

Towards Tractable and Practical ABox Abduction over Inconsistent Description Logic Ontologies

Jianfeng Du

Guangdong University of
Foreign Studies,
Guangzhou 510006, China
jfd@gdufs.edu.cn

Kewen Wang

Griffith University,
Brisbane, QLD 4111, Australia
k.wang@griffith.edu.au

Yi-Dong Shen

State Key Laboratory of
Computer Science,
Institute of Software,
Chinese Academy of Sciences,
Beijing 100190, China
ydshen@ios.ac.cn

Abstract

ABox abduction plays an important role in reasoning over description logic (DL) ontologies. However, it does not work with inconsistent DL ontologies. To tackle this problem while achieving tractability, we generalize ABox abduction from the classical semantics to an inconsistency-tolerant semantics, namely the Intersection ABox Repair (IAR) semantics, and propose the notion of *IAR-explanations* in inconsistent DL ontologies. We show that computing all minimal IAR-explanations is tractable in data complexity for first-order rewritable ontologies. However, the computational method may still not be practical due to a possibly large number of minimal IAR-explanations. Hence we propose to use preference information to reduce the number of explanations to be computed. In particular, based on the specificity of explanations, we introduce the notion of \subseteq_{cps} -*cminimal* IAR-explanations, which can be computed in a highly efficient way. Accordingly, we propose a tractable level-wise method for computing all \subseteq_{cps} -*cminimal* IAR-explanations in a first-order rewritable ontology. Experimental results on benchmarks of inconsistent ontologies show that the proposed method scales to tens of millions of assertions and can be of practical use.

1 Introduction

ABox abduction (Klarman, Endriss, and Schlobach 2011; Du et al. 2011a; Du, Wang, and Shen 2014), an adaptation of abductive reasoning (Eiter and Gottlob 1995) to description logics (DLs), has gained much attention recently since many ontology-based applications raise a requirement for explaining why an observation cannot be entailed by a DL ontology. Given a consistent DL ontology and an observation, such as a Boolean conjunctive query (BCQ), ABox abduction computes sets of assertions (namely *explanations*) whose appending to the ontology enforces the entailment of the observation while keeping the ontology consistent. The usefulness of ABox abduction has been identified in medical diagnosis (Elsenbroich, Kutz, and Sattler 2006), ontology quality control (Bada, Mungall, and Hunter 2008), conjunctive query answering (Borgida, Calvanese, and Rodriguez-Muro 2008), semantic matchmaking (Du et al. 2011b), etc.

ABox abduction does not work under inconsistency since anything can be entailed from an inconsistent DL ontology.

As pointed out in the literature e.g. (Du, Qi, and Shen 2013), inconsistency may often occur in some important ontology-based applications such as data integration and ontology population. In these applications it is also needed to explain why an observation cannot be a meaningful entailment of an inconsistent DL ontology under some inconsistency-tolerant semantics. Thus it is useful to extend ABox abduction to handle inconsistency. However, to the best of our knowledge this extension has not been explored in the literature.

To make ABox abduction works with inconsistency, we consider how to define explanations for observations in an inconsistent DL ontology. We focus on a general problem for ABox abduction, the query abduction problem (QAP) (Calvanese et al. 2013), where observations are BCQs and explanations may contain *fresh individuals* neither in the ontology nor in the observation. Since QAP often requires to compute all minimal explanations, we consider two requirements for the notion of minimal explanations under inconsistency. First, minimal explanations in an inconsistent DL ontology degenerate into traditional minimal explanations when the given ontology is consistent. Second, computing all minimal explanations is tractable at least in terms of *data complexity* which is measured by the size of the ABox only.

It has been shown (Du, Wang, and Shen 2014) that computing all traditional minimal explanations in a consistent first-order rewritable ontology is tractable in data complexity. We expect that this tractable result still holds when the ontology is inconsistent. Thus, we adapt an inconsistency-tolerant semantics, namely the Intersection ABox Repair (IAR) semantics, to defining meaningful explanations in inconsistent DL ontologies. The IAR semantics, first proposed in (Lembo et al. 2010), has been shown to be computable in polynomial time in data complexity for first-order rewritable ontologies (Lukasiewicz, Martinez, and Simari 2013).

Based on the IAR semantics, we require that an explanation \mathcal{E} for a BCQ in a possibly inconsistent ontology is a set of assertions such that appending it to the given ontology \mathcal{O} (with TBox \mathcal{T} and ABox \mathcal{A}) makes the union of \mathcal{T} and the intersection of all *ABox repairs* of $\mathcal{O} \cup \mathcal{E}$ entail the BCQ, where an ABox repair of $\mathcal{O} \cup \mathcal{E}$ is a maximal subset of the ABox of $\mathcal{O} \cup \mathcal{E}$ (namely $\mathcal{A} \cup \mathcal{E}$) that is consistent with \mathcal{T} . However, this condition alone does not ensure a minimal explanation to degenerate into a traditional one when \mathcal{O} is consistent. Hence we further require that no assertion in the

explanation \mathcal{E} is also in the ABox \mathcal{A} while appending \mathcal{E} to an arbitrary ABox repair R of \mathcal{O} does not introduce inconsistency, i.e., $R \cup \mathcal{E}$ is consistent with the TBox \mathcal{T} .

We call an explanation satisfying all the above conditions an *IAR-explanation*. We show that computing all minimal IAR-explanations for a BCQ in a first-order rewritable ontology is tractable in data complexity. However, the computational method may still not be practical because there can be a large number of minimal IAR-explanations for a BCQ, as shown in (Du, Wang, and Shen 2014) for traditional minimal explanations. Hence we propose to use preference information to reduce the number of explanations to be computed. Given a precedence relation \preceq between IAR-explanations, we call a minimal IAR-explanation \mathcal{E} \preceq -*minimal* if for all minimal IAR-explanations \mathcal{E}' , $\mathcal{E}' \preceq \mathcal{E}$ implies $\mathcal{E} \preceq \mathcal{E}'$. The representative explanations proposed in (Du, Wang, and Shen 2014) amount to \subseteq_s -minimal IAR-explanations in consistent DL ontologies, where for two explanations \mathcal{E}' and \mathcal{E} , we say $\mathcal{E}' \subseteq_s \mathcal{E}$ if there is a substitution θ for \mathcal{E}' , which replaces fresh individuals in \mathcal{E}' with fresh or existing individuals, such that $\mathcal{E}'\theta$ is a subset of \mathcal{E} .

