Trust-Sensitive Evolution of DL-Lite Knowledge Bases

Anonymised

Abstract

Evolution of Knowledge Bases (KBs) consists of incorporating new information in an existing KB. Previous studies assume that the new information should be fully trusted and thus completely incorporated in the old knowledge. We suggest a setting where the new knowledge can be partially trusted and develop model-based approaches (MBAs) to KB evolution that rely on this assumption. Under MBAs the result of evolution is a set of interpretations and thus two core problems for MBAs are closure, i.e., whether evolution result can be axiomatised with a KB, and approximation, i.e., whether it can be (maximally) approximated with a KB. We show that DL-Lite is not closed under a wide range of trust-sensitive MBAs. We introduce a notion of s-approximation that improves the previously proposed approximations and show how to compute it for various trust-sensitive MBAs.

1 Introduction

Recent years have witnessed a strong and increasing interest in Description Logic (DL) knowledge bases (KBs) [Baader *et al.*, 2003] as a mechanism for representing structured knowledge; in particular, DLs became the foundation for OWL 2, the standard ontology language of the Semantic Web. A DL KB $\mathcal{K}$ consists of a TBox $\mathcal{T}$ that models at the intensional level the static and structural aspects of an application domain, and an ABox $\mathcal{A}$ that models at the extensional level the current state of affairs or data about individuals.

In many applications KBs are subject to changes, for instance, when they are constructed from evolving Web pages [Suchanek and Weikum, 2013] or databases [Furche *et al.*, 2012], or created collaboratively [Bollacker *et al.*, 2008; Stearns *et al.*, 2001]. A typical scenario for such applications is to incorporate in a given KB $\mathcal{K}$ an acquired KB $\mathcal{N}$ that expresses new information. In the case where $\mathcal{N}$ interacts with $\mathcal{K}$ in an undesirable way, e.g., by causing the KB or relevant parts of it to become unsatisfiable, $\mathcal{N}$ cannot simply be added to $\mathcal{K}$. Different ways to address this problem are possible, corresponding to different approaches for *KB evolution*.

Knowledge evolution in the context of DL KBs has recently attracted a lot of attention from both practical and foundational perspectives, see e.g. [Fridman Noy *et al.*, 2004; Haase and Stojanovic, 2005; Flouris *et al.*, 2008; Konev *et*

al., 2008; Cuenca Grau *et al.*, 2012; Wu and Lécué, 2014]. Among the foundational, the most common are *model-based approaches* (MBAs) where the result of evolution is the *set of first-order interpretations* $\mathcal{M}$ that are models of $\mathcal{N}$ and that are minimally distant from the models of $\mathcal{K}$. The latter condition corresponds to a widely accepted *principle of minimal change* [Eiter and Gottlob, 1992]. Depending on how minimality and distance are defined, one can obtain various evolution semantics and a number of them have been introduced and studied [Qi and Du, 2009; Calvanese *et al.*, 2010; Wang *et al.*, 2010; Kharlamov and Zheleznyakov, 2011; Liu *et al.*, 2011; Kharlamov *et al.*, 2013; Zhuang *et al.*, 2014; Qi *et al.*, 2015; Wang *et al.*, 2015; Zhuang *et al.*, 2016].

To the best of our knowledge, all previous studies of KB evolution assume that the new knowledge $\mathcal{N}$ should be *fully* trusted and thus completely taken on board (see related work in Section 6). However, this assumption does not hold in a wide range of important applications [Suchanek and Weikum, 2013] where $\mathcal{N}$ comes from a *partially trusted* source, e.g., from the Web or from a source with a limited expertise.

In this work we address this issue for MBAs and study how an external notion of trust could be used in order to determine how new knowledge should be integrated with existing knowledge. Following Hunter and Booth (2015), who studied trust in the context of propositional belief revision, we assume that the knowledge provider has expertise that is restricted to a particular area and thus cannot distinguish between certain states of the application domain—first-order interpretations in our case. We formalise such a notion of trust as an equivalence relation on first-order interpretations and introduce four natural classes of trust. Then we use trust as an external mechanism to relativise arbitrary interpretations (not necessarily models of $\mathcal{N}$) to models of $\mathcal{N}$ by considering equivalence classes of the latter's models. This allows us to define the result of evolution as a set of minimally distinct interpretations $\mathcal{M}$ selected from these equivalence classes instead of just models of $\mathcal{N}$ as in classical MBAs. Our trust-sensitive evolution is generic in the sense that it is applicable to KBs of *any* DL, and is backwards-compatible with classical MBAs in the sense that it coincides with them whenever the model of trust assumes that the knowledge provider is an expert in everything.

Since evolution under MBAs is defined as a set of interpretations $\mathcal{M}$, in practice one would want to efficiently axiomatise this $\mathcal{M}$ as a KB or, whenever this is impossible, to efficiently 'closely' approximate it with a KB. Thus, the two

core KB evolution problems for a given DL $\mathcal{L}$ and MBA semantics are *closure* of $\mathcal{L}$ under the semantics, i.e., whether and how for every $\mathcal{K}$ and $\mathcal{N}$ in $\mathcal{L}$, the corresponding $\mathcal{M}$ under the semantics can be axiomatised in $\mathcal{L}$, and *approximation* of the semantics in $\mathcal{L}$ whenever $\mathcal{L}$ is not closed.

We study the closure and approximation problems for DL-Lite—a tractable DL behind the QL profile of OWL 2—under trust-sensitive MBAs for various models of trust. Firstly, we show that DL-Lite is not closed under *any* trust-sensitive semantics. It was known that DL-Lite is not closed under many classical MBAs [Calvanese *et al.*, 2010; Kharlamov *et al.*, 2013; Qi *et al.*, 2015] and our results in particular imply the non-closure of DL-Lite under those classical MBAs for which this problem remained open.

We next turn our attention to the approximation problem for DL-Lite and an important practical setting of ABox evolution where the TBox is static and only the ABox evolves. A widely studied approach for this setting is sound approximation, where $\mathcal{M}$ is approximated with a KB whose set of models contains $\mathcal{M}$. For classical MBAs De Giacomo *et al.* (2009), Kharlamov *et al.* (2013), and Qi *et al.* (2015) proposed algorithms to compute maximal sound approximations for various semantics. Here we propose the notion of *s-approximation*—a KB that may use special predicates and constants—and show that in general it improves sound approximations by better capturing $\mathcal{M}$ for both classical and trust-sensitive MBAs. Moreover, we show that s-approximations are also better in preserving Boolean queries satisfied by $\mathcal{M}$ and we determine an important class of such queries. Finally, we develop polynomial time algorithms to compute maximal sound s-approximations for several trust-sensitive and classical evolution semantics.

2 Preliminaries

Description Logics. We assume standard definitions of first-order logic signature, sentences, interpretations, satisfiability, and entailment. We further assume a fixed signature with disjoint countable sets of unary and binary predicates and constants, and that all interpretations are over this signature and the same countable domain Δ. Let PW (*Possible Worlds*) denote the class of all such interpretations. Whenever convenient we treat interpretations as sets of atoms.

In DLs [Baader *et al.*, 2003], the doman of interest is modelled by means of *concepts*, that are formulae with one free variable, denoting sets of objects, *roles*, that are formulae with two free variables, denoting binary relations between objects, and *constants*, denoting objects. In order to support such modelling, DLs provide a specialised variable-free syntax and operators for constructing concepts and roles from unary predicates (called *atomic concepts*) and binary predicates (called *atomic roles*). A DL KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ consists of a *TBox* $\mathcal{T}$ that is a finite set of sentences (called *TBox assertions*) over concepts and roles, and an *ABox* $\mathcal{A}$ that is a finite set of sentences (called *ABox* or *membership assertions*) of the form $C(a)$ and $R(a, b)$, where C is a concept, R is a role and a, b are constants. A DL $\mathcal{L}$ is a recursive set of KBs closed under renaming of constants and the subset relation.

All the logics of the DL-Lite family have the following constructs for complex concepts and roles [Calvanese *et al.*, 2007]: *(i)* $B ::= A \mid \exists R$, *(ii)* $C ::= B \mid \neg B$, and *(iii)* $R ::= P \mid P^-$, where A and P are an atomic concept and

role, B and C are *basic* and *general* concepts, and R is a *basic* role. A DL-Lite$_{core}$ TBox consists of *concept inclusions* assertions $B \sqsubseteq C$. DL-Lite extends DL-Lite$_{core}$ by allowing in a TBox *role inclusion* assertions $R_1 \sqsubseteq R_2$ and *functionality* assertions $(\texttt{funct } R)$ in a way that if $R_1 \sqsubseteq R_2$ appears in a TBox, then neither $(\texttt{funct } R_2)$ nor $(\texttt{funct } R_2^-)$ appears in the TBox.[1] This syntactic restriction ensures the tractability of the logic. ABoxes in DL-Lite$_{core}$ and DL-Lite consist of membership assertions of the form $C(a)$ and $P(a, b)$.

