

Instance-Driven Ontology Evolution in DL-Lite

Zhe Wang and Kewen Wang and Zhiqiang Zhuang

School of Information and Communication Technology, Griffith University, Australia

Guilin Qi

School of Computer Science and Engineering, Southeast University, China
 State Key Lab for Novel Software Technology, Nanjing University, Nanjing, China

Abstract

The development and maintenance of large and complex ontologies are often time-consuming and error-prone. Thus, automated ontology learning and evolution have attracted intensive research interest. In data-centric applications where ontologies are designed from the data or automatically learnt from it, when new data instances are added that contradict the ontology, it is often desirable to incrementally revise the ontology according to the added data. In description logics, this problem can be intuitively formulated as the operation of TBox contraction, i.e., rational elimination of certain axioms from the logical consequences of a TBox, and it is w.r.t. an ABox. In this paper we introduce a model-theoretic approach to such a contraction problem by using an alternative semantic characterisation of DL-Lite TBoxes. We show that entailment checking (without necessarily first computing the contraction result) is in coNP, which does not shift the corresponding complexity in propositional logic, and the problem is tractable when the size of the new data is bounded.

Introduction

With the Web Ontology Language (OWL) and its latest version OWL 2 being standardised by the WWW Consortium (W3C), a large number of professional and comprehensive ontologies have been developed for data modelling and access. Existing large ontologies like NCI (Hartel et al. 2005) are mostly hand-crafted by human experts. The development and maintenance of such large ontologies are often time-consuming and error-prone. Recently, automated ontology learning (Baader et al. 2007; Lehmann and Hitzler 2010; Konev, Lutz, and Wolter 2013) and evolution (Qi and Yang 2008; Qi et al. 2008; Ribeiro and Wassermann 2009; Qi and Du 2009; Wang, Wang, and Topor 2010; Zheleznyakov et al. 2010; Cuenca Grau et al. 2012; Qi et al. 2014) have attracted intensive interest.

Although it is often conceived that ontologies, once established, are more stable and reliable than data, it is shown in (Pesquita and Couto 2012) that automated ontology enrichment and refinement are desired in applications such as biomedical ontologies, which are constant evolving due to new understandings of the domain. One typical approach to the automation of the process is ontology learn-

ing (Lehmann and Hitzler 2010), where terminological relations are learnt from example data provided by domain experts. The development of biomedicine ontologies is described as an evolution process (Pesquita and Couto 2012)—as new knowledge is discovered daily, the initial ontology needs to be enriched with or refined by the new knowledge continuously. Hence, automated learning and incremental revision techniques can largely benefit the ontology engineering. Several ontology learning approaches have been proposed in the literature and effective learning tools have been developed (Lehmann and Hitzler 2010; Ma and Distel 2013). An obvious requirement in ontology learning is that the learnt ontology must be consistent with the example data. However, when new data instances become available, it is possible that the new data contradicts some previously learnt axioms, which indicates that mistakes were made in previous learning processes. Thus, a challenge is how to revise the existing ontology according to the new data. In DLs, it amounts to the question of how to effectively revise a TBox w.r.t. an ABox.

TBox debugging (Kalyanpur et al. 2006; Schlobach et al. 2007) and interactive revision approaches (Nikitina, Rudolph, and Glimm 2012) can pinpoint the problematic axioms, and most of TBox repair approaches seek to eliminate a minimal number of axioms to restore consistency. Similar syntax-based approaches were adopted in revision (Haase and Stojanovic 2005; Qi et al. 2008; Ribeiro and Wassermann 2009). However, most syntax-based approaches to TBox change lack of a suitable semantic justification, which is partially reflected in their inability of preserving implicit knowledge. For example, given $\mathcal{T} = \{\text{Bird} \sqsubseteq \text{CanFly}, \text{Mammal} \sqsubseteq \neg\text{CanFly}\}$, the revision of \mathcal{T} by $\mathcal{A} = \{\text{CanFly}(\text{bat}), \text{Mammal}(\text{bat})\}$ using syntax-based approaches in (Haase and Stojanovic 2005; Ribeiro and Wassermann 2009) will discard implicit information $\text{Bird} \sqsubseteq \neg\text{Mammal}$, though it is not involved in the contradiction. To resolve this issue, some more recent works (Horridge, Parsia, and Sattler 2008; Calvanese et al. 2010; Cuenca Grau et al. 2012) advocate to apply syntax-based TBox change on the deductive closure $cl(\mathcal{T})$ of the initial TBox \mathcal{T} instead of \mathcal{T} itself. However, there may not exist a unique optimal solution to such change (Calvanese et al. 2010), since often multiple minimal sets of axioms in \mathcal{T} (or $cl(\mathcal{T})$) exist that are responsible for the inconsistency.

Seeing the TBox as a formalisation of the conceptual model of the ABox, which is a common perspective in ontology learning, the TBox essentially represents a collection of possible DL models. Taking this perspective, one can argue that in TBox evolution, it is critical to ensure minimal changes to the prospective models. Hence, we adopt the classical model-based belief change approach in defining our TBox change operator, with the aim of fulfilling the minimal change principle in a model-theoretic manner. In classical belief change, the operation of eliminating problematic axioms from the logical consequences of a knowledge base is formalised as *contraction*, and that of consistently incorporating newly formed axioms is formalised as *revision*. Extant model-based TBox change literature only cover the revision and/or contraction w.r.t. TBox axioms (Qi and Du 2009; Zheleznyakov et al. 2010; Zhuang et al. 2014), with a few exceptions on revision w.r.t. a combination of a TBox and an ABox (Wang, Wang, and Topor 2010).

In this paper, we develop a novel approach to TBox revision w.r.t. an ABox through TBox contraction, that is, by first contracting the “negation” of the ABox and then incorporating the ABox. To the best of our knowledge, it is the first work in this direction, and this is challenging due to the fact that negation of an ABox was previously undefined. Moreover, as witnessed by previous model-based TBox revision approaches, working directly with classical DL models is difficult. Thus, we adopt the type semantics for DL-Lite, a family of logics underlying the OWL 2 QL profile. We define the negation of an ABox as a DL-Lite TBox, and provide a concrete TBox contraction operator based on type semantics. We show that our operator possesses desired properties. In particular, our operator is *fine-grained* in that it supports modification of axioms, and the result of contraction is uniquely defined and thus solves the problem of non-determinism. For the computational aspects, we transform the entailment problem of TBox contraction in DL-Lite into a corresponding problem in propositional logic. We show that entailment checking in DL-Lite does not shift the corresponding complexity in propositional logic. We also provide a tractable algorithm for contraction.