As empirically found by Du, Wang, and Shen (2014), sometimes there can still be too many representative explanations to be computed. Hence we consider computing some but not all \preceq -minimal IAR-explanations. According to Occam's razor, we propose \preceq -*cminimal* IAR-explanations, which are \preceq -minimal IAR-explanations with the minimum cardinality. One may consider to compute all \subseteq_s -cminimal IAR-explanations. However, the minimality checking for these explanations is rather inefficient since it cannot be done without computing IAR-explanations that have larger cardinalities. For example, consider an ontology $\mathcal{O} = \emptyset$ and a BCQ $Q = \{r(a, x), r(x, y)\}$ where a is an individual, and x and y are variables. There are four minimal IAR-explanations for Q in \mathcal{O} , namely $\mathcal{E}_1 = \{r(a, a)\}$, $\mathcal{E}_2 = \{r(a, u_1), r(u_1, a)\}$, $\mathcal{E}_3 = \{r(a, u_1), r(u_1, u_1)\}$ and $\mathcal{E}_4 = \{r(a, u_1), r(u_1, u_2)\}$, where u_1 and u_2 are fresh individuals. When checking if \mathcal{E}_1 is \subseteq_s -cminimal, we need to consider \mathcal{E}_2 , \mathcal{E}_3 or \mathcal{E}_4 , which have larger cardinalities. In fact, since $\mathcal{E}_4 \subseteq_s \mathcal{E}_1$ but $\mathcal{E}_1 \not\subseteq_s \mathcal{E}_4$, \mathcal{E}_1 cannot be \subseteq_s -cminimal.

To achieve a more efficient (thus more practical) computational method, we introduce another concrete precedence relation \subseteq_{cps} , where for two explanations \mathcal{E}' and \mathcal{E} , we say $\mathcal{E}' \subseteq_{cps} \mathcal{E}$ if there is a substitution θ for \mathcal{E}' such that $\mathcal{E}'\theta$ has the same cardinality as \mathcal{E}' while $\mathcal{E}'\theta$ is a subset of \mathcal{E} . For the aforementioned example, both \mathcal{E}_1 and \mathcal{E}_4 are \subseteq_{cps} -minimal, but only \mathcal{E}_1 is \subseteq_{cps} -cminimal. We show that \subseteq_{cps} -cminimal IAR-explanations are computable without computing IAR-explanations that have larger cardinalities. Moreover, we propose a tractable (in data complexity) method for computing all \subseteq_{cps} -cminimal IAR-explanations in a first-order rewritable ontology. It works in an efficient level-wise manner. We also propose an important optimization to make \subseteq_{cps} -cminimal IAR-explanations computed as early as possible. Experimental results on benchmarks of inconsistent ontologies show that the proposed method is rather efficient and scales to tens of millions of assertions.

Due to the space limitation, proofs in this paper are only provided in our technical report (Du, Wang, and Shen 2015).

2 Preliminaries

We assume that the reader is familiar with DLs (Baader et al. 2003). A DL ontology consists of a TBox and an ABox, where the TBox consists of a finite set of axioms declaring the relations between concepts and roles, and the ABox consists of a finite set of assertions mainly declaring which individuals (resp. individual pairs) are instances of a concept (resp. a role). We assume that the *Unique Name Assumption* (Baader et al. 2003) is adopted and that an ABox contains only *basic* assertions which are concept assertions of the form $A(a)$ or role assertions of the form $r(a, b)$, where A is a concept name, r is a role name, and a and b are individuals. Other concept assertions and role assertions can be normalized to basic ones in a standard way. Let Σ be a set of concept names and role names. An ABox that contains only concept names and role names from Σ is called a Σ -ABox.

We use the traditional semantics for DLs e.g. given in (Baader et al. 2003), which coincides with the classical first-order semantics. A DL ontology \mathcal{O} is said to be *consistent*, denoted by $\mathcal{O} \not\models \perp$, if it has at least one model; otherwise, it is *inconsistent*, denoted by $\mathcal{O} \models \perp$. An ABox \mathcal{A} is said to be *consistent with* a TBox \mathcal{T} if $\mathcal{T} \cup \mathcal{A}$ is consistent.

A Boolean conjunctive query (BCQ) $\exists \vec{x} \Phi(\vec{x}, \vec{c})$ is made up of a conjunction of atoms $\Phi(\vec{x}, \vec{c})$ over concept names and role names, where \vec{x} are variables and \vec{c} are individuals. In this paper a BCQ is usually written and treated as a set of atoms. For example, the BCQ $\exists x A(x) \wedge B(x)$ is written as $\{A(x), B(x)\}$. The cardinality of a set S is denoted by $|S|$. A *substitution* θ for a BCQ Q is a mapping from variables in Q to individuals or variables. θ is *ground* if it maps variables to individuals only; in this case $Q\theta$ is called a *ground instance* of Q . A BCQ Q is *entailed* by a DL ontology \mathcal{O} if Q is satisfied by all models of \mathcal{O} , written $\mathcal{O} \models Q$.

The *query abduction problem* (simply *QAP*) (Calvanese et al. 2013) is a general problem for ABox abduction, where observations are BCQs and explanations may contain *fresh* individuals neither in the ontology nor in the observation.

Definition 1. Given a DL ontology with TBox \mathcal{T} and ABox \mathcal{A} , a BCQ Q and a set Σ of concept names and role names (called *abducible predicates*). We call $\mathcal{P} = (\mathcal{T}, \mathcal{A}, Q, \Sigma)$ an instance of QAP. An *explanation* for \mathcal{P} is a Σ -ABox \mathcal{E} such that $\mathcal{T} \cup \mathcal{A} \cup \mathcal{E} \models Q$ and $\mathcal{T} \cup \mathcal{A} \cup \mathcal{E} \not\models \perp$.

QAP often requires to compute all *minimal* explanations for an instance $\mathcal{P} = (\mathcal{T}, \mathcal{A}, Q, \Sigma)$. To compare with explanations containing different fresh individuals, fresh individuals are treated as variables. A *substitution* for an explanation \mathcal{E} is a mapping from fresh individuals in \mathcal{E} to existing or fresh individuals. A *renaming* for \mathcal{E} is a substitution for \mathcal{E} that maps different fresh individuals to different fresh individuals. We say $\mathcal{E}' \subset_r \mathcal{E}$ for two explanations \mathcal{E}' and \mathcal{E} , if there exists a renaming ρ for \mathcal{E}' such that $\mathcal{E}'\rho \subset \mathcal{E}$. Then minimal explanations for \mathcal{P} can be defined below.