The semantics for concepts and roles is defined in the standard way under an assumption that $a^\mathcal{I} = a$ for each constant a. That is, $A^\mathcal{I} \subseteq \Delta$, $P^\mathcal{I} \subseteq \Delta \times \Delta$, $(P^-)^\mathcal{I} = \{(b, a) \mid (a, b) \in P^\mathcal{I}\}$, $(\neg B)^\mathcal{I} = \Delta \setminus B^\mathcal{I}$, and $(\exists R)^\mathcal{I} = \{a \mid \text{there exists } b \text{ s.t } (a, b) \in R^\mathcal{I}\}$. The semantics of assertions is also defined in the standard way: $\mathcal{I} \models D_1 \sqsubseteq D_2$ if $D_1^\mathcal{I} \subseteq D_2^\mathcal{I}$, $\mathcal{I} \models (\texttt{funct } R)$ if the relation $R^\mathcal{I}$ is a function, $\mathcal{I} \models C(a)$ if $a^\mathcal{I} \in C^\mathcal{I}$, and $\mathcal{I} \models P(a, b)$ if $(a^\mathcal{I}, b^\mathcal{I}) \in P^\mathcal{I}$.

For ξ a concept, role, (set of) assertions(s), or (set of) interpretation(s), $\texttt{pred}(\xi)$ denotes the set of atomic concepts and roles in ξ. $\text{Mod}(\mathcal{K})$ denotes the set of all interpretations $\mathcal{I}$ that are models of $\mathcal{K}$, i.e., $\mathcal{I} \models \varphi$ for each assertion φ in $\mathcal{K}$.

Classical Model-Based Evolution. A classical evolution setting consists of an old KB $\mathcal{K}$ and a new KB $\mathcal{N}$. Under classical MBAs, the evolution result is the subset of $\text{Mod}(\mathcal{N})$ that is minimally distant from $\text{Mod}(\mathcal{K})$. We now formally introduce classical MBAs [Calvanese *et al.*, 2010].

Let $\mathcal{I}$ and $\mathcal{J}$ be interpretations. Recall that $\mathcal{I} \ominus \mathcal{J}$ denotes the symmetric difference $(\mathcal{I} \setminus \mathcal{J}) \cup (\mathcal{J} \setminus \mathcal{I})$. A *distance function* dist between $\mathcal{I}$ and $\mathcal{J}$ can be defined in one of the following ways. Distances based on atoms are defined as *(i)* a set $\mathcal{I} \ominus \mathcal{J}$, denoted $\text{dist}^\texttt{a}_{\{\}}(\mathcal{I}, \mathcal{J})$, or *(ii)* a number $|\mathcal{I} \ominus \mathcal{J}|$, denoted $\text{dist}^\texttt{a}_{\#}(\mathcal{I}, \mathcal{J})$. Distances based on predicates are defined as *(i)* a set $\{\alpha \mid \alpha \text{ is a predicate and } \alpha^\mathcal{I} \neq \alpha^\mathcal{J}\}$, denoted $\text{dist}^\texttt{p}_{\{\}}(\mathcal{I}, \mathcal{J})$, or *(ii)* a number $|\text{dist}^\texttt{p}_{\{\}}(\mathcal{I}, \mathcal{J})|$ denoted $\text{dist}^\texttt{p}_{\#}(\mathcal{I}, \mathcal{J})$. Distances returned by $\text{dist}^x_{\{\}}$, where $x \in \{\texttt{a}, \texttt{p}\}$ are sets and thus can be compared via set inclusion: $\text{dist}^x_{\{\}}(\mathcal{I}_1, \mathcal{J}_1) \leq \text{dist}^x_{\{\}}(\mathcal{I}_2, \mathcal{J}_2)$ if $\text{dist}^x_{\{\}}(\mathcal{I}_1, \mathcal{J}_1) \subseteq \text{dist}^x_{\{\}}(\mathcal{I}_2, \mathcal{J}_2)$. Distances returned by $\text{dist}^x_{\#}$ are natural numbers and thus can be compared numerically.[2]

Let $\mathcal{S}$ and $\mathcal{S}'$ be sets of interpretations and dist a distance function. The subset $\min_{\text{dist}, \mathcal{S}}(\mathcal{S}')$ of $\mathcal{S}'$ that consists of interpretations minimally distant from $\mathcal{S}$ is defined as follows:

$$\{\mathcal{J} \in \mathcal{S}' \mid \text{there is } \mathcal{I} \in \mathcal{S} \text{ s.t. for each } \mathcal{I}' \in \mathcal{S}, \text{ and}$$
$$\text{for each } \mathcal{J}' \in \mathcal{S}' \text{ it holds } \text{dist}(\mathcal{I}', \mathcal{J}') \not< \text{dist}(\mathcal{I}, \mathcal{J})\}.$$

We now define *selectors* that choose those interpretations of $\mathcal{S}'$ that are minimally distant from $\mathcal{S}$.

Definition 1. *A* selector*, denoted* $\mathbb{S}$*, is a function that maps each pair* $(\mathcal{S}, \mathcal{S}')$ *of sets of models into* $2^{\mathcal{S}'}$*. We consider the following selectors, where* $x \in \{\texttt{a}, \texttt{p}\}$*, and* $y \in \{_{\{\}}, _{\#}\}$*:*
- *a* global *selector induced by* dist^x_y*, denoted* $\mathbb{G}^x_y$*, is defined as* $\min_{\text{dist}^x_y, \mathcal{S}}(\mathcal{S}')$*;*

---

[1]In [Calvanese *et al.*, 2010] this DL is referred to as DL-Lite$_{\mathcal{FR}}$.
[2]Note that for infinite $\text{dist}^x_{\{\}}(\mathcal{I}, \mathcal{J})$, we assume that: for any $\mathcal{I}'$ and $\mathcal{J}'$ if $\text{dist}^x_{\{\}}(\mathcal{I}', \mathcal{J}')$ is *(i)* finite, then $\text{dist}^x_{\{\}}(\mathcal{I}', \mathcal{J}') < \text{dist}^x_{\{\}}(\mathcal{I}, \mathcal{J})$, or *(ii)* infinite, then $\text{dist}^x_{\{\}}(\mathcal{I}', \mathcal{J}') = \text{dist}^x_{\{\}}(\mathcal{I}, \mathcal{J})$.

- *a* local selector induced by dist_y^x, denoted $\mathbb{L}_y^x$, is defined as $\cup_{\mathcal{I}\in\mathcal{S}}\min_{\text{dist}_y^x,\{\mathcal{I}\}}(\mathcal{S}')$.

Finally, *classical evolution semantics* for $\mathcal{K}$ and $\mathcal{N}$ is defined as $\mathbb{S}(\text{Mod}(\mathcal{K}),\text{Mod}(\mathcal{N}))$. Note that, in terms of Katsuno and Mendelzon (1991), semantics based on local selectors correspond to *knowledge update*, and semantics based on global selectors correspond to *knowledge revision*.

3 Trust-Sensitive Model-Based Evolution

In this section we introduce four models of trust and define how they can be incorporated in MBAs.

Models of Trust. Our models of trust reflect the assumption that the knowledge provider has a restricted area of expertise, and thus we do not trust facts that are outside this area. In terms of interpretations this means that if two interpretations disagree only on such facts, then the provider cannot distinguish between them.