Preliminaries

DL-Lite_{horn}

A *signature* is a union of three disjoint (possibly infinite) sets N_C , N_R , and N_I , where N_C is the set of atomic concepts, N_R is the set of atomic roles, and N_I is the set of individuals. Concepts and roles in DL-Lite_{horn} are built on atomic ones as follows (Calvanese et al. 2007; Artale et al. 2009):

$$B \rightarrow A \mid \exists R \quad C \rightarrow B \mid \neg B \mid \perp \quad R \rightarrow P \mid P^-$$

where $A \in N_C$ and $P \in N_R$. B is a *basic concept*, C is a *general concept*, and R is a *basic role*. In what follows, we often use A for an atomic concept, B for a basic concept, C for a general concept, P for an atomic role, and R for a basic role. \mathcal{B} and \mathcal{R} denote the sets of basic concepts and basic roles, respectively. We write R^- for P if $R = P^-$. A DL-Lite_{horn} *concept* is of the form $\prod_k B_k$ or C . An *axiom* in DL-Lite_{horn} is of the form $\prod_k B_k \sqsubseteq C$, whereas

in DL-Lite_{core} it is of the form $B \sqsubseteq C$. A DL-Lite_{horn} (DL-Lite_{core}) *TBox* is a finite set of axioms in the logic. An *ABox* is a finite set of *assertions* of the form $A(a)$ or $P(a, b)$, where $a, b \in N_I$. For the convenience of presentation, we sometimes write $P^-(a, b) \in \mathcal{A}$ meaning $P(b, a) \in \mathcal{A}$. A knowledge base (KB) $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ consists of a TBox \mathcal{T} and an ABox \mathcal{A} . In this paper, an ontology is a DL-Lite TBox.

The semantics of DL-Lite_{horn} is defined as usual using interpretations $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, and we refer to (Calvanese et al. 2007) for the details. \mathcal{I} is a *model* of \mathcal{T} (\mathcal{A} , or $\mathcal{T} \cup \mathcal{A}$) if \mathcal{I} satisfies \mathcal{T} (resp., \mathcal{A} , or both \mathcal{T} and \mathcal{A}). $\mathcal{T} \cup \mathcal{A}$ is *consistent* if it has at least one model. We also say that \mathcal{T} is consistent with \mathcal{A} . A concept or role E is *satisfiable* w.r.t. \mathcal{T} if a model \mathcal{I} of \mathcal{T} exists such that $E^{\mathcal{I}} \neq \emptyset$; otherwise, E is *unsatisfiable*. \mathcal{T} is *coherent* if all atomic concepts and all atomic roles in \mathcal{T} are satisfiable. \mathcal{T} *entails* an axiom α , written $\mathcal{T} \models \alpha$, if all models of $\mathcal{T} \cup \mathcal{A}$ satisfy α . $cl(\mathcal{T})$ denotes the set of axioms entailed by \mathcal{T} . Two TBoxes $\mathcal{T}_1, \mathcal{T}_2$ (two ABoxes $\mathcal{A}_1, \mathcal{A}_2$) are *equivalent*, written $\mathcal{T}_1 \equiv \mathcal{T}_2$ ($\mathcal{A}_1 \equiv \mathcal{A}_2$), if they have the same models.

Type Semantics

To develop a model-theoretic approach to TBox change in DL-Lite, we note that DL models may have complex structures and the number of models is always infinite. This makes it hard to handle TBox change directly through classical DL models. Hence, we use the *type semantics* introduced for DL-Lite TBoxes in (Kontchakov, Wolter, and Zakharyashev 2010; Zhuang et al. 2014), which provides an alternative semantic characterisation using structures that are finite and much simpler than classical DL models. We first introduce the definition and then shows that the nice properties of the type semantics extend to DL-Lite_{horn}. In TBox change, it suffices to consider a finite signature that covers the initial TBox and the new ABox. In this case, \mathcal{B} and \mathcal{R} are both finite.

Formally, a *type* $\tau \subseteq \mathcal{B}$ is a (possibly empty) set of basic concepts. For a DL-Lite_{horn} TBox \mathcal{T} (or concept E), $PM(\mathcal{T})$ ($PM(E)$) denotes the set of types corresponding to the propositional models of \mathcal{T} (E), i.e., by seeing a type τ as a propositional interpretation (with basic concepts seen as propositional atoms) and \mathcal{T} (or E) as a propositional theory.

Definition 1 (Zhuang et al. 2014) A type τ satisfies a DL-Lite_{core} TBox \mathcal{T} if both of the following conditions hold:

1. $\tau \in PM(\mathcal{T})$, and
2. if $\mathcal{T} \models \exists R \sqsubseteq \perp$ then $\exists R^- \notin \tau$.

In this case, τ is a type model (T-model) of \mathcal{T} . The set of T-models of \mathcal{T} is denoted as $TM(\mathcal{T})$.

For example, let $\mathcal{T} = \{\exists P \sqsubseteq \perp\}$ and $\tau = \{\exists P^-\}$. Then, $\tau \in PM(\mathcal{T})$ and $\tau \notin TM(\mathcal{T})$. For brevity, we write $TM(\alpha)$ for singleton TBox $\mathcal{T} = \{\alpha\}$.

The definition of T-models extends directly to DL-Lite_{horn}, and type semantics characterise the semantics of DL-Lite_{horn} TBoxes.

Proposition 1 Given TBoxes $\mathcal{T}, \mathcal{T}'$, a concept E , and an axiom α in DL-Lite_{horn}, the following results hold:

1. \mathcal{T} is consistent iff there exists a T-model of \mathcal{T} .
2. E is satisfiable w.r.t. \mathcal{T} iff $TM(\mathcal{T}) \cap PM(E) \neq \emptyset$.
3. $\mathcal{T} \models \alpha$ iff $TM(\mathcal{T}) \subseteq TM(\alpha)$.
4. $\mathcal{T} \equiv \mathcal{T}'$ iff $TM(\mathcal{T}) = TM(\mathcal{T}')$.