Definition 2 (Du, Wang, and Shen 2014). A *minimal explanation* \mathcal{E} for $\mathcal{P} = (\mathcal{T}, \mathcal{A}, Q, \Sigma)$ is an explanation for \mathcal{P} such that there is no explanation \mathcal{E}' for \mathcal{P} fulfilling $\mathcal{E}' \subset_r \mathcal{E}$.

As shown in (Du, Wang, and Shen 2014), computing all minimal explanations for \mathcal{P} is tractable in data complexity

for *first-order rewritable* ontologies that are given by Definition 3. A first-order rewritable ontology has a TBox that can be translated to a Datalog[±] (Calì, Gottlob, and Lukasiewicz 2012) ontology. Such an ontology consists of finitely many *tuple generating dependencies* (TGDs) $\forall \vec{x} \forall \vec{y} \Phi(\vec{x}, \vec{y}) \rightarrow \exists \vec{z} \varphi(\vec{x}, \vec{z})$, *constraints* $\forall \vec{x} \Phi'(\vec{x}) \rightarrow \perp$, as well as *equality generating dependencies* (EGDs) $\forall \vec{x} \Phi'(\vec{x}) \rightarrow x_1 = x_2$, where $\Phi(\vec{x}, \vec{y})$, $\varphi(\vec{x}, \vec{z})$ and $\Phi'(\vec{x})$ are conjunctions of atoms, x_1 and x_2 occur in \vec{x} , and \perp denotes the truth constant false. By \mathcal{T}_D , \mathcal{T}_C and \mathcal{T}_E we denote the portions of a TBox \mathcal{T} that are translated to TGDs, constraints and EGDs, respectively.

Definition 3. A DL ontology is said to be *first-order rewritable* if its TBox \mathcal{T} can be translated to a Datalog[±] ontology and satisfies the following three conditions for an arbitrary BCQ Q and an arbitrary $\Sigma_{\mathcal{T}}$ -ABox \mathcal{A} , where $\Sigma_{\mathcal{T}}$ is the set of concept names and role names in \mathcal{T} :

- (1) $\mathcal{T} \cup \mathcal{A} \models Q$ if and only if $\mathcal{T}_D \cup \mathcal{A} \models Q$ or $\mathcal{T} \cup \mathcal{A} \models \perp$;
- (2) $\mathcal{T}_C \cup \mathcal{T}_E$ can be rewritten to a finite set of BCQs according to \mathcal{T}_D , denoted by $\gamma(\mathcal{T}_C \cup \mathcal{T}_E, \mathcal{T}_D)$, such that $\mathcal{T} \cup \mathcal{A} \models \perp$ if and only if $\mathcal{A} \models Q'$ for some $Q' \in \gamma(\mathcal{T}_C \cup \mathcal{T}_E, \mathcal{T}_D)$;
- (3) Q can be rewritten to a finite set of BCQs according to \mathcal{T}_D , denoted by $\tau(Q, \mathcal{T}_D)$, such that $\mathcal{T}_D \cup \mathcal{A} \models Q$ if and only if $\mathcal{A} \models Q'$ for some $Q' \in \tau(Q, \mathcal{T}_D)$.

An ontology expressed in most DLs in the DL-Lite family (Calvanese et al. 2007) is first-order rewritable (Calì, Gottlob, and Lukasiewicz 2012; Calì, Gottlob, and Pieris 2012). Since the DL-Lite family has become popular in many ontology-based applications and first-order rewritable ontologies are sufficient for these applications, we only consider first-order rewritable ontologies in this work.

3 Defining Explanations under Inconsistency

An inconsistent DL ontology has no model and entails anything. To define meaningful entailments, inconsistency-tolerant semantics have been proposed e.g. in (Lembo et al. 2010; Lukasiewicz, Martinez, and Simari 2013). Regarding QAP, it is helpful to see why a BCQ cannot be a meaningful entailment of an inconsistent DL ontology. Hence it is reasonable to adapt some inconsistency-tolerant semantics to defining explanations for QAP under inconsistency. We consider two requirements for the notion of minimal explanations under inconsistency. Firstly, to ensure that any computational method designed for these minimal explanations is still correct for traditional minimal explanations given by Definition 2, minimal explanations should degenerate into traditional ones when the given ontology is consistent. Secondly, to guarantee efficiency, computing all these minimal explanations should be tractable for a large class of ontologies e.g. the first-order rewritable ontologies.

To seek an adaptation that meets the above two requirements, we consider inconsistency-tolerant semantics proposed in (Lembo et al. 2010; Lukasiewicz, Martinez, and Simari 2013). These semantics are rather popular because they are based on the classical first-order semantics and do not depend on extra information (such as priorities and weights) about assertions. The most popular semantics is the *ABox Repair* (AR) semantics (Lembo et al. 2010). Let

\mathcal{O} be a possibly inconsistent DL ontology that has a consistent TBox \mathcal{T} and an ABox \mathcal{A} . The AR semantics defines that a BCQ Q is a meaningful entailment of \mathcal{O} , written $\mathcal{O} \models_{AR} Q$, if $\mathcal{T} \cup R \models Q$ for all *ABox repairs* R of \mathcal{O} , where an ABox repair of \mathcal{O} is a maximal subset of \mathcal{A} that is consistent with \mathcal{T} . However, it has been shown (Lembo et al. 2010) that deciding if $\mathcal{O} \models_{AR} Q$ is coNP-complete in data complexity even when \mathcal{O} is expressed in DL-Lite_{core}, the least expressive DL in the DL-Lite family. This implies that adapting the AR semantics to QAP cannot meet the second requirement. Hence we turn to another popular semantics, the *Intersection ABox Repair* (IAR) semantics (Lembo et al. 2010). Let $AR(\mathcal{T}, \mathcal{A})$ denote the set of ABox repairs of \mathcal{O} . The IAR semantics defines that a BCQ Q is a meaningful entailment of \mathcal{O} , written $\mathcal{O} \models_{IAR} Q$, if $\mathcal{T} \cup \bigcap_{R \in AR(\mathcal{T}, \mathcal{A})} R \models Q$. It has been shown (Lukasiewicz, Martinez, and Simari 2013) that deciding if $\mathcal{O} \models_{IAR} Q$ is tractable in data complexity for first-order rewritable ontologies. Hence we choose the IAR semantics for adaptation.

