $\mathcal{K}^a$ is a $($*DL-Lite*$_\mathcal{F}, \mathcal{T})$-approximation of $\mathcal{K} \bullet_\mathcal{T} \mathcal{F}$.

Now, suppose $\mathcal{K}^a$ is not the maximal $($*DL-Lite*$_\mathcal{F}, \mathcal{T})$-approximation of $\mathcal{K} \bullet_\mathcal{T} \mathcal{F}$. Then, there exists a *DL-Lite*$_\mathcal{F}$ KB $\mathcal{K}' = \langle \mathcal{T}, \mathcal{A}' \rangle$ such that $\mathcal{K} \bullet_\mathcal{T} \mathcal{F} \subseteq Mod(\mathcal{K}') \subset Mod(\mathcal{K}^a)$. Since $Mod(\mathcal{K}') \subset Mod(\mathcal{K}^a)$, there exists at least a (membership) assertion $\alpha \in \mathcal{A}' - \mathcal{A}^a$ such that $\mathcal{K}^a \not\models \alpha$, and since $\mathcal{K} \bullet_\mathcal{T} \mathcal{F} \subseteq Mod(\mathcal{K}')$ and $\mathcal{K} \bullet_\mathcal{T} \mathcal{F} = Mod(\mathcal{K}) \cup \mathcal{K} \circ_\mathcal{T} \neg\mathcal{F}$, it follows that $\mathcal{K}^p \models \alpha$. Now, by Lemma 1 it follows that there exists a membership assertion $\alpha' \in \mathcal{A}^p$ such that $\langle \mathcal{T}, \{\alpha'\} \rangle \models \alpha$. Hence, by definition of $ComputeErasure^{app}$, it follows that $\alpha' \in \mathcal{A}^a$, consequently $\mathcal{K}^a \models \alpha$. Contradiction. Therefore, $\mathcal{K}^a$ is the maximal $($*DL-Lite*$_\mathcal{F}, \mathcal{T})$-approximation of $\mathcal{K} \bullet_\mathcal{T} \mathcal{F}$. □

Finally, from Theorem 8 and Theorem 2, it immediately follows that the algorithm $ComputeErasure^{app}$ captures, in a sound and complete way, logical implication of *DL-Lite*$_\mathcal{F}$ assertions after erasure.

Theorem 9 *Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a DL-Lite$_\mathcal{F}$ KB, let $\mathcal{F}$ be a finite set of DL-Lite$_\mathcal{F}$ membership assertions such that $Mod(\mathcal{T} \cup \neg\mathcal{F}) \neq \emptyset$, and let $\mathcal{K}^a = \langle \mathcal{T}, \mathcal{A}^a \rangle$, where $\mathcal{A}^a$ is the ABox returned by $ComputeErasure^{app}(\mathcal{T}, \mathcal{A}, \mathcal{F})$. Then, for every membership assertion α in DL-Lite$_\mathcal{F}$, we have that $\mathcal{K} \bullet_\mathcal{T} \mathcal{F} \models \alpha$ iff $\mathcal{K}^a \models \alpha$.*

Conclusion

We have investigated the notion of evolution of a DL KB, under three basic assumptions: formal semantics, fixed TBox, and fixed DL. In order to cope with the non-expressibility problem, we have introduced the notion of maximal approximation of an update and of an erasure. Finally we have presented efficient algorithms for computing such approximations in the tractable DL *DL-Lite*$_\mathcal{F}$. In the future, we will incorporate the algorithms presented here in the QuOnto reasoning system, and we will try to develop algorithms for computing approximations in the context of other DLs used in the current reasoning systems.

Acknowledgments This research has been partially supported by the FET project TONES (Thinking ONtologiES), funded by the EU under contract number FP6-7603, and by the MIUR FIRB 2005 project "Tecnologie Orientate alla Conoscenza per Aggregazioni di Imprese in Internet" (TOCAI.IT).

References

Baader, F.; Calvanese, D.; McGuinness, D.; Nardi, D.; and Patel-Schneider, P. F., eds. 2003. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press.

Calvanese, D.; De Giacomo, G.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2005. DL-Lite: Tractable description logics for ontologies. In *Proc. of AAAI 2005*, 602–607.

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

De Giacomo, G.; Lenzerini, M.; Poggi, A.; and Rosati, R. 2006. On the update of description logic ontologies at the instance level. In *Proc. of AAAI 2006*.

Gutiérrez, C.; Hurtado, C.; and Vaisman, A. 2006. The meaning of erasing in RDF under the Katsuno-Mendelzon approach.

Haase, P., and Stojanovic, L. 2005. Consistent evolution of owl ontologies. In *Proc. of ESWC 2005*, 182–197.

Herzig, A., and Rifi, O. 1999. Propositional belief update and minimal change. *Artificial Intelligence* 115(1):107–138.

Katsuno, H., and Mendelzon, A. 1991. On the difference between updating a knowledge base and revising it. In *Proc. of KR'91*, 387–394.

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

Poggi, A. 2006. *Structured and Semi-Structured Data Integration*. Ph.D. Dissertation, Dip. di Inf. e Sist., Univ. di Roma "La Sapienza".

Schaerf, M., and Cadoli, M. 1995. Tractable reasoning via approximation. *Artificial Intelligence* 74(2):249–310.

Winslett, M. 1988. Reasoning about action using a possible models approach. In *Proc. of AAAI'98*.

Winslett, M. 1990. *Updating Logical Databases*. Cambridge University Press.