An arbitrary set M of types may not be always expressible in DL-Lite_{horn}, that is, M may not be the set of T-models of any DL-Lite_{horn} TBox. Let $\mathcal{L} \in \{\text{DL-Lite}_{core}, \text{DL-Lite}_{horn}\}$, a \mathcal{L} -TBox \mathcal{T} corresponds to M in \mathcal{L} if $M \subseteq TM(\mathcal{T})$ and there is no \mathcal{L} -TBox \mathcal{T}' such that $M \subseteq TM(\mathcal{T}') \subset TM(\mathcal{T})$. In general, there may be multiple TBoxes that correspond to M in \mathcal{L} . The following proposition shows a sufficient condition for the uniqueness.

Proposition 2 *Let $\mathcal{L} \in \{\text{DL-Lite}_{core}, \text{DL-Lite}_{horn}\}$ and M be a set of types. A unique TBox exists that corresponds to M in \mathcal{L} (up to semantic equivalence) if M includes a type containing $\exists R$ whenever it includes a type containing $\exists R^-$ for all $R \in \mathcal{R}$.*

When a unique corresponding TBox exists, $\mathcal{T}_{\mathcal{L}}(M)$ denotes the TBox corresponding to M in \mathcal{L} . $\mathcal{T}_{\mathcal{L}}(M)$ has the same set of logical consequences in \mathcal{L} as M .

Corollary 1 *Let $\mathcal{L} \in \{\text{DL-Lite}_{core}, \text{DL-Lite}_{horn}\}$ and M be a set of types. If $\mathcal{T}_{\mathcal{L}}(M)$ exists then for each axiom α in \mathcal{L} , $M \subseteq TM(\alpha)$ iff $\mathcal{T}_{\mathcal{L}}(M) \models \alpha$.*

Expressing Negation of ABoxes

When a set \mathcal{A} of new assertions contradict the existing TBox \mathcal{T} , we want to change \mathcal{T} to a new TBox \mathcal{T}' , such that \mathcal{T}' is consistent with \mathcal{A} and that \mathcal{T}' preserves as much terminological knowledge as possible from \mathcal{T} . In classical belief change, revision of a propositional knowledge base K by a formula ϕ can be achieved through contraction via the *Levi identity* (Levi 1991). In particular, let $\dot{-}$ be a contraction operator, the revision of K by ϕ can be defined as $(K \dot{-} \neg\phi) \cup \{\phi\}$. However, DL-Lite (and other common DLs) does not permit negation of ABoxes. Although negation of assertions are allowed in DL-Lite, the negation of an ABox would require expressing the disjunction of negated assertions and would be an unnatural extension to DL-Lite. For example, the ABox $\{A(a), A(b)\}$ corresponds to first-order formula $A(a) \wedge A(b)$, whose negation is $\neg A(a) \vee \neg A(b)$. Even if disjunction is allowed, it remains unclear how to define contraction of assertions from a TBox. On the other hand, it is natural to contract TBox axioms (that are conflicting with the new ABox) from the TBox before adding the new ABox. To define such TBox contraction, we first show how to define the negation of an ABox as a DL-Lite_{horn} TBox.

As suggested in (Flouris et al. 2006), negation can be defined through inconsistency. That is, we define the negation of an ABox \mathcal{A} , denoted $\mathcal{T}_-(\mathcal{A})$, to be a TBox that is inconsistent with \mathcal{A} . Yet different from (Flouris et al. 2006), our negation is over sets of assertions instead of single axioms. Also, the following conditions should be satisfied:

- (i) For each TBox \mathcal{T} , the result of contracting $\mathcal{T}_-(\mathcal{A})$ from \mathcal{T} must be consistent with \mathcal{A} ; that is informally, $\mathcal{T}_-(\mathcal{A})$ contains all potential contradictions towards \mathcal{A} that would need to be removed from TBoxes.

- (ii) $\mathcal{T}_-(\mathcal{A})$ should represent potential contradictions without any redundancy.

To formalise “the contradictions” to be removed from TBoxes, we adapt the notion of (*fine-grained*) justifications (Horridge, Parsia, and Sattler 2008) for inconsistency.

Definition 2 *For a TBox \mathcal{T} that is inconsistent with an ABox \mathcal{A} , a TBox justification for the inconsistency is a TBox $\mathcal{T}' \subseteq cl(\mathcal{T})$ such that \mathcal{T}' is inconsistent with \mathcal{A} and no strict subset of \mathcal{T}' is inconsistent with \mathcal{A} .*

Intuitively, a TBox justification is a minimal set of logical consequences of the TBox that contradicts the ABox. A TBox justification contains only TBox axioms from $cl(\mathcal{T})$ that are involved in a contradiction, and can be seen as a projection of a fine-grained justification on $cl(\mathcal{T})$ ¹. The following definition resembles the notion of *hitting sets* of justifications but quantifies over all TBoxes.

Definition 3 *For an ABox \mathcal{A} , a TBox is a negation of \mathcal{A} , denoted $\mathcal{T}_-(\mathcal{A})$, if*

1. for each TBox that is inconsistent with \mathcal{A} , there exists a TBox justification \mathcal{T}' for the inconsistency such that $\mathcal{T}_-(\mathcal{A}) \cap \mathcal{T}' \neq \emptyset$; and
2. any strict subset of $\mathcal{T}_-(\mathcal{A})$ does not satisfy condition 1.

Condition 1 in the definition corresponds to the pre-specified condition (i), and condition 2 is the minimality condition corresponding to (ii). For example, let $\mathcal{A} = \{A(a), B(a)\}$, then $\mathcal{T} = \{A \sqsubseteq \neg B\}$ is a negation of \mathcal{A} , as informally any TBox that contradicts \mathcal{A} implies A and B are disjoint (including the cases where A or B is empty). If the disjointness is contracted from a given TBox, the TBox will be consistent with \mathcal{A} . On the other hand, $\mathcal{T}' = \{A \sqsubseteq \perp, B \sqsubseteq \perp\}$ is not a negation of \mathcal{A} , as it does not intersect with a TBox justification for the inconsistency between \mathcal{T} and \mathcal{A} . Intuitively, contracting \mathcal{T}' makes both A and B satisfiable but cannot remove disjointness.