Based on the IAR semantics, an explanation \mathcal{E} for $\mathcal{P} = (\mathcal{T}, \mathcal{A}, Q, \Sigma)$ where \mathcal{T} is consistent should be an Σ -ABox such that $\mathcal{T} \cup \mathcal{A} \cup \mathcal{E} \models_{IAR} Q$, i.e., $\mathcal{T} \cup \bigcap_{R \in AR(\mathcal{T}, \mathcal{A} \cup \mathcal{E})} R \models Q$. However, this condition alone does not ensure a minimal explanation to degenerate into a traditional one when $\mathcal{T} \cup \mathcal{A}$ is consistent. To meet the first requirement, we compensate another condition that $\mathcal{E} \cap \mathcal{A} = \emptyset$ and $\mathcal{T} \cup \mathcal{E} \cup R \not\models \perp$ for all $R \in AR(\mathcal{T}, \mathcal{A})$. This condition is justifiable in the sense that \mathcal{E} is more doubtful than the ABox, while an arbitrary ABox repair is a subset of the ABox whose consistency should not be violated by more doubtful assertions. From the point of view of minimal conflicts, this condition is also justifiable. A *conflict* in a DL ontology \mathcal{O} with TBox \mathcal{T} and ABox \mathcal{A} is a subset of \mathcal{A} that is inconsistent with \mathcal{T} ; it is *minimal* if all its proper subsets are consistent with \mathcal{T} . By $MC(\mathcal{T}, \mathcal{A})$ we denote the set of minimal conflicts in \mathcal{O} . By the following lemma, we can interpret the condition as that appending \mathcal{E} to \mathcal{O} does not introduce new minimal conflicts.

Lemma 1. *Let \mathcal{O} be a DL ontology with a consistent TBox \mathcal{T} and an ABox \mathcal{A} . For any set \mathcal{E} of assertions such that $\mathcal{E} \cap \mathcal{A} = \emptyset$, we have $\mathcal{T} \cup \mathcal{E} \cup R \not\models \perp$ for all $R \in AR(\mathcal{T}, \mathcal{A})$ if and only if $MC(\mathcal{T}, \mathcal{A} \cup \mathcal{E}) = MC(\mathcal{T}, \mathcal{A})$.*

The following lemma further shows that checking the aforementioned condition in a first-order rewritable ontology is tractable in data complexity, since $\gamma(\mathcal{T}_C \cup \mathcal{T}_E, \mathcal{T}_D)$ is independent of \mathcal{A} and computing $MC(\mathcal{T}, \mathcal{A})$ is tractable in data complexity.

Lemma 2. *For a first-order rewritable ontology \mathcal{O} with TBox \mathcal{T} and ABox \mathcal{A} , let $\Lambda(\mathcal{O}) = \{Q\theta \mid Q \in \gamma(\mathcal{T}_C \cup \mathcal{T}_E, \mathcal{T}_D), \theta \text{ is a ground substitution for } Q \text{ such that } Q\theta \subseteq \mathcal{A}\}$. Then $MC(\mathcal{T}, \mathcal{A}) = \{S \in \Lambda(\mathcal{O}) \mid \nexists S' \in \Lambda(\mathcal{O}) : S' \subset S\}$.*

Lemma 2 also shows that checking if $\mathcal{T} \cup \mathcal{A} \cup \mathcal{E} \models_{IAR} Q$ is also tractable in data complexity for first-order rewritable ontologies, because when \mathcal{T} is consistent, we have
$$\bigcap_{R \in AR(\mathcal{T}, \mathcal{A} \cup \mathcal{E})} R = (\mathcal{A} \cup \mathcal{E}) \setminus \bigcup_{R \in AR(\mathcal{T}, \mathcal{A} \cup \mathcal{E})} ((\mathcal{A} \cup \mathcal{E}) \setminus R) = (\mathcal{A} \cup \mathcal{E}) \setminus \bigcup_{S \in MC(\mathcal{T}, \mathcal{A} \cup \mathcal{E})} S.$$

By combining the above conditions, we propose the following notion of *IAR-explanations*.

Definition 4. Given an instance $\mathcal{P} = (\mathcal{T}, \mathcal{A}, Q, \Sigma)$ of QAP where \mathcal{T} is consistent, an IAR-explanation \mathcal{E} for \mathcal{P} is a Σ -ABox such that $\mathcal{T} \cup \mathcal{A} \cup \mathcal{E} \models_{IAR} Q$, $\mathcal{E} \cap \mathcal{A} = \emptyset$ and $\mathcal{T} \cup \mathcal{E} \cup R \not\models \perp$ for all $R \in \text{AR}(\mathcal{T}, \mathcal{A})$.

We call an IAR-explanation \mathcal{E} for \mathcal{P} *minimal* if there is no IAR-explanation \mathcal{E}' for \mathcal{P} such that $\mathcal{E}' \subset_r \mathcal{E}$. Below we show that minimal IAR-explanations degenerate into traditional ones when the given ontology is consistent.

Proposition 1. Let $\mathcal{P} = (\mathcal{T}, \mathcal{A}, Q, \Sigma)$ where $\mathcal{T} \cup \mathcal{A}$ is consistent. Then the set of minimal IAR-explanations for \mathcal{P} is the same as the set of minimal explanations for \mathcal{P} given by Definition 2.

For a BCQ Q , let $\text{pred}(Q)$ denote the set of concept names and role names appearing in Q , and $\text{bipart}(Q)$ denote the set of different bipartitions of Q , where a bipartition of Q is a tuple of two BCQs (Q_1, Q_2) such that $Q_1 \cap Q_2 = \emptyset$ and $Q_1 \cup Q_2 = Q$. Given an arbitrary set \mathcal{S} of IAR-explanations, let $\text{reduce}_r(\mathcal{S})$ denote the set of IAR-explanations obtained from $\{\mathcal{E} \in \mathcal{S} \mid \nexists \mathcal{E}' \in \mathcal{S} : \mathcal{E}' \subset_r \mathcal{E}\}$ by deleting all duplicate IAR-explanations up to renaming of fresh individuals. By using bipartitions we develop a method for computing all minimal IAR-explanations in a first-order rewritable ontology. The following theorem shows this method together with its correctness and tractability in data complexity.

Theorem 1. Given $\mathcal{P} = (\mathcal{T}, \mathcal{A}, Q, \Sigma)$ where $\mathcal{T} \cup \mathcal{A}$ is a first-order rewritable ontology and \mathcal{T} is consistent, let $\Xi(\mathcal{P}) = \{Q_1\sigma\theta \mid Q' \in \tau(Q, \mathcal{T}_D), (Q_1, Q_2) \in \text{bipart}(Q'), \text{pred}(Q_1) \subseteq \Sigma, \sigma \text{ is a ground substitution for } Q_2 \text{ such that } Q_2\sigma \subseteq \mathcal{A} \setminus \bigcup_{S \in \text{MC}(\mathcal{T}, \mathcal{A})} S, \text{ and } \theta \text{ is a ground substitution for } Q_1\sigma \text{ such that } Q_1\sigma\theta \cap \mathcal{A} = \emptyset \text{ and } \text{MC}(\mathcal{T}, \mathcal{A} \cup Q_1\sigma\theta) = \text{MC}(\mathcal{T}, \mathcal{A})\}$. Then the set of minimal IAR-explanations for \mathcal{P} is $\text{reduce}_r(\Xi(\mathcal{P}))$ up to renaming of fresh individuals, and can be computed in polynomial time in data complexity.