Definition 2. *A* model of trust *is an equivalence relation $\equiv$ on models. We consider four classes of models of trust:*
- total trust, *denoted* TT, *consists of one equivalence relation $\equiv_{\text{TT}}$ defined as $\mathcal{I}_1 \equiv_{\text{TT}} \mathcal{I}_2$ iff $\mathcal{I}_1 = \mathcal{I}_2$;*
- total distrust, *denoted* TD, *consists of one equivalence relation $\equiv_{\text{TD}}$ defined as $\mathcal{I}_1 \equiv_{\text{TD}} \mathcal{I}_2$ for each $\mathcal{I}_1$ and $\mathcal{I}_2$;*
- assertion trust, *denoted* AT, *consists of one equivalence relation $\equiv_\Phi$ for each finite set of assertions Φ which is defined as $\mathcal{I}_1 \equiv_\Phi \mathcal{I}_2$ iff either $\mathcal{I}_1 \models \varphi$ and $\mathcal{I}_2 \models \varphi$ or $\mathcal{I}_1 \not\models \varphi$ and $\mathcal{I}_2 \not\models \varphi$ for each $\varphi \in \Phi$.*
- predicate trust, *denoted* PT, *consists of one equivalence relation $\equiv_\mathcal{P}$ for each finite set of predicates $\mathcal{P}$ which is defined as $\mathcal{I}_1 \equiv_\mathcal{P} \mathcal{I}_2$ iff $p^{\mathcal{I}_1} = p^{\mathcal{I}_2}$ for each $p \in \mathcal{P}$.*

Example 3. *Consider a scenario about places where famous researchers Einstein (*Ein*) and Mendeleev (*Men*) live (*livesIn*). Consider the two following models of trust: $\equiv_{\mathcal{P}_{\text{ex}}} \in$ PT and $\equiv_{\Phi_{\text{ex}}} \in$ AT, where $\Phi_{\text{ex}} = \{\texttt{livesIn(Ein,us)}, \exists\texttt{livesIn(Men)}\}$ and $\mathcal{P}_{\text{ex}} = \{\texttt{livesIn}\}$. In $\equiv_{\mathcal{P}_{\text{ex}}}$ we trust that the knowledge provider is an expert in places of residence in general, while in $\equiv_{\Phi_{\text{ex}}}$ we trust that they can tell whether or not Einstein lives in the USA and Mendeleev lives somewhere. Consider two interpretations $\mathcal{I}_{\text{ex}}^1 = \{\texttt{livesIn(Ein,us)}\}$ and $\mathcal{I}_{\text{ex}}^2 = \{\texttt{livesIn(Ein,us)}, \texttt{livesIn(Ein,ru)}\}$. It is easy to see that $\mathcal{I}_{\text{ex}}^1 \not\equiv_{\mathcal{P}_{\text{ex}}} \mathcal{I}_{\text{ex}}^2$, while $\mathcal{I}_{\text{ex}}^1 \equiv_{\Phi_{\text{ex}}} \mathcal{I}_{\text{ex}}^2$.* ∎

We will use models of trust to relativise interpretations to a given one using the following *extender* function.

Definition 4. *An* extender, *denoted $\mathbb{E}$, is a function that maps each pair $(\equiv, \mathcal{I})$ where $\equiv$ is a model of trust and $\mathcal{I}$ is an interpretation, into 2^{PW} in the following way:*
$$\mathbb{E}(\equiv, \mathcal{I}) = \{\mathcal{J} \in \text{PW} \mid \mathcal{J} \equiv \mathcal{I}\}.$$

For a set of interpretations $\mathcal{S}$, $\mathbb{E}(\equiv, \mathcal{S}) = \bigcup_{\mathcal{I}\in\mathcal{S}} \mathbb{E}(\equiv, \mathcal{I})$.

Clearly, for each $\mathcal{S}$ it holds that $\mathcal{S} \subseteq \mathbb{E}(\equiv, \mathcal{S})$, while $\mathbb{E}(\equiv, \mathcal{S}) \subseteq \mathcal{S}$ does not hold in general.

Example 5. $\mathbb{E}(\equiv_{\mathcal{P}_{\text{ex}}}, \mathcal{I}_{\text{ex}}^1)$ *and* $\mathbb{E}(\equiv_{\Phi_{\text{ex}}}, \mathcal{I}_{\text{ex}}^1)$ *are, resp.:*
$\{\mathcal{I} \in \text{PW} \mid \{(\texttt{Ein,us})\} = \texttt{livesIn}^{\mathcal{I}}\}$ *and*
$\{\mathcal{I} \in \text{PW} \mid \mathcal{I} \models \texttt{livesIn(Ein,us)}, \text{ and } \mathcal{I} \not\models \exists\texttt{livesIn(Men)}\}.$∎

Trust-Sensitive Evolution Settings and Semantics. We distinguish between KB and ABox evolution. In the former case the whole KB changes, while in the latter case the TBox is fixed and only the ABox evolves. The following definition of evolution settings reflects this distinction.

Definition 6. *Let $\mathcal{L}$ be a DL and $\mathcal{C}$ a class of models of trust.*
- *An $(\mathcal{L}, \mathcal{C})$-setting $\mathcal{E}$ for KB evolution is a quadruple $(\mathcal{T}, \mathcal{A}, \mathcal{N}, \equiv)$, where $(\mathcal{T}, \mathcal{A})$ and $\mathcal{N}$ are satisfiable $\mathcal{L}$-KBs and $\equiv$ is a model of trust in $\mathcal{C}$.*
- *An $(\mathcal{L}, \mathcal{C})$-setting $\mathcal{E}$ for ABox evolution is a quadruple $(\mathcal{T}, \mathcal{A}, \mathcal{N}, \equiv)$, where $\mathcal{N}$ is an $\mathcal{L}$-ABox, $(\mathcal{T}, \mathcal{A})$ and $(\mathcal{T}, \mathcal{N})$ are satisfiable $\mathcal{L}$-KBs, $\equiv$ is a model of trust in $\mathcal{C}$.*

We will refer to $\mathcal{E}$ as just a $\mathcal{C}$-*setting* (resp., *setting*) when $\mathcal{L}$ is (resp., $\mathcal{L}$ and $\mathcal{C}$ are) clear or not important.

Example 7. *Consider $\mathcal{E}_{\text{ex}} = (\mathcal{T}_{\text{ex}}, \mathcal{A}_{\text{ex}}, \mathcal{N}_{\text{ex}}, \equiv_{\Phi_{\text{ex}}})$, a* (DL-Lite, AT)-*setting for ABox evolution, where the TBox is $\mathcal{T}_{\text{ex}} = \emptyset$, ABox is $\mathcal{A}_{\text{ex}} = \{\texttt{livesIn(Men,ru)}\}$, and the new ABox is $\mathcal{N}_{\text{ex}} = \{\texttt{livesIn(Ein,us)}, \texttt{livesIn(Men,us)}\}$.* ∎

We are now ready to show how models of trust can provide an external mechanism to guide evolution semantics. Intuitively, trust models work like filters that are applied to the (models of the) new knowledge $\mathcal{N}$ before performing the evolution. Recall that the classical MBAs 'pick' interpretations $\mathcal{J}$ from $\text{Mod}(\mathcal{N})$ that comes from the knowledge provider. In our case, however, we know that the knowledge provider cannot distinguish between any two $\equiv$-equivalent interpretations, i.e., any $\mathcal{J}'$ that is $\equiv$-equivalent to $\mathcal{J}$ is as 'good' as $\mathcal{J}$, and therefore, trust-sensitive evolution 'picks' interpretations from $\mathbb{E}(\equiv, \text{Mod}(\mathcal{N}))$ that extends $\text{Mod}(\mathcal{N})$ with all such $\mathcal{J}'$s. This approach corresponds to how Hunter and Booth (2015) introduced trust in the evolution of propositional theories.

Definition 8. *Let $\mathbb{S}$ be a selector. Then a* trust-sensitive evolution semantics $\text{sem}_\mathbb{S}$ *maps each setting $\mathcal{E} = (\mathcal{T}, \mathcal{A}, \mathcal{N}, \equiv)$ to a set of interpretations $\mathbb{S}(\text{Mod}(\mathcal{T}, \mathcal{A}), \mathbb{E}(\equiv, \mathcal{M}'))$, where $\mathcal{M}'$ is equal to $\text{Mod}(\mathcal{N})$ if $\mathcal{E}$ is for KB evolution and to $\text{Mod}(\mathcal{T}, \mathcal{N})$ if $\mathcal{E}$ is for ABox evolution.*

Example 9. *Consider two sets of interpretations: $\mathcal{M}_{\text{ex}} = \text{Mod}(\mathcal{T}_{\text{ex}}, \mathcal{A}_{\text{ex}})$, $\mathcal{M}'_{\text{ex}} = \text{Mod}(\mathcal{T}_{\text{ex}}, \mathcal{N}_{\text{ex}})$. Then, the evolution result $\text{sem}_{\mathbb{G}_{\prime\prime}^{\text{a}}}(\mathcal{E}_{\text{ex}}) = \mathbb{G}_{\prime\prime}^{\text{a}}(\mathcal{M}_{\text{ex}}, \mathbb{E}(\equiv_{\Phi_{\text{ex}}}, \mathcal{M}'_{\text{ex}}))$ is equal to*
$\{\mathcal{J} \in \text{PW} \mid \{\texttt{livesIn(Ein,us)}, \texttt{livesIn(Men,ru)}\} \subseteq \mathcal{J}\}.$
For classical MBAs, the evolution result for $\mathcal{E}_{\text{ex}}$ under $\mathbb{G}_{\prime\prime}^{\text{a}}$ is:

$\{\mathcal{J} \in \text{PW} \mid \{\texttt{livesIn(Ein,us)}, \texttt{livesIn(Men,ru)},$
$$\texttt{livesIn(Men,us)}\} \subseteq \mathcal{J}\}.$$ ∎

In practice one would expect the result of evolution to be a KB. Thus, a natural problem to study for MBAs is how evolution results can be axiomatised. Observe that the result of trust-sensitive evolution from Example 9 can be axiomatised respectively as

$$(\emptyset, \{\texttt{livesIn(Ein,us)}, \texttt{livesIn(Men,ru)}\}),$$

while the evolution result from Example 9 under the classical MBA $\mathbb{G}_{\{\}}^{\text{a}}$ can be axiomatised as

$$(\emptyset, \{\texttt{livesIn(Ein,de)}, \texttt{livesIn(Men,ru)}, \texttt{livesIn(Men,us)}\}).$$

In the classical case, the resulting KB is the union of the old $\mathcal{A}_{\text{ex}}$ and the new knowledge $\mathcal{N}_{\text{ex}}$. In the trust-sensitive case, the semantics *rejects* the new knowledge about Mendeleev since there is no trust in the fact that he is a US born.

4 Closure of DL-Lite Under Evolution

We now turn our attention to DL-Lite and show that evolution results in general cannot be axiomatised as DL-Lite KBs.We start with a definition of the closure problem for DLs.

Definition 10. *Let $\mathcal{L}$ be a DL, $\mathcal{C}$ a class of models of trust, and* sem *a trust-sensitive evolution semantics. Then, $\mathcal{L}$ is* closed under sem *for $\mathcal{C}$ if for every $(\mathcal{L},\mathcal{C})$-setting $\mathcal{E}$ there is an $\mathcal{L}$-KB $\mathcal{K}$ such that* $\mathsf{Mod}(\mathcal{K}) = \mathtt{sem}(\mathcal{E})$.

Total Trust and Total Distrust. The main reason why we introduce TT and TD is to verify on these extreme cases whether trust-sensitive evolution semantics behave in an intuitive way. In particular, we expect backward compatibility of trust-sensitive MBAs with the classical ones, that is, $\mathtt{sem}_{\mathbb{S}}$ should coincide with the corresponding classical semantics $\mathbb{S}$ in the case of TT. The following proposition confirms that this is indeed the case.

Proposition 11. *Let $\mathcal{E}=(\mathcal{T},\mathcal{A},\mathcal{N},\equiv_{\mathtt{TT}})$ be a $(\mathcal{L},\mathtt{TT})$-setting for some DL $\mathcal{L}$. Then for any selector $\mathbb{S}$, it holds that $\mathtt{sem}_{\mathbb{S}}(\mathcal{E}) = \mathbb{S}(\mathsf{Mod}(\mathcal{T},\mathcal{A}),\mathcal{M}')$, where $\mathcal{M}'$ is equal to $\mathsf{Mod}(\mathcal{N})$ if $\mathcal{E}$ is for KB and $\mathsf{Mod}(\mathcal{T},\mathcal{N})$ if $\mathcal{E}$ is for ABox evolution.*

The proposition implies that all non-closure results for DL-Lite under classical MBAs are inherited by trust-sensitive MBAs for TT. In particular, it is known that for ABox evolution DL-Lite is not closed under six out of eight MBAs: Calvanese *et al.* (2010) showed the non-closure under $\mathbb{L}^{\mathtt{a}}_{\{\}}$ and $\mathbb{L}^{\mathtt{a}}_{\#}$, Kharlamov *et al.* (2013) under $\mathbb{L}^{\mathtt{p}}_{\{\}}$ and $\mathbb{L}^{\mathtt{p}}_{\#}$, and finally Qi *et al.* (2015) under $\mathbb{G}^{\mathtt{a}}_{\{\}}$ and $\mathbb{G}^{\mathtt{a}}_{\#}$. For KB evolution Calvanese *et al.* (2010) showed the non-closure under all eight MBAs. Thus, the remaining open problem for trust-sensitive MBAs for TT is the closure under $\mathbb{G}^{\mathtt{p}}_{\{\}}$ and $\mathbb{G}^{\mathtt{p}}_{\#}$ for ABox evolution. The following theorem closes this gap.

Theorem 12. *For ABox evolution, DL-Lite is not closed under $\mathtt{sem}_{\mathbb{S}}$ for TT, where $\mathbb{S} \in \{\mathbb{G}^{\mathtt{p}}_{\{\}},\mathbb{G}^{\mathtt{p}}_{\#}\}$.*

Proof (Sketch). Regarding $\mathbb{G}^{\mathtt{p}}_{\{\}}$, one can check that for the TT-setting with $\mathcal{A} = \{\neg\exists R^-(a)\}$, $\mathcal{N} = \{\exists R^-(a)\}$, and $\mathcal{T} = \{A \sqsubseteq \exists R, \exists R \sqsubseteq A\}$, the set of interpretations obtained by evolution satisfies $\forall x.R(x,a) \to \exists y.(y \neq a \land R(x,y))$. One can show that this set is not axiomatisable in DL-Lite. The non-closure for the case of $\mathbb{G}^{\mathtt{p}}_{\#}$ can be shown similarly. $\square$

In the case of TD, regardless of the DL $\mathcal{L}$, selector $\mathbb{S}$, and $(\mathcal{L},\mathtt{TD})$-setting $\mathcal{E} = (\mathcal{T},\mathcal{A},\mathcal{N},\equiv_{\mathtt{TD}})$, it is easy to see that $\mathtt{sem}_{\mathbb{S}}(\mathcal{E}) = \mathsf{Mod}(\mathcal{T},\mathcal{A})$ for both KB and ABox evolution. Thus, $\mathtt{sem}_{\mathbb{S}}$ satisfies our intuition: it rejects the new information $\mathcal{N}$ as it is distrusted.

Assertion and Predicate Trust. We denote with S the set of all semantics introduced in Section 3, i.e., $\mathsf{S} = \bigcup_{x,y,\mathbb{Z}} \{\mathtt{sem}_{\mathbb{Z}^x_y}\}$, where $x \in \{\mathtt{a},\mathtt{p}\}$, $y \in \{\{\},\#\}$, and $\mathbb{Z} \in \{\mathbb{L},\mathbb{G}\}$.

Observe that for each $(\mathcal{L},\mathtt{TT})$-setting $\mathcal{E}_{\mathtt{TT}}=(\mathcal{T},\mathcal{A},\mathcal{N},\equiv_{\mathtt{TT}})$ one can construct an $(\mathcal{L},\mathtt{PT})$-setting $\mathcal{E}_{\mathtt{PT}} = (\mathcal{T},\mathcal{A},\mathcal{N},\equiv_{\mathcal{P}})$, where $\mathcal{P} = \mathtt{pred}(\mathcal{T} \cup \mathcal{A} \cup \mathcal{N})$, such that for each sem $\in \mathsf{S}$ we have $\mathtt{sem}(\mathcal{E}_{\mathtt{TT}}) = \mathtt{sem}(\mathcal{E}_{\mathtt{PT}})$. Therefore, all the non-closure results for TT are inherited by PT.

Finally, we turn our attention to AT and show the non-closure of DL-Lite under various trust-sensitive semantics.

Theorem 13. *For* AT *it holds that:*
- *For KB evolution and each* sem $\in \mathsf{S}$, DL-Lite *is not closed under* sem*; this holds already in the case when the new information consists of one TBox assertion.*
- *For ABox evolution and each* sem $\in \mathsf{S}$ DL-Lite *is not closed under* sem.

In order to prove these results one can show that for each semantics sem considered in the theorem, there is a setting $\mathcal{E}$, such that $\mathtt{sem}(\mathcal{E})$ is a set of models that satisfies a so-called *genuine disjunction*. That is, $\mathtt{sem}(\mathcal{E})$ satisfies $\varphi \lor \psi$, for some ABox assertions φ and ψ, but does not satisfy either φ or ψ. By Lemma 1 from [Calvanese *et al.*, 2010] such a set of interpretations is not axiomatisable in DL-Lite.