In what follows, we show that a negation of an ABox always exists in DL-Lite_{horn}, which is unique (up to semantic equivalence) and can be computed efficiently. In particular, although a negation of an ABox is defined via TBox justifications, it can be computed directly without first computing justifications. For an ABox \mathcal{A} and an individual a in \mathcal{A} , define the *type* of a in \mathcal{A} as $\tau(a, \mathcal{A}) = \{A \mid A(a) \in \mathcal{A}\} \cup \{\exists R \mid R(a, b) \in \mathcal{A}, b \in N_I\}$. Let $\Gamma(\mathcal{A}) = \max_{\subseteq}(\{\tau(a, \mathcal{A}) \mid a \in N_I \text{ in } \mathcal{A}\})$ be the set of maximal types (w.r.t. set containment) of individuals in \mathcal{A} . Before showing the construction of ABox negation in DL-Lite_{horn}, we first show how the consistency between a TBox and an ABox can be characterised by type semantics.

Lemma 1 *Given a TBox \mathcal{T} and an ABox \mathcal{A} in DL-Lite_{horn}, $\mathcal{T} \cup \mathcal{A}$ is consistent iff for each type $\tau \in \Gamma(\mathcal{A})$, a T-model $\tau' \in TM(\mathcal{T})$ exists such that $\tau \subseteq \tau'$.*

Based on the above characterisation, we can establish the following construction of ABox negation in DL-Lite_{horn}.

¹Note that a fine-grained justification is also required to be minimal in size and be logically weakest (Horridge, Parsia, and Sattler 2008), which are not enforced in our definition.

Proposition 3 *Let \mathcal{A} be a DL-Lite_{horn} ABox. Then a unique negation of \mathcal{A} exists (up to semantic equivalence) $\mathcal{T}_-(\mathcal{A}) \equiv \{\bigwedge_{B \in \tau} B \sqsubseteq \perp \mid \tau \in \Gamma(\mathcal{A})\}$.*

Proof We first show that $\mathcal{T}_-(\mathcal{A})$ satisfies the two conditions in Definition 3. For each TBox \mathcal{T} that is inconsistent with \mathcal{A} , by Lemma 1, there exists a type $\tau \in \Gamma(\mathcal{A})$ such that no T-model $\tau' \in TM(\mathcal{T})$ exists such that $\tau \subseteq \tau'$. Let α be the axiom $\bigwedge_{B \in \tau} B \sqsubseteq \perp$. Then, it is clear that $TM(\mathcal{T}) \subseteq TM(\alpha)$. From Proposition 1, $\mathcal{T} \models \alpha$. Also, from the definition of $\Gamma(\mathcal{A})$, there is an individual a occurring in \mathcal{A} such that $\tau(a, \mathcal{A}) = \tau$. Then, $\{\alpha\} \cup \mathcal{A}$ is inconsistent. Hence, $\{\alpha\}$ is a TBox justification for \mathcal{T} . By the construction of $\mathcal{T}_-(\mathcal{A})$, $\alpha \in \mathcal{T}_-(\mathcal{A})$. That is, $\mathcal{T}_-(\mathcal{A})$ satisfies condition 1 in Definition 3. For condition 2, consider an arbitrary axiom $\alpha \in \mathcal{T}_-(\mathcal{A})$ and the TBox $\mathcal{T} = \{\alpha\}$, then \mathcal{T} is inconsistent with \mathcal{A} . It is not hard to see that $\mathcal{T}_-(\mathcal{A}) \setminus \{\alpha\}$ does not contain any axiom from $cl(\mathcal{T})$. That is, $\mathcal{T}_-(\mathcal{A}) \setminus \{\alpha\}$ does not satisfy condition 1, and hence $\mathcal{T}_-(\mathcal{A})$ is minimal.

We have shown that $\mathcal{T}_-(\mathcal{A})$ is a negation of \mathcal{A} , and now we show that a negation of \mathcal{A} must be equivalent to $\mathcal{T}_-(\mathcal{A})$. Let \mathcal{T} be a negation of \mathcal{A} . For each axiom $\alpha \in \mathcal{T}_-(\mathcal{A})$, consider TBox $\{\alpha\}$ that is inconsistent with \mathcal{A} . By condition 1, \mathcal{T} must intersect with a TBox justification for $\{\alpha\}$. From the deductive closure of α and Lemma 1, it is not hard to see that such a TBox justification is a singleton TBox equivalent to $\{\alpha\}$. Thus, \mathcal{T} contains axiom equivalent to α . By the minimality condition of \mathcal{T} , \mathcal{T} is equivalent to $\mathcal{T}_-(\mathcal{A})$. ■

In the following sections, we use $\mathcal{T}_-(\mathcal{A})$ to denote the TBox in Proposition 3. It is clear that $\mathcal{T}_-(\mathcal{A})$ can be computed in linear time to the size of \mathcal{A} . The following result shows that $\mathcal{T}_-(\mathcal{A})$ consists of exactly the axioms to be contracted in order to consistently incorporate \mathcal{A} .

Proposition 4 *Given a TBox \mathcal{T} and an ABox \mathcal{A} in DL-Lite_{horn}, $\mathcal{T} \cup \mathcal{A}$ is consistent iff $\mathcal{T} \not\models \alpha$ for each $\alpha \in \mathcal{T}_-(\mathcal{A})$.*

TBox Contraction

In order to incorporate an ABox \mathcal{A} into a TBox \mathcal{T} in a consistent manner, each axiom in $\mathcal{T}_-(\mathcal{A})$ needs to be contracted from \mathcal{T} . We adapt the TBox contraction approach in (Zhuang et al. 2014), and characterise minimal change from a model-theoretic perspective, that is, the T-models of the result of contraction can be obtained from the T-models of the initial TBox with minimal change. In contrast to (Zhuang et al. 2014), which introduces a generic approach for contracting single and conjunctions of axioms, we define a concrete TBox contraction operator for contracting a set of axioms. In classical belief contraction literature, *multiple contraction* (Fuhrmann and Hansson 1993), that is contracting a set S of formulas, is shown to have distinct nature from contracting the conjunction of the formulas in S and from iterated contraction of each formula in S —contracting the conjunction is insufficient for contracting all the conjuncts, and iterated contraction is sensitive to the order of contraction. Also, to show how our contraction operator can resolve incoherence in TBoxes, we relax the coherence assumption in (Zhuang et al. 2014) and assume the initial TBox is possibly incoherent.