4 Exploiting Preference Information

There may be a large number of minimal IAR-explanations for a BCQ, as shown in (Du, Wang, and Shen 2014) for traditional minimal explanations. For practicality, it is needed to reduce the number of explanations to be computed. Thus, we propose to use preference information on minimal (IAR-) explanations. Given a precedence relation \preceq between explanations, we consider computing \preceq -minimal ones, defined below, among minimal explanations.

Definition 5. Let \preceq be a precedence relation between (IAR-) explanations for \mathcal{P} . A minimal (IAR-) explanation \mathcal{E} for \mathcal{P} is said to be \preceq -minimal if for all minimal (IAR-) explanations \mathcal{E}' for \mathcal{P} , $\mathcal{E}' \preceq \mathcal{E}$ implies $\mathcal{E} \preceq \mathcal{E}'$.

One concrete precedence relation, written \subseteq_s here, is proposed in (Du, Wang, and Shen 2014), where we say $\mathcal{E}' \subseteq_s \mathcal{E}$ for two explanations \mathcal{E}' and \mathcal{E} if there is a substitution θ for \mathcal{E}' such that $\mathcal{E}'\theta \subseteq \mathcal{E}$. The \subseteq_s -minimal explanations in a consistent DL ontology are called *representative explanations* in (Du, Wang, and Shen 2014). It has been shown that the number of representative explanations can often be much smaller than that of minimal explanations. However, sometimes the number of representative explanations can still be

so large that computing all representative explanations cannot be done in hours (Du, Wang, and Shen 2014). Hence we further consider special \preceq -minimal explanations. According to Occam's razor, it is natural to consider \preceq -cminimal explanations defined below.

Definition 6. A \preceq -minimal (IAR-) explanation \mathcal{E} for \mathcal{P} is said to be \preceq -cminimal if for all \preceq -minimal (IAR-) explanations \mathcal{E}' for \mathcal{P} , $|\mathcal{E}| \leq |\mathcal{E}'|$.

By Proposition 1, \preceq -minimal IAR-explanations and \preceq -cminimal IAR-explanations also degenerate into \preceq -minimal explanations and \preceq -cminimal explanations, respectively, when the given ontology is consistent.

The simple example given in Section 1 shows that it is hard to compute all \subseteq_s -cminimal IAR-explanations without computing all \subseteq_s -minimal ones beforehand. To achieve more efficient computational methods, we consider other \preceq . We observe that if $\mathcal{E}_2 \not\preceq \mathcal{E}_1$ whenever $|\mathcal{E}_1| < |\mathcal{E}_2|$, then checking if an IAR-explanation is \preceq -cminimal need not consider IAR-explanations with larger cardinalities. Consider why $\mathcal{E}_4 \subseteq_s \mathcal{E}_1$ for $\mathcal{E}_4 = \{r(a, u_1), r(u_1, u_2)\}$ and $\mathcal{E}_1 = \{r(a, a)\}$ given by the example in Section 1. The proof of $\mathcal{E}_4 \subseteq_s \mathcal{E}_1$ relies on a substitution $\theta = \{u_1 \mapsto a, u_2 \mapsto a\}$ such that $|\mathcal{E}_4\theta| < |\mathcal{E}_4|$, which is not as intuitive as substitutions σ such that $|\mathcal{E}_4\sigma| = |\mathcal{E}_4|$. Hence we introduce the notion of *cardinality-preserving* substitutions. A cardinality-preserving substitution θ for an explanation \mathcal{E} is a substitution for \mathcal{E} such that $|\mathcal{E}\theta| = |\mathcal{E}|$. We propose another concrete precedence relation \subseteq_{cps} , where we say $\mathcal{E}' \subseteq_{\text{cps}} \mathcal{E}$ if there is a cardinality-preserving substitution θ for \mathcal{E}' such that $\mathcal{E}'\theta \subseteq \mathcal{E}$. It can be seen that $\mathcal{E}' \not\subseteq_{\text{cps}} \mathcal{E}$ if $|\mathcal{E}'| < |\mathcal{E}'|$. The following lemma shows that checking if an IAR-explanation is \subseteq_{cps} -minimal needs to consider only \subseteq_{cps} -minimal IAR-explanations with equal or smaller cardinalities.

Lemma 3. By $\mathcal{E}' \subseteq_{\text{cps}} \mathcal{E}$ we simply denote $\mathcal{E} \subseteq_{\text{cps}} \mathcal{E}'$ and $\mathcal{E}' \not\subseteq_{\text{cps}} \mathcal{E}$. An IAR-explanation \mathcal{E} for \mathcal{P} is \subseteq_{cps} -minimal if and only if there is no \subseteq_{cps} -minimal IAR-explanation \mathcal{E}' for \mathcal{P} such that $|\mathcal{E}'| \leq |\mathcal{E}|$ and $\mathcal{E}' \subset_{\text{cps}} \mathcal{E}$.

Checking if an IAR-explanation is \subseteq_{cps} -cminimal also needs to consider only \subseteq_{cps} -minimal IAR-explanations with equal or smaller cardinalities, because a \subseteq_{cps} -minimal IAR-explanation \mathcal{E} is \subseteq_{cps} -cminimal if and only if there is no \subseteq_{cps} -minimal IAR-explanation \mathcal{E}' such that $|\mathcal{E}'| < |\mathcal{E}|$.

5 Computing all \subseteq_{cps} -cminimal Explanations

Given an instance $\mathcal{P} = (\mathcal{T}, \mathcal{A}, Q, \Sigma)$ of QAP, where $\mathcal{T} \cup \mathcal{A}$ is a first-order rewritable ontology and \mathcal{T} is consistent, we develop a level-wise method for computing all \subseteq_{cps} -cminimal IAR-explanations for \mathcal{P} . At level k , this method computes a set of IAR-explanations for \mathcal{P} containing all \subseteq_{cps} -minimal IAR-explanations for \mathcal{P} that have the same cardinality k . We exploit *restrictive substitutions* to avoid generating IAR-explanations that are not \subseteq_{cps} -minimal. A ground substitution σ for a BCQ Q is called a *restrictive substitution* for Q in \mathcal{P} if it maps variables to fresh individuals not in \mathcal{P} or existing individuals in Q only. Let $\text{reduce}_{\text{cps}}(\mathcal{S})$ denote the set of IAR-explanations obtained from $\{\mathcal{E} \in \mathcal{S} \mid \nexists \mathcal{E}' \in \mathcal{S} : \mathcal{E}' \subset_{\text{cps}} \mathcal{E}\}$ by deleting all duplicate IAR-explanations up to

renaming of fresh individuals. The following lemma shows the process performed at level k and its correctness.