5 Approximation of Evolution in DL-Lite

Since DL-Lite is not closed under the trust-sensitive MBAs, we turn our attention to approximation of evolution results. In this section we focus on ABox evolution and thus all settings are for ABox evolution.

A *sound approximation* of $\mathtt{sem}(\mathcal{E})$ is a KB $\mathcal{K}$ such that $\mathtt{sem}(\mathcal{E}) \subseteq \mathsf{Mod}(\mathcal{K})$, and it is *maximal* if no other sound approximation $\mathcal{K}'$ exists s.t. $\mathsf{Mod}(\mathcal{K}') \subset \mathsf{Mod}(\mathcal{K})$. Sound approximation of evolution in the context of DLs has been studied for classical MBAs by De Giacomo *et al.* (2009), Kharlamov *et al.* (2013), and Qi *et al.* (2015). We extend the notion of sound approximation by considering s-approximations which we introduce next. In order to define them we use the following notation. Let Σ be a signature, then $\mathcal{I}|_{\Sigma}$ is a sub-interpretation of $\mathcal{I}$ consisting of all atoms of $\mathcal{I}$ whose predicates are in Σ, and for a set of models $\mathcal{S}$, we define $\mathcal{S}|_{\Sigma}$ as $\{\mathcal{I}|_{\Sigma} \mid \mathcal{I} \in \mathcal{S}\}$. Finally, $\mathcal{S} \subseteq_{\Sigma} \mathcal{S}'$, if $\mathcal{S}|_{\Sigma} \subseteq \mathcal{S}'|_{\Sigma}$.

Definition 14. *Let $\mathcal{S}$ be a set of interpretations, and $\mathcal{K}$ a knowledge base. Then, $\mathcal{K}$ is a* sound s-approximation *of $\mathcal{S}$ if $\mathcal{S} \subseteq_{\mathtt{pred}(\mathcal{S})} \mathsf{Mod}(\mathcal{K})$. Moreover, $\mathcal{K}$ is a* maximal *sound s-approximation of $\mathcal{S}$ if no other sound s-approximation $\mathcal{K}'$ of $\mathcal{S}$ exists such that $\mathsf{Mod}(\mathcal{K}') \subset_{\mathtt{pred}(\mathcal{S})} \mathsf{Mod}(\mathcal{K})$. Finally, $\mathcal{K}$ is an* s-axiomatisation *of $\mathcal{S}$ if $\mathcal{S} = \mathsf{Mod}(\mathcal{K})|_{\mathtt{pred}(\mathcal{S})}$.*

Note that s-approximations coincide with sound approximations when $\mathtt{pred}(\mathcal{S}) = \mathtt{pred}(\mathcal{K})$.

5.1 Total Trust

In this section we will study $\mathtt{sem}_{\mathbb{S}}$ evolution in case of TT and $\mathbb{S} \in \{\mathbb{G}^{\mathtt{a}}_{\{\}},\mathbb{G}^{\mathtt{a}}_{\#}\}$ and use the following notations: $\mathtt{cl}_{\mathcal{T}}(\mathcal{A})$ is the set of membership assertions φ s.t. $\mathcal{A} \cup \mathcal{T} \models \varphi$ and $\mathtt{cl}(\mathcal{T})$ is the set of TBox assertions φ s.t. $\mathcal{T} \models \varphi$. It is known [Calvanese *et al.*, 2007] that in DL-Lite $\mathtt{cl}_{\mathcal{T}}(\mathcal{A})$ and $\mathtt{cl}(\mathcal{T})$ are finite and can be computed in polynomial time.

Let $\mathcal{E} = (\mathcal{T},\mathcal{A},\mathcal{N},\equiv_{\mathtt{TT}})$ be a TT-setting. Qi *et al.* (2015) showed that the algorithm AtAlg, introduced by Kharlamov and Zheleznyakov (2011), computes a maximal sound approximation $\mathcal{K}'$ of $\mathbb{S}(\mathsf{Mod}(\mathcal{T},\mathcal{A}),\mathsf{Mod}(\mathcal{T},\mathcal{N}))$, where $\mathbb{S} \in \{\mathbb{G}^{\mathtt{a}}_{\{\}},\mathbb{G}^{\mathtt{a}}_{\#}\}$. Given $\mathcal{E}$, AtAlg returns an ABox $\mathcal{N} \cup \mathcal{A}'$, where $\mathcal{A}'$ is the maximal subset of $\mathtt{cl}_{\mathcal{T}}(\mathcal{A})$ such that $(\mathcal{T},\mathcal{N} \cup \mathcal{A}')$ is satisfiable. By Proposition 11, $\mathcal{K}'$ is also a maximal sound approximation of $\mathtt{sem}_{\mathbb{S}}(\mathcal{E})$. However, $\mathcal{K}'$ is not necessarily their maximal sound s-approximation, as illustrated next.

Example 15. *Consider a* TT-*setting $\mathcal{E}^2_{\mathrm{ex}} = (\mathcal{T}^2_{\mathrm{ex}}, \mathcal{A}^2_{\mathrm{ex}}, \mathcal{N}^2_{\mathrm{ex}}, \equiv_{\mathtt{TT}})$ over the signature $\Sigma_{\mathrm{ex}} = \{\mathtt{livesIn}, \mathtt{place}\}$, where $\mathcal{T}^2_{\mathrm{ex}} = \{\exists\mathtt{livesIn}^- \sqsubseteq \mathtt{place}\}$, $\mathcal{A}^2_{\mathrm{ex}} =*

$\{\exists \texttt{livesIn}(\texttt{Men})\}$, and $\mathcal{N}_{\text{ex}}^2 = \{\neg\exists\texttt{livesIn}(\texttt{Men})\}$. *The maximal sound approximation obtained with* `AtAlg` *is* $(\mathcal{T}_{\text{ex}}^2, \mathcal{N}_{\text{ex}}^2)$. *However, any model* $\mathcal{J}$ *from* $\mathsf{Mod}(\mathcal{T}_{\text{ex}}^2, \mathcal{N}_{\text{ex}}^2)$ *with* $\texttt{place}^{\mathcal{J}} = \emptyset$ *is not in* $\mathcal{M} = \mathsf{sem}_{\mathbb{G}_{/\!/}^{\text{a}}}(\mathcal{E}_{\text{ex}}^2)$. *We can rule out such models by introducing a fresh role* P *that would 'enforce' the existence of an element from* Δ *in the interpretation of* `place`. *Indeed, consider a KB* $\mathcal{K}^s = (\mathcal{T}^s, \mathcal{A}^s)$, *where* $\mathcal{T}^s = \mathcal{T}_{\text{ex}}^2 \cup \{\exists P^- \sqsubseteq \texttt{place}\}$ *and* $\mathcal{A}^s = \{\exists P(a^*)\}$ *with* P *and* a^* *a fresh role and constant, respectively. Note also that* $\mathcal{K}^s$ *is a sound s-approximation of* $\mathcal{M}$ *and* $\mathsf{Mod}(\mathcal{K}^s) \subset_{\Sigma_{\text{ex}}} \mathsf{Mod}(\mathcal{K})$ *holds as no model* $\mathcal{J}$ *of* $\mathcal{K}^s$ *is such that* $\texttt{place}^{\mathcal{J}} = \emptyset$.

In contrast to `AtAlg`, for a given TT-setting $\mathcal{E}$, Algorithm 1 (TT-SApprox) provides a maximal sound s-approximation for $\mathsf{sem}_{\mathbb{S}}(\mathcal{E})$, where $\mathbb{S} \in \{\mathbb{G}_{/\!/}^{\text{a}}, \mathbb{G}_{\#}^{\text{a}}\}$. One can follow the steps of the algorithm in Example 15. TT-SApprox first computes the maximal sound approximation using `AtAlg` (Line 1). Then, in the spirit of Example 15, the algorithm finds general concepts C whose interpretations should not be empty in the resulting set of models (Line 8) and ensures that this will not happen via introducing new TBox and ABox assertions with fresh roles and constants (Lines 6 and 9) as in the example. Finally, if an interpretation of C should contain at least n constants, then the algorithm ensures that (Lines 10-11).

The following theorem shows that TT-SApprox efficiently computes maximal sound s-approximations.