A classical model-based contraction operator contracts a propositional formula ϕ from knowledge base K by (i) extending the models of K with counter-models of ϕ that have minimal distances to the models of K , and (ii) defining the knowledge base corresponding to the extended set of models to be the result of contraction. To contract a set S of formulas, one needs to extend the models of K with counter-models of each formula in S . Inspired by classical belief change, we first introduce a notion of distance between T-models, and then define a selection function that selects from a set of types those having minimal distances to the T-models of the initial TBox. We adopt Satoh's distance (Satoh 1988), that is the symmetric difference between two (propositional) models. The *distance* between two types τ and τ' is their symmetric difference $sd(\tau, \tau') = (\tau \setminus \tau') \cup (\tau' \setminus \tau)$. Intuitively, a pair of types τ_1, τ'_1 are considered to be (strictly) closer than another pair τ_2, τ'_2 if $sd(\tau_1, \tau'_1) \subset sd(\tau_2, \tau'_2)$. While we adopt Satoh's distance, we note that other notions of distance can also be applied in our contraction framework. Note that while contraction has been studied for (propositional) Horn logic (Zhuang and Pagnucco 2012; Delgrande and Peppas 2015), none of these works focused on defining a concrete operator. Further, given two sets M and M' of types, a selection function γ selects the types in M that are closest to those in M' . In particular,

$$\gamma(M, M') = \{\tau_0 \in M \mid \tau'_0 \in M' \text{ exists s.t.} \\ \text{for all } \tau \in M \text{ and } \tau' \in M', sd(\tau, \tau') \not\subset sd(\tau_0, \tau'_0)\}$$

Let Ω be the set of all types (which is finite as \mathcal{B} is finite). For an axiom α , $TM(\alpha) = \Omega \setminus TM(\alpha)$ is the set of counter-models of α . For a TBox \mathcal{T} and an ABox \mathcal{A} , to contract $\mathcal{T}_-(\mathcal{A})$ from \mathcal{T} , we could simply add $\gamma(TM(\alpha), TM(\mathcal{T}))$ to $TM(\mathcal{T})$ for each axiom $\alpha \in \mathcal{T}_-(\mathcal{A})$. However, a result of contraction defined this way may not be coherent. A simple case is when \mathcal{T} is incoherent and \mathcal{A} is empty. In general, suppose concept A is unsatisfiable in \mathcal{T} , i.e., $\mathcal{T} \models A \sqsubseteq \perp$, if $A(a) \in \mathcal{A}$ for some a then contracting $\mathcal{T}_-(\mathcal{A})$ is sufficient to make A satisfiable (as otherwise the result of contraction is inconsistent with \mathcal{A}). Otherwise, to make A satisfiable it suffices to contract additionally $A \sqsubseteq \perp$ from \mathcal{T} . Similarly, to resolve an unsatisfiable role P , one only needs to contract $\exists P \sqsubseteq \perp$. Based on this intuition, we introduce the *coherence closure* $\mathcal{T}_-^*(\mathcal{A})$ of $\mathcal{T}_-(\mathcal{A})$, which extends $\mathcal{T}_-(\mathcal{A})$ with $A \sqsubseteq \perp$ for each $A \in N_C$ not occurring in \mathcal{A} and with $\exists P \sqsubseteq \perp$ and $\exists P^- \sqsubseteq \perp$ for each $P \in N_R$ not occurring in \mathcal{A} . Note that if A (or P) is satisfiable in \mathcal{T} , contracting $A \sqsubseteq \perp$ ($\exists P \sqsubseteq \perp$) has no effect, as we will see later. Also, $\exists P^- \sqsubseteq \perp$ is added to guarantee uniqueness of contraction.

We define our TBox contraction as follows.

Definition 4 *Let $\mathcal{L} \in \{DL\text{-Lite}_{core}, DL\text{-Lite}_{horn}\}$, \mathcal{T} be a TBox and \mathcal{A} be an ABox in \mathcal{L} . The (coherent) contraction of $\mathcal{T}_-(\mathcal{A})$ from \mathcal{T} is $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A}) = \mathcal{T}_{\mathcal{L}}(TM(\mathcal{T}) \cup M_{\mathcal{A}}^*)$ where*

$$M_{\mathcal{A}}^* = \bigcup_{\alpha \in \mathcal{T}_-^*(\mathcal{A})} \gamma(TM(\alpha), TM(\mathcal{T})).$$

Intuitively, $M_{\mathcal{A}}^*$ is constituted by, for each axiom α in $\mathcal{T}_-^*(\mathcal{A})$, a set of counter-models of α that are closest to the

T-models of \mathcal{T} —it is constructed in a way that the result of contraction is coherent and is consistent with \mathcal{A} ; and the result of contraction is defined to be the TBox corresponding to $TM(\mathcal{T}) \cup M_{\mathcal{A}}^*$. $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})$ is well defined and unique (up to semantic equivalence), as $TM(\mathcal{T}) \cup M_{\mathcal{A}}^*$ satisfies the condition in Proposition 2: For each $P \in N_R$, either $P(b, c)$ occurs in \mathcal{A} for some b, c or $\{\exists P \sqsubseteq \perp, \exists P^- \sqsubseteq \perp\} \subseteq \mathcal{T}_-(\mathcal{A})$, and as a result, a type containing $\exists P$ and a type containing $\exists P^-$ are contained in $M_{\mathcal{A}}^*$.

Consider the following example adapted from the NCI ontology (Hartel et al. 2005).

Example 1 Let TBox \mathcal{T} consists of the following axioms

$$\text{Heart_Disease} \sqsubseteq \exists \text{has_Site}, \quad (1)$$

$$\exists \text{has_Site}^- \sqsubseteq \text{Cardiovascular_System}, \quad (2)$$

$$\exists \text{has_Site}^- \sqsubseteq \text{Respiratory_System}, \quad (3)$$

$$\text{Respiratory_System} \sqsubseteq \text{Organ_System}, \quad (4)$$

$$\text{Cardiovascular_System} \sqsubseteq \text{Organ_System}, \quad (5)$$

$$\text{Respiratory_System} \sqsubseteq \neg \text{Cardiovascular_System}, \quad (6)$$

$$\text{Heart_Disease} \sqsubseteq \neg \text{Organ_System}. \quad (7)$$

The TBox \mathcal{T} is incoherent, as both concept `Heart_Disease` and role `has_Site` are unsatisfiable.

Let ABox $\mathcal{A} = \{\text{Heart_Disease}(\text{hd1}), \text{has_Site}(\text{hd1}, \text{s1}), \text{Organ_System}(\text{s1})\}$, by Definition 4, $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})$ consists of axioms (1), (4)–(7), and the following axiom

$$\exists \text{has_Site}^- \sqsubseteq \text{Organ_System}, \quad (8)$$

which can be seen as a revision from axioms (2)–(5).