Lemma 4. *Let $\Phi(\mathcal{P}, k) = \{Q_1\sigma\theta \mid Q' \in \tau(Q, \mathcal{T}_D), (Q_1, Q_2) \in \text{bipart}(Q'), \text{pred}(Q_1) \subseteq \Sigma, \sigma \text{ is a ground substitution for } Q_2 \text{ such that } Q_2\sigma \subseteq \mathcal{A} \setminus \bigcup_{S \in \text{MC}(\mathcal{T}, \mathcal{A})} S, \text{ and } \theta \text{ is a restrictive substitution for } Q_1\sigma \text{ in } \mathcal{P} \text{ such that } |Q_1\sigma\theta| = k, Q_1\sigma\theta \cap \mathcal{A} = \emptyset \text{ and } \text{MC}(\mathcal{T}, \mathcal{A} \cup Q_1\sigma\theta) = \text{MC}(\mathcal{T}, \mathcal{A})\}$. Then $\text{reduce}_{\text{cps}}(\Phi(\mathcal{P}, k))$ is a set of IAR-explanations for \mathcal{P} containing all different \subseteq_{cps} -minimal IAR-explanations \mathcal{E} for \mathcal{P} such that $|\mathcal{E}| = k$ up to renaming of fresh individuals.*

Let $\text{maxc}(\mathcal{P})$ denote the maximum cardinality of BCQs in $\tau(Q, \mathcal{T}_D)$. To compute all \subseteq_{cps} -minimal IAR-explanations for \mathcal{P} , we can compute $\text{reduce}_{\text{cps}}(\Phi(\mathcal{P}, k))$ for k increasing from 0 to $\text{maxc}(\mathcal{P})$. Once we find some k such that $\text{reduce}_{\text{cps}}(\Phi(\mathcal{P}, k)) \neq \emptyset$, we can output $\text{reduce}_{\text{cps}}(\Phi(\mathcal{P}, k))$ as the set of \subseteq_{cps} -minimal IAR-explanations for \mathcal{P} . The correctness and tractability (in data complexity) of this level-wise method are shown in the following theorem.

Theorem 2. *Let $k_m = \min\{k \mid 1 \leq k \leq \text{maxc}(\mathcal{P}), \text{reduce}_{\text{cps}}(\Phi(\mathcal{P}, k)) \neq \emptyset\}$. Then the set of \subseteq_{cps} -minimal IAR-explanations for \mathcal{P} is $\text{reduce}_{\text{cps}}(\Phi(\mathcal{P}, k_m))$ up to renaming of fresh individuals, and can be computed in polynomial time in data complexity.*

In the level-wise method, any element of $\Phi(\mathcal{P}, k_m)$ cannot be output before the final $\text{reduce}_{\text{cps}}$ filter. To make explanations output as early as possible, we introduce an important optimization called the *prompt-output* optimization. At level k , we compute $\Phi(\mathcal{P}, k)$ in a BCQ-wise manner. For a rewritten BCQ Q' in $\tau(Q, \mathcal{T}_D)$, we first compute the set \mathcal{S} of IAR-explanations in $\Phi(\mathcal{P}, k)$ that are derived from Q' , then compute $\text{reduce}_{\text{cps}}(\mathcal{S})$ and for every $\mathcal{E} \in \text{reduce}_{\text{cps}}(\mathcal{S})$, check if \mathcal{E} is a ground instance of some subset of some BCQ Q'' in $\tau(Q, \mathcal{T}_D)$ other than Q' ; if so, we store \mathcal{E} , otherwise we promptly output \mathcal{E} as a \subseteq_{cps} -minimal IAR-explanation. At last, a final $\text{reduce}_{\text{cps}}$ filter is performed on all stored IAR-explanations. This optimization together with its correctness are formalized in the following theorem.

Theorem 3. *Suppose there is no \subseteq_{cps} -minimal IAR-explanation \mathcal{E} for \mathcal{P} such that $|\mathcal{E}| < k$. Let Q' be a BCQ in $\tau(Q, \mathcal{T}_D)$ and $\Psi(\mathcal{T}, \mathcal{A}, Q', \Sigma, k) = \{Q_1\sigma\theta \mid (Q_1, Q_2) \in \text{bipart}(Q'), \text{pred}(Q_1) \subseteq \Sigma, \sigma \text{ is a ground substitution for } Q_2 \text{ such that } Q_2\sigma \subseteq \mathcal{A} \setminus \bigcup_{S \in \text{MC}(\mathcal{T}, \mathcal{A})} S, \text{ and } \theta \text{ is a restrictive substitution for } Q_1\sigma \text{ in } \mathcal{P} \text{ such that } |Q_1\sigma\theta| = k, Q_1\sigma\theta \cap \mathcal{A} = \emptyset \text{ and } \text{MC}(\mathcal{T}, \mathcal{A} \cup Q_1\sigma\theta) = \text{MC}(\mathcal{T}, \mathcal{A})\}$. Then, $\mathcal{E} \in \text{reduce}_{\text{cps}}(\Psi(\mathcal{T}, \mathcal{A}, Q', \Sigma, k))$ is a \subseteq_{cps} -minimal IAR-explanation for \mathcal{P} if there is no $Q'' \in \tau(Q, \mathcal{T}_D) \setminus \{Q'\}$ and ground substitution θ for Q'' such that $\mathcal{E} \subseteq Q''\theta$.*

6 Experimental Evaluation

We implemented the proposed method with the prompt-output optimization in Java, using the Requiem (Pérez-Urbina, Motik, and Horrocks 2010) API for query rewriting and the MySQL system to store and access ABoxes. Six benchmark ontologies that are almost first-order rewritable were used. One is Semintec and the others are LUBMn

Table 1: The statistics about test ontologies

Ontology	#C	#R	#TA	#AA	#I
Semintec+0 \sim	60	16	203	65,240 \sim	17,941 \sim
Semintec+400				65,795	18,096
LUBM10+0 \sim	43	32	158	1,311,409 \sim	207,426 \sim
LUBM10+400				1,311,982	207,599
LUBM1+400 \sim	43	32	158	102,707 \sim	17,174 \sim
LUBM100+400				13,825,027	2,179,956

Note: #C/#R/#TA/#AA/#I is the number of concept names/role names/TBox axioms/ABox assertions/individuals.