Theorem 16. *Let* $\mathcal{E} = (\mathcal{T}, \mathcal{A}, \mathcal{N}, \equiv_{\text{TT}})$ *be a* (DL-Lite$_{core}$, TT)-*setting. Then,* TT-SApprox$(\mathcal{E})$ *is a maximal sound s-approximation of* $\mathsf{sem}_{\mathbb{S}}$, *where* $\mathbb{S} \in \{\mathbb{G}_{/\!/}^{\text{a}}, \mathbb{G}_{\#}^{\text{a}}\}$. *Moreover,* TT-SApprox *runs in time polynomial in* $|\mathcal{T} \cup \mathcal{A} \cup \mathcal{N}|$.

A practical benefit of sound s-approximations is that they preserve important queries that may be lost by sound approximations. We will now introduce a class of such queries. An example query from this class is:

> *Is it true that a university has at least five (distinct) researchers* $r_1, \ldots, r_5$ *such that* r_1, r_2, r_3 *work on the first project,* r_2, r_3, r_4 *on the second one, and* r_1, r_5 *on the third one?*

Let $\Theta(x)$ be a formula recursively defined as follows:

$$\Psi(x) ::= A(x) \mid \exists x' R(x, x') \mid \exists x' R(x', x),$$
$$\Theta(x) ::= \Psi(x) \mid \neg\Psi(x).$$

Then, $\mathcal{Q}$ is the class of queries of the following form:

$$\exists x_1, \ldots, x_n \Big(\bigwedge_{i \neq j} (x_i \neq x_j) \wedge \bigwedge_{k=1}^{m} \bigwedge_{\substack{i \in X_k \\ X_k \subseteq \{1, \ldots, n\}}} (\Theta_k(x_i)) \Big).$$

Theorem 17. *Let* $q \in \mathcal{Q}$ *and* $\mathsf{sem} = \mathsf{sem}_{\mathbb{G}_y^{\text{a}}}$ *be an evolution semantics, where* $y \in \{/\!/, \#\}$. *Then, there exists a* TT-*setting* $\mathcal{E}_q$ *such that* $\mathsf{sem}(\mathcal{E}_q) \models q$ *and* TT-SApprox$(\mathcal{E}_q) \models q$, *but* AtAlg$(\mathcal{E}_q) \not\models q$.

5.2 Predicate Trust

Let $(\mathcal{T}, \mathcal{A})$ be a DL-Lite KB and $\equiv_{\mathcal{P}} \in$ PT for some $\mathcal{P}$. A natural approach to this case would be to capture $\mathbb{E}(\equiv_{\mathcal{P}}, \mathsf{Mod}(\mathcal{T}, \mathcal{A}))$ with some theory, thus reducing the

---

Algorithm 1: TT-SApprox

INPUT : a (DL-Lite$_{core}$, TT)-setting $\mathcal{E} = (\mathcal{T}, \mathcal{A}, \mathcal{N}, \equiv_{\text{TT}})$
OUTPUT: a DL-Lite$_{core}$ KB $(\mathcal{T}', \mathcal{A}')$

1   set $\mathcal{T}' = \mathcal{T}$ and $\mathcal{A}' = \texttt{AtAlg}(\mathcal{E})$
2   **introduce** a fresh constant a^* not occurring in $\mathcal{T}, \mathcal{A}$, nor $\mathcal{N}$
3   **for each** *basic role* R *over* $\mathsf{pred}(\mathcal{A}, \mathcal{T})$ **do**
4   **introduce** a fresh atomic role $P_R \notin \mathsf{pred}(\mathcal{E})$
5   **if** $\exists R(a) \in \mathtt{cl}_{\mathcal{T}}(\mathcal{A})$ *for some* a **then**
6   set $\mathcal{A}' = \mathcal{A}' \cup \{\exists P_R(a^*)\}$
7   **for each** $\exists R^- \sqsubseteq C \in \mathtt{cl}(\mathcal{T})$ *for some general concept* C **do**
8   **if** $\exists R(a) \in \mathtt{cl}_{\mathcal{T}}(\mathcal{A})$ *for some* a **and** *either* $\neg\exists R(a) \notin \mathtt{cl}_{\mathcal{T}}(\mathcal{N})$ *or there is no constant* b *s.t.* $C(b) \in \mathtt{cl}_{\mathcal{T}}(\mathcal{A})$ *and* $\neg C(b) \in \mathtt{cl}_{\mathcal{T}}(\mathcal{N})$ **then**
9   set $\mathcal{T}' = \mathcal{T}' \cup \{\exists P_R^- \sqsubseteq C\}$
10   **for each** *pair of introduced roles* P_R *and* P_S **do**
11   **if** $\exists R^- \sqsubseteq \neg\exists S^-$ *is in* $\mathtt{cl}(\mathcal{T})$ **then set** $\mathcal{T}' = \mathcal{T}' \cup \{\exists P_R^- \sqsubseteq \neg\exists P_S^-\}$
12   **return** $(\mathcal{T}', \mathcal{A}' \cup \mathcal{N})$

---

problem to the TT case. On the first glance, it would be sufficient just to remove assertions that are over the 'wrong' signature from $\mathtt{cl}_{\mathcal{T}}(\mathcal{A})$. This, however, is not enough as shown in the following example.

Example 18. *Consider* $\equiv_{\mathcal{P}_{\text{ex}}^3} \in$ PT, *where* $\mathcal{P}_{\text{ex}}^3 = \{\texttt{person}\}$ *and* $(\mathcal{T}_{\text{ex}}^3, \mathcal{A}_{\text{ex}}^3)$, *where* $\mathcal{T}_{\text{ex}}^3 = \{\exists\texttt{livesIn}^- \sqsubseteq \neg\texttt{person}\}$, $\mathcal{A}_{\text{ex}}^3 = \{\texttt{person}(\texttt{Men}), \exists\texttt{livesIn}(\texttt{Men})\}$. *Consider also* $\mathcal{A}' = \{\texttt{person}(\texttt{Men})\}$, *a subset of* $\mathtt{cl}_{\mathcal{T}}(\mathcal{A}_{\text{ex}}^3)$ *consisting of all the membership assertions over* $\mathcal{P}_{\text{ex}}^3$. *One can see that* $\mathcal{I} = \{\texttt{person}(x) \mid x \in \Delta\}$ *is a model of* $(\mathcal{T}_{\text{ex}}^3, \mathcal{A}')$, *while* $\mathcal{I} \notin \mathbb{E}(\equiv_{\mathcal{P}_{\text{ex}}^3}, \mathsf{Mod}(\mathcal{T}_{\text{ex}}^3, \mathcal{A}_{\text{ex}}^3))$. *However, we can still capture* $\mathbb{E}(\equiv_{\mathcal{P}_{\text{ex}}^3}, \mathsf{Mod}(\mathcal{T}_{\text{ex}}^3, \mathcal{A}_{\text{ex}}^3))$ *with* $(\mathcal{T}^s, \mathcal{A}^s)$, *where* $\mathcal{T}^s = \mathcal{T}_{\text{ex}}^3 \cup \{\exists P^- \sqsubseteq \neg\texttt{person}\}$ *and* $\mathcal{A}^s = \mathcal{A}' \cup \{\exists P(a^*)\}$ *for a fresh role* P *and constant* a^*. *Indeed, one can check that* $\mathbb{E}(\equiv_{\mathcal{P}_{\text{ex}}^3}, \mathsf{Mod}(\mathcal{T}_{\text{ex}}^3, \mathcal{A}_{\text{ex}}^3)) = \mathsf{Mod}(\mathcal{T}^s, \mathcal{A}^s)|_{\mathsf{pred}(\mathcal{T}_{\text{ex}}^3, \mathcal{A}_{\text{ex}}^3)}$.

$\mathbb{E}(\equiv_{\mathcal{P}}, \mathsf{Mod}(\mathcal{K}))$ can still be captured using s-axiomatisation. Consider Algorithm 2 (PT-ExtendSAx) that, for a given $(\mathcal{T}, \mathcal{A})$ and $\mathcal{P}$, returns a maximal sound s-approximation of $\mathbb{E}(\equiv_{\mathcal{P}}, \mathsf{Mod}(\mathcal{T}, \mathcal{A}))$. One can follow the steps of the algorithm in Example 18: it finds concepts C whose interpretation should not be empty in each $\mathcal{I}$ of $\mathbb{E}(\equiv_{\mathcal{P}}, \mathsf{Mod}(\mathcal{T}, \mathcal{A}))$ (Line 10) and ensures their non-emptiness (Lines 8 and 9). Finally, as in Algorithm 1, PT-ExtendSAx guarantees that the minimal number of constants in C is as required (Lines 11-12).