The example shows that our TBox contraction operator allows modification of axioms, in contrast to most existing TBox repair approaches that only support axiom removal.

Furthermore, the contraction operator satisfies the following desired properties presented as AGM-style postulates (Katsuno and Mendelzon 1991), specially tailored to TBox contraction w.r.t. ABoxes.

Proposition 5 Let $\mathcal{L} \in \{\text{DL-Lite}_{\text{core}}, \text{DL-Lite}_{\text{horn}}\}$, \mathcal{T} be a TBox and \mathcal{A} an ABox in \mathcal{L} . The following properties hold.

(C1) $\mathcal{T} \models \mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})$.

(C2) $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A}) \equiv \mathcal{T}$ if \mathcal{T} is coherent and $\mathcal{T} \cup \mathcal{A}$ is consistent.

(C3) $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A}) \not\models \alpha$ for each $\alpha \in \mathcal{T}_-(\mathcal{A})$.

(C4) If $\mathcal{T}_1 \equiv \mathcal{T}_2$ and $\mathcal{A}_1 \equiv \mathcal{A}_2$ then $\mathcal{T}_1 \dot{-} \mathcal{T}_-(\mathcal{A}_1) \equiv \mathcal{T}_2 \dot{-} \mathcal{T}_-(\mathcal{A}_2)$.

(C5) $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})$ is coherent.

(C6) $(\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})) \cup \mathcal{A}$ is consistent.

Proof (Sketch) (C1) and (C3)–(C6) are not hard to see, and we only show the proof for (C2). For (C2), we only need to show that in this case $M_{\mathcal{A}}^* \subseteq TM(\mathcal{T})$. For each axiom α , if α is in $\mathcal{T}_-(\mathcal{A})$ then it is of the form $\prod_k B_k \sqsubseteq \perp$. By the construction of $\mathcal{T}_-(\mathcal{A})$, there is an individual a occurring in \mathcal{A} such that $\tau(a, \mathcal{A}) = \{B_k\}$. Since \mathcal{T} is consistent with \mathcal{A} , by Lemma 1, there exists a T-model $\tau \in TM(\mathcal{T})$ such that $\tau(a, \mathcal{A}) \subseteq \tau$, that is, $\{B_k\} \subseteq \tau$. Hence, $\tau \in \overline{TM(\alpha)}$ and thus $TM(\mathcal{T}) \cap \overline{TM(\alpha)} \neq \emptyset$. Thus, $\gamma(\overline{TM(\alpha)}, TM(\mathcal{T})) =$

$TM(\mathcal{T}) \cap \overline{TM(\alpha)} \subseteq TM(\mathcal{T})$. If α is not in $\mathcal{T}_-(\mathcal{A})$ then it is of the form $B \sqsubseteq \perp$ added for the coherence closure. Since \mathcal{T} is coherent, $\mathcal{T} \not\models B \sqsubseteq \perp$. By Proposition 1, there is at least one T-model of \mathcal{T} containing B . Let M_B be the set of T-models of \mathcal{T} that contain B . Then, $\gamma(\overline{TM(\alpha)}, TM(\mathcal{T})) = M_B \subseteq TM(\mathcal{T})$. We have shown that $M_{\mathcal{A}}^* \subseteq TM(\mathcal{T})$. ■

Properties (C1)–(C4) are adapted from the first four AGM postulates for belief contraction, whereas a meaningful adaptation of the fifth (namely the *recovery*) postulate is missing in our scenario. (C1) states that the contraction is a weakening of the initial TBox, and (C2) says that no change is needed if the initial TBox is already coherent and is consistent with the new ABox. Then, (C3) corresponds to the *success* postulate for multiple contraction (note that $\mathcal{T}_-(\mathcal{A})$ does not contain a tautology axiom), and (C4) shows that the contraction operator is syntax-independent, *i.e.*, contracting the negations of equivalent ABoxes from equivalent TBoxes have equivalent results. Finally, the last two properties (C5) and (C6) connect contraction to revision, showing that the result of contraction is coherent and is consistent with the new ABox as required.

As an additional evidence for the rationality of our TBox contraction, we show a connection to the KB revision approach introduced in (Wang, Wang, and Topor 2010), where a more complex (but finite) structure is used as the model-theoretic characterisation of revision. Let $*_f$ be the second revision operator (namely *f-revision*) in (Wang, Wang, and Topor 2010).

Proposition 6 Let $\mathcal{L} \in \{\text{DL-Lite}_{\text{core}}, \text{DL-Lite}_{\text{horn}}\}$, \mathcal{T} be a coherent TBox and \mathcal{A} be an ABox in \mathcal{L} . If $\mathcal{T} *_f \mathcal{A}$ can be expressed in \mathcal{L} then $\mathcal{T} *_f \mathcal{A} \equiv (\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})) \cup \mathcal{A}$.

Computational Aspects

In this section, we first look into the computational complexity of the *entailment problem* for contraction, *i.e.*, the problem of deciding whether $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})$ entails a given TBox axiom. We show that the entailment can be checked without first computing $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})$, and is achieved through a reduction to propositional belief revision. When the size of \mathcal{A} is bounded, the reasoning can be done in polynomial time. After that, we will provide a tractable algorithm to compute $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})$ in $\text{DL-Lite}_{\text{core}}$ through entailment checking.

Firstly, the entailment problem in DL-Lite can be reduced to a corresponding one in propositional logic. To this end, we assign each basic concept in \mathcal{B} a distinct propositional atom. Function $\phi(\cdot)$ maps each basic concept in \mathcal{B} to its corresponding propositional atom, and a concept description or an axiom to a propositional formula as follows.

$$\phi(\perp) = \perp, \quad \phi(\neg B) = \neg\phi(B), \quad \phi\left(\prod_k B_k\right) = \bigwedge_k \phi(B_k)$$

$$\phi\left(\prod_k B_k \sqsubseteq C\right) = \phi\left(\prod_k B_k\right) \rightarrow \phi(C).$$

Then, $\phi(\cdot)$ can be extended to map a $\text{DL-Lite}_{\text{horn}}$ TBox \mathcal{T} to a propositional horn formula as follows

$$\phi(\mathcal{T}) = \bigwedge_{\alpha \in \mathcal{T}} \phi(\alpha) \wedge \bigwedge_{R \in \mathcal{R}, \mathcal{T} \models \exists R \sqsubseteq \perp} (\neg\phi(\exists R) \wedge \neg\phi(\exists R^-)).$$

Further, to encode the counter-models of axioms in $\mathcal{T}_-^*(\mathcal{A})$, note that each axiom in $\mathcal{T}_-^*(\mathcal{A})$ is of the form $\bigwedge_k B_k \sqsubseteq \perp$. Intuitively, a counter model of the axiom must (and only needs to) satisfy each B_k . Formally, for each axiom β of the above form, define $\phi_-(\beta) = \bigwedge_k \phi(B_k)$.