($n = 1, 5, 10, 50, 100$) from the Lehigh University Benchmark (Guo, Pan, and Heflin 2005), where n is the number of universities. We removed a few axioms that Requiem cannot handle, making all ontologies first-order rewritable. Since LUBMn has no constraints in the TBox and cannot be made inconsistent by adding assertions, we added to LUBMn disjointness axioms for every two sibling concept names in the concept hierarchy of the LUBM TBox whenever these two concept names have no common instances in any LUBMn. All modified ontologies are still consistent, hence we used the `Injector` tool provided by (Du, Qi, and Shen 2013) to insert conflicts. By $\mathcal{O}+m$ we denote the ontology obtained from \mathcal{O} by inserting m conflicts. We generated Semintec+ m and LUBM10+ m ($m = 0, 100, 200, 300, 400$) to test the method against different number of conflicts, and generated LUBMn+400 ($n = 1, 5, 10, 50, 100$) to test the method against different number of universities. The statistics of test ontologies are given in Table 1. All experiments were conducted on a laptop having Dual-Core 2.20GHz CPU and 4GB RAM, with the maximum Java heap size set to 1GB.

For each Semintec+ m ontology, we used the same 50 BCQs as observations, where ten BCQs were randomly generated from each of the five benchmark conjunctive queries (CQs) of Semintec given in (Du, Qi, and Shen 2013). None of the generated BCQs is entailed by Semintec+0. Moreover, appending any of them to Semintec+0 does not render Semintec+0 inconsistent. For each LUBMn+ m ontology, we used the same 140 BCQs as observations, where ten BCQs were randomly generated from each of the 14 benchmark CQs of LUBM given in (Guo, Pan, and Heflin 2005). None of the generated BCQs is entailed by LUBM1+0. Moreover, appending any of them to LUBM1+0 does not render LUBM1+0 inconsistent. Let Gk denote the set of BCQs generated from the k^{th} benchmark CQ. We set all concept names and role names as abducible predicates.

Our implementation handles one ontology by two phases. In the first phase, it computes all minimal conflicts. In the second phase, it handles all generated BCQs in turn by reusing minimal conflicts. We set a time limit of 1000 seconds for handling one BCQ. The main results are provided in Figure 1, where the curves named MC show the execution time for computing all minimal conflicts in the first phase, and the other curves show the statistics in the second phase.

A curve named Gk in Part (a) shows the average execution time for computing all \subseteq_{cps} -minimal IAR-explanations for a BCQ in Gk on Semintec+ m against increasing number m of conflicts. While the execution time for computing all minimal conflicts increases with the number of conflicts, the

execution time for handling a BCQ does not; especially, for each of G3, G4 and G5 the average execution time is rather stable when the number of conflicts increases. Each curve Gk in Part (b) shows the average execution time for computing all \subseteq_{cps} -minimal IAR-explanations for a BCQ in Gk on LUBM10+ m against increasing number m of conflicts, where the execution time for a BCQ in G9 does not include that for rewriting the BCQ. The time for rewriting a BCQ in G9 is quite long (about 2100 seconds) possibly due to some issues in Requiem. In general handling a BCQ either takes a rather stable time or has a similar scalability as computing all minimal conflicts. The failure cases for G6 and G8 are caused by a large number of \subseteq_{cps} -minimal IAR-explanations. This can be verified by Part (c), which shows the average number of \subseteq_{cps} -minimal IAR-explanations output within 1000 seconds in failure cases. These results show that the proposed method is efficient in computing \subseteq_{cps} -minimal IAR-explanations, while this efficiency is only slightly affected by the number of conflicts.

Each curve Gk in Part (d) shows the average execution time for computing all \subseteq_{cps} -minimal IAR-explanations for a BCQ in Gk on LUBM n +400 against increasing number n of universities, where the execution time for a BCQ in G9 does not include that for rewriting the BCQ. In G6, G8 and G9, there exist BCQs that cannot be handled in 1000 seconds on some test ontologies. Part (e) shows the average number of \subseteq_{cps} -minimal IAR-explanations output within 1000 seconds in failure cases. It can be seen that all failure cases are caused by too many \subseteq_{cps} -minimal IAR-explanations to be computed. But in all failure cases, our method still outputs tens of thousands of \subseteq_{cps} -minimal IAR-explanations in 1000 seconds. In general, handling a BCQ has a similar scalability as computing all minimal conflicts and scales well to large ontologies. In particular, most BCQs are handled in 100 seconds even on LUBM100+400. These results show that the proposed method is rather efficient and scales to tens of millions of assertions.

7 Related Work

By now there are only a few studies on ABox abduction. The complexity for QAP is systematically studied in (Calvanese et al. 2013), but no method for computing all minimal explanations is provided there. The problem of computing all minimal explanations is initially addressed in (Klarman, Endriss, and Schlobach 2011), but the method proposed there may not terminate since the size of a minimal explanation in their sense can be infinite. This termination issue is addressed in (Du et al. 2011a) by introducing abducible predicates that are concept or role names to guarantee finiteness of minimal explanations. Accordingly, a sound and possibly incomplete method is proposed there to compute all minimal explanations. The method is then extended to handle more kinds of abducible predicates that can be arbitrary concepts (Du et al. 2012). A first tractable (in data complexity) method is proposed in (Du, Wang, and Shen 2014) to compute exactly all minimal explanations for QAP. The method brings some basic ideas to this work, such as restricting the given ontology to be first-order rewritable, but it is unable to handle inconsistency. Differing from existing approaches to

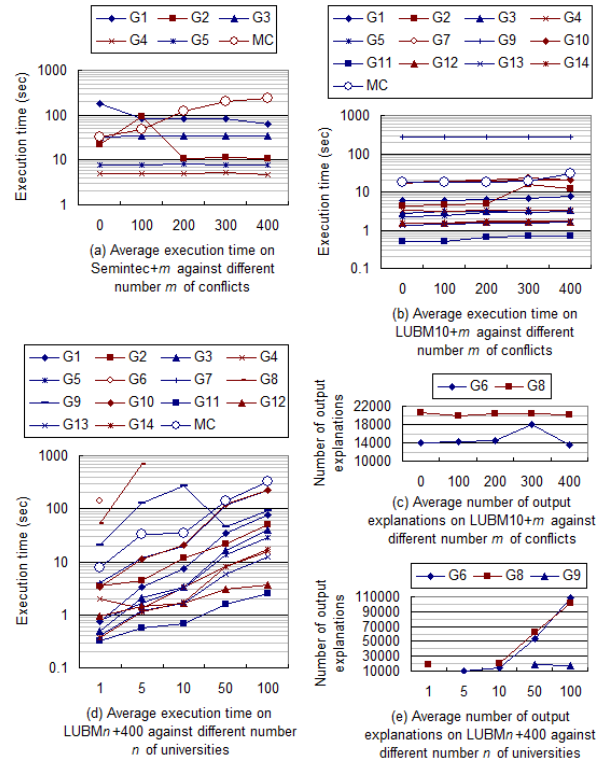


Figure 1: Experimental results for the proposed method

ABox abduction, our proposal is able to coherently handle both consistent and inconsistent ontologies while still allowing efficient computation of all minimal explanations.