Theorem 19. *Let* $(\mathcal{T}, \mathcal{A})$ *be a* DL-Lite$_{core}$ *KB and let* $\equiv_{\mathcal{P}} \in$ PT *for some finite set of predicates* $\mathcal{P}$. *Then* PT-ExtendSAx$(\mathcal{T}, \mathcal{A}, \mathcal{P})$ *is an s-axiomatisation of* $\mathbb{E}(\equiv_{\mathcal{P}}, \mathsf{Mod}(\mathcal{T}, \mathcal{A}))$. *Moreover,* PT-ExtendSAx *runs in time polynomial in* $|\mathcal{T} \cup \mathcal{A} \cup \mathcal{P}|$.

Finally, we will show that PT-ExtendSAx can be used to compute maximal sound s-approximations of $\mathsf{sem}_{\mathbb{S}}$ for $\mathbb{S} = \mathbb{L}_{()}^{\text{a}}$, that corresponds to a widely accepted Winslett's evolution semantics [De Giacomo *et al.*, 2006; Winslett, 1990].

Theorem 20. *Let* $\mathcal{P}$ *be a finite set of predicates and* $\mathcal{E}_{\mathcal{P}} = (\mathcal{T}, \mathcal{A}, \mathcal{N}, \equiv_{\mathcal{P}})$ *a* (DL-Lite$_{core}$, PT)-*setting. Let also* $(\mathcal{T}', \mathcal{N}') = $ PT-ExtendSAx$((\mathcal{T}, \mathcal{N}), \mathcal{P})$ *and* $\mathcal{E}_{\text{TT}} = $

Algorithm 2: `PT-ExtendSAx`

INPUT : a DL-Lite$_{core}$ KB $(\mathcal{T}, \mathcal{A})$,
a finite set of predicates $\mathcal{P}$
OUTPUT: a DL-Lite$_{core}$ KB $(\mathcal{T}', \mathcal{A}')$

1 **set** $\mathcal{T}' = \mathcal{T}$ **and** $\mathcal{A}' = \emptyset$
2 **introduce** a fresh constant a^* not occurring in $\mathcal{T}$ nor $\mathcal{A}$
3 **for each** $\alpha \in \text{cl}_{\mathcal{T}}(\mathcal{A})$ **do**
4   **if** *the predicate of α is in $\mathcal{P}$* **then set** $\mathcal{A}' = \mathcal{A}' \cup \{\alpha\}$
5 **for each** *basic role R over* $\text{pred}(\mathcal{T}, \mathcal{A})$ **do**
6   **introduce** a fresh atomic role $P_R \notin \text{pred}(\mathcal{T}, \mathcal{A})$
7   **if** $\exists R(a) \in \text{cl}_{\mathcal{T}}(\mathcal{A})$ *for some constant a* **then**
8   **set** $\mathcal{A}' = \mathcal{A}' \cup \{\exists P_R(a^*)\}$
9   **for each** $\exists R^- \sqsubseteq C$ in $\text{cl}(\mathcal{T})$ *with* $C \in \mathcal{P}$ **do**
10   **set** $\mathcal{T}' = \mathcal{T}' \cup \{\exists P_R^- \sqsubseteq C\}$
11 **for each** *pair of fresh atomic roles P_R and P_S* **do**
12   **if** $\exists R^- \sqsubseteq \neg\exists S^-$ *is in* $\text{cl}(\mathcal{T})$ **then set**
     $\mathcal{T}' = \mathcal{T}' \cup \{\exists P_R^- \sqsubseteq \neg\exists P_S^-\}$
13 **return** $(\mathcal{T}', \mathcal{A}')$

$(\mathcal{T}', \mathcal{A}, \mathcal{N}', \equiv_{\text{TT}})$ *be* $(\text{DL-Lite}_{core}, \text{TT})$-*setting. Then, if $\mathcal{K}$ is a maximal sound s-approximation (resp., s-axiomatisation) of* $\text{sem}_{\mathbb{L}_{\{\}}^{\text{a}}}(\mathcal{E}_{\text{TT}})$, *then $\mathcal{K}$ is also a maximal sound s-approximation (resp., s-axiomatisation) of* $\text{sem}_{\mathbb{L}_{\{\}}^{\text{a}}}(\mathcal{E}_{\mathcal{P}})$.

5.3 Assertion Trust

Let $(\mathcal{T}, \mathcal{A})$ be a DL-Lite KB and $\equiv_{\Phi} \in \text{AT}$ for some finite set of membership assertions Φ. Firstly, observe that the set of models $\mathbb{E}(\equiv_{\Phi}, \text{Mod}(\mathcal{T}, \mathcal{A}))$ can be axiomatised and Algorithm 3 (`AT-ExtendAx`), where $\neg\alpha = \neg B(a)$ if $\alpha = B(a)$ and $\neg\alpha = B(a)$ if $\alpha = \neg B(a)$ for some basic concept B, provides this axiomatisation. The algorithm keeps assertions of $\text{cl}_{\mathcal{T}}(\mathcal{A})$ such that they or their negations are in Φ. The following theorem shows the correctness of `AT-ExtendAx`.

Theorem 21. *Let $(\mathcal{T}, \mathcal{A})$ be a DL-Lite KB and $\equiv_{\Phi} \in \text{AT}$ for some finite set of membership assertions Φ. Then, $\mathbb{E}(\equiv_{\Phi}, \text{Mod}(\mathcal{T}, \mathcal{A})) = \text{Mod}(\texttt{AT-ExtendAx}(\mathcal{T}, \mathcal{A}, \Phi))$, and* `AT-ExtendAx` *runs in time polynomial in* $|\mathcal{T} \cup \mathcal{A} \cup \Phi|$.

An immediate consequence of the theorem is that each AT-setting $\mathcal{E}$ can be transformed into a TT setting $\mathcal{E}_{\text{TT}}$ with the same evolution result under $\text{sem}_{\mathbb{S}}$ with any selector $\mathbb{S}$.

Corollary 22. *Let $\mathbb{S}$ be a selector and Φ a finite set of membership assertions. Let $\mathcal{E}_{\Phi} = (\mathcal{T}, \mathcal{A}, \mathcal{N}, \equiv_{\Phi})$ and $\mathcal{E}_{\text{TT}} = (\mathcal{T}, \mathcal{A}, \mathcal{N}_{\text{TT}}, \equiv_{\text{TT}})$ be $(\text{DL-Lite}, \text{AT})$ and $(\text{DL-Lite}, \text{TT})$-settings, respectively, where $(\mathcal{T}, \mathcal{N}_{\text{TT}}) = \texttt{AT-ExtendAx}(\mathcal{T}, \mathcal{N}, \Phi)$. Then $\text{sem}_{\mathbb{S}}(\mathcal{E}_{\Phi}) = \text{sem}_{\mathbb{S}}(\mathcal{E}_{\text{TT}})$.*

6 Related Work and Discussions

Related Work. To the best of our knowledge, this is the first work that combines trust and evolution in the context of DLs. The closest research to ours is knowledge management with preferences where either logical formulae or predicates are ordered since less preferred elements can be seen as less-trusted. However, this rather corresponds to defining levels of importance than trust. Bienvenu *et al.* (2014) studied inconsistency-tolerant semantics for querying inconsistent KBs. They rely on KB repairs which are subsets of the ABox that are consistent with the TBox, and use various models of preferences to determine the most important repairs. Since they do not select repairs that are of low importance,

---

Algorithm 3: `AT-ExtendAx`

INPUT : a DL-Lite KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$,
a finite set of membership assertions Φ
OUTPUT: a DL-Lite KB $(\mathcal{T}, \mathcal{A}')$

1 **set** $\mathcal{A}' = \emptyset$
2 **for each** $\alpha \in \text{cl}_{\mathcal{T}}(\mathcal{A})$ **do**
3   **if** *either α or $\neg\alpha$ is in Φ* **then set** $\mathcal{A}' = \mathcal{A}' \cup \{\alpha\}$
4 **return** $(\mathcal{T}, \mathcal{A}')$

---

this can be seen as a trust-based KB repairing. However, their approach is based on formulae and thus closer to so-called *formula-based evolution* [Eiter and Gottlob, 1992] rather than to the model-based approach that we study in this paper.

Qi and Du (2009) studied evolution under a modified version of $\mathbb{G}_{\#}^{\text{p}}$ selector that relies on predicate-based preferences in selecting models of $\text{Mod}(\mathcal{N})$. The crucial difference between their and our work is that their evolution result $\mathcal{M}$ is a subset of $\text{Mod}(\mathcal{N})$ that consists of the most important models, while in our case we construct $\mathcal{M}$ by first extending $\text{Mod}(\mathcal{N})$ according with the model of trust and then choosing minimally distant elements from this extended set regardless their importance. Note that our approach can be combined with the selector of [Qi and Du, 2009], but this requires further investigation.