Let $*_s$ be Satoh's revision operator in propositional logic. The following connection between our contraction operator and Satoh's revision holds.

Proposition 7 *Let $\mathcal{L} \in \{DL\text{-Lite}_{core}, DL\text{-Lite}_{horn}\}$, \mathcal{T} be a TBox and \mathcal{A} be an ABox in \mathcal{L} . Given a TBox axiom α in \mathcal{L} , $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A}) \models \alpha$ iff $\phi(\mathcal{T}) \vee \bigvee_{\beta \in \mathcal{T}_-^*(\mathcal{A})} (\phi(\mathcal{T}) *_s \phi_-(\beta)) \models \phi(\alpha)$.*

Proof (Sketch) First, we can show that the models of $\phi(\mathcal{T})$ correspond exactly to the T-models of \mathcal{T} , and the models of $\phi_-(\beta)$ to the types in $\overline{TM}(\beta)$. Let $\omega_\tau = \{\phi(B) \mid B \in \tau\}$ for a type τ . $\text{mod}(\phi)$ denotes the set of models of formula ϕ . Then, from the correspondence between our selection function and that of Satoh's, it is clear that

$$\text{mod}(\phi(\mathcal{T}) *_s \phi_-(\beta)) = \{\omega_\tau \mid \tau \in \gamma(\overline{TM}(\beta), TM(\mathcal{T}))\}.$$

Let $\varphi = \bigvee_{\beta \in \mathcal{T}_-^*(\mathcal{A})} (\phi(\mathcal{T}) *_s \phi_-(\beta))$. Then, $\text{mod}(\varphi) = \{\omega_\tau \mid \tau \in M_{\mathcal{A}}^*\}$. If α is of the form $\exists R \sqsubseteq \perp$ then neither entailment holds, as from the construction of $M_{\mathcal{A}}^*$, there must exist a type in $M_{\mathcal{A}}^*$ containing $\exists R$; and similarly, there exists a model of φ containing $\phi(\exists R)$. If α is not of the form $\exists R \sqsubseteq \perp$, by Corollary 1, $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A}) \models \alpha$ iff each type in $TM(\mathcal{T})$ and each type in $M_{\mathcal{A}}^*$ satisfies α , iff each model of $\phi(\mathcal{T}) \vee \varphi$ satisfies $\phi(\alpha)$. ■

Proposition 7 reduces the entailment problem of TBox contraction in DL-Lite to that of propositional belief revision, and in particular, to checking $\phi(\mathcal{T}) \models \phi(\alpha)$ and $\phi(\mathcal{T}) *_s \phi_-(\beta) \models \phi(\alpha)$ for each $\beta \in \mathcal{T}_-^*(\mathcal{A})$. Note that $\phi(\mathcal{T})$, $\phi(\alpha)$ and $\phi_-(\beta)$ are all Horn propositional formulas and in sizes polynomial to the size of $\mathcal{T} \cup \mathcal{A}$. From the complexity results in (Eiter and Gottlob 1992), we can conclude the following results, showing that our contraction does not shift the complexity of the propositional case.

Theorem 1 *For a TBox \mathcal{T} , an ABox \mathcal{A} , and a TBox axiom α in $DL\text{-Lite}_{horn}$, deciding $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A}) \models \alpha$ is in coNP.*

If the size of \mathcal{A} is bounded, i.e., a constant k exists s.t. $|\mathcal{A}| \leq k$, then deciding $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A}) \models \alpha$ is in PTIME.

Having a bound on the size of \mathcal{A} and the PTIME complexity are practically relevant, as in ontology learning, the newly added example data is often small. Indeed, the PTIME complexity result holds as long as the number of assertions about each individual in \mathcal{A} is bounded.

For a practical algorithm to decide the entailment problem, we can further encode it to SAT. The main idea is to first encode the entailment problem of propositional revision as a Horn 2-QBF, and then eliminate universal quantifiers by encoding unit propagation of Horn-SAT into a SAT formula.

In what follows, we provide an algorithm to compute $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})$ in $DL\text{-Lite}_{core}$ using entailment checking as a sub-method. The algorithm verifies each possible entailment and add those entailed axioms to the result.

Algorithm 1: Contraction

Input: TBox \mathcal{T} and ABox \mathcal{A} in $DL\text{-Lite}_{core}$
Output: a $DL\text{-Lite}_{core}$ TBox \mathcal{T}'
 $\mathcal{T}' := \emptyset$;
foreach $B_1, B_2 \in \mathcal{B}$ s.t. $B_1 \neq B_2$ **do**
 if $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A}) \models B_1 \sqsubseteq B_2$ **then**
 $\mathcal{T}' := \mathcal{T}' \cup \{B_1 \sqsubseteq B_2\}$;
 end
 if $\mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A}) \models B_1 \sqsubseteq \neg B_2$ **then**
 $\mathcal{T}' := \mathcal{T}' \cup \{B_1 \sqsubseteq \neg B_2\}$;
 end
end
return \mathcal{T}' ;

Proposition 8 *Let \mathcal{T} be a TBox, \mathcal{A} be an ABox in $DL\text{-Lite}_{core}$, and \mathcal{T}' be the TBox returned by the algorithm Contraction. Then, $\mathcal{T}' \equiv \mathcal{T} \dot{-} \mathcal{T}_-(\mathcal{A})$.*

If the size of \mathcal{A} is bounded, the algorithm Contraction runs in polynomial time in the size of $\mathcal{T} \cup \mathcal{A}$. In particular, the number of possible axioms in $DL\text{-Lite}_{core}$ is quadratic to the number of basic concepts in $\mathcal{T} \cup \mathcal{A}$, and the entailment check is tractable as discussed before.