8 Conclusion and Future Work

In this work we have proposed an approach to ABox abduction over inconsistent DL ontologies. There are three main contributions. First of all, we proposed the notion of IAR-explanations for BCQs in a possibly inconsistent DL ontology whose TBox is consistent. Minimal IAR-explanations degenerate into traditional ones when the given ontology is consistent, while computing all of them is tractable in data complexity for first-order rewritable ontologies. Secondly, for practicality we proposed to use preference information on IAR-explanations and introduced a new type of preferred explanations, called \subseteq_{cps} -minimal IAR-explanations. These explanations can be computed in a more efficient way than \subseteq_{s} -minimal IAR-explanations, which correspond to representative explanations (Du, Wang, and Shen 2014) that have the minimum cardinality. Finally, we proposed a level-wise method for computing all \subseteq_{cps} -minimal IAR-explanations in a first-order rewritable ontology. The method is tractable in data complexity. It is also shown to be rather efficient and scalable to tens of millions of assertions in our experiments. For future work, in order to make ABox abduction practical for more applications, we plan to identify some classes of DL ontologies that are not first-order rewritable but still guarantee the tractability in computing all minimal IAR-explanations.

Acknowledgements

This work is partly supported by NSFC grants (61375056, 61005043 and 61379043), Business Intelligence Key Team of Guangdong University of Foreign Studies (TD1202), Guangdong Natural Science Foundation (S2013010012928), China National 973 program (2014CB340301) as well as Australian Research Council (DP110101042 and DP130102302).

References

- Baader, F.; Calvanese, D.; McGuinness, D. L.; Nardi, D.; and Patel-Schneider, P. F., eds. 2003. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.
- Bada, M.; Mungall, C.; and Hunter, L. 2008. A call for an abductive reasoning feature in owl-reasoning tools toward ontology quality control. In *Proceedings of the 5th OWLED Workshop on OWL: Experiences and Directions*.
- Borgida, A.; Calvanese, D.; and Rodriguez-Muro, M. 2008. Explanation in the DL-Lite family of description logics. In *Proceedings of the 7th International Conference On Ontologies, DataBases, and Applications of Semantics (ODBASE)*, 1440–1457.
- Calì, A.; Gottlob, G.; and Lukasiewicz, T. 2012. A general datalog-based framework for tractable query answering over ontologies. *Journal of Web Semantics* 14:57–83.
- Calì, A.; Gottlob, G.; and Pieris, A. 2012. Towards more expressive ontology languages: The query answering problem. *Artificial Intelligence* 193:87–128.
- Calvanese, D.; Giacomo, G.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2007. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *Journal of Automated Reasoning* 39(3):385–429.
- Calvanese, D.; Ortiz, M.; Simkus, M.; and Stefanoni, G. 2013. Reasoning about explanations for negative query answers in DL-Lite. *Journal of Artificial Intelligence Research* 48:635–669.
- Du, J.; Qi, G.; Shen, Y.; and Pan, J. Z. 2011a. Towards practical ABox abduction in large OWL DL ontologies. In *Proceedings of the 25th National Conference on Artificial Intelligence (AAAI)*, 1160–1165.
- Du, J.; Wang, S.; Qi, G.; Pan, J. Z.; and Hu, Y. 2011b. A new matchmaking approach based on abductive conjunctive query answering. In *Proceedings of the 1st Joint International Semantic Technology Conference (JIST)*, 144–159.
- Du, J.; Qi, G.; Shen, Y.; and Pan, J. Z. 2012. Towards practical ABox abduction in large description logic ontologies. *International Journal on Semantic Web and Information Systems* 8(2):1–33.
- Du, J.; Qi, G.; and Shen, Y. 2013. Weight-based consistent query answering over inconsistent *SHIQ* knowledge bases. *Knowledge and Information Systems* 34(2):335–371.
- Du, J.; Wang, K.; and Shen, Y. 2014. A tractable approach to ABox abduction over description logic ontologies. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 1034–1040.
- Du, J.; Wang, K.; and Shen, Y. 2015. Towards tractable and practical ABox abduction over inconsistent description logic ontologies. Technical report, Guangdong University of Foreign Studies. <http://www.dataminingcenter.net/abduction/AAAI15-TR.pdf>.
- Eiter, T., and Gottlob, G. 1995. The complexity of logic-based abduction. *Journal of the ACM* 42(1):3–42.
- Elsenbroich, C.; Kutz, O.; and Sattler, U. 2006. A case for abductive reasoning over ontologies. In *Proceedings of the 3rd OWLED Workshop on OWL: Experiences and Directions*.
- Guo, Y.; Pan, Z.; and Heflin, J. 2005. LUBM: A benchmark for OWL knowledge base systems. *Journal of Web Semantics* 3(2–3):158–182.
- Klarman, S.; Endriss, U.; and Schlobach, S. 2011. ABox abduction in the description logic *ALC*. *Journal of Automated Reasoning* 46(1):43–80.
- Lembo, D.; Lenzerini, M.; Rosati, R.; Ruzzi, M.; and Savo, D. F. 2010. Inconsistency-tolerant semantics for description logics. In *Proceedings of the 4th International Conference on Web Reasoning and Rule Systems*, 103–117.
- Lukasiewicz, T.; Martinez, M. V.; and Simari, G. I. 2013. Complexity of inconsistency-tolerant query answering in datalog+/- . In *OTM Conferences - Confederated International Conferences: CoopIS, DOA-Trusted Cloud, and ODBASE*, 488–500.
- Pérez-Urbina, H.; Motik, B.; and Horrocks, I. 2010. Tractable query answering and rewriting under description logic constraints. *Journal of Applied Logic* 8(2):186–209.