Conclusion. We have formalised the notion of trust and evolution semantics as operators $\mathbb{E}$ and $\mathbb{S}$ and have shown how they can be composed to obtain trust-sensitive evolution semantics. This approach can be generalised to any notions of trust and evolution semantics that can be captured via operators. We have applied trust-sensitive MBAs to DL-Lite and have shown that under all of them DL-Lite is not closed. On the one hand, this is expected since DL-Lite has a limited expressive power, and already in the case of classical MBAs can return sets of models $\mathcal{M}$ whose axiomatisation requires syntactic constructs beyond DL-Lite. On the other hand, this is not selfevident since trust-sensitive MBAs in general return sets of models that are very different from the ones of classical MBAs. Thus, for both classical and trust-sensitive MBAs the main challenge is to find a 'good' notion of evolution approximation and to develop (efficient) algorithms to compute 'optimal' approximations. We have proposed a novel notion of sound s-approximation that captures evolution results better than sound approximations previously considered in the evolution context, and that preserve queries that may be lost by regular approximations. We have provided polynomial-time algorithms that compute maximal sound s-approximations for classical and several trust-sensitive MBAs. Our algorithms work either directly on trust-sensitive evolution settings or they can be utilised as subroutines to reduce evolution from trust-sensitive to classical settings. We see these algorithms as a starting point for developing efficient procedures for automated knowledge update and revision. We also believe this work is a timely contribution for the Semantic Web, where applications may depend on third-party information that is only partly trusted.

Future Work. We plan to address the following problems. It is open whether our results for total trust (Theorem 16) and assertion trust (Theorem 19) can be extended from DL-Lite$_{core}$ to DL-Lite, while we conjecture that they hold for DL-Lite without functionality assertions. It is also open whether Theorem 20 holds for selectors other than $\mathbb{L}_{\{\}}^{\text{a}}$.

References

[Baader *et al.*, 2003] Franz Baader, Diego Calvanese, Deborah L McGuinness, Daniele Nardi, and Peter F Patel-Schneider. The Description Logic Handbook: Theory, Implementation, and Applications. *Description Logic Handbook*, 2003.

[Bienvenu *et al.*, 2014] Meghyn Bienvenu, Camille Bourgaux, and François Goasdoué. Querying inconsistent description logic knowledge bases under preferred repair semantics. In *AAAI*, pages 996–1002, 2014.

[Bollacker *et al.*, 2008] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. Freebase: A Collaboratively Created Graph Database for Structuring Human Knowledge. In *SIGMOD*, pages 1247–1250, 2008.

[Calvanese *et al.*, 2007] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tractable Reasoning and Efficient Query Answering in Description Logics: The *DL-Lite* Family. *J. Autom. Reasoning*, 39(3):385–429, 2007.

[Calvanese *et al.*, 2010] Diego Calvanese, Evgeny Kharlamov, Werner Nutt, and Dmitriy Zheleznyakov. Evolution of DL-Lite Knowledge Bases. In *ISWC*, pages 112–128, 2010.

[Cuenca Grau *et al.*, 2012] Bernardo Cuenca Grau, Ernesto Jiménez-Ruiz, Evgeny Kharlamov, and Dmitriy Zheleznyakov. Ontology Evolution Under Semantic Constraints. In *KR*, 2012.

[De Giacomo *et al.*, 2006] Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, and Riccardo Rosati. On the Update of Description Logic Ontologies at the Instance Level. In *AAAI*, pages 1271–1276, 2006.

[De Giacomo *et al.*, 2009] Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, and Riccardo Rosati. On Instance-level Update and Erasure in Description Logic Ontologies. *JLC*, 19(5):745–770, 2009.

[Eiter and Gottlob, 1992] Thomas Eiter and Georg Gottlob. On the Complexity of Propositional Knowledge Base Revision, Updates, and Counterfactuals. In *PODS*, pages 261–273, 1992.

[Flouris *et al.*, 2008] Giorgos Flouris, Dimitris Manakanatas, Haridimos Kondylakis, Dimitris Plexousakis, and Grigoris Antoniou. Ontology Change: Classification and Survey. *Knowledge Eng. Review*, 23(2):117–152, 2008.

[Fridman Noy *et al.*, 2004] Natalya Fridman Noy, Sandhya Kunnatur, Michel Klein, and Mark Musen. Tracking Changes During Ontology Evolution. In *ISWC*, pages 259–273, 2004.

[Furche *et al.*, 2012] Tim Furche, Georg Gottlob, Giovanni Grasso, Omer Gunes, Xiaoanan Guo, Andrey Kravchenko, Giorgio Orsi, Christian Schallhart, Andrew Sellers, and Cheng Wang. DIADEM: Domain-Centric, Intelligent, Automated Data Extraction Methodology. In *WWW Demo*, pages 267–270, 2012.

[Haase and Stojanovic, 2005] Peter Haase and Ljiljana Stojanovic. Consistent Evolution of OWL Ontologies. In *ESWC*, pages 182–197, 2005.

[Hunter and Booth, 2015] Aaron Hunter and Richard Booth. Trust-Sensitive Belief Revision. In *IJCAI*, pages 3062–3068, 2015.

[Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto Mendelzon. On the Difference between Updating a Knowledge Base and Revising It. In *KR*, pages 387–394, 1991.

[Kharlamov and Zheleznyakov, 2011] Evgeny Kharlamov and Dmitriy Zheleznyakov. Capturing Instance Level Ontology Evolution for DL-lite. In *ISWC*, pages 321–337, 2011.

[Kharlamov *et al.*, 2013] Evgeny Kharlamov, Dmitriy Zheleznyakov, and Diego Calvanese. Capturing Model-Based Ontology Evolution at the Instance Level: The Case of DL-Lite. *J. Comput. Syst. Sci.*, 79(6):835–872, 2013.

[Konev *et al.*, 2008] Boris Konev, Dirk Walther, and Frank Wolter. The Logical Difference Problem for Description Logic Terminologies. In *IJCAR*, pages 259–274, 2008.

[Liu *et al.*, 2011] Hongkai Liu, Carsten Lutz, Maja Miličić, and Frank Wolter. Foundations of Instance Level Updates in Expressive Description Logics. *Artif. Intell.*, 175(18):2170–2197, 2011.

[Qi and Du, 2009] Guilin Qi and Jianfeng Du. Model-based Revision Operators for Terminologies in Description Logics. In *IJCAI*, pages 891–897, 2009.

[Qi *et al.*, 2015] Guilin Qi, Zhe Wang, Kewen Wang, Xuefeng Fu, and Zhiqiang Zhuang. Approximating Model-Based ABox Revision in DL-Lite: Theory and Practice. In *AAAI*, pages 254–260, 2015.

[Stearns *et al.*, 2001] Michael Q. Stearns, Colin Price, Kent A. Spackman, and Amy Y. Wang. SNOMED Clinical Terms: Overview of the Development Process and Project Status. In *AMIA*, pages 662–666, 2001.

[Suchanek and Weikum, 2013] Fabian M Suchanek and Gerhard Weikum. Knowledge Harvesting from Text and Web Sources. *ICDE*, pages 1250–1253, 2013.

[Wang *et al.*, 2010] Zhe Wang, Kewen Wang, and Rodney Topor. A New Approach to Knowledge Base Revision in DL-Lite. In *AAAI*, 2010.

[Wang *et al.*, 2015] Zhe Wang, Kewen Wang, Zhiqiang Zhuang, and Guilin Qi. Instance-Driven Ontology Evolution in DL-Lite. In *AAAI*, pages 1656–1662, 2015.

[Winslett, 1990] Marianne Winslett. *Updating Logical Databases*. Cambridge University Press, 1990.

[Wu and Lécué, 2014] Jiewen Wu and Freddy Lécué. Towards Consistency Checking over Evolving Ontologies. In *CIKM*, pages 909–918, 2014.

[Zhuang *et al.*, 2014] Zhiqiang Zhuang, Zhe Wang, Kewen Wang, and Guilin Qi. Contraction and Revision Over DL-Lite TBoxes. In *AAAI*, pages 1149–1156, 2014.

[Zhuang *et al.*, 2016] Zhiqiang Zhuang, Zhe Wang, Kewen Wang, and Guilin Qi. DL-Lite Contraction and Revision. *JAIR*, 56:329–378, 2016.