Conclusion

We have presented a novel approach for instance-driven ontology evolution in DL-Lite through TBox contraction according to newly added ABox instances, and the contraction incurs minimal change to the initial TBox measured in a model-theoretic manner. To the best of our knowledge, our work is the first attempt to address particularly the TBox contraction w.r.t. ABox assertions. To tackle this problem, we introduced the notion of ABox negation, which is defined as a TBox. Based on type semantics, we have developed an operator for TBox contraction. We showed that the proposed operator possesses several desired properties and have also developed efficient algorithms for reasoning with and computing the contraction.

We are currently working on a graph-based algorithm for entailment checking and computation of contraction in $DL\text{-Lite}_{core}$ with unbounded inputs. We have also obtained some preliminary results in extending our approach to $DL\text{-Lite}_{\mathcal{R}}$. For the latter, a key task is to extend the type semantics to $DL\text{-Lite}_{\mathcal{R}}$, which needs to take into account the additional non-propositional inferences introduced by role inclusions. Finally, it would be also interesting to look into practical applications of our approach in ontology learning.

Acknowledgments

This work was partially supported by Australian Research Council (ARC) under DP130102302 and DP1093652. Guilin Qi was partially supported by the NSFC grant 61272378.

References

- Artale, A.; Calvanese, D.; Kontchakov, R.; and Zakharyashev, M. 2009. The DL-Lite family and relations. *Journal of Artificial Intelligence Research* 36:1–69.
- Baader, F.; Ganter, B.; Sertkaya, B.; and Sattler, U. 2007. Completing description logic knowledge bases using formal concept analysis. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, 230–235.
- Calvanese, D.; Giacomo, G. D.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2007. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *Journal of Automatic Reasoning* 39(3):385–429.
- Calvanese, D.; Kharlamov, E.; Nutt, W.; and Zheleznyakov, D. 2010. Evolution of DL-Lite knowledge bases. In *Proceedings of the 9th International Semantic Web Conference (ISWC-10)*, 112–128.
- Cuenca Grau, B.; Ruiz, E. J.; Kharlamov, E.; and Zheleznyakov, D. 2012. Ontology evolution under semantic constraints. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR-12)*, 137–147.
- Delgrande, J. P., and Peppas, P. 2015. Belief revision in horn theories. *Artificial Intelligence* 218:1–22.
- Eiter, T., and Gottlob, G. 1992. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence* 57:227–270.
- Flouris, G.; Huang, Z.; Pan, J. Z.; Plexousakis, D.; and Wache, H. 2006. Inconsistencies, negations and changes in ontologies. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-06)*, 1295–1300.
- Fuhrmann, A., and Hansson, S. O. 1993. A survey of multiple contractions. *Journal of Logic Language and Information* 3(1):39–75.
- Haase, P., and Stojanovic, L. 2005. Consistent evolution of OWL ontologies. In *Proceedings of the 2nd European Semantic Web Conference (ESWC-05)*, 182–197.
- Hartel, F. W.; de Coronado, S.; Dionne, R.; Fragoso, G.; and Golbeck, J. 2005. Modeling a description logic vocabulary for cancer research. *Journal of Biomedical Informatics* 38(2):114–129.
- Horridge, M.; Parsia, B.; and Sattler, U. 2008. Laconic and precise justifications in OWL. In *Proceedings of the 7th International Semantic Web Conference (ISWC-08)*, 323–338.
- Kalyanpur, A.; Parsia, B.; Sirin, E.; and Cuenca Grau, B. 2006. Repairing unsatisfiable concepts in OWL ontologies. In *Proceedings of the 3rd European Semantic Web Conference (ESWC-06)*, 170–184.
- Katsuno, H., and Mendelzon, A. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3):263–294.
- Konev, B.; Lutz, C.; and Wolter, F. 2013. Exact learning of TBoxes in EL and DL-Lite. In *Proceedings of the 26th International Workshop on Description Logics*, 341–352.
- Kontchakov, R.; Wolter, F.; and Zakharyashev, M. 2010. Logic-based ontology comparison and module extraction, with an application to DL-Lite. *Artificial Intelligence* 174(15):1093–1141.
- Lehmann, J., and Hitzler, P. 2010. Concept learning in description logics using refinement operators. *Machine Learning* 78(1-2):203–250.
- Levi, I. 1991. *The Fixation of Belief and its Undoing*. Cambridge University Press.
- Ma, Y., and Distel, F. 2013. Concept adjustment for description logics. In *Proceedings of the 7th International Conference on Knowledge Capture (K-CAP'13)*, 65–72.
- Nikitina, N.; Rudolph, S.; and Glimm, B. 2012. Interactive ontology revision. *Journal of Web Semantics* 12:118–130.
- Pesquita, C., and Couto, F. M. 2012. Predicting the extension of biomedical ontologies. *PLoS Computational Biology* 8(9).
- Qi, G., and Du, J. 2009. Model-based revision operators for terminologies in description logics. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI-09)*, 891–897.
- Qi, G., and Yang, F. 2008. A survey of revision approaches in description logics. In *Proceedings of the 2nd International Conference of Web Reasoning and Rule Systems*, 74–88.
- Qi, G.; Haase, P.; Huang, Z.; Ji, Q.; Pan, J. Z.; and Völker, J. 2008. A kernel revision operator for terminologies - algorithms and evaluation. In *Proceedings of the 7th International Semantic Web Conference (ISWC-08)*, 419–434.
- Qi, G.; Wang, Z.; Wang, K.; Fu, X.; and Zhuang, Z. 2014. Approximating model-based ABox revision in DL-Lite: Theory and practice. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-15)*, (to appear).
- Ribeiro, M. M., and Wassermann, R. 2009. Base revision for ontology debugging. *Journal of Logic Computation* 19(5):721–743.
- Satoh, K. 1988. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, 455–462.
- Schlobach, S.; Huang, Z.; Cornet, R.; and van Harmelen, F. 2007. Debugging incoherent terminologies. *Journal of Automated Reasoning* 39(3):317–349.
- Wang, Z.; Wang, K.; and Topor, R. 2010. A new approach to knowledge base revision in DL-Lite. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI-10)*, 369–374.
- Zheleznyakov, D.; Calvanese, D.; Kharlamov, E.; and Nutt, W. 2010. Updating TBoxes in DL-Lite. In *Proceedings of the 23rd International Workshop on Description Logics*.
- Zhuang, Z. Q., and Pagnucco, M. 2012. Model based horn contraction. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR-12)*.
- Zhuang, Z.; Wang, Z.; Wang, K.; and Qi, G. 2014. Contraction and revision over DL-Lite tboxes. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI-14)*, 1149–